Evidenční číslo smlouvy: MUCHP00SV07C

Smlouva o dílo

město Cheb
Sídlo: náměstí Krále Jiřího z Poděbrad 1/14, 350 20 Cheb
IČO: 00253979
jednající: Ing. Jan Vrba, starosta města Chebu
Bankovní spojení:
Číslo účtu:

kontaktní osoba:
email:
na straně jedné jako „**Objednatel**“

a

Aricoma Systems a.s.
Sídlo: Hornopolní 3322/34, Moravská Ostrava, 702 00 Ostrava
Korespondenční adresa: Chebská 73/48, 360 06 Karlovy Vary
IČO: 04308697
DIČ: CZ04308697
jednající: Zdeněk Chobot, na základě plné moci
Bankovní spojení:
Číslo účtu:
kontaktní osoba:
tel./fax kontaktní osoby:
e-mail:
zápis ve veřejném rejstříku: OR vedený Krajským soudem v Ostravě, sp. zn. B 11012
na straně druhé jako „**Zhotovitel**“, nebo „dodavatel“
(Objednatel a Zhotovitel jsou dále společně též označováni jako „**Strany**“ nebo „**Smluvní strany“** nebo kdokoli z nich jednotlivě též „**Strana**“ nebo „**Smluvní strana**“)

uzavírají v souladu s § 2358 a § 2371 a násl. a § 2586 a násl. z. č. 89/2012 Sb., občanského zákoníku, ve znění pozdějších předpisů (dále jen „**občanský zákoník**“), a zákonem č. 121/2000 Sb., o právu autorském, o právech souvisejících s právem autorským a o změně některých zákonů, ve znění pozdějších předpisů (dále jen „**autorský zákon**“) tuto Smlouvu o dílo (dále jen „**Smlouva**“).

Preambule

Předmět veřejné zakázky je spolufinancován Evropskou unií z Národního plánu obnovy ČR, výzva č. 41 Kybernetická bezpečnost – obce, název projektu je „Kybernetická bezpečnost města Cheb“, registrační číslo CZ.31.2.0/0.0/0.0/23_093/0008706.

Výběr Zhotovitele plnění dle této Smlouvy byl proveden Objednatelem v nadlimitním zadávacím řízení realizovaného dle zákona č. 134/2016 Sb., o zadávání veřejných zakázek, ve znění pozdějších předpisů (dále jen „**ZZVZ**“).

Objednatel vybral v zadávacím řízení veřejné zakázky s názvem „Kybernetická bezpečnost města Cheb 2024+“ a uveřejněné na Věstníku veřejných zakázek dne 11. 12. 2024 pod ev. číslem Z2024-062084 (dále jen „**Veřejná zakázka**“) nabídku Zhotovitele na realizaci zakázky vyhodnocenou jako nejvýhodnější.

0. Definice a Úvodní ustanovení

Definice.

Není-li dále výslovně uvedeno jinak, následující termíny jsou definovány v této Smlouvě takto:

„**Nabídka**“ znamená nabídku Zhotovitele doručenou Objednateli v rámci Zadávacího řízení;

„**Dodávky**“ znamenají dodávky a služby poskytované Zhotovitelem Objednateli dle této Smlouvy, specifikované níže v čl. II této Smlouvy;

„Software“ znamená veškeré systémové a aplikační programové vybavení, potřebné k řádnému, plně funkčnímu, nepřetržitému a bezporuchovému fungování předmětu plnění, které bude předmětem Dodávek.

„Právní předpisy“ znamená všechny platné a účinné obecně závazné právní předpisy České republiky a Evropské unie, a to zejména předpisy související s poskytováním Dodávek dle této Smlouvy;

„Spor“ znamená jakýkoliv spor vzniklý ze Smlouvy nebo v souvislosti s ní;

„Poddodavatel“ znamená jakoukoli právnickou nebo fyzickou osobu, s níž Zhotovitel uzavřel smlouvu, na jejímž základě bude taková osoba provádět plnění předmětu této Smlouvy nebo její části;

„Vyšší moc“ znamená mimořádnou událost nebo okolnost, kterou nemohla žádná ze Stran před uzavřením Smlouvy předvídat, která je mimo kontrolu kterékoliv Strany a nebyla způsobena úmyslně nebo z nedbalosti jednáním nebo opomenutím kterékoliv Strany a která podstatným způsobem ztěžuje nebo znemožňuje plnění povinností dle Smlouvy kteroukoliv ze Stran. Takovými událostmi nebo okolnostmi jsou zejména, nikoliv však výlučně, válka, teroristický útok, občanské nepokoje, vzpoura, přítomnost ionizujícího nebo radioaktivního záření, požár, výbuch, záplava či jiné živelné nebo přírodní katastrofy, pandemie. Výslovně se stanoví, že Vyšší mocí není stávka personálu Zhotovitele ani hospodářské poměry Stran;

„Zadávací dokumentace“ znamená zadávací dokumentaci Veřejné zakázky;

(B) Výklad

Slova v jednotném čísle rovněž zahrnují množné číslo a slova v množném čísle zahrnují i číslo jednotné;

Ustanovení obsahující slovo „souhlasit“, „souhlas“ nebo „dohoda“ nebo slova podobného významu vyžadují, aby souhlas nebo dohoda byly učiněny písemně.

„Písemný“ nebo „písemně“ znamená psaný rukou, strojem, tištěný, případně zhotovený elektronicky a existující ve formě trvalého záznamu (tím se rozumí i elektronická pošta).

Pokud se v textu této Smlouvy vyskytuje spojení „poskytování Dodávek“ a z příslušného ustanovení nevyplývá jinak, rozumí se Dodávkou i zajištění služeb nezbytných pro zajištění funkčnosti předmětu díla dle požadavků Zadávací dokumentace (tj. zejména po dobu udržitelnosti).

Výklad veškerých pojmů a ujednání bude prováděn s ohledem na účel a cíle Veřejné zakázky, na jejímž základě byla uzavřena tato Smlouva, které přímo či nepřímo vyplývají ze Zadávací dokumentace nebo této Smlouvy.

(C) Komunikace mezi Stranami

Kdykoliv se v této Smlouvě vyžaduje vyhotovení nebo vystavení souhlasů, osvědčení, svolení, rozhodnutí, oznámení a žádosti jakoukoliv osobou, tato sdělení musejí být doručena na kontaktní adresy uvedené v čl. XIII. a způsobem uvedeným v čl. XIV. této Smlouvy.

Veškerá komunikace podle Smlouvy bude probíhat výlučně v českém nebo slovenském jazyce.

(D) Licence

Zhotovitel prohlašuje, že je oprávněn poskytnout Objednateli licenci k Software dle této Smlouvy.

Zhotovitel dále prohlašuje, že poskytnutím licence Objednateli neporušuje práva duševního vlastnictví třetích osob a že nejsou třetí osoby, které by mohly oprávněně uplatňovat své nároky z těchto práv vůči Objednateli. V případě, že Zhotovitel nedodrží toto ustanovení, zavazuje se uhradit veškeré nároky třetích osob z důvodu porušení práv duševního vlastnictví třetích osob a dále náhradu škody způsobenou tím Objednateli.

Zhotovitel rovněž poskytuje Objednateli vlastnické právo k případným instalačním médiím, na nichž je software poskytován a k dalšímu příslušenství (návody v jakékoli podobě, hardwarové klíče apod.). Instalačním médiem se rozumí flash paměť (flash disk), magnetické diskové paměťové jednotky (přenosné disky) a médium pro ukládání dat (cd-rom,dvd-rom) (dále jen "instalační médium").

I. Předmět díla

1.1. Zhotovitel se touto Smlouvou zavazuje provést pro Objednatele řádně a včas, na svůj náklad a nebezpečí dílo specifikované v čl. II. této Smlouvy (dále jen „dílo“) a Objednatel se zavazuje za řádně provedené dílo zaplatit Zhotoviteli cenu ve výši a za podmínek sjednaných v této Smlouvě.

1.2. Zhotovitel poskytuje Objednateli práva duševního vlastnictví dle čl. XVI této Smlouvy.

1.3. Objednatel je povinen dodaný Software užívat v souladu s touto Smlouvou, v souladu s licenčními podmínkami vlastníka autorských práv k Software, a dle platných zákonných norem. Dodaný Software musí umožňovat zpřístupnění programových produktů za účelem integrace s jinými informačními systémy a to obvyklou formou komunikačního rozhraní například API, webové služby, atp. včetně potřebné dokumentace komunikačního rozhraní. Zhotovitel jako součást plnění zajistí, aby licenční ani technické podmínky možností integrace s dalšími systémy nevytvořily jakékoliv další požadavky na Objednatele.

II. Specifikace díla

2.1. Předmětem plnění je dodávka a implementace technologií pro zvýšení kybernetické bezpečnosti informačních systémů (IS) a komunikačních systémů (KS) zadavatele v souladu se standardy kybernetické bezpečnosti včetně nezbytných služeb (dále jen **„Dodávka“**).

2.2. Smluvní strany se dohodly, že dílem je provedení všech plnění dle specifikace v Zadávací dokumentaci, zejména pak v Příloze č. 1 této Smlouvy (Technická specifikace) směřujících k vytvoření díla samotného.

Předmětem díla jsou rovněž výkony či plnění, které jsou dle zavedené odborné praxe nedílnou a přirozenou součástí plnění vymezených v této Smlouvě.

Podkladem pro realizaci díla bude popis nabízeného technického řešení zpracovaný Zhotovitelem v souladu se Zadávací dokumentací, který tvoří Přílohu č. 2 této Smlouvy. Zhotovitel se zavazuje při plnění Díla postupovat v souladu s Přílohou č. 2.

Specifikace předmětu díla je obsažena zejména v Příloze č. 1 této Smlouvy (Technická specifikace).

2.3. Předmět díla bude proveden v rozsahu, způsobem a v jakosti stanovené:

(a) touto Smlouvou;

(b) technickými podmínkami, které jsou jako Příloha č. 1 součástí této Smlouvy;

(c) Návrhem Zhotovitele, který je Přílohou č. 2 této Smlouvy;

(d) písemnými pokyny Objednatele řádně podepsanými oprávněným zástupcem Objednatele;

(e) obecně závaznými právními předpisy, normami, zvyklostmi v příslušné oblasti a veškerými podklady předanými Objednatelem Zhotoviteli podle této Smlouvy a případnými pozdějšími změnami shora uvedené dokumentace, které byly vyvolány potřebami zjištěnými v průběhu provádění předmětu díla nebo okolnostmi Smluvními stranami nepředvídanými, rozhodnutími, resp. vyjádřeními veřejnoprávních orgánů s tím, že Objednatel je oprávněn upravit způsob provádění předmětu díla; veškeré požadované změny se však musí týkat následné funkčnosti předmětu díla v kontextu původních požadavků na funkčnost díla ze strany Objednatele a závazných právních předpisů.

2.4. Nepředvídanými okolnostmi se rozumí:

a) plnění svým rozsahem nebo povahou nad rámec plnění dle této Smlouvy, tj. takové plnění Zhotovitele, které nebylo součástí řešení provedení předmětu díla vyplývajícího z této Smlouvy, obecně závazných právních předpisů na provedení předmětu díla touto Smlouvou dohodnutého rozsahu a kvality či ověřené technické praxe; nebo

b) plnění vyvolané zásadní změnou dodávky předmětu díla provedené na základě zvláštního požadavku Objednatele, a to pouze a výlučně po uzavření písemného dodatku k této Smlouvě

Za nepředvídané plnění se nepovažují zejména:

(i) plnění jinak splňující podmínky této Smlouvy na nepředvídané práce, o kterých prokazatelně Zhotovitel při podpisu této Smlouvy věděl nebo měl při vynaložení odborné péče vědět; nebo

(ii) plnění, jejichž provedení bylo vyvoláno prodlením Zhotovitele s prováděním předmětu díla nebo prodlením s poskytováním s ním spojených plnění, za které Zhotovitel odpovídá; nebo

(iii) plnění, která jsou důsledkem vadného plnění Zhotovitele, dále i plnění, která jsou v souladu s řešením provedení předmětu díla, a tato pouze zpřesňují.

2.5. Zhotovitel se zavazuje naplánovat realizaci díla tak, aby minimalizoval možnost prodlení z důvodu vyšší moci.

2.6. Zhotovitel je povinen při svém plnění dodržovat a splňovat požadavky všech platných a účinných právních předpisů a technických norem, které se vztahují k předmětu této Smlouvy, a to zejména:

- občanský zákoník,
- Zákon č. 320/2001 Sb., o finanční kontrole ve veřejné správě a o změně některých zákonů (zákon o finanční kontrole), ve znění pozdějších předpisů,
- Zákon č. 262/2006 Sb., zákoník práce, ve znění pozdějších předpisů,
- Zákon č. 110/2019 Sb., o zpracování osobních údajů, ve znění pozdějších předpisů,
- Nařízení Evropského parlamentu a Rady (EU) 2016/679 ze dne 27. dubna 2016 o ochraně fyzických osob v souvislosti se zpracováním osobních údajů a o volném pohybu těchto údajů a o zrušení směrnice 95/46/ES (obecné nařízení o ochraně osobních údajů).

III. Doba a místo plnění

3.1. Smluvní strany se dohodly, že předmět díla bude proveden dle následujícího harmonogramu, a to v následujících termínech:

Poř. č.	Aktivita projektu	Termín pro dokončení aktivity
E1.1	Předimplementační analýza a zhotovení Prováděcí dokumentace	70 *kalendářních dnů ode dne nabytí účinnosti této Smlouvy*
E1.2	Předání Prováděcí dokumentace Objednateli, připomínkové řízení	70 *kalendářních dnů ode dne nabytí účinnosti této Smlouvy*
E1.3	Zapracování připomínek a předání finální verze Prováděcí dokumentace – akceptace Objednatelem	90 *kalendářních dnů ode dne nabytí účinnosti této Smlouvy*
E1.4	Dodávky a implementace	240 *kalendářních dnů ode dne nabytí účinnosti této Smlouvy*
E1.5	Školení uživatelů a administrátorů	240 *kalendářních dnů ode dne nabytí účinnosti této Smlouvy*
E1.6	Zkušební provoz	240 *kalendářních dnů ode dne nabytí účinnosti této Smlouvy*
E1.7	Akceptační testy	240 *kalendářních dnů ode dne nabytí účinnosti této Smlouvy*

Požadované zahájení plnění díla jako celku je bezprostředně po nabytí účinnosti Smlouvy.

3.2. Provedením díla se rozumí úplné dodání a dokončení díla prostého všech vad a současně řádné protokolární předání díla Objednateli formou dle čl. X. této Smlouvy.

3.3. Smluvní strany se dohodly, že předmět díla bude proveden v termínech uvedených v odst. 3.1. tohoto článku Smlouvy. Předmět díla bude proveden nejpozději do 240 dnů od účinnosti této Smlouvy.

3.4. Smluvní strany se dohodly, že celková doba úplného dokončení předmětu díla stanovená touto Smlouvou je konečná a nelze ji prodlužovat vyjma případů popsaných v této Smlouvě.

3.5. Změna etap a termínů realizace předmětu díla je možná pouze na základě schválení Objednatelem za podmínek stanovených v této Smlouvě provedeného písemnou formou v listinné podobě.

3.6. Před dobou sjednanou pro předání a převzetí řádně provedeného díla dle čl. III. odst. 3.1. Smlouvy není Objednatel povinen od Zhotovitele předmět díla či kteroukoli jeho část převzít.

3.7. Místem plnění je sídlo Objednatele a pokud to povaha plnění umožňuje a není-li to v rozporu s požadavky Objednatele, může Zhotovitel poskytovat plnění i prostřednictvím vzdáleného přístupu (off-site).

3.8. Místem předání a převzetí díla je sídlo Objednatele.

IV. Cena a způsob plnění, platební podmínky

4.1. Smluvní strany se dohodly na celkové ceně za provedení díla, ve výši **19 985 069,06 Kč (slovy: devatenáct milionů devět set osmdesát pět tisíc šedesát devět korun českých**

a šest haléřů) včetně DPH, (dále jen „cena díla“), tj. cena předmětu díla bez DPH 16 516 586,00 Kč (slovy: šestnáct milionů pět set šestnáct tisíc pět set osmdesát šest korun českých) a DPH 3 468 483, 06 Kč (slovy: tři miliony čtyři sta šedesát osm tisíc čtyři sta osmdesát tři korun českých a šest haléřů). Uvedená cena bez DPH je cenou pevnou a nejvýše přípustnou po celou dobu trvání Smlouvy. V případě změny legislativy bude účtována DPH podle platných předpisů. Podrobný rozpis položek dodávky a implementace je uveden jako příloha č. 4 této Smlouvy.

4.2. V ceně díla jsou zahrnuty veškeré náklady Zhotovitele, které při plnění svého závazku dle této Smlouvy vynaloží. Zhotovitel prohlašuje, že všechny technické, finanční, věcné a ostatní podmínky díla zahrnul do kalkulace ceny díla. Zhotovitel výslovně prohlašuje, že součástí ceny díla jsou i veškeré náklady spojené se splněním podmínek pro realizaci předmětu díla dle obecně závazných právních předpisů a přebírá na sebe nebezpečí změny okolností dle § 1765 odst. 2 občanského zákoníku.

4.3. Objednatel uhradí cenu díla v souladu s odst. 4.5 Požadavků na provedení akceptačních testů a přechod do zkušebního provozu dle přílohy č. 1 této Smlouvy následovně:

50 % ceny komodity dle odst. 4.1 této Smlouvy (dokončena etapa E1.4 Dodávka a implementace)

30 % ceny komodity dle odst. 4.1 této Smlouvy (dokončena etapa E1.6 Zkušební provoz).

20 % celkové ceny za provedení díla dle odst. 4.1 této Smlouvy po dokončení a akceptaci celého díla (dokončení etapy E1.7 Akceptační testy).

4.4. Cena, resp. ceny dle předchozího odstavce budou uhrazeny na základě Zhotovitelem vystaveného daňového dokladu – faktury.

Faktura bude vystavena se splatností 30 kalendářních dní ode dne doručení Objednateli. Smluvní strany se vzájemně dohodly, že DPH bude Zhotovitelem účtována v sazbách dle právních předpisů platných v době uskutečnitelného zdanitelného plnění pro to které účtované dílčí plnění dle předchozího odstavce.

Každá faktura vystavená Zhotovitelem dle této Smlouvy musí obsahovat pojmové náležitosti daňového dokladu stanovené zákonem č. 235/2004 Sb., o dani z přidané hodnoty, ve znění pozdějších předpisů a zákonem č. 563/1991 Sb., o účetnictví, ve znění pozdějších předpisů, a dále následující údaje:

- název a registrační číslo projektu dle Preambule této Smlouvy,
- číslo Smlouvy,
- identifikaci Objednatele podle Smlouvy,
- identifikaci Zhotovitele podle Smlouvy,
- označení banky a číslo účtu, na který má být platba zaplacena, včetně konstantního a variabilního symbolu,
- den splatnosti a den uskutečnění zdanitelného plnění,
- název a popis poskytnutých Dodávek s odkazem na Smlouvu,
- účtovanou částku bez DPH,

- vyčíslenou částku DPH,
- celkovou částku včetně DPH,
- jakékoliv další údaje vyžadované pro účetní a daňový doklad příslušnými Právními předpisy.

Přílohou každé faktury dále bude akceptační protokol v souladu s čl. X. této Smlouvy.

V případě, že daňový doklad nebude obsahovat uvedené údaje či bude neúplný či nebude mít všechny přílohy, není Objednatel povinen na jeho základě plnit a nedostává se do prodlení. Zhotovitel je povinen takový daňový doklad opravit, aby splňoval podmínky stanovené touto Smlouvou. Lhůta splatnosti běží znovu od doručení nové nebo opravené faktury.

Objednatelem podepsaný předávací protokol nezbavuje Zhotovitele odpovědnosti za řádné provedení předmětu díla jako celku bez vad a nedodělků.

4.5. Strany se dohodly, že Objednatel je oprávněn požadovat po Zhotoviteli bližší vysvětlení, objasnění nebo zdůvodnění částek obsažených ve fakturách, a to na základě písemné výzvy adresované Zhotoviteli. Od okamžiku odeslání písemné výzvy k objasnění do prokázání oprávněnosti požadovaných plateb se lhůta splatnosti faktury prodlužuje.

4.6. Objednatel je oprávněn ponížit Zhotovitelem fakturovanou úhradu ceny o jakékoliv případné smluvní pokuty, náhrady škod a další platby splatné ve prospěch Objednatele vůči Zhotoviteli. Pouze Objednatel je oprávněn započíst jakékoliv své splatné pohledávky dle Smlouvy vůči pohledávkám Zhotovitele. Zhotovitel dává k zápočtu vzájemných pohledávek výslovný souhlas.

4.7. Pokud Zhotovitel poruší povinnosti ze Smlouvy podstatným způsobem, je Objednatel oprávněn pozastavit jakoukoliv platbu na základě faktury až do odstranění prodlení nebo porušení povinnosti Zhotovitele.

4.8. Veškeré vícepráce, změny, doplňky nebo rozšíření, které nejsou součástí předmětu díla dle Smlouvy, musí být vždy před jejich realizací písemně odsouhlaseny Objednatelem včetně jejich ocenění ve formě písemného dodatku (v listinné podobě) k této Smlouvě. Pokud Zhotovitel provede některé z těchto prací bez předchozího písemného odsouhlasení Objednatelem, má Objednatel právo odmítnout jejich úhradu a cena za jejich provedení je součástí ceny za provedení předmětu díla.

4.9. Úhrada ceny za provedení předmětu díla, ať již jako celku či dílčích plnění, nemá vliv na možnost uplatnění práva Objednatele z vad předmětu díla.

V. Součinnost Smluvních stran

5.1. Smluvní strany se zavazují vyvinout veškeré úsilí k vytvoření potřebných podmínek pro realizaci díla dle podmínek stanovených touto Smlouvou, které vyplývají z jejich smluvního postavení. To platí i v případech, kde to není výslovně stanoveno ustanovením této Smlouvy.

5.2. Pokud jsou kterékoli ze Smluvních stran známy skutečnosti, které jí budou bránit, aby dostála svým smluvním povinnostem, sdělí tuto skutečnost neprodleně písemně druhé

smluvní straně. Smluvní strany se dále zavazují neprodleně odstranit v rámci svých možností všechny okolnosti, bránící z její strany splnění jejich smluvních povinností.

5.3. Zhotovitel se zavazuje, že na základě skutečností zjištěných v průběhu plnění povinností dle této Smlouvy navrhne a provede opatření směřující k dodržení podmínek stanovených touto Smlouvou pro naplnění Smlouvy, k ochraně Objednatele před škodami, ztrátami a zbytečnými výdaji a že poskytne Objednateli, zástupci Objednatele jednajícímu ve věcech technických a jiným osobám zúčastněným na provádění díla veškeré potřebné doklady, konzultace, pomoc a jinou součinnost.

5.4 Zhotovitel je podle ustanovení § 2 písm. e) zákona č. 320/2001 Sb., o finanční kontrole ve veřejné správě a o změně některých zákonů, ve znění pozdějších předpisů, osobou povinou spolupůsobit při výkonu finanční kontroly prováděné v souvislosti s úhradou zboží nebo služeb z veřejných výdajů.

5.5 Zhotovitel je povinen uchovávat veškerou dokumentaci související s realizací projektu včetně účetních dokladů minimálně do konce roku 2034. Pokud je v českých právních předpisech stanovena lhůta delší, musí zhotovitel uvedenou dokumentaci uchovávat po dobu trvání takové lhůty. Zhotovitel je povinen minimálně do konce roku 2034 poskytovat požadované informace a dokumentaci související s realizací projektu zaměstnancům nebo zmocněncům pověřených orgánů (CRR, MMR ČR, MF ČR, Evropské komise, Evropského účetního dvora, Nejvyššího kontrolního úřadu, příslušného orgánu finanční správy a dalších oprávněných orgánů státní správy) a je povinen vytvořit výše uvedeným osobám podmínky k provedení kontroly vztahující se k realizaci projektu a poskytnout jim při provádění kontroly součinnost.

VI. Prohlášení, práva a závazky smluvních stran

6.1. Zhotovitel prohlašuje, že ke dni podpisu Smlouvy:

(a) není jako právnická osoba v likvidaci;

(b) není proti němu vedeno konkursní řízení ani vyrovnací řízení ve smyslu zákona č. 182/2006 Sb., o úpadku a způsobech jeho řešení (insolvenční zákon), ve znění pozdějších předpisů (dále jen „insolvenční zákon) a takové řízení nebylo zastaveno či zrušeno z důvodu nedostatku majetku Zhotovitele a dále není předlužen či neschopen plnit své splatné závazky vůči svým věřitelům;

(c) uzavření/m této Smlouvy:

- neporuší správní rozhodnutí orgánu státní správy České republiky;
- neporuší ustanovení žádné dohody, smlouvy či jiného ujednání, které uzavřel se třetí osobou;
- nebude mít za následek újmu nebo požadavek na splacení jakéhokoli správního poplatku, dotací nebo jiného závazku Zhotovitele;

(d) neučinil nic, ať již sám anebo za spolupráce či prostřednictvím třetí osoby, co by omezilo či znemožnilo dosažení účelu této Smlouvy, a to v České republice, a pokud dodavatel má sídlo v jiné zemi, než v České republice, tak i ve vztahu k zemi, v níž má sídlo.

(e) není osobou, na kterou by se vztahovaly sankce související s nařízením Rady (EU) 2022/576 ze dne 8. dubna 2022 kterým se mění nařízení (EU) č. 833/2014 o omezujících opatřeních vzhledem k činnostem Ruska destabilizujícím situaci na Ukrajině a že splňuje požadavky uvedené v článku 5k odst. 1 písm. a) – c) výše uvedeného Nařízení.

6.2. Zhotovitel se zavazuje, že Objednateli bezodkladně po vzniku takové skutečnosti písemně oznámí:

(a) podání návrhu na prohlášení konkursu na majetek Zhotovitele dle insolvenčního zákona; nebo

(b) podání návrhu na vyrovnání na majetek Zhotovitele dle insolvenčního zákona; nebo

(c) vstup Zhotovitele do likvidace; nebo

(d) splnění podmínek prohlášení konkursu na majetek Zhotovitele, tj. zejména že Zhotovitel je předlužen anebo insolventní; nebo

(e) rozhodnutí o provedení přeměny Zhotovitele, zejména fúzí, převodem jmění na společníka či rozdělením, provedení změny právní formy Zhotovitele či provedení jiných organizačních změn; nebo

(f) omezení či ukončení činnosti Zhotovitele, která bezprostředně souvisí s předmětem této Smlouvy; nebo

(g) všechny skutečnosti, které by mohly mít vliv na přechod či vypořádání závazků Zhotovitele vůči Objednateli vyplývajících z této Smlouvy či s touto Smlouvou souvisejících; nebo

(h) rozhodnutí o zrušení Zhotovitele.

6.3. Zhotovitel prohlašuje, že

(a) je odborně způsobilý ke splnění všech svých závazků podle této Smlouvy, a to s ohledem na předmět plnění, se kterým se náležitě seznámil, a že

(b) před podpisem této Smlouvy se řádně seznámil a překontroloval předané materiální podklady a dokumentaci a řádně prověřil místní podmínky a všechny nejasné podmínky pro realizaci díla či jeho části si vyjasnil s Objednatelem nebo místním šetřením,

(c) Smlouva byla Zhotovitelem řádně schválena a podepsána a zakládá platný závazek Zhotovitele, vynutitelný vůči němu v souladu s podmínkami v ní uvedenými,

(d) podpisem ani plněním Smlouvy Zhotovitel neporušuje žádné ustanovení svých zakladatelských dokumentů ani žádnou jinou smlouvu nebo ujednání, jehož je Zhotovitel stranou, nebo kterým je Zhotovitel nebo jeho majetek vázán, ani žádný zákon či jiný právní předpis nebo rozhodnutí státního orgánu,

(e) podle nejlepšího vědomí Zhotovitele proti němu neprobíhá žádné soudní, rozhodčí ani správní řízení, které by mohlo negativně ovlivnit platnost, účinnosti

nebo vymahatelnost Smlouvy nebo plnění jakýchkoliv povinností Zhotovitele podle této Smlouvy, ani nehrozí zahájení žádného takového řízení.

6.4. Zhotovitel se zavazuje:

(a) při provádění předmětu díla postupovat s odbornou péčí a dodržovat Právní předpisy a rozhodnutí orgánů veřejné správy,

(b) udržovat a obnovovat po celou dobu účinnosti této Smlouvy veškeré nezbytné souhlasy, povolení, oprávnění či licence potřebné k řádnému poskytování Dodávek v souladu s Právními předpisy, přičemž Zhotovitel odškodní Objednatele v případě, že tak Zhotovitel opomněl nebo opomene kdykoliv v průběhu trvání Smlouvy učinit.

6.5. Objednatel je oprávněn postoupit jakákoliv práva a povinnosti z této Smlouvy na kteroukoliv třetí osobu, s čímž Zhotovitel podpisem Smlouvy vyslovuje svůj souhlas.

6.6. Zhotovitel se zavazuje uhradit Objednateli do deseti dnů poté, kdy k tomu bude Objednatelem písemně vyzván, veškeré pokuty či další sankce, které byly Objednateli vyměřeny (pravomocným rozhodnutím) státními orgány v souvislosti s porušením povinností Zhotovitele stanovených touto Smlouvou či obecně závaznými právními předpisy při provádění předmětu díla nebo uspokojit veškeré nároky třetích osob, o kterých bude pravomocně rozhodnuto příslušným orgánem veřejné moci, jež vznikly v souvislosti s porušením povinností Zhotovitele stanovených touto Smlouvou či obecně závaznými právními předpisy při provádění předmětu díla. Úhrada bude provedena na účet Objednatele uvedený v záhlaví této Smlouvy.

6.7. Objednatel neudělil Zhotoviteli žádné oprávnění uzavírat pracovně právní či jiné vztahy jménem Objednatele nebo jednat jménem Objednatele.

6.8. Zhotovitel se zavazuje, že pokud pro plnění díla použije třetí osoby v jiném než pracovněprávním vztahu, tak s takovými osobami ošetří veškeré vztahy a zejména autorská práva tak, aby tyto třetí osoby nemohly vznášet jakékoli nároky vůči Objednateli. Zhotovitel je povinen na základě výzvy Objednatele předložit seznam osob, které se na plnění díla podíleli spolu se specifikací právního vztahu, na základě, kterého tak činily, a současně prokázat splnění povinnosti podle předchozí věty. V případě, že Zhotovitel tuto skutečnost na základě písemné výzvy Objednatele do 30 dní od doručení výzvy Zhotoviteli nedoloží, zavazuje se zaplatit Objednateli smluvní pokutu ve výši 1 000 Kč za každý i započatý den prodlení.

6.9. Objednatel prohlašuje, že podpisem ani plněním Smlouvy Objednatel neporušuje žádné ustanovení svých zakladatelských dokumentů ani žádnou jinou smlouvu nebo ujednání, jehož je Objednatel stranou, nebo kterým je Objednatel nebo jeho majetek vázán, ani žádný zákon či jiný právní předpis nebo rozhodnutí státního orgánu.

VII. Předání díla, přechod vlastnictví a nebezpečí škody

7.1. Zhotovitel se zavazuje realizovat předmět díla průběžně a předat výstupy z plnění předmětu díla v termínech definovaných harmonogramem v čl. III. odst. 3.1 této Smlouvy.

7.2. O předání a převzetí bude sepsán písemný protokol dle čl. X. této Smlouvy podepsaný oprávněnými osobami obou Smluvních stran. V rámci předání díla nebo jednotlivé etapy je Zhotovitel povinen předvést, že dílo je způsobilé sloužit svému účelu.

7.3. Písemný protokol dle čl. X. této Smlouvy bude sepsán po ukončení aktivity E1.7 – Akceptační testy.

7.4. Veškerá dokumentace bude Zhotovitelem Objednateli předávána v originálech, a to jak ve formě listinných dokumentů, tak v elektronické editovatelné podobě.

7.5. Objednatel nabývá právo užívat předmět plnění a přechází na něj nebezpečí škody k předmětu plnění okamžikem jejich předání Zhotovitelem, resp. převzetím na základě písemného protokolu podepsaného Objednatelem i Zhotovitelem.

7.6. Po datu předání odpovídá Zhotovitel za ztrátu nebo škodu pouze v případě, že je tato ztráta nebo škoda způsobena zaviněním Zhotovitele nebo když má Zhotovitel předměty Dodávek v držení z důvodu poskytování záručního servisu.

VIII. Podmínky provádění předmětu díla

8.1. Zhotovitel se zavazuje:

(a) zajistit provádění předmětu díla tak, aby provádění předmětu díla v co nejmenší míře omezovalo činnost Objednatele;

(b) zajistit provádění předmětu díla tak, aby provádění předmětu díla bylo prováděno pod odborným dozorem Zhotovitele, který bude garantovat dodržování postupů nabídnutých Zhotovitelem v Nabídce nebo postupů dohodnutých s Objednatelem v průběhu plnění; totéž platí pro práce Poddodavatelů;

(c) neprodleně, nejpozději však do tří (3) dnů, písemně oznámit Objednateli veškeré skutečnosti a okolnosti, které při poskytování dodávek zjistil nebo se o nich dozvěděl a které mohou mít vliv na poskytování dodávek;

(d) vyvstane-li v průběhu provádění předmětu díla nutnost upřesnění způsobu jeho provedení, neprodleně si vyžádat předchozí písemný souhlas či pokyn Objednatele;

(e) písemně upozornit Objednatele na nevhodnost, případně nepřípustnost podkladových materiálů, pokynů a věcí, které mu byly předány Objednatelem nebo Objednatelem požadovaných změn, ať již z hlediska důsledků pro jakost a provedení předmět díla či rozporu s podklady pro uzavření této Smlouvy, ustanoveními nebo rozhodnutími orgánů veřejné správy či obecně závaznými právními předpisy či jinými normami, a to bezodkladně poté, co tuto skutečnost zjistí či mohl zjistit. V případě, že Objednatel bude, i přes upozornění Zhotovitele, písemně trvat na užití podkladových materiálů, pokynů a věcí, které byly Zhotoviteli předány Objednatelem, je Zhotovitel oprávněn odmítnout jejich plnění pouze tehdy, pokud by se jejich splněním mohl vystavit správnímu či trestnímu postihu;

(f) vždy předkládat návrhy veškerých písemných podkladů a dokumentů souvisejících s poskytováním dodávek, nestanovuje-li Zadávací dokumentace či dohoda Stran jinak.

8.2. Zhotovitel bude svým jménem projednávat a hradit náklady vyplývající z projednaných záležitostí přímo souvisejících s jeho činností při realizaci předmětu díla a dokončení předmětu díla, které jsou v jeho kompetenci a za které plně odpovídá.

Zhotovitel na sebe přejímá zodpovědnost za škody způsobené všemi osobami zúčastněnými na provádění předmětu díla na straně Zhotovitele, stejně tak za škody způsobené svou činností Objednateli nebo třetím osobám.

Zhotovitel prohlašuje, že zejména ve smyslu § 2914 občanského zákoníku odpovídá za veškeré škody, které způsobí osoby, které při plnění předmětu této Smlouvy použije. Pro případ, že by Zhotovitel porušil svoji povinnost podle této Smlouvy, a pro plnění předmětu této Smlouvy použil pomocníka, jenž by jednal samostatně, přebírá ručení za jakoukoli škodu, kterou by tento pomocník způsobil.

Zhotovitel není oprávněn postoupit jakákoliv práva, povinnosti, závazky a pohledávky z této Smlouvy na třetí osoby bez předchozího písemného souhlasu Objednatele.

Zhotovitel je povinen:

a) zajistit a financovat veškeré poddodavatelské práce a nese za ně záruku vůči Objednateli v plném rozsahu dle této Smlouvy,

b) zajistit, aby všichni poddodavatelé měli platná příslušná oprávnění, koncese, certifikace, licence a rovněž odbornou kvalifikaci a dostatek odborných zkušeností, jež jsou nezbytné pro poskytování příslušných částí dodávek dle jejich smluv se Zhotovitelem,

c) jednat s poddodavateli v souladu se zásadami poctivého obchodního styku tzn. zejména uhradit poddodavatelům sjednanou cenu za řádné a včasné poskytnutí příslušných částí dodávek,

d) zajistit, aby příslušné plnění prováděly výhradně osoby tvořící realizační tým Zhotovitele uvedené v příloze č. 3 této Smlouvy; doplnění, nebo změny členů realizačního týmu Zhotovitele budou možné pouze po předchozím schválení Objednatelem, přičemž jednotliví členové realizačního týmu na stanovených pozicích uvedené pro splnění kvalifikace mohou být nahrazeny pouze osobami splňujícími pro nahrazovanou pozici kvalifikační požadavky dle zadávací dokumentace k Veřejné zakázce,

e) dodržet požadavek týkající se identifikace poddodavatelů, kteří se mají zapojit do realizace předmětu plnění, a to před zahájením plněním služeb ze strany těchto poddodavatelů.

8.3. V případě zjištění závad či nedostatků v průběhu provádění díla musí být o těchto zjištěných skutečnostech sepsán zápis a stanoveny termíny jejich odstranění. Nedohodnou-li se smluvní strany na lhůtě k odstranění vad, musí být vady odstraněny nejpozději do 14 dní od oznámení vady Zhotoviteli, nejpozději však do předání díla Objednateli.

8.4. Objednatel je oprávněn:

(a) sám či prostřednictvím třetí osoby vykonávat v místě provádění předmětu díla dozor Zhotovitele a v jeho průběhu zejména sledovat, zda jsou práce prováděny podle Smlouvy a právních předpisů;

(b) pokud Zhotovitel nesplní jakoukoliv povinnost podle této Smlouvy a nesplní ji ani v dodatečné lhůtě stanovené touto Smlouvou, jinak v dodatečně přiměřené lhůtě stanovené Objednatelem, jež však nebude delší než třicet dnů, je Objednatel, aniž by tím byla dotčena jakákoliv jiná práva a nároky Objednatele dle této Smlouvy, oprávněn, nikoliv však povinen, podle svého uvážení splnit povinnost Zhotovitele nebo pověřit splněním této povinnosti jiné osoby na náklady Zhotovitele,

(c) po Zhotoviteli požadovat, aby pro splnění Veřejné zakázky nevyužíval člena týmu Zhotovitele, který prokazatelně:

- plní své povinnosti nekompetentně nebo nedbale, nebo

- neplní nebo porušuje některá ustanovení této Smlouvy nebo právních předpisů,

přičemž takový člen týmu Zhotovitele musí být po výzvě Objednatele bez zbytečného odkladu nahrazen jiným členem s odpovídající kvalifikací.

8.5. Objednatel si jako zadavatel v Zadávací dokumentaci Veřejné zakázky stanovil požadavky na vzdělání a odbornou kvalifikaci ve vztahu k fyzickým osobám, které se mají na plnění předmětu Smlouvy podílet (příloha č. 3 této Smlouvy). Požadavek na změnu člena realizačního týmu je přípustný pouze ve výjimečném a řádně zdůvodněném případě. Tato změna je možná pouze se souhlasem Objednatele, přičemž fyzické osoby, které se budou na plnění předmětu Smlouvy nově podílet, musí splňovat stejné minimálně kvalifikační předpoklady kladené v Zadávací dokumentaci na pozici nahrazované fyzické osoby. Objednatel nebude udělení souhlasu bezdůvodně odpírat.

8.6. Zhotovitel výslovně prohlašuje, že se Objednatel provedením přímé platby poddodavateli nemůže dostat do prodlení s platbou cen, neboť provedením přímé platby poddodavateli závazek Objednatele v rozsahu částky fakturované poddodavatelem Zhotovitele jeho splněním zaniká. Objednatel není povinen provést žádnou platbu poddodavatelům dle tohoto článku, jedná se výhradně o možnost Objednatele, nikoli povinnost ze strany poddodavatelů Zhotovitele jakkoli vymahatelnou.

8.7. Zhotovitel se zavazuje po celou dobu trvání smluvního vztahu založeného touto Smlouvou sjednat a dodržovat srovnatelné smluvní podmínky v oblasti rozdělení rizika a smluvních pokut se svými poddodavateli s ohledem na charakter, rozsah a cenu plnění poddodavatele, jako jsou sjednané v této Smlouvě.

8.8. Zhotovitel se zavazuje po celou dobu trvání smluvního vztahu založeného touto Smlouvou zajistit dodržování veškerých právních předpisů, zejména pak pracovněprávních (odměňování, pracovní doba, doba odpočinku mezi směnami, placené přesčasy), dále předpisů týkajících se oblasti zaměstnanosti a BOZP, tj. zejména zákona č. 435/2004 Sb., o zaměstnanosti, ve znění pozdějších předpisů,

a zákona č. 262/2006 Sb., zákoníku práce, ve znění pozdějších předpisů, příp. dle odpovídajících předpisů právního řádu jiného státu, a to vůči všem osobám, které se na plnění Smlouvy podílejí (a bez ohledu na to, zda budou činnosti prováděny Zhotovitelem či jeho poddodavateli). Zhotovitel se také zavazuje zajistit, že všechny osoby, které se na plnění Smlouvy podílejí (bez ohledu na to, zda budou činnosti prováděny Zhotovitelem či jeho poddodavateli), jsou vedeny v příslušných registrech, jako například v registru pojištěnců apod.

IX. Záruka za jakost

9.1. Zhotovitel se zavazuje, že předaný předmět díla bude prostý jakýchkoli vad a bude mít vlastnosti dle obecně závazných právních předpisů, této Smlouvy a Zadávací dokumentace a bude proveden v souladu s ověřenou technickou praxí. Zhotovitel poskytuje Objednateli záruku za jakost díla, a to ve struktuře a délce dle požadavků Zadávacích podmínek, zejména Přílohy č. 1 této Smlouvy (Technická specifikace) a v délce 24 měsíců v případech v Zadávacích podmínkách neuvedených.

9.2. Zhotovitelem bude Objednateli poskytován bezplatný záruční servis na Objednatelem reklamované vady předmětu díla vzniklé v době trvání záruční doby určené v čl. IX. odst. 9.1. této Smlouvy.

9.3. Objednatel je oprávněn reklamovat v záruční době dle čl. IX. odst. 9.1. této Smlouvy vady předmětu díla u Zhotovitele, a to písemnou formou. V reklamaci musí být popsána vada předmětu díla, určen nárok Objednatele z vady předmětu díla, případně požadavek na způsob odstranění vad, a to včetně termínu pro odstranění vad Zhotovitelem. Objednatel má právo volby způsobu odstranění důsledku vadného plnění.

9.4. Zhotovitel se zavazuje zahájit odstraňování vady předmětu díla či jeho části i tehdy, neuznává-li svou odpovědnost za vady či příčiny, které ji vyvolaly, a vady odstranit v technicky co nejkratší lhůtě, a současně zahájit reklamační řízení. O reklamačním řízení budou Objednatelem pořizovány písemné zápisy ve dvojím vyhotovení, z nichž jeden stejnopis obdrží každá ze Smluvních stran. Reklamační řízení musí být ukončeno do čtyřiceti osmi pracovních hodin po jeho zahájení. Bude-li v reklamačním řízení vada uznána jako reklamační vada, bude odstranění vady předmětu díla či jeho části provedeno bezúplatně. Nebude-li v reklamačním řízení vada uznána jako reklamační vada, bude odstranění vady předmětu díla či jeho části provedeno úplatně.

9.5. Práva a povinnosti ze Zhotovitelem poskytnuté záruky vznikají okamžikem provedení a předání Objednateli té části díla, ke které se poskytnuté záruky vztahují, a nezanikají ani odstoupením kterékoli ze Smluvních stran od Smlouvy.

X. Protokol o předání a převzetí předmětu díla

10.1. Zhotovitel se zavazuje předmět díla, resp. jeho části odpovídající samostatně fakturovaným částem podle čl. IV. odst. 4.3 řádně provést a protokolárně předat Objednateli. O předání předmětu díla Zhotovitelem Objednateli bude sepsán písemný protokol. Dílo a jeho jednotlivé části se považují za řádně předané, pokud je plněno řádně, včas a bez vad a tato skutečnost je vyznačena v předávacím protokolu.

Osobou oprávněnou k podpisu předávacího protokolu za Objednatele je XXX,

Osobou oprávněnou k podpisu předávacího protokolu za Zhotovitele je XXX.

10.2. Nejpozději na poslední den provedení předmětu díla, resp. jeho části, Zhotovitel svolá do místa předání a převzetí díla předávací řízení. Na předávací řízení přizve Zhotovitel Objednatele, a to písemným oznámením, které musí být doručeno Objednateli alespoň pět pracovních dnů předem včetně návrhu předávacího protokolu a případně další příslušné dokumentace. Objednatel do dvou pracovních dnů buď potvrdí navržený termín, nebo požádá o stanovení nového termínu; posunutím předávacího řízení na žádost Objednatele se Zhotovitel nemůže dostat do prodlení s plněním předmětu této Smlouvy. V případě, že nebude Objednateli řádně a včas doručena výzva k účasti na předávacím řízení, může dojít k předávacímu řízení nejdříve po uplynutí pátého pracovního dne ode dne doručení písemné výzvy k zahájení předávacího řízení.

10.3. K předání předmětu díla, resp. jeho části Zhotovitelem Objednateli dojde na základě předávacího řízení, a to formou písemného předávacího protokolu (jehož součástí budou příslušné výstupy i příslušná dokumentace, pokud je to stanoveno touto Smlouvou či obvyklé), který bude podepsán oprávněnými zástupci obou Smluvních stran.

Předávací protokol musí obsahovat alespoň předmět a charakteristiku předmětu díla, resp. jeho části, soupis zjištěných vad předmětu díla stanovených Zhotovitelem či Objednatelem, vyjádření Zhotovitele k vadám předmětu díla vytčeným Objednatelem, lhůty pro odstranění vad předmětu díla, zhodnocení jakosti předmětu díla a jeho částí, dohodu o lhůtách a opatřeních k odstranění vad či jeho části, záznam o nutných dodatečně požadovaných pracích, případnou dohodu o slevě z ceny za provedení předmětu díla, stanovisko Objednatele, zda předmět díla přejímá či nikoli a soupis příloh. Nedohodnou-li se smluvní strany na lhůtě k odstranění vad, musí být vady odstraněny nejpozději do 14 dní od oznámení vady Zhotoviteli. Předávací protokol bude vyhotoven ve třech stejnopisech, z nichž jeden obdrží Zhotovitel a dva Objednatel. Každý stejnopis bude podepsán oběma Stranami a má právní sílu originálu.

10.4. V případě, že je Objednatelem přebírán dokončený předmět díla, skutečnost, že předmět díla je dokončen co do množství, jakosti, kompletnosti a schopnosti trvalého užívání, prokazuje zásadně Zhotovitel a za tím účelem předkládá nezbytné písemné doklady Objednateli. V případě, že nedojde k předložení a předání Objednateli shora uvedených dokladů nejpozději při předávacím řízení, nepovažuje se předmět díla za řádně předaný.

10.5. V případě, že se při přejímání předmětu díla Objednatelem ukáže, že je Zhotovitelem předáván předmět díla, který nese vady nebo není schopen sloužit svému účelu, není Objednatel povinen předávaný předmět díla převzít. Tato skutečnost bude uvedena v předávacím protokole. Po odstranění vad předmětu díla či jeho části, pro které Objednatel odmítl od Zhotovitele předmět díla převzít, se opakuje předávací řízení analogicky dle tohoto článku Smlouvy. V takovém případě bude k původnímu předávacímu protokolu sepsán dodatek, ve kterém bude uvedeno převzetí předmětu díla. Dodatek obsahuje veškeré náležitosti stanovené pro předávací protokol v tomto článku Smlouvy.

10.6. Objednatel má právo ve lhůtě 10 dnů od předání díla či části díla vznést výhrady nebo připomínky k předávanému dílu či části díla; v takovém případě se Strany zavazují zahájit společné jednání za účelem odstranění veškerých vzájemných rozporů a nalezení shody nad předávaným dílem či částí díla, a to nejpozději do pěti (5) pracovních dnů od výzvy kterékoliv Strany. V takovém případě se dílo či část díla nepovažují za převzaté se všemi důsledky z toho vyplývajícími a předávací řízení se v takovém případě zopakuje analogicky dle tohoto článku Smlouvy.

10.7. Zhotovitel je povinen bez zbytečného odkladu odstranit vady předmětu díla, i když se za ně necítí odpovědný. Náklady na odstranění těchto vad nese Zhotovitel, a to až do účinnosti dohody Smluvních stran o jejich úhradě nebo do právní moci rozhodnutí příslušného soudu ve věci úhrady těchto nákladů.

XI. Smluvní pokuty a úrok z prodlení, odpovědnost za škodu

11.1. Smluvní strany se dohodly na tom, že v případě porušení ustanovení čl. III. odst. 3.1. a v případě porušení čl. IX. odst. 9.4. Smlouvy Zhotovitelem je Zhotovitel povinen uhradit Objednateli smluvní pokutu ve výši 0,1 % (slovy: jedna desetina procenta) z celkové ceny v Kč bez DPH, kterou je Objednatel podle této Smlouvy povinen Zhotoviteli uhradit, a to za každý započatý den prodlení.

11.2 Smluvní strany se dohodly, že v případě nepravdivosti prohlášení v ustanovení čl. VI. odst. 6.1. nebo porušení některé povinnosti sjednané v čl. V. této Smlouvy Zhotovitelem je Zhotovitel povinen uhradit Objednateli smluvní pokutu ve výši 0,1 % (slovy: jedna desetina procenta) z celkové ceny ujednané v čl. IV, a to za každé porušení Smlouvy zvlášť.

11.3. V případě, kdy nastane některá ze situací uvedených v čl. XVII. odst. 17.6. je Zhotovitel povinen zaplatit smluvní pokutu ve výši 100 000 Kč (slovy: stotisíc korun českých), a to za každý jednotlivý případ. Úhradou smluvní pokuty není dotčen nárok Objednatele na náhradu škody v plné výši.

11.4. Smluvní strany se dohodly na tom, že v případě prodlení s úhradou odměny dle ustanovení čl. IV této Smlouvy je Objednatel povinen uhradit Zhotoviteli úrok z prodlení ve 0,1 % (slovy: jedna desetina procenta) z nezaplacené částky v Kč bez DPH za každý den prodlení.

11.5. Smluvní pokuta je splatná do 21 dní ode dne, kdy byla povinné straně doručena písemná výzva k jejímu zaplacení ze strany oprávněné strany, a to na účet oprávněné strany uvedený v písemné výzvě. Ustanovením o smluvní pokutě není dotčeno právo oprávněné strany na náhradu škody v plné výši s tím, že zaplacená smluvní pokuta se na úhradu škody nezapočítává. Případným odstoupení od Smlouvy nárok na úhradu smluvní pokuty nezaniká.

11.6. V případě, že porušením povinnosti Zhotovitele podle této Smlouvy vznikne Objednateli škoda, jejímž důsledkem bude neproplacení dotace nebo její části poskytovatelem dotačního titulu, odpovídá Zhotovitel Objednateli za škodu až do výše finančního postihu ze strany poskytovatele dotačního titulu uplatněného vůči Objednateli a Zhotovitel se zavazuje tuto škodu Objednateli nahradit, a to na písemnou výzvu Objednatele se splatností 21 dní ode dne doručení výzvy Zhotoviteli. Případným

odstoupení od Smlouvy nárok na odškodnění dle tohoto odstavce nezaniká. Zhotovitel dává Objednateli výslovný souhlas k případnému zápočtu vzájemných pohledávek.

11.7 Objednatel prohlašuje, že neprovedení díla ve smyslu odst. 3.2 Smlouvy v termínu do 31. 12. 2025 může způsobit neposkytnutí finančních prostředků od dotačního orgánu. V případě neposkytnutí těchto prostředků z důvodu nesplnění termínu dokončení díla vinou Zhotovitele bude výše těchto prostředků považována za škodu vzniklou Objednateli, kterou se Zhotovitel zavazuje uhradit Objednateli. V případě neposkytnutí finančních prostředků od dotačního orgánu z důvodu nedodržení podmínek způsobené činností nebo nečinností Zhotovitele bude výše těchto prostředků považována za škodu vzniklou Objednateli, kterou se Zhotovitel zavazuje uhradit Objednateli.

11.8. Pro případ porušení povinností podle odst. 14.1. této Smlouvy si smluvní strany sjednávají smluvní pokutu ve výši 1.000 Kč za každé jednotlivé porušení. Smluvní strany dále ujednaly, že vůči sobě neuplatní právo namítat nepřiměřenost výše smluvní pokuty dle této Smlouvy u soudu ve smyslu § 2051 občanského zákoníku.

XII. Ukončení Smlouvy

12.1. Smluvní strany se dohodly, že tuto Smlouvu mohou ukončit pouze za podmínek dále upravených v této Smlouvě a nebo v případech, které stanoví zákon.

12.2 Odstoupení od Smlouvy musí být provedeno písemnou formou a je účinné okamžikem jeho doručení druhé straně smluvního vztahu. Odstoupením od Smlouvy se tato Smlouva od okamžiku doručení projevu vůle směřujícího k odstoupení od Smlouvy druhé smluvní straně ruší; tím není dotčeno právo Objednatele dle odst. 12.3. této Smlouvy.

12.3 Dojde-li k odstoupení od Smlouvy ze strany Objednatele, má Objednatel právo společně s projevem vůle směřujícího k odstoupení od Smlouvy sdělit Zhotoviteli, že odstoupením se Smlouva ruší od počátku nebo že odstoupení má účinky pouze do budoucna, a tudíž si Objednatel ponechává práva k částem díla, jež byla zhotovena před zánikem Smlouvy. V případě odstoupení s účinky pouze do budoucna Zhotoviteli náleží poměrná část odměny určená podle poměru části díla již zhotoveného vůči dílu jako celku. Zhotovitel v takovém případě předá Objednateli již zhotovenou část díla bez zbytečného odkladu, nejpozději do 14 dní od zániku Smlouvy. Ustanovení této Smlouvy o předání díla, záruce za jakost, prohlášení smluvních stran, ustanovení o právní odpovědnosti vč. závazku odškodnění, ustanovení o právech duševního vlastnictví, vč. předání související dokumentace a zdrojových kódů, se pro již zhotovenou část díla použijí obdobně. Ke Smlouvě o zabezpečení podpory provozu, kterou uzavírají Smluvní strany v souvislosti s touto Smlouvou, sdělí Objednatel Zhotoviteli, zda trvá na jejím plnění ve vztahu k již zhotovené části díla, přičemž Zhotoviteli náleží za její plnění poměrná část odměny určená podle poměru části díla již zhotoveného vůči dílu jako celku; nesdělí-li Objednatel Zhotoviteli, že na plnění Smlouvy o zabezpečení podpory provozu trvá, zaniká Smlouva o zabezpečení podpory provozu jako celek k okamžiku doručení projevu vůle směřujícího k odstoupení od Smlouvy. Zhotovitel se zavazuje plnit své povinnosti ve vztahu k již zhotovené části díla

ve smyslu tohoto odstavce a případně rovněž své povinnosti k již zhotovení části díla dle Smlouvy o zabezpečení podpory provozu v souladu s pravidly dle tohoto odstavce.

12.4. Odstoupení od Smlouvy ze strany Objednatele – Objednatel je oprávněn odstoupit od této Smlouvy v těchto případech:

(a) Zhotovitel poruší povinnost z této Smlouvy zvlášť závažným způsobem, a to zejména pro neplnění harmonogramu;

(b) jestliže se Zhotovitel dostane do prodlení s prováděním předmětu díla, ať již jako celku či jeho jednotlivých částí, ve vztahu k termínům provádění předmětu díla dle čl. III. odst. 3.1. této Smlouvy, které bude delší než čtrnáct kalendářních dnů;

(c) Zhotovitel porušil některou ze svých povinností uvedených v čl. VIII. Smlouvy;

(d) Zhotovitel porušil některý ze svých závazků dle čl. VI. odst. 6.2. Smlouvy nebo se ukáže nepravdivým, neúplným či zkresleným některé z prohlášení Zhotovitele dle článku VI. odst. 6.1. této Smlouvy;

(e) Zhotovitel poruší povinnost mlčenlivosti dle čl. XVII odst. 17.6. této Smlouvy;

(f) Zhotovitel přestane být subjektem oprávněným poskytovat dodávky dle této Smlouvy;

(f) Zhotovitel nebo jeho poddodavatel bude identifikován jako osoba definovaná v článku 5k odst. 1 písm. a) – c) Nařízení Rady (EU) 2022/576 ze dne 8. dubna 2022 kterým se mění nařízení (EU) č. 833/2014 o omezujících opatřeních vzhledem k činnostem Ruska destabilizujícím situaci na Ukrajině.

12.5. V případě odstoupení od Smlouvy ze strany Objednatele vzniká Objednateli vůči Zhotoviteli nárok na úhradu prokázaných vícenákladů (tj. nákladů vynaložených Objednatelem nad cenu za provedení předmětu díla) vynaložených na dokončení předmětu díla třetí osobou a na úhradu škod vzniklých prodlením se splněním předmětu díla. Povinnost Zhotovitele zaplatit smluvní pokuty, k jejichž úhradě vznikla povinnost před odstoupením, nezaniká.

12.6. Rozhodnutí Objednatele vypovědět tuto Smlouvu není na újmu jakýmkoli dalším právům Objednatele vyplývajícím ze Smlouvy, právních předpisů nebo vzniklým z jiného titulu.

12.7. Rozhodnutí Zhotovitele vypovědět tuto Smlouvu není na újmu jakýmkoli dalším právům Zhotovitele vyplývajícím ze Smlouvy, právních předpisů nebo vzniklým z jiného titulu.

12.8. Objednatel je dále oprávněn odstoupit od této Smlouvy bez naplnění jakéhokoliv důvodu do 5 pracovních dnů od ukončení jednotlivých etap vymezených v odst. 3.1. (vč. předání veškeré související dokumentace). V případě, část díla odpovídající již provedeným etapám bude splňovat požadavky dle této Smlouvy, pak má Zhotovitel právo na vyplacení části ceny z díla dle čl. IV. odst. 4.1. této Smlouvy.

12.9. Zhotovitel je oprávněn odstoupit od této Smlouvy v případě, že Objednatel poruší povinnost z této Smlouvy zvlášť závažným způsobem.

XIII. Adresy pro doručování

13.1. Smluvní strany této Smlouvy se dohodly následujícím způsobem na adrese pro doručování písemné korespondence:

(a) adresa pro doručování Objednateli je: **v záhlaví této Smlouvy**, datová schránka: a8gbnyc

(b) adresa pro doručování Zhotoviteli je: Chebská 73/48, 360 06 Karlovy Vary, datová schránka: ctb7phe

13.2. Smluvní strany se dohodly, že v případě změny sídla, a tím i adresy pro doručování, budou písemné informovat o této skutečnosti bez zbytečného odkladu druhou smluvní stranu. Do doby nové adresy doručování se doručuje na stávající adresy.

XIV. Doručování

14.1. Smluvní strany se dohodly, že doručovat si budou zejména prostřednictvím datových schránek. Jiným způsobem (osobně nebo prostřednictvím držitele poštovní licence) je doručování možné pouze v případě, že je to vzhledem ke všem okolnostem vhodnější a doručování prostřednictvím datové schránky není možné (z důvodu času nebo věcně). Smluvní strany jsou povinné udržovat nastavení své datové schránky tak, aby doručování běžných písemností v souvislosti s touto Smlouvou umožňovaly (viz § 18a odst. 1 zákona č. 300/2008 Sb.). Smluvní strany jsou dále povinny zajistit, aby se do datové schránky přihlásil oprávněna osoba od podpisu této Smlouvy minimálně každé tři pracovní dny. Porušení této povinnosti má pro účely této Smlouvy za následek, že zásilka platí za odmítnutou, resp. že bylo doručení zmařeno.

14.2. Aniž by tím byly dotčeny další prostředky, kterými lze prokázat doručení, má se za to, že oznámení bylo řádně doručené:

(a) při doručování osobně:

- dnem faktického přijetí oznámení příjemcem; nebo
- dnem, v němž bylo doručeno osobě na příjemcově adrese určené k přebírání listovních zásilek; nebo
- dnem, kdy bylo doručováno osobě na příjemcově adrese určené k přebírání listovních zásilek, a tato osoba odmítla listovní zásilku převzít; nebo
- dnem, kdy příjemce při prvním pokusu o doručení zásilku z jakýchkoli důvodů nepřevzal či odmítl zásilku převzít, a to i přesto, že se v místě doručení nezdržuje, pokud byla na zásilce uvedena adresa pro doručování dle čl. XIII. odst. 13.1., resp. 13.2. této Smlouvy.

(b) při doručování prostřednictvím držitele poštovní licence:

- se má za to, že došlá zásilka odeslaná s využitím provozovatele poštovních služeb došla třetí pracovní den po odeslání, byla-li však odeslána na adresu v jiném státu, pak patnáctý pracovní den po odeslání, a to doručování na adresy pro doručování dle čl. XIII. odst. 13.1., resp. 13.2. této Smlouvy.

(c) při doručování do datové schránky:

- okamžikem přihlášení oprávněné osoby do datové schránky,
- pro případ, že se do datové schránky oprávněná osoba nepřihlásí, ani čtvrtý pracovní den od dodání zprávy do datové schránky platí, že zásilka je doručena pátým pracovním dnem od odeslání analogicky podle § 570 věta za středníkem občanského zákoníku pro zmaření doručení.

XV. Společná ustanovení

Pokud není v předchozích částech této Smlouvy uvedeno něco jiného, vztahují se na ně příslušné články společných ustanovení.

15.1. Smluvní strany se dohodly na tom, že jakákoliv peněžitá plnění dle Smlouvy jsou řádně a včas splněna, pokud byla příslušná částka odepsána z účtu povinné strany ve prospěch účtu oprávněné smluvní strany (věřitele) nejpozději v poslední den splatnosti.

15.2. Všechny spory, které vzniknou ze Smlouvy a v souvislosti s ní, se přednostně pokusí vyřešit prokazatelnou dohodou Smluvních stran. Objednatel se Zhotovitelem ujednali, v souladu s ustanovením § 89a zákona č. 99/1963 Sb., občanský soudní řád, ve znění pozdějších předpisů, že v případě jejich sporu, který by byl řešen soudní cestou, je místně příslušným soudem místně příslušný soud Objednatele. Rozhodným právem je právo České republiky s výjimkou kolizních norem mezinárodního soukromého práva a Úmluvy OSN o smlouvách o mezinárodní koupi zboží; rozhodným jazykem je český jazyk.

15.3. Smluvní strany se zavazují:

(a) vzájemně včas a řádně informovat o všech podstatných skutečnostech, které mohou mít vliv na plnění dle této Smlouvy,

(b) vyvinout potřebnou součinnost k plnění této Smlouvy.

15.4. Tato Smlouva může být měněna nebo doplňována pouze písemnými oboustranně odsouhlasenými, a průběžně číslovanými dodatky, podepsanými oprávněnými zástupci obou smluvních stran.

15.5. Přílohy uvedené v textu této Smlouvy a sumarizované v závěrečných ustanoveních Smlouvy tvoří součást Smlouvy.

15.6. Žádná Strana neuděluje druhé Straně právo užívat její ochranné známky či jiná označení (včetně ochranných známek či označení v rámci Podniku) pro účely propagace nebo publikování bez předchozího písemného souhlasu druhé Strany.

15.7. Smlouva nezakládá žádné zastoupení, společný podnik nebo partnerství mezi Objednatelem a Zhotovitelem. Obě Strany mohou svobodně uzavírat obdobné Smlouvy s jinými stranami za účelem vývoje, nákupu či poskytování konkurenčních produktů a služeb.

15.8. Žádný ze zaměstnanců nebo ostatních osob kterékoliv z obou Stran není oprávněn poskytovat záruky třetím stranám, které nejsou součástí Smlouvy a obě strany prohlašují, že se nespoléhaly na žádná taková ústní či písemná prohlášení při poskytování záruk, s výjimkou oprávněných statutárních zástupců obou Stran.

15.9. Obě Strany svým podpisem potvrzují, že tuto Smlouvu četly, rozumí jí a souhlasí s tím, že budou jejími podmínkami vázány. Dále souhlasí, že tato Smlouva nahrazuje jakékoliv předchozí dohody mezi Stranami a je nadřazena všem předchozím návrhům ústním či písemným a veškeré další komunikaci mezi oběma Stranami vztahující se k předmětu Smlouvy.

15.10. Pokud není uvedeno jinak, není ani jedna ze Stran oprávněna jednat jménem druhé Strany či zastupovat druhou Stranu jakýmkoliv způsobem při smluvních jednáních.

XVI. Autorské právo a ochrana duševního vlastnictví

16.1. Veškerá data zpracovávaná při poskytování Dodávek dle této Smlouvy jsou ve vlastnictví Objednatele; tedy Objednatel je dle dohody stran pořizovatelem příslušných databází ve smyslu § 89 autorského zákona.

16.2. Pro účely této Smlouvy se rozlišuje předmět práva duševního vlastnictví na:

a) plnění (zejm. počítačové programy a databáze) dodávané Zhotovitelem dle této Smlouvy, které byly vytvořeny nebo upraveny výhradně pro účely této Smlouvy (dále jen „**Individualizovaný software**“); a

b) plnění (zejm. počítačové programy a databáze) dodávané Zhotovitelem dle této Smlouvy, které nebyly vytvořeny nebo upraveny výhradně pro účely této Smlouvy (dále jen „**Standardizovaný software**“).

16.3. Zhotovitel poskytuje Objednateli oprávnění k výkonu práv duševního vlastnictví k Individualizovanému software (dále jen „**Licence k Individualizovanému software**“). Licence se týká veškerých autorských nebo jiných duševních práv (jejichž povaha umožňuje licenci v dále uvedeném rozsahu poskytnout) k Individualizovanému software. Dodavatel uděluje Objednateli Licenci k Individualizovanému software v následujícím rozsahu:

16.3.1. Licence je nevýhradní k veškerým známým způsobům užití Individualizovaného software, zejména (nikoliv však výlučně) k účelu, ke kterému byl příslušný Individualizovaný software Objednateli poskytnut v souladu se Smlouvou, a to v rozsahu minimálně nezbytném pro řádné užívání příslušného Individualizovaného software Objednatelem v souladu s touto Smlouvou.

16.3.2. Licence je ze strany Zhotovitele neodvolatelná a nevypověditelná.

16.3.3. Licence je neomezená územním či množstevním rozsahem, dále je Licence neomezená způsobem nebo rozsahem užití, zejména neomezená počtem uživatelů či mírou využívání.

16.3.4. Licence je převoditelná a postupitelná, tj. udělena s právem udělení sublicence či postoupení licence jakékoliv třetí osobě; s tímto vyslovuje Zhotovitel souhlas.

16.3.5. Licenci není Objednatel povinen využít.

16.3.6. Licence se poskytuje na dobu trvání autorských práv či jiných práv duševního vlastnictví k příslušnému dílu.

16.4. Zhotovitel se zavazuje společně s předáním díla předat Objednateli veškerou související dokumentaci a zdrojový kód k Individualizovanému software, a to tak, že

zdrojové kódy a dokumentace budou uloženy na k tomu vyhrazených datových prostředcích Objednatele nebo budou Objednateli předány na datovém nosiči s tím, že:

i. předaný kód musí být čitelný a komentovaný;

ii. předaný kód a dokumentace musí být vytvořeny podle nejnovějších a nejlepších standardů v oboru (tzv. best-practice), zejm. bude v přehledné a strukturované podobě;

iii. Zhotovitel musí poskytnout rovněž všechny nástroje a komponenty pro korektní editaci a kompilaci kódu (resp. musí specifikovat, které běžně dostupné nástroje byly použity, v jaké verzi a konfiguraci).

16.5. Zhotovitel uděluje Objednateli souhlas k tomu, aby byl Objednatel (či Objednatelem pověřená třetí osoba) oprávněn Individualizovaný software zveřejnit, upravovat, zpracovávat, překládat, měnit jeho název, spojit jej s dílem jiným a zařadit jej do díla souborného. Ustanovení tohoto odstavce neomezuje práva Objednatele plynoucí z § 66 autorského zákona.

16.6. Zhotovitel se zavazuje zajistit oprávnění k užití předmětu práv duševního vlastnictví jiných osob, který představuje Individualizovaný software, a souhlas se zásahy do takového předmětu práv duševního vlastnictví v takovém rozsahu, aby Objednateli mohl udělit Licenci k Individualizovanému software v rozsahu odst. 16.3. této Smlouvy, souhlas k zásahům ve smyslu odst. 16.5. této Smlouvy a aby mohl Objednateli předat zdrojový kód k takovému předmětu práv duševního vlastnictví v souladu s odst. 16.4. této Smlouvy.

16.7. Zhotovitel je povinen zajistit, aby Objednatel nabyl veškerá oprávnění k užití Standardizovanému software nezbytná k řádnému užívání díla Objednatelem a zachování jeho funkčnosti (dále jen „**Licence ke Standardizovanému software**“).

16.8. Zhotovitel je povinen poskytnout Licenci ke Standardizovanému software nebo zajistit její poskytnutí Objednateli třetí osobou v takovém rozsahu, že Standardizovaný software bude oprávněn užívat libovolný počet uživatelů po celou dobu, co bude Standardizovaný software Objednatelem využíván; ostatní parametry Licenci ke Standardizovanému software musí být uděleny v rozsahu, který zajistí plnou využitelnost Standardizovaného software při užívání díla bez nutnosti platit Zhotoviteli nebo třetí osobě odměnu nad rámec odměny dle této Smlouvy.

16.9. Zhotovitel je povinen společně s předáním příslušné části díla předat Objednateli rovněž informace o tom, jaký Standardizovaný software je součástí předávané části díla a jaké jsou licenční podmínky a jiné podmínky užití se k němu vztahující.

16.10. Zhotovitel musí předat Objednateli všechny využité komponenty, resp. součásti díla, které představují Standardizovaný software, a to ve stavu, který umožní jejich využití pro další rozvoj díla.

16.11. Cena za plný rozsah licencí dle tohoto článku (vč. případného poskytnutí uživatelské a technické dokumentace) je obsažena v odměně za příslušné plnění dle této Smlouvy. Zhotovitel prohlašuje, že před podáním nabídky do Zadávacího řízení pečlivě zvážil veškeré přínosy, které může poskytnutí těchto licencí Objednateli přinést, a že úplata za licence, která je zahrnuta v odměně za příslušné plnění dle této Smlouvy,

představuje adekvátní protiplnění Zhotoviteli za poskytnutí licencí. Zhotovitel dále prohlašuje a zavazuje se zajistit, že nositelům práv duševního vlastnictví k předmětu práv duševního vlastnictví, které je plněním dle této Smlouvy, nepřísluší a nebude příslušet vůči Objednateli žádné právo na odměnu, či jakékoliv jiné plnění v souvislosti s užitím příslušného plnění.

16.12. Udělení veškerých práv Objednateli na základě Licence k Individualizovanému software, případně Licencí ke Standardizovanému software nebo jiných oprávnění k užití předmětu práv duševního vlastnictví dle tohoto článku nelze ze strany Zhotovitele vypovědět nebo jinak jednostranně zrušit a ukončení závazku z této Smlouvy nemá vliv na udělení těchto práv.

XVII. Ochrana informací

17.1. Smluvní strany jsou si vědomy toho, že v rámci plnění této Smlouvy:

(a) si mohou vzájemně úmyslně nebo i opominutím poskytnout informace, které budou považovány za neveřejné (dále „neveřejné informace“),

(b) mohou jejich zaměstnanci získat vědomou činností druhé Strany nebo i jejím opominutím přístup k důvěrným informacím druhé Strany.

17.2. Strany se zavazují, že žádná z nich nezpřístupní třetí osobě neveřejné informace, které při plnění této Smlouvy nebo v souvislosti s plněním Smlouvy získala od druhé Strany.

17.3. Za třetí osoby se nepovažují:

(a) zaměstnanci Stran a osoby v obdobném postavení,

(b) orgány Stran a jejich členové a

(c) Poddodavatelé Zhotovitele,

za předpokladu, že se podílejí na plnění Smlouvy. Neveřejné informace jsou jim zpřístupněny výhradně za tímto účelem a zpřístupnění neveřejných informací je v rozsahu nezbytně nutném pro naplnění jeho účelu a za stejných podmínek, jaké jsou stanoveny Stranám ve Smlouvě.

17.4. Veškeré neveřejné informace zůstávají výhradním vlastnictvím předávající strany a přijímající strana vyvine pro zachování jejich neveřejnosti a pro jejich ochranu stejné úsilí, jako by se jednalo o její vlastní neveřejné informace. S výjimkou plnění této Smlouvy se obě strany zavazují neduplikovat žádným způsobem neveřejné informace druhé strany, nepředat je třetí straně ani svým vlastním zaměstnancům a zástupcům s výjimkou těch, kteří s nimi potřebují být seznámeni, aby mohli splnit tuto Smlouvu. Obě strany se zároveň zavazují nepoužít neveřejné informace druhé strany jinak než za účelem plnění této Smlouvy.

17.5 Smluvní strany se výslovně dohodly, že za neveřejné informace nejsou považovány informace poskytnuté v rámci Veřejné zakázky tzn., Zadávací dokumentace, nabídka Zhotovitele, smluvní dokumentace jakož i informace a dokumentace předané Zhotovitelem v rámci realizace předmětu plnění. Smluvní strany prohlašují, že skutečnosti uvedené v této Smlouvě nepovažují za obchodní tajemství ve smyslu § 504 občanského zákoníku, tímto výslovně souhlasí se zveřejněním veškerých náležitostí a

podmínek této Smlouvy a/nebo souvisejících dokumentů a informací, včetně zveřejnění této Smlouvy jako celku, v rámci informací zpřístupňovaných veřejnosti bez stanovení jakýchkoli dalších podmínek, a to i prostřednictvím dálkového přístupu, zejména na webových stránkách města. V případě utajovaných příloh (například podléhající obchodnímu tajemství) poskytovatel při podpisu Smlouvy předal nabyvateli verzi strany nebo přílohy, která zůstane neveřejná – z této listiny musí být patrný alespoň obsah tohoto dokumentu.

17.6. Strany se zavazují v plném rozsahu zachovávat povinnost mlčenlivosti a povinnost chránit neveřejné informace způsobem vyplývajícím ze Smlouvy a též z příslušných právních předpisů, zejména povinností vyplývajících ze zákona č. 110/2019 Sb., o zpracování osobních údajů, ve znění pozdějších předpisů,. Strany se v této souvislosti zavazují poučit veškeré osoby, které se budou podílet na plnění Smlouvy, o výše uvedených povinnostech mlčenlivosti a ochrany neveřejných informací a dále se zavazují vhodným způsobem zajistit dodržování těchto povinností všemi osobami podílejícími se na plnění Smlouvy.

17.7 Budou-li informace poskytnuté Objednatelem či třetími stranami, které jsou nezbytné pro plnění Smlouvy, obsahovat data podléhající režimu zvláštní ochrany podle zákona č. 110/2019 Sb., o zpracování osobních údajů, ve znění pozdějších předpisů, zavazuje se Zhotovitel zabezpečit splnění všech ohlašovacích povinností, které citovaný zákon vyžaduje po zpracovateli osobních údajů, a v případě, že v rámci plnění povinností dle této Smlouvy je Zhotovitel povinen údaje od subjektů údajů též získat, pak je povinen obstarat předepsané souhlasy subjektů osobních údajů předaných ke zpracování.

17.8. Pokud jsou neveřejné informace poskytovány v písemné podobě nebo ve formě textových souborů na počítačových médiích, je předávající strana povinna upozornit přijímající stranu na neveřejnost takového materiálu jejím vyznačením alespoň na titulní stránce.

17.9. Bez ohledu na výše uvedená ustanovení se za neveřejné nepovažují informace, které:

(a) se staly veřejně známými, aniž by to zavinila záměrně či opominutím přijímající strana,

(b) měla přijímající strana legálně k dispozici před uzavřením této Smlouvy, pokud takové informace nebyly předmětem jiné, dříve mezi Smluvními stranami uzavřené smlouvy o ochraně informací,

(c) jsou výsledkem postupu, při kterém k nim přijímající strana dospěje nezávisle a je to schopna doložit svými záznamy nebo neveřejnými informacemi třetí strany,

(d) po podpisu této Smlouvy poskytne přijímající straně třetí osoba, jež takové informace přitom nezíská přímo ani nepřímo od strany, jež je jejich vlastníkem.

17.10. Ustanovení tohoto článku není dotčeno ukončením účinnosti této Smlouvy z jakéhokoliv důvodu po dobu dalších 5 let od ukončení účinnosti Smlouvy. Ochrana osobních údajů třetích osob není lhůtou omezena.

17.11. Zhotovitel bere na vědomí, že Objednatel v Software dodaném na základě této Smlouvy bude zpracovávat osobní údaje. Software musí respektovat Nařízení EU a právní předpisy České republiky s ochranou osobních údajů související, zejména

a) Nařízení Evropského parlamentu a Rady 2016/679 Obecné nařízení na ochranu osobních údajů neboli GDPR (General Data Protection Regulation) a

b) Nařízení Evropského parlamentu a Radyč. 910/2014 o elektronické identifikaci a důvěryhodných službách pro elektronické transakce na vnitřním evropském trhu EIDAS

a prohlašuje, že dodaný Software odpovídá těmto a dalším předpisům týkajících se ochrany osobních údajů.

17.12. Zhotovitel je dále povinen:

a) Zajistit, aby osobní údaje nebyly uchovávány mimo území České republiky, resp. mimo EU;

b) bez zbytečného odkladu informovat Objednatele o porušení zabezpečení osobních údajů a být nápomocen při plnění povinností Objednatele, zejména povinnosti reagovat na žádosti o výkon práv subjektů údajů, povinnosti ohlašovat případy porušení zabezpečení osobních údajů dozorovému úřadu dle čl. 33 obecného nařízení o ochraně osobních údajů, povinnosti oznamovat případy porušení zabezpečení osobních údajů subjektu údajů dle čl. 34 obecného nařízení o ochraně osobních údajů, povinnosti posoudit vliv na ochranu osobních údajů dle čl. 35 obecného nařízení o ochraně osobních údajů a povinnosti provádět předchozí konzultace dle čl. 36 obecného nařízení o ochraně osobních údajů, a že za tímto účelem zajistí nebo přijme vhodná technická a organizační opatření, o kterých ihned informuje Objednatele;

c) po ukončení Smlouvy řádně naložit se zpracovávanými osobními údaji, např. že všechny osobní údaje vymaže, nebo je vrátí Objednateli a vymaže existující kopie, dle požadavku Objednatele;

d) při porušení povinností stanovených touto Smlouvou v případech, kdy byla v důsledku nedodržení zásad a pravidel ochrany osobních údajů uložena Objednateli pokuta ze strany příslušných dozorujících a kontrolních orgánů, je Zhotovitel povinen uhradit Objednateli smluvní pokutu ve výši odpovídající vyměřené pokutě. Úhradou této smluvní pokuty není dotčeno právo na náhradu škody vzniklé Objednateli nad rámec takto vyměřené pokuty.

XVIII. Závěrečná ustanovení

18.1. Práva a povinnosti z této Smlouvy vyplývající a ve Smlouvě neupravené se řídí příslušnými ustanoveními občanského zákoníku, autorským zákonem, a s obsahem Smlouvy souvisejících předpisů. Pokud by bylo jedno z výše uvedených ustanovení zcela nebo zčásti právně neúčinné, zůstává tím nedotčena právní účinnost ostatních ustanovení. Totéž platí i pro případ smluvní mezery.

18.2. Pokud jakýkoliv závazek vyplývající z této Smlouvy, avšak netvořící její podstatnou náležitost je nebo se stane neplatným nebo nevymahatelným jako celek nebo jeho část, je plně oddělitelným od ostatních ustanovení této Smlouvy a taková neplatnost nebo nevymahatelnost nebude mít žádný vliv na platnost a vymahatelnost jakýchkoliv

ostatních závazků z této Smlouvy. Strany se zavazují v rámci této Smlouvy nahradit formou dodatku k této Smlouvě tento neplatný nebo nevymahatelný oddělený závazek takovým novým platným a vymahatelným závazkem, jehož předmět bude v nejvyšší možné míře odpovídat předmětu původního odděleného závazku. Pokud však jakýkoliv závazek vyplývající z této Smlouvy a tvořící její podstatnou náležitost je nebo kdykoliv se stane neplatným nebo nevymahatelným jako celek nebo jeho část, strany nahradí neplatný nebo nevymahatelný závazek v rámci nové smlouvy takovým novým platným a vymahatelným závazkem, jehož předmět bude v nejvyšší možné míře odpovídat předmětu původního závazku obsaženém v této Smlouvě. Totéž platí i pro případ smluvní mezery.

18.3. Strany tímto prohlašují, že si nejsou vědomy, že by kterákoliv Strana při sjednávání této Smlouvy zneužila svou kvalitu odborníka či své hospodářské postavení, přičemž Strany prohlašují, že vzájemná práva a povinnosti sjednané v této Smlouvě považují za rovnovážná.

18.4. Zhotovitel bere na vědomí a výslovně souhlasí, že se podpisem této Smlouvy stává v souladu s ustanovením § 2 písm. e) zákona č. 320/2001 Sb., o finanční kontrole ve veřejné správě a o změně některých zákonů, ve znění pozdějších předpisů, osobou povinnou spolupůsobit při výkonu finanční kontroly prováděné v souvislosti s úhradou zboží nebo služeb z veřejných výdajů nebo z veřejné finanční podpory. Zhotovitel je povinen při kontrole poskytnout na vyžádání kontrolnímu orgánu daňovou evidenci v plném rozsahu. Zhotovitel je rovněž povinen smluvně zajistit, aby totožným způsobem byli povinni působit i všichni z jeho poddodavatelů.

18.5. Uzavření této Smlouvy bylo schváleno Radou města Cheb dne 28. 11. 2024, číslo usnesení 690/16/2024.

18.6. Tuto Smlouvu je možné měnit pouze písemnou dohodou Smluvních stran ve formě vzestupně číslovaných dodatků této Smlouvy, podepsaných oprávněnými zástupci Smluvních stran. Žádné úkony či jednání ze strany jedné Smluvní strany nelze považovat za příslib uzavření Smlouvy nebo dodatku k této Smlouvě. V souladu s ustanovením § 1740 odst. 3 občanského zákoníku, Smluvní strany nepřipouští přijetí návrhu na uzavření Smlouvy s dodatkem nebo odchylkou, čímž druhá Smluvní strana podpisem Smlouvy souhlasí. Smluvní strany se dále dohodly, že možnost zhojení nedostatku písemné formy právního jednání se vylučuje, a že neplatnost právního jednání, pro nějž si Smluvní strany sjednaly písemnou formu, lze namítnout kdykoliv, tedy že mezi Smluvními stranami neplatí ustanovení § 582 odst. 1 první věta a odst. 2 občanského zákoníku.

18.7. Smluvní strany souhlasí s tím, že text Smlouvy je veřejně přístupnou listinou ve smyslu zákona č. 106/1999 Sb. o svobodném přístupu k informacím, ve znění pozdějších předpisů, a že Objednatel jako povinný subjekt má povinnost na žádost poskytnout informace o tomto smluvním vztahu včetně poskytnutí kopie Smlouvy.

18.8. Tato Smlouva nabývá platnosti dnem jejího podpisu oběma Smluvními stranami a účinnosti dnem jejího uveřejnění v Registru smluv dle zákona č. 340/2015 Sb., o zvláštních podmínkách účinnosti některých smluv, uveřejňování těchto smluv a o registru smluv (zákon o registru smluv). Smlouvu se zavazuje bez zbytečného odkladu v Registru smluv uveřejnit Objednatel.

18.9. Smluvní strany se dohodly, že v případě zániku právního vztahu založeného touto Smlouvou zůstávají v platnosti a účinnosti i nadále ustanovení, z jejichž povahy vyplývá, že mají zůstat nedotčena zánikem právního vztahu založeného touto Smlouvou.

18.10. Součást této Smlouvy tvoří:

Příloha č. 1: Technická specifikace

Příloha č. 2: Návrh Zhotovitele – Popis nabízeného technického řešení

Příloha č. 3: Realizační tým

Příloha č. 4: Podrobná kalkulace

V případě rozporu mezi různými částmi této Smlouvy, není-li určeno jinak, mají přednost dokumenty této Smlouvy v následujícím pořadí:

- Technická specifikace
- Návrh Zhotovitele
- očíslované články této Smlouvy
- ostatní přílohy.

V Chebu, dne (dle el. podpisu)

_____	_____
za Objednatele	za Zhotovitele

BitLocker overview

- Article
- 12/05/2024
- 3 contributors
- Applies to:

✅ Windows 11, ✅ Windows 10, ✅ Windows Server 2025, ✅ Windows Server 2022, ✅ Windows Server 2019, ✅ Windows Server 2016

Feedback

In this article

1. Practical applications
2. BitLocker and TPM
3. System requirements
4. Windows edition and licensing requirements

Show 2 more

BitLocker is a Windows security feature that provides encryption for entire volumes, addressing the threats of data theft or exposure from lost, stolen, or inappropriately decommissioned devices.

Practical applications

Data on a lost or stolen device is vulnerable to unauthorized access, either by running a software-attack tool against it, or by transferring the device's hard drive to a different device. BitLocker helps mitigate unauthorized data access by enhancing file and system protections, rendering data inaccessible when BitLocker-protected devices are decommissioned or recycled.

BitLocker and TPM

BitLocker provides maximum protection when used with a Trusted Platform Module (TPM), which is a common hardware component installed on Windows devices. The TPM works with BitLocker to ensure that a device hasn't been tampered with while the system is offline.

In *addition* to the TPM, BitLocker can lock the normal startup process until the user supplies a personal identification number (PIN) or inserts a removable device that contains a *startup key*. These security measures provide multifactor authentication and assurance that the device can't start or resume from hibernation until the correct PIN or startup key is presented.

On devices that don't have a TPM, BitLocker can still be used to encrypt the operating system drive. This implementation requires the user to either:

- use a startup key, which is a file stored on a removable drive that is used to start the device, or when resuming from hibernation

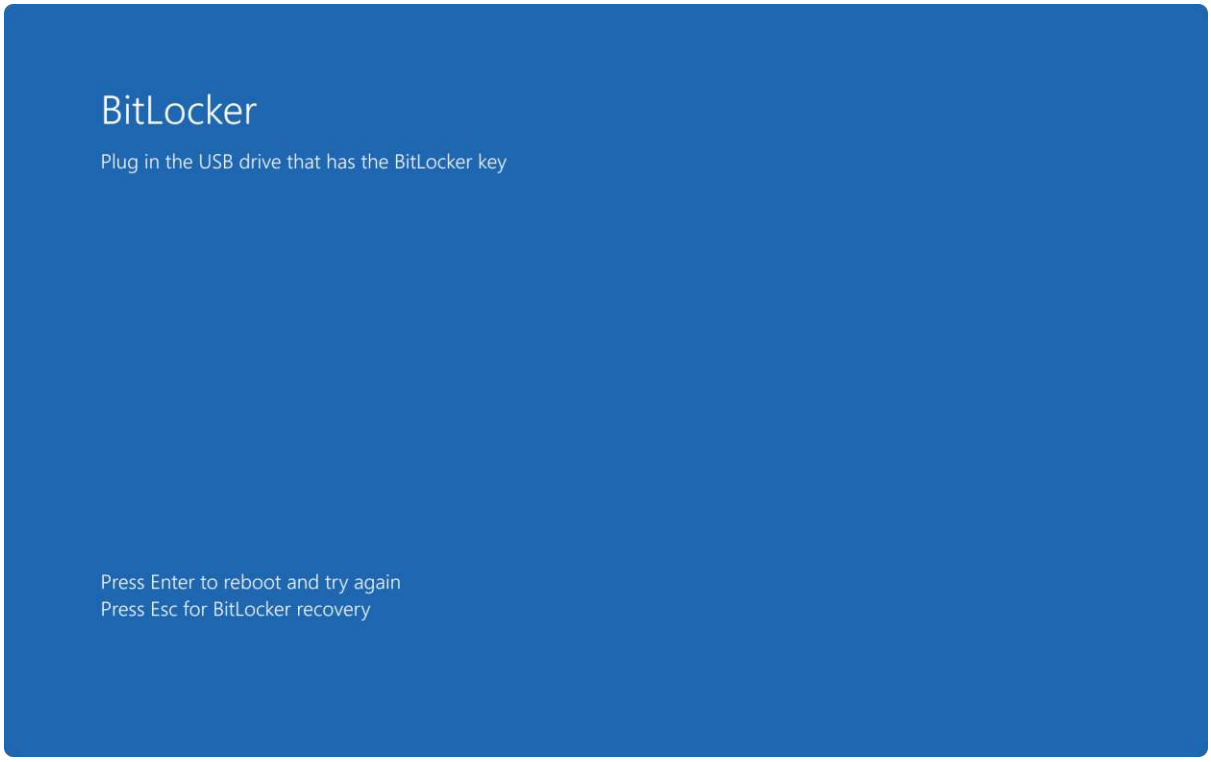
- use a password. This option isn't secure since it's subject to brute force attacks as there isn't a password lockout logic. As such, the password option is discouraged and disabled by default

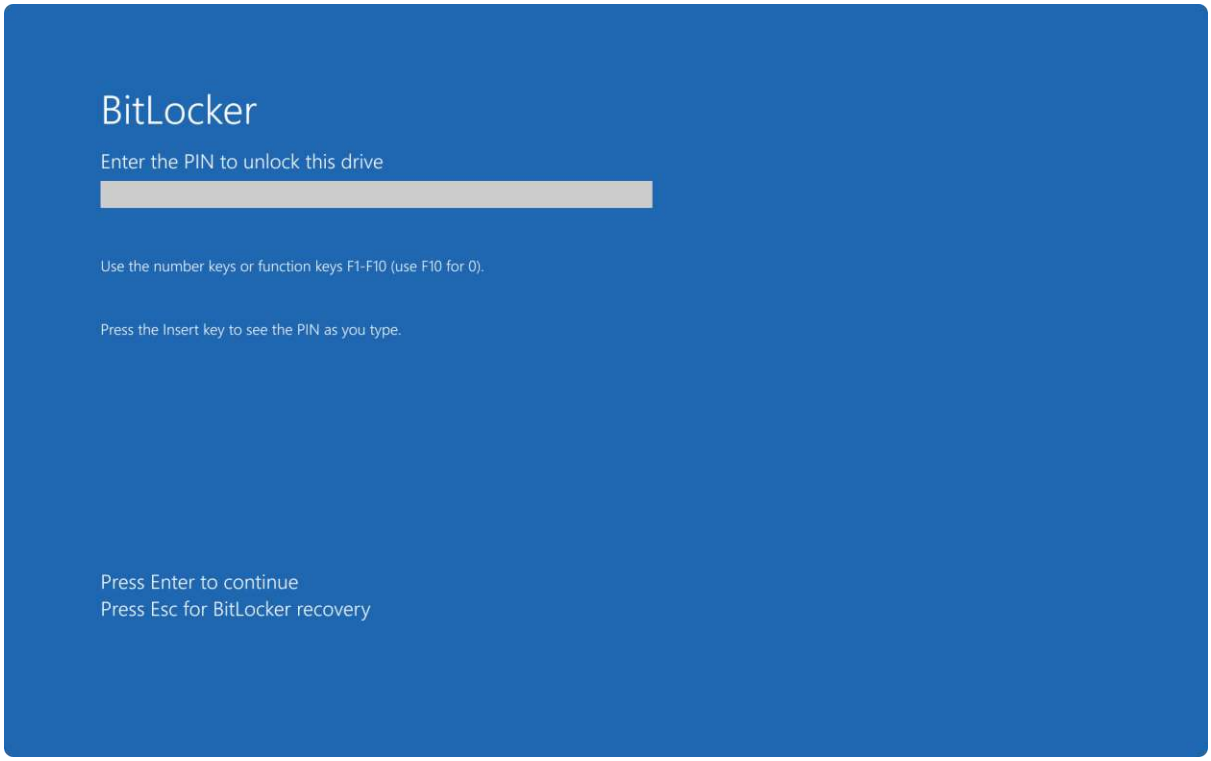
Both options don't provide the preboot system integrity verification offered by BitLocker with a TPM.

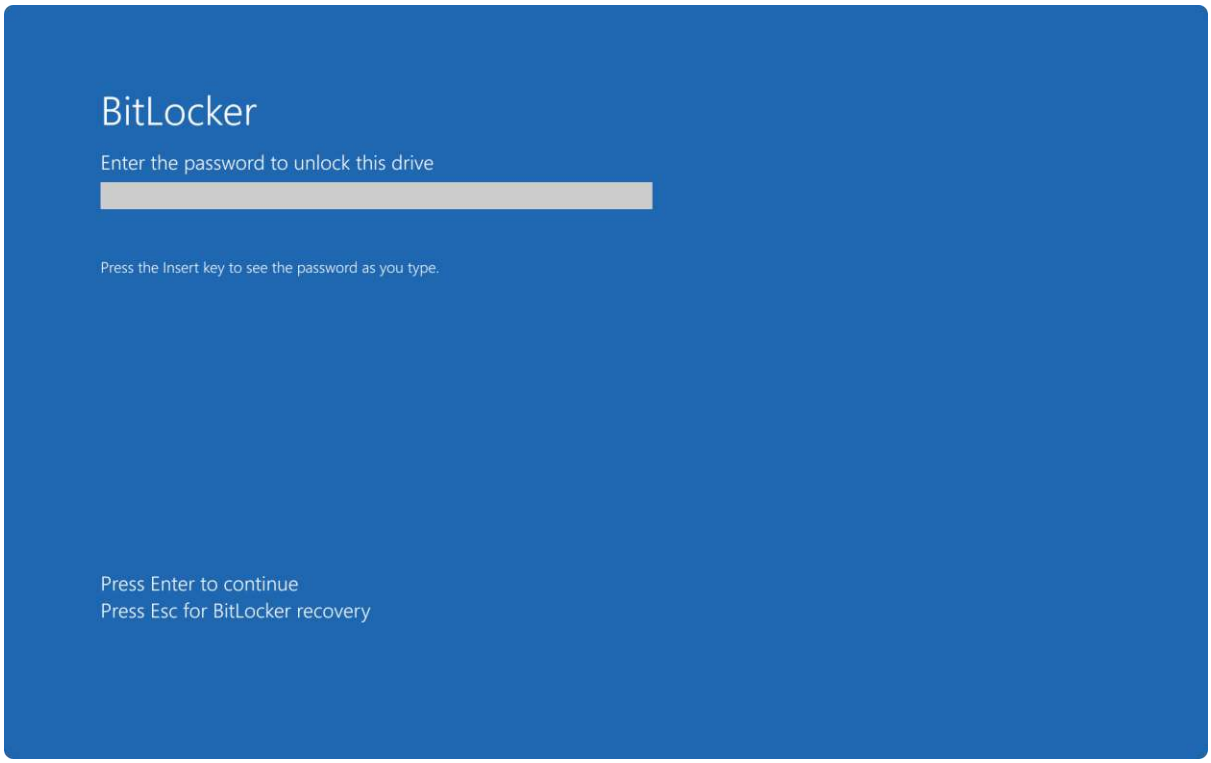
BitLocker preboot screen with startup key:

BitLocker preboot screen with PIN:

BitLocker preboot screen with password:







System requirements

BitLocker has the following requirements:

- For BitLocker to use the system integrity check provided by a TPM, the device must have TPM 1.2 or later versions. If a device doesn't have a TPM, saving a startup key on a removable drive is mandatory when enabling BitLocker

- A device with a TPM must also have a *Trusted Computing Group* (TCG)-compliant BIOS or UEFI firmware. The BIOS or UEFI firmware establishes a chain of trust for the preboot startup, and it must include support for *TCG-specified Static Root of Trust Measurement*. A computer without a TPM doesn't require TCG-compliant firmware
- The system BIOS or UEFI firmware (for TPM and non-TPM devices) must support the USB mass storage device class, and reading files on a USB drive in the preboot environment

Note

TPM 2.0 is not supported in *Legacy* and *Compatibility Support Module (CSM)* modes of the BIOS. Devices with TPM 2.0 must have their BIOS mode configured as native UEFI only. The Legacy and CSM options must be disabled. For added security, enable the *secure boot* feature.

Installed operating system on hardware in Legacy mode stops the OS from booting when the BIOS mode is changed to UEFI. Use the tool **mbr2gpt.exe** before changing the BIOS mode, which prepares the OS and the disk to support UEFI.

- The hard disk must be partitioned with at least two drives:
    - The *operating system drive* (or boot drive) contains the OS and its support files. It must be formatted with the NTFS file system
    - The *system drive* contains files required to boot, decrypt, and load the operating system. BitLocker isn't enabled on this drive. For BitLocker to work, the system drive:
        - must not be encrypted
        - must differ from the operating system drive
        - must be formatted with the FAT32 file system on computers that use UEFI-based firmware, or with the NTFS file system on computers that use BIOS firmware
        - it's recommended that to be approximately 350 MB in size. After BitLocker is turned on, it should have approximately 250 MB of free space

Important

When installed on a new device, Windows automatically creates the partitions that are required for BitLocker.

If the drive was prepared as a single contiguous space, BitLocker requires a new volume to hold the boot files. BdeHdCfg.exe can create the volume. For more information about using the tool, see **Bdehdcfg** in the Command-Line Reference.

Note

When installing the BitLocker optional component on a server, the *Enhanced Storage* feature must be installed. The feature is used to support hardware encrypted drives.

Windows edition and licensing requirements

The following table lists the Windows editions that support BitLocker enablement:

Expand table

Windows Pro	Windows Enterprise	Windows Pro Education/SE	Windows Education
Yes	Yes	Yes	Yes

BitLocker enablement license entitlements are granted by the following licenses:

Expand table

Windows Pro/Pro Education/SE	Windows Enterprise E3	Windows Enterprise E5	Windows Education A3	Windows Education A5
Yes	Yes	Yes	Yes	Yes

For more information about Windows licensing, see Windows licensing overview.

Note

Licensing requirements for BitLocker enablement are different from the licensing requirements for BitLocker *management*. To learn more, review the how-to guide: **configure BitLocker**.

Device encryption

Device encryption is a Windows feature that provides a simple way for some devices to enable BitLocker encryption automatically. Device encryption is available on all Windows versions, and it requires a device to meet either Modern Standby or HSTI security requirements. Device encryption can't have externally accessible ports that allow DMA access. Device encryption encrypts only the OS drive and fixed drives, it doesn't encrypt external/USB drives.

Important

Starting in Windows 11, version 24H2, the prerequisites of DMA and HSTI/Modern Standby are removed. As a result, more devices are eligible for automatic and manual device encryption. For more information, see **BitLocker drive encryption in Windows 11 for OEMs**.

Unlike a standard BitLocker implementation, device encryption is enabled automatically so that the device is always protected. When a clean installation of Windows is completed and the out-of-box experience is finished, the device is prepared for first use. As part of this preparation, device encryption is initialized on the OS drive and fixed data drives on the computer with a clear key that is the equivalent of standard BitLocker suspended state. In this state, the drive is shown with a warning icon in Windows Explorer. The yellow warning icon is removed after the TPM protector is created and the recovery key is backed up.

- If the device is Microsoft Entra joined or Active Directory domain joined, the clear key is removed once the recovery key is successfully backed up to Microsoft Entra ID or Active Directory Domain Services (AD DS). The following policy settings must be enabled for the recovery key to be backed up: Choose how BitLocker-protected operating system drives can be recovered
    - For Microsoft Entra joined devices: the recovery password is created automatically when the user authenticates to Microsoft Entra ID, then the recovery key is backed up to Microsoft Entra ID, the TPM protector is created, and the clear key is removed

    - For AD DS joined devices: the recovery password is created automatically when the computer joins the domain. The recovery key is then backed up to AD DS, the TPM protector is created, and the clear key is removed

- If the device isn't Microsoft Entra joined or Active Directory domain joined, a Microsoft account with administrative privileges on the device is required. When the administrator uses a Microsoft account to sign in, the clear key is removed, a recovery key is uploaded to the online Microsoft account, and a TPM protector is created. Should a device require the recovery key, the user is guided to use an alternate device and navigate to a recovery key access URL to retrieve the recovery key by using their Microsoft account credentials
- If a device uses only local accounts, then it remains unprotected even though the data is encrypted

Important

Device encryption uses the XTS-AES 128-bit encryption method, by default. In case you configure a policy setting to use a different encryption method, you can use the Enrollment Status Page to avoid the device to begin encryption with the default method. BitLocker has a logic that doesn't start encrypting until the end of OOBE, after the Enrollment Status Page device configuration phase is complete. This logic gives a device enough time to receive the BitLocker policy settings before starting encryption.

If a different encryption method and/or cipher strength is needed but the device is already encrypted, it must first be decrypted before the new encryption method and/or cipher strength can be applied. After the device is decrypted, you can apply different BitLocker settings.

If a device doesn't initially qualify for device encryption, but then a change is made that causes the device to qualify (for example, by turning on *Secure Boot*), device encryption enables BitLocker automatically as soon as it detects it.

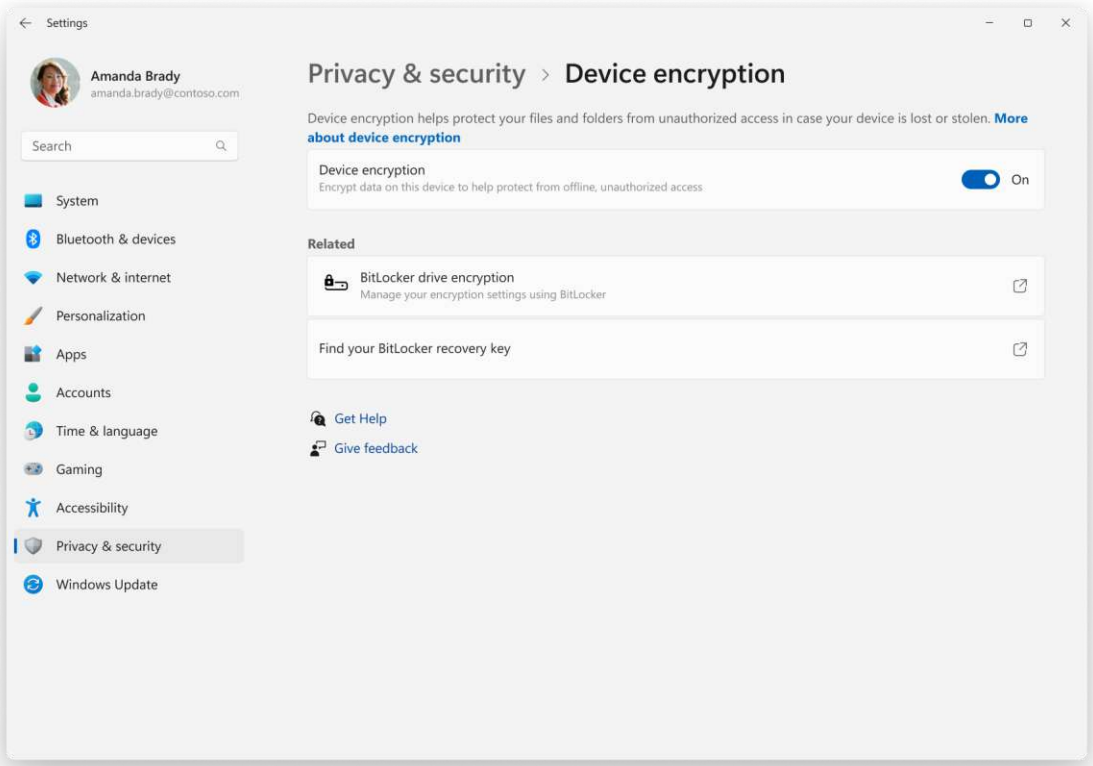
You can check whether a device meets requirements for device encryption in the System Information app (msinfo32.exe). If the device meets the requirements, System Information shows a line that reads:

Expand table

Item	**Value**
Device Encryption Support	Meets prerequisites

Difference between BitLocker and device encryption

- Device encryption turns on BitLocker automatically on device encryption-qualifying devices, with the recovery key automatically backed up to Microsoft Entra ID, AD DS, or the user's Microsoft account
- Device encryption adds a device encryption setting in the Settings app, which can be used to turn device encryption on or off
    - The Settings UI doesn't show device encryption enabled until encryption is complete



Note

If device encryption is turned off, it will no longer automatically enable itself in the future. The user must enable it manually in Settings

Disable device encryption

It's recommended to keep device encryption on for any systems that support it. However, you can prevent the automatic device encryption process by changing the following registry setting:

Expand table

Path	Name	Type
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BitLocker	PreventDeviceEncryption	REG_DWORD

For more information about device encryption, see BitLocker device encryption hardware requirements.

Next steps

Learn about technologies and features to protect against attacks on the BitLocker encryption key:

BitLocker countermeasures >

Group Policy overview

- Article
- 04/22/2024
- Applies to:

✅ Windows Server 2025, ✅ Windows Server 2022, ✅ Windows Server 2019, ✅ Windows Server 2016

Feedback

In this article

1. What is a Group Policy
2. How Group Policy works
3. What is an Organizational Unit (OU)
4. Related content

Group Policy enables configuration and settings management of user and computer settings on computers running Windows Server and Windows Client operating systems. In addition to using Group Policy to define configurations for groups of users and client computers, you can also use Group Policy to help manage server computers, by configuring many server-specific operational and security settings.

What is a Group Policy

Group policy can represent policy settings in the locally in the file system or in Active Directory Domain Services. When used with Active Directory, Group Policy settings are contained in a Group Policy Object (GPO). A GPO is a virtual collection of policy settings, security permissions, and scope of management (SOM) that you can apply to users and computers in Active Directory. A GPO has a unique name, such as a GUID. Clients evaluate GPO settings using the hierarchical nature of Active Directory.

Policy settings are divided into policy settings that affect a computer and policy settings that affect a user. Computer-related policies specify system behavior, application settings, security settings, assigned applications, and computer startup and shutdown scripts. User-related policies specify system behavior, application settings, security settings, assigned and published applications, user logon and logoff scripts, and folder redirection. Computer settings override user-related settings.

To create Group Policy, an administrator can use the Local Group Policy Editor (gpedit.msc), which can be a stand-alone tool and the settings stored locally. We recommend that you use the Group Policy Object Editor as an extension to an Active Directory-related MMC snap-in. The Group Policy Object Editor allows you to link GPOs to selected Active Directory sites, domains, and organizational units (OUs). Linking applies the policy settings in the GPO to the users and computers in those Active Directory objects. GPOs are stored in both Active Directory and in the SYSVOL folder on each domain controller.

How Group Policy works

For computers, Group Policy is applied when the computer starts. For users, Group Policy is applied at sign in. This initial processing of policy can also be referred to as a foreground policy application.

The foreground application of Group Policy can be synchronous or asynchronous. In synchronous mode, the computer doesn't complete the system start until computer policy is applied successfully. The user logon process doesn't complete until user policy is applied successfully. In asynchronous mode, if there are no policy changes that require synchronous processing, the computer can complete the start sequence before the application of computer policy is complete. The shell can be available to the user before the application of user policy is complete. The system then periodically applies (refreshes) Group Policy in the background. During a refresh, policy settings are applied asynchronously.

To learn more about how Group Policies work, see Group Policy Processing.

What is an Organizational Unit (OU)

An OU is the lowest-level Active Directory container to which you can assign Group Policy settings. Typically, you assign most GPOs at the OU level, so make sure that your OU structure supports your Group Policy-based client-management strategy. You might also apply some Group Policy settings at the domain level, particularly password policies. Few policy settings are applied at the site level. A well-designed OU structure that reflects the administrative structure of your organization and takes advantage of GPO inheritance simplifies the application of Group Policy. For example, a well-designed OU structure can prevent duplicating certain GPOs so that you can apply these GPOs to different parts of the organization. If possible, create OUs to delegate administrative authority and to help implement Group Policy.

OU design requires balancing requirements for delegating administrative rights independent of Group Policy needs, and the need to scope the application of Group Policy. You can create OUs within a domain and delegate administrative control for specific OUs to particular users or groups. By using a structure in which OUs contain homogeneous objects, such as either user or computer objects but not both, you can easily disable those sections of a GPO that don't apply to a particular type of object. This approach to OU design reduces complexity and improves the speed at which Group Policy is applied. GPOs linked to the higher layers of the OU structure are inherited by default for OUs at the lower layer, reducing the need to duplicate GPOs or to link a GPO to multiple containers.

Related content

- Group Policy Processing
- Backing up, restoring, migrating, and copying Group Policy Objects (GPOs)
- Group Policy Modeling and Group Policy Results
- Group Policy Management Console (GPMC)
- Group Policy Settings Used in Windows Authentication

Manage Device Installation with Group Policy

- Article
- 07/01/2024
- 1 contributor
- Applies to:

✅ Windows 11, ✅ Windows 10

Feedback

In this article

1. Introduction
2. Scenario Overview
3. Technology Review
4. Requirements for completing the scenarios

Show 7 more

By using Windows operating systems, administrators can determine what devices can be installed on computers they manage. This guide summarizes the device installation process and demonstrates several techniques for controlling device installation by using Group Policy.

Introduction

General

This step-by-step guide describes how you can control device installation on the computers that you manage, including designating which devices users can and can't install. This guide applies to all Windows versions starting with Windows 10, version 1809. The guide includes the following scenarios:

- Prevent users from installing devices that are on a "prohibited" list. If a device isn't on the list, then the user can install it.
- Allow users to install only devices that are on an "approved" list. If a device isn't on the list, then the user can't install it.

This guide describes the device installation process and introduces the device identification strings that Windows uses to match a device with the device-driver packages available on a machine. The guide also illustrates two methods of controlling device installation. Each scenario shows, step by step, one method you can use to allow or prevent the installation of a specific device or a class of devices.

The example device used in the scenarios is a USB storage device. You can perform the steps in this guide using a different device. However, if you use a different device, then the instructions in the guide won't exactly match the user interface that appears on the computer.

It's important to understand that the Group Policies that are presented in this guide are only applied to machines/machine-groups, not to users/user-groups.

Important

The steps provided in this guide are intended for use in a test lab environment. This step-by-step guide isn't meant to be used to deploy Windows Server features without accompanying documentation and should be used with discretion as a stand-alone document.

Who Should Use This Guide?

This guide is targeted at the following audiences:

- Information technology planners and analysts who are evaluating Windows 10, Windows 11 or Windows Server 2022
- Enterprise information technology planners and designers
- Security architects who are responsible for implementing trustworthy computing in their organization
- Administrators who want to become familiar with the technology

Benefits of Controlling Device Installation Using Group Policy

Restricting the devices that users can install reduces the risk of data theft and reduces the cost of support.

Reduce the risk of data theft

It's more difficult for users to make unauthorized copies of company data if users' computers can't install unapproved devices that support removable media. For example, if users can't install a USB thumb-drive device, they can't download copies of company data onto a removable storage. This benefit can't eliminate data theft, but it creates another barrier to unauthorized removal of data.

Reduce support costs

You can ensure that users install only those devices that your technical support team is trained and equipped to support. This benefit reduces support costs and user confusion.

Scenario Overview

The scenarios presented in this guide illustrate how you can control device installation and usage on the computers that you manage. The scenarios use Group Policy on a local machine to simplify using the procedures in a lab environment. In an environment where you manage multiple client computers, you should apply these settings using Group Policy. With Group Policy deployed by Active Directory, you can apply settings to all computers that are members of a domain or an organizational unit in a domain. For more information about how to create a Group policy object to manage your client computers, see Create a Group Policy Object.

Expand table

Scenario	Description
Scenario #1: Prevent installation of all printers	In this scenario, the administrator wants to prevent users from installing any printers. Thus is a basic scenario to introduce you

Scenario	Description
	to the 'prevent/allow' functionality of Device Installation policies in Group Policy.
Scenario #2: Prevent installation of a specific printer	In this scenario, the administrator allows standard users to install all printers while but preventing them from installing a specific one.
Scenario #3: Prevent installation of all printers while allowing a specific printer to be installed	In this scenario, you combine what you learned from both scenario #1 and scenario #2. The administrator wants to allow standard users to install only a specific printer while preventing the installation of all other printers. This scenario is a more realistic one and brings you a step farther in understanding of the Device Installation Restrictions policies.
Scenario #4: Prevent installation of a specific USB device	This scenario, although similar to scenario #2, brings another layer of complexity-how does device connectivity work in the PnP tree. The administrator wants to prevent standard users from installing a specific USB device. By the end of the scenario, you should understand the way devices are nested in layers under the PnP device connectivity tree.
Scenario #5: Prevent installation of all USB devices while allowing an installation of only an authorized USB thumb drive	In this scenario, combining all previous four scenarios, you learn how to protect a machine from all unauthorized USB devices. The administrator wants to allow users to install only a small set of authorized USB devices while preventing any other USB device from being installed. In addition, this scenario includes an explanation of how to apply the 'prevent' functionality to existing USB devices that have already been installed on the machine, and the administrator likes to prevent any farther interaction with them (blocking them all together). This scenario builds on the policies and structure we introduced in the first four scenarios and therefore it's preferred to go over them first before attempting this scenario.

Technology Review

The following sections provide a brief overview of the core technologies discussed in this guide and give background information that is necessary to understand the scenarios.

Device Installation in Windows

A device is a piece of hardware with which Windows interacts to perform some function, or in a more technical definition-it's a single instance of a hardware component with a unique representation in the Windows Plug and Play subsystem. Windows can communicate with a device only through a piece of software called a device-driver (also known as a *driver*). To install a driver, Windows detects the device, recognizes its type, and then finds the driver that matches that type.

When Windows detects a device that has never been installed on the computer, the operating system queries the device to retrieve its list of device identification strings. A device usually has

multiple device identification strings, which the device manufacturer assigns. The same device identification strings are included in the .inf file (also known as an *INF*) that is part of the driver package. Windows chooses which driver package to install by matching the device identification strings retrieved from the device to those strings included with the driver packages.

Windows uses four types of identifiers to control device installation and configuration. You can use the Group Policy settings in Windows to specify which of these identifiers to allow or block.

The four types of identifiers are:

- Device Instance ID
- Device ID
- Device setup classes
- 'Removable Devices' device type

Device Instance ID

A device instance ID is a system-supplied device identification string that uniquely identifies a device in the system. The Plug and Play (PnP) manager assigns a device instance ID to each device node (devnode) in a system's device tree.

Device ID

Windows can use each string to match a device to a driver package. The strings range from the specific, matching a single make and model of a device, to the general, possibly applying to an entire class of devices. There are two types of device identification strings: hardware IDs and compatible IDs.

Hardware IDs

Hardware IDs are the identifiers that provide the exact match between a device and a driver package. The first string in the list of hardware IDs is referred to as the device ID, because it matches the exact make, model, and revision of the device. The other hardware IDs in the list match the details of the device less exactly. For example, a hardware ID might identify the make and model of the device but not the specific revision. This scheme allows Windows to use a driver for a different revision of the device if the driver for the correct revision isn't available.

Compatible IDs

Windows uses these identifiers to select a driver if the operating system can't find a match with the device ID or any of the other hardware IDs. Compatible IDs are listed in the order of decreasing suitability. These strings are optional, and, when provided, they're generic, such as Disk. When a match is made using a compatible ID, you can typically use only the most basic functions of the device.

When you install a device, such as a printer, a USB storage device, or a keyboard, Windows searches for driver packages that match the device you're attempting to install. During this search, Windows assigns a "rank" to each driver package it discovers with at least one match to a hardware or compatible ID. The rank indicates how well the driver matches the device. Lower rank numbers indicate better matches between the driver and the device. A rank of zero

represents the best possible match. A match with the device ID to one in the driver package results in a lower (better) rank than a match to one of the other hardware IDs. Similarly, a match to a hardware ID results in a better rank than a match to any of the compatible IDs. After Windows ranks all of the driver packages, it installs the one with the lowest overall rank. For more information about the process of ranking and selecting driver packages, see How Windows selects a driver package for a device.

Note

For more information about the driver installation process, see the "Technology review" section of the Step-by-Step Guide to Driver Signing and Staging.

Some physical devices create one or more logical devices when they're installed. Each logical device might handle part of the functionality of the physical device. For example, a multi-function device, such as an all-in-one scanner/fax/printer, might have a different device identification string for each function.

When you use Device Installation policies to allow or prevent the installation of a device that uses logical devices, you must allow or prevent all of the device identification strings for that device. For example, if a user attempts to install a multifunction device and you didn't allow or prevent all of the identification strings for both physical and logical devices, you could get unexpected results from the installation attempt. For more detailed information about hardware IDs, see Device identification strings.

Device setup classes

Device setup classes (also known as *Class*) are another type of identification string. The manufacturer assigns the Class to a device in the driver package. The Class groups devices that are installed and configured in the same way. For example, all Biometric devices belong to the Biometric Class (ClassGuid = {53D29EF7-377C-4D14-864B-EB3A85769359}), and they use the same co-installer when installed. A long number called a globally unique identifier (GUID) represents each device setup class. When Windows starts, it builds an in-memory tree structure with the GUIDs for all of the detected devices. Along with the GUID for the Class of the device itself, Windows may need to insert into the tree the GUID for the Class of the bus to which the device is attached.

When you use device Classes to allow or prevent users from installing drivers, you must specify the GUIDs for all of the device's device setup classes, or you might not achieve the results you want. The installation might fail (if you want it to succeed) or it might succeed (if you want it to fail).

For example, a multi-function device, such as an all-in-one scanner/fax/printer, has a GUID for a generic multi-function device, a GUID for the printer function, a GUID for the scanner function, and so on. The GUIDs for the individual functions are "child nodes" under the multi-function device GUID. To install a child node, Windows must also be able to install the parent node. You must allow installation of the device setup class of the parent GUID for the multi-function device in addition to any child GUIDs for the printer and scanner functions.

For more information, see Device Setup Classes.

This guide doesn't depict any scenarios that use device setup classes. However, the basic principles demonstrated with device identification strings in this guide also apply to device

setup classes. After you discover the device setup class for a specific device, you can then use it in a policy to either allow or prevent installation of drivers for that class of devices.

The following two links provide the complete list of Device Setup Classes. 'System Use' classes are mostly referred to devices that come with a computer/machine from the factory, while 'Vendor' classes are mostly referred to devices that could be connected to an existing computer/machine:

- System-Defined Device Setup Classes Available to Vendors - Windows drivers
- System-Defined Device Setup Classes Reserved for System Use - Windows drivers

'Removable Device' Device type

Some devices could be classified as *Removable Device*. A device is considered *removable* when the driver for the device to which it's connected indicates that the device is removable. For example, a USB device is reported to be removable by the drivers for the USB hub to which the device is connected.

Group Policy Settings for Device Installation

Group Policy is an infrastructure that allows you to specify managed configurations for users and computers through Group Policy settings and Group Policy Preferences.

Device Installation section in Group Policy is a set of policies that control which device could or couldn't be installed on a machine. Whether you want to apply the settings to a stand-alone computer or to many computers in an Active Directory domain, you use the Group Policy Object Editor to configure and apply the policy settings. For more information, see Group Policy Object Editor.

The following passages are brief descriptions of the Device Installation policies that are used in this guide.

Note

Device Installation control is applied only to machines ('computer configuration') and not users ('user configuration') by the nature of the Windows OS design. These policy settings affect all users who log on to the computer where the policy settings are applied. You can't apply these policies to specific users or groups except for the policy Allow administrators to override device installation policy. This policy exempts members of the local Administrators group from any of the device installation restrictions that you apply to the computer by configuring other policy settings as described in this section.

Allow administrators to override Device Installation Restriction policies

This policy setting allows members of the local Administrators group to install and update the drivers for any device, regardless of other policy settings. If you enable this policy setting, administrators can use the Add Hardware Wizard or the Update Driver Wizard to install and update the drivers for any device. If you disable or don't configure this policy setting, administrators are subject to all policy settings that restrict device installation.

Allow installation of devices that match any of these device IDs

This policy setting specifies a list of Plug and Play hardware IDs and compatible IDs that describe devices that users can install. This setting is intended to be used only when the Prevent installation of devices not described by other policy settings policy setting is enabled and doesn't take precedence over any policy setting that would prevent users from installing a device. If you enable this policy setting, users can install and update any device with a hardware ID or compatible ID that matches an ID in this list if that installation hasn't been prevented by the Prevent installation of devices that match these device IDs policy setting, the Prevent installation of devices for these device classes policy setting, or the Prevent installation of removable devices policy setting. If another policy setting prevents users from installing a device, users can't install it even if the device is also described by a value in this policy setting. If you disable or don't configure this policy setting and no other policy describes the device, the Prevent installation of devices not described by other policy settings policy setting determines whether users can install the device.

Allow installation of devices that match any of these device instance IDs

This policy setting allows you to specify a list of Plug and Play device instance IDs for devices that Windows is allowed to install. Use this policy setting only when the "Prevent installation of devices not described by other policy settings" policy setting is enabled. Other policy settings that prevent device installation take precedence over this one. If you enable this policy setting, Windows is allowed to install or update any device whose Plug and Play device instance ID appears in the list you create, unless another policy setting specifically prevents that installation (for example, the "Prevent installation of devices that match any of these device IDs" policy setting, the "Prevent installation of devices for these device classes" policy setting, the "Prevent installation of devices that match any of these device instance IDs" policy setting, or the "Prevent installation of removable devices" policy setting). If you enable this policy setting on a remote desktop server, the policy setting affects redirection of the specified devices from a remote desktop client to the remote desktop server.

Allow installation of devices using drivers that match these device setup classes

This policy setting specifies a list of device setup class GUIDs that describe devices that users can install. This setting is intended to be used only when the Prevent installation of devices not described by other policy settings policy setting is enabled and doesn't take precedence over any policy setting that would prevent users from installing a device. If you enable this setting, users can install and update any device with a hardware ID or compatible ID that matches one of the IDs in this list if that installation hasn't been prevented by the Prevent installation of devices that match these device IDs policy setting, the Prevent installation of devices for these device classes policy setting, or the Prevent installation of removable devices policy setting. If another policy setting prevents users from installing a device, users can't install it even if the device is also described by a value in this policy setting. If you disable or don't configure this policy setting and no other policy setting describes the device, the Prevent installation of devices not described by other policy settings policy setting determines whether users can install the device.

Prevent installation of devices that match these device IDs

This policy setting specifies a list of Plug and Play hardware IDs and compatible IDs for devices that users can't install. If you enable this policy setting, users can't install or update the driver for a device if its hardware ID or compatible ID matches one in this list. If you disable or don't

configure this policy setting, users can install devices and update their drivers, as permitted by other policy settings for device installation. Note: This policy setting takes precedence over any other policy settings that allow users to install a device. This policy setting prevents users from installing a device even if it matches another policy setting that would allow installation of that device.

Prevent installation of devices that match any of these device instance IDs

This policy setting allows you to specify a list of Plug and Play device instance IDs for devices that Windows is prevented from installing. This policy setting takes precedence over any other policy setting that allows Windows to install a device. If you enable this policy setting, Windows is prevented from installing a device whose device instance ID appears in the list you create. If you enable this policy setting on a remote desktop server, the policy setting affects redirection of the specified devices from a remote desktop client to the remote desktop server. If you disable or don't configure this policy setting, devices can be installed and updated as allowed or prevented by other policy settings.

Prevent installation of devices using drivers that match these device setup classes

This policy setting specifies a list of Plug and Play device setup class GUIDs for devices that users can't install. If you enable this policy setting, users can't install or update devices that belong to any of the listed device setup classes. If you disable or don't configure this policy setting, users can install and update devices as permitted by other policy settings for device installation. Note: This policy setting takes precedence over any other policy settings that allow users to install a device. This policy setting prevents users from installing a device from being installed even if it matches another policy setting that would allow installation of that device.

Apply layered order of evaluation for Allow and Prevent device installation policies across all device match criteria

This policy setting changes the evaluation order in which Allow and Prevent policy settings are applied when more than one install policy setting is applicable for a given device. Enable this policy setting to ensure that overlapping device match criteria is applied based on an established hierarchy where more specific match criteria supersedes less specific match criteria. The hierarchical order of evaluation for policy settings that specify device match criteria is as follows:

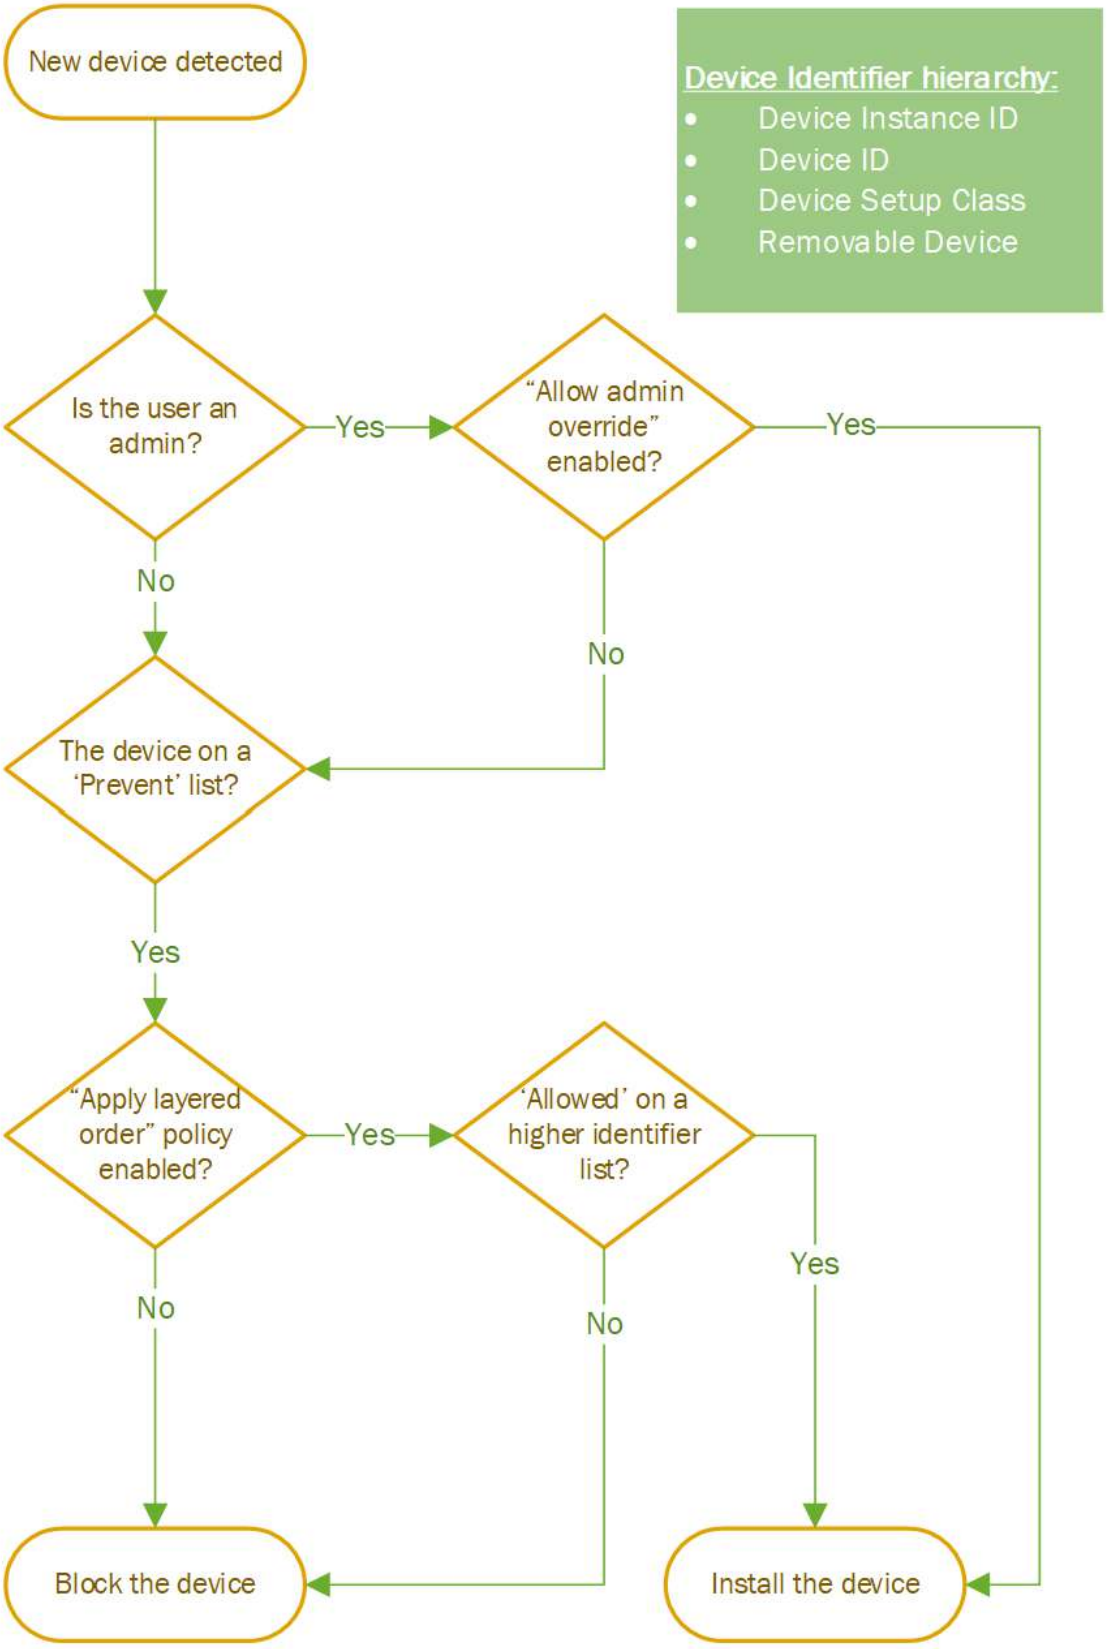
Device instance IDs > **Device IDs** > **Device setup class** > **Removable devices**

Note

This policy setting provides more granular control than the "Prevent installation of devices not described by other policy settings" policy setting. If these conflicting policy settings are enabled at the same time, the "Apply layered order of evaluation for Allow and Prevent device installation policies across all device match criteria" policy setting will be enabled and the other policy setting will be ignored.

If you disable or don't configure this policy setting, the default evaluation is used. By default, all "Prevent installation..." policy settings have precedence over any other policy setting that allows Windows to install a device.

Some of these policies take precedence over other policies. The following flowchart illustrates how Windows processes them to determine whether a user can install a device or not.



Device Installation policies flow chart

Requirements for completing the scenarios

General

To complete each of the scenarios, ensure you have:

- A client computer running Windows.
- A USB thumb drive. The scenarios described in this guide use a USB thumb drive as the example device (also known as a "removable disk drive", "memory drive," a "flash drive," or a "keyring drive"). Most USB thumb drives don't require any manufacturer-provided drivers, and these devices work with the inbox drivers provided with the Windows build.
- A USB/network printer pre-installed on the machine.
- Access to the administrator account on the testing machine. The procedures in this guide require administrator privileges for most steps.

Understanding implications of applying 'Prevent' policies retroactive

All 'Prevent' policies can apply the block functionality to already installed devices-devices that have been installed on the machine before the policy took effect. Using this option is recommended when the administrator isn't sure of the installation history of devices on the machine and would like to make sure the policy applies to all devices.

For example: A printer is already installed on the machine, preventing the installation of all printers will block any future printer from being installed while keeping only the installed printer usable. To apply the block retroactive, the administrator should check mark the "apply this policy to already installed devices" option. Marking this option will prevent access to already installed devices in addition to any future ones.

This option is a powerful tool, but as such it has to be used carefully.

Important

Applying the 'Prevent retroactive' option to crucial devices could render the machine useless/unacceptable! For example: Preventing retroactive all 'Disk Drives' could block the access to the disk on which the OS boots with; Preventing retroactive all 'Net' could block this machine from accessing network and to fix the issue the admin will have to have a direct connection.

Determine device identification strings

By following these steps, you can determine the device identification strings for your device. If the hardware IDs and compatible IDs for your device don't match those IDs shown in this guide, use the IDs that are appropriate to your device (this policy applies to Instance IDs and Classes, but we aren't going to give an example for them in this guide).

You can determine the hardware IDs and compatible IDs for your device in two ways. You can use Device Manager, a graphical tool included with the operating system, or PnPUtil, a command-line tool available for all Windows versions. Use the following procedure to view the device identification strings for your device.

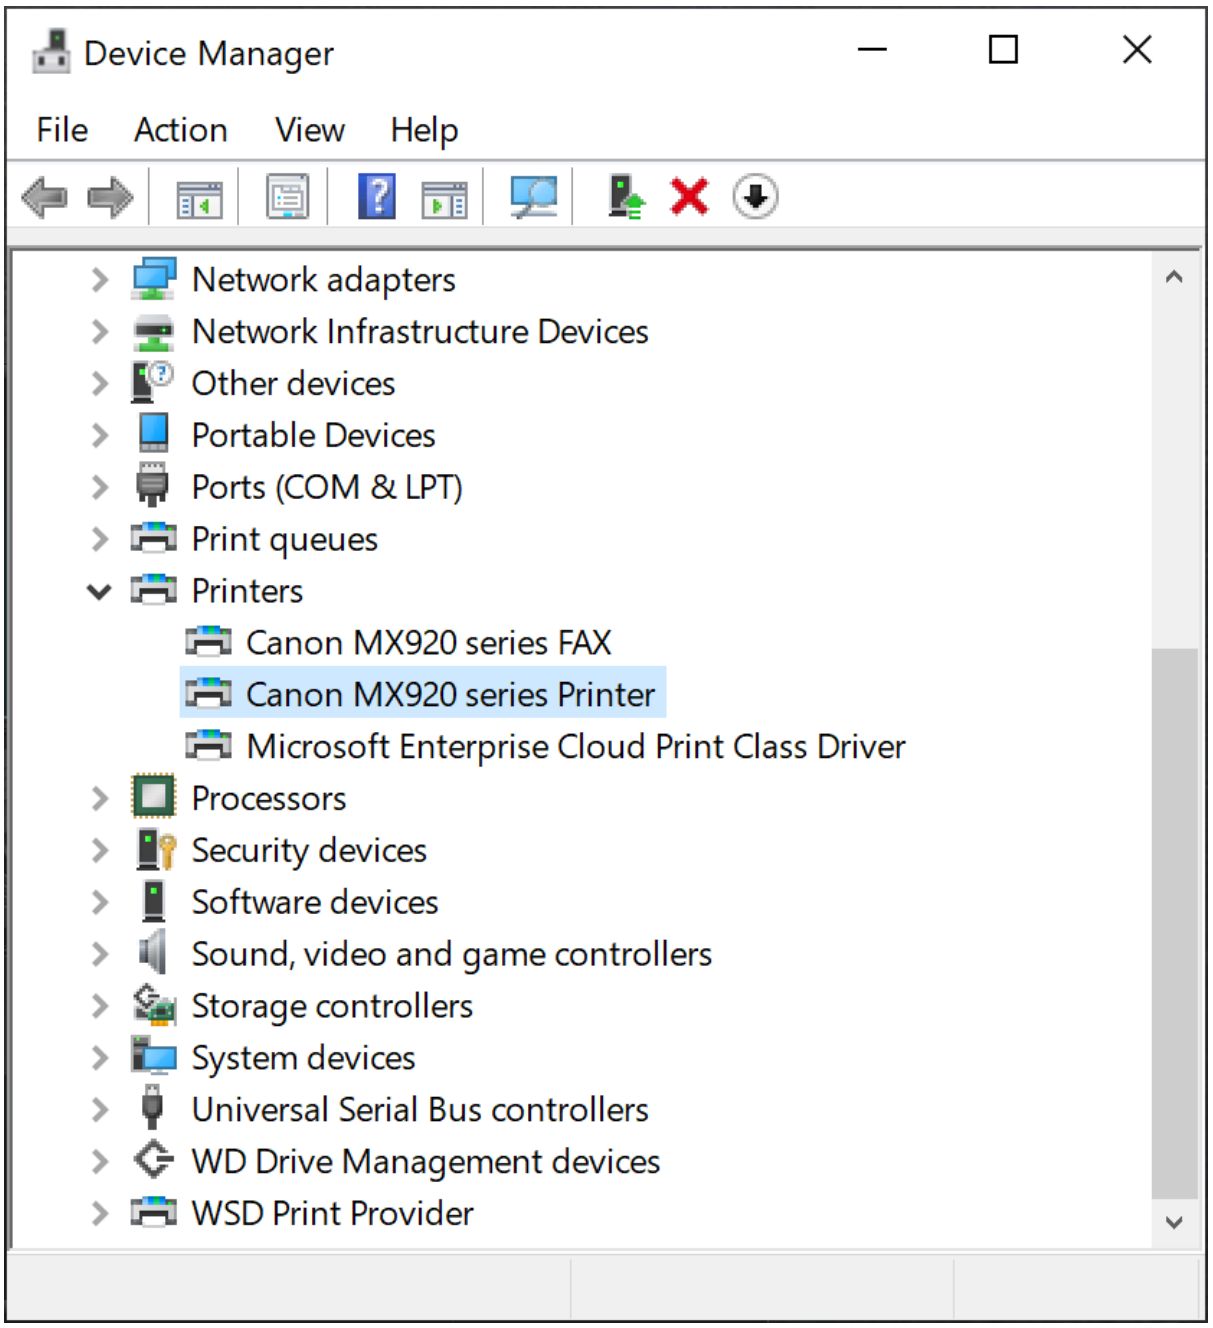
Note

These procedures are specific to a Canon printer. If you are using a different type of device, you must adjust the steps accordingly. The significant difference will be the location of the device in

the Device Manager hierarchy. Instead of being located in the Printers node, you must locate your device in the appropriate node.

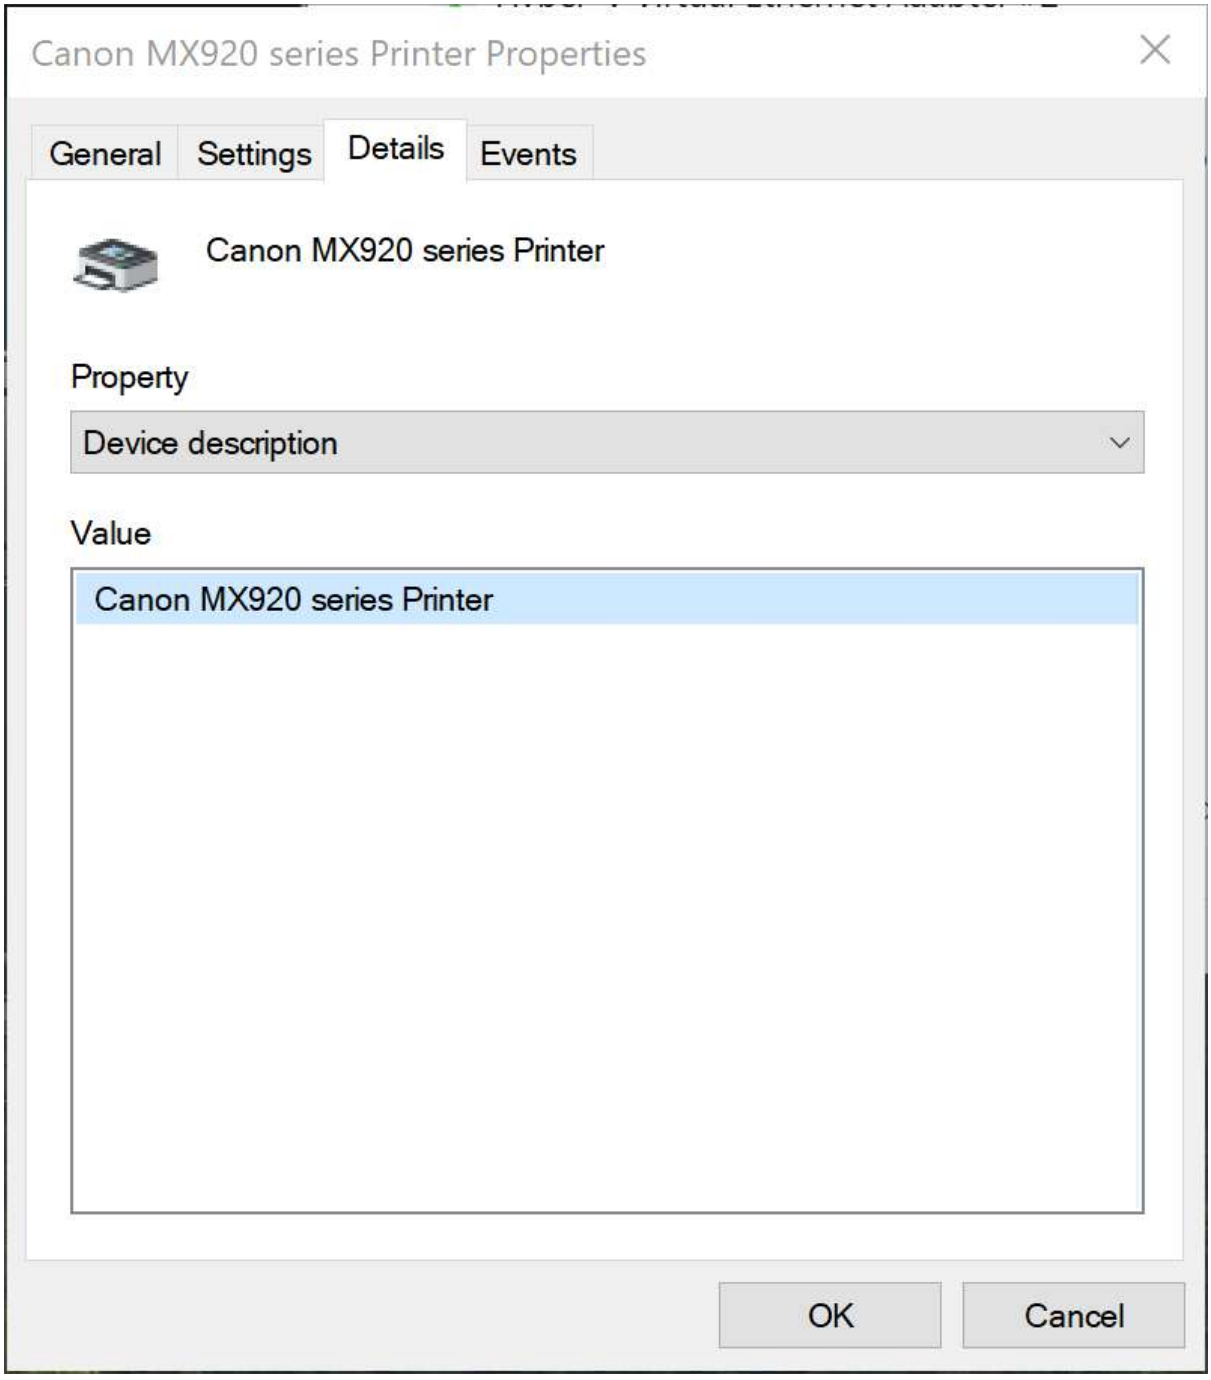
To find device identification strings using Device Manager

1. Make sure your printer is plugged in and installed.
2. To open Device Manager, select the Start button, type mmc devmgmt.msc in the Start Search box, and then press ENTER; or search for Device Manager as application.
3. Device Manager starts and displays a tree representing all of the devices detected on your computer. At the top of the tree is a node with your computers name next to it. Lower nodes represent the various categories of hardware into which your computers devices are grouped.
4. Find the "Printers" section and find the target printer



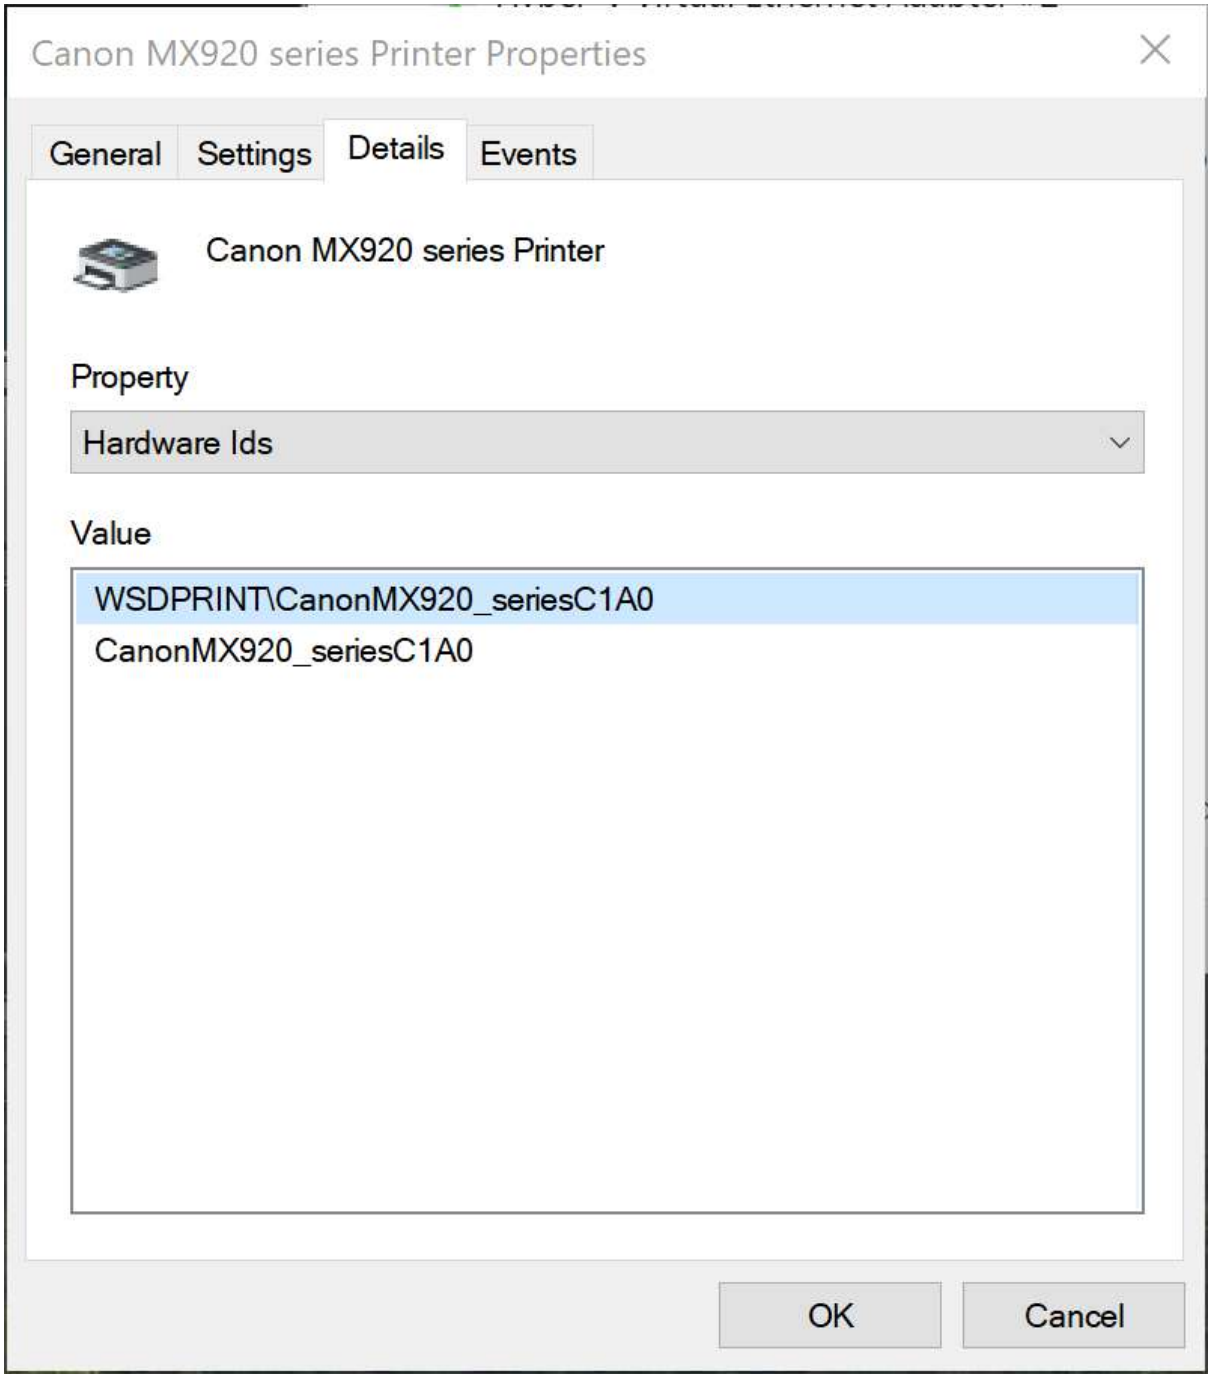
Selecting the printer in Device Manager

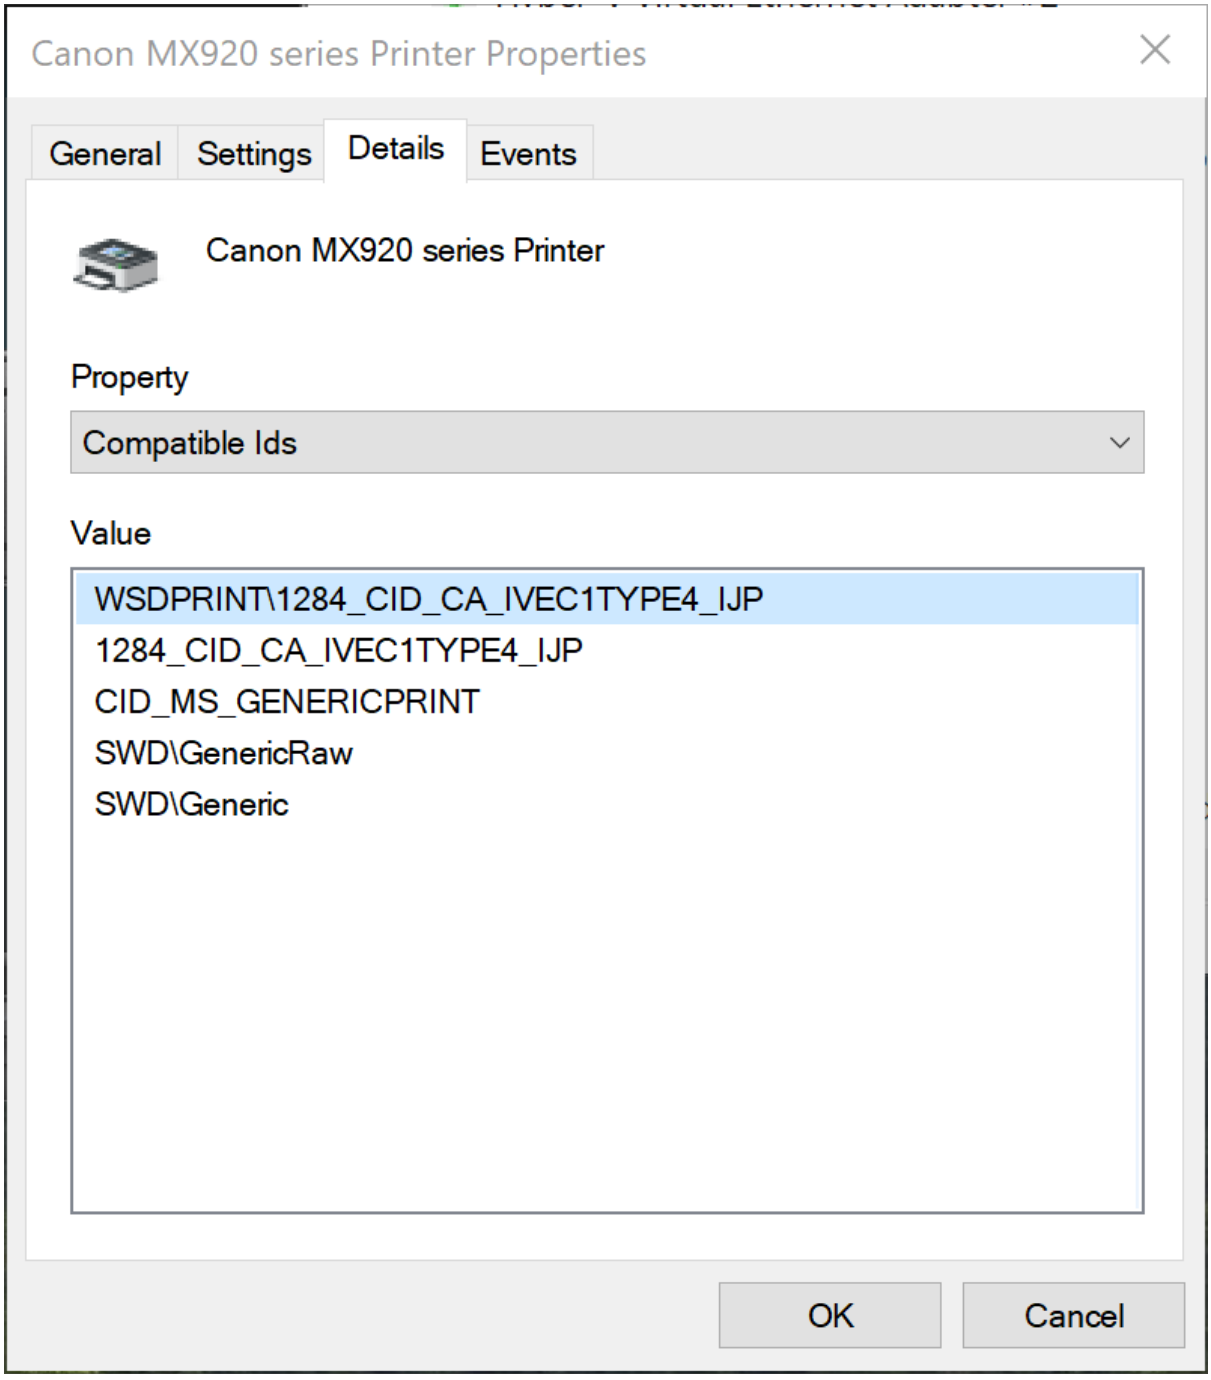
5. Double-click the printer and move to the 'Details' tab.



Open the 'Details' tab to look for the device identifiers

6. From the 'Value' window, copy the most detailed Hardware ID-we'll use this value in the policies.


Canon MX920 series Printer Properties
General
Settings
Details
Events
Canon MX920 series Printer
Property
Hardware Ids
Value
WSDPRINT\CanonMX920_seriesC1A0
CanonMX920_seriesC1A0
OK
Cancel



HWID and Compatible ID

Tip

You can also determine your device identification strings by using the PnPUtil command-line utility. For more information, see **PnPUtil - Windows drivers**.

Getting device identifiers using PnPUtil

ConsoleCopy

pnputil /enum-devices /ids

Here's an example of an output for a single device on a machine:

ConsoleCopy

<snip>

Instance ID:
PCI\VEN_8086&DEV_2F34&SUBSYS_2F348086&REV_02\3&103a9d54&0&81

Device Description: Intel(R) Xeon(R) E7 v3/Xeon(R) E5 v3/Core i7 PCIe Ring Interface - 2F34

Class Name: System

Class GUID: {4d36e97d-e325-11ce-bfc1-08002be10318}

Manufacturer Name: INTEL

Status: Stopped

Driver Name: oem6.inf

Hardware IDs: PCI\VEN_8086&DEV_2F34&SUBSYS_2F348086&REV_02

PCI\VEN_8086&DEV_2F34&SUBSYS_2F348086

PCI\VEN_8086&DEV_2F34&CC_110100

PCI\VEN_8086&DEV_2F34&CC_1101

Compatible IDs: PCI\VEN_8086&DEV_2F34&REV_02

PCI\VEN_8086&DEV_2F34

PCI\VEN_8086&CC_110100

PCI\VEN_8086&CC_1101

PCI\VEN_8086

PCI\CC_110100

PCI\CC_1101

<snip>

Scenario #1: Prevent installation of all printers

In this simple scenario, you'll learn how to prevent the installation of an entire Class of devices.

Setting up the environment

Setting up the environment for the scenario with the following steps:

1. Open Group Policy Editor and navigate to the Device Installation Restriction section.
2. Disable all previous Device Installation policies, except 'Apply layered order of evaluation'-although the policy is disabled in default, this policy is recommended to be enabled in most practical applications.
3. If there are any enabled policies, changing their status to 'disabled', would clear them from all parameters

4. Have a USB/network printer available to test the policy with

Scenario steps - preventing installation of prohibited devices

Getting the right device identifier to prevent it from being installed:

1. If you have on your system a device from the class you want to block, you could follow the steps in the previous section to find the Device Class identifier through Device Manager or PnPUtil (Class GUID).
2. If you don't have such device installed on your system or know the name of the class, you can check the following two links:
    - System-Defined Device Setup Classes Available to Vendors - Windows drivers
    - System-Defined Device Setup Classes Reserved for System Use - Windows drivers
3. Our current scenario is focused on preventing all printers from being installed, as such here's the Class GUID for most of printers in the market:

Printers
Class = Printer
ClassGuid = {4d36e979-e325-11ce-bfc1-08002be10318}
This class includes printers.

Note

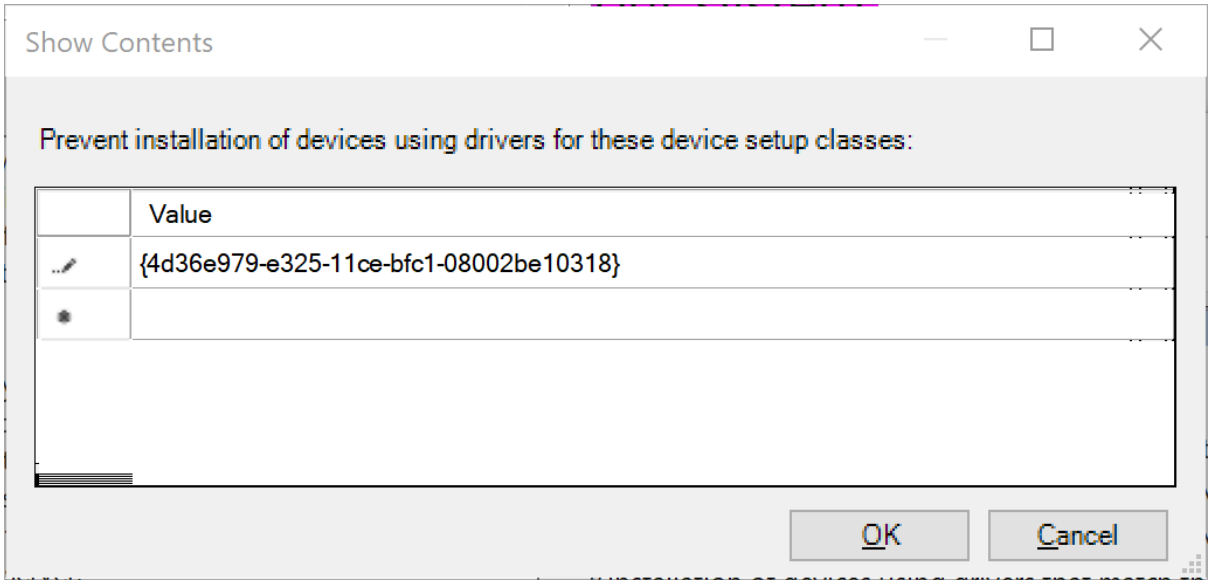
As mentioned before, preventing an entire Class could block you from using your system completely. Please make sure you understand which devices are going to be blocked when specifying a Class. For our scenario, there are other classes that relate to printers but before you apply them, make sure they're not blocking any other existing device that is crucial to your system.

Creating the policy to prevent all printers from being installed:

1. Open Group Policy Object Editor-either click the Start button, type mmc gpedit.msc in the Start Search box, and then press ENTER; or type in the Windows search "Group Policy Editor" and open the UI.
2. Navigate to the Device Installation Restriction page:

Computer Configuration > Administrative Templates > System > Device Installation > Device Installation Restrictions

3. Make sure all policies are disabled (recommended to keep 'applied layered order of evaluation' policy enabled).
4. Open **Prevent installation of devices using drivers that match these device setup classes** policy and select the 'Enable' radio button.
5. In the lower left side, in the 'Options' window, click the 'Show...' box. This option takes you to a table where you can enter the class identifier to block.
6. Enter the printer class GUID you found with the curly braces: {4d36e979-e325-11ce-bfc1-08002be10318}.



List of prevent Class GUIDs

7. Click 'OK'.
8. Click 'Apply' on the bottom right of the policy's window-this option pushes the policy and blocks all future printer installations, but doesn't apply to existing installs.
9. Optional-if you would like to apply the policy to existing installs: Open the **Prevent installation of devices using drivers that match these device setup classes** policy again; in the 'Options' window mark the checkbox that says 'also apply to matching devices that are already installed'

Important

Using a Prevent policy (like the one we used in scenario #1 above) and applying it to all previously installed devices (see step #9) could render crucial devices unusable; hence, use with caution. For example: If an IT admin wants to prevent all removable storage devices from being installed on the machine, using 'Disk Drive' class for blocking and applying it retroactive could render the internal hard-drive unusable and to break the machine.

Testing scenario 1

1. If you haven't completed step #9, follow these steps:

a. Uninstall your printer: Device Manager > Printers > right click the Canon Printer > click "Uninstall device".

b. For USB printer-unplug and plug back the cable; for network device-make a search for the printer in the Windows Settings app.

c. You shouldn't be able to reinstall the printer.

2. If you completed step #9 above and restarted the machine, look for your printer under Device Manager or the Windows Settings app and see that it's no-longer available for you to use.

Scenario #2: Prevent installation of a specific printer

This scenario builds upon scenario #1, Prevent installation of all printers. In this scenario, you target a specific printer to prevent from being installed on the machine.

Setting up the environment

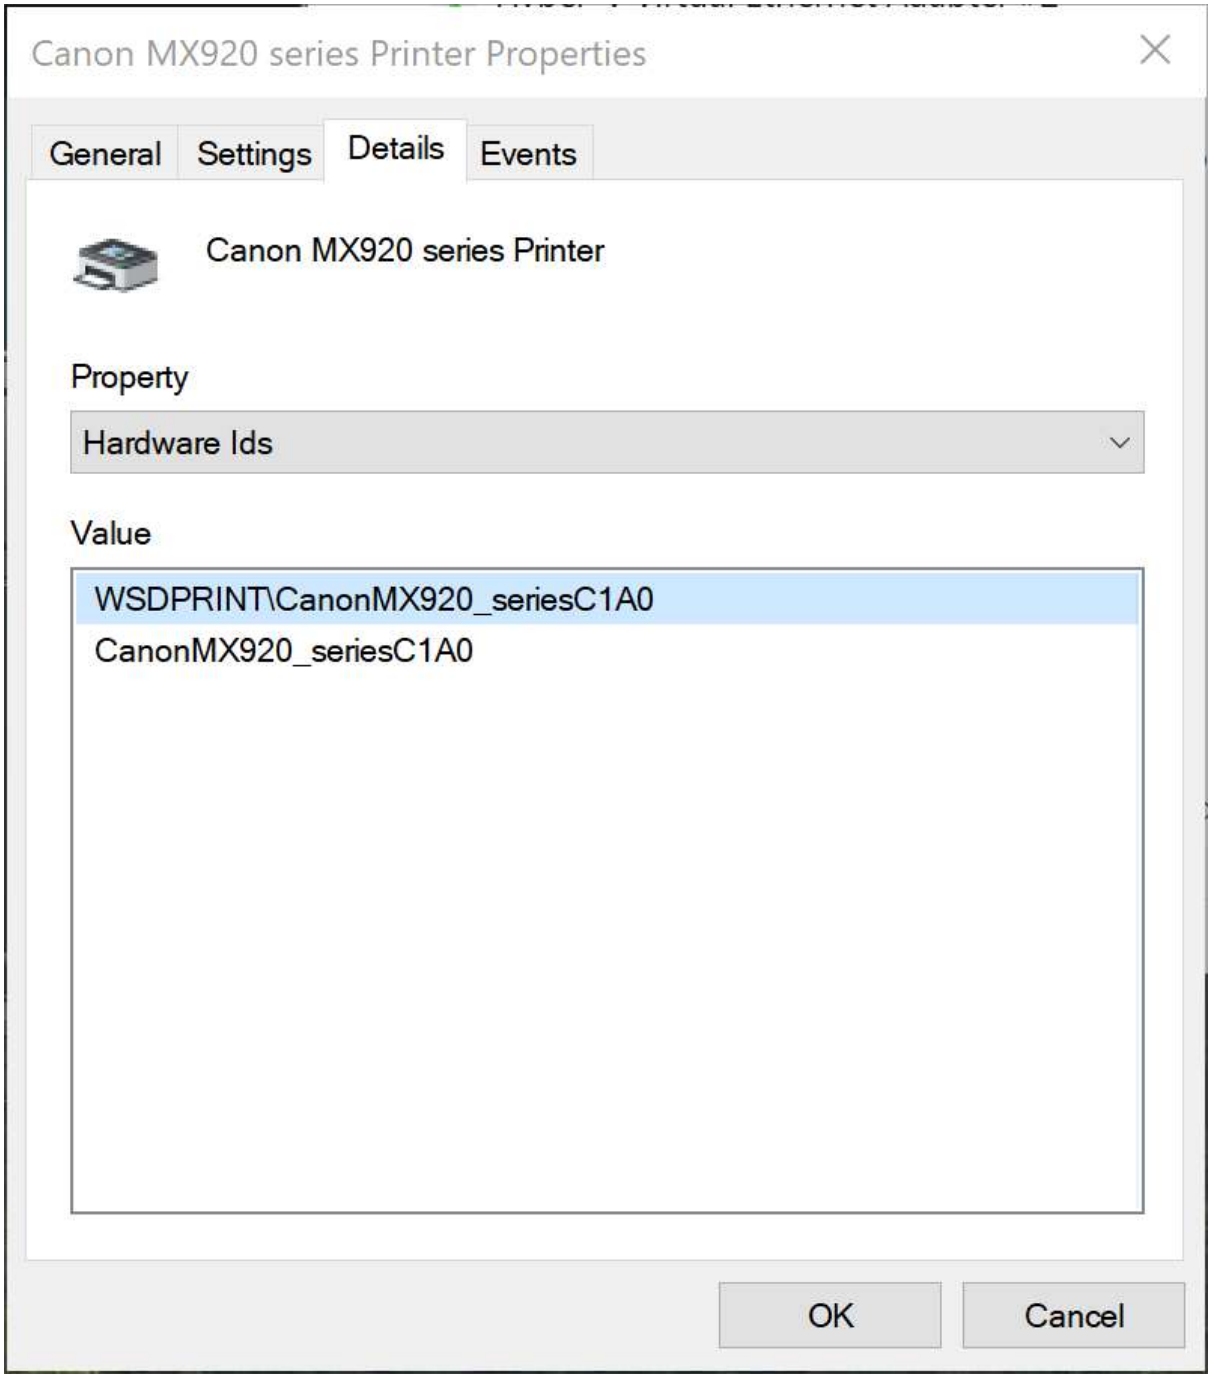
Setting up the environment for the scenario with the following steps:

1. Open Group Policy Editor and navigate to the Device Installation Restriction section.
2. Ensure all previous Device Installation policies are disabled except 'Apply layered order of evaluation' (this prerequisite is optional to be On/Off this scenario). Although the policy is disabled in default, it's recommended to be enabled in most practical applications. For scenario #2, it's optional.

Scenario steps - preventing installation of a specific device

Getting the right device identifier to prevent it from being installed:

1. Get your printer's Hardware ID. In this example we'll use the identifier we found previously.



Printer Hardware ID

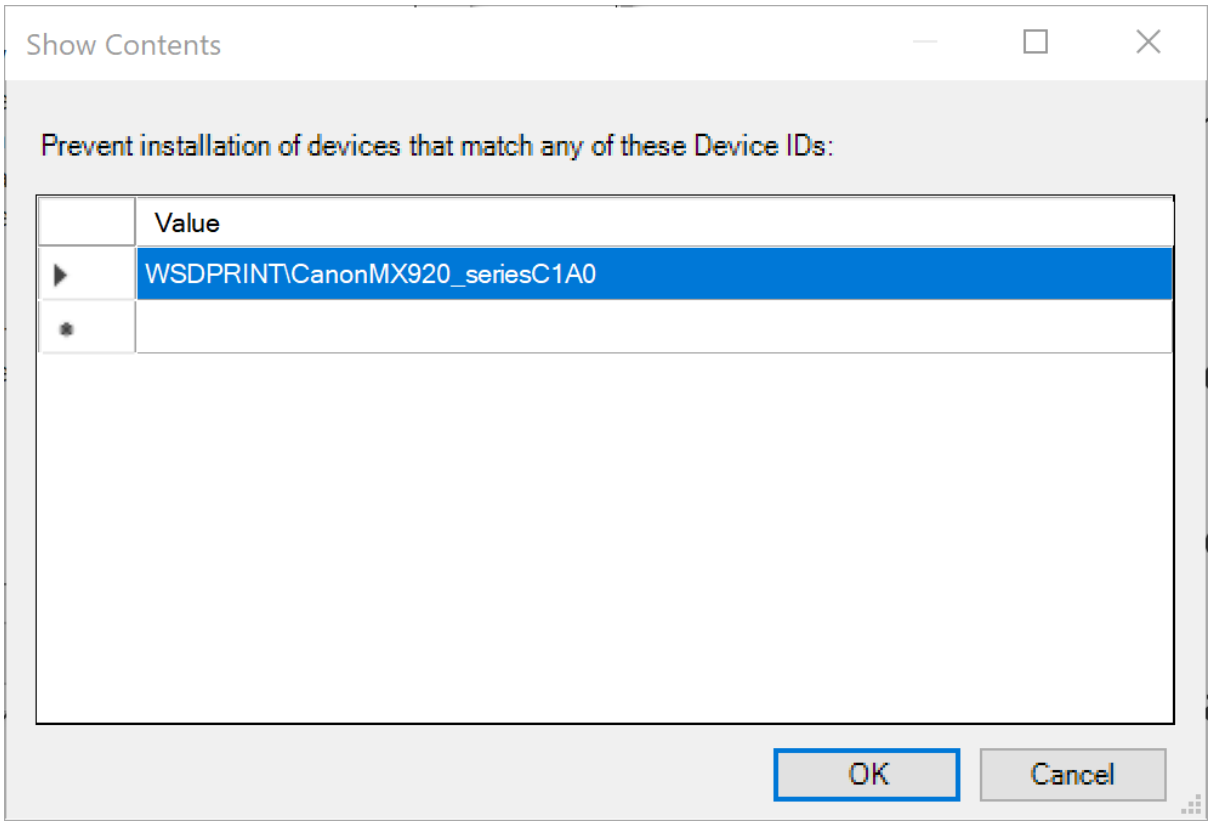
2. Write down the device ID (in this case Hardware ID): WSDPRINT\CanonMX920_seriesC1A0;. Take the more specific identifier to make sure you block a specific printer and not a family of printers

Creating the policy to prevent a single printer from being installed:

1. Open Group Policy Object Editor.
2. Navigate to the Device Installation Restriction page:

Computer Configuration > Administrative Templates > System > Device Installation > Device Installation Restrictions

3. Open **Prevent installation of devices that match any of these device IDs** policy and select the 'Enable' radio button.
4. In the lower left side, in the 'Options' window, click the 'Show...' box. This option will take you to a table where you can enter the device identifier to block.
5. Enter the printer device ID you found above: WSDPRINT\CanonMX920_seriesC1A0.



Prevent Device ID list

6. Click 'OK'.
7. Click 'Apply' on the bottom right of the policy's window. This option pushes the policy and blocks the target printer in future installations, but doesn't apply to an existing install.
8. Optionally, if you would like to apply the policy to an existing install, open the **Prevent installation of devices that match any of these device IDs** policy again. In the 'Options' window, mark the checkbox that says 'Also apply to matching devices that are already installed'.

Testing scenario 2

If you completed step #8 above and restarted the machine, look for your printer under Device Manager or the Windows Settings app and see that it's no-longer available for you to use.

If you haven't completed step #8, follow these steps:

1. Uninstall your printer: Device Manager > Printers > right click the Canon Printer > click "Uninstall device".

2. For USB printer, unplug and plug back the cable; for network device, make a search for the printer in the Windows Settings app.
3. You shouldn't be able to reinstall the printer.

Scenario #3: Prevent installation of all printers while allowing a specific printer to be installed

Now, using the knowledge from both previous scenarios, you'll learn how to prevent the installation of an entire Class of devices while allowing a single printer to be installed.

Setting up the environment

Setting up the environment for the scenario with the following steps:

1. Open Group Policy Editor and navigate to the Device Installation Restriction section.
2. Disable all previous Device Installation policies, and enable 'Apply layered order of evaluation'.
3. If there are any enabled policies, changing their status to 'disabled', would clear them from all parameters.
4. Have a USB/network printer available to test the policy with.

Scenario steps - preventing installation of an entire class while allowing a specific printer

Getting the device identifier for both the Printer Class and a specific printer-following the steps in scenario #1 to find Class identifier and scenario #2 to find Device identifier you could get the identifiers you need for this scenario:

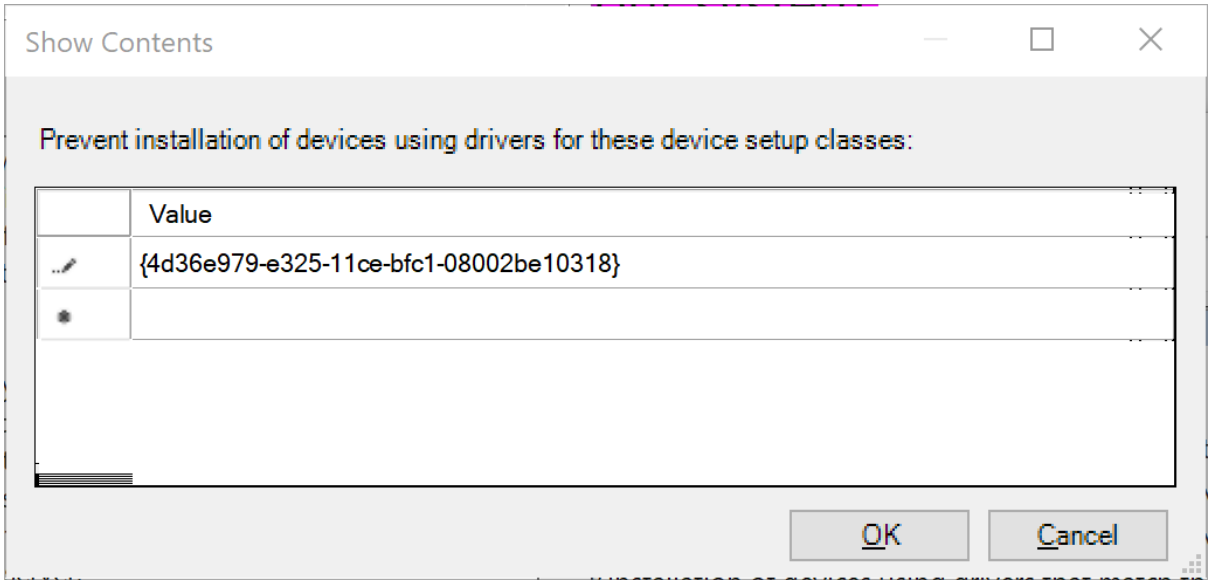
- ClassGuid = {4d36e979-e325-11ce-bfc1-08002be10318}
- Hardware ID = WSDPRINT\CanonMX920_seriesC1A0

First create a 'Prevent Class' policy and then create 'Allow Device' one:

1. Open Group Policy Object Editor-either click the Start button, type mmc gpedit.msc in the Start Search box, and then press ENTER; or type in the Windows search "Group Policy Editor" and open the UI.
2. Navigate to the Device Installation Restriction page:

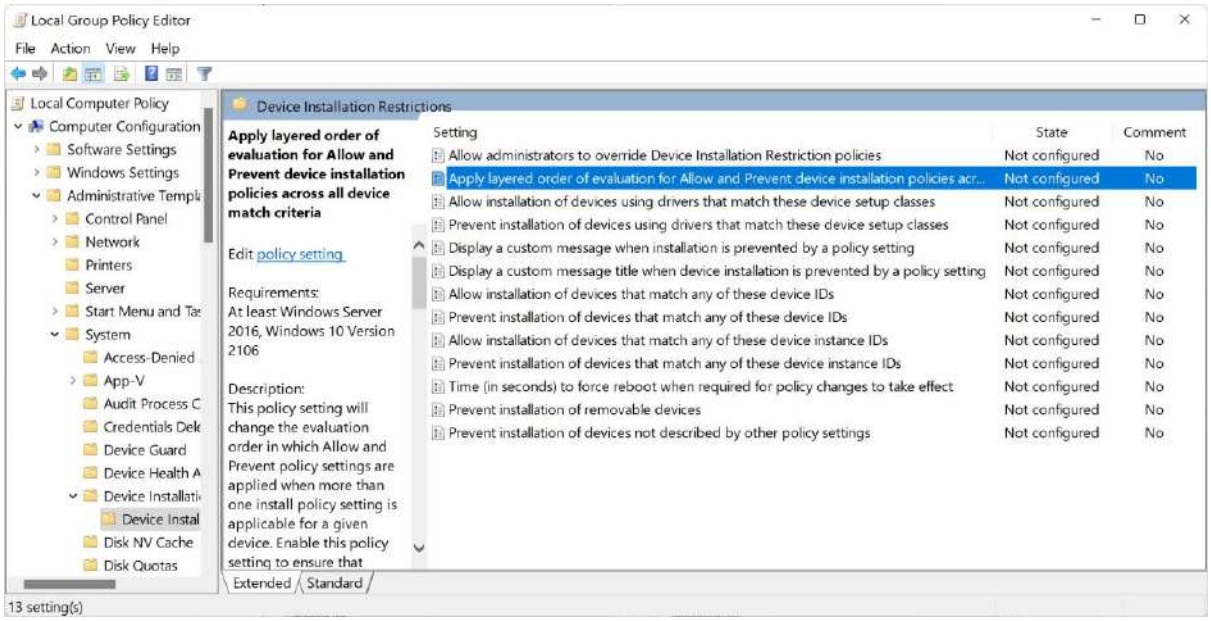
Computer Configuration > Administrative Templates > System > Device Installation > Device Installation Restrictions

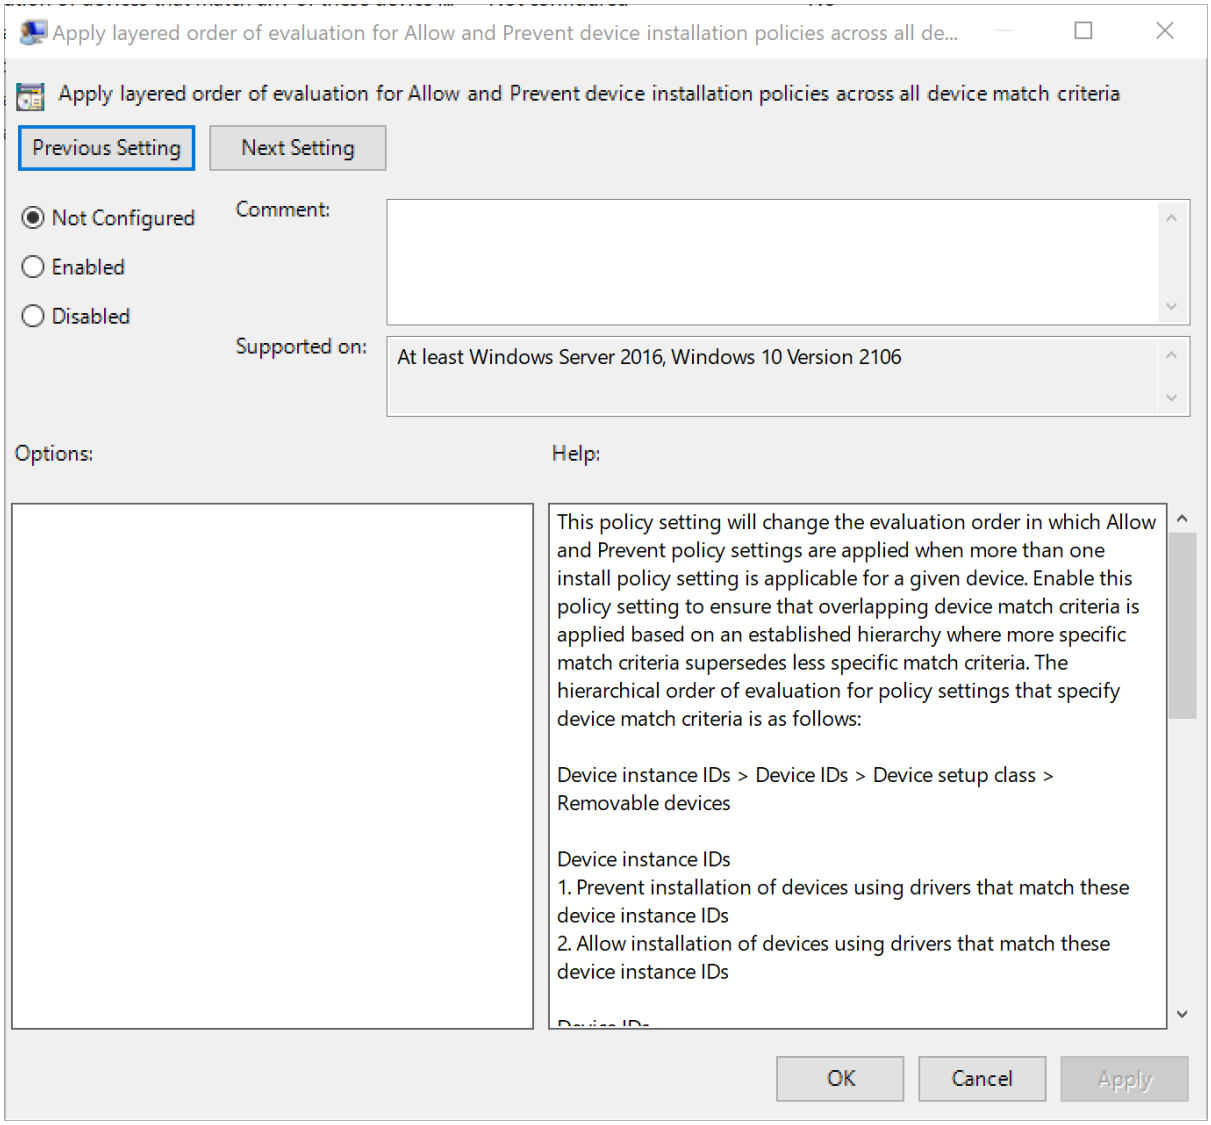
3. Make sure all policies are disabled
4. Open **Prevent installation of devices using drivers that match these device setup classes** policy and select the 'Enable' radio button.
5. In the lower left side, in the 'Options' window, click the 'Show...' box. This option will take you to a table where you can enter the class identifier to block.
6. Enter the printer class GUID you found above with the curly braces (this value is important! Otherwise, it won't work): {4d36e979-e325-11ce-bfc1-08002be10318}



List of prevent Class GUIDs

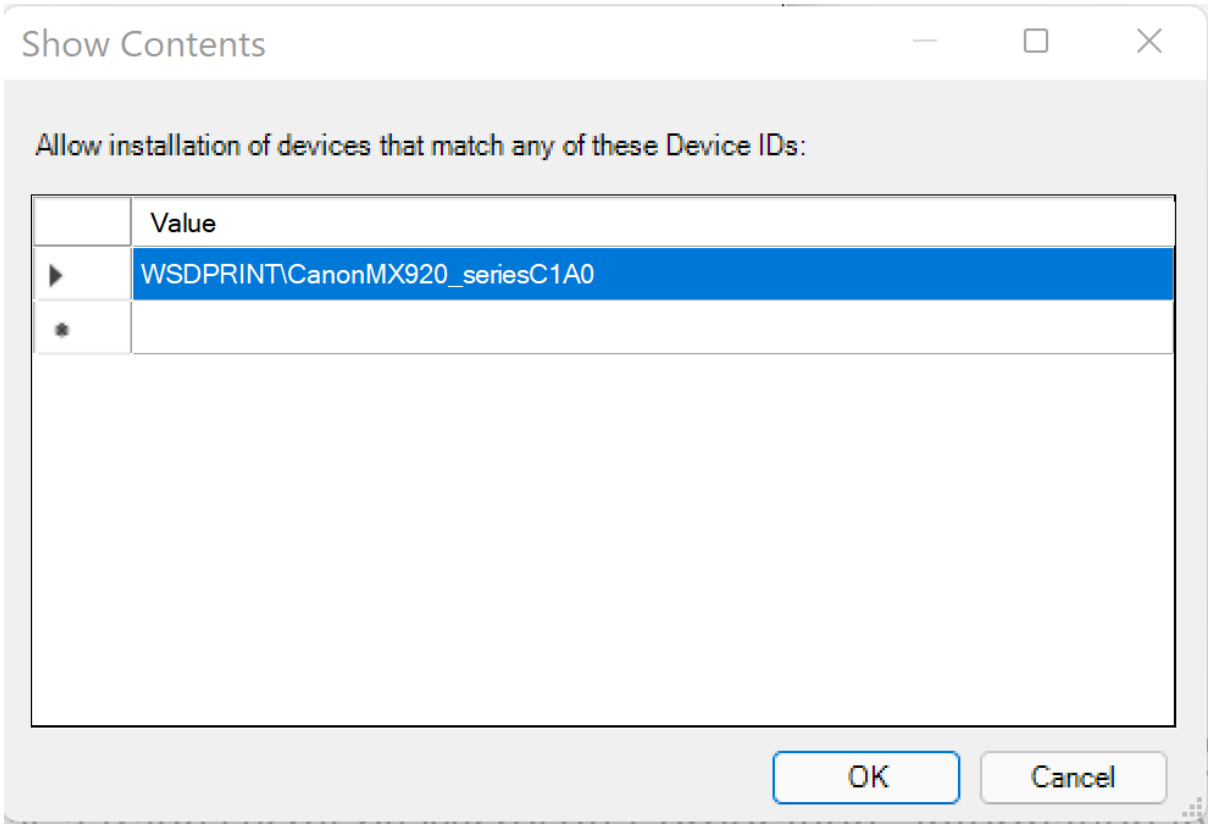
7. Click 'OK'.
8. Click 'Apply' on the bottom right of the policy's window-this option pushes the policy and blocks all future printer installations, but doesn't apply to existing installs.
9. To complete the coverage of all future and existing printers, open the **Prevent installation of devices using drivers that match these device setup classes** policy again; in the 'Options' window mark the checkbox that says 'also apply to matching devices that are already installed' and click 'OK'
10. Open the **Apply layered order of evaluation for Allow and Prevent device installation policies across all device match criteria** policy and enable it-this policy will enable you to override the wide coverage of the 'Prevent' policy with a specific device.





Apply layered order of evaluation policy

11. Now Open **Allow installation of devices that match any of these device IDs** policy and select the 'Enable' radio button.
12. In the lower left side, in the 'Options' window, click the 'Show...' box. This option will take you to a table where you can enter the device identifier to allow.
13. Enter the printer device ID you found above: WSDPRINT\CanonMX920_seriesC1A0.



Allow Printer Hardware ID

14. Click 'OK'.
15. Click 'Apply' on the bottom right of the policy's window-this option pushes the policy and allows the target printer to be installed (or stayed installed).

Testing scenario 3

1. Look for your printer under Device Manager or the Windows Settings app and see that it's still there and accessible. Or just print a test document.
2. Go back to the Group Policy Editor, disable **Apply layered order of evaluation for Allow and Prevent device installation policies across all device match criteria** policy and test again your printer-you shouldn't be bale to print anything or able to access the printer at all.

Scenario #4: Prevent installation of a specific USB device

The scenario builds upon the knowledge from scenario #2, Prevent installation of a specific printer. In this scenario, you'll gain an understanding of how some devices are built into the PnP (Plug and Play) device tree.

Setting up the environment

Setting up the environment for the scenario with the following steps:

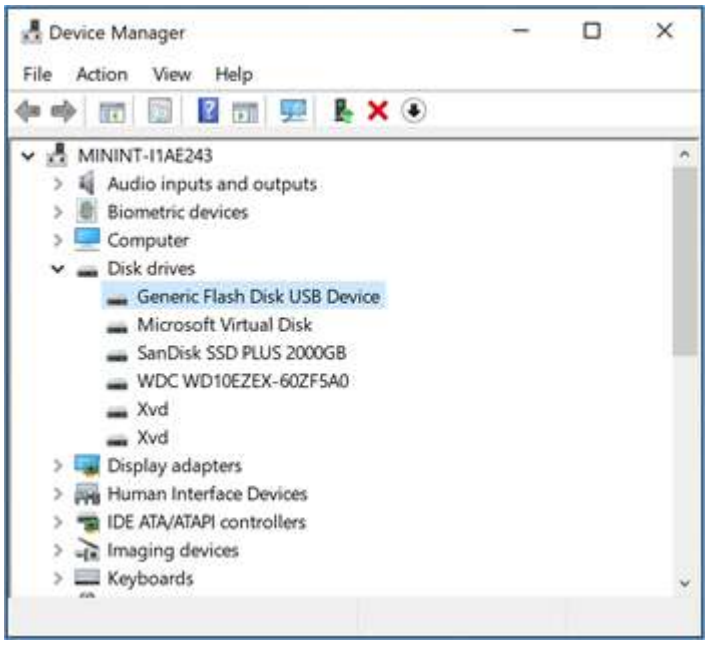
1. Open Group Policy Editor and navigate to the Device Installation Restriction section
2. Ensure all previous Device Installation policies are disabled except 'Apply layered order of evaluation'. This prerequisite is optional to be On/Off this scenario. Although the

policy is disabled in default, it's recommended to be enabled in most practical applications.

Scenario steps - preventing installation of a specific device

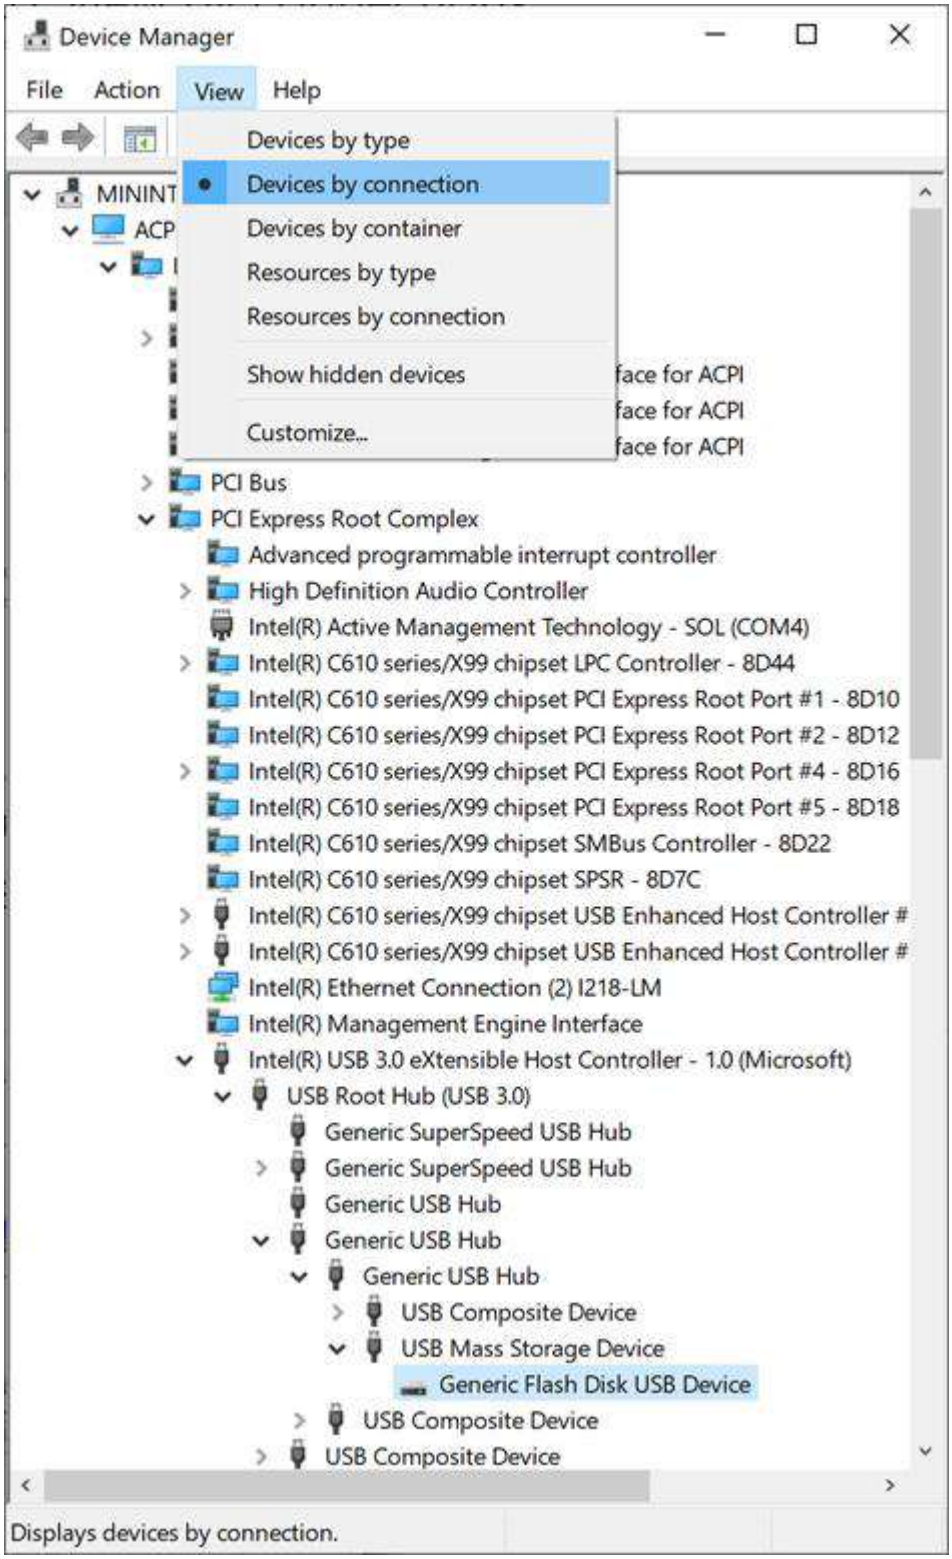
Getting the right device identifier to prevent it from being installed and its location in the PnP tree:

1. Connect a USB thumb drive to the machine
2. Open Device Manager
3. Find the USB thumb-drive and select it.



Selecting the usb thumb-drive in Device Manager

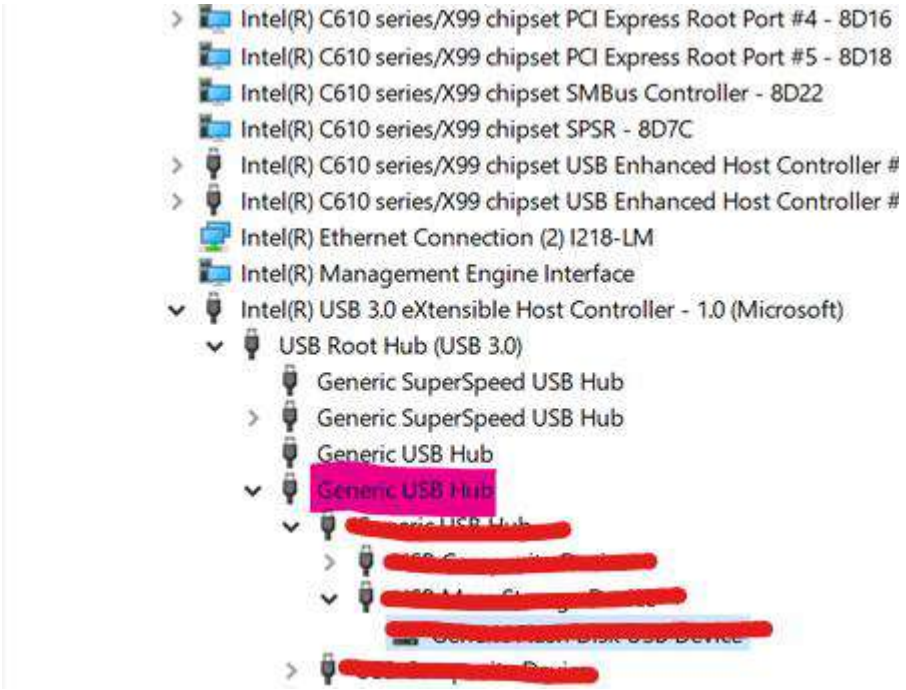
4. Change View (in the top menu) to 'Devices by connections'. This view represents the way devices are installed in the PnP tree.



Changing view in Device Manager to see the PnP connection tree

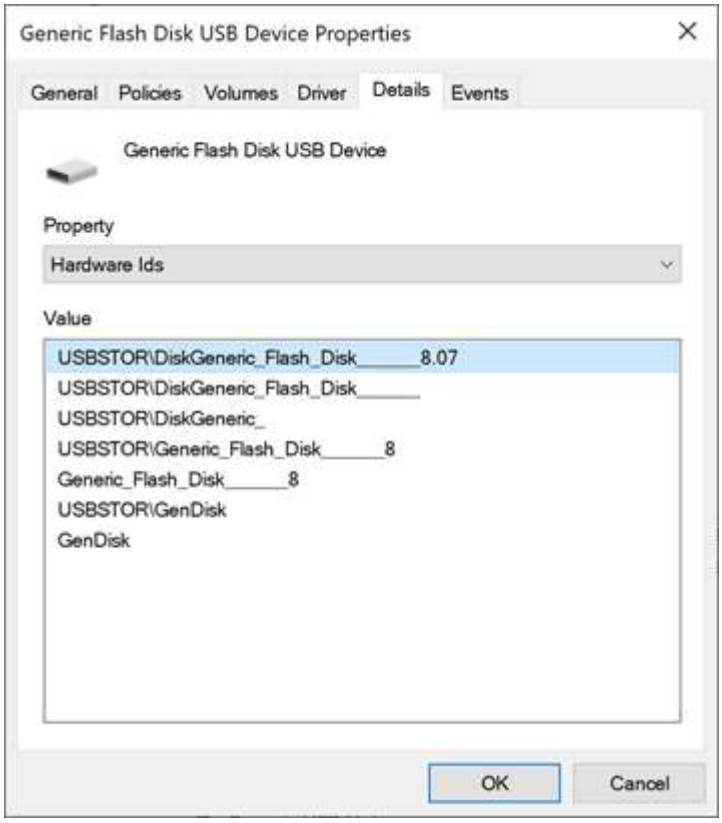
Note

When blocking\Preventing a device that sits higher in the PnP tree, all the devices that sit under it will be blocked. For example: Preventing a "Generic USB Hub" from being installed, all the devices that lay below a "Generic USB Hub" will be blocked.



When blocking one device, all the devices that are nested below it will be blocked as well

5. Double-click the USB thumb-drive and move to the 'Details' tab.
6. From the 'Value' window, copy the most detailed Hardware ID-we'll use this value in the policies. In this case Device ID = USBSTOR\DiskGeneric_Flash_Disk_____8.07



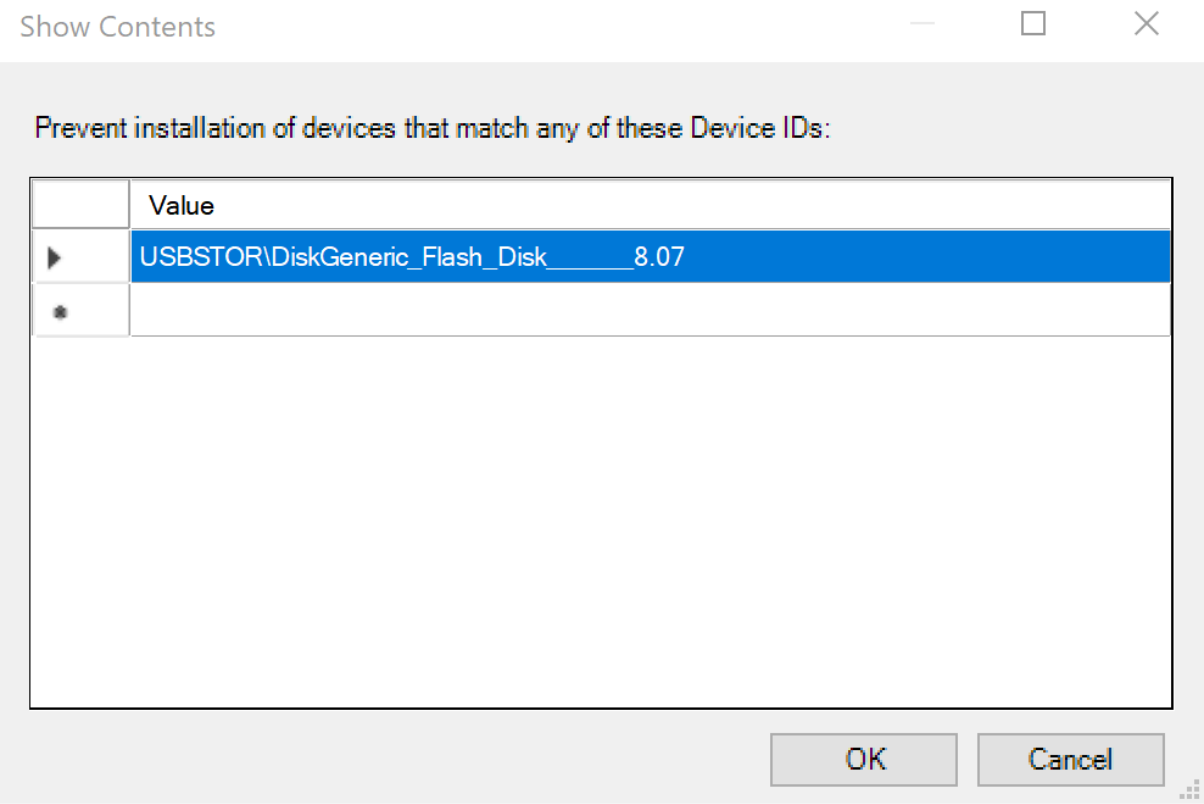
USB device hardware IDs

Creating the policy to prevent a single USB thumb-drive from being installed:

1. Open Group Policy Object Editor and either click the Start button, type mmc gpedit.msc in the Start Search box, and then press ENTER; or type in the Windows search "Group Policy Editor" and open the UI.
2. Navigate to the Device Installation Restriction page:

Computer Configuration > Administrative Templates > System > Device Installation > Device Installation Restrictions

3. Open **Prevent installation of devices that match any of these device IDs** policy and select the 'Enable' radio button.
4. In the lower left side, in the 'Options' window, click the 'Show' box. This option will take you to a table where you can enter the device identifier to block.
5. Enter the USB thumb-drive device ID you found above- USBSTOR\DiskGeneric_Flash_Disk_____8.07.



Prevent Device IDs list

6. Click 'OK'.
7. Click 'Apply' on the bottom right of the policy's window. This option pushes the policy and blocks the target USB thumb-drive in future installations, but doesn't apply to an existing install.
8. Optional - if you would like to apply the policy to an existing install, open the **Prevent installation of devices that match any of these device IDs** policy again. In the 'Options' window, mark the checkbox that says 'also apply to matching devices that are already installed'.

Testing scenario 4

1. If you haven't completed step #8, follow these steps:
    - Uninstall your USB thumb-drive: Device Manager > Disk drives > right click the target USB thumb-drive > click "Uninstall device".
    - You shouldn't be able to reinstall the device.
2. If you completed step #8 above and restarted the machine, look for your Disk drives under Device Manager and see that it's no-longer available for you to use.

Scenario #5: Prevent installation of all USB devices while allowing an installation of only an authorized USB thumb-drive

Now, using the knowledge from all the previous four scenarios, you'll learn how to prevent the installation of an entire Class of devices while allowing a single authorized USB thumb-drive to be installed.

Setting up the environment

Setting up the environment for the scenario with the following steps:

1. Open Group Policy Editor and navigate to the Device Installation Restriction section.
2. Disable all previous Device Installation policies, and **enable** 'Apply layered order of evaluation'.
3. If there are any enabled policies, changing their status to 'disabled', would clear them from all parameters.
4. Have a USB thumb-drive available to test the policy with.

Scenario steps - preventing installation of all USB devices while allowing only an authorized USB thumb-drive

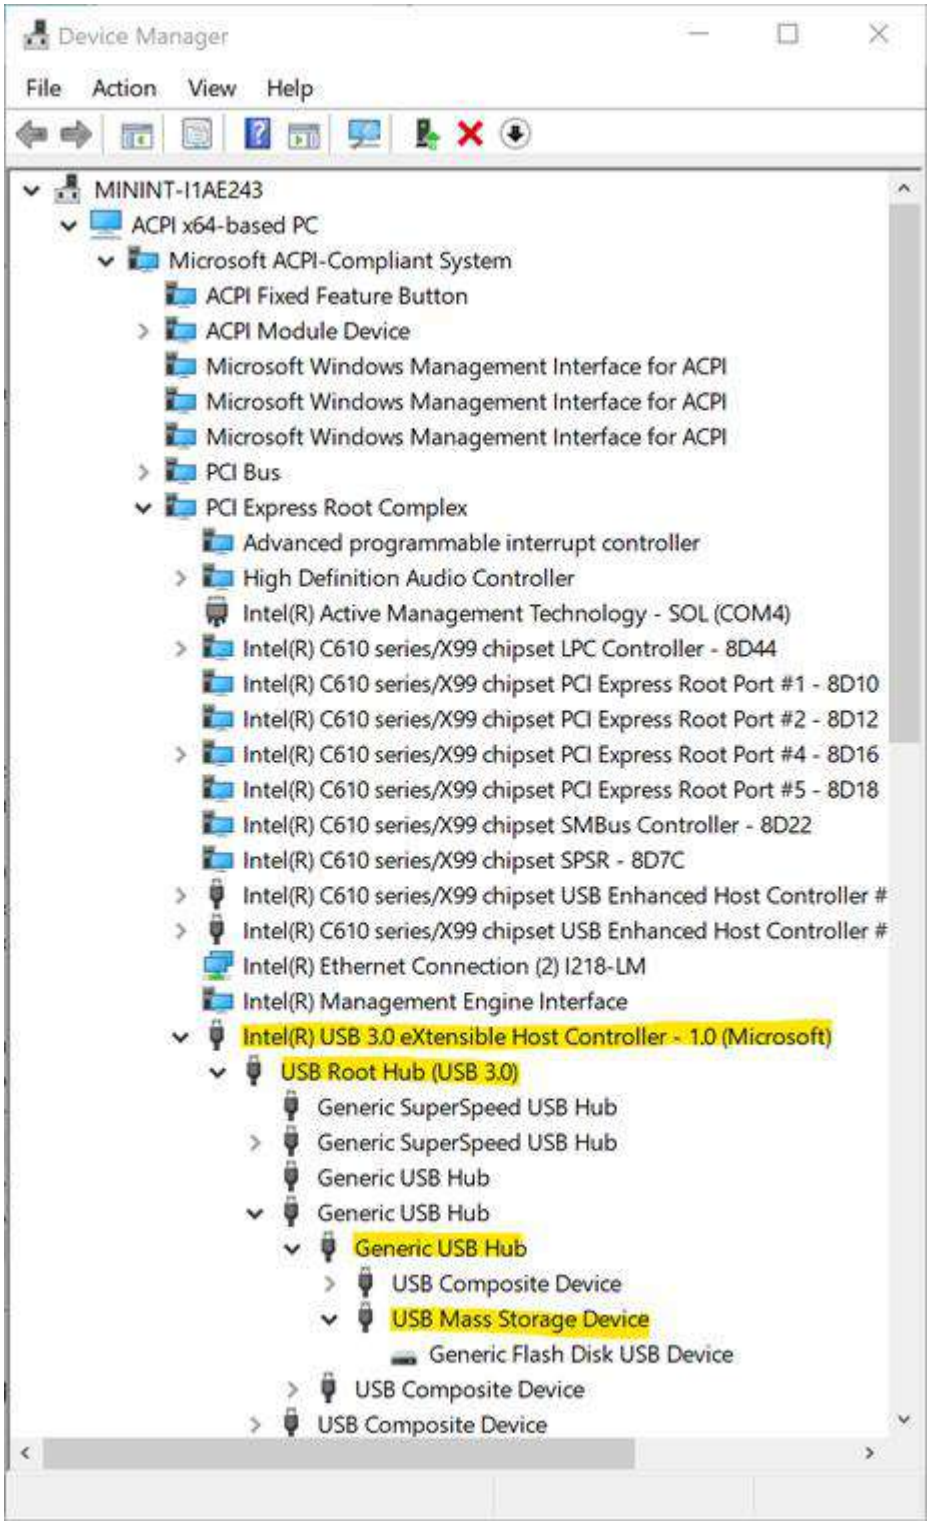
Getting the device identifier for both the USB Classes and a specific USB thumb-drive and following the steps in scenario #1 to find Class identifier and scenario #4 to find Device identifier you could get the identifiers you need for this scenario:

- USB Bus Devices (hubs and host controllers)
    - Class = USB
    - ClassGuid = {36fc9e60-c465-11cf-8056-444553540000}
    - This class includes USB host controllers and USB hubs, but not USB peripherals. Drivers for this class are system-supplied.
- USB Device
    - Class = USBDevice
    - ClassGuid = {88BAE032-5A81-49f0-BC3D-A4FF138216D6}
    - USBDevice includes all USB devices that don't belong to another class. This class isn't used for USB host controllers and hubs.

- Hardware ID = USBSTOR\DiskGeneric_Flash_Disk_____8.07

As mentioned in scenario #4, it's not enough to enable only a single hardware ID in order to enable a single USB thumb-drive. The IT admin has to ensure all the USB devices that preceding the target one aren't blocked (allowed) as well. In Our case the following devices has to be allowed so the target USB thumb-drive could be allowed as well:

- "Intel(R) USB 3.0 eXtensible Host Controller - 1.0 (Microsoft)" -> PCI\CC_0C03
- "USB Root Hub (USB 3.0)" -> USB\ROOT_HUB30
- "Generic USB Hub" -> USB\USB20_HUB



USB devices nested under each other in the PnP tree

These devices are internal devices on the machine that define the USB port connection to the outside world. Enabling them shouldn't enable any external/peripheral device from being installed on the machine.

Important

Some device in the system have several layers of connectivity to define their installation on the system. USB thumb-drives are such devices. Thus, when looking to either block or allow them

on a system, it's important to understand the path of connectivity for each device. There are several generic Device IDs that are commonly used in systems and could provide a good start to build an 'Allow list' in such cases. See below for the list:

PCI\CC_0C03; PCI\CC_0C0330; PCI\VEN_8086; PNP0CA1; PNP0CA1&HOST (for Host Controllers)/ USB\ROOT_HUB30; USB\ROOT_HUB20 (for USB Root Hubs)/ USB\USB20_HUB (for Generic USB Hubs)/

Specifically for desktop machines, it's very important to list all the USB devices that your keyboards and mice are connected through in the above list. Failing to do so could block a user from accessing its machine through HID devices.

Different PC manufacturers sometimes have different ways to nest USB devices in the PnP tree, but in general this is how it's done.

First create a 'Prevent Class' policy and then create 'Allow Device' one:

1. Open Group Policy Object Editor: either click the Start button, type mmc gpedit.msc in the Start Search box, and then press ENTER; or type in the Windows search "Group Policy Editor" and open the UI.
2. Navigate to the Device Installation Restriction page:

Computer Configuration > Administrative Templates > System > Device Installation > Device Installation Restrictions

3. Make sure all policies are disabled
4. Open **Prevent installation of devices using drivers that match these device setup classes** policy and select the 'Enable' radio button.
5. In the lower left side, in the 'Options' window, click the 'Show...' box. This option will take you to a table where you can enter the class identifier to block.
6. Enter both USB classes GUID you found above with the curly braces:

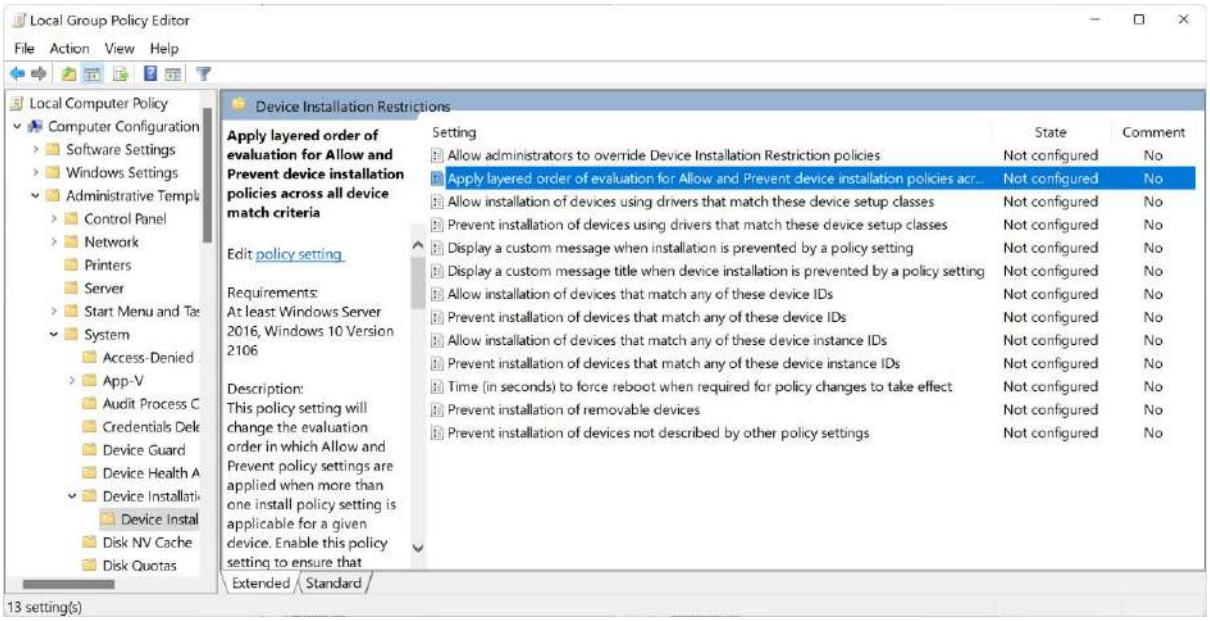
{36fc9e60-c465-11cf-8056-444553540000}/ {88BAE032-5A81-49f0-BC3D-A4FF138216D6}

7. Click 'OK'.
8. Click 'Apply' on the bottom right of the policy's window. This option pushes the policy and blocks all future USB device installations, but doesn't apply to existing installs.

Important

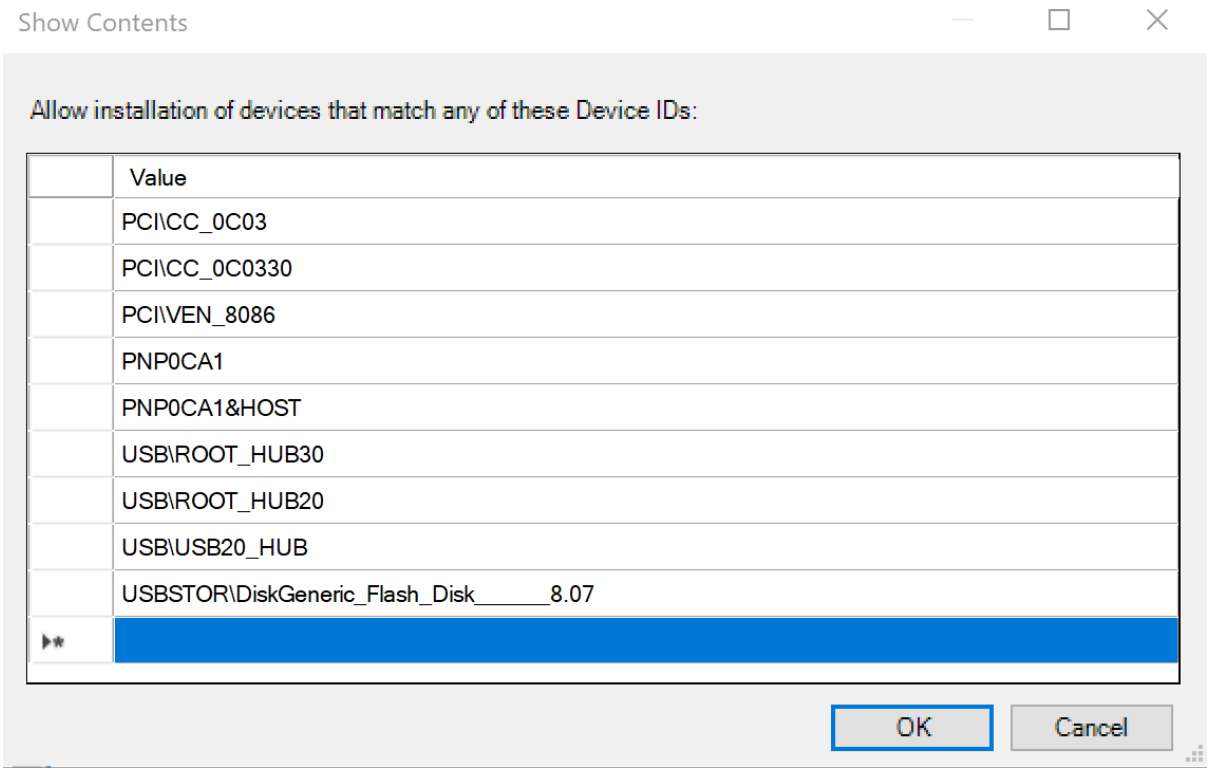
The previous step prevents all future USB devices from being installed. Before you move to the next step make sure you have as complete list as possible of all the USB Host Controllers, USB Root Hubs and Generic USB Hubs Device IDs available to prevent blocking you from interacting with your system through keyboards and mice.

9. Open the **Apply layered order of evaluation for Allow and Prevent device installation policies across all device match criteria** policy and enable it. This policy will enable you to override the wide coverage of the 'Prevent' policy with a specific device.



Apply layered order of evaluation policy

10. Now Open **Allow installation of devices that match any of these device IDs** policy and select the 'Enable' radio button.
11. In the lower left side, in the 'Options' window, click the 'Show...' box. This option will take you to a table where you can enter the device identifier to allow.
12. Enter the full list of USB device IDs you found above including the specific USB Thumb-drive you would like to authorize for installation- USBSTOR\DiskGeneric_Flash_Disk_____8.07.



Allowed USB Device IDs list

13. Click 'OK'.
14. Click 'Apply' on the bottom right of the policy's window.
15. To apply the 'Prevent' coverage of all currently installed USB devices, open the **Prevent installation of devices using drivers that match these device setup classes** policy again; in the 'Options' window mark the checkbox that says 'also apply to matching devices that are already installed' and click 'OK'.

Testing scenario 5

You shouldn't be able to install any USB thumb-drive, except the one you authorized for usage.

SecureAnyBox5 – souhrn produktových informací

https://secureanybox5.com/

Přehled edicí SecureAnyBox5

Entry level – základní úroveň produktu pro malé instalace do 200 uživatelů. Neobsahuje Audit log archiving. Limitace v počtu zakládání bezpečnostních schránek, rekordů a External Files na uživatele. Není v ní k dispozici SIEM integrace.

Advanced – rozšířená edice pro střední zákazníky do 1000 uživatelů. Limitace v počtu zakládání bezpečnostních schránek, rekordů a External Files na uživatele.

Enterprise – edice určená pro velké zákazníky a pro ty, kteří potřebují dosažení a udržení shody s požadavky NIS2. Je bez omezení. Obsahuje Comliance report. Obsahuje KeyShield SSO pro SSO k SAB5.

EDU – edice určená pro školství. Je bez omezení. Obsahuje Comliance report. Obsahuje KeyShield SSO pro SSO k SAB5.

Home – edice určená pro domácí použití do 15 uživatelů.

Kdo je zákazníkem SAB5

Jakákoliv organizace, firma, škola apod., pokud má více než jeden počítač a více než jednoho uživatele. Pro úřad malé obce, která má jeden počítač na stole pana starosty má SecureAnyBox5 smysl pouze v případě, že mají účet v SecureAnyBox5, který pro ně provozuje krajský úřad. Mezi zákazníky SAB5 logicky patří mnoho úřadů a jiných složek státu, protože podléhají regulacím.

Popis řešení

SecureAnyBox5 má dvě části.

Tou první jsou agenti, které se instalují na počítače, notebooky a servery s OS Windows, Linux nebo MacOSX.

Agent se pak stará o změny hesel vybraných (privilegovaných) účtů. Agent nepotřebuje registraci v doméně nebo komunikaci se serverem. Může pracovat zcela autonomně od začátku až na věky. I tak, když se potřebujete autentizovat k tomu kterému stroji, musíte se zeptat serveru, jaké máte použít heslo. Server vám ho sdělí, pokud k tomu máte oprávnění a zároveň vás zapíše do logu.

Tou druhou je nutnost uchovávat mnoho citlivých informací, které souhrně nazýváme "tajemství".

Ty jsou uchovávány v úložišti SAB, jehož klíčovými vlastnostmi jsou:

- ochrana kryptováním (šifrováním)
- důvěryhodný audit přístupu
- sdílení interně v rámci pracovních skupin, oddělení apod.
- odolnost proti přístupu správce
- řešení situací jako odchod nebo ztráta uživatele, příchod nového uživatele, změna pracovního zařazeníuživatele
- sdílení s externisty – podpora, právní služby, vedení účetnictví, daňové a jiné poradenství
- odolnost proti získání databáze, zálohy databáze nebo snapshot databáze
- sada speciálních klientů, kteří uživatelům umožní chránit citlivá data v různých situacích - hesla, dokumenty, soubory, platební karty, klíče, certifikáty, licenční klíče atd..

Ochrana dat při ztrátě zaměstnance – trezor „Bílá obálka“

SecureAnyBox5 umožňuje řešit i méně standardní situace – například to, jak organizace získá přístup k datům uživatele, o kterého přišla nebo je nedostupný. Měl ale přistup k důležitým datům a ten je nezbytné co nejrychleji obnovit. Proto je součástí SecureAnyBox5 i tzv. Bílá obálka – speciální trezor, který má uživatel standardně vyhrazený jen pro sebe. Tento trezor ale mohou otevřít tzv. Bezpečnostní důstojníci – ne samostatně kterýkoliv z nich, ale vždy určený nejmenší počet. Nelze

nastavit méně než 2. Tím je obsah „Bílé obálky“ chráněn před neautorizovaným přístupem i před ztrátou.

Compliance Report

Reporting zahrnuje i naprosto unikátní řešení, které nazýváme Compliance Report - hesla, která jsou uložená, jsou ve shodě s regulací.
Nastavíte si profily pro různé typy tajemství – například WIndowsServerPrivilegedUser nebo Printer nebo NetworkAppliance. Ne moc, stačí tolik, kolik různých požadavků máte –jak dlouhé a komplexní heslo třeba na daném typu zařízení jde nastavit. Co je možné na Windows serveru, není možné pro Wifi AP. Označíte záznamy a nastavíte jim profil. To lze i hromadně, jedním kliknutím. Pak definujete Scope pro report – proskenování všech záznamů typu WindowsPrivilegedUser. A získáte report, který můžete bez dalších úprav předložit auditorovi. Nebo projít výsledek a zjednat nápravu dříve, než audit proběhne.
Compliance report může spustit každý uživatel a zkontrolovat, zda jeho hesla nebo hesla, za která odpovídá, jsou ve shodě.

External Files

SecureAnyBox5 chrání a řeší sdílení a přístup a audit přístupu pro libovolně velké soubory v libovolném úložišti. Soubory se šifrují, ale nejsou uložené v databázi, mohou být sdíleny například i přes cloud.

Výhody používání SAB5 a důvody k přechodu na něj

- Enterprise Class produkt
- skutečná bezpečnost s využitím kryptografie – celosvětově ojedinělý koncept – každý záznam je kryptovaný zvlášť a správce ani nikdo jiný Vám ani sobě nemůže umožnit přístup k datům, pokud nemá klíč, kterým jsou jednotlivé záznamy resp. jejich chráněné části zakryptovány
- dosažení a udržení shody s požadavky NIS2 – požadavky na hesla uživatelů, správců, technických a privilegovaných účtů – pro dané zařízení nebo účet zajistí používání nejlepších možných hesel, tj. hesel, která dosáhnou požadované úrovně entropie
- audit
- SIEM integrace
- hesla, klíče, dokumenty a soubory v interním nebo externím úložišti
- automatické zálohování a archivace se stejným zabezpečením jako živý systém
- hromadná správa uživatelů, integrace pro AD, AAD, EDIR, LDAP
- podpora jiných OS než jen Windows – Linux či MacOSX
- reporting, řešení situací jako nástup, odchod, přechod, externista atd.
- Compliance, generování a kontrola kvality hesel
- flexibilita
- uživatelsky intuitivní systém na používání

Další informace

SAB obsahuje plně veřejně dokumentované API pro integraci s dalšími produkty - https://secureanybox5.com/api-documentation/
Podrobné informace o funkčních vlastnostech, jejich nastavení a používání jsou veřejně dostupné formou uživatelského manuálu - https://secureanybox5.com/documentation/SecureAnyBox_en.html

FEATURES COMPARISON SHEET

Safetica ONE vs. Safetica

Feature	Safetica ONE	Safetica
DEPLOYMENT & HOSTING		
On-prem and Cloud deployment	On-Prem only Microsoft SQL Server 2012+	Microsoft SQL Server 2016+
Support for older operating systems (Windows 7, Windows 8.1, Windows 10 32-bit)	✓	✗
Support for macOS	✓	✓
Support of 1,000+ seat customer environments	✓	✓
Integration with third-party analytics tools (i.e. PowerBI)	✓	✓
DATA LOSS PREVENTION		
Data-in-use classification	✓	✓
Data-at-rest discovery	✓	✓
Safetica Unified Data Classification – combines content, context, file types, and the option to use existing 3rd-party data classifications	✗	✓
Protection policies	✓	✓
INSIDER RISK MANAGEMENT		
Risk-based approach	✗	✓
Application control, web control, device control	✓	✓
UEBA and employee activity audit	✓	✓
BitLocker encryption management [On-Prem Exclusive]	✓	✓
Real-time alerts	✓	✓
CLOUD DATA PROTECTION		
Microsoft 365 integration	✓	✓
REPORTS		
Automated Security audit report	✓	✓
Customizable and scheduled reports	✓	✓
Modern management console	✗	✓
New functionality and continuous innovation updates	✗ *still supported for bugfixes and security updates	✓

FEATURES COMPARISON SHEET

Safetica NXT vs. Safetica

Feature	Safetica NXT	Safetica
DEPLOYMENT & HOSTING		
On-prem and Cloud deployment	Cloud only	Microsoft SQL Server 2016+
Support for older operating systems (Windows 7, Windows 8.1, Windows 10 32-bit)	✓	✗
Configurable user license management	✗	✓
Admin account permissions	✗	✓
Support for macOS	✓	✓
Support of 1,000+ seat customer environments	✗	✓
DATA LOSS PREVENTION		
Data-in-use classification	✓	✓
Data-at-rest discovery	✗	✓
Safetica Unified Data Classification – combines content, context, file types, and the option to use existing 3rd-party data classifications	✗	✓
Enhanced OCR – ability to activate or deactivate OCR for specific devices and to configure 2nd language	✗	✓
Auditing policies – ability to decide which data-related action to record	✗	✓
Protection policies	✓	✓ with enhanced capabilities
INSIDER RISK MANAGEMENT		
Risk-based approach	✓	✓
Application control, web control, device control	✗	✓
Complex data destination control	◐	✓
UEBA and detailed user behavior records	◐	✓
Real-time alerts for incidents	✓	✓
CLOUD DATA PROTECTION		
Microsoft 365 integration	✗	✓
REPORTS		
Automated Security audit report	✓	✓
Customizable and scheduled reports	✓	✓
New functionality and continuous innovation updates	✗ *still supported for bugfixes and security updates	✓

Overview

HPE Aruba Networking CX 6300 Switch Series

The HPE Aruba Networking CX 6300 Switch Series is a modern, flexible, and intelligent family of stackable switches ideal for enterprise network access, aggregation, core, and data center top of rack (ToR) deployments. Created for game- changing operational efficiency with built-in security and resiliency, the CX 6300 switches provide the foundation for high-performance networks supporting IoT, mobile and cloud applications.

Built from the ground up with a combination of cutting-edge hardware, software and analytics and automation tools, the stackable CX 6300 switches are part of the CX switching portfolio, designed for today's enterprise campus, branch and data center networks. By combining a modern, fully programmable OS with the Network Analytics Engine, the CX 6300 switches provide industry leading monitoring and troubleshooting capabilities for the access layer.

A powerful Gen7 ASIC architecture delivers performance and robust feature support with flexible programmability for tomorrow's applications. The Virtual Stacking Framework (VSF) allows for stacking of up to 10 switches, providing scale and simplified management. This flexible series has built-in wirespeed 1/10/25/50GbE uplinks and supports high density IEEE 802.3bt high power PoE. HPE Smart Rate multi-gigabit Ethernet paves the way for high-speed access points and IoT devices by delivering fast connectivity and high power PoE using existing cabling. Modular models offer redundancy and PoE customization with hot-swappable power supplies and fans. Back-to-front airflow available in switch bundle for hot-cold aisle top-of-rack (TOR) and out-of-band-management (OOBM) data center deployments.

Dynamic Segmentation extends foundational wireless role-based policy capability to wired switches. What this means is that the same security, user experience and simplified IT management can be enjoyed throughout the network. Regardless of how users and IoT devices connect, consistent policies are enforced across wired and wireless networks, keeping traffic secure and separate.



HPE Aruba Networking CX 6300 Switch Series

Overview

Key Features

- Stackable Layer 3 switches with BGP, EVPN, VXLAN, VRF, and OSPF with robust security and QoS
- High performance 880 Gbps system switching capacity, 660 MPPS of system throughput and up to 200 Gbps stacking bandwidth (400 Gbps at full duplex)
- Compact 1U switches with full density HPE Smart Rate (1G/2.5G/5G/10GbE) multi-gigabit, up to 90W PoE (Class 8) and 10G LRM SFP+ available on select models
- Power-to-port switch bundle with back-to-front airflow ideal for data center 1GbE ToR and OOBM deployments
- Three stackable, high-performance Layer 2-only CX 6300L access switcheswith HPE Smart Rate Multi-Gigabit and MACsec encryption
- Built-in high speed 1/10/25/50GbE uplinks[1]
- 50GbE connectivity with 50GbE DACs[1]
- Intelligent monitoring, visibility, and remediation with Network Analytics Engine
- Manage via single pane of glass with HPE Aruba Networking Central across wired, wireless, and WAN
- NetEdit support for automated configuration and verification
- Dynamic Segmentation enables secure and simple access for users and IoT

Notes: [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G transceivers and DACs are not supported on S0E91A and S0X44A switch models, which requires QSFP to SFP56 DAC cable for VSF stacking with other CX 6300F and CX 6300M switch models only. VSF stacking not supported on 1G ports.

What's New

- Stackable Layer 3 switches with BGP, EVPN, VXLAN, VRF, and OSPF with robust security and QoS, now supporting MACsec 256, PTP, AVB and LRM.
- Stackable Layer 3 switches with BGP, EVPN, VXLAN, VRF, and OSPF with robust security and QoS, now supporting MACsec 256, PTP, AVB and LRM..
- Series includes full density HPE Smart Rate (1/2.5/5/10GbE) multi-gigabit and up to 90W PoE, SFP+ model and switch bundle with power-to-port airflow for data center 1GbE ToR and OOBM deployments.
- Built-in high speed 10/25/40/50/100GbE uplinks (50GbE for DAC connectivity only) with intelligent monitoring and visibility via Network Analytics Engine.
- Manage via single pane of glass with Central across wired, wireless, and WAN. Dynamic Segmentation enables enhanced security and simple access for users and IoT.
- HPE Data Center Networking Solution SKU (R9F63A) for HPE integrated and mixed use HPE Compute, HPE Storage and HPE Aruba Networking configuration and deployments.

Standard Features

AOS-CX - A Modern Operating System

The CX 6300 Switch Series is based on AOS-CX, a modern, database-driven operating system that automates and simplifies many critical and complex network tasks. A built-in time series database enables customers and developers to utilize software scripts for historical troubleshooting, as well as analysis of past trends. This helps predict and avoid future problems due to scale, security, and performance bottlenecks. Because AOS-CX is built on a modular architecture with a stateful database, our operating system provides the following unique capabilities:

- Easy access to all network state information allows unique visibility and analytics
- REST APIs and Python scripting for fine-grained programmability of network tasks
- A micro-services architecture that enables full integration with other workflow systems and services
- Continuous telemetry data with WebSocket subscriptions for event driven automation
- Continual state synchronization that provides superior fault tolerance and high availability
- All software processes communicate with the database rather than each other, ensuring near real-time state and resiliency and allowing individual software modules to be independently upgraded for higher availability.

Every CX switch includes AOS-CX at no cost and with an active, perpetual set of native features which has everything needed to deploy, connect, and troubleshoot an enterprise network, including:

- Network Analytics Engine (NAE)
- Dynamic Segmentation
- Switch Stacking
- High Availability and Resiliency
- Quality of Service (QoS)
- Layer 2 Switching
- Layer 3 Services and Routing
- IP Multicast
- Network Security

In addition to the native features available in AOS-CX, we offer an optional, term-based CX Advanced Feature Pack that unlocks visibility and advanced security use cases.

For more information, read the **CX Feature Pack Ordering Guide.**

HPE Aruba Networking Central - Unified Single Pane of Glass Management

HPE Aruba Networking Central is an AI-powered solution that simplifies IT operations, improves agility, and reduces costs by unifying management of all network infrastructure. Built for enterprise-grade resiliency and security, while simple enough for smaller businesses with limited IT staff, Central is your single point of visibility and control that spans the entire network --from branch to data center, wired and wireless LAN to WAN.

Available as a cloud-based or on-premises solution, HPE Aruba Networking Central is designed to simplify day zero through day two operations with streamlined workflows for tasks such as virtual switch stack creation, automated monitoring using AI-powered insights and NAE, as well as a unified view of all devices and users, both wired and wireless. Comprehensive switch management capabilities include configuration, on-boarding, monitoring, troubleshooting, and reporting.

A Central Foundation license subscription enables comprehensive switch management capabilities that include configuration, onboarding, monitoring, troubleshooting, and reporting. A Central Advanced license expands these capabilities with premium security and AIOps, including the Central NetConductor Fabric Wizard and Policy Manager to enable dynamic segmentation and distributed enforcement at a global scale. Additionally, an Central Advanced subscription enables the CX Advanced Feature Pack so there is no need to separately purchase a CX Advanced Feature Pack. This streamlines operational efficiency, reducing the need for your IT team to keep track of multiple subscriptions, active terms, and renewal dates.
For more information on HPE Aruba Networking Central subscriptions, see the **HPE Aruba Networking Central SaaS Subscription Ordering Guide.**

Standard Features

HPE Aruba Networking Network Analytics Engine - Advanced Monitoring and Diagnostics

For enhanced visibility and troubleshooting, Network Analytics Engine (NAE) automatically monitors and analyzes events that can impact network health. Advanced telemetry and automation provide the ability to easily identify and troubleshoot networks, system, application, and security related issues easily, using python agents, CLI-based agents, and REST APIs.

Advanced telemetry and automation provide the ability to easily identify and troubleshoot network, system, application, and security related issues easily, through the use of python agents, CLI-based agents, CLI-based agents and REST APIs.

The Time Series Database (TSDB) stores configuration and operational state data, making it available to quickly resolve network issues. The data may also be used to analyze trends, identify anomalies, and predict future capacity requirements.

HPE Aruba Networking Central uses NAE and agents to deliver switch monitoring, analytics, and enhanced troubleshooting for wired assurance. Switch Multi-Edit Software and third-party tools such as ServiceNow and Slack provide the intelligence to integrate NAE alerts into IT service management processes, speeding problem resolution.

NetEdit - Automated Switch Configuration and Management

The entire CX portfolio empowers IT teams to orchestrate multiple switch configuration changes for smooth end-to-end service rollouts. NetEdit introduces automation that allows for rapid network-wide changes and ensures policy conformance post network updates. Intelligent capabilities include search, editing, validation (including conformance checking), deployment and audit features. Capabilities include:

- Centralized configuration with validation for consistency and compliance
- Time savings via simultaneous viewing and editing of multiple configurations
- Customized validation tests for corporate compliance and network change analysis
- Automated large-scale configuration deployment without programming
- Network health and topology visibility via NAE integration

Notes: A separate software license is required to use NetEdit.

CX Mobile App - True Deployment Convenience

An easy-to-use mobile app simplifies connecting and managing CX 6300 Switch Series for any size project. Switch information can also be imported into NetEdit for simplified configuration management and to continuously validate the conformance of configurations anywhere in the network. The CX Mobile App is available for download.

ASICs - Programmable Innovation

Based on over 30 years of continuous investment, ASICs create the basis for innovative and agile software feature advancements, unparalleled performance, and deep visibility. These programmable ASICs are purpose-built to allow for a tighter integration of switch hardware and software within campus and data center architectures to optimize performance and capacity. Virtual Output Queuing (VOQ) isolates congestion prevents Head of Line Blocking (HOLB) and allows full line rate on outgoing (egress) ports. Flexible ASIC resources enable NAE solution to inspect all data, which allows for industry-leading analytics capabilities. The CX 6300 Switch Series is based on the Gen7 ASIC architecture.

Dynamic Segmentation - Campus and Branch Fabric

Standard Features

The Dynamic Segmentation solution enables seamless mobility, consistent policy enforcement, and automated configurations for wired and wireless clients across networks of all sizes. It unifies role-based access and policy enforcement across LAN, WLAN, and SD-WAN networks with centralized policy definition and dedicated enforcement points, ensuring that users and devices can only communicate with destinations consistent with their role-keeping traffic secure and separate.

Dynamic Segmentation is based on establishing least privilege access to IT resources by segmenting traffic based on identity, a fundamental concept of both Zero Trust and SASE frameworks where trust is based on roles and policies, not on where and how a user or device connects.

This innovation begins with colorless ports and role-based micro-segmentation technologies. Colorless ports allow wired clients to connect to any switch port, with the configuration automated using RADIUS-based access control. This eliminates the need for manual on-boarding of clients, including IoT devices, onto the network.

Role-based micro-segmentation delivers benefits of reduced subnet and VLAN sprawl, simplified policy definition, and scalable policy enforcement by introducing the concept of client user roles. Independent of network constructs such as VLANs and VRFs, clients can be grouped into a user role based on their identity, allowing the colorless ports technology to be extended to the centralized overlay fabric, as clients are on-boarded with automatic tunnel creation based on the associated user roles policy. The user roles policy offers the choice between micro-segmentation using centralized and unified policy enforcement for wireless and wired traffic with Layer 7 stateful firewall on gateways or a distributed approach with a Layer 4 role-role ACL on switches.

Dynamic Segmentation provides scale and flexibility in network design by allowing the stretching of VLANs and subnets across the entire network with an EVPN/VXLAN-based distributed overlay fabric. Fabric overlays use VXLAN or VXLAN-GBP tunnels on the data plane and provide the option of a Multi-Protocol BGP EVPN control plane for large deployments, or a static Layer 2 control plane for simplified deployments.

Mobility and IoT Performance

The CX 6300 Switch Series uses a fully distributed architecture that utilizes the Gen7 ASICs. This ensures that our switches offer very low latency, increased packet buffering, and adaptive power consumption. All switching and routing are wire-speed to meet the demands of bandwidth-intensive applications today and in the future. Each switch includes the following:

- Up to 1760 Gbps in non-blocking bandwidth and up to 1310 Mpps for forwarding
- 1/10/25/40/50/100G uplinks[1] and large TCAM sizes ideal for mobility and IoT deployments in large campuses with several thousand clients
- Selectable queue configurations that allow for increased performance by defining a number of queues and associated memory buffering to best meet the requirements of network applications

Notes: [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G transceivers and DACs are not supported on S0E91A and S0X44A switch models, which requires QSFP to SFP56 DAC cable for VSF stacking with other CX 6300F and CX 6300M switch models only. VSF stacking not supported on 1G ports.

VSF Stacking - Scale and Simplicity

The Virtual Switching Framework (VSF) allows you to quickly grow your network using high performance front plane stacking. Additional features include:

- Support for up to 10 switches (or members) in a stack via chain or ring topology
- Flexibility to create stacks that span longer distances such as hundreds of meters across campuses to

Standard Features

kilometers between sites using long-range 10GbE/25GbE transceivers
- Flexibility to mix both modular and fixed CX 6300 Switch Series models within a single stack to meet your deployment requirements
- Simplified configuration and management as the switches act as a single chassis when stacked
- High availability by design using VSF in-service software upgrades for ISSU orchestration and no downtime or restart when upgrading within the same major release (requires at least a 2 member VSF stack)
- The CX Mobile app provides support for a validated stack deployment that ensure that all stack links and uplinks are connected properly

---

An HPE Aruba Networking CX 6300 Switch Series Switch for Any Enterprise Environment

Whether in the branch office or a small to large enterprise environment, you can choose from 24 and 48 ports of HPE Smart Rate Multi-Gigabit. Each switch includes four high-speed built-in uplinks that auto-negotiate from 1GbE, 10GbE to 50GbE[1] to deliver non-blocking performance. Fixed format (F) models include built-in power supplies. The modular (M) models have rear slots for hot swappable power supplies that allow you to customize your PoE requirements, and its fans are field replaceable. Additional highlights:

- Compact 1U models support:

  24 and 48 ports of HPE Smart Rate Multi-gigabit Ethernet IEEE 802.3bz (100M/1GbE/2.5GbE/5GbE/10GbE) supporting high power IEEE 802.3bt Class 6 (60W) to Class 8 (90W)

  High density 24 port SFP+ model which is ideal for aggregation

  1/10/25/50GbEuplink[1] port connectivity

- HPE Smart Rate Multi-Gigabit (IEEE 802.3bz) Ethernet supports high speed wireless access points
- For deployments that need higher port and PoE density, the 6300 supports up to 90W of PoE in a 48-port switch for a total of 2880W of PoE.
- Industry standard IEEE 802.3bt High Power PoE support (Class 8) provides up to 90W per port for support of the latest IoT devices and APs. PoE support for IEEE 802.3at Power over Ethernet (PoE+) provides up to 30W per port as well as any IEEE 802.3af-compliant end device
- Support for pre-standard PoE detection provides power to legacy PoE devices
- High availability with always-on PoE that supplies PoE power even during scheduled reboots and firmware upgrades
- Quick PoE supplies PoE power to powered devices as soon as the switch is plugged into AC power so device can initialize at the same time as the switch OS boots up
- Support for Energy Efficient Ethernet IEEE 802.3az reduces power consumption during periods of low network traffic
- Support for top-of-rack (ToR) and out-of-band management (OOBM) data center deployments with CX 6300M power-to-port bundle that delivers required power-to-port (back to front) airflow
- Auto-MDIX provides automatic adjustments for straight-through or crossover cables on all 10M/100M/1G and Smart Rate ports
- Unsupported Transceiver Mode (UTM) allows to insert and enable all unsupported 1/10/25/50GbE transceivers and cables. Note that there is no warranty nor support for the transceiver/cable when this feature is used.
- IPv6 capabilities include:

  IPv6 host enables switches to be managed in an IPv6 network

  Dual stack (IPv4 and IPv6) transitions from IPv4 to IPv6, supporting connectivity for both protocols

Standard Features

MLD snooping forwards IPv6 multicast traffic to the appropriate interface

IPv6 ACL/QoS supports ACL and QoS for IPv6 network traffic

IPv6 routing supports Static and OSPFv3 protocols

Security provides RA guard, DHCPv6 protection, dynamic IPv6 lockdown, ND snooping, IPv6 Destination Guard, IPv6 DHCP Guard, and IPv6 Router Advertisement Guard

- Jumbo frames allow for high-performance backups and disaster-recovery systems; provides a maximum frame size of 9198 bytes
- Packet storm protection against broadcast and multicast storms with user-defined thresholds
- Smart link enables simple, fast converging link redundancy and load balancing with dual uplinks avoiding Spanning Tree complexities

Notes: [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G transceivers and DACs are not supported on S0E91A and S0X44A switch models, which requires QSFP to SFP56 DAC cable for VSF stacking with other CX 6300F and CX 6300M switch models only. VSF stacking not supported on 1G ports.

CX 6300L layer 2 switches

Three CX 6300L switch models are available for customers needing scalability, high-capacity, and cost-effective connectivity in the access layer. Features include:

- 24 ports of SR10 (1G/2.5G/5G/10G) or 48 ports of SR5 (1G/2.5G/5G) HPE Smart Rate Multi-gigabit Ethernet downlinks that support high power IEEE 802.3bt Class 6 (60W) to Class 8 (90W) PoE and MACsec 256 data link layer encryption
- Scalability with VSF front-plane stacking up to 10 CX 6300L switch members (does not stack with CX 6300F or CX 6300M switches)
- Layer 2 switching with support for IPv4 based static routing, quality of service (QoS), access control lists (ACL), and User-Based Tunneling for Dynamic Segmentation
- Modular, hot-swappable power supplies and fans that allow you to customize for PoE requirements and field replacement needs

Notes:[1]150G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G transceivers and DACs are not supported on S0E91A and S0X44A switch models, which requires QSFP to SFP56 DAC cable for VSF stacking with other CX 6300F and CX 6300M switch models only. VSF stacking not supported on 1G ports.

CX 6300M bundle for data centers

The CX 6300M 48 port power-to-port switch bundle serves as a top of rack (ToR) switch for 1GbE servers and also as a 1GbE out-of-band management (OOBM) switch for data centers server racks. Features include:

- Power-to-port bundle (JL762A) includes 48 port 1GbE switch with 2 x Fan Trays (JL761A) and 1 x power supply (JL760A)
- Back (power-side) to front (1GbE port side) airflow
- 1/10/25/50GbE[1] SFP uplinks

Notes: [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G transceivers and DACs are not supported on S0E91A and S0X44A switch models, which requires QSFP to SFP56 DAC cable for VSF stacking with other CX 6300F and CX 6300M switch models only. VSF stacking not supported on 1G ports.

High Availability and Resiliency

To ensure a high degree of up-time we offer high availability and multicast features needed for a full Layer 3 deployment at access and aggregation such as PBR, BFD, MSDP, BSR, and IP SLA without the need for software licenses. This includes:

- Hot Swappable Power Supplies available in the CX 6300 "M" models

  Provides N+1 and N+N redundancy for high reliability in the event of power line or supply failures

  Optional secondary power supplies to increase the total available PoE power

  Fixed power supplies in CX 6300 "F" models

- Bidirectional Forward Detection (BFD) enables sub-second failure detection for rapid routing protocol re-balancing, supporting both IPV4 and IPv6 networks.
- Virtual Router Redundancy Protocol (VRRP) allows groups of two routers to dynamically create highly available routed environments in IPV4 and IPV6 networks
- Uni-directional Link Detection (UDLD) to monitor link connectivity and shut down ports at both ends if uni- directional traffic is detected, preventing loops in STP- based networks
- EEE 802.3ad LACP supports up to 256 LAGs, each with up to 8 links per LAG; and provides support for static or dynamic groups and a user-selectable hashing algorithm
- IEEE 802.1s Multiple Spanning Tree provides high link availability in VLAN environments where multiple spanning trees are required; and legacy support for IEEE 802.1d and IEEE 802.1w
- IEEE 802.3ad link-aggregation-control protocol (LACP) and port trunking support static and dynamic trunks where each trunk supports up to eight links (ports) per static trunk
- Support for Microsoft Network Load Balancer (NLB) for server applications
- Ethernet Ring Protection Switching (ERPS) supports rapid protection and recovery in a ring topology.
- Hot-Patching support and ISSU for standalone CX 6300 and for CX 6300 with VSF Stacking

Quality of Service (QoS) features

To support congestion actions and traffic prioritization, the CX 6300 Series includes the following:

- Strict priority (SP) queuing and Deficit Weighted Round Robin (DWRR)
- Traffic prioritization (IEEE 802.1p) for real-time classification into 8 priority levels that are mapped to 8 queues
- Layer 4 prioritization based on TCP/UDP port numbers
- Class of Service (CoS) sets the IEEE 802.1p priority tag based on IP address, IP Type of Service (ToS), Layer 3 protocol, TCP/UDP port number, source port, and DiffServ
- Rate limiting sets per-port ingress enforced maximums and per-port, per-queue minimums
- Transmission rates of egressing frames can be limited on a per-queue basis using Egress Queue Shaping (EQS)
- Large buffers for graceful congestion management

Simplified Configuration and Management

In addition to HPE Aruba Networking Central, the CX Mobile App, NetEdit and Network Analytics Engine, the CX 6300 series offers the following:

- Built-in programmable and easy to use REST API interface

Standard Features

- AirWave on-premises and Central cloud- based management
- Zero-Touch Provisioning (ZTP) simplifies installation of switching infrastructure using DHCP-based or Activate-based process with AirWave and HPE Aruba Networking Central
- Scalable ASIC-based wire speed network monitoring and accounting with no impact on network performance; network operators can gather a variety of network statistics and information for capacity planning and real- time network monitoring purposes
- Management interface control enables or disables each of the following depending on security preferences, console port, or reset button
- Industry-standard CLI with a hierarchical structure for reduced training time and expense. Delivers increased productivity in multivendor environments
- Management security restricts access to critical configuration commands, provides multiple privilege levels with password protection and local and remote syslog capabilities allow logging of all access
- SNMP v2c/v3 provides SNMP read and trap support of industry standard Management Information Base (MIB), and private extensions
- SNMP support includes: Write Set Speed and Duplex, Write Port Security, Write POE Priority, Write Config Mgmt, SNMP-Read single OID for average CPU and memory, SNMP MIB View
- SNMP Trap include: Transceiver Traps (insertion/removal), SNMP Trap, SNMP MIB-SNMB Authentication, SNMPv2 MIB, Port Sec MIB-Port Sec, Config MIB-Running Config Change, Config MIB, AAA Server MIB, AAA Server State
- Remote monitoring (RMON) with standard SNMP to monitor essential network functions. Supports events, alarms, history, and statistics groups as well as a private alarm extension group; RMON and sFlow provide advanced monitoring and reporting capabilities for statistics, history, alarms and events
- IP Flow Information Export (IPFix) enables client flow information collection to enhance visibility
- Simplifies configuration while onboarding switches with Zero Touch Provisioning by using Dynamic Border Gateway Protocol (BGP) peering to establish a peer group of switches within an IP range
- Enhanced visibility during client onboarding, providing insights on latency, failures, and error events
- Client telemetries and application visibility using IP Flow Information Export (IPFix), Deep Packet Inspection (DPI) and traffic insights
- Simplifies configuration while onboarding switches with Zero Touch Provisioning by using Dynamic Border Gateway Protocol (BGP) peering to establish a peer group of switches within an IP range
- Enhanced visibility during client onboarding, providing insights on latency, failures, and error events
- TFTP and SFTP support offers different mechanisms for configuration updates; trivial FTP (TFTP) allows bidirectional transfers over a TCP/ IP network; Secure File Transfer Protocol (SFTP) runs over an SSH tunnel to provide additional security
- Debug and sampler utility supports ping and traceroute for IPv4 and IPv6
- Network Time Protocol (NTP) synchronizes timekeeping among distributed time servers and clients; keeps timekeeping consistent among all clock-dependent devices within the network so the devices can provide diverse applications based on the consistent time
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP) advertises and receives management information from adjacent devices on a network, facilitating easy mapping by network management applications
- Dual flash images provides independent primary and secondary operating system files for backup while upgrading
- Assignment of descriptive names to ports for easy identification
- Multiple configuration files can be stored to a flash image
- Ingress and egress port monitoring enable more efficient network problem solving
- Unidirectional link detection (UDLD) monitors the link between two switches and blocks the ports on both ends of the link if the link goes down at any point between the two devices
- Power down mode delivers energy savings by allowing the switch to power down most of the switch, except a clock which will boot up the switch when scheduled
- IP SLA for Voice monitors quality of voice traffic using the UDP Jitter and UDP Jitter for VoIP tests
- Precision Time Protocol (PTP) allows for precise clock synchronization across distributed network switches as defined in IEEE 1588. Transparent Clock (PTP-TC) and Boundary Clock (PTP-BC) are needed for time critical applications like Audio Video Bridging (AVB), smart grid power automation, financial systems and more. Boundary Clock makes use of 2-Step time stamping mode.

Standard Features

Layer 2 Switching

The following layer 2 services are supported:

- VLAN support and tagging for IEEE 802.1Q (4094 VLAN IDs)
- Jumbo packet support improves the performance of large data transfers; supports frame size of up to 9198 bytes
- IEEE 802.1v protocol VLANs isolate select non-IPv4 protocols automatically into their own VLANs
- Rapid Per-VLAN Spanning Tree (RPVST+) allows each VLAN to build a separate spanning tree to improve link bandwidth usage; is compatible with PVST+
- MVRP allows automatic learning and dynamic assignment of VLANs
- VXLAN encapsulation (tunnelling) protocol for overlay network that enables a more scalable virtual network deployment
- Bridge Protocol Data Unit (BPDU) tunnelling Transmits STP BPDUs transparently, allowing correct tree calculations across service providers, WANs, or MANs
- Port mirroring duplicates port traffic (ingress and egress) to a monitoring port; supports 4 mirroring groups
- STP supports standard IEEE 802.1D STP, IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) for faster convergence, and IEEE 802.1s Multiple Spanning Tree Protocol (MSTP)
- Internet Group Management Protocol (IGMP) Controls and manages the flooding of multicast packets in a Layer 2 network
- IPv4 Multicast in VXLAN/EVPN Overlay support allows PIM-SM/IGMP snooping in the VXLAN Overlay
- IPv6 VXLAN/EVPN Overlay support, allows IPv6 traffic over the VXLAN overlay
- VXLAN ARP/ND suppression allows minimization of ARP and ND traffic flooding within individual VXLAN segments, thus optimizing the VXLAN network
- QinQ support to improve the VLAN utilization by adding another 802.1Q tag to tagged packets

Layer 3 Services

The following layer 3 services are supported:

- Bidirectional Forwarding Detection (BFD) enables link connectivity monitoring and reduces network convergence time for static route, OSPFv2 and VRRP
- User Datagram Protocol (UDP) helper function allows UDP broadcasts to be directed across router interfaces to specific IP unicast or subnet broadcast addresses and prevents server spoofing for UDP services such as DHCP
- Loopback interface address defines an address in Open Shortest Path First (OSPF), improving diagnostic capability
- Route maps provide more control during route redistribution; allow filtering and altering of route metrics
- Address Resolution Protocol (ARP) determines the MAC address of another IP host in the same subnet; supports static ARPs; gratuitous ARP allows detection of duplicate IP addresses; proxy ARP allows normal ARP operation between subnets or when subnets are separated by a Layer 2 network
- Dynamic Host Configuration Protocol (DHCP) simplifies the management of large IP networks and supports client; DHCP Relay enables DHCP operation across subnets
- DHCP server centralizes and reduces the cost of IPv4 address management
- Domain Name System (DNS) provides a distributed database that translates domain names and IP addresses, which simplifies network design; supports client and server
- mDNS (Multicast Domain Name System) Gateway enables discovery of mDNS groups across L3 boundaries
- Generic Routing Encapsulation (GRE) enables tunneling traffic from site-to-site over a Layer 3 path
- Supports internal loopback testing for maintenance purposes and increased availability; loopback detection protects against incorrect cabling or network configurations and can be enabled on a per-port or per-VLAN basis for added flexibility
- IP sub-interface is a virtual interface created by dividing physical interface into multiple logical

Standard Features

interfaces tagged using different VLAN-IDs. A physical interface can be a regular physical, Split port or LAG L3 interface. A sub-interface is used for many uses-cases such as VRF-lite interconnection and inter-vlan routing (router on-a-stick)

Layer 3 Routing

The following layer 3 routing services are supported:

- Border Gateway Protocol (BGP) provides IPv4 and IPv6 routing, which is scalable, robust, and flexible
- Border Gateway Protocol 4 (BGP-4) delivers an implementation of the Exterior Gateway Protocol (EGP) utilizing path vectors; uses TCP for enhanced reliability for the route discovery process; reduces bandwidth consumption by advertising only incremental updates; supports extensive policies for increased flexibility; scales to very large networks with graceful restart capability
- Equal-Cost Multipath (ECMP) enables multiple equal-cost links in a routing environment to increase link redundancy and scale bandwidth
- Multi-protocol BGP (MP-BGP) enables sharing of IPv6 routes using BGP and connections to BGP peers using IPv6
- Routing Information Protocol version 2 (RIPv2) provides an easy to configure routing protocol for small networks as while RIPng provides support for small IPv6 networks
- Open shortest path first (OSPF) delivers faster convergence; uses link-state routing Interior Gateway Protocol (IGP), which supports ECMP, NSSA, and MD5 authentication for increased security and graceful restart for faster failure recovery
- OSPF provides OSPFv2 for IPv4 routing and OSPFv3 for IPv6 routing
- Static IP routing provides manually configured routing; includes ECMP capability
- Policy-based routing uses a classifier to select traffic that can be forwarded based on policy set by the network administrator
- Static IPv4 and IPv6 routing provides simple manually configured IPv4 and IPv6 routes
- IP performance optimization provides a set of tools to improve the performance of IPv4 networks; includes directed broadcasts, customization of TCP parameters, support of ICMP error packets, and extensive display capabilities
- Dual IP stack maintains separate stacks for IPv4 and IPv6 to ease the transition from an IPv4-only network to an IPv6-only network design

Security

The CX 6300 Switch Series comes with an integrated trusted platform module (TPM) for platform integrity. This ensures the boot process started from a trusted combination of AOS-CX switches. Other security features include:

- AOS-CX uses FIPS 140-2 validated cryptography for protection of sensitive information.
- Access control list (ACL) support for both IPv4 and IPv6; allows for filtering traffic to prevent unauthorized users from accessing the network, or for controlling network traffic to save resources; rules can either deny or permit traffic to be forwarded; rules can be based on a Layer 2 header or a Layer 3 protocol header
- ACLs also provide filtering based on the IP field, source/ destination IP address/subnet, and source/ destination TCP/UDP port number on a per-VLAN or per-port basis
- Enrollment over Secure Transport (EST) enables secure certificate enrollment, allowing for easier enterprise management of PKI
- Remote Authentication Dial-In User Service (RADIUS)
- Terminal Access Controller Access-Control System (TACACS+) delivers an authentication tool using TCP with encryption of the full authentication request, providing additional security
- Management access security for both on- and off- box authentication for administrative access. RADIUS or TACACS+ can be used to provide encrypted user authentication. Additionally, TACACS+ can also provide admin authorization services
- Control Plane Policing sets rate limit on control protocols to protect CPU overload from DOS attacks

Standard Features

- Supports multiple user authentication methods. Uses an IEEE 802.1X supplicant on the client in conjunction with a RADIUS server to authenticate in accordance with industry standards
- Supports MAC-based client authentication
- Concurrent IEEE 802.1X, Web, and MAC authentication schemes per switch port accepts up to 32 sessions of IEEE 802.1X, Web, and MAC authentications
- DHCP protection blocks DHCP packets from unauthorized DHCP servers, preventing denial-of-service attacks
- Secure management access delivers secure encryption of all access methods (CLI, GUI, or MIB) through SSHv2, SSL, and/or SNMPv3
- Switch CPU protection provides automatic protection against malicious network traffic trying to shut down the switch
- ICMP throttling defeats ICMP denial-of-service attacks by enabling any switch port to automatically throttle ICMP traffic
- Identity-driven ACL enables implementation of a highly granular and flexible access security policy and VLAN assignment specific to each authenticated network user
- STP BPDU port protection blocks Bridge Protocol Data Units (BPDUs) on ports that do not require BPDUs, preventing forged BPDU attacks
- Dynamic IP lockdown works with DHCP protection to block traffic from unauthorized hosts, preventing IP source address spoofing
- Dynamic ARP protection blocks ARP broadcasts from unauthorized hosts, preventing eavesdropping or theft of network data
- STP root guard protects the root bridge from malicious attacks or configuration mistakes
- Port security allows access only to specified MAC addresses, which can be learned or specified by the administrator
- MAC address lockout prevents particular configured MAC addresses from connecting to the network
- Source-port filtering allows only specified ports to communicate with each other
- Secure shell encrypts all transmitted data for secure remote CLI access over IP networks
- Secure Sockets Layer (SSL) encrypts all HTTP traffic, allowing secure access to the browser-based management GUI in the switch
- Secure FTP allows secure file transfer to and from the switch; protects against unwanted file downloads or unauthorized copying of a switch configuration file
- Critical Authentication Role ensures that important infrastructure devices such as IP phones are allowed network access even in the absence of a RADIUS server
- MAC Pinning allows non-chatty legacy devices to stay authenticated by pinning client MAC addresses to the port until the clients logoff or get disconnected
- Security banner displays a customized security policy when users log in to the switch
- RadSec enables RADIUS authentication and accounting data to be passed safely and reliably across insecure networks
- Private VLAN (PVLAN) provides traffic isolation between users on the same VLAN; typically, a switch port can only communicate with other ports in the same community and/or an uplink port, regardless of VLAN ID or destination MAC address. This extends network security by restricting peer-peer communication to prevent variety of malicious attacks.
- Auto VLAN Creation automates VLAN creation on access switches for authenticated clients.
- DHCP smart relay allows the DHCP relay agent to use secondary IP addresses when the DHCP server does not reply the DHCP-OFFER message
- IEEE 802.1AE MACsec provides switch-to-switch and switch-to-host security on a link between two ports using standard encryption and authentication, available on uplink and downlink ports

---

Visibility and Advanced Security

Customers can choose to upgrade the active, embedded CX Foundation license to the term based CX Advanced license to unlock the following benefits for their business:

- Delivers deep visibility with CX Edge Insights for application recognition, identification, and flow capture from layer 4 to layer 7. CX Edge Insights enables granular datapoint collection with search, sort, and reporting as well as the ability to recognize 22 categories and more than 3700 applications.

Multicast

- IGMP Snooping allows multiple VLANs to receive the same IPv4 multicast traffic, lessening network bandwidth demand by reducing multiple streams to each VLAN
- Multicast Listener Discovery (MLD) enables discovery of IPv6 multicast listeners; support MLD v1 and v2
- Protocol Independent Multicast (PIM) defines modes of IPv4 and IPv6 multicasting to allow one-to-many and many-to-many transmission of information; supports PIM Sparse Mode (SM), Source-Specific Multicast (SSM), and Dense Mode (DM) for both IPv4 and IPv6
- Internet Group Management Protocol (IGMP) utilizes Any-Source Multicast (ASM) to manage IPv4 multicast networks; supports IGMPv1, v2, and v3
- Multicast Service Discovery Protocol (MSDP) efficiently routes multicast traffic through core networks
- MSDP for Anycast RP is an intra-domain feature that provides redundancy and load-sharing capabilities

Convergence

- IP multicast routing includes PIM Sparse, Source Specific Multicast, and Dense modes to route IP multicast traffic
- IP multicast snooping (data-driven IGMP) prevents flooding of IP multicast traffic
- Protocol Independent Multicast for IPv6 supports one-to- many and many-to-many media casting use cases such as IPTV over IPv6 networks
- LLDP-MED (Media Endpoint Discovery) defines a standard extension of LLDP that stores values for parameters such as QoS and VLAN to automatically configure network devices such as IP phones
- PoE allocations supports multiple methods (allocation by usage or class, with LLDP and LLDP-MED) to allocate PoE power for more efficient power management and energy savings.
- Auto VLAN configuration for voice RADIUS VLAN uses a standard RADIUS attribute and LLDP-MED to automatically configure a VLAN for IP phones
- CDPv2 uses CDPv2 to configure legacy IP phones

Additional information

- Green initiative support for RoHS (EN 50581:2012) and WEEE regulations
- TAA-compliant CX 6300 switches are available for order

Customer First, Customer Last Support

When your network is important to your business, then your business needs the backing of Support Services. Partner with HPE Aruba Networking product experts to increase your team productivity, keep pace with technology advances, software releases, and obtain break-fix support.

- HPE Aruba Networking Pro Care adds fast access to senior HPE Aruba Networking TAC engineers, who are assigned as a single point of contact for case management, reducing the time spent addressing and resolving issues.

For complete details on Foundation Care and HPE Aruba Networking Pro Care, please visit: **https://www.hpe.com/us/en/networking/hpe-aruba-networking-support-services.html**

Warranty, Services and Support

- Limited Lifetime Warranty, see **https://www.hpe.com/us/en/networking/hpe-aruba-networking-support-services.html** for warranty and support information included with your product purchase

Standard Features

- For Software Releases and Documentation, refer to **https://networkingsupport.hpe.com/**
- For more detailed information on HPE Aruba Networking AOS-CX software release and features, please visit the **AOS-CX Switch Software Documentation Portal**
- Explore and compare switch features for each platform and software release on the **HPE Aruba Networking Switch Feature Navigator**
- For support and services information, visit **https://www.hpe.com/us/en/networking/hpe-aruba-networking-support-services.html**

Configuration Information

BTO Models

HPE Aruba Networking CX 6300M

Rule #	Description	SKU
1, 2, 3, 4, 6, 8	HPE Aruba Networking CX 6300M 24-port SFP+ and 4-port SFP56 Switch • HPE Aruba Networking 6300M 24-port SFP+ and 4-port SFP56 Switch • Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max= 24 SFP/SFP+ 100M/1/10G Transceivers • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G+I43+I69:I76+I69:I80+I43+I69:I69:I104	JL658A
5, 8, 9, 11, 12, 13, 14	HPE Aruba Networking CX 6300M 48p SR10 1G/2.5G/5G/10G PTP/AVB Class8 PoE and 4p 100G MACsec Switch • HPE Aruba Networking 6300M 48p SR10 1G/2.5G/5G/10G PTP/AVB Class 8 PoE and 4p 100G MACsec Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 4 QSFP56 50G Transceiver • QSA28 Adapter Min=0 \ Max=4, ports 49-52 (rule11) • 1U - Height	S0E91A
1, 2, 3, 4, 8	HPE Aruba Networking CX 6300M 48-port HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4-port SFP56 Switch • HPE Aruba Networking 6300M 48-port HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4-port SFP56 Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL659A
2, 3, 4, 8	HPE Aruba Networking CX 6300M 48p HPE Smart Rate 1G/2.5G/5G Class8 PoE 2p 50G 2p 25G Switch • HPE Aruba Networking 6300M 48p HPE Smart Rate 1G/2.5G/5G CL8 PoE and 2p 50G and 2p 25G Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 2 SFP+/SFP28 10/25 Transceiver (Ports 49/50) • Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25/50G Transceiver (Ports 51/52) • 1U - Height	R8S90A
1, 2, 3, 4, 8	HPE Aruba Networking CX 6300M 24-port HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4-port SFP56 Switch • HPE Aruba Networking 6300M 24-port HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4-port SFP56 Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes 1 Fan tray (JL669B), with 1 open slot with blank cover • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL660A
2, 3, 4, 8	HPE Aruba Networking CX 6300M 24p Smart Rate 1G/2.5G/5G/10G Class6 PoE 2p 50G 2p 25G Switch • HPE Aruba Networking 6300M 24p HPE Smart Rate 1G/2.5G/5G/10G CL6 PoE	R8S89A

Configuration Information

	and 2p 50G and 2p 25G Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 2 SFP+/SFP28 10/25 Transceiver (Ports 25/26) • Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25/50G Transceiver (Ports 27/28) • 1U - Height	
1, 2, 3, 4, 8	HPE Aruba Networking CX 6300M 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch • HPE Aruba Networking 6300M 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes 1 Fan tray (JL669B), with 1 open slot with blank cover • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL661A
1, 2, 3, 4, 7, 8	HPE Aruba Networking CX 6300M 12p CL8 PoE 36p CL6 PoE Smart Rate 1G/2.5G/5G 2p 50G 2p 10G Switch • HPE Aruba Networking 6300M 12p Class8 PoE and 36p Class6 PoE HPE Smart Rate 1G/2.5G/5G and 2p 50G and 2p 25G Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 2 SFP+ 10/25/50G Transceiver (LRM Suported) (Ports 49/50) • Min=0 \ Max = 2 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver (Ports 51/52) • 1U - Height	R8S91A
1, 2, 3, 4, 8	HPE Aruba Networking CX 6300M 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch • HPE Aruba Networking 6300M 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes 1 Fan tray (JL669B), with 1 open slot with blank cover • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL662A
1, 2, 3, 4, 7, 8. 14	HPE Aruba Networking CX 6300M 24p SFP+ LRM support and 2p 50G and 2p 25G MACSec Switch • HPE Aruba Networking 6300M 24p SFP+ LRM support and 2p 50G and 2p 25G MACsec Switch • Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK • Includes Fantrays Min2 / Max2 (JL669B) • Min=0 \ Max = 24 SFP/SFP+ Transceiver (LRM Supported Ports 1-24) • Min=0 \ Max = 2 SFP+/SFP28 10/25/50G Transceiver (Ports 25/26) • Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25/50G Transceiver (Ports 27/28) • 1U - Height	R8S92A
1, 2, 3, 4, 8	HPE Aruba Networking CX 6300M 48-port 1GbE and 4-port SFP56 Switch • HPE Aruba Networking 6300M 48-port 1GbE and 4-port SFP56 Switch • Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK • Includes 1 Fan tray (JL669B), with 1 open slot with blank cover • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver (Ports 25/26) • 1U - Height	JL663A
1, 2, 3, 4, 8	HPE Aruba Networking CX 6300M 24-port 1GbE and 4-port SFP56 Switch	JL664A

Configuration Information

- HPE Aruba Networking 6300M 24-port 1GbE and 4-port SFP56 Switch
- Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 5, 8 HPE Aruba Networking CX 6300M 48p 1GbE 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle — JL762A

- HPE Aruba Networking 6300M 48-port 1GbE and 4-port SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle
- Includes 1 Pwr2Prt PSU, can select Min0 / Max1 (250W JL760A)
- Includes 2 Pwr2Prt Fan trays (JL761A), with no open slots
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

HPE Aruba Networking CX 6300M 48p 1GbE 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle PDU — JL762A#B2B

- C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)

HPE Aruba Networking CX 6300M 48p 1GbE 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle PDU — JL762A#B2C

- C15 PDU Jumper Cord (ROW) (J9944A)

HPE Aruba Networking CX 6300M 48p 1GbE 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle 220v — JL762A#B2E

- HPE 2.5m C15 to NEMA 6-20P Pwr. Cord (JL336A)

HPE Aruba Networking CX 6300M 48p 1GbE 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle NoLoc — JL762A#AC3

- No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)

HPE Aruba Networking CX 6300M TAA

1, 2, 3, 4, 8 HPE Aruba Networking CX 6300M 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch — S0F99A

- HPE Aruba Networking 6300M 24p 1GbE Class 4 PoE and 4p SFP56 TAA Switch
- Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 8 HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G TAA Switch — S0G00A

- HPE Aruba Networking 6300M 48p 1GbE and 4p SFP56 TAA Switch
- Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 8 HPE Aruba Networking CX 6300M 24p 10M/100M/1G 4p SFP56 50G TAA Switch — S0G01A

- HPE Aruba Networking 6300M 24p 1GbE and 4p SFP56 TAA Switch
- Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 6, 8 HPE Aruba Networking CX 6300M 24p SFP+ 1G/10G 4p SFP56 50G TAA Switch — S0G03A

- HPE Aruba Networking 6300M 24p SFP+ and 4p SFP56 TAA Switch
- Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK

Configuration Information

- Includes Fantrays Min2 / Max 2 (JL669B)
- Min=0 \ Max= 24 SFP/SFP+ 100M/1/10G Transceivers
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 8	HPE Aruba Networking CX 6300M 48p Smart Rate 1G/2.5G/5G Class6 PoE 4p SFP56 50G TAA Switch • HPE Aruba Networking 6300M 48p HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4p SFP56 TAA Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	S0G04A
1, 2, 3, 4, 8	HPE Aruba Networking CX 6300M 24p Smart Rate 1G/2.5G/5G Class6 PoE 4p SFP56 50G TAA Switch • HPE Aruba Networking 6300M 24p HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4p SFP56 TAA Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes 1 Fan tray (JL669B), with 1 open slot with blank cover • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	S0G05A
1, 2, 3, 4, 8	HPE Aruba Networking CX 6300M 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch • HPE Aruba Networking 6300M 48p 1GbE Class 4 PoE and 4p SFP56 TAA Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes 1 Fan tray (JL669B), with 1 open slot with blank cover • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	S0G06A
8, 9, 11, 12, 13, 14	HPE Aruba Networking CX 6300M 48p SR10 1G/2.5G/5G/10G PTP/AVB Class8 PoE 4p 100G MACsec TAA Switch • HPE Aruba Networking 6300M 48p SR10 1G/2.5G/5G/10G PTP/AVB Class 8 PoE and 4p 100G MACsec TAA Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 4 QSFP+/QSFP56/QSFP28 50G Transceiver • QSA28 Adapter Min=0 \ Max=4, ports 49-52 (rule11) • 1U - Height	S0X44A
1, 2, 3, 4, 5, 8	HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G Power-to-Port 2xFan PSU TAA Bundle • HPE Aruba Networking 6300M 48p 1GbE and 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU TAA Switch Bundle • Includes 1 Pwr2Prt PSU, can select Min0 / Max1 (250W JL760A) • Includes 2 Pwr2Prt Fan trays (JL761A), with no open slots • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height"	S0G02A
	HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G Back-to-Front 2xFan PSU TAA Bundle PDU • C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	S0G02A#B2B
	HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G Back-to-Front 2xFan PSU TAA Bundle PDU • C15 PDU Jumper Cord (ROW) (J9944A)	S0G02A#B2C

Configuration Information

	HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G Back-to-Front 2xFan PSU TAA 220v Bundle • HPE 2.5m C15 to NEMA 6-20P Pwr Cord(JL336A)	S0G02A#B2E
	HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G Back-to-Front 2xFan PSU TAA Bundle No Loc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	S0G02A#AC3
	6300F TAA	
1, 2, 3, 4, 5, 8	HPE Aruba Networking CX 6300F 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch • HPE Aruba Networking 6300F 48-port 1GbE Class 4 PoE and 4-port SFP56 TAA Switch • Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame • Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	S0G95A
	HPE Aruba Networking CX 6300F 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA PDU Switch • C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	S0G95A#B2B
	HPE Aruba Networking CX 6300F 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA PDU Switch • C15 PDU Jumper Cord (ROW) (J9944A)	S0G95A#B2C
	HPE Aruba Networking CX 6300F 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA 220v Switch • HPE 2.5m C15 to NEMA 6-20P Pwr Cord(JL336A)	S0G95A#B2E
	HPE Aruba Networking CX 6300F 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch No Loc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	S0G95A#AC3
1, 2, 3, 4, 5, 8	HPE Aruba Networking CX 6300F 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch • HPE Aruba Networking 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 TAA Switch • Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame • Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	S0G96A
	HPE Aruba Networking CX 6300F 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA PDU Switch • C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	S0G96A#B2B
	HPE Aruba Networking CX 6300F 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA PDU Switch • C15 PDU Jumper Cord (ROW) (J9944A)	S0G96A#B2C
	HPE Aruba Networking CX 6300F 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA 220v Switch • HPE 2.5m C15 to NEMA 6-20P Pwr Cord(JL336A)	S0G96A#B2E
	HPE Aruba Networking CX 6300F 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch No Loc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	S0G96A#AC3
1, 2, 3, 4, 5, 8	HPE Aruba Networking CX 6300F 48p 10M/100M/1G 4p SFP56 50G TAA Switch	S0G97A

Configuration Information

- HPE Aruba Networking 6300F 48-port 1GbE and 4-port SFP56 TAA Switch
- Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame
- Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

HPE Aruba Networking CX 6300F 48p 10M/100M/1G 4p SFP56 50G TAA PDU Switch — S0G97A#B2B

- C13 PDU Jumper Cord (NA/MEX/TW/JP)

HPE Aruba Networking CX 6300F 48p 10M/100M/1G 4p SFP56 50G TAA PDU Switch — S0G97A#B2C

- C13 PDU Jumper Cord (ROW)

HPE Aruba Networking CX 6300F 48p 10M/100M/1G 4p SFP56 50G TAA 220v Switch — S0G97A#B2E

- HPE 2.3m C13 to NEMA 6-15P Pwr Cord(J9936A)

HPE Aruba Networking CX 6300F 48p 10M/100M/1G 4p SFP56 50G TAA Switch No Loc — S0G97A#AC3

- No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)

Rule #	Description	SKU
1, 2, 3, 4, 5, 8	HPE Aruba Networking CX 6300F 24p 10M/100M/1G 4p SFP56 50G TAA Switch	S0G98A

- HPE Aruba Networking 6300F 24-port 1GbE and 4-port SFP56 TAA Switch
- Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame
- Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

HPE Aruba Networking CX 6300F 24p 10M/100M/1G 4p SFP56 50G TAA PDU Switch — S0G98A#B2B

- C13 PDU Jumper Cord (NA/MEX/TW/JP)

HPE Aruba Networking CX 6300F 24p 10M/100M/1G 4p SFP56 50G TAA PDU Switch — S0G98A#B2C

- C13 PDU Jumper Cord (ROW)

HPE Aruba Networking CX 6300F 24p 10M/100M/1G 4p SFP56 50G TAA 220v Switch — S0G98A#B2E

- HPE 2.3m C13 to NEMA 6-15P Pwr Cord(J9936A)

HPE Aruba Networking CX 6300F 24p 10M/100M/1G 4p SFP56 50G TAA Switch No Loc — S0G98A#AC3

- No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)

HPE Aruba Networking CX 6300F

Rule #	Description	SKU
1, 2, 3, 4, 5, 8	HPE Aruba Networking CX 6300F 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch	JL665A

- Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame
- Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

HPE Aruba Networking CX 6300F 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch PDU — JL665A#B2B

- C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)

HPE Aruba Networking CX 6300F 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch PDU — JL665A#B2C

- C15 PDU Jumper Cord (ROW) (J9944A)

Configuration Information

	HPE Aruba Networking CX 6300F 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch 220v • HPE 2.5m C15 to NEMA 6-20P Pwr. Cord (JL336A)	JL665A#B2E
	HPE Aruba Networking CX 6300F 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch NoLoc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	JL665A#AC3
1, 2, 3, 4, 5, 8	HPE Aruba Networking CX 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch • Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame • Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL666A
	HPE Aruba Networking CX 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch PDU • C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	JL666A#B2B
	HPE Aruba Networking CX 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch PDU • C15 PDU Jumper Cord (ROW) (J9944A)	JL666A#B2C
	HPE Aruba Networking CX 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch 220v • HPE 2.5m C15 to NEMA 6-20P Pwr. Cord (JL336A)	JL666A#B2E
	HPE Aruba Networking CX 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch NoLoc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	JL666A#AC3
1, 2, 3, 4, 5, 8	HPE Aruba Networking CX 6300F 48-port 1GbE and 4-port SFP56 Switch • Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame • Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL667A
	HPE Aruba Networking CX 6300F 48-port 1GbE and 4-port SFP56 Switch PDU • C13 PDU Jumper Cord (NA/MEX/TW/JP)	JL667A#B2B
	HPE Aruba Networking CX 6300F 48-port 1GbE and 4-port SFP56 Switch PDU • C13 PDU Jumper Cord (ROW)	JL667A#B2C
	HPE Aruba Networking CX 6300F 48-port 1GbE and 4-port SFP56 Switch 220v • HPE 2.3m C13 to NEMA 6-15P Pwr. Cord (J9936A)	JL667A#B2E
	HPE Aruba Networking CX 6300F 48-port 1GbE and 4-port SFP56 Switch NoLoc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	JL667A#AC3
1, 2, 3, 4, 5, 8	HPE Aruba Networking CX 6300F 24-port 1GbE and 4-port SFP56 Switch • Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame • Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL668A
	HPE Aruba Networking CX 6300F 24-port 1GbE and 4-port SFP56 Switch PDU	JL668A#B2B

Configuration Information

	• C13 PDU Jumper Cord (NA/MEX/TW/JP)	
	HPE Aruba Networking CX 6300F 24-port 1GbE and 4-port SFP56 Switch PDU	JL668A#B2C
	• C13 PDU Jumper Cord (ROW)	
	HPE Aruba Networking CX 6300F 24-port 1GbE and 4-port SFP56 Switch 220v	JL668A#B2E
	• HPE 2.3m C13 to NEMA 6-15P Pwr. Cord (J9936A)	
	HPE Aruba Networking CX 6300F 24-port 1GbE and 4-port SFP56 Switch NoLoc	JL668A#AC3
	• No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	

HPE Aruba Networking CX 6300L

2, 3, 4, 8	HPE Aruba Networking CX 6300L 24p Smart Rate 1G/2.5G/5G/10G CL6 2p SFP56 50G 2p SFP28 25G L2 Switch • HPE Aruba Networking 6300L 24p Smart Rate 1G/2.5G/5G/10G Class6 PoE 2p 50G and 2p 25G L2 Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25/50G Transceiver (Ports 25/26) • Min=0 \ Max = 2 SFP+ or 2 SFP28 10/25 Transceiver (Ports 27/28) • 1U - Height	S3L75A
2, 3, 4, 8	HPE Aruba Networking CX 6300L 48p Smart Rate 100M/1G/2.5G/5G CL8 2p SFP56 50G 2p SFP28 25G L2 Switch • HPE Aruba Networking 6300L 48p Smart Rate 1G/2.5G/5G Class8 PoE 2p 50G and 2p 25G L2 Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25/50G Transceiver (Ports 49/50) • Min=0 \ Max = 2 SFP+ or 2 SFP28 10/25 Transceiver (Ports 51/52) • 1U - Height	S3L76A
1, 2, 3, 4, 7, 8	HPE ANW CX 6300L 48p Smart Rate 100M/1G/2.5G/5G Class6/8 2p SFP56 50G 2p SFP+ 1G/10G LRM Layer2 Sw • HPE Aruba Networking 6300L 48p Smart Rate 1G/2.5G/5G Class8/6 PoE 2p 50G and 2p 10G LRM L2 Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25/50G Transceiver (Ports 49/50) • Min=0 \ Max = 2 SFP or 2 SFP+ 1/10 Transceiver (LRM Supported) (Ports 51/52) • 1U - Height	S3L77A

Configuration Rules

Rule #	Description	SKU
1	The following Transceivers install into this Switch: (Use BTO only when adding to switch)	
	HPE Aruba Networking 1G SFP LC SX 500m OM2 MMF Transceiver	J4858D
	HPE Aruba Networking 1G SFP LC LX 10km SMF Transceiver	J4859D
	HPE Aruba Networking 1G SFP LC LH 70km SMF Transceiver	J4860D
	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e Transceiver	J8177D
	HPE Aruba Networking 1G SFP LC SX 500m MMF TAA Transceiver	JL745A
	HPE Aruba Networking 1G SFP LC LX 10km SMF TAA Transceiver	JL746A
	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e TAA Transceiver	JL747A

Configuration Information

2	The following Transceivers install into this Switch: (Use BTO only when adding to switch)	
	HPE Aruba Networking 10GBASE-T SFP+ RJ45 30m Cat6A Transceiver	JL563C
	HPE Aruba Networking 10G SFP+ LC SR 300m OM3 MMF Transceiver	J9150D
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF Transceiver	J9151E
	HPE Aruba Networking 10G SFP+ LC ER 40km SMF Transceiver	J9153D
	HPE Aruba Networking 10G SFP+ LC SR 300m MMF TAA Transceiver	JL748A
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF TAA Transceiver	JL749A
	HPE Aruba Networking 10G SMF Simplex LC BiDi 40km-Downstream 1330/1270 Transceiver	R9X54A
	HPE Aruba Networking 10G SMF Simplex LC BiDi 40km-Upstream 1270/1330 Transceiver	R9X55A
	HPE Aruba Networking 10G SR SFP+ LC 400m OM4 MMF C-Class Transceiver	S2P30A
	HPE Aruba Networking 10G LR SFP+ LC 10km SMF C-Class Transceiver	S2P31A
	HPE Aruba Networking 10G ER SFP+ LC 40km SMF C-Class Transceiver	S2P32A
	HPE Aruba Networking 10G SFP+ to SFP+ 1m Direct Attach Copper Cable	J9281D
	HPE Aruba Networking 10G SFP+ to SFP+ 3m Direct Attach Copper Cable	J9283D
3	The following Transceivers install into this Switch: (Use BTO only when adding to switch)	
	HPE Aruba Networking 25G ER LC 40km SMF Transceiver	S0V69A
	HPE Aruba Networking 25G BiDi 10km-Downstream 1330/1270 Transceiver	S1C96A
	HPE Aruba Networking 25G BiDi 10km-Upstream 1270/1330 Transceiver	S1C98A
	HPE Aruba Networking 25G SFP28 LC SR 100m MMF Transceiver	JL484A
	HPE Aruba Networking 25G SFP28 LC eSR 400m MMF Transceiver	JL485A
	HPE Aruba Networking 25G SFP28 LC LR 10km SMF Transceiver	JL486A
	HPE Aruba Networking 25G SR SFP28 LC 100m MMF C-Class Transceiver	S2P33A
	HPE Aruba Networking 25G LR SFP28 LC 10km SMF C-Class Transceiver	S2P34A
	HPE Aruba Networking 25G SFP28 to SFP28 0.65m Direct Attach Cable	JL487A
	HPE Aruba Networking 25G SFP28 to SFP28 3m Direct Attach Copper Cable	JL488A
	HPE Aruba Networking 25G SFP28 to SFP28 5m Direct Attach Copper Cable	JL489A
	HPE Aruba Networking 25G SFP28 to SFP28 3m Active Optical Cable	R0M44A
	HPE Aruba Networking 25G SFP28 to SFP28 7m Active Optical Cable	R0M45A
	HPE Aruba Networking 25G SFP28 to SFP28 15m Active Optical Cable	R0Z21A
4	The following Transceivers install into this Switch: (Use BTO only when adding to switch)	
	HPE Aruba Networking 50G eSR 300m MMF Transceiver	S0V64A
	HPE Aruba Networking 50G LR 10km SMF Transceiver	S0V65A
	HPE Aruba Networking 50G ER 40km SMF Transceiver	S0V66A
	HPE Aruba Networking 50G BiDi 10km-Downstream 1330/1270 Transceiver	S1C92A
	HPE Aruba Networking 50G BiDi 10km-Upstream 1270/1330 Transceiver	S1C94A
	HPE Aruba Networking 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable	R0M46A
	HPE Aruba Networking 50G SFP56 to SFP56 3m Direct Attach Copper Cable	R0M47A
	HPE Aruba Networking 50G SFP56 LC SR 100m MMF Transceiver	R0M48A
	HPE Aruba Networking 50G QSFP56 to SFP56 0.65m DAC Cable	S1J07A
	HPE Aruba Networking 50G QSFP56 to SFP56 3m DAC Cable	S1J08A
5	Localization required on orders without #B2B, #B2C, #B2E or #AC3 options.	
6	The following Transceivers install into this Switch: (Use BTO only when adding to switch)	
	HPE Aruba Networking 100M SFP LC FX 2km MMF Transceiver	J9054D
7	The following Transceivers install into this Switch and is only available on LRM Supported ports. See AOS-Switch and AOS-CX Transceiver Guide (Edition 12) for port compatibility: (Use BTO only when adding to switch)	

Configuration Information

	HPE Aruba Networking 10G SFP+ LC LRM 220m OM2 MMF Transceiver	J9152D
8	If ANY Option is integrated 0D1 to this Switch, then the Switch requires 0D1. (Box level integration is not allowed)	
9	The following Transceivers install into this Switch using QSFP56 side of cable: (Use BTO only when adding to switch)	
	HPE Aruba Networking 50G QSFP56 to SFP56 0.65m DAC Cable	S1J07A
	HPE Aruba Networking 50G QSFP56 to SFP56 3m DAC Cable	S1J08A
11	If qty1 of the folowing QSA28 Adapter(845970-B21) is selected, then increase max SFP28 Port qty by 1 and allow user selection of the following SFP Transceivers. Refer to qty and port restrictions for individual Switch in the "Additional Info" sections: (Use BTO only when adding to switch)	
	HPE Aruba Networking 10G SFP+ LC SR 300m OM3 MMF Transceiver	J9150D
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF Transceiver	J9151E
	HPE Aruba Networking 10G SFP+ LC ER 40km SMF Transceiver	J9153D
	HPE Aruba Networking 10G SFP+ LC SR 300m MMF TAA Transceiver	JL748A
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF TAA Transceiver	JL749A
	HPE Aruba Networking 25G SFP28 LC SR 100m MMF Transceiver	JL484A
	HPE Aruba Networking 25G SFP28 LC eSR 400m MMF Transceiver	JL485A
	HPE Aruba Networking 25G SFP28 LC LR 10km SMF Transceiver	JL486A
	HPE Aruba Networking 25G ER LC 40km SMF Transceiver	S0V69A
	HPE Aruba Networking 25G BiDi 10km-Downstream 1330/1270 Transceiver	S1C96A
	HPE Aruba Networking 25G BiDi 10km-Upstream 1270/1330 Transceiver	S1C98A
	HPE Aruba Networking 25G SFP28 to SFP28 3m Active Optical Cable	R0M44A
	HPE Aruba Networking 25G SFP28 to SFP28 7m Active Optical Cable	R0M45A
	HPE Aruba Networking 25G SFP28 to SFP28 15m Active Optical Cable	R0Z21A
12	The following Transceivers install into this Switch: (Use BTO only when adding to switch)	
	HPE Aruba Networking 100G SR1.2 QSFP28 LC 100m MMF Transceiver	S4B44A
	HPE 100Gb QSFP28 Bidirectional Transceiver	845972-B21
	HPE QSFP28 to SFP28 Adapter	845970-B21
	HPE Aruba Networking 100G QSFP28 LC CWDM4 2km SMF Transceiver	R0Z30A
	HPE Aruba Networking 100G QSFP28 LC LR4 10km SMF 2-strand Transceiver	JL310A
	HPE Aruba Networking 100G QSFP28 MPO SR4 100m 12-fiber MPO OM3 MMF Transceiver	JL309A
	HPE Aruba Networking 100G QSFP28 LC ER4L 40km SMF Transceiver	JL743A
	HPE Aruba Networking 100G QSFP28 LC FR1 SMF 2km Transceiver	R9B63A
	HPE Aruba Networking 100G QSFP28 to QSFP28 3m Direct Attach Copper Cable	JL307A
	HPE Aruba Networking 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable	R0Z25A
	HPE Aruba Networking 100G QSFP28 to QSFP28 5m Direct Attach Copper Cable	R0Z26A
	HPE Aruba Networking 100G QSFP28 to QSFP28 2m Active Optical Cable	JL856A
	HPE Aruba Networking 100G QSFP28 to QSFP28 7m Active Optical Cable	R0Z27A
	HPE Aruba Networking 100G QSFP28 to QSFP28 15m Active Optical Cable	R0Z28A
	HPE Aruba Networking 100G QSFP28 to QSFP28 30m Active Optical Cable	R0Z29A
13	The following Transceivers install into this Switch(Use BTO only when adding to switch)	
	HPE Networking X142 40G QSFP+ MPO SR4 Transceiver	JH231A
	HPE Networking X142 40G QSFP+ LC LR4 SM Transceiver	JH232A
	HPE Networking X142 40G QSFP+ MPO eSR4 300M Transceiver	JH233A
	HPE Aruba Networking 40G QSFP+ LC Bidirectional 150m MMF 2-strand Transceiver	JL308A
	HPE Aruba Networking 40G QSFP+ LC ER4 40km SMF Transceiver	Q9G82A
	HPE Networking X242 40G QSFP+ to QSFP+ 1m Direct Attach Copper Cable	JH234A
	HPE Networking X242 40G QSFP+ to QSFP+ 3m Direct Attach Copper Cable	JH235A
	HPE Networking X242 40G QSFP+ to QSFP+ 5m Direct Attach Copper Cable	JH236A

Configuration Information

	HPE Aruba Networking 40G QSFP+ to QSFP+ 7m Active Optical Cable	R0Z22A
	HPE Aruba Networking 40G QSFP+ to QSFP+ 15m Active Optical Cable	R0Z23A
	HPE Aruba Networking 40G QSFP+ to QSFP+ 30m Active Optical Cable	R0Z24A
14	The following Transceivers install into this Switch (Use BTO only when adding to switch)	
	HPE Aruba Networking 25G LR SFP28 LC 10km SMF TAA Transceiver	S2N63A

Notes:

Drop down under power supply should offer the following options and results:

Switch/Router/Power Supply to PDU Power Cord - #B2B in North America, Mexico, Taiwan, and Japan or #B2C ROW. (Watson Default B2B or B2C for Rack Level CTO)

Switch/Router/Power Supply to Wall Power Cord - Localized Option (Watson Default for BTO and Box Level CTO)

High Volt Switch/Router/Power Supply to Wall Power Cord - #B2E Option. (Offered only in North America, Mexico, Taiwan, and Japan)

No Power Cord - #AC3 Option

Locking Power Cord (J9955A) L6-20P is available through the OCA Accessories tab

OCA Only Model Selection Form>HPE Aruba Networking > Switches > HPE Aruba Networking OS > AOS-CX: HPE Aruba Networking 6300 Switch Series

---

Rack Level Integration CTO Models

HPE Aruba Networking CX 6300M

Rule #	Description	SKU
1, 2, 3, 4, 6, 8	HPE Aruba Networking CX 6300M 24-port SFP+ and 4-port SFP56 Switch • HPE Aruba Networking 6300M 24-port SFP+ and 4-port SFP56 Switch • Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max= 24 SFP/SFP+ 100M/1/10G Transceivers • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL658A
5, 6, 7, 10, 11, 12, 13, 14	HPE Aruba Networking CX 6300M 48p SR10 1G/2.5G/5G/10G PTP/AVB Class8 PoE and 4p 100G MACsec Switch • HPE Aruba Networking 6300M 48p SR10 1G/2.5G/5G/10G PTP/AVB Class 8 PoE and 4p 100G MACsec Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 4 QSFP56 50G Transceiver • QSA28 Adapter Min=0 \ Max=4, ports 49-52 (rule11) • 1U - Height	S0E91A
1, 2, 3, 4, 6	HPE Aruba Networking CX 6300M 48-port HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4-port SFP56 Switch • HPE Aruba Networking 6300M 48-port HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4-port SFP56 Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL659A
2, 3, 4, 6	HPE Aruba Networking CX 6300M 48p HPE Smart Rate 1G/2.5G/5G Class8 PoE 2p 50G 2p 25G Switch • HPE Aruba Networking 6300M 48p HPE Smart Rate 1G/2.5G/5G CL8 PoE and 2p 50G and 2p 25G Switch	R8S90A

Configuration Information

- Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK
- Includes Fantrays Min2 / Max 2 (JL669B)
- Min=0 \ Max = 2 SFP+/SFP28 10/25/50G Transceiver (Ports 49/50)
- Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25 Transceiver (Ports 51/52)
- 1U - Height

1, 2, 3, 4, 6 HPE Aruba Networking CX 6300M 24-port HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4-port SFP56 Switch — JL660A

- HPE Aruba Networking 6300M 24-port HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4-port SFP56 Switch
- Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

2, 3, 4, 6 HPE Aruba Networking CX 6300M 24p Smart Rate 1G/2.5G/5G/10G Class6 PoE 2p 50G 2p 25G Switch — R8S89A

- HPE Aruba Networking 6300M 24p HPE Smart Rate 1G/2.5G/5G/10G CL6 PoE and 2p 50G and 2p 25G Switch
- Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK
- Includes Fantrays Min2 / Max 2 (JL669B)
- Min=0 \ Max = 2 SFP+/SFP28 10/25/50G Transceiver (Ports 25/26)
- Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25 Transceiver (Ports 27/28)
- 1U - Height

1, 2, 3, 4, 6 HPE Aruba Networking CX 6300M 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch — JL661A

- HPE Aruba Networking 6300M 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch
- Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 6, 9 HPE Aruba Networking CX 6300M 12p CL8 PoE 36p CL6 PoE Smart Rate 1G/2.5G/5G 2p 50G 2p 10G Switch — R8S91A

- HPE Aruba Networking 6300M 12p Class8 PoE and 36p Class6 PoE HPE Smart Rate 1G/2.5G/5G and 2p 50G and 2p 10G Switch
- Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK
- Includes Fantrays Min2 / Max 2 (JL669B)
- Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25/50G Transceiver (Ports 49/50)
- Min=0 \ Max = 2 SFP/SFP+ 1/10 Transceiver (LRM Supported) (Ports 51/52)
- 1U - Height

1, 2, 3, 4, 6 HPE Aruba Networking CX 6300M 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch — JL662A

- HPE Aruba Networking 6300M 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch
- Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 6, 9, 14 HPE Aruba Networking CX 6300M 24p SFP+ LRM support and 2p 50G and 2p 25G MACSec Switch — R8S92A

Configuration Information

- HPE Aruba Networking 6300M 24p SFP+ LRM support and 2p 50G and 2p 25G MACSec Switch
- Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK
- Includes Fantrays Min2 / Max 2 (JL669B)
- Min=0 \ Max = 24 SFP/SFP+ Transceiver (LRM Supported Ports 1-24)
- Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25/50G Transceiver (Ports 25/26)
- Min=0 \ Max = 2 SFP+/SFP28 1/10/25 Transceiver (Ports 27/28)
- 1U - Height

1, 2, 3, 4, 6 HPE Aruba Networking CX 6300M 48-port 1GbE and 4-port SFP56 Switch — JL663A

- HPE Aruba Networking 6300M 48-port 1GbE and 4-port SFP56 Switch
- Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 6 HPE Aruba Networking CX 6300M 24-port 1GbE and 4-port SFP56 Switch — JL664A

- HPE Aruba Networking 6300M 24-port 1GbE and 4-port SFP56 Switch
- Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height"

1, 2, 3, 4, 5, 6, 7 HPE Aruba Networking CX 6300M 48p 1GbE 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle — JL762A

- HPE Aruba Networking 6300M 48-port 1GbE and 4-port SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle
- Includes 1 Pwr2Prt PSU, can select Min0 / Max1 (250W JL760A)
- Includes 2 Pwr2Prt Fan trays (JL761A), with no open slots
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

HPE Aruba Networking CX 6300M 48p 1GbE 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle PDU — JL762A#B2B

- C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)

HPE Aruba Networking CX 6300M 48p 1GbE 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle PDU — JL762A#B2C

- C15 PDU Jumper Cord (ROW) (J9944A)

HPE Aruba Networking CX 6300M 48p 1GbE 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle 220v — JL762A#B2E

- HPE 2.5m C15 to NEMA 6-20P Pwr. Cord (JL336A)

HPE Aruba Networking CX 6300M 48p 1GbE 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle NoLoc — JL762A#AC3

- No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)

HPE Aruba Networking CX 6300M TAA

1, 2, 3, 4, 6 HPE Aruba Networking CX 6300M 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch — S0F99A

- HPE Aruba Networking 6300M 24p 1GbE Class 4 PoE and 4p SFP56 TAA Switch
- Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 6 HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G TAA Switch — S0G00A

Configuration Information

- HPE Aruba Networking 6300M 48p 1GbE and 4p SFP56 TAA Switch
- Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 6 HPE Aruba Networking CX 6300M 24p 10M/100M/1G 4p SFP56 50G TAA Switch S0G01A

- HPE Aruba Networking 6300M 24p 1GbE and 4p SFP56 TAA Switch
- Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 6, 8 HPE Aruba Networking CX 6300M 24p SFP+ 1G/10G 4p SFP56 50G TAA Switch S0G03A

- HPE Aruba Networking 6300M 24p SFP+ and 4p SFP56 TAA Switch
- Must Select PSU Min1 / Max2 (250W JL085A, 250W JL757A) Mix OK
- Includes Fantrays Min2 / Max 2 (JL669B)
- Min=0 \ Max= 24 SFP/SFP+ 100M/1/10G Transceivers
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 6 HPE Aruba Networking CX 6300M 48p Smart Rate 1G/2.5G/5G Class6 PoE 4p SFP56 50G TAA Switch S0G04A

- HPE Aruba Networking 6300M 48p HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4p SFP56 TAA Switch
- Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK
- Includes Fantrays Min2 / Max 2 (JL669B)
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 6 HPE Aruba Networking CX 6300M 24p Smart Rate 1G/2.5G/5G Class6 PoE 4p SFP56 50G TAA Switch S0G05A

- HPE Aruba Networking 6300M 24p HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4p SFP56 TAA Switch
- Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

1, 2, 3, 4, 6 HPE Aruba Networking CX 6300M 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch S0G06A

- HPE Aruba Networking 6300M 48p 1GbE Class 4 PoE and 4p SFP56 TAA Switch
- Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK
- Includes 1 Fan tray (JL669B), with 1 open slot with blank cover
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

6, 7, 10, 11, 12, 13 HPE Aruba Networking CX 6300M 48p SR10 1G/2.5G/5G/10G PTP/AVB Class8 PoE 4p 100G MACsec TAA Switch S0X44A

- HPE Aruba Networking 6300M 48p SR10 1G/2.5G/5G/10G PTP/AVB Class 8 PoE and 4p 100G MACsec TAA Switch
- Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK
- Includes Fantrays Min2 / Max 2 (JL669B)

Configuration Information

- Min=0 \ Max = 4 QSFP+/QSFP56/QSFP28 50G Transceiver
- QSA28 Adapter Min=0 \ Max=4, ports 49-52 (rule11)
- 1U - Height

1, 2, 3, 4, 5, 6, 7 HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G Power-to-Port 2xFan PSU TAA Bundle — S0G02A

- HPE Aruba Networking 6300M 48p 1GbE and 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU TAA Switch Bundle
- Includes 1 Pwr2Prt PSU, can select Min0 / Max1 (250W JL760A)
- Includes 2 Pwr2Prt Fan trays (JL761A), with no open slots
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G Back-to-Front 2xFan PSU TAA Bundle PDU — S0G02A#B2B

- C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)

HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G Back-to-Front 2xFan PSU TAA Bundle PDU — S0G02A#B2C

- C15 PDU Jumper Cord (ROW) (J9944A)

HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G Back-to-Front 2xFan PSU TAA 220v Bundle — S0G02A#B2E

- HPE 2.5m C15 to NEMA 6-20P Pwr Cord(JL336A)

HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G Back-to-Front 2xFan PSU TAA Bundle No Loc — S0G02A#AC3

- No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)

6300F TAA

1, 2, 3, 4, 5, 6, 7 HPE Aruba Networking CX 6300F 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch — S0G95A

- HPE Aruba Networking 6300F 48-port 1GbE Class 4 PoE and 4-port SFP56 TAA Switch
- Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame
- Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver
- 1U - Height

HPE Aruba Networking CX 6300F 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA PDU Switch — S0G95A#B2B

- C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)

HPE Aruba Networking CX 6300F 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA PDU Switch — S0G95A#B2C

- C15 PDU Jumper Cord (ROW) (J9944A)

HPE Aruba Networking CX 6300F 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA 220v Switch — S0G95A#B2E

- HPE 2.5m C15 to NEMA 6-20P Pwr Cord(JL336A)

HPE Aruba Networking CX 6300F 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch No Loc — S0G95A#AC3

- No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)

1, 2, 3, 4, 5, 6, 7 HPE Aruba Networking CX 6300F 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch — S0G96A

- HPE Aruba Networking 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 TAA Switch
- Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame
- Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame
- Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver

Configuration Information

	• 1U - Height	
	HPE Aruba Networking CX 6300F 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA PDU Switch	S0G96A#B2B
	• C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	
	HPE Aruba Networking CX 6300F 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA PDU Switch	S0G96A#B2C
	• C15 PDU Jumper Cord (ROW) (J9944A)	
	HPE Aruba Networking CX 6300F 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA 220v Switch	S0G96A#B2E
	• HPE 2.5m C15 to NEMA 6-20P Pwr Cord(JL336A)	
	HPE Aruba Networking CX 6300F 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch No Loc	S0G96A#AC3
	• No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	
1, 2, 3, 4, 5, 6, 7	HPE Aruba Networking CX 6300F 48p 10M/100M/1G 4p SFP56 50G TAA Switch	S0G97A
	• HPE Aruba Networking 6300F 48-port 1GbE and 4-port SFP56 TAA Switch • Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame • Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	
	HPE Aruba Networking CX 6300F 48p 10M/100M/1G 4p SFP56 50G TAA PDU Switch	S0G97A#B2B
	• C13 PDU Jumper Cord (NA/MEX/TW/JP)	
	HPE Aruba Networking CX 6300F 48p 10M/100M/1G 4p SFP56 50G TAA PDU Switch	S0G97A#B2C
	• C13 PDU Jumper Cord (ROW)	
	HPE Aruba Networking CX 6300F 48p 10M/100M/1G 4p SFP56 50G TAA 220v Switch	S0G97A#B2E
	• HPE 2.3m C13 to NEMA 6-15P Pwr Cord(J9936A)	
	HPE Aruba Networking CX 6300F 48p 10M/100M/1G 4p SFP56 50G TAA Switch No Loc	S0G97A#AC3
	• No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	
1, 2, 3, 4, 5, 6, 7	HPE Aruba Networking CX 6300F 24p 10M/100M/1G 4p SFP56 50G TAA Switch	S0G98A
	• HPE Aruba Networking 6300F 24-port 1GbE and 4-port SFP56 TAA Switch • Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame • Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	
	HPE Aruba Networking CX 6300F 24p 10M/100M/1G 4p SFP56 50G TAA PDU Switch	S0G98A#B2B
	• C13 PDU Jumper Cord (NA/MEX/TW/JP)	
	HPE Aruba Networking CX 6300F 24p 10M/100M/1G 4p SFP56 50G TAA PDU Switch	S0G98A#B2C
	• C13 PDU Jumper Cord (ROW)	
	HPE Aruba Networking CX 6300F 24p 10M/100M/1G 4p SFP56 50G TAA 220v Switch	S0G98A#B2E
	• HPE 2.3m C13 to NEMA 6-15P Pwr Cord(J9936A)	
	HPE Aruba Networking CX 6300F 24p 10M/100M/1G 4p SFP56 50G TAA Switch No Loc	S0G98A#AC3
	• No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	

Configuration Information

HPE Aruba Networking CX 6300F

Rule #	Description	SKU
1, 2, 3, 4, 5, 6, 7	HPE Aruba Networking CX 6300F 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch • Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame • Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL665A
	HPE Aruba Networking CX 6300F 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch PDU • C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	JL665A#B2B
	HPE Aruba Networking CX 6300F 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch PDU • C15 PDU Jumper Cord (ROW) (J9944A)	JL665A#B2C
	HPE Aruba Networking CX 6300F 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch 220v • HPE 2.5m C15 to NEMA 6-20P Pwr. Cord (JL336A)	JL665A#B2E
	HPE Aruba Networking CX 6300F 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch NoLoc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	JL665A#AC3
1, 2, 3, 4, 5, 6, 7	HPE Aruba Networking CX 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch • Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame • Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL666A
	HPE Aruba Networking CX 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch PDU • C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	JL666A#B2B
	HPE Aruba Networking CX 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch PDU • C15 PDU Jumper Cord (ROW) (J9944A)	JL666A#B2C
	HPE Aruba Networking CX 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch 220v • HPE 2.5m C15 to NEMA 6-20P Pwr. Cord (JL336A)	JL666A#B2E
	HPE Aruba Networking CX 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch NoLoc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	JL666A#AC3
1, 2, 3, 4, 5, 6, 7	HPE Aruba Networking CX 6300F 48-port 1GbE and 4-port SFP56 Switch • Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame • Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL667A
	HPE Aruba Networking CX 6300F 48-port 1GbE and 4-port SFP56 Switch PDU • C13 PDU Jumper Cord (NA/MEX/TW/JP)	JL667A#B2B
	HPE Aruba Networking CX 6300F 48-port 1GbE and 4-port SFP56 Switch PDU • C13 PDU Jumper Cord (ROW)	JL667A#B2C

Configuration Information

	HPE Aruba Networking CX 6300F 48-port 1GbE and 4-port SFP56 Switch 220v • HPE 2.3m C13 to NEMA 6-15P Pwr. Cord (J9936A)	JL667A#B2E
	HPE Aruba Networking CX 6300F 48-port 1GbE and 4-port SFP56 Switch NoLoc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	JL667A#AC3
1, 2, 3, 4, 5, 6, 7	HPE Aruba Networking CX 6300F 24-port 1GbE and 4-port SFP56 Switch • Includes Non-Pluggable, Internal PSU behind sheetmetal Chassis Frame • Includes Non-Pluggable, Internal Fans behind sheetmetal Chassis Frame • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	JL668A
	HPE Aruba Networking CX 6300F 24-port 1GbE and 4-port SFP56 Switch PDU • C13 PDU Jumper Cord (NA/MEX/TW/JP)	JL668A#B2B
	HPE Aruba Networking CX 6300F 24-port 1GbE and 4-port SFP56 Switch PDU • C13 PDU Jumper Cord (ROW)	JL668A#B2C
	HPE Aruba Networking CX 6300F 24-port 1GbE and 4-port SFP56 Switch 220v • HPE 2.3m C13 to NEMA 6-15P Pwr. Cord (J9936A)	JL668A#B2E
	HPE Aruba Networking CX 6300F 24-port 1GbE and 4-port SFP56 Switch NoLoc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	JL668A#AC3
	HPE Aruba Networking CX 6300F	
2, 3, 4, 6	HPE Aruba Networking CX 6300L 24p Smart Rate 1G/2.5G/5G/10G CL6 2p SFP56 50G 2p SFP28 25G L2 Switch • HPE Aruba Networking 6300L 24p Smart Rate 1G/2.5G/5G/10G Class6 PoE 2p 50G and 2p 25G L2 Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25/50G Transceiver (Ports 25/26) • Min=0 \ Max = 2 SFP+ or 2 SFP28 10/25 Transceiver (Ports 27/28) • 1U - Height	S3L75A
2, 3, 4, 6	HPE Aruba Networking CX 6300L 48p Smart Rate 100M/1G/2.5G/5G CL8 2p SFP56 50G 2p SFP28 25G L2 Switch • HPE Aruba Networking 6300L 48p Smart Rate 1G/2.5G/5G Class8 PoE 2p 50G and 2p 25G L2 Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B) • Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25/50G Transceiver (Ports 49/50) • Min=0 \ Max = 2 SFP+ or 2 SFP28 10/25 Transceiver (Ports 51/52) • 1U - Height	S3L76A
1, 2, 3, 4, 6, 9	HPE ANW CX 6300L 48p Smart Rate 100M/1G/2.5G/5G Class6/8 2p SFP56 50G 2p SFP+ 1G/10G LRM Layer2 Sw • HPE Aruba Networking 6300L 48p Smart Rate 1G/2.5G/5G Class8/6 PoE 2p 50G and 2p 10G LRM L2 Switch • Must Select PSU Min1 / Max2 (680W JL086A,1050W JL087A, 1600W JL670A, 1050W JL758A) Mix OK • Includes Fantrays Min2 / Max 2 (JL669B)	S3L77A

Configuration Information

- Min=0 \ Max = 2 SFP+/SFP28/SFP56 10/25/50G Transceiver (Ports 49/50)
- Min=0 \ Max = 2 SFP or 2 SFP+ 1/10 Transceiver (LRM Supported) (Ports 51/52)
- 1U - Height

Configuration Rules

Rule #	Description	SKU
1	The following Transceivers install into this Switch (Use #0D1 quoted to switch if switch is CTO) - if applicable:	
	HPE Aruba Networking 100M SFP LC FX 2km MMF Transceiver	J9054D
Notes:	Not supported in the SFP56 Ports	
	HPE Aruba Networking 1G SFP LC SX 500m OM2 MMF Transceiver	J4858D
	HPE Aruba Networking 1G SFP LC LX 10km SMF Transceiver	J4859D
	HPE Aruba Networking 1G SFP LC LH 70km SMF Transceiver	J4860D
	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e Transceiver	J8177D
	HPE Aruba Networking 1G SFP LC SX 500m MMF TAA Transceiver	JL745A
	HPE Aruba Networking 1G SFP LC LX 10km SMF TAA Transceiver	JL746A
	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e TAA Transceiver	JL747A
2	The following Transceivers install into this Switch (Use #0D1 quoted to switch if switch is CTO) - if applicable:	
	HPE Aruba Networking 10GBASE-T SFP+ RJ45 30m Cat6A Transceiver	JL563C
	HPE Aruba Networking 10G SFP+ LC SR 300m OM3 MMF Transceiver	J9150D
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF Transceiver	J9151E
	HPE Aruba Networking 10G SFP+ LC ER 40km SMF Transceiver	J9153D
	HPE Aruba Networking 10G SFP+ LC SR 300m MMF TAA Transceiver	JL748A
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF TAA Transceiver	JL749A
	HPE Aruba Networking 10G SMF Simplex LC BiDi 40km-Downstream 1330/1270 Transceiver	R9X54A
	HPE Aruba Networking 10G SMF Simplex LC BiDi 40km-Upstream 1270/1330 Transceiver	R9X55A
	HPE Aruba Networking 10G SR SFP+ LC 400m OM4 MMF C-Class Transceiver	S2P30A
	HPE Aruba Networking 10G LR SFP+ LC 10km SMF C-Class Transceiver	S2P31A
	HPE Aruba Networking 10G ER SFP+ LC 40km SMF C-Class Transceiver	S2P32A
	HPE Aruba Networking 10G SFP+ to SFP+ 1m Direct Attach Copper Cable	J9281D
	HPE Aruba Networking 10G SFP+ to SFP+ 3m Direct Attach Copper Cable	J9283D
3	The following Transceivers install into this Switch (Use #0D1 quoted to switch if switch is CTO) - if applicable:	
	HPE Aruba Networking 25G ER LC 40km SMF Transceiver	S0V69A
	HPE Aruba Networking 25G BiDi 10km-Downstream 1330/1270 Transceiver	S1C96A
	HPE Aruba Networking 25G BiDi 10km-Upstream 1270/1330 Transceiver	S1C98A
	HPE Aruba Networking 25G SFP28 LC SR 100m MMF Transceiver	JL484A
	HPE Aruba Networking 25G SFP28 LC eSR 400m MMF Transceiver	JL485A
	HPE Aruba Networking 25G SFP28 LC LR 10km SMF Transceiver	JL486A
	HPE Aruba Networking 25G SR SFP28 LC 100m MMF C-Class Transceiver	S2P33A
	HPE Aruba Networking 25G LR SFP28 LC 10km SMF C-Class Transceiver	S2P34A
	HPE Aruba Networking 25G SFP28 to SFP28 0.65m Direct Attach Cable	JL487A
	HPE Aruba Networking 25G SFP28 to SFP28 3m Direct Attach Copper Cable	JL488A
	HPE Aruba Networking 25G SFP28 to SFP28 5m Direct Attach Copper Cable	JL489A
	HPE Aruba Networking 25G SFP28 to SFP28 3m Active Optical Cable	R0M44A
	HPE Aruba Networking 25G SFP28 to SFP28 7m Active Optical Cable	R0M45A
	HPE Aruba Networking 25G SFP28 to SFP28 15m Active Optical Cable	R0Z21A
4	The following Transceivers install into this Switch (Use #0D1 quoted to switch if switch is CTO) - if applicable:	
	HPE Aruba Networking 50G eSR 300m MMF Transceiver	S0V64A

Configuration Information

	HPE Aruba Networking 50G LR 10km SMF Transceiver	S0V65A
	HPE Aruba Networking 50G ER 40km SMF Transceiver	S0V66A
	HPE Aruba Networking 50G BiDi 10km-Downstream 1330/1270 Transceiver	S1C92A
	HPE Aruba Networking 50G BiDi 10km-Upstream 1270/1330 Transceiver	S1C94A
	HPE Aruba Networking 50G SFP56 LC SR 100m MMF Transceiver	R0M48A
	HPE Aruba Networking 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable	R0M46A
	HPE Aruba Networking 50G SFP56 to SFP56 3m Direct Attach Copper Cable	R0M47A
	HPE Aruba Networking 50G QSFP56 to SFP56 0.65m DAC Cable	S1J07A
	HPE Aruba Networking 50G QSFP56 to SFP56 3m DAC Cable	S1J08A
5	Localization required on orders without #B2B, #B2C, #B2E or #AC3 options.	
6	If the CTO Switch Chassis needs to be racked, Then the CTO Base Model needs to integrate (with #0D1) to the HPE Network Rack.	
7	When Switches are Factory Racked with this power supply, Then B2B, or B2C should be the Defaulted Power Cable option on the Power Supplies. (See Drop down remark in "Internal Power Supplies" section.)	
8	The following Transceivers install into this Switch: (Use 0D1 quoted to switch if switch is CTO) - if applicable:	
9	The following Transceivers install into this Switch and is only available on LRM Supported ports. See AOS-Switch and AOS-CX Transceiver Guide (Edition 12) for port compatibility: (Use BTO only when adding to switch) (Use 0D1 quoted to switch if switch is CTO) - if applicable:	
	HPE Aruba Networking 10G SFP+ LC LRM 220m OM2 MMF Transceiver	J9152D
10	The following Transceivers install into this Switch using QSFP56 side of cable(Use 0D1 quoted to switch if switch is CTO) - if applicable:	
11	If qty1 of the following QSA28 Adapter(845970-B21) is selected, then increase max SFP28 Port qty by 1 and allow user selection of the following SFP Transceivers. Refer to qty and port restrictions for individual Switch in the "Additional Info" sections: (Use BTO only when adding to switch)	
	HPE Aruba Networking 10G SFP+ LC SR 300m OM3 MMF Transceiver	J9150D
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF Transceiver	J9151E
	HPE Aruba Networking 10G SFP+ LC ER 40km SMF Transceiver	J9153D
	HPE Aruba Networking 10G SFP+ LC SR 300m MMF TAA Transceiver	JL748A
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF TAA Transceiver	JL749A
	HPE Aruba Networking 25G SFP28 LC SR 100m MMF Transceiver	JL484A
	HPE Aruba Networking 25G SFP28 LC eSR 400m MMF Transceiver	JL485A
	HPE Aruba Networking 25G SFP28 LC LR 10km SMF Transceiver	JL486A
	HPE Aruba Networking 25G ER LC 40km SMF Transceiver	S0V69A
	HPE Aruba Networking 25G BiDi 10km-Downstream 1330/1270 Transceiver	S1C96A
	HPE Aruba Networking 25G BiDi 10km-Upstream 1270/1330 Transceiver	S1C98A
	HPE Aruba Networking 25G SFP28 to SFP28 3m Active Optical Cable	R0M44A
	HPE Aruba Networking 25G SFP28 to SFP28 7m Active Optical Cable	R0M45A
	HPE Aruba Networking 25G SFP28 to SFP28 15m Active Optical Cable	R0Z21A
12	The following Transceivers install into this Switch: (Use BTO only when adding to switch)	
	HPE Aruba Networking 100G SR1.2 QSFP28 LC 100m MMF Transceiver	S4B44A
	HPE 100Gb QSFP28 Bidirectional Transceiver	845972-B21
	HPE Aruba Networking 100G SR2 MPO QSFP28 100m MMF Transceiver	S1C93A
	HPE QSFP28 to SFP28 Adapter	845970-B21
	HPE Aruba Networking 100G QSFP28 LC CWDM4 2km SMF Transceiver	R0Z30A
	HPE Aruba Networking 100G QSFP28 LC LR4 10km SMF 2-strand Transceiver	JL310A
	HPE Aruba Networking 100G QSFP28 MPO SR4 100m 12-fiber MPO OM3 MMF Transceiver	JL309A

Configuration Information

	HPE Aruba Networking 100G QSFP28 LC ER4L 40km SMF Transceiver	JL743A
	HPE Aruba Networking 100G QSFP28 LC FR1 SMF 2km Transceiver	R9B63A
	HPE Aruba Networking 100G QSFP28 to QSFP28 3m Direct Attach Copper Cable	JL307A
	HPE Aruba Networking 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable	R0Z25A
	HPE Aruba Networking 100G QSFP28 to QSFP28 5m Direct Attach Copper Cable	R0Z26A
	HPE Aruba Networking 100G QSFP28 to QSFP28 2m Active Optical Cable	JL856A
	HPE Aruba Networking 100G QSFP28 to QSFP28 7m Active Optical Cable	R0Z27A
	HPE Aruba Networking 100G QSFP28 to QSFP28 15m Active Optical Cable	R0Z28A
	HPE Aruba Networking 100G QSFP28 to QSFP28 30m Active Optical Cable	R0Z29A
13	The following Transceivers install into this Switch(Use BTO only when adding to switch)	
	HPE Networking X142 40G QSFP+ MPO SR4 Transceiver	JH231A
	HPE Networking X142 40G QSFP+ LC LR4 SM Transceiver	JH232A
	HPE Networking X142 40G QSFP+ MPO eSR4 300M Transceiver	JH233A
	HPE Aruba Networking 40G QSFP+ LC Bidirectional 150m MMF 2-strand Transceiver	JL308A
	HPE Aruba Networking 40G QSFP+ LC ER4 40km SMF Transceiver	Q9G82A
	HPE Networking X242 40G QSFP+ to QSFP+ 1m Direct Attach Copper Cable	JH234A
	HPE Networking X242 40G QSFP+ to QSFP+ 3m Direct Attach Copper Cable	JH235A
	HPE Networking X242 40G QSFP+ to QSFP+ 5m Direct Attach Copper Cable	JH236A
	HPE Aruba Networking 40G QSFP+ to QSFP+ 7m Active Optical Cable	R0Z22A
	HPE Aruba Networking 40G QSFP+ to QSFP+ 15m Active Optical Cable	R0Z23A
	HPE Aruba Networking 40G QSFP+ to QSFP+ 30m Active Optical Cable	R0Z24A
14	The following Transceivers install into this Switch(Use BTO only when adding to switch)	
	HPE Aruba Networking 25G LR SFP28 LC 10km SMF TAA Transceiver	S2N63A

Notes: Drop down under power supply should offer the following options and results:

- Switch/Router/Power Supply to PDU Power Cord - #B2B in North America, Mexico, Taiwan, and Japan or #B2C ROW. (OCA Default B2B or B2C for Rack Level CTO)
- Switch/Router/Power Supply to Wall Power Cord - Localized Option (OCA Default for BTO)
- High Volt Switch/Router/Power Supply to Wall Power Cord - B2E Option. (Offered only in North America, Mexico, Taiwan, and Japan)
- No Power Cord - #AC3 Option

Locking Power Cord (J9955A) L6-20P is available through the OCA Accessories tab

Transceivers

Remarks:	Description	SKU
	SFP Transceivers	
	HPE Aruba Networking 100M SFP LC FX 2km MMF Transceiver	J9054D
Notes:	Not supported on SFP56 Ports	
	HPE Aruba Networking 1G SFP LC SX 500m OM2 MMF Transceiver	J4858D
	HPE Aruba Networking 1G SFP LC LX 10km SMF Transceiver	J4859D
	HPE Aruba Networking 1G SFP LC LH 70km SMF Transceiver	J4860D
	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e Transceiver	J8177D
	HPE Aruba Networking 1G SFP LC SX 500m MMF TAA Transceiver	JL745A
	HPE Aruba Networking 1G SFP LC LX 10km SMF TAA Transceiver	JL746A
	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e TAA Transceiver	JL747A
	SFP+ Transceivers	
	HPE Aruba Networking 10GBASE-T SFP+ RJ45 30m Cat6A Transceiver	JL563C
	HPE Aruba Networking 10G SFP+ LC SR 300m OM3 MMF Transceiver	J9150D
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF Transceiver	J9151E
	HPE Aruba Networking 10G SFP+ LC ER 40km SMF Transceiver	J9153D

Configuration Information

HPE Aruba Networking 10G SFP+ LC SR 300m MMF TAA Transceiver	JL748A
HPE Aruba Networking 10G SFP+ LC LR 10km SMF TAA Transceiver	JL749A
HPE Aruba Networking 10G SR SFP+ LC 400m OM4 MMF C-Class Transceiver	S2P30A
HPE Aruba Networking 10G LR SFP+ LC 10km SMF C-Class Transceiver	S2P31A
HPE Aruba Networking 10G ER SFP+ LC 40km SMF C-Class Transceiver	S2P32A
HPE Aruba Networking 10G SFP+ to SFP+ 1m Direct Attach Copper Cable • Not compatible with COTA(S0E91A, S0X44A)	J9281D
HPE Aruba Networking 10G SFP+ to SFP+ 3m Direct Attach Copper Cable • Not compatible with COTA(S0E91A, S0X44A)	J9283D

SFP28 Transceivers

HPE Aruba Networking 25G ER LC 40km SMF Transceiver • All models/Ports that support 25G • NOT R8S91A: 51-52 • COTA: YES thru QSA28 (Min0-Max4)	S0V69A
HPE Aruba Networking 25G BiDi 10km-Downstream 1330/1270 Transceiver • All models/Ports that support 25G • NOT R8S91A: 51-52 • COTA: YES thru QSA28 (Min0-Max4)	S1C96A
HPE Aruba Networking 25G BiDi 10km-Upstream 1270/1330 Transceiver • All models/Ports that support 25G • NOT R8S91A: 51-52 • COTA: YES thru QSA28 (Min0-Max4)	S1C98A
HPE Aruba Networking 25G SFP28 LC SR 100m MMF Transceiver	JL484A
HPE Aruba Networking 25G SFP28 LC eSR 400m MMF Transceiver	JL485A
HPE Aruba Networking 25G LR SFP28 LC 10km SMF TAA Transceiver	S2N63A
HPE Aruba Networking 25G SFP28 LC LR 10km SMF Transceiver	JL486A
HPE Aruba Networking 25G SR SFP28 LC 100m MMF C-Class Transceiver	S2P33A
HPE Aruba Networking 25G LR SFP28 LC 10km SMF C-Class Transceiver	S2P34A
HPE Aruba Networking 25G SFP28 to SFP28 0.65m Direct Attach Cable	JL487A
HPE Aruba Networking 25G SFP28 to SFP28 3m Direct Attach Copper Cable	JL488A
HPE Aruba Networking 25G SFP28 to SFP28 5m Direct Attach Copper Cable	JL489A
HPE Aruba Networking 25G SFP28 to SFP28 3m Active Optical Cable	R0M44A
HPE Aruba Networking 25G SFP28 to SFP28 7m Active Optical Cable	R0M45A
HPE Aruba Networking 25G SFP28 to SFP28 15m Active Optical Cable	R0Z21A

QSFP + Transceivers

HPE Networking X142 40G QSFP+ MPO SR4 Transceiver	JH231A
HPE Networking X142 40G QSFP+ LC LR4 SM Transceiver	JH232A
HPE Networking X142 40G QSFP+ MPO eSR4 300M Transceiver	JH233A
HPE Aruba Networking 40G QSFP+ LC Bidirectional 150m MMF 2-strand Transceiver	JL308A
HPE Aruba Networking 40G QSFP+ LC ER4 40km SMF Transceiver	Q9G82A
HPE Networking X242 40G QSFP+ to QSFP+ 1m Direct Attach Copper Cable	JH234A
HPE Networking X242 40G QSFP+ to QSFP+ 3m Direct Attach Copper Cable	JH235A
HPE Networking X242 40G QSFP+ to QSFP+ 5m Direct Attach Copper Cable	JH236A
HPE Aruba Networking 40G QSFP+ to QSFP+ 7m Active Optical Cable	R0Z22A
HPE Aruba Networking 40G QSFP+ to QSFP+ 15m Active Optical Cable	R0Z23A
HPE Aruba Networking 40G QSFP+ to QSFP+ 30m Active Optical Cable	R0Z24A

Notes: Only COTA ports 49-52 (Min0-Max4)

SFP56 Transceivers

HPE Aruba Networking 50G eSR 300m MMF Transceiver	S0V64A

Configuration Information

- All models/Ports that support 50G
- NOT R8S91A: 51-52
- NO Cota models

HPE Aruba Networking 50G LR 10km SMF Transceiver S0V65A

- All models/Ports that support 50G
- NOT R8S91A: 51-52
- NO Cota models

HPE Aruba Networking 50G ER 40km SMF Transceiver S0V66A

- All models/Ports that support 50G
- NOT R8S91A: 51-52
- NO Cota models

HPE Aruba Networking 50G BiDi 10km-Downstream 1330/1270 Transceiver S1C92A

- All models/Ports that support 50G
- NOT R8S91A: 51-52
- NO Cota models

HPE Aruba Networking 50G BiDi 10km-Upstream 1270/1330 Transceiver S1C94A

- All models/Ports that support 50G
- NOT R8S91A: 51-52
- NO Cota models

HPE Aruba Networking 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable R0M46A

- Not supported on SFP

HPE Aruba Networking 50G SFP56 to SFP56 3m Direct Attach Copper Cable R0M47A

HPE Aruba Networking 50G SFP56 LC SR 100m MMF Transceiver R0M48A

HPE Aruba Networking 50G QSFP56 to SFP56 0.65m DAC Cable S1J07A

HPE Aruba Networking 50G QSFP56 to SFP56 3m DAC Cable S1J08A

QSFP28 Transceivers

HPE Aruba Networking 100G SR1.2 QSFP28 LC 100m MMF Transceiver S4B44A

HPE Aruba Networking 100G SR2 MPO QSFP28 100m MMF Transceiver S1C93A

HPE 100Gb QSFP28 Bidirectional Transceiver 845972-B21

HPE Aruba Networking 100G QSFP28 LC CWDM4 2km SMF Transceiver R0Z30A

HPE Aruba Networking 100G QSFP28 LC LR4 10km SMF 2-strand Transceiver JL310A

HPE Aruba Networking 100G QSFP28 MPO SR4 100m 12-fiber MPO OM3 MMF Transceiver JL309A

HPE Aruba Networking 100G QSFP28 LC ER4L 40km SMF Transceiver JL743A

HPE Aruba Networking 100G QSFP28 LC FR1 SMF 2km Transceiver R9B63A

HPE Aruba Networking 100G QSFP28 to QSFP28 3m Direct Attach Copper Cable JL307A

HPE Aruba Networking 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable R0Z25A

HPE Aruba Networking 100G QSFP28 to QSFP28 5m Direct Attach Copper Cable R0Z26A

HPE Aruba Networking 100G QSFP28 to QSFP28 2m Active Optical Cable JL856A

HPE Aruba Networking 100G QSFP28 to QSFP28 7m Active Optical Cable R0Z27A

HPE Aruba Networking 100G QSFP28 to QSFP28 15m Active Optical Cable R0Z28A

HPE Aruba Networking 100G QSFP28 to QSFP28 30m Active Optical Cable R0Z29A

Notes: Only COTA ports 49-52 (Min0-Max4)

QSFP56 Transceivers

HPE Aruba Networking 50G QSFP56 to SFP56 0.65m DAC Cable S1J07A

HPE Aruba Networking 50G QSFP56 to SFP56 3m DAC Cable S1J08A

Notes: Only COTA ports 49-52 (Min0-Max4)

Transceivers for LRM Adapter

HPE Aruba Networking 10G SFP+ LC LRM 220m OM2 MMF Transceiver J9152D

Notes: Compatible on LRM supported ports for R8S91A and R8S92A

Configuration Information

QSA28 Adapter

1	HPE QSFP28 to SFP28 Adapter	845970-B21

Configuration Rules

Rule #	Description
1	If selecting the 845970-B21 - QSFP28 to SFP28 Adapter, then see Tranceiver Guide for details.

Power Supplies

System (std 0 // max 2) User Selection (min 1 // max 2) per enclosure
JL762A, S0G02A System (std 1 // max 2) User Selection (min 0 // max 1) per enclosure

Rule #	Description	SKU
1, 3	HPE Aruba Networking X372 54VDC 1600W 110-240VAC Power Supply • Uses 1 x C15, 1600w	JL670A
	HPE Aruba Networking X372 54VDC 1600W 110-240VAC Power Supply PDU • C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	JL670A#B2B
	HPE Aruba Networking X372 54VDC 1600W 110-240VAC Power Supply PDU • C15 PDU Jumper Cord (ROW) (J9944A)	JL670A#B2C
	HPE Aruba Networking X372 54VDC 1600W 110-240VAC Power Supply 220v • HPE 2.5m C15 to NEMA 6-20P Pwr Cord (JL336A)	JL670A#B2E
	HPE Aruba Networking X372 54VDC 1600W 110-240VAC Power Supply NoLoc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	JL670A#AC3
1, 2	HPE Aruba Networking X371 12VDC 250W 100-240VAC Power Supply • Uses 1 x C13, 250w	JL085A
	HPE Aruba Networking X371 12VDC 250W 100-240VAC Power Supply PDU • C13 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	JL085A#B2B
	HPE Aruba Networking X371 12VDC 250W 100-240VAC Power Supply PDU • C13 PDU Jumper Cord (ROW) (JL697A)	JL085A#B2C
	HPE Aruba Networking X371 12VDC 250W 100-240VAC Power Supply 220v • HPE 2.5m C13 to NEMA 6-20P Pwr Cord (JL336A)	JL085A#B2E
	HPE Aruba Networking X371 12VDC 250W 100-240VAC Power Supply • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	JL085A#AC3
2	HPE Aruba Networking CX 6300M 250W 36-72VDC Input Non-PoE Power Supply • DC supply, comes with DC power cable, only mix with AC PSU JL085A 250W	JL757A
1, 4	HPE Aruba Networking X371 12VDC 250W 100-240VAC Power-to-Port Power Supply • Uses 1 x C13, 250w	JL760A
	HPE Aruba Networking X371 12VDC 250W 100-240VAC Power-to-Port Power Supply PDU • C13 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	JL760A#B2B
	HPE Aruba Networking X371 12VDC 250W 100-240VAC Power-to-Port Power Supply PDU	JL760A#B2C

Configuration Information

	• C13 PDU Jumper Cord (ROW) (J9944A)	
	HPE Aruba Networking X371 12VDC 250W 100-240VAC Power-to-Port Power Supply 220v	JL760A#B2E
	• HPE 2.5m C13 to NEMA 6-20P Pwr Cord (JL336A)	
	HPE Aruba Networking X371 12VDC 250W 100-240VAC Power-to-Port Power Supply NoLoc	JL760A#AC3
	• No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	
1, 3	HPE Aruba Networking X372 54VDC 680W 100-240VAC Power Supply	JL086A
	• Uses 1 x C15, 680w	
	HPE Aruba Networking X372 54VDC 680W 100-240VAC Power Supply PDU	JL086A#B2B
	• C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	
	HPE Aruba Networking X372 54VDC 680W 100-240VAC Power Supply PDU	JL086A#B2C
	• C15 PDU Jumper Cord (ROW) (J9944A)	
	HPE Aruba Networking X372 54VDC 680W 100-240VAC Power Supply 220v	JL086A#B2E
	• HPE 2.5m C15 to NEMA 6-20P Pwr Cord (JL336A)	
	HPE Aruba Networking X372 54VDC 680W 100-240VAC Power Supply	JL086A#AC3
	• No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	
3, 5	HPE Aruba Networking CX 6300M 1050W 36-72VDC Input PoE Power Supply	JL758A
	• DC supply, comes with DC power cable, only mix with AC PSU JL087A 1050W	
1, 3	HPE Aruba Networking X372 54VDC 1050W 110-240VAC Power Supply	JL087A
	• Uses 1 x C15, 1050w	
	HPE Aruba Networking X372 54VDC 1050W 110-240VAC Power Supply PDU	JL087A#B2B
	• C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	
	HPE Aruba Networking X372 54VDC 1050W 110-240VAC Power Supply PDU	JL087A#B2C
	• C15 PDU Jumper Cord (ROW) (J9944A)	
	HPE Aruba Networking X372 54VDC 1050W 110-240VAC Power Supply 220v	JL087A#B2E
	• HPE 2.5m C15 to NEMA 6-20P Pwr Cord (JL336A)	
	HPE Aruba Networking X372 54VDC 1050W 110-240VAC Power Supply	JL087A#AC3
	• No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	

Configuration Rules

Rule #	Description	SKU
1	Localization (Wall Power Cord) required on orders without #B2B, #B2C, (PDU Power Cord) or #B2E. (See Localization Menu)	
2	The Following Switches are compatible with this PSU	
	HPE Aruba Networking CX 6300M 24-port SFP+ and 4-port SFP56 Switch	JL658A
	HPE Aruba Networking CX 6300M 48-port 1GbE and 4-port SFP56 Switch	JL663A
	HPE Aruba Networking CX 6300M 24-port 1GbE and 4-port SFP56 Switch	JL664A
	HPE Aruba Networking CX 6300M 24p SFP+ LRM support and 2p 50G and 2p 25G MACSec Switch	R8S92A
	HPE Aruba Networking CX 6300M 24p SFP+ 1G/10G 4p SFP56 50G TAA Switch	S0G03A

Configuration Information

	HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G TAA Switch	S0G00A
	HPE Aruba Networking CX 6300M 24p 10M/100M/1G 4p SFP56 50G TAA Switch	S0G01A
3	The Following Switches are compatible with this PSU	
	HPE Aruba Networking CX 6300M 48-port HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4-port SFP56 Switch	JL659A
	HPE Aruba Networking CX 6300M 24-port HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4-port SFP56 Switch	JL660A
	HPE Aruba Networking CX 6300M 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch	JL661A
	HPE Aruba Networking CX 6300M 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch	JL662A
	HPE Aruba Networking CX 6300M 48p HPE Smart Rate 1G/2.5G/5G Class8 PoE 2p 50G 2p 25G Switch	R8S90A
	HPE Aruba Networking CX 6300M 24p Smart Rate 1G/2.5G/5G/10G Class6 PoE 2p 50G 2p 25G Switch	R8S89A
	HPE Aruba Networking CX 6300M 12p CL8 PoE 36p CL6 PoE Smart Rate 1G/2.5G/5G 2p 50G 2p 10G Switch	R8S91A
	HPE Aruba Networking CX 6300M 48p Smart Rate 1G/2.5G/5G Class6 PoE 4p SFP56 50G TAA Switch	S0G04A
	HPE Aruba Networking CX 6300M 24p Smart Rate 1G/2.5G/5G Class6 PoE 4p SFP56 50G TAA Switch	S0G05A
	HPE Aruba Networking CX 6300M 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch	S0G06A
	HPE Aruba Networking CX 6300M 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch	S0F99A
	HPE Aruba Networking CX 6300L 48p Smart Rate 100M/1G/2.5G/5G CL8 2p SFP56 50G 2p SFP28 25G L2 Switch	S3L76A
	HPE Aruba Networking CX 6300L 24p Smart Rate 1G/2.5G/5G/10G CL6 2p SFP56 50G 2p SFP28 25G L2 Switch	S3L75A
	HPE ANW CX 6300L 48p Smart Rate 100M/1G/2.5G/5G Class6/8 2p SFP56 50G 2p SFP+ 1G/10G LRM Layer2 Sw	S3L77A
	HPE Aruba Networking CX 6300M 24p SFP+ LRM support and 2p 50G and 2p 25G MACSec Switch	R8S92A
	HPE Aruba Networking CX 6300M 48p SR10 1G/2.5G/5G/10G PTP/AVB Class8 PoE and 4p 100G MACsec Switch	S0E91A
	HPE Aruba Networking CX 6300M 48p SR10 1G/2.5G/5G/10G PTP/AVB Class8 PoE 4p 100G MACsec TAA Switch	S0X44A
4	The Following Switch is only compatible with this Power to Port PSU;	
	HPE Aruba Networking CX 6300M 48p 1GbE 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle	JL762A
	HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G Power-to-Port 2xFan PSU TAA Bundle	S0G02A
5	It is recommended that the JL758A 1050W PSU NOT be mixed with the JL670A 1600W PSU.	

Notes:

If you want the Locking Power Cord (J9955A) L6-20P, then you must order this power cord through the Accessories tab

Drop down under power supply should offer the following options and results:

Switch/Router to PDU Power Cord - #B2B in NA, Mexico, Taiwan, and Japan or #B2C ROW. (OCA Default B2B or B2C for Rack Level CTO)

Switch/Router/Power Supply to Wall Power Cord - Localized Option (OCA Default for BTO and Box Level CTO)

High Volt Switch/Router/Power Supply to Wall Power Cord - #B2E Option. (Offered only in North America, Mexico, Taiwan, and Japan)

No Localized Power Cord Selected - #AC3 Option

PSU Options

Configuration Information

For JL670A, JL085A, JL760A, JL086A, JL087A (std 0 // max 1) User Selection (min 0 // max 1) per PSU

HPE Networking 2.0m C13 to C14 PDU India Power Cord JL671A

- C13 India PDU Cable for Factory Racked Systems Only

HPE Networking 2.5m C15 to C14 PDU India Power Cord JL672A

- C15 India PDU Cable for Factory Racked Systems Only

Notes:
– This cable is intended for India use only. Typically, power cord is ordered when power supply option #AC3 is selected.
– These PDU cables are for Solutions shipping to India.

Switch Options

Fan Trays

- JL660A, JL661A, JL662A, JL663A, JL664A System (std 1 // max 2) User Selection (min 0 // max 1) per enclosure
- S0F99A, S0G00A, S0G01A, S0G03A, S0G04A, S0G05A, S0G06A System (std 1 // max 2) User Selection (min 0 // max 1) per enclosure

Rule#	Description	SKU
	HPE Aruba Networking X751 Front to Back Fan Tray	JL669B
Notes:	When configuring JL658A, JL659A, JL762A, R8S90A, R8S89A, R8S91A, R8S92A, S0E91A, S0X44A, S3L75A, S3L76A, S3L77A, S0G02A	
	When configuring JL660A, JL661A, JL662A, JL663A, JL664A, S0G05A, S0G06A, S0F99A, S0G00A, S0G01A	

Rack Mount Kits

System (std 0 // max 1) User Selection (min 0 // max 1) per enclosure

HPE Aruba Networking X414 1U Universal 4-post Rack Mount Kit J9583B

Notes: If the switch will be factory racked into an HPE Universal Rack, then (Min 1) of the 4 Post Rack Mount kit is required.

Air Duct Kit

For System (std 0 // max 1) User Selection (min 0 // max 1) per Switch

HPE Aruba Networking X544 Universal 4-post Duct Kit (Must order 4-post rack mount kit) JL716A

Notes:
Only for Power to Port Bundles
If the Switch Bundle will be Factory Racked, then this Duct Kit is required with #0D1 for the Power to Port Switch Bundles: JL762A
For optimal performance, it is recommended that the user select the Duct Kit for Power to Port Switch Bundles
If this Air Duct Kit is selected, then the following 4 Post Rack Mount kit must be selected:
J9583B - HPE Aruba Networking X414 1U Universal 4-post Rack Mount Kit

India PDU Cable

For JL665A, JL666A, JL667A, JL668A, JL762A (std 0 // max 1) User Selection (min 0 // max 1) per enclosure

For S0G95A, S0G96A, S0G97A, S0G98A, S0G02A (std 0 // max 1) User Selection (min 0 // max 1) per enclosure

3 HPE Networking 2.0m C13 to C14 PDU India Power Cord JL671A

- C13 India PDU Cable for Factory Racked Systems Only

2 HPE Networking 2.5m C15 to C14 PDU India Power Cord JL672A

- C15 India PDU Cable for Factory Racked Systems Only

Configuration Rules

Rule #	Description	SKU
2	This Power Cord is compatible with the following Switches: JL665A, JL666A, S0G95A, S0G96A, JL762A, S0G02A	

Configuration Information

3	This Power Cord is compatible with the following Switches: JL667A, JL668A, S0G97A, S0G98A	
Notes:	These PDU cables are for Solutions shipping to India.	
	USB Console Cables	
	System (std 0 // max 99) User Selection (min 0 // max 99) per switch	
1	HPE Aruba Networking USBA-RJ45 PIN3TX-6RX 2.5m Cable	R8Z87A
	HPE Aruba Networking USBA-RJ45 PC-to-Switch PIN6TX-3RX 2.5m Cable	R9G48B
1	HPE Aruba Networking USB-A reversible to USB-C PC-to-Switch 3m Cable	R9J32A
1	HPE Aruba Networking USB-C to USB-C PC-to-Switch 3m Cable	R9J33A
	Configuration Rules	
Rule #	**Description**	**SKU**
1	This cable is only compatible with the following Switches; JL658A, JL659A, R8S90A, JL660A, R8S89A, JL661A, R8S91A, JL662A, R8S92A, JL663A, JL664A, JL762A, JL665A, JL666A, JL667A, JL668A, S0E91A, S0F99A, S0G00A, S0G01A, S0G03A, S0G04A, S0G05A, S0G06A, S0X44A, S3L75A, S3L76A, S3L77A, S0G02A, S0G95A, S0G96A, S0G97A, S0G98A	
	Switch Options	
	System (std 0 // max 99) User Selection (min 0 // max 99) per switch	
	HPE Aruba Networking CX Switch Bluetooth Adapter	S1H23A
Notes:	This cable is only compatible with the following Switches; JL658A, JL659A, R8S90A, JL660A, R8S89A, JL661A, R8S91A, JL662A, R8S92A, JL663A, JL664A, JL762A, JL665A, JL666A, JL667A, JL668A, S0E91A, S0F99A, S0G00A, S0G01A, S0G03A, S0G04A, S0G05A, S0G06A, S0X44A, S3L75A, S3L76A, S3L77A, S0G02A, S0G95A, S0G96A, S0G97A, S0G98A	
	Spares	
	System (Std 0 // max 99) User selection (min 0 // max 99) per switch	
1	HPE Aruba Networking CX 6300M Power-to-Port Fan Tray	JL761A
2	HPE Aruba Networking X741 Port-to-Power Fan	JL714A
	Configuration Rules	
Rule #	**Description**	**SKU**
1	This Fan tray is only compatible with the following switches:	
	HPE Aruba Networking CX 6300M 48p 1GbE 4p SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle	JL762A
	HPE Aruba Networking CX 6300M 48p 10M/100M/1G 4p SFP56 50G Power-to-Port 2xFan PSU TAA Bundle	S0G02A
2	This Fan tray is only compatible with the following switches:	
	HPE Aruba Networking CX 6300M 48p SR10 1G/2.5G/5G/10G PTP/AVB Class8 PoE and 4p 100G MACsec Switch	S0E91A
	HPE Aruba Networking CX 6300M 48p SR10 1G/2.5G/5G/10G PTP/AVB Class8 PoE 4p 100G MACsec TAA Switch	S0X44A

Software

HPE Aruba Networking CX Mobile App **https://www.hpe.com/us/en/aruba-cx-switches.html**

Remarks	**Description**	**SKU**
	HPE Aruba Networking AOS-CX Software	
	HPE Aruba Networking CX Advanced Software Licenses	
	HPE Aruba Networking CX Software 63xx Switch Advanced 1-year Subscription E-STU	S0T77AAE
	HPE Aruba Networking CX Software 63xx Switch Advanced 3-year Subscription E-STU	S0T78AAE
	HPE Aruba Networking CX Software 63xx Switch Advanced 5-year Subscription E-STU	S0T79AAE

Configuration Information

	HPE Aruba Networking CX Software 63xx Switch Advanced 7-year Subscription E-STU	S0T80AAE
	HPE Aruba Networking CX Software 63xx Switch Advanced 10-year Subscription E-STU	S0T76AAE
	HPE Aruba Networking NetEdit	
	Single Node Subscription	
	HPE Aruba Networking NetEdit Single Node 1yr Subscription E-STU	JL639AAE
	HPE Aruba Networking NetEdit Single Node 3yr Subscription E-STU	JL640AAE
	HPE Aruba Networking Central	
Notes:	For details and complete listing of HPE Aruba Networking Central licensing options, please see **https://www.hpe.com/psnow/doc/a00049630enw**	
	Cloud Services / 63XX/38XX Switch Foundation Subscriptions	
2	HPE Aruba Networking Central Switch Class-3 Foundation 1 year Subscription E-STU	Q9Y78AAE
2	HPE Aruba Networking Central Switch Class-3 Foundation 3 year Subscription E-STU	Q9Y79AAE
2	HPE Aruba Networking Central Switch Class-3 Foundation 5 year Subscription E-STU	Q9Y80AAE
2	HPE Aruba Networking Central Switch Class-3 Foundation 7 year Subscription E-STU	Q9Y81AAE
2	HPE Aruba Networking Central Switch Class3 Foundation 10-year Subscription E-STU	R3K02AAE
	On-Prem Services / 63XX/38XX Switch Foundation Subscriptions	
3	HPE Aruba Networking Central on Prem Switch Class-3 Foundation 1 year Subscription E-STU	R6U83AAE
3	HPE Aruba Networking Central on Prem Switch Class-3 Foundation 3 year Subscription E-STU	R6U84AAE
3	HPE Aruba Networking Central on Prem Switch Class-3 Foundation 5 year Subscription E-STU	R6U85AAE
3	HPE Aruba Networking Central on Prem Switch Class-3 Foundation 7 year Subscription E-STU	R6U86AAE
3	HPE Aruba Networking Central on Prem Switch Class-3 Foundation 10 year Subscription E-STU	R6U87AAE
	On-Prem Services / 63XX/38XX Switch Foundation Subscriptions	
3	HPE Aruba Networking Central on Prem Switch Class-3 Foundation 1 year Subscription E-STU	R6U83AAE
3	HPE Aruba Networking Central on Prem Switch Class-3 Foundation 3 year Subscription E-STU	R6U84AAE
3	HPE Aruba Networking Central on Prem Switch Class-3 Foundation 5 year Subscription E-STU	R6U85AAE
3	HPE Aruba Networking Central on Prem Switch Class-3 Foundation 7 year Subscription E-STU	R6U86AAE
3	HPE Aruba Networking Central on Prem Switch Class-3 Foundation 10 year Subscription E-STU	R6U87AAE
	On-Prem Services / 63XX/38XX Switch Advanced Subscriptions	
3	HPE Aruba Networking Central On-Premises Switch Class3 Advanced 1-year Subscription E-STU	R6V03AAE
3	HPE Aruba Networking Central On-Premises Switch Class3 Advanced 3-year Subscription E-STU	R6V04AAE
3	HPE Aruba Networking Central On-Premises Switch Class3 Advanced 5-year Subscription E-STU	R6V05AAE
3	HPE Aruba Networking Central On-Premises Switch Class3 Advanced 7-year Subscription E-STU	R6V06AAE

Configuration Information

3	HPE Aruba Networking Central On-Premises Switch Class3 Advanced 10-year Subscription E-STU	R6V07AAE
	FedRAMP Services / 63XX/38XX Switch Foundation Subscriptions	
6	HPE Aruba Networking Central 63xx/38xx Switch Foundation Government 1-year Subscription E-STU	R8K99AAE
6	HPE Aruba Networking Central 63xx/38xx Switch Foundation Government 3-year Subscription E-STU	R8L00AAE
6	HPE Aruba Networking Central 63xx/38xx Switch Foundation Government 5-year Subscription E-STU	R8L01AAE
6	HPE Aruba Networking Central 63xx/38xx Switch Foundation Government 7-year Subscription E-STU	R8L02AAE
6	HPE Aruba Networking Central 63xx/38xx Switch Foundation Government 10-year Subscription E-STU	R8L03AAE
	Configuration Rules	
Rule #	**Description**	**SKU**
2	Add the Central Cloud Skus to the HPE Aruba Networking Catalog as Standalone: HPE Aruba Networking > Network Management > Central > Cloud Services	
3	Add the Central On-Prem Skus to the HPE Aruba Networking Catalog as Standalone: HPE Aruba Networking > Network Management > Central > On-Prem Services	
6	Add the Central FedRAMP Service Skus to the HPE Aruba Networking Catalog as Standalone: HPE Aruba Networking > Network Management > Central > FedRAMP	
	HPE Aruba Networking Fabric Composer	
	Single Node Subscription	
	HPE Aruba Networking Fabric Composer Device Management Service Tier 3 Switch 1y Subscription E-STU	R8D18AAE
	HPE Aruba Networking Fabric Composer Device Management Service Tier 3 Switch 3y Subscription E-STU	R8D19AAE
	HPE Aruba Networking Fabric Composer Device Management Service Tier 3 Switch 5y Subscription E-STU	R8D20AAE

As-a-Service

	HPE Aruba Networking Central	
	Cloud Services / 63XX/38XX Switch Foundation Subscriptions	
	HPE Aruba Networking Central Switch Class-3 Foundation 1 year Subscription SaaS	Q9Y78AAS
	HPE Aruba Networking Central Switch Class-3 Foundation 3 year Subscription SaaS	Q9Y79AAS
	HPE Aruba Networking Central Switch Class-3 Foundation 5 year Subscription SaaS	Q9Y80AAS
	HPE Aruba Networking Central Switch Class-3 Foundation 7 year Subscription SaaS	Q9Y81AAS
	HPE Aruba Networking Central Switch Class-3 Foundation 10 year Subscription SaaS	R3K02AAS
Notes:	Add the Central Cloud Skus to the HPE Aruba Networking Catalog as Standalone: HPE Aruba Networking > Network Management > Central > Cloud Services	
	Cloud Services / Switch Advanced AAS Licenses	
	HPE Aruba Networking Central Switch Class-3 Advanced 1 year Subscription SaaS	S0W52AAS
	HPE Aruba Networking Central Switch Class-3 Advanced 3 year Subscription SaaS	S0W53AAS
	HPE Aruba Networking Central Switch Class-3 Advanced 5 year Subscription SaaS	S0W54AAS
	HPE Aruba Networking Central Switch Class-3 Advanced 7 year Subscription SaaS	S0W55AAS
	HPE Aruba Networking Central Switch Class-3 Advanced 10 year Subscription SaaS	S0W56AAS

Configuration Information

	HPE Aruba Networking Central Switch Class-3 Advanced 1 year Subscription SaaS	S0W77AAS
	HPE Aruba Networking Central Switch Class-3 Advanced 3 year Subscription SaaS	S0W78AAS
	HPE Aruba Networking Central Switch Class-3 Advanced 5 year Subscription SaaS	S0W79AAS
	HPE Aruba Networking Central Switch Class-3 Advanced 7 year Subscription SaaS	S0W80AAS
	HPE Aruba Networking Central Switch Class-3 Advanced 10 year Subscription SaaS	S0W81AAS
Notes:	For IRIS reference only. No action required for OCX and Clic	

Technical Specifications

HPE Aruba Networking CX 6300M 48p SR10 PTP/AVB Class8 PoE 4p 100G MACsec Switch (S0E91A)		
Description	48x ports SmartRate 100M[2]/1G/2.5G/5G/10G BaseT Class 8 PoE ports supporting up to 90W per port (MACsec) 4x 10G/25G/40G/100G QSFP/QSFP28 ports (MACsec) Supports PoE Standards IEEE 802.3af, 802.3at and 802.3bt (up to 90W) 1x USB-C Console Port (higher priority than RJ45 console port) 1x RJ45 console port 1x OOBM 1x USB Type A Host port **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G transceivers and DACs are not supported on S0E91A and S0X44A switch models, which requires QSFP to SFP56 DAC cable for VSF stacking with other CX 6300F and CX 6300M switch models only. VSF stacking not supported on 1G ports.	
Power Supplies	2 field-replaceable, hot-swappable power supply slots 1 minimum power supply required (ordered separately) Supported PSUs JL086A JL087A JL670A JL758A Max PoE Power: 2640W	
Fans	Switch has three fan tray slots and comes with three JL714A fan trays installed • Min 3 fan trays required. • Fan trays are field replaceable and hot-swappable. • Each fan tray contains two fans.	
Physical Characteristics	**Dimensions**	(H) 4.4 cm x (W) 44.2 cm x (D) 47.2 cm (1.73"" x 17.4"" x 18.6"")
	Weight	7.75 kg (17.09 lbs)
Additional Specifications	**CPU**	Quad Core ARM Cortex™ A72 @ 1.8GHz
	Memory and Flash	8 GB DDR4 32 GB eMMC
	Packet Buffer	-
Performance	**Model switching capacity**	1760 Gbps
	Model throughput capacity	1310 Mpps
	Average latency (LIFO, 64-byte packets)	-
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	400 Gbps
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768

Technical Specifications

	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VFR	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C) up to 5,000 ft. Derate -1 degree C for every 1,000 ft from 5,000 ft to 10,000 ft. Can support excursion to 131°F (55°C) for short periods[1] of time.
	Operating relative humidity	5% to 95% @ 104°F (40°C) non-condensing
	Non-operating temperature	-40°F to 158°F (-40°C to 70°C) up to 15,000 ft
	Non-operating humidity	5% to 90% @ 149°F (65°C) non-condensing
	Max operating altitude	10,000 feet (3.04 km) Max
	Max non-operating altitude	15,000 feet (4.6 km) Max
	Acoustic	Sound Power, LWAd = 5.8 Bel Sound Pressure, LpAm (Bystander) = 41.7 dB
	Primary airflow	Front and side-to-back
Electrical Characteristics	**Frequency**	50Hz/60Hz
	AC voltage	JL670A PSU: 110V-120V/200V/208V-240V; AC input JL086A PSU: 100V-240V; AC input JL087A PSU: 110V-240V; AC input JL758A PSU: 36-72VDC; DC input
	Current (for voltages listed above)	JL670A PSU: 11A/9A/8A JL086A PSU: 8A/3.5A JL087A PSU: 12A/5A JL758A PSU: 16.6 - 34.3A
	Power consumption (230VAC)	With single JL086A PSU: Idle: 207W 100% Traffic Rate: 283W With single JL087A PSU: Idle: 208W 100% Traffic Rate: 282W With single JL670A PSU: Idle: 211W 100% Traffic Rate: 283W
Safety	Include US, Canada, Europe, Worldwide	Europe: EN 62368-1:2014 +A11:2017 EN 62368-1:2020 +A11:2020 US: UL 62368-1 2nd Ed. CAN: CSA-C22.2 No. 62368-1-14 2nd Ed. Worldwide: IEC 62368-1:2014 (2nd Ed) IEC 62368-1:2018 (3rd Ed) Taiwan: CNS 15598-1:2020
Emissions	Include US, Canada, Europe, Worldwide	Europe: EN 55032:2015+A11:2020, Class A

		EN 55035:2017+A11:2020 EN IEC 61000-3-2:2019+A1:2021 EN 61000-3-3:2013+A2:2021 US: FCC 47 CFR part 15 subpart B, Class A CAN: ICES-003 Issue 7:2020, Class A Japan: VCCI-CISPR 32:2016, Class A Taiwan: CNS 15936:2016, Class A AUS/NZ: AS/NZS CISPR 32:2015+A1:2020, Class A Worldwide: CISPR 32:2015/AMD1:2019, Class A CISPR 35:2016
Lasers	Include US, Canada, Europe, Worldwide	EN 60825-1:2014 +A11:2021 / IEC 60825-1:2014 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)
Immunity	**Generic**	CISPR 35
	EN	EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2:2018+A1:2020; EN IEC 61000-3-2:2019+A1:2021
	Flicker	IEC/EN 61000-3-3:2013+A2:2021
Mounting and Enclosure	Mounts in an EIAstandard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	

HPE Aruba Networking 6300M 24-port SFP+ and 4-port SFP56 Switch (R8S89A)	
Description	24x ports SmartRate 100M2/1G/2.5G/5G/10G BaseT Class 6 PoE ports supporting up to 60W per port (MACsec) 2x 10G/25G/50G1 SFP ports 2x 10G/25G SFP ports (MACsec) Supports PoE Standards IEEE 802.3af, 802.3at and 802.3bt (up to 60W) 1x USB-C Console Port 1x RJ Console Port 1x OOBM port 1x USB Type A Host port **Notes:** 150G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G transceivers and DACs are not supported on S0E91A and S0X44A switch models, which requires QSFP to SFP56 DAC cable for VSF stacking with other CX 6300F and CX 6300M switch models only. VSF stacking not supported on 1G ports.
Power Supplies	2 field-replaceable, hotswappable power supply slots 1 minimum power supply required (ordered separately) Supported PSUs JL086A JL087A JL670A JL758A

Technical Specifications

	Max PoE Power: 2880W	
Fans	The switch has two fan tray slots and comes with two fan trays installed. Min 2 fan trays required. Fan trays are field replaceable and hotswappable. Each fan tray contains two fans.	
Physical Characteristics	**Dimensions**	(H) 4.4 cm x (W) 44.2 cm x (D) 38.5 cm (1.73" x 17.4" x 15.2")
	Weight	5.26 kg (11.60 lbs)

Mounting and Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
Additional Specifications	**CPU**	Quad Core ARM Cortex™ A72 @ 1.8GHz
	Memory and Flash	8 GB DDR4 32 GB eMMC
	Packet Buffer	16 MB
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	880 Gbps
	Model throughput capacity	654 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 4.24µSec 10Gbps: 1.50µSec 25Gbps: 2.91µSec 50Gbps[1]: 3.49µSec **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VFR	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C) up to 5,000 ft. Derate -1 degree C for every 1,000 ft from 5,000 ft to 10,000 ft. Can support excursion to 131°F (55°C) for short periods[1] of time. 55C excursion not supported when 10G LRM/LR/ER

		inserted: • When 10G BT and 10G LRM/LR/ER transceivers are installed together, fan redundancy is only supported up to 104°F (40°C), 5,000ft **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.
	Operating relative humidity	15% to 95% @ 104°F (40°C) non-condensing
	Non-operating temperature	-40°F to 158°F (-40°C to 70°C) up to 15,000 ft
	Non-operating humidity	15% to 90% @ 149°F (65°C) non-condensing
	Max operating altitude	10,000 feet (3.04 km) Max
	Max non-operating altitude	15,000 feet (4.6 km) Max
	Acoustic	Sound Power, LWAd = 4.9 Bel Sound Pressure, LpAm (Bystander) = 33.0 dB
	Primary airflow	Front and side-to-back
Electrical Characteristics	**Frequency**	50Hz/60Hz
	AC voltage	JL670A PSU: 110V-120V/208V-240V JL086A PSU: 100V-240V JL087A PSU: 110V-240V
	Current (for voltages listed above)	JL670A PSU: 11A/8A JL086A PSU: 8A/3.5A JL087A PSU: 12A/5A
	Power consumption (230VAC)	With JL086A PSU: Idle: 90W 100% Traffic Rate: 143W Idle: 90W 100% Traffic Rate: 140W With JL670A PSU: Idle: 101W 100% Traffic Rate: 152W
Safety	Include US, Canada, Europe, Worldwide	Europe: EN 62368-1:2014 +A11:2017 2nd Ed. EN 62368-1:2020 +A11:2020 3rd Ed. UK: BS EN 62368-1:2014 + A11:2017 2nd Ed BS EN 62368-1:2020 + A11:2020 3rd Ed US/Canada: UL 62368-1 2nd Ed. CAN/CSA-C22.2 No. 62368-1-14 2nd Ed. Worldwide: IEC 60950-1:2005 + Am1:2009 + Am2:2013 w/all known National Deviations IEC 62368-1:2014 2nd Ed. w/all known National Deviations IEC 62368-1:2018 3rd Ed. w/all known National Deviations
Emissions	Include US, Canada,	Europe:

Technical Specifications

	Europe, Worldwide	EN 55032:2015 +A11:2020, Class A EN 55035:2017 +A11:2020 EN 61000-3-2:2014, Class A EN 61000-3-3:2013 US/Canada: FCC CFR47 Part 15:2014, Class A ICES-003 Class A Worldwide: VCCI Class A CISPR 32 Class A CISPR 35:2016
Lasers	Include US, Canada, Europe, Worldwide	EN 60825-1:2014 / IEC 60825-1:2014 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)

Immunity	**Generic**	CISPR 35
	EN	EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000- 3-3
Mounting and Enclosure	Mounts in an EIAstandard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	

HPE Aruba Networking 6300M 48p HPE Smart Rate 1G/2.5G/5G Class8 PoE and 2p 50G and 2p 25G Switch (R8S90A)	
Description	48x ports SmartRate 100M2/1G/2.5G/5G BaseT Class 8 PoE ports supporting up to 90W per port (MACsec) 2x 10G/25G/50G1 SFP ports 2x 10G/25G SFP ports (MACsec) Supports PoE Standards IEEE 802.3af, 802.3at and 802.3bt (up to 90W) 1x USB-C Console Port 1x RJ Console Port 1x OOBM port 1x USB Type A Host port **Notes:** 150G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G transceivers and DACs are not supported on S0E91A and S0X44A switch models, which requires QSFP to SFP56 DAC cable for VSF stacking with other CX 6300F and CX 6300M switch models only. VSF stacking not supported on 1G ports.
Power Supplies	2 field-replaceable, hotswappable power supply slots 1 minimum power supply required (ordered separately) Supported PSUs JL086A JL087A

Technical Specifications

	JL670A JL758A Max PoE Power: 2880W	
Fans	The switch has two fan tray slots and comes with two fan trays installed. Min 2 fan trays required. Fan trays are field replaceable and hotswappable. Each fan tray contains two fans.	
Physical Characteristics	**Dimensions**	(H) 4.4 cm x (W) 44.2 cm x (D) 38.5 cm (1.73" x 17.4" x 15.2")
	Weight	5.48 kg (12.08 lbs)
Mounting and Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
Additional Specifications	**CPU**	Quad Core ARM Cortex™ A72 @ 1.8GHz
	Memory and Flash	8 GB DDR4 32 GB eMMC
	Packet Buffer	16 MB
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	880 Gbps
	Model throughput capacity	654 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 4.24µSec 10Gbps: 1.50µSec 25Gbps: 2.91µSec 50Gbps[1]: 3.49µSec **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VFR	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C) up to 5,000 ft. Derate -1 degree C for every 1,000 ft from 5,000 ft to 10,000 ft. Can support excursion to 131°F (55°C) for short periods[1] of time.

		55C excursion not supported when 10G LRM/LR/ER inserted: • When 10G BT and 10G LRM/LR/ER transceivers are installed together, fan redundancy is only supported up to 104°F (40°C), 5,000ft **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.
	Operating relative humidity	15% to 95% @ 104°F (40°C) non-condensing
	Non-operating temperature	-40°F to 158°F (-40°C to 70°C) up to 15,000 ft
	Non-operating humidity	15% to 90% @ 149°F (65°C) non-condensing
	Max operating altitude	10,000 feet (3.04 km) Max
	Max non-operating altitude	15,000 feet (4.6 km) Max
	Acoustic	Sound Power, LWAd = 5.0 Bel Sound Pressure, LpAm (Bystander) = 33.4 dB
	Primary airflow	Front and side-to-back

Electrical Characteristics	**Frequency**	50Hz/60 Hz
	AC voltage	JL670A PSU: 110V-120V/208V-240V JL086A PSU: 100V-240V JL087A PSU: 110V-240V
	Current (for voltages listed above)	JL670A PSU: 11A/8A JL086A PSU: 8A/3.5A JL087A PSU: 12A/5A
	Power consumption (230VAC)	With JL086A PSU: Idle: 90W 100% Traffic Rate: 143W Idle: 90W 100% Traffic Rate: 140W With JL670A PSU: Idle: 101W 100% Traffic Rate: 152W
Safety	Include US, Canada, Europe, Worldwide	Europe: EN 62368-1:2014 +A11:2017 2nd Ed. EN 62368-1:2020 +A11:2020 3rd Ed. UK: BS EN 62368-1:2014 + A11:2017 2nd Ed BS EN 62368-1:2020 + A11:2020 3rd Ed US/Canada: UL 62368-1 2nd Ed. CAN/CSA-C22.2 No. 62368-1-14 2nd Ed. Worldwide: IEC 60950-1:2005 + Am1:2009 + Am2:2013 w/all known National Deviations IEC 62368-1:2014 2nd Ed. w/all known National Deviations IEC 62368-1:2018 3rd Ed. w/all known National Deviations

Technical Specifications

Emissions	Include US, Canada, Europe, Worldwide	Europe: EN 55032:2015 +A11:2020, Class A EN 55035:2017 +A11:2020 EN 61000-3-2:2014, Class A EN 61000-3-3:2013 US/Canada: FCC CFR47 Part 15:2014, Class A ICES-003 Class A Worldwide: VCCI Class A CISPR 32 Class A CISPR 35:2016
Lasers	Include US, Canada, Europe, Worldwide	EN 60825-1:2014 / IEC 60825-1:2014 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)
Immunity	**Generic**	CISPR 35
	EN	EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000- 3-3
Mounting and Enclosure	Mounts in an EIAstandard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	

HPE Aruba Networking 6300M 48SR5 12p Class8 PoE and 36p Class6 PoE HPE Smart Rate 1G/2.5G/5G and 2p 50G and 2p 10G LRM support Switch (R8S91A)	
Description	48x ports SmartRate 100M/1G/2.5G/5G BaseT Class 8 PoE ports supporting up to 90W per port on ports 1-12, and up to 60W per port on ports 13-48 (MACsec) 2x 10G/25G/50G[1] SFP ports 2x 1G/10G SFP ports (LRM + MACsec) Supports PoE Standards IEEE 802.3af, 802.3at and 802.3bt (up to 90W) 1x USB-C Console Port 1x RJ Console Port 1x OOBM port 1x USB Type A Host port **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.
Power Supplies	2 field-replaceable, hotswappable power supply slots 1 minimum power supply required (ordered separately) Supported PSUs JL086A JL087A JL670A JL758A Max PoE Power: 2880W

Technical Specifications

Fans	The switch has two fan tray slots and comes with two fan trays installed. Min 2 fan trays required. Fan trays are field replaceable and hotswappable. Each fan tray contains two fans.	
Physical Characteristics	**Dimensions**	(H) 4.4 cm x (W) 44.2 cm x (D) 38.5 cm (1.73" x 17.4" x 15.2")
	Weight	5.47 kg (12.06 lbs)
Mounting and Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
Additional Specifications	**CPU**	Quad Core ARM Cortex™ A72 @ 1.8GHz
	Memory and Flash	8 GB DDR4 32 GB eMMC
	Packet Buffer	16 MB
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	880 Gbps
	Model throughput capacity	654 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 4.24µSec 10Gbps: 1.50µSec 25Gbps: 2.91µSec 50Gbps[1]: 3.49µSec **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VFR	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C) up to 5,000 ft. Derate -1 degree C for every 1,000 ft from 5,000 ft to 10,000 ft. Can support excursion to 131°F (55°C) for short periods[1] of time. 55C excursion not supported when 10G LRM/LR/ER inserted:

Technical Specifications

		• When 10G BT and 10G LRM/LR/ER transceivers are installed together, fan redundancy is only supported up to 104°F (40°C), 5,000ft **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.
	Operating relative humidity	15% to 95% @ 104°F (40°C) non-condensing
	Non-operating temperature	-40°F to 158°F (-40°C to 70°C) up to 15,000 ft
	Non-operating humidity	15% to 90% @ 149°F (65°C) non-condensing
	Max operating altitude	10,000 feet (3.04 km) Max
	Max non-operating altitude	15,000 feet (4.6 km) Max
	Acoustic	Sound Power, LWAd = 4.9 Bel Sound Pressure, LpAm (Bystander) = 32.6 dB
	Primary airflow	Front and side-to-back
Electrical Characteristics	**Frequency**	50Hz/60Hz
	AC voltage	JL670A PSU: 110V-120V/208V-240V JL086A PSU: 100V-240V JL087A PSU: 110V-240V
	Current (for voltages listed above)	JL670A PSU: 11A/8A JL086A PSU: 8A/3.5A JL087A PSU: 12A/5A
	Power consumption (230VAC)	With JL086A PSU: Idle: 90W 100% Traffic Rate: 143W With JL087A PSU: Idle: 90W 100% Traffic Rate: 140W With JL670A PSU: Idle: 101W 100% Traffic Rate: 152W
Safety	Include US, Canada, Europe, Worldwide	Europe: EN 62368-1:2014 +A11:2017 2nd Ed. EN 62368-1:2020 +A11:2020 3rd Ed. UK: BS EN 62368-1:2014 + A11:2017 2nd Ed BS EN 62368-1:2020 + A11:2020 3rd Ed US/Canada: UL 62368-1 2nd Ed. CAN/CSA-C22.2 No. 62368-1-14 2nd Ed. Worldwide: IEC 60950-1:2005 + Am1:2009 + Am2:2013 w/all known National Deviations IEC 62368-1:2014 2nd Ed. w/all known National Deviations IEC 62368-1:2018 3rd Ed. w/all known National Deviations
Emissions	Include US, Canada, Europe, Worldwide	Europe: EN 55032:2015 +A11:2020, Class A EN 55035:2017 +A11:2020

		EN 61000-3-2:2014, Class A EN 61000-3-3:2013 US/Canada: FCC CFR47 Part 15:2014, Class A ICES-003 Class A Worldwide: VCCI Class A CISPR 32 Class A CISPR 35:2016
Lasers	Include US, Canada, Europe, Worldwide	EN 60825-1:2014 / IEC 60825-1:2014 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)
Immunity	**Generic**	CISPR 35
	EN	EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000- 3-3
Mounting and Enclosure	Mounts in an EIAstandard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	

HPE Aruba Networking 6300M 24p SFP+ LRM support and 2p 50G and 2p 25G MACsec Switch (R8S92A)		
Description	24x 1G/10G SFP+ ports (LRM + MACsec) 2x 10G/25G/50G[1] SFP ports 2x 10G/25G SFP ports (MACsec) 1x USB-C Console Port 1x RJ Console Port 1x OOBM port 1x USB Type A Host port **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.	
Power Supplies	2 field-replaceable, hotswappable power supply slots 1 minimum power supply required (ordered separately) Supported PSUs JL085A JL757A PSU	
Fans	The switch has two fan tray slots and comes with two fan trays installed. Min 2 fan trays required. Fan trays are field replaceable and hotswappable. Each fan tray contains two fans.	
Physical Characteristics	**Dimensions**	(H) 4.4 cm x (W) 44.2 cm x (D) 38.5 cm (1.73" x 17.4" x 15.2")
	Weight	5.47 kg (12.06 lbs)

Technical Specifications

Mounting and Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
Additional Specifications	**CPU**	Quad Core ARM Cortex™ A72 @ 1.8GHz
	Memory and Flash	8 GB DDR4 32 GB eMMC
	Packet Buffer	16 MB
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	880 Gbps
	Model throughput capacity	654 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 4.24µSec 10Gbps: 1.50µSec 25Gbps: 2.91µSec 50Gbps[1]: 3.49µSec **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VFR	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C) up to 5,000 ft. Derate -1 degree C for every 1,000 ft from 5,000 ft to 10,000 ft. Can support excursion to 131°F (55°C) for short periods[2] of time. 55C excursion not supported when 10G LRM/LR/ER inserted: • When 10G BT and 10G LRM/LR/ER transceivers are installed together, fan redundancy is only supported up to 104°F (40°C), 5,000ft **Notes:** [2]No more than 96 consecutive hours and no more than 360 hours total (15 days) in 1 year.
	Operating relative humidity	15% to 95% @ 104°F (40°C) non-condensing

Technical Specifications

	Non-operating temperature	-40°F to 158°F (-40°C to 70°C) up to 15,000 ft
	Non-operating humidity	15% to 90% @ 149°F (65°C) non-condensing
	Max operating altitude	10,000 feet (3.04 km) Max
	Max non-operating altitude	15,000 feet (4.6 km) Max
	Acoustic	Sound Power, LWAd = 4.6 Bel Sound Pressure, LpAm (Bystander) = 30.1 dB
	Primary airflow	Front and side-to-back
Electrical Characteristics	**Frequency**	50Hz/60Hz
	AC voltage	JL085A PSU: 100V-240V
	Current (for voltages listed above)	JL085A PSU: 3A/1.2A
	Power consumption (230VAC)	Idle: 87W 100% Traffic Rate: 131W
Safety	Include US, Canada, Europe, Worldwide	Europe: EN 62368-1:2014 +A11:2017 2nd Ed. EN 62368-1:2020 +A11:2020 3rd Ed. UK: BS EN 62368-1:2014 + A11:2017 2nd Ed BS EN 62368-1:2020 + A11:2020 3rd Ed US/Canada: UL 62368-1 2nd Ed. CAN/CSA-C22.2 No. 62368-1-14 2nd Ed. Worldwide: IEC 60950-1:2005 + Am1:2009 + Am2:2013 w/all known National Deviations IEC 62368-1:2014 2nd Ed. w/all known National Deviations IEC 62368-1:2018 3rd Ed. w/all known National Deviations
Emissions	Include US, Canada, Europe, Worldwide	Europe: EN 55032:2015 +A11:2020, Class A EN 55035:2017 +A11:2020 EN 61000-3-2:2014, Class A EN 61000-3-3:2013 US/Canada: FCC CFR47 Part 15:2014, Class A ICES-003 Class A Worldwide: VCCI Class A CISPR 32 Class A CISPR 35:2016
Lasers	Include US, Canada, Europe, Worldwide	EN 60825-1:2014 / IEC 60825-1:2014 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)
Immunity	**Generic**	CISPR 35
	EN	EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6

	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000- 3-3
Mounting and Enclosure	Mounts in an EIAstandard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	

HPE Aruba Networking 6300M 24-port SFP+ and 4-port SFP56 Switch (JL658A) **HPE Aruba Networking 6300M 24p SFP+ 1G/10G 4p SFP56 50G TAA Switch** (S0G03A)		
Description	24x 1G/10G SFP+ ports 4x 1G/10G/25G SFP ports	
Additional Ports and Slots	1x USB-C console port 1x OOBM 1x USB Type A host port	
Power Supplies	2 field-replaceable, hot-swappable power supply slots. 1 minimum power supply required (ordered separately) Supports JL085A PSU	
Fan Tray	The switch has two fan tray slots and comes with two fan trays installed. Fan trays are field replaceable and hot-swappable. Minimum 2 fan trays required. Each fan tray is comprised of two fans.	
Physical Characteristics	**Dimensions**	17.4 (w) x 15.2 (d) x 1.73 (h) in (44.2 x 38.5 x 4.4 cm)
	Weight	12.78 lbs (5.8 Kg)
Mounting and Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory and Flash	8 GBytes DDR4 32 GBytes eMMC	
Packet Buffer	8 MB packet buffer memory	
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	880 Gbps
	Model throughput capacity	654 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 1.99µSec 10Gbps: 1.49µSec 25Gbps: 2.85µSec 50Gbps: 2.82µSec
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192

Technical Specifications

	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VFR	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet Can support excursion to 131°F (55°C) for short periods[1] of time. Operating temperature is reduced to 32°F (0°C) to 104°F (40°C) up to 5000ft when 10G SFP+ LR or ER Transceivers are installed.
	Operating relative humidity	15% to 95% relative humidity at 104°F (40°C), non-condensing
	Non-operating temperature	-40°F to 158°F ('-40°C to 70°C)
	Non-operating humidity	15% to 95% relative humidity at 149°F(65°C), non-condensing
	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)
	Acoustic	Sound power, LWAd = 4.9 Bel Sound pressure, LpAm (bystander) = 31.0 dB
	Primary airflow	Front and side-to-back
Electrical Characteristics	**Frequency**	50-60 Hz
	AC voltage	JL085A PSU: 100V-240V
	Current (for voltages listed above)	JL085A PSU: 3A/1.2A
	Power consumption (230VAC)	Idle: 51W 100% traffic rate: 85W
Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013 **US:** UL 60950-1 2nd Ed **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations	
Emissions	**Europe:** EN 55022:2010, Class A EN 55032:2012, Class A EN 55024:2010, EN 61000-3-2:2014, EN 61000-3-3:2013 **US:** FCC part 15 Class A **Canada:** ICES-003 Class A **Worldwide:** VCCI Class A, CISPR 22 Class A CISPR 32 Class A CISPR 24:2010	
Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 24 / CISPR 35

	EN	EN 55024:2010 / EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3
Notes: [1] Not more than 96 consecutive hours or 360 hours total or 15 occurrences in a 1-year period.		

HPE Aruba Networking 6300M 48- port HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4-port SFP56 Switch (JL659A) **HPE Aruba Networking 6300M 48p Smart Rate 1G/2.5G/5G Class6 PoE 4p SFP56 50G TAA Switch** (S0G04A)		
Description	48x ports SmartRate 100M/1G/2.5G/5G BASE-T Class 6 PoE ports supporting up to 60W per port 4x 1G/10G/25G SFP ports	
	Supports PoE Standards IEEE 802.3af, 802.3at and 802.3bt (up to 60W)	
Additional Ports and Slots	1x USB-C console port 1x OOBM 1x USB Type A host port	
Power Supplies	2 field-replaceable, hot-swappable power supply slots 1 minimum power supply required (ordered separately) Supported PSUs: JL086A, JL087A, JL670A Max PoE power: 2880W	
Fan Tray	The switch has two fan tray slots and comes with two fan trays installed. Fan trays are field replaceable and hot-swappable. Minimum 2 fan trays required. Each fan tray is comprised of two fans.	
Physical Characteristics	**Dimensions**	17.4 (w) x 15.2 (d) x 1.73 (h) in (44.2 x 38.5 x 4.4 cm)
	Weight	14.8 lbs (6.71 kg)
Mounting and Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory and Flash	8 GBytes DDR4 32 GBytes eMMC	
Packet Buffer	8 MB packet buffer memory	
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	880 Gbps
	Model throughput capacity	654 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 4.24µSec 10Gbps: 1.50µSec 25Gbps: 2.91µSec 50Gbps: 3.49µSec
	Stack size	10 members

Technical Specifications

	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VRF	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet
	Operating humidity	15% to 95% relative humidity at 104°F (40°C), non-condensing
	Non-operating temperature	-40°F to 158°F ('-40°C to 70°C)
	Non-operating humidity	15% to 95% relative humidity at 149°F(65°C), non-condensing
	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)
	Acoustic	Sound power, LWAd = 4.8 Bel Sound pressure, LpAm (bystander) = 30.6 dB
	Primary airflow	Front and side-to-back
Electrical Characteristics	**Frequency**	50-60 Hz
	AC voltage	JL670A PSU: 110V-120V/208V-240V JL086A PSU: 100V-240V JL087A PSU: 110V-240V
	Current (for voltages listed above)	JL670A PSU: 11A/8A JL086A PSU: 8A/3.5A JL087A PSU: 12A/5A
	Power consumption (230VAC)	With JL086A PSU: Idle: 133W 100% traffic rate: 199W With JL087A PSU: Idle: 138W 100% traffic rate: 193W With JL670A PSU: Idle: 140W 100% traffic rate: 201W

Technical Specifications

Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013 **US:** UL 60950-1 2nd Ed.. **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations	
Emissions	**Europe:** EN 55022:2010, Class A EN 55032:2012, Class A EN 55024:2010, EN 61000-3-2:2014, EN 61000-3-3:2013 **US:** FCC part 15 Class A **Canada:** ICES-003 Class A **Worldwide:** VCCI Class A, CISPR 22 Class A CISPR 32 Class A CISPR 24:2010	
Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 24 / CISPR 35
	EN	EN 55024:2010 / EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3

HPE Aruba Networking 6300M 24- port HPE Smart Rate 1/2.5/5GbE Class 6 PoE and 4-port SFP56 Switch (JL660A) **HPE Aruba Networking 6300M 24p Smart Rate 1G/2.5G/5G Class6 PoE 4p SFP56 50G TAA Switch** (S0G05A)		
Description	24x ports SmartRate 100M/1G/2.5G/5G BASE-T Class 6 PoE ports supporting up to 60W per port 4x 1G/10G/25G SFP ports Supports PoE Standards IEEE 802.3af, 802.3at and 802.3bt (up to 60W)	
Additional Ports and Slots	1x USB-C console port 1x OOBM 1x USB Type A host port	
Power Supplies	2 field-replaceable, hot-swappable power supply slots 1 minimum power supply required (ordered separately) Supported PSUs: JL086A, JL087A, JL670A Max PoE power: 1440W	
Fan Tray	The switch has two fan tray slots and comes with one fan tray installed. Fan trays are field replaceable and hot-swappable. Minimum 1 fan tray required. Second fan tray ordered separately. Each fan tray is comprised of two fans.	
Physical Characteristics	**Dimensions**	17.4 (w) x 15.2 (d) x 1.73 (h) in (44.2 x 38.5 x 4.4 cm)
	Weight	13.36 lbs (6.06 kg)
Mounting and Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory and Flash	8 GBytes DDR4 32 GBytes eMMC	

Technical Specifications

Packet Buffer	8 MB packet buffer memory	
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	640 Gbps
	Model throughput capacity	476 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 4.24µSec 10Gbps: 1.50µSec 25Gbps: 2.91µSec 50Gbps: 3.49µSec
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (egress)	8,192/2,048/8,192
	VRF	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet Can support excursion to 131°F (55°C) for short periods[1] of time. Requires two fan trays to support excursion.
	Operating humidity	15% to 95% relative humidity at 104°F (40°C), non-condensing
	Non-operating temperature	-40°F to 158°F ('-40°C to 70°C)
	Non-operating humidity	15% to 95% relative humidity at 149°F(65°C), non-condensing
	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)
	Acoustic	Sound power, LWAd = 5.2 Bel Sound pressure, LpAm (bystander) = 34.2 dB
	Primary airflow	Front and side-to-back

Technical Specifications

Electrical Characteristics	**Frequency**	50-60 Hz
	AC voltage	JL670A PSU: 110V-120V/208V-240V JL086A PSU: 100V-240V JL087A PSU: 110V-240V
	Current (for voltages listed above)	JL670A PSU: 11A/8A JL086A PSU: 8A/3.5A JL087A PSU: 12A/5A
	Power consumption (230VAC)	With JL086A PSU: Idle: 93W 100% traffic rate: 137W With JL087A PSU: Idle: 91W 100% traffic rate: 131W With JL670A PSU: Idle: 98W 100% traffic rate: 139W
Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013 **US:** UL 60950-1 2nd Ed.. **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations	
Emissions	**Europe:** EN 55022:2010, Class A EN 55032:2012, Class A EN 55024:2010, EN 61000-3-2:2014 EN 61000-3-3:2013 **US:** FCC part 15 Class A **Canada:** ICES-003 Class A **Worldwide:** VCCI Class A, CISPR 22 Class A CISPR 32 Class A CISPR 24:2010	
Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 24 / CISPR 35
	EN	EN 55024:2010 / EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
Immunity	**Voltage dips and interruptions**	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3
Notes: [1] Not more than 96 consecutive hours or 360 hours total or 15 occurrences in a 1-year period.		

HPE Aruba Networking 6300M 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch (JL661A) **HPE Aruba Networking 6300M 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch** (S0G06A)	
Description	48x 10/100/1000 BASE-T PoE+ ports supporting up to 30W per port 4x 1G/10G/25G SFP ports Supports PoE Standards IEEE 802.3af, 802.3at

Technical Specifications

Additional Ports and Slots	1x USB-C console port 1x OOBM 1x USB Type A host port	
Power Supplies	2 field-replaceable, hot-swappable power supply slots 1 minimum power supply required (ordered separately) Supported PSUs: JL086A, JL087A, JL670A Max PoE power: 1600W	
Fan Tray	The switch has two fan tray slots and comes with one fan tray installed. Fan trays are field replaceable and hot-swappable. Minimum 1 fan tray required. Second fan tray ordered separately. Each fan tray is comprised of two fans.	
Physical Characteristics	**Dimensions**	17.4 (w) x 15.2 (d) x 1.73 (h) in (44.2 x 38.5 x 4.4 cm)
	Weight	12.61 lbs (5.72 kg)
Mounting and Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory and Flash	8 GBytes DDR4 32 GBytes eMMC	
Packet Buffer	8 MB packet buffer memory	
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	496 Gbps
	Model throughput capacity	369 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 2.28µSec 10Gbps: 1.46µSec 25Gbps: 1.90µSec 50Gbps: 3.49µSec
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VRF	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet Can support excursion to 131°F (55°C) for short

		periods[1] of time.
	Operating humidity	15% to 95% relative humidity at 104°F (40°C), non-condensing
	Non-operating temperature	-40°F to 158°F ('-40°C to 70°C)
	Non-operating humidity	15% to 95% relative humidity at 149°F(65°C), non-condensing
	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)
	Acoustic	Sound power, LWAd = 4.7 Bel Sound pressure, LpAm (bystander) = 29.8 dB
	Primary airflow	Front and side-to-back
Electrical Characteristics	**Frequency**	50-60 Hz
	AC voltage	JL670A PSU: 110V-120V/208V-240V JL086A PSU: 100V-240V JL087A PSU: 110V-240V
	Current (for voltages listed above)	JL670A PSU: 11A/8A JL086A PSU: 8A/3.5A JL087A PSU: 12A/5A

Electrical Characteristics	**Power consumption (230VAC)**	With JL086A PSU: Idle: 70W 100% traffic rate: 90W With JL087A PSU: Idle: 71W 100% traffic rate: 88W With JL670A PSU: Idle: 73W 100% traffic rate: 96W
Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013 **US:** UL 60950-1 2nd Ed.. **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations	
Emissions	**Europe:** EN 55022:2010, Class A EN 55032:2012, Class A EN 55024:2010, EN 61000-3-2:2014 EN 61000-3-3:2013 **US:** FCC part 15 Class A **Canada:** ICES-003 Class A **Worldwide:** VCCI Class A, CISPR 22 Class A CISPR 32 Class A CISPR 24:2010	
Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 24 / CISPR 35
	EN	EN 55024:2010 / EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4

Technical Specifications

Immunity	**Surge**	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3
Notes: [1] Not more than 96 consecutive hours or 360 hours total or 15 occurrences in a 1-year period.		

HPE Aruba Networking 6300M 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch (JL662A) **HPE Aruba Networking 6300M 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch** (S0F99A)		
Description	24x ports 10/100/1000 BaseT PoE+ ports supporting up to 30W per port 4x 1G/10G/25G/50G[1] SFP ports Supports PoE Standards IEEE 802.3af, 802.3at 1x USB-C Console Port 1x OOBM port 1x USB Type A Host port **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.	
Power Supplies	2 field-replaceable, hot-swappable power supply slots	
	1 minimum power supply required (ordered separately)	
	Supported PSUs: JL086A, JL087A, JL670A	
	Max PoE power: 720W	

Fan Tray	The switch has two fan tray slots and comes with one fan tray installed. Fan trays are field replaceable and hot-swappable. Minimum 1 fan tray required. Second fan tray ordered separately. Each fan tray is comprised of two fans.	
Physical Characteristics	**Dimensions**	17.4 (w) x 15.2 (d) x 1.73 (h) in (44.2 x 38.5 x 4.4 cm)
	Weight	12.23 lbs (5.55 kg)
Mounting and Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory and Flash	8 GBytes DDR4 32 GBytes eMMC	
Packet Buffer	8 MB packet buffer memory	
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	448 Gbps
	Model throughput capacity	334 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 2.28µSec 10Gbps: 1.46µSec 25Gbps: 1.90µSec 50Gbps: 3.49µSec
	Stack size	10 members

Technical Specifications

	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	VRF	256
Performance	**MAC table capacity**	32,768
	IGMP Groups	4,096
	MLD Groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet Can support excursion to 131°F (55°C) for short periods[1] of time.
	Operating humidity	15% to 95% relative humidity at 104°F (40°C), non-condensing
	Non-operating temperature	-40°F to 158°F ('-40°C to 70°C)
	Non-operating humidity	15% to 95% relative humidity at 149°F(65°C), non-condensing
	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)

	Acoustic	Sound power, LWAd = 4.7 Bel Sound pressure, LpAm (bystander) = 29.4 dB
	Primary airflow	Front and side-to-back
Electrical Characteristics	**Frequency**	50-60 Hz
	AC voltage	JL670A PSU: 110V-120V/208V-240V JL086A PSU: 100V-240V JL087A PSU: 110V-240V
	Current (for voltages listed above)	JL670A PSU: 11A/8A JL086A PSU: 8A/3.5A JL087A PSU: 12A/5A
	Power consumption (230VAC)	With JL086A PSU: 100% traffic rate: 76W With JL087A PSU: Idle: 59W 100% traffic rate: 74W With JL670A PSU:

Technical Specifications

		Idle: 62W 100% traffic rate: 81W
Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013 **US:** UL 60950-1 2nd Ed.. **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations	
Emissions	**Europe:** EN 55022:2010, Class A EN 55032:2012, Class A EN 55024:2010, EN 61000-3-2:2014 EN 61000-3-3:2013 **US:** FCC part 15 Class A **Canada: I**CES-003 Class A **Worldwide:** VCCI Class A, CISPR 22 Class A CISPR 32 Class A CISPR 24:2010	
Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 24 / CISPR 35
	EN	EN 55024:2010 / EN 55035:2017
	ESD	IEC 61000-4-2
Immunity	**Radiated**	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3
Notes: [1] Not more than 96 consecutive hours or 360 hours total or 15 occurrences in a 1-year period.		

HPE Aruba Networking 6300M 48-port 1GbE and 4-port SFP56 Switch (JL663A) **HPE Aruba Networking 6300M 48p 10M/100M/1G 4p SFP56 50G TAA Switch** (S0G00A)		
Description	24x ports 10/100/1000 BaseT PoE+ ports supporting up to 30W per port 4x 1G/10G/25G/50G[1] SFP ports 1x USB-C Console Port 1x OOBM port 1x USB Type A Host port **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.	
Power Supplies	2 field-replaceable, hot-swappable power supply slots 1 minimum power supply required (ordered separately) Supports JL085A PSU	
Fan Tray	The switch has two fan tray slots and comes with one fan tray installed. Fan trays are field replaceable and hot-swappable. Minimum 1 fan tray required. The second fan tray ordered separately. Each fan tray is comprised of two fans.	
Physical Characteristics	**Dimensions**	17.4 (w) x 15.2 (d) x 1.73 (h) in (44.2 x 38.5 x 4.4 cm)
	Weight	12.14 lbs (5.51 kg)
Mounting and Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	

Technical Specifications

CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory and Flash	8 GBytes DDR4 32 GBytes eMMC	
Packet Buffer	8 MB packet buffer memory	
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	496 Gbps
	Model throughput capacity	369 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 2.28µSec 10Gbps: 1.46µSec 25Gbps: 1.90µSec 50Gbps: 3.49µSec
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
Performance	**MAC table capacity**	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VRF	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet Can support excursion to 131°F (55°C) for short periods[1] of time.
	Operating humidity	15% to 95% relative humidity at 104°F (40°C), non-condensing
	Non-operating temperature	-40℉ to 158℉ ('-40℃ to 70℃)
	Non-operating humidity	15% to 95% relative humidity at 149°F(65°C), non-condensing
	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)
	Acoustic	Sound power, LWAd = 4.6 Bel Sound pressure,

Technical Specifications

		LpAm (bystander) = 28.7 dB
	Primary airflow	Front and side-to-back
Electrical characteristics	**Frequency**	50-60 Hz
	AC voltage	JL085A PSU: 100V-240V
	Current (for voltages listed above)	JL085A PSU: 3A/1.2A
	Power Consumption (230VAC)	Idle: 56W 100% traffic rate: 75W
Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013 **US:** UL 60950-1 2nd Ed.. **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations	
Emissions	**Europe:** EN 55022:2010, Class A EN 55032:2012, Class A EN 55024:2010, EN 61000-3-2:2014 EN 61000-3-3:2013 **US:** FCC part 15 Class A **Canada:** ICES-003 Class A **Worldwide:** VCCI Class A, CISPR 22 Class A CISPR 32 Class A CISPR 24:2010	
Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 24 / CISPR 35
	EN	EN 55024:2010 / EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3
Notes: [1] Not more than 96 consecutive hours or 360 hours total or 15 occurrences in a 1-year period.		

HPE Aruba Networking 6300M 24-port 1GbE and 4-port SFP56 Switch (JL664A) **HPE Aruba Networking 6300M 24p 10M/100M/1G 4p SFP56 50G TAA Switch** (S0G01A)	
Description	24x ports 10/100/1000 BaseT PoE+ ports supporting up to 30W per port 4x 1G/10G/25G/50G[1] SFP ports 1x USB-C Console Port 1x OOBM port 1x USB Type A Host port **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.
Power supplies	2 field-replaceable, hot-swappable power supply slots 1 minimum power supply required (ordered separately) Supports JL085A PSU

Technical Specifications

Fan tray	The switch has two fan tray slots and comes with one fan tray installed. Fan trays are field replaceable and hot-swappable. Minimum 1 fan tray required. Second fan tray ordered separately. Each fan tray is comprised of two fans.	
Physical characteristics	**Dimensions**	17.4 (w) x 15.2 (d) x 1.73 (h) in (44.2 x 38.5 x 4.4 cm)
	Weight	11.97 lbs (5.43 kg)
Mounting and enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory and Flash	8 GBytes DDR4 32 GBytes eMMC	
Packet Buffer	8 MB packet buffer memory	
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	448 Gbps
	Model throughput capacity	334 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 2.28µSec 10Gbps: 1.46µSec 25Gbps: 1.90µSec 50Gbps: 3.49µSec
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VRF	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet Can support excursion to 131°F (55°C) for short periods[1] of time.
	Operating humidity	15% to 95% relative humidity at 104°F (40°C), non-condensing
Environment	**Non-operating temperature**	-40°F to 158°F ('-40°C to 70°C)

Technical Specifications

	Non-operating humidity	15% to 95% relative humidity at 149°F(65°C), non-condensing
	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)
	Acoustic	Sound power, LWAd = 4.6 Bel Sound pressure, LpAm (bystander) = 28.6 dB
	Primary airflow	Front and side-to-back
Electrical Characteristics	**Frequency**	50-60 Hz
	AC voltage	JL085A PSU: 100V-240V
	Current (for voltages listed above)	JL085A PSU: 3A/1.2A
	Power Consumption (230VAC)	Idle: 49W 100% traffic rate: 64W
Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013 **US:** UL 60950-1 2nd Ed. **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations	
Emissions	**Europe:** EN 55022:2010, Class A EN 55032:2012, Class A EN 55024:2010, EN 61000-3-2:2014 EN 61000-3-3:2013 **US:** FCC part 15 Class A **Canada:** ICES-003 Class A **Worldwide:** VCCI Class A, CISPR 22 Class A CISPR 32 Class A CISPR 24:2010	
Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 24 / CISPR 35
	EN	EN 55024:2010 / EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3
Notes: [1] Not more than 96 consecutive hours or 360 hours total or 15 occurrences in a 1-year period.		

HPE Aruba Networking 6300M 48-port 1GbE and 4-port SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle (JL762A)
HPE Aruba Networking 6300M 48p 10M/100M/1G 4p SFP56 50G Power-to-Port 2xFan PSU TAA Bundle (S0G02A)

Technical Specifications

Description	24x ports 10/100/1000 BaseT PoE+ ports supporting up to 30W per port 4x 1G/10G/25G/50G[1] SFP ports 1x USB-C Console Port 1x OOBM port 1x USB Type A Host port **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.	
Additional ports and slots	1x USB-C console port 1x OOBM 1x USB Type A host portJL	
Power supplies	2 field-replaceable, hot-swappable power supply slots Comes with 1 power-to-port power supply pre-installed Additional power-to-port power supply can be ordered separately Supports JL760A HPE Aruba Networking X371 12VDC 250W 100-240VAC Power-to-Port Power Supply only	
Fan tray	Switch has two fan tray slots and comes with two fan trays installed Fan trays are field replaceable and hot-swappable. Minimum 2 fan trays required. Second fan tray ordered separately Each fan tray is comprised of two fans Supports JL761A HPE Aruba Networking CX 6300M Power-to-Port Fan Tray only	
Physical characteristics	**Dimensions**	17.4 (w) x 15.2 (d) x 1.73 (h) in (44.2 x 38.5 x 4.4 cm)
	Weight	12.5 lbs (5.7 kg), with 1 PSU 13.8 lbs (6.27kg), with 2 PSUs
Mounting and enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory and Flash	8 GBytes DDR4 32 GBytes eMMC	
Packet Buffer	8 MB shared packet buffer memory	
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	496 Gbps
	Model throughput capacity	369 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 2.28µSec 10Gbps: 1.46µSec 25Gbps: 1.90µSec 50Gbps: 3.49µSec
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192

	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VRF	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet Can support excursion to 131°F (55°C) for short periods[1] of time.
	Operating humidity	15% to 95% relative humidity at 104°F (40°C), non-condensing
	Non-operating temperature	-40°F to 158°F ('-40°C to 70°C)
	Non-operating humidity	15% to 90% relative humidity at 149°F(65°C), non-condensing
	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)
	Acoustic	Sound power, LWAd = 4.6 Bel Sound pressure, LpAm (bystander) = 28.7 dB
	Primary airflow	Back-to-front and side
Electrical Characteristics	**Frequency**	50-60 Hz
	AC voltage	JL760A PSU: 100V-240V
	Current (for voltages listed above)	JL760A PSU: 3A/1.2A
	80plus.org certification	TBA for JL760A PS.
	Power Consumption (230VAC)	Idle: 56W 100% traffic rate: 75W
Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013, EN 62368-1:2014 +A11:2017 **US:** UL 60950-1 2nd Ed.. **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations **Taiwan:** CNS-14336-1	
Emissions	**Europe:** EN 55032:2015 +AC:2016, Class A, EN 55035:2017, EN 61000-3-2:2014, EN 61000-3-3:2013 **US:** FCC 47 CFR part 15B, Class A **Canada:** ICES-003 Class A **Worldwide:** VCCI Class A, CISPR 32 Ed 2.0: 2015 + COR1:2016, Class A, CISPR 35:2016	

Technical Specifications

Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 35
	EN	EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3
Notes: [1] Not more than 96 consecutive hours or 360 hours total or 15 occurrences in a 1-year period.		

HPE Aruba Networking 48-port 1GbE Class 4 PoE and 4-port SFP56 Switch (JL665A) **HPE Aruba Networking 6300F 48p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch** (S0G95A)		
Description	48x ports 10/100/1000 BaseT PoE+ ports supporting up to 30W per port 1G/10G/25G/50G[1] SFP ports Supports PoE Standards IEEE 802.3af, 802.3at 1x USB-C Console Port 1x OOBM port 1x USB Type A Host port **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.	
Power Supplies	Internal (fixed) power supply (950W) Max PoE power: 740W	
Fan Tray	Fixed fans	
Physical Characteristics	**Dimensions**	17.4 (w) x 12.9 (d) x 1.73 (h) in (44.2 cm x 32.7 x 4.39 cm))
	Weight	11.24 lbs (5.10 kg)
Mounting And Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory And Flash	8 GBytes DDR4 32 GBytes eMMC	
Packet Buffer	8 MB packet buffer memory	
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	496 Gbps
	Model throughput capacity	369 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 2.28µSec 10Gbps: 1.46µSec 25Gbps: 1.90µSec 50Gbps: 3.49µSec

Technical Specifications

	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VRF	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet Can support excursion to 131°F (55°C) for short periods[1] of time.
	Operating	15% to 95% relative humidity at 104°F (40°C), non-condensing
	Non-operating temperature	40℉ to 158℉ ('-40℃ to 70℃)
	Non-operating humidity	15% to 95% relative humidity at 149°F(65°C), non-condensing
	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)
	Acoustic	Sound power, LWAd = 5.2 Bel Sound pressure, LpAm (bystander) = 34.9 dB
	Primary airflow	Front and side-to-back
Electrical Characteristics	**Frequency**	50-60 Hz
	AC voltage	Fixed PSU: 100V-120V/200V-240V
	Current (for voltages listed above)	Fixed PSU: 11A/6A
	Power Consumption (230VAC)	Hibernation (0 rpm fan): 12W Idle: 63W 100% traffic rate: 86W

Technical Specifications

Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013 **US:** UL 60950-1 2nd Ed.. **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations	
Emissions	**Europe:** EN 55022:2010, Class A EN 55032:2012, Class A EN 55024:2010, EN 61000-3-2:2014 EN 61000-3-3:2013 **US:** FCC part 15 Class A **Canada:** ICES-003 Class A **Worldwide:** VCCI Class A, CISPR 22 Class A CISPR 32 Class A CISPR 24:2010	
Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 24 / CISPR 35
	EN	EN 55024:2010 / EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3
Notes: [1] Not more than 96 consecutive hours or 360 hours total or 15 occurrences in a 1-year period.		

HPE Aruba Networking 6300F 24-port 1GbE Class 4 PoE and 4-port SFP56 Switch (JL666A) **HPE Aruba Networking 6300F 24p 10M/100M/1G Class4 PoE 4p SFP56 50G TAA Switch** (S0G96A)		
Description	24x ports 10/100/1000 BaseT PoE+ ports supporting up to 30W per port 1G/10G/25G/50G[1] SFP ports 1x USB-C Console Port 1x OOBM port 1x USB Type A Host port **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.	
Power Supplies	Internal (fixed) power supply (950W)	
	Max PoE power: 370W	
Fan Tray	Fixed fans	
Physical Characteristics	**Dimensions**	17.4 (w) x 12.9 (d) x 1.73 (h) in (44.2 cm x 32.7 x 4.39 cm))
	Weight	10.91 lbs (4.95 kg)
Mounting And Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory And Flash	8 GBytes DDR4 32 GBytes eMMC	
Packet Buffer	8 MB packet buffer memory	
Performance	**System switching**	880 Gbps

	capacity	
	System throughput capacity	660 Mpps
	Model switching capacity	448 Gbps
	Model throughput capacity	334 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 2.28µSec 10Gbps: 1.46µSec 25Gbps: 1.90µSec 50Gbps: 3.49µSec
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VRF	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet Can support excursion to 131°F (55°C) for short periods[1] of time.
	Operating humidity	15% to 95% relative humidity at 104°F (40°C), non-condensing
	Non-operating temperature	-40℉ to 158℉ ('-40℃ to 70℃)
	Non-operating humidity	15% to 95% relative humidity at 149°F(65°C), non-condensing
	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)
	Acoustic	Sound power, LWAd = 5.0 Bel Sound pressure, LpAm (bystander) = 32.3 dB
	Primary airflow	Front and side-to-back

Technical Specifications

Electrical Characteristics	**Frequency**	50-60 Hz
	AC voltage	Fixed PSU: 100V-120V/200V-240V
	Current (for voltages listed above)	Fixed PSU: 11A/6A
	Power Consumption (230VAC)	Idle: 52W 100% traffic rate: 67W
Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013 **US:** UL 60950-1 2nd Ed **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations	
Emissions	**Europe:** EN 55022:2010, Class A EN 55032:2012, Class A EN 55024:2010, EN 61000-3-2:2014 EN 61000-3-3:2013 **US:** FCC part 15 Class A **Canada:** ICES-003 Class A **Worldwide:** VCCI Class A, CISPR 22 Class A CISPR 32 Class A CISPR 24:2010	
Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 24 / CISPR 35
	EN	EN 55024:2010 / EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3

Notes: [1] Not more than 96 consecutive hours or 360 hours total or 15 occurrences in a 1-year period.

HPE Aruba Networking 6300F 48-port 1GbE and 4-port SFP56 Switch (JL667A) **HPE Aruba Networking 6300F 48p 10M/100M/1G 4p SFP56 50G TAA Switch** (S0G97A)		
Description	48x ports 10/100/1000 BaseT PoE+ ports supporting up to 30W per port 1G/10G/25G/50G[1] SFP ports 1x USB-C Console Port 1x OOBM port 1x USB Type A Host port **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.	
Power Supplies	Internal (fixed) power supply (200W)	
Fan Tray	Fixed fans	
Physical Characteristics	**Dimensions**	17.4 (w) x 12.9 (d) x 1.73 (h) in (44.2 cm x 32.7 x 4.39 cm))
	Weight	9.83 lbs (4.46 kg)
Mounting And Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	

Technical Specifications

CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory And Flash	8 GBytes DDR4 32 GBytes eMMC	
Packet Buffer	8 MB packet buffer memory	
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	496 Gbps
	Model throughput capacity	369 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 2.28µSec 10Gbps: 1.46µSec 25Gbps: 1.90µSec 50Gbps: 3.49µSec
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VRF	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet Can support excursion to 131°F (55°C) for short periods[1] of time.
	Operating humidity	15% to 95% relative humidity at 104°F (40°C), non-condensing
	Non-operating temperature	-40℉ to 158℉ ('-40℃ to 70℃)
	Non-operating humidity	15% to 95% relative humidity at 149°F(65°C), non-condensing
	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)
	Acoustic	Sound power, LWAd = 4.9 Bel

Technical Specifications

		Sound pressure, LpAm (bystander) = 31.5 dB
	Primary airflow	Front and side-to-back
Electrical Characteristics	**Frequency**	50-60 Hz
	AC voltage	Fixed PSU: 100V-120V/200V-240V
	Current (for voltages listed above)	Fixed PSU: 2.5A/1.4A
	Power Consumption (230VAC)	Idle: 52W 100% traffic rate: 74W
Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013 **US:** UL 60950-1 2nd Ed.. **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations	
Emissions	**Europe:** EN 55022:2010, Class A EN 55032:2012, Class A EN 55024:2010, EN 61000-3-2:2014 EN 61000-3-3:2013 **US:** FCC part 15 Class A **Canada:** ICES-003 Class A **Worldwide:** VCCI Class A, CISPR 22 Class A CISPR 32 Class A CISPR 24:2010	
Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 24 / CISPR 35
	EN	EN 55024:2010 / EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3
Notes: [1] Not more than 96 consecutive hours or 360 hours total or 15 occurrences in a 1-year period.		

HPE Aruba Networking 6300F 24-port 1GbE and 4-port SFP56 Switch (JL668A) **HPE Aruba Networking 6300F 24p 10M/100M/1G 4p SFP56 50G TAA Switch** (S0G98A)	
Description	24x ports 10/100/1000 BaseT PoE+ ports supporting up to 30W per port 1G/10G/25G/50G[1] SFP ports 1x USB-C Console Port 1x OOBM port 1x USB Type A Host port **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking. 50G SR transceivers have been added with a minimum software of 10.09.1010. VSF stacking not supported on 1G ports.
Additional Ports And Slots	1x USB-C console port 1x OOBM 1x USB Type A host port

Technical Specifications

Power Supplies	Internal (fixed) power supply (200W)	
Fan Tray	Fixed fans	
Physical Characteristics	**Dimensions**	17.4 (w) x 12.9 (d) x 1.73 (h) in (44.2 cm x 32.7 x 4.39 cm))
	Weight	9.61 lbs (4.36 kg)
Mounting And Enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.	
CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory And Flash	8 GBytes DDR4 32 GBytes eMMC	
Packet Buffer	8 MB packet buffer memory	
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	448 Gbps
	Model throughput capacity	334 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 2.28µSec 10Gbps: 1.46µSec 25Gbps: 1.90µSec 50Gbps: 3.49µSec
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VRF	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet Can support excursion to 131°F (55°C) for short periods[1] of time.
	Operating humidity	15% to 95% relative humidity at 104°F (40°C), non-condensing
	Non-operating temperature	-40℉ to 158℉ ('-40℃ to 70℃)
	Non-operating humidity	15% to 95% relative humidity at 149°F(65°C), non-condensing

Technical Specifications

	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)
	Acoustic	Sound power, LWAd = 4.9 Bel Sound pressure, LpAm (bystander) = 31.6 dB
	Primary airflow	Front and side-to-back
Electrical Characteristics	**Frequency**	50-60 Hz
	AC voltage	Fixed PSU: 100V-120V/200V-240V
	Current (for voltages listed above)	Fixed PSU: 2.5A/1.4A
	Power Consumption (230VAC)	Idle: 49W 100% traffic rate: 63W
Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013 **US:** UL 60950-1 2nd Ed.. **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations	
Emissions	**Europe:** EN 55022:2010, Class A EN 55032:2012, Class A EN 55024:2010, EN 61000-3-2:2014 EN 61000-3-3:2013 **US:** FCC part 15 Class A **Canada:** ICES-003 Class A **Worldwide:** VCCI Class A, CISPR 22 Class A CISPR 32 Class A CISPR 24:2010	
Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 24 / CISPR 35
	EN	EN 55024:2010 / EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3
Notes: [1] Not more than 96 consecutive hours or 360 hours total or 15 occurrences in a 1-year period.		

Standards and Protocols

Applies to all products in series

- ANSI/TIA-1057 LLDP Media Endpoint Discovery (LLDP-MED)
- CPU DoS Protection
- Bootstrap Router (BSR) Mechanism for PIM, PIM WG

Technical Specifications

- Draft-ietf-savi-mix
- IEEE 802.1AB-2005
- IEEE 802.1ak-2007
- IEEE 802.1AX-2008 Link Aggregation
- IEEE 802.1D MAC Bridges
- IEEE 802.1p Priority
- IEEE 802.1Q VLANs
- IEEE 802.1s Multiple Spanning Trees
- IEEE 802.1t-2001
- IEEE 802.1v VLAN classification by Protocol and Port
- IEEE 802.1w Rapid Reconfiguration of Spanning Tree
- IEEE 802.3ab 1000BASE-T
- IEEE 802.3ad Link Aggregation Control Protocol (LACP)
- IEEE 802.3ae 10-Gigabit Ethernet
- IEEE 802.3af Power over Ethernet
- IEEE 802.3at Power over Ethernet
- IEEE 802.3bt Power over Ethernet
- IEEE 802.3az Energy Efficient Ethernet (EEE)
- IEEE 802.3x Flow Control
- IEEE 802.3z 1000BASE-X
- RFC 783 TFTP Protocol (revision 2)
- RFC 791 IP
- RFC 792 ICMP
- RFC 793 TCP
- RFC 813 Window and Acknowledgement Strategy in TCP
- RFC 815 IP datagram reassembly algorithms
- RFC 826 ARP
- RFC 879 TCP maximum segment size and related topics
- RFC 896 Congestion control in IP/TCP internetworks
- RFC 917 Internet subnets
- RFC 919 Broadcasting Internet Datagrams
- RFC 922 Broadcasting Internet Datagrams in the Presence of Subnets (IP_BROAD)
- RFC 925 Multi-LAN address resolution
- RFC 951 BOOTP
- RFC 1027 Proxy ARP
- RFC 1122 Requirements for Internet Hosts - Communications Layers
- RFC 1215 Convention for defining traps for use with the SNMP
- RFC 1256 ICMP Router Discovery Messages
- RFC 1350 TFTP Protocol (revision 2)
- RFC 1393 Traceroute Using an IP Option
- RFC 1403 BGP OSPF Interaction
- RFC 1519 CIDR
- RFC 1542 BOOTP Extensions
- RFC 1583 OSPF Version 2
- RFC 1591 Domain Name System Structure and Delegation
- RFC 1657 Definitions of Managed Objects for BGP-4 using SMIv2
- RFC 1772 Application of the Border Gateway Protocol on the Internet
- RFC 1757 Remote Network Monitoring Management Information Base
- RFC 1812 Requirements for IP Version 4 Router
- RFC 1918 Address Allocation for Private Internet
- RFC 1997 BGP Communities Attribute
- RFC 1998 An Application of the BGP Community Attribute in Multi-home Routing
- RFC 2131 DHCP
- RFC 2132 DHCP Options and BOOTP Vendor Extensions
- RFC 2236 IGMP
- RFC 2328 OSPF Version 2
- RFC 2375 IPv6 Multicast Address Assignments

Technical Specifications

- RFC 2385 Protection of BGP Sessions via the TCP MD5 Signature Option
- RFC 2401 Security Architecture for the Internet Protocol
- RFC 2402 IP Authentication Header
- RFC 2439 BGP Route Flap Damping
- RFC 2460 Internet Protocol, Version 6 (IPv6) Specification
- RFC 2464 Transmission of IPv6 over Ethernet Networks
- RFC 2545 Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing
- RFC 2576 (Coexistence between SNMP V1, V2, V3)
- RFC 2579 (SMIv2 Text Conventions)
- RFC 2580 (SMIv2 Conformance)
- RFC 2710 Multicast Listener Discovery (MLD) for IPv6
- RFC 2711 IPv6 Router Alert Option
- RFC 2787 Definitions of Managed Objects for the Virtual Router Redundancy Protocol
- RFC 2918 Route Refresh Capability for BGP-4
- RFC 2925 Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations (Ping only)
- RFC 2934 Protocol Independent Multicast MIB for IPv4
- RFC 3019 MLDv1 MIB
- RFC 3046 DHCP Relay Agent Information Option
- RFC 3056 Connection of IPv6 Domains via IPv4 Clouds
- RFC 3065 Autonomous System Confederation for BGP
- RFC 3068 An Anycast prefix for 6to4 Relay Route
- RFC 3101 OSPF Not-so-stubby-area option
- RFC 3137 OSPF Stub Router Advertisement sFlow
- RFC 3376 IGMPv3
- RFC 3416 (SNMP Protocol Operations v2)
- RFC 3417 (SNMP Transport Mappings)
- RFC 3418 Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)
- RFC 3484 Default Address Selection for IPv6
- RFC 3509 Alternative Implementations of OSPF Area Border Routers
- RFC 3575 IANA Considerations for RADIUS
- RFC 3623 Graceful OSPF Restart
- RFC 3768 VRRP
- RFC 3810 Multicast Listener Discovery Version 2 (MLDv2) for IPv6
- RFC 3973 PIM Dense Mode
- RFC 4022 MIB for TCP
- RFC 4113 MIB for UDP
- RFC 4213 Basic Transition Mechanisms for IPv6 Hosts and Routers
- RFC 4251 The Secure Shell (SSH) Protocol
- RFC 4252 SSHv6 Authentication
- RFC 4253 SSHv6 Transport Layer
- RFC 4254 SSHv6 Connection
- RFC 4271 A Border Gateway Protocol 4 (BGP-4)
- RFC 4273 Definitions of Managed Objects for BGP-4
- RFC 4291 IP Version 6 Addressing Architecture
- RFC 4292 IP Forwarding Table MIB
- RFC 4293 Management Information Base for the Internet Protocol (IP)
- RFC 4360 BGP Extended Communities Attribute
- RFC 4419 Key Exchange for SSH
- RFC 4443 ICMPv6
- RFC 4456 BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)
- RFC 4486 Subcodes for BGP Cease Notification Message
- RFC 4541 IGMP & MLD Snooping Switch
- RFC 4552 Authentication/Confidentiality for OSPFv3
- RFC 4601 PIM Sparse Mode
- RFC 4607 Source-Specific Multicast for IP

Technical Specifications

- RFC 4675 RADIUS VLAN & Priority
- RFC 4724 Graceful Restart Mechanism for BGP
- RFC 4750 OSPFv2 MIB partial support no SetMIB
- RFC 4760 Multiprotocol Extensions for BGP-4
- RFC 4861 IPv6 Neighbor Discovery
- RFC 4862 IPv6 Stateless Address Auto-configuration
- RFC 4940 IANA Considerations for OSPF
- RFC 5065 Autonomous System Confederation for BGP
- RFC 5095 Deprecation of Type 0 Routing Headers in IPv6
- RFC 5187 OSPFv3 Graceful Restart
- RFC 5340 OSPFv3 for IPv6
- RFC 5424 Syslog Protocol
- RFC 5492 Capabilities Advertisement with BGP-4
- RFC 5519 Multicast Group Membership Discovery MIB (MLDv2 only)
- RFC 5701 IPv6 Address Specific BGP Extended Community Attribute
- RFC 5722 Handling of Overlapping IPv6 Fragments
- RFC 5798 VRRP (exclude Accept Mode and sub-sec timer)
- RFC 5880 Bidirectional Forwarding Detection
- RFC 5905 Network Time Protocol Version 4: Protocol and Algorithms Specification
- RFC 6620 FCFS SAVI
- RFC 6987 OSPF Stub Router Advertisement
- RFC 7047 The Open vSwitch Database Management Protocol
- RFC 7313 Enhanced Route Refresh Capability for BGP-4
- RFC 8201 Path MTU Discovery for IP version 6
- SNMPv1/v2c/v3
- ITU-T Rec G.8032/Y.1344 Mar. 2010
- 2.5G/5GBASE-T (IEEE 802.3bz-2016), 2.5G/5G NBASE-T
- 10GBASE-T (IEEE 802.3an-2006)
- 25-Gigabit Ethernet (IEEE 802.3by-2016, 802.3cc-2017)
- 40-Gigabit Ethernet (IEEE 802.3ba-2010)
- 50-Gigabit Ethernet (IEEE 802.3cd-2018)
- 100-Gigabit Ethernet (IEEE 802.3ba-2010, 802.3bj-2014, 802.3bm-2014)

HPE Aruba Networking Data Center Solution

New HPE Aruba Networking Data Center Solution SKUs

HPE and HPE Aruba Networking offer customers highly differentiated pre-engineered IT infrastructure solutions that span a wide variety **HPE compute, HPE storage, HPE Aruba Networking** that span virtualization, vSAN, HCI, HPC, MCS, Microsoft, SAP HANA, Vmware, Nutanix application and IaaS service offerings. HPE Aruba Networking 10/25 and 40/100G CX switches can be deployed as part of these solutions and is often designed into these integrated solutions along with HPE ProLiant DL/DX servers, SimpliVity, Nimble, Synergy, Cray Shasta, Cray ClusterStor, Superdome Flex and HPE GreenLake. These ready-to-deploy, integrated IT data center solutions help simplify and speed IT service delivery while reducing the time, risk, and expertise needed to deploy complex solutions.

To ensure that these HPE and HPE Aruba Networking integrated solutions receive simplified ordering and the highest-level of customer service and support, HPE Aruba Networking has created a special tracking HPE Aruba Networking Data Center SKUs for HPE deployments that identifies these integrated solutions to ensure they receive rapid support triage and streamlines escalation through HPE Services. Please use these new tracking SKUs when CX switches are included in HPE integrated and mixed HPE compute, HPE storage and HPE Aruba Networking configuration and deployments.

Notes: Current HPE Aruba Networking "J#" SKUs should still be used for all data center network centric (HPE Aruba Networking "only", Non HPE environments). Please contact your sales representative for additional information and ordering guidance.

HPE Aruba Networking CX 6300 DC Switch Series

BTO Models

Rule #	Description	SKU
1, 2, 3, 4, 5, 6	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle • HPE Aruba Networking 6300M 48-port 1GbE and 4-port SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle • Includes 1 Pwr2Prt PSU, can select Min0 / Max1 (250W R9F61A) • Includes 2 Pwr2Prt Fan trays (R9F62A), with no open slots • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	R9F63A
	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle PDU • C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	R9F63A#B2B
	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle PDU • C15 PDU Jumper Cord (ROW) (J9944A)	R9F63A#B2C
	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle 220v • HPE 2.5m C15 to NEMA 6-20P Pwr Cord (JL336A)	R9F63A#B2E
	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle NoLoc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	R9F63A#AC3

Configuration Rules

Rule #	Description	
1	The following Transceivers install into this Switch: (Use BTO only when adding to switch)	
	HPE Aruba Networking 1G SFP LC SX 500m OM2 MMF Transceiver	R9F86A
	HPE Aruba Networking 1G SFP LC LX 10km SMF Transceiver	J4859D
	HPE Aruba Networking 1G SFP LC LH 70km SMF Transceiver	J4860D
	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e Transceiver	R9F87A
	HPE Aruba Networking 1G SFP LC SX 500m MMF TAA Transceiver	R9Q43A
	HPE Aruba Networking 1G SFP LC LX 10km SMF TAA Transceiver	R9Q44A
	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e TAA Transceiver	R9Q45A

HPE Aruba Networking Data Center Solution

2	The following Transceivers install into this Switch: (Use BTO only when adding to switch)	
	HPE Aruba Networking 10GBASE-T SFP+ RJ45 30m Cat6A Transceiver	R9F85A
	HPE Aruba Networking 10G SFP+ LC SR 300m OM3 MMF Transceiver	R9F82A
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF Transceiver	J9151E
	HPE Aruba Networking 10G SFP+ LC ER 40km SMF Transceiver	J9153D
	HPE Aruba Networking 10G SFP+ LC SR 300m MMF TAA Transceiver	R9Q46A
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF TAA Transceiver	R9Q47A
	HPE Aruba Networking 10G SFP+ to SFP+ 1m Direct Attach Copper Cable	R9F83A
	HPE Aruba Networking 10G SFP+ to SFP+ 3m Direct Attach Copper Cable	R9F84A
3	The following Transceivers install into this Switch: (Use BTO only when adding to switch)	
	HPE Aruba Networking 25G SFP28 LC SR 100m MMF Transceiver	R9F89A
	HPE Aruba Networking 25G SFP28 LC eSR 400m MMF Transceiver	R9F90A
	HPE Aruba Networking 25G SFP28 LC LR 10km SMF Transceiver	JL486A
	HPE Aruba Networking 25G SFP28 to SFP28 0.65m Direct Attach Copper Cable	R9F91A
	HPE Aruba Networking 25G SFP28 to SFP28 3m Direct Attach Copper Cable	R9F92A
	HPE Aruba Networking 25G SFP28 to SFP28 5m Direct Attach Copper Cable	R9F93A
4	The following Transceivers install into this Switch: (Use BTO only when adding to switch)	
	HPE Aruba Networking 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable	R9G06A
	HPE Aruba Networking 50G SFP56 to SFP56 3m Direct Attach Copper Cable	R9G07A
5	Localization required on orders without B2B, B2C, B2E or AC3 options.	
6	The following Transceivers install into this Switch: (Use BTO only when adding to switch)	
	HPE Aruba Networking 100M SFP LC FX 2km MMF Transceiver	R9F88A

Notes:

- Drop down under power supply should offer the following options and results:
  - Switch/Router/Power Supply to PDU Power Cord - B2B in North America, Mexico, Taiwan, and Japan or B2C ROW. (Watson Default B2B or B2C for Rack Level CTO)
  - Switch/Router/Power Supply to Wall Power Cord - Localized Option (Watson Default for BTO and Box Level CTO)
- High Volt Switch/Router/Power Supply to Wall Power Cord - B2E Option. (Offered only in North America, Mexico, Taiwan, and Japan)
- No Power Cord - AC3 Option
- Locking Power Cord (J9955A) L6-20P is available through the OCA Accessories tab
- OCA Only Model Selection Form - HPE Offering > HPE Aruba Networking > DC Solutions Switches for HPE > AOS-CX: HPE Aruba Networking 6300 DC Switch Series

Rack Level Integration CTO Models

Rule #	Description	SKU
1, 2, 3, 4, 5, 6, 7	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle • HPE Aruba Networking 6300M 48-port 1GbE and 4-port SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle • Includes 1 Pwr2Prt PSU, can select Min0 / Max1 (250W R9F61A) • Includes 2 Pwr2Prt Fan trays (R9F62A), with no open slots • Min=0 \ Max = 4 SFP/SFP+/SFP28/SFP56 1/10/25/50G Transceiver • 1U - Height	R9F63A
	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle PDU • C15 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	R9F63A#B2B
	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle PDU • C15 PDU Jumper Cord (ROW) (J9944A)	R9F63A#B2C
	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle 220v	R9F63A#B2E

HPE Aruba Networking Data Center Solution

- HPE 2.5m C15 to NEMA 6-20P Pwr Cord (JL336A)

HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle NoLoc R9F63A#AC 3

- No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)

Configuration Rules

Rule #	Description	
1	The following Transceivers install into this Switch: (Use 0D1 or B01 quoted to switch if switch is CTO) - if applicable:	
	HPE Aruba Networking 1G SFP LC SX 500m OM2 MMF Transceiver	R9F86A
	HPE Aruba Networking 1G SFP LC LX 10km SMF Transceiver	J4859D
	HPE Aruba Networking 1G SFP LC LH 70km SMF Transceiver	J4860D
	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e Transceiver	R9F87A
	HPE Aruba Networking 1G SFP LC SX 500m MMF TAA Transceiver	R9Q43A
	HPE Aruba Networking 1G SFP LC LX 10km SMF TAA Transceiver	R9Q44A
	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e TAA Transceiver	R9Q45A
2	The following Transceivers install into this Switch: (Use 0D1 quoted to switch if switch is CTO) - if applicable:	
	HPE Aruba Networking 10GBASE-T SFP+ RJ45 30m Cat6A Transceiver	R9F85A
	HPE Aruba Networking 10G SFP+ LC SR 300m OM3 MMF Transceiver	R9F82A
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF Transceiver	J9151E
	HPE Aruba Networking 10G SFP+ LC ER 40km SMF Transceiver	J9153D
	HPE Aruba Networking 10G SFP+ LC SR 300m MMF TAA Transceiver	R9Q46A
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF TAA Transceiver	R9Q47A
	HPE Aruba Networking 10G SFP+ to SFP+ 1m Direct Attach Copper Cable	R9F83A
	HPE Aruba Networking 10G SFP+ to SFP+ 3m Direct Attach Copper Cable	R9F84A
3	The following Transceivers install into this Switch: (Use 0D1 or B01 quoted to switch if switch is CTO) - if applicable:	
	HPE Aruba Networking 25G SFP28 LC SR 100m MMF Transceiver	R9F89A
	HPE Aruba Networking 25G SFP28 LC eSR 400m MMF Transceiver	R9F90A
	HPE Aruba Networking 25G SFP28 LC LR 10km SMF Transceiver	JL486A
	HPE Aruba Networking 25G SFP28 to SFP28 0.65m Direct Attach Copper Cable	R9F91A
	HPE Aruba Networking 25G SFP28 to SFP28 3m Direct Attach Copper Cable	R9F92A
	HPE Aruba Networking 25G SFP28 to SFP28 5m Direct Attach Copper Cable	R9F93A
4	The following Transceivers install into this Switch: (Use 0D1 or B01 quoted to switch if switch is CTO) - if applicable:	
	HPE Aruba Networking 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable	R9G06A
	HPE Aruba Networking 50G SFP56 to SFP56 3m Direct Attach Copper Cable	R9G07A
5	Localization required on orders without B2B, B2C, B2E or AC3 options.	
6	If the CTO Switch Chassis needs to be racked, Then the CTO Base Model needs to integrate (with 0D1) to the HPE Network Rack.	
7	When Switches are Factory Racked with this power supply, Then B2B, or B2C should be the Defaulted Power Cable option on the Power Supplies. (See Drop down remark in "Power Supplies" section.)	
8	The following Transceivers install into this Switch: (Use BTO only when adding to switch) R9F88A - HPE Aruba Networking 100M SFP LC FX 2km MMF Transceiver for HPE	
	HPE Aruba Networking 100M SFP LC FX 2km MMF Transceiver	R9F88A

Notes:

Drop down under power supply should offer the following options and results:

- o Switch/Router/Power Supply to PDU Power Cord - B2B in North America, Mexico, Taiwan, and Japan or B2C ROW. (OCA Default B2B or B2C for Rack Level CTO)
- o Switch/Router/Power Supply to Wall Power Cord - Localized Option (OCA Default for BTO)

High Volt Switch/Router/Power Supply to Wall Power Cord - B2E Option. (Offered only in North America, Mexico, Taiwan, and Japan)

No Power Cord - AC3 Option

Locking Power Cord (J9955A) L6-20P is available through the OCA

Transceivers

SFP Transceivers

Rule #	Description	SKU
	HPE Aruba Networking 100M SFP LC FX 2km MMF Transceiver	R9F88A
Notes:	Not supported in the SFP56 Ports	
	HPE Aruba Networking 1G SFP LC SX 500m OM2 MMF Transceiver	R9F86A
	HPE Aruba Networking 1G SFP LC LX 10km SMF Transceiver	J4859D
	HPE Aruba Networking 1G SFP LC LH 70km SMF Transceiver	J4860D
	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e Transceiver	R9F87A
	HPE Aruba Networking 1G SFP LC SX 500m MMF TAA Transceiver	R9Q43A
	HPE Aruba Networking 1G SFP LC LX 10km SMF TAA Transceiver	R9Q44A
	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e TAA Transceiver	R9Q45A

SFP+ Transceivers

Rule #	Description	SKU
1	HPE Aruba Networking 10GBASE-T SFP+ RJ45 30m Cat6A Transceiver	R9F85A
	HPE Aruba Networking 10G SFP+ LC SR 300m OM3 MMF Transceiver	R9F82A
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF Transceiver	J9151E
	HPE Aruba Networking 10G SFP+ LC ER 40km SMF Transceiver	J9153D
	HPE Aruba Networking 10G SFP+ LC SR 300m MMF TAA Transceiver	R9Q46A
	HPE Aruba Networking 10G SFP+ LC LR 10km SMF TAA Transceiver	R9Q47A
	HPE Aruba Networking 10G SFP+ to SFP+ 1m Direct Attach Copper Cable	R9F83A
	HPE Aruba Networking 10G SFP+ to SFP+ 3m Direct Attach Copper Cable	R9F84A

SFP28 Transceivers

Rule #	Description	SKU
	HPE Aruba Networking 25G SFP28 LC SR 100m MMF Transceiver	R9F89A
	HPE Aruba Networking 25G SFP28 LC eSR 400m MMF Transceiver	R9F90A
	HPE Aruba Networking 25G SFP28 LC LR 10km SMF Transceiver	JL486A
	HPE Aruba Networking 25G SFP28 to SFP28 0.65m Direct Attach Copper Cable	R9F91A
	HPE Aruba Networking 25G SFP28 to SFP28 3m Direct Attach Copper Cable	R9F92A
	HPE Aruba Networking 25G SFP28 to SFP28 5m Direct Attach Copper Cable	R9F93A

SFP56 Transceivers

Rule #	Description	SKU
	HPE Aruba Networking 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable	R9G06A
	HPE Aruba Networking 50G SFP56 to SFP56 3m Direct Attach Copper Cable	R9G07A

Switch Options

Fan Trays

R9F63A System (std 2 // max 2) User Selection (min 0 // max 0) per enclosure

Rule #	Description	SKU
1	HPE Aruba Networking CX 6300M Power-to-Port Airflow Fan Tray	R9F62A
Notes:	Spare only	

Configuration Rules

Rule #	Description	
1	The Following Switch is only compatible with this Power to Port FanTray;	
	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle	R9F63A
Notes:	This switch includes 2 Fan Trays. No additional Fan Trays necessary"	

Rack Mount Kits

System (std 0 // max 1) User Selection (min 0 // max 1) per enclosure

	HPE Aruba Networking 1U Universal 4-post Rack Mount Kit	R9F57A
Notes:	If the switch will be factory racked into an HPE Universal Rack, then (Min 1) of the 4 Post Rack Mount kit is required and should nest to Rack.	

Air Duct Kit

HPE Aruba Networking Data Center Solution

For System (std 0 // max 1) User Selection (min 0 // max 1) per Switch

1, 2, 3	HPE Aruba Networking Universal 4-post Duct Kit	R9F60A
Notes:	Only for Power to Port Bundles	

Configuration Rules

Rule #	Description
1	The Following Switch is only compatible with this Power to Port FanTray; HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle R9F63A
2	For optimal performance, it is recommended that the user select the Duct Kit for Power to Port Switch Bundles
3	If this Air Duct Kit is selected, then the following 4 Post Rack Mount kit must be selected: HPE Aruba Networking 1U Universal 4-post Rack Mount Kit R9F57A

India PDU Cable

For R9F63A (std 0 // max 1) User Selection (min 0 // max 1) per enclosure

1	HPE Networking 2.0m C13 to C14 PDU India Power Cord • C13 India PDU Cable for Factory Racked Systems Only	JL671A
1	HPE Networking 2.5m C15 to C14 PDU India Power Cord • C15 India PDU Cable for Factory Racked Systems Only	JL672A

Configuration Rules

Rule #	Description
1	This Power Cord is only valid when the #AC3 option is selected for the supported Switch Enclosure
Notes:	This PDU cable is for Solutions shipping to India.

Power Supplies

Power Supply Units

R9F63A System (std 1 // max 2) User Selection (min 0 // max 1) per enclosure

1, 2	HPE Aruba Networking CX 6300M 12VDC 250W 100-240VAC Power-to-Port Airflow Power Supply Unit • Uses 1 x C13, 250w	R9F61A
	HPE Aruba Networking CX 6300M 12VDC 250W 100-240VAC Power-to-Port Airflow Power Supply Unit PDU • C13 PDU Jumper Cord (NA/MEX/TW/JP) (J9943A)	R9F61A#B2B
	HPE Aruba Networking CX 6300M 12VDC 250W 100-240VAC Power-to-Port Airflow Power Supply Unit PDU • C13 PDU Jumper Cord (ROW) (J9944A)	R9F61A#B2C
	HPE Aruba Networking CX 6300M 12VDC 250W 100-240VAC Power-to-Port Airflow Power Supply Unit 220v • HPE 2.5m C13 to NEMA 6-20P Pwr Cord (JL336A)	R9F61A#B2E
	HPE Aruba Networking CX 6300M 12VDC 250W 100-240VAC Power-to-Port Airflow Power Supply Unit NoLoc • No Localized Power Cord Selected. Use J9955A to obtain a Locking Plug Power Cord (L6-20P)	R9F61A#AC3

Configuration Rules

Rule #	Description
1	Localization (Wall Power Cord) required on orders without B2B, B2C, (PDU Power Cord) or B2E. (See Localization Menu)
2	The Following Switch is only compatible with this Power to Port PSU: HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle R9F63A
Notes:	– Drop down under power supply should offer the following options and results: o Switch/Router to PDU Power Cord - B2B in NA, Mexico, Taiwan, and Japan or B2C ROW. (OCA Default B2B or B2C for Rack Level CTO) o Switch/Router/Power Supply to Wall Power Cord - Localized Option (OCA Default for BTO and Box Level CTO) – High Volt Switch/Router/Power Supply to Wall Power Cord - B2E Option. (Offered only in North America, Mexico, Taiwan, and Japan)

HPE Aruba Networking Data Center Solution

– No Localized Power Cord Selected - AC3 Option
– If you want the Locking Power Cord (J9955A) L6-20P, then you must order this power cord through the Accessories tab

PSU Options

For R9F61A (std 0 // max 1) User Selection (min 0 // max 1) per PSU

HPE Networking 2.0m C13 to C14 PDU India Power Cord JL671A
- C13 India PDU Cable for Factory Racked Systems Only

HPE Networking 2.5m C15 to C14 PDU India Power Cord JL672A
- C15 India PDU Cable for Factory Racked Systems Only

Notes:
– This Power Cord is only valid when the #AC3 option is selected for the supported Power Supply
– This PDU cable is for Solutions shipping to India.

Software

HPE Aruba Networking Net Edit

Single Node Subscription for Data Center Solutions

HPE Aruba Networking NetEdit Single Node 1-year Subscription E-STU R9G36AAE
HPE Aruba Networking NetEdit Single Node 3-year Subscription E-STU R9G37AAE

HPE Aruba Networking Fabric Composer

Single Node Subscription for Data Center Solutions

HPE Aruba Networking Fabric Composer Device Management Service Tier 3 Switch 1y Subscription E-STU R9G31AAE
HPE Aruba Networking Fabric Composer Device Management Service Tier 3 Switch 3y Subscription E-STU R9G32AAE
HPE Aruba Networking Fabric Composer Device Management Service Tier 3 Switch 5y Subscription E-STU R9G33AAE

Technical Specifications

HPE Aruba Networking 6300M 48-port 1GbE and 4-port SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle (R9F63A)		
Description	48x 10/100/1000 BASE-T PoE+ ports 1G/10G/25G[1] SFP ports **Notes:** [1]50G capability is for use with 50G DACs for both interconnect and VSF stacking.	
Additional ports and slots	1x USB-C console port 1x OOBM 1x USB Type A host portJL	
Power supplies	2 field-replaceable, hot-swappable power supply slots Comes with 1 power-to-port power supply pre-installed Additional power-to-port power supply can be ordered separately Supports R9F61A HPE Aruba Networking X371 12VDC 250W 100-240VAC Power-to-Port Power Supply only	
Fan tray	Switch has two fan tray slots and comes with two fan trays installed Fan trays are field replaceable and hot-swappable. Minimum 2 fan trays required. Second fan tray ordered separately Each fan tray is comprised of two fans Supports R9F62A HPE Aruba Networking CX 6300M Power-to-Port Fan Tray only	
Physical characteristics	**Dimensions**	17.4 (w) x 15.2 (d) x 1.73 (h) in (44.2 x 38.5 x 4.4 cm)
	Weight	12.5 lbs (5.7 kg), with 1 PSU 13.8 lbs (6.27kg), with 2 PSUs

Mounting and enclosure	Mounts in an EIA- standard 19 in. telco rack or equipment cabinet. Horizontal surface mounting only. 2-post rack kit included.

HPE Aruba Networking Data Center Solution

CPU	Quad Core ARM Cortex™ A72 @ 1.8GHz	
Memory and Flash	8 GBytes DDR4 32 GBytes eMMC	
Packet Buffer	8 MB shared packet buffer memory	
Performance	**System switching capacity**	880 Gbps
	System throughput capacity	660 Mpps
	Model switching capacity	496 Gbps
	Model throughput capacity	369 Mpps
	Average latency (LIFO, 64-byte packets)	1Gbps: 2.28µSec 10Gbps: 1.46µSec 25Gbps: 1.90µSec 50Gbps: 3.49µSec
	Stack size	10 members
	Max stacking distance	Up to 10 kms with long range transceivers
	Stacking bandwidth	200 Gbps (400 Gbps at full duplex)
	Switched virtual interfaces (dual stack)	1,024
	IPv4 host table (ARP)	49,152
	IPv6 host table (ND)	49,152
	IPv4 unicast routes	61,000
	IPv6 unicast routes	61,000
	IPv4 multicast routes	8,192
	IPv6 multicast routes	8,192
	MAC table capacity	32,768
	IGMP groups	4,096
	MLD groups	8,192
	IPv4/IPv6/MAC ACL entries (ingress)	20,480/5,120/20,480
	IPv4/IPv6/MAC ACL entries (engress)	8,192/2,048/8,192
	VRF	256
Environment	**Operating temperature**	32°F to 113°F (0°C to 45°C), up to 5,000 feet 32°F to 104°F (0°C to 40°C), 5,001 to 10,000 feet 1°C de-rating per 1,000 feet above 5,000 feet Can support excursion to 131°F (55°C) for short periods[1] of time.
	Operating humidity	15% to 95% relative humidity at 104°F (40°C), non-condensing
	Non-operating temperature	-40°F to 158°F ('-40°C to 70°C)
	Non-operating humidity	15% to 90% relative humidity at 149°F(65°C), non-condensing
	Max operating altitude	Up to 10,000ft (3.048 Km)
	Max non-operating altitude	Up to 15,000ft (3.048 Km)
	Acoustic	Sound power, LWAd = 4.6 Bel

		Sound pressure, LpAm (bystander) = 28.7 dB
	Primary airflow	Back-to-front and side

Electrical Characteristics	**Frequency**	50-60 Hz
	AC voltage	R9F61A PSU: 100V-240V
	Current (for voltages listed above)	R9F61A PSU: 3A/1.2A
	Power Consumption (230VAC)	Hibernation (0 rpm fan): 9W Idle: 56W 100% traffic rate: 75W
Safety	**Europe:** EN 60950-1:2006 +A11:2009 +A1:2010 +A12:2011 + A2:2013, EN 62368-1:2014 +A11:2017 **US:** UL 60950-1 2nd Ed.. **Canada:** CAN/CSA-C22.2 No. 60950-1-07 **Worldwide:** IEC 60950-1:2005 w/all known National Deviations **Taiwan:** CNS-14336-1	
Emissions	**Europe:** EN 55032:2015 +AC:2016, Class A, EN 55035:2017, EN 61000-3-2:2014, EN 61000-3-3:2013 **US:** FCC 47 CFR part 15B, Class A **Canada:** ICES-003 Class A **Worldwide:** VCCI Class A, CISPR 32 Ed 2.0: 2015 + COR1:2016, Class A, CISPR 35:2016	
Lasers	EN 60825-1:2007 / IEC 60825-1:2007 Class 1 Class 1 Laser Products / Laser Klasse 1 (Applicable for accessories - Optical Transceivers only)	
Immunity	**Generic**	CISPR 35
	EN	EN 55035:2017
	ESD	IEC 61000-4-2
	Radiated	IEC 61000-4-3
	EFT/Burst	IEC 61000-4-4
	Surge	IEC 61000-4-5
	Conducted	IEC 61000-4-6
	Power frequency magnetic field	IEC 61000-4-8
	Voltage dips and interruptions	IEC 61000-4-11
	Harmonics	IEC 61000-3-2, EN 61000-3-2
	Flicker	IEC 61000-3-3, EN 61000-3-3

Summary of Changes

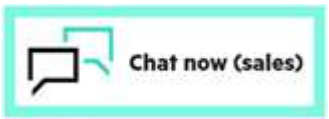
Date	Version History	Action	Description of Change
06-Jan-2025	Version 35	Changed	Configuration Information section was updated.
18-Nov-2024	Version 34	Changed	QuickSpecs was updated.
03-Sep-2024	Version 33	Changed	Configuration Information section was updated.
03-Jun-2024	Version 32	Changed	Configuration Information section was updated.
01-Apr-2024	Version 31	Changed	Configuration Information section was updated.
04-Dec-2023	Version 30	Changed	Obsolete SKU was removed. Configuration Information section was updated. Series name was updated.
06-Nov-2023	Version 29	Changed	Overview, Standard Features, Configuration Information, Technical Specifications sections were updated
07-Aug-2023	Version 28	Changed	Configuration Information section was updated.
10-Jul-2023	Version 27	Changed	Configuration Information section was updated.
15-May-2023	Version 26	Changed	Configuration Information section was updated.
13-Mar-2023	Version 25	Changed	Configuration Information section was updated.
06-Feb-2023	Version 24	Changed	Standard Features, Configuration Information and Technical Specifications sections were updated.
05-Dec-2022	Version 23	Changed	Configuration Information section was updated, and new SKUs were added.
07-Nov-2022	Version 22	Changed	Standard Features, Configuration Information, Technical Specifications sections were updated.
10-Oct-2022	Version 21	Changed	Configuration Information section was updated.
03-Oct-2022	Version 20	Changed	Configuration Information and HPE Aruba Networking Data Center Networking Solution for HPE sections were updated.
15-Aug-2022	Version 19	Changed	Overview, Standard Features, Configuration Information, and Technical Specifications sections were updated.
06-Jun-2022	Version 18	Changed	Standard Features, Configuration Information, Technical Specifications sections were updated.
02-May-2022	Version 17	Changed	Overview, Standard Features, Configuration Information, Technical Specifications and HPE Aruba Networking Data Center Networking Solution for HPE sections were updated.
04-Apr-2022	Version 16	Changed	Configuration Information section was updated.
07-Feb-2022	Version 15	Changed	Configuration Information section was updated.
06-Dec-2021	Version 14	Changed	New HPE Aruba Networking Data Center Networking Solution for HPE section was added to QuickSpecs
02-Aug-2021	Version 13	Changed	Standard Features, Configuration Information, and Technical Specifications section were updated. SKUs added as well.
07-Jun-2021	Version 12	Changed	Overview, Standard Features, and Configuration Information sections were updated.
04-May-2021	Version 11	Changed	Standard Features, Configuration Information, and Technical Specifications sections were updated.
08-Mar-2021	Version 10	Changed	SKUs were added in Configuration Information section.
14-Dec-2020	Version 9	Changed	Overview Standard Features, Configuration information and Technical Specifications sections were updated.
07-Dec-2020	Version 8	Changed	Standard Features, Configuration information and Technical Specifications sections were updated.
08-Sep-2020	Version 7	Changed	Configuration Information section was updated.
10-Aug-2020	Version 6	Changed	Standard Features and Configuration information sections were updated.
06-Jul-2020	Version 5	Changed	Overview Standard Features, Configuration information and Technical Specifications sections were updated.
03-Feb-2020	Version 4	Changed	Configuration information section was updated.
20-Jan-2020	Version 3	Changed	Configuration information and Technical Specifications sections were updated.
06-Jan-2020	Version 2	Changed	Standard Features, Configuration information and Technical Specifications sections were updated.

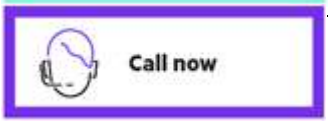
Summary of Changes

01-Nov-2019	Version 1	New	New QuickSpecs

Copyright

Make the right purchase decision. Contact our presales specialists.







© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Hewlett Packard Enterprise

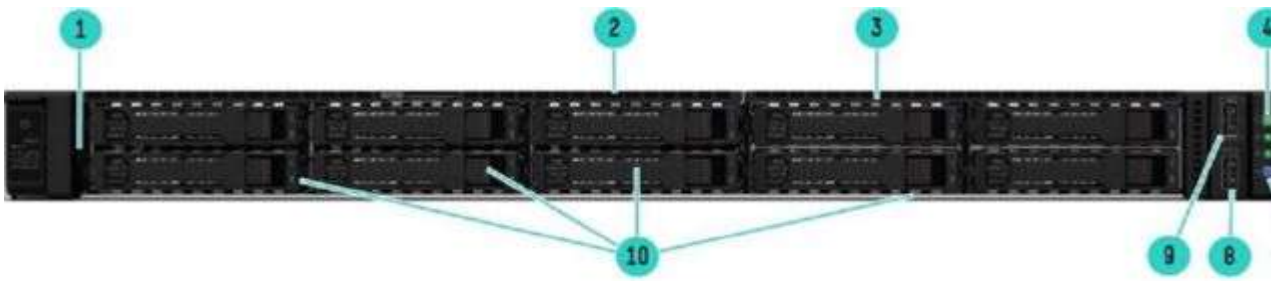
To learn more, visit: **http://www.hpe.com/networking**

a00073540enw - 16489 - Worldwide - V35 - 06-January-2025

Overview

HPE ProLiant DL325 Gen11

Are you looking for a scalable, low-cost performance server solution for your virtualized and software-defined co
The HPE ProLiant DL325 Gen11 server is a low-cost 1U 1P solution that delivers exceptional value balancing co
network bandwidth at 1P economics. Powered by 4th and 5th Generation AMD EPYC™ Processors with up to 16
memory bandwidth (up to 3 TB), high-speed PCIe Gen5 I/O and EDSFF storage, and supporting up to 2 double-
this server is a superb low-cost, 1U 1P, performance solution for your virtualized workloads. The silicon root of t
firmware, creating a fingerprint for the AMD Secure Processor that must be matched exactly before the server bo
DL325 Gen11 server is an excellent choice for virtualized workloads such as software-defined compute, CDN, V
that require balancing processor, memory, and network bandwidth.



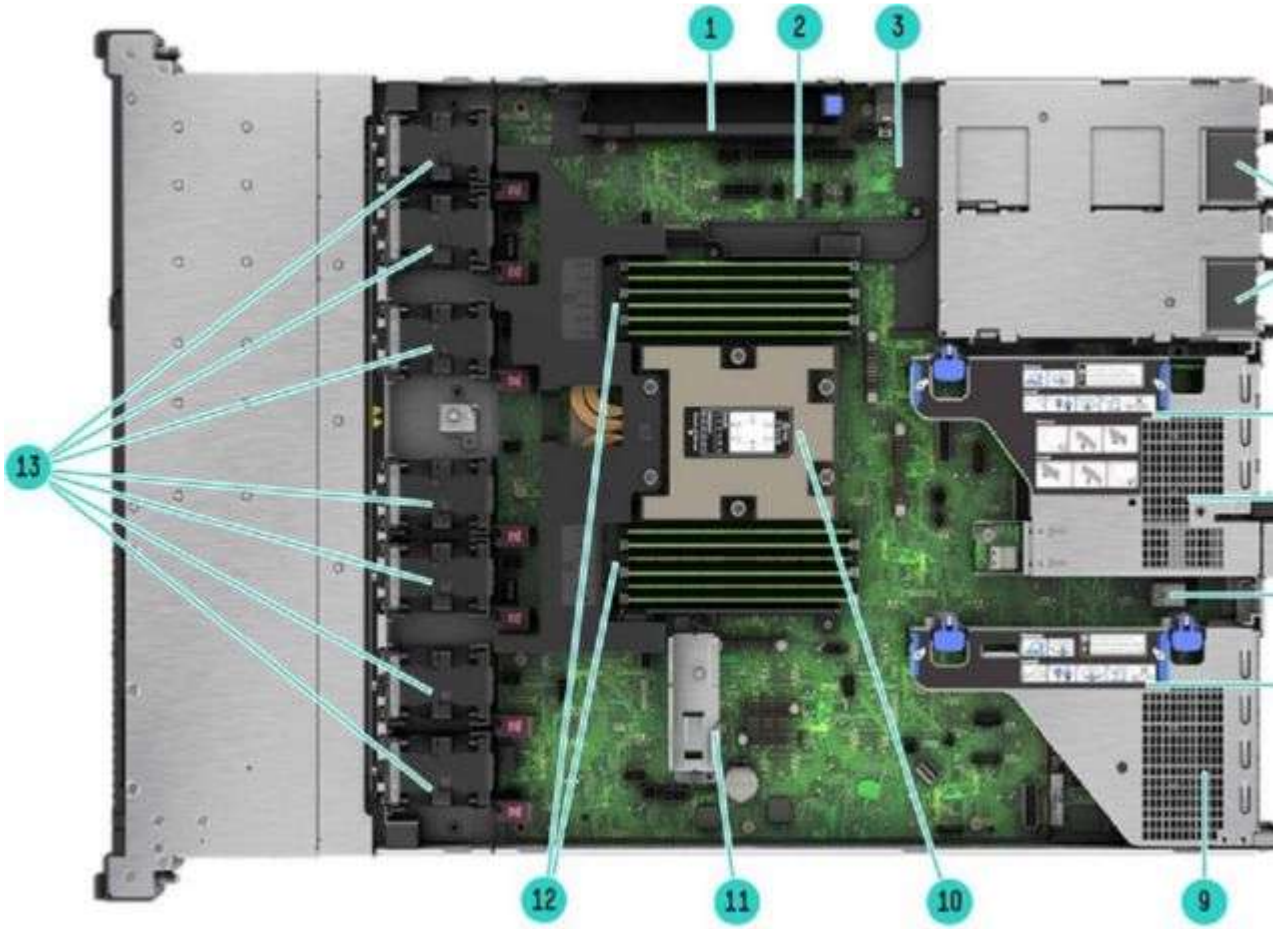
Front View - 8 SFF + optional 2 SFF Drive Bay shown

1. Serial number pull tab
2. Quick removal access panel
3. 2 SFF Cage Bay (Optional - shown) [1]
4. Power On/Standby button and system power LED
5. Health LED
6. NIC status LED [2]
7. Unit ID button/ LED
8. USB 3.2 Gen1 port
9. iLO Service Port
10. 8 SFF Cage Bay

Notes:

- [1] Optional: Optical Drives
- [2]Front NIC LED display doesn't support NIC LED ACT/LINK indication from PCIE NIC's

Overview



Internal View - Standard for all HPE ProLiant DL325 Gen 11

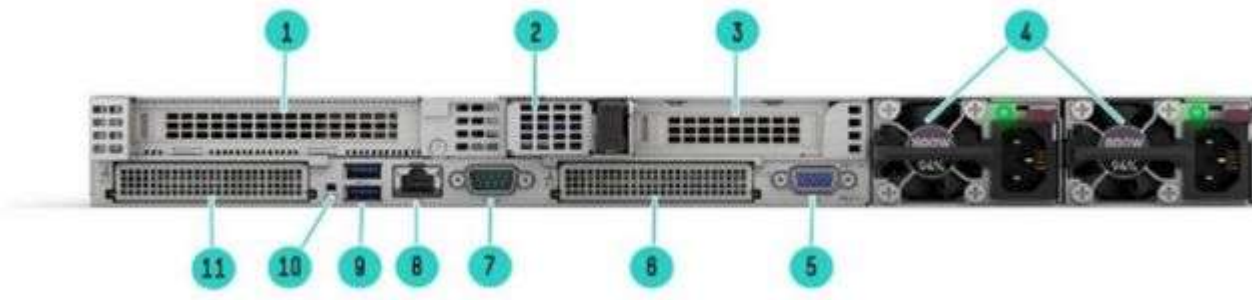
1. Megacell battery holder
2. Hard drive backplane power connectors
3. Chassis intrusion detection connector
4. Up to 2 Hot Plug redundant HPE Flexible Slot Power supplies
5. Secondary PCIe 5.0 riser
6. OCP 3.0 Slot 22 (Under)
7. Internal Dual USB 3.2 Gen1 port
8. Primary PCIe 5.0 riser
9. OCP 3.0 Slot 21 (Under)
10. Processor is shown with Perform 1)
11. FHFL PCIe card holder
12. DDR5 DIMM slots[2]
13. Hot-plug fans[3]

Notes:

– [1]Optional: Standard Heat Sink and Closed-Loop Liquid Cooling Heat Sink

– [2]Fully populated 12 DIMMs shown.

– [3]7 dual-rotor standard fans shown. Optional: Performance Fans and Liquid Cooling Fans



Rear View - Secondary Low Profile Riser Shown

1. Slot 1 Primary PCIe 5.0 Riser
7. Optional Serial port

Overview

2. Optional NS204i-u hot-plug NVMe boot device
3. Slot 2 Secondary PCIe 5.0 Riser[1]
4. Hot-plug Power Supply 1 and 2[2]
5. Video (VGA) port
6. OCP 3.0 Slot 22
8. Dedicated iLO management port
9. USB 3.1 Gen1 Ports (2)
10. Unit ID LED
11. OCP 3.0 Slot 21

Notes:

– [1]Low profile and full height options

– [2]Hot-plug Power Supply 2 is optional

What's New

- All new DL325 Gen11
- New 4th and 5th Generation AMD EPYC™ Processors, up to 160 cores, 400W, and 1150MB of L3 Cache.
- New DDR5 Smart Memory - up to 6000MT/s.
- New PCIe Gen5 support.
- New HPE Integrated Lights-Out 6 (iLO 6) server management software.
- New hot-pluggable NS204i-u Boot Device.
- New 20 EDSFF E3.S 1T Drive bays.
- New GPU support, up to four single-width or two double-width GPUs.
- OpenBMC Capable through iLO6 Transfer of Ownership Process

Platform Information

Form Factor

- 1U rack

Chassis Types

- 8 SFF with optional 2 SFF drive bay or optical drive.
- 4 LFF with an optional optical drive
- 20 EDSFF E3.S 1T drive bay.
- 2 Single-Width or 2 Double-Width GPUs with 8 EDSFF or 4 SFF drive bay.

System Fans

- Choice of Standard Fan Kit, Performance Fan, and Liquid cooling Fan Kit

Notes:

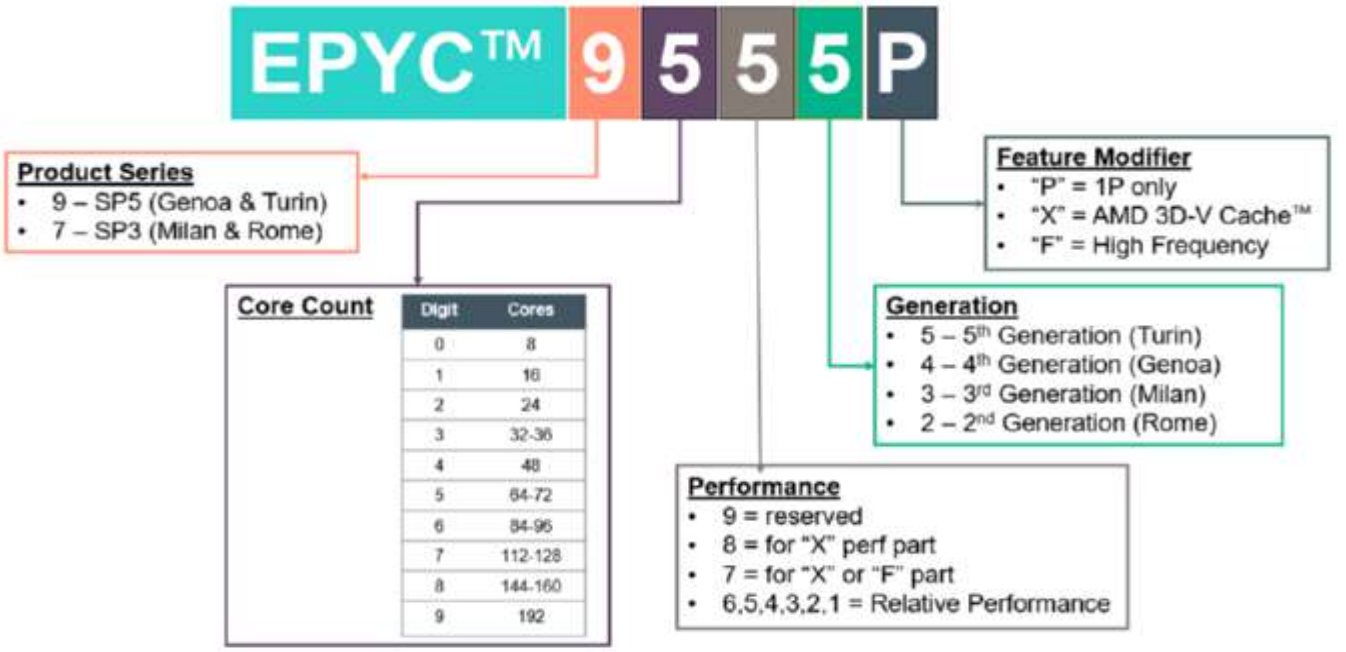
– The DL325 Gen11 supports up to 7 fans with fan redundancy built in. One fan rotor failure will place the server in de
fully functional. Two fan rotor failures could provide a warning and imminent server shutdown.

– Each Fan kit is designated to operate under different configurations. Please refer to the cooling option message in t
section for more information.

Standard Features

Processors - One of the following, depending on the model.

Notes: For more information regarding AMD EPYC processors, please see the following:
https://www.amd.com/en/processors/epyc-server-cpu-family



5th Gen AMD EPYC Processor	Cores	Base Frequency	Max Frequency	Max Memory	Wattage	L3 Cache (MB)	Memory
EPYC 9845	160	2.1 GHz	3.7 GHz	3TB	390	320	6000MT/s
EPYC 9825	144	2.2 GHz	3.7 GHz	3TB	390	384	6000MT/s
EPYC 9745	128	2.4 GHz	3.7 GHz	3TB	400	256	6000MT/s
EPYC 9645	96	2.3 GHz	3.7 GHz	3TB	320	256	6000MT/s
EPYC 9655P	96	2.6 GHz	4.5 GHz	3TB	400	384	6000MT/s
EPYC 9565	72	3.15 GHz	4.3 GHz	3TB	400	384	6000MT/s
EPYC 9535	64	2.4 GHz	4.3 GHz	3TB	300	256	6000MT/s
EPYC 9575F	64	3.3 GHz	5 GHz	3TB	400	256	6000MT/s
EPYC 9555P	64	3.2 GHz	4.4 GHz	3TB	360	256	6000MT/s
EPYC 9475F	48	3.65 GHz	4.8 GHz	3TB	400	256	6000MT/s
EPYC 9455P	48	3.15 GHz	4.4 GHz	3TB	300	256	6000MT/s
EPYC 9365	36	3.4 GHz	4.3 GHz	3TB	300	192	6000MT/s
EPYC 9335	32	3 GHz	4.4 GHz	3TB	210	128	6000MT/s
EPYC 9375F	32	3.8 GHz	4.8 GHz	3TB	320	256	6000MT/s
EPYC 9355P	32	3.55 GHz	4.4 GHz	3TB	280	256	6000MT/s
EPYC 9255	24	3.25 GHz	4.3 GHz	3TB	200	128	6000MT/s
EPYC 9275F	24	4.1 GHz	4.8 GHz	3TB	320	256	6000MT/s
EPYC 9135	16	3.65 GHz	4.3 GHz	3TB	200	64	6000MT/s
EPYC 9115	16	2.6 GHz	4.1 GHz	3TB	125	64	6000MT/s
EPYC 9175F	16	4.2 GHz	5 GHz	3TB	320	512	6000MT/s
EPYC 9015	8	3.6 GHz	4.1 GHz	3TB	125	64	6000MT/s

Standard Features

4th Gen AMD EPYC Processor	Cores	Base Frequency	Max Frequency	Max Memory	Wattage	L3 Cache (MB)	Memory
EPYC 9754	128	2.25 GHz	3.1 GHz	3TB	360	256	4800MT/s
EPYC 9734	112	2.2 GHz	3.0 GHz	3TB	340	256	4800MT/s
EPYC 9654P	96	2.4 GHz	3.7 GHz	3TB	360	384	4800MT/s
EPYC 9684X	96	2.55 GHz	3.7 GHz	3TB	400	1150	4800MT/s
EPYC 9634	84	2.25 GHz	3.7 GHz	3TB	290	384	4800MT/s
EPYC 9554P	64	3.1 GHz	3.75 GHz	3TB	360	256	4800MT/s
EPYC 9534	64	2.45 GHz	3.7 GHz	3TB	280	256	4800MT/s
EPYC 9454P	48	2.75 GHz	3.8 GHz	3TB	290	256	4800MT/s
EPYC 9474F	48	3.6 GHz	4.1 GHz	3TB	360	256	4800MT/s
EPYC 9354P	32	3.25 GHz	3.8 GHz	3TB	280	256	4800MT/s
EPYC 9334	32	2.7 GHz	3.9 GHz	3TB	210	128	4800MT/s
EPYC 9374F	32	3.85 GHz	4.3 GHz	3TB	320	256	4800MT/s
EPYC 9384X	32	3.1 GHz	3.9 GHz	3TB	320	768	4800MT/s
EPYC 9254	24	2.9 GHz	4.15 GHz	3TB	200	128	4800MT/s
EPYC 9224	24	2.5 GHz	3.7 GHz	3TB	200	64	4800MT/s
EPYC 9274F	24	4.05 GHz	4.3 GHz	3TB	320	256	4800MT/s
EPYC 9124	16	3 GHz	3.7 GHz	3TB	200	64	4800MT/s
EPYC 9174F	16	4.1 GHz	4.4 GHz	3TB	320	256	4800MT/s
EPYC 9184X	16	3.55 GHz	4.2 GHz	3TB	320	768	4800MT/s

Notes:
- 6096pin LGA SP5 socket type, 128 PCIe 5.0 Lanes per processor.
- All 4th and 5th generation AMD EPYC processors can support up to 3TB of memory each under 1DPC, 12 channels per processor.
- The wattage information indicates the processor's default cTDP (Configurable TDP).

Chipset

No chipset - System on Chip (SoC) design.

On System Management Chipset

HPE iLO 6 ASIC

Notes: Read and learn more in the **iLO QuickSpecs**.

Memory

Type	HPE DDR5 Smart Memory Registered (RDIMM)
DIMM Slots Available	12 12 DIMM slots per processor, 12 channels per processor, 1 DIMM per channel
Maximum capacity (RDIMM)	3.0 TB 12 x 256 GB RDIMM @ 4800 MT/s at 1DPC for 4th Gen EPYC Processors 12 x 256 GB RDIMM @ 6000 MT/s at 1DPC for 5th Gen EPYC Processors

Notes:
- All processors support up to 3TB of memory per server.
- LRDIMM and Persistent Memory are not supported.
- For additional information, please see the **HPE DDR5 Smart Memory QuickSpecs**
- For the Memory Population Rules and Guidelines with AMD EPYC 9004 and 9005 series processors, see

Standard Features

details here: https://www.hpe.com/psnow/doc/a50007481enw

Memory Protection

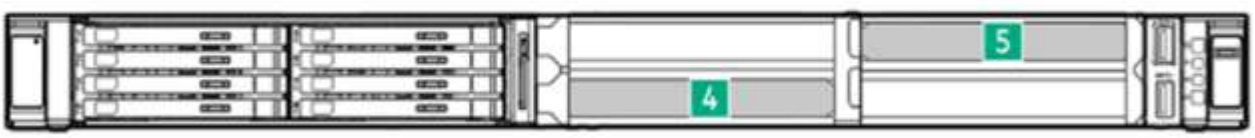
Advanced ECC
Advanced ECC uses single-device data correction to detect and correct single and all multibit error that occurs within a single DRAM chip.
Online Spare
Memory online spare mode detects a rank that is degrading and switches operation to the spare rank.
Notes: For more information see our **Memory RAS feature technical whitepaper.**

Expansion Slots

Slots #	Technology	Bus Width	Connector Width	Slot Form Factor
1 (Default Primary Riser)	PCIe 5.0	X16	X16	Full-height, Full-length slot
2 (Secondary Riser)	PCIe 5.0	X16	X16	Low Profile or Full-height, Half-length slot
21	PCIe 5.0	X8	X16	OCP 3.0
22	PCIe 5.0	X8	X16	OCP 3.0

Notes:

- Both OCP slots (slot 21 and 22) support shared NIC and WOL (wake on LAN) functions.
- If NS204i-u Boot Device is selected then low profile secondary riser (P55029-B21) must be selected.
- Requires a FHFL card holder to support the full-length cards at primary riser.



Front risers of GPU CTO server

Front Riser				
Slots #	Technology	Bus Width	Connector Width	Slot Form Factor
4	PCIe 5.0	X16	X16	Full-height, Full-length slot
5	PCIe 5.0	X16	X16	Full-height, Full-length slot

Notes:

- When supporting Slot4 & Slot21 scenario, Slot4 & OCP 21 slot combined can support up to 112GB/s bandwidth due to AMD CPU limitation.
- When supporting Slot5 & Slot1 scenario, Slot 5 & Slot1 combined can support up to 112GB/s bandwidth due to AMD CPU limitation.
- The extension slots at the front of the GPU CTO server do not support external cabling.

Storage Controllers

Boot Device

- HPE NS204i-u Gen11 NVMe Hot Plug Boot Optimized Storage Device

Notes:

– Can only be selected without M.2 enablement kit.
– Includes Hot Plug capable dual 480GB NVMe M.2 automatically configured into a RAID 1 Mirror
– Externally accessible but does not occupy a PCIe slot
– Requires specific cable kit and secondary low-profile riser along with specific cooling selections based on configuration

Essential RAID Controller

- HPE Smart Array E208e-p SR Gen10 Controller

MR Gen11 Storage Controller

- HPE MR216i-p Gen11 x16 Lanes without Cache PCI SPDM Plug-in Storage Controller
- HPE MR216i-o Gen11 x16 Lanes without Cache OCP SPDM Storage Controller
- HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller
- HPE MR416i-o Gen11 x16 Lanes 8GB Cache OCP SPDM Storage Controller
- HPE MR416i-p Gen11 x16 Lanes 8GB Cache PCI SPDM Plug-in Storage Controller

SR Gen11 Storage Controller

- HPE SR932i-p Gen11 x32 Lanes 8GB Wide Cache PCI SPDM Plug-in Storage

Notes: For additional details, please visit:

HPE Compute MR Gen11 Controllers QuickSpecs

HPE Compute SR Gen11 Controllers QuickSpecs

Internal Storage Devices

Optical Drive

- Available on 8SFF and 4LFF CTO Servers as an option (DVD-ROM or DVD-RW)

Drives

- None ship standard

Maximum Storage		
	Capacity	**Configuration**
Hot Plug LFF SAS HDD	80 TB	4 x 20 TB
Hot Plug LFF SATA HDD	80 TB	4 x 20 TB
Hot Plug SFF SAS SSD	76.8 TB	10 x 7.68 TB
Hot Plug SFF SATA SSD	76.8 TB	10 x 7.68 TB
Hot Plug SFF NVMe PCIe U.3 SSD	153.6 TB	10 x 15.36 TB
Hot Plug EDSFF E3.S 1T NVMe SSD	307.2 TB	20 x 15.36 TB
M.2 22110 NVMe SSD	3.84 TB	2 x 1.92 TB (via M.2 enablement Kit)
M.2 2280 SATA SSD	960 GB	2 x 480 GB (via M.2 enablement Kit)

Standard Features

Interfaces

Serial	1 optional port - rear
Video Port	1 standard VGA Port - rear
Network Ports	None. Choice of OCP or stand-up card, supporting a wide arrange of NIC adapters BTO models will come pre-selected with a primary networking card.
HPE iLO Remote Mgmt Port	1 1Gb Dedicated - rear
Front iLO Service Port	1 standard
USB 3.2 Gen1	5 standard on all models: 1 front, 2 rear, 2 internal

Graphics

Integrated Video Standard

- Video modes up to 1920 x 1200@60Hz (32 bpp)
- 16MB Video Memory

HPE iLO 6 on system management memory

- 64 MB Flash
- 8 Gbit DDR 4 with ECC protection

Power Supply

- HPE 500W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
  Notes: Available in 94% Power Efficiency.
- HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
  Notes: Available in 94% Power Efficiency.
- HPE 1000W Flex Slot Titanium Hot Plug Low Halogen Power Supply Kit
  Notes: Available in 96% Power Efficiency.
- HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
  Notes: Available in 94% Power Efficiency.
- HPE 1600W Flex Slot -48VDC Hot Plug Power Supply Kit
  Notes: Available in 94% Power Efficiency. 200-240VAC power input only.
- HPE 1800W-2200W Flex Slot Titanium Hot Plug Power Supply Kit

HPE Flexible Slot (Flex Slot) Power Supplies share a common electrical and physical design that allows for hot plug, and tool-less installation into HPE ProLiant Servers. Flex Slot power supplies are certified for high-efficiency operation and offer multiple power output options, allowing users to "right-size" a power supply for specific server configurations. This flexibility helps to reduce power waste, lower overall energy costs, and avoid "trapped" power capacity in the data center.

All pre-configured servers ship with a standard 6-foot IEC C-13/C-14 jumper cord (A0K02A). This jumper cord is also included with each standard AC power supply option kit. If a different power cord is required, please check the **ProLiant Power Cables** web page.

To review the power requirements for your selected system, please use the **HPE Power Advisor Tool**.

Standard Features

For information on power specifications and technical content visit **HPE Server power supplies**.

Operating Systems and Virtualization Software Support for HPE Servers

HPE servers are designed for seamless integration with partner Operating Systems and Virtualization Software. By collaborating closely with our partners, we ensure that their products are optimized, certified, and fully supported within your HPE server environment.

Access the certified and supported servers for each of the OS and Virtualization software: **HPE Servers Support & Certification Matrices**

Notes: Minimum required version includes all future updates of the indicated release unless a maximum is listed in the Notes

---

Industry Standard Compliance

- ACPI 6.1 Compliant
- PCIe 5.0 Compliant
- WOL Support
- Microsoft® Logo certifications
- PXE Support
- VGA/Display Port
- USB 3.2 Gen1 Compliant
- USB 2.0 Compliant
- Energy Star 4.0
- SMBIOS 3.1
- UEFI 2.7
- UEFI Class 3
- Redfish API
- IPMI 2.0
- Secure Digital 2.0
- Advanced Encryption Standard (AES)
- Triple Data Encryption Standard (3DES)
- SNMP v3
- TLS 1.2
- DMTF Systems Management Architecture for Server Hardware Command Line Protocol (SMASH CLP)
- Active Directory v1.0
- ASHRAE A3/A4

  Notes: For additional technical thermal details regarding ambient temperatures, humidity, and features support please visit: **Extended Ambient Temperature Guidelines for HPE Gen11 servers**

- UEFI (Unified Extensible Firmware Interface Forum)
- APML 1.0

Embedded Management

HPE Integrated Lights-Out (HPE iLO)

Monitor your servers for ongoing management, service alerting, reporting, and remote management with HPE iLO.
Learn more at **http://www.hpe.com/info/ilo**.

Standard Features

UEFI
Configure and boot your servers securely with industry-standard Unified Extensible Firmware Interface (UEFI).

Intelligent Provisioning
Hassle-free server and OS provisioning for 1 or more servers with Intelligent Provisioning.
Learn more at https://support.hpe.com/hpesc/public/docDisplay?docId=c04465280&docLocale=en_US

iLO RESTful API
iLO RESTful API is DMTF Redfish API implementation and offers simplified server management automation such as configuration and maintenance tasks based on modern industry standards. Learn more at **http://www.hpe.com/info/restfulapi**.

OpenBMC Support
OpenBMC Capable through iLO6 Transfer of Ownership Process.
Learn more at **OpenBMC Support**

HPE Server UEFI

Unified Extensible Firmware Interface (UEFI) is an industry standard that provides better manageability and more secure configuration than the legacy ROM while interacting with your server at boot time. HPE ProLiant Gen11 servers have a UEFI Class 3 implementation.
UEFI enables numerous new capabilities specific to HPE ProLiant servers such as

- Secure Boot and Secure Start enable enhanced security
- Operating system specific functionality
- Support for > 2.2 TB (using GPT) boot drives
- USB 3.1 Gen1 Stack
- Embedded UEFI Shell
- Mass Configuration Deployment Tool using iLO RESTful API that is Redfish API Conformant
- PXE boot support for IPv6 networks
- Workload Profiles for simple performance optimization
- Embedded TPM Support

UEFI Boot Mode only

- NVMe Boot Support
- iSCSI Software Initiator Support.
- HTTP/HTTPs Boot support as a PXE alternative.
- Boot support for option cards that only support a UEFI option ROM

Notes:
– For UEFI Boot Mode, boot environment and OS image installation should be configured properly to support UEFI
– TPM is embedded on the DL325 Gen11 mainboard and does not require additional option kit selection to enable this function.

Server Utilities

Active Health System
The HPE Active Health System (AHS) is an essential component of the iLO management portfolio that provides continuous, proactive health monitoring of HPE servers. Learn more at **http://www.hpe.com/servers/ahs**.

Standard Features

Smart Update

Keep your servers up to date with the HPE Smart Update solution by using Smart Update Manager (SUM) to optimize the firmware and driver updates of the Service Pack for ProLiant (SPP). Learn more at **https://www.hpe.com/us/en/servers/smart-update.html**

iLO Amplifier Pack

Designed for large enterprise and service provider environments with hundreds of HPE servers, the iLO Amplifier Pack is a free, downloadable open virtual application (OVA) that delivers the power to discover, inventory, and update Gen8, Gen9, Gen10, and Gen10 Plus HPE servers. Use an iLO Advanced License to unlock full capabilities.

Learn more at **http://www.hpe.com/servers/iLOamplifierpack**.

RESTful Interface Tool
RESTful Interface tool (iLOREST) is a single scripting tool to provision using iLO RESTful API to discover and deploy servers at scale. Learn more at **http://www.hpe.com/info/resttool**.

Scripting Tools
Provision one to many servers using your own scripts to discover and deploy with Scripting Tool (STK) for Windows and Linux or Scripting Tools for Windows PowerShell. Learn more at **http://www.hpe.com/servers/powershell**.

HPE OneView Standard
HPE OneView is an on-premises, multi-generational server monitoring, and management solution. HPE OneView Standard can be used for inventory, health monitoring, alerting, and reporting without additional fees. Customers can upgrade their management experience with an HPE OneView Advanced license, all provided by the same tool. Learn more at **http://www.hpe.com/info/oneview**.

HPE GreenLake for Compute Ops Management

HPE is intelligently transforming compute management with an intuitive cloud operating experience through HPE GreenLake cloud platform to streamline and secure operations from edge-to-cloud. Automated key lifecycle tasks, for onboarding, updating, managing, and monitoring HPE servers, brings agility and greater efficiencies to wherever compute devices reside via a unified single browser-based interface. Manage single locations or multiple, distributed sites. Keep tens to thousands of servers secure with batch policy controls and automated updates.
Compute Ops Management is cloud-native software that is continually updated with new services, features, patches, and fixes. The management application resides in the HPE GreenLake cloud platform (access via **https://console.greenlake.hpe.com**) and leverages the HPE GreenLake architecture, security, and unified operations.
A 3-year subscription to HPE GreenLake for Compute Ops Management is added by default when ordering an HPE ProLiant Gen11 rack, tower, or micro server.
For more information, visit the HPE GreenLake for Compute Ops Management QuickSpecs: **https://www.hpe.com/psnow/doc/a50004263enw**

Security

- UEFI Secure Boot and Secure Start support
- Immutable Silicon Root of Trust
- FIPS 140-3 validation (iLO 6 certification in progress)
- Common Criteria certification (iLO 6 certification in progress)
- Configurable for PCI DSS compliance
- Advanced Encryption Standard (AES) and Triple Data Encryption Standard (3DES) on browser
- Support for Commercial National Security Algorithms (CNSA)

Standard Features

- Tamper-free updates - components digitally signed and verified
- Secure Recovery - recover critical firmware to a known good state on detection of compromised firmware
- Ability to rollback firmware
- Secure erase of NAND/User data
- TPM (Trusted Platform Module) 2.0 option

  Notes: TPM is embedded on the DL325 Gen11 mainboard and does not require additional option kit selection to enable this function.

- Bezel Locking Kit option
- Chassis Intrusion detection option

Server Management

HPE iLO Advanced

HPE iLO Advanced licenses offer smart remote functionality without compromise, for all HPE ProLiant servers. The license includes the fully integrated remote console, virtual keyboard, video, and mouse (KVM), multi-user collaboration, console record and replay, and GUI-based and scripted virtual media and virtual folders. You can also activate the enhanced security and power management functionality.

HPE OneView Advanced

HPE OneView Advanced offers a sophisticated level of automation to infrastructure management by taking a template driven approach to provisioning, updating, and integrating compute, storage, and networking infrastructure. It builds upon the base features of HPE OneView Standard and provides full-featured licenses which can be purchased for managing multiple HPE server generations. To learn more visit **http://www.hpe.com/info/oneview**.

Warranty

This product is covered by a global limited warranty and supported by HPE Services and a worldwide network of HPE Authorized Channel Partners resellers. Hardware diagnostic support and repair are available for three years from the date of purchase. Support for software and initial setup is available for 90 days from the date of purchase. Enhancements to warranty services are available through HPE Services operational services or customized service agreements. Hard drives have either a one year or three year warranty; refer to the specific hard drive QuickSpecs for details.

Notes: Server Warranty includes 3-Year Parts, 3-Year Labor, and 3-Year Onsite support with next business day response. Warranty repairs may be accomplished through the use of Customer Self Repair (CSR) parts. These parts fall into two categories: 1) Mandatory CSR parts are designed for easy replacement. A travel and labor charge will result when customers decline to replace a Mandatory CSR part; 2) Optional CSR parts are also designed for easy replacement but may involve added complexity. Customers may choose to have Hewlett Packard Enterprise replace Optional CSR parts at no charge. Additional information regarding worldwide limited warranty and technical support is available at:

https://www.hpe.com/support/ProLiantServers-Warranties

Optional Features

Rack and Power Infrastructure

The story may end with servers, but it starts with the foundation that makes compute go - and business grow. We've reinvented our entire portfolio of rack and power products to make IT infrastructure more secure, more practical, and more efficient. In other words, we've created a stronger, smarter, and simpler infrastructure to help you get the most out of your IT equipment. As an industry leader, Hewlett Packard Enterprise is uniquely positioned to address the key concerns of power, cooling, cable management, and system access.

HPE G2 Advanced and Enterprise Racks are perfect for the server room or today's modern data center with the enhanced airflow and thermal management, flexible cable management, and a 10-year Warranty to support higher-density computing.

HPE G2 PDUs offer reliable power in flexible form factors that operate at temperatures up to 60°C, include color-coded outlets and load segments, and a low-profile design for optimal access to the rack and support for dense rack environments.

HPE Uninterruptible Power Systems are cost-effective power protection for any type of workload. Some UPSs include options for remote management and extended runtime modules so you're critical dense data center is covered in power outages.

HPE KVM Solutions include a console and switches designed to work with your server and IT equipment reliably. We've got a cost-effective KVM switch for your first rack and multiple connection IP switches with remote management and security capabilities to keep your data center rack up and running.

Learn more about HPE Racks, KVM, PDUs, and UPSs at **HPE Rack and Power Infrastructure**.

One Config Simple (OCS/SCE)

OCS/SCE is a guided self-service tool to help sales and non-technical people provide customers with initial configurations in 3 to 5 minutes. You may then send the configuration on for configuration help, or use it in your existing ordering processes. If you require "custom" rack configuration or configuration for products not available in SCE, please contact Hewlett Packard Enterprise Customer Business Center or an Authorized Partner for assistance. **https://h22174.www2.hpe.com/SimplifiedConfig/Welcome**

Service and Support

HPE Services

No matter where you are in your digital transformation journey, you can count on HPE Services to deliver the expertise you need when, where and how you need it. From planning to deployment, ongoing operations and beyond, our experts can help you realize your digital ambitions.

https://www.hpe.com/services

Consulting Services

No matter where you are in your journey to hybrid cloud, experts can help you map out your next steps. From determining what workloads should live where, to handling governance and compliance, to managing costs, our experts can help you optimize your operations.

https://www.hpe.com/services/consulting

HPE Managed Services

HPE runs your IT operations, providing services that monitor, operate, and optimize your infrastructure and applications, delivered consistently and globally to give you unified control and let you focus on innovation.

HPE Managed Services | HPE

Operational services

Optimize your entire IT environment and drive innovation. Manage day-to-day IT operational tasks while freeing up valuable time and resources. Meet service-level targets and business objectives with features designed to drive better business outcomes.

https://www.hpe.com/services/operational

HPE Complete Care Service

HPE Complete Care Service is a modular, edge-to-cloud IT environment service designed to help optimize your entire IT environment and achieve agreed upon IT outcomes and business goals through a personalized experience. All delivered by an assigned team of HPE Services experts. HPE Complete Care Service provides:

- A complete coverage approach -- edge to cloud
- An assigned HPE team
- Modular and fully personalized engagement
- Enhanced Incident Management experience with priority access
- Digitally enabled and AI driven customer experience

https://www.hpe.com/services/completecare

HPE Tech Care Service

HPE Tech Care Service is the operational support service experience for HPE products. The service goes beyond traditional support by providing access to product specific experts, an AI driven digital experience, and general technical guidance to not only reduce risk but constantly search for ways to do things better. HPE Tech Care Service delivers a customer-centric, AI driven, and digitally enabled customer experience to move your business forward. HPE Tech Care Service is available in three response levels. Basic, which provides 9x5 business hour availability and a 2-hour response time. Essential which provides a 15-minute response time 24x7 for most enterprise level customers, and Critical which includes a 6-hour repair commitment where available and outage management response for severity 1 incidents.

https://www.hpe.com/services/techcare

Service and Support

HPE Lifecycle Services

HPE Lifecycle Services provide a variety of options to help maintain your HPE systems and solutions at all stages of the product lifecycle. A few popular examples include:

- Lifecycle Install and Startup Services: Various levels for physical installation and power on, remote access setup, installation and startup, and enhanced installation services with the operating system.
- HPE Firmware Update Analysis Service: Recommendations for firmware revision levels for selected HPE products, taking into account the relevant revision dependencies within your IT environment.
- HPE Firmware Update Implementation Service: Implementation of firmware updates for selected HPE server, storage, and solution products, taking into account the relevant revision dependencies within your IT environment.
- Implementation assistance services: Highly trained technical service specialists to assist you with a variety of activities, ranging from design, implementation, and platform deployment to consolidation, migration, project management, and onsite technical forums.
- HPE Service Credits: Access to prepaid services for flexibility to choose from a variety of specialized service activities, including assessments, performance maintenance reviews, firmware management, professional services, and operational best practices.

Notes: To review the list of Lifecycle Services available for your product go to:
https://www.hpe.com/services/lifecycle
For a list of the most frequently purchased services using service credits, see the **HPE Service Credits Menu**

Other Related Services from HPE Services:

HPE Education Services

Training and certification designed for IT and business professionals across all industries. Broad catalogue of course offerings to expand skills and proficiencies in topics ranging from cloud and cybersecurity to AI and DevOps. Create learning paths to expand proficiency in a specific subject. Schedule training in a way that works best for your business with flexible continuous learning options.
https://www.hpe.com/services/training

Defective Media Retention

An option available with HPE Complete Care Service and HPE Tech Care Service and applies only to Disk or eligible SSD/Flash Drives replaced by HPE due to malfunction.
Consult your HPE Sales Representative or Authorized Channel Partner of choice for any additional questions and services options.

Parts and Materials

HPE will provide HPE-supported replacement parts and materials necessary to maintain the covered hardware product in operating condition, including parts and materials for available and recommended engineering improvements.
Parts and components that have reached their maximum supported lifetime and/or the maximum usage limitations as set forth in the manufacturer's operating manual, product quick-specs, or the technical product data sheet will not be provided, repaired, or replaced as part of these services.

How to Purchase Services

Services are sold by Hewlett Packard Enterprise and Hewlett Packard Enterprise Authorized Service Partners:

- Services for customers purchasing from HPE or an enterprise reseller are quoted using HPE order configuration tools.
- Customers purchasing from a commercial reseller can find services at **https://ssc.hpe.com/portal/site/ssc/**

Service and Support

AI Powered and Digitally Enabled Support Experience

Achieve faster time to resolution with access to product-specific resources and expertise through a digital and data driven customer experience

Sign into the HPE Support Center experience, featuring streamlined self-serve case creation and management capabilities with inline knowledge recommendations. You will also find personalized task alerts and powerful troubleshooting support through an intelligent virtual agent with seamless transition when needed to a live support agent.

https://support.hpe.com/hpesc/public/home/signin

Consume IT On Your Terms

HPE GreenLake edge-to-cloud platform brings the cloud experience directly to your apps and data wherever they are-the edge, colocations, or your data center. It delivers cloud services for on-premises IT infrastructure specifically tailored to your most demanding workloads. With a pay-per-use, scalable, point-and-click self-service experience that is managed for you, HPE GreenLake edge-to-cloud platform accelerates digital transformation in a distributed, edge-to-cloud world.

- Get faster time to market
- Save on TCO, align costs to business
- Scale quickly, meet unpredictable demand
- Simplify IT operations across your data centers and clouds

To learn more about HPE Services, please contact your Hewlett Packard Enterprise sales representative or Hewlett Packard Enterprise Authorized Channel Partner. Contact information for a representative in your area can be found at "Contact HPE" **https://www.hpe.com/us/en/contact-hpe.html**

For more information

http://www.hpe.com/services

Pre-Configured Models

Pre-Configured models ship with the configurations below.

- Pre-Configured models ship with the configurations below. Options can be selected from the Core or Additional options section of this QuickSpecs.
- Hewlett Packard Enterprise does not allow factory integration of options into pre-configured models. Any additional options purchased will not be shipped inside the server.
- Network Choice models do not include embedded LOM.

Base Models		
SKU Number	P66775-B21 P66775-291	P58690-421
Model Name	HPE ProLiant DL325 Gen11 9124 3.0GHz 16-core 1P 32GB-DR MR408i-o 8SFF 800W PS Server	HPE ProLiant DL325 Gen11 9124 3.0GHz 16-core 1P 32GB-R MR408i-o 8SFF 1000W PS EU Server
Chassis	HPE ProLiant DL325 Gen11 8SFF Configure-to-order Server	
Processor	9124 (16 core, 3.0 GHz, 200W)	
Number of Processors	One with standard heatsink	
Memory	32 GB (1x32 GB, 4800 MT/s)	
Network Controller	Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE	
Storage Controller	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller	
Included Hard Drives	None ship standard, 8 SFF supported	
Internal Storage	8 SFF Chassis (upgradeable to 10 SFF front)	
Optical Drive	Optional, None ship standard	
Expansion Slots	1 PCIe x16 Primary Riser	
Power Supply	1x HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit
Fans	7x Standard Fans	
Management	Default: HPE iLO Standard with Intelligent Provisioning, HPE OneView Standard (requires download)	
Rail Kit	HPE ProLiant DL3XX Gen11 Easy Install Rail 2 Kit	
Security	TPM (Trusted Platform Module)	
Energy Star	4.0 certified	
Form Factor	1U Rack	
Warranty	Server warranty includes 3-year parts, 3-year labor, 3-year onsite support with next business day response.	

Pre-Configured Models

Performance Models		
SKU Number	P66776-B21 P66776-291	P58691-421
Model Name	HPE ProLiant DL325 Gen11 9354P 3.25GHz 32-core 1P 32GB-DR MR408i-o 8SFF 800W PS Server	HPE ProLiant DL325 Gen11 9354P 3.25GHz 32-core 1P 32GB-R MR408i-o 8SFF 1000W PS EU Server
Chassis	HPE ProLiant DL325 Gen11 8SFF Configure-to-order Server	
Processor	9354P (32 core, 3.25 GHz, 280W)	
Number of Processors	One with performance heatsink	
Memory	32 GB (1x32 GB, 4800 MT/s)	
Network Controller	Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T OCP3 Adapter for HPE	
Storage Controller	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller	
Included Hard Drives	None ship standard, 8 SFF supported	
Internal Storage	8 SFF Chassis (upgradeable to 10 SFF front)	
Optical Drive	Optional, None ship standard	
Expansion Slots	1 PCIe x16 Primary Riser	
Power Supply	1x HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit
Fans	7x Performance Fans	
Management	Default: HPE iLO Standard with Intelligent Provisioning, HPE OneView Standard (requires download)	
Rail Kit	HPE ProLiant DL3XX Gen11 Easy Install Rail 2 Kit	
Security	TPM (Trusted Platform Module)	
Energy Star	4.0 certified	
Form Factor	1U Rack	
Warranty	Server warranty includes 3-year parts, 3-year labor, 3-year onsite support with next business day response.	

Pre-Configured Models

Country Code Key

- -B21 = Worldwide
- -291 = Japan
- -421 = Europe, the Middle East and Africa

HPE Smart Choice purchase program

The HPE Smart Choice purchase program features popular fully configured products that can be quoted in minutes and shipped quickly through HPE Authorized Partners. Products are configured and tested in an HPE factory and stocked at HPE Authorized Distributors and Partners. The products arrive in a single box, making onsite integration easier and more efficient for partners and customers. Additionally, there are aggressively priced HPE Tech Care Services available only through the HPE Smart Choice program when you purchase an HPE Smart Choice product. For additional information on the HPE Smart Choice purchase program, please visit: **https://www.hpe.com/psnow/doc/a50009219enw**

Configuration Information

Smart Templates from HPE

HPE is releasing new Smart Template technology in the One Config Advanced (OCA) configurator. These Templates represent the CTO equivalents of the top-selling BTO configurations. They are intended to provide simple starting points to assist you in easily creating and customizing your desired Server solutions. HPE Servers that have Platform Templates, developed by HPE Product Managers, will have a separate tab in the HPE OCA configurator.

Workload Solutions Templates from HPE

The Workload Solutions Templates build on the Smart Templates technology to easily develop working configurations of the most compelling Workload Solutions. The templates complement the Reference Builds developed by HPE. Workload Solutions templates preconfigure some of the key architecture decisions and make it easier for Sellers to get started and complete a differentiated server solution for your customer's specific workload.

Mainstream SKUs

HPE launched the Mainstream SKU initiative as a market-driven approach to Demand Steering. It is a simplified portfolio of our top selling options that meet the current and future market trends. HPE has committed to providing a more predictable and faster experience for these options. Mainstream SKUs enjoy higher safety stock levels and have high fulfillment service levels than non-Mainstream SKUs. Mainstream orders are fulfilled +30% faster than non-Mainstream orders, have fewer shortages, and better recovery dates. This platform has Mainstream SKUs in the options portfolio and is eligible for an improved Mainstream experience. Mainstream SKUs are designated with a Mainstream symbol in our configurators.

Mainstream Configurations

HPE is using the new Smart Templates technology to present Mainstream configurations. All the options in a Mainstream configuration are pre-selected Mainstream SKUs to optimize the performance, predictability, and fulfillment experience. Check the Template section in our configurators for eligible Mainstream configurations.

European Union Erp Lot 9 Regulation

Beginning on January 1st, 2024, units sold into the European Union (EU), European Economic Area (EEA), the United Kingdom, or Switzerland must include more efficient AC power supplies: 94% for multi-output and 96% for single-output. HPE Flexible Slot power supplies are single-output, and part numbers 865438-B21, P03178-B21, and P44712-B21 are 96% efficient, thus meeting requirements.
HPE is on target to fulfil compliant systems ahead of time and will begin enforcing these requirements in advance to satisfy requests with the current power supplies by the set deadline.

This section lists some of the steps required to configure a Factory Integrated Model. To ensure only valid configurations are ordered, Hewlett Packard Enterprise recommends the use of an HPE approved configurator. Contact your local sales representative for information on configurable product offerings and requirements.

- Factory Integrated Models must start with a CTO Server.
- FIO indicates that this option is only available as a factory installable option.
- All Factory Integrated Models will be populated with sufficient hard drive blanks based on the number of initial hard drives ordered with the server.
- Some options may not be integrated at the factory. Contact your local sales representative for additional information.

Configuration Information

Step 1: Base Configuration (choose one (1) of the following configurable server models from the tables below)

CTO Server	HPE ProLiant DL325 Gen11 8SFF Configure-to-order Server	HPE ProLiant DL325 Gen11 4LFF Configure-to-order Server	HPE ProLiant DL325 Gen11 EDSFF Configure-to-order Server	HPE ProLiant DL325 Gen11 GPU Configure-to-order Server
SKU Number	P54199-B21	P54200-B21	P54201-B21	P54202-B21
TAA SKU	P54199-B21#GTA	P54200-B21#GTA	P54201-B21#GTA	P54202-B21#GTA
HPE Trusted Supply Chain	P36394-B21 - Optional			
Processor	Not included as standard			
DIMM Slots	12-DIMM slots			
Storage Controller	Choice of HPE storage controllers			
PCIe	1 PCIe 5.0 x16 Primary Riser			3 PCIe 5.0 x16 Risers (Slot 1,4,5)
OCP3.0 Slot	2 PCIe 5.0 x8			
Drive Cage - included	Not included	4 LFF	20 EDSFF E3.S 1T	Not included
Network Controller	Choice of either OCP 3.0 or select stand-up network adapters for primary networking selection plus additional/optional stand-up network adapters **Notes:** No embedded networking			
Cooling	Choice of Standard, Performance, or Closed-Loop Liquid Cooling Heat Sink Choice of Standard, Performance, or Liquid Cooling Fan Kit			
Management	Default: HPE iLO Standard with Intelligent Provisioning, HPE OneView Standard (requires download), HPE GreenLake for Compute Ops Management (subscription included)			
Video	1 VGA rear			
USB	Front: 1 USB 3.2 Gen1 + iLO service port Rear: 2 USB 3.2 Gen1 Internal: 2 USB 3.2 Gen1			
Security	TPM2.0 (Trusted Platform Module) embedded			
Rail Kit	Optional Easy Install rails and CMA			
Form Factor	1U Rack			
Warranty	3-year parts, 3-year labor, 3-year onsite support with next business day response.			

Notes:

- HPE offers multiple Trade Agreement Act (TAA) compliant configurations to meet the needs of US Federal Government customers. These products are either manufactured or substantially transformed in a designated country. TAA compliance is only provided when HPE options are included as part of factory integrated orders (CTO).
- TAA compliant configuration requires TAA versions of the CTO Server SKUs.
- HPE Trusted Supply Chain (P36394-B21) is an optional security upgrade intended for agencies and regulated industries needing enhanced security and compliance needs. Applying this option to a DL325 Gen11 CTO server ensures it is built in the USA in a secured facility by vetted HPE personnel assigned to the manufacturing processes. A multitude of checkpoints/inspections for malicious microcode and counterfeit parts are performed throughout the server build, and additional safeguards are put in place against cyber-exploits throughout the server lifecycle. See "HPE Security" section within this document for more detail and learn more at **http://www.hpe.com/security**
- All CTO servers are Energy Star 4.0 compliant.

Configuration Information

CTO Server	8SFF CTO server	4LFF CTO server	EDSFF CTO server	GPU CTO server
Included Drive Cage	Not available	4 LFF backplane	20 EDSFF backplane	Not available
Universal Media Bay	1 Optional	Not Available	Not Available	Not available
ODD	1 Optional	1 Optional	Not Available	Not available
4 LFF SAS/SATA	Not Available	1 Optional	Not Available	Not available
8 SFF SAS/SATA	1 Optional	Not Available	Not Available	Not available
8 SFF NVMe	1 Optional	Not Available	Not Available	Not available
2 SFF SAS/SATA	1 Optional	Not Available	Not Available	Not available
2 SFF NVMe	1 Optional	Not Available	Not Available	Not available
20 EDSFF NVMe	Not Available	Not Available	1 Optional	Not available
4 SFF NVMe	Not Available	Not Available	Not Available	1 Optional
8 EDSFF NVMe	Not Available	Not Available	Not Available	1 Optional

Notes:

– This applies to CTO configurations, field upgrades may differ depending on field configuration.

– Drive cage kits need to be ordered separately for the 8SFF CTO server and GPU CTO server.

Step 2: Choose Core Options

- Choice of 1 Processor model and Heat Sink Kit
  Requires necessary Heat Sink for different processor wattage.
- Choice of DDR5 memory options.
  Requires necessary Fan Kits for different memory configurations and subjects to the recommended system ambient temperature.
- Choice of Drive cage, Storage Controllers, and Storage Controller Cables
- Choice of SSD, HDD, and Optical Drive
- Choice of OS Boot Devices
- Choice of Riser Cards
- Choice of Networking options
  PCIe standup or OCP 3.0. Requires necessary Fan Kits and subjects to the recommended system ambient temperature.
- Choice of Accelerator options
- Choice of Power and Cooling options
- Choice of Security options
- Choice of Software as a Service Management - HPE GreenLake for Compute Ops Management and HPE OneView

Step 3: Choose Additional Options

- Choice of Embedded Management
- Choice of Rail Kits
- Choice of Rack options
- Choice of Support Services

Choice of Core Options

Processor

Please select ONE 4th or 5th Generation AMD EPYC Processor

5th Generation AMD EPYC Processor

AMD EPYC 9845 2.1GHz 160-core 390W Processor for HPE	P72646-B21
AMD EPYC 9825 2.2GHz 144-core 390W Processor for HPE	P72647-B21
AMD EPYC 9745 2.4GHz 128-core 400W Processor for HPE	P72648-B21
AMD EPYC 9645 2.3GHz 96-core 320W Processor for HPE	P72649-B21
AMD EPYC 9655P 2.6GHz 96-core 400W Processor for HPE	P72662-B21
AMD EPYC 9565 3.15GHz 72-core 400W Processor for HPE	P72651-B21
AMD EPYC 9535 2.4GHz 64-core 300W Processor for HPE	P72652-B21
AMD EPYC 9575F 3.3GHz 64-core 400W Processor for HPE	P72758-B21
AMD EPYC 9555P 3.2GHz 64-core 360W Processor for HPE	P72663-B21
AMD EPYC 9475F 3.65GHz 48-core 400W Processor for HPE	P72666-B21
AMD EPYC 9455P 3.15GHz 48-core 300W Processor for HPE	P72664-B21
AMD EPYC 9365 3.4GHz 36-core 300W Processor for HPE	P72655-B21
AMD EPYC 9335 3GHz 32-core 210W Processor for HPE	P72656-B21
AMD EPYC 9375F 3.8GHz 32-core 320W Processor for HPE	P72667-B21
AMD EPYC 9355P 3.55GHz 32-core 280W Processor for HPE	P72665-B21
AMD EPYC 9255 3.25GHz 24-core 200W Processor for HPE	P72658-B21
AMD EPYC 9275F 4.1GHz 24-core 320W Processor for HPE	P72668-B21
AMD EPYC 9135 3.65GHz 16-core 200W Processor for HPE	P72660-B21
AMD EPYC 9115 2.6GHz 16-core 155W Processor for HPE	P72659-B21
AMD EPYC 9175F 4.2GHz 16-core 320W Processor for HPE	P72669-B21
AMD EPYC 9015 3.6GHz 8-core 155W Processor for HPE	P72661-B21

4th Generation AMD EPYC Processor

AMD EPYC 9754 2.25GHz 128-core 360W Processor for HPE	P60463-B21
AMD EPYC 9734 2.2GHz 112-core 340W Processor for HPE	P60465-B21
AMD EPYC 9654P 2.4GHz 96-core 360W Processor for HPE	P53697-B21
AMD EPYC 9684X 2.55GHz 96-core 400W Processor for HPE	P63493-B21
AMD EPYC 9634 2.25GHz 84-core 290W Processor for HPE	P53705-B21
AMD EPYC 9534 2.45GHz 64-core 280W Processor for HPE	P53699-B21
AMD EPYC 9554P 3.1GHz 64-core 360W Processor for HPE	P53703-B21
AMD EPYC 9454P 2.75GHz 48-core 290W Processor for HPE	P53709-B21
AMD EPYC 9474F 3.6GHz 48-core 360W Processor for HPE	P53706-B21
AMD EPYC 9334 2.7GHz 32-core 210W Processor for HPE	P53712-B21
AMD EPYC 9354P 3.25GHz 32-core 280W Processor for HPE	P53704-B21
AMD EPYC 9374F 3.85GHz 32-core 320W Processor for HPE	P53710-B21
AMD EPYC 9384X 3.1GHz 32-core 320W Processor for HPE	P63492-B21
AMD EPYC 9254 2.9GHz 24-core 200W Processor for HPE	P53707-B21
AMD EPYC 9224 2.5GHz 24-core 200W Processor for HPE	P58540-B21
AMD EPYC 9274F 4.05GHz 24-core 320W Processor for HPE	P53711-B21
AMD EPYC 9124 3.0GHz 16-core 200W Processor for HPE	P53702-B21
AMD EPYC 9174F 4.1GHz 16-core 320W Processor for HPE	P53698-B21
AMD EPYC 9184X 3.55GHz 16-core 320W Processor for HPE	P63491-B21

Notes:

– Processors less than or equal to 240W require Standard Heat Sink (P58456-B21).

Core Options

– Processors more than 240W and less than or equal to 300W require Performance Heat Sink (P58457-B21)
– Processors more than or equal to 320W require Closed-Loop Liquid Cooling Heat Sink (P58463-B21).
– The supported system ambient temperature of EPYC 9254 is 25C
– The supported system ambient temperature of EPYC 9384X is 25C and cannot support with EDSFF CTO server

Memory

Please select one or more memory from below.
For new DDR5 memory, please go to **HPE DDR5 Smart Memory QuickSpecs**
For details on the Memory Population Rules and Guidelines with AMD EPYC 9004 series processors, please go to: **https://www.hpe.com/psnow/doc/a50007481enw**

Notes:

– Quantity of memory DIMMs selected per socket must be 1, 2, 4, 6, 8, 10, or 12.
– The maximum memory speed and capacity is a function of the memory type, memory configuration, and processor model
– The maximum speed capability of the memory system is governed by the combination of the CPU and any other DIMMs installed in the server. If higher speed DIMMs are installed with a CPU that only supports a lower memory speed, the DIMMs will only run at the (lower) memory speed supported by the processor. Likewise, if memory DIMMs are mixed with slower DIMMs within a server, all DIMMs will run at the slower memory speed. For further information please refer to the Memory Population Rules for your specific server.

Registered DIMMs DDR5 (RDIMMs)

DDR5-6400 (applies to the 5th Generation AMD® EPYC® Processors)

HPE 16GB (1x16GB) Single Rank x8 DDR5-6400 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64984-B21
HPE 32GB (1x32GB) Dual Rank x8 DDR5-6400 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64985-B21
HPE 64GB (1x64GB) Dual Rank x4 DDR5-6400 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64986-B21
HPE 96GB (1x96GB) Dual Rank x4 DDR5-6400 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64987-B21
HPE 128GB (1x128GB) Dual Rank x4 DDR5-6400 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64988-B21

DDR5-4800 (applies to the 4th Generation AMD® EPYC® Processors)

HPE 16GB (1x16GB) Single Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P50309-B21
HPE 32GB (1x32GB) Dual Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P50311-B21
HPE 64GB (1x64GB) Dual Rank x4 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P50312-B21
HPE 96GB (1x96GB) Dual Rank x4 DDR5-4800 CAS-46-45-45 EC8 Registered Smart Memory Kit	P66676-B21
HPE 128GB (1x128GB) Quad Rank x4 DDR5-4800 CAS-46-39-39 EC8 Registered 3DS Smart Memory Kit	P50313-B21
HPE 128GB (1x128GB) Dual Rank x4 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P69982-B21

Core Options

HPE 256GB (1x256GB) Octal Rank x4 DDR5-4800 CAS-46-39-39 EC8 Registered 3DS Smart Memory Kit — P50314-B21

Notes:

- DDR5-6400 Memory SKUs offer a transfer rate of 6000MT/s with 5th Gen EPYC Processors.
- Mixing of x4 memory and x8 memory is not supported
- Mixing of 3DS memory and non-3DS memory is not supported.
- Supported memory configuration and recommended system ambient temperature:

	SFF/LFF/GPU CTO server			EDSFF CTO server	
Memory	**Std Fans (P58461-B21)**	**Perf Fans (P58462-B21)**	**LC Fans (P59668-B21)**	**Perf Fans (P58462-B21)**	**LC Fans (P59668-B21)**
<= 64GB DIMM	30C	30C	30C	25C	25C
96GB DIMM 128GB DIMM	Not Support	30C	25C	25C	25C Max = 8 for 128G DIMM
256GB DIMM	Not Support	25C	25C Max = 4	25C	Not Support

Notes:

- Not Support = Configuration not allowed because of thermal limitation.
- Requires Performance or Liquid Cooling Fan Kit for 96GB, 128GB and 256GB DIMMs.
- Max=4 of 256GB DIMM memory can be selected if the Liquid Cooling Heat Sink/ Liquid Cooling Fan Kit is selected.
- Max=8 of 96/128GB DIMM memory can be selected if the Liquid Cooling Heat Sink/ Liquid Cooling Fan Kit is selected along with the EDSFF CTO server.

Storage

Drive cages

Notes:

- For the 8SFF CTO server, If 8SFF Backplane is not selected then Internal Controllers, Controller cables and Drives must not be allowed for selection. This config will be shipped as a driveless config.
- Maximum one (1) 2SFF backplane kits can be selected together with 8SFF backplane kit, to support up to 10SFF in total.
- The type of drives that each drive cage supports are listed in the below table.

PN	Description	SATA	SAS	NVMe U.3 Static SSD	NVMe U.3 SSD	NVMe U.2 SSD
P54999-B21	HPE DL325 Gen11 8SFF x1 TM BP Kit	X	X	X	X	Not Support
P55000-B21	HPE DL325 Gen11 8SFF x4 TM BP Kit	X	X	X	X	Not Support
P56652-B21	HPE DL325 Gen11 2SFF x4 TM BP Kit	X	X	X	X	Not Support
P64521-B21	HPE DL325 Gen11 4SFF x4 NVMe Kit	X	X	X	X	Not Support

HPE ProLiant DL325 Gen11 8SFF x1 Tri-Mode U.3 Backplane Kit — P54999-B21

Core Options

Notes:
- Supports 8 SFF SAS/SATA/ NVMe (U.3) Basic Carrier (BC) Drives.
- This drive cage can only be selected with 8SFF CTO Server.
- Max = 1.
- Requires Tri-Mode controllers if NVMe u.3 drives are selected with this backplane kit.
- if this Backplane kit is selected then one of the following cable options is supported:
  - with PCIe controllers: 8SFF x1 Tri-Mode Secondary Cable Kit (P57009-B21).
  - with OCP controllers: 8SFF x1 OCP2 Tri-Mode Cable Kit (P59619-B21).
  - Onboard SATA: no cable kit selection required.

HPE ProLiant DL325 Gen11 8SFF x4 Tri-Mode U.3 BC Backplane Kit P55000-B21

Notes:
- Supports 8 SFF SAS/SATA/ NVMe (U.3) Basic Carrier (BC) Drives.
- This drive cage can only be selected with 8SFF CTO Server.
- Max = 1.
- Requires Tri-Mode controllers if SAS/SATA SFF drives are selected with this backplane kit.
- if this 8SFF x4 U.3 Backplane kit is selected then one of the following cable options is supported:
  - with SR932i-p: 8SFF x4 Primary SR932i-p Tri-Mode Cable Kit (P57004-B21) or 8SFF x4 Secondary SR932i-p Tri-Mode Cable Kit (P57005-B21).
  - with PCIe controllers: 8SFF x2 Tri-Mode Secondary Cable Kit (P57006-B21).
  - with OCP controllers: 8SFF x2 Tri-Mode OCP2 Cable Kit (P57008-B21).
  - NVMe Direct Attach: no cable kit selection required.
- Requires Performance Fan Kit (P58462-B21) or Liquid Cooling Fan Kits (P59668-B21).

HPE ProLiant DL325 Gen11 2SFF x4 Tri-Mode U.3 BC Backplane Kit P56652-B21

Notes:
- Supports 2 SFF SAS/SATA/ NVMe (U.3) Basic Carrier (BC) Drives.
- This drive cage can only be selected with 8SFF CTO Server.
- Max = 1.
- if this 2SFF U.3 Backplane kit is selected then one of the following cable options is supported:
  - with PCIe controllers: 2SFF x4 Secondary Tri-Mode Cable Kit (P59621-B21).
  - with OCP controllers: 2SFF x4 OCP2 Tri-Mode Cable Kit (P59620-B21).
  - NVMe Direct Attach: no cable kit selection required.
  - Onboard SATA: 2SFF SATA Direct Attach Cable Kit (P59617-B21).
- Requires 8SFF x1 U.3 Backplane Kit (P54999-B21) or 8SFF x4 U.3 Backplane Kit (P55000-B21) in the order.
- If this drive cage is selected then optical drives (726536-B21 & 726537-B21) cannot be selected.

HPE ProLiant DL325 Gen11 4SFF x4 NVMe Drive Cage Kit P64521-B21

Notes:
- Supports 4 SFF NVMe (U.3) Basic Carrier (BC) Drives.
- This drive cage can only be selected with GPU CTO Server.

Core Options

– Max = 1.
– if this 4SFF U.3 Backplane kit is selected then one of the following cable options is supported:
  - with PCIe controllers: 4SFF x4 Secondary Tri-Mode Cable Kit (P70318-B21).
  - with OCP controllers: 4SFF x4 OCP2 Tri-Mode Cable Kit (P69876-B21).
  - NVMe Direct Attach: no cable kit selection required.

HPE ProLiant DL325 Gen11 8EDSFF x4 Drive Cage Kit P64522-B21

Notes:
– Supports 8 EDSFF NVMe Drives direct attach. No additional cable kit selection required
– This drive cage can only be selected with GPU CTO Server.
– Max = 1.

HPE ProLiant DL325 Gen11 GPU 4SFF x2 OCP Tri-Mode Backplane Kit P70287-B21

Notes:
– Supports 4 SFF NVMe (U.3) Basic Carrier (BC) Drives.
– This drive cage can only be selected with GPU CTO Server.
– Max = 1.
– This drive cage can only support connection to OCP controllers in x2 bandwidth

Storage Controller

The Gen11 storage controller portfolio has been updated to include new technology like OCP3.0 as well as PCIe adapters. For a more detailed breakout of the available Gen11 controllers visit the storage controllers QuickSpecs site:

HPE Compute MR Gen11 Controllers QuickSpecs

HPE Compute SR Gen11 Controllers QuickSpecs

Notes:
– When selecting SR RAID controllers for external storage (E208e-p, 804398-B21) and MR RAID controllers for internal storage, please be aware these two products use different RAID configuration tools.
– Mixing of MR (MegaRAID) series controllers and SR (SmartRAID) series controllers is not allowed.

HPE Smart Array E208e-p SR Gen10 (8 External Lanes/No Cache) 12G SAS PCIe Plug-in Controller 804398-B21

Notes:
– This controller supports up to 8 SAS/SATA Drives (external).
– Controller Based Encryption (CBE) with a remote key management server is not supported. Local key management(LKM) is supported.
– One Button Secure Erase (OBSE) used to sanitize drives and factory reset the controller is not supported.

HPE MR216i-o Gen11 x16 Lanes without Cache OCP SPDM Storage Controller P47789-B21

Notes: This controller supports up to 16 SAS/SATA/NVMe Drives.

HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller P58335-B21

Notes:
– This controller supports up to 8 SAS/SATA/NVMe Drives.
– Requires 96W Smart Stg Li-ion Batt 145mm Kit (P01366-B21) or Smart Hybrid Capacitor w/ 145mm Cable (P02377-B21).

HPE MR416i-o Gen11 x16 Lanes 8GB Cache OCP SPDM Storage Controller P47781-B21

Core Options

Notes:
– This controller supports up to 16 SAS/SATA/NVMe Drives.
– Requires 96W Smart Stg Li-ion Batt 145mm Kit (P01366-B21) or Smart Hybrid Capacitor w/ 145mm Cable (P02377-B21).

HPE MR216i-p Gen11 x16 Lanes without Cache PCI SPDM Plug-in Storage Controller P47785-B21

Notes: This controller supports up to 16 SAS/SATA/NVMe Drives.

HPE MR416i-p Gen11 x16 Lanes 8GB Cache PCI SPDM Plug-in Storage Controller P47777-B21

Notes:
– This controller supports up to 16 SAS/SATA/NVMe Drives.
– Requires 96W Smart Stg Li-ion Batt 145mm Kit (P01366-B21) or Smart Hybrid Capacitor w/ 145mm Cable (P02377-B21).

HPE SR932i-p Gen11 x32 Lanes 8GB Wide Cache PCI SPDM Plug-in Storage Controller P47184-B21

Notes:
– This controller supports up to 32 SAS/SATA/NVMe Drives.
– Requires 96W Smart Stg Li-ion Batt 145mm Kit (P01366-B21) or Smart Hybrid Capacitor w/ 145mm Cable (P02377-B21).

Battery and Hybrid Capacitor

HPE Smart Storage Hybrid Capacitor with 145mm Cable Kit P02377-B21

HPE 96W Smart Storage Lithium-ion Battery with 145mm Cable Kit P01366-B21

HPE ProLiant DL325 Gen11 Megacell Extension Cable Kit P56659-B21

Notes:
– If HPE 96W Smart Stg Li-ion Batt 145mm Kit is selected then HPE Smart Hybrid Capacitor 145mm kit cannot be selected and vice versa.
– If M.2 enablement Kit and "96W Smart Stg Li-ion Batt 145mm Kit OR Smart Hybrid Capacitor w/ 145mm Kit" are selected then Megacell Ext Cable Kit must be selected.

Storage Controller Cables

HPE ProLiant DL325 Gen11 8SFF x4 Primary SR932i-p Tri-Mode Cable Kit P57004-B21

Notes: Supports 8 SFF U.3 SAS/SATA/NVMe connecting to SR932i-p controllers at the primary riser slot with up to x4 speed.

HPE ProLiant DL325 Gen11 8SFF x4 Secondary SR932i-p Tri-Mode Cable Kit P57005-B21

Notes: Supports 8 SFF U.3 SAS/SATA/NVMe connecting to SR932i-p controllers at the secondary riser slot with up to x4 speed.

HPE ProLiant DL325 Gen11 8SFF x2 Secondary Tri-Mode Cable Kit P57006-B21

Notes: Supports 8 SFF U.3 SAS/SATA/NVMe connecting to storage controllers at primary riser slot with x2 speed.

HPE ProLiant DL325 Gen11 8SFF x2 OCP2 Tri-Mode Cable Kit P57008-B21

Notes: Supports 8 SFF U.3 SAS/SATA/NVMe connecting to storage controllers at OCP22 slot with x2 speed.

HPE ProLiant DL325 Gen11 8SFF x1 Secondary Tri-Mode Cable Kit P57009-B21

Notes: Supports 8 SFF U.3 SAS/SATA/NVMe connecting to storage controllers at secondary riser slot with x1 speed.

HPE ProLiant DL325 Gen11 8SFF x1 OCP2 Tri-Mode Cable Kit P59619-B21

Notes: Supports 8 SFF U.3 SAS/SATA/NVMe connecting to storage controllers at OCP22 slot with x1 speed.

Core Options

HPE ProLiant DL325 Gen11 2SFF SATA Direct Attach Cable Kit **Notes:** Supports 2 SFF SATA direct attach.	P59617-B21
HPE ProLiant DL325 Gen11 2SFF x4 OCP2 Tri-Mode Cable Kit **Notes:** Supports 2 SFF U.3 SAS/SATA/NVMe connecting to storage controllers at OCP22 slot with up to x4 speed.	P59620-B21
HPE ProLiant DL325 Gen11 2SFF x4 Secondary Tri-Mode Cable Kit **Notes:** Supports 2 SFF U.3 SAS/SATA/NVMe connecting to storage controllers at secondary riser slot with up to x4 speed.	P59621-B21
HPE ProLiant DL325 Gen11 20EDSFF x2 NVMe Direct Attach Cable Kit **Notes:** Supports 20 EDSFF E3.S 1T NVMe direct attach with x2 speed.	P57010-B21
HPE ProLiant DL3X5 Gen11 16EDSFF x2 PCIe Tri-Mode Cable Kit **Notes:** Supports 16 EDSFF NVMe connecting to storage controllers at Primary riser with up to x2 speed.	P69878-B21
HPE ProLiant DL3X5 Gen11 GPU 4SFF x4 PCIe Tri-Mode Cable Kit **Notes:** Supports 4 SFF SAS/SATA/NVMe connecting to storage controllers at secondary riser slot with up to x4 speed in GPU CTO server.	P70318-B21
HPE ProLiant DL365 Gen11 GPU 4SFF x4 OCP Tri-Mode Cable Kit **Notes:** Supports 4 SFF SAS/SATA/NVMe connecting to storage controllers at OCP22 slot with up to x4 speed in GPU CTO server.	P69876-B21
HPE ProLiant DL3X5 Gen11 1P GPU 8SFF/EDSFF x4 Tri-Mode PCIe Cable Kit **Notes:** Supports 8 EDSFF NVMe connecting to storage controllers at Primary riser with up to x4 speed in GPU CTO server.	P70406-B21

Supported Storage Configurations

8SFF CTO server

Drives				Backplane		Storage Controller + Cable Kit
Max Qty	SAS	SATA	U.3 NVMe	Box1	Box2	
8	-	8	-	P54999-B21	-	8SFF DA (SATA)
8	8	8	8	P54999-B21	-	OCP Ctrlr + P59619-B21
8	8	8	8	P54999-B21	-	PCIe Ctrlr + P57009-B21
8	-	-	8	P55000-B21	-	8SFF DA (NVMe x4)
8	8	8	8	P55000-B21	-	8SFF x4 SR932i-p + P57004-B21 (Pri.)
8	8	8	8	P55000-B21	-	8SFF x4 SR932i-p + P57005-B21 (Sec.)
8	8	8	8	P55000-B21	-	8SFF x2 PCIe Ctrlr + P57006-B21 (Pri.)
8	8	8	8	P55000-B21	-	8SFF x2 OCP Ctrlr + P57008-B21
10	-	8	2	P54999-B21	P56652-B21	8SFF DA (SATA); 2SFF DA (NVMe x4)
10	-	10	-	P54999-B21	P56652-B21	10SFF DA (SATA) + P59617-B21

Core Options

10	2	10	2	P54999-B21	P56652-B21	8SFF DA (SATA); 2SFF OCP Ctrlr + P59620-B21
10	8	8	10	P54999-B21	P56652-B21	8SFF x1 OCP Ctrlr + P59619-B21; 2SFF DA (NVMe x4)
10	8	10	8	P54999-B21	P56652-B21	8SFF x1 OCP Ctrlr + P59619-B21; 2SFF DA (SATA) + P59617-B21
10	10	10	10	P54999-B21	P56652-B21	8SFF x1 OCP Ctrlr + P59619-B21; 2SFF x4 PCIe Ctrlr + P59621-B21 (Sec.)
10	10	10	10	P54999-B21	P56652-B21	10SFF x1 OCP Ctrlr + P59619-B21 & P59620-B21
10	10	10	10	P54999-B21	P56652-B21	10SFF x1 PCIe Ctrlr + P57009-B21 & P59621-B21 (Sec.)
10	-	-	10	P55000-B21	P56652-B21	10SFF DA (NVMe x4)
10	-	10	8	P55000-B21	P56652-B21	8SFF DA (NVMe x4); 2SFF DA (SATA) + P59617-B21
10	2	10	10	P55000-B21	P56652-B21	8SFF DA (NVMe x4); 2SFF x4 PCIe Ctrlr + P59621-B21 (Sec.)
10	2	10	10	P55000-B21	P56652-B21	8SFF DA (NVMe x4); 2SFF x4 OCP Ctrlr + P59620-B21
10	10	10	10	P55000-B21	P56652-B21	8SFF x4 SR932i-p + P57004-B21 (Pri.); 2SFF OCP Ctrlr + P59620-B21
10	8	8	10	P55000-B21	P56652-B21	8SFF x4 SR932i-p + P57004-B21 (Pri.); 2SFF DA (NVMe x4)
10	8	10	8	P55000-B21	P56652-B21	8SFF x4 SR932i-p + P57004-B21 (Pri.); 2SFF DA (SATA) + P59617-B21
10	10	10	10	P55000-B21	P56652-B21	8SFF x4 SR932i-p + P57005-B21 (Sec.); 2SFF OCP Ctrlr + P59620-B21
10	8	8	10	P55000-B21	P56652-B21	8SFF x4 SR932i-p + P57005-B21 (Sec.); 2SFF DA (NVMe x4)
10	8	10	8	P55000-B21	P56652-B21	8SFF x4 SR932i-p + P57005-B21 (Sec.); 2SFF DA (SATA) + P59617-B21
10	10	10	10	P55000-B21	P56652-B21	10SFF x2 SR932i-p + P57006-B21 & P59621-B21 (Sec.)
10	10	10	10	P55000-B21	P56652-B21	8SFF x2 OCP Ctrlr + P57008-B21; 2SFF PCIe Ctrlr + P59621-B21 (Sec.)
10	10	10	10	P55000-B21	P56652-B21	8SFF x2 PCIe Ctrlr + P57006-B21 (Sec.); 2SFF OCP Ctrlr + P59620-B21

Notes:

– DA = Direct Attach; Ctrlr = controller

– If no controller or cable kit information in the table then cable kit selection is not required.

GPU CTO server

Core Options

Drives					Backplane		Storage Controller + Cable Kit
Max Qty	SAS	SATA	U.3 NVMe	EDSFF	Box1	Box2	
4	-	-	-	-	P64521-B21	-	4SFF NVMe x4 DA
4	4	4	4		P64521-B21	-	PCIe Ctrlr + P70318-B21
4	4	4	4		P64521-B21	-	OCP Ctrlr + P69876-B21 (x4 speed)
4	4	4	4		P70287-B21	-	OCP Ctrlr (x2 speed)
8	-	-	-	8	P64522-B21	-	8EDSFF NVMe x4 DA
8	-	-	-	8	P64522-B21	-	SR932i-p + P70406-B21

EDSFF CTO server

Drives					Backplane		Storage Controller + Cable Kit
Max Qty	SAS	SATA	U.3 NVMe	EDSFF	Box1	Box2	
20	-	-	-	20	Included	-	20EDSFF NVMe x4 DA
20	-	-	-	20	Included	-	P57010-B21 (x2 DA)
16	-	-	-	16	Included	-	SR932i-p + P69878-B21

Notes:

– DA = Direct Attach; Ctrlr = controller; Included = item included in the CTO server or option kit.

– If no controller or cable kit information in the table then cable kit selection is not required.

HPE Drives

Solid State Drives

For SSD selection guidance, please visit **https://ssd.hpe.com/**

Read Intensive - 12G SAS - SFF

HPE 960GB SAS 12G Read Intensive SFF BC Value SAS Multi Vendor SSD	P40506-B21
HPE 1.92TB SAS 12G Read Intensive SFF BC Value SAS Multi Vendor SSD	P40507-B21
HPE 3.84TB SAS 12G Read Intensive SFF BC Value SAS Multi Vendor SSD	P40508-B21
HPE 7.68TB SAS 12G Read Intensive SFF BC Value SAS Multi Vendor SSD	P40509-B21
HPE 1.92TB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49031-B21
HPE 3.84TB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49035-B21
HPE 7.68TB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49041-B21
HPE 15.36TB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49045-B21

Mixed Use - 12G SAS - SFF

HPE 960GB SAS 12G Mixed Use SFF BC Value SAS Multi Vendor SSD	P40510-B21
HPE 1.92TB SAS 12G Mixed Use SFF BC Value SAS Multi Vendor SSD	P40511-B21
HPE 3.84TB SAS 12G Mixed Use SFF BC Value SAS Multi Vendor SSD	P40512-B21
HPE 1.6TB SAS 24G Mixed Use SFF BC Multi Vendor SSD	P49049-B21
HPE 800GB SAS 24G Mixed Use SFF BC Multi Vendor SSD	P49047-B21
HPE 3.2TB SAS 24G Mixed Use SFF BC Multi Vendor SSD	P49053-B21
HPE 6.4TB SAS 24G Mixed Use SFF BC Multi Vendor SSD	P49057-B21

Core Options

HPE 960GB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49029-B21
Mixed Use - 12G SAS- LFF	
HPE 960GB SAS 12G Mixed Use LFF LPC Value SAS Multi Vendor SSD	P37009-B21
Read Intensive - 6G SATA - SFF	
HPE 240GB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40496-B21
HPE 480GB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40497-B21
HPE 960GB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40498-B21
HPE 1.92TB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40499-B21
HPE 3.84TB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40500-B21
HPE 7.68TB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40501-B21
HPE 480GB SATA 6G Read Intensive SFF BC PM893a SSD	P63886-B21
HPE 3.84TB SATA 6G Read Intensive SFF BC PM893a SSD	P63910-B21
Mixed Use - 6G SATA - SFF	
HPE 480GB SATA 6G Mixed Use SFF BC Multi Vendor SSD	P40502-B21
HPE 960GB SATA 6G Mixed Use SFF BC Multi Vendor SSD	P40503-B21
HPE 1.92TB SATA 6G Mixed Use SFF BC Multi Vendor SSD	P40504-B21
HPE 3.84TB SATA 6G Mixed Use SFF BC Multi Vendor SSD	P40505-B21
Read Intensive - 6G SATA - LFF	
HPE 960GB SATA 6G Read Intensive LFF LPC Multi Vendor SSD	P47808-B21
Read Intensive - NVMe - SFF	
HPE 1.92TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PM1733a SSD	P50216-B21
HPE 3.84TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PM1733a SSD	P50219-B21
HPE 7.68TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PM1733a SSD	P50222-B21
HPE 15.36TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PM1733a SSD	P50224-B21
HPE 1.92TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 CM7 SSD	P63829-B21
HPE 3.84TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 CM7 SSD	P63833-B21
HPE 7.68TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 CM7 SSD	P63837-B21
HPE 15.36TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 CM7 SSD	P63841-B21
HPE 960GB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static V2 Multi Vendor SSD	P64842-B21
HPE 1.92TB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static V2 Multi Vendor SSD	P64844-B21
HPE 3.84TB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static V2 Multi Vendor SSD	P64846-B21
HPE 7.68TB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static V2 Multi Vendor SSD	P64848-B21
HPE 15.36TB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static SPDM Multi Vendor SSD	P69255-B21
HPE 7.68TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PS1010 SSD	P70434-B21
HPE 15.36TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PS1010 SSD	P70436-B21
Mixed Use - NVMe - SFF	
HPE 1.6TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD	P50227-B21
HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD	P50230-B21
HPE 6.4TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD	P50233-B21
HPE 1.6TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 CM7 SSD	P63845-B21

Core Options

HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 CM7 SSD	P63849-B21
HPE 6.4TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 CM7 SSD	P63853-B21
HPE 800GB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static V2 Multi Vendor SSD	P64999-B21
HPE 1.6TB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static V2 Multi Vendor SSD	P65007-B21
HPE 3.2TB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static V2 Multi Vendor SSD	P65015-B21
HPE 6.4TB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static V2 Multi Vendor SSD	P65023-B21
HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PS1030 SSD	P70426-B21
HPE 6.4TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PS1030 SSD	P70428-B21
SED (Self-Encryption Drive) - SATA- SFF	
HPE 960GB SATA 6G Mixed Use SFF BC Self-encrypting 5400M SSD	P58244-B21
HPE 480GB SATA 6G Read Intensive SFF BC Self-encrypting 5400P SSD	P58236-B21
SED (Self-Encryption Drive) - SAS SFF	
HPE 1.6TB SAS Mixed Use SFF BC Self-encrypting FIPS 140-2 PM7 SSD	P63871-B21
HPE 3.84TB SAS Read Intensive SFF BC Self-encrypting FIPS 140-2 PM7 SSD	P63875-B21
SED (Self-Encryption Drive) - NVMe SFF	
HPE 1.6TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 Self-encrypting FIPS 140-3 CM7 SSD	P61043-B21
HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 Self-encrypting FIPS 140-3 CM7 SSD	P61051-B21
HPE 6.4TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 Self-encrypting FIPS 140-3 CM7 SSD	P61059-B21
HPE 1.92TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 Self-encrypting FIPS 140-3 CM7 SSD	P61019-B21
HPE 3.84TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 Self-encrypting FIPS 140-3 CM7 SSD	P61027-B21
HPE 7.68TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 Self-encrypting FIPS 140-3 CM7 SSD	P61035-B21
Read Intensive - NVMe - EDSFF E3.S 1T	
HPE 3.84TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM PM1743 SSD	P57799-B21
HPE 7.68TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM PM1743 SSD	P57803-B21
HPE 15.36TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM PM1743 SSD	P57807-B21
HPE 3.84TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM CM7 SSD	P61179-B21
HPE 7.68TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM CM7 SSD	P61183-B21
HPE 15.36TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM CM7 SSD	P61187-B21
HPE 1.92TB NVMe Gen5 Mainstream Performance Read Intensive E3S EC1 CD8P SSD	P69234-B21
HPE 3.84TB NVMe Gen5 Mainstream Performance Read Intensive E3S EC1 CD8P SSD	P69237-B21
HPE 7.68TB NVMe Gen5 Mainstream Performance Read Intensive E3S EC1 CD8P SSD	P69239-B21
HPE 15.36TB NVMe Gen5 Mainstream Performance Read Intensive E3S EC1 CD8P SSD	P69546-B21

Core Options

HPE 3.84TB NVMe Gen5 High Performance Read Intensive E3S EC1 PS1010 SSD	P70392-B21
HPE 7.68TB NVMe Gen5 High Performance Read Intensive E3S EC1 PS1010 SSD	P70395-B21
HPE 15.36TB NVMe Gen5 High Performance Read Intensive E3S EC1 PS1010 SSD	P70397-B21
Mixed Use - NVMe - EDSFF E3.S 1T	
HPE 3.2TB NVMe Gen5 High Performance Mixed Use E3S EC1 EDSFF SPDM CM7 SSD	P61191-B21
HPE 6.4TB NVMe Gen5 High Performance Mixed Use E3S EC1 EDSFF SPDM CM7 SSD	P61195-B21
HPE 1.6TB NVMe Gen5 Mainstream Performance Mixed Use E3S EC1 CD8P SSD	P69241-B21
HPE 3.2TB NVMe Gen5 Mainstream Performance Mixed Use E3S EC1 CD8P SSD	P69243-B21
HPE 6.4TB NVMe Gen5 Mainstream Performance Mixed Use E3S EC1 CD8P SSD	P69245-B21
HPE 3.2TB NVMe Gen5 High Performance Mixed Use E3S EC1 PS1030 SSD	P70399-B21
HPE 6.4TB NVMe Gen5 High Performance Mixed Use E3S EC1 PS1030 SSD	P70401-B21
HPE 12.8TB NVMe Gen5 High Performance Mixed Use E3S EC1 PS1030 SSD	P70403-B21
Very Read Optimized - NVMe - EDSFF E3.S 1T	
HPE 3.84TB NVMe Gen4 Mainstream Performance Very Read Optimized E3S EC1 EDSFF P5430 SSD	P63930-B21
HPE 7.68TB NVMe Gen4 Mainstream Performance Very Read Optimized E3S EC1 EDSFF P5430 SSD	P63934-B21
HPE 15.36TB NVMe Gen4 Mainstream Performance Very Read Optimized E3S EC1 EDSFF P5430 SSD	P63938-B21
SED (Self-Encryption Drive) - NVMe - EDSFF E3.S 1T	
HPE 3.2TB NVMe Gen5 High Performance Mixed Use E3S EC1 Self-encrypting FIPS 140-2 CM7 SSD	P70669-B21
HPE 6.4TB NVMe Gen5 High Performance Mixed Use E3S EC1 Self-encrypting FIPS 140-2 CM7 SSD	P70672-B21
HPE 7.68TB NVMe Gen5 High Performance Read Intensive E3S EC1 Self-encrypting FIPS 140-2 CM7 SSD	P70674-B21

Hard Disk Drive

Enterprise - 12G SAS - SFF Drives

HPE 2.4TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty 512e Multi Vendor HDD	P28352-B21
HPE 1.2TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Multi Vendor HDD	P28586-B21
HPE 300GB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Multi Vendor HDD	P40430-B21
HPE 600GB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Multi Vendor HDD	P53561-B21
HPE 1.8TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty 512e Multi Vendor HDD	P53562-B21

Notes: If 15K drives are selected then 25C is the recommended system ambient temperature.

Midline - 12G SAS - LFF Drives

HPE 12TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e Multi Vendor HDD	881781-B21
HPE 14TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e Multi Vendor HDD	P09155-B21
HPE 8TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty 512e Multi Vendor HDD	834031-B21
HPE 2TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	833926-B21

Core Options

HPE 4TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	833928-B21
HPE 6TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty 512e Multi Vendor HDD	861746-B21
HPE 16TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e ISE Multi Vendor HDD	P23608-B21
HPE 20TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e ISE Multi Vendor HDD	P53553-B21
Midline - 12G SAS - SFF Drives	
HPE 1TB SAS 12G Business Critical 7.2K SFF BC 1-year Warranty HDD	P53563-B21
Midline - 6G SATA - LFF Drives	
HPE 2TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	861681-B21
HPE 4TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	861683-B21
HPE 1TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	861686-B21
HPE 12TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e Multi Vendor HDD	881787-B21
HPE 8TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty 512e Multi Vendor HDD	834028-B21
HPE 6TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty 512e Multi Vendor HDD	861742-B21
HPE 14TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e Multi Vendor HDD	P09165-B21
HPE 16TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e ISE Multi Vendor HDD	P23449-B21
HPE 20TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e ISE Multi Vendor HDD	P53554-B21
Midline - 6G SATA - SFF Drives	
HPE 2TB SATA 6G Business Critical 7.2K SFF BC 1-year Warranty 512e HDD	P28500-B21
SED (Self-Encryption Drive)	
HPE 2.4TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty 512e Self-encrypting FIPS HDD	P28618-B21
HPE 1.2TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Self-encrypting FIPS HDD	P28622-B21
Optical Drive	
HPE 9.5mm SATA DVD-ROM Optical Drive	726536-B21
HPE 9.5mm SATA DVD-RW Optical Drive	726537-B21
HPE Mobile USB DVD-RW Optical Drive	701498-B21
HPE ProLiant DL325 Gen11 8SFF Display Port/USB/Optical Drive Blank Kit	P56654-B21
HPE ProLiant DL325 Gen11 4LFF Display Port/USB/Optical Drive Blank Kit	P56655-B21

Notes:

– If the 2SFF drive cage (P56652-B21) is selected then optical drives cannot be selected and vice versa.

– If the optical drive is selected along with the 8SFF CTO server (P54199-B21), then the 8SFF ODD blank kit (P56654-B21) must be selected.

– If the optical drive is selected along with the 4LFF CTO server (P54200-B21), then the 4LFF ODD blank kit
(P56655-B21) must be selected.

– Both 8SFF ODD blank kit (P56654-B21) and 4LFF ODD blank kit (P56655-B21) support one (1) Display Port and one (1) USB 2.0 port.

Boot Controllers

HPE NS204i-u Gen11 NVMe Hot Plug Boot Optimized Storage Device	P48183-B21

Core Options

Notes:
- RAID 1 is preconfigured on this option and additional RAID cannot be applied on this Boot Device
- Requires Performance Fan Kits (P58462-B21) or Liquid Cooling Fan Kits (P59668-B21)
- If this NS204i-u boot device is selected along with the SFF/LFF CTO servers and Liquid Cooling Fan Kits, then the 2SFF drive cage (P56652-B21) cannot be selected, and recommended system ambient temperature is 25C.
- Not allowed If this NS204i-u boot device is selected along with the EDSFF CTO servers and Liquid Cooling Fan Kit.
- If this NS204i-u boot device is selected then the Secondary Low Profile riser (P55029-B21) and NS204i-u Cable Kit

(P57013-B21) must be selected.
- For additional information, please visit **HPE OS Boot Device QuickSpecs**

HPE ProLiant DL3X5 Gen11 NS204i-u NVMe Hot Plug Boot Device Cable Kit	P57013-B21
HPE ProLiant DL325 Gen11 NVMe/SATA M.2 Enablement Kit	P57014-B21

Notes:
- Requires two (2) M.2 SSD Drives In the same interface (SATA or NVMe).
- No RAID is supported on this M.2 enablement kit.
- If this M.2 enablement kit is selected along with the SFF/LFF CTO servers and Liquid Cooling Fan Kit (P59668-B21), then the 2SFF drive cage (P56652-B21) cannot be selected and recommended system ambient temperature is 25C.
- Not allowed If this M.2 enablement kit is selected along with the EDSFF CTO servers and Liquid Cooling Fan Kit.

Read Intensive - 6G SATA - M.2 - Solid State Drives

HPE 480GB SATA 6G Read Intensive M.2 Multi Vendor SSD	P47818-B21
HPE 480GB NVMe Gen3 Mainstream Performance Read Intensive M.2 Multi Vendor SSD	P40513-B21
HPE 960GB NVMe Gen3 Mainstream Performance Read Intensive M.2 Multi Vendor SSD	P40514-B21
HPE 1.92TB NVMe Gen3 Mainstream Performance Read Intensive M.2 Multi Vendor SSD	P40515-B21
HPE 480GB NVMe Gen4 Mainstream Performance Read Intensive M.2 PM9A3 SSD	P69543-B21

Risers

Notes: The Primary riser shipping default in ALL CTO server is PCIe Gen5 x16 FH HL.

HPE ProLiant DL3X5 Gen11 1U x16 Low Profile Secondary Riser Kit	P55029-B21
HPE ProLiant DL3X5 Gen11 1U x16 Riser Kit	P56915-B21

Notes:
- Both riser kits are in the secondary slot.
- Requires Low Profile Secondary riser kit if NS204i-u (P48183-B21) is selected.

HPE ProLiant DL325 Gen11 FHFL Add-on Cards Support Kit	P64520-B21

Notes: this kit supports single-width FHFL add-on PCIe cards at the primary riser position

HPE Networking

Notes: Direct Attach Cable (DAC) for copper environments or fiber transceivers and cables for fiber-optic environments must be purchased separately. Please see the related NIC QuickSpecs for Technical Specifications and additional information: **https://h20195.www2.hpe.com/v2/getpdf.aspx/A00002507ENW.pdf**

PCIe Adapters

Core Options

1 Gigabit Ethernet adapter

Intel I350-T4 Ethernet 1Gb 4-port BASE-T Adapter for HPE	P21106-B21
Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T Adapter for HPE	P51178-B21

10 Gigabit Ethernet adapters

Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T Adapter for HPE	P26253-B21
Broadcom BCM57412 Ethernet 10Gb 2-port SFP+ Adapter for HPE	P26259-B21

10/25 Gigabit Ethernet adapters

Notes: Require Performance Fan Kits (P58462-B21) or Liquid Cooling fan kits (P59668-B21) and subject to the recommended system ambient temperature

Broadcom BCM57414 Ethernet 10/25Gb 2-port SFP28 Adapter for HPE	P26262-B21
Broadcom BCM57504 Ethernet 10/25Gb 4-port SFP28 Adapter for HPE	P26264-B21
Intel E810-XXVDA2 Ethernet 10/25Gb 2-port SFP28 Adapter for HPE	P08443-B21
Mellanox MCX631102AS-ADAT Ethernet 10/25Gb 2-port SFP28 Adapter for HPE	P42044-B21
Intel E810-XXVDA4 Ethernet 10/25Gb 4-port SFP28 Adapter for HPE	P08458-B21
HPE Ethernet 10/25Gb 2-port Secure Network Adapter	S2A69A

100/200 Gigabit Ethernet adapters

Notes: Require Performance Fan Kits (P58462-B21) or Liquid Cooling fan kits (P59668-B21) and subject to the recommended system ambient temperature

Mellanox MCX623106AS-CDAT Ethernet 100Gb 2-port QSFP56 Adapter for HPE	P25960-B21
Intel E810-CQDA2 Ethernet 100Gb 2-port QSFP28 Adapter for HPE	P21112-B21
Mellanox MCX623105AS-VDAT Ethernet 200Gb 1-port QSFP56 Adapter for HPE	P10180-B21
HPE NV60100M 100Gb 2-port Storage Offload Adapter	R8M41A
HPE Slingshot SA210S Ethernet 200Gb 1-port PCIe NIC	R4K46A

OCP 3.0 Adapter

1 Gigabit Ethernet OCP adapters

Intel I350-T4 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE	P08449-B21
Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE	P51181-B21

10 Gigabit Ethernet OCP Adapters

Broadcom BCM57412 Ethernet 10Gb 2-port SFP+ OCP3 Adapter for HPE	P26256-B21
Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T OCP3 Adapter for HPE	P10097-B21

10/25 Gigabit Ethernet OCP adapters

Notes: Require Performance Fan Kits (P58462-B21) or Liquid Cooling fan kits (P59668-B21) and subject to the recommended system ambient temperature

Broadcom BCM57414 Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE	P10115-B21
Broadcom BCM57504 Ethernet 10/25Gb 4-port SFP28 OCP3 Adapter for HPE	P26269-B21
Intel E810-XXVDA2 Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE	P10106-B21
Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE	P42041-B21

100/200 Gigabit Ethernet adapters

Intel E810-CQDA2 Ethernet 100Gb 2-port QSFP28 OCP3 Adapter for HPE	P22767-B21

Notes:

- Require Performance Fan Kits (P58462-B21) or Liquid Cooling fan kits (P59668-B21) and subject to the recommended system ambient temperature
- Requires OCP1 upgrade cable kit (P56658-B21) to support PCIe Gen5 x16 bandwidth on OCP21 slot

Recommended System Ambient Temperature

Core Options

P/N	SFF/LFF CTO servers				EDSFF CTO server			
	Perf Fans (P58462-B21)		LC Fans (P59668-B21)		Perf Fans (P58462-B21)		LC Fans (P59668-B21)	
	Pri. Riser	Sec. Riser	Pri. Riser	Sec. Riser	Pri. Riser	Sec. Riser	Pri. Riser	Sec. Riser
P08443-B21	30C	30C	30C	30C	30C	25C	25C	Not support
P26264-B21	30C	25C	30C	25C	30C	Not support	25C	Not support
P42044-B21	30C	25C	30C	25C	30C	Not support	25C	Not support
S2A69A	30C	25C	30C	25C	30C	Not support	25C	Not support
P08458-B21	30C	25C	30C	25C	30C	Not support	25C	Not support
P21112-B21	30C	25C	30C	25C	30C	Not support	25C	Not support
P10180-B21	30C	Not support	25C	Not support	30C	Not support	Not support	Not support
P25960-B21	30C	Not support	25C	Not support	30C	Not support	Not support	Not support
R8M41A	30C	Not support	25C	Not support	30C	Not support	Not support	Not support
P/N	**OCP21**	**OCP22**	**OCP21**	**OCP22**	**OCP21**	**OCP22**	**OCP21**	**OCP22**
P10106-B21	30C	25C	30C	25C	30C	25C	25C	Not support
P42041-B21	30C	25C	30C	25C	30C	Not support	Not support	Not support
P26269-B21	30C	Not support	30C	Not support	30C	Not support	Not support	Not support
P22767-B21	30C	Not support	30C	Not support	30C	Not support	Not support	Not support

P/N	GPU CTO server*			
	Perf Fans (P58462-B21)		LC Fans (P59668-B21)	
	Pri. Riser	Sec. Riser	Pri. Riser	Sec. Riser
P08443-B21	30C	25C	30C	25C
P26264-B21	30C	Not support	25C	Not support
P42044-B21	30C	Not support	25C	Not support
S2A69A	30C	Not support	25C	Not support
P08458-B21	30C	Not support	25C	Not support
P21112-B21	30C	Not support	25C	Not support
P10180-B21	30C	Not support	Not support	Not support
P25960-B21	30C	Not support	Not support	Not support
R8M41A	30C	Not support	Not support	Not support
P/N	**OCP21**	**OCP22**	**OCP21**	**OCP22**
P10106-B21	30C	25C	30C	25C
P42041-B21	30C	25C	30C	25C
P26269-B21	30C	Not support	25C	Not support
P22767-B21	30C	Not support	25C	Not support

Notes:

- Not support = configuration not allowed because of thermal limitation.
- The thermal condition of GPU CTO server is based on 2pcs 72W GPUs installed at the front cage.

HPE InfiniBand

Notes:

– Requires Performance Fan Kit (P58462-B21) or Liquid Cooling fan kits (P59668-B21) and subject to the recommended system ambient temperature

– Requires OCP upgrade cable kit (P56658-B21) for 200Gb OCP adapters (P31323-B21 or P31348-B21)

– For more information, please visit: **HPE InfiniBand Options for HPE ProLiant and Apollo Servers**

HPE InfiniBand HDR100/Ethernet 100Gb 1-port QSFP56 PCIe4 x16 MCX653105A-ECAT Adapter	P23665-B21
HPE InfiniBand HDR100/Ethernet 100Gb 2-port QSFP56 PCIe4 x16 MCX653106A-ECAT Adapter	P23666-B21
HPE InfiniBand HDR/Ethernet 200Gb 1-port QSFP56 PCIe4 x16 MCX653105A-HDAT Adapter	P23664-B21
HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 MCX653106A-HDAT Adapter	P31324-B21
HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 OCP3 MCX653436A-HDAI Adapter	P31348-B21
HPE InfiniBand NDR 1-port OSFP PCIe5 x16 MCX75310AAS-NEAT Adapter	P45641-B21
HPE InfiniBand NDR200/Ethernet 200GbE 2-port QSFP112 PCIe5 x16 MCX755106AC-HEAT Adapter	P65333-B21
HPE InfiniBand NDR200/Ethernet 200Gb 1-port OSFP PCIe5 x16 MCX75310AAS-HEAT Adapter	P45642-B22

Recommended System Ambient Temperature

P/N	SFF/LFF CTO servers				EDSFF CTO server			
	Perf Fans (P58462-B21)		LC Fans (P59668-B21)		Perf Fans (P58462-B21)		LC Fans (P59668-B21)	
	Pri. Riser	Sec. Riser	Pri. Riser	Sec. Riser	Pri. Riser	Sec. Riser	Pri. Riser	Sec. Riser
P23665-B21	30C	30C	30C	30C	30C	25C	25C	Not support
P23664-B21	30C	25C	30C	25C	30C	Not support	25C	Not support
P23666-B21	30C	25C	30C	25C	30C	Not support	25C	Not support
P31324-B21	30C	Not support	25C	Not support	30C	Not support	Not support	Not support
P45641-B21/23	30C	Not support	25C	Not support	30C	Not support	Not support	Not support
P45642-B22	30C	Not support	25C	Not support	30C	Not support	Not support	Not support
P65333-B21	30C	Not support	25C	Not support	30C	Not support	Not support	Not support
P/N	**OCP21**	**OCP22**	**OCP21**	**OCP22**	**OCP21**	**OCP22**	**OCP21**	**OCP22**
P31323-B21	30C	Not support	25C	Not support	25C	Not support	Not support	Not support
P31348-B21	30C	Not support	Not support	Not support	Not support	Not support	Not support	Not support

Core Options

P/N	GPU CTO server*			
	Perf Fans(P58462-B21)		LC Fans(P59668-B21)	
	Pri. Riser	Sec. Riser	Pri. Riser	Sec. Riser
P23665-B21	30C	25C	30C	25C
P23664-B21	30C	Not support	25C	Not support
P23666-B21	30C	Not support	25C	Not support
P31324-B21	30C	Not support	Not support	Not support
P45641-B21/23	30C	Not support	Not support	Not support
P45642-B22	30C	Not support	Not support	Not support
P65333-B21	30C	Not support	Not support	Not support
P/N	**OCP21**	**OCP22**	**OCP21**	**OCP22**
P31323-B21	25C	Not support	Not support	Not support
P31348-B21	25C	Not support	Not support	Not support

Notes:

– Not support = configuration not allowed because of thermal limitation.

– The thermal condition of GPU CTO server is based on 2pcs 72W GPUs installed at the front cage.

Accelerators

NVIDIA A2 16GB PCIe Non-CEC Accelerator for HPE R9H23C

Notes:

– This is a PCIe Gen4 x 8 single-width HHHL GPU card.

– Max = 2 at the rear.

– This GPU can only be selected with 8SFF/4LFF/EDSFF CTO Server.

– If this GPU is installed in PCIe Slot2 with Performance Fan kits (P58462-B21), the recommended system ambient temperature is 25C.

– If this GPU is installed on either PCIe Slot1 or Slot2 with Liquid Cooling Fan kits (P59668-B21), the recommended system ambient temperature is 25C.

NVIDIA L4 24GB PCIe Accelerator for HPE S0K89C

Notes:

– This is a PCIe Gen4 x 16 single-width HHHL GPU card.

– Max = 4, 2 at the front and 2 at the rear.

– This GPU can only be selected with GPU CTO Server.

– If this GPU is installed in PCIe Slot2 with Performance Fan kits (P58462-B21), the recommended system ambient temperature is 25C.

– This GPU cannot be selected with Liquid Cooling Fan kits (P59668-B21) on either PCIe Slot1 or Slot2 due to thermal limitation.

HPE Storage Options

Emulex Fibre Channel HBAs

HPE SN1610E 32Gb 1-port Fibre Channel Host Bus Adapter R2J62A

HPE SN1610E 32Gb 2-port Fibre Channel Host Bus Adapter R2J63A

HPE SN1700E 64Gb 1-port Fibre Channel Host Bus Adapter R7N77A

HPE SN1700E 64Gb 2-port Fibre Channel Host Bus Adapter R7N78A

QLogic Fibre Channel HBAs

HPE SN1610Q 32Gb 1-port Fibre Channel Host Bus Adapter R2E08A

HPE SN1610Q 32Gb 2-port Fibre Channel Host Bus Adapter R2E09A

HPE SN1700Q 64Gb 1-port Fibre Channel Host Bus Adapter R7N86A

HPE SN1700Q 64Gb 2-port Fibre Channel Host Bus Adapter R7N87A

Core Options

Power and Cooling

Cooling

Notes: Requires one (1) Heat Sink and Seven (7) Fan Kit in the order.

HPE ProLiant DL3X5 Gen11 1U CPU Standard Heat Sink Kit P58456-B21

Notes: Required for processors less than or equal to 240W

HPE ProLiant DL3X5 Gen11 1U CPU Performance Heat Sink Kit P58457-B21

Notes: Required for processors more than or equal to 260W and less than or equal to 300W

HPE ProLiant DL325 Gen11 Closed-loop Liquid Cooling FIO Heat Sink Kit P58463-B21

Notes:

- This Closed-loop liquid cooling Heat Sink FIO kit is designed for processors higher than or equal to 320W.
- Requires Liquid Cooling Fan Kits (P59668-B21).
- The HPE DL325 Gen11 Closed-Loop Liquid Cooling Heat Sink FIO kit is subject to a Maximum Usage Limitation of not exceeding five (5) years of operation and is required to be replaced when reaching limitation. Parts and components that Hewlett Packard Enterprise determines have reached or exceeded their Maximum Usage limitations will not be provided, repaired, or replaced under warranty or service contract. Contact your local sales representative for additional information.
- For more information see our **HPE ProLiant Gen11 Closed-Loop Liquid Cooling Heat Sink FAQs**

HPE ProLiant DL3XX Gen11 1U Standard Fan Kit P58461-B21

HPE ProLiant DL3XX Gen11 1U Performance Fan Kit P58462-B21

HPE ProLiant DL325 Gen11 Liquid Cooling Fan Kit P59668-B21

CTO	Drive Cage	CPU	Heat Sink	Fan	Sys Temp	96/128G DIMM	256G DIMM	NS204i-u/ M.2 Kit
SFF	8SFF x1	<=240W	Standard	Standard*	30C	Not Support	Not Support	Not Support
	10SFF x1	<=240W	Standard	Performance	30C	30C	25C	30C
	8SFF x4	<=300W	Performance	Performance	30C	30C	25C	25C
	10SFF x4	<=300W	Performance	Performance	30C	30C	25C	Not Support
	8SFF x4	>300W	Liquid Cool	Liquid Cool	30C	25C	25C Max=4	25C
	10SFF x4	>300W	Liquid Cool	Liquid Cool	30C	25C	25C Max=4	Not Support
LFF	4LFF x1	<=240W	Standard	Standard*	28C	Not Support	Not Support	30C
	4LFF x1	<=240W	Standard	Performance	30C	30C	25C	30C
	4LFF x1	<=300W	Performance	Performance	30C	30C	25C	25C
	4LFF x1	>300W	Liquid Cool	Liquid Cool	30C	25C	25C Max =4	25C
EDSFF	20EDSFF	<=300W	Performance	Performance	25C	25C	25C	25C
	20EDSFF	>300W	Liquid Cool	Liquid Cool	25C	25C Max =8	Not Support	Not Support
GPU	4SFF or 8EDSFF	<=240W	Standard	Performance	30C	30C	25C	30C
		<=300W	Performance	Performance	30C	30C	25C	30C
		>300W	Liquid Cool	Liquid Cool	25C	25C	25C Max =4	Not Support

Core Options

Notes:

- Require Performance Fan with <=240W CPU if any of the below options are selected with 8SFF/4LFF CTO server
  - 8SFF x4 U.3 backplane kit (P55000-B21)
  - 2SFF x4 U.3 backplane kit (P56652-B21)
  - 96GB, 128GB, or 256GB DIMM
  - NS204i-u (P48183-B21) or M.2 enablement kit (P57014-B21)
  - Networking options: 10/25G, 100/200G, and InfiniBand options.
  - Graphic options
- Liquid cooling fan (P59668-B21) can only be selected with liquid cooling heat sink (P58463-B21)

Power Supplies

HPE Flexible Slot (Flex Slot) Power Supplies share a common electrical and physical design that allows for hot plug, and tool-less installation into HPE ProLiant Servers. Flex Slot power supplies are certified for high-efficiency operation and offer multiple power output options, allowing users to "right-size" a power supply for specific server configurations. This flexibility helps to reduce power waste, lower overall energy costs, and avoid "trapped" power capacity in the data center.

Notes:

- Select a minimum (1), maximum (2) power supplies
- All power supplies in a server should match. Mixing Power Supplies is not supported.
- 1600W Power supplies only support high line voltage (200VAC to 240VAC).
- Before making a power supply selection it is highly recommended that the HPE Power Advisor is run to determine the right size power supply for your server configuration. The HPE Power Advisor is located at: **https://poweradvisorext.it.hpe.com/?Page=Index**
- HPE ProLiant servers ship with an IEC-IEC power cord used for rack mounting with Power Distribution Units (PDUs). Visit **HPE Power Cords and Cables** for a full list of optional power cords

HPE 500W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	865408-B21
HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P38995-B21
HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit	P03178-B21
HPE 1600W Flex Slot -48VDC Hot Plug Power Supply Kit	P17023-B21
HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P38997-B21
HPE 1800W-2200W Flex Slot Titanium Hot Plug Power Supply Kit	P44712-B21
HPE 1600W -48VDC Power Cable Lug Kit	P36877-B21

Notes: Must be selected along with HPE 1600W Flex Slot -48VDC Hot Plug Power Supply Kit (P17023-B21)

HPE Security

HPE Trusted Supply Chain for HPE ProLiant	P36394-B21

Notes:

- HPE Trusted Supply Chain is an optional security upgrade intended for agencies and regulated industries needing enhanced security and compliance needs. Applying this option to a DL325 Gen11 CTO server ensures it is built in the USA in a secured facility by vetted HPE personnel assigned to

Core Options

the manufacturing processes. A multitude of checkpoints/inspections for malicious microcode and counterfeit parts are performed throughout the server build, and additional safeguards are put in place against cyber-exploits throughout the server lifecycle. Learn more at **http://www.hpe.com/security**

- This option requires the selection of HPE Gen11 Intrusion Detection Kit (P48922-B21)
- This option requires the selection of either HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features (BD505A) or HPE iLO Advanced 1-server License with 1yr Support on iLO Licensed Features (512485-B21)
- This option is limited to stand-alone DL325Gen11 CTO servers only. The HPE Trusted Supply Chain configuration will not be available if the server is ordered as factory integrated into a rack
- One instance of the following Electronic License to Use is required per order (not per server): R6X85AAE (HPE Trusted Supply Chain E-LTU)
- This option cannot be selected with TAA instruction SKU or TAA CTO Models.

HPE ProLiant DL3XX Gen11 Intrusion Cable Kit P48922-B21

Notes: This provides a physical connection from the chassis board and hood and detects any physical intrusion into the chassis, providing security during the entire supply chain process of shipping, receiving distribution, and operation.

HPE ProLiant Gen11 1U Common Bezel Kit P50450-B21

HPE Bezel Lock Kit 875519-B21

Notes: The Bezel lock kit (875519-B21) must be selected with the bezel kit (P50450-B21)

HPE iLO Common Password FIO Setting P08040-B21

Notes:

- Replaces iLO default randomized password by an HPE defined common password. HPE highly recommends changing this password immediately after the initial onboarding process.
- Customers who want to choose their own custom iLO default password should use the HPE Factory Express Integration Services

---

Additional Cable Options

HPE ProLiant DL3X5 Gen11 OCP1 Upgrade Cable Kit P56658-B21

Notes: Supports PCIe x16 bandwidth at OCP slot 21. Required if one of the following options is in the order

- OCP InfiniBand network adapters (P31323-B21, P31348-B21)
- BCM 57504 10/25GbE 4p SFP28 Adaptor (P26269-B21)
- Intel E810 100GbE 2p QSFP28 OCP3 Adptr (P22767-B21)

HPE ProLiant DL3X5 Serial Port Enablement Kit P50887-B21

Notes: This cable kit supports an optional serial port at the rear of the server.

Software as a Service Management

HPE Compute Ops Management

HPE Compute Ops Management Enhanced 3-year Upfront ProLiant SaaS	R7A11AAE
HPE Compute Ops Management Enhanced 5-year Upfront ProLiant SaaS	R7A12AAE
HPE Compute Cloud Management Server FIO Enablement	S1A05A

HPE OneView

HPE OneView w/o iLO including 3yr 24x7 Support Track 1-server LTU	P8B25A
HPE OneView w/o iLO including 3yr 24x7 Support Flexible Quantity E-LTU	P8B26AAE
HPE OneView w/o iLO including 3yr 24x7 Support 1-server FIO LTU	P8B31A
HPE OneView including 3yr 24x7 Support Flexible Quantity E-LTU	E5Y35AAE
HPE OneView including 3yr 24x7 Support Track 1-server LTU	E5Y36A

Core Options

HPE OneView for ProLiant DL Server including 3yr 24x7 Support FIO Bundle Physical 1-server LTU	E5Y43A
HPE OneView for ProLiant DL Server including 3yr 24x7 Support Bundle Track 1-server LTU	E5Y44A

For more information, visit the HPE Compute Ops Management QuickSpecs:
https://www.hpe.com/psnow/doc/a50004263enw
Supported Servers - CTO only. No OEM. - Complete list can be found here: Latest Supported Server List: **https://www.hpe.com/info/com-supported-servers**

Additional Options

Some options may not be integrated at the factory. To ensure only valid configurations are ordered, Hewlett Packard Enterprise recommends the use of an HPE approved configurator. Contact your local sales representative for additional information.

Embedded Management

HPE iLO Advanced

HPE iLO Advanced Electronic License with 1yr Support on iLO Licensed Features	E6U59ABE
HPE iLO Advanced Electronic License with 3yr Support on iLO Licensed Features	E6U64ABE
HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features	BD505A
HPE iLO Advanced AKA Tracking License with 3yr Support on iLO Licensed Features	BD507A
HPE iLO Advanced 1-server License with 1yr Support on iLO Licensed Features	512485-B21
HPE iLO Advanced AKA Tracking License with 1yr Support on iLO Licensed Features	512487-B21

Rail Kits

Easy Install rail kits contain telescoping rails which allow for in-rack serviceability.
To assist in the installation of the server into the rack an optional installation tool is available by contacting your local services representative.

Notes:

- HPE rail kits are designed to work with HPE racks in compliance with industry standard EIA-310-E. In the event a customer elects to purchase a third-party rack for use with an HPE rail kit, any such use is at customer's own risk. HPE makes no express or implied warranties with respect to such third-party racks and specifically disclaims any implied warranties of merchantability and fitness for a particular purpose. Furthermore, HPE has no obligation and assumes no liability for the materials, design, specifications, installation, safety, and compatibility of any such third-party racks with any rail kits, including HPE rail kits.
- Hewlett Packard Enterprise recommends that a minimum of two people are required for all Rack Server installations. Please refer to your installation instructions for proper tools and the number of people to use for any installation.

HPE ProLiant DL3XX Gen11 Easy Install Rail 3 Kit	P52341-B21

Notes: This Rail kit can be selected only with the 4LFF/EDSFF CTO server.

HPE DL3XX Gen11 Easy Install Rail 2 Kit	P52351-B21

Notes: This Rail kit can be selected only with the 8SFF CTO server.

HPE Easy Install Rail 7 Kit	P52339-B21

Notes: This Rail kit can be selected only with the GPU CTO server.

HPE ProLiant DL300 Gen10 Plus 1U Cable Management Arm for Rail Kit	P26489-B21

Notes: CMA can be selected only with the Rail kit.

HPE Tape Backup

For the complete range of tape drives, autoloaders, libraries and media see:
https://www.hpe.com/us/en/storage/storeever-tape-storage.html
For hardware and software compatibility of Hewlett Packard Enterprise tape backup products please visit the StoreEver Tape Solutions in SPOCK (requires registration/login)
https://h20272.www2.hpe.com/SPOCK/default.aspx
Only external drives supported
All libraries and autoloaders supported via compatible FC or SAS controller. Refer to the StoreEver Tape Soutions Compatibilty Matrix link above.

Additional Options

HPE Racks

- Please see the HPE Advanced Series Racks QuickSpecs for information on additional racks options and rack specifications. **HPE G2 Advanced Series Racks**
- Please see the HPE Enterprise Series Racks QuickSpecs for information on additional racks options and rack specifications. **HPE G2 Enterprise Series Racks**

HPE Power Distribution Units (PDUs)

- Please see the **HPE Basic Power Distribution Units (PDU) QuickSpecs** for information on these products and their specifications.
- Please see the **HPE Metered Power Distribution Units (PDU) QuickSpecs** for information on these products and their specifications.
- Please see the **HPE Intelligent Power Distribution Unit (PDU) QuickSpecs** for information on these products and their specifications.
- Please see the **HPE Metered and Switched Power Distribution Units (PDU) QuickSpecs** for information on these products and their specifications.

HPE Uninterruptible Power Systems (UPS)

- To learn more, please visit the **HPE Uninterruptible Power Systems (UPS) web page.**
- Please see the **HPE DirectFlow Three Phase Uninterruptible Power System QuickSpecs** for information on these products and their specifications.
- Please see the **HPE Line Interactive Single Phase UPS QuickSpecs** for information on these products and their specifications.

HPE Rack Options

- Please see the **HPE KVM Switches web page** for information on these products and their specifications

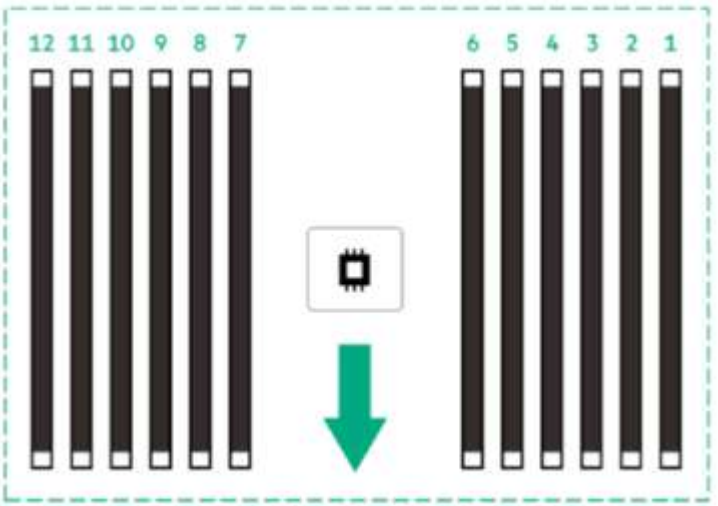
---

HPE Support Service

Tech Care

HPE 3 Year Tech Care Essential DL325 Gen11 Service	H78S6E
HPE 3 Year Tech Care Essential wDMR DL325 Gen11 Service	H78S7E
HPE 5 Year Tech Care Essential DL325 Gen11 Service	H78V0E
HPE 5 Year Tech Care Essential wDMR DL325 Gen11 Service	H78V1E

Memory

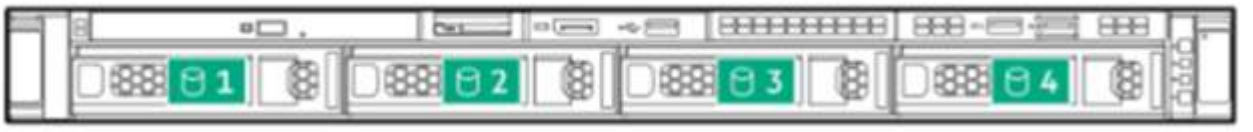


The arrow points to the front of the server

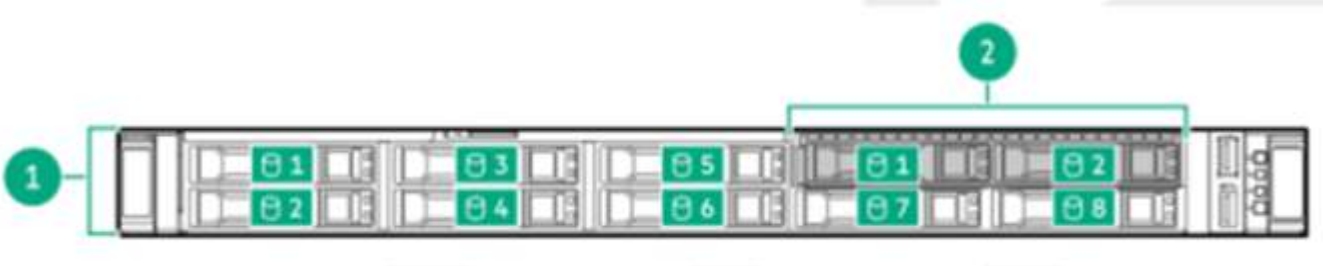
General Memory Population Rules and Guidelines:

- Install DIMMs only after the corresponding processor is installed.
- To maximize performance, it is recommended to balance the total memory capacity between all installed processors.
- The maximum memory speed is a function of the memory type, memory configuration, and processor model.
- The maximum memory capacity is a function of the number of DIMM slots on the platform, the largest DIMM capacity qualified on the platform, and the number and model of installed processors qualified on the platform.
- To realize the performance memory capabilities listed in this document, HPE DDR5 Smart Memory is required. For additional information, please see the: **HPE DDR5 Smart Memory QuickSpecs**
- For details on the Memory Population Rules and Guidelines with AMD EPYC 9004 and 9005 series processors, please go to: https://www.hpe.com/psnow/doc/a50007481enw

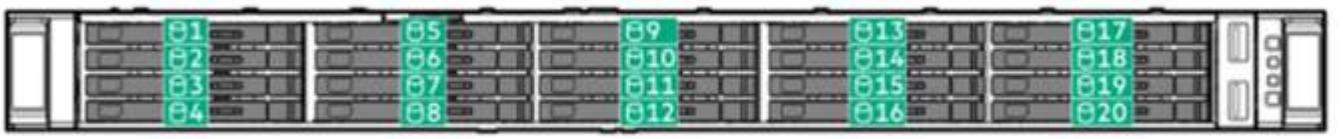
Storage



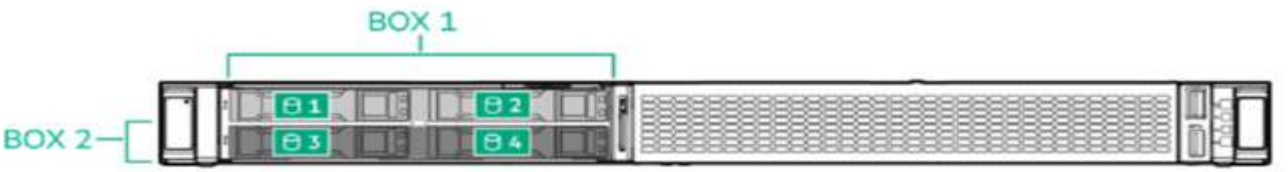
4LFF drives



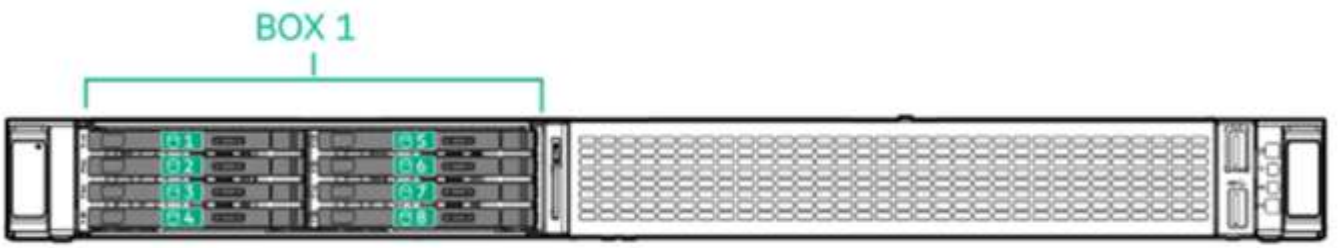
8SFF + Optional 2SFF (SAS/SATA/ NVMe)



20 EDSFF E3.S 1T Drives



4 SFF Drives in GPU CTO server



8 EDSFF Drives in GPU CTO server

System Unit

Dimensions (Height x Width x Depth)

- 8SFF chassis:
  - 4.29 X 43.46 X 64.94 cm
  - 1.69 X 17.11 X 25.57 In
- 4LFF & EDSFF chassis:
  - 4.29 X 43.46 X 70.89 cm
  - 1.69 X 17.11 X 27.91 In
- GPU Chassis
  - 4.29 X 43.46 X 81.84 cm
  - 1.69 X 17.11 X 32.22 In
- Package
  - 24.2 X 60 X 91.6 cm
  - 9.53 X 23.6 X 36.06 In

Weight (approximate)

- 8SFF chassis:
  - Minimum: 8 SFF chassis with 0 drives, 1 processor, 1 power supply, 1 standard heatsink, 1 DIMM, 1 Smart Array controller, and 7 standard fans.
    - 12.56 kg
    - 27.69 lb
  - Maximum: 8 SFF chassis with 8 drives, 1 processor, 2 power supply, 1 standard heatsink, 12 DIMM, 1 Smart Array controller, and 7 standard fans.
    - 15.54 kg
    - 34.27 lb
  - Package
    - 4.21 kg
    - 9.281 lb
- 4LFF chassis:
  - Minimum: 4 LFF chassis with 0 drives, 1 processor, 1 power supply, 1 performance heatsink, 1 DIMM, 1 Smart Array controller, and 7 performance fans.
    - 14.31 kg
    - 31.54 lb
  - Maximum: 4 LFF chassis with 4 drives, 1 processor, 1 power supply, 1 performance heatsink, 1 DIMM, 1 Smart Array controller, and 7 performance fans.
    - 17.07 kg
    - 37.63 lb
  - Package
    - 4.145 kg
    - 9.138 lb
- EDSFF chassis:
  - Minimum: EDSFF chassis with 1 drive, 1 processor, 1 power supply, 1 standard heatsink, 2 DIMM, and 7 performance fans.
    - 13.71 kg
    - 30.23 lb
  - Maximum: EDSFF chassis with 20 drives, 1 processor, 1 power supply, 1 performance heatsink, 12 DIMM, and 7 performance fans.
    - 17.76 kg
    - 39.15 lb

Technical Specifications

- GPU Chassis
  - Minimum: GPU chassis with 2 EDSFF drives, 1 double-width accelerators, 1 processor, 1 power supply, 1 standard heatsink, 2 DIMM, and 7 performance fans.
    - 16.59 kg
    - 36.58 lb
  - Maximum GPU chassis with 8 EDSFF drives, 2 double-width accelerators, 1 processor, 1 power supply, 1 performance heatsink, 12 DIMM, and 7 performance fans.
    - 21.05 kg
    - 46.41. lb

Input Requirements(per power supply)

Rated Line Voltage

- 100 to 120 VAC
- 200 to 240 VAC

BTU Rating

Maximum

- For 1600W Power Supply: 5918 BTU/hr (at 200 VAC), 5884 BTU/hr (at 240 VAC) for China
- For 800W Power Supply: 3207 BTU/hr (at 100 VAC), 3071 BTU/hr (at 200 VAC), 3112 BTU/hr (at 240 VAC) for China Only
- For 500W Power Supply: 1979 BTU/hr (at 100 VAC), 1911 BTU/hr (at 200 VAC), 1965 BTU/hr (at 240 VAC) for China Only

---

Power Supply Output(per power supply)

- **Rated Steady-State Power**
  - For 1600W Power Supply: 1600W (at 240 VAC), 1600W (at 240 VAC)
  - For 800W Power Supply: 800W (at 100 VAC), 800W (at 240 VAC), 800W (at 240 VAC) input for China only
  - For 500W Power Supply: 500W (at 100 VAC), 500W (at 240 VAC), 500W (at 240 VAC) input for China only

- **Maximum Peak Power**
  - For 1600W Power Supply: 1600W (at 200 to 240 1VAC), 1600W (at 240 VAC) input for China only
  - For 800W Power Supply: 800W (at 100 to 127 VAC), 800W (at 200 to 240 1VAC), 800W (at 240 VAC) input for China only
  - For 500W Power Supply: 500W (at 100 to 127 VAC), 500W (at 200 to 240 VAC), and 500W (at 240 VAC) input for China only

---

System Inlet Temperature

- **Standard Operating Temperature**

10° to 35°C (50° to 95°F) at sea level with an altitude derating of 1.0°C per every 305 m (1.8°F per every 1000 ft) above sea level to a maximum of 3050 m (10,000 ft), no direct sustained sunlight. The maximum rate of change is 20°C/hr (36°F/hr). The upper limit and rate of change may be limited by the type and number of options installed.

System performance during standard operating support may be reduced if operating with a fan fault or above 30°C (86°F).

Technical Specifications

- **Extended Ambient Operating Temperature**

For approved hardware configurations, the supported system inlet range is extended to be: 5° to 10°C (41° to 50°F) and 35° to 40°C (95° to 104°F) at sea level with an altitude derating of 1.0°C per every 175 m (1.8°F per every 574 ft) above 900 m (2953 ft) to a maximum of 3050 m (10,000 ft). The approved hardware configurations for this system are listed at the URL: **http://www.hpe.com/servers/ashrae**

For approved hardware configurations, the supported system inlet range is extended to be: 40° to 45°C (104° to 113°F) at sea level with an altitude derating of 1.0°C per every 125 m (1.8°F per every 410 ft) above 900 m (2953 ft) to a maximum of 3050 m (10,000 ft). The approved hardware configurations for this system are listed at the URL:
http://www.hpe.com/servers/ashrae
System performance may be reduced if operating in the extended ambient operating range or with a fan fault.

- **Non-operating**

-30° to 60°C (-22° to 140°F). The maximum rate of change is 20°C/hr (36°F/hr).

Relative Humidity(non-condensing)

- **Operating**

8% to 90% - Relative humidity (Rh), 28°C maximum wet bulb temperature, non-condensing.

- **Non-operating**

5 to 95% relative humidity (Rh), 38.7°C (101.7°F) maximum wet bulb temperature, non-condensing.

Altitude

- **Operating**

3050 m (10,000 ft). This value may be limited by the type and number of options installed. The maximum allowable altitude change rate is 457 m/min (1500 ft/min).

- **Non-operating**

9144 m (30,000 ft). The maximum allowable altitude change rate is 457 m/min (1500 ft/min).

Emissions Classification (EMC) - Regulatory Information

To view the regulatory information for your product, view the Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products, available at the Hewlett Packard Enterprise Support Center:

https://support.hpe.com/hpesc/public/docDisplay?docLocale=en_US&docId=c03471072

Acoustic Noise

Listed are the declared mean A-Weighted sound power levels (LwAm), declared average bystander position A-Weighted sound pressure levels (LpAm), and the statistical adder for verification, Kv, is a quantity to be added to the declared mean A-weighted sound power level, LwA,m when the product is operating in a 23°C ambient environment. Noise emissions were measured under ISO 7779 (ECMA 74) and declared under ISO 9296 (ECMA 109). The listed sound levels apply to standard shipping configurations. Additional options may result in increased sound levels. Please have your HPE representative provide information from the HPE EMESC website for further technical details regarding the configurations listed below.

Technical Specifications

Idle	
LWA,m	5.1 B Perf 4.7 B Value
LpAm	37 dBA Perf 35 dBA Value
Kv	0.4 B Perf 0.4 B Value
Operating	
LWA,m	5.9 B Perf 5.7 B Value
LpAm	47 dBA Perf 42 dBA Value
Kv	0.4 B Perf 0.4 B Value

Notes:

- The declared mean A-weighted sound power level, LWA,m, is computed as the arithmetic average of the measured.
- A-weighted sound power levels for a randomly selected sample, rounded to the nearest 0,1 B.
- The declared mean A-weighted emission sound pressure level, LpA,m, is computed as the arithmetic average of the measured A-weighted emission sound pressure levels at the bystander positions for a randomly selected sample, rounded to the nearest 1 dB.
- The statistical adder for verification, Kv, is a quantity to be added to the declared mean A-weighted sound power level, LWA,m, such that there will be a 95 % probability of acceptance when using the verification procedures of ISO 9296, if no more than 6,5 % of the batch of new equipment, has A-weighted sound power levels greater than (LWA,m + Kv).
- The quantity, LWA,c (formerly called LWAd), can be computed from the sum of LWA,m, and Kv.
- All measurements made to conform to ISO 7779 / ECMA-74 and declared to conform to ISO 9296 / ECMA-109.
- B, dB, abbreviations for bels and decibels, respectively, where 1 B = 10 dB.
- The results in this declaration apply only to the model numbers listed above when operating and tested according to the indicated modes and standards. A system with additional configuration components or increased operating functionality may increase the noise emission values.
- System under abnormal conditions may increase the noise level, persons in the vicinity of the product [cabinet] for extended periods should consider wearing hearing protection or using other means to reduce noise exposure.

Environment-friendly Products and Approach - End-of-life Management and Recycling

Hewlett Packard Enterprise offers **end-of-life product return, trade-in, and recycling programs**, in many geographic areas, for our products. Products returned to Hewlett Packard Enterprise will be recycled, recovered, or disposed of in a responsible manner.

The EU WEEE directive (2002/95/EC) requires manufacturers to provide treatment information for each product type for use by treatment facilities. This information (product disassembly instructions) is posted on the Hewlett Packard Enterprise website. These instructions may be used by recyclers and other WEEE treatment facilities as well as Hewlett Packard Enterprise OEM customers who integrate and re-sell Hewlett Packard Enterprise equipment.

Summary of Changes

Date	Version History	Action	Description of Change
06-Jan-2025	Version 26	Changed	Core Optiosn section was updated.
02-Dec-2024	Version 25	Changed	Core Options and Additional Options sections were updated. (OBS SKUs were removed).
04-Nov-2024	Version 24	Changed	Standard Features and Core Options sections were updated.
10-Oct-2024	Version 23	Changed	Overview, Standard Features and Core Options sections were updated.
26-Sep-2024	Version 22	Changed	**Standard Features (Operating Systems and Virtualization Software Support for HPE Servers)**
05-Aug-2024	Version 21	Changed	Configuration Information (TPM China) and Core Options sections were updated.
15-Jul-2024	Version 20	Changed	Pre-Configured Models section was updated.
01-Jul-2024	Version 19	Changed	Core Options section was updated.
03-Jun-2024	Version 18	Changed	Pre-Configured Models and Core Options sections were updated.
20-May-2024	Version 17	Changed	Configuration Information and Core Options sections were updated.
15-Apr-2024	Version 16	Changed	Pre-Configured Models section was updated.
01-Apr-2024	Version 15	Changed	Core Options section was updated.
04-Mar-2024	Version 14	Changed	Overview, Standard Features, Pre-Configured Models, Configuration Information, Core Options and Additional Options sections were updated.
04-Dec-2023	Version 13	Changed	Service and Support and Core Options sections were updated.
02-Oct-2023	Version 12	Changed	Overview, Standard Features, Pre-Configured Models and Core Options sections were updated.
05-Sep-2023	Version 11	Changed	Overview, Standard Features, Pre-configured Models, Configuration Information, and Core Options sections were updated.
07-Aug-2023	Version 10	Changed	Overview, Standard Features, Pre-Configured Models, Configuration Information, Core Options, Additional Options, Storage, and Technical Specifications sections were updated.
10-Jul-2023	Version 9	Changed	Overview, Standard Features, Service and Support, Pre-Configured Models, Configuration Information, Core Options and Memory sections were updated.
13-Jun-2023	Version 8	Changed	Overview, Standard Features, Service and Support, Pre-Configured Models and Core Options sections were updated.
01-May-2023	Version 7	Changed	Standard Features and Core Options sections were updated
03-Apr-2023	Version 6	Changed	Overview, Standard Features, Configuration Information, Core Options, Additional Options, Memory, Storage and Technical Specifications sections were updated.
06-Mar-2023	Version 5	Changed	Overview, Standard Features, Configuration Information, additional Options and Technical Specifications sections were updated.
06-Feb-2023	Version 4	Changed	Overview, Standard Features, Configuration Information, additional Options and Technical Specifications sections were updated.
19-Dec-2022	Version 3	Changed	Overview and Standard Features sections were updated.
05-Dec-2022	Version 2	Changed	All sections were updated.
10-Nov-2022	Version 1	New	New QuickSpecs.

Copyright

Make the right purchase decision. Contact our presales specialists.



© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

AMD® and EPYC® are registered trademarks of Advanced Micro Devices Corporation in the U.S. and other countries.

Microsoft®, Windows®, and Windows Server® are U.S. registered trademarks of the Microsoft group of companies.

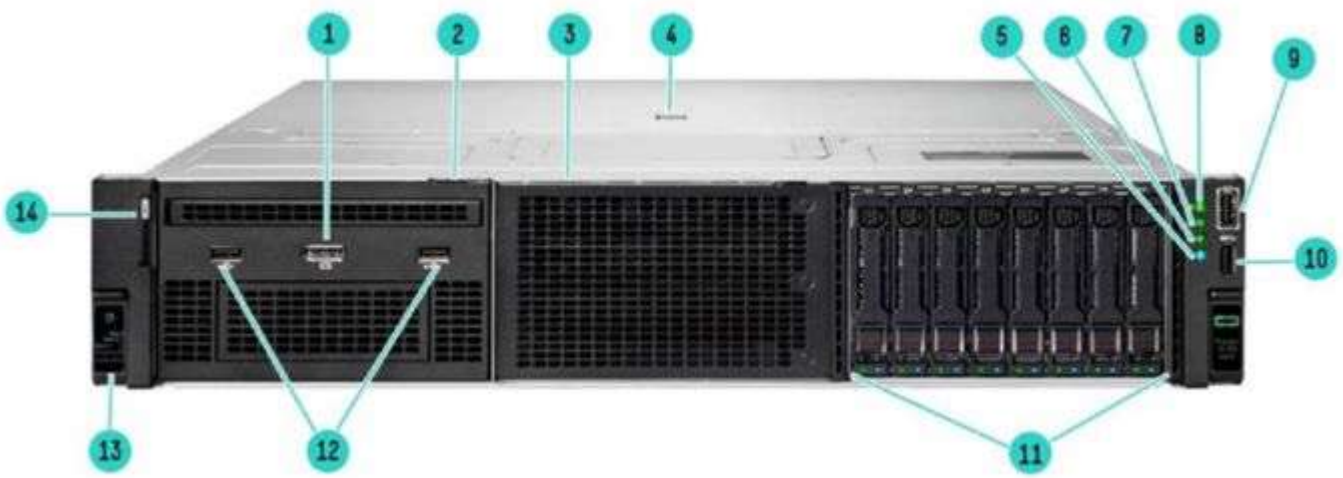
Hewlett Packard Enterprise

For hard drives, 1GB = 1 billion bytes. Actual formatted capacity is less

a50004297enw - 16901 - Worldwide - V26 - 06-January-2025

HPE ProLiant DL380 Gen11

Adaptable for diverse workloads and environments, the secure 2P 2U HPE ProLiant DL380 Gen11 Servers delivers world-class performance with the right balance of expandability and scalability. Designed for supreme versatility and resiliency while being backed by a comprehensive warranty make it ideal for multiple environments from Containers to Cloud to Big Data. Standardize on the industry's most trusted compute platform.

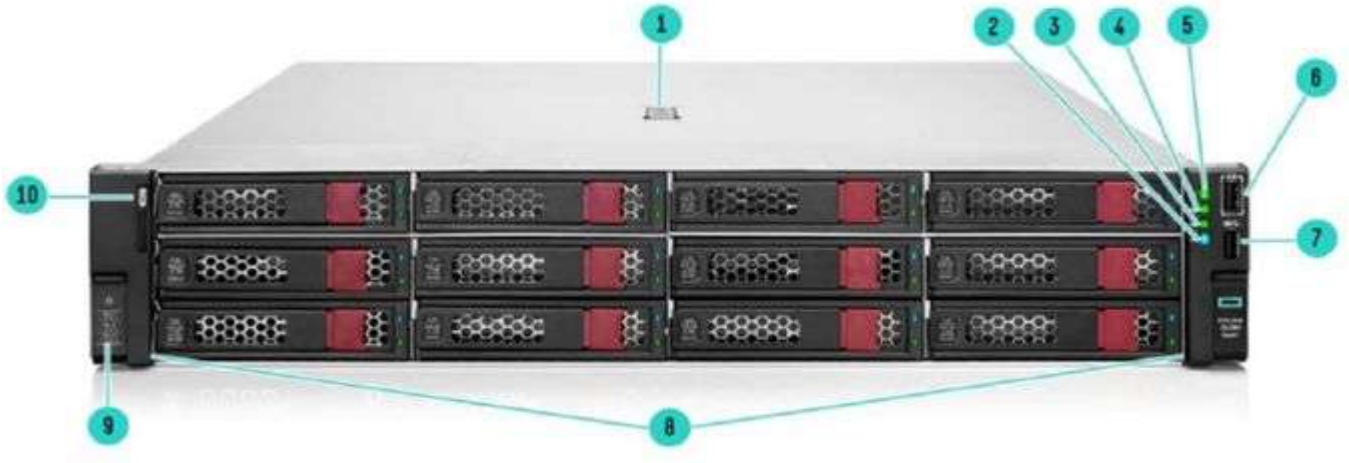


Front View - 8SFF chassis with optional Universal Media Bay shown

1.	Optional Front Display Port (via Universal Media Bay)	8.	Power On / Standby button / LED
2.	Box 1 (shown with optional Universal Media Bay installed)	9.	iLO Service Port
3.	Box 2 (Empty)	10.	USB 3.2 Gen1
4.	Quick removal access panel	11.	Box 3 (shown with 8SFF drives populated)
5.	Unit Identification button/LED	12.	Optional USB 2.0 port (via Universal Media Bay)
6.	NIC Status[1]	13.	Drive Support Label
7.	Health LED	14.	Serial Number Label Pull Tab

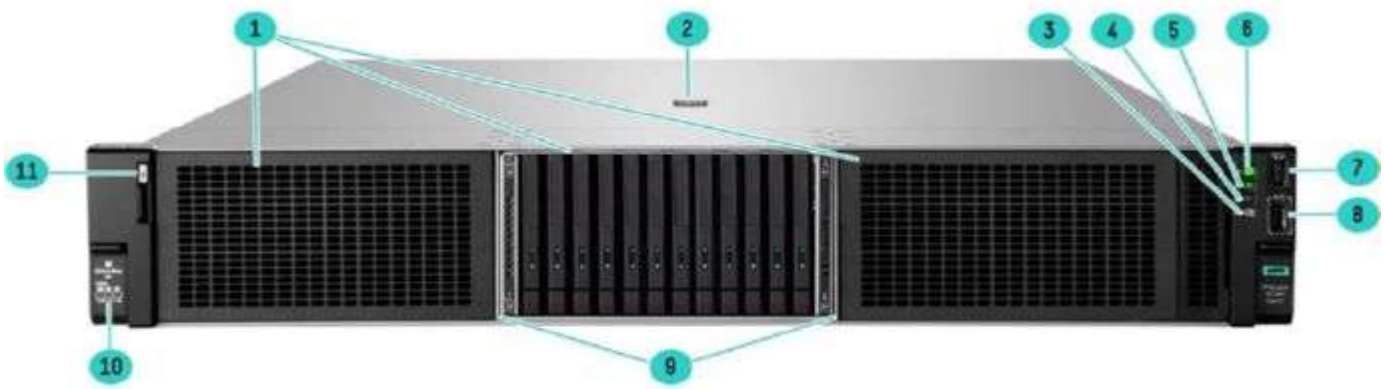
Notes: [1] Front NIC LED display doesn't support NIC LED ACT/LINK indication from ALOM/PCIE/FLOM NIC's

Overview



Front View - 12LFF chassis shown

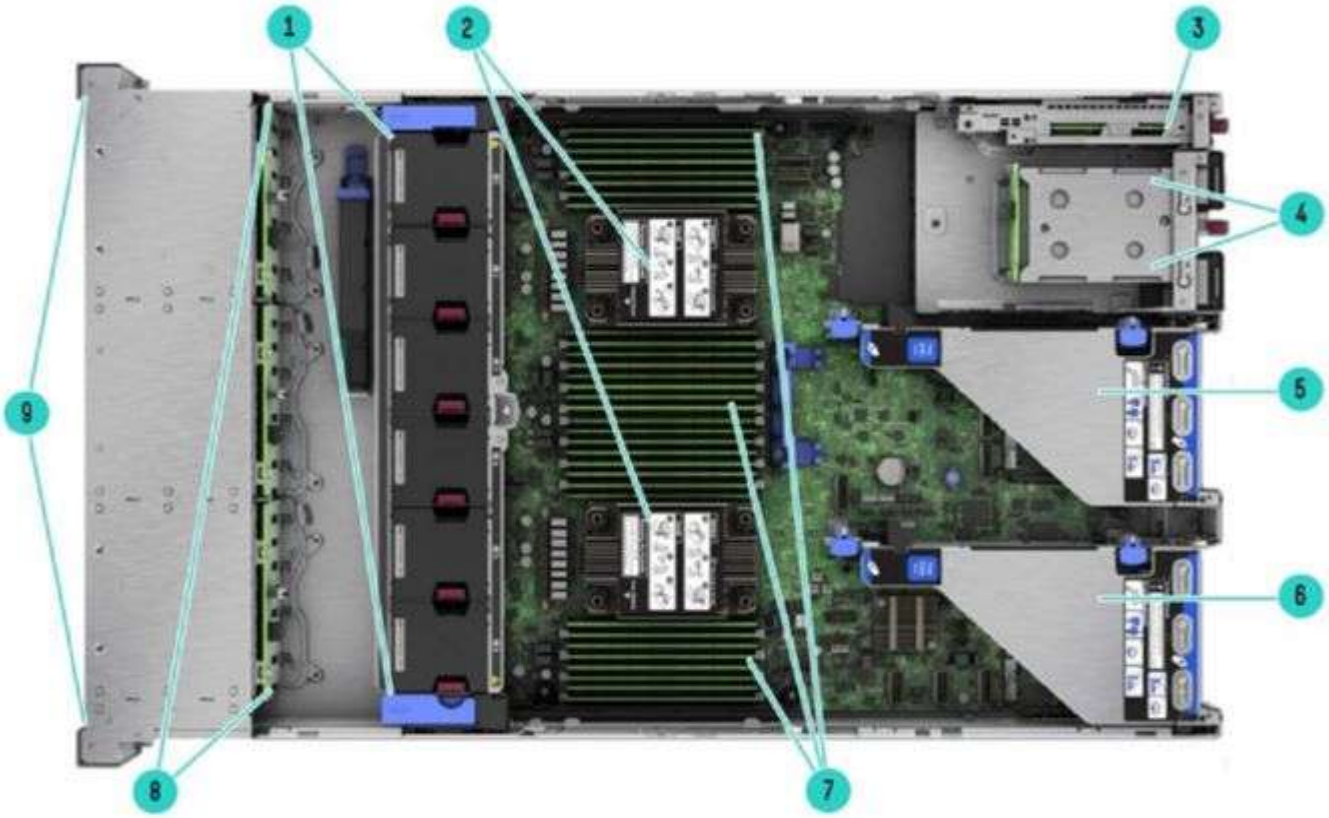
1.	Quick removal access panel	6.	iLO Service Port
2.	Unit Identification Button / LED	7.	USB 3.2 Gen1 Port
3.	NIC Status	8.	12 x LFF Media
4.	Health LED	9.	Drive support label
5.	Power On / Standby button / LED	10.	Serial Number Label Pull Tab



Front View -12EDSFF chassis shown

1	12EDSFF drive bays optical drive	7	iLO Service Port
2	Quick removal access panel	8	USB 3.2 Gne1
3	UID Button / LED	9	12x EDSFF Media
4	NIC Status	10	Drive support label
5	Health LED	11	Serial Number Label Pull Tab
6	Power On / Standby button and LED		

Overview



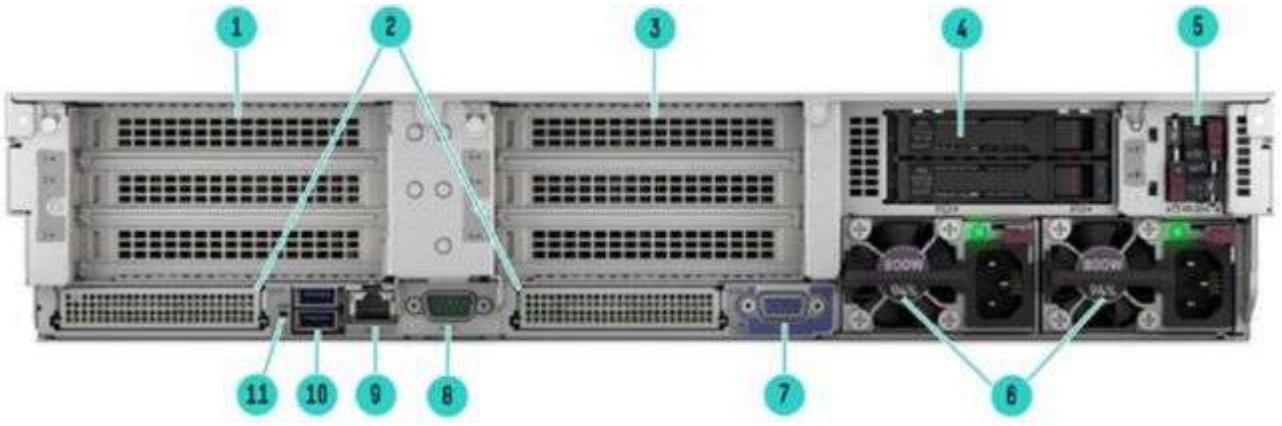
Internal View 8SFF chassis

1. Hot Plug Fans[1]
2. Processors, heatsinks showing
3. Optional NS204i-u Boot Device
4. Hot Plug redundant HPE Flexible Slot Power Supplies
5. Secondary Riser (Optional) (Requires second processor)
6. Primary Riser
7. DIMM slots, shown fully populated in 32 slots[2]
8. Drive Backplanes
9. Drive Cages

Notes:

– [1]High performance temperature fans optional
– [2]Shown fully populated in 32 slots (16 per processor)

Overview



Rear View - Standard for all DL380 Gen11

1.	Primary Riser. PCIe 5.0 Slots (Slots 1-3)	6.	Power Supply
2.	OCP NIC 3.0 Slots, shown covered	7.	VGA Connector
3..	Secondary Riser. PCIe 5.0 Slots (Slots 4-6)	8.	Serial Port Optional
4.	Tertiary Riser (Slots 7-8) shown with optional 2SFF drive cage installed	9.	Dedicated iLO Management Port
5.	Optional NS204i-u Boot Device	10	USB 3.2 Gen1 Connectors (2)
		11	UID Indicator LED

Notes: [1] Supports various NICs, and Storage controllers.

Platform Information

Form Factor

- 2U rack

Chassis Types

- 8SFF (SAS/SATA/NVMe) with optional 6SFF Universal Media Bay (P50728-B21), and/or up to 6SFF rear drive bay options.
- 24SFF bay (SAS/SATA/NVMe) with up to 6SFF rear drive bay options for a total 30 SFF drives.
- 8LFF supporting 2SFF front, and up to 4LFF rear or 2SFF rear drive bay options.
- 12LFF with optional 4LFF rear for a total of 16LFF drives.
- 12EDSFF (NVMe) supporting up to 36 EDSFF in front and up to 2SFF rear drive bay option

Notes:

– The 8SFF chassis is upgraded to support up to 24SFF (front) with a variety of 8SFF Drive Cages to select from, including 8SFF U.3 x4/x2 Trimode, 8SFF U.3 (x1 Trimode), and 8SFF SAS/SATA. See "Drive Cages" section within this document for options.

– The 8SFF chassis comes with an 8SFF U.3 x1 Trimode drive bay by default in bay 3.

– The Universal Media Bay (P50728-B21) is only available as an option for the 8SFF chassis and can only be populated in Box 1.

– The 2LFF primary and 2LFF secondary rear cages will consume all PCIe slots for the primary and secondary riser, respectively

– The 8LFF chassis cannot be upgraded to 12LFF front in the field.

– The 2LFF primary and 2LFF secondary rear cages supported in LFF chassis only.

Overview

System Fans

- High Performance Fan Kit - required for all CPUs over 205W TDP.

Notes:

- On 8SFF CTO server models ship with 4 standard fans.
- The 12LFF and 8LFF CTO server models ship with 4 standard fans.
- The 24SFF CTO server model ships with 6 high performance fans.
- The High-Performance fan kit (P48820-B21) is available to meet ambient temperature requirements.
- In general, the Maximum Performance fan kit is required when rear drives, or >205W Processors SKUs, or High Performance NVMe drives, three drive cages, mid-tray, GPU card, or certain backplanes are populated. See notes under each option category or each individual option for specifics.

Standard Features

Processors - Up to 2 of the following depending on model.

- The 2nd digit of the processor model number "x4xx" is used to denote the processor generation (i.e. 4=4th generation Intel Scalable Series Processors)
- The required firmware for the 5th Generation Intel Xeon Processors is not compatible with the DL380 Gen11 system board shipped with the 4th Generation Intel Xeon Processors. A new server order (latest firmware system board) is required for the activation of the 5th Generation Intel Xeon Processors. Field upgrade from the 4th Generation Intel Xeon Processors to the 5th Generation is not supported.

For more information regarding Intel Xeon processors, please see the following **https://www.intel.com/xeon.**

This table covers the public Intel offering only.

Processor Suffix	Description	Offering
H	Database and Analytics	Highest core counts. Database and Analytics usages benefit from Intel Data Streaming Accelerator (DSA) and Intel In-Memory Accelerator Analytics (IAA)
M	Media Transcode	Optimized for Intel Advanced Vector Extensions (AVX) frequencies to deliver better performance/watt for Media, AI, and HPC workloads.
N	Network/5G/Edge (High Throughput / Low Latency)	Designed for Network Function Visualization and networking workloads, such as: L3 forwarding, 5G User Plane Function, Open vSwitch Data Plane Development Kit, Vector Packet Processing FIB router, VPP IPsec, web server/NGINX, vEPC, vBNG, and vCMTS.
S	Storage and HCI	Optimized for Storage UMA use cases with increased UPI Bandwidth for vs Mainline SKUs.
P	Cloud - IAAS	Designed for cloud IaaS environments to deliver higher frequencies at constrained TDPs.
Q	Liquid Cooling	Liquid cooled processors with higher frequency and performance at same TDP.
U	Single Socket Optimized	Optimized for targeted platforms adequately served by the cores, memory bandwidth and IO capacity available from a single processor
V	Cloud - SAAS	Optimized for orchestration efficiency that delivers higher core counts and VMs per rack.
Y	Intel SST-PP (Speed Select Technology - Performance Profile)	Intel® SST-PP increases base frequency when fewer cores are enabled. Allows greater flexibility, deployment options and platform longevity.

Standard Features

5th Generation Intel® Xeon® Scalable Processor Family (Platinum)								
Intel® Xeon® Models	**Frequency**	**Cores**	**L3 Cache**	**Power**	**UPI**	**DDR5**	**SGX Enclave size**	**Die**
Platinum 8593Q Processor	2.2 GHz	64	320 MB	385W	4	5600 MT/s	512 GB	XCC
Platinum 8592+ Processor	1.9 GHz	64	320 MB	350W	4	5600 MT/s	512 GB	XCC
Platinum 8592V Processor	2.0 GHz	64	320 MB	330W	3	4800 MT/s	512 MB	XCC
Platinum 8581V[1] Processor	2.0 GHz	60	300 MB	270W	0	4800 MT/s	512 GB	XCC
Platinum 8580 Processor	2.0 GHz	60	300 MB	350W	4	5600 MT/s	512 GB	XCC
Platinum 8570 Processor	2.1 GHz	56	300 MB	350W	4	5600 MT/s	512 GB	XCC
Platinum 8568Y Processor	2.3 GHz	48	300 MB	350W	4	5600 MT/s	512 GB	XCC
Platinum 8562Y+ Processor	2.8 GHz	32	60.0 MB	300W	3	5600 MT/s	512 GB	MCC
Platinum 8558P Processor	2.7 GHz	48	260 MB	350W	3	5600 MT/s	512 GB	XCC
Platinum 8558 Processor	2.1 GHz	48	260 MB	330W	4	5200 MT/s	512 GB	XCC
Platinum 8558U1 Processor	2.0 GHz	48	260 MB	300W	0	4800 MT/s	512 GB	XCC

Notes:

- In 300Watt processor one socket configuration, the air cooling with Performance Heatsinks & Performance Fan Kits can be supported. Field upgrade to two socket is not supported with air cooling solution.
- [1] Single socket only, no dual socket support
- Intel® Speed Select enabled processors: Platinum 8593Q, 8592V, 8581V, 8568Y+, 8562Y+, 8558P, 8558 and 8558U

5th Generation Intel® Xeon® Scalable Processor Family (Gold 6)								
Intel® Xeon® Models	**Frequency**	**Cores**	**L3 Cache**	**Power**	**UPI**	**DDR5**	**SGX Enclave size**	**Die**
Gold 6558Q Processor	3.2 GHz	32	60.0 MB	350W	3	5200 MT/s	128 GB	MCC
Gold 6554S Processor	2.2 GHz	36	180 MB	270W	4	5200 MT/s	128 GB	XCC
Gold 6548N Processor	2.8 GHz	32	60.0 MB	300W	3	5200 MT/s	128 GB	MCC
Gold 6548Y+ Processor	2.5 GHz	32	60.0 MB	250W	3	5200 MT/s	128 GB	MCC
Gold 6544Y Processor	3.6 GHz	16	45.0 MB	270W	3	5200 MT/s	128 GB	MCC
Gold 6542Y Processor	2.9 GHz	24	60.0 MB	250W	3	5200 MT/s	128 GB	MCC
Gold 6538N Processor	2.1 GHz	32	60.0 MB	205W	3	5200 MT/s	128 GB	MCC
Gold 6538Y+ Processor	2.2 GHz	32	60.0 MB	225W	3	5200 MT/s	128 GB	MCC
Gold 6534 Processor	3.9 GHz	8	22.5 MB	195W	3	4800 MT/s	128 GB	MCC

Standard Features

Gold 6530 Processor	2.1 GHz	32	160 MB	270W	3	4800 MT/s	128 GB	XCC
Gold 6526Y Processor	2.8 GHz	16	37.5 MB	195W	3	5200 MT/s	128 GB	MCC

Notes:

– One or two processor(s)

– Intel® Speed Select enabled processors: Gold 6558Q, 6554S, 6548N, 6548Y+, 6544Y, 6542Y, 6538N, 6538Y+ and 6526Y

5th Generation Intel® Xeon® Scalable Processor Family (Gold 5)								
Intel® Xeon® Models	**Frequency**	**Cores**	**L3 Cache**	**Power**	**UPI**	**DDR5**	**SGX Enclave size**	**Die**
Gold 5520+ Processor	2.2 GHz	28	52. 5MB	205W	3	4800 MT/s	128 GB	MCC
Gold 5515+Processor	3.2 GHz	8	22.5 MB	165W	3	4800 MT/s	128 GB	MCC

Notes:

– One or two processor(s)

– Intel® Speed Select enabled processors: N.A.

5th Generation Intel® Xeon® Scalable Processor Family (Silver)								
Intel® Xeon® Models	**Frequency**	**Cores**	**L3 Cache**	**Power**	**UPI**	**DDR5**	**SGX Enclave size**	**Die**
Silver 4516+ Processor	2.2 GHz	24	45.0 MB	185W	2	4400 MT/s	64 GB	MCC
Silver 4514Y Processor	2.0 GHz	16	30.0 MB	150W	2	4400 MT/s	64 GB	MCC
Silver 4510 Processor	2.4 GHz	12	30.0 MB	150W	2	44000 MT/s	64 GB	EE LCC
Silver 4509Y Processor	2.6 GHz	8	22.5 MB	125W	2	4400 MT/s	64 GB	EE LCC

Notes:

– One or two processor(s)

– Intel® Speed Select enabled processors: 4516Y+, 4514Y and 4509Y

5th Generation Intel® Xeon® Scalable Processor Family (Bronze)								
Intel® Xeon® Models	**Frequency**	**Cores**	**L3 Cache**	**Power**	**UPI**	**DDR5**	**SGX Enclave size**	**Die**
Bronze 3508U Processor1	2.1 GHz	8	22.5 MB	125W	N/A	4400 MT/s	64 GB	EE LCC

Notes:

– [1]Single socket capable, no dual socket support

– Intel® Speed Select enabled processors: N.A.

– If 3508U is selected, then 96GB 5600MT/s Memory cannot be selected

– PCIe4.0 only

Standard Features

4th Generation Intel® Xeon® Scalable Processor Family (Platinum)							
Intel Xeon Model	**CPU Frequency**	**Cores**	**L3 Cache (MB)**	**Power**	**UPI Links**	**DDR5**	**SGX Enclave size (GB)**
Platinum 9462 Processor[2]	2.7GHz	32	75	350W	3	4800 MT/s	128
Platinum 8490H Processor	1.9GHz	60	112.5	350W	4	4800 MT/s	512
Platinum 8480+ Processor	2.0GHz	56	105	350W	4	4800 MT/s	512
Platinum 8470 Processor	2.0GHz	52	105	350W	4	4800 MT/s	512
Platinum 8470N Processor	1.7GHz	52	97.5	300W	4	4800 MT/s	128
Platinum 8470Q Processor[1]	2.1GHz	52	105	350W	4	4800 MT/s	512
Platinum 8468 Processor	2.1GHz	48	105	350W	4	4800 MT/s	512
Platinum 8468V Processor	2.4GHz	48	97.5	330W	3	4800 MT/s	128
Platinum 8462Y+ Processor	2.8GHz	32	60	300W	3	4800 MT/s	128
Platinum 8460Y+ Processor	2.0GHz	40	105	300W	4	4800 MT/s	128
Platinum 8458P Processor	2.7GHz	44	82.5	350W	3	4800 MT/s	512
Platinum 8452Y Processor	2.0GHz	36	67.5	300W	3	4800 MT/s	128
Platinum 8444H Processor	2.9GHz	16	45	270W	4	4800 MT/s	512

Notes:

- Processors do not ship with heatsinks or fan kits, these must be ordered separately.
- Processors with TDP equal to or greater than 150W through 350W require High Performance Heatsink (P48818-B21)
- Processors with TDP greater than 205W through 350W and mid-tray drive cage require HPE DL3xx/560 Gen11 High Performance Heatsink (P48905-B21)
- "Q" processors require Max Performance Heatsink (P48817-B21)
- Processors with TDP equal to or less than 150W require Standard Heatsink (P49145-B21)
- [1]Liquid cooled CPUs require Maximum Performance Heat Sink (P48817-B21). One heatsink covers both CPUs.
- [2] This is Intel High Bandwidth Memory (HBM) CPU.

Standard Features

4th Generation Intel® Xeon® Scalable Processor Family (Gold)							
Intel Xeon Models	**CPU Frequency**	**Cores**	**L3 Cache (MB)**	**Power**	**UPI Links**	**DDR5**	**SGX Enclave size (GB)**
Gold 6454S Processor	2.2GHz	32	60	270W	4	4800 MT/s	128
Gold 6448H Processor	2.4GHz	32	60	250W	3	4800 MT/s	512
Gold 6430 Processor	2.1GHz	32	60	270W	3	4800 MT/s	128
Gold 6414U Processor[1]	2.0GHz	32	60	250W	0	4800 MT/s	128
Gold 6458Q Processor	3.1GHz	32	60	350W	3	4800 MT/s	128
Gold 6448Y Processor	2.1GHz	32	60	225W	3	4800 MT/s	128
Gold 6444Y Processor	3.6GHz	16	45	270W	3	4800 MT/s	128
Gold 6442Y Processor	2.6GHz	24	60	225W	3	4800 MT/s	128
Gold 6438N Processor	2.0GHz	32	60	205	3	4800 MT/s	128
Gold 6438Y+ Processor	2.0GHz	32	60	205W	3	4800 MT/s	128
Gold 6434 Processor	3.7GHz	8	22.5	195W	3	4800 MT/s	128
Gold 6426Y Processor	2.5GHz	16	37.5	185W	3	4800 MT/s	128
Gold 6421N Processor	1.8GHz	32	60	185	0	4400 MT/s	128
Gold 6418H	2.1GHz	24	60	185W	3	4800 MT/s	512
Gold 6416H	2.2GHz	18	45	165W	3	4800 MT/s	512
Gold 5415+ Processor	2.9GHz	8	22.5	150W	3	4400 MT/s	128
Gold 5416S Processor	2.0GHz	16	30	150W	3	4400 MT/s	128
Gold 5418N Processor	1.8GHz	24	45	165W	3	4000 MT/s	128
Gold 5418Y Processor	2.0GHz	24	45	185W	3	4400 MT/s	128
Gold 5420+ Processor	2.0GHz	28	52.5	205W	3	4400 MT/s	128
Gold 5411N Processor	1.9GHz	24	45	165W	0	4400 MT/s	128

Notes:

- Processors do not ship with heatsinks or fan kits, these must be ordered separately.
- Processors with TDP greater than 150W through 350W require High Performance Heatsink (P48818-B21)
- Processors with TDP greater than 205W through 350W and mid-tray drive cage require HPE DL3xx/560 Gen11 High Performance Heatsink (P48905-B21)
- "Q" processors require Max Performance Heatsink (P48817-B21)
- Processors with TDP equal to or less than 150W require Standard Heatsink (P49145-B21)

Standard Features

- 8-Channel DDR5 @ 4800 MT/s
- [1]Single socket processor. No dual socket support.

4th Generation Intel® Xeon® Scalable Processor Family (Silver)							
Intel Xeon Models	**CPU Frequency**	**Cores**	**L3 Cache (MB)**	**Power**	**UPI Links**	**DDR5**	**SGX Enclave size (GB)**
Silver 4410Y Processor	2.0GHz	12	30	150W	2	4000 MT/s	64
Silver 4416+ Processor	2.0GHz	20	37.5	165W	2	4000 MT/s	64

Notes:

- Processors do not ship with heatsinks or fan kits, these must be ordered separately.
- Processors with TDP greater than 150W through 350W require High Performance Heatsink (P48818-B21)
- Processors with TDP greater than 205W through 350W and mid-tray drive cage require HPE DL3xx/560 Gen11 High Performance Heatsink (P48905-B21)
- Processors with TDP equal to or less than 150W require Standard Heatsink (P49145-B21)

4th Generation Intel® Xeon® Scalable Processor Family (Bronze)							
Intel Xeon Models	**CPU Frequency**	**Cores**	**L3 Cache (MB)**	**Power**	**UPI Links**	**DDR5**	**SGX Enclave size (GB)**
Bronze 3408U	1.8GHz	8	22.5	125W	0	4000 MT/s	64

Notes:

- Processors do not ship with heatsinks or fan kits, these must be ordered separately.
- Processors with TDP equal to or less than 150W require Standard Heatsink (P49145-B21)

Chipset

Intel C741 Chipset

Notes: For more information regarding Intel® chipsets, please see the following URL:
https://www.intel.com/content/www/us/en/products/chipsets/server-chipsets.html

---

On System Management Chipset

HPE iLO 6 ASIC

Read and learn more in the **iLO QuickSpecs**.

Memory

One of the following depending on model.

Standard Features

Type	HPE DDR5 Smart Memory, Registered (RDIMM)
DIMM Slots Available	32 16 DIMM slots per processor, 8channels per processor, 2 DIMMs per channel
Maximum capacity	8.0 TB 32 x 256 GB RDIMM @ 4800 MT/s (32 DIMMs only with 8SFF or 16SFF, 16 DIMMs maximum with 24SFF)

Notes: The maximum memory speed is limited by the processor selection.

Memory Protection

One of the following depending on model.

Advanced ECC

Advanced ECC uses single device data correction to detect and correct single and all multibit error that occurs within a single DRAM chip.

Expansion Slots

Primary Riser

Notes:

– Bus width indicates the number of physical electrical lanes running to the connector.
– There are 2 types of risers supported on Primary Slot
– x16 cards installed on x8 slots could observe sub-optimal performance.

Primary Riser1					
Slots #	**Technology**	**Bus Width**	**Connector Width**	**Slot Form Factor**	**Notes**
1	PCIe 5.0	X8	X16	Full-height, full-length slot	Proc 1
2	PCIe 5.0	X16	X16	Full-height, full-length slot	Proc 1
3	PCIe 5.0	X8	X16	Full-height, half-length slot	Proc 1

Standard Features

Primary Riser2					
Slots #	**Technology**	**Bus Width**	**Connector Width**	**Slot Form Factor**	**Notes**
1**	NA	NA	NA	NA	NA
1	PCIe 5.0	X16	X16	Full-height, full-length slot	Proc 1
2	PCIe 5.0	X16	X16	Full-height, full-length slot	Proc 1
3	PCIe 5.0	X16	X16	Full-height, half-length slot	Proc 1

**Notes: ** If Slot 1 of HPE DL380 Gen11 2U 3x16 Prim Riser Kit needs to be enabled then 3 x16 Primary Cable Kit (P56073-B21) must be selected.

Secondary Riser:

Notes:

– Bus Width Indicates the number of physical electrical lanes running to the connector.

– There are 2 types of risers support on Secondary Slot

– x16 cards installed on x8 slots could observe sub-optimal performance.

Secondary Riser1					
Slots #	**Technology**	**Bus Width**	**Connector Width**	**Slot Form Factor**	**Notes**
4	PCIe 5.0	X8	X16	Full-height, full-length slot	Proc 2
5	PCIe 5.0	X16	X16	Full-height, full-length slot	Proc 2
6	PCIe 5.0	X8	X16	Full-height, half-length slot	Proc 2

Secondary Riser2					
Slots #	**Technology**	**Bus Width**	**Connector Width**	**Slot Form Factor**	**Notes**
4*	NA	NA	NA	NA	NA
4	PCIe 5.0	X16	X16	Full-height, full-length slot	Proc 2
5	PCIe 5.0	X16	X16	Full-height, full-length slot	Proc 2
6	PCIe 5.0	X16	X16	Full-height, half-length slot	Proc 2

Standard Features

Notes: * If Slot 4 of HPE DL380 Gen11 2U 3x16 Sec Riser Kit needs to be enabled then 3 x16 Secondary Cable Kit (P56074-B21) must be selected.

Tertiary Riser

Notes:

- Bus Width Indicates the number of physical electrical lanes running to the connector.
- There is 1 type of riser supported on the Tertiary Slot
- x16 cards installed on x8 slots could observe sub-optimal performance.

Tertiary Riser1 (default)					
Slots #	**Technology**	**Bus Width**	**Connector Width**	**Slot Form Factor**	**Notes**
7	PCIe 5.0	X16	X16	Full-height, full-length slot	Proc 2
8	PCIe 4.0	X16	X16	Full-height, full-length slot	Proc 2

Tertiary Riser1 (with Optional Tertiary Riser FIO x8 Enablement Kit P53632-B21)					
Slots #	**Technology**	**Bus Width**	**Connector Width**	**Slot Form Factor**	**Notes**
7	PCIe 5.0	X16	X16	Full-height, full-length slot	Proc 2
8	PCIe 5.0	X8	X16	Full-height, full-length slot	Proc 2

Graphics

Integrated Video Standard

- Video modes up to 1920 x 1200@60Hz (32 bpp)
- 16MB Video Memory

Maximum Internal Storage

Drive	Capacity	Configuration
Hot Plug SFF SAS HDD	91.2 TB	24+8+6 x 2.4TB
Hot Plug SFF SAS SSD	583.3 TB	24 +8+6 15.35TB
Hot Plug SFF SATA HDD	76 TB	24+8+6 x 2 TB
Hot Plug SFF SATA SSD	291.84 TB	24 +8+ 6 x 7.68 TB
Hot Plug LFF SAS HDD	360 TB	12+4+4x 18 TB (with optional rear LFF drive cage)
Hot Plug LFF SATA HDD	360 TB	12+4+4 x 18 TB (with optional rear LFF drive cage)
Hot Plug SFF NVMe PCIe SSD	374.4 TB	24+ x 15.36TB + 6 x 960GB<10W (with optional rear Primary and Secondary 2SFF and rear 2SFF drive cages)
EDSFF NVMe	550.8 TB	36EDSFF Drives x 15.3 TB

Standard Features

Notes:

– If all seven Drive cages are selected, then maximum SAS Drive support will be 30 drives (16+8+6).

– Maximum SAS Drive support can be 36 Drives (24+8+4), if one 2SFF Primary / Secondary cage is not selected.

Internal Storage Devices

- **Optical Drive**

  Optional: DVD-ROM, DVD-RW

- **Hard Drives**

  None ship standard

Power Supply

- HPE 1800W-2200W Flex Slot Titanium Hot Plug Low Halogen Power Supply Kit
  Notes: 1 available in 96% efficiency.
- HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
  Notes: 1 available in 94% efficiency.
- HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
- HPE 1000W Flex Slot Titanium Hot Plug Low Halogen Power Supply Kit
  Notes: 1 available in 96% efficiency.
- HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
  Notes: 1 available in 94% efficiency.

HPE Flexible Slot (Flex Slot) Power Supplies share a common electrical and physical design that allows for hot plug, tool-less installation into HPE ProLiant Gen11 Performance Servers. Flex Slot power supplies are certified for high-efficiency operation and offer multiple power output options, allowing users to "right-size" a power supply for specific server configurations. This flexibility helps to reduce power waste, lower overall energy costs, and avoid "trapped" power capacity in the data center.

A standard 6-foot IEC C-13/C-14 jumper cord (A0K02A) is included with each standard AC power supply option kit. If a different power cord is required, please check the **ProLiant Power Cables** web page. review the power requirements for your selected system, please use the **HPE Power Advisor Tool**.

For information on power specifications and technical content visit **HPE Server power supplies**.

Storage Controllers

The available Gen11 controllers are depicted below.

Software RAID Controller

- Intel VROC SATA for HPE ProLiant Gen11

Standard Features

Notes:

- All models feature an embedded storage controller, with embedded software SATA RAID support for up to 14 bays.
- Intel VROC for HPE ProLiant Gen11 is an enterprise, hybrid Software RAID solution specifically designed for SSDs.
- Intel VROC is a software-based solution utilizing Intel CPU to RAID or HBA direct connected drives.
- RAID Support- 0/1/5/10.
- Windows and Linux OS support.
- Host Tools- Windows GUI/CLI, Linux CLI.
- UEFI Support- HII Utility, OBSE.
- iLO Support- IML, Alert, SNMP, AHS.
- iLO Redfish- Redfish Read.
- Intel VROC SATA for HPE ProLiant Gen11 will operate in UEFI mode only. For legacy support an additional storage controller will be needed.
- Intel VROC SATA is off by default and must be enabled.
- VROC cannot support mix among SATA/sSATA/tSATA due to VROC limitation. Visit **Intel Virtual RAID on CPU for HPE Gen11 User Guide**.
  - 12LFF: Mix of raid is not allowed across boxes/drive bays. Each LFF drive cage (box1/2/3) has four bays and is separated with different drive group with Intel VROC SATA RAID configuration.
  - 8SFF: Mix of raid is not allowed across boxes/drive bays. One box contains 8 bays and is separated into two drive groups (Bay1-4 and Bay5-8) with Intel VROC SATA RAID configuration.

- Intel VROC NVMe for HPE ProLiant Gen11

Notes:

- All models feature 4 x8 PCIe 5.0 connectors per socket for NVMe connectivity, provides support for up to 8 direct attach x4 NVMe bays.
- Only supported on SFF models.
- Intel VROC for HPE ProLiant Gen11 is an enterprise, hybrid Software RAID solution specifically designed for NVMe SSDs connected directly to the CPU. Intel VROC is a software-based solution utilizing Intel CPU to RAID or HBA direct connected drives.
- Intel Virtual RAID on CPU Standard for RAID 0/1/10 (S3Q19A/S3Q39AAE) or Premium SKU for RAID 0/1/5/10 (R7J57A/R7J59AAE) must be ordered to enable RAID support.
- Windows, Linux, VMware OS support.
- Host Tools- Windows GUI/CLI, Linux CLI.
- UEFI Support- HII Utility, OBSE.
- Active health monitoring of NVMe M.2 drives requires use of SMART tools.
- Intel VROC NVMe for HPE ProLiant Gen11 will operate in UEFI mode only. For legacy support an additional Tri-
- Mode controller will be needed.
- For NVMe SSDs only, no PCIe card support.

Essential RAID Controller

- HPE Smart Array E208e-p SR Gen10 Controller

Tri-Mode Controller

- HPE MR416i-p Gen11 Controller
- HPE MR416i-o Gen11 Controller
- HPE MR216i-p Gen11 Controller
- HPE MR216i-o Gen11 Controller
- HPE MR408i-o Gen11 Controller

Standard Features

- HPE SR932i-p Gen11 Controller[1,2]

Notes:

- PE80xx NVMe drives are not supported.
- [1]Requires x16 physical and electrical riser slot
- [2]If second controller is required, must select secondary riser
- Controllers with cache require either P02377-B21 HPE Smart Storage Hybrid Capacitor with 145mm Cable Kit or P01366-B21 HPE 96W Smart Storage Lithium-ion Battery with 145mm Cable Kit.

Interfaces

Serial	Optional Rear
Display Port	1 optional front display port via Universal Media Bay
VGA Port	1 standard, rear for all chassis. 1 Optional front display port (Via Universal Media Bay) **Notes:** Both ports are not active simultaneously.
Network Ports	None standard. Choice of OCP networking card or stand-up networking card required. BTO models will come pre-selected with a primary networking card.
HPE iLO Remote Management Network Port	1 Gb Dedicated, rear
Front iLO Service Port	1 standard (Not available when System Insight Display Kit is ordered)
USB 3.2 Port Gen1	Up to 4 total: 1 front (3.2 Gen1), 2 rear (3.0), 2 internal (secure - 1 - 3.2 Gen1, 1 - 2.0), 1 optional USB 2.0 front via Universal Media Bay
System Insight Display **(SID)**	Optional **Notes:** Not shipping as standard. Available as a CTO option or as a field upgrade (P48819-B21).

Operating Systems and Virtualization Software Support for HPE Servers

HPE servers are designed for seamless integration with partner Operating Systems and Virtualization Software. By collaborating closely with our partners, we ensure that their products are optimized, certified, and fully supported within your HPE server environment.

Access the certified and supported servers for each of the OS and Virtualization software: **HPE Servers Support & Certification Matrices**

HPE Server UEFI

Unified Extensible Firmware Interface (UEFI) is an industry standard that provides better manageability and more secured configuration than the legacy ROM while interacting with your server at boot time. HPE ProLiant Gen11 servers have a UEFI Class 3 implementation to support UEFI Mode.

Notes: The UEFI System Utilities tool is analogous to the HPE ROM-Based Setup Utility (RBSU) of legacy BIOS. For more information, please visit **https://www.hpe.com/servers/uefi**.

Standard Features

UEFI enables numerous new capabilities specific to HPE ProLiant servers such as:

- Secure Boot and Secure Start enable for enhanced security
- Embedded UEFI Shell
- Operating system specific functionality
- Mass Configuration Deployment Tool using iLO RESTful API that is Redfish API Conformant
- Support for > 2.2 TB (using GPT) boot drives
- PXE boot support for IPv6 networks
- USB 3.0 Stack
- Workload Profiles for simple performance optimization

UEFI Boot Mode only

- TPM 2.0 Support

Notes: Enabling TPM 2.0 no longer requires TPM module option kit for Gen11. It is an embedded feature.

- iSCSI Software Initiator Support.
- NVMe Boot Support
- HTTP/HTTPs Boot support as a PXE alternative.
- Platform Trust Technology (PTT) can be enabled.
- Boot support for option cards that only support a UEFI option ROM

Notes: For UEFI Boot Mode, boot environment and OS image installations should be configured properly to support UEFI.

Industry Standard Compliance

- ACPI 6.4 Compliant
- PCIe 5.0 Compliant
- Wake on LAN (WoL) Support
- Microsoft® Logo certifications
- PXE Support
- VGA
- Display Port

Notes: This support is on the optional Universal Media Bay.

- USB 3.2 Gen1 Compliant
- USB 2.0 Compliant (via Universal Media Bay)

Notes: This support is on the optional Universal Media Bay.

- Energy Star
- SMBIOS 3.4
- Redfish API
- IPMI 2.0
- Secure Digital 4.0
- 2.0 Support

Notes: Enabling TPM 2.0 no longer requires TPM module option kit for Gen11. It is an embedded feature.

- Advanced Encryption Standard (AES)
- Triple Data Encryption Standard (3DES)
- SNMP v3
- TLS 1.2
- DMTF Systems Management Architecture for Server Hardware Command Line Protocol (SMASH CLP)
- Active Directory v1.0
- ASHRAE A3/A4

Standard Features

Notes: For additional technical, thermal details regarding ambient temperature, humidity, and feature support, please visit **DL380 Gen11 Extended Ambient Temperature Guidelines**

- European Union Erp Lot 9 Regulation

Notes:

- Beginning on January 1st, 2024, units sold into the European Union (EU), European Economic Area (EEA), the United Kingdom, or Switzerland must include more efficient AC power supplies: 94% for multi-output and 96% for single-output. HPE Flexible Slot power supplies are single-output, and part numbers 865438-B21, P03178-B21, and P44712-B21 are 96% efficient, thus meeting requirements.
- HPE is on target to fulfil compliant systems ahead of time and will begin enforcing these requirements in advance to satisfy requests with the current power supplies by the set deadline.
- Please visit: **https://www.hpe.com/us/en/about/environment/msds-specs-more.html** for more information regarding HPE Lot 9 conformance.
- UEFI (Unified Extensible Firmware Interface Forum) 2.8

Embedded Management

HPE Integrated Lights-Out (HPE iLO)
Monitor your servers for ongoing management, service alerting, reporting and remote management with HPE iLO.
Learn more at **https://www.hpe.com/info/iLO**.

UEFI

Configure and boot your servers securely with industry standard Unified Extensible Firmware Interface (UEFI).

Intelligent Provisioning
Hassle free server and OS provisioning for 1 or more servers with Intelligent Provisioning.

Learn more at **https://support.hpe.com/hpesc/public/docDisplay?docId=c04465280&docLocale=en_US**

iLO RESTful API

iLO RESTful API is DMTF Redfish API implementation and offers simplified server management automation such as configuration and maintenance tasks based on modern industry standards. Learn more at **https://www.hpe.com/info/restfulapi.**

OpenBMC Support
OpenBMC Capable through iLO6 Transfer of Ownership Process.
Learn more at **OpenBMC Support**

Active Health System

The HPE Active Health System (AHS) is an essential component of the iLO management portfolio that provides continuous, proactive health monitoring of HPE servers. Learn more at **https://www.hpe.com/servers/ahs**.

Smart Update

Standard Features

Keep your servers up to date with the HPE Smart Update solution by using Smart Update Manager (SUM) to optimize the firmware and driver updates of the Service Pack for ProLiant (SPP).

iLO Amplifier Pack

Designed for large enterprise and service provider environments with hundreds of HPE servers, the iLO Amplifier Pack is a free, downloadable open virtual application (OVA) that delivers the power to discover, inventory and update Gen8, Gen9 and Gen10 HPE servers. Use with an iLO Advanced License to unlock full capabilities.

Learn more at **https://www.hpe.com/servers/iLOamplifierpack**

RESTful Interface Tool

RESTful Interface tool (iLOREST) is a single scripting tool to provision using iLO RESTful API to discover and deploy servers at scale. Learn more at **https://www.hpe.com/info/resttool**.

Scripting Tools
Provision one to many servers using your own scripts to discover and deploy with Scripting Tool (STK) for Windows and Linux or Scripting Tools for Windows PowerShell. Learn more at **https://www.hpe.com/servers/powershell**.

HPE OneView Standard
HPE OneView is an on premise, multi-generational server monitoring and management solution. HPE OneView Standard can be used for inventory, health monitoring, alerting, and reporting without additional fees. Customers can upgrade their management experience with an HPE OneView Advanced license all provided by the same tool. Learn more at **https://www.hpe.com/info/oneview**.

HPE Compute Ops Management
HPE is intelligently transforming compute management with an intuitive cloud operating experience through HPE GreenLake cloud platform to streamline and secure operations from edge-to-cloud. Automated key lifecycle tasks, for onboarding, updating, managing, and monitoring HPE servers, brings agility and greater efficiencies to wherever compute devices reside via a unified single browser-based interface. Manage single locations or multiple, distributed sites. Keep tens to thousands of servers secure with batch policy controls and automated updates.
Compute Ops Management is cloud-native software that is continually updated with new services, features, patches, and fixes. The management application resides in the HPE GreenLake cloud platform (access via **https://console.greenlake.hpe.com**) and leverages the HPE GreenLake architecture, security, and unified operations.
A 3-year subscription to Compute Ops Management is added by default when ordering an HPE ProLiant Gen11 rack, tower, or micro server.
For more information, visit the Compute Ops Management QuickSpecs:
https://www.hpe.com/psnow/doc/a50004263enw

Security

- UEFI Secure Boot and Secure Start support
- Tamper-free updates - components digitally signed and verified
- Immutable Silicon Root of Trust
- Ability to rollback firmware
- FIPS 140-2 validation
- Secure erase of NAND/User data
- Configurable for PCI DSS compliance
- TPM (Trusted Platform Module) 2.0 option

Notes: Enabling TPM 2.0 no longer requires TPM module option kit for Gen11. It is an embedded feature.

Standard Features

- Advanced Encryption Standard (AES) and Triple Data Encryption Standard (3DES) on browser
- Bezel Locking Kit option
- Support for Commercial National Security Algorithms (CNSA)
- Chassis Intrusion detection option
- Secure Recovery - recover critical firmware to known good state on detection of compromised firmware

Warranty

This product is covered by a global limited warranty and supported by HPE Services and a worldwide network of Hewlett Packard Enterprise Authorized Channel Partners resellers. Hardware diagnostic support and repair is available for three years from date of purchase. Support for software and initial setup is available for 90 days from date of purchase. Enhancements to warranty services are available through HPE Services operational services or customized service agreements. Hard drives have either a one year or three year warranty; refer to the specific hard drive QuickSpecs for details.

Notes: Server Warranty includes 3-Year Parts, 3-Year Labor, 3-Year Onsite support with next business day response. Warranty repairs may be accomplished through the use of Customer Self Repair (CSR) parts. These parts fall into two categories: 1) Mandatory CSR parts are designed for easy replacement. A travel and labor charge will result when customers decline to replace a Mandatory CSR part; 2) Optional CSR parts are also designed for easy replacement but may involve added complexity. Customers may choose to have Hewlett Packard Enterprise replace Optional CSR parts at no charge. Additional information regarding worldwide limited warranty and technical support is available.

https://www.hpe.com/support/ProLiantServers-Warranties.

Server Management

HPE iLO Advanced

HPE iLO Advanced licenses offer smart remote functionality without compromise, for all HPE ProLiant servers. The license includes the full integrated remote console, virtual keyboard, video, and mouse (KVM), multi-user collaboration, console record and replay, and GUI-based and scripted virtual media and virtual folders. You can also activate the enhanced security and power management functionality.

HPE OneView Advanced

HPE OneView Advanced offers a sophisticated level of automation to infrastructure management by taking a template driven approach to provisioning, updating, and integrating compute, storage, and networking infrastructure. It builds upon the base features of HPE OpenView Standard, provides full-featured licenses which can be purchased for managing multiple HPE server generations.

To learn more visit **https://www.hpe.com/info/oneview**.

HPE Insight Cluster Management Utility (CMU)

HPE Insight Cluster Management Utility is a HyperScale management framework that includes software for the centralized provisioning, management and monitoring of nodes and infrastructure. Learn more at **https://www.hpe.com/info/cmu**.

One Config Simple (OCS/SCE)

OCS/SCE is a guided self-service tool to help sales and non-technical people provide customers with initial configurations in 3 to 5 minutes. You may then send the configuration on for configuration help, or use in your existing ordering processes. If you require "custom" rack configuration or configuration for products not available in SCE, please contact Hewlett Packard Enterprise Customer Business Center or an Authorized Partner for assistance. **https://h22174.www2.hpe.com/SimplifiedConfig/Welcome#**

Rack and Power Infrastructure

The story may end with servers, but it starts with the foundation that makes compute go - and business grow. We've reinvented our entire portfolio of rack and power products to make IT infrastructure more secure, more practical, and more efficient. In other words, we've created a stronger, smarter, and simpler infrastructure to help you get the most out of your IT equipment. As an industry leader, Hewlett Packard Enterprise is uniquely positioned to address the key concerns of power, cooling, cable management and system access.

HPE G2 Advanced and Enterprise Racks are perfect for the server room or today's modern data center with enhanced airflow and thermal management, flexible cable management, and a 10 year Warranty to support higher density computing.

HPE G2 PDUs offer reliable power in flexible form factors that operate at temperatures up to 60°C, include color-coded outlets and load segments and a low-profile design for optimal access to the rack and support for dense rack environments.

HPE Uninterruptible Power Systems are cost-effective power protection for any type workload. Some UPSs include options for remote management and extended runtime modules so you're critical dense data center is covered in power outages.

Optional Features

HPE KVM Solutions include a console and switches designed to work with your server and IT equipment reliably. We've got a cost-effective KVM switch for your first rack and multiple connection IP switches with remote management and security capabilities to keep your data center rack up and running.

Learn more about HPE Racks, KVM, PDUs and UPSs at **HPE Rack and Power Infrastructure**.

Service and Support

HPE Services

No matter where you are in your digital transformation journey, you can count on HPE Services to deliver the expertise you need when, where and how you need it. From planning to deployment, ongoing operations and beyond, our experts can help you realize your digital ambitions.

https://www.hpe.com/services

Consulting Services

No matter where you are in your journey to hybrid cloud, experts can help you map out your next steps. From determining what workloads should live where, to handling governance and compliance, to managing costs, our experts can help you optimize your operations.

https://www.hpe.com/services/consulting

HPE Managed Services

HPE runs your IT operations, providing services that monitor, operate, and optimize your infrastructure and applications, delivered consistently and globally to give you unified control and let you focus on innovation.

HPE Managed Services | HPE

Operational services

Optimize your entire IT environment and drive innovation. Manage day-to-day IT operational tasks while freeing up valuable time and resources. Meet service-level targets and business objectives with features designed to drive better business outcomes.

https://www.hpe.com/services/operational

HPE Complete Care Service

HPE Complete Care Service is a modular, edge-to-cloud IT environment service designed to help optimize your entire IT environment and achieve agreed upon IT outcomes and business goals through a personalized experience. All delivered by an assigned team of HPE Services experts. HPE Complete Care Service provides:

- A complete coverage approach -- edge to cloud
- An assigned HPE team
- Modular and fully personalized engagement
- Enhanced Incident Management experience with priority access
- Digitally enabled and AI driven customer experience

https://www.hpe.com/services/completecare

HPE Tech Care Service

HPE Tech Care Service is the operational support service experience for HPE products. The service goes beyond traditional support by providing access to product specific experts, an AI driven digital experience, and general technical guidance to not only reduce risk but constantly search for ways to do things better. HPE Tech Care Service delivers a customer-centric, AI driven, and digitally enabled customer experience to move your business forward. HPE Tech Care Service is available in three response levels. Basic, which provides 9x5 business hour availability and a 2-hour response time. Essential which provides a 15-minute response time 24x7 for most enterprise level customers, and Critical which includes a 6-hour repair commitment where available and outage management response for severity 1 incidents.

https://www.hpe.com/services/techcare

Service and Support

HPE Lifecycle Services

HPE Lifecycle Services provide a variety of options to help maintain your HPE systems and solutions at all stages of the product lifecycle. A few popular examples include:

- Lifecycle Install and Startup Services: Various levels for physical installation and power on, remote access setup, installation and startup, and enhanced installation services with the operating system.
- HPE Firmware Update Analysis Service: Recommendations for firmware revision levels for selected HPE products, taking into account the relevant revision dependencies within your IT environment.
- HPE Firmware Update Implementation Service: Implementation of firmware updates for selected HPE server, storage, and solution products, taking into account the relevant revision dependencies within your IT environment.
- Implementation assistance services: Highly trained technical service specialists to assist you with a variety of activities, ranging from design, implementation, and platform deployment to consolidation, migration, project management, and onsite technical forums.
- HPE Service Credits: Access to prepaid services for flexibility to choose from a variety of specialized service activities, including assessments, performance maintenance reviews, firmware management, professional services, and operational best practices.

Notes: To review the list of Lifecycle Services available for your product go to:
https://www.hpe.com/services/lifecycle
For a list of the most frequently purchased services using service credits, see the **HPE Service Credits Menu**

Other Related Services from HPE Services:

HPE Education Services

Training and certification designed for IT and business professionals across all industries. Broad catalogue of course offerings to expand skills and proficiencies in topics ranging from cloud and cybersecurity to AI and DevOps. Create learning paths to expand proficiency in a specific subject. Schedule training in a way that works best for your business with flexible continuous learning options.
https://www.hpe.com/services/training

Defective Media Retention

An option available with HPE Complete Care Service and HPE Tech Care Service and applies only to Disk or eligible SSD/Flash Drives replaced by HPE due to malfunction.
Consult your HPE Sales Representative or Authorized Channel Partner of choice for any additional questions and services options.

Parts and Materials

HPE will provide HPE-supported replacement parts and materials necessary to maintain the covered hardware product in operating condition, including parts and materials for available and recommended engineering improvements.
Parts and components that have reached their maximum supported lifetime and/or the maximum usage limitations as set forth in the manufacturer's operating manual, product quick-specs, or the technical product data sheet will not be provided, repaired, or replaced as part of these services.

How to Purchase Services

Services are sold by Hewlett Packard Enterprise and Hewlett Packard Enterprise Authorized Service Partners:

- Services for customers purchasing from HPE or an enterprise reseller are quoted using HPE order configuration tools.
- Customers purchasing from a commercial reseller can find services at **https://ssc.hpe.com/portal/site/ssc/**

Service and Support

AI Powered and Digitally Enabled Support Experience

Achieve faster time to resolution with access to product-specific resources and expertise through a digital and data driven customer experience

Sign into the HPE Support Center experience, featuring streamlined self-serve case creation and management capabilities with inline knowledge recommendations. You will also find personalized task alerts and powerful troubleshooting support through an intelligent virtual agent with seamless transition when needed to a live support agent.

https://support.hpe.com/hpesc/public/home/signin

Consume IT On Your Terms

HPE GreenLake edge-to-cloud platform brings the cloud experience directly to your apps and data wherever they are-the edge, colocations, or your data center. It delivers cloud services for on-premises IT infrastructure specifically tailored to your most demanding workloads. With a pay-per-use, scalable, point-and-click self-service experience that is managed for you, HPE GreenLake edge-to-cloud platform accelerates digital transformation in a distributed, edge-to-cloud world.

- Get faster time to market
- Save on TCO, align costs to business
- Scale quickly, meet unpredictable demand
- Simplify IT operations across your data centers and clouds

To learn more about HPE Services, please contact your Hewlett Packard Enterprise sales representative or Hewlett Packard Enterprise Authorized Channel Partner. Contact information for a representative in your area can be found at "Contact HPE" **https://www.hpe.com/us/en/contact-hpe.html**

For more information

https://www.hpe.com/services

Pre-Configured Models

Pre-Configured models ship with the configurations below.

- Pre-Configured models ship with the configurations below. Options can be selected from the Core or Additional options section of this QuickSpecs.
- Hewlett Packard Enterprise does not allow factory integration of options into pre-configured models. Any additional options purchased will not be shipped inside the server.
- Network Choice models do not include embedded LOM.

SKU Number	P52560-B21 P52560-291 P52560-421	P52561-B21 P52561-291 P52561-421	P52562-B21 P52562-291 P52562-421
Model Name	P52560-B21 & P52560-291: HPE ProLiant DL380 Gen11 4410Y 2.0GHz 12-core 1P 32GB-R MR408i-o NC 8SFF 800W PS Server	P52561-B21 & P52561-291: HPE ProLiant DL380 Gen11 5416S 2.0GHz 16-core 1P 32GB-R MR408i-o NC 8SFF 800W PS Server	P52562-B21 & P52562-291: HPE ProLiant DL380 Gen11 4410Y 2.0GHz 12-core 1P 32GB-R NC 12LFF 800W PS Server
	P52560-421: HPE ProLiant DL380 Gen11 4410Y 2.0GHz 12-core 1P 32GB-R MR408i-o NC 8SFF 1000W PS Server	P52561-421: HPE ProLiant DL380 Gen11 5416S 2.0GHz 16-core 1P 32GB-R MR408i-o NC 8SFF 1000W PS Server	P52562-421: HPE ProLiant DL380 Gen11 4410Y 2.0GHz 12-core 1P 32GB-R NC 12LFF 1000W PS Server
Chassis	HPE ProLiant DL380 Gen11 8SFF NC Configure-to-order Server	HPE ProLiant DL380 Gen11 8SFF NC Configure-to-order Server	HPE ProLiant DL380 Gen11 12LFF NC Configure-to-order Server
Processor	4410Y (12 core, 2.0 GHz, 150W)	5416S (16 core, 2.0 GHz, 150W)	4410Y (12 core, 2.0 GHz, 150W)
Number of Processors	One		
Memory	32 GB (1x32 GB, 4800 MT/s) **Notes:** Runs at 4000 MT/s due to processor limitation.	32 GB (1x32 GB, 4800 MT/s) **Notes:** Runs at 4400 MT/s due to processor limitation.	32 GB (1x32 GB, 4800 MT/s) **Notes:** Runs at 4000 MT/s due to processor limitation.
Network Controller	Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE **Notes:** No embedded networking	Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T Adapter for HPE **Notes:** No embedded networking	Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE **Notes:** No embedded networking
Storage Controller	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller	Embedded SATA controller (AHCI or Intel SATA software RAID controller)
Included Hard Drives	None ship standard, 8 SFF supported	None ship standard, 8 SFF supported	None ship standard, 12 LFF supported
Optical Drive	Optional DVD-ROM Optional via Universal Media Bay External support only	Optional DVD-ROM Optional via Universal Media Bay External support only	Not supported
Expansion Slots	3-slots (x8, x16, x8) as standard; upgradeable to 8-slots in a 2 processor configuration		

Pre-Configured Models

Power Supply	P52560-B21 & P52560-291: 1x HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P52561-B21 & P52561-291: 1x HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P52562-B21 & P52562-291: 1x HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
	P52560-421: 1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit	P52561-421: 1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit	P52562-421: 1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit
Fans	4x Standard Fans		
Management	HPE iLO 6		
Rail Kit	HPE ProLiant DL3XX Gen11 Easy Install Rail 3 Kit		
Security	TPM (Trusted Platform Module)		
Form Factor	2U Rack		
Warranty	Server warranty includes 3-year parts, 3-year labor, 3-year onsite support with next business day response.		

Pre-Configured Models

SKU Number	P52564-B21 P52564-291 P52564-421	P58417-B21 P58417-291	P60636-B21 P60636-291 P60636-421
Model Name	P52564-B21 & P52564-291: HPE ProLiant DL380 Gen11 5415+ 2.9GHz 8-core 1P 32GB-R MR408i-o NC 8SFF 800W PS Server	P58417-B21: HPE ProLiant DL380 Gen11 6430 2.1GHz 32-core 1P 64GB-R NC 8SFF 1000W PS Server	P60636-B21 & P60636-291: HPE ProLiant DL380 Gen11 4416+ 2.0GHz 20-core 1P 32GB-R MR408i-o NC 8SFF 800W PS Server
	P52564-421: HPE ProLiant DL380 Gen11 5415+ 2.9GHz 8-core 1P 32GB-R MR408i-o NC 8SFF 1000W PS Server	P58417-291 & **P58417**-421: HPE ProLiant DL380 Gen11 6430 2.1GHz 32-core 1P 32GB-R NC 8SFF 1000W PS Server	P60636-421: HPE ProLiant DL380 Gen11 4416+ 2.0GHz 20-core 1P 32GB-R MR408i-o NC 8SFF 1000W PS Server
Chassis	HPE ProLiant DL380 Gen11 8SFF NC Configure-to-order Server		
Processor	5415+ (8 core, 2.9 GHz, 150W)	6430 (32 core, 2.1 GHz, 270W)	4416+ (20 core, 2.0 GHz, 165W)
Number of Processors	One	One with high performance heatsink	One with high performance heatsink
Memory	32 GB (1x32 GB, 4800 MT/s) **Notes:** Runs at 4400 MT/s due to processor limitation.	P58417-B21: 64 GB (2x32 GB, 4800 MT/s) **Notes:** Runs at 4400 MT/s due to processor limitation.	32 GB (1x32 GB, 4800 MT/s) **Notes:** Runs at 4000 MT/s due to processor limitation.
		P58417-291 & P58417-421: 32 GB (2x32 GB, 4800 MT/s) **Notes:** Runs at 4400 MT/s due to processor limitation.	
Network Controller	Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T Adapter for HPE **Notes:** No embedded networking	Broadcom BCM57414 Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE **Notes:** No embedded networking	Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T Adapter for HPE **Notes:** No embedded networking
Storage Controller	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller	Embedded SATA controller (AHCI or Intel SATA software RAID controller)	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller
Included Hard Drives	None ship standard, 8 SFF supported		
Optical Drive	Optional DVD-ROM Optional via Universal Media Bay External support only		
Expansion Slots	3-slots (x8, x16, x8) as standard; upgradeable to 8-slots in a 2 processor configuration		

Pre-Configured Models

Power Supply	P52564-B21 & P52564-291: 1x HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit	P60636-B21 & P60636-291: 1x HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
	P52564-421: 1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit		P60636-421: 1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit
Fans	4x Standard Fans	6x High Performance Fan	4x Standard Fans
Management	HPE iLO 6		
Rail Kit	HPE ProLiant DL3XX Gen11 Easy Install Rail 3 Kit		
Security	TPM (Trusted Platform Module)		
Form Factor	2U Rack		
Warranty	Server warranty includes 3-year parts, 3-year labor, 3-year onsite support with next business day response.		

Pre-Configured Models

SKU Number	P60637-B21 P60637-291 P60637-421	P60638-B21 P60638-291 P60638-421	P60740-AA1
Model Name	P60637-B21 & P60637-291: HPE ProLiant DL380 Gen11 6426Y 2.5GHz 16-core 1P 32GB-R MR408i-o NC 8SFF 800W PS Server	P60638-B21 & P60638-291: HPE ProLiant DL380 Gen11 5418Y 2.0GHz 24-core 1P 64GB-R MR408i-o NC 8SFF 800W PS Server	HPE ProLiant DL380 Gen11 4416+ 2.1GHz 20-core 1P 32GB-R MR408i-o NC 8SFF 800W PS Server
	P60637-421: HPE ProLiant DL380 Gen11 6426Y 2.5GHz 16-core 1P 32GB-R MR408i-o NC 8SFF 1000W PS Server	P60638-421: HPE ProLiant DL380 Gen11 5418Y 2.0GHz 24-core 1P 64GB-R MR408i-o NC 8SFF 1000W PS Server	
Chassis	HPE ProLiant DL380 Gen11 8SFF NC Configure-to-order Server		
Processor	6426Y (16 core, 2.5 GHz, 185W)	5418Y (24 core, 2.0 GHz, 185W)	4416+ (20 core, 2.0 GHz, 165W)
Number of Processors	One with high performance heatsink		
Memory	32 GB (1x32 GB, 4800 MT/s)	64 GB (2x32 GB, 4800 MT/s) **Notes:** Runs at 4400 MT/s due to processor limitation.	32 GB (1x32 GB, 4800 MT/s) **Notes:** Runs at 4000 MT/s due to processor limitation.
Network Controller	Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T Adapter for HPE **Notes:** No embedded networking	Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T Adapter for HPE **Notes:** No embedded networking	Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE **Notes:** No embedded networking
Storage Controller	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller		
Included Hard Drives	None ship standard, 8 SFF supported		
Optical Drive	Optional DVD-ROM Optional via Universal Media Bay External support only		
Expansion Slots	3-slots (x8, x16, x8) as standard; upgradeable to 8-slots in a 2 processor configuration		
Power Supply	P60637-B21 & P60637-291: 1x HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P60638-B21 & P60638-291: 1x HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	1x HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
	P60637-421: 1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit	P60638-421: 1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit	
Fans	4x Standard Fans		
Management	HPE iLO 6		
Rail Kit	HPE ProLiant DL3XX Gen11 Easy Install Rail 3 Kit		

Pre-Configured Models

Security	TPM (Trusted Platform Module)	TPM (Trusted Platform Module)	TPM (Trusted Platform Module)
Form Factor	2U Rack		
Warranty	Server warranty includes 3-year parts, 3-year labor, 3-year onsite support with next business day response.		

Pre-Configured Models

SKU Number	P70456-291 P70456-421	P70457-291 P70457-421	P70458-421
Model Name	P70456-291: HPE ProLiant DL380 Gen11 4509Y 2.6GHz 8-core 1P 32GB-R MR408i-o NC 8SFF 1000W PS JP Server	P70457-291: HPE ProLiant DL380 Gen11 4514Y 2.0GHz 16-core 1P 32GB-R MR408i-o NC 8SFF 1000W PS JP Server	HPE ProLiant DL380 Gen11 4514Y 2.0GHz 16-core 1P 32GB-R MR416i-p NC 12LFF 1000W PS EMEA Server
	P70456-421: HPE ProLiant DL380 Gen11 4509Y 2.6GHz 8-core 1P 32GB-R MR408i-o NC 8SFF 1000W PS EMEA Server	P70457-421: HPE ProLiant DL380 Gen11 4514Y 2.0GHz 16-core 1P 32GB-R MR408i-o NC 8SFF 1000W PS EMEA Server	
Chassis	HPE ProLiant DL380 Gen11 8SFF NC Configure-to-order Server	HPE ProLiant DL380 Gen11 8SFF NC Configure-to-order Server	HPE ProLiant DL380 Gen11 12LFF NC Configure-to-order Server
Processor	4509Y (8 core, 2.6 GHz, 125W)	4514Y (16 core, 2.0 GHz, 150W)	4514Y (16 core, 2.0 GHz, 150W)
Number of Processors	One		
Memory	32 GB (1x32 GB, 5600 MT/s) **Notes:** Runs at 4400 MT/s due to processor limitation.		
Network Controller	Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE **Notes:** No embedded networking	Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T Adapter for HPE **Notes:** No embedded networking	Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T Adapter for HPE **Notes:** No embedded networking
Storage Controller	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller	HPE MR416i-p Gen11 x16 Lanes 8GB Cache PCI SPDM Plug-in Storage Controller
Included Hard Drives	None ship standard, 8 SFF supported	None ship standard, 8 SFF supported	None ship standard, 12 LFF supported
Optical Drive	Optional DVD-ROM Optional via Universal Media Bay External support only		
Expansion Slots	3-slots (x8, x16, x8) as standard; upgradeable to 8-slots in a 2 processor configuration		
Power Supply	1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit		
Fans	4x Standard Fans		
Management	HPE iLO 6		
Rail Kit	HPE ProLiant DL3XX Gen11 Easy Install Rail 3 Kit		
Security	TPM (Trusted Platform Module)		
Form Factor	2U Rack		
Warranty	Server warranty includes 3-year parts, 3-year labor, 3-year onsite support with next business day response.		

Pre-Configured Models

SKU Number	P70579-D61	P70580-D61
Model Name	HPE ProLiant DL380 Gen11 6430 2.1GHz 32-core 1P 32GB-R NC 8SFF 800W PS Server	HPE ProLiant DL380 Gen11 6442Y 2.6GHz 24-core 1P 32GB-R NC 8SFF 800W PS Server
Chassis	HPE ProLiant DL380 Gen11 8SFF NC Configure-to-order Server	
Processor	6430 (32 core, 2.1 GHz, 270W)	6442Y (24 core, 2.6 GHz, 225W)
Number of Processors	One with high performance heatsink	
Memory	32 GB (1x32 GB, 4800 MT/s) **Notes:** Runs at 4400 MT/s due to processor limitation.	32 GB (1x32 GB, 4800 MT/s)
Network Controller	Broadcom BCM57414 Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE **Notes:** No embedded networking	
Storage Controller	Embedded SATA controller (AHCI or Intel SATA software RAID controller)	
Included Hard Drives	None ship standard, 8 SFF supported	
Optical Drive	Optional DVD-ROM Optional via Universal Media Bay External support only	
Expansion Slots	3-slots (x8, x16, x8) as standard; upgradeable to 8-slots in a 2 processor configuration	
Power Supply	2x HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	
Fans	6x High Performance Fan	
Management	HPE iLO 6	
Rail Kit	HPE ProLiant DL3XX Gen11 Easy Install Rail 3 Kit	
Security	TPM (Trusted Platform Module)	
Form Factor	2U Rack	
Warranty	Server warranty includes 3-year parts, 3-year labor, 3-year onsite support with next business day response.	

Pre-Configured Models

Country Code Key

- -B21 = Worldwide
- -291 = Japan
- -421 = Europe, the Middle East and Africa
- -AA1 = China
- -D61 = India

HPE Smart Choice purchase program

The HPE Smart Choice purchase program features popular fully configured products that can be quoted in minutes and shipped quickly through HPE Authorized Partners. Products are configured and tested in an HPE factory and stocked at HPE Authorized Distributors and Partners. The products arrive in a single box, making onsite integration easier and more efficient for partners and customers. Additionally, there are aggressively priced HPE Tech Care Services available only through the HPE Smart Choice program when you purchase an HPE Smart Choice product.

For additional information on the HPE Smart Choice purchase program, please visit:
https://www.hpe.com/psnow/doc/a50009219enw

Smart Templates from HPE

HPE is releasing new Smart Template technology in the One Config Advanced (OCA) configurator. These Templates represent the CTO equivalents of the top-selling BTO configurations. They are intended to provide simple starting points to assist you in easily creating and customizing your desired Server solutions. HPE Servers that have Platform Templates, developed by HPE Product Managers, will have a separate tab in the HPE OCA configurator.

Workload Solutions Templates from HPE

The Workload Solutions Templates build on the Smart Templates technology to easily develop working configurations of the most compelling Workload Solutions. The templates complement the Reference Builds developed by HPE. Workload Solutions templates preconfigure some of the key architecture decisions and make it easier for Sellers to get started and complete a differentiated server solution for your customer's specific workload.

Mainstream SKUs

HPE launched the Mainstream SKU initiative as a market-driven approach to Demand Steering. It is a simplified portfolio of our top selling options that meet the current and future market trends. HPE has committed to provide a more predictable and faster experience for these options. Mainstream SKUs enjoy higher safety stock levels and have higher fulfilment service levels than non-Mainstream SKUs. Mainstream orders are fulfilled +30% faster than non-Mainstream orders, have fewer shortages and better recovery dates. This platform has Mainstream SKUs in the options portfolio, and is eligible for the improved Mainstream experience. Mainstream SKUs are designated with a Mainstream symbol in our configurators.

Mainstream Configurations

HPE is using the new Smart Templates technology to present Mainstream configurations. All the options in a Mainstream configuration are pre-selected Mainstream SKUs to optimize the performance, predictability and fulfilment experience. Check the Template section in our configurators for eligible Mainstream configurations.

This section lists some of the steps required to configure a Factory Integrated Model. To ensure only valid configurations are ordered, Hewlett Packard Enterprise recommends the use of an HPE approved configurator. Contact your local sales representative for information on configurable product offerings and requirements.

- Factory Integrated Models must start with a CTO Server.
- FIO indicates that this option is only available as a factory installable option.
- All Factory Integrated Models will be populated with sufficient hard drive blanks based on the number of initial hard drives ordered with the server.
- Some options may not be integrated at the factory. Contact your local sales representative for additional information.

Step 1: Base Configuration (choose one (1) of the following four (4) configurable server models from the tables below)

The below (4) CTO server models, denoted with "NC" in the SKU description, provide flexibility in the networking choice and require a network adapter from the "HPE Networking" section be selected.

Configuration Information

Networking Choice CTO Server Models	HPE ProLiant DL380 Gen11 8LFF NC CTO Server	HPE ProLiant DL380 Gen11 12LFF NC CTO Server	HPE ProLiant DL380 Gen11 8SFF NC CTO Server	HPE ProLiant DL380 Gen11 24SFF NC CTO Server
SKU Number	P52532-B21	P52533-B21	P52534-B21	P52535-B21
TAA SKU*	P52532-B21#GTA	P52533-B21#GTA	P52534-B21#GTA	P52535-B21#GTA
HPE Trusted Supply Chain	P36394-B21 - Optional			
Processor	Not included as standard	Not included as standard	Not included as standard	Not included as standard
DIMM Slots	32-DIMM slots	32-DIMM slots	32-DIMM slots	32-DIMM slots
Storage Controller	Embedded SW RAID with 14 SATA ports (12-ports accessible), optional choice of HPE modular MegaRAID controller.			
PCIe	Three standard in primary riser			
Drive Cage - available	8LFF	12LFF	8 SFF (8SFF U.3 x1)	24SFF (3x 8SFF U.3 x1)
Network Controller	Choice of either OCP 3.0 or select stand-up network adapters for primary networking selection plus additional/optional stand-up network adapters. **Notes:** No embedded networking			
Fans	4-Standard	4-Standard	4-Standard	6-High Performance
Management	HPE iLO with Intelligent Provisioning (standard), iLO Advanced and OneView (optional)			
USB	4x 3.0 standard plus iLO front service port	4x 3.0 standard plus iLO front service port	4x 3.0 standard plus iLO front service port	4x 3.0 standard plus iLO front service port
Networking Choice CTO Server Models	HPE ProLiant DL380 Gen11 12EDSFF NC Configure-to-order Server			
SKU Number	P52536-B21			
TAA SKU*	P52536-B21#GTA			
HPE Trusted Supply Chain	P36394-B21 - Optional			
Processor	Not included as standard			
DIMM Slots	32-DIMM slots			
Storage Controller	Embedded SW RAID with 14 SATA ports (12-ports accessible), and optional choice of HPE modular MegaRAID controller.			
PCIe	Three standard in primary riser			
Drive Cage - available	12 EDSFF			
Network Controller	Choice of either OCP 3.0 or select stand-up network adapters for primary networking selection plus additional/optional stand-up network adapters **Notes:** No embedded networking			
Fans	6 -Standard			
Management	HPE iLO with Intelligent Provisioning (standard), iLO Advanced and OneView (optional)			
USB	4x 3.0 standard plus iLO front service port			

Notes:

- Network Choice (NC) server models require a networking selection of a network adapters in the "HPE Networking" section.
- HPE Trusted Supply Chain (P36394-B21) is an optional security upgrade intended for agencies and regulated industries needing enhanced security and compliance needs. Applying this option to a DL380 Gen11 CTO server ensures it is built in the USA in a secured facility by vetted HPE personnel assigned to the manufacturing processes. A multitude of checkpoints/inspections for malicious microcode and counterfeit parts are performed throughout the server build, and additional safeguards are put in place against cyber-exploits throughout the server lifecycle. The HPE ProLiant DL380 Gen11 is re-branded as a HPE ProLiant DL380T Gen11 to denote the HPE Trusted Supply Chain security enhancements. The DL380T is Trade Agreement Act (TAA) compliant. See "HPE Security" section within this document for more detail and learn more at **https://www.hpe.com/security**
- *HPE offers multiple Trade Agreement Act (TAA) compliant configurations to meet the needs of US Federal Government customers. These products are either manufactured or substantially transformed in a designated country. TAA compliance is only provided when HPE options are included as part of factory integrated orders (CTO).
- All CTO servers are Energy Star 3.0 compliant.

CTO Server	8 SFF CTO Chassis	24SFF CTO Chassis	8LFF CTO Chassis	12LFF CTO Chassis
Available Drive Cage	8SFF U.3 x1 Drive Cage	3x 8SFF U.3 x1 Drive Cage	8LFF (2x 4LFF Drive Cages)	12LFF (3x 4LFF Drive Cages)
Universal Media Bay	1 Optional	Not available	Not available	Not available
Optical Disk Drive	1 Optional with UMB	Not available	1 Optional with ODD Enable Kit	Not available
8 SFF NVMe/SAS/SATA Drive Cage (mid-tray)	1 Optional	1 Optional	1 Optional	1 Optional
2 SFF NVMe/SAS/SATA (Front)	1 Optional with UMB	Not available	1 Optional with Side-by-Side Drive Cage	Not available
2 SFF NVMe/SAS/SATA (Rear)	3 Optional	3 Optional	1 Optional	1 Optional
2LFF SAS/SATA (Rear riser)	Not available	Not available	2 Optional	2 Optional
CTO Server	**EDSFF CTO Chassis**			
Available Drive Cage	12 EDSFF Drive Cage			
Universal Media Bay	Not supported			
Optical Disk Drive	Not supported			
8 SFF NVMe/SAS/SATA Drive Cage	Up to 1 Optional			
2 SFF NVMe/SAS/SATA (Front)	Not supported			

Configuration Information

2 SFF NVMe/SAS/SATA (Rear)	1 Optional			
2LFF SAS/SATA (Rear riser)	Not available			

Notes: This applies to CTO configurations; field upgrades may differ depending on field configuration.

Step 2: Choose Required Options

Please select up to two processors required below.

Notes:

- 8SFF, 8LFF, and 12LFF CTO models ship with 4 standard fans.
- 24SFF CTO Servers ship with 6 High performance fans included. Maximum Performance fan kit is available to meet ambient temperature environments and are required for rear drives or NVMe configurations.
- Maximum memory capacity per processor is dependent on processor models. All processors support up to 4TB max memory per processor.
- Mixing of 2 different processor models are NOT allowed.
- Processors with TDP greater than 150W require High Performance Heatsink (P48818-B21).
- Q series processors require Max Performance Heat Sink (P48817-B21)
- DDR5 speed is the maximum memory speed of the processor. Actual memory speed may depend on the quantity and type of DIMMs installed.

Step 2a: Choose Processors

Processor Option Kits (Required Processor)

5th Generation Intel Xeon-Platinum

All SKUs below ship with processors only. Adequate fans and heatsinks must be selected.

Supports "HPE DDR5 Smart Memory - Registered (RDIMM), 5600MT/s".

Intel Xeon-Platinum 8593Q 2.2GHz 64-core 385W Processor for HPE P68449-B21

Notes:

- XCC die
- Max 2
- Qty 2 of processor must be selected.
- Max Performance Heat Sink or DLC Module must be selected
- High Performance Fan Kit must be selected and defaulted
- Max Performance Heat Sink is defaulted however customer is allowed to deselect and select DLC Module.
- Without DLC Module, drives above 11W cannot be populated in Rear Drive cages
- If Full Length (FL) GPU is selected with this Processor then DLC Module must be selected for Processors.
- Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Platinum 8592+ 1.9GHz 64-core 350W Processor for HPE P67089-B21

Notes:

- XCC die

Configuration Information

– Max 2
– High Performance Heat Sink or DLC Module must be selected.
– High Performance Fan Kit must be selected and defaulted
– Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Platinum 8592V 2.0GHz 64-core 330W Processor for HPE P67107-B21

Notes:

– XCC die
– Max 2
– High Performance Heat Sink or DLC Module must be selected.
– High Performance Fan Kit must be selected and defaulted
– Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Platinum 8581V 2.0GHz 60-core 270W Processor for HPE P67109-B21

Notes:

– XCC die
– Max 1
– High Performance Heat Sink must be selected.
– DLC Module cannot be selected.
– High Performance Fan Kit must be selected
– Cannot be selected with EDSFF CTO Server.
– Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Platinum 8580 2.0GHz 60-core 350W Processor for HPE P67088-B21

Notes:

– XCC die
– Max 2
– High Performance Heat Sink or DLC Module must be selected.
– High Performance Fan Kit must be selected
– Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Platinum 8570 2.1GHz 56-core 350W Processor for HPE P67087-B21

Notes:

– XCC die
– Max 2
– High Performance Heat Sink or DLC Module must be selected.
– High Performance Fan Kit must be selected
– Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Platinum 8568Y+ 2.3GHz 48-core 350W Processor for HPE P67086-B21

Notes:

– XCC die
– Max 2
– High Performance Heat Sink or DLC Module must be selected.
– High Performance Fan Kit must be selected
– Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Platinum 8562Y+ 2.8GHz 32-core 300W Processor for HPE P67085-B21

Notes:

Configuration Information

– MCC die
– Max 2
– High Performance Heat Sink or DLC Module must be selected.
– High Performance Fan Kit must be selected
– Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Platinum 8558P 2.7GHz 48-core 350W Processor for HPE P67108-B21

Notes:

– XCC die
– Max 2
– High Performance Heat Sink or DLC Module must be selected.
– High Performance Fan Kit must be selected
– Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Platinum 8558 2.1GHz 48-core 330W Processor for HPE P67097-B21

Notes:

– XCC die
– Max 2
– High Performance Heat Sink or DLC Module must be selected.
– High Performance Fan Kit must be selected
– Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Platinum 8558U 2.0GHz 48-core 300W Processor for HPE P67102-B21

Notes:

– XCC die
– Max 1
– High Performance Heat Sink or DLC Module must be selected.
– High Performance Fan Kit must be selected
– Single socket processor and DLC Module cannot be selected together.
– Cannot be selected with EDSFF CTO Server.
– Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Gold 6558Q 3.2GHz 32-core 350W Processor for HPE P67098-B21

Notes:

– MCC die
– Qty 2 of processor must be selected.
– Max Performance Heat Sink or DLC Module must be selected.
– High Performance Fan Kit must be selected
– Supports only PC5-5600B/ PC5-5200B Memory.
– If Full Length (FL) GPU is selected with this Processor then DLC Module must be selected for Processors.

Intel Xeon-Gold 6554S 2.2GHz 36-core 270W Processor for HPE P67110-B21

Notes:

– XCC die
– Max 2
– High Performance Heat Sink or DLC Module must be selected.
– High Performance Fan Kit must be selected.
– Supports only PC5-5600B/ PC5-5200B Memory.

Configuration Information

Intel Xeon-Gold 6548N 2.8GHz 32-core 250W Processor for HPE P67105-B21

Notes:

- MCC die
- Max 2
- High Performance Heat Sink or DLC Module must be selected.
- High Performance Fan Kit must be selected.
- Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Gold 6548Y+ 2.5GHz 32-core 250W Processor for HPE P67082-B21

Notes:

- MCC die
- Max 2
- High Performance Heat Sink or DLC Module must be selected.
- High Performance Fan Kit must be selected.
- Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Gold 6544Y 3.6GHz 16-core 270W Processor for HPE P67084-B21

Notes:

- MCC die
- Max 2
- High Performance Heat Sink or DLC Module must be selected.
- High Performance Fan Kit must be selected.
- Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Gold 6542Y 2.9GHz 24-core 250W Processor for HPE P67081-B21

Notes:

- MCC die
- Max 2
- High Performance Heat Sink or DLC Module must be selected.
- High Performance Fan Kit must be selected.
- Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Gold 6538N 2.1GHz 32-core 205W Processor for HPE P67104-B21

Notes:

- MCC die
- Max 2
- High Performance Heat Sink or DLC Module must be selected.
- Standard Fan kit or High Performance Fan Kit must be selected.
- Standard Fan kit or High Performance Fan Kit must be selected.
- Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Gold 6538Y+ 2.2GHz 32-core 225W Processor for HPE P67096-B21

Notes:

- MCC die
- Max 2
- High Performance Heat Sink or DLC Module must be selected.
- High Performance Fan Kit must be selected.
- Supports only PC5-5600B/ PC5-5200B Memory.

Configuration Information

Intel Xeon-Gold 6534 3.9GHz 8-core 195W Processor for HPE P67083-B21

Notes:

- MCC die
- Max 2
- High Performance Heat Sink or DLC Module must be selected.
- Standard Fan kit or High Performance Fan Kit must be selected.
- Standard Fan kit or High Performance Fan Kit must be selected.
- Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Gold 6530 2.1GHz 32-core 270W Processor for HPE P67095-B21

Notes:

- XCC die
- Max 2
- High Performance Heat Sink or DLC Module must be selected.
- High Performance Fan Kit must be selected.
- Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Gold 6526Y 2.8GHz 16-core 195W Processor for HPE P67080-B21

Notes:

- MCC die
- Max 2
- High Performance Heat Sink or DLC Module must be selected.
- Standard Fan kit or High Performance Fan Kit must be selected.
- Standard Fan kit or High Performance Fan Kit must be selected.
- Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Gold 5520+ 2.2GHz 28-core 205W Processor for HPE P67094-B21

Notes:

- MCC die
- Max 2
- High Performance Heat Sink or DLC Module must be selected.
- Standard Fan kit or High Performance Fan Kit must be selected.
- Standard Fan kit or High Performance Fan Kit must be selected.
- Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Gold 5515+ 3.2GHz 8-core 165W Processor for HPE P67079-B21

Notes:

- MCC die
- Max 2
- High Performance Heat Sink or DLC Module must be selected.
- Standard Fan kit or High Performance Fan Kit must be selected.
- Standard Fan kit or High Performance Fan Kit must be selected.
- Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Silver 4516Y+ 2.2GHz 24-core 185W Processor for HPE P67093-B21

Notes:

- MCC die
- Max 2
- High Performance Heat Sink or DLC Module must be selected.

Configuration Information

– Standard Fan kit or High Performance Fan Kit must be selected.
– Standard Fan kit or High Performance Fan Kit must be selected.
– Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Silver 4514Y 2.0GHz 16-core 150W Processor for HPE P67092-B21

Notes:

– MCC die
– Max 2
– Standard Heat Sink or High Performance Heat Sink or DLC Module must be selected.
– Standard Fan kit or High Performance Fan Kit must be selected.
– Supports only PC5-5600B/ PC5-5200B Memory.

Intel Xeon-Silver 4510 2.4GHz 12-core 150W Processor for HPE P67091-B21

Notes:

– EE LCC die
– Max 2
– Standard Heat Sink or High Performance Heat Sink or DLC Module must be selected.
– Standard Fan kit or High Performance Fan Kit must be selected.
– Supports only PC5-5600B/ PC5-5200B Memory.
– 96GB 5600B Memory cannot be selected.

Intel Xeon-Silver 4509Y 2.6GHz 8-core 125W Processor for HPE P67090-B21

Notes:

– EE LCC die
– Max 2
– Standard Heat Sink or High Performance Heat Sink or DLC Module must be selected.
– Standard Fan kit or High Performance Fan Kit must be selected.
– Supports only PC5-5600B/ PC5-5200B Memory.
– If "HPE Energy Star4.0 configuration FIO kit" is selected then Qty2 of this Processor cannot be selected.
– 96GB 5600B Memory cannot be selected.

Intel Xeon-Bronze 3508U 2.1GHz 8-core 125W Processor for HPE P67100-B21

Notes: 18, 282, 283, 284, 285, 288, 309, 319, 320, 322, 323, 430, 436, 499, 500, 501, 510

– EE LCC die
– Max 1
– Standard Heat Sink or High Performance Heat Sink or DLC Module must be selected.
– Standard Fan kit or High Performance Fan Kit must be selected.
– Supports only PC5-5600B/ PC5-5200B Memory.
– If "HPE Energy Star4.0 configuration FIO kit" is selected then Qty2 of this Processor cannot be selected.
– 96GB 5600B Memory cannot be selected.
– Supports PCIe4.0 Only29, 30, 31, 32, 34, 259, 445, 449, 458, 475, 499

4th Generation Intel Xeon-Platinum

Notes:

– All SKUs below ship with processor only. Adequate fans and heatsinks must be selected.
– Processors with TDP equal to or greater than 150W require High Performance Heatsink (P48818-B21)

Configuration Information

– Q series processors require Max Performance Heat Sink (P48817-B21)
– 8470Q processor is not supported with 12LFF CTO Server and 24SFF CTO Server.

Intel Xeon-Platinum 9462 2.7GHz 32-core 350W Processor for HPE P49645-B21

Notes:

– This is Intel High Bandwidth Memory (HBM) CPU.
– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Platinum 8490H 1.9GHz 60-core 350W Processor for HPE P49630-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Platinum 8480+ 2.0GHz 56-core 350W Processor for HPE P49607-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Platinum 8470Q 2.1GHz 52-core 350W Processor for HPE P49609-B21

Notes:

– Requires selection of Max Performance Heat Sink or DLC Module
– Requires High Performance Fan Kit
– If Full Length (FL) GPU is selected with this Processor then DLC Module must be selected for Processors.

Intel Xeon-Platinum 8470 2.0GHz 52-core 350W Processor for HPE P49606-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Platinum 8470N 1.7GHz 52-core 300W Processor for HPE P49649-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Platinum 8468V 2.4GHz 48-core 330W Processor for HPE P49631-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Platinum 8468 2.1GHz 48-core 350W Processor for HPE P49605-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Platinum 8462Y+ 2.8GHz 32-core 300W Processor for HPE P49603-B21

Configuration Information

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Platinum 8458P 2.7GHz 44-core 350W Processor for HPE P49632-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Platinum 8460Y+ 2.0GHz 40-core 300W Processor for HPE P49604-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Platinum 8452Y 2.0GHz 36-core 300W Processor for HPE P49616-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Platinum 8444H 2.9GHz 16-core 270W Processor for HPE P49625-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

4th Generation Intel Xeon-Gold

Intel Xeon-Gold 6458Q 3.1GHz 32-core 350W Processor for HPE P49608-B21

Notes:

– Requires selection of Max Performance Heat Sink or DLC Module
– Requires High Performance Fan Kit
– If Full Length (FL) GPU is selected with this Processor then DLC Module must be selected for Processors.

Intel Xeon-Gold 6448H 2.4GHz 32-core 250W Processor for HPE P49622-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Gold 6448Y 2.1GHz 32-core 225W Processor for HPE P49600-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Gold 6444Y 3.6GHz 16-core 270W Processor for HPE P49602-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Gold 6442Y 2.6GHz 24-core 225W Processor for HPE P49599-B21

Configuration Information

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Gold 6438N 2.0GHz 32-core 205W Processor for HPE P49638-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Gold 6438Y+ 2.0GHz 32-core 205W Processor for HPE P49615-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Gold 6434 3.7GHz 8-core 195W Processor for HPE P49601-B21

Notes: Requires High Performance Heat Sink

Intel Xeon-Gold 6430 2.1GHz 32-core 270W Processor for HPE P49614-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Gold 6426Y 2.5GHz 16-core 185W Processor for HPE P49598-B21

Notes: Requires High Performance Heat Sink

Intel Xeon-Gold 6454S 2.2GHz 32-core 270W Processor for HPE P49654-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Gold 6421N 1.8GHz 32-core 185W Processor for HPE P49641-B21

Notes: Requires High Performance Heat Sink

Intel Xeon-Gold 6418H 2.1GHz 24-core 185W Processor for HPE P49621-B21

Notes: Requires High Performance Heat Sink

Intel Xeon-Gold 6416H 2.2GHz 18-core 165W Processor for HPE P49620-B21

Notes: Requires High Performance Heat Sink

Intel Xeon-Gold 6414U 2.0GHz 32-core 250W Processor for HPE P49619-B21

Notes:

– This is a single socket CPU, max allowed = 1
– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Gold 5415+ 2.9GHz 8-core 150W Processor for HPE P49597-B21

Notes: Requires Standard Heat Sink

Intel Xeon-Gold 5416S 2.0GHz 16-core 150W Processor for HPE P49653-B21

Notes: Requires Standard Heat Sink

Intel Xeon-Gold 5418N 1.8GHz 24-core 165W Processor for HPE P49640-B21

Notes: Requires High Performance Heat Sink

Intel Xeon-Gold 5418Y 2.0GHz 24-core 185W Processor for HPE P49612-B21

Configuration Information

Notes: Requires High Performance Heat Sink

Intel Xeon-Gold 5420+ 2.0GHz 28-core 205W Processor for HPE P49613-B21

Notes:

– Requires High Performance Heat Sink
– Requires High Performance Fan Kit

Intel Xeon-Gold 5415+ 2.9GHz 8-core 150W Processor for HPE P49597-B21

Notes:

– Requires Standard Heat Sink
– 4th Generation Intel Xeon-Silver

Intel Xeon-Gold 5411N 1.9GHz 24-core 165W Processor for HPE P49639-B21

Notes: Requires High Performance Heat Sink

Intel Xeon-Silver 4410Y 2.0GHz 12-core 150W Processor for HPE P49610-B21

Notes: Requires Standard Heatsink

Intel Xeon-Silver 4416+ 2.0GHz 20-core 165W Processor for HPE P49611-B21

Notes:

– Requires Standard Heatsink
– 4th Generation Intel Xeon-Bronze

Intel Xeon-Bronze 3408U 1.8GHz 8-core 125W Processor for HPE P49617-B21

Notes: Requires Standard Heatsink

Step 2b: Choose Memory Options

Please select one or more memory from below.

For new Gen11 memory population rule whitepaper and optimal memory performance guidelines, please go to:

HPE Memory Population Rules

For details on the HPE Server Memory Options Population Rules, please go to:

Memory population rules for HPE Gen11 servers with 4th Generation Intel Scalable Processors

Notes:

– HPE Server Memory compatibility for a specific server platform may vary or be limited within a server platform depending upon the specific configuration being requested. Because each server environment and requirements can vary, memory compatibility is based not only upon the server family, but may also be affected by the amount and type of additional hardware options installed within a specific server configuration. For this reason, some HPE memory DIMMs may be qualified for an HPE server model or family and yet occasionally not be supported with some configurations within that server family
– Memory should be installed in even quantity of DIMMs
– The maximum memory speed and capacity is a function of the memory type, memory configuration, and processor model.
– DDR5-4800 Memory Kits are only supported with 4th Generation Intel Xeon Scalable Series Processors.

Configuration Information

- Memory compatibility may vary or be limited within a specific server family depending upon the specific configuration being requested. Because each server environment and requirements can vary, memory compatibility is based not only upon the server family, but may also be affected by the amount and type of additional hardware options installed within a specific server configuration. For this reason, some HPE memory DIMMs may be qualified for a server model or family and yet occasionally not be supported with limited configurations within that server family.
- Please consult with the HPE server QuickSpecs or your HPE representative if you have any questions regarding memory compatibility with a specific HPE server configuration.

Registered DIMMs DDR5 (RDIMMs)

HPE 16GB (1x16GB) Single Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P43322-B21
HPE 32GB (1x32GB) Dual Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P43328-B21
HPE 64GB (1x64GB) Dual Rank x4 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P43331-B21
HPE 96GB (1x96GB) Dual Rank x4 DDR5-4800 CAS-46-45-45 EC8 Registered Smart Memory Kit	P66675-B21

Notes: 96GB DIMMs supported only on XCC and MCC CPUs.

HPE 128GB (1x128GB) Dual Rank x4 DDR5-5600 CAS-46-45-45 EC8 Registered Smart Memory Kit	P69976-B21
HPE 128GB (1x128GB) Quad Rank x4 DDR5-4800 CAS-46-39-39 EC8 Registered 3DS Smart Memory Kit	P43334-B21
HPE 256GB (1x256GB) Octal Rank x4 DDR5-4800 CAS-46-39-39 EC8 Registered 3DS Smart Memory Kit	P43337-B21

Notes:

- 4800 MT/s memory SKUs offer a transfer rate of 4800 MT/s at 1 DIMM per channel and 4400 MT/s at 2 DIMMs per channel
- Mixing of 3DS memory and non-3DS memory is not supported
- Mixing of x4 and x8 memory is not allowed
- Memory with larger than 128GB capacity will need High Performance Fan Kit (P48820-B21) and ambient limitation. 256GB DIMM will also need to limit the maximum front-end cage to two.
- If 96GB or higher density memory is selected then High Performance Fan Kit must be selected.
- 96GB memory cannot be mixed with any other memory.
- 96GB memory configuration must be either 8 or 16 only on 1 processor unit.
- 96GB memory configuration must be either 16 or 32only on 2 processor unit.

Memory Blank Kit

HPE DDR4 DIMM Blank Kit	P07818-B21

Notes: DIMM Blanks are optional and not required.

Step 2c: Choose Power Supplies

Select one or two power supplies from below.

Notes: Mixing of 2 different power supplies is NOT allowed.

HPE Flex Slot Power Supplies

HPE 1800W-2200W Flex Slot Titanium Hot Plug Power Supply Kit	P44712-B21

Configuration Information

HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P38997-B21
HPE 1600W Flex Slot -48VDC Hot Plug Power Supply Kit	P17023-B21
HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit	P03178-B21
HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P38995-B21

Notes:

- Select a minimum (1), maximum (2) power supplies.
- 1600W Power supplies only support high line voltage (200VAC to 240VAC).
- Prior to making a power supply selection it is highly recommended that the HPE Power Advisor is run to determine the right size power supply for your server configuration. The HPE Power Advisor is located at:

https://www.hpe.com/info/hppoweradvisor.

- All power supplies in a server should match.
- Mixing Power Supplies is not supported.
- HPE ProLiant servers ship with an IEC-IEC power cord used for rack mounting with Power Distribution Units (PDUs). Visit **HPE power cords** for a full list of optional power cords.
- Beginning on January 1st, 2024, units sold into the European Union (EU), European Economic Area (EEA), the United Kingdom, or Switzerland must include more efficient AC power supplies: 94% for multi-output and 96% for single-output.
- HPE Flexible Slot power supplies are single-output, and part numbers 865438-B21, P03178-B21, and P44712-B21 are 96% efficient, thus meeting requirements. HPE is on target to fulfil compliant systems ahead of time and will begin enforcing these requirements in advance to satisfy requests with the current power supplies by the set deadline.

Step 3: Choose Additional Factory Integrable Options

One of the following from each list may be selected if desired at time of factory integration

HPE Security Options

HPE iLO Common Password FIO Setting	P08040-B21

Notes:

- Replaces iLO default randomized password by an HPE defined common password. HPE highly recommends changing this password immediately after the initial onboarding process.
- Customers who want to choose their own custom iLO default password should use the HPE Factory Express Integration Services

HPE Bezel Lock Kit	875519-B21

Notes: This option can be selected only if HPE Gen11 2U Bezel Kit (P50400-B21) is selected.

HPE ProLiant DL3XX Gen11 Intrusion Cable Kit	P48922-B21

Notes: This option must be selected if HPE Trusted Supply Chain SKU (P36394-B21) is selected.

HPE Gen11 2U Bezel Kit	P50400-B21

Configuration Information

Factory Instructions and Server Settings

HPE ProLiant DL380 Gen11 8NVMe 1P FIO Bundle Kit P69548-B21

Notes:

- This bundle provides Direct Attach to 8NVMe Drives through First/ Single Processor
- If NVMe bundle is selected then defaulted Controller has to be deselected.
- If this NVMe Bundle is selected then Qty 1 of HPE DL380 Gen11 8NVMe CPU1/2 Cable Kit must be selected and defaulted.
- If this NVMe Bundle/ EDSFF Bundle is selected then Qty 1 of High Performance Fan Kit must be selected and defaulted.
- If 8NVMe Bundle is selected then Qty 1 (min/max) of 8SFF U.3 x4 cage OR 8SFF U.2 x4 cage must be selected.
- NVMe Bundle is supported with 8SFF CTO Server Only.
- If 8SFF U.3 x1 Cage is selected along with 8 OR 16 NVMe bundle then controller must be selected except SR308i-o/SR308i-p.
- If 8NVMe 1P Bundle is selected then Primary 3 x16 Cable Kit cannot be selected and vice-versa.

HPE ProLiant DL380 Gen11 36EDSFF FIO Bundle Kit P56075-B21

HPE ProLiant DL380 Gen11 20EDSFF FIO Bundle Kit P56076-B21

Notes:

- If this NVMe Bundle/ EDSFF Bundle is selected then Qty 1 of High Performance Fan Kit must be selected.
- If NVMe Bundle/ EDSFF Bundle is selected then Second Processor must be selected.
- For 12EDSFF CTO Server, If EDSFF Bundle is selected then Max Quantity of 256GB Memory is limited to 16.
- EDSFF Drive cage supports Direct attach Only and requires selection of EDSFF Bundle.
- If EDSFF Bundle is selected then Qty1 of HPE DL380 Gen11 12EDSFF CPU1/2 Cable Kit must be selected
- If EDSFF Bundle is selected then Qty2 of HPE DL380/DL560 G11 2U 12EDSFF NVMe Kit must be selected.
- EDSFF Bundle is supported with 12EDSFF CTO Server Only.
- If EDSFF Bundle is selected then OCP1 x16 Enablement Kit AND OCP2 x16 Enablement Kit AND CPU1 OCP2 x8 Enable kit cannot be selected.
- If EDSFF Bundle is selected then Tertiary Riser cannot be selected.
- If EDSFF Bundle is selected then HPE DL380 Gen11 x16/x16/x16 Prim Cable Kit AND HPE DL380 Gen11 x16/x16/x16 Sec Cable Kit cannot be selected.
- If 20EDSFF Bundle is selected then Maximum 20 EDSFF Drives only can be selected.

HPE ProLiant DL380 Gen11 32NVMe Balanced FIO Bundle Kit P53639-B21

Notes:

- If NVMe bundle is selected then default storage controller cannot be selected.
- If this NVMe Bundle is selected then Qty 1 of HPE DL380 Gen11 8NVMe CPU1/2 Cable Kit must be selected.
- If this NVMe Bundle/ EDSFF Bundle is selected then Qty 1 of High Performance Fan Kit must be selected.
- If 24NVMe Bundle/ 32NVMe Bundle is selected then Qty 3 (min/max) of 8SFF U.3 x4 cage must be selected.
- If 24NVMe Bundle/ 32NVMe Bundle is selected then Qty 1 of HPE DL380 Gen11 3x16 Primary riser must be selected.
- If 24NVMe Bundle/ 32NVMe Bundle is selected then Qty 1 of HPE DL380 Gen11 3x16 SEC Riser must be selected.

Configuration Information

- If 24NVMe Bundle/ 32NVMe Bundle is selected then Qty 1 of OCP1 x16 Enablement Kit must be selected.
- If 24NVMe Bundle/ 32NVMe Bundle is selected then Qty 1 of OCP2 x16 Enablement Kit must be selected.
- If 24NVMe Bundle/ 32NVMe Bundle is selected then Qty 1 of HPE Gen4 Re-timer/-p Cable Kit must be selected.
- NVMe Bundle is supported with 8SFF CTO Server Only.
- If NVMe Bundle/ EDSFF Bundle is selected, then Second Processor must be selected.
- If 24NVMe/ 32NVMe Bundle is selected, then Tertiary Riser cannot be selected.
- For 8SFF CTO Server, 8SFF U.3 x4 Mid Tray requires selection of 32NVMe Bundle or SR932i controller.
- If 32NVMe Bundle is selected, then Qty 1 (min/max) of HPE DL380 G11 2U 8SFF x4 U.3 Mid TM Kit must be selected.
- If 32NVMe Bundle is selected, then Qty 4 of HPE DL385 Gen10+ 12Gb NVMe 2p Adapter must be selected.

HPE ProLiant DL380 Gen11 8NVMe Balanced FIO Bundle Kit P53633-B21

Notes:

- If this NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 8NVMe CPU1/2 Cable Kit (P48825-B21) must be selected and defaulted.
- If this NVMe Bundle is selected, then Qty 1 of High Performance Fan Kit (P48820-B21) must be selected and defaulted.
- If 8NVMe Bundle is selected, then Qty 1 (min) of 8SFF U.3 x4 drive must be selected.
- If NVMe Bundle is selected, then 8SFF U.3 x4 Cage is defaulted.
- NVMe Bundle is supported with 8SFF CTO Server Only.
- If NVMe Bundle is selected, then Second Processor must be selected.
- If 8SFF U.3 x1 Cage is selected along with 8 OR 16 NVMe bundle, then controller must be selected.

HPE ProLiant DL380 Gen11 16NVMe Balanced FIO Bundle Kit P53634-B21

Notes:

- If NVMe bundle is selected, then defaulted Controller has to be deselected.
- If this NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 8NVMe CPU1/2 Cable Kit
(P48825-B21) must be selected and defaulted.
- If this NVMe Bundle is selected, then Qty 1 of High Performance Fan Kit (P48820-B21) must be selected and defaulted.
- If 16NVMe Bundle is selected, then Qty 2 (min) of 8SFF U.3 x4 drive cage must be selected.
- If NVMe Bundle is selected, then 8SFF U.3 x4 drive cage will be defaulted.
- NVMe Bundle is supported with 8SFF CTO Server Only.
- If 16NVMe/ 24NVMe Bundle is selected, then Universal Media Bay kit cannot be selected.
- If NVMe Bundle is selected, then Second Processor must be selected.
- If Tertiary Riser is selected along with 16NVMe Bundle, then 2x16 Tertiary FIO x8 Enable Kit (P53632-B21) must be selected
- If 8SFF U.3 x1 drive cage (P48813-B21) is selected along with 8 OR 16 NVMe bundle, then controller must be selected.

HPE ProLiant DL380 Gen11 24NVME Balanced I/O FIO Bundle Kit P53635-B21

Notes:

- If NVMe bundle is selected, then defaulted Controller has to be deselected.
- If NVMe Bundle is selected, then 8SFF U.3 x1 drive cage (P48813-B21) cannot be

Configuration Information

selected.
– If this NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 8NVMe CPU1/2 Cable Kit
(P48825-B21) must be selected and defaulted.
– If this NVMe Bundle is selected, then Qty 1 of High Performance Fan Kit (P48820-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 3 (min) of 8SFF U.3 x4 drives must be selected.
– If 24NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 3x16 Primary Riser (P48803-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 3x16 Secondary Riser (P51083-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 1 of OCP1 x16 Enablement Kit (P48827-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 1 of OCP2 x16 Enablement Kit (P48828-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 2 of HPE DL385 Gen10+ 12Gb NVMe 2p Adapter (P25527-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 1 of HPE Gen4 Re-timer/-p Cable Kit must be selected and defaulted.
– If NVMe Bundle is selected, then 8SFF U.3 x4 drive cage will be defaulted.
– NVMe Bundle is supported with 8SFF CTO Server Only.
– If 16NVMe/ 24NVMe Bundle is selected, then Universal Media Bay kit cannot be selected.
– If NVMe Bundle is selected, then Second Processor must be selected.
– If 24NVMe Bundle is selected, then Tertiary Riser cannot be selected.
– If 24NVMe Balanced or 24NVMe DIFF IO-3 Bundle is selected, then CPU1 OCP2 x8 enablement OR CPU2 OCP2 x8 enablement kit cannot be selected.
– If 24NVMe Bundle is selected, then Primary 3 x16 Cable OR Secondary 3 x16 Cable cannot be selected.

HPE ProLiant DL380 Gen11 24NVME Differential I/O OCP1/2 x8 FIO Bundle Kit P53636-B21

Notes:

– If NVMe bundle is selected, then defaulted Controller has to be deselected.
– If NVMe Bundle is selected, then 8SFF U.3 x1 drive cage (P48813-B21) cannot be selected.
– If this NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 8NVMe CPU1/2 Cable Kit
(P48825-B21) must be selected and defaulted.
– If this NVMe Bundle is selected, then Qty 1 of High Performance Fan Kit (P48820-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 3 (min) of 8SFF U.3 x4 drive cage must be selected.
– If 24NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 3x16 Primary Riser (P48803-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected then Qty 1 of HPE DL380 Gen11 x8x16x8 Secondary Riser (P48802-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 1 of HPE DL300 Gen11 CPU1 OCP2 x8 Enable kit (P51911-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 1 of HPE DL385 Gen10+ 12Gb NVMe 2p Adapter (P25527-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 1 of HPE Gen4 Re-timer/-p Cable Kit (P54874-B21) must be selected and defaulted.
– If NVMe Bundle is selected, then 8SFF U.3 x4 drive cage will be defaulted.
– NVMe Bundle is supported with 8SFF CTO Server Only.

Configuration Information

- If 16NVMe/ 24NVMe Bundle is selected, then Universal Media Bay kit cannot be selected.
- If NVMe Bundle is selected, then Second Processor must be selected.
- If Tertiary Riser is selected along with 16NVMe Bundle then 2x16 Tertiary FIO x8 Enable Kit (P53632-B21) must be selected
- If 24NVMe Diff IO-1 bundle is selected, then OCP1 x16 enablement Kit and OCP2 x16 enablement kit and CPU2 OCP2 x8 enablement Kit cannot be selected.
- If 24NVMe Bundle is selected then primary 3 x16 Cable OR Sec 3 x16 Cable cannot be selected.

HPE ProLiant DL380 Gen11 24NVME Differential I/O OCP1 x16 FIO Bundle Kit P53637-B21

Notes:

- If NVMe bundle is selected, then defaulted Controller has to be deselected.
- If NVMe Bundle is selected, then 8SFF U.3 x1 drive cage (P48813-B21) cannot be selected.
- If this NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 8NVMe CPU1/2 Cable Kit (P48825-B21) must be selected and defaulted.
- If this NVMe Bundle is selected, then Qty 1 of High-Performance Fan Kit (P48820-B21) must be selected and defaulted.
- If 24NVMe Bundle is selected, then Qty 3 (min) of 8SFF U.3 x4 drive cage must be selected.
- If 24NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 3x16 Primary Riser (P48803-B21) must be selected and defaulted.
- If 24NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 x8x16x8 Secondary Riser (P48802-B21) must be selected and defaulted.
- If 24NVMe Bundle is selected, then Qty 1 of OCP1 x16 Enablement Kit (P48827-B21) must be selected and defaulted.
- If 24NVMe Bundle is selected, then Qty 1 of HPE DL385 Gen10+ 12Gb NVMe 2p Adapter (P25527-B21) must be selected and defaulted.
- If 24NVMe Bundle is selected, then Qty 1 of HPE Gen4 Re-timer/-p Cable Kit (P54874-B21) must be selected and defaulted.
- If NVMe Bundle is selected, then 8SFF U.3 x4 drive cage will be defaulted.
- NVMe Bundle is supported with 8SFF CTO Server Only.
- If 16NVMe/ 24NVMe Bundle is selected, then Universal Media Bay kit cannot be selected.
- If NVMe Bundle is selected, then Second Processor must be selected.
- If 24NVMe Bundle is selected, then Tertiary Riser cannot be selected.
- If 24NVMe Diff IO-2 bundle is selected, then OCP2 x16 enablement kit and CPU2 OCP2 enablement Kit OR CPU1 OCP2 x8 enablement Kit OR OCP2 x16 enablement cannot be selected.
- If 24NVMe Bundle is selected, then primary 3 x16 Cable OR Sec 3 x16 Cable cannot be selected.

HPE ProLiant DL380 Gen11 24NVME Differential I/O OCP1/2 x16 FIO Bundle Kit P53638-B21

Notes:

- If NVMe bundle is selected, then defaulted Controller has to be deselected.
- If NVMe Bundle is selected, then 8SFF U.3 x1 drive cage (P48813-B21) cannot be selected.
- If this NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 8NVMe CPU1/2 Cable Kit (P48825-B21) must be selected and defaulted.
- If this NVMe Bundle is selected, then Qty 1 of High Performance Fan Kit (P48820-B21) must be selected and defaulted.
- If 24NVMe Bundle is selected, then Qty 3 (min) of 8SFF U.3 x4 drive cage must be selected.
- If 24NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 3x16 Primary Riser

Configuration Information

(P48803-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 x8x16x8 Secondary Riser (P51083-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 1 of OCP1 x16 Enablement Kit (P48827-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 1 of OCP2 x16 Enablement Kit (P48828-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 2 of HPE DL385 Gen10+ 12Gb NVMe 2p Adapter (P25527-B21) must be selected and defaulted.
– If 24NVMe Bundle is selected, then Qty 1 of HPE Gen4 Re-timer/-p Cable Kit (P54874-B21) must be selected and defaulted.
– If NVMe Bundle is selected, then 8SFF U.3 x4 drive cage will be defaulted.
– NVMe Bundle is supported with 8SFF CTO Server Only.
– If 16NVMe/ 24NVMe Bundle is selected, then Universal Media Bay kit cannot be selected.
– If NVMe Bundle is selected, then Second Processor must be selected.
– If 24NVMe Bundle is selected, then Tertiary Riser cannot be selected.
– If 24NVMe Balanced or 24NVMe DIFF IO-3 Bundle is selected, then CPU1 OCP2 x8 enablement OR CPU2 OCP2 x8 enablement kit cannot be selected
– If 24NVMe Bundle is selected, then primary 3 x16 Cable OR Sec 3 x16 Cable cannot be selected.

HPE iLO Common Password FIO Setting P08040-B21

Notes: Sets common iLO password, instead of randomly generated password for each server during Factory Diagnostics.

HPE ProLiant Platform Certificate and IDevID iLO FIO Setting P42104-B21

Notes:

– Initial Device Identity (IDevID) certificates are part of a Zero Trust Architecture. This SKU instructs factory to provision IDevID on HPE iLO.
– Directs HPE manufacturing site to create, digitally sign and store a platform certificate on the server.
– Requires HPE Trusted Platform Module (TPM).

HPE Converged Infrastructure Management Software

HPE OneView w/o iLO including 3yr 24x7 Support Flexible Quantity E-LTU P8B26AAE

HPE OneView including 3yr 24x7 Support Flexible Quantity E-LTU E5Y35AAE

vSAN ReadyNode

- 3, 6, 8 or 16 node vSAN Clusters (3 node minimum)
- HW is optimized for vSAN
- VMware vSAN Advanced LTU bundled

Step 4: Choose additional options for Factory Integration from Core and Additional Options sections below

---

Core Options

Some options may not be integrated at the factory. To ensure only valid configurations are ordered, Hewlett Packard Enterprise recommends the use of an HPE approved configurator. Contact your local sales representative for additional information.

Software as a Service Management

HPE Compute Ops Management

Base SKU

HPE Compute Ops Management Enhanced 3-year Upfront ProLiant SaaS R7A11AAE

Upgrade SKUS

HPE Compute Ops Management Enhanced 1-year Upfront ProLiant SaaS R7A10AAE

HPE Compute Ops Management Enhanced 5-year Upfront ProLiant SaaS R7A12AAE

HPE OneView

HPE OneView including 3yr 24x7 Support Flexible Quantity E-LTU E5Y35AAE

HPE OneView w/o iLO including 3yr 24x7 Support Flexible Quantity E-LTU P8B26AAE

Notes: For customers purchasing HPE Compute Ops Management, without a hardware purchase or a BTO purchase, use this base SKU within ASQ order:

HPE Compute Ops Management Base SaaS R6Z73AAE

For more information, visit the Compute Ops Management QuickSpecs:
https://www.hpe.com/psnow/doc/a50004263enw
Supported Servers - CTO only. No OEM. - Complete list can be found here: Latest Supported Server List:
https://www.hpe.com/info/com-supported-servers

HPE Unique Options

HPE ProLiant DL380 Gen11 12EDSFF Drive Cage Kit P48808-B21

Notes:

- This is the 12EDSFF drive cage.
- This drive cages holds a maximum of 12 single thickness EDSFF drives
- Available for Factory Install Only orders and not available for standalone or upgrade ordering.

HPE ProLiant DL380 Gen11 2SFF U.3 Primary/Secondary Riser Cage Kit P48810-B21

Notes:

- 2SFF drive cage for rear in Primary or Secondary riser position.
- This Drive cage can be selected with 8SFF CTO Server and 24SFF CTO Server Only.
- For 8SFF/ 24SFF CTO Server, Max = 2.

HPE ProLiant DL380 Gen11 2SFF U.3 HDD Stacking Drive Cage Kit P48811-B21

Notes:

- This is a 2SFF drive cage for front or rear. For a front mount it installs into Universal Media Bay.
- This drive cage supports controller and Direct Attach. If Direct Attached then it will support NVMe Drives only
- For 8SFF CTO Server, Max = 2.

HPE ProLiant DL380 Gen11 2SFF U.3 Side-by-Side Drive Cage Kit P48812-B21

Notes:

- This is 2SFF side-by-side drive cage for the 8LFF CTO server only.
- This Drive cage supports controller and Direct Attach. If Direct Attached then it will support NVMe drives only.
- Max = 1

Core Options

HPE ProLiant DL380 Gen11 2U 8SFF x1 Tri-Mode U.3 Drive Cage Kit P48813-B21

Notes:

- This is an 8SFF U.3 x1 front drive cage.
- This drive cage can be selected with 8SFF CTO Server Only.
- Max =3
- If Qty3 of 8SFF Front cage is selected then High Performance Fan Kit (P48820-B21) must be selected.
- This Drive cage supports controller and Direct Attach. If Direct Attached then it will support SATA Drives only and connects to SATA port on motherboard.

HPE ProLiant DL380 Gen11 8SFF U.3 Premium Drive Cage Kit P48814-B21

Notes:

- This is an 8SFF U.3 x4 front drive cage.
- This Drive cage can be selected with 8SFF CTO Server Only.
- Max = 3
- If Qty3 is selected then High Performance Fan Kit (P48820-B21) must be selected.
- This drive cage supports controller and Direct Attach. If Direct Attached then it will support NVMe Drives only.

HPE ProLiant DL380 Gen11 4LFF SAS/SATA 12G LP Midplane Drive Cage Kit P48809-B21

Notes:

- This is the 4LFF mid-tray drive cage.
- This drive cage holds a maximum of 4LFF drives.
- Max 16LFF SAS/SATA is possible when storage controller is selected
- Max 14LFF SATA is possible if mid-tray is direct connected to motherboard

HPE ProLiant DL380 Gen11 8SFF x1 Tri-Mode 24G U.3 BC Midplane Drive Cage Kit P48815-B21

Notes:

- This is the 8SFF U.3 x1 mid-tray drive cage.
- This drive cage holds a maximum of 8 SAS/SATA/NVMe drives.

HPE ProLiant DL380 Gen11 2U 8SFF x4 U.3 Mid Tray Tri-Mode Drive Cage Kit P48816-B21

Notes:

- This is the 8SFF U.3 x4 mid-tray drive cage.
- This drive cage holds a maximum of 8 SAS/SATA/NVMe drives.

HPE ProLiant DL380 Gen11 2LFF Primary Riser Cage Kit P48823-B21

Notes:

- This is a 2LFF drive cage for the rear Primary Riser position.
- This Drive cage supports controller and Direct Attach. If Direct Attached then it will support SATA Drives only and connects to SATA port on motherboard.
- Max = 1
- This drive cage can be selected with 8LFF CTO Server and 12LFF CTO Server Only.

HPE ProLiant DL380 Gen11 2LFF Tertiary Riser Cage Kit P48826-B21

Notes:

- This is a 2LFF drive cage for the rear Secondary + Tertiary Riser position.
- Max = 1

HPE ProLiant DL380 Gen11 SFF Universal Media Bay Kit P50728-B21

Core Options

Notes:

– This is the Universal Media Bay, it occupies an 8SFF drive cage slot on the front of the 8SFF CTO server.
– The Universal Media Bay can be selected with the 8SFF CTO Server only.

HPE ProLiant DL380 Gen11 2LFF LP Secondary Riser Cage Kit P51095-B21

Notes:

– This is a 2LFF drive cage for the rear Primary Riser position.
– Max = 1
– This drive cage can be selected with 8LFF CTO Server and 12LFF CTO Server Only.

HPE ProLiant DL380/DL560 Gen11 2U High Performance Fan Kit P48820-B21

Notes:

– If Processor above 205W is selected then High Performance Fan Kit is required.
– If 128GB or higher memory is selected then High Performance Fan Kit is required.
– If quantity 3 of front drive cage is selected, then High Performance Fan Kit is required.
– If NVMe is selected, then High Performance Fan Kit is required.
– 24SFF CTO server comes with High Performance Fan Kit installed.

HPE ProLiant DL380 Gen11 Max Performance Heat Sink Kit P48817-B21

Notes:

– If Q series Processor is selected then Max Performance Heat Sink is required.
– Max = 1

HPE ProLiant DL380/DL560 Gen11 High Performance 2U Heat Sink Kit P48818-B21

Notes:

– If Processor above 150W is selected then High Performance Heat Sink is required.
– Number of Heat Sinks selected must match number of processor(s) selected.

HPE ProLiant DL3XX/560 Gen11 High Performance Heat Sink Kit P48905-B21

Notes: This is a low profile high performance heatsink, this is only used when storage mid-tray is selected.

HPE ProLiant DL380 Gen11 Standard Heat Sink Kit P49145-B21

Notes: Standard heatsink is for CPUs with 150W or lower TDP.

HPE ProLiant DL380/DL560 Gen11 2U Rear Serial Port Cable Kit P48824-B21

HPE ProLiant DL380 Gen11 System Insight Display Kit P48819-B21

Notes: Max = 1

HPE ProLiant DL380 Compute Gen11 2U 8-pin GPU Power Cable Kit P77896-B21

HPE Processors

Processor Option Kits

5th Generation Intel Xeon-Platinum

All SKUs below ship with processors only. Adequate fans and heatsinks must be selected.

Supports "HPE DDR5 Smart Memory - Registered (RDIMM), 5600MT/s".

Intel Xeon-Platinum 8593Q 2.2GHz 64-core 385W Processor for HPE P68449-B21

Core Options

Notes:

– XCC die
– Max 2
– Requires selection of Max Performance Heat Sink or DLC Module.

Intel Xeon-Platinum 8592+ 1.9GHz 64-core 350W Processor for HPE P67089-B21

Notes:

– XCC die
– Max 2

Intel Xeon-Platinum 8592V 2.0GHz 64-core 330W Processor for HPE P67107-B21

Notes:

– XCC die
– Max 2

Intel Xeon-Platinum 8581V 2.0GHz 60-core 270W Processor for HPE P67109-B21

Notes:

– XCC die
– Max 1

Intel Xeon-Platinum 8580 2.0GHz 60-core 350W Processor for HPE P67088-B21

Notes:

– XCC die
– Max 2

Intel Xeon-Platinum 8570 2.1GHz 56-core 350W Processor for HPE P67087-B21

Notes:

– XCC die
– Max 2

Intel Xeon-Platinum 8568Y+ 2.3GHz 48-core 350W Processor for HPE P67086-B21

Notes:

– XCC die
– Max 2

Intel Xeon-Platinum 8562Y+ 2.8GHz 32-core 300W Processor for HPE P67085-B21

Notes:

– MCC die
– Max 2

Intel Xeon-Platinum 8558P 2.7GHz 48-core 350W Processor for HPE P67108-B21

Notes:

– XCC die
– Max 2

Intel Xeon-Platinum 8558 2.1GHz 48-core 330W Processor for HPE P67097-B21

Core Options

Notes:

– XCC die
– Max 2

Intel Xeon-Platinum 8558U 2.0GHz 48-core 300W Processor for HPE P67102-B21

Notes:

– XCC die
– Max 1

Intel Xeon-Gold 6558Q 3.2GHz 32-core 350W Processor for HPE P67098-B21

Notes:

– MCC die
– Max 2
– Requires selection of Max Performance Heat Sink or DLC Module

Intel Xeon-Gold 6554S 2.2GHz 36-core 270W Processor for HPE P67110-B21

Notes:

– XCC die
– Max 2

Intel Xeon-Gold 6548N 2.8GHz 32-core 250W Processor for HPE P67105-B21

Notes:

– MCC die
– Max 2

Intel Xeon-Gold 6548Y+ 2.5GHz 32-core 250W Processor for HPE P67082-B21

Notes:

– MCC die
– Max 2

Intel Xeon-Gold 6544Y 3.6GHz 16-core 270W Processor for HPE P67084-B21

Notes:

– MCC die
– Max 2

Intel Xeon-Gold 6542Y 2.9GHz 24-core 250W Processor for HPE P67081-B21

Notes:

– MCC die
– Max 2

Intel Xeon-Gold 6538N 2.1GHz 32-core 205W Processor for HPE P67104-B21

Notes:

– MCC die
– Max 2

Intel Xeon-Gold 6538Y+ 2.2GHz 32-core 225W Processor for HPE P67096-B21

Core Options

Notes:

– MCC die
– Max 2

Intel Xeon-Gold 6534 3.9GHz 8-core 195W Processor for HPE P67083-B21

Notes:

– MCC die
– Max 2

Intel Xeon-Gold 6530 2.1GHz 32-core 270W Processor for HPE P67095-B21

Notes:

– XCC die
– Max 2

Intel Xeon-Gold 6526Y 2.8GHz 16-core 195W Processor for HPE P67080-B21

Notes:

– MCC die
– Max 2

Intel Xeon-Gold 5520+ 2.2GHz 28-core 205W Processor for HPE P67094-B21

Notes:

– MCC die
– Max 2

Intel Xeon-Gold 5515+ 3.2GHz 8-core 165W Processor for HPE P67079-B21

Notes:

– MCC die
– Max 2

Intel Xeon-Silver 4516Y+ 2.2GHz 24-core 185W Processor for HPE P67093-B21

Notes:

– MCC die
– Max 2

Intel Xeon-Silver 4514Y 2.0GHz 16-core 150W Processor for HPE P67092-B21

Notes:

– MCC die
– Max 2

Intel Xeon-Silver 4510 2.4GHz 12-core 150W Processor for HPE P67091-B21

Notes:

– EE LCC die
– Max 2
– If this Processor is selected, then 96GB 5600B Memory cannot be selected.

Intel Xeon-Silver 4509Y 2.6GHz 8-core 125W Processor for HPE P67090-B21

Core Options

Notes:

– EE LCC die
– Max 2
– If this Processor is selected, then 96GB 5600B Memory cannot be selected.

Intel Xeon-Bronze 3508U 2.1GHz 8-core 125W Processor for HPE P67100-B21

Notes:

– EE LCC die
– Max 1
If this Processor is selected, then 96GB 5600B Memory cannot be selected.
– Supports PCIe4.0 Only

4th Generation Intel Xeon-Platinum

Notes: All SKUs below ship with processor only. Adequate fans and heatsinks must be selected.

Intel Xeon-Platinum 9462 2.7GHz 32-core 350W Processor for HPE P49645-B21

Notes: This is Intel High Bandwidth Memory (HBM) CPU.

Intel Xeon-Platinum 8490H 1.9GHz 60-core 350W Processor for HPE P49630-B21
Intel Xeon-Platinum 8480+ 2.0GHz 56-core 350W Processor for HPE P49607-B21
Intel Xeon-Platinum 8470Q 2.1GHz 52-core 350W Processor for HPE P49609-B21
Intel Xeon-Platinum 8470 2.0GHz 52-core 350W Processor for HPE P49606-B21
Intel Xeon-Platinum 8470N 1.7GHz 52-core 300W Processor for HPE P49649-B21
Intel Xeon-Platinum 8468V 2.4GHz 48-core 330W Processor for HPE P49631-B21
Intel Xeon-Platinum 8468 2.1GHz 48-core 350W Processor for HPE P49605-B21
Intel Xeon-Platinum 8458P 2.7GHz 44-core 350W Processor for HPE P49632-B21
Intel Xeon-Platinum 8460Y+ 2.0GHz 40-core 300W Processor for HPE P49604-B21
Intel Xeon-Platinum 8452Y 2.0GHz 36-core 300W Processor for HPE P49616-B21
Intel Xeon-Platinum 8444H 2.9GHz 16-core 270W Processor for HPE P49625-B21

4thGeneration Intel Xeon-Gold

Notes: All SKUs below ship with processor only. Adequate fans and heatsinks must be selected.

Intel Xeon-Gold 6458Q 3.1GHz 32-core 350W Processor for HPE P49608-B21
Intel Xeon-Gold 6448H 2.4GHz 32-core 250W Processor for HPE P49622-B21
Intel Xeon-Gold 6448Y 2.1GHz 32-core 225W Processor for HPE P49600-B21
Intel Xeon-Gold 6444Y 3.6GHz 16-core 270W Processor for HPE P49602-B21
Intel Xeon-Gold 6442Y 2.6GHz 24-core 225W Processor for HPE P49599-B21
Intel Xeon-Gold 6438N 2.0GHz 32-core 205W Processor for HPE P49638-B21
Intel Xeon-Gold 6438Y+ 2.0GHz 32-core 205W Processor for HPE P49615-B21
Intel Xeon-Gold 6434 3.7GHz 8-core 195W Processor for HPE P49601-B21
Intel Xeon-Gold 6430 2.1GHz 32-core 270W Processor for HPE P49614-B21
Intel Xeon-Gold 6454S 2.2GHz 32-core 270W Processor for HPE P49654-B21
Intel Xeon-Gold 6426Y 2.5GHz 16-core 185W Processor for HPE P49598-B21
Intel Xeon-Gold 6414U 2.0GHz 32-core 250W Processor for HPE P49619-B21

Notes: Single socket capable, no dual socket support.

Intel Xeon-Gold 6421N 1.8GHz 32-core 185W Processor for HPE P49641-B21
Intel Xeon-Gold 6418H 2.1GHz 24-core 185W Processor for HPE P49621-B21
Intel Xeon-Gold 6416H 2.2GHz 18-core 165W Processor for HPE P49620-B21
Intel Xeon-Gold 5415+ 2.9GHz 8-core 150W Processor for HPE P49597-B21

Core Options

Intel Xeon-Gold 5416S 2.0GHz 16-core 150W Processor for HPE	P49653-B21
Intel Xeon-Gold 5418N 1.8GHz 24-core 165W Processor for HPE	P49640-B21
Intel Xeon-Gold 5418Y 2.0GHz 24-core 185W Processor for HPE	P49612-B21
Intel Xeon-Gold 5420+ 2.0GHz 28-core 205W Processor for HPE	P49613-B21
Intel Xeon-Gold 5411N 1.9GHz 24-core 165W Processor for HPE	P49639-B21

4thGeneration Intel Xeon-Silver

Notes: All SKUs below ship with processor only. Adequate fans and heatsinks must be selected.

Intel Xeon-Silver 4410Y 2.0GHz 12-core 150W Processor for HPE	P49610-B21
Intel Xeon-Silver 4416+ 2.0GHz 20-core 165W Processor for HPE	P49611-B21

4thGeneration Intel Xeon-Bronze

Notes: All SKUs below ship with processor only. Adequate fans and heatsinks must be selected.

Intel Xeon-Bronze 3408U 1.8GHz 8-core 125W Processor for HPE	P49617-B21

Memory Selection

To streamline the configuration process for HPE ProLiant Gen11 servers and to provide the best product availability, HPE recommends memory from the list located here: **https://www.hpe.com/products/recommend**.

Best product availability is limited to US, Canada, and Latin America at this time.

Notes:

- HPE Server Memory compatibility for a specific server platform may vary or be limited within a server platform depending upon the specific configuration being requested. Because each server environment and requirements can vary, memory compatibility is based not only upon the server family, but may also be affected by the amount and type of additional hardware options installed within a specific server configuration. For this reason, some HPE memory DIMMs may be qualified for an HPE server model or family and yet occasionally not be supported with some configurations within that server family
- Maximum memory capacity and speed per processor is dependent on processor model selection or limitation.

HPE DDR5 Memory

Registered DIMMs (RDIMMs)

HPE 16GB (1x16GB) Single Rank x8 DDR5-5600 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64705-B21
HPE 32GB (1x32GB) Dual Rank x8 DDR5-5600 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64706-B21
HPE 64GB (1x64GB) Dual Rank x4 DDR5-5600 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64707-B21

Notes: If Memory Fault Tolerance is selected then only x4 memory options can be selected.

HPE 96GB (1x96GB) Dual Rank x4 DDR5-5600 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64708-B21

Notes:

Core Options

– 96GB Memory cannot be mixed with any other Memory.
– 96GB 5600 Memory cannot be selected with EE LCC (Edge Enhanced Low Core Count) die of the New 5th Generation Intel® Xeon® Scalable Processors.
– For 1 Processor, the allowed qty of 96GB 5600 Memory is 1, 6, 8, 12,16 only.
– For 2 Processor, the allowed qty of 96GB 5600 Memory is 2, 12, 16, 24, 32 only.
– If Memory Fault Tolerance is selected, then only x4 memory options can be selected.

HPE 128GB (1x128GB) Quad Rank x4 DDR5-5600 CAS-52-45-45 EC8 Registered 3DS Smart Memory Kit P64709-B21

Notes: If Memory Fault Tolerance is selected then only x4 memory options can be selected.

HPE 16GB (1x16GB) Single Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit P43322-B21

HPE 32GB (1x32GB) Dual Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit P43328-B21

HPE 64GB (1x64GB) Dual Rank x4 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit P43331-B21

Notes: If Memory Fault Tolerance is selected then only x4 memory options can be selected.

HPE 96GB (1x96GB) Dual Rank x4 DDR5-4800 CAS-46-45-45 EC8 Registered Smart Memory Kit P66675-B21

Notes: 96GB DIMMs supported on XCC and MCC processors only.

HPE 128GB (1x128GB) Quad Rank x4 DDR5-4800 CAS-46-39-39 EC8 Registered 3DS Smart Memory Kit P43334-B21

HPE 256GB (1x256GB) Octal Rank x4 DDR5-4800 CAS-46-39-39 EC8 Registered 3DS Smart Memory Kit P43337-B21

HPE 256GB (1x256GB) Octal Rank x4 DDR5-5600 CAS-52-45-45 EC8 Registered 3DS Smart Memory Kit P64710-B21

Notes:

– Memory should be installed in even quantity of DIMMs
– 4800 MT/s memory SKUs offer a transfer rate of 4800 MT/s at 1 DIMM per channel and 4400 MT/s at 2 DIMMs per channel
– Mixing of 3DS memory and non-3DS memory is not supported
– Mixing of x4 and x8 memory is not allowed
– If 96GB or higher density memory is selected then High Performance Fan Kit must be selected.
– 96GB memory cannot be mixed with any other memory.
– 96GB memory configuration must be either 8 or 16 only on 1 processor unit.
– 96GB memory configuration must be either 16 or 32only on 2 processor unit.

HPE DDR Blank Kit

HPE DDR4 DIMM Blank Kit P07818-B21

Notes: DIMM Blanks are optional and not required.

HPE Optical Drives

HPE 9.5mm SATA DVD-ROM Optical Drive 726536-B21

Notes: HPE DL38X Gen11 Universal Media Bay Kit (P50728-B21) is required for this option on a SFF model. No support in 12LFF or 24SFF models.

HPE 9.5mm SATA DVD-RW Optical Drive 726537-B21

Core Options

Notes: HPE DL38X Gen11 Universal Media Bay Kit (P50728-B21) is required for this option on a SFF model. No support in 12LFF or 24SFF models.

HPE Mobile USB DVD-RW Optical Drive 701498-B21

Notes: This is only supported on USB 3.0 ports.

Media Bay Kits

HPE ProLiant DL380 Gen11 SFF Universal Media Bay Kit P50728-B21

Notes:

- The HPE DL380 Gen11 Universal Media Bay provides front Display Port and 2xUSB 2.0; plus support for 2x SFF front drives or 2 NVME front drives (riser required) and ODD support (Not included); and can only be located in Box1 in either an 8 SFF or 8+8 SFF model.
- This is a SFF model option only.
- Not supported on 12EDSFF, 24SFF, 12FF, or 8LFF CTO servers.

HPE Hard Disk Drives

Mission Critical - 12G SAS - SFF Drives

HPE 2.4TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty 512e Self-encrypting FIPS HDD P28618-B21

HPE 1.2TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Self-encrypting FIPS HDD P28622-B21

Enterprise - 12G SAS - SFF Drives

HPE 2.4TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty 512e Multi Vendor HDD P28352-B21

HPE 1.8TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty 512e Multi Vendor HDD P53562-B21

HPE 1.2TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Multi Vendor HDD P28586-B21

HPE 600GB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Multi Vendor HDD P53561-B21

HPE 300GB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Multi Vendor HDD P40430-B21

Midline - 6G SATA - SFF Drives

HPE 2TB SATA 6G Business Critical 7.2K SFF BC 1-year Warranty 512e HDD P28500-B21

Midline - 12G SAS - LFF Drives

HPE 20TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e ISE Multi Vendor HDD P53553-B21

HPE 16TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e ISE Multi Vendor HDD P23608-B21

HPE 14TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e Multi Vendor HDD P09155-B21

HPE 12TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e Multi Vendor HDD 881781-B21

HPE 8TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty 512e Multi Vendor HDD 834031-B21

HPE 6TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty 512e Multi Vendor HDD 861746-B21

HPE 4TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD 833928-B21

Core Options

HPE 2TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	833926-B21
Midline - 6G SATA - LFF Drives	
HPE 20TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e ISE Multi Vendor HDD	P53554-B21
HPE 16TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e ISE Multi Vendor HDD	P23449-B21
HPE 12TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e Multi Vendor HDD	881787-B21
HPE 8TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty 512e Multi Vendor HDD	834028-B21
HPE 6TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty 512e Multi Vendor HDD	861742-B21
HPE 4TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	861683-B21
HPE 2TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	861681-B21
HPE 1TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	861686-B21

SSD Selection

For SSD selection guidance, please visit **https://ssd.hpe.com/**

Read Intensive - NVMe - EDSFF - Solid State Drives	
HPE 15.36TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM CM7 SSD	P61187-B21
HPE 15.36TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM PM1743 SSD	P57807-B21
HPE 7.68TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM PM1743 SSD	P57803-B21
HPE 7.68TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM CM7 SSD	P61183-B21
HPE 7.68TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM PM1743 SSD	P57803-B21
HPE 3.84TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM CM7 SSD	P61179-B21
HPE 3.84TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM PM1743 SSD	P57799-B21
HPE 3.84TB NVMe Gen5 High Performance Read Intensive E3S EC1 PS1010 SSD	P70392-B21
HPE 1.92TB NVMe Gen5 Mainstream Performance Read Intensive E3S EC1 CD8P SSD	P69234-B21
Mixed Use - NVMe - EDSFF - Solid State Drives	
HPE 3.2TB NVMe Gen5 High Performance Mixed Use E3S EC1 PS1030 SSD	P70399-B21
HPE 3.2TB NVMe Gen5 High Performance Mixed Use E3S EC1 EDSFF SPDM CM7 SSD	P61191-B21
HPE 6.4TB NVMe Gen5 High Performance Mixed Use E3S EC1 EDSFF SPDM CM7 SSD	P61195-B21
HPE 1.6TB NVMe Gen5 Mainstream Performance Mixed Use E3S EC1 CD8P SSD	P69241-B21
Read Intensive - 24G SAS - SFF - Solid State Drives	
HPE 15.36TB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49045-B21
HPE 3.84TB SAS Read Intensive SFF BC Self-encrypting FIPS 140-2 PM7 SSD	P63875-B21
HPE 7.68TB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49041-B21

Core Options

HPE 3.84TB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49035-B21
HPE 1.92TB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49031-B21
HPE 960GB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49029-B21
Read Intensive - 12G SAS - SFF - Solid State Drives	
HPE 7.68TB SAS 12G Read Intensive SFF BC Value SAS Multi Vendor SSD	P40509-B21
HPE 3.84TB SAS 12G Read Intensive SFF BC Value SAS Multi Vendor SSD	P40508-B21
HPE 1.92TB SAS 12G Read Intensive SFF BC Value SAS Multi Vendor SSD	P40507-B21
HPE 960GB SAS 12G Read Intensive SFF BC Value SAS Multi Vendor SSD	P40506-B21
HPE 960GB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49029-B21
Mixed Use - 24G SAS - SFF - Solid State Drives	
HPE 1.6TB SAS Mixed Use SFF BC Self-encrypting FIPS 140-2 PM7 SSD	P63871-B21
HPE 6.4TB SAS 24G Mixed Use SFF BC Multi Vendor SSD	P49057-B21
HPE 3.2TB SAS 24G Mixed Use SFF BC Multi Vendor SSD	P49053-B21
HPE 1.6TB SAS 24G Mixed Use SFF BC Multi Vendor SSD	P49049-B21
HPE 800GB SAS 24G Mixed Use SFF BC Multi Vendor SSD	P49047-B21
Mixed Use - 12G SAS - SFF - Solid State Drives	
HPE 3.84TB SAS 12G Mixed Use SFF BC Value SAS Multi Vendor SSD	P40512-B21
HPE 1.92TB SAS 12G Mixed Use SFF BC Value SAS Multi Vendor SSD	P40511-B21
HPE 960GB SAS 12G Mixed Use SFF BC Value SAS Multi Vendor SSD	P40510-B21
Read Intensive - 6G SATA - SFF - Solid State Drives	
HPE 480GB SATA 6G Read Intensive SFF BC PM893a SSD	P63886-B21
HPE 480GB SATA 6G Read Intensive SFF BC PM893a SSD	P63886-B21
HPE 3.84TB SATA 6G Read Intensive SFF BC PM893a SSD	P63910-B21
HPE 3.84TB SATA 6G Read Intensive SFF BC PM893a SSD	P63910-B21
HPE 7.68TB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40501-B21
HPE 3.84TB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40500-B21
HPE 1.92TB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40499-B21
HPE 480GB SATA 6G Read Intensive SFF BC Self-encrypting 5400P SSD	P58236-B21
HPE 480GB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40497-B21
HPE 960GB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40498-B21
HPE 240GB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40496-B21
Mixed Use - 6G SATA - SFF - Solid State Drives	
HPE 3.84TB SATA 6G Mixed Use SFF BC Multi Vendor SSD	P40505-B21
HPE 1.92TB SATA 6G Mixed Use SFF BC Multi Vendor SSD	P40504-B21
HPE 960GB SATA 6G Mixed Use SFF BC Self-encrypting 5400M SSD	P58244-B21
HPE 960GB SATA 6G Mixed Use SFF BC Multi Vendor SSD	P40503-B21
HPE 480GB SATA 6G Mixed Use SFF BC Multi Vendor SSD	P40502-B21
Mixed Use - 12G SAS - LFF -Solid State Drives	
HPE 960GB SAS 12G Mixed Use LFF LPC Value SAS Multi Vendor SSD	P37009-B21
Read Intensive - 6G SATA - LFF - Solid State Drives	
HPE 960GB SATA 6G Read Intensive LFF LPC Multi Vendor SSD	P47808-B21
Read Intensive - NVMe - SFF - Solid State Drives	
HPE 15.36TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PS1010 SSD	P70436-B21
HPE 15.36TB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static SPDM Multi Vendor SSD	P69255-B21
HPE 15.36TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 CM7 SSD	P63841-B21
HPE 15.36TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PM1733a SSD	P50224-B21

Core Options

HPE 7.68TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PS1010 SSD	P70434-B21
HPE 7.68TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 CM7 SSD	P63837-B21
HPE 7.68TB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static V2 Multi Vendor SSD	P64848-B21
HPE 7.68TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PM1733a SSD	P50222-B21
HPE 3.84TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 CM7 SSD	P63833-B21
HPE 3.84TB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static V2 Multi Vendor SSD	P64846-B21
HPE 3.84TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PM1733a SSD	P50219-B21
HPE 1.92TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 CM7 SSD	P63829-B21
HPE 1.92TB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static V2 Multi Vendor SSD	P64844-B21
HPE 1.92TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PM1733a SSD	P50216-B21
HPE 960GB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static V2 Multi Vendor SSD	P64842-B21
Mixed Use - NVMe - SFF - Solid State Drives	
HPE 6.4TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PS1030 SSD	P70428-B21
HPE 6.4TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 CM7 SSD	P63853-B21
HPE 6.4TB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static V2 Multi Vendor SSD	P65023-B21
HPE 6.4TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD	P50233-B21
HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PS1030 SSD	P70426-B21
HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 CM7 SSD	P63849-B21
HPE 3.2TB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static V2 Multi Vendor SSD	P65015-B21
HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD	P50230-B21
HPE 1.6TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 CM7 SSD	P63845-B21
HPE 1.6TB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static V2 Multi Vendor SSD	P65007-B21
HPE 1.6TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD	P50227-B21
HPE 800GB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static V2 Multi Vendor SSD	P64999-B21

Hard Drive Blank Kits

HPE Gen9 LFF HDD Spade Blank Kit	807878-B21
HPE Small Form Factor Hard Drive Blank Kit	666987-B21
Hard Drive Cage Kits	
HPE ProLiant DL380 Gen11 12EDSFF Drive Cage Kit	P48808-B21
Notes: Currently available as Factory Install Only and not available for standalone or upgrade ordering.	
HPE ProLiant DL380 Gen11 2U 8SFF x1 Tri-Mode U.3 Drive Cage Kit	P48813-B21
HPE ProLiant DL380 Gen11 8SFF U.3 Premium Drive Cage Kit	P48814-B21
HPE ProLiant DL380 Gen11 SFF Universal Media Bay Kit	P50728-B21
HPE ProLiant DL380 Gen11 2SFF U.3 HDD Stacking Drive Cage Kit	P48811-B21
HPE ProLiant DL380 Gen11 2SFF U.3 Side-by-Side Drive Cage Kit	P48812-B21
HPE ProLiant DL380 Gen11 2SFF U.3 Primary/Secondary Riser Cage Kit	P48810-B21

Core Options

HPE ProLiant DL380 Gen11 8SFF x1 Tri-Mode 24G U.3 BC Midplane Drive Cage Kit	P48815-B21
HPE ProLiant DL380 Gen11 2U 8SFF x4 U.3 Mid Tray Tri-Mode Drive Cage Kit	P48816-B21
HPE ProLiant DL380 Gen11 4LFF SAS/SATA 12G LP Midplane Drive Cage Kit	P48809-B21
HPE ProLiant DL380 Gen11 2LFF Primary Riser Cage Kit	P48823-B21
HPE ProLiant DL380 Gen11 2LFF LP Secondary Riser Cage Kit	P51095-B21
HPE ProLiant DL380 Gen11 2LFF Tertiary Riser Cage Kit	P48826-B21

HPE Networking

1 Gigabit Ethernet adapters

Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T Adapter for HPE	P51178-B21
Intel I350-T4 Ethernet 1Gb 4-port BASE-T Adapter for HPE	P21106-B21

10 Gigabit Ethernet adapters

Notes: Unless otherwise noted, one of the below 10Gb networking adapters below can be selected as the primary networking choice when configuring a Networking Choice (NC) Configure-to-Order (CTO) chassis. The DL380 Gen11 NC CTO chassis does not come with embedded networking, hence the requirement to configure with either a FlexibleLOM or select PCIe networking adapter.

Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T Adapter for HPE	P26253-B21
Broadcom BCM57412 Ethernet 10Gb 2-port SFP+ Adapter for HPE	P26259-B21

25 Gigabit Ethernet adapters

Notes: Unless otherwise noted, one of the below 10/25Gb networking adapters below can be selected as the primary networking choice when configuring a Networking Choice (NC) Configure-to-Order (CTO) chassis.

The DL380 Gen11 NC CTO chassis does not come with embedded networking, hence the requirement to configure with either an OCP3 or select PCIe networking adapter.

Broadcom BCM57504 Ethernet 10/25Gb 4-port SFP28 Adapter for HPE	P26264-B21
Broadcom BCM57414 Ethernet 10/25Gb 2-port SFP28 Adapter for HPE	P26262-B21
HPE Ethernet 10/25Gb 2-port Secure Network Adapter	S2A69A
Intel E810-XXVDA2 Ethernet 10/25Gb 2-port SFP28 Adapter for HPE	P08443-B21
Intel E810-XXVDA4 Ethernet 10/25Gb 4-port SFP28 Adapter for HPE	P08458-B21
Xilinx X2522-25G-PLUS Ethernet 10/25Gb 2-port SFP28 Adapter for HPE	P21109-B21
Mellanox MCX631102AS-ADAT Ethernet 10/25Gb 2-port SFP28 Adapter for HPE	P42044-B21
Mellanox MCX623106AS-CDAT Ethernet 100Gb 2-port QSFP56 Adapter for HPE	P25960-B21

100 Gigabit Ethernet Adapters

Mellanox MCX623106AS-CDAT Ethernet 100Gb 2-port QSFP56 Adapter for HPE	P25960-B21

Recommended System Ambient Temperature	
System Config	**P25960-B21**
8LFF	25C
24SFF	Not support
16SFF	25C
8SFF	25C

Core Options

Other Restrictions

1. These cards are not supported with 12LFF CTO server and 24SFF CTO server config.
2. Required to use High Performance Fan Kit (P48820-B21)
3. Only supported on x16 physical and electrical slots.

Intel E810-CQDA2 Ethernet 100Gb 2-port QSFP28 Adapter for HPE P21112-B21

Recommended System Ambient Temperature	
System Config	**P21112-B21**
8LFF	25C
24SFF	Not supported
16SFF	25C
8SFF	30C

Notes:

– This adapter requires High Performance Fan Kit (P48820-B21).
– Not supported on 8SFF CTO server with 3x drive cages.
– Not Supported with 24SFF and 12LFF CTO Servers.
– Only supported on x16 physical and electrical slots.

200 Gigabit Ethernet Adapters

Mellanox MCX623105AS-VDAT Ethernet 200Gb 1-port QSFP56 Adapter for HPE P10180-B21

Recommended Ambient Temperature	
System Config	**P10180-B21**
8LFF	25C
24SFF	25C (3 max)
16SFF	25C
8SFF	30C

200 Gigabit Slingshot Adapters

HPE Slingshot SA210S Ethernet 200Gb 1-port PCIe NIC R4K46A

Notes:

– Can only be selected or configured for a Cray or Slingshot Solution. Not allowed for a Non-Cray or Non-Slingshot Solution.
– Cannot have the following networking options configured within the same server: Slingshot 11 or Slingshot 22.

Storage Offload Adapters

HPE NV60100M 100Gb 2-port Storage Offload Adapter R8M41A

Notes: This storage offload adapter requires selection 2 each of either 100GbE QSFP28 PSM4 500m XCVR OR 100GbE QSFP28 SR4 100m XCVR transceivers OR HPE 100GbE QSFP28 to QSFP28 5m AOC.

Core Options

OCP 3.0 Adapters

Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE	P51181-B21
Intel I350-T4 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE	P08449-B21
Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T OCP3 Adapter for HPE	P10097-B21
Broadcom BCM57412 Ethernet 10Gb 2-port SFP+ OCP3 Adapter for HPE	P26256-B21
Broadcom BCM57504 Ethernet 10/25Gb 4-port SFP28 OCP3 Adapter for HPE	P26269-B21
Broadcom BCM57414 Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE	P10115-B21
Intel E810-XXVDA4 Ethernet 10/25Gb 4-port SFP28 OCP3 Adapter for HPE	P41614-B21
Intel E810-XXVDA2 Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE	P10106-B21
Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE	P42041-B21
Intel E810-CQDA2 Ethernet 100Gb 2-port QSFP28 OCP3 Adapter for HPE	P22767-B21

OCP 3.0 Enablement

HPE ProLiant DL360 Gen11 CPU1 to OCP2 x8 Enablement Kit	P51911-B21
HPE ProLiant DL3XX Gen11 OCP1 x16 Enablement Kit	P48827-B21
HPE ProLiant DL3XX Gen11 OCP2 x16 Enablement Kit	P48828-B21
HPE ProLiant DL3XX Gen11 CPU2 to OCP2 x8 Enablement Kit	P48830-B21

DL380 Gen11 OCP 1 and OCP 2 Priority Support Matrix

OCP Slot Location	1 OCP Storage Controller (OROC) + 1OCP NIC	1 OCP NIC	2 OCP NICs	1 OCP Storage Controller (OROC)	2 OCP Storage Controllers (OROC)
OCP 1	OROC	N/A	OCP NIC	OROC (Higher priority)	OROC (Higher priority)
OCP 2 (with shared NIC and WoL)	OCP NIC	NIC (higher priority)	OCP NIC (higher priority)	N/A	OROC

HPE InfiniBand

HPE InfiniBand NDR/Ethernet 400Gb 1-port OSFP PCIe5 x16 MCX75310AAS-NEAT Adapter	P45641-B23
HPE InfiniBand NDR200/Ethernet 200Gb 1-port OSFP PCIe5 x16 MCX75310AAS-HEAT Adapter	P45642-B22
HPE InfiniBand NDR200/Ethernet 200GbE 2-port QSFP112 PCIe5 x16 MCX755106AC-HEAT Adapter	P65333-B21
HPE InfiniBand NDR 1-port OSFP PCIe5 x16 MCX75310AAS-NEAT Adapter	P45641-B21

Notes:

- High Performance Fan Kit is required (P48820-B21).
- Must be populated in x16 physical and electrical slot.
- Ambient temperature should not exceed 25C.

Recommended Ambient Temperature	
System Config	**P31323-B21**
8LFF	25C (only to OCP2)
24SFF	Not supported
16SFF	25C (only to OCP2)
8SFF	30C (only to OCP2)

Core Options

HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 OCP3 MCX653436A-HDAI Adapter P31348-B21

Recommended Ambient Temperature	
System Config	**P31323-B21**
8LFF	Not supported
24SFF	Not supported
16SFF	Not supported
8SFF	25C (only to OCP2)

Other Restrictions

1. High Performance Fan Kit is required (P48820-B21).
2. Not supported on 24SFF CTO server or 12LFF CTO server.
3. Ambient temperature should not exceed 25C.
4. OCP2 x16 Enablement Kit (P48828-B21) is required.
5. 256GB DIMMs not supported if these adapters are selected.
6. Max = 1Could observe sub-optimal performance if installed in x8 slot.

HPE InfiniBand HDR/Ethernet 200Gb 1-port QSFP56 PCIe4 x16 MCX653105A-HDAT Adapter P23664-B21

Recommended Ambient Temperature	
System Config	**P23664-B21**
8LFF	25C
24SFF	25C
16SFF	25C
8SFF	30C

Other Restrictions

1. High Performance Fan Kit is required (P48820-B21).
2. Must be populated in x16 physical and electrical slot.
3. Ambient temperature should not exceed 25C.

If configured for a Cray or Slingshot Solution, this option is to be used as the Slingshot 10 networking card.

HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 MCX653106A-HDAT Adapter P31324-B21

Recommended System Ambient Temperature	
System Config	**P31324-B21**
8LFF	25C
24SFF	Not supported
16SFF	25C
8SFF	25C

Other Restrictions

1. High Performance Fan Kit is required (P48820-B21).
2. Not supported on 24SFF CTO server or 12LFF CTO server.
3. Must be populated in x16 physical and electrical slot.
4. Ambient temperature should not exceed 25C.

If configured for a Cray or Slingshot Solution, this option is to be used as the Slingshot 10 networking card.

HPE 100Gb 1-port OP101 QSFP28 x16 PCIe Gen3 with Intel Omni-Path Architecture Adapter 829335-B21

Core Options

Notes:

– High Performance Fan Kit is required (P48820-B21).
– Must be populated in x16 physical and electrical slot.

HPE InfiniBand HDR100/Ethernet 100Gb 1-port QSFP56 PCIe4 x16 MCX653105A-ECAT Adapter P23665-B21

Notes: High Performance Fan Kit is required (P48820-B21).

HPE InfiniBand HDR100/Ethernet 100Gb 2-port QSFP56 PCIe4 x16 MCX653106A-ECAT Adapter P23666-B21

Recommended Ambient Temperature		
System Config	**P23665-B21**	**P23666-B21**
8LFF	30C	25C
24SFF	30C	Not supported
16SFF	30C	25C
8SFF	30C	30C

Notes:
– High Performance Fan Kit is required (P48820-B21).
– Must be populated in x16 physical and electrical slot.
– If 2SFF drive cage is selected then Max = 2
– Max = 4

HPE I/O Expansion Options

Notes:

– The Primary Riser shipping default in the CTO server is a x8 FH, FL, x16 FH, FL and x8 FH, HL.
– For a Secondary/Tertiary riser, the second processor is required.
– x16 cards installed on x8 slots could observe sub-optimal performance.

HPE ProLiant DL380 Gen11 2U x16/x16/x16 Primary Riser Kit P48803-B21

Notes:
– Slot 1 - PCIe 5.0 x16 Full Height and Full Length
– Slot 2 - PCIe 5.0 x16 Full Height and Full Length
– Slot 3 - PCIe 5.0 x16 Full Height and Half Length
– If 24NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 3x16 Primary riser must be selected and defaulted.
– If this Primary Riser is selected, then default Primary Riser is replaced with this riser.
– If Slot 1 of HPE DL380 Gen11 2U 3x16 Primary Riser Kit needs to be enabled then 3 x16 Primary Cable Kit (P56073-B21) must be selected.
– If Primary 3 x16 Cable Kit is NOT selected, then only Slot 2 and Slot 3 will be available for PCIe card selection and no PCIe cards can be selected for Slot 1.

HPE ProLiant DL380 Gen11 2U x16/x16/x16 Secondary Riser Kit P51083-B21

Notes:

– Slot 4 - PCIe 5.0 x16 Full Height and Full Length
– Slot 5 - PCIe 5.0 x16 Full Height and Full Length
– Slot 6 - PCIe 5.0 x16 Full Height and Half Length
– When 2LFF Tertiary Cage is selected then Secondary and Tertiary Riser cannot be selected.

Core Options

– When 2LFF Secondary Cage is selected then Secondary Riser cannot be selected.
– If 24NVMe Bundle is selected, then Qty 1 of HPE DL380 Gen11 3x16 SEC Riser must be selected and defaulted.
– If Slot 4 of HPE DL380 Gen11 2U 3x16 Secondary Riser Kit needs to be enabled then 3 x16 Secondary Cable Kit (P56074-B21) must be selected.
– If Secondary 3 x16 Cable Kit is NOT selected, then only Slot 5 and Slot 6 will be available and no PCIe cards can be selected for Slot4.
– If Secondary OR Tertiary Riser is selected, then Second Processor must be selected.
– Tertiary Riser and Secondary 3 x16 Riser cannot be selected together.

HPE ProLiant DL380 Gen11 2U Primary/Secondary NEBS-compliant Riser Kit P48805-B21

Notes:

– GPU and NEBS Riser cannot be selected together.
– When this riser is selected it replaces the default Primary Riser.
– When Primary and Secondary NEBS risers are selected the Tertiary Riser cannot be selected.
– NEBS risers cannot be mixed with other non-NEBS risers.

HPE ProLiant DL380 Gen11 2U Secondary/Tertiary NEBS-compliant Riser Kit P48806-B21

Notes:

– This riser requires selection of second processor.
– GPU and NEBS Riser cannot be selected together.
– When Primary and Secondary NEBS risers are selected the Tertiary Riser cannot be selected.
– NEBS risers cannot be mixed with other non-NEBS risers.
– When 24 NVMe or 32 NVMe bundles are selected the Tertiary Riser cannot be selected.

HPE ProLiant DL380 Gen11 2U x8/x16/x8 Secondary Riser Kit P48802-B21

Notes:

– Slot 4 - PCIe 5.0 x8 Full Height and Full Length
– Slot 5 - PCIe 5.0 x16 Full Height and Full Length
– Slot 6 - PCIe 5.0 x8 Full Height and Half Length
– If quantity 1 of 2SFF Primary/Secondary Drive Cage is selected, then HPE DL380 Gen11 x8x16x8 secondary riser is required.
– If quantity 1 of 2SFF Primary/Secondary Drive Cage is selected, then top two slots (Slot 4 and Slot 5) of HPE DL380 Gen11 x8x16x8 Secondary Riser will be blocked by Drive Cage.
– If 2LFF Tertiary Drive Cage is selected, then Secondary and Tertiary Riser cannot be selected.
– if 2LFF Secondary Drive Cage is selected then Secondary Riser cannot be selected.
– If 24NVMe Bundle is selected, then quantity 1 of HPE DL380 Gen11 x8x16x8 secondary riser must be selected.
– If Secondary OR Tertiary Riser is selected, then Second Processor must be selected.

HPE ProLiant DL380 Gen11 2U x16/x16 Tertiary Riser Kit P48804-B21

Notes:

– This is the tertiary riser.
– Slot 7 - PCIe 5.0 x16 Full Height and Full Length
– Slot 8 - PCIe 4.0 x16 Full Height and Full Length

HPE ProLiant DL380 Gen11 x16/x16/x16 Primary Cable Kit P56073-B21

Core Options

HPE ProLiant DL380 Gen11 x16/x16/x16 Secondary Cable Kit	P56074-B21
HPE ProLiant DL380 Gen11 2x16 Tertiary Riser x8 Enablement FIO Bundle Kit	P53632-B21

Risers

Riser Information*							
Part number	**Description**	**Riser position**			**Bus width (Gen5 lanes)**		
		Primary	**Secondary**	**Tertiary**	**Top slot**	**Middle Slot**	**Bottom slot**
N/A	This is the default riser in the chassis	D	N	N	x8	x16	x8
P48803-B21	HPE DL380 Gen11 x16/x16/x16 Primary Riser Kit	O	N	N	x16	x16	x16[1]
P51083-B21	HPE DL380 Gen11 x16/x16/x16 Secondary Riser Kit	N	O	N	x16	x16	x16[2]
P48802-B21	HPE DL38X Gen11 x8/x16/x8 Sec Riser Kit	N	O	N	x8	x16	x8
P48804-B21	HPE DL38X Gen11 2x16 Tertiary Riser Kit	N	N	O	x16	x16[3]	
P48805-B21	HPE ProLiant DL380 Gen11 2U Primary/Secondary NEBS-compliant Riser Kit	O	N	N	x8	x16	x8
P48806-B21	HPE ProLiant DL380 Gen11 2U Secondary/Tertiary NEBS-compliant Riser Kit	N	O	N	x8	x16	x8

Notes:

- D = Default on chassis; O = Optional; N = not supported or slot/connector not present.
- [1]Requires HPE DL380 Gen11 x16/x16/x16 Primary Cable Kit (P56073-B21)
- [2]Requires HPE DL380 Gen11 x16/x16/x16 Secondary Cable Kit (P56074-B21)
- [3]PCIe Gen4 lanes.
- x16 cards installed on x8 slots could observe sub-optimal performance.

HPE Power Supplies

HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P38995-B21
Notes: Flex Slot Platinum power supplies support power efficiency of up to 94% and include a standard C-14 power inlet connector.	
HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit	P03178-B21
Notes: Flex Slot Titanium power supplies support power efficiency of up to 96% and include a standard C-14 power inlet connector.	
HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P38997-B21
Notes: Flex Slot Platinum Plus power supplies support power efficiency of up to 94% and include a C-14 power inlet connector that can support HPE Power Discovery Services (blue connector).	
HPE 1600W Flex Slot -48VDC Hot Plug Power Supply Kit	P17023-B21
Notes: Requires selection of HPE 1600W DC PSU power lug option kit OR HPE 1600W DC PSU Power Cable Kit.	
HPE 1600W -48VDC Power Cable Lug Kit	P36877-B21

Core Options

HPE 1800W-2200W Flex Slot Titanium Hot Plug Power Supply Kit P44712-B21

Notes:

- Flex Slot Platinum power supplies support power efficiency of up to 94% and are EU Lot 9 compliant.
- Beginning on January 1st, 2024, units sold into the European Union (EU), European Economic Area (EEA), the United Kingdom, or Switzerland must include more efficient AC power supplies: 94% for multi-output and 96% for single output. HPE Flexible Slot power supplies are single-output, and part numbers 865438-B21, P03178-B21, and P44712-B21 are 96% efficient, thus meeting requirements. HPE is on target to fulfil compliant systems ahead of time and will begin enforcing these requirements in advance to satisfy requests with the current power supplies by the set deadline.

HPE Cooling Options

HPE ProLiant DL380/DL560 Gen11 2U High Performance Fan Kit P48820-B21

Notes:

- This kit is required for specific **Ambient temperature environments**.
- High Performance fan kit consists of 6 fans, these will need to replace all the standard fans in the unit and fill all 6 fan cages.
- The 24SFF CTO server will already include 6 High Performance fans.
- The High-Performance fan kit is needed to support certain ASHRAE operating environments.
- For elevated ambient temperature support please see:

DL380 Gen11 Extended Ambient Temperature Guidelines.

HPE ProLiant DL380 Gen11 Standard Fan Kit P49146-B21

Notes:

- Includes two standard fans.
- Not supported with 24SFF and 12EDSFF CTO Server.
- High Performance Fan Kit and Standard Fan Kit cannot be selected together.
- High Performance Fan Kit can be selected over Standard Fan Kit in all configurations.

HPE ProLiant DL380 Gen11 Max Performance Heat Sink Kit P48817-B21

Notes: This kit is required for "Q" series processors.

HPE ProLiant DL380/DL560 Gen11 High Performance 2U Heat Sink Kit P48818-B21

Notes: This kit is CPUs with TDP over 150W.

HPE ProLiant DL3XX/560 Gen11 High Performance Heat Sink Kit P48905-B21

Notes: This kit is a low-profile high performance heatsink. This heatsink is required when a mid-tray storage cage is selected.

HPE ProLiant DL380 Gen11 Standard Heat Sink Kit P49145-B21

Notes: The standard heatsink is for CPUs with TDP equal to or lower than 150W.

HPE Direct Liquid Cooling Options for HPE ProLiant DL380 Gen11

HPE ProLiant DL380 Gen11 Cold Plate Module NS204 Quick Disconnect Tube Set FIO Kit P62023-B21

Notes:

- This kit is the Direct Liquid Cooling Kit that uses the NS204i-u slot for connections

Core Options

– If this option is selected the HPE NS204i-u Gen11 NVMe Hot Plug Boot Optimized Storage Device(P48183-B21) cannot be selected.
– This kit contains 2 Cold Plate Modules and 1 Quick Disconnect Module.

HPE ProLiant DL3XX Gen11 Cold Plate Module FIO Kit from PCIe P62029-B21

Notes:

– This kit is the Direct Liquid Cooling Kit that uses a PCIe slot on the Primary Riser.
– If this option is selected the Primary Riser will have one less PCIe slot available for PCIe adapters. Please keep this in mind when considering total number of PCIe adapters required.
– This kit contains 2 Cold Plate Modules and 1 Quick Disconnect Module.

HPE ProLiant DL3XX Gen11 Direct Liquid Cooling 600mm FIO Hose Kit P62038-B21

Notes:

– This kit includes the 65CM tube kit for Direct Liquid Cooling.
– This kit must be selected when using the Direct Liquid Colling Kit in the NS204i-u slot (P62023-B21).

HPE ProLiant DL3XX Gen11 Direct Liquid Cooling 55cm Quick Disconnect Tube Set FIO Kit P62042-B21

Notes:

– This kit includes the 55CM tube kit for Direct Liquid Cooling.
– This kit must be selected when using the Direct Liquid Colling Kit in the Primary Riser slot (P62029-B21).

Additional Options

Some options may not be integrated at the factory. To ensure only valid configurations are ordered, Hewlett Packard Enterprise recommends the use of an HPE approved configurator. Contact your local sales representative for additional information.

HPE Computation and Graphics Accelerators

NVIDIA H100 NVL 94GB PCIe Accelerator for HPE S2D86C

Notes:

- Max =3
- This GPU is PCIe x16, so it must be populated on a x16 slot only. The "PCIe Card Capacity Limits" must account for this limitation.
- Thermal constraints for this GPU are based on ambient temperature of 25 degrees Centigrade.
- Mixing of Graphics options is not allowed.
- High Performance Fan Kit is required.
- The system memory must be at least twice the memory of all GPUs.
- When this GPU is selected then the top x16 PCIe slot (FH FL) connector on the riser kit CANNOT be utilized as the GPU occupies this space. The "PCIe Card Capacity Limits" must account for this limitation.
- If this GPU is selected, then HPE DL380/DL560 Gen11 2U GPU Power Cable Kit (P56072-B21) must be selected.
- For best performance across common workloads, HPE recommends system main memory at least twice the memory of all GPU.
- Full Length (FL) PCIe Card and Mid Tray cannot be selected together.
- GPU and NEBS Riser cannot be selected together.

NVIDIA RTX 4000 Ada Graphics Accelerator for HPE S3T54C

Notes:

- Max = 5
- When this GPU is selected then the top x16 PCIe slot (FH FL) connector on the riser kit CANNOT be utilized as the GPU occupies this space. The "PCIe Card Capacity Limits" must account for this limitation.
- Thermal constraints for this GPU are based on ambient temperature of 25 degrees Centigrade.
- Mixing of Graphics options is not allowed.
- High Performance Fan Kit is required.
- The system memory must be at least twice the memory of all GPUs.
- If this GPU is selected, then HPE DL380/DL560 Gen11 2U GPU Power Cable Kit (P56072-B21) must be selected.
- GPU and NEBS Riser cannot be selected together.

Intel Data Center GPU Max 1100 48GB Accelerator for HPE S1T66C

Notes:

- Max = 3
- This GPU is PCIe x16, so it must be populated on a x16 slot only. The "PCIe Card Capacity Limits" must account for this limitation.
- Thermal constraints for this GPU are based on ambient temperature of 25 degrees Centigrade.
- Mixing of GPU types is not allowed.
- High Performance Fan Kit is required.
- This GPU is not supported with 12LFF CTO Server and 24SFF CTO Server.
- The system memory must be at least twice the memory of all GPUs.

Additional Options

- When this GPU is selected then the top x16 PCIe slot (FH FL) connector on the riser kit CANNOT be utilized as the GPU occupies this space. The "PCIe Card Capacity Limits" must account for this limitation.
- For the DL380 Gen11 12EDSFF CTO Server, If EDSFF Bundle is selected then this GPU cannot be selected.
- If this GPU is selected, then HPE DL380/DL560 Gen11 2U GPU Power Cable Kit (P56072-B21) must be selected.
- If this GPU is selected, then High Performance Heatsink or Direct Liquid Cooling (DLC) must be selected.
- If this GPU is selected with Intel Liquid Cooled Processors (Q) then Direct Liquid Cooling (DLC) must be selected.

NVIDIA H100 80GB PCIe Accelerator for HPE R9S41C

Notes:

- Max = 3
- Must be populated in x16 slot.
- System memory should be 2x of GPU memory.
- Requires selection of High-Performance Fan Kit.
- This GPU requires HPE ProLiant DL380/DL560 Gen11 2U GPU Power Cable Kit - P56072-B21
- Not supported with 12LFF CTO Server, 24SFF CTO Server, or 8SFF CTO Server with 3x 8SFF drive cages.
- For 12EDSFF CTO Server, If no additional cage is selected then Max of 2 GPU can be selected per Server.
- On 12EDSFF CTO Server, when one additional 8SFF cage is selected with DLC Component then Max 1 GPU can be selected per sever.
- On 8SFF CTO Server, If Qty2 of 8SFF Front cage is selected with DLC Component then Max 1 GPU can be selected per sever.
- On 8SFF CTO Server, If Qty2 of 8SFF Front cage is selected without DLC Component then this GPU is not supported.
- For 12EDSFF CTO Server, when additional cages are selected without DLC Component then this GPU is not supported
- On 8LFF CTO Server, If DLC component is selected then Max 1 GPU can be selected per sever.
- If Qty1 of this GPU is selected, then Secondary OR Tertiary Riser Must be selected.
- If Qty2 of this GPU is selected, then Secondary AND Tertiary Riser Must be selected.

NVIDIA L40S 48GB PCIe Accelerator S2L70C

Notes:

- Max = 2
- Must be populated in x16 slot.
- If Qty2 of DW GPU is selected with no tertiary riser, then primary 3 x16 riser cable and secondary 3 x16 Riser Cable cannot be selected.
- Thermal constraints are based on Ambient Temp 25C.
- Mixing of GPU types is not allowed.
- For 8SFF/ 8LFF/ 12LFF/ 12EDSFF CTO Server, If GPU is selected then High-Performance Fan Kit must be selected and defaulted.
- This GPU is not supported with 12LFF CTO Server and 24SFF CTO Server.
- For 8SFF CTO Server, If Qty3 of 8SFF Front cage is selected then this GPU cannot be selected.
- For best performance across common workloads, HPE recommends system main memory at least twice the memory of all GPU.
- If doublewide GPU is selected on this riser, then the top x16 PCIe slot (FH FL) connector on

Additional Options

this riser kit cannot be utilized as the doublewide GPU occupies this space. The "PCIe Card Capacity Limits" must account for this limitation.

- If Qty1 of 2SFF primary/secondary cage is selected, then Doublewide GPU Cannot be populated on HPE DL380 G11 2U x8/x16/x8 Sec Riser Kit.
- If Qty2 of 2SFF Primary/secondary cage is selected, then Doublewide GPU Cannot be populated on HPE DL380 G11 2U x8/x16/x8 Sec Riser Kit and Default Primary Riser.
- Full Length (FL) PCIe Card and Mid Tray cannot be selected together.
- For 12EDSFF CTO Server, If EDSFF Bundle is selected then this GPU cannot be selected.
- If this GPU is selected, then HPE DL380/DL560 Gen11 2U GPU Power Cable Kit must be selected and defaulted.
- For 8SFF CTO Server, If Qty1 of 8SFF Front cage is selected then Max of 2 GPU can be selected per Server.
- For 12EDSFF CTO Server, if no additional cage is selected then Max of 2 GPU can be selected per Server.
- For 8SFF CTO Server, If Qty2 of 8SFF Front cage is selected with NO DLC Component then this GPU is not supported.
- For 12EDSFF CTO Server, if additional cage is selected with NO DLC Component then this GPU is not supported
- For 8LFF CTO Server, If NO DLC component is selected then this GPU is not supported.
- For 8SFF CTO Server, If Qty2 of 8SFF Front cage is selected with DLC Component then Max 1 GPU can be selected per sever.
- For 12EDSFF CTO Server, if one additional 8SFF cage is selected with DLC Component then Max 1 GPU can be selected per sever.
- For 8LFF CTO Server, If DLC component is selected then Max 1 GPU can be selected per sever.
- For 8SFF CTO Server with two 8SFF cages/ 8LFF CTO Server/ 12EDSFF CTO Server with 8SFF Cage or EDSFF Bundle/ 12LFF CTO Server/ 24SFF CTO Server,
- If Qty1 or 2 of this GPU is selected with NO DLC Component then Secondary 3x16 cable cannot be selected
- If Qty1 of this GPU is selected then Secondary OR Tertiary Riser Must be selected. If Qty2 of this GPU is selected then Secondary AND Tertiary Riser Must be selected.
- GPU and NEBS Riser cannot be selected together.
- If Full Length (FL) GPU is selected then High Performance HS or DLC must be selected for Processors.
- If Full Length (FL) GPU is selected with this Processor then DLC Module must be selected for Processors.
- If Doublewide GPU is selected on default Primary Riser then the top x16 PCIe slot (FH FL) connector CANNOT be utilized as the Doublewide GPU occupies this space. The "PCIe Card Capacity Limits" must account for this limitation.
- If this GPU is selected then Secondary 3x16 cable cannot be selected.
- This GPU cannot be populated on primary Riser. The "PCIe Card Capacity Limits" must account for this limitation.
- For 8SFF CTO Server with two 8SFF cages/ 8LFF CTO Server/ 12EDSFF Server with one 8SFF Cage, If Qty1 this GPU is selected with DLC Module then Tertiary Riser must be selected. GPU is populated on Tertiary Riser.

NVIDIA L40 48GB PCIe Accelerator for HPE S0K90C

Notes:

- Max = 3
- Must be populated in x16 slot.
- System memory should be 2x of GPU memory.
- Requires selection of High-Performance Fan Kit.
- This GPU requires HPE ProLiant DL380/DL560 Gen11 2U GPU Power Cable Kit - P56072-

Additional Options

B21
- On 8SFF CTO Server, if Qty2 of 8SFF Front cage is selected without DLC component then Max of 2 GPU can be selected per Server.
- On 8SFF CTO Server, if Qty1 of 8SFF Front cage is selected then Max of 3 GPU can be selected per Server.
- On 8LFF CTO Server, if NO DLC component is selected then Max of 2 GPU can be selected per Server
- On 8LFF CTO Server, if DLC component is selected then Max of 3 GPU can be selected per Server.
- On 12LFF CTO Server, 24SFF CTO Server, or 8SFF CTO Server with Qty3 8SFF drive cages if NO DLC component is selected then Max of 1 GPU can be selected per Server
- On 12LFF CTO Server, 24SFF CTO Server, or 8SFF CTO Server with Qty3 8SFF drive cages and DLC component is selected then Max of 2 GPU can be selected per Server.
- On 12EDSFF CTO Server, if one additional 8SFF cage is selected without DLC Component then Max 2 GPU can be selected per sever.
- On 12EDSFF CTO Server, if one additional 8SFF cage is selected with DLC Component then Max 3 GPU can be selected per sever.
- On 8SFF CTO Server, if Qty2 of 8SFF Front cage is selected with DLC component then Max of 3 GPU can be selected per Server.

NVIDIA L4 24GB PCIe Accelerator for HPE S0K89C

Notes:

- Max = 8
- System memory should be 2x of GPU memory.
- Requires selection of High-Performance Fan Kit
- On 12LFF CTO Server, 24SFF CTO Server, or 8SFF CTO Server with Qty3 8SFF drive cages are selected without DLC component then Max of 5 can be selected per server

NVIDIA A16 64GB PCIe Non-CEC Accelerator for HPE R8T26C

Notes:

- Max = 3
- System memory should be 2x of GPU memory.
- This GPU requires power Cable Kit (P39102-B21) to also be selected.
- This GPU requires HPE DL380/DL560 G11 2U High Perf Fan Kit P48820-B21.

Intel Data Center GPU Max 1100 48GB Accelerator for HPE S1T66C

Notes:

- Max = 3
- This GPU is PCIe x16, so it must be populated on a x16 slot only. The "PCIe Card Capacity Limits" must account for this limitation.
- Mixing of GPU types is not allowed.
- Requires selection of High-Performance Fan Kit.
- System memory should be 2x of GPU memory.
- Quantity1 of this cable can support max of 3 double wide GPUs.
- When this GPU is selected then the top x16 PCIe slot (FH FL) connector on the riser kit CANNOT be utilized as the GPU occupies this space. The "PCIe Card Capacity Limits" must account for this limitation.
- Full Length (FL) PCIe Card and Mid Tray cannot be selected together.
- This GPU requires HPE DL380/DL560 Gen11 2U GPU Power Cable Kit (P56072-B21)
- If Qty2 of this GPU is selected, then Secondary Riser must be selected.
- For 12LFF CTO Server/ 24SFF CTO Server/ 8SFF CTO Server with front 3 8SFF Cage/ 12EDSFF CTO Server with EDSFF Bundle, if this GPU is selected then DLC Must be

Additional Options

selected.
– If Qty3 of DW GPU is selected, then primary 3 x16 Riser cable and secondary 3 x16 Riser Cable cannot be selected.

HPE ProLiant DL300 Gen10 Plus GPU 8-pin Keyed Cable Kit P39102-B21

Notes:

– This GPU power cable is used for NVIDIA A16.
– One power cable supports up to 3x GPUs

HPE ProLiant DL380/DL560 Gen11 2U GPU Power Cable Kit P56072-B21

Notes:

– This GPU power cable is used for NVIDIA H100 and NVIDIA L40 GPUs.
– One power cable supports up to 3x GPUs

GPU Information

Notes: The listed ambient temperatures represent the maximum levels allowed without any impact on system performance (e.g., throttling). It is possible to exceed these temperatures by an additional +5°C without causing a system shutdown, but performance degradation might be expected.

HPE DL380 Gen11 Configuration						
Part number	**Card**	**Qty Supported**	**PCIe**	**8SFF**	**16SFF/8LFF**	**24SFF/12LFF**
S1T66C	Intel® Data Center GPU Max 1100	2 or 3	Gen5	3@30C(Air)	3@25C(Air) 3@30C(Air) if CPU = 185W	Not Supported
R9S41C	NVIDIA H100 80GB PCIe Accelerator	2 or 3	Gen5	2@25C (Air) 3@25C (DLC)	2 @ 20C (Air) 2 @ 25C (DLC) 1 @ 25C (Air) if CPU = 185W	Not Supported
S2D86C	NVIDIA H100 NVL 94GB PCIe Accelerator	1, 2 or 3	Gen5	3@27C (Air) 2@30C (Air) 3@30C (DLC)	3@18C (Air) 2@27C (Air) 3@25C (DLC)	1@25C (Air) 2@25C (DLC)
S3T54C	NVIDIA RTX 4000 Ada Graphics Accelerator for HPE	up to 5	Gen4	3@30C(Air) 3@30C(DLC) 5@27C(Air) 5@27C(DLC)	3@30C(Air) 3@30C(DLC) 5@27C(Air) 5@27C(DLC)	3@30C(Air) 3@30C(DLC) 5@27C(Air) 5@27C(DLC)
S0K90C	NVIDIA L40 48GB PCIe Accelerator	1, 2 or 3	Gen4	3@25C(Air) 2@30C(Air) 3@30C(DLC)	2@30C(Air) 3@25C(DLC) 2@30C(Air) if CPU = 185W 3@25C(Air) if CPU = 185W	1@25C(Air) 2@25C(DLC) 2@25C(Air) if CPU = 185W
S2L70C	NVIDIA L40S 48GB	1 or 2	Gen5	2@25C (Air)	2 @ 20C (Air)	Not Supported

Additional Options

	PCIe Accelerator			3@25C (DLC)	2 @ 25C (DLC) 1 @ 25C (Air) if CPU = 185W 3 @ 25C (DLC) if CPU = 185W	
S0K89C	NVIDIA L4 24GB PCIe Accelerator	Up to 8	Gen4	8@30C(Air) 8@30C(Air)	5@30C(Air) 8@25C(Air) 8@30C(DLC)	5@25C(Air) 8@25C(DLC)
R8T26C	NVIDIA A16 64GB PCIe Non-CEC Accelerator for HPE	1, 2 or 3	Gen4	30C	2@30C(Air) 3@30C(DLC) 3@25C(Air) if CPU= 185W	1@30C(Air) 2@25C(Air) 3@25C(DLC) 3@25C(Air) if CPU= 185W

Qty SFF Drive Cages	Cooling Method	NVIDIA A16
1 drive cage	Air Cooled + Heatsink	3
2 drive cages	Air Cooled + Heatsink	2
3 drive cages	Air Cooled + Heatsink	1
1 drive cage	Direct Liquid Cooling	3
2 drive cages	Direct Liquid Cooling	3
3 drive cages	Direct Liquid Cooling	2

Embedded Management

HPE iLO Common Password FIO Setting

HPE iLO Common Password FIO Setting P08040-B21

Notes:

- Replaces iLO default randomized password by an HPE defined common password. HPE highly recommends changing this password immediately after the initial onboarding process.
- Customers who want to choose their own custom iLO default password should use the HPE Factory Express Integration Services

HPE ILO Advanced

HPE iLO Advanced Electronic License with 1yr Support on iLO Licensed Features E6U59ABE
HPE iLO Advanced 1-server License with 1yr Support on iLO Licensed Features 512485-B21
HPE iLO Advanced AKA Tracking License with 1yr Support on iLO Licensed Features 512487-B21
HPE iLO Advanced Electronic License with 3yr Support on iLO Licensed Features E6U64ABE
HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features BD505A
HPE iLO Advanced AKA Tracking License with 3yr Support on iLO Licensed Features BD507A

Additional Options

HPE Converged Infrastructure Management Software

HPE OneView including 3yr 24x7 Support Flexible Quantity E-LTU E5Y35AAE

HPE OneView w/o iLO including 3yr 24x7 Support Flexible Quantity E-LTU P8B26AAE

Notes: Licenses ship without media. The HPE OneView Media Kit can be ordered separately, or can be **downloaded**.

HPE Security

HPE Trusted Supply Chain for HPE ProLiant P36394-B21

Notes:

- HPE Trusted Supply Chain (P36394-B21) is an optional security upgrade intended for agencies and regulated industries needing enhanced security and compliance needs. Applying this option to a DL380 Gen11 CTO server ensures it is built in the USA in a secured facility by vetted HPE personnel assigned to the manufacturing processes. A multitude of checkpoints/inspections for malicious microcode and counterfeit parts are performed throughout the server build, and additional safeguards are put in place against cyber-exploits throughout the server lifecycle. The HPE ProLiant DL380 Gen11 is re-branded as a HPE ProLiant DL380T Gen11 to denote the HPE Trusted Supply Chain security enhancements. The DL380T Gen11 is Trade Agreement Act (TAA) compliant. Learn more at **https://www.hpe.com/security**
- This option requires the selection of HPE Gen11 Intrusion Detection Kit (P48922-B21)
- This option requires the selection of either HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features (BD505A) or HPE iLO Advanced 1-server License with 1yr Support on iLO Licensed Features (512485-B21)
- This option is limited to stand-alone DL380 Gen11 CTO servers only. The HPE Trusted Supply Chain configuration will not be available if the server is ordered as factory integrated into a rack
- One instance of the following Electronic License to Use is required per order (not per server): R6X85AAE
- HPE Trusted Supply Chain E-LTU
- Logistics delivery speeds and services are available and selectable within Next Gen Quoter.
- This option cannot be selected with TAA instruction SKU nor TAA CTO Models

HPE iLO Common Password FIO Setting P08040-B21

Notes:

- Replaces ILO default randomized password by an HPE defined common password. HPE highly recommends changing this password immediately after the initial onboarding process.
- Customers who want to choose their own custom iLO default password should use the HPE Factory Express Integration Services

HPE Gen11 2U Bezel Kit P50400-B21

HPE Bezel Lock Kit 875519-B21

Notes: Requires the bezel kit

HPE Boot Controllers

HPE NS204i-u Gen11 NVMe Hot Plug Boot Optimized Storage Device P48183-B21

Notes:

Additional Options

– This is the NS204i-u hot pluggable boot device
– Default is NVMe are internal to system and not hot pluggable
– If external accessible drives are needed please add trigger SKU P54542-B21 HPE ProLiant DL380 Gen11 NS204i-u FIO Bundle Kit. This trigger SKU allows NVMe drives to be externally accessible and hot pluggable.
– Max = 1

HPE ProLiant DL380 Gen11 NS204i-u Internal Cable Kit P52152-B21

Notes:

– If NS204i-u Gen11 Hot Plug Boot Option Device is selected then HPE DL380 Gen11 NS204i-u Internal Cable Kit is required.
– Max = 1

HPE ProLiant DL380 Gen11 NS204i-u FIO Bundle Kit P54542-B21

Notes:

– This SKU is required only when external accessible drives are required for the NS204i-u.
– Max = 1

HPE Storage Controllers

The Gen11 storage controller portfolio has been updated to include new technology like OCP3.0 as well as PCIe adapters. For a more detailed breakout of the available Gen11 controllers visit the storage controllers QuickSpecs site:

HPE MegaRAID Storage Controllers

HPE Tri-Mode Controllers

HPE MR416i-p Gen11 x16 Lanes 8GB Cache PCI SPDM Plug-in Storage Controller P47777-B21
HPE MR416i-o Gen11 x16 Lanes 8GB Cache OCP SPDM Storage Controller P47781-B21
HPE MR216i-p Gen11 x16 Lanes without Cache PCI SPDM Plug-in Storage Controller P47785-B21
HPE MR216i-o Gen11 x16 Lanes without Cache OCP SPDM Storage Controller P47789-B21
HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller P58335-B21
HPE ProLiant DL320/DL380 Gen11 PCIe Gen5 Retimer Card P48833-B21

Notes: When adding this Retimer as a field upgrade, it is recommended to update the UBM5 PIC to a minimum of version 1.10. This firmware can be downloaded at: **https://support.hpe.com/connect/s/softwaredetails?language=en_US&softwareId=MTX_23c46bfbef394b5e9353075c5e**

HPE SR932i-p Gen11 x32 Lanes 8GB Wide Cache PCI SPDM Plug-in Storage Controller P47184-B21

Notes: Requires x16 riser slot

Essential RAID Controllers

HPE Smart Array E208e-p SR Gen10 (8 External Lanes/No Cache) 12G SAS PCIe Plug-in Controller 804398-B21

NVMe Adapter

HPE DL385 Gen10 Plus 12Gb NVMe 2-port Adapter P25527-B21

Additional Options

HPE Cable Options

HPE ProLiant DL380 Gen11 12EDSFF CPU1/2 Cable Kit	P52153-B21
HPE ProLiant DL360 Gen11 Storage Controller Enablement Cable Kit	P48918-B21
HPE ProLiant DL380 Gen11 8SFF CPU1/2 NVMe Cable Kit	P48825-B21
HPE ProLiant DL380 Gen11 8SFF OROC1/2 x2 Cable Kit	P48829-B21
HPE ProLiant DL380 Gen11 2U Tri-Mode Premium Cable Kit	P48831-B21
HPE ProLiant DL380 Gen11 Tri-Mode Splitter Cable Kit	P48832-B21
HPE ProLiant DL380 Gen11 PCIe Gen5 Retimer Card Cable Kit	P52154-B21
HPE ProLiant DL380 Gen11 8SFF to Retimer/-P Controller Cable Kit	P54874-B21
HPE ProLiant DL380 Gen11 LFF Front Tri-Mode Cable Kit	P56995-B21

Optional Upgrades

HPE 96W Smart Storage Lithium-ion Battery with 145mm Cable Kit	P01366-B21
HPE Smart Storage Hybrid Capacitor with 145mm Cable Kit	P02377-B21

Notes: Provides backup power for multiple HPE storage controllers or other devices.

HPE Tape Backup

For the complete range of tape drives, autoloaders, libraries and media see:

https://www.hpe.com/us/en/storage/storeever-tape-storage.html.

For hardware and software compatibility of Hewlett Packard Enterprise tape backup products please visit the StoreEver Tape Solutions in SPOCK (requires registration/login)
https://h20272.www2.hpe.com/SPOCK/default.aspx.

Only external drives supported

All libraries and autoloaders supported via compatible FC or SAS controller. Refer to the StoreEver Tape Solutions Compatibility Matrix link above.

HPE Storage Options

Emulex Fibre Channel HBAs

HPE SN1610E 32Gb 1-port Fibre Channel Host Bus Adapter	R2J62A
HPE SN1610E 32Gb 2-port Fibre Channel Host Bus Adapter	R2J63A
HPE SN1700E 64Gb 1-port Fibre Channel Host Bus Adapter	R7N77A
HPE SN1700E 64Gb 2-port Fibre Channel Host Bus Adapter	R7N78A

QLogic Fibre Channel HBAs

HPE SN1610Q 32Gb 1-port Fibre Channel Host Bus Adapter	R2E08A
HPE SN1610Q 32Gb 2-port Fibre Channel Host Bus Adapter	R2E09A
HPE SN1700Q 64Gb 1-port Fibre Channel Host Bus Adapter	R7N86A
HPE SN1700Q 64Gb 2-port Fibre Channel Host Bus Adapter	R7N87A

HPE Racks

- Please see the HPE Advanced Series Racks QuickSpecs for information on additional racks options and rack specifications. **HPE G2 Advanced Series Racks**
- Please see the HPE Enterprise Series Racks QuickSpecs for information on additional racks options and rack specifications. **HPE G2 Enterprise Series Racks**

HPE Power Distribution Units (PDUs)

- Please see the **HPE Basic Power Distribution Units (PDU) QuickSpecs** for information on these products and their specifications.
- Please see the **HPE Metered Power Distribution Units (PDU) QuickSpecs** for information on these products and their specifications. Please see the **HPE Intelligent Power Distribution Unit (PDU)** QuickSpecs for information on these products and their specifications.
- Please see the **HPE Metered and Switched Power Distribution Units (PDU)** QuickSpecs for information on these products and their specifications.

HPE Uninterruptible Power Systems (UPS)

- To learn more, please visit the **HPE Uninterruptible Power Systems (UPS)** web page.
- Please see the **HPE Direct Flow Three Phase Uninterruptible Power System QuickSpecs** for information on these products and their specifications.
- Please see the **HPE Line Interactive Single Phase UPS QuickSpecs** for information on these products and their specifications.

HPE Rack Options

Please see the **HPE KVM Switches web page** for information on these products and their specifications.

Easy Install Rail Kits

Easy Install rail kits contain telescoping rails which allow for in-rack serviceability.

To assist in the installation of the server into the rack an optional installation tool is available by contacting your local services representative.

Notes:

- Hewlett Packard Enterprise recommends that a minimum of two people are required for all Rack Server installations. Please refer to your installation instructions for proper tools and number of people to use for any installation.
- HPE rail kits are designed to work with HPE racks in compliance with industry standard EIA-310-E. In the event a customer elects to purchase a third-party rack for use with an HPE rail kit, any such use is at customer's own risk. HPE makes no express or implied warranties with respect to such third-party racks and specifically disclaims any implied warranties of merchantability and fitness for a particular purpose. Furthermore, HPE has no obligation and assumes no liability for the materials, design, specifications, installation, safety, and compatibility of any such third-party racks with any rail kits, including HPE rail kits.

Additional Options

HPE ProLiant DL3XX Gen11 Easy Install Rail 3 Kit	P52341-B21

Notes: Does not include Cable Management Arm (CMA) (P22020-B21).

HPE DL38X Gen10 Plus 2U Cable Management Arm for Rail Kit	P22020-B21

HPE USB and SD Options

Notes:

- In vSphere 7.0, VMware made changes that impact the use of an SD Card/USB media as a standalone boot device and will be removing support for them after version 7.x.
- SD Card/USB media can still be used as a standalone boot option through all 7.x releases via published Customer Advisory **Usage of SD Card/USB Devices As Standalone Boot Devices Has Changed Due to System Storage Changes For VMware ESXi 7.0 (Or Later)**.
- For any major release beyond VMware ESXi 7.x, VMware will require M.2 or another local persistent device as the standalone boot option.

HPE USB Keyboard/Mouse Kits

HPE USB CN Keyboard/Mouse Kit	631364-B21
HPE USB KR Keyboard/Mouse Kit	672097-KD3

HPE Support Services

Installation & Startup Services

HPE ProLiant DL/ML Install Service	U4554E
HPE ProLiant DL/ML Startup Service	U4555E

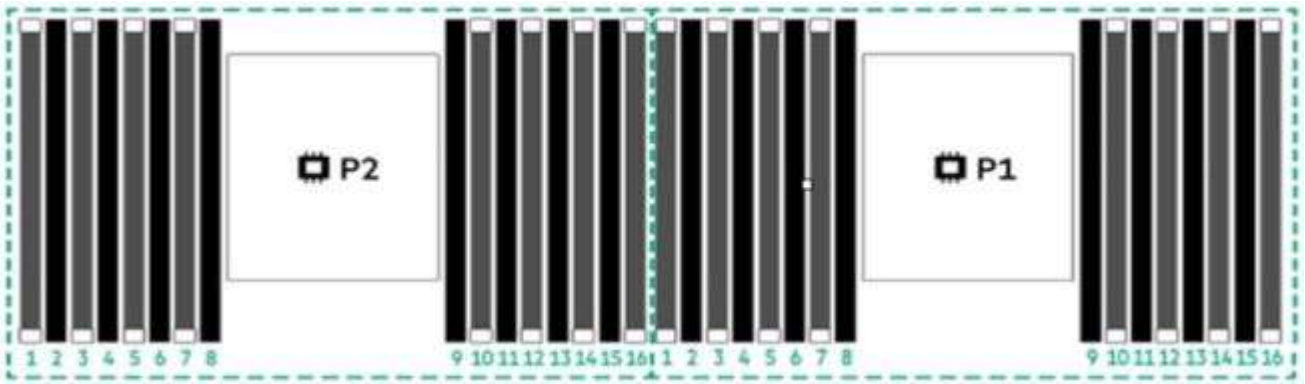
Tech Care Services

HPE 3 Year Tech Care Essential DL380 Gen11 HW Service	H93G4E
HPE 3 Year Tech Care Essential wDMR DL380 Gen11 HW Service	H93G5E
HPE 5 Year Tech Care Essential DL380 Gen11 HW Service	H93J8E
HPE 5 Year Tech Care Essential wDMR DL380 Gen11 HW Service	H93J9E

Notes: For a full listing of support services available for this server, please visit **https://www.hpe.com/services**.

Memory

Memory Population guidelines



HPE ProLiant DL380 Gen11

HPE ProLiant Gen11 16 slot per CPU DIMM population order																
DIMM population order																
DIMM slot	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**
1 DIMM										10						
2 DIMMs[2]			3							10						
4 DIMMs[2]			3				7			10				14		
6 DIMMs			3		5		7			10				14		16
8 DIMMs[1,2]	1		3		5		7			10		12		14		16
12 DIMMs	1	2	3		5	6	7			10	11	12		14	15	16
16 DIMMs[1,2]	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16

Notes:

– Ommited DIMM counts/socket not qualified by Intel.
– [1] Supports SGX (Software Guard Extensions)
– [2] Support Hemi (hemisphere mode).

General Memory Population Rules and Guidelines:

- DIMMs should be installed in quantities of even numbers.
- Install DIMMs only if the corresponding processor is installed.
- If only one processor is installed in a two-processor system, only half of the DIMM slots are available.
- To maximize performance, it is recommended to balance the total memory capacity between all installed processors.
- When two processors are installed, balance the DIMMs across the two processors.
- White DIMM slots denote the first slot to be populated in a channel.
- Mixing of DIMM types (UDIMM, RDIMM, and LRDIMM) is not supported.
- Mixing of x4 and x8 memory is not allowed.
- The maximum memory speed is a function of the memory type, memory configuration, and processor model.
- The maximum memory capacity is a function of the number of DIMM slots on the platform, the largest DIMM capacity qualified on the platform, and the number and model of installed processors qualified on the platform.
- For details on the HPE Server Memory Options Population Rules, visit:

Memory

Server memory populations rules for HPE Gen11 servers with 4th Gen Intel Xeon Scalable processors

- To realize the performance memory capabilities listed in this document, HPE DDR5 Smart Memory is required.
- For additional information, please see the **HPE DDR5 Smart Memory QuickSpecs**.

HPE SKU P/N	P43322-B21	P43328-B21	P43331-B21
SKU Description	HPE 16GB (1x16GB) Single Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	HPE 32GB (1x32GB) Dual Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	HPE 64GB (1x64GB) Dual Rank x4 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit
DIMM Capacity	16GB	32GB	64GB
DIMM Rank	Single Rank (1R)	Dual Rank (2R)	Dual Rank (2R)
Voltage	1.1 V	1.1 V	1.1 V
DRAM Depth [bit]	2G	2G	4G
DRAM Width [bit]	x8	x8	x4
DRAM Density	16Gb	16Gb	16Gb
CAS Latency	40-39-39	40-39-39	40-39-39
DIMM Native Speed	4800 MT/s	4800 MT/s	4800 MT/s

HPE SKU P/N	P43334-B21	P43337-B21
SKU Description	HPE 128GB (1x128GB) Quad Rank x4 DDR5-4800 CAS-46-39-39 EC8 Registered 3DS Smart Memory Kit	HPE 256GB (1x256GB) Octal Rank x4 DDR5-4800 CAS-46-39-39 EC8 Registered 3DS Smart Memory Kit
DIMM Capacity	128GB	256GB
DIMM Rank	Quad Rank (4R)	Octal Rank (8R)
Voltage	1.1 V	1.1 V
DRAM Depth [bit]	4G	4G
DRAM Width [bit]	x4	x4
DRAM Density	16Gb	16Gb
CAS Latency	40-39-39	40-39-39
DIMM Native Speed	4800 MT/s	4800 MT/s

Notes: The maximum memory speed is a function of the memory type, memory configuration, and processor model.

For details on the HPE Server Memory speed, visit: **https://www.hpe.com/docs/memory-speed-table**

DDR5 memory options part number decoder

Notes:

– Capacity references are rounded to the common gigabyte (GB) values.

- 8GB = 8,192 MB
- 16GB = 16,384 MB
- 32GB = 32,768 MB
- 64GB = 65,536 MB

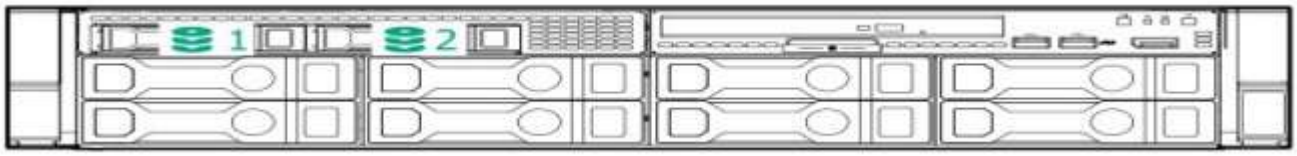
Memory

- 96GB = 98,304 MB
- 128GB = 131072 MB
- 256GB = 262144 MB
- 512GB = 524288 MB

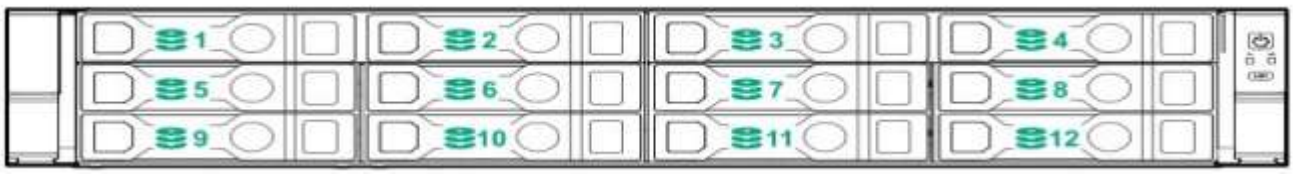
For more information on memory, please see the Memory QuickSpecs: **HPE DDR5 Smart Memory**

Memory Speed Table for HPE ProLiantDL380 Gen Gen11

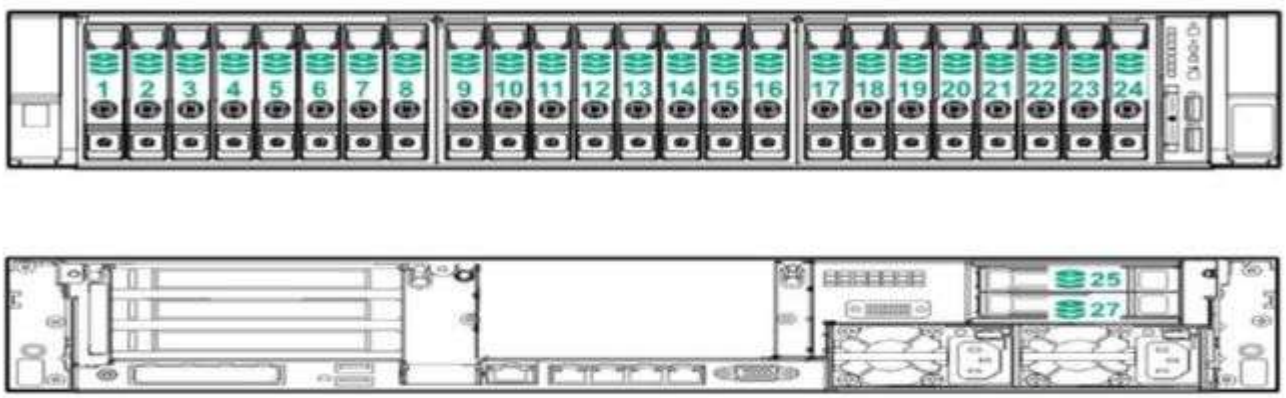
For details on the HPE Server Memory speed, please visit: **https://www.hpe.com/docs/server-memory**



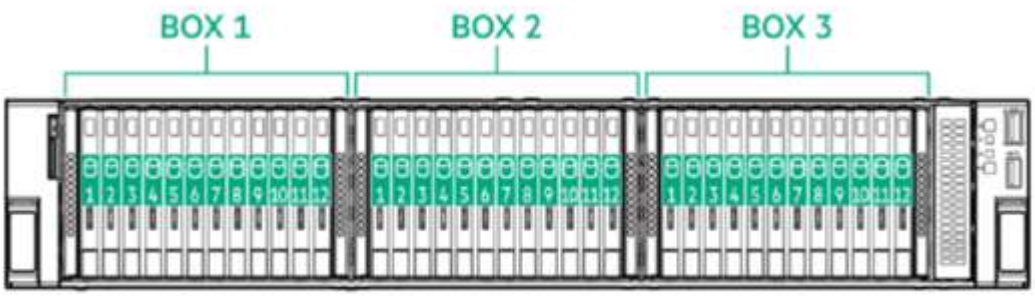
8LFF chassis with Universal media bay and optional 2SFF and optical drive shown



12LFF chassis



24SFF + rear 2 SFF drives



12EDSFF chassis

System Unit

Dimensions

- **SFF CTO servers:**
  8.75 x 44.8 x 72.7. cm / 3.44 x 17.64 x 28.62 in
- **LFF CTO servers:**
  8.75 x 44.8 x 73.25 cm / 3.44 x 17.64 x 28.84 in

Weight (approximate)

- **Maximum:** 8 SFF hard drives (no rear drives), 2x processors, 2x power supplies, 1x RAID controller, 2x Risers installed)
  - **Maximum:** 33kg/72.75 lbs.
  - **Minimum:** 16kg/35.27 lbs.
- **Maximum:** 12LFF hard drives (no rear drives), 2x processors, 2x power supplies, 1x RAID controller, 2x Risers installed)
  - **Maximum:** 37kg/81.57 lbs.
  - **Minimum:** 18kg/39.68 lbs.

Input Requirements (per power supply)

Rated Line Voltage

- For 1800W-2200W (Titanium) Power Supply: 200-240 VAC
- For 1600W (Platinum) Power Supply: 200-240 VAC
- For 1000W (Titanium): 100 to240 VAC
- For 800W (Titanium) Power Supply: 200-240 VAC
- For 800W (Platinum) Power Supply: 100-240 VAC
- For 800W (Universal) Power Supply: 200-277 VAC

BTU Rating

Maximum

- For 1800W-2200W Power Supply: 6497 BTU/hr. (at 200 VAC), 6868 BTU/hr. (at 208 VAC), 7230 BTU/hr. (at 220 VAC), 7596 BTU/hr. (at 230VAC), 7962 BTU/hr. (at 240VAC)
- For 1600W Power Supply: 5918 BTU/hr. (at 200 VAC), 5888 BTU/hr. (at 220 VAC), 5884 BTU/hr. (at 240 VAC)
- For 1000W (Titanium) Power Supply: 3741 BTU/hr. (at 100 VAC), 3596 BTU/hr. (at 200 VAC), 3582 BTU/hr. (at 240 VAC)
- For 800W (Titanium) Power Supply: 2905 BTU/hr. (at 200 VAC), 2899 BTU/hr. (at 220 VAC), 2893 BTU/hr. (at 240 VAC)
- For 800W (Platinum) Power Supply: 3067 BTU/hr. (at 100 VAC), 2958 BTU/hr. (at 200 VAC), 2949 BTU/hr. (at 240 VAC)
- For 800W (Universal) Power Supply: 2964 BTU/hr. (at 200 VAC), 2951 BTU/hr. (at 230 VAC), 2936 BTU/hr. (at 277 VAC)

Relative Humidity (non-condensing)

- **Operating**

8% to 90% - Relative humidity (Rh), 28°C maximum wet bulb temperature, non-condensing.

- **Non-operating**

5 to 95% relative humidity (Rh), 38.7°C (101.7°F) maximum wet bulb temperature, non-condensing.

Power Supply Output

(per power supply)

Rated Steady-State Power

- For 1800W-2200W Power Supply: 1800W (at 200 VAC), 1900W (at 208 VAC), 2000W (at 220 VAC), 2100W (at 230VAC), 2200W (at 240VAC)
- For 1600W Power Supply: 1600W (at 240 VAC), 1600W (at 240 VDC) for China only
- For 1000W (Titanium) Power Supply: 1097W (at 100 VAC), 1000W (at 240 VAC), 1000W (at 240 VDC) input for China only
- For 800W (Titanium) Power Supply: 800W (at 200 VAC), 800W (at 240 VAC), 800W (at 240 VDC) for China only
- For 800W (Platinum) Power Supply: 800W (at 100 VAC), 800W (at 240 VAC), 800W (at 240 VDC) input for China only
- For 800W (Universal) Power Supply: 800W (at 200 VAC), 800W (at 277 VAC)

Maximum Peak Power

- For 1800W-2200W Power Supply: 1800W (at 200 VAC), 1900W (at 208 VAC), 2000W (at 220 VAC), 2100W (at 230VAC), 2200W (at 240VAC)
- For 1600W Power Supply: 1600W (at 240 VAC), 1600W (at 240 VDC) for China only
- For 1000W (Titanium) Power Supply: 1000W (at 100 VAC), 1000W (at 240 VAC), 1000W (at 240 VDC) input for China only
- For 800W (Titanium) Power Supply: 800W (at 200 VAC), 800W (at 240 VAC), 800W (at 240 VDC) for China only
- For 800W (Platinum) Power Supply: 800W (at 100 VAC), 800W (at 240 VAC), 800W (at 240 VDC) input for China only
- For 800W (Universal) Power Supply: 800W (at 200 VAC), 800W (at 277 VAC)

System Inlet Temperature

- **Standard Operating Temperature**

10° to 35°C (50° to 95°F) at sea level with an altitude derating of 1.0°C per every 305 m (1.8°F per every 1000 ft.) above sea level to a maximum of 3050 m (10,000 ft.), no direct sustained sunlight. Maximum rate of change is 20°C/hr. (36°F/hr.). The upper limit and rate of change may be limited by the type and number of options installed.

System performance during standard operating support may be reduced if operating with a fan fault or above 30°C (86°F).

Technical Specifications

- **Extended Ambient Operating Temperature**

For approved hardware configurations, the supported system inlet range is extended to be: 5° to 10°C (41° to 50°F) and 35° to 40°C (95° to 104°F) at sea level with an altitude derating of 1.0°C per every 175 m (1.8°F per every 574 ft.) above 900 m (2953 ft.) to a maximum of 3050 m (10,000 ft.). The approved hardware configurations for this system are listed at the URL: **DL380 Gen11 Extended Ambient Temperature Guidelines**

For approved hardware configurations, the supported system inlet range is extended to be: 40° to 45°C (104° to 113°F) at sea level with an altitude derating of 1.0°C per every 125 m (1.8°F per every 410 ft.) above 900 m (2953 ft.) to a maximum of 3050 m (10,000 ft.). The approved hardware configurations for this system are listed at the URL:
DL380 Gen11 Extended Ambient Temperature Guidelines

System performance may be reduced if operating in the extended ambient operating range or with a fan fault.

- **Non-operating**

-30° to 60°C (-22° to 140°F). Maximum rate of change is 20°C/hr. (36°F/hr.).

Altitude

- **Operating**

3050 m (10,000 ft.). This value may be limited by the type and number of options installed. Maximum allowable altitude change rate is 457 m/min (1500 ft./min).

- **Non-operating**

9144 m (30,000 ft.). Maximum allowable altitude change rate is 457 m/min (1500 ft./min).

Acoustic Noise

Listed are the declared A-Weighted sound power levels (LwA,m) and declared average bystander position A-Weighted sound pressure levels (LpA,m) when the product is operating in a 23°C ambient environment. Noise emissions were measured in accordance with ISO 7779 (ECMA 74) and declared in accordance with ISO 9296 (ECMA 109). The listed sound levels apply to standard shipping configurations. Additional options may result in increased sound levels. Please have your HPE representative provide information from the HPE EMESC website for further technical details regarding the configurations listed below.

Technical Specifications

Acoustic Noise	
Idle	
LwA,m	4.2 B Entry 4.2 B Base 4.2 B Performance
LpAm	28 dBA Entry 27 dBA Base 30 dBA Performance
Operating	
LwA,m	4.2 B Entry 4.2 B Base 4.2 B Performance
LpAm	29 dBA Entry 27 dBA Base 29 dBA Performance
Kv	0.4 B Entry 0.4 B Base 0.4 B Performance

Notes:

- The declared mean A-weighted sound power level, LwA,m, is computed as the arithmetic average of the measured.
- A-weighted sound power levels for a randomly selected sample, rounded to the nearest 0,1 B.
- The declared mean A-weighted emission sound pressure level, LpA,m, is computed as the arithmetic average of the measured A-weighted emission sound pressure levels at the bystander positions for a randomly selected sample, rounded to the nearest 1 dB.
- The statistical adder for verification, Kv, is a quantity to be added to the declared mean A-weighted sound power level, LwA,m, such that there will be a 95 % probability of acceptance, when using the verification procedures of ISO 9296, if no more than 6,5 % of the batch of new equipment, has A-weighted sound power levels greater than (LwA,m + Kv).
- The quantity, LwA,c (formerly called LwAd), can be computed from the sum of LWA,m and Kv.
- All measurements made to conform to ISO 7779 / ECMA-74 and declared to conform to ISO 9296 / ECMA-109.
- B, dB, abbreviations for bels and decibels, respectively, where 1 B = 10 dB.
- The results in this declaration apply only to the model numbers listed above when operating and tested according to the indicated modes and standards. A system with additional configuration components or increased operating functionality may increase the noise emission values.
- System under abnormal conditions may increase the noise level, persons in the vicinity of the product [cabinet] for extended periods of time should consider wearing hearing protection or using other means to reduce noise exposure.

Emissions Classification (EMC) - Regulatory Information

To view the regulatory information for your product, view the Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products, available at the Hewlett Packard Enterprise Support Center:

https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts

---

Environment-friendly Products and Approach End-of-life Management and Recycling

Hewlett Packard Enterprise offers **end-of-life product return, trade-in, and recycling programs**, in many geographic areas, for our products. Products returned to Hewlett Packard Enterprise will be recycled, recovered or disposed of in a responsible manner.

The EU WEEE directive (2002/95/EC) requires manufacturers to provide treatment information for each product type for use by treatment facilities. This information (product disassembly instructions) is posted on the Hewlett Packard Enterprise web site. These instructions may be used by recyclers and other WEEE treatment facilities as well as Hewlett Packard Enterprise OEM customers who integrate and re-sell Hewlett Packard Enterprise equipment.

Summary of Changes

Date	Version History	Action	Description of Change
06-Jan-2025	Version 29	Changed	Overview, Core Options and Storage sections were updated.
02-Dec-2024	Version 28	Changed	Core Options and Additional Options sections were updated. (Removed OBS items: HDDs, SSDs, iLO Licenses, and Mouse/Keyboard options).
04-Nov-2024	Version 27	Changed	Standard Features, Configuration Information and Additional Options sections were updated.
07-Oct-2024	Version 26	Changed	Standard Features, Core Options and Additional Options sections were updated.
03-Sep-2024	Version 25	Changed	Standard Features (Operating Systems and Virtualization Software Support for HPE Servers) and Pre-Configured Models sections were updated.
05-Aug-2024	Version 24	Changed	Standard Features, Core Options and Additional Options sections were updated.
15-Jul-2024	Version 23	Changed	Pre-Configured Models section was updated.
01-Jul-2024	Version 22	Changed	Standard Features, Core Options and Additional Options sections were updated.
17-Jun-2024	Version 21	Changed	Pre-Configured Models section was updated.
03-Jun-2024	Version 20	Changed	Standard Features, Configuration Information, Core Options, Additional Options and Technical Specifications sections were updated.
29-Apr-2024	Version 19	Changed	Core Options section was updated.
15-Apr-2024	Version 18	Changed	Pre-Configured Models section was updated.
01-Apr-2024	Version 17	Changed	Standard Features, Pre-Configured Models, Configuration Information, Core Options and Additional Options sections were updated.
18-Mar-2024	Version 16	Changed	Pre- Configured Models section was updated.
04-Mar-2024	Version 15	Changed	Configuration Information and Core Options sections were updated
05-Feb-2024	Version 14	Changed	Overview, Standard Features, Configuration Information, Core Options and Technical Specifications sections were updated
14-Dec-2023	Version 13	Changed	Overview, Configuration Information and Additional Options sections were updated
06-Nov-2023	Version 12	Changed	Overview, Configuration Information and Additional Options sections were updated
02-Oct-2023	Version 11	Changed	Overview, Configuration Information and Additional Options sections were updated
11-Sep-2023	Version 10	Changed	Standard Features section was updated
05-Sep-2023	Version 9	Changed	Overview, Standard Features, Configuration Information, Core Options and Memory sections were updated
24-Jul-2023	Version 8	Changed	Standard Features section was updated
10-Jul-2023	Version 7	Changed	Overview, Standard Features, Pre- Configured Models, Configuration Information and Core Options sections were updated
05-Jun-2023	Version 6	Changed	Overview, Standard Features, Pre- Configured Models, Configuration Information and Core Options sections were updated
01-May-2023	Version 5	Changed	Standard Features, Configuration Information, Core Options and Technical Specifications sections were updated.
03-Apr-2023	Version 4	Changed	Standard Features, Pre- Configured Models, Configuration Information, Core Options and Technical Specifications sections were updated.
06-Mar-2023	Version 3	Changed	Standard Features and Core Options sections were updated.
15-Feb-2023	Version 2	Changed	Standard Features section was updated
10-Jan-2023	Version 1	New	New QuickSpecs

Copyright

Make the right purchase decision. Contact our presales specialists.



© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Intel® and Xeon® are registered trademarks of Intel Corporation in the U.S. and other countries.

Microsoft®, Windows®, and Windows Server® are U.S. registered trademarks of the Microsoft group of companies.

For hard drives, 1GB = 1 billion bytes. Actual formatted capacity is less

Hewlett Packard Enterprise

a50004307enw - 16911 - Worldwide - V29 - 06-January-2025

Remote Desktop sign-in with Windows Hello for Business

Article • 06/11/2024 • Applies to: ✅ Windows 11, ✅ Windows 10

You can use Windows Hello for Business to sign in to a remote desktop session, using the redirected smart card capabilities of the Remote Desktop Protocol (RDP). This is possible by deploying a certificate to the user's device, which is then used as the supplied credential when establishing the RDP connection to another Windows device.

This article describes two certificate deployment approaches, where authentication certificates are deployed to the Windows Hello for Business container:

- Using Microsoft Intune with SCEP or PKCS connectors
- Using an Active Directory Certificate Services (AD CS) enrollment policy

> **Important**
>
> If you deploy the certificate using Microsoft Intune, and you have User Account Control configured to *Prompt for credentials on the secure desktop*, you won't be able to use the *run as* feature. In such scenario, when you try to execute an application with elevated privileges and choose the Windows Hello for Business credential, you'll receive the error message: **The username or password is incorrect**.

> **Tip**
>
> Consider using Remote Credential Guard instead of Windows Hello for Business for RDP sign-in. Remote Credential Guard provides single sign-on (SSO) to RDP sessions using Kerberos authentication, and doesn't require the deployment of certificates. For more information, see Remote Credential Guard.

How it works

Windows generates and stores cryptographic keys using a software component called a *key storage provider* (KSP):

- Software-based keys are created and stored using the *Microsoft Software Key Storage Provider*
- Smart card keys are created and stored using the *Microsoft Smart Card Key Storage Provider*
- Keys created and protected by Windows Hello for Business are created and stored using the *Microsoft Passport Key Storage Provider*

A certificate on a smart card starts with creating an asymmetric key pair using the Microsoft Smart Card KSP. Windows requests a certificate based on the key pair from your enterprises issuing certificate authority, which returns a certificate that is stored in the user's Personal certificate store. The private key remains on the smart card and the public key is stored with the certificate. Metadata on the certificate (and the key) stores the key storage provider used to create the key (remember the certificate contains the public key).

The same concept applies to Windows Hello for Business, except that the keys are created using the Microsoft Passport KSP. The user's private key remains protected by the device's security module (TPM) and the user's gesture (PIN/biometric). The certificate APIs hide the complexity. When an application uses a certificate, the certificate APIs locate the keys using the saved key storage provider. The key storage providers direct the certificate APIs on which provider they use to find the private key associated with the certificate. This is how Windows knows you have a smart card certificate without the smart card inserted, and prompts you to insert the smart card.

Windows Hello for Business emulates a smart card for application compatibility, and the Microsoft Passport KSP prompts the user for their biometric gesture or PIN.

> ⓘ **Note**
>
> Remote Desktop with biometric doesn't work with **Dual Enrollment** or scenarios where the user provides alternative credentials.

Requirements

Here's a list of requirements to enable RDP sign-in with Windows Hello for Business:

- ✓ A PKI infrastructure based on AD CS or third-party
- ✓ Windows Hello for Business deployed to the clients

- ✓ If you plan to support Microsoft Entra joined devices, the domain controllers must have a certificate, which serves as a *root of trust* for the clients. The certificate ensures that clients don't communicate with rogue domain controllers

If you plan to deploy certificates using Microsoft Intune, here are more requirements:

- ✓ Ensure you have the infrastructure to support either SCEP or PKCS deployment
- ✓ Deploy the root CA certificate and any other intermediate certificate authority certificates to Microsoft Entra joined Devices using a Trusted root certificate policy

Create a certificate template

The process of creating a certificate template is applicable to scenarios where you use an on-premises Active Directory Certificate Services (AD CS) infrastructure.
You must first create a certificate template, and then deploy certificates based on that template to the Windows Hello for Business container.

The certificate template configuration is different depending on whether you deploy certificates using Microsoft Intune or an Active Directory enrollment policy. Select the option that best suits your needs.

Microsoft Intune

1. Sign in to your issuing certificate authority (CA) and open *Server Manager*
2. Select **Tools** > **Certification Authority**. The Certification Authority Microsoft Management Console (MMC) opens
3. In the MMC, expand the CA name and right-click **Certificate Templates** > **Manage**
4. The Certificate Templates console opens. All of the certificate templates are displayed in the details pane
5. Right-click the **Smartcard Logon** template and select **Duplicate Template**
6. Use the following table to configure the template:

Expand table

Tab Name	Configurations
Compatibility	• Clear the **Show resulting changes** check box • Select **Windows Server 2012 or Windows Server 2012 R2** from the *Certification Authority list* • Select **Windows Server 2012 or Windows Server 2012 R2** from the *Certification Recipient list*
General	• Specify a **Template display name**, for example *WHfB Certificate Authentication* • Set the validity period to the desired value • Take note of the template name for later, which should be the same as the Template display name minus spaces (*WHfBCertificateAuthentication* in this example)
Extensions	Verify the **Application Policies** extension includes **Smart Card Logon**.
Subject Name	Select **Supply in the request**.
Request Handling	• Set the Purpose to **Signature and smartcard logon** and select **Yes** when prompted to change the certificate purpose • Select the **Renew with same key** check box • Select **Prompt the user during enrollment** **Note:** If you deploy certificates with a PKCS profile, select the option **Allow private key to be exported**
Cryptography	• Set the Provider Category to **Key Storage Provider** • Set the Algorithm name to **RSA** • Set the minimum key size to **2048** • Select **Requests must use one of the following providers** • Select **Microsoft Software Key Storage Provider** • Set the Request hash to **SHA256**
Security	Add the security principal used for SCEP or PKCS **Enroll** access

7. Select **OK** to finalize your changes and create the new template. Your new template should now appear in the list of Certificate Templates

8. Close the Certificate Templates console

Issue the certificate template

1. In the Certificate Authority console, right-click **Certificate Templates**, select **New** > **Certificate Template to Issue**
2. From the list of templates, select the template you previously created (**WHFB Certificate Authentication**) and select **OK**. It can take some time for the template to replicate to all servers and become available in this list
3. After the template replicates, in the MMC, right-click in the Certification Authority list, select **All Tasks** > **Stop Service**. Right-click the name of the CA again, select **All Tasks** > **Start Service**

Deploy certificates

The process of deploying certificates is different depending on whether you use Microsoft Intune or an Active Directory enrollment policy. Select the option that best suits your needs.

Microsoft Intune

This section describes how to configure a SCEP policy in Intune. Similar steps can be followed to configure a PKCS policy.

1. Go to the Microsoft Intune admin center
2. Select **Devices** > **Configuration profiles** > **Create profile**
3. Select **Platform** > **Windows 10 and later** and **Profile type** > **Templates** > **SCEP Certificate**
4. Select **Create**
5. In the *Basics* panel, provide a **Name** and, optionally, a **Description** > **Next**
6. In the *Configuration settings* panel, use the following table to configure the policy:

Expand table

Setting	Configurations
Certificate Type	User
Subject name format	`CN={{UserPrincipalName}}` **Note:** if there's a mismatch between the user UPN suffix and the Active Directory domain FQDN, use `CN={{OnPrem_Distinguished_Name}}` instead.
Subject alternative name	From the dropdown, select **User principal name (UPN)** with a value of `{{UserPrincipalName}}`
Certificate validity period	Configure a value of your choosing
Key storage provider (KSP)	**Enroll to Windows Hello for Business, otherwise fail**
Key usage	**Digital Signature**
Key size (bits)	**2048**
For Hash algorithm	**SHA-2**
Root Certificate	Select **+Root Certificate** and select the trusted certificate profile created earlier for the Root CA Certificate
Extended key usage	• *Name:* **Smart Card Logon** • *Object Identifier:* `1.3.6.1.4.1.311.20.2.2` • *Predefined Values:* **Not configured** • *Name:* **Client Authentication** • *Object Identifier:* `1.3.6.1.5.5.7.3.2` • *Predefined Values:* **Client Authentication**
Renewal threshold (%)	Configure a value of your choosing
SCEP Server URLs	Provide the public endpoint(s) that you configured during the deployment of your SCEP infrastructure

7. Select **Next**

8. In the *Assignments* panel, assign the policy to a security group that contains as members the devices or users that you want to configure and select **Next**
9. In the *Applicability Rules* panel, configure issuance restrictions, if needed, and select **Next**
10. In the *Review + create* panel, review the policy configuration and select **Create**

For more information how to configure SCEP policies, see Configure SCEP certificate profiles in Intune. To configure PKCS policies, see Configure and use PKCS certificate with Intune.

> **⊗ Caution**
>
> If you deploy certificates via Intune and configure Windows Hello for Business via group policy, the devices will fail to obtain a certificate, logging the error code `0x82ab0011` in the `DeviceManagement-Enterprise-Diagnostic-Provider` log.
> To avoid the error, configure Windows Hello for Business via Intune instead of group policy.

Use non-Microsoft certification authorities

If you're using a non-Microsoft PKI, the certificate templates published to the on-premises Active Directory may not be available. For guidance with integration of Intune/SCEP with non-Microsoft PKI deployments, refer to Use non-Microsoft certification authorities (CA) with SCEP in Microsoft Intune.

As an alternative to using SCEP, or if none of the previously covered solutions work in your environment, you can manually generate Certificate Signing Requests (CSR) for submission to your PKI. To assist with this approach, you can use the Generate-CertificateRequest PowerShell commandlet.

The `Generate-CertificateRequest` commandlet generates an `.inf` file for a pre-existing Windows Hello for Business key. The `.inf` can be used to generate a certificate request manually using `certreq.exe`. The commandlet also generates a `.req` file, which can be submitted to your PKI for a certificate.

Verify that the certificate is deployed

To verify that the certificate is correctly deployed to the Windows Hello for Business container, use the following command:

Windows Command Prompt

```
certutil -store -user my
```

The output lists keys and certificates stored in the user store. If a certificate issued from your CA is deployed to the Windows Hello for Business container, the output displays the certificate with a `Provider` value of `Microsoft Passport Key Storage Provider`.

For example:

Windows Command Prompt

```
C:\Users\amanda.brady>certutil -store -user my
my "Personal"
================ Certificate 0 ================
Serial Number: 110000001f4c4eccc46fc8f93a00000000001f
Issuer: CN=Contoso - Issuing CA, DC=CONTOSO, DC=COM
 NotBefore: 12/8/2023 6:16 AM
 NotAfter: 12/7/2024 6:16 AM
Subject: CN=amanda.brady@contoso.com
Non-root Certificate
Template:
1.3.6.1.4.1.311.21.8.2835349.12167323.7094945.1118853.678601.83.11484210.800573
9
Cert Hash(sha1): 63c6ce5fc512933179d3c0a5e94ecba98092f93d
Key Container = S-1-12-1-../../login.windows.net/../amanda.brady@contoso.com
Provider = Microsoft Passport Key Storage Provider
Private key is NOT exportable
Encryption test passed
```

User experience

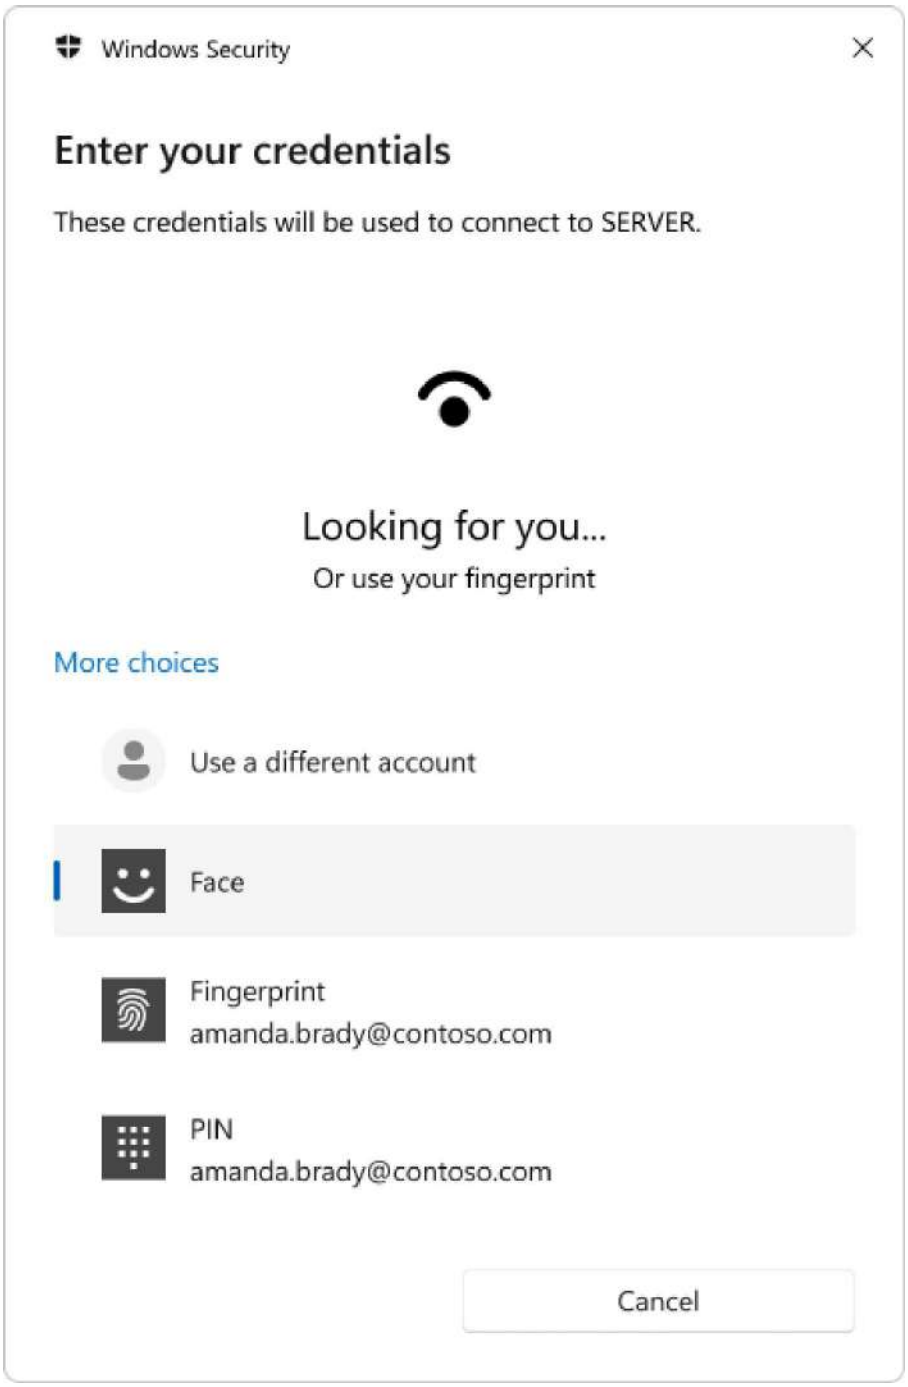
Once users obtain their certificate, they can RDP to any Windows devices in the same Active Directory forest as the users' Active Directory account by opening Remote Desktop Connection (`mstsc.exe`). When connecting to the remote host, they're prompted to use Windows Hello for Business to unlock the private key of the certificate.

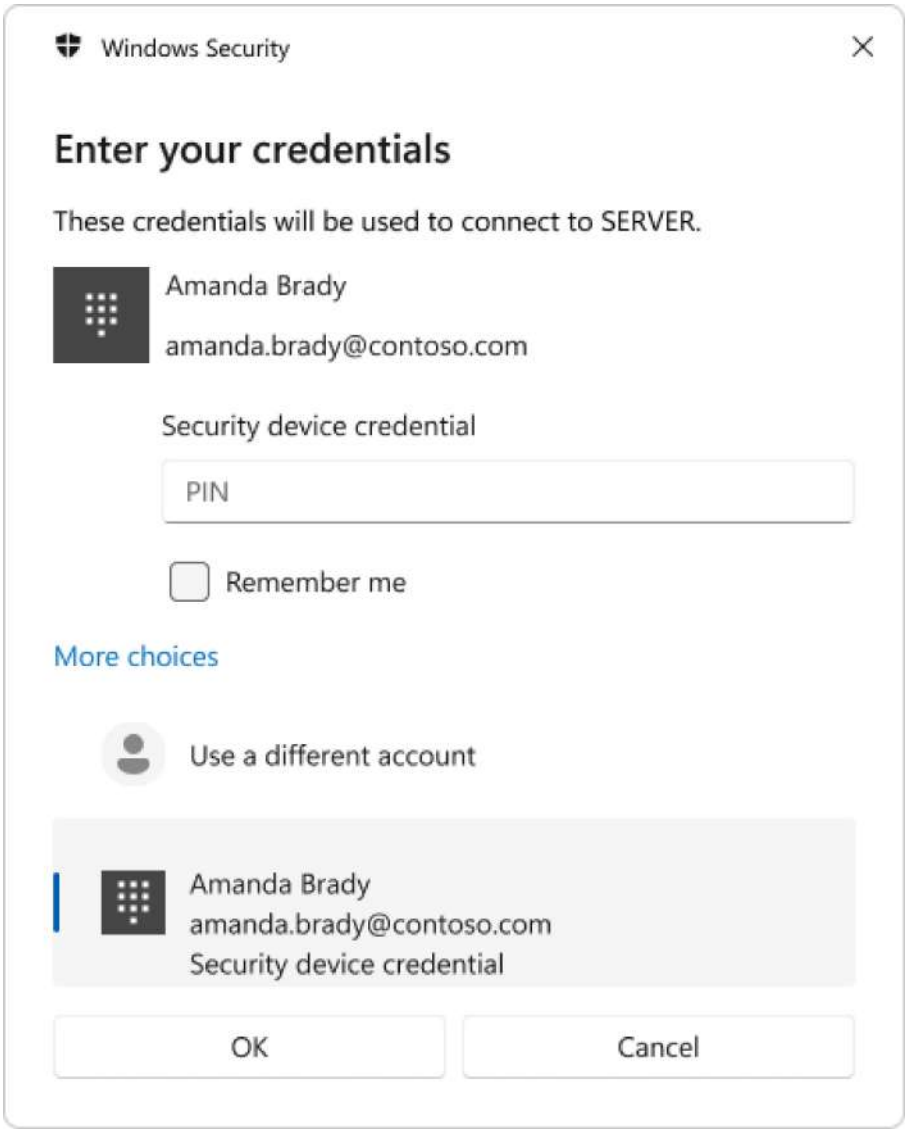
Microsoft Entra joined device

The user can authenticate using any available Windows Hello unlock gestures, including biometrics.

Microsoft Entra hybrid joined device

The credential prompt identifies the Windows Hello credential provider as *Security device credential*. The user must use the PIN credential provider to unlock.





Here's a brief video showing the user experience from a Microsoft Entra joined device using fingerprint as unlock factor: https://learn-video.azurefd.net/vod/player?id=b6e1038d-98b5-48dc-8afb-65523d12cfaf&locale=en-us&embedUrl=%2Fwindows%2Fsecurity%2Fidentity-protection%2Fhello-for-business%2Frdp-sign-in

> **Note**
>
> The user must be authorized to connect to the remote server using the Remote Desktop protocol, for example by being a member of the Remote Desktop Users local group on the remote host.

Compatibility

While users appreciate the convenience of biometrics, and administrators value the security, you might experience compatibility issues with applications and Windows Hello for Business certificates. In such scenarios, you can deploy a policy setting to revert to the previous behavior for the users needing it.

For more information, see Use Windows Hello for Business certificates as smart card certificate

Known issues

There's a known issue when attempting to perform TLS 1.3 client authentication with a Hello certificate via RDP. The authentication fails with the error: `ERR_SSL_CLIENT_AUTH_SIGNATURE_FAILED`. Microsoft is investigating possible solutions.

Feedback

Was this page helpful? Yes No

Provide product feedback

SFF Committee documentation may be purchased in hard copy or electronic form.
SFF specifications are available at ftp://ftp.seagate.com/sff

SFF Committee

SFF-8431 Specifications for

Enhanced Small Form Factor Pluggable Module SFP+

Revision 4.1

6th of July 2009

Secretariat: SFF Committee

Abstract: This document defines the low speed electrical and management interface specifications for enhanced Small Form Factor Pluggable (SFP+) modules and hosts. The SFP+ module is a hot pluggable small footprint serial-to-serial data-agnostic optical transceiver. This document defines the high speed electrical interface specifications for 10 Gigabit/s SFP+ modules and hosts. The 8.5 Gigabit/s high speed electrical interface specifications are defined in FC-PI-4. The modules may optionally support lower signalling rates as well. The modules may be used to implement single-mode or multimode serial optical interfaces at 850 nm, 1310 nm, or 1550 nm. The SFP+ module design may use one of several different optical connectors.

This specification provides a common reference for system manufacturers, system integrators, and suppliers. This is an internal working specification of the SFF Committee, an industry ad hoc group.

This specification is made available for public review, and written comments are solicited from readers. Comments received by the members will be considered for inclusion in future revisions of this specification.

Support: This specification is supported by the identified member companies of the SFF Committee.

POINTS OF CONTACT:

Technical Editor:
Ali Ghiasi
Broadcom Corporation
3151 Zanker Road
San Jose, CA 95134
Voice: (408) 922-7423
EMail: aghiasi@broadcom.com

Chairman SFF Committee
I. Dal Allan
ENDL
14426 Black Walnut Court
Saratoga, CA 95070
Voice: (408) 867-6630
EMail: endlcom@acm.org

SSWG Leader:
Tom Lindsay
ClariPhy Communications
16 Technology Drive, Suite 165
Irvine, CA 92618
Voice: (425)608-0209
EMail: tom.lindsay@clariphy.com

EXPRESSION OF SUPPORT BY MANUFACTURERS

The following member companies of the SFF Committee voted in favor of this industry specification:

AMCC
Amphenol
Arista Networks
Avago
Broadcom
Cinch
Clariphy
Cortina Systems
EMC
Emulex
ETRI
Finisar
Foxconn
Fujitsu CPA
Hewlett Packard
Hitachi GST
JDS Uniphase
Luxtera
Molex
NetLogic uSyst
OpNext
Panduit
Picolight
QLogic
Samsung
Sumitomo
Sun Microsystems
Tyco
Vitesse Semiconductor
W L Gore

The following member companies of the SFF Committee voted to abstain on this industry specification.

Comax
FCI
ICT Solutions
Leoni Cables
Meritec
Seagate
Toshiba

The user's attention is called to the possibility that implementation to this Specification may require use of an invention covered by patent rights. By distribution of this specification, no position is taken with respect to the validity of a claim or claims of any patent rights in connection therewith. Members of the SFF Committee which advise that a patent exists are required to provide a statement of willingness to grant a license under these rights on reasonable and non-discriminatory terms and conditions to applicants desiring to obtain such a license.

Foreword

The development work on this specification was done by the SFF Committee, an industry group. The membership of the committee since its formation in August 1990 has included a mix of companies which are leaders across the industry.

When 2 1/2" diameter disk drives were introduced, there was no commonality on external dimensions e.g. physical size, mounting locations, connector type, connector location, between vendors.

The first use of these disk drives was in specific applications such as laptop portable computers and system integrators worked individually with vendors to develop the packaging. The result was wide diversity, and incompatibility.

The problems faced by integrators, device suppliers, and component suppliers led to the formation of the SFF Committee as an industry ad hoc group to address the marketing and engineering considerations of the emerging new technology.

During the development of the form factor definitions, other activities were suggested because participants in the SFF Committee faced more problems than the physical form factors of disk drives. In November 1992, the charter was expanded to address any issues of general interest and concern to the storage industry. The SFF Committee became a forum for resolving industry issues that are either not addressed by the standards process or need an immediate solution.

Those companies which have agreed to support a specification are identified in the first pages of each SFF Specification. Industry consensus is not an essential requirement to publish an SFF Specification because it is recognized that in an emerging product area, there is room for more than one approach. By making the documentation on competing proposals available, an integrator can examine the alternatives available and select the product that is felt to be most suitable.

SFF Committee meetings are held during T10 weeks (see www.t10.org), and Specific Subject Working Groups are held at the convenience of the participants. Material presented at SFF Committee meetings becomes public domain, and there are no restrictions on the open mailing of material presented at committee meetings.

Most of the specifications developed by the SFF Committee have either been incorporated into standards or adopted as standards by EIA (Electronic Industries Association), ANSI (American National Standards Institute) and IEC (International Electrotechnical Commission).

The SFF Committee activities support the requirements of the storage industry, and it is involved with several standards.

If you are interested in participating or wish to follow the activities of the SFF Committee, the signup for membership and/or documentation can be found at:

www.sffcommittee.com/ie/join.html

The complete list of SFF Specifications which have been completed or are currently being worked on by the SFF Committee can be found at:

ftp://ftp.seagate.com/sff/SFF-8000.TXT

If you wish to know more about the SFF Committee, the principles which guide the activities can be found at:

ftp://ftp.seagate.com/sff/SFF-8032.TXT

Suggestions for improvement of this specification will be welcome. They should be sent to the SFF Committee, 14426 Black Walnut Ct, Saratoga, CA 95070.

Industry Documents

The following interface standards are relevant to SFP+ Specifications.
SFF-8432 - Improved Pluggable Formfactor (ftp://ftp.seagate.com/sff/)
SFF-8083 - Improved 0.8mm Card Edge Connector (ftp://ftp.seagate.com/sff/)
SFF-8089 - SFP Rate and Application codes (ftp://ftp.seagate.com/sff/)
SFF-8079 - SFP Rate and Application Selection (ftp://ftp.seagate.com/sff/)
SFF-8472 - Diagnostic Monitoring Interface for Optical Transceivers (ftp://ftp.seagate.com/sff/)
INF-8074i - SFP (Small Form Factor) Transceiver (ftp://ftp.seagate.com/sff/)
INF-8077i - 10 Gigabit Small Form Factor Pluggable Module (XFP MSA) (ftp://ftp.seagate.com/sff/)
FC-PI-4 - Fibre Channel - Physical Interface-4
10GFC - Fibre Channel - 10 Gigabit
FC-MJSQ - Methodologies for Jitter and Signal Quality Specifications FC INCITS Project 1316-DT Rev 14.1, June 5, 2005
IEEE 802.3 - IEEE Standard 802.3, Relevant 10 Gigabit Ethernet clauses are 49, 10GBASE-R LAN PHY; 50, 10GBASE-W WAN PHY; 52, 10 Gigabit Ethernet serial PMDs; and 68, 10GBASE-LRM)
OIF CEI - Optical Internetworking Forum - Implementation Agreement # OIF-CEI-02.0 Common Electrical I/O (CEI) (http://www.oiforum.com/public/impagreements.html)

Acronyms and other abbreviations

64B/66B	Data encoded with 64B/66B encoder as defined by the IEEE Std. 802.3 CL 49.
BER	bit error ratio
CDR	clock and data recovery
CRU	clock recovery unit
dB	decibel. 10*log10(ratio of power quantities). Powers can be electrical or optical. Conventional usage. See also dBe and dBo.
dBe	Specific case of dB where signals are electrical. 10*log10(ratio of electrical power quantities). 20*log10(ratio of voltage quantities) can be used if reference impedances are equal.
dBm	decibel (relative to 1 mW)
dBo	Specific case of dB where the signals are in optical power. 10*log10(ratio of optical power quantities). Also, in certain cases with electrical signals relating to linear optical modules, where it is expected that electrical voltage is in proportion to optical power, 10*log10(ratio of voltage quantities).
DCD	Duty cycle distortion
DDPWS	Data Dependent Pulse Width Shrinkage
DDJ	Data Dependent Jitter
dRN	Difference of Relative noise see Appendix D
DUT	device under test
dWDP	Difference of the waveform distortion penalty of an optical receiver
dWDPc	Difference of the waveform distortion penalty of an electrical cable assembly
EMC	electromagnetic compatibility
EMI	electromagnetic Interference
FC	Fibre Channel
h	hexadecimal notation
HCB	Host Compliance Board
IEEE	Institute of Electrical and Electronics Engineers
ITU-T	ITU Telecommunication Standardization Sector
Gbit	gigabit = 10^9 bits
GBd	Gigabaud
J2	99% Jitter
LRM	IEEE 802.3 CL68 Physical Layer Specifications for 10Gb/s using 10GBASE-R encoding and long wavelength optics for multimode fiber

MCB	Module Compliance Board
OMA	optical modulation amplitude
PCB	printed circuit board
PRBS9	Pseudo-Random Bit Sequence 2^9-1, see D.1.1
PRBS31	Pseudo-Random Bit Sequence 2^{31}-1, see D.1.1
Q_{sq}	Q_{sq} a measure of SNR, see D.8 and IEEE 802.3.68.6.7
RI	random interference
RMS	root mean square
RN	relative noise
Rx	receiver
Rx_LOS	Loss of signal same as defined in FC PI-4 and the inverse of signal detect (SD) in 802.3
RSS	"Root Sum of Squares"
SD	Signal Detect
SerDes	Serializer/Deserializer
SFI	SFP+ high speed serial electrical interface
SNR	signal-to-noise ratio
VccT	Module positive power supply rail for the transmitter
VccR	Module positive power supply rail for the receiver
VMA	voltage modulation amplitude
Tx	transmitter
TWDP	Transmitter Waveform Distortion Penalty for an optical transmitter
TWDPc	Transmitter Waveform Distortion Penalty of a host transmitter supporting an electrical cable assembly
UI	unit interval = 1 symbol period
UJ	Uncorrelated Jitter
WDP	Waveform distortion penalty
WDPc	Waveform distortion penalty for an electrical cable assembly

SFF Specifications

There are several projects active within the SFF Committee. The complete list of specifications which have been completed or are still being worked on are listed in the specification at ftp://ftp.seagate.com/sff/SFF-8000.TXT

Document Sources

Those who join the SFF Committee as an Observer or Member receive electronic copies of the minutes and SFF specifications (http://www.sffcommittee.com/ie/join.html).

Copies of ANSI standards may be purchased from the InterNational Committee for Information Technology Standards (http://tinyurl.com/c4psg).

Copies of SFF, T10 (SCSI), T11 (Fibre Channel) and T13 (ATA) standards and standards still in development are available on the HPE version of CD_Access (http://tinyurl.com/85fts).

Conventions

The American convention of numbering is used i.e., the thousands and higher multiples are separated by a comma and a period is used as the decimal point. This is equivalent to the ISO/IEC convention of a space and comma.

American:	ISO:
0.6	0,6
1,000	1 000
1,323,462.9	1 323 462,9

SFP+ Publication History

Revision Number	Description	Date
0.1	Initial Publication of Document, Preliminary	May 9, 2006
0.5	Initial Public Review Draft	June 29, 2006,
1.0	2nd Public Draft	August 28, 2006
1.1	3rd Public Draft	October 10, 2006
1.2	4rd Public Draft	December 21, 2006
1.3	5rd Public Draft	February 16, 2007
2.0	6th Public Draft and the 1st SFF A ballot	April 26, 2007
2.1	7th Public Draft	August 30, 2007
2.2	8th Public Draft	December 19, 2007
3.0	9th Public Draft and the 2nd SFF A ballot	May 8, 2008
3.1	Editor review draft	November 11, 2008
3.2	10th Public Draft and the 3rd SFF A ballot	November 12, 2008
3.3	11th Public Draft and the 4rd SFF A ballot	April 8, 2009
4.0	12th Public Draft and the 1st SFF Publication Ballot Valid comments were made on draft revision 3.3 that mask hit ratios of 1e-12 at C' and C" compliance points are too time consuming for normal testing. While changes to the document to enable extrapolation from 1e-12 to a higher hit ratio would be desirable, due to concerns that developing a better specification would delay the schedule, no solutions are provided in this revision of SFF-8431. Improved testing methods should be considered for future projects.	June 11, 2009
4.1	Published It was proposed to truncate the diamond mask into a hexagonal mask while maintaining the current slope of the diamond mask. The underlying reason for this comment was that the downstream CDR due to setup and hold times could not use the upper and lower apexes of the diamond mask. It was considered with interest but it could not be accommodated with the schedule.	July 6, 2009

Contents

Chapter 1 Scope of SFP+ Specification 1
1.1 Introduction 1
1.2 The SFP+ Supported Standards 2
1.3 SFI Typical PCB Reach (Informative) 3

Chapter 2 Low Speed Electrical and Power Specifications 4
2.1 Introduction 4
2.2 General Requirements 4
2.3 SFP+ Host Connector Definition 4
2.4 Low Speed Electrical Control Contacts and 2-wire Interface 7
2.4.1 Tx_Fault 7
2.4.2 Tx_Disable 7
2.4.3 RS0/RS1 7
2.4.4 Mod_ABS 8
2.4.5 SCL/SDA 8
2.4.6 Rx_LOS 8
2.5 Rate Select Hardware Control 8
2.6 Low Speed Electrical Specifications 9
2.6.1 Module Low Speed Electrical Specifications 9
2.6.2 Host Low speed Electrical Specifications 10
2.7 Timing Requirement of Control and Status I/O 11
2.7.1 Module power on initialization procedure, Tx_Disable negated 12
2.7.2 Module power on initialization procedure, Tx_Disable asserted. 12
2.7.3 Initialization During Hot Plugging 13
2.7.4 Transmitter management 14
2.7.5 Transmitter safety detection and presentation 14
2.7.6 Module fault recovery 15
2.7.7 Module loss of signal indication 16
2.8 SFP+ Power Requirement 17
2.8.1 Module Power Supply Requirements 17
2.8.2 Host Power Supply Noise Output 17
2.8.3 Module Power Supply Noise Output 17
2.8.4 Power Supply Noise Tolerance 17
2.9 ESD 19

Chapter 3 High Speed Electrical Specification SFI 20
3.1 Introduction 20
3.2 SFI Applications Definition 20
3.3 SFI Test Points Definition and Measurements 21
3.3.1 Host Compliance Points 21
3.3.2 Module Compliance Points 22
3.3.3 ASIC/SerDes Test Points (Informative) 23
3.3.4 Host Input Calibration Point 24
3.3.5 Module Input Calibration Point 24
3.4 SFI Termination and DC Blocking 25
3.5 SFP+ Host System Specifications 27
3.5.1 Host Transmitter Output Specifications at B 27
3.5.2 Host Receiver Input Specifications at C and C″ 30
3.6 SFP+ Module Specifications 35
3.6.1 Module Transmitter Input Specifications at B′ and B″ 35
3.6.2 Module Receiver Output Specifications at C′ 37

Chapter 4 SFP+ 2-wire Interface 41
4.1 Introduction 41
4.2 2-wire Electrical Specifications 41
4.3 SFP+ 2-wire Timing Diagram 42
4.4 Memory Transaction Timing 43
4.5 Device Addressing and Operation 44
4.6 Read/Write Functionality 45
4.6.1 SFP+ Memory Address Counter (Read and Write Operations) 45
4.6.2 Read Operations (Current Address Read) 46
4.6.3 Read Operations (Random Read) 46
4.6.4 Read Operations (Sequential Read) 47
4.6.5 Write Operations (Byte Write) 48
4.6.6 Write Operations (Sequential Write) 49
4.6.7 Write Operations (Acknowledge Polling) 49

Appendix A SFI Channel Recommendation (Informative) 50
A.1 SFI Host Channel General Recommendations 50
A.2 SFI Channel Transfer Recommendations 50
A.3 SFI Channel Return Loss Recommendations 52
A.4 SFI Channel Ripple Recommendations 52

Appendix B SFI ASIC/SerDes Specification (Informative) 54
B.1 Introduction 54
B.2 SFI ASIC/SerDes Transmitter Output Specifications at A (Informative) 54

B.3 SFI ASIC/SerDes Receiver Input Specifications at D (Informative) 54

Appendix C Application Reference Boards (Normative) 56

C.1 Compliance Boards 56
C.1.1 Host Compliance Board Transfer Characteristics 56
C.1.2 Module Compliance Board Transfer Characteristics 57
C.1.3 ASIC/SerDes Test Board Transfer Characteristics 58
C.2 Host Compliance Board 59
C.2.1 Host Compliance Board Material and Layer Stack-up 59
C.2.2 Host Compliance Board Partlist 60
C.2.3 HCB Gerber Files 60
C.2.4 Schematic of Host Compliance Board 60
C.3 Module Compliance Board 62
C.3.1 Module Compliance Board Material and Layer Stack-up 62
C.3.2 Schematic of Module Compliance Board 63
C.3.3 Module Compliance Board Partlist 65
C.3.4 MCB Gerber files 65
C.4 Specifications for Mated Host and Module Compliance Boards 66

Appendix D Test Methodology and Measurement (Normative) 72

D.1 Introduction 72
D.1.1 Test Patterns 72
D.2 Eye Mask Compliance 72
D.2.1 Example Calculations for 5x10-5 Hit Ratio 73
D.3 Data Dependent Jitter (DDJ) and Pulse Width Shrinkage (DDPWS) 74
D.3.1 Duty Cycle Distortion (DCD) 75
D.4 Uncorrelated Jitter (UJ) 75
D.5 99% Jitter (J2) and Total Jitter (TJ) 76
D.6 Rise and Fall Times 77
D.7 Voltage Modulation Amplitude (VMA) 77
D.8 Relative Noise (RN) 78
D.9 Waveform Distortion Penalty (WDP) 79
D.10 Electrical compliance signal at B″ for the SFP+ Module transmitter 80
D.11 Test Method for a host receiver for a limiting module 82
D.11.1 Test Equipment & Setup 83
D.11.2 Stressed-Eye Jitter Characteristics 83
D.11.3 Calibration 85
D.11.4 Calibration Procedure 86
D.11.5 Test Procedure 87
D.12 Limiting Module Receiver Compliance Tests 87
D.13 Test Method for a Host Receiver with a Linear Module 88

D.13.1 Test Description and Procedure for Host Receiver for Linear Module 89
D.13.2 Host Linear Tester calibration 90
D.14 Linear Module Receiver Compliance Tests 91
D.14.1 Linear Module receiver noise compliance test 91
D.14.2 linear Module Receiver Distortion Penalty Compliance Test 93
D.14.3 Linear Module Receiver Output Differential Peak-Peak Voltage 94
D.15 AC Common Mode Voltage 95
D.15.1 Definition of AC Common Mode Voltage 95
D.15.2 AC Common Mode Generation Test 95
D.15.3 AC Common Mode Tolerance Test 95
D.16 Termination Mismatch 95
D.17 Power Supply Testing methodology 96
D.17.1 Host Power Supply Noise Output 97
D.17.2 SFP+ Module Power Supply Noise Output 99
D.17.3 Module Power Supply Tolerance Testing 99

Appendix E SFP+ Direct Attach Cable Specifications "10GSFP+Cu" (Optional) 101
E.1 10GSFP+Cu Direct Attach Construction 101
E.2 SFP+ Host Output Specifications for Passive Direct Attach Cables 102
E.2.1 Transmitter Stressor 102
E.3 SFP+ Host Receiver Supporting 10GSFP+Cu Input Compliance Test Signal Calibrated at C"104
E.3.1 Copper Host Receiver Specifications 104
E.3.2 Copper Host Stress Generator 1 UI Pulse Response 105
E.4 SFP+ Passive Direct Attach cable Assembly Specifications 107
E.4.1 SFP+ Direct Attach Cable Test Setup 108
E.4.2 Cable dWDP Test Procedure 109
E.4.3 Cable NEXT Measurement Procedure 109
E.4.4 VMA to Crosstalk Ratio (VCR) 110

Appendix F 1.25 GBd Operation Support (Optional) 112
F.1 Introduction 112
F.2 SFP+ Host Operation Guideline for Supporting Classic SFP 112

Appendix G Matlab Code for TWDP 115

List of Tables

Table 1 SFP+ Standard Compliance 2
Table 2 Host Board Achievable Trace Length 3
Table 3 SFP+ Module and Host Electrical contact definition 6
Table 4 Rate Select Hardware Control Contacts 9
Table 5 Low Speed Module Electrical Specifications 10
Table 6 Low Speed Host Electrical Specifications 10
Table 7 Timing Parameters for SFP+ Management 11
Table 8 SFP+ Module Power Supply Requirements 18
Table 9 SFI Supported signalling Rates 20
Table 10 SFI Reference Points 21
Table 11 Host Transmitter Output Electrical Specifications at B 28
Table 12 Host Transmitter Output Jitter and Eye Mask Specifications at B 29
Table 13 Host Receiver Input Electrical Specifications at C and C" 30
Table 14 Host receiver supporting limiting module input compliance test signal calibrated at C" .. 31
Table 15 Host receiver supporting linear module input compliance test signal calibrated at C" 34
Table 16 Module Transmitter Input Electrical Specifications at B' 35
Table 17 Module Transmitter Input Tolerance Signal Calibrated at B" 36
Table 18 Module Receiver Output Electrical Specifications at C' 37
Table 19 Limiting Module Receiver Output Jitter and Eye Mask Specifications at C' 38
Table 20 Linear Module Receiver Specifications at C' 39
Table 21 2-Wire Interface Electrical Specifications 41
Table 22 SFP+ 2-wire Timing Specifications 43
Table 23 SFP+ Memory Specifications 43
Table 24 SFP+ Device Address Word 45
Table 25 SFI Host Interconnect Budget 50
Table 26 ASIC/SerDes Transmitter Output Electrical Specifications at A 55
Table 27 ASIC/SerDes Receiver Electrical Input Specifications at D 55
Table 28 Host Compliance Board Part List 60
Table 29 Module Compliance Board Part List 65
Table 30 Estimated parameter values for an ideal stressed signal generator 82
Table 31 Target RNi Values 93
Table 32 Truncated Filter Response Coefficients 98
Table 33 SFP+ Host Transmitter Output Specifications at B for Cu 102
Table 34 10GSFP+Cu TWDPc Stressor 103
Table 35 10GSFP+ Host receiver input stress Generator at C" 104
Table 36 Stress Generator 1 UI Pulse Response with 8x Over-sampling 106
Table 37 10GSFP+Cu Cable Assembly Specifications at B' and C' 107

Table 38 INF-8074i Voltage Levels for Reference Only 112
Table 39 SFP+ Host Transmitter Requirements to Support 1.25 GBd Mode 113
Table 40 SFP+ Host Receiver Requirements to Support 1.25 GBd Mode 113
Table 41 SFP Module input and output ranges that can be supported by the SFP+ Host 114

List of Figures

Figure 1 Host PCB SFP+ pad assignment top view 5
Figure 2 SFP+ module contact assignment 5
Figure 3 Power on initialization of module, Tx_Disable negated 12
Figure 4 Power on initialization of module, Tx_Disable asserted 13
Figure 5 Example of initialization during hot plugging, Tx_Disable negated. 13
Figure 6 Management of module during normal operation, Tx_Disable implemented 14
Figure 7 Occurrence of condition generating Tx_Fault 14
Figure 8 Successful recovery from transient safety fault condition 15
Figure 9 Unsuccessful recovery from safety fault condition 16
Figure 10 Timing of Rx_LOS detection 16
Figure 11 Instantaneous and sustained peak current for VccT or VccR 19
Figure 12 SFI Application Reference Model 20
Figure 13 Host Compliance Board 22
Figure 14 Module Compliance Board 23
Figure 15 ASIC/SerDes Test Board 23
Figure 16 Host input calibration point C" and crosstalk source calibration point B" 24
Figure 17 Module input calibration point B" and crosstalk source calibration point C" 25
Figure 18 SFI Termination and AC Coupling 26
Figure 19 Transmitter Differential Output Compliance Mask at B and B" 29
Figure 20 Host Receiver Input Compliance Mask at C" Supporting Limiting Module 32
Figure 21 SR and LR Host Sinusoidal Jitter Tolerance Mask 32
Figure 22 Module Transmitter Differential Input Compliance Mask at B" 36
Figure 23 Limiting Module Receiver Differential Output Compliance Mask at C' 38
Figure 24 LRM module receiver RN and dWDP compliance and host receiver test calibration 40
Figure 25 2-wire Timing Diagram 42
Figure 26 Detail of Clock Stretching 42
Figure 27 Current Address Read Operation 46
Figure 28 Random Read 47
Figure 29 Sequential Address Read Starting at Current Word Address 47
Figure 30 Sequential Address Read Starting with Random Read 48
Figure 31 SFP+ Write Byte Operation 48
Figure 32 Sequential Write Operation 49
Figure 33 Example of SFI Host Recommended Channel 51
Figure 34 Approximate Response of Host Compliance Board 57
Figure 35 Approximate Transfer Response of Module Compliance Board 58
Figure 36 Host Compliance Board stack-up 59
Figure 37 Schematic of the Host Compliance Board 61
Figure 38 Module Compliance Board stack up 62
Figure 39 Schematic of The Module Compliance Board 64
Figure 40 Measurements Port Definition 66
Figure 41 Maximum Differential Response of mated MCB and HCB 67
Figure 42 Mated MCB-HCB Differential Through Response Limits 68

Figure 43 Maximum Common Mode Response of mated MCB and HCB 69
Figure 44 Maximum Differential to Common Mode Response of mated MCB and HCB 70
Figure 45 Maximum Differential NEXT Response of mated MCB and HCB 71
Figure 46 Eye mask measurement setup - block diagram 73
Figure 47 DDJ Test Method 74
Figure 48 Example xMA waveform showing xMA measurement windows 78
Figure 49 Compliance signal generator for module transmitter 81
Figure 50 Jitter Tolerance Test Configuration 84
Figure 51 Stressed eye jitter components 84
Figure 52 TP3 to electrical adaptor for host that operates with linear modules 89
Figure 53 Linear Module Receiver Noise Test 91
Figure 54 Module receiver waveform penalty compliance test 93
Figure 55 AC Termination Mismatch Measurement 96
Figure 56 Module Compliance Board Power Supply Filters 97
Figure 57 Reference filter response 98
Figure 58 Power Supply Noise Tolerance Test Setup 100
Figure 59 10GSFP+Cu Direct Attach Block Diagram 102
Figure 60 10GSFP+Cu TWDPc Stressor Impulse Response 103
Figure 61 Block Diagram of Copper Stressor Noise Model 105
Figure 62 Stress Generator 1UI Pulse Response with 8x over-sampling 105
Figure 63 10GSFP+ Cable Test Setup 108
Figure 64 10GSFP+Cu Cable NEXT dWDP Test Setup 109
Figure 65 SFP+ Host Transmitter Output Mask for 1.25 GBd Operation 113
Figure 66 SFP+ Host Receiver Input Mask for 1.25 GBd Operation 114

CHAPTER 1 SCOPE OF SFP+ SPECIFICATION

1.1 INTRODUCTION

This specification defines the electrical interfaces and their test methods between the SFP+ module and host board for operation up to 11.1 GBd. The high speed electrical interface between the host and SFP+ module is called "SFI". SFI simplifies the module and leverages host based transmit pre-emphasis and host based receive equalization to overcome PCB and external media impairments.

SFI typically operates with one connector at the module interface and up to about 200 mm of improved FR4 material or 150 mm of standard FR4, see 1.3. The electrical interface is based on high speed, low voltage AC coupled logic with a nominal differential impedance of 100 Ω.

The SFP+ specifications includes management, connector[1], mechanical[2], low speed signalling, high speed signalling, and appendices providing parameter and test board definitions, and implementation and measurement descriptions.

SFP+ modules are hot pluggable and active connections are powered by individual power connections for the transmitter (VccT) and the receiver (VccR). Multiple modules can share a single 3.3 V power supply with individual filtering for each VccT and VccR. Detailed power supply specifications are given in 2.8.

All SFP+ module compliance points are defined and measured through the mated reference test card as defined by C.3. All SFP+ host compliance points are defined and measured through the mated reference test card as defined by C.2.

The SFP+ module could be an electrical-to-optical or an electrical-to-electrical device intended to support one or more of the applications listed in Table 1.

It is expected that a range of SFP+ modules will operate on single-mode fiber, multimode fiber, and SFP+ electrical cable assemblies.

SFP+ compliant hosts are permitted to support just linear modules, just limiting modules, or both linear and limiting modules. Linear modules are modules which contain a linear receiver. Limiting modules are modules which contain a limiting receiver. Although not required, host supporting linear spec-

1. Defined in SFF-8083
2. Defined in SFF-8432

ifications are encouraged to support 10GSFP+Cu direct attach cables (Appendix E). For other copper variants see SFF-8461.

1.2 The SFP+ Supported Standards

An SFP+ module may comply with any combination of the standards shown in Table 1, and may be suitable for other or future standards. This specification does not preclude operation at other signalling rates not listed in this table, such as 2.125 GBd for 2GFC, or 4.25 GBd for 4GFC.

Due to the possibility of insertion of classic SFP modules into a host designed for SFP+ the damage threshold of the host for the input signal at C (see Figure 13) shall be at least 2000 mV peak to peak differential.

Table 1 SFP+ Standard Compliance

Standard	*Signalling Rate (GBd)*	*High Speed Serial Interface*	*High Speed Serial Test Method*	*Low Speed Electrical Definitions*	*Low Speed Test Methods*	*Management*	*Mechanical/ Connector*
IEEE 802.3 Clause 38 or Clause 59 (1 Gb/s Ethernet)	1.25	802.3 Clause 38 or 59 Appendix F		Chapter 2	Appendix D	Chapter 4 SFF-8472, SFF-8079, SFF-8089	SFF-8432 SFF-8083
8 GFC	8.5	FC-PI-4	FC-PI-4				
10GSFP+Cu	10.3125	Chapter 3 Appendix E	Appendix D Appendix E				
IEEE 802.3 Clause 52 (10 Gb/s Ethernet LAN PHY)	10.3125	Chapter 3	Appendix D				
IEEE 802.3 Clause 52 (10 Gb/s Ethernet WAN PHY)	9.95328						
IEEE 802.3 Clause 68 (LRM)	10.3125						
10 GFC	10.51875						
10GBASE-R (IEEE 802.3 Clause 49) Encapsulated in G.709 ODU-2 Frame (FEC)	11.10						

1.3 SFI TYPICAL PCB REACH (INFORMATIVE)

The SFI channel may be implemented with either microstrip or stripline structures. Example host board designs with typical PCB trace reaches are shown in Table 2. Detailed channel properties and recommendations are documented in Appendix A.

Table 2 Host Board Achievable Trace Length

Type	*Material*	*Trace Width (mm)*	*Loss Tangent*	*Copper Thickness (oz) see 1*	*Copper Thickness (μm)*	*Trace Length (mm)*
Microstrip	FR4-6/8	0.3	0.022	1	35	200
	Nelco 4000-13	0.3	0.016	1	35	300
Stripline	FR4-6/8	0.125	0.022	0.5	17.5	150
	Nelco 4000-13	0.125	0.016	0.5	17.5	200

1. Copper (oz) is defined as an ounce of copper over one square foot of laminate.

CHAPTER 2 LOW SPEED ELECTRICAL AND POWER SPECIFICATIONS

2.1 INTRODUCTION

The SFP+ low speed electrical interface has several enhancements over the classic SFP interface (INF-8074i), but the SFP+ host can be designed to also support most legacy SFP modules. SFP+ 2-wire interface electrical and timing specifications are defined in Chapter 4, and the SFP+ 2-wire interface management and register map are defined by SFF-8472.

2.2 GENERAL REQUIREMENTS

The SFP+ modules are hot-pluggable. Hot pluggable refers to plugging in or unplugging a module while the host board is powered.

The module signal ground contacts VeeR and VeeT should be isolated from module case.

All electrical specifications shall be met over the entire specified range of power supplies given in section 2.8.

2.3 SFP+ HOST CONNECTOR DEFINITION

The SFP+ host connector is a 0.8 mm pitch 20 position improved connector specified by SFF-8083, or stacked connector with equivalent electrical performance. Host PCB contact assignment is shown in Figure 1 and contact definitions are given in Table 3. SFP+ module contacts mate with the host in the order of ground, power, followed by signal as illustrated by Figure 2 and the contact sequence order listed in Table 3.

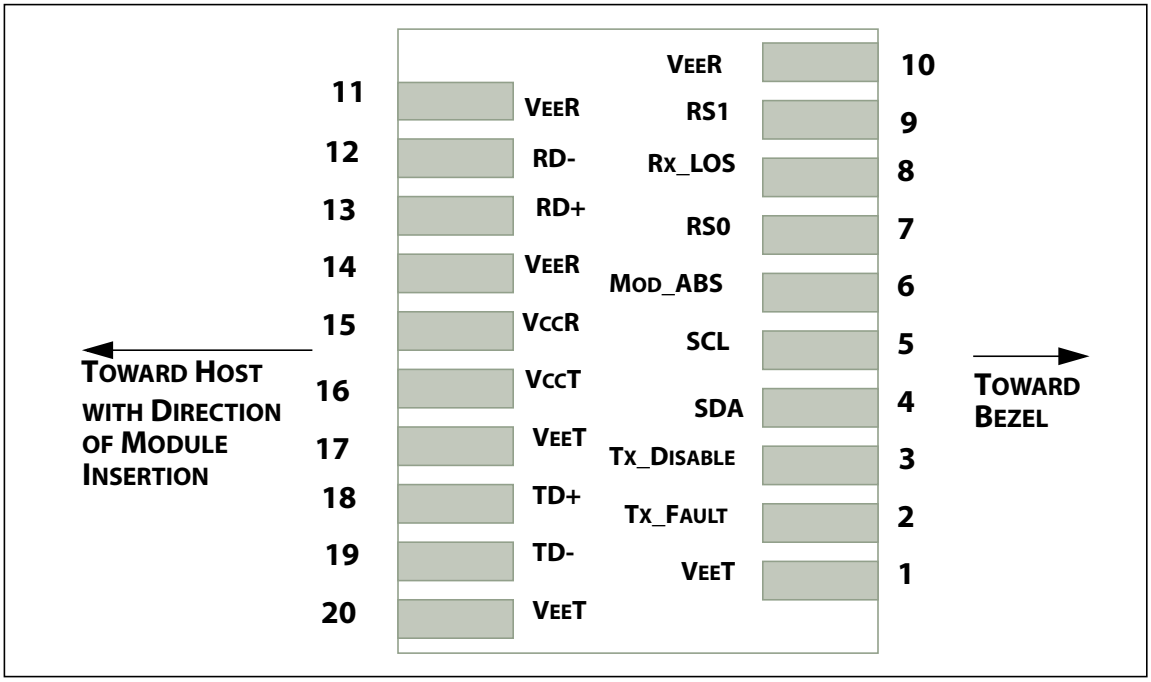


Figure 1 Host PCB SFP+ pad assignment top view

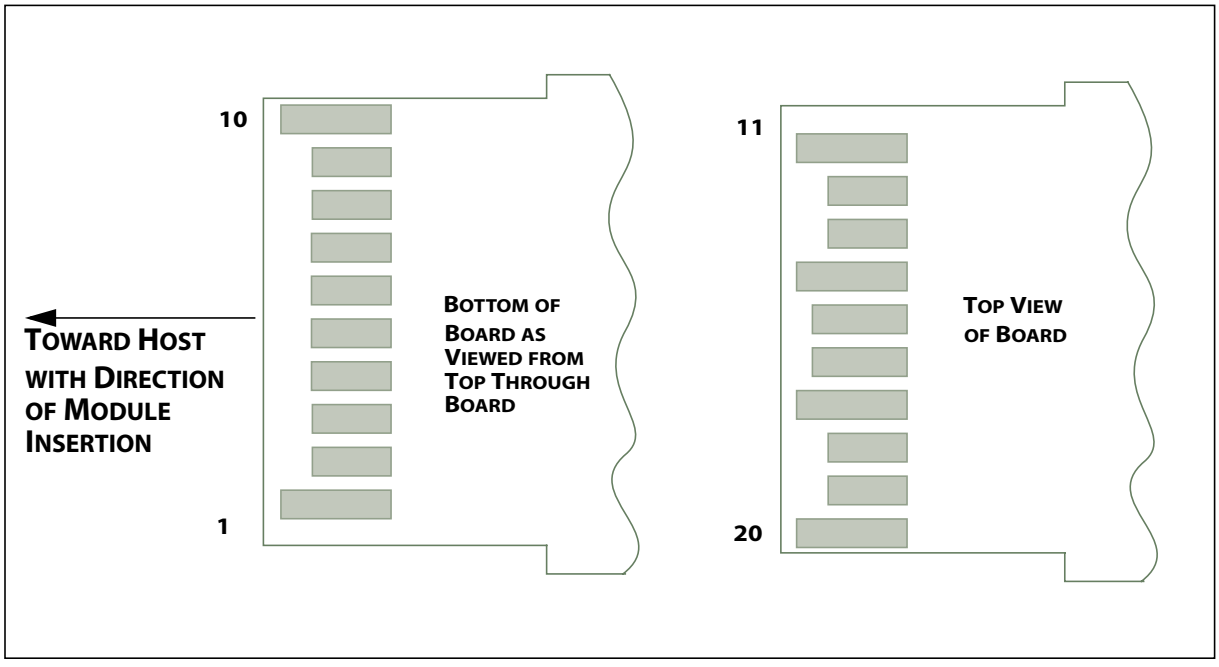


Figure 2 SFP+ module contact assignment

Table 3 SFP+ Module and Host Electrical contact definition

Contacts	*Logic* [1]	*Symbol*	*Power Sequence Order*	*Name/Description*	*Note*
case		case	See 2	Module case	
1		VeeT	1st	Module Transmitter Ground	3
2	LVTTL-O	Tx_Fault	3rd	Module Transmitter Fault	4
3	LVTTL-I	Tx_Disable	3rd	Transmitter Disable; Turns off transmitter laser output	5
4	LVTTL-I/O	SDA	3rd	2-wire Serial Interface Data Line (Same as MOD-DEF2 in INF-8074i)	6
5	LVTTL-I/O	SCL	3rd	2-wire Serial Interface Clock (Same as MOD-DEF1 in INF-8074i)	6
6		Mod_ABS	3rd	Module Absent, connected to VeeT or VeeR in the module	7
7	LVTTL-I	RS0	3rd	Rate Select 0, optionally controls SFP+ module receiver.	8
8	LVTTL-O	Rx_LOS	3rd	Receiver Loss of Signal Indication (In FC designated as Rx_LOS and in Ethernet designated as $\overline{\text{Signal Detect}}$)	4
9	LVTTL-I	RS1	3rd	Rate Select 1, optionally controls SFP+ module transmitter	8
10		VeeR	1st	Module Receiver Ground	3
11		VeeR	1st	Module Receiver Ground	3
12	CML-O	RD-	3rd	Receiver Inverted Data Output	
13	CML-O	RD+	3rd	Receiver Non-Inverted Data Output	
14		VeeR	1st	Module Receiver Ground	3
15		VccR	2nd	Module Receiver 3.3 V Supply	
16		VccT	2nd	Module Transmitter 3.3 V Supply	
17		VeeT	1st	Module Transmitter Ground	3
18	CML-I	TD+	3rd	Transmitter Non-Inverted Data Input	
19	CML-I	TD-	3rd	Transmitter Inverted Data Input	
20		VeeT	1st	Module Transmitter Ground	3

1. Labeling as inputs (I) and outputs (O) are from the perspective of the module

2. The case makes electrical contact to the cage before any of the board edge contacts are made.

3. The module signal ground contacts, VeeR and VeeT, should be isolated from the module case.

4. This contact is an open collector/drain output contact and shall be pulled up on the host see 2.4.1 and 2.4.6. Pull ups can be connected to one of several power supplies, however the host board design shall ensure that no module contact has voltage exceeding module VccT/R + 0.5 V.

5. Tx_Disable is an input contact with a 4.7 kΩ to 10 kΩ pullup to VccT inside the module.

6. See 4.2.

7. See 2.4.4.

8. For SFF-8431 rate select definition see section 2.4.3 and 2.5. (If implementing SFF-8079 contact 7 and 9 in SFF-8431 are used for AS0 and AS1 respectively).

2.4 LOW SPEED ELECTRICAL CONTROL CONTACTS AND 2-WIRE INTERFACE

In addition to the 2-wire serial interface, the SFP+ module has the following low speed contacts for control and status:

- Tx_Fault
- Tx_Disable
- RS0/RS1
- Mod_ABS
- Rx_LOS

2.4.1 TX_FAULT

Tx_Fault is a module output that when high, indicates that the module transmitter has detected a fault condition related to laser operation or safety. If Tx_Fault is not implemented, the Tx_Fault contact signal shall be held low by the module and may be connected to Vee within the module.

The Tx_Fault output is an open drain/collector and shall be pulled up to the Vcc_Host in the host with a resistor in the range 4.7 kΩ to 10 kΩ, or with an active termination according to Table 6.

2.4.2 TX_DISABLE

When Tx_Disable is asserted high or left open, the SFP+ module transmitter output shall be turned off unless the module is a passive cable assembly (see Appendix E) in which case this signal may be ignored. This contact shall be pulled up to VccT with a 4.7 kΩ to 10 kΩ resistor in modules and cable assemblies. Tx_Disable is a module input contact.

When Tx_Disable is asserted low or grounded the module transmitter is operating normally.

2.4.3 RS0/RS1

RS0 and RS1 are module inputs and are pulled low to VeeT with > 30 kΩ resistors in the module. RS0 optionally selects the optical receive signaling rate coverage. RS1 optionally selects the optical transmit signaling rate coverage. For logical definitions of RS0/RS1 see 2.5.

These contacts can also be used for AS0 and AS1 if implementing SFF8079. See SFF8079 for details.

RS1 is commonly connected to VeeT or VeeR in the classic SFP modules. The host needs to ensure that it will not be damaged if this contact is connected to VeeT or VeeR in the module.

2.4.4 MOD_ABS

Mod_ABS is connected to VeeT or VeeR in the SFP+ module. The host may pull this contact up to Vcc_Host with a resistor in the range 4.7 kΩ to10 kΩ. Mod_ABS is asserted "High" when the SFP+ module is physically absent from a host slot. In the SFP MSA (INF-8074i) this contact has the same function but is called MOD_DEF0.

2.4.5 SCL/SDA

SCL is the 2-wire interface clock and SDA is the 2-wire interface data line. SCL and SDA are pulled up to Vcc_Host_2w by resistors in the host. For full specifications see Chapter 4.

2.4.6 RX_LOS

Rx_LOS when high indicates an optical signal level below that specified in the relevant standard. Rx_LOS is an open drain/collector output, but may also be used as an input by supervisory circuitry in the module. For a nominally 3.3 V Vcc_Host using a resistive pull up to Vcc_Host the resistor value shall be in the range 4.7 kΩ to 10 kΩ. For a nominally 2.5 V Vcc_Host using a resistive pull up to Vcc_Host the resistor value shall be in the range 4.7 kΩ to 7.2 kΩ. Alternatively, an active termination according to Table 6 may be used.

The Rx_LOS signal is intended as a preliminary indication to the host in which the module is installed that the received signal strength is below the specified range. Such an indication typically points to non-installed cables, broken cables, or a disabled, failing or a powered off transmitter at the far end of the cable. Additional indications are provided by the host in which the module is installed to verify that the information being transmitted is valid, correctly encoded, and in the correct format. Such additional indications are outside the scope of the module specification.

Rx_LOS may be an optional function depending on the supported standard. If the Rx_LOS function is not implemented, or is reported via the two-wire interface only, the Rx_LOS contact shall be held low by the module and may be connected to Vee within the module.

Rx_LOS assert min and de-assert max are defined in the relevant standard. To avoid spurious transition of Rx_LOS a minimum hysteresis of 0.5 dBo is recommended.

2.5 RATE SELECT HARDWARE CONTROL

The SFP+ module provides two inputs RS0 and RS1 that can optionally be used for rate selection. RS0 controls the receive path signalling rate capability, and RS1 controls the transmit path signalling rate capability, as defined in Table 4. The host and module may choose to use either, both, or none of these functions. Because contact 9 in the classic SFP INF-8074i is connected to VeeR, an SFP+ host utilizing RS1 must provide short circuit protection.

This rate select functionality can also be controlled by software as defined by SFF-8472.

Optionally the rate select methods of Part 2 of SFF-8079 may be used instead of the method described here by setting the management declaration bit (A0h byte 93 bit 2) to 1, see SFF-8472.

Table 4 Rate Select Hardware Control Contacts

Parameter	*State*	*Conditions*
RS0	Low	RX signalling rate less than or equal to 4.25 GBd
	High	RX signalling rate greater than 4.25 GBd
RS1	Low	TX signalling rate less than or equal to 4.25 GBd
	High	TX signalling rate greater than 4.25 GBd

2.6 LOW SPEED ELECTRICAL SPECIFICATIONS

SFP+ low speed signalling is based on Low Voltage TTL (LVTTL) operating with a module supply of 3.3 V +/-5% and with a host supply range of 2.38 to 3.46 V.

The 2-wire interface protocol and electrical specifications are defined in Chapter 4.

2.6.1 MODULE LOW SPEED ELECTRICAL SPECIFICATIONS

The SFP+ module low speed electrical specifications are given in Table 5. All I/O powered by VccT is referenced to VeeT and similarly VccR is referenced to VeeR.

Table 5 Low Speed Module Electrical Specifications

Parameter	*Symbol*	*Min.*	*Max.*	*Unit*	*Conditions*
Module Vcc	VccT, VccR	3.14	3.46	V	
Tx_Fault, Rx_LOS	V_{OL}	-0.3	0.40	V	At 0.7 mA [1]
	I_{OH} [1]	-50	37.5	µA	Measured with a 4.7 kΩ load pulled up to Vcc_Host where Vcc_Host_min<Vcc_host<Vcc_Host_max
Tx_Disable, RS0, RS1	V_{IL}	-0.3	0.8	V	
	V_{IH}	2.0	VccT + 0.3	V	
1. Positive values indicate current flowing into the module.					

2.6.2 HOST LOW SPEED ELECTRICAL SPECIFICATIONS

The SFP+ Host low speed electrical specifications are given in Table 6. All I/O powered by VccT is referenced to VeeT and similarly VccR is referenced to VeeR.

Table 6 Low Speed Host Electrical Specifications

Parameter	*Symbol*	*Min.*	*Max.*	*Unit*	*Conditions*
Host Vcc Pullup	Vcc_Host	2.38	3.46	V	
Tx_Fault, Rx_LOS	V_{IL}	see 1	see 1	V	
	V_{IH}	see 1	see 1	V	
Tx_Disable	V_{OL}	-0.3	0.4	V	V_{OL} measured with 4.7 kΩ to 10 kΩ pull up to VccT max
	V_{OH}	VccT-0.5	VccT + 0.3	V	V_{OH} measured with 10 kΩ pull up to VccT min
RS0, RS1	V_{OL}	-0.3	0.4	V	V_{OL} measured with no load
	V_{OH}	2.2	VccT + 0.3	V	V_{OH} measured with 30 kΩ to VeeR.
1. Determined by host design, such that VIH > 2.1 V for the range of IOH in Table 5. One option is using standard LVTTL input with a pull-up to Vcc_Host in the range 4.7 kΩ to 10 kΩ					

2.7 TIMING REQUIREMENT OF CONTROL AND STATUS I/O

The timing requirements of control and status I/O are defined in Table 7.

Table 7 Timing Parameters for SFP+ Management

Parameter	*Symbol*	*Min.*	*Max.*	*Unit*	*Conditions*
Tx_Disable assert time	t_off		100	µs	Rising edge of Tx_Disable to fall of output signal below 10% of nominal
Tx_Disable negate time	t_on		2	ms	Falling edge of Tx_Disable to rise of output signal above 90% of nominal. This only applies in normal operation, not during start up or fault recovery.
Time to initialize 2-wire interface	t_2w_start_up		300	ms	From power on or hot plug after the supply meeting Table 8.
Time to initialize	t_start_up		300	ms	From power supplies meeting Table 8 or hot plug or Tx disable negated during power up, or Tx_Fault recovery, until non-cooled power level I part (or non-cooled power level II part already enabled at power level II for Tx_Fault recovery) is fully operational.
Time to initialize cooled module and time to power up a cooled module to Power Level II	t_start_up_cooled		90	s	From power supplies meeting Table 8 or hot plug, or Tx disable negated during power up or Tx_Fault recovery, until cooled power level I part (or cooled power level II part during fault recovery) is fully operational. Also, from stop bit low-to-high SDA transition enabling Power Level II until cooled module is fully operational
Time to Power Up to Level II	t_power_level2		300	ms	From stop bit low-to-high SDA transition enabling power level II until non-cooled module is fully operational
Time to Power Down from Level II	t_power_down		300	ms	From stop bit low-to-high SDA transition disabling power level II until module is within power level I requirements
Tx_Fault assert	Tx_Fault_on		1	ms	From occurrence of fault to assertion of Tx_Fault
Tx_Fault assert for cooled module	Tx_Fault_on_cooled		50	ms	From occurrence of fault to assertion of Tx_Fault
Tx_Fault Reset	t_reset	10		µs	Time Tx_Disable must be held high to reset Tx_Fault
RS0, RS1 rate select timing for FC	t_RS0_FC, t_RS1_FC		500	µs	From assertion till stable output
RS0, RS1 rate select timing non FC	t_RS0, t_RS1		24	ms	From assertion till stable output
Rx_LOS assert delay	t_los_on		100	µs	From occurrence of loss of signal to assertion of Rx_LOS
Rx_LOS negate delay	t_los_off		100	µs	From occurrence of presence of signal to negation of Rx_LOS

2.7.1 MODULE POWER ON INITIALIZATION PROCEDURE, TX_DISABLE NEGATED

During power on of the module, Tx_Fault, if implemented, may be asserted (high) as soon as power supply voltages are within specification. For module initialization with Tx_Disable negated, Tx_Fault shall be negated when the transmitter safety circuitry, if implemented, has detected that the transmitter is operating in its normal state. If a transmitter fault has not occurred, Tx_Fault shall be negated within a period t_start_up from the time that VccT exceeds the specified minimum operating voltage (see Table 8). If the Tx_Fault remains asserted after t_start_up, the host shall determine whether the module is cooled by reading the status bit over 2-wire interface. If the module is not cooled, the host may assume that a transmission fault has occurred. If the module is cooled, the host may assume that a transmission fault has occurred if Tx_Fault remains asserted beyond t_start_up_cooled.

The power on initialization timing for a module with Tx_Disable negated is shown in Figure 3.

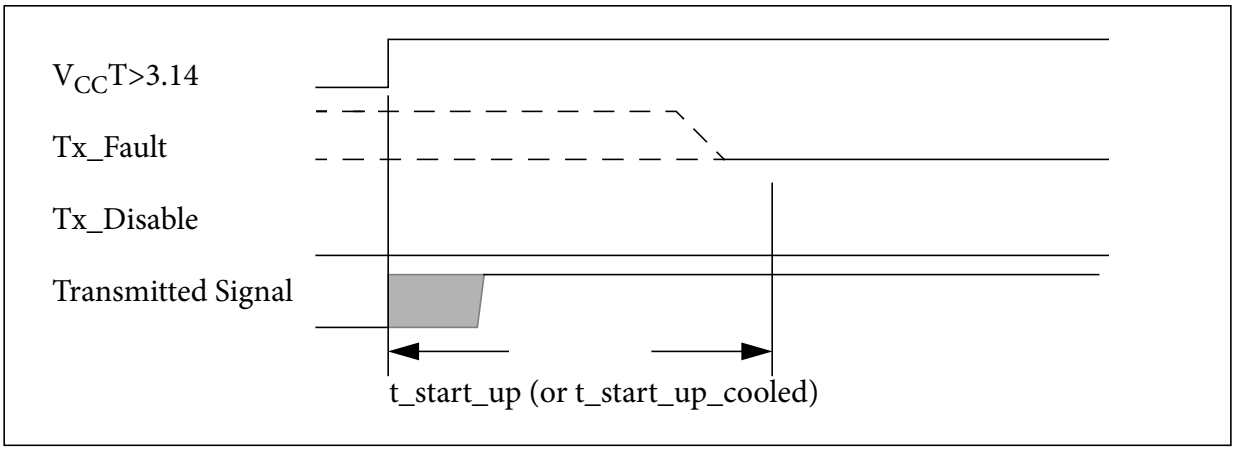


Figure 3 Power on initialization of module, Tx_Disable negated

2.7.2 MODULE POWER ON INITIALIZATION PROCEDURE, TX_DISABLE ASSERTED.

For module power on initialization with Tx_Disable asserted, the state of Tx_Fault is not defined while Tx_Disable is asserted. After Tx_Disable is negated, Tx_Fault may be asserted while safety circuit initialization is performed. Tx_Fault shall be negated when the transmitter safety circuitry, if implemented, has detected that the transmitter is operating in its normal state. If a transmitter fault has not occurred, Tx_Fault shall be negated within a period t_start_up from the time that Tx_Disable is negated. If Tx_Fault remains asserted beyond the period t_start_up, the host may assume that a transmission fault has been detected by the module.

If no transmitter safety circuitry is implemented, the Tx_Fault signal may be tied to its negated state.

The power on initialization timing for a module with Tx_Disable asserted is shown in Figure 4.

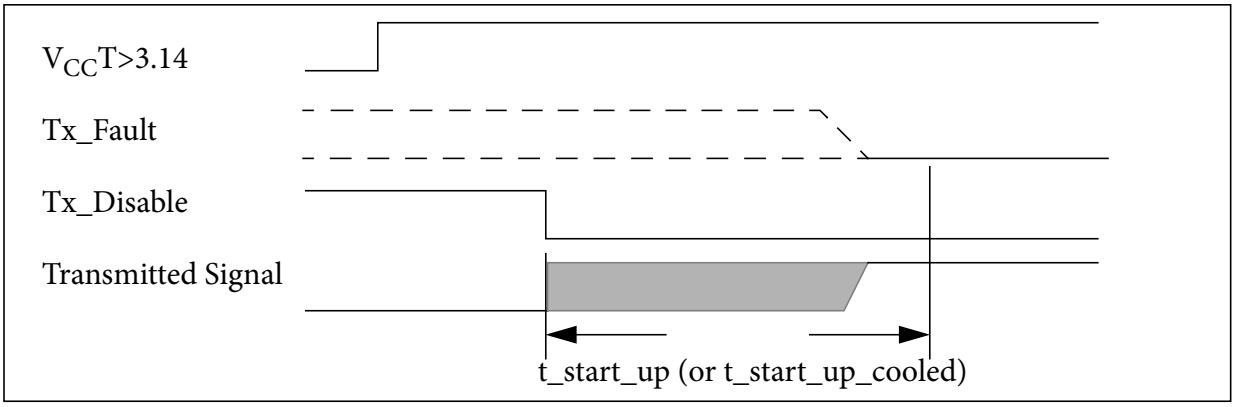


Figure 4 Power on initialization of module, Tx_Disable asserted

2.7.3 INITIALIZATION DURING HOT PLUGGING

When a module is not installed, Tx_Fault is held to the asserted state by the pull up circuits on the host. As the module is installed, contact is made with the ground, voltage, and signal contacts in the specified order. After the module has determined that $V_{CC}T$ has reached the specified value, the power on initialization takes place as described in the previous section. An example of initialization during hot plugging is provided in Figure 5.

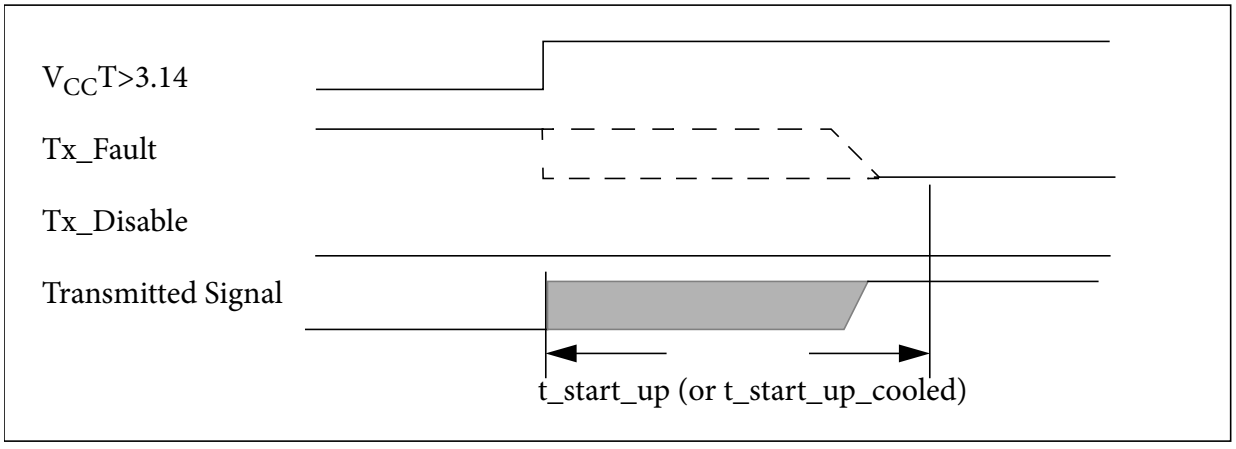


Figure 5 Example of initialization during hot plugging, Tx_Disable negated.

2.7.4 TRANSMITTER MANAGEMENT

The timing requirements for the management of optical outputs from the module using the Tx_Disable signal are shown in Figure 6. Note that t_on time refers to the maximum delay until the modulated optical signal reaches 90% of the final value, not just the average optical power.

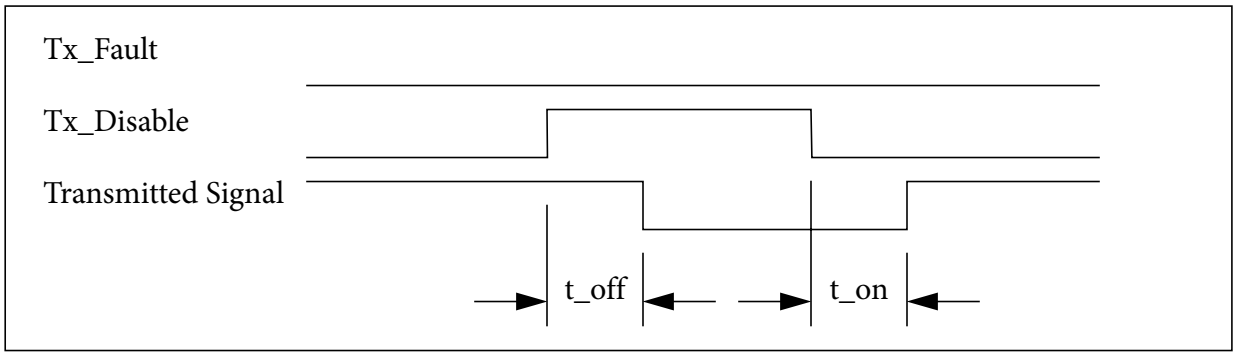


Figure 6 Management of module during normal operation, Tx_Disable implemented

2.7.5 TRANSMITTER SAFETY DETECTION AND PRESENTATION

If Tx_Fault is implemented it shall meet the timing requirements of Figure 7.

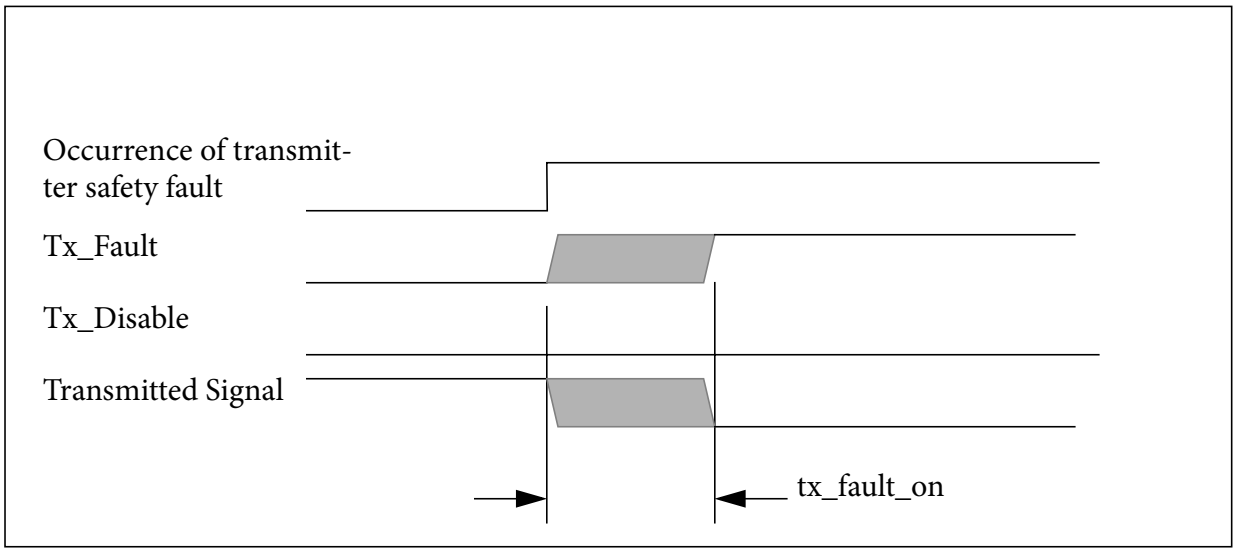


Figure 7 Occurrence of condition generating Tx_Fault

2.7.6 MODULE FAULT RECOVERY

The detection of a safety-related transmitter fault condition presented by Tx_Fault shall be latched. The following protocol may be used to reset the latch in case the transmitter fault condition is transient.

To reset the fault condition and associated detection circuitry, Tx_Disable shall be asserted for a minimum of t_reset. Tx_Disable shall then be negated. Alternatively the Software Tx disable is asserted and negated. In less than the maximum value of t_start_up the optical transmitter will correctly reinitialize the laser circuits, negate Tx_Fault, and begin normal operation if the fault condition is no longer present. If a fault condition is detected during the reinitialization, Tx_Fault shall again be asserted, the fault condition again latched, and the optical transmitter circuitry will again be disabled until the next time a reset protocol is attempted. The manufacturer of the module shall ensure that the optical power emitted from an open connector or fiber is compliant with applicable eye safety requirements during all reset attempts, during normal operation or upon the occurrence of reasonable single fault conditions. The module may require internal protective circuitry to prevent the frequent assertion of the Tx_Disable signal from generating frequent pulses of energy that violate the safety requirements. The timing for successful recovery from a transient safety fault condition is shown in Figure 8.

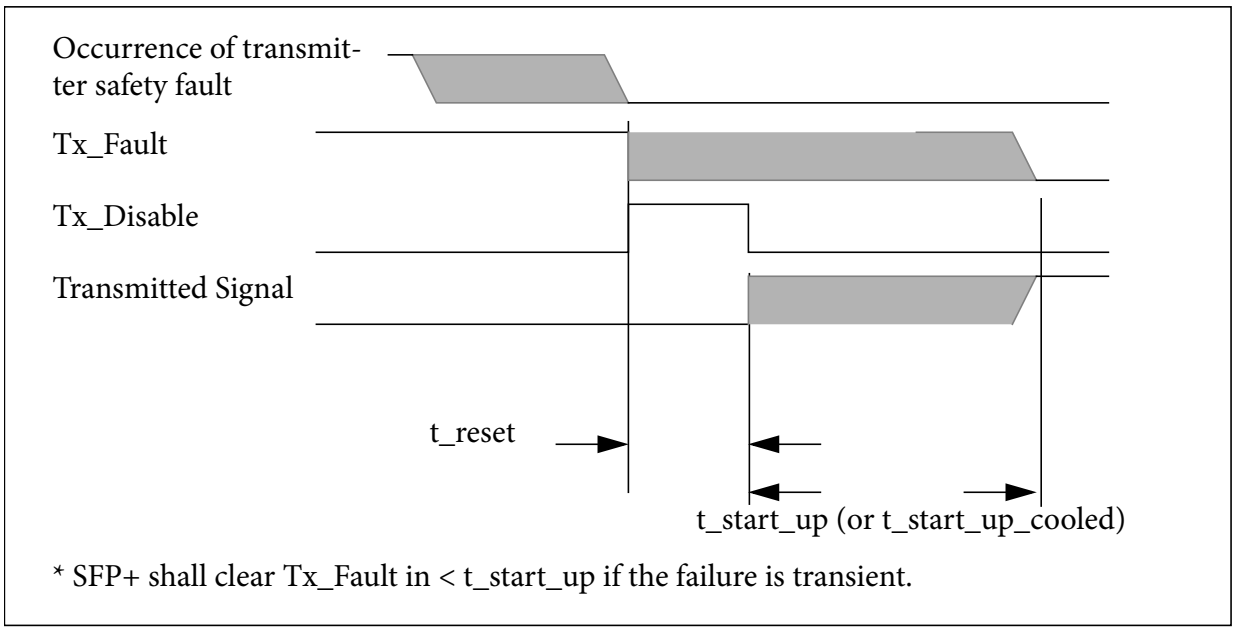


Figure 8 Successful recovery from transient safety fault condition

An example of an unsuccessful recovery, where the fault condition was not transient, is shown in Figure 9.

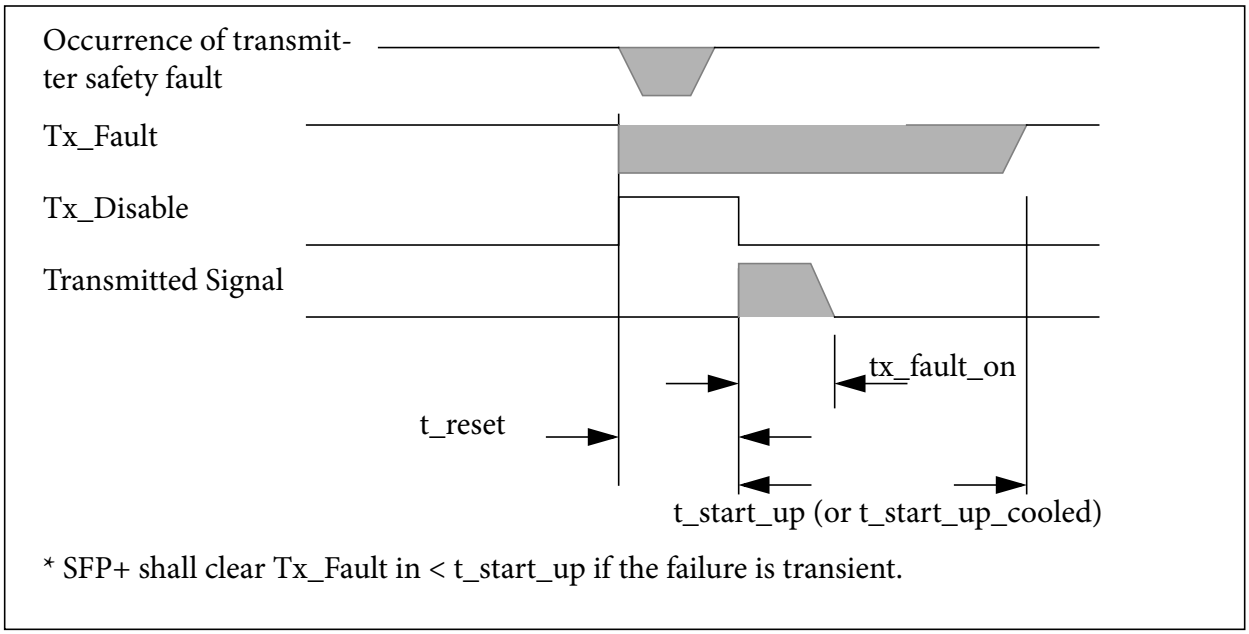


Figure 9 Unsuccessful recovery from safety fault condition

2.7.7 MODULE LOSS OF SIGNAL INDICATION

If the module is specified as implementing Rx_LOS, the timing is specified in Figure 10.

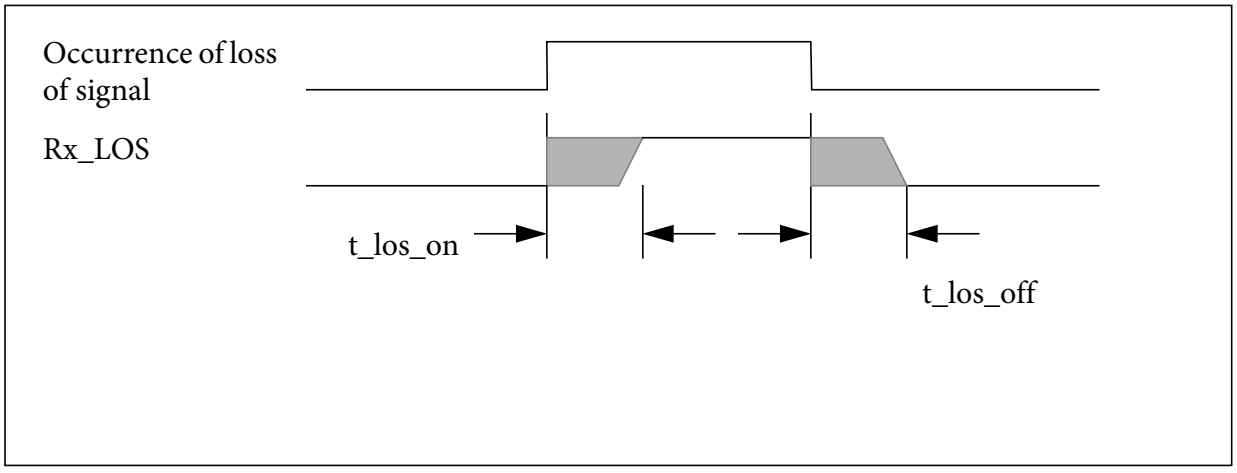


Figure 10 Timing of Rx_LOS detection

2.8 SFP+ POWER REQUIREMENT

The module host has two 3.3 V power contacts, one supplying the module transmitter voltage (VccT) and the other supplying the module receiver voltage (VccR). The maximum current capacity, both continuous and peak, for each connector contact is 500 mA.

SFP+ module maximum power consumption shall meet one of the following power classes:

- Power Level I modules – Up to 1.0 W
- Power Level II modules – Up to 1.5 W

To avoid exceeding system power supply limits and cooling capacity, all modules at power up by default shall operate with ≤1.0 W. Hosts supporting Power Level II operation may enable a Power Level II module through the 2-wire interface. Power Level II modules shall assert the power level declaration bit of SFF-8472.

The maximum power level is allowed to exceed the classified power level for 500 ms following hot insertion or power up, or Power Level II authorization, however the current is limited to values given by Table 8 and illustrated in Figure 11.

At host power up the host shall supply VccT and VccR to the module within 100 ms of each other.

2.8.1 MODULE POWER SUPPLY REQUIREMENTS

SFP+ module operates from the host supplied VccT and VccR. To protect the host and system operation, each SFP+ module during hot plug and normal operation shall follow the requirements listed in Table 8 and illustrated by Figure 11. The requirements for current apply to the current through each inductor of Figure 56 while the power supply voltages are defined at the SFP+ connector.

2.8.2 HOST POWER SUPPLY NOISE OUTPUT

The host shall generate an effective weighted integrated spectrum RMS noise less than 25 mV in the frequency range 10 Hz to 10 MHz, according to the methods of D.17.1

2.8.3 MODULE POWER SUPPLY NOISE OUTPUT

The module shall generate less than 15 mV RMS noise at point X of Figure 56 in the frequency range 10Hz to 10MHz, according to the methods of D.17.2.

2.8.4 POWER SUPPLY NOISE TOLERANCE

SFP+ modules shall meet all electrical requirements and remain fully operational in the presence of a sinusoidal tolerance signal of amplitude given by

Table 8 swept from 10 Hz to 10 MHz according to the methods of D.17.3. This emulates the worst case noise of the host.

It is also desirable for a module and host to each tolerate a degree of random or semi-random noise on both VccT and VccR simultaneously, but the characteristics of this noise are beyond the scope of this document.

Table 8 SFP+ Module Power Supply Requirements

Parameters	*Symbol*	*Conditions*	*Min*	*Max*	*Units*
		Power Level I Module			
Power supply noise tolerance including ripple [peak-to-peak]		see D.17.3		66	mV
Power supply voltages including ripple, droop and noise below 100 kHz	VccT, VccR	Note 1	3.14	3.46	V
Instantaneous peak current at hot plug		Note 2, 3		400	mA
Sustained peak current at hot plug		Note 2, 3, 5		330	mA
Module maximum power consumption				1.0	W
		Power Level II Module			
Power supply noise tolerance including ripple [peak-to-peak]		see D.17.3		66	mV
Power supply voltages including ripple, droop and noise below 100 kHz	VccT, VccR	Note 1	3.14	3.46	V
Instantaneous peak current at hot plug		Note 2, 3		400	mA
Sustained peak current at hot plug		Note 2, 3, 5		330	mA
Instantaneous peak current on enabling Power Level II		Note 2, 3, 5		600	mA
Module sustained peak current on enabling Power Level II		Note 2, 3, 5		500	mA
Module maximum power consumption				1.5	W
Maximum power consumption at power up		Note 4		1	W

1. Set point is measured at the input to the connector on the host board reference to Vee. Droop is any temporary drop in voltage of the power supply such as that caused by plugging in another module or when enabling another module to Power Level II.
2. The requirements for current apply to the current through each inductor of Figure 56.
3. The maximum currents are the allowed currents for each power supply VccT or VccR, therefore the total module peak currents can be twice this value. The instantaneous peak current is allowed to exceed the specified maximum current capacity of the connector contact for a short period, see Figure 11.
4. Maximum module power consumption shall not exceed 1.0 W from 500ms after power up until level II operation is enabled.
5. Not to exceed the sustained peak limit for more than 50 µs; may exceed this limit for shorter durations.

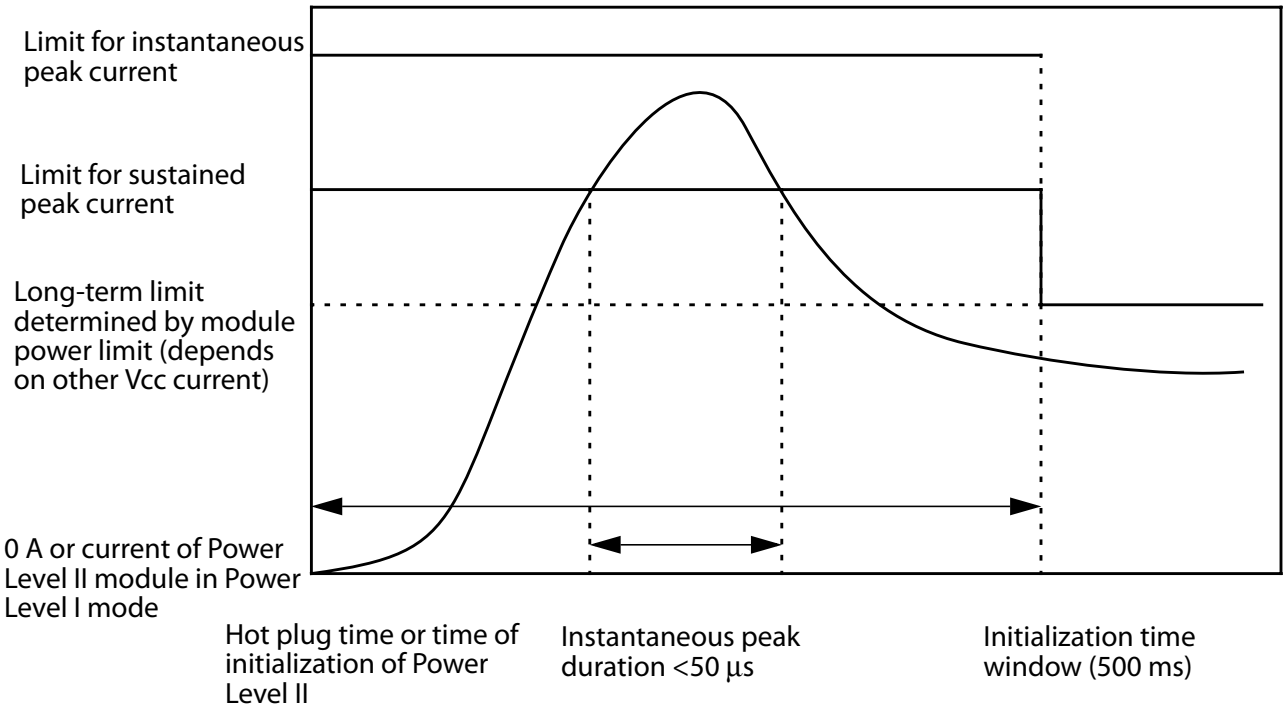


Figure 11 Instantaneous and sustained peak current for VccT or VccR

2.9 ESD

The SFP+ module and host SFI contacts (High Speed Contacts) shall withstand 1000 V electrostatic discharge based on Human Body Model per JEDEC JESD22-A114-B.

The SFP+ module and all host contacts with exception of the SFI contacts (High Speed Contacts) shall withstand 2 kV electrostatic discharge based on Human Body Model per JEDEC JESD22-A114-B.

The SFP+ module shall meet ESD requirements given in EN61000-4-2, criterion B test specification such that units are subjected to 15 kV air discharges during operation and 8 kV direct contact discharges to the case.

CHAPTER 3 HIGH SPEED ELECTRICAL SPECIFICATION SFI

3.1 INTRODUCTION

SFI signalling is based on differential high speed low voltage logic with AC-coupling in the module. SFI was developed with the primary goal of low power and low electromagnetic interference (EMI). To satisfy this requirement the nominal differential signal levels are ~500 mV p-p with edge speed control to reduce EMI. SFP+ compliant hosts are allowed to support just linear modules, just limiting modules, or both linear and limiting modules.

3.2 SFI APPLICATIONS DEFINITION

The application reference model for SFI connects a high speed ASIC/SERDES to the SFP+ module as shown in Figure 12. The SFI interface is designed to support IEEE 802.3 10Gig standards Clauses 49, 50, and 51, and 10GFC. For all other FC signalling rates see FC-PI-4. SFI supported signalling rates are listed in Table 9. SFP+ compliant modules and hosts may support one or more of the signalling rates listed in Table 9. For 10GSFP+Cu (direct attach copper) specifications and applications reference model, see Appendix E.

Table 9 SFI Supported signalling Rates

Standard	*Description*	*signalling Rate*	*Units*
IEEE std-802.3 Clause 50	10GBASE-W WAN PHY	9.95328	GBd
IEEE std-802.3 Clause 49	10GBASE-R LAN PHY	10.3125	GBd
Fibre Channel - 10 Gigabit (10GFC)	10GFC	10.51875	GBd
10Gig Ethernet with FEC	10GBASE-R over G.709	11.10	GBd

The SFI interface operates from 9.95 to 11.1 GBd.

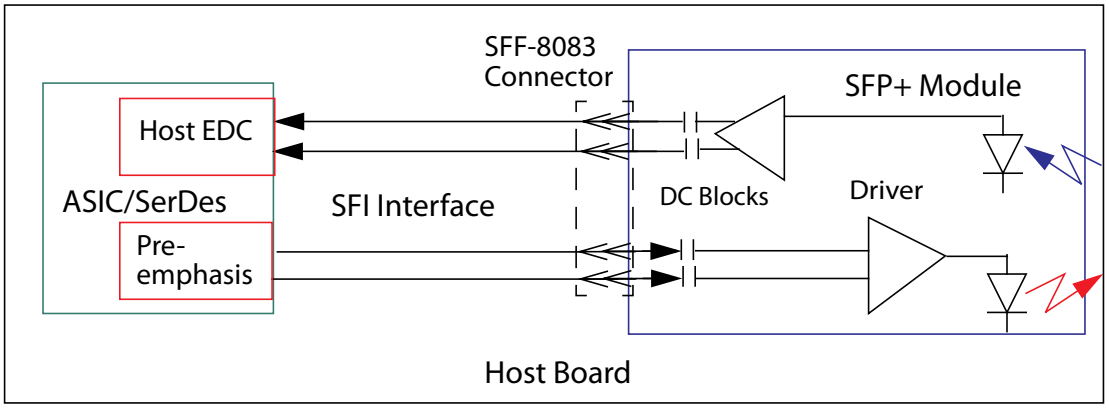


Figure 12 SFI Application Reference Model

3.3 SFI TEST POINTS DEFINITION AND MEASUREMENTS

SFI reference compliance test points are defined with the Host Compliance Board and the Module Compliance Board for measurement consistency, see Appendix C. The reference test boards provide a set of overlapping measurements for ASIC/SerDes, module, and host validation to ensure interoperability. For improved measurement accuracy the actual reference test card responses may be calibrated out of the measurements and replaced with functions that represent the ideal responses defined in Appendix C for the reference test cards.

Points A, B, C, and D require AC coupled test equipment. All SFI test equipment must have 50 Ω single ended impedance on all test ports.

The reference impedance for differential measurements and S-parameters is 100 Ω, and the reference impedance for common mode measurements and S-parameters is 25 Ω.

The bandwidth of measurement instrument shall be 12 GHz unless specified otherwise.

SFI reference points are listed in Table 10.

Table 10 SFI Reference Points

Compliance point	*Designation*
ASIC/SerDes output	A
Host output	B
Host input	C
ASIC/SerDes input	D
Module input	B’
Module output	C’
Module input calibration	B” (double quotation)
Host input calibration	C” (double quotation)

3.3.1 HOST COMPLIANCE POINTS

Host system transmitter and receiver compliance are defined by tests in which a Host Compliance Board is inserted as shown in Figure 13 in place of the SFP+ module. The Host Compliance Board meets the specifications of Appendix C. The compliance points are B and C.

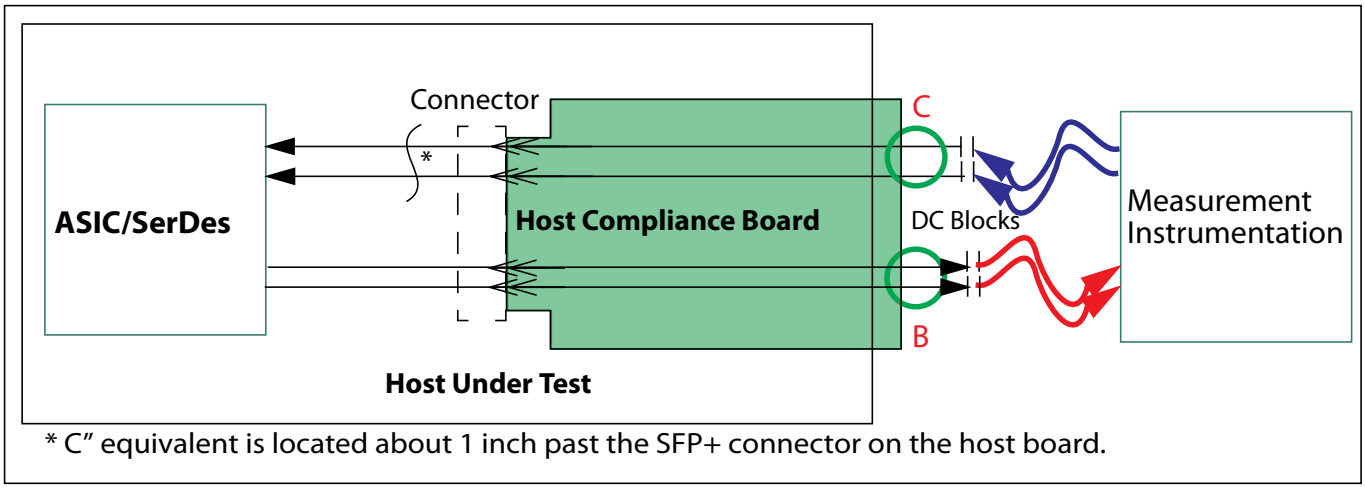


Figure 13 Host Compliance Board

SFP+ host compliance points are defined as the following:

- B: Host transmitter output at the output of the Host Compliance Board. Specifications for B are given in 3.5.1.
- C: Host receiver input at the input of the Host Compliance Board. Specifications for C are given in 3.5.2.

3.3.2 MODULE COMPLIANCE POINTS

Module transmitter and receiver compliance are defined by tests in which the module is inserted into the Module Compliance Board as shown in Figure 14. The Module Compliance Board meets the specifications of Appendix C. The compliance points for the module are B′ and C′.

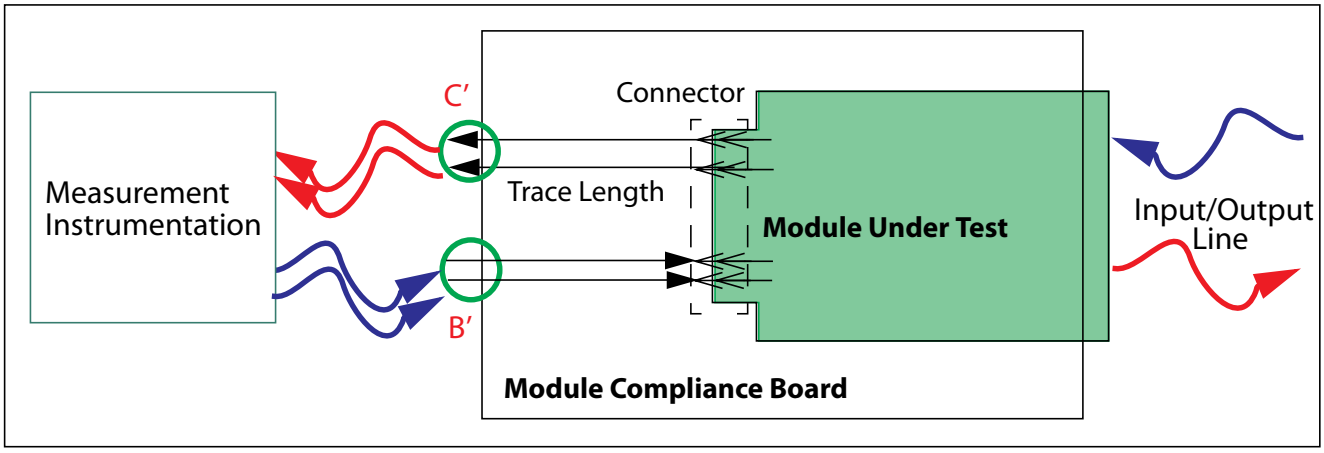


Figure 14 Module Compliance Board

SFP+ module compliance points are defined as the following:

- B′: SFP+ module transmitter input at the input of the Module Compliance Board. Specifications for B′ are given in 3.6.1.
- C′: SFP+ module receiver output at the output of the Module Compliance Board. Specifications for C′ are given in 3.6.2.

3.3.3 ASIC/SERDES TEST POINTS (INFORMATIVE)

ASIC/SerDes transmitter and receiver may be tested on a test board as shown in Figure 15 with nominal trace response as specified by C.1.3 to avoid degradation due to excessive trace loss and to ensure consistent measurements.

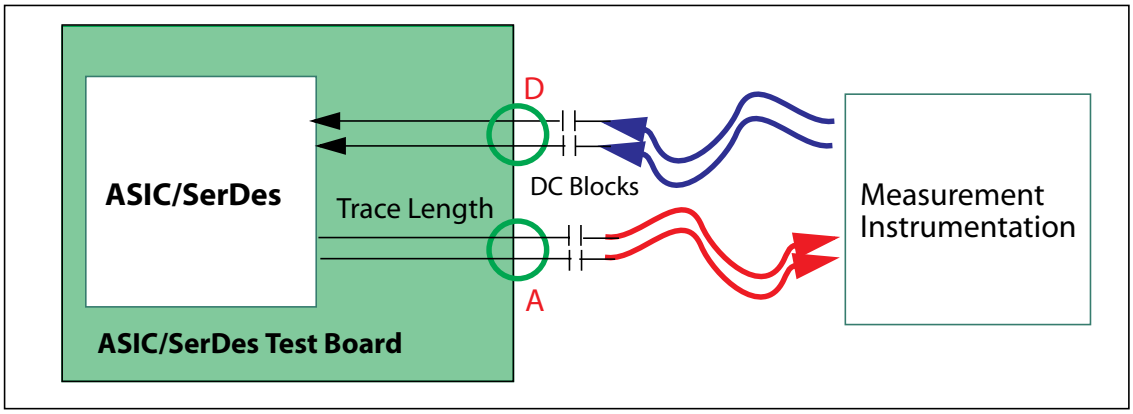


Figure 15 ASIC/SerDes Test Board

SFI ASIC/SerDes test points are defined as the following:

- A: SerDes transmitter output at the output of the ASIC/SerDes Test Board. Recommendations for A are given in B.2.
- D: ASIC/SerDes receiver input at the input of the ASIC/SerDes Test Board. Recommendations for D are given in B.3.

3.3.4 HOST INPUT CALIBRATION POINT

Host receiver input tolerance signals are calibrated through the Host Compliance Board at the output of the Module Compliance Board as shown in Figure 16. The host input calibration point is at C″ with specifications for C″ given in 3.5.2. The response between the connector and C″ is specified by C.1.2.

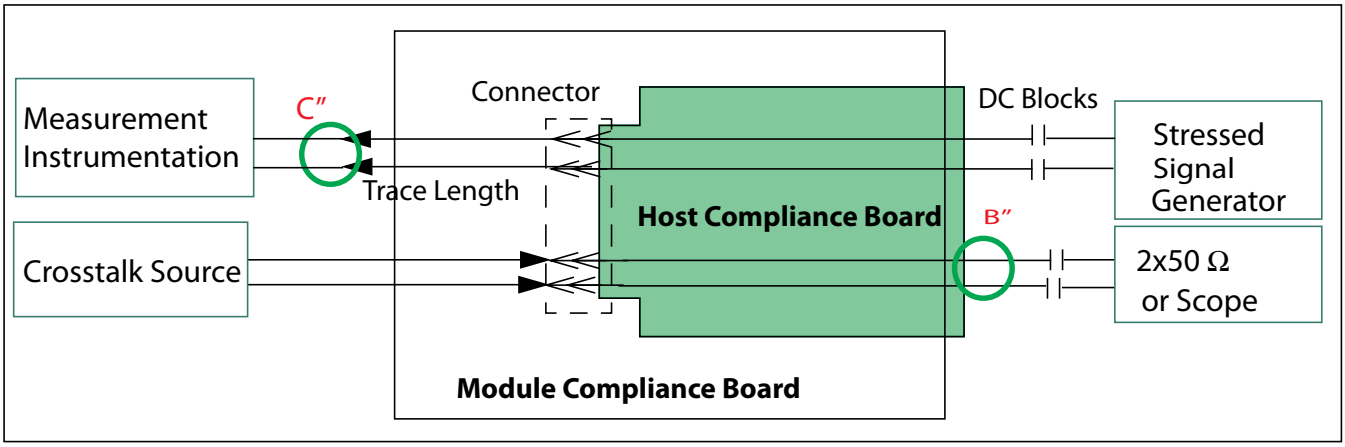


Figure 16 Host input calibration point C″ and crosstalk source calibration point B″

3.3.5 MODULE INPUT CALIBRATION POINT

Module transmitter input tolerance signals are calibrated through the Module Compliance Board at the output of the Host Compliance Board as shown in Figure 17. The module input calibration point is at B″ with specifications for B″ given in 3.6.1. The response between the connector and B″ is specified by C.1.1.

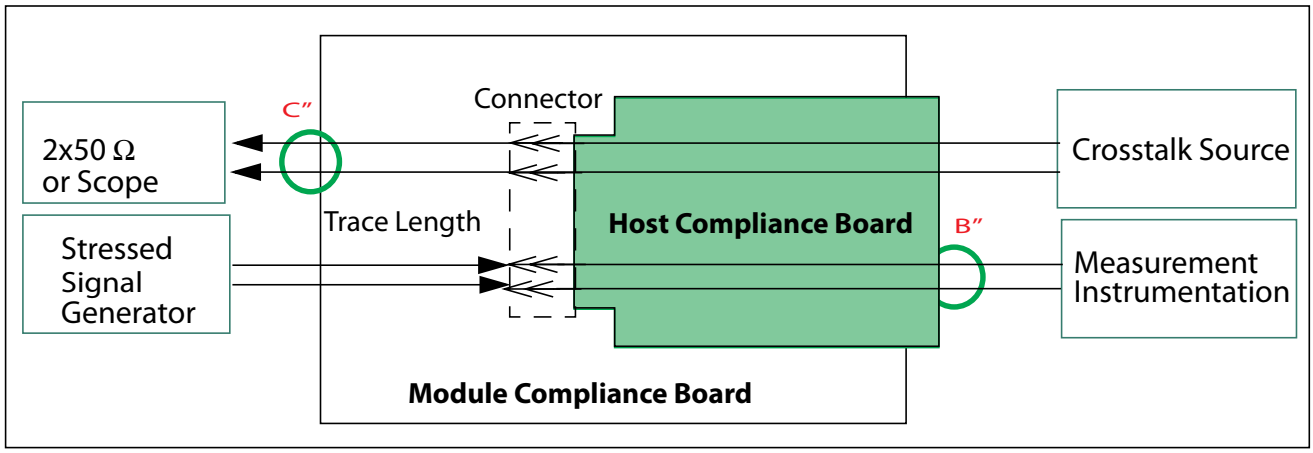


Figure 17 Module input calibration point B″ and crosstalk source calibration point C″

3.4 SFI TERMINATION AND DC BLOCKING

The SFI link uses nominal 100 Ω differential source and load terminations on both the host board and the module. The SFI transmitter provides both differential and common mode termination. The SFI transmitter and receiver termination specifications for each of the compliance points are given by:

- Host – 3.5 SFP+ Host System Specifications
- Module – 3.6 SFP+ Module Specifications .

Host SerDes termination recommendations are given by:

- ASIC/SerDes – Appendix B

SFP+ modules shall incorporate blocking capacitors or equivalent on all SFI inputs and outputs as shown in Figure 18. The SFI transmitter is represented by terminations Z_p and Z_n which form a 100 Ω differential source. Each termination has a nominal value of 50 Ω, and therefore the common mode impedance is 25 Ω. The SFI receiver is represented with termination Z_{diff} with nominal 100 Ω value. This representation is not intended to preclude the use of other implementations which may provide common mode termination, however the SFI specification does not require any common mode termination at the receiver. If common mode terminations are provided, it may reduce common mode voltage and EMI.

It is recommended that both the module and the host use transmission lines targeted to have 100 Ω differential impedance with about 7% coupling. SFP+ percent differential coupling is defined by the following equation:

$$Coupling = \frac{Zcm \times 4 - Zdiff}{Zcm \times 4 + Zdiff} \times 100$$

Where Zcm is the common mode impedance and Zdiff is the differential impedance.

Differential traces with nominal 7% coupling offer a good compromise between reasonable common mode match and practical transmission line geometries. These are the targets for the module and host Compliance Boards described in Appendix C.

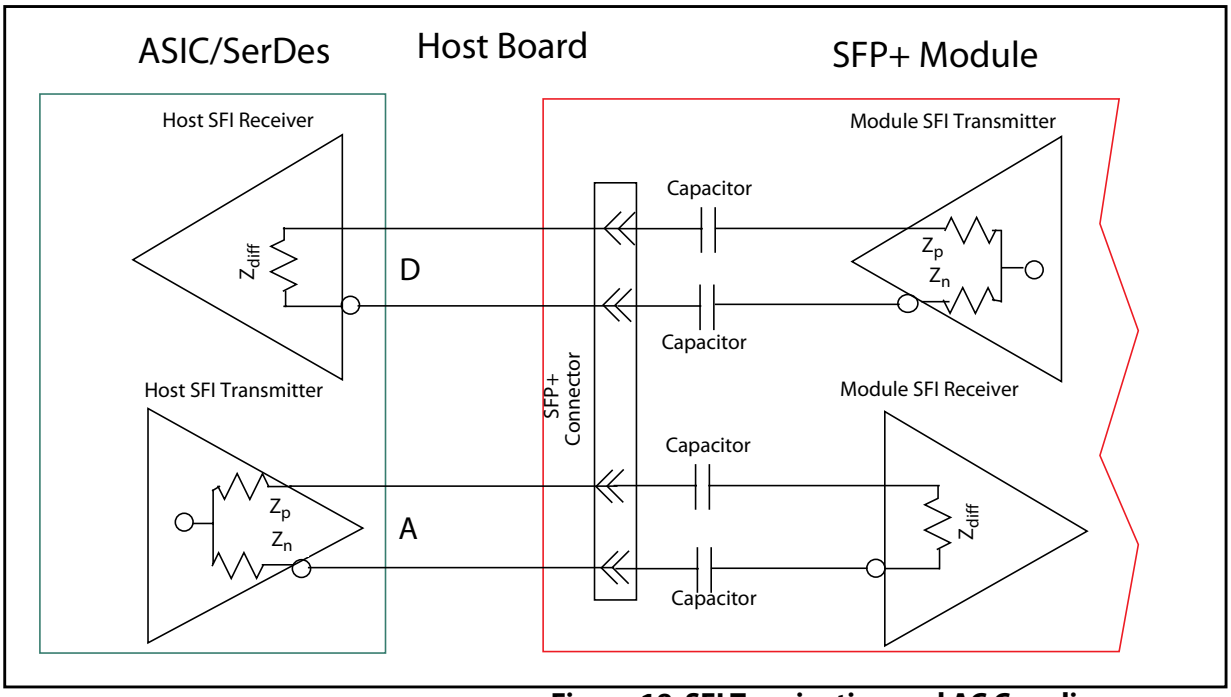


Figure 18 SFI Termination and AC Coupling

3.5 SFP+ HOST SYSTEM SPECIFICATIONS

SFP+ host system transmitter specifications at compliance point B are given in 3.5.1. SFP+ Host system receiver specifications at compliance point C are given in 3.5.2.

All specifications are to be met at the host compliance test points defined in 3.3.1.

The solder pads for the high speed traces in the SFF-8431 Module Compliance Board are 1.1x0.4 mm to improve high frequency performance instead of 2.0x0.5 mm as defined in the SFF-8083 for improved manufacturability. Trade-off between host performance and manufacturability are left to the host designer. For detailed geometry of the Module Compliance Board, see the Gerber files in C.3.4.

Warning: The host expects DC blocking in the module, and for improved performance the Host Compliance Board is not required to incorporate DC blocks. DC blocking within the test equipment or between the host and the equipment is necessary for all host SFI signals.

3.5.1 HOST TRANSMITTER OUTPUT SPECIFICATIONS AT B

SFP+ host transmitter electrical specifications defined at compliance point B are given in Table 11 and Table 12. These specifications are defined at the output of the Host Compliance Board specified in C.2. Host transmitters must provide adequate low frequency signal response for the applications supported.

Table 11 Host Transmitter Output Electrical Specifications at B

Parameter - B	*Symbol*	*Conditions*	*Min*	*Max*	*Units*
Termination Mismatch at 1 MHz	ΔZ_M	See D.16, Figure 55		5	%
Single Ended Output Voltage Range			-0.3	4.0	V
Output AC Common Mode Voltage		See D.15		15	mV (RMS)
Differential Output S-parameter	SDD22	0.01 to 2 GHz		-12	dB
		2 to 11.1 GHz		see 1	dB
Common Mode Output S-parameter	SCC22	0.01 to 2.5 GHz		see 2	dB
		2.5 to 11.1 GHz		-3	dB
1. Reflection coefficient given by equation SDD22(dB) < -6.68 + 12.1 × log10(f/5.5), with f in GHz. 2. Reflection coefficient given by equation SCC22(dB) < -7 + 1.6 × f, with f in GHz.					

The specification of common mode output return loss reduces EMI and noise by absorbing common mode reflections and noise.

The SFI jitter specifications at reference point B are listed in Table 12 and the compliance mask is shown in Figure 19. As baseline wander can create low probability eye closure which is not detected by the $5x10^{-5}$ mask hit ratio, baseline wander must be controlled so as not to significantly degrade the signal at B.

Table 12 Host Transmitter Output Jitter and Eye Mask Specifications at B

Parameters- B	*Symbol*	*Conditions*	*Min*	*Target Value*	*Max*	*Units*
Crosstalk Source Rise/Fall time (20% to 80%)	Tr, Tf	See 1, 2, D.6		34		ps
Crosstalk Source Amplitude (p-p differential)		See 1, 2, D.7		1000		mV
Signal Rise/Fall time (20% to 80%)	Tr, Tf	See D.6	34			ps
Total Jitter	TJ	See D.5			0.28	UI(p-p)
Data Dependent Jitter	DDJ	See D.3			0.1	UI(p-p)
Data Dependent Pulse Width Shrinkage	DDPWS				0.055	UI (p-p)
Uncorrelated Jitter	UJ	See 3 and D.4			0.023	UI (RMS)
Transmitter Q_{sq}	Q_{sq}	See 4	50			
Parameters- B	*Symbol*	*Conditions*	*Value*			*Units*
Eye Mask	X1	Mask hit ratio of 5×10^{-5}, See D.2 and Figure 19	0.12			UI
Eye Mask	X2		0.33			UI
Eye Mask	Y1		95			mV
Eye Mask	Y2		350			mV

1. Measured at C" with Host Compliance Board and Module Compliance Board pair, see Figure 17.
2. Since the minimum module output transition time is faster than the crosstalk transition time the amplitude of crosstalk source is increased to achieve the same slew rate.
3. It is not possible to have the maximum UJ and meet the TJ specifications if the UJ is all Gaussian.
4. Q_{sq}=1/RN if the one level and zero level noises are identical and see D.8.

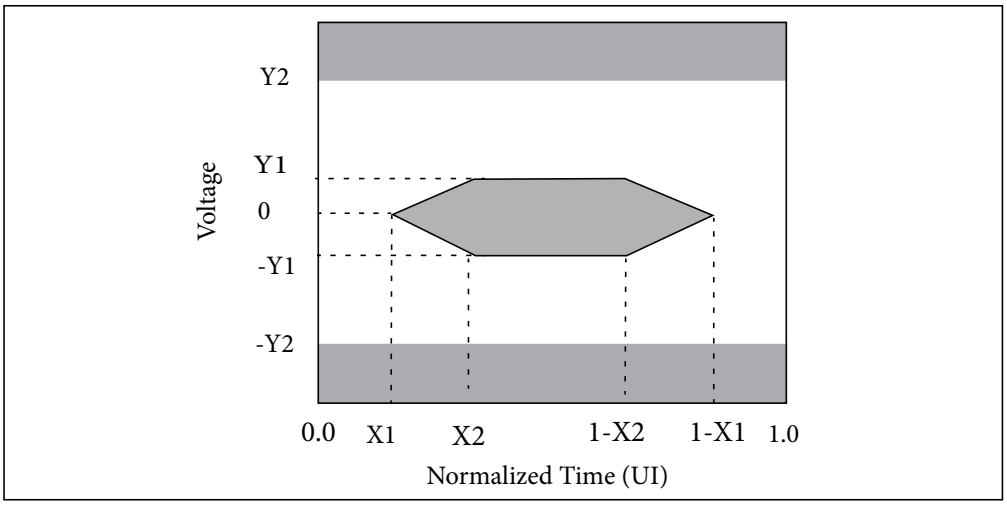


Figure 19 Transmitter Differential Output Compliance Mask at B and B"

3.5.2 HOST RECEIVER INPUT SPECIFICATIONS AT C AND C"

The SFP+ Host receiver electrical specifications at compliance point C and C" for both linear and limiting modules are given in Table 13. The host shall provide differential termination and must constrain differential to common mode conversion for quality signal termination and low EMI, as given in Table 13. Common mode termination on the receiver is not required see Figure 18.

Signals used as input tolerance test conditions are calibrated at C" with the Host Compliance Board connected through a Module Compliance Board to measurement instrumentation. Specifications at C" supporting limiting modules are given in Table 14. Specifications at C" supporting linear module are given in Table 15.

SFP+ compliant hosts are allowed to support just linear modules, just limiting modules, or both linear and limiting modules.

Table 13 Host Receiver Input Electrical Specifications at C and C"

Parameters - C and C"	*Symbol*	*Conditions*	*Min*	*Target*	*Max*	*Units*
Single Ended Input Voltage Range		Referenced to VeeR	-0.3		4.0	V
Input AC Common Mode Voltage Tolerance		See 1 and D.15.3	7.5			mV (RMS)
Damage Threshold (p-p differential)		See 1	2000			mV
Differential Input S-parameter	SDD11	0.01 to 2 GHz			-12	dB
		2 to 11.1 GHz			see 2	dB
Reflected Differential to Common Mode Conversion	SCD11	0.1 to 11.1 GHz			-10	dB

1. Measured at C" with the Module Compliance Board.
2. Reflection Coefficient given by equation SDD11(dB)< -6.68 + 12.1 × $\log_{10}$(f/5.5), with f in GHz .

Jitter specifications to support the limiting module are listed in Table 14. Figure 20 gives the host compliance eye mask requirements to support the limiting module. The host shall operate at and between the sensitivity and overload limits. The SFP+ limiting host shall tolerate sinusoidal jitter given by Figure 21. Test procedures for the host for limiting module are given in D.11.

Table 14 Host receiver supporting limiting module input compliance test signal calibrated at C″

Parameters - C″	*Symbol*	*Conditions*	*Target Value*	*Max*	*Units*
Crosstalk Source Rise/Fall time (20% to 80%)	Tr, Tf	See D.6	34		ps
Crosstalk Source Amplitude (p-p differential)		See 1	700		mV
AC Common Mode Voltage		See 2 and D.15		7.5	mV (RMS)
99% Jitter	J2	See 3, D.5, D.11	0.42		UI (p-p)
Pulse Width Shrinkage Jitter	DDPWS	See 4, D.3	0.3		UI (p-p)
Total Jitter	TJ	BER 1×10^{-12} see D.5, D.11	0.70		UI (p-p)
Eye Mask	X1	Mask hit ratio of 1×10^{-12}, See D.2, D.11	0.35		UI
Eye Mask Amplitude Sensitivity [5,8]	Y1		150		mV
Eye Mask Amplitude Overload [6,7,8]	Y2		425		mV

1. Measured at B" with Host Compliance Board and Module Compliance Board pair, see Figure 16.
2. The tester is not expected to generate this common mode voltage however its output must not exceed this value.
3. Includes sinusoidal jitter, per Figure 21, when measured with the reference PLL specified by the given standard.
4. In practice the test implementer may trade DDPWS with other pulse width shrinkage from the sinusoidal interferer.
5. Eye mask amplitude sensitivity tests the host receiver with the minimum eye opening expected from a module within the constraint set by Y2.
6. Eye mask amplitude overload tests the host receiver tolerance to the largest peak signal levels expected from the module within the constraint set by Y1.
7. It is not expected that module Rx output will exhibit both maximum peak level and minimum eye opening.
8. Sensitivity and overload are tested separately, see D.11.

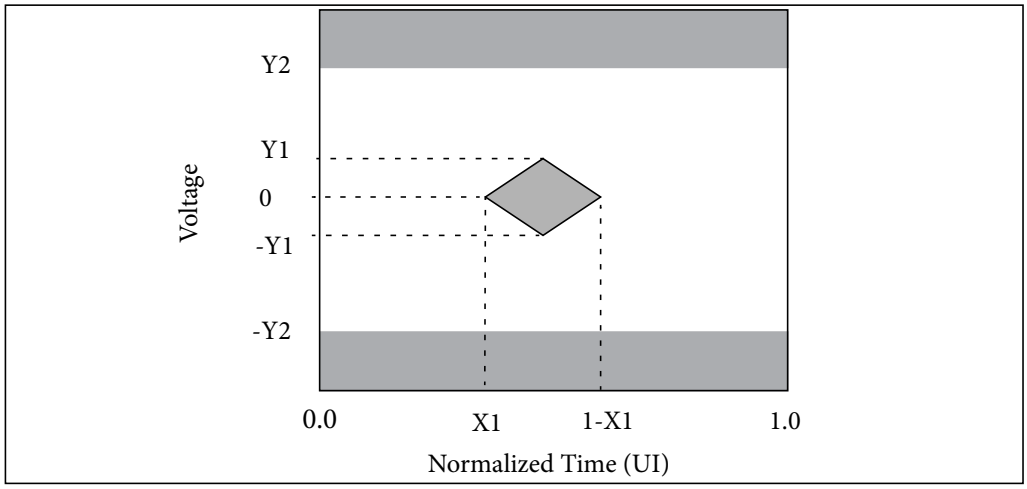


Figure 20 Host Receiver Input Compliance Mask at C" Supporting Limiting Module

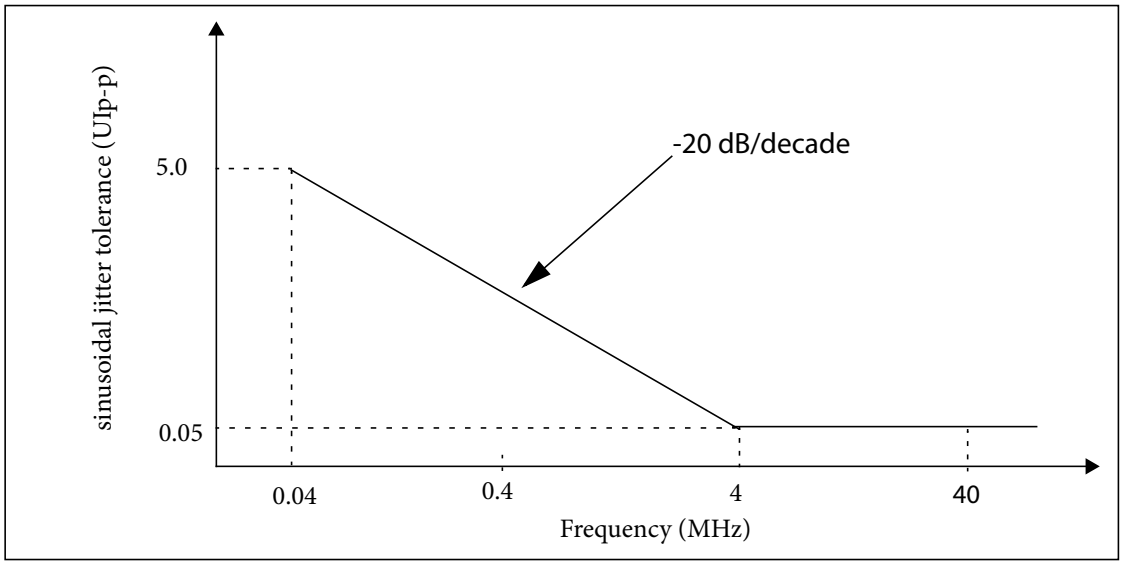


Figure 21 SR and LR Host Sinusoidal Jitter Tolerance Mask

Table 15 defines the input compliance test signal as calibrated at C" for a host that supports linear modules. The parameters in Table 15 include the effects of a worst case module that operates in conjunction with optical TP3 tester(s) defined for the LRM and LR standards. SR specifications are covered by the fact that LR links have high noise, and on the other extreme, LRM links have high distortion. Test procedures for the linear host are given in D.13. For illustrative purpose, Figure 24 shows the host test calibration line along which specific host test points for LRM are defined.

For LR test conditions, the SFP+ linear host shall operate with sinusoidal jitter given by Figure 21 while the stress conditions given in Table 15 are applied. For LRM test conditions, the host shall operate with sinusoidal jitter as defined in IEEE802.3, clause 68, with the stressors and noises in Figure 52 including those in the TP3 tester turned off.

Only two specific test conditions for each LRM stressor are defined in Table 15. In general, however, a host must meet operational requirements with any compliant module. It is expected that lower dWDP modules will exist. However, this specification has not defined host test conditions below dWDP of 0.6 to 0.8 dB. At low dWDP values, guard bands between module specifications and host requirements are left to the host implementation.

.

Table 15 Host receiver supporting linear module input compliance test signal calibrated at C”

Parameters - C"	*Symbol*	*Conditions*	*Min*	*Target*	*Max*	*Units*
Crosstalk Source Rise/Fall time (20% to 80%)	Tr/Tf	See 1, D.6		34		ps
Crosstalk Source Amplitude (p-p differential)		See 1		700		mV
AC Common Mode Voltage		See 2 and D.15			7.5	mV (RMS)
Differential Voltage Modulation Amplitude	VMA	for LRM, See 3	180		600	mV
Differential Voltage Modulation Amplitude	VMA	for SR and LR, See 3	150		600	mV
Applications	*Symbol*	*Compliance stress test conditions* ***See 4, 7, and D.13***	*Target WDP (dBo)*	*Target RN, (RMS)*		*WDPi (dBo)*
				m	*b*	
LRM	WDP	High WDP & pre-cursor stressor	Approximately 5.1, see 5	-0.0148	0.0477	4.1
	WDP	High WDP & split-symmetric stressor	5.4			3.9
	WDP	High WDP & post-cursor stressor	Approximately 5.2, see 5			4.2
	WDP	Low WDP & pre-cursor stressor	Approximately 4.7, see 6			4.1
	WDP	Low WDP & split-symmetric stressor	Approximately 4.7, see 6			3.9
	WDP	Low WDP & post-cursor stressor	Approximately 4.8, see 6			4.2
LR	WDP	Low WDP	Approximately 2.6, see 6	-0.02	0.096	1.9

1. Measured at B” with Host Compliance Board and Module Compliance Board pair, see Figure 16.

2. The tester is not expected to generate this common mode voltage, however its output must not exceed this value.

3. Peak levels of received signals in service may exceed their VMA due to overshoot of the far end transmitter and/or the module receiver.

4. Target WDP is calibrated with a reference receiver with 14 T/2 spaced FFE taps and 5 T spaced DFE taps.

5. The filter bandwidth in the TP3 to electrical adapter in Figure 52 is set to produce 5.4 dBo for WDP for the split-symmetrical TP3 stressor. The same filter is to be used for high WDP pre-cursor and post-cursor LRM stressors - their approximate target WDP values are given only for guidance. WDP is to be measured for each stressor, and target RN is determined by the relevant equation in note 7.

6. The filter bandwidth in the TP3 to electrical adapter in Figure 52 is set to 7.5 GHz for all three LRM low WDP conditions and for the LR condition. The approximate target WDP values are given for guidance. WDP is to be measured for each stressor, and target RN is determined by the relevant equation in note 7.

7. Target RN rms values are given by the following equation: $RN = m \times (WDP - WDP_i) + b$, where WDP is the actual value of the tester, and WDP_i values are based on waveshapes expected at TP3.

3.6 SFP+ MODULE SPECIFICATIONS

SFP+ module transmitter specifications at compliance point B′ are given in 3.6.1. SFP+ module receiver specifications at compliance point C′ are given in 3.6.2.

3.6.1 MODULE TRANSMITTER INPUT SPECIFICATIONS AT B′ AND B″

The SFP+ module transmitter electrical specifications, given in Table 16, at compliance point B′ are measured with the Module Compliance Board as shown in 3.3.2. The transmitter input impedance is 100 Ω differential. The module must provide differential termination and limit differential to common mode conversion for quality signal termination and low EMI.

Signals used as input conditions for testing the transmitter input tolerance are calibrated at B″ with the Module Compliance Board connected through a Host Compliance Board to appropriate instrumentation. This is further described in D.10. The specifications used for this calibration are listed in Table 17. The test signal at B″ as illustrated by Figure 17 shall comply with the mask defined in Table 17 and illustrated in Figure 19.

Table 16 Module Transmitter Input Electrical Specifications at B′

Parameters - B'	*Symbol*	*Conditions*	*Min*	*Target*	*Max*	*Units*
Single Ended Input Voltage Tolerance		Referenced to VeeT	-0.3		4.0	V
AC common mode voltage tolerance		See 1, D.15.3	15			mV
Differential Input S-parameter	SDD11	0.01 to 4.1 GHz			See 2	dB
		4.1 to 11.1 GHz			See 3	dB
Reflected Differential to Common Mode Conversion	SCD11	0.01 to 11.1 GHz			-10	dB

1. Measured at B" with Host Compliance Board and Module Compliance Board pair, see Figure 17.
2. Reflection Coefficient given by equation SDD11(dB)< -12 + 2 × SQRT(f), with f in GHz.
3. Reflection Coefficient given by equation SDD11(dB)< -6.3 + 13 × log10(f/5.5), with f in GHz

Table 17 Module Transmitter Input Tolerance Signal Calibrated at B"

Parameters- B"	*Symbol*	*Conditions*	*Target Value*	*Max*	*Units*
Crosstalk Source Rise/Fall time (20% to 80%)	Tr, Tf	See 1, 2 and D.6	34		ps
Crosstalk Source Amplitude (p-p differential)		See 1, 2	1000		mV
AC Common Mode Voltage		See 3 and D.15.2		15	mV (RMS)
Total Jitter	TJ	See D.5		0.28	UI (p-p)
Data Dependent Jitter	DDJ	See D.3	0.10		UI (p-p)
Pulse Width Shrinkage Jitter	DDPWS		0.055		UI (p-p)
Uncorrelated Jitter	UJ	See 4 and D.4	0.023		UI (RMS)
Parameters- B"	*Symbol*	*Conditions*	*Value*		*Units*
Eye Mask	X1	Mask hit ratio of 5×10^{-5}, See D.2	0.12		UI
Eye Mask	X2		0.33		UI
Eye Mask	Y1		95		mV
Eye Mask	Y2		350		mV

1. Measured at C" with Host Compliance Board and Module Compliance Board pair, see Figure 17.
2. Since the minimum module output transition time is faster than the crosstalk transition time the amplitude of crosstalk source is increased to achieve the same slew rate.
3. The tester is not expected to generate this common mode voltage however its output must not exceed this value
4. It is not possible to have the maximum UJ and meet the TJ specifications if the UJ is all Gaussian.

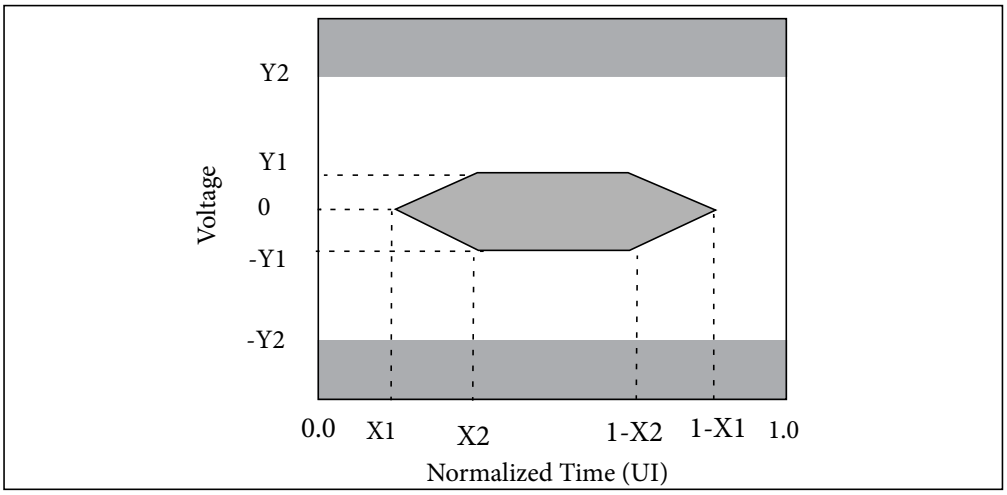


Figure 22 Module Transmitter Differential Input Compliance Mask at B"

3.6.2 MODULE RECEIVER OUTPUT SPECIFICATIONS AT C'

The SFP+ receiver electrical output specifications at compliance point C' are given in Table 18. The module must provide differential termination and common mode termination for quality signal termination and low EMI, as given in Table 18.

Table 18 Module Receiver Output Electrical Specifications at C'

Parameters - C'	*Symbol*	*Conditions*	*Min*	*Target*	*Max*	*Units*
Crosstalk source rise/fall time (20% to 80%)	Tr, Tf	See 1, D.6,		34		ps
Crosstalk Source Amplitude (p-p differential)		See 1		700		mV
Termination Mismatch at 1 MHz	ΔZ_M	See D.16, Figure 55			5	%
Single Ended Output Voltage Tolerance			-0.3		4.0	V
Output AC Common Mode Voltage		See D.15			7.5	mV (RMS)
Differential Output S-parameter	SDD22	0.01 to 4.1 GHz			See 2	dB
		4.1 to 11.1 GHz			See 3	dB
Common Mode Output Reflection Coefficient	SCC22	0.01 to 2.5 GHz			See 4	dB
		2.5 to 11.1 GHz			-3	dB

1. Measured at B" with the Host Compliance Board and Module Compliance Board pair, see Figure 16.
2. Reflection Coefficient given by equation SDD22(dB)< -12 + 2 × SQRT(f), with f in GHz.
3. Reflection Coefficient given by equation SDD22(dB)< -6.3 + 13 × log10(f/5.5), with f in GHz
4. Reflection coefficient given by equation SCC22(dB) < -7 + 1.6 × f, with f in GHz.

Common Mode Output Reflection Coefficient helps absorb reflection and noise improving EMI.

Jitter specifications for limiting modules are listed in Table 19. Figure 23 gives the compliance eye mask for limiting modules output. Requirements for linear modules are given in Table 20.

Both limiting and linear modules must provide adequate low frequency signal response for the applications supported, to control the effects of baseline wander.

Table 19 Limiting Module Receiver Output Jitter and Eye Mask Specifications at C'

Parameters - C'	*Symbol*	*Conditions*	*Min*	*Target*	*Max*	*Units*
Output Rise and Fall time (20% to 80%)	Tr, Tf	See D.6	28			ps
Total Jitter	TJ	See D.5, D.12			0.70	UI (p-p)
99% Jitter	J2	See D.5, D.12			0.42	UI (p-p)
Parameters - C'	*Symbol*	*Conditions*	*Value*			*Units*
Eye Mask	X1	Mask hit ratio of 1×10^{-12} See D.2, D.12	0.35			UI
Eye Mask	Y1		150			mV
Eye Mask	Y2		425			mV

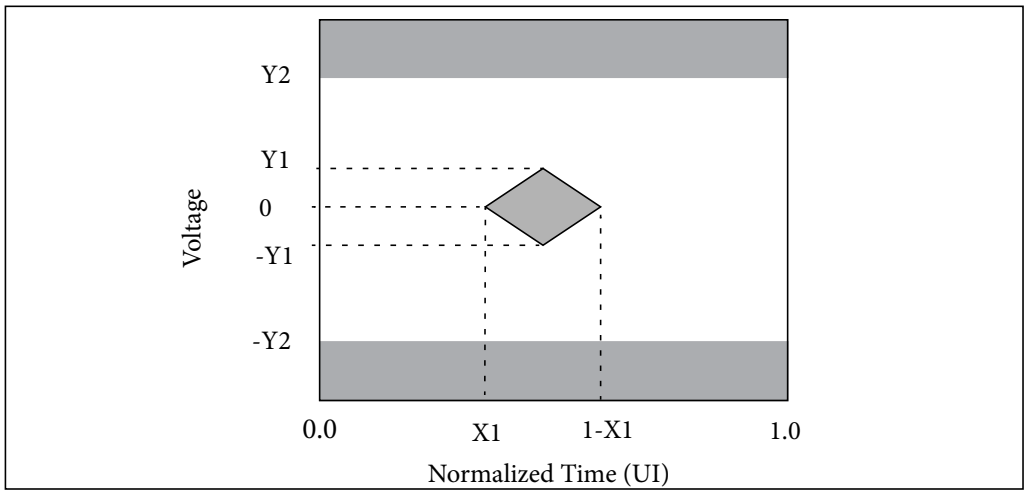


Figure 23 Limiting Module Receiver Differential Output Compliance Mask at C'

Linear module test parameters are given by Table 20. Compliance methods for a linear module are given in Appendix D.14.

Table 20 Linear Module Receiver Specifications at C'

Parameters - C'	*Symbol*	*Conditions*	*RN*				
			m1	*b1*	*m2*	*b2*	*RNmax*
Relative Noise SR	RN	See D.14.1	-0.02	0.078	-0.02	0.078	0.078
Relative Noise LR			-0.02	0.083	-0.02	0.083	0.083
Relative Noise LRM with pre-cursor stressor			-0.0153	0.0475	-0.0092	0.0457	0.0475
Relative Noise LRM with split-symmetrical stressor			-0.017	0.0475	-0.01	0.0455	0.0475
Relative Noise LRM with post-cursor stressor			-0.0153	0.0475	-0.0092	0.0457	0.0475
Parameters - C'	*Symbol*	*Conditions*		*Min*		*Max*	*Unit*
Difference Waveform Distortion Penalty for SR and LR	dWDP	See 1 and D.14.2				2.7	dBo
Difference Waveform Distortion Penalty for LRM	dWDP					1.5	dBo
Differential Voltage Modulation Amplitude For SR and LR	VMA	See D.7		150		600	mV
Differential Voltage Modulation Amplitude For LRM	VMA			180		600	mV
Differential peak to peak voltage	Vpk-pk	See D.14.3				600	mV
1. Defined with reference receiver with 14 T/2 spaced FFE taps and 5 T spaced DFE taps.							

Appendix D.14.2 defines RN for a linear module receiver. The limits for RN are functions of measured dWDP for the module, expressed in optical decibels. As an example, the trade-off between the parameters for LRM are shown in Figure 24. To pass, RN must be below the respective limit line.

dWDP and RN shall meet the specifications in Table 20 and can be calculated by

$$RN \leq min[(m1 \times dWDP + b1), (m2 \times dWDP + b2), RNmax]$$

for each TP3 test condition for which compliance is required. For example, if compliance is required for LRM, the module must meet specifications with all three stressors and under the sensitivity and overload test conditions specified in IEEE Std 802.3 68.6.9.

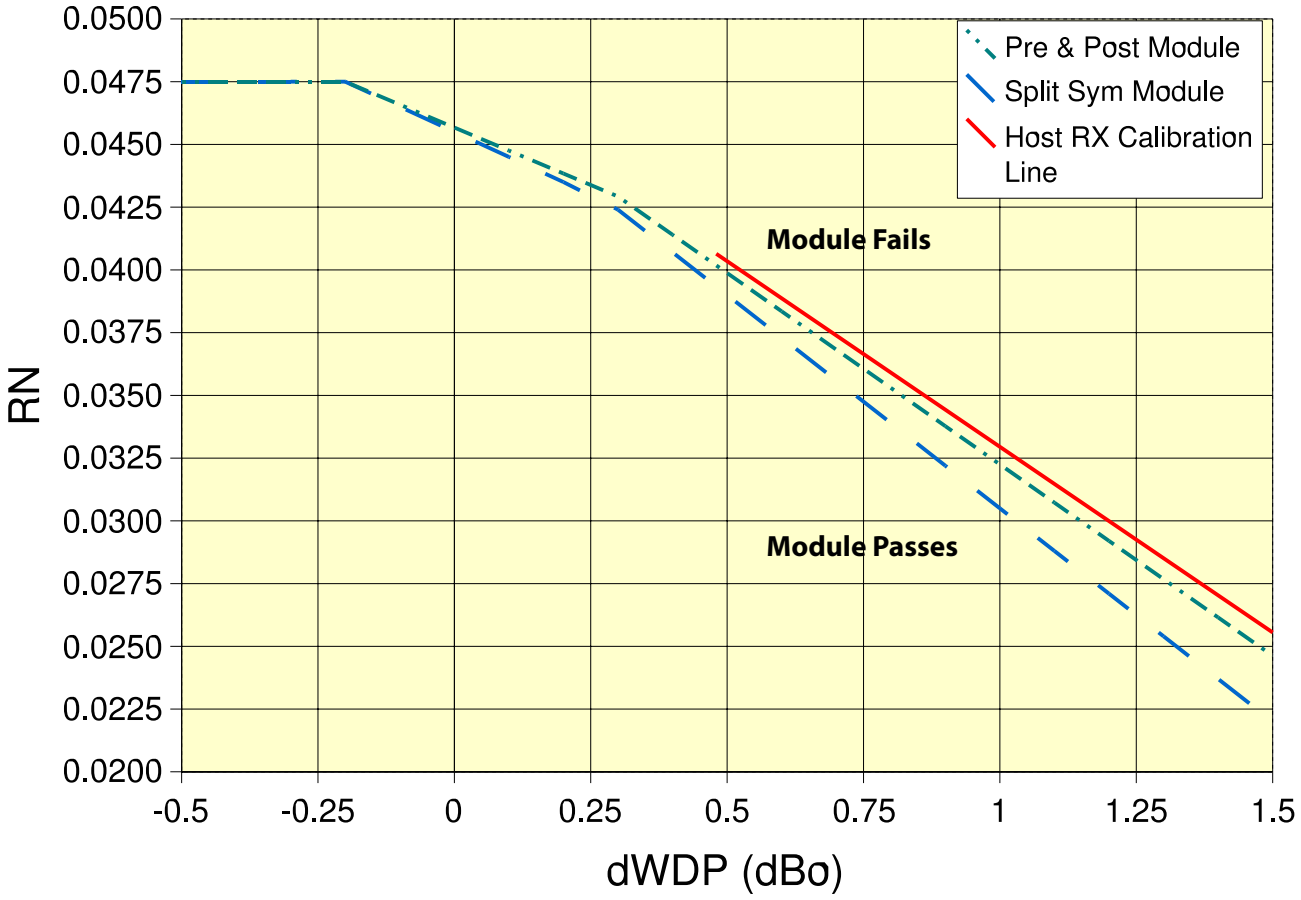


Figure 24 LRM module receiver RN and dWDP compliance and host receiver test calibration

For illustrative purpose, Figure 24 shows the host test calibration line along which specific host test points for LRM are defined. The gap between the host and module lines is because the host is tested with linear impairments, which given the same dWDP, are more benign to a host than non-linear impairments which are possible from a module.

CHAPTER 4 SFP+ 2-WIRE INTERFACE

4.1 INTRODUCTION

The SFP+ management interface is a two-wire interface, similar to I^2C. SFP+ management memory map is specified by SFF-8472. Nomenclature for all registers more than 1 bit long are MSB...LSB (MSB transmitted first).

4.2 2-WIRE ELECTRICAL SPECIFICATIONS

The SFP+ 2-wire interface specifications are given in Table 21. This specification ensures compatibility between host masters and SFP+ SCL/SDA lines and compatibility with I^2C. All voltages are referenced to VeeT.

Table 21 2-Wire Interface Electrical Specifications

Parameter	*Symbol*	*Min.*	*Max.*	*Unit*	*Conditions*
Host 2-wire Vcc	Vcc_Host_2w	3.14	3.46	V	see 1
SCL and SDA	V_{OL}	0.0	0.40	V	Rp2w [2] pulled toVcc_Host_2w, see 3
	V_{OH}	Vcc_Host_2w - 0.5	Vcc_Host_2w + 0.3	V	Rp2w [2] pulled to Vcc_Host_2w, see 3
SCL and SDA	V_{IL}	-0.3	VccT*0.3	V	see 3
	V_{IH}	VccT*0.7	VccT + 0.5	V	see 3
Input current on the SCL and SDA contacts	I_I	-10	10	µA	
Capacitance on SCL and SDA I/O contact	C_i [4]		14	pF	
Total bus capacitance for SCL and for SDA	C_b [5]		100	pF	At 400 kHz, 3.0 kΩ Rp2w, max At 100 kHz, 8.0 kΩ Rp2w, max
			290	pF	At 400 kHz, 1.1 kΩ Rp2w, max At 100 kHz, 2.75 kΩ Rp2w, max

1. The Host 2-wire Vcc is the voltage used for resistive pull ups for the 2 wire interface
2. Rp2w is the pull up resistor. Active bus termination may be used by the host in place of a pullup resistor. Pull ups can be connected to any one of several power supplies, however the host board design shall ensure that no module contact has voltage exceeding module VccT/R + 0.5 V nor requires the module to sink more than 3.0 mA current.
3. These voltages are measured on the other side of the connector to the device under test.
4. C_i is the capacitance looking into the module SCL and SDA contacts
5. C_b is the total bus capacitance on the SCL or SDA bus.

4.3 SFP+ 2-WIRE TIMING DIAGRAM

SFP+ 2-wire bus timing is shown in Figure 25 and the detail of clock stretching is shown in Figure 26. SFP+ 2-wire timing specifications are given in Table 22.

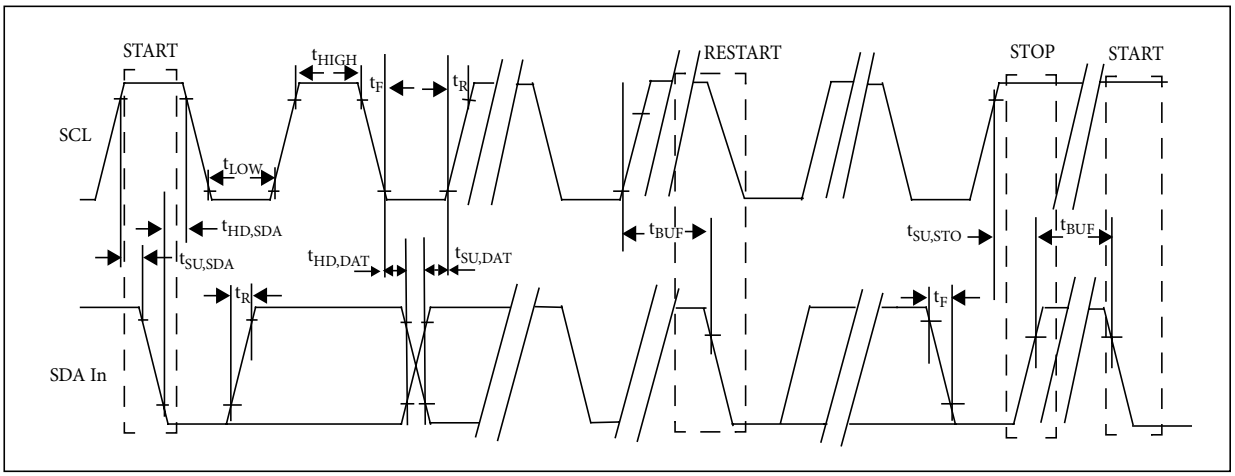


Figure 25 2-wire Timing Diagram

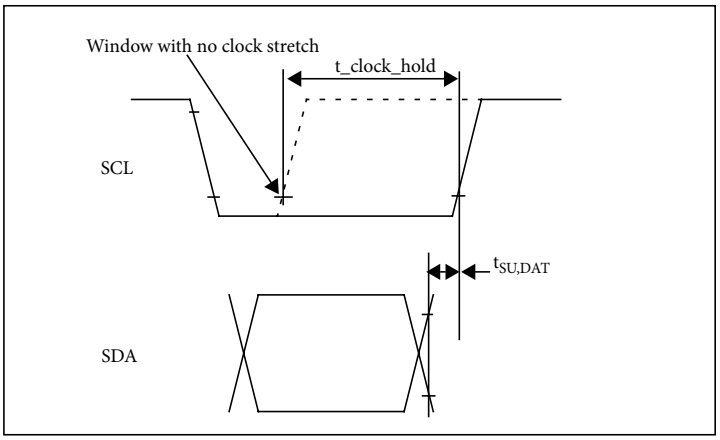


Figure 26 Detail of Clock Stretching

The 2-wire serial interface addresses of the SFP+ module are 1010000x (A0h) and 1010001x (A2h).

Table 22 SFP+ 2-wire Timing Specifications

Parameter	*Symbol*	*Min.*	*Max.*	*Unit*	*Conditions*
Clock Frequency	f_{SCL}	0	400	kHz	Module shall operate with f_{SCL} up to 100 kHz without requiring clock stretching. The module may clock stretch with f_{SCL} greater than 100 kHz and up to 400 kHz.
Clock Pulse Width Low	t_{LOW}	1.3		µs	
Clock Pulse Width High	t_{HIGH}	0.6		µs	
Time bus free before new transmission can start	t_{BUF}	20		µs	Between STOP and START and between ACK and ReSTART
START Hold Time	$t_{HD,STA}$	0.6		µs	
START Set-up Time	$t_{SU,STA}$	0.6		µs	
Data In Hold Time	$t_{HD,DAT}$	0		µs	
Data In Set-up Time	$t_{SU,DAT}$	0.1		µs	
Input Rise Time (100 kHz)	$t_{R,100}$		1000	ns	From ($V_{IL,MAX}$ - 0.15) to ($V_{IH,MIN}$ + 0.15)
Input Rise Time (400 kHz)	$t_{R,400}$		300	ns	From ($V_{IL,MAX}$ - 0.15) to ($V_{IH,MIN}$ + 0.15)
Input Fall Time (100 kHz)	$t_{F,100}$		300	ns	From ($V_{IH,MIN}$ + 0.15) to ($V_{IL,MAX}$ - 0.15)
Input Fall Time (400 kHz)	$t_{F,400}$		300	ns	From ($V_{IH,MIN}$ + 0.15) to ($V_{IL,MAX}$ - 0.15)
STOP Set-up Time	$t_{SU,STO}$	0.6		µs	

4.4 MEMORY TRANSACTION TIMING

SFP+ memory transaction timings are given in Table 23.

Table 23 SFP+ Memory Specifications

Parameter	*Symbol*	*Min.*	*Max.*	*Unit*	*Conditions*
Serial Interface Clock Holdoff "Clock Stretching"	T_clock_hold		500	µs	Maximum time the SFP+ module may hold the SCL line low before continuing with a read or write operation
Complete Single or Sequential Write up to 4 Byte	t_{WR}		40	ms	
Complete Sequential Write of 5-8 Byte	t_{WR}		80	ms	
Endurance (Write Cycles)		10 k		cycles	

4.5 DEVICE ADDRESSING AND OPERATION

Serial Clock (SCL): The host supplied SCL input to SFP+ transceivers is used to positively edge clock data into each SFP+ device and negative edge clock data out of each device. The SCL line may be pulled low by an SFP+ module during clock stretching.

Serial Data (SDA): The SDA contact is bi-directional for serial data transfer. This contact is open-drain or open-collector driven and may be wire-ORed with other open-drain or open collector devices with different device addresses, provided the total bus capacitance meets the requirement of Table 21 and the Serial Clock (SCL) is also wire-ORed.

Master/Slave: SFP+ transceivers operate only as slave devices. The host must provide a bus master for SCL and initiate all read/write communication.

Device Address: Each SFP+ is hard wired at the device addresses A0h and A2h. See SFF-8472 for memory structure within each transceiver.

Clock and Data Transitions: The SDA contact is normally pulled high with an external device. Data on the SDA contact may change only during SCL low time periods. Data changes during SCL high periods indicate a START or STOP condition. All addresses and data words are serially transmitted to and from the SFP+ in 8-bit words. Every byte on the SDA line must be 8-bits long. Data is transferred with the most significant bit (MSB) first.

START Condition: A high-to-low transition of SDA with SCL high is a START condition, which must precede any other command.

STOP Condition: A low-to-high transition of SDA with SCL high is a STOP condition.

Acknowledge: After sending each 8-bit word, the transmitter releases the SDA line for one bit time, during which the receiver is allowed to pull SDA low (zero) to acknowledge (ACK) that it has received each word. Device address bytes and write data bytes initiated by the host shall be acknowledged by SFP+ transceivers. Read data bytes transmitted by SFP+ transceivers shall be acknowledged by the host for all but the final byte read, for which the host shall respond with a STOP instead of an ACK.

Non-acknowledge (NACK): When a slave is unable to receive or transmit, because, e.g., it is performing a higher priority function, the data line shall be left high by the slave. A NACK is generated when the slave leaves the data line high during the ACK clock pulse. The master can then generate either a STOP condition to abort the transfer or a repeated START condition to start a new transfer.

When in a transfer, a master-receiver must signal the end of data to the slave-transmitter by not generating an acknowledge on the last byte clocked out of the slave. A NACK is generated when the master leaves the data line High during the ACK clock pulse. The slave-transmitter must release the data line to permit the master to generate a STOP or repeated START condition.

Memory (Management Interface) Reset: After an interruption in protocol, power loss or system reset the SFP+ management interface can be reset. Memory reset is intended only to reset the SFP+ transceiver management interface (to correct a hung bus). No other transceiver functionality is implied.

1) Clock up to 9 cycles.
2) Look for SDA high in each cycle while SCL is high.
3) Create a START condition as SDA is high.

Device Addressing: SFP+ devices require an 8 bit device address word following a start condition to enable a read or write operation. The device addresses to select A0h or A2h are shown in Table 24. This is common to all SFP+ devices.

Table 24 SFP+ Device Address Word

Address	(MSB)						Address Select	R/W select (LSB)
A0h	1	0	1	0	0	0	0	x
A2h	1	0	1	0	0	0	1	x

The LSB of the device address word is the read/write operating select bit. A read operation is initiated if this bit is set high and a write operation is initiated if this bit is set low.

4.6 READ/WRITE FUNCTIONALITY

The methods for reading from and writing to the two different SFP+ addresses A0h and A2h are described in this section. They are identical for the two different addresses except that the appropriate address is used for each read and write. For simplicity in the figures the address is labelled 101000x where the x is 0 for the A0h address and 1 for the A2h address. Note that the address here is only seven bits. In order to complete the full 8 bit byte a one or zero is added to the end of the address depending on whether a read or a write operation is taking place.

4.6.1 SFP+ MEMORY ADDRESS COUNTER (READ AND WRITE OPERATIONS)

SFP+ devices maintain two internal data word address counters one for each address. These counters contain the last address accessed during the latest

read or write operation, incremented by one. The address counter is incremented whenever a data word is received or sent by the transceiver. This address stays valid between operations as long as SFP+ power is maintained. The address "roll over" during read and write operations is from the last byte of the 256 byte memory page to the first byte of the same page.

4.6.2 READ OPERATIONS (CURRENT ADDRESS READ)

A current address read operation requires only the device address read word (10100001 or 10100011) be sent, Figure 27. Once acknowledged by the SFP+, the current address data word is serially clocked out. The host does not respond with an acknowledge, but does generate a STOP condition once the data word is read.

		< -- SFP+ ADDRESS ->																			
HOST	START	MSB						LSB	READ										NACK	STOP
		1	0	1	0	0	0	x	1	0	x	x	x	x	x	x	x	x	1	
SFP+										ACK	MSB							LSB		
											<----- DATA WORD ----->									

Figure 27 Current Address Read Operation

4.6.3 READ OPERATIONS (RANDOM READ)

A random read operation requires a "dummy" write operation to load in the target byte address Figure 28. This is accomplished by the following sequence: The target 8-bit data word address is sent following the device address write word (10100000 or 10100010) and acknowledged by the SFP+. The host then generates another START condition (aborting the dummy write without incrementing the counter) and a current address read by sending a device read address (10100001 or 10100011). The SFP+ acknowledges the device address and serially clocks out the requested data word. The host does not respond with an acknowledge, but does generate a STOP condition once the data word is read.

		< -- SFP+ ADDRESS ->									<- MEMORY ADDRESS ->										< -- SFP+ ADDRESS -->																		
HOST	START	MSB						LSB	WRITE		MSB							LSB		START	MSB						LSB	READ										NACK	STOP
		1	0	1	0	0	0	x	0	0	x	x	x	x	x	x	x	x	0		1	0	1	0	0	0	x	1	0	x	x	x	x	x	x	x	x	1	
SFP+										ACK									ACK										ACK	MSB							LSB		
																														<---- DATA WORD ---->									

Figure 28 Random Read

4.6.4 READ OPERATIONS (SEQUENTIAL READ)

Sequential reads are initiated by either a current word address read Figure 29 or a random address read Figure 30. To specify a sequential read, the host responds with an acknowledge (instead of a STOP) after each data word. As long as the SFP+ receives an acknowledge, it shall serially clock out sequential data words. The sequence is terminated when the host responds with a NACK and a STOP instead of an acknowledge.

		< -- SFP+ ADDRESS ->																																			
HOST	START	MSB						LSB	READ										ACK									ACK									NACK	STOP
		1	0	1	0	0	0	x	1	0	x	x	x	x	x	x	x	x	0	x	x	x	x	x	x	x	x	0	x	x	x	x	x	x	x	x	1	
SFP+										ACK	MSB							LSB		MSB							LSB		MSB							LSB		
											<---- DATA WORD n -->									<--- DATA WORD n+1 --->									<-- DATA WORD n+x -->									

Figure 29 Sequential Address Read Starting at Current Word Address

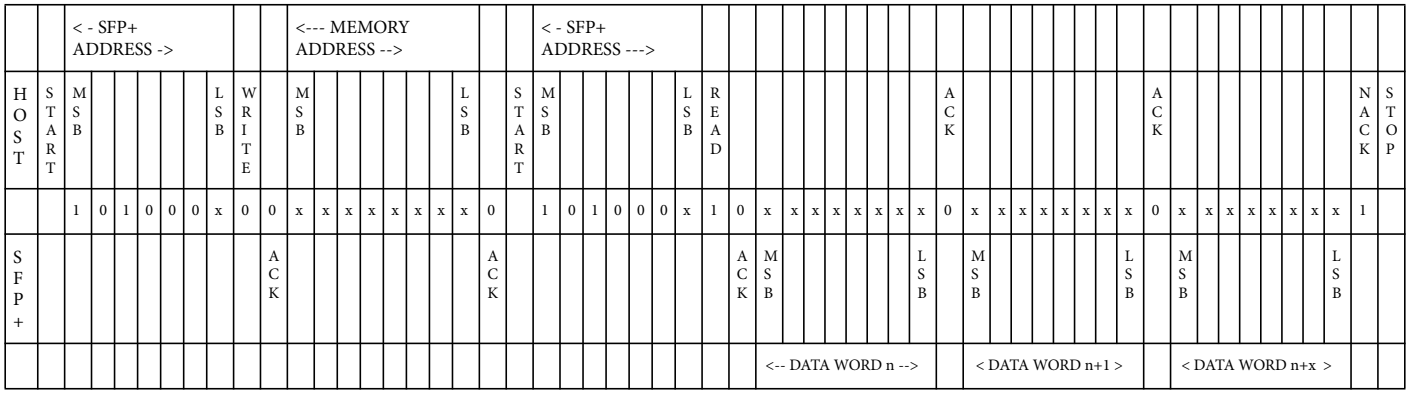


Figure 30 Sequential Address Read Starting with Random Read

4.6.5 WRITE OPERATIONS (BYTE WRITE)

A write operation requires an 8-bit data word address following the device address write word (10100000 or 10100010) and acknowledgement Figure 31. Upon receipt of this address, the SFP+ shall again respond with a zero (ACK) to acknowledge and then clock in the first 8 bit data word. Following the receipt of the 8 bit data word, the SFP+ shall output a zero (ACK) and the host master must terminate the write sequence with a STOP condition for the write cycle to begin. If a START condition is sent in place of a STOP condition (i.e. a repeated START per the I^2C specification) the write is aborted and the data received during that operation is discarded. Upon receipt of the proper STOP condition, the SFP+ enters an internally timed write cycle, t_{WR}, to internal memory. The SFP+ disables its management interface input during this write cycle and shall not respond or acknowledge subsequent commands until the write is complete. Note that I^2C "Combined Format" using repeated START conditions is not supported on SFP+ write commands.

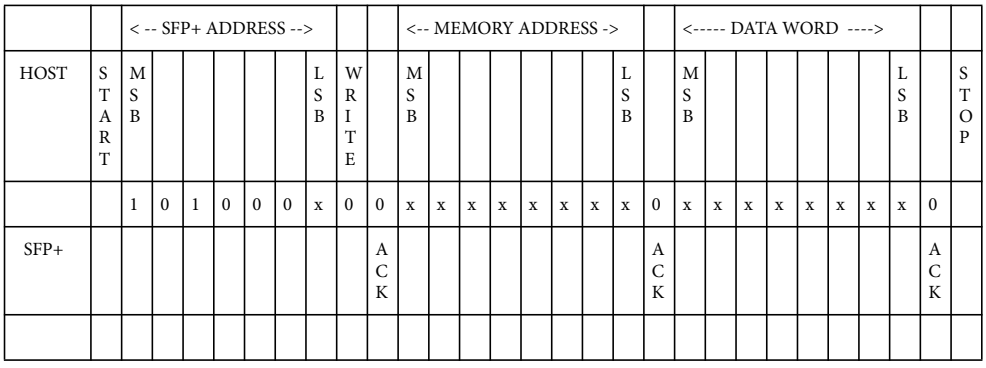


Figure 31 SFP+ Write Byte Operation

4.6.6 WRITE OPERATIONS (SEQUENTIAL WRITE)

SFP+ shall support up to an 8 sequential byte write without repeatedly sending SFP+ address and memory address information. A “sequential” write is initiated the same way as a single byte write, but the host master does not send a stop condition after the first word is clocked in. Instead, after the SFP+ acknowledges receipt of the first data word, the host can transmit up to seven more data words. The SFP+ shall send an acknowledge after each data word received. The host must terminate the sequential write sequence with a STOP condition or the write operation shall be aborted and data discarded. Note that I^2C “combined format” using repeated START conditions is not supported on SFP+ write commands.

		< --- SFP+ ADDRESS --->									<--- MEMORY ADDRESS -->									<--DATA WORD 1 -->									<--- DATA WORD 2 -->									<------ DATA WORD 3 ------>									<------ DATA WORD 4 ------>									
HOST	START	MSB						LSB	WRITE		MSB							LSB		MSB							LSB		MSB							LSB		MSB							LSB		MSB							LSB		STOP
		1	0	1	0	0	0	x	0	0	x	x	x	x	x	x	x	x	0	x	x	x	x	x	x	x	x	0	x	x	x	x	x	x	x	x	0	x	x	x	x	x	x	x	x	0	x	x	x	x	x	x	x	x	0	
SFP+										ACK									ACK									ACK									ACK									ACK									ACK	

Figure 32 Sequential Write Operation

4.6.7 WRITE OPERATIONS (ACKNOWLEDGE POLLING)

Once the SFP+ internally timed write cycle has begun (and inputs are being ignored on the bus) acknowledge polling can be used to determine when the write operation is complete. This involves sending a START condition followed by the device address word. Only if the internal write cycle is complete shall the SFP+ respond with an acknowledge to subsequent commands, indicating read or write operations can continue.

APPENDIX A SFI CHANNEL RECOMMENDATION (INFORMATIVE)

A.1 SFI HOST CHANNEL GENERAL RECOMMENDATIONS

The purpose of the recommended SFI channel is to provide guidelines for host designers. The recommended SFI host channel consists of PCB traces, vias, and the 20 position enhanced connector defined by SFF-8083. The PCB traces are recommended to meet 100 ±10 Ω differential impedance with nominal 7% differential coupling.

SFI channel S-parameters are defined from ASIC transmitter pads to Host Compliance Board output at B and from Host Compliance Board input at C to ASIC input pads.

Please see SFF INF-8077i for differential S-parameters measurements and conversions.

A.2 SFI CHANNEL TRANSFER RECOMMENDATIONS

The SFI maximum channel transfer budget is 9.0 dB allocated as shown in Table 25.

Table 25 SFI Host Interconnect Budget

Parameter	*Symbol*	*Conditions*	*Min*	*Max*	*Units*
Channel Transfer Including Connector measured with Host Compliance Board (see Appendix C)	SDD21	at 5.5 GHz, see 1	-6.5	-2.25	dB
Penalty for reflections and other impairments			-2.5		dB
Total Channel Link Budget When Measured with HCB			-9.0	-2.25	dB
1.SFI channel response (SDD21) is defined from chip pads to compliance point B or C.					

To mitigate multiple reflections, SFI also recommends a minimum channel attenuation. This requirement for both a minimum and maximum channel attenuation results in a mask that is shown approximately by Figure 33. The response including ripple should be within the channel response mask.

The SFI recommended channel is measured with the ASIC removed and measured with the Host Compliance Board of section C.2. The S-parameters are measured by connecting a 4-port network analyzer to the ASIC pads and the SMA connectors on the Host Compliance Board.

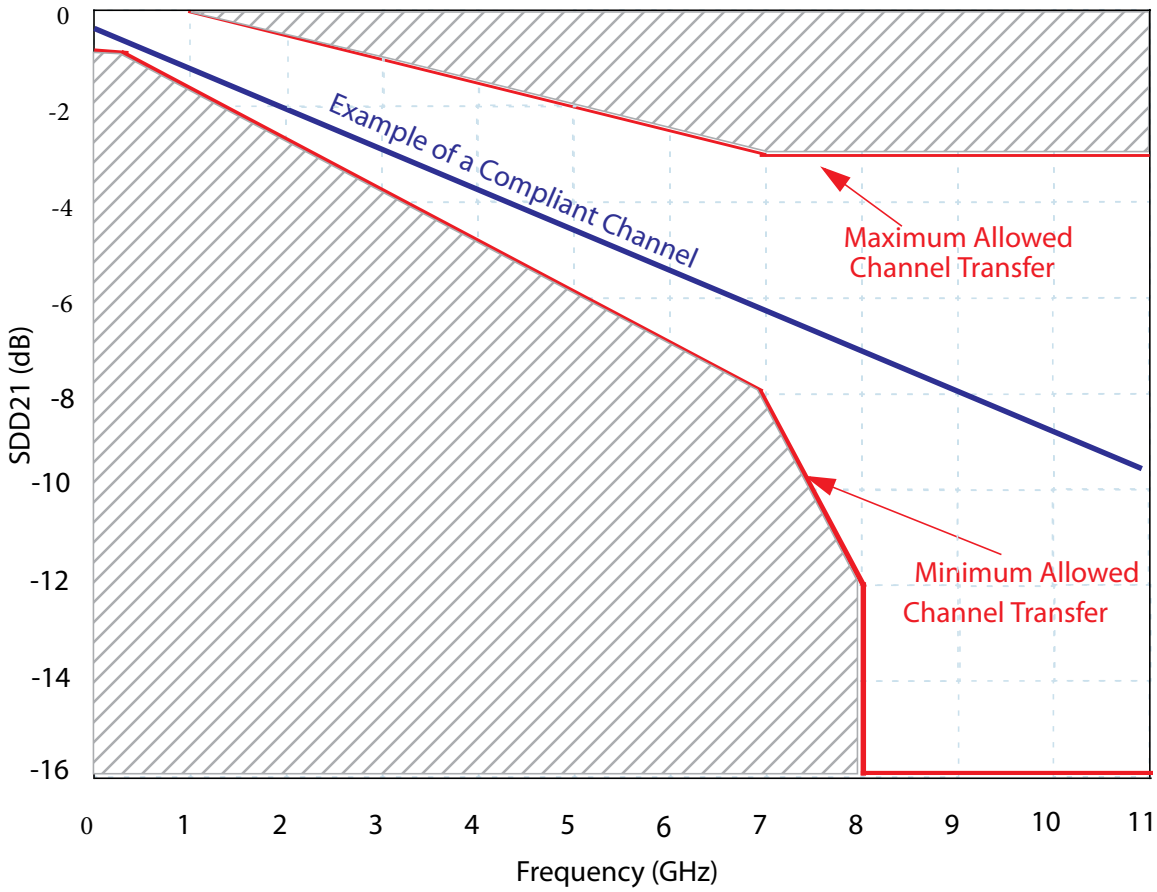


Figure 33 Example of SFI Host Recommended Channel

The minimum channel transfer SDD21 (maximum loss) mask contour is given by:

$$SDD21(dB) = -0.73 \qquad \text{f from 0.01 GHz to 0.25 GHz}$$

$$SDD21(dB) = (-0.108 - 0.845 \times \sqrt{f} - 0.802 \times f) \qquad \text{f from 0.25 GHz to 7 GHz}$$

$$SDD21(dB) = 20 - 4 \times f \qquad \text{f from 7 GHz to 8 GHz}$$

$$SDD21(dB) \geq -16 \qquad \text{f from 8 GHz to 11.1 GHz}$$

where *f* is the frequency in GHz.

The SFI channel maximum transfer is given by:.

$$SDD21(dB) = 0 \qquad \text{f from 0.25 GHz to 1.0 GHz}$$

$$SDD21(dB) = 0.5 \times (1 - f) \qquad \text{f from 1 GHz to 7 GHz}$$

$$SDD21(dB) = -3 \qquad \text{f from 7 GHz to 11.1 GHz}$$

where f is the frequency in GHz.

A.3 SFI Channel Return Loss Recommendations

The reflection coefficients, SDD11 and SDD22, of the SFI channel are recommended to meet the following equations:

$$SDDxx(dB) \leq -14.5 \qquad \text{f from 0.01 to 5 GHz}$$

$$SDDxx(dB) \leq -23.25 + 8.75 \times \left(\frac{f}{5}\right) \qquad \text{f from 5 to 11.1 GHz}$$

where *f* is the frequency in GHz and SDDxx is either SDD11 or SDD22.

A.4 SFI Channel Ripple Recommendations

SFI channel ripple is defined as the difference between the measured insertion response ($SDD21_m$) and the fitted transfer response (SDD21f), all in dB magnitude:

$$Ripple(dB) = SDD21_m - SDD21_f$$

The channel ripple magnitude should conform to the equation:

$$|Ripple(dB)| \leq 0.15 + 0.1 \times f$$

where the variable *f* (frequency) is in GHz. The above equation must be satisfied over the frequency range of 0.25 GHz to 5.5 GHz.

$SDD21_m$ is the measured channel differential transfer response. $SDD21_f$ is the fitted channel differential transfer response and is given by

$$SDD21_f = [-a - b \times \sqrt{f} - c \times f]$$

Where a, b, and c are determined by the least squares fit over the frequency range of 250 MHz to 5.5 GHz as defined below. Frequency steps should be of equal size and not greater than 50 MHz.

Measured data will provide a frequency vector, *f*, and gain vector, *G* defined by

$$G = 20 \times log_{10}[|SDD21|]$$

Create an input vector array called X from frequency variable f

$$X = \begin{bmatrix} 1 & \sqrt{f0} & f0 \\ 1 & \sqrt{f1} & f1 \\ . & . & . \\ 1 & \sqrt{fn} & fn \end{bmatrix}$$

Next calculate the coefficient vector using matrix math

$$C = [X^T \times X]^{-1} X^T \times G$$

Where the calculated coefficient values are given by

- a = -C(1)
- b = -C(2)
- c = -C(3).

APPENDIX B SFI ASIC/SERDES SPECIFICATION (INFORMATIVE)

B.1 INTRODUCTION

SFI ASIC/SerDes specifications are informative. SFI ASIC/SerDes Transmitter specifications at reference point A are given in B.2. SFI ASIC/SerDes Receiver specifications at reference point D are given in B.3. ASIC/SerDes meeting the specifications in this appendix when used with the recommended channel of Appendix A are expected to meet the host specifications at B 3.5.1 and C 3.5.2, however any implementation that meets those host specifications is a compliant SFP+ implementation, independent of whether the ASIC/SerDes and/or channel meet the specifications in Appendix A and this appendix. This allows flexibility between channel and SerDes performances and costs.

B.2 SFI ASIC/SERDES TRANSMITTER OUTPUT SPECIFICATIONS AT A (INFORMATIVE)

The driver is based on low voltage high speed driver logic with a nominal differential impedance of 100 Ω. The SFI transmitter electrical specifications at reference point A are given in Table 26. The source must provide both differential and common mode termination for quality signal termination and low EMI.

Pre-compensation such as de-emphasis may be required to mitigate data dependent jitter at compliance point B.

All parameters at A are measured with the ASIC/SerDes Test Board as shown in C.1.3.

Jitter specifications at A are not provided, the host transmitter in conjunction with the host SFP+ channel must deliver jitter specifications as given by reference point B, Table 12.

B.3 SFI ASIC/SERDES RECEIVER INPUT SPECIFICATIONS AT D (INFORMATIVE)

SFI ASIC/SerDes receiver electrical specifications are given in Table 27 and measured at reference point D. All specifications at D are measured with the SerDes on a ASIC/SerDes Test Board C.1.3. The nominal receiver input impedance is 100 Ω differential. The load must provide differential termination and avoid significant differential to common mode conversion for high quality signal termination and low EMI.

The necessary jitter performance at D is to be determined by the implementer based on the specifications at C.

Table 26 ASIC/SerDes Transmitter Output Electrical Specifications at A

Parameter - A	*Symbol*	*Conditions*	*Min*	*Typ*	*Max*	*Units*
Differential Output Voltage	V_{diff}				see 1	mV (p-p)
Termination Mismatch at 1 MHz	ΔZ_M	See D.16			5	%
Single Ended Output Voltage Range			-0.3		4.0	V
Output Rise and Fall time (20% to 80%)	Tr,Tf	See D.6	24			ps
Output AC Common Mode Voltage		See D.15			12	mV (RMS)
Differential Output S-parameter [2]	SDD22	0.01 to 2.8 GHz			-12	dB
		2.8 to 11.1 GHz			see 3	dB
Common Mode Output S-parameter [4]	SCC22	0.01 to 4.74 GHz			-9	dB
		4.74 to 11.1 GHz			see 5	dB

1. Host ASIC output must be set in combination of host channel to meet Y1 and Y2 levels of Table 12.
2. Reference differential impedance is 100 Ω
3. Differential Output S-parameter is given by equation SDD22(dB)= -8.15 + 13.33 $\log_{10}$(f/5.5), with f in GHz.
4. Reference common mode impedance is 25 Ω
5. Common mode output S-parameter is given by equation SCC22(dB)= -8.15 + 13.33 $\log_{10}$(f/5.5), with f in GHz.

Table 27 ASIC/SerDes Receiver Electrical Input Specifications at D

Parameter - D	*Symbol*	*Conditions*	*Min*	*Typ*	*Max*	*Units*
Differential Input Voltage Swing Supporting Limiting Module	Vdiff	See 1			850	mV (p-p)
Differential Input Voltage Modulation Amplitude Supporting Linear Module	VMA	See 1, 5, D.7			600	mV
AC Common Mode Voltage Tolerance		see D.15.3	15			mV (RMS)
Differential Input S-parameter [2]	SDD11	0.01 to 2.8 GHz			-12	dB
		2.8 to 11.1 GHz			see 3	dB
Differential to Common Mode Input Conversion S-parameter [4]	SCD11	0.01 to 11.1 GHz			-15	dB

1. Maximum value represents maximum input level to be tolerated by receiver.
2. Reference differential impedance is 100 Ω.
3. Response is given by equation SDD11(dB)= -8.15 + 13.33 $\log_{10}$(f/5.5), with f in GHz.
4. The test set common mode reference impedance is 25 Ω.
5. Peak levels may exceed VMA due to overshoot of the far end transmitter.

APPENDIX C APPLICATION REFERENCE BOARDS (NORMATIVE)

In order to provide test results that are reproducible and easily measured, this document defines 3 test boards that have SMA interfaces for easy connection to test equipment. One reference board is designed for testing the ASIC/SerDes, one reference board (HCB) for testing the host, and another reference board (MCB) for testing the module. All host, module and ASIC/SerDes specifications and recommendations in this document, and the specifications for the mated pair of compliance boards, are defined at the SMA interfaces. This appendix describes these test cards in detail. The reference test boards' objectives are:

- Satisfy the need for interoperablity at the electrical level.
- Allow for independent validation of ASIC/SerDes, host, and module.
- The PCB traces are targeted at 100 Ω differential impedance with nominal 7% differential coupling.

Testing compliance to specifications in a high-speed system is delicate and requires thorough consideration. Using common test boards that allow predictable, repeatable and consistent results among vendors will help to ensure consistency and true compliance in the testing.

C.1 COMPLIANCE BOARDS

The Host Compliance Board, the Module Compliance Board, and the ASIC test board are made with defined losses of PCB trace with specific high performance properties. Compliance boards are intended to ease building practical test boards with non-zero loss. SFI specifications incorporate the effect of non-zero loss reference test boards which improve the return loss and slightly slows down edges.

Measurements made using non-compliant test boards are invalid and no valid conclusions can be drawn from the results.

C.1.1 HOST COMPLIANCE BOARD TRANSFER CHARACTERISTICS

SDD21 is defined from the SFF-8083 mating pads, excluding these mating pads, to the mating interface of the SMA connector. The Host Compliance Board is a passive test board and SDD21 and SDD12 should be identical. The recommended response of the Host Compliance Board PCB excluding the SFF-8083 connector is given by.

$$SDD21(dB) = (-0.01 - 0.25 \times \sqrt{f} - 0.0916 \times f) \quad \text{from 0.01 to 15 GHz}$$

where f is the frequency in GHz. From 10 MHz to 11.1 GHz the discrepancy between the measured transfer response and the specified SDD21(dB) shall be <±15% of the transfer response in dB or ±0.1 dB, whichever is larger. For frequencies > 11.1 GHz and up to 15 GHz the discrepancy between measured transfer response and the specified SDD21(dB) shall be less than ±25% transfer response in dB.

The channel transfer characteristic is shown approximately in Figure 34.

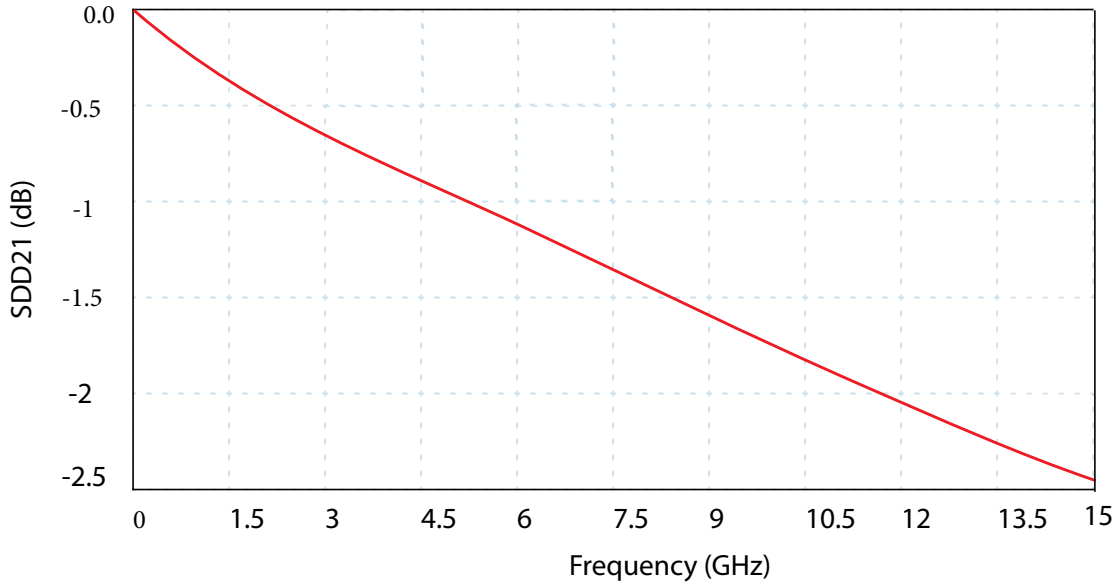


Figure 34 Approximate Response of Host Compliance Board

SFF-8083 connector response is defined by SFF-8083.

C.1.2 MODULE COMPLIANCE BOARD TRANSFER CHARACTERISTICS

SDD21 is defined from the SFF-8083 connector, excluding its solder pads, to the mating interface of the SMA connector. The Host Compliance Board is a passive test board and SDD21 and SDD12 should be identical. The recommended response of the Module Compliance Board PCB excluding the SFF-8083 connector is given by:

$$SDD21(dB) = (-0.00045 - 0.1135 \times \sqrt{f} - 0.04161 \times f) \quad \text{from 0.01 to 15 GHz}$$

where f is the frequency in GHz. Over the range of frequencies specified (10 MHz to 11.1 GHz) any discrepancy between measured transfer response and the specified SDD21(dB) shall be <±15% of the transfer response in dB or ±0.1 dB, whichever is larger. For frequencies > 11.1 GHz and up to 15 GHz the

discrepancy between measured transfer response and the specified SDD21(dB) shall be less than ±25% of the transfer response in dB.

The channel transfer response is shown approximately in Figure 35.

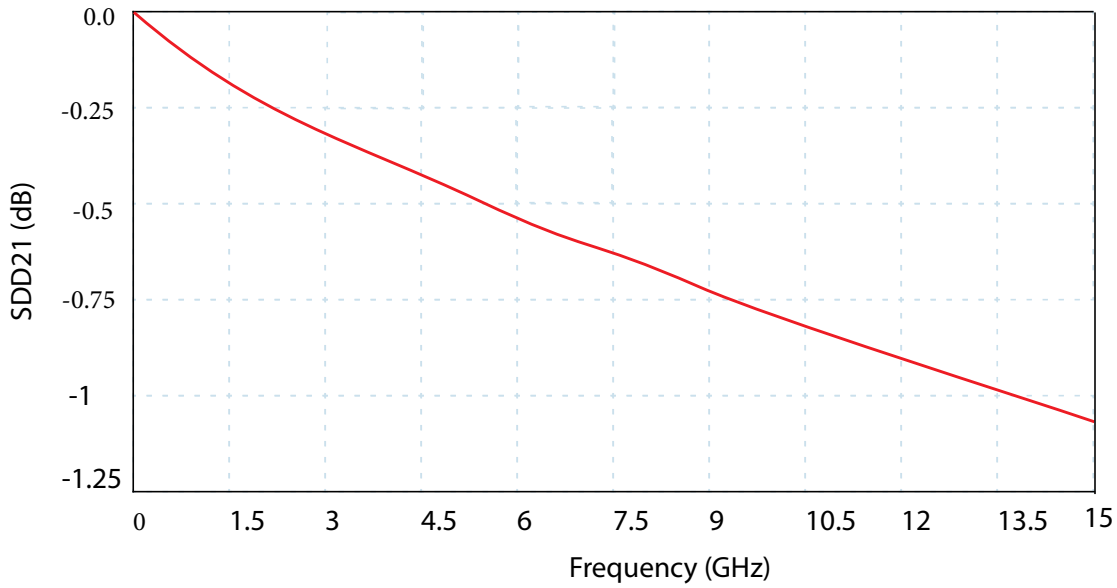


Figure 35 Approximate Transfer Response of Module Compliance Board

SFP+ connector response is defined by SFF-8083.

C.1.3 ASIC/SERDES TEST BOARD TRANSFER CHARACTERISTICS

The recommended response of the ASIC/SerDes test board PCB is the same as for the Module Compliance Board (see C.1.2).

C.2 HOST COMPLIANCE BOARD

The Host Compliance Board allows predictable, repeatable and consistent results among Host vendors and will help to ensure consistency and true compliance in the testing of Hosts. Host Compliance Boards are provided by Spirent Communication.

C.2.1 HOST COMPLIANCE BOARD MATERIAL AND LAYER STACK-UP

Host Compliance Board stack-up shown in Figure 36 is on six metal layers Rogers RO4350B© / FR4-6 material. The board is compliant with requirements of SFF-8432 and SFF-8083. SFI signals are routed on signal layer 1, low speed signals and controls are routed on signal layer 6.

1. Top Layer	Signal	17 µm/0.5 oz Copper plated to 1 oz min+ 1.25 µm Nickel + 2.5 µm Gold
0.168 mm / 6.6 mils Rogers RO4350B		
2. Layer	Vee	34 µm/1 oz Copper
0.14 mm / 5.5 mils FR4-6		
3. Layer	Signal 1	17 µm/0.5 oz Copper
0.178 mm / 7 mils FR4-6		
4. Layer	Signal 2	17 µm /0.5 oz Copper
0.14 mm / 5.5 mils FR4-6		
5. Layer	Power	34 µm/1 oz Copper
0.168 mm / 6.6 mils Rogers RO4350B		
6. Bottom Layer	Signal	17 µm/0.5 oz Copper plated to 1 oz min + 1.25 µm Nickel + 0.25 µm Gold

Figure 36 Host Compliance Board stack-up

C.2.2 HOST COMPLIANCE BOARD PARTLIST

The Host Compliance Board part list is given below.

Table 28 Host Compliance Board Part List

Qty	*RefDes*	*Value*	*Description*	*Example Part Number*
2	C5, C6	0.1 uF	Ceramic Capacitor	10% X7R 10V 0402 SMT LFR
3	D1, D2, D3	GREEN	LED Single Green	120 DEG 0603 SMT LFR
2	D4, D5	Blue	LED Single Blue	120 DEG 0603 SMT LFR
1	J1	Conn3	Connector Header 3 Pins Straight	Tyco PN#3-644695-3
4	J2, J3, J4, J5	EDGE SMA	SMA Connector Jack R/A	Rosenberger PN# 32K243-40ME3
1	J6	CONN1X3P	Connector Header 3 Pins 100 mil Pitch	Molex PN# 22-23-203
5	R1, R2, R3, R4, R5	1.0 kΩ	Resistor	RES 1.00K 1% 1/10W 0603 SMT LFR
1	SW1	SPST	SW 4 Position Dip Switch SMT	ITT Cannon PN# TDA04H0SB1

Note: Table 28 does not use all in-sequence part numbers.

C.2.3 HCB GERBER FILES

The Gerber file for the Host Compliance Board is available in SFF-8434.

C.2.4 SCHEMATIC OF HOST COMPLIANCE BOARD

The schematic of Host Compliance Board is shown in Figure 37.

Mod-DEF0 in the schematic is Mod_ABS as defined by Table 3 and AS0/AS1 in the schematic are RS0/RS1 as defined by Table 3.

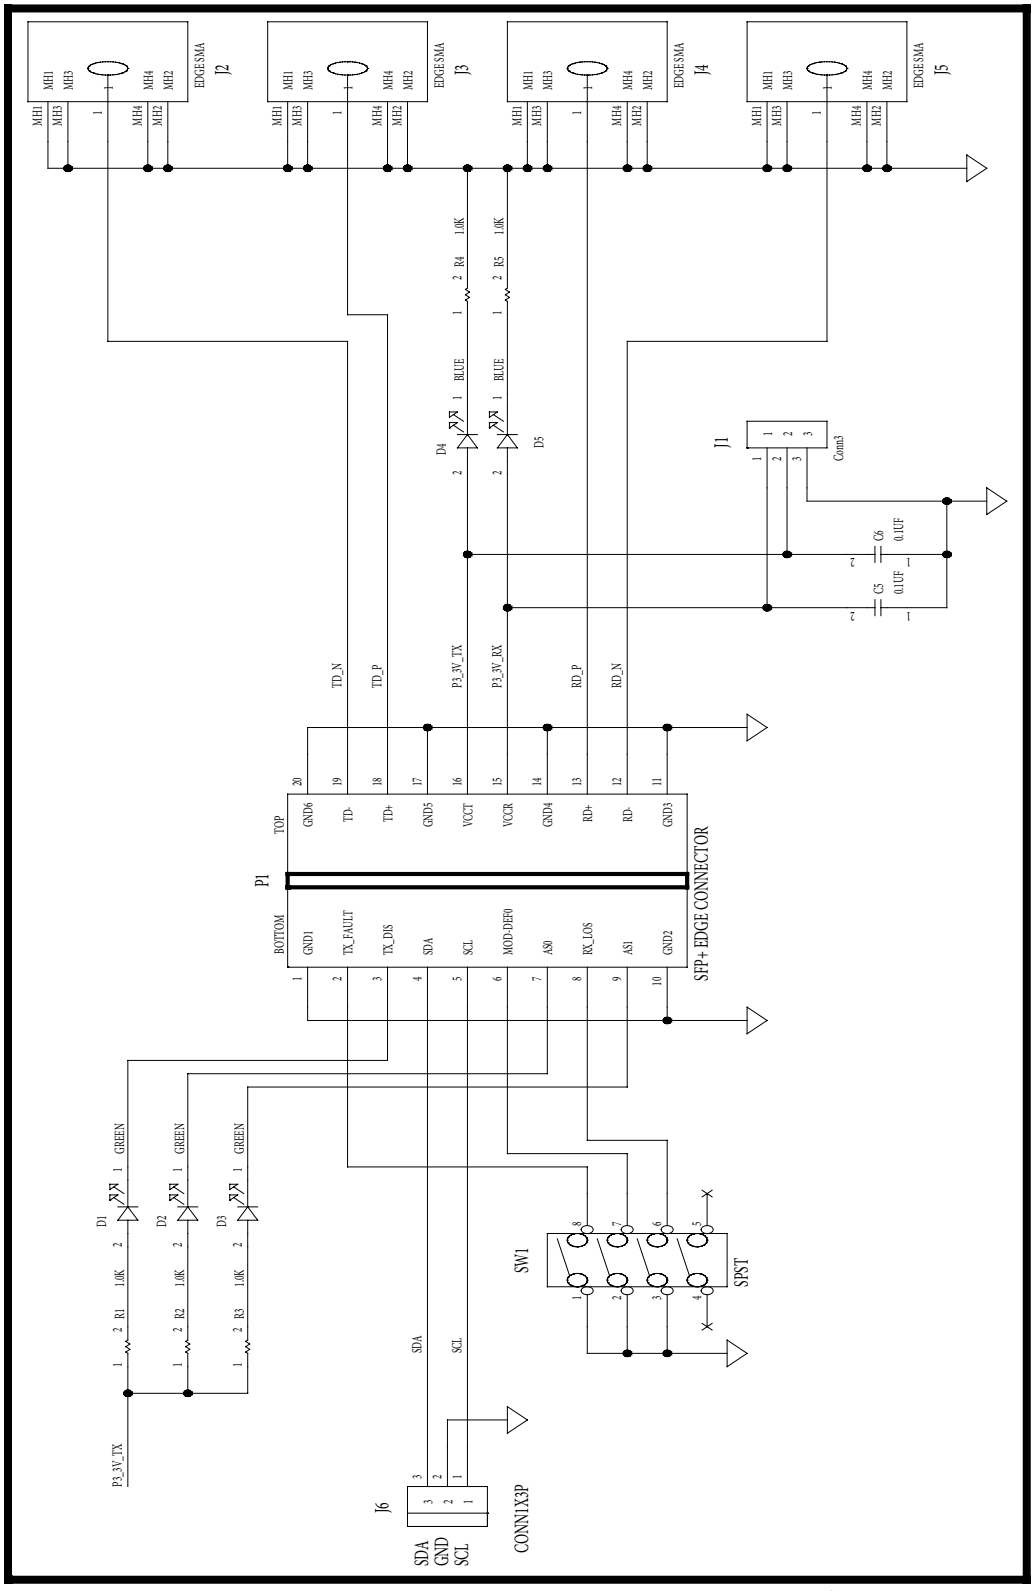


Figure 37 Schematic of the Host Compliance Board

C.3 MODULE COMPLIANCE BOARD

The Module Compliance Board allows predictable, repeatable and consistent results among module vendors and will help to ensure consistency and true compliance in the testing of modules. Module Compliance Boards are provided by Broadcom Corporation.

The solder pads for the high speed traces in the Module Compliance Board are 1.1x0.4 mm to improve high frequency performance instead of 2.0x0.5 mm as defined in the SFF-8083 for improved manufacturability. For detailed geometry, see the Gerber files in C.3.4.

C.3.1 MODULE COMPLIANCE BOARD MATERIAL AND LAYER STACK-UP

Module Compliance Board stack-up shown in Figure 38 is based on a laminate of Rogers RO4350B/ FR4-6 with ten metal layers. SFI signals are routed on signal layer 1, low speed signals and controls are routed on signal layers 8 and 10.

1. Top Layer	Signal	17 µm/0.5 oz Copper + 1.25 µm Nickel + 2.5 µm Gold
0.168 mm / 6.6 mils Rogers RO4350B		
2. Layer	Vee	17 µm /0.5 oz Copper
0.382 mm / 15 mils FR4-6		
3. Layer	Vee	34 µm/1 oz Copper
0.076 mm / 3 mils FR4-6		
4. Layer	VccR	34 µm/1 oz Copper
0.076 mm / 3 mils FR4-6		
5. Layer	Vee	34 µm /1 oz Copper
0.076 mm / 3 mils FR4-6		
6. Layer	VccT	34 µm /1 oz Copper
0.076 mm / 3 mils FR4-6		
7. Layer	Vee	34 µm/1 oz Copper
0.076 mm / 3 mils FR4-6		
8. Layer	Signal	34 µm/1 oz Copper
0.382 mm / 15 mils FR4-6		
9. Layer	Vee	17 µm/ 0.5 oz Copper
0.168 mm / 6.6 mils Rogers RO 4350B		
10. Bottom Layer	Signal	17 µm Cu / 0.5 oz Copper+ 1.25 µm Nickel + 0.25 µm Gold

Figure 38 Module Compliance Board stack up

C.3.2 SCHEMATIC OF MODULE COMPLIANCE BOARD

Schematic of Module Compliance Board is shown in Figure 39.

Mod-DEF0 in the schematic is Mod_ABS as defined by Table 3 and AS0/AS1 in the schematic are RS0/RS1 as defined by Table 3.

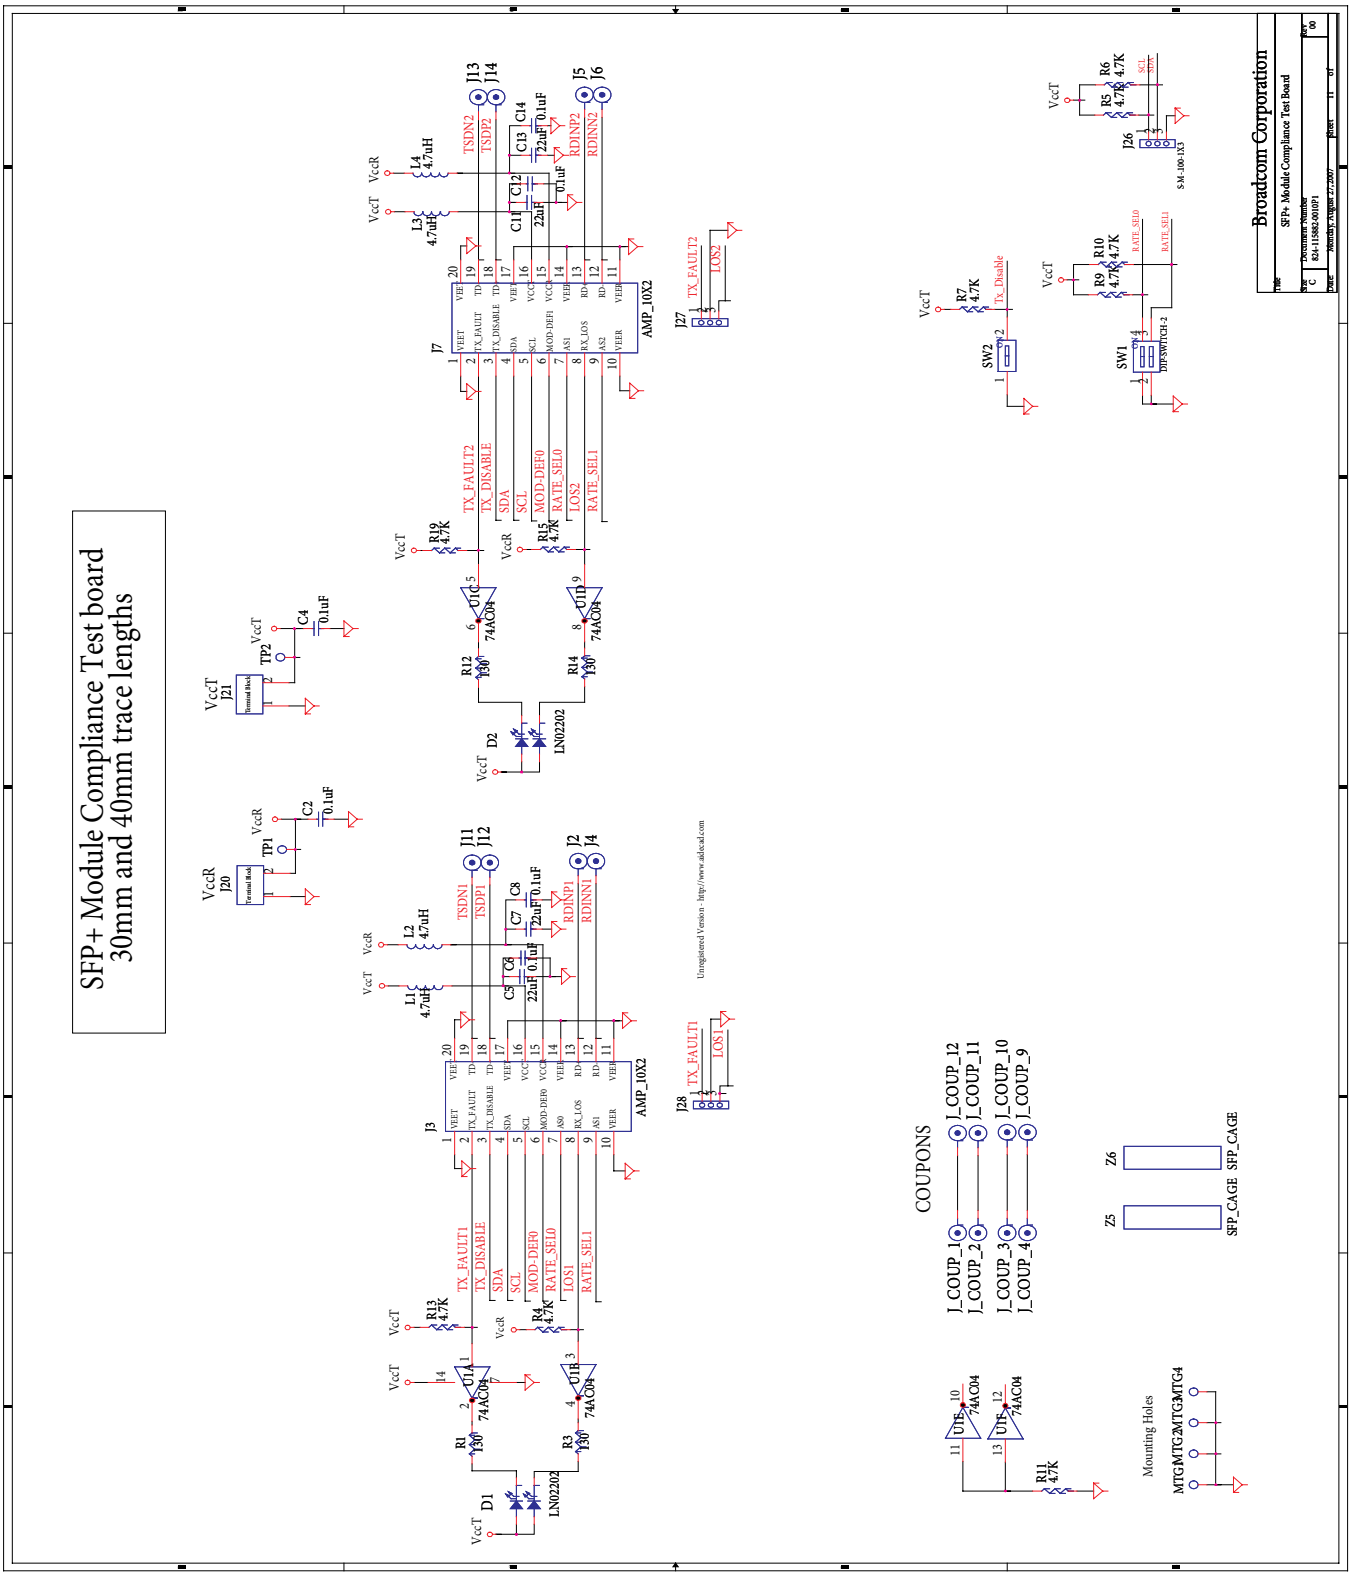


Figure 39 Schematic of The Module Compliance Board

C.3.3 MODULE COMPLIANCE BOARD PARTLIST

Component part list for the Module Compliance Board is given below.

Table 29 Module Compliance Board Part List

Qty	*RefDes*	*Value*	*Description*	*Example Part Numbers*
6	C2, C4, C6, C8, C12, C14	0.1uF	Ceramic Capacitors	Murata/GRM188R71C104MA01D
4	C5, C7, C11, C13	22 uF	Ceramic Capacitors	Murata/GRM21BR60J226ME39K
4	D1, D2, D4, D5	RED	LED	Panasonic/LNJ208R8ARA
12	J_COUP_2, J2, J_COUP_4, J4, J5,J6, J_COUP_9, J_COUP_11, J12, J14J_COUP_1, J_COUP_3, J_COUP_10, J11, J_COUP_12, J13	SMA	SMA Connector R/A	Huber&Suhner/92_SK-U50-0-3/199_NE
2	J3, J7	Con_10x2	SFF-8083 Connector	Tyco 1888247 or Molex 74441
2	J20, J21	Terminal Block	Terminal Bloc	On-Shore-Tech/EDZ5002DS
3	J26, J27, J28	S-M-.100-1X3	PCB Header	Molex/22-10-2031
4	L1, L2, L3, L4	4.7 uH	Inductor	Toko/A914BYW-4R7M
4	R1, R3, R12, R14	130 Ω	Resistors	Walsin/WR06X131JTL
10	R4, R5, R6, R7, R9, R10, R11, R13, R15, R19	4.7 kΩ	Resistors	Walsin/WR06X472JTL
1	SW1	DIP-SWITCH-2	DipSwitch	CT2062-ND
1	SW2	sw_pb_ck-k	Toggle Switch	C&K/ET01MD1AVBE
1	U1	74AC04	Inverter	Fairchild/530438-00
2	Z5, Z6	SFP_CAGE	SFP Cage	Tyco 1489962-1

Note: Table 29 does not use all in-sequence part numbers.

C.3.4 MCB GERBER FILES

The Gerber file for the Module Compliance Board is available in SFF-8434.

C.4 SPECIFICATIONS FOR MATED HOST AND MODULE COMPLIANCE BOARDS

Based on measurements of the Module Compliance Board (MCB) mated with the Host Compliance Board (HCB) the following specifications have been derived for the mated pair. Compliance to these limits help ensure the module and host specifications can be met.

S-parameters are defined based on two ports mixed mode differential definition [see INF-8077i Appendix C], see Figure 40. All single port measurements are listed on the figure.

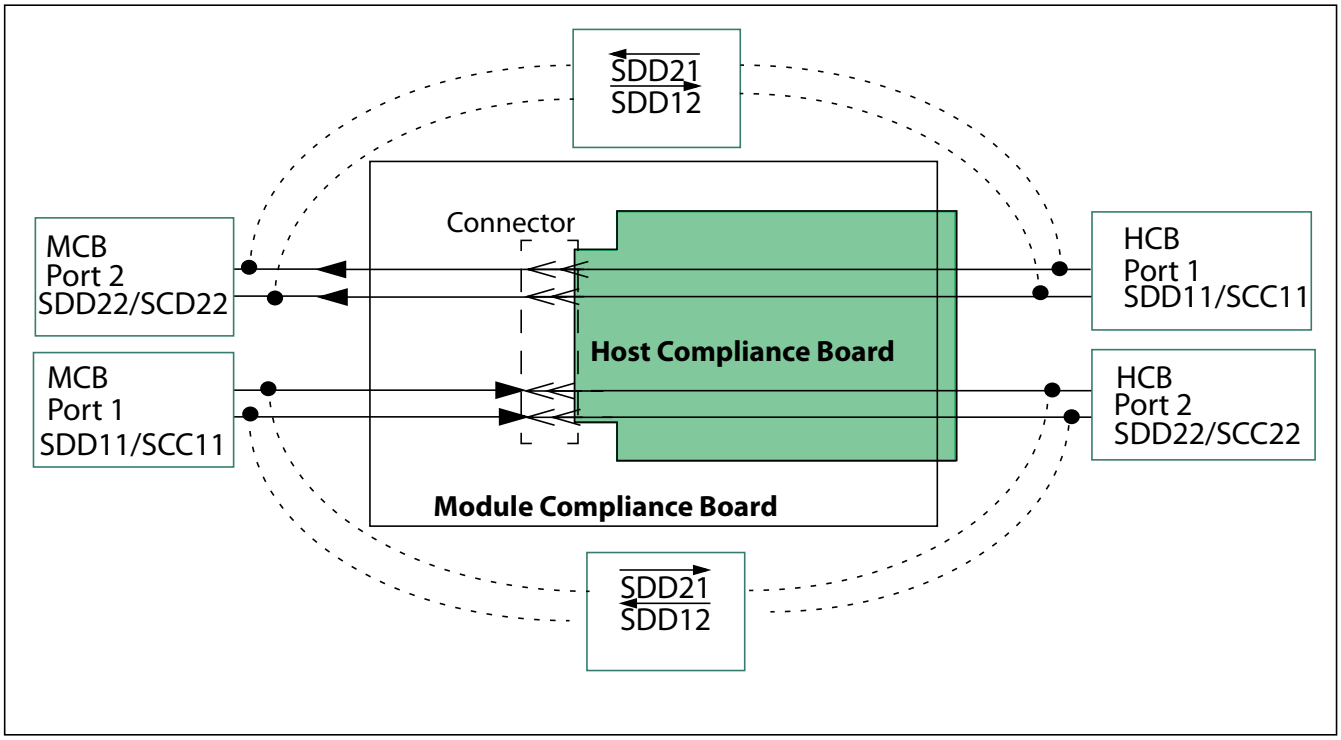


Figure 40 Measurements Port Definition

The maximum values of SDD11 or SDD22 looking into the Module Compliance Board and Host Compliance Board are illustrated in Figure 41.

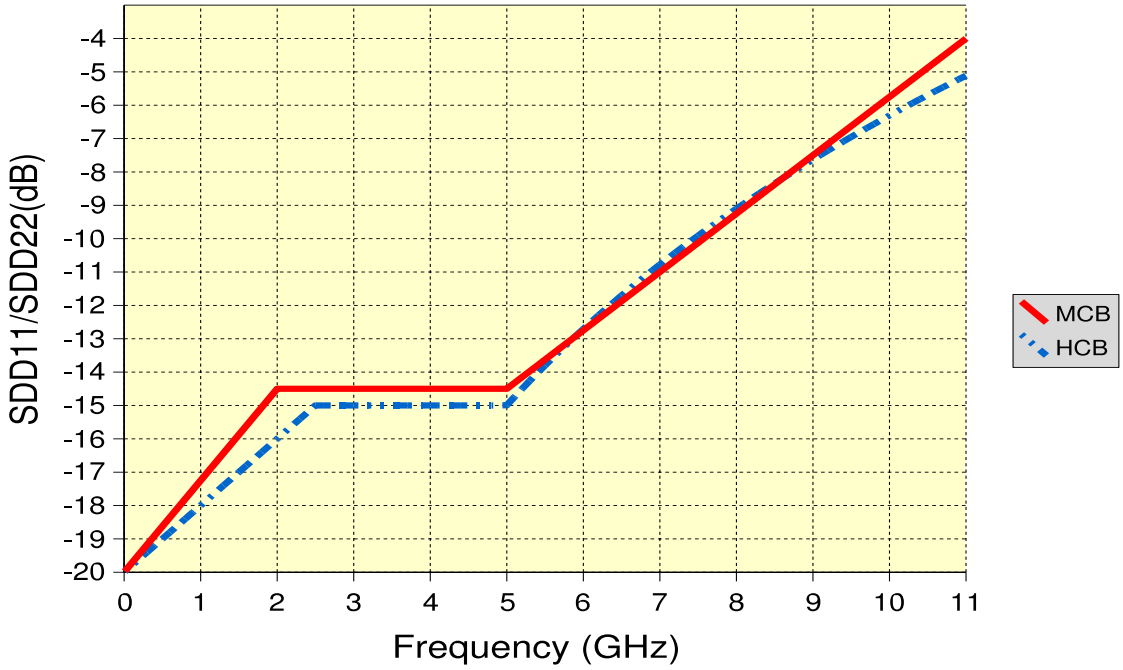


Figure 41 Maximum Differential Response of mated MCB and HCB

The maximum values of SDD11 or SDD22 looking into the Module Compliance Board are given by the following equations:

$$SDDxx(dB) \leq -20 + 2.75 \times f \qquad \text{f in GHz from 0.01 to 2}$$

$$SDDxx(dB) \leq -14.5 \qquad \text{f in GHz from 2 to 5}$$

$$SDDxx(dB) \leq -23.25 + 8.75 \times \left(\frac{f}{5}\right) \qquad \text{f in GHz from 5 to 11.1}$$

The maximum values of SDD11 or SDD22 looking into the Host Compliance Board are given by the following equations:

$$SDDxx(dB) \leq -20 + 2 \times f \qquad \text{f in GHz from 0.01 to 2.5}$$

$$SDDxx(dB) \leq -15 \qquad \text{f in GHz from 2.5 to 5}$$

$$SDDxx(dB) \leq -13.8 + 28.85 \times \log_{10}\left(\frac{f}{5.5}\right) \qquad \text{f in GHz from 5 to 11.1}$$

The maximum and the minimum values of SDD21 or SDD12 looking into either the Module Compliance Board or Host Compliance Board are illustrated in Figure 42 and given by the equations below.

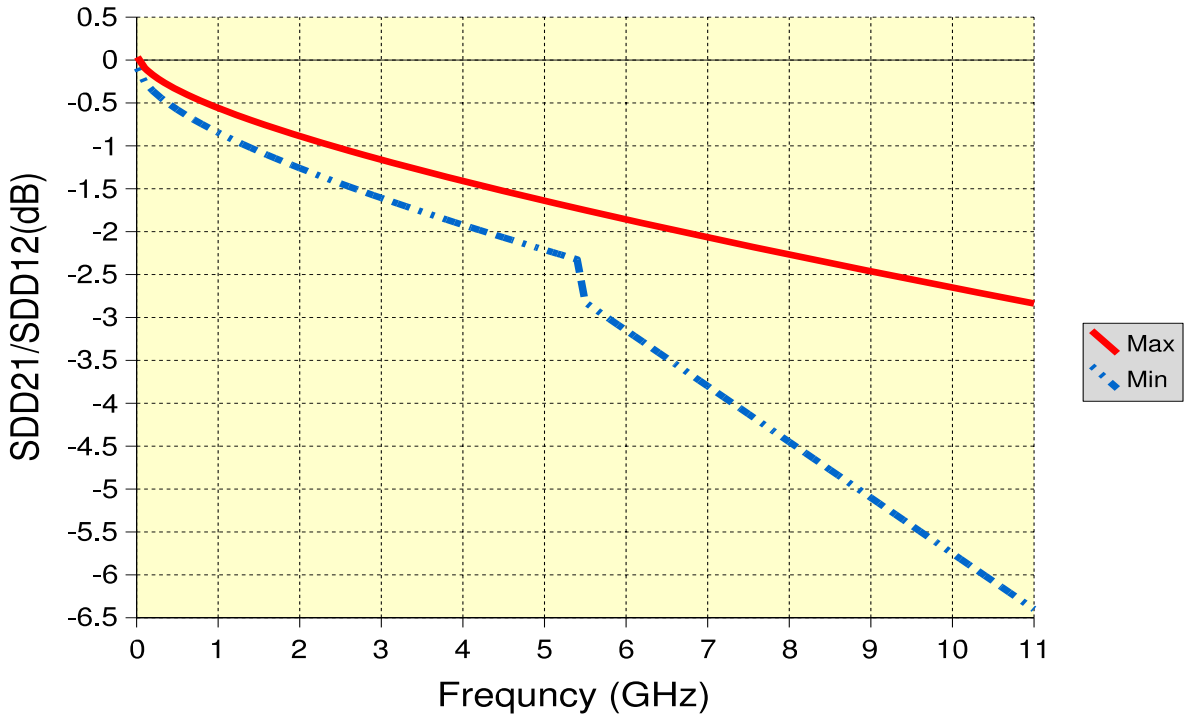


Figure 42 Mated MCB-HCB Differential Through Response Limits

$$SDDxx(dB) \leq (-0.012 - 0.694 \times \sqrt{f} - 0.127 \times f) \qquad \text{f in GHz from 0.01 to 5.5}$$

$$SDDxx(dB) \leq 0.75 - 0.65 \times f \qquad \text{f in GHz from 5.5 to 11.1}$$

$$SDDxx(dB) \leq 0.0915 - 0.549 \times \sqrt{f} - 0.101 \times f \qquad \text{f in GHz from 0.01 to 11.1}$$

The maximum values of SCC11 or SCC22 looking into either the Module Compliance Board or Host Compliance Board are illustrated in Figure 43.

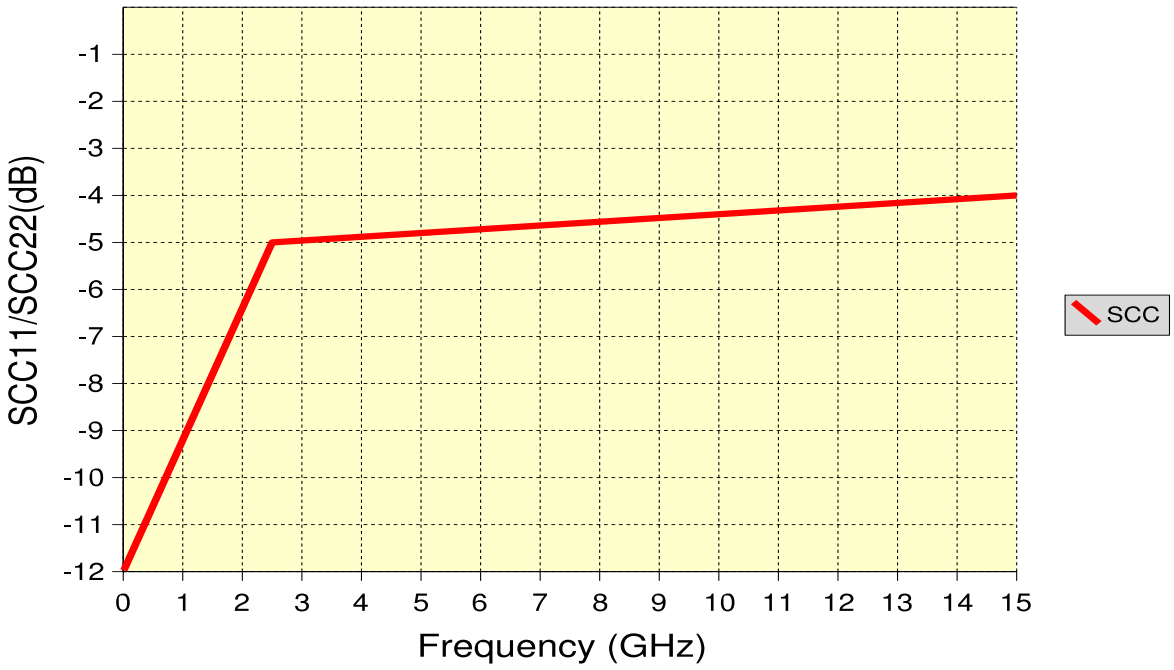


Figure 43 Maximum Common Mode Response of mated MCB and HCB

The SCC11 and SCC22 are also given by the following equations:

$$SCCxx(dB) \leq -12 + 2.8 \times f \qquad \text{f in GHz from 0.01 to 2.5}$$

$$SCCxx(dB) \leq -5.2 + 0.08 \times f \qquad \text{f in GHz from 2.5 to 15.}$$

The maximum values of Differential to Common Mode Response SCD21 and SCD12 looking into either the Module Compliance Board or Host Compliance Board are illustrated in Figure 44.

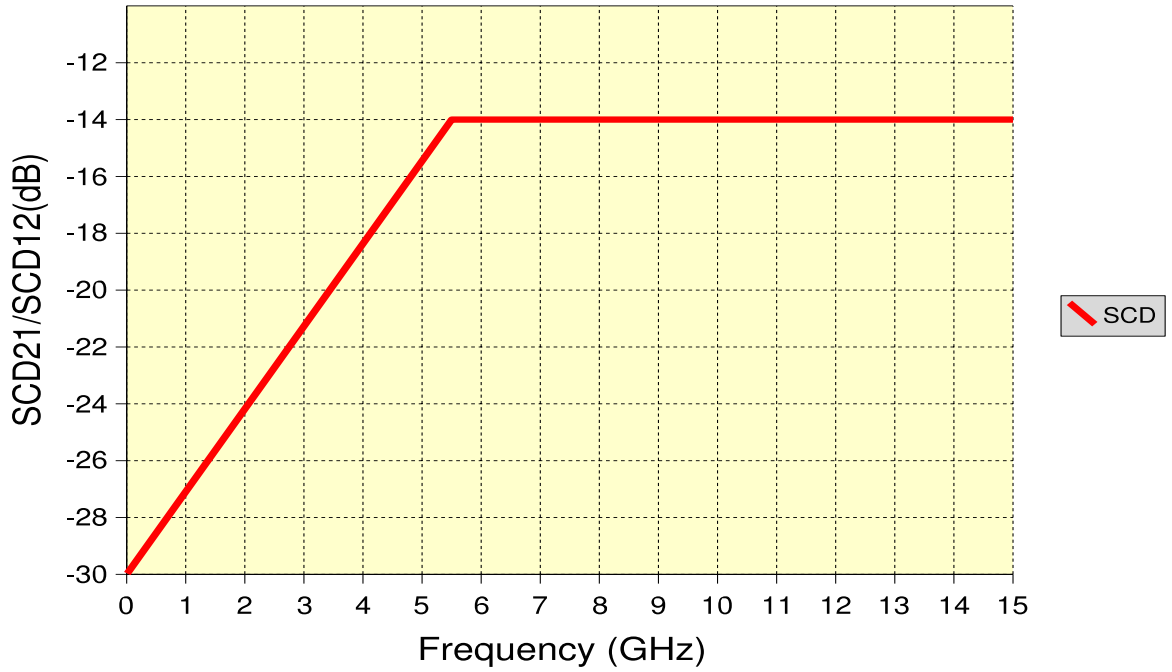


Figure 44 Maximum Differential to Common Mode Response of mated MCB and HCB

Mated response SCD21 and SCD12 of the mated Module and Host Compliance Board are given by:

$$SCDxx(dB) \leq -30 + 2.91 \times f \qquad \text{f in GHz from 0.01 to 5.5}$$

$$SCDxx(dB) \leq -14 \qquad \text{f in GHz from 5.5 to 15.}$$

When MCB Port 1 of the mated Module Compliance Board and Host Compliance Board is excited by the crosstalk source defined in Table 12, the RMS differential NEXT voltage at MCB Port 2 shall be less than 1 mV when measured in a 12 GHz bandwidth. The differential NEXT voltage from HCB Port 1 to HCB Port 2 is expected to be about the same. Compliance boards meeting this response are expected to pass the integrated NEXT requirement, however it is also expected that the curve can be exceeded due to frequency resonances while still passing the integrated NEXT requirement. The frequency domain curve shown in Figure 45 shows a recommended response which is described in the equations following the figure. Compliance boards meeting this NEXT response are expected to pass the NEXT voltage requirement, however it is

also expected that the curve can be exceeded due to frequency resonances while still passing the NEXT voltage requirement.

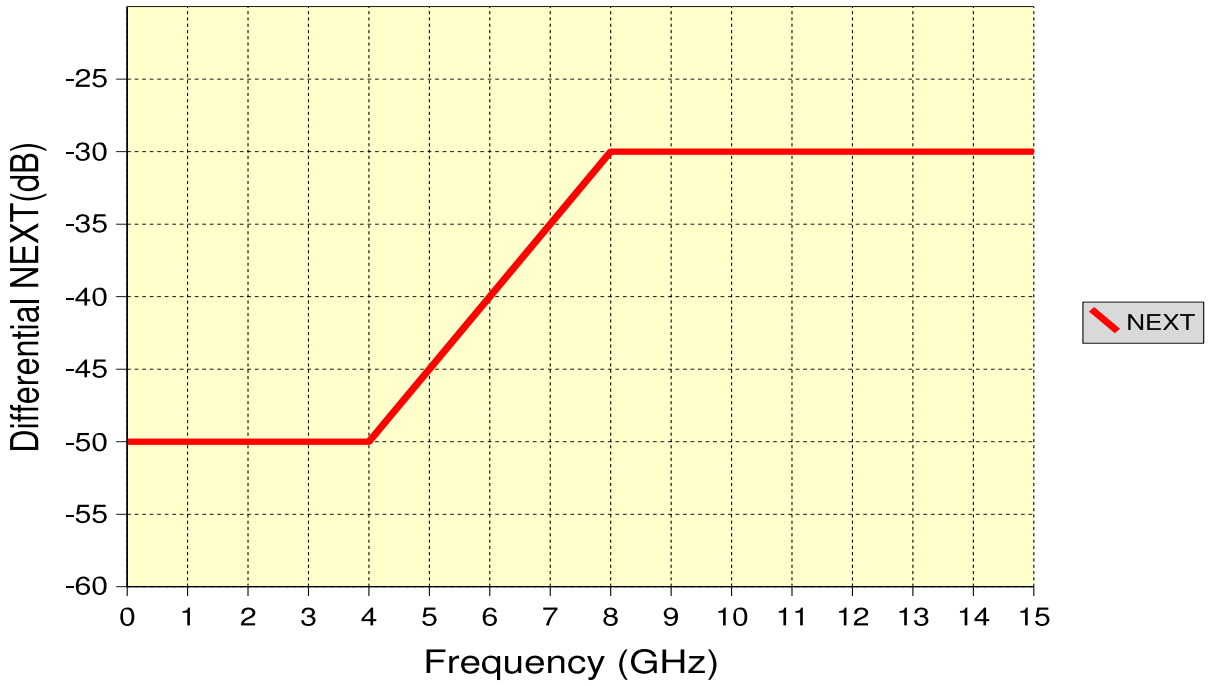


Figure 45 Maximum Differential NEXT Response of mated MCB and HCB

The recommended NEXT response is also given by the following equations:

$$NEXT(dB) \leq -50 \qquad \text{f in GHz from 0.01 to 4}$$

$$NEXT(dB) \leq -70 + 5 \times f \qquad \text{f in GHz from 4 to 8}$$

$$NEXT(dB) \leq -30 \qquad \text{f in GHz from 8 to 15.}$$

APPENDIX D TEST METHODOLOGY AND MEASUREMENT (NORMATIVE)

D.1 INTRODUCTION

This appendix defines metrics for SFP+ high-speed and power electrical interfaces and provides practical guidance for test implementation. Each parameter is defined in terms of a measurement procedure. The instruments for measurement are assumed to be ideal: accurate, precise, with infinite or defined bandwidth, zero or defined noise and so on. In practice, the necessary level of instrument performance and the approach to calibration and margining must be considered. Some guidance is given in the following sections.

All measurements are made differentially, with the exception of AC Common Mode Generation Test D.15.2, Common Mode Tolerance Test D.15.3, Termination Mismatch D.16, Module Power Supply Tolerance Filtering D.17, and Power Supply Noise Testing Methodology D.17.

Accurate calibration of test equipment is assumed for all measurements. To avoid pessimistic WDP and jitter results, the scope may require correction for time base linearity errors.

D.1.1 TEST PATTERNS

Test patterns used in this specification include the 8+8 square wave, PRBS9, IEEE 802.3 test patterns1, 2 and 3, and any valid 64B/66B signal. PRBS9 is defined in IEEE Std 802.3, 68.6.1 and a file for the sequence can be found at http://ieee802.org/3/aq/public/tools/TWDP/prbs9_950.txt. Test patterns 1, 2 and 3 are defined in IEEE Std 802.3, 52.9.1.1. Test pattern 3 is PRBS31 as defined by ITU-T or in IEEE Std 802.3, 49.2.8.

D.2 EYE MASK COMPLIANCE

This section defines what is meant by eye mask compliance and gives guidance for its determination. Mask templates and coordinates are given in subclauses in 3.5 SFP+ Host System Specifications and 3.6 SFP+ Module Specifications .

- The pattern(s) for eye mask testing is according to the relevant standard(s) listed in Table 1.

- The output being tested should comply over the range of operating conditions while the opposing direction bit stream, operates with the target crosstalk rise and fall and amplitude given in Table 12, Table 14, Table 17, and Table 18. The opposing direction bit stream (than the one being tested) shall be asynchronous PRBS31 or valid 64B/66B bit stream.

Testing may include guard banding, extrapolation, or other methods, but must ensure that mask violations do not occur at a rate exceeding the hit ratio limit given in the appropriate table.

- An AC coupling 3 dB corner frequency of 20 kHz is expected to be adequate to eliminate baseline wander effects, however high frequency performance is critical and must not be sacrificed by the AC coupling.
- All loads are specified at 100 Ω differential.
- 0.0 UI and 1.0 UI on the time axis are defined by the eye crossing means at the average value (zero volts if AC coupled) of the signal. The average value might not be at the jitter waist.

A clock recovery unit (CRU) is used to trigger the scope for mask measurements as shown in Figure 46. The reference CRU has a high frequency corner bandwidth of 4 MHz and a slope of -20 dB/decade with peaking of 0.1 dB or less.

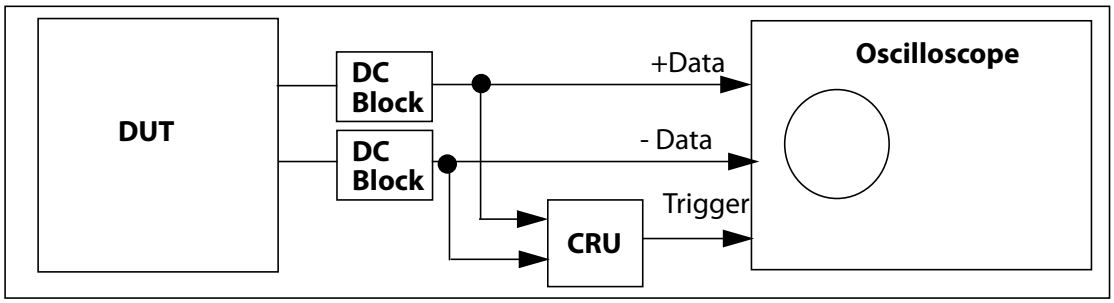


Figure 46 Eye mask measurement setup - block diagram.

D.2.1 EXAMPLE CALCULATIONS FOR 5X10^{-5} HIT RATIO

If an oscilloscope records 1350 samples/screen, and the time-base is set to 0.2 UI per division with 10 divisions across the screen, and the measurement is continued for 200 waveforms, then a transmitter with repeated measurement averaging to less than 6.75 hits is compliant. i.e.,

$$HitRatio = \frac{5 \times 10^{-5} \times 1350}{0.2 \times 10} = 6.75$$

Likewise, if a measurement is continued for 1000 waveforms, then repeated measurement averaging to less than 33.75 hits is compliant. An extended measurement is expected to give a more repeatable result, whereas a single reading of 6 hits in 200 waveforms would not give a statistically significant pass or fail.

D.3 DATA DEPENDENT JITTER (DDJ) AND PULSE WIDTH SHRINKAGE (DDPWS)

A high-resolution oscilloscope, time interval analyzer, or other instrument with equivalent capability may be used to measure DDJ and DDPWS. A repeating PRBS9 pseudo-random test pattern, 511 bits long, is used. For electrical jitter measurements, the measurement bandwidth is 12 GHz. If the measurement bandwidth affects the result, it can be corrected for by post-processing. However, a bandwidth above 12 GHz is expected to have little effect on the results.

DCD and Pulse Width Shrinkage (DDPWS) are components of DDJ.

Establish a crossing level equal to the average value of the entire waveform being measured. Synchronize the instrument to the pattern repetition frequency and average the waveforms or the crossing times sufficiently to remove the effects of random jitter and noise in the system. The PRBS9 pattern has 128 positive-going transitions and 128 negative-going transitions. The mean time of each crossing is then compared to the expected time of the crossing, and a set of 256 timing variations is determined. DDJ is the range (max-min) of the timing variations. Keep track of the signs (early/late) of the variations. Note, it may be convenient to align the expected time of one of the crossings with the measured mean crossing.

The following Figure 47 illustrates the method. The vertical axis is in arbitrary units, and the horizontal axis is plotted in UI. The waveform is AC coupled to an average value of 0, therefore 0 is the appropriate crossing level. The rectangular waveform shows the ideal crossing times, and the other is the waveform with jitter that is being measured. Only 32 UI are shown (out of 511). The waveforms have been arbitrarily aligned with ($\Delta t_2 = 0$) at 14 UI.

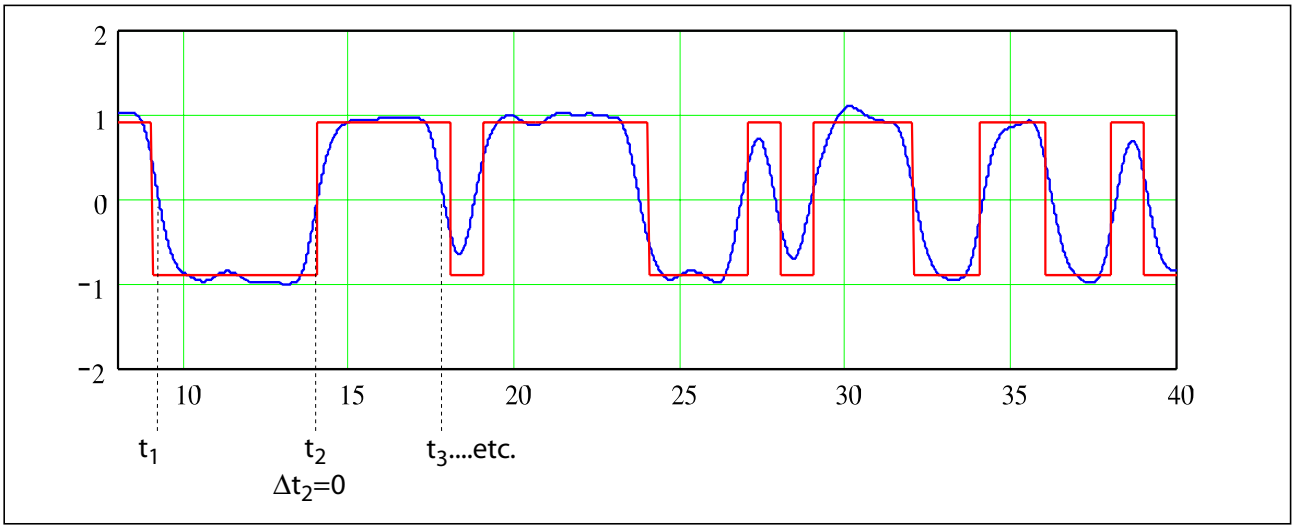


Figure 47 DDJ Test Method

DDJ is defined as

$$DDJ = max(\Delta t_1, \Delta t_2, \ldots \Delta t_n) - min(\Delta t_1, \Delta t_2, \ldots \Delta t_n)$$

Every edge, 1...n, in a complete repetition of the pattern is measured (n = 256 in a PRBS9 pattern).

DDPWS is determined as the difference between one symbol period and the minimum of all the differences between pairs of adjacent edges

DDPWS = T - min(t_2-t_1, t_3-t_2,..... t_{n+1}-t_n)

where T is one symbol period. Note that the difference from the next edge in the repeating sequence, t_{n+1}, is also considered.

D.3.1 DUTY CYCLE DISTORTION (DCD)

DCD represents a deviation from the intended duty cycle. It is the difference between the mean position of all falling edges and the mean position of all rising edges with uncorrelated effects minimized through averaging. DCD is measured at the average value of the waveform.

D.4 UNCORRELATED JITTER (UJ)

UJ as defined by IEEE 802.3 CL 68 is a measure of any jitter that is un-correlated to the 64B/66B bit stream. The definition and test procedure for UJ are identical to those defined in IEEE 802.3 CL 68.6.8 with following considerations:

- The host transmitter shall comply while the host receiver is operating with asynchronous PRBS31 or valid 64B/66B signal and all other ports operating as in normal operation, including proper termination.
- The receive path input of the Host Compliance Board is connected to a pattern generator and calibrated through a Module Compliance Board. The amplitude and rise time are set to the target values stated in Table 12 at C".
- For the purposes of this document the procedures defined for optical testing also apply to electrical testing. Optical terms (such as power) and units, such as in Figure 68-9 in IEEE 802.3, can be converted to corresponding electrical terms (such as voltage) and units, etc.

- The 4th-order Bessel-Thomson response is to be used only for optical measurements of UJ. UJ in the electrical domain is defined in a bandwidth of 12 GHz, unless specified by the application standard.
- PRBS9 is suitable as a test sequence for all applications unless specified otherwise.
- The bandwidth of the CRU is defined in IEEE 802.3 clause 68.6.8 or in the relevant standard for the application.

D.5 99% Jitter (J2) and Total Jitter (TJ)

Jitter is a property of the timing of a signal's edges. The time of occurrence of an edge is defined as when the signal crosses its average level (e.g., 0 V for A.C. coupled, ground terminated measurements). Jitter is defined using the CRU of section D.2. The test pattern for Total Jitter (TJ) and 99% Jitter (J2) testing shall be either PRBS31 or a valid 64B/66B signal. These metrics of jitter are measured without averaging.

J2 is the same as J, "all but 1% for jitter", used in IEEE 802.3 Clause 52.9.9. It is defined as the time interval that includes all but 10^{-2} of the jitter distribution. If measured using an oscilloscope, it is the time interval from the 0.5th to the 99.5th percentile of the jitter distribution measured on the histogram.

TJ, as used in this document, is the Level 1 definition for TJ as described in the FC-MJSQ, where TJ is the crossing width, defined as the late time at which the BER is 10^{-12} minus the early time at which the BER is 10^{-12}. This is one unit interval (UI), minus the "jitter eye opening" defined in FC-MJSQ. TJ can be expressed as:

- TJ=T-t1

Where t1 is the jitter eye opening at the CDF = 10^{-12}, and T is one symbol period.

The CDF is a cumulative distribution function of the timings of the edges with a maximum close to 0.5 because the transition density is close to 50%.

A measurement using the BERT bathtub method must be corrected for the instrument's setup-and-hold time and noise. As PRBS31 is more demanding than a 64B/66B signal, a 10GBASE-R instance whose TJ is compliant using a 64B/66B signal is considered compliant even if it does not meet the required limit using PRBS31. A 10GBASE-W instance shall be compliant with PRBS31. It is not expected that the J2 value will differ between these patterns.

Both J2 and TJ are measured from side to side of the CDF, not from median to side of the CDF.

D.6 Rise and Fall Times

In this document, rise and fall times are defined as the time between the 20% and 80% times, or 80% and 20% times, respectively, of isolated edges. The normative test pattern is the OMA test pattern (eight ones, eight zeros). The 0% level and the 100% level are as defined by the xMA measurement procedure (see D.7 and IEEE Std 802.3, 68.6.2).

Alternatively, suitable edges exist in the PRBS9, within sequences of five zeros and four ones, and nine ones and five zeros, respectively. These are bits 10 to 18 and 1 to 14, respectively. In this case, the 0% level and the 100% level may be estimated as ZeroLevel and ZeroLevel + MeasuredxMA in the xWDP code (see Appendix G), or by the average signal within windows from -3 to -2 UI and from 2 to 3 UI relative to the edge. The PRBS9 methods are inaccurate for rise and fall times above 1.5 UI.

For electrical signals, the waveform is observed through a 12 GHz low pass filter response. For optical signals, the rise and fall times may be defined either without a filter response or through the standard 7.5 GHz Bessel-Thomson response; one or the other option is specified in each case.

NOTE -- The rise and fall definition in this document is not the same as the rise and fall times typically reported by an oscilloscope from an eye diagram derived from a mixed frequency signal such as PRBS or a 64B/66B signal, which takes all the edges into account.

D.7 Voltage Modulation Amplitude (VMA)

VMA is the difference between the nominal one and zero levels of an electrical signal. It is analogous to the OMA of an optical signal (see IEEE Std 802.3 52.9.5 and 68.6.2). VMA is defined with the square wave test pattern of eight ones and eight zeros defined in IEEE Std 802.3, 68.6.1 (this is a subset of the square waves allowed in IEEE 802.3, 52.9.5), or in the case of a non-802.3 application, a test pattern defined by the relevant standard.

It can be measured as follows:

- The signal under test is set to carry the square wave pattern and is observed, typically with an oscilloscope triggered to the pattern. The bandwidth of this measurement system is at least 3/T, where T is the period between transitions. For the square wave test pattern (0000000011111111) this gives approximately 4 GHz at 10.3125 GBd; the 12 GHz bandwidth defined for other electrical quantities is convenient. Electrical measurements of VMA do not require a 7.5 GHz Bessel-Thomson filter.

- The square wave being measured is divided into two equal time intervals, 8 UI long, aligned to the average time of both edges.
- The time of occurrence of an edge is defined as when the square-wave signal crosses its average level (0 V for A.C. coupled measurements).
- The average voltage level in the central 20% of each time interval is measured.
- The difference between the two levels (a positive voltage) is the VMA.
- An estimate of the OMA or VMA of a PRBS9 waveform is provided by the variable MeasuredxMA calculated by the algorithm in Appendix G.

An example square wave signal with eight zeros and eight ones with the two measurement windows is shown in Figure 48.

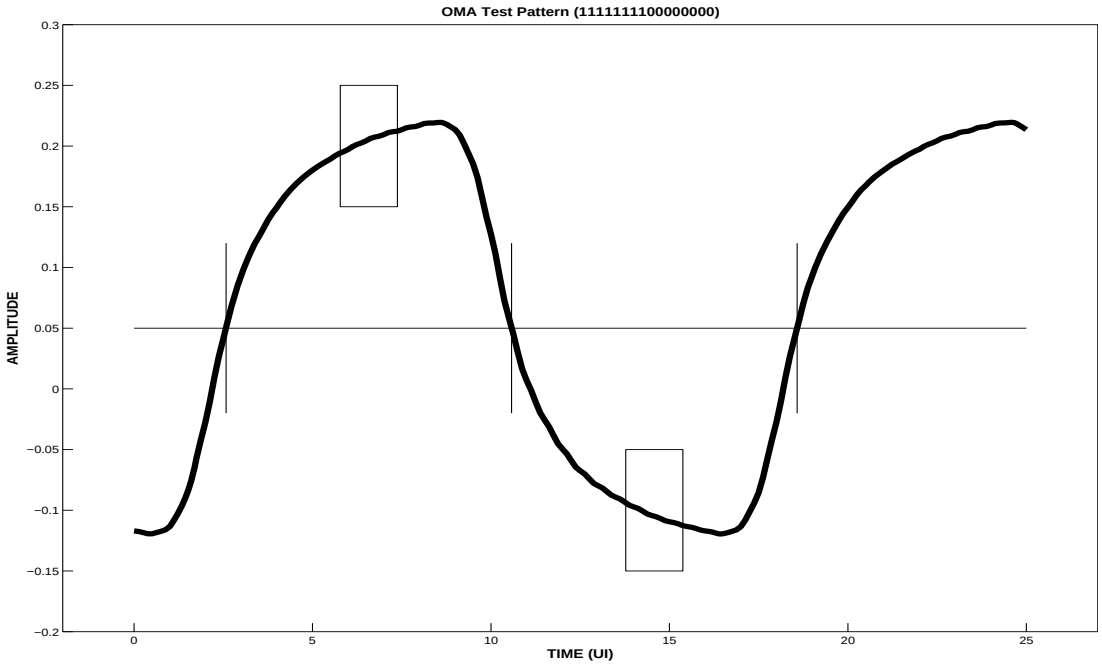


Figure 48 Example xMA waveform showing xMA measurement windows

D.8 RELATIVE NOISE (RN)

RN is a measure of reciprocal SNR for a signal. RN is given by:

$$RN = \frac{2 \times noise(RMS)}{(xMA)}$$

where for this document, xMA is OMA if an optical signal is being measured, or VMA if an electrical signal is being measured, and noise(RMS) is measured on the same optical signal or electrical signal, respectively.

Important parts of the measurement procedure for RN can be found in IEEE Std. 802.3 CL 68.6.7 (LRM). Some comments:

- For purposes of this document, the definitions and procedures generally apply to both optical and electrical signals. Optical terms (such as power) and units can be converted to corresponding electrical terms (such as voltage) and units.
- The test pattern defined for OMA in IEEE 802.3 Clause 68, or other standard relevant for the application, shall be used regardless if the RN measurement is being done on an optical or an electrical signal.

- The 4th-order Bessel-Thomson response is to be used only for optical measurements of RN. The bandwidth of the Bessel-Thomson response is called out in the relevant standard for the application. RN in the electrical domain is defined in a bandwidth of 12 GHz.
- Location of histograms are shown in Figure 68-4 in 802.3 Clause 68.
- Noises at both logic levels should be measured: logicONEnoise(rms) and logicZEROnoise(rms). Apply the rms technique according to the equation:

$$noise(RMS) = \sqrt{(logicONEnoise(RMS)^2 + logicZEROnoise(RMS)^2)/2}$$

- The equation for RN is given above. A calculation of Q_{sq} is not required, nor is a calculation in units of dB/Hz, such as for transmitter RIN. If logicONEnoise(RMS) equals logicZEROnoise(RMS) then RN equals $1/Q_{sq}$.

D.9 WAVEFORM DISTORTION PENALTY (WDP)

WDP is a waveshape metric for waveform filtering and/or nonlinear distortion. WDP uses the same procedure as defined for TWDP in IEEE 802.3 Clause 68.6.6 (LRM).

- For purposes of this document, the definitions and procedures generally apply to both optical and electrical signals. Optical terms (such as power) and units can be converted to corresponding electrical terms (such as voltage) and units, etc.

- WDP is not restricted to transmitter measurements (hence, the "T" is dropped).
- The 4th-order Bessel-Thomson response is to be used only for optical measurements of WDP, such as calibration of an optical receiver test system. The bandwidth of the Bessel-Thomson response is called out in the relevant standard for the application.
- The definition of electrical WDP assumes a measurement bandwidth of 12 GHz. A different measurement bandwidth can be corrected for by processing the captured waveform before the WDP calculation. However, a higher bandwidth is expected to have little effect on the result.
- PRBS9 is the normative test sequence for this specification.
- To improve measurement accuracy, uncorrelated jitter and noise should be reduced. For IEEE 802.3 CL 52, sinusoidal interference and sinusoidal jitter are turned off.
- Averaging should be used to further reduce instrumentation and measurement noise so their effect on the results are negligible.
- Specific code for calculating WDP is found in Appendix G.

D.10 Electrical compliance signal at B" for the SFP+ Module transmitter

Figure 49 shows the test configuration for testing SFP+ transmitters. It applies to all SFP+ transmitter types.

The receive channel of the calibration setup is exercised by the upper crosstalk generator in Figure 47 to ensure that the crosstalk within the setup is acceptable. The crosstalk specifications of Table 17 are to be achieved through the mated host and module compliance boards and into appropriate test equipment.

The compliance signal at B" has deliberate ISI and sinusoidal jitter. It is calibrated through the Host Compliance Board to deliver the DDJ or DDPWS, UJ, and Y1 or Y2 specified in Table 17. The compliance signal is applied to the module under test in place of the Host Compliance Board, with receive side active, so that the transmitted signal can be assessed as specified by the supported transmission standard e.g. 10GBASE-SR, 10GBASE-LR or 10GBASE-LRM. There are four conditions in all: large and small signals, under-compensated and over-compensated. The opposing direction bit stream shall be asynchronous PRBS31 or valid 64B/66B bit stream.

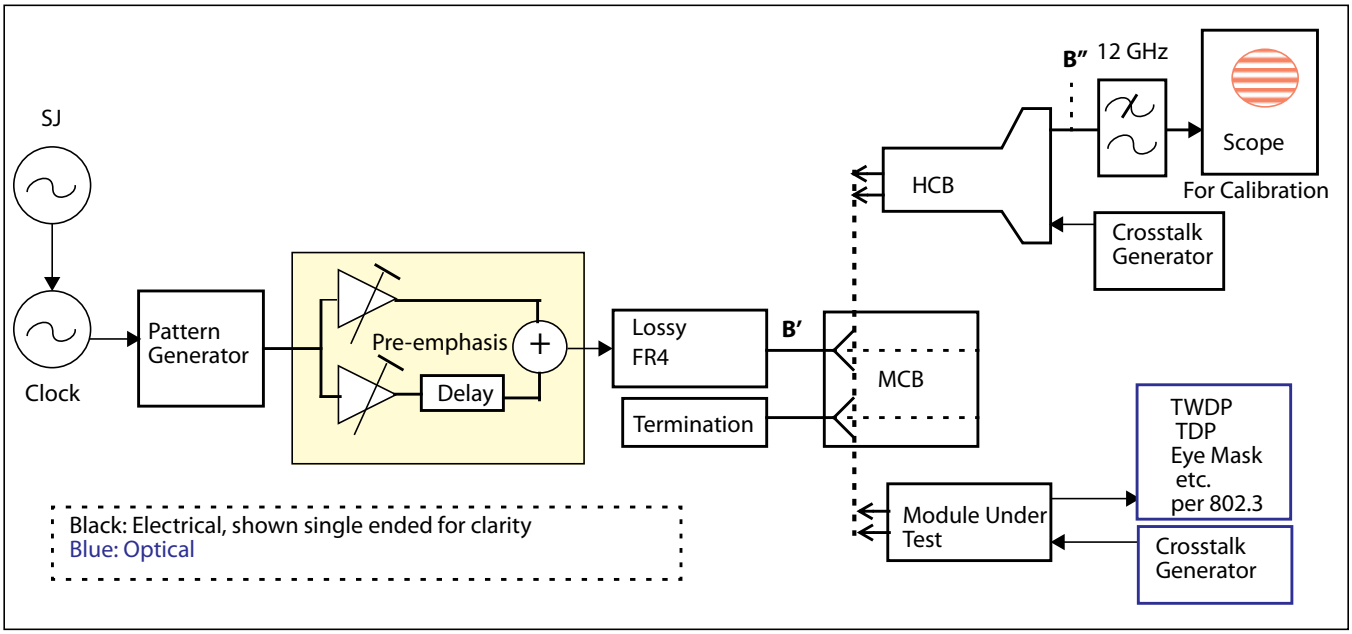


Figure 49 Compliance signal generator for module transmitter

The emphasis settings are adjusted to give the specified DDJ (over-compensated) and DDPWS (under-compensated) at B", in two test conditions. In the over-compensated condition the DDJ shall be equal to the target value in Table 17 while the DDPWS is between 0.045 UI and 0.055 UI. In the under-compensated condition DDPWS shall be equal to the target value in Table 17 while the DDJ is between 0.075 UI and 0.1 UI. The amplitude is adjusted so that an eye mask measurement shows that the compliance signal meets the specified Y1 or Y2 at a hit ratio of 5×10^{-5}. The sinusoidal jitter (SJ) is adjusted to give the specified UJ. Otherwise, the compliance signal is clean and low noise. There are no deliberate Gaussian or "random" impairments other than crosstalk.

The single ended reflection coefficients looking to the right of the HCB and the single ended reflection coefficients looking to the left of the MCB as shown in Figure 49, shall be according to:

$$Sxx(dB) \leq 20 \qquad \text{f in GHz from 0.01 to 5.5}$$

$$Sxx(dB) \leq -25.8 + 1.053 \times f \qquad \text{f in GHz from 5.5 to 15.}$$

The compliance signal complies to the mask in 3.6.1, and has margin to the dimensions given by X1, X2 (jitter margin). The large signal has margin to the dimension given by Y1 and approaches Y2 closely, while the small signal approaches Y1 closely and has margin toY2.

The frequency of the SJ is significantly higher than the bandwidth of the clock recovery unit used to assess the signal transmitted by the module (specified as 4 MHz). Care should be taken that this frequency does not beat against the sampling frequency used to measure the averaged waveform in a TWDP measurement. It must not have a harmonic relationship to the pattern repetition frequency.

The patterns to be used for calibration are specified by the appropriate appendix, e.g. D.3. The patterns to be used with the module, both transmitted and received, are defined by the supported transmission standard. Other characteristics of the compliance signal are defined by the supported transmission standard.

Note that TJ is not intended to be near the maximum TJ allowed in Table 17, and apart from deliberate SJ, there should be much less UJ than the maximum allowed in Table 17. It is recommended that adequate averaging be used in TWDP, DDJ and DDPWS measurements to average the effect of the uncorrelated jitter. Table 30 lists the estimated parameter values for an ideal stressed signal generator.

Table 30 Estimated parameter values for an ideal stressed signal generator

Parameter	*Value*	*Unit*
Delay	1	UI
Filter bandwidth	*For Further Study*	GHz
VMA Min at B”	*For Further Study*	mV
VMA Max at B”	*For Further Study*	mV
Rise times at B”	*For Further Study*	ps

D.11 TEST METHOD FOR A HOST RECEIVER FOR A LIMITING MODULE

This clause provides guidance for jitter tolerance testing at the RX host compliance point C. Compliance is required with input jitter, vertical eye opening (Y1), and vertical peak level (Y2) as specified in Table 14. Compliance is defined at the error rate(s) set by the appropriate optical standard. There are two test conditions; once each for the sensitivity and overload vertical eye parameters conditions.

Further information on definitions and test methods for stressed-eye jitter tolerance are contained in the references (FC-MJSQ and OIF-CEI).

D.11.1 TEST EQUIPMENT & SETUP

A test source is used to continuously generate an appropriate test signal. The test signal shall be appropriately conditioned within the guidelines outlined in D.11.2 to exhibit the appropriate jitter stress.

An RF attenuator or other output amplitude control of the test source may be required to set the vertical eye opening of the stressed eye.

The test equipment measured at C looking into the low pass filter shall have better than 20 dB return loss up to 12 GHz.

The output return loss properties of the test system when measured at C″ with the Module Compliance Board shall be 2 dB better than the specifications of Table 18 up to 8 GHz and 1 dB better up to 11 GHz.

It is required that the receiver under test include a mechanism to allow measurement of BER performance.

D.11.2 STRESSED-EYE JITTER CHARACTERISTICS

This section describes required test signal characteristics along with considerations and suggested approaches for test signal generation. The test signal is generated by the functions shown in Figure 50 or by equivalent means. Figure 51 illustrates how the jitter parameters in Table 14 map to the jitter components in the stressed-eye test signal.

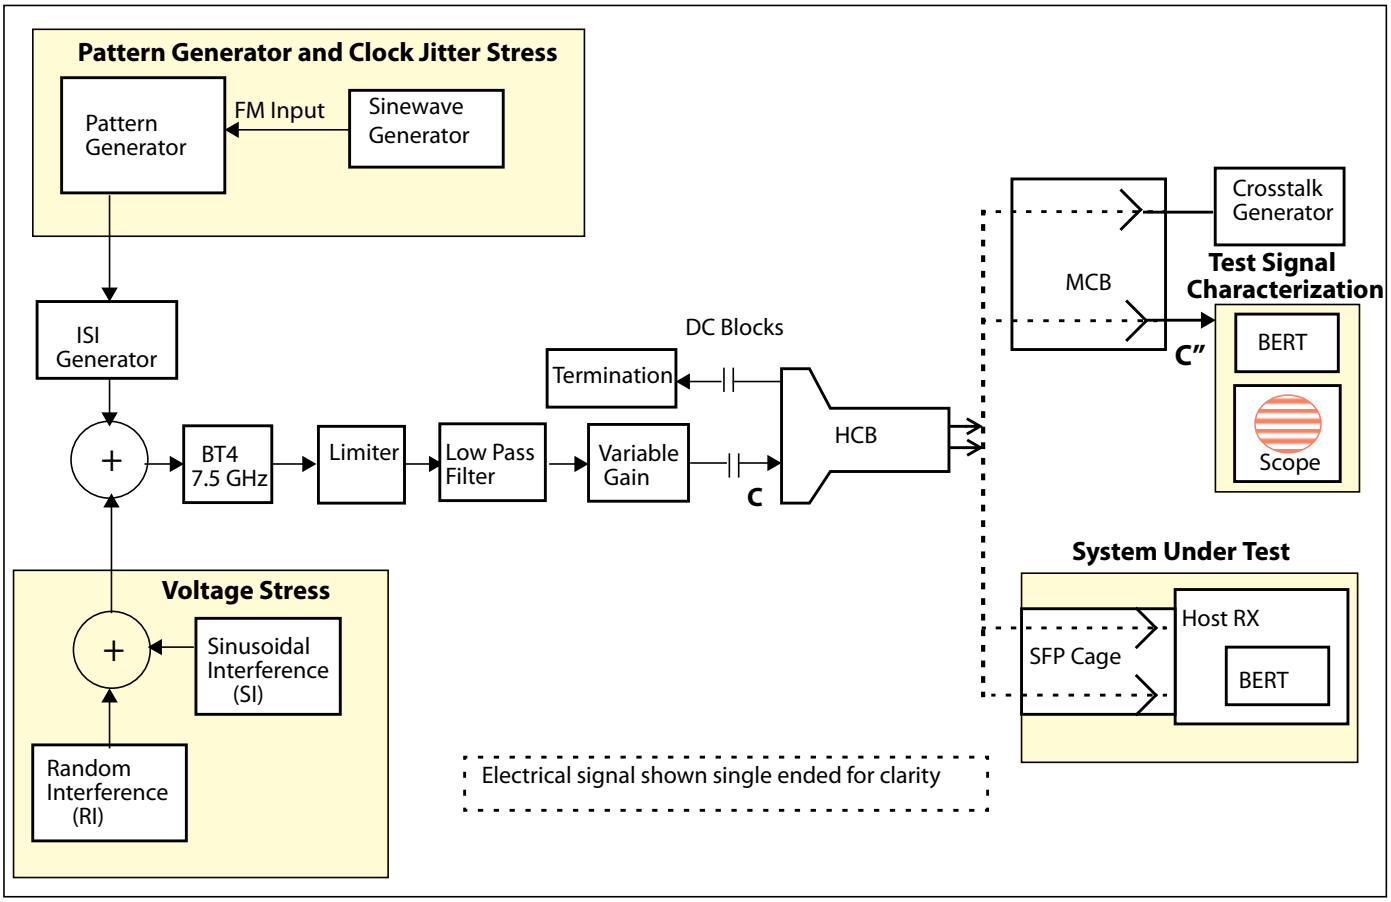


Figure 50 Jitter Tolerance Test Configuration

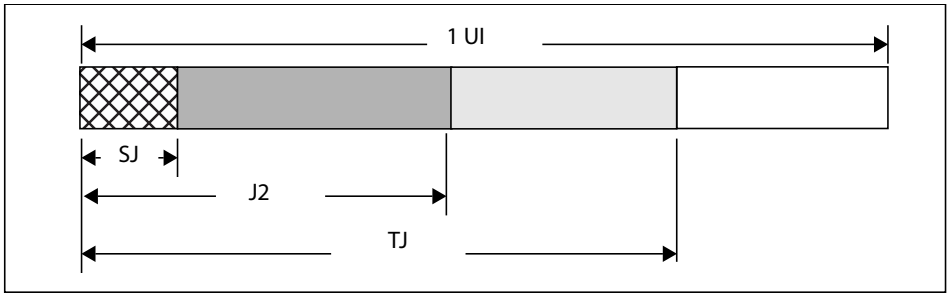


Figure 51 Stressed eye jitter components

The 0.05 UI SJ component of 99% Jitter (J2) is defined for frequencies much higher than the CDR bandwidth (e.g. ~ 20 MHz). At lower frequencies the CDR must track additional applied SJ as detailed in the relevant specifications [1].

1. See Figure 21 and IEEE 802.3 CL52.8.1.

The balance of the J2 is composed of a combination of the following forms of jitter: ISI, sinusoidal interference (SI), and random interference (RI) all passed through a limiting function.

The signal at C″ shall have DDPWS as defined by Table 14. Magnitude of any DCD (see D.3.1) in the test shall not exceed 0.02 UI.

ISI jitter creation may be achieved by the ISI generator through the use of a low pass filter, length of FR4 trace, length of coax cable or other equivalent method. It is required that this signal be passed through a limiter function to ensure that the resulting jitter is not totally equalizable jitter. A suitable limiter function may be implemented using a discrete limiting amplifier followed by a low pass filter and an attenuator. The low pass filter emulates the bandwidth and/or slew rate of a practical limiter. The attenuator is used to set the output amplitude to minimum and maximum values allowed by the eye mask coordinates of Table 14.

A voltage stress is to be applied before the limiter. This stress is composed of a single tone sinusoidal interferer (SI) in the frequency range 100 MHz to 2 GHz and a broadband noise source (RI) with a minimum power spectrum of -3 dB at 6 GHz and minimum 7 crest factor. It is the intent that this combination of voltage stress and limiting function introduce pulse-shrinkage jitter behavior. However no more than 20% of the J2 is created by the sinusoidal interferer.

Jitter generation mechanisms for the pattern generator are typically based on phase modulation of the clock source, edge modulation of a variable delay line or a combination thereof.

Any approach that modulates or creates the appropriate levels and frequencies of the jitter components is acceptable.

D.11.3 CALIBRATION

Calibration of the test signal is to be performed using the guidelines for test setup in D.11.1 and illustrated in Figure 50. The aim of the calibration is to achieve a test signal exhibiting jitter stress in accordance with Table 14.

The test signal should be calibrated differentially into standard instrumentation loads. If complementary single-ended signals are used they should be carefully matched in both amplitude and phase.

For improved visibility for calibration, it is imperative that all elements in the signal path (cables, DC blocks, etc.) have wide and flat frequency response as well as linear phase response throughout the spectrum of interest. Baseline wander and overshoot/undershoot should be minimized.

An AC coupling 3 dB corner frequency of 20 kHz is expected to be adequate to eliminate baseline wander effects, however high frequency performance is critical and must not be sacrificed by the AC coupling.

Jitter requirements are defined for a probability level of $1x10^{-12}$. To calibrate the jitter, methods given in CEI 2.C Annex and MJSQ Chap 8 are recommended. Given random jitter and the nature of the long test patterns, low probability jitter events will likely be present. It is recommended for jitter calibration that a technique that can accurately measure low probability events should be used to avoid overly stressful test conditions.

It is recommended that the actual compliance test pattern be used during calibration. For jitter stress calibration it is permissible, however, to use any appropriate test pattern which still results in the creation of a compliance test pattern with the appropriate jitter stress.

D.11.4 CALIBRATION PROCEDURE

The vertical eye opening and peak level should be set approximately to the levels specified in Table 14.

With an applied calibration test pattern and no additional jitter stress applied; the intrinsic jitter of the test source due to intrinsic noise and finite bandwidth effects should be measured and calibrated. The 99% jitter (J2) shall be <0.15 UI and TJ <0.25 UI.

SJ should be added until the J2 component of jitter increases by 0.05 UI above the measured reference level. This should be high frequency SJ well above the CDR bandwidth. The SJ frequency should be asynchronous to the characteristic frequency of the signal.

Next, additional high probability jitter as specified in D.11.2 should be added by the ISI generator until at least 80% of the J2 has been created. The Sine Interferer amplitude should then be turned on and adjusted until the required level of J2 is achieved. The frequency of any Sine interferer should be asynchronous to the characteristic frequency of the signal.

A compliant test signal exhibits data dependent pulse width shrinkage as specified in Table 14. Data dependent pulse width shrinkage is defined in D.3. This is measured with noise and clock-jitter sources turned off.

Once the required level of J2 has been achieved turn on the crosstalk source that should be set such that at the output of the Host Compliance Board the amplitude and the rise and fall times should be as given in Table 14. The crosstalk pattern should be PRBS31 or valid 64B/66B signal and should be asynchronous with the data. Then the RI (random interference) should be added until the required value of TJ is achieved at a probability of $1x10^{-12}$.

If necessary the sine interferer should be readjusted to obtain the required level of J2 and if the sine interferer is changed then the random interferer should be readjusted to obtain the required level of TJ. Iterative adjustments of the sine interferer and random interferer should be made until the required values of both J2 and TJ are achieved.

If necessary, the vertical eye opening should be readjusted to required levels.

It should be verified that the vertical eye opening and peak level specification is met.

Care must be taken when characterizing the signal used to make receiver tolerance measurements. The intrinsic noise and jitter introduced by the calibration measurement equipment (e.g. filters, oscilloscope and BERT) must be accounted for and controlled. If equipment imperfections affect the results materially, corrections such as RSS deconvolution of Gaussian noise and jitter should be used.

D.11.5 TEST PROCEDURE

Testing should be performed differentially through a Host Compliance Board (see C.2).

Using a test signal calibrated conforming to D.11.1 and calibrated as per D.11.4, operate the system with an appropriate compliance test pattern for the relevant application (10G Ethernet, 10GFC, or 10G Ethernet with FEC).

All signals and reference clocks that operate during normal operation shall be active during the test including the other host signal path in the duplex pair. The other signal path shall be asynchronous.

The opposing direction bit stream (than the one being tested) shall be asynchronous PRBS31 or valid 64B/66B signal.

The sinusoidal jitter is stepped across frequency and amplitude range according to Figure 21 while monitoring the BER. The BER shall remain $< 1x10^{-12}$.

D.12 LIMITING MODULE RECEIVER COMPLIANCE TESTS

Compliance to the specifications at C' Table 18 and Table 19 must be met over the range of input optical signals specified by standards supported e.g. IEEE 802.3 Clause 52 and calibration procedure defined in Clause 52.9.9.

This test includes the effects of crosstalk within the module and within the Module Compliance Board. The module transmit path is operational. The transmit path input of the Module Compliance Board is connected to a pattern generator and calibrated through a Host Compliance Board. The amplitude

and rise/fall times are given in Table 18. Testing for compliance at point C' is done through a Module Compliance Board.

The pattern for the crosstalk source is PRBS31 or a valid 64B/66B sequence. The crosstalk source is asynchronous to the TP3 test source.

The minimum test conditions (vertical eye closure penalty, VECP[min] and stressed eye jitter, J(min) for stressed receiver sensitivity defined in IEEE 802.3 clause 52 were chosen as sufficient to ensure compliant receivers. Consequently, test conditions more severe than the minimum requirements represent an overstress condition for which compensation is appropriate.

Compensation for overly stressful VECP is straightforward; the stressed receiver sensitivity (SRS) maximum can be adjusted one-for-one for any VECP overstress (or a small amount of under-stress).

SRS[compensated, dBm] = SRS[max, dBm] + dVECP[overstress, dBo],

where dVECP[overstress] = VECP[measured, dBo] - VECP[min, dBo].

Compensation for overly stressful jitter is less straightforward since definition permits compositions over a trade-off range of deterministic and random jitter. Further, there is no generally accepted practice for compensating deterministic jitter and the only recourse is re-calibrating the test source. Fortunately, most cases of overly stressful jitter are expected to be due to excessive random jitter.

Where the jitter composition is known or can be measured, any excess random jitter can be backed out of the measured result, or specifications in Table 19, Total Jitter and Eye Mask X1 coordinate, can be adjusted to accommodate the excess input signal.

D.13 Test Method for a Host Receiver with a Linear Module

A compliance setup for a host for use with a linear module receiver is shown in Figure 52. The host input at point C is tested for BER compliance with test signals that represent the worst case waveshape and noise properties expected from the output of a module during compliant operation.

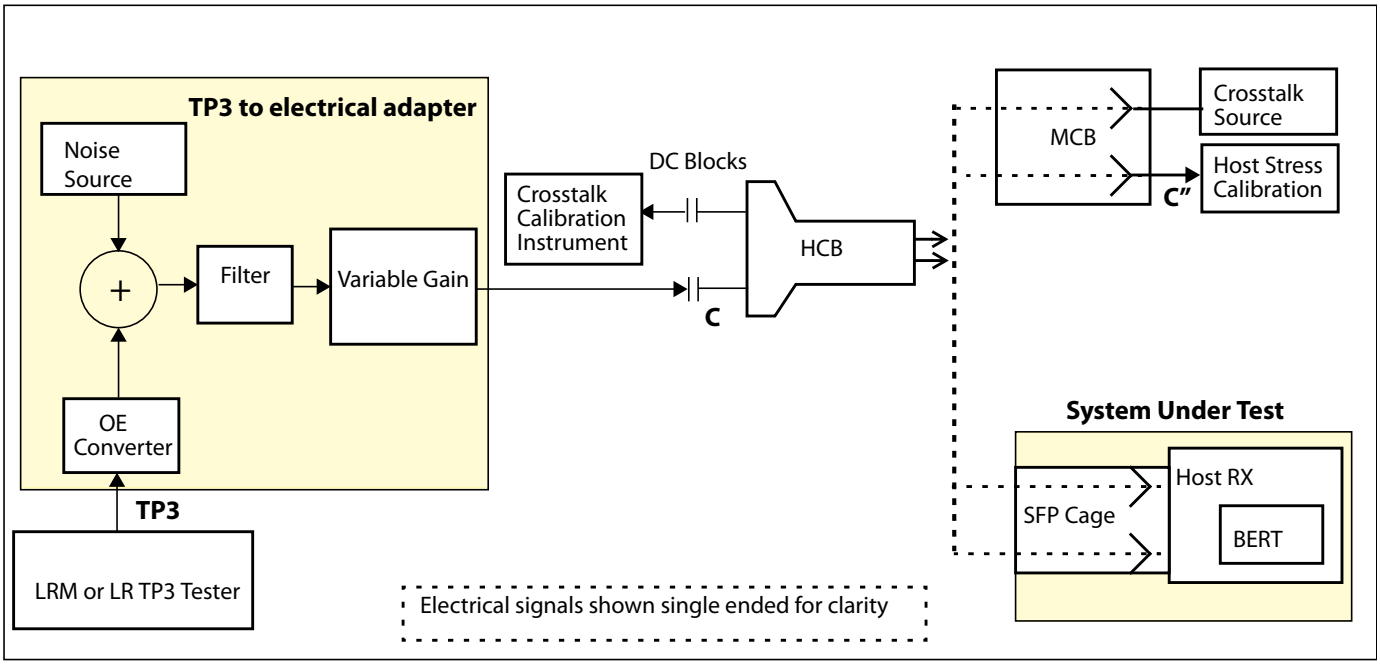


Figure 52 TP3 to electrical adaptor for host that operates with linear modules

D.13.1 TEST DESCRIPTION AND PROCEDURE FOR HOST RECEIVER FOR LINEAR MODULE

Compliance shall be achieved for each of the three TP3 pulse shapes defined for 10GBASE-LRM in IEEE 802.3 Clause 68.6.9 and for the one 10GBASE-LR stressed receiver conformance test signal defined in IEEE Std 802.3 Clause 52.9.9. Compliance shall be achieved over the range of VMA in Table 15. The TP3 tester block is the same test system as defined by the LRM or LR standard for testing the TP3 compliance point. LRM and LR are chosen because this combination of tests includes both high distortion with low noise, and low distortion with high noise. Testing with an SR equivalent input is not required as the noise and distortion are between those for LR and LRM.

The TP3 to electrical adapter as shown in Figure 52 converts the TP3 test signal(s) into electrical signal(s) with output VMA, noise (RN) and distortion (WDP) properties defined in Table 15.

The specifications given in Table 15 are as measured during calibration at C″ through the Module Compliance Board.

The noise source, in conjunction with the other blocks, is intended to represent the additive noise properties of a worst-case linear module. The magnitude of the noise is calibrated such that the RN values at C″ are consistent with Table 15. The spectrum of the noise source at the summing point is white with a 3 dB frequency of at least 10 GHz. The noise measured at C″ represents the

noise of the module and the optical signal combined. The noise source crest factor should be at least 6.

The filter and gain blocks are intended to represent the deterministic dWDP and gain properties of a worst-case linear module. For the low WDP cases in Table 15, including LR, the filter has a bandwidth of 7.5 GHz. For the high WDP cases in Table 15, the frequency response of the filter is set such that the WDP value specified in Table 15 at C" for the split-symmetrical LRM stressor is achieved. This bandwidth is expected to be approximately 4.5 GHz. In all cases, the overall response of the adapter has a Bessel Thomson response.

The gain block and/or the input optical power level can be used to adjust VMA.

During calibration and host compliance testing, crosstalk source see Figure 52 shall be an asynchronous PRBS31 or 64B/66B signal.

Care must be taken to not induce greater than 0.02 UI of DCD at C"

A balun or other means provides a differential signal.

The test signal output shall be AC coupled. An AC coupling 3 dB corner frequency of 20 kHz is expected to be adequate to eliminate baseline wander effects, however high frequency performance is critical and must not be sacrificed by the AC coupling.

The output return loss properties of the test system when measured with Module Compliance Board shall be at least 2 dB better than the specifications of Table 18 up to 8GHz and 1 dB better up to 11GHz.

Any implementation of the measurement configuration may be used, provided that the resulting signal and noise match those defined in Table 15.

Under all specified test conditions, a BER of better than $1x10^{-12}$ shall be achieved. The transmitter of the port under test and all other ports operate in normal operation, including termination. The transmitter of the port being tested is terminated through the Host Compliance Board with a DC block and 50 Ω at each Tx SMA connector.

D.13.2 HOST LINEAR TESTER CALIBRATION

The output of the Host Compliance Board is plugged through the Module Compliance Board into laboratory equipment for calibration.

Calibration should be done with all tester elements in place, although some components may be shut down, such as jitter and noise, while other elements are being calibrated - see below. After calibration is completed, all components are set to their calibrated levels for testing.

RN of the host test system is adjusted via the magnitude of the adapter's noise source. Calibration should use the RN measurement methods given in section D.8. RN values are given in Table 15 for each test condition. The crosstalk source must be calibrated to the requirements in Table 13 and running during calibration of RN. After calibration and during host compliance testing, the crosstalk calibration instrument can be removed and replaced with 50 Ohm terminations, although DC blocking must be maintained.

WDP of the host test system is set via the filter in the adapter. If the calibration is off by a small amount, the ISI generator in the TP3 tester can be adjusted to obtain the required values.

Although WDP is a characteristic of an electrical signal in this case, its units are in dBo to better align with WDPo out of a linear optical module, which is also given in dBo.

After calibration, the Host Compliance Board is plugged into the host receiver under test for compliance testing.

D.14 LINEAR MODULE RECEIVER COMPLIANCE TESTS

Linear module receiver compliance tests ensure that noise generation, waveform filtering and other distortion due to the module are kept within acceptable bounds when tested with the optical input signals as specified in the standards supported by the module, e.g. IEEE 802.3 CL 52 and/or CL 68.

D.14.1 LINEAR MODULE RECEIVER NOISE COMPLIANCE TEST

The module receiver can be tested for noise compliance by measuring how much noise it passes and adds to an input test signal. Figure 53 is a block diagram of a test system that defines the module receiver noise test.

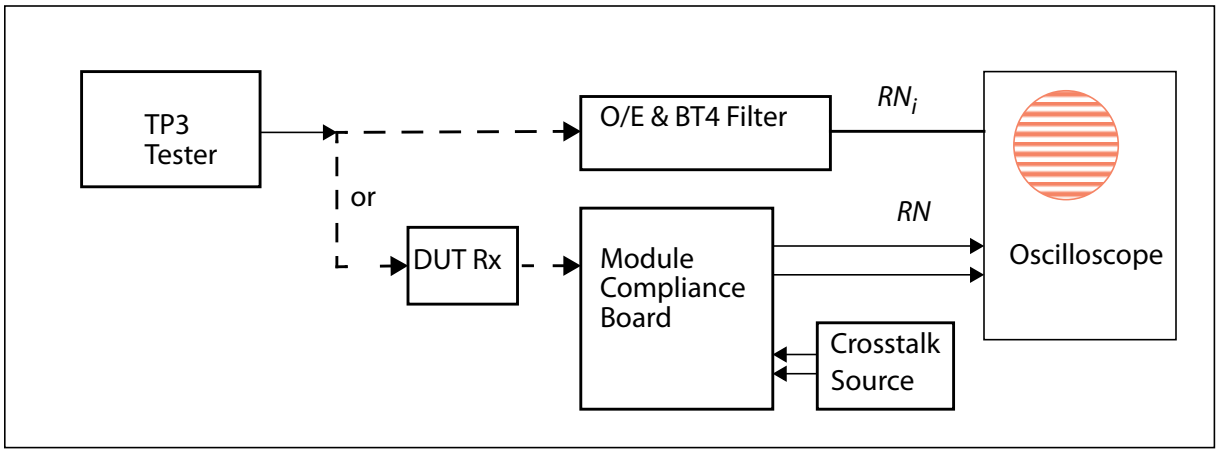


Figure 53 Linear Module Receiver Noise Test

This test includes the effects of crosstalk within the module and the Module Compliance Board. The transmit path input of the Module Compliance Board is connected to a crosstalk source and calibrated through a Host Compliance Board. The crosstalk amplitude and rise/fall times are set to the values given in Table 18. The pattern for the crosstalk source is PRBS31 or a valid 64B/66B signal. The crosstalk source is asynchronous to the TP3 test source. After calibration, the Host Compliance Board is replaced with the module under test. The module transmit path is operational during compliance testing.

The TP3 tester should be set to the OMA/VMA pattern for this test as defined in D.7.

The waveform shaping stress of the TP3 tester is enabled. The sinusoidal jitter and/or sinusoidal interference of the TP3 tester should be disabled or set to very low magnitudes for this test.

RN_i of the TP3 tester is set to the level specified by the Table 31.

The TP3 tester is connected into the module under test. The module is plugged into the Module Compliance Board, which in turn is connected to the oscilloscope. The relative noise of the module output signal, RN, is then measured. The relative noise measurement method is described in D.8.

Relative noise of the TP3 test signal RN_i is characterized through a reference O/E converter and 4th-order Bessel Thomson filter and a digital oscilloscope. If the noise of the TP3 test source does not match the target value in Table 31, RN can be corrected using the following equation:

$$RN = \sqrt{(RN_{measured})^2 - 1.24 \times RN_i(target) \times (RN_i - RN_i(target))}$$

where $RN_{measured}$ includes the effect of actual TP3 tester noise at the module output, RN_i is the actual TP3 tester noise, and RN_i(target) is the target test noise given in Table 31 for the test conditions.The resulting noise result is to be compared against the compliance limit specified in Table 20. Compliance must be met over the range of optical power specified by the standards supported by the module.

.

Table 31 Target RN_i Values

Application	*RN_i (target)*
LRM pre-cursor	0.0219
LRM split-symmetrical	0.0269
LRM post-cursor	0.0213
LR	0.014
SR	0.020

For LRM, RN_i should be within 1 dBo of the appropriate value given in Table 31. For LR and SR, RN_i should be no more than 1 dBo greater than the appropriate value in Table 31; any lower value is allowable.

This procedure is described for an oscilloscope as the measuring instrument. However, noise generated by a practical scope can affect the result. The noise due to the scope is calibrated out of the result by subtracting the square of the scope's noise from the noise of the RN measurement as appropriate, so as to obtain the relative noise associated with the signal under test. For electrical scope noise measurement, the scope inputs are terminated with 50 Ω termination. For optical scope noise measurement, the scope input should have zero light.

D.14.2 LINEAR MODULE RECEIVER DISTORTION PENALTY COMPLIANCE TEST

This section defines dWDP, a measure of waveform filtering and other distortion associated with the linear optical receiver. The block diagram dWDP test system that defines linear module receiver distortion test is shown in Figure 54.

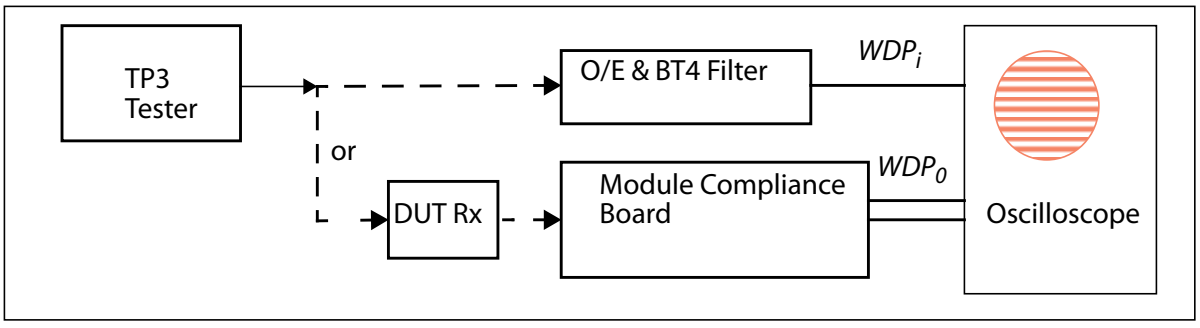


Figure 54 Module receiver waveform penalty compliance test

- WDP_i and WDP_o in Figure 54 are measured using the WDP method defined in D.9. WDP_i of the TP3 test signal is first characterized through an O/E converter and 4th-order Bessel Thomson filter and a digital oscilloscope. For 10GBASE-LRM, this signal should represent the waveforms described in IEEE Std. 802.3 CL 68.6.9, and for 10GBASE-LR, this signal represents the waveform described in IEEE Std. 802.3 CL 52.9.9.
- The TP3 tester is removed from the O/E converter and connected into the module under test. The module in turn is plugged into a Module Compliance Board which in turn is connected to the oscilloscope. WDP_o of the module output signal is then measured.

Although WDPo is based upon measurements of an electrical signal, its units for a linear optical module output are in dBo to allow a direct comparison with the optical input signal in the equation for dWDP below.

The distortion contributed by the module is determined by the following equation:

$$dWDP = WDP_o - WDP_i$$

dWDP is to be compared against the compliance limit specified in Table 20. Each dWDP must comply for each specified TP3 condition. The TP3 tester is the same test system as defined by the relevant standard for testing the TP3 compliance point.

D.14.3 LINEAR MODULE RECEIVER OUTPUT DIFFERENTIAL PEAK-PEAK VOLTAGE

A compliant TP3 stress receiver tester for the relevant application (SR, LR, or LRM) is connected to the module receiver input. The OMA test pattern for the application should be used, and all stress impairments such as sine jitter, sine interference, ISI, and noise should be turned off. The rise/fall time should be 47 ps 20-80%. When observed through a 7.5 GHz reference O/E converter, the input waveform should have no overshoot or ripple.

The output of the module is measured with a Module Compliance Board connected into an oscilloscope. The measurement bandwidth is 12 GHz. A wider measurement bandwidth is expected to have only a minor effect on the result. If the measurement bandwidth affects the results, it can be corrected for by post processing. Averaging is used to eliminate noise from the measurement. The peak to peak swing of the differential signal is measured and compared against the limit in Table 20.

D.15 AC COMMON MODE VOLTAGE

The SFI transmitter and channel limit but do not eliminate AC common mode voltage generation. SFI receivers, both module and host, must operate fully with the maximum allowed input common mode voltage. Common mode voltage often gets generated due to the crossing points of the driver outputs (P and N) being shifted from 50%, impedance mismatch, mismatch of the PCB traces, or mode conversion.

D.15.1 DEFINITION OF AC COMMON MODE VOLTAGE

The common mode voltage at any time is the average of signal+ and signal- at that time. The RMS AC common mode voltage is calculated by applying the histogram function over one UI to the common mode signal. As AC common mode generation is very sensitive to the cable or scope delay mismatch, it is recommended to delay match the scope inputs for any measurements.

D.15.2 AC COMMON MODE GENERATION TEST

The test pattern for AC common mode generation is either pattern 1 (BnBi) or pattern 3 (PRBS31) as defined in IEEE CL 52.9.1.1. It is expected that any 64B/66B scrambled signal should give a similar result.

D.15.3 AC COMMON MODE TOLERANCE TEST

The test pattern for AC common mode tolerance is either pattern 1 (BnBi) or pattern 3 (PRBS31) as defined in IEEE CL 52.9.1.1. It is expected that any 64B/66B coded signal should give a similar result.

If the transmitter output does not generate a sufficient amount of AC common mode then the method to generate additional amounts is by adjusting the P and N delay until the right amounts is generated.

D.16 TERMINATION MISMATCH

Termination mismatch is defined as the percent difference between the complimentary Z_p and Z_n resistors as shown in Figure 18. Termination mismatch is defined as:

$$\Delta Z_M = 2 \times \frac{Z_p - Z_n}{Z_p + Z_n} \times 100$$

Alternatively, the termination mismatch can be measured by applying a low frequency test tone to the differential inputs as shown in Figure 55. The test frequency must be high enough to overcome the high pass effects of the AC

coupling capacitor. The measured differential output or input impedance is designated by Z_{diff}.

Low frequency termination mismatch is then given by:

$$\Delta Z_M = 2 \times \frac{I_p - I_n}{I_p + I_n} \cdot \frac{Z_{diff} + 100}{Z_{diff}} \cdot 100$$

where I_p and I_n are the current flowing into the SFI port as shown in Figure 55. Z_s is the effective series impedance between the driver terminations Z_p and Z_n and the AC Ground.

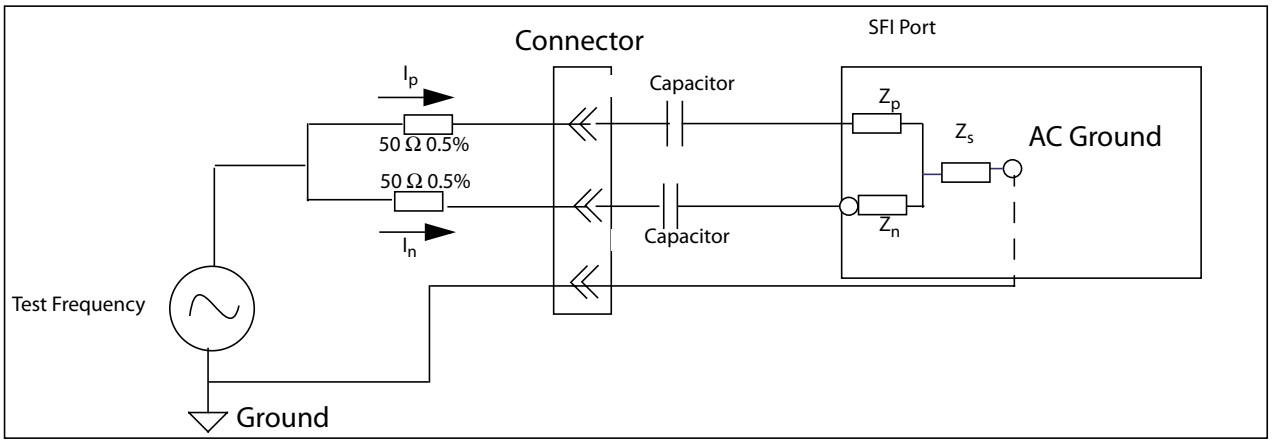


Figure 55 AC Termination Mismatch Measurement

D.17 POWER SUPPLY TESTING METHODOLOGY

This section defines power supply noise output as given in 2.8.2 and 2.8.3, and power supply noise tolerance as in 2.8.4.

The reference power supply filter shown in Figure 56 is provided for module testing, including power supply tolerance testing. This filter will meet the noise filtering requirements in most host systems. Other filtering implementations or local regulation may be used to meet the power noise output requirements described in 2.8.2 and 2.8.3.

For each Vcc, the sum of the equivalent series resistances of the 4.7uH inductor, the 22uF capacitor and the damping resistor is 0.5 Ω. This resistance is desirable in actual host filters as well as in the reference filter; however, any voltage drop across a filter network on the host is counted against the host VccT and VccR accuracy specification in Table 8.

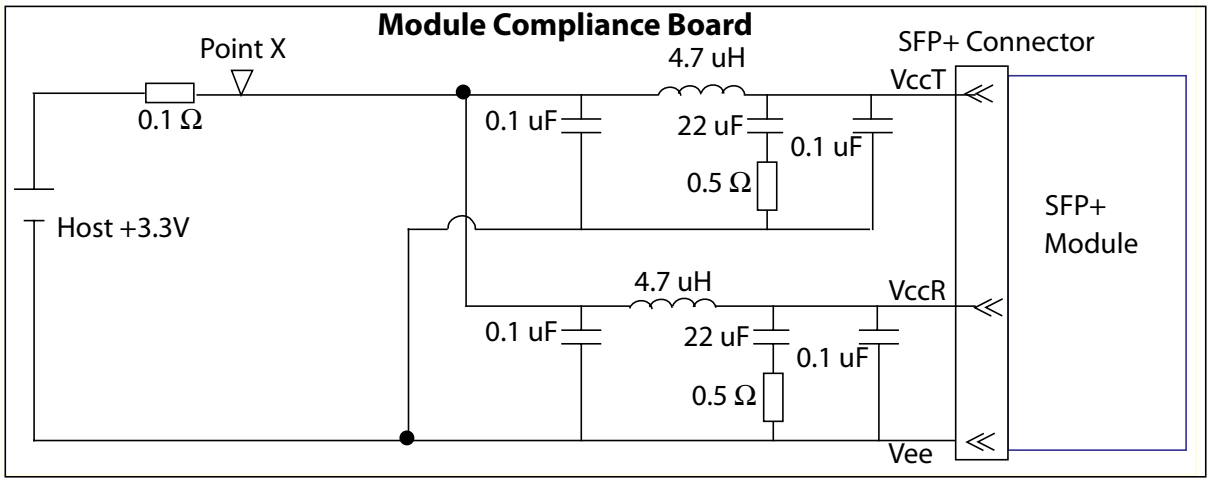


Figure 56 Module Compliance Board Power Supply Filters

D.17.1 HOST POWER SUPPLY NOISE OUTPUT

The noise output of a Vcc supply of a host is defined with a resistive load that draws the maximum rated power (1 W or 1.5 W) connected between one Vcc contact and Vee, in place of the SFP+ module. When the noise on VccT is being measured, VccR is left open circuit, and vice versa. 8 Ω is used for a host capable of supporting Power Level II, and 12 Ω otherwise. The AC voltage spectrum is measured at the module side of the SFP+ connector. The noise power spectrum is divided by the truncated response of the reference filter and then integrated from 10 Hz to 10 MHz and converted to a voltage. This function is illustrated in the equation below and Figure 57. The specification limit is given in 2.8.2. The test is performed with all other portions of the host board/system active. Hosts with multiple SFP+ modules shall test ports one at a time, with active SFP+ in all the remaining ports.

$$H(f) = a \times (log10(f))^4 + b \times (log10(f))^3 + c \times (log10(f))^2 + d \times (log10(f)) + e$$

The reference filter response *H(f)* shown in Figure 57 and the coefficients a, b, c, d, and e for the 5 frequency bands are defined in Table 32.

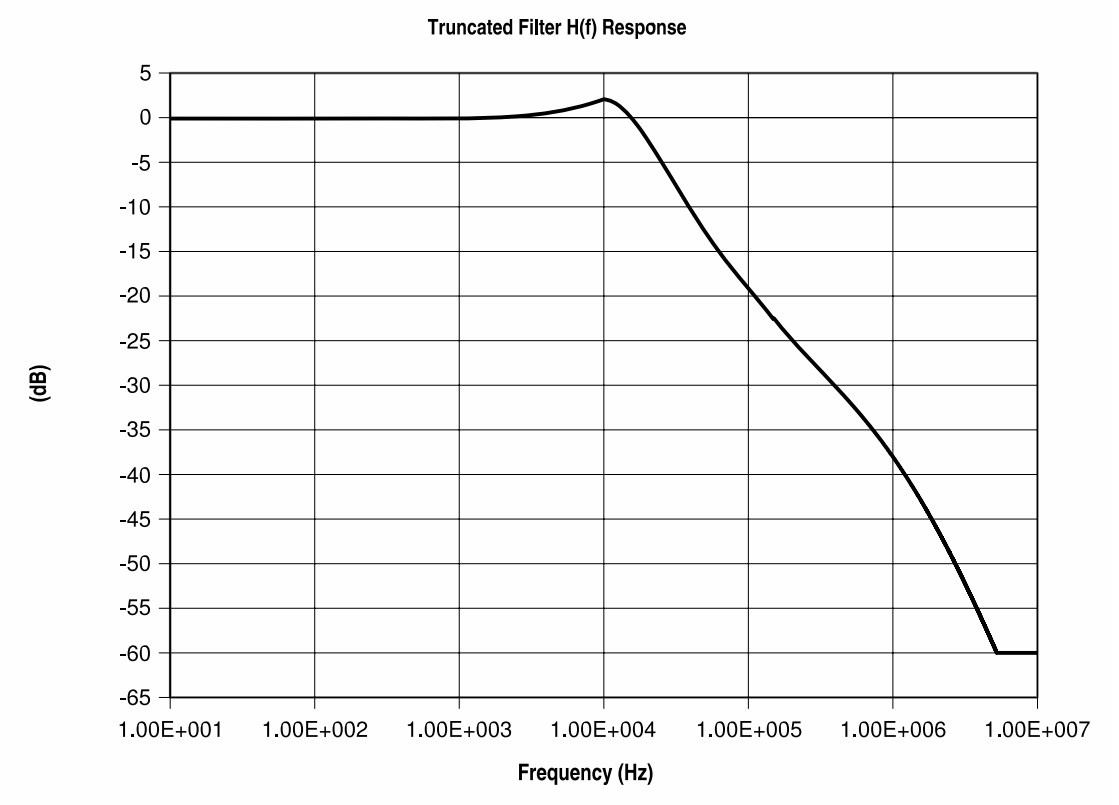


Figure 57 Reference filter response

Table 32 Truncated Filter Response Coefficients

Frequency (f)	*a*	*b*	*c*	*d*	*e*
10 Hz ≤ f < 100 Hz	0	0	0	0	-0.1
100 Hz ≤ f < 10 kHz	0.3784	-3.6045	12.694	-19.556	11.002
10 kHz ≤ f < 150 kHz	-22.67038	430.392	-3053.779	9574.26	-11175.98
150 kHz ≤ f ≤ 5.243 MHz	3.692166	91.467	838.80	-3400.38	5139.285
5.243 MHz ≤ f ≤ 10 MHz	0	0	0	0	-60

NOTES - As a lightly loaded power supply might generate more noise than a fully loaded supply, the host implementer may wish to assess the host power supply noise output at less than maximum current draw also. Because a small measured noise signal at high frequencies is multiplied up to give the inferred noise at virtual point X, care should be taken over the noise floor of the spec-

trum analyzer. Other measurement methods could be used, e.g. a measurement at a point inside the host, with appropriate consideration to any difference between the reference filter and a host's actual filter.

D.17.2 SFP+ Module Power Supply Noise Output

The module noise voltage output is defined in the frequency band 10 Hz to 10 MHz at point X in Figure 56.

The module must pass module power supply noise output test in all operating modes. This test ensures the module will not couple excessive noise from inside the module back onto the host board. A power meter technique, or a spectrum analyzer technique with integration of the spectrum, may be used. The maximum allowed noise amplitude is given in 2.8.3.

D.17.3 Module Power Supply Tolerance Testing

In this test, a swept sinusoidal tone is applied at point X of Figure 58 with the tolerance signal amplitude and frequency range given in 2.8.4. The AC tolerance signal is created by a circuit such as a low impedance buffer amplifier between the power supply and point X. The impedance of the PSU and sine wave generator is less than 0.5 Ω. The amplitude of the sine wave is calibrated at each frequency at point X with the module replaced with a 12 Ω load between Vcc and Vee.

NOTES -- It may be desirable to remove the 0.1 uF capacitors on the host side of the reference filters for this test, to reduce the power needed by the sine wave generator. The calibration of the sine wave is not expected to be significantly different if the module were in place rather than the test resistor.

Alternatively, the test may be performed separately for VccT and VccR with the other supply filter connected directly to the power supply. It is not necessary to show compliance with both separate and common Vcc modulation.

This test applies at minimum and maximum DC setpoint levels. Note that the DC level is inset to the limits in Table 5 by the peak of the sinusoidal voltage at the input to the module (which is frequency dependent).

The source frequency is varied over the range specified by 2.8.4 to determine if any frequency causes a parameter to fall out of the specification limit. In all cases, the parameters measured shall pass the optical standards with the tone present over all frequencies specified. Parameters of interest for the transmit may include UJ, Q_{sq} and TDP see [IEEE 802.3]. For the receive side, they include stressed sensitivity, overload, RN and Rx_LOS function.

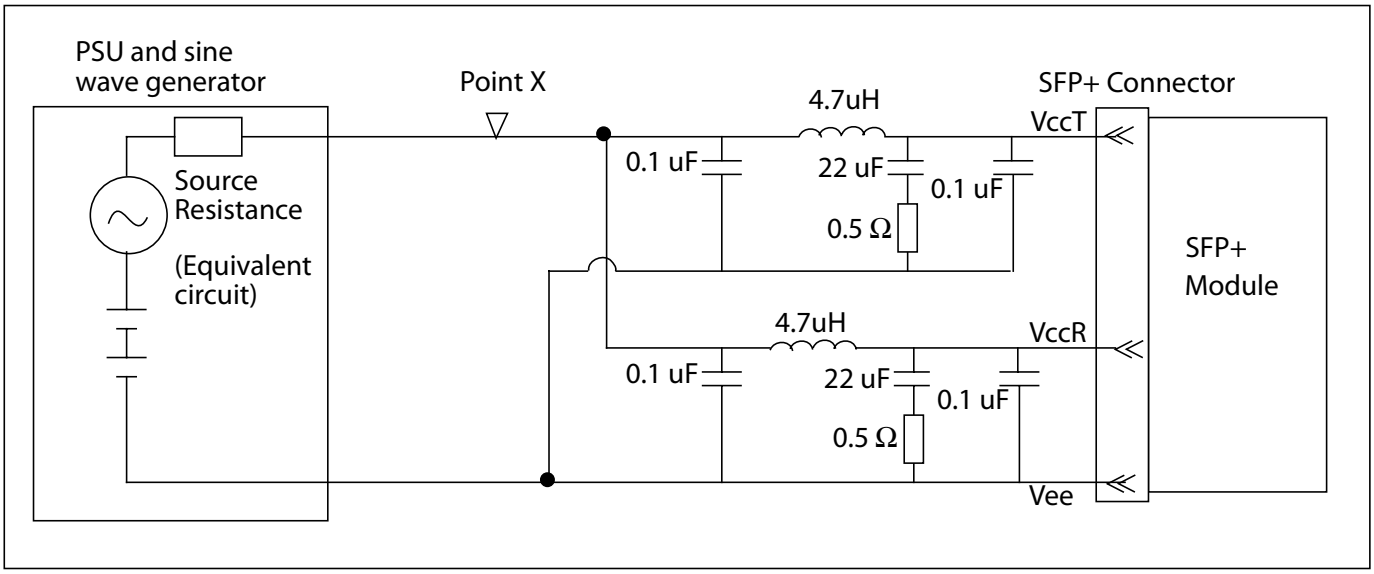


Figure 58 Power Supply Noise Tolerance Test Setup

The AC voltage at node X is defined with reference to Vee. The DC voltage specification including ripple, droop and noise below 100 kHz is met at both VccT and VccR (at the SFP+ connector).

Appendix E SFP+ Direct Attach Cable Specifications "10GSFP+Cu" (Optional)

A passive copper cable compliant to this appendix is identified using the 2-wire management defined in chapter 4 and memory map of SFF-8472.

This appendix describes additional requirements or exceptions to the linear host specification of Chapter 3 to implement passive direct attach SFP+ cable assemblies.

The compliance points for SFP+ Direct Attach Cable (10GSFP+Cu) are the same as host compliance test points 3.3.1 and the module compliance test points in 3.3.2.

All SFI test equipment must have 50 Ohms single ended impedance on all test ports.

Each Tx_Disable contacts of 10GSFP+Cu passive cable assemblies shall be pulled to VccT with a 4.7 kΩ to 10 kΩ in the module. The Rx_LOS contacts in the module shall be pulled low in the module for 10GSFP+Cu passive cable assemblies. Direct connection of Rx_LOS to VeeR is allowed.

Active cable assemblies must operate with existing linear or limiting specifications of Chapter 3.

This specification does not assume additional transmit pre-emphasis beyond the level required to meet the jitter specifications at point B (see Table 12) and TWDPc specification (see Table 33) at point B. Increasing the transmit pre-emphasis may increase cable reach, however it may increase transmitter DDJ and is outside the scope of this specification.

Warning: 10GSFP+Cu can only be used on systems with common grounds. Connecting systems with different ground potential with SFP+ direct attach cable results in a short and may cause damage.

E.1 10GSFP+Cu Direct Attach Construction

10GSFP+Cu cable assemblies are effectively constructed out of a pair of SFP+ modules with the OE components replaced with copper cabling as shown in Figure 59. SFP+ Edge card connector contacts are defined in Table 3. The cable assembly shall incorporate DC blocking capacitors with at least 4.3 V rating on the RX side and with high pass pole of between 20 kHz and 100 kHz. The drain wire is connected to VeeT and to VeeR. The cable shield directly connects the module A and B cases.

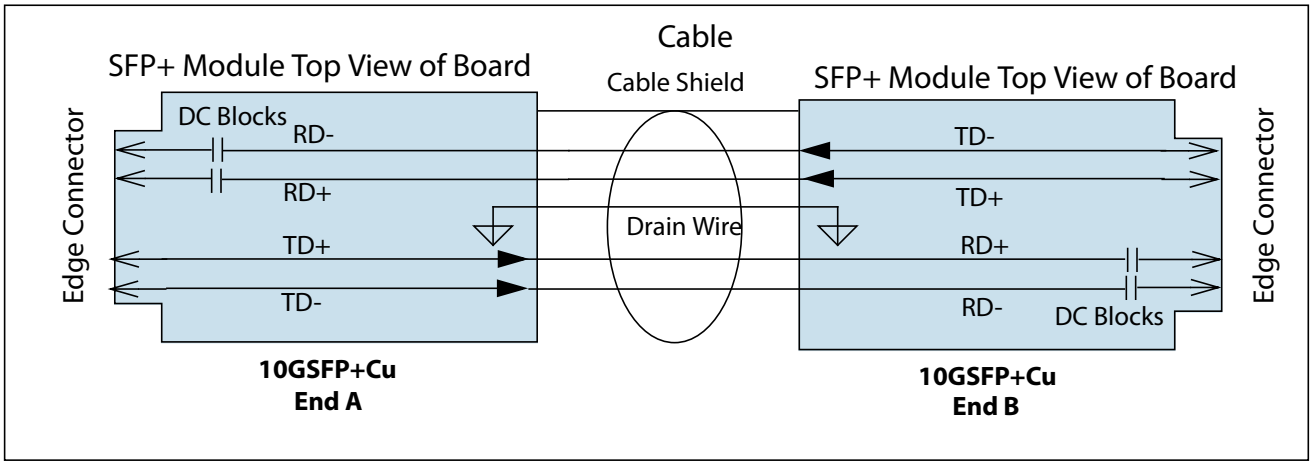


Figure 59 10GSFP+Cu Direct Attach Block Diagram

E.2 SFP+ HOST OUTPUT SPECIFICATIONS FOR PASSIVE DIRECT ATTACH CABLES

SFP+ host supporting direct attach cables must meet transmitter output specifications in Table 11 and jitter specifications in Table 12 at reference point B. In addition SFP+ host transmitter must meet the specifications in Table 33.

Table 33 SFP+ Host Transmitter Output Specifications at B for Cu

Parameters- B	*Symbol*	*Conditions*	*Min*	*Target*	*Max*	*Units*
Voltage Modulation Amplitude (p-p)	VMA	See D.7	300			mV
Transmitter Q_{sq}	Q_{sq}	See 1	63.1			
Output AC Common Mode Voltage		See D.15			12.0	mV (RMS)
Host Output TWDPc	TWDPc	See 2, 3			10.7	dBe

1. Q_{sq}= 1/RN if the one level and zero level noises are identical and see D.8.
2. Host electrical output measured with LRM 14 taps FFE and 5 taps DFE Equalizer with PRBS9 for copper direct attach stressor, see Appendix G.
3. The stressor for TWDPc is given in Table 34 and is included in the code in Appendix G.

TWDPc is the host transmitter penalty for copper cable stressor shown in Figure 60 and given in Table 34. Code to calculate TWDPc using this stressor is given in Appendix G.

E.2.1 TRANSMITTER STRESSOR

For TWDPc compliance, a simulated cable response is required. The response is modeled as a set of delta functions with specific amplitudes and delays. The

copper stressor was created from measurements of commonly available direct attach SFP+ cables with the transmitter response de-convolved. The stressor is shown in Figure 60 and the values are listed in Table 34. The sum of all stressor components is normalized to an approximate value of 1.

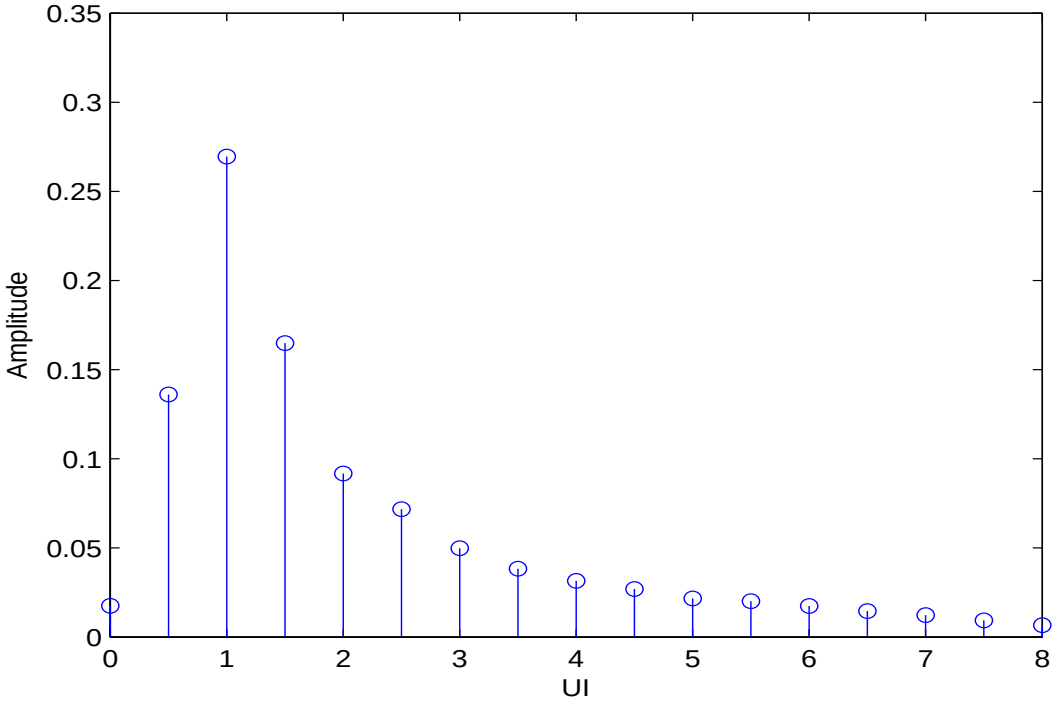


Figure 60 10GSFP+Cu TWDPc Stressor Impulse Response

Table 34 10GSFP+Cu TWDPc Stressor

Delay (UI)	*Delay (ns)*	*Amplitude*	*Delay (UI)*	*Delay (ns)*	*Amplitude*
0	0	0.0175	4.5	0.43637	0.0270
0.5	0.04849	0.1360	5.0	0.48485	0.0216
1	0.09697	0.2695	5.5	0.53334	0.0202
1.5	0.14546	0.1649	6.0	0.58182	0.0174
2	0.19394	0.0917	6.5	0.63031	0.0146
2.5	0.24243	0.0717	7.0	0.67879	0.0123
3.0	0.29091	0.0498	7.5	0.72728	0.0094
3.5	0.33940	0.0383	8.0	0.77576	0.0066
4.0	0.38788	0.0315			

E.3 SFP+ HOST RECEIVER SUPPORTING 10GSFP+CU INPUT COMPLIANCE TEST SIGNAL CALIBRATED AT C"

A host that is to support the direct attach copper option is to meet the required 1×10^{-12} BER when tested with the stressed signal described in Section E.3.1 in addition to the requirements of section 3.5.2 relating to a host receiver supporting linear module.

E.3.1 COPPER HOST RECEIVER SPECIFICATIONS

The SFP+ host receiver stress generator is described by a set of tapped delay lines described in E.3.2, a suitable length of copper cable is expected to generate the stressor described here. The stress generator must meet the target WDPc (Waveform Distortion Penalty for copper) as given in Table 35. The stressor generator shall implement the noise model as captured in Figure 61 using the parameters given in Table 35. The noise model contains two noise sources: Q_{sq} noise which is relative to the transmitter signal level and shaped by the channel response and N_o fixed noise (modeling cable NEXT) added post channel. The added noise sources Q_{sq} and N_o are white and Gaussian in this test.

The sensitivity test shall be made with the minimum VMA and the overload test shall be made with the maximum p-p voltage as given in Table 35.

Table 35 10GSFP+ Host receiver input stress Generator at C"

Parameters- C"	*Symbol*	*Conditions*	*Min*	*Target*	*Max*	*Units*
Waveform Distortion Penalty of the ISI Generator	WDPc	See 1, 2		9.3		dBe
Transmitter Q_{sq}	Q_{sq}	See 4, 5		63.1		
Post channel fixed noise source	N_o	See 3		2.14		mV (RMS)
Differential Voltage Modulation Amplitude	VMA	See 4, D.7		180		mV
Differential Peak-Peak Voltage Overload				700		mV
Input AC Common Mode Voltage		See 6, D.15.2			13.5	mV(RMS)

1. Copper stressor as defined in Table 36. WDPc is measured with reference receiver with 14 FFE taps and with 5 DFE taps, see Appendix G.
2. WDPc for the stress is smaller than the transmitter TWDPc due to the VMA loss in the host stressor.
3. N_o is the RMS voltage measured over one symbol period at the output of the MCB in a 12 GHz bandwidth. The source for Q_{sq} should be disabled during this calibration.
4. Square pattern with eight ONEs and eight ZEROs.
5. Q_{sq}= 1/RN if the one level and zero level noises are identical and see D.8. Q_{sq} is calibrated at the output of the MCB in a 12 GHz bandwidth with the ISI of the channel model in Figure 61 disabled. The source for N_o should be disabled during this calibration.
6. AC common mode target value is achieved by adjusting relative delay of the P and N signals.

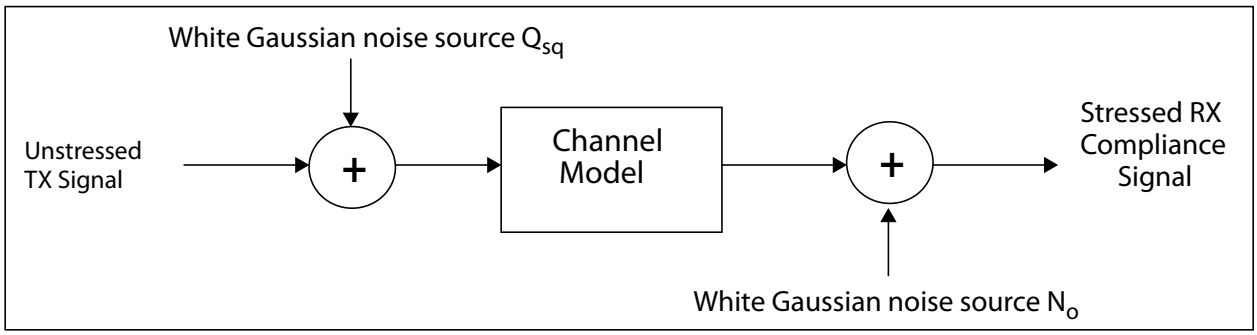


Figure 61 Block Diagram of Copper Stressor Noise Model

E.3.2 COPPER HOST STRESS GENERATOR 1 UI PULSE RESPONSE

Copper host stressor was created from measurements of commonly available direct attach SFP+ cables. The response of the copper host stress generator 1 UI pulse response is shown in Figure 62 and the pulse response values are listed in Table 36.

A suitable length of copper cable is an acceptable substitute to the stressor of Table 36 provided it has the same WDPc. The RMS fit between the tabulated pulse response in Table 36 and the measured isolated pulse response should be minimized to get the target WDPc values as listed in Table 35.

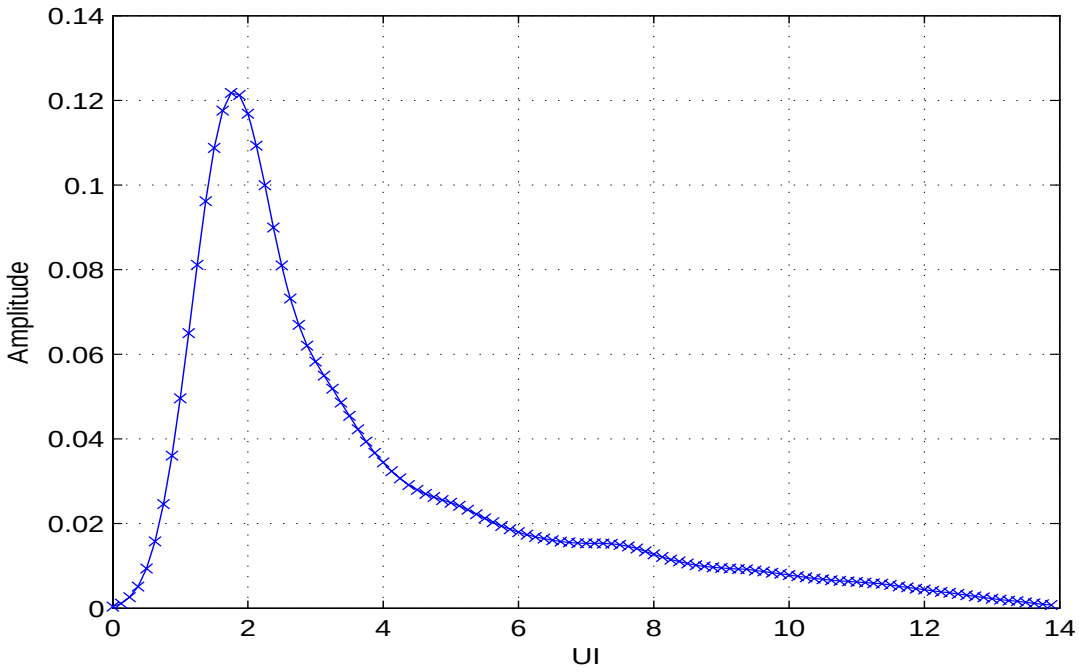


Figure 62 Stress Generator 1UI Pulse Response with 8x over-sampling

Table 36 Stress Generator 1 UI Pulse Response with 8x Over-sampling

Delay (UI)	*Delay (ns)*	*Amplitude*	*Delay (UI)*	*Delay (ns)*	*Amplitude*	*Delay (UI)*	*Delay (ns)*	*Amplitude*
0	0.0000	0.0004	4.625	0.4485	0.0270	9.25	0.8970	0.0093
0.125	0.0121	0.0011	4.75	0.4606	0.0263	9.375	0.9091	0.0092
0.25	0.0242	0.0026	4.875	0.4727	0.0256	9.5	0.9212	0.0089
0.375	0.0364	0.0051	5	0.4848	0.0249	9.625	0.9333	0.0087
0.5	0.0485	0.0094	5.125	0.4970	0.0242	9.75	0.9455	0.0084
0.625	0.0606	0.0158	5.25	0.5091	0.0233	9.875	0.9576	0.0082
0.75	0.0727	0.0246	5.375	0.5212	0.0222	10	0.9697	0.0079
0.875	0.0848	0.0360	5.5	0.5333	0.0212	10.125	0.9818	0.0076
1	0.0970	0.0496	5.625	0.5455	0.0204	10.25	0.9939	0.0073
1.125	0.1091	0.0650	5.75	0.5576	0.0194	10.375	1.0061	0.0070
1.25	0.1212	0.0811	5.875	0.5697	0.0187	10.5	1.0182	0.0068
1.375	0.1333	0.0962	6	0.5818	0.0180	10.6250	1.0303	0.0066
1.5	0.1455	0.1088	6.125	0.5939	0.0174	10.7500	1.0424	0.0065
1.625	0.1576	0.1176	6.25	0.6061	0.0169	10.8750	1.0545	0.0063
1.75	0.1697	0.1218	6.375	0.6182	0.0165	11.0000	1.0667	0.0062
1.875	0.1818	0.1213	6.5	0.6303	0.0161	11.1250	1.0788	0.0061
2	0.1939	0.1169	6.625	0.6424	0.0157	11.2500	1.0909	0.0059
2.125	0.2061	0.1093	6.75	0.6545	0.0155	11.3750	1.1030	0.0058
2.25	0.2182	0.1000	6.875	0.6667	0.0154	11.5000	1.1152	0.0055
2.375	0.2303	0.0899	7	0.6788	0.0153	11.6250	1.1273	0.0052
2.5	0.2424	0.0810	7.125	0.6909	0.0153	11.7500	1.1394	0.0049
2.625	0.2545	0.0732	7.25	0.7030	0.0153	11.8750	1.1515	0.0046
2.75	0.2667	0.0670	7.375	0.7152	0.0153	12.0000	1.1636	0.0044
2.875	0.2788	0.0620	7.5	0.7273	0.0150	12.1250	1.1758	0.0041
3	0.2909	0.0583	7.625	0.7394	0.0146	12.2500	1.1879	0.0039
3.125	0.3030	0.0549	7.75	0.7515	0.0141	12.3750	1.2000	0.0037
3.25	0.3152	0.0519	7.875	0.7636	0.0134	12.5000	1.2121	0.0034
3.375	0.3273	0.0486	8	0.7758	0.0128	12.6250	1.2242	0.0030
3.5	0.3394	0.0455	8.125	0.7879	0.0121	12.7500	1.2364	0.0028
3.625	0.3515	0.0423	8.25	0.8000	0.0115	12.8750	1.2485	0.0026
3.75	0.3636	0.0394	8.375	0.8121	0.0110	13.0000	1.2606	0.0022
3.875	0.3758	0.0367	8.5	0.8242	0.0106	13.1250	1.2727	0.0020
4	0.3879	0.0345	8.625	0.8364	0.0101	13.2500	1.2848	0.0018
4.125	0.4000	0.0324	8.75	0.8485	0.0099	13.3750	1.2970	0.0016
4.25	0.4121	0.0307	8.875	0.8606	0.0097	13.5000	1.3091	0.0014
4.375	0.4242	0.0291	9	0.8727	0.0095	13.6250	1.3212	0.0012
4.5	0.4364	0.0280	9.125	0.8848	0.0094	13.7500	1.3333	0.0009
						13.8750	1.3455	0.0008

E.4 SFP+ Passive Direct Attach cable Assembly Specifications

Passive direct attach cables are tested with a pair of Module Compliance Boards at compliance point B′ and C′. SFP+ passive cable assemblies need to meet specification in Table 37.

VCR, VMA, Vcm, and dWDP may be derived using frequency based methodologies that yield equivalent results e.g., utilizing frequency dependent crosstalk and insertion loss transfer functions with transmitter behavioral models.

Table 37 10GSFP+Cu Cable Assembly Specifications at B′ and C′

Parameter - C' (Cable Output)	*Symbol*	*Conditions*	*Min*	*Target*	*Max*	*Units*
Single Ended Input and Output Voltage Tolerance			-0.3		4.0	V
Output AC Common Mode Voltage	Vcm	See 1			13.5	mV (RMS)
Difference Waveform Distortion Penalty	dWDPc	See 2, 9, E.4.1, E.4.2 and D.14.2			6.75	dBe
VMA Loss	L	See 3, 9, D.7, E.4.4			4.4	dBe
VMA Loss to Crosstalk Ratio	VCR	See 1, D.7, E.4.1, E.4.4	32.5			dB
Differential Output/Input Reflection Coefficient [4]	SDDxx	0.01-4.1 GHz			See 5	dB
		4.1-11.1 GHz			See 6	dB
Common Mode Output/Input Reflection Coefficient [7]	SCCxx	0.01-2.5 GHz			See 10	dB
		2.5-11.1 GHz			-3	dB
Parameter - B" (Input Test Conditions)	*Symbol*	*Conditions*	*Min*	*Target*	*Max*	*Units*
Input AC Common Mode Voltage	Vcm	See 1, D.15.2		12		mV (RMS)
Signal Rise and fall time Time	Tr/tf	See D.6		34		ps
Crosstalk Source Rise/Fall time (20% to 80%)	Tr, Tf	See D.6		34		ps
Crosstalk Source Amplitude Differential (p-p)				700		mV
WDP_i		See 8		2.4		dBe

1. When input common mode voltage is 12.0 mV RMS and when input rise and fall times are 34ps and the amplitude is the max amplitude allowed by Table 12.
2. Defined with reference receiver with 14 T/2 spaced FFE taps and 5 T spaced DFE taps, see Appendix G.
3. VMA loss is the ratio of VMA measured at input and output, respectively.
4. Reference differential impedance is 100 Ω. The dB value listed here are the same as dBe.
5. Reflection Coefficient given by equation SDDxx(dB)= -12 + 2 × SQRT(f), with f in GHz.
6. Reflection Coefficient given by equation SDDxx(dB)= -6.3 + 13 × log10(f/5.5), with f in GHz.
7. Common mode reference impedance is 25 Ω. The dB value listed here are the same as dBe
8. Adjust DDJ and/or DDPWS by adjusting pre-emphasis until the target WDP_i is achieved.
9. With input test condition given by parameters B" given in this table.
10. Reflection coefficient given by equation SCCxx(dB) < -7 + 1.6 ×f, with f in GHz.

E.4.1 SFP+ DIRECT ATTACH CABLE TEST SETUP

Direct attach cable testing methodology is based on the SFP+ test methodology as defined in section 3.3. The cable is measured through a pair of Module Compliance Boards as shown in Figure 63. This diagram shows the block diagram for testing NEXT on cable A end and for measuring WDP on path 1. To measure NEXT on B end and WDP on path 2 the cable end A and B are reversed. The Compliance Signal Generator is described in Figure 49.

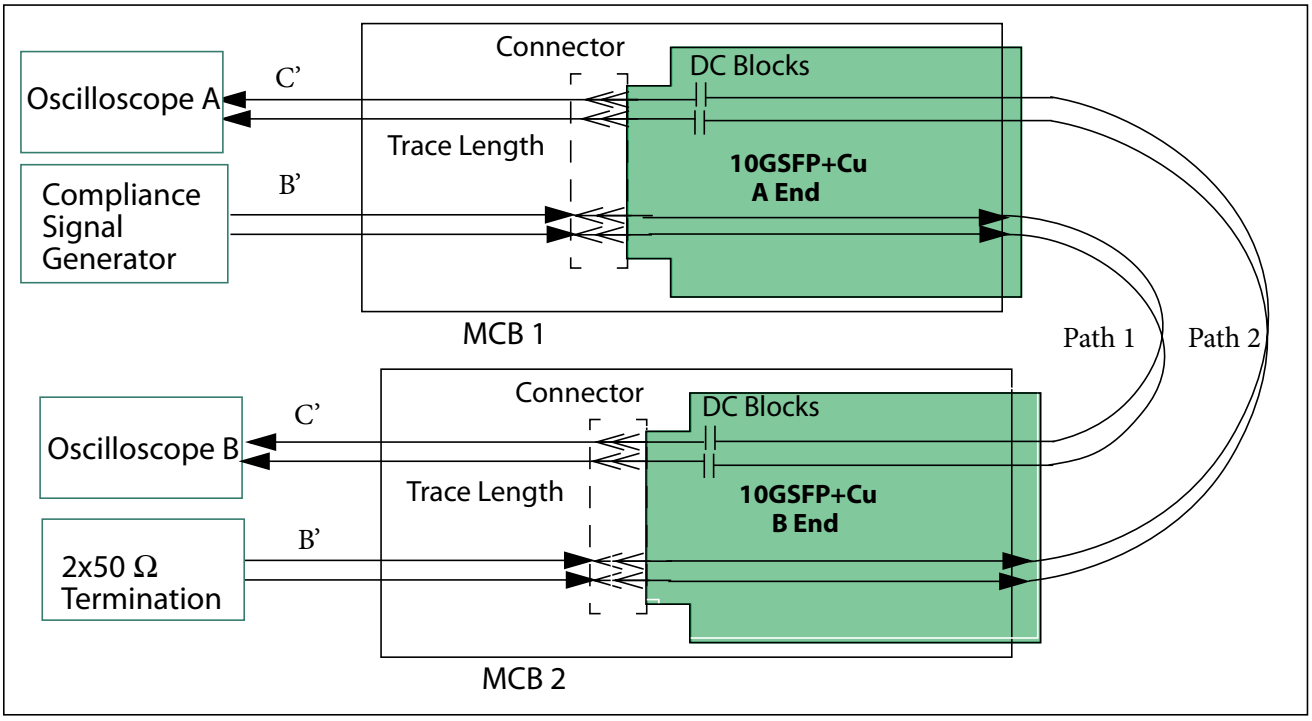


Figure 63 10GSFP+ Cable Test Setup

WDP_i and WDP_0 in Figure 64 use the WDP method defined in D.9. WDPi for copper is measured by plugging Host Compliance Board into the Module Compliance Board 1 and then meeting the target WDPi as listed in Table 37. WDP_0 is measured by plugging one end of the cable in to Module Compliance Board 1 and the other end in to the Module Compliance Board 2. B" Stress Generator can be the test system described in D.10.

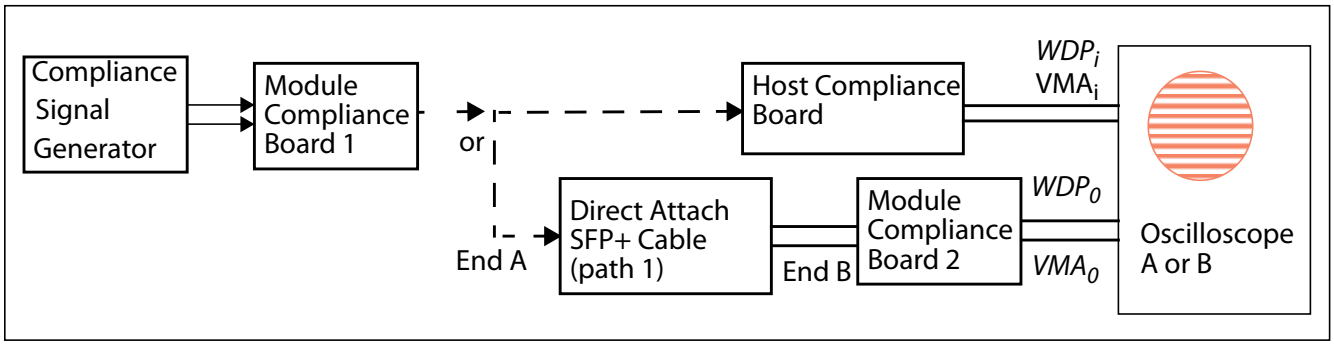


Figure 64 10GSFP+Cu Cable NEXT dWDP Test Setup

E.4.2 CABLE DWDP TEST PROCEDURE

The measurement procedure for dWDP is described below:

- The compliance signal generator is set to the PRBS9.
- To improve measurement accuracy, uncorrelated jitter and noise should be reduced.
- Averaging should be used to further reduce instrumentation and measurement noise so their effect on the results are negligible.
- To calibrate WDP_i per Table 37, refer to Figure 64. Plug a Host Compliance Board into the Module Compliance Board connected to the pattern generator. Adjust input rise and fall times to the target value as in Table 37. Adjust DDJ and DDPWS to obtain WDP_i given by Table 37. Varying pre-emphasis as described in D.10 is an acceptable method.
- Unplug the Host Compliance Board and connect the cable assembly to the Module Compliance Board as shown in Figure 64. Measure WDP_o.
- $dWDP = WDP_o - WDP_i$.

E.4.3 CABLE NEXT MEASUREMENT PROCEDURE

Cable NEXT is measured based on the following procedure using the test setup shown in Figure 63:

- The Compliance Signal Generator should be calibrated via a Host Compliance Board inserted into the Module Compliance Board. The output of the Host Compliance Board is point B".
- The Compliance Signal Generator amplitude and rise and fall times at B" are calibrated to the crosstalk target values as defined in Table 37.

- The Compliance Signal Generator DDJ and DDPWS at B" should meet or be less than the target specified in Table 17.
- The pattern for the Compliance Signal Generator is PRBS31.
- Module Compliance Board B outputs and inputs are terminated in 50 Ω.
- NEXT is the RMS voltage measured by Oscilloscope A in a bandwidth of 12 GHz. Oscilloscope A should be free running (not triggered).
- The inherent Oscilloscope noise may be corrected by the RSS of Gaussian noise from the measured NEXT result.
- The far end Module Compliance Board outputs and input are terminated in to 50 Ω.
- This measurement is then repeated for the other cable end.

E.4.4 VMA TO CROSSTALK RATIO (VCR)

Cable VMA loss (L) for cable path 1 can be measured using the test setup shown in Figure 64. VMA loss (L) for cable path 2 is measured by reversing cable end A with B.

$$L(dBe) = 20\log\left(\frac{VMA_i}{VMA_o}\right)$$

Where VMA_i is the measured VMA at B" and VMA_o is measured at C'.

VMA/2 to crosstalk ratio (VCR) is the ratio of the transmitter minimum VMA at B" Table 33 divided by the cable NEXT which already incorporates reflective FEXT. The factor 0.3 in the VCR [1] equation accounts for SFP+ finite host return loss.

$$VCR(dBe) = VNR - L - K - 20\log 10(1 + C)$$

where

$$C = 0.3 \times 10^{\left(-\frac{2L}{20}\right)}$$

1. VCR equation can be as the following 20*log10(VMAMIN/(2*NEXT*(1+C)))-L.

$$VNR = 20\log 10\left[\frac{\frac{(NEXTaggressorVMA)}{2}}{NEXT}\right]$$

.

$$K(dBe) = 20\log 10\left(\frac{VMAmax}{VMAmin}\right) = 20\log 10\left(\frac{700}{300}\right) = 7.36$$

The procedure to measure NEXT is described in E.4.3.

Appendix F 1.25 GBd Operation Support (Optional)

F.1 Introduction

SFP+ host may be designed to operate at 1.25 GBd Ethernet rate using the classic SFP modules based on INF-8074i. Although IEEE Std 802.3, clauses 38 and 59 (1000BASE-SX, 1000BASE-LX and 1000BASE-LX10 PMDs) do not define the electrical levels for the module, INF-8074i specifies those levels which are reproduced here for reference Table 38.

Host transmitter output levels B and host receiver input tolerance levels at C are respectively given by Table 39 and Table 40 for SFP+ host operating at 1.25 GBd. Note: levels specified here may not be fully compliant with all classic SFP modules, but are expected to include a large percentage of existing 1.25 GBd classic modules. In order to be fully compliant to all classic SFP modules the max host receiver input tolerance level has to be 2000 mV, however this is not considered practical for modern 10Gb/s SerDes.

Table 38 INF-8074i Voltage Levels for Reference Only

Parameters - B'	*Symbol*	*Conditions*	*Min*	*Max*	*Units*
SFP Module Input Differential at B'	Vin	See 1	500	2400	mV (p-p)
SFP Module Output Differential at C'	Vout		370	2000	mV (p-p)
1. INF-8074i recommends value between 500-1200 mV differential p-p be used for best EMI performance.					

F.2 SFP+ Host Operation Guideline for Supporting Classic SFP

SFP+ host must support TP1 and TP4 jitter specifications per IEEE CL 38.5 and CL 59.6 at point B and C respectively.

SFP+ host output pre-emphasis level may need to be adjusted for optimum output eye diagram for 1.25 GBd operation.

The module maximum input at B is given in Table 39 and is the same value as in INF-8074i. However, to provide compatibility with SFP+ hosts the module maximum output in Table 40 is much lower than the value specified in INF-8074i. An SFP module meeting the specifications of Table 41 will interoperate with SFP+ hosts.

Table 39 SFP+ Host Transmitter Requirements to Support 1.25 GBd Mode

Parameters - B	*Symbol*	*Conditions*	*Min*	*Units*
Host Output VMA Differential	Vout		500	mV
Parameters- B	***Symbol***	***Conditions***	***Value***	***Units***
Eye Mask	Y1	See D.2 and Figure 65	150	mV
Eye Mask	Y2		500	mV

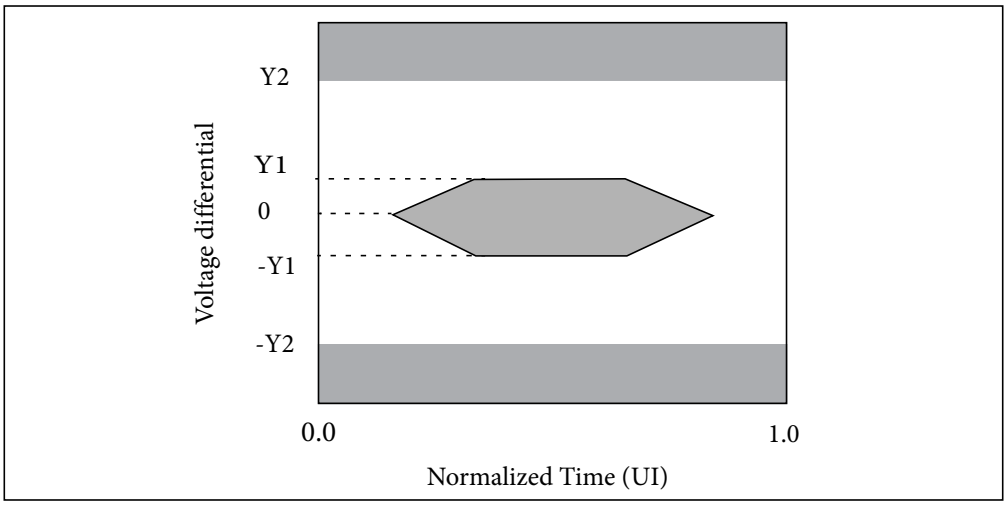


Figure 65 SFP+ Host Transmitter Output Mask for 1.25 GBd Operation

Table 40 SFP+ Host Receiver Requirements to Support 1.25 GBd Mode

Parameters - C	*Symbol*	*Conditions*	*Min*	*Units*
Host Input VMA Differential	Vin		370	mV
Parameters- C	***Symbol***	***Conditions***	***Value***	***Units***
Eye Mask	Y1	See D.2 and Figure 66	125	mV
Eye Mask	Y2		600	mV

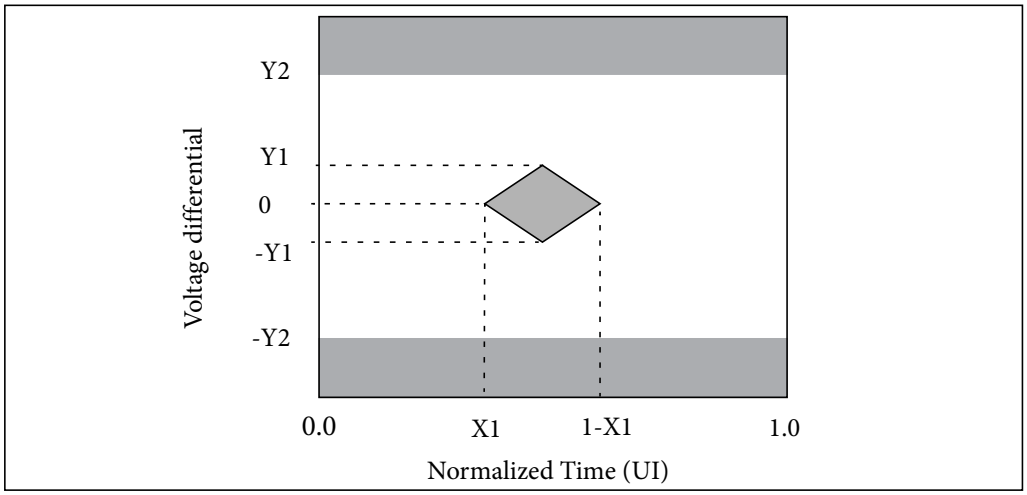


Figure 66 SFP+ Host Receiver Input Mask for 1.25 GBd Operation

.

Table 41 SFP Module input and output ranges that can be supported by the SFP+ Host

Parameters - Module	*Symbol*	*Conditions*	*Min*	*Max*	*Units*
SFP Module Input at B'	Vin		500 [1]	2400 [1]	mV (p-p)
SFP Module Output at C'	Vout		370 [1]	1200	mV (p-p)
1. Values are identical to the SFP classic INF-8074i specification.					

APPENDIX G MATLAB CODE FOR TWDP

```
%%%%%%%%%%%%%%%%% MATLAB (R) Code for xWDP Computation %%%%%%%%%%%%%%%%%

%% SFF-8431 TWDP Code - PAlloc for Cu is now 14.0 dBe.
%% Based on original TWDP methodology described in IEEE Std 802.3aq(TM)-2006

%% Reference: N. L. Swenson, P. Voois, T. Lindsay, and S. Zeng, "Standards
%% compliance testing of optical transmitters using a software-based equalizing
%% reference receiver", paper NWC3, Optical Fiber Communication Conference and
%% Exposition and The National Fiber Optic Engineers Conference on CD-ROM
%% (Optical Society of America, Washin[gton, DC), Feb. 2007.

function [xWDP,MeasuredxMA]=SFF8431xWDP(WaveformFile,EqNf,EqNb,SymbolRate,Usage)
%% Example calling syntax:
%%   [xWDP,MeasuredxMA]=SFF8431xWDP('wavefile.txt',14,5,10.3125,'Optical_WDP')

%% The fields in the example given above should be replaced by the actual values
%% being used. WaveformFile should be the actual path\filename for each waveform
%% tested. The waveform consists of exactly N samples per unit interval T, where
%% N is the oversampling rate. The waveform must be circularly shifted to align
%% with the data sequence. The file format for the measured waveform is ASCII
%% with a single column of chronological numerical samples, in signal level,
%% with no headers or footers.
%% EqNf is the # of T/2-spaced feedforward equalizer taps; EqNb is the # of
%% T-spaced feedback equalizer taps.
%% SymbolRate is in gigabaud.
%% Options for Usage are 'Optical_WDP', 'Copper_WDP', and 'Copper_TWDP'.
%% 'Optical_WDP' is used in support of Chapter 3 for
%%    measuring WDPi at the output of an optical TP3 tester,
%%    measuring WDPo at C' of a linear optical module receiver, and
%%    calibrating WDP at C'' for testing a host that supports linear optical modules.
%% 'Copper_WDP' is used in support of Annex E for
%%    measuring WDPi and calibrating WDP at B'' for testing a copper cable assembly,
%%    measuring WDPo at C' of a copper cable assembly (C), and
%%    calibrating WDP at C'' for testing a host that supports copper cable assemblies
%% 'Copper_TWDP' is used for measuring TWDP at B of a host that supports copper
%% cable assemblies.

%% Transmit data file: The transmit data sequence is the 511 bit PRBS9 TWDP test
%% patterns defined in Table 686. The file format is ASCII with a single column
%% of chronological ones and zeros with no headers or footers.
TxDataFile = 'prbs9_950.txt';

%% Program constants %%
OverSampleRate = 16; % Oversampling rate, must be even
SymbolPeriod = 1/SymbolRate; % Symbol period is in ns
Q0 = 7.03; % BER = 10^(-12)

%% Load input waveform and data sequence, generate filter and other matrices
yout0 = load(WaveformFile);
XmitData = load(TxDataFile);
PtrnLength = length(XmitData);
TotLen = PtrnLength*OverSampleRate;
Fgrid = [-TotLen/2:TotLen/2-1].'/(PtrnLength*SymbolPeriod);
%% Compute response of 7.5 GHz 4th order Butterworth antialiasing filter
```

```
a = [1 123.1407 7581.811 273453.7 4931335]; % Denominator polynomial
b = 4931335; % Numerator for frequency response
ExpArg = -j*2*pi*Fgrid;
H_r = b./polyval(a,-ExpArg);

%% Get usage parameters for the application
[H_chan,Delays,PAlloc,dBscale] = GetParams(Usage,ExpArg);
N0 = SymbolPeriod/2 / (Q0 * 10^(PAlloc/dBscale))^2;

%% Set search range for equalizer delay, specified in symbol periods. Lower end
%% of range is minimum channel delay. Upper end of range is the sum of the
%% lengths of the FFE and channel. Round up and add 5 to account for the
%% antialiasing filter.
EqDelMin = floor(min(Delays)/SymbolPeriod);
EqDelMax = ceil(EqNf/2 + max(Delays)/SymbolPeriod);

ONE=ones(PtrnLength,1);
%% Normalize the received xMA (OMA or VMA) to 1. Estimate the xMA of the captured
%% waveform by using a linear fit to estimate a pulse response, synthesize a
%% square wave, and calculate the xMA of the synthesized square wave per IEEE
%% 802.3, clause 52.9.5.
ant=4; mem=40; % Anticipation and memory parameters for linear fit
X=zeros(ant+mem+1,PtrnLength); % Size data matrix for linear fit
Y=zeros(OverSampleRate,PtrnLength); % Size observation matrix for linear fit
for ind=1:ant+mem+1
    X(ind,:)=circshift(XmitData,ind-ant-1)'; % Wrap appropriately for lin fit
end
X=[X;ones(1,PtrnLength)]; % The all-ones row is included to compute the bias
for ind=1:OverSampleRate
    Y(ind,:)=yout0([0:PtrnLength-1]*OverSampleRate+ind)'; % Each column is 1 bit
end
Qmat=Y*X'*(X*X')^(-1); % Coefficient matrix resulting from linear fit. Each
%% column (except the last) is one bit period of the pulse response. The last
%% column is the bias.
SqWvPer=16; % Even number; sets the period of the sq wave used to compute xMA
SqWv=[zeros(SqWvPer/2,1);ones(SqWvPer/2,1)]; % One period of sq wave (column)
X=zeros(ant+mem+1,SqWvPer); % Size data matrix for synthesis
for ind=1:ant+mem+1
    X(ind,:)=circshift(SqWv,ind-ant-1)'; % Wrap appropriately for synthesis
end
X=[X;ones(1,SqWvPer)]; % Include the bias
Y=Qmat*X;Y=Y(:); % Synthesize the modulated square wave, put into one column
Y=AlignY(Y,SqWvPer,OverSampleRate);
avgpos=[0.4*SqWvPer/2*OverSampleRate:0.6*SqWvPer/2*OverSampleRate];
ZeroLevel=mean(Y(round(avgpos),:)); % Average over middle 20% of "zero" run
% Average over middle 20% of "one" run, compute xMA
MeasuredxMA=mean(Y(round(SqWvPer/2*OverSampleRate+avgpos),:))-ZeroLevel;
%% Subtract zero level and normalize xMA
yout0 = (yout0-ZeroLevel)/MeasuredxMA;

%% Compute the noise autocorrelation sequence at the output of the front-end
%% antialiasing filter and rate-2/T sampler.
Snn = N0/2 * fftshift(abs(H_r).^2) * 1/SymbolPeriod * OverSampleRate;
Rnn = real(ifft(Snn));
```

```
Corr = Rnn(1:OverSampleRate/2:end);
C = toeplitz(Corr(1:EqNf));

%% Compute the minimum slicer MSE and corresponding xWDP
X = toeplitz(XmitData, [XmitData(1); XmitData(end:-1:end+1-EqNb)]);
Xtil = toeplitz(circshift(XmitData,EqDelMin), ...
    XmitData(mod(-EqDelMin:-1:-(EqDelMax+EqNb),PtrnLength)+1));
Rxx = X'*X; % Used in MSE calculation
%% Propagate the waveform through channel.
yout = real(ifft(fft(yout0) .* fftshift(H_chan)));
%% Process signal through front-end antialiasing filter %%%%%%%%%%%%%%%%%%
yout = real(ifft(fft(yout) .* fftshift(H_r)));
%% Compute MMSE-DFE %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%% The MMSE-DFE filter coefficients computed below minimize mean-squared error
%% at the slicer input. The derivation follows from the fact that the slicer
%% input over the period of the data sequence can be expressed as Z = (R+N)*W -
%% X*[0 B]', where R and N are Toeplitz matrices constructed from the signal and
%% noise components, respectively, at the sampled output of the antialiasing
%% filter, W is the feedforward filter, X is a Toeplitz matrix constructed from
%% the input data sequence, and B is the feedback filter. The computed W and B
%% minimize the mean square error between the input to the slicer and the
%% transmitted sequence due to residual ISI and Gaussian noise. Minimize MSE
%% over 2/T sampling phase and FFE delay and determine BER.
MseOpt = Inf;
for jj= [0:OverSampleRate-1]-OverSampleRate/2 % sampling phase
    %% Sample at rate 2/T with new phase (wrap around as required)
    yout_2overT = yout(mod([1:OverSampleRate/2:TotLen]+jj-1,TotLen)+1);
    Rout = toeplitz(yout_2overT, [yout_2overT(1); yout_2overT(end:-1:end-EqNf+2)]);
    R = Rout(1:2:end, :);
    RINV = inv([R'*R+PtrnLength*C R'*ONE;ONE'*R PtrnLength]);
    R=[R ONE]; % Add all-ones column to compute optimal offset
    Rxr = Xtil'*R; Px_r = Rxr*RINV*Rxr';
    %% Minimize MSE over equalizer delay
    for kk = 1:EqDelMax-EqDelMin+1
        SubRange = [kk:kk+EqNb];
        SubRange = mod(SubRange-1,PtrnLength)+1;
        P = Rxx - Px_r(SubRange,SubRange);
        P00 = P(1,1); P01 = P(1,2:end); P11 = P(2:end,2:end);
        Mse = P00 - P01*inv(P11)*P01';
        if (Mse<MseOpt)
            MseOpt = Mse;
            B = -inv(P11)*P01'; % Feedback filter
            XSel = Xtil(:,SubRange);
            W = RINV*R'*XSel*[1;B]; % Feedforward filter
            Z = R*W - XSel*[0;B]; % Input to slicer
            %% Compute BER using semi-analytic method %%%%%%%%%%%%%%%%%%
            MseGaussian = W(1:end-1)'*C*W(1:end-1);
            Ber = mean(0.5*erfc((abs(Z-0.5)/sqrt(MseGaussian))/sqrt(2)));
        end
    end
end

%% Compute equivalent SNR %%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%% This function computes the inverse of the Gaussian error probability
```

```
%% function. The built-in function erfcinv() is not sensitive enough for low
%% probability of error cases.
if Ber>10^(-12) Q = sqrt(2)*erfinv(1-2*Ber);
elseif Ber>10^(-323) Q = 2.1143*(-1.0658-log10(Ber)).^0.5024;
else Q = inf;
end

%% Compute penalty %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
RefSNR = dBscale * log10(Q0) + PAlloc;
xWDP = RefSNR-dBscale*log10(Q);
%% End of main function

%% GetParams subFunction
function [H_chan,Delays,PAlloc,dBscale] = GetParams(Usage,ExpArg);
    switch Usage
        case 'Optical_WDP' % Identity channel for optical
            Delays = 0;
            H_chan = 1;
            PAlloc = 6.5; % Total allocated dispersion penalty (dBo)
            dBscale = 10;
        case 'Copper_WDP' % Identity channel for copper
            Delays = 0;
            H_chan = 1;
            PAlloc = 14.0;% Total allocated dispersion penalty (dBe)
            dBscale = 20;
        case 'Copper_TWDP' % Cu TWDP stressor
            ChanResp = [...
                .0 .04849 .09697 .14546 .19394 .24243 .29091 .33940 .38788, ...
                .43637 .48485 .53334 .58182 .63031 .67879 .72728 .77576;
                .0175 .136 .2695 .1649 .0917 .0717 .0498 .0383 .0315, ...
                .027 .0216 .0202 .0174 .0146 .0123 .0094 .0066];
            Delays = ChanResp(1,:);
            PCoefs = ChanResp(2,:)';
            H_chan = exp(ExpArg*Delays)*PCoefs/sum(PCoefs); %With normalization
            PAlloc = 14.0;% Total allocated dispersion penalty (dBe)
            dBscale = 20;
    end
%% End of GetParams function

%% AlignY subFunction
function Y = AlignY(Y0,SqWvPer,OverSampleRate)
    % Aligns the mid crossing of the xMA square waveform to its ideal position.

    Y = Y0-mean(Y0); % AC-couple so crossings are at 0.

    % Look only for the crossing in the middle by ignoring any within ~2 UI from
    % its beginning. Due to possible misalignment of the captured waveform, this
    % is the only crossing that is certain.
    x = find(sign(Y(2*OverSampleRate:end-1))~= ...
        sign(Y(2*OverSampleRate+1:end)),1)+2*OverSampleRate-1;
    % Find a more exact crossing point.
    xinterp = interp1([Y(x),Y(x+1)],[x,x+1],0);

    % Shift to create the aligned square waveform
    Y = circshift(Y0,SqWvPer/2*OverSampleRate-x); % Coarse shift.
    X = [1:length(Y)].'; Y = interp1(X,Y,X+xinterp-x,'spline'); % Fine shift.
%% End of AlignY function
```

SFP+ 10G LR (10km)

10Gb/s 1310nm Single-mode SFP+ Transceiver

PRODUCT FEATURES

- 10Gb/s serial optical interface compliant to 802.3ae 10GBASE-LR
- 1310nm DFB transmitter, PIN photo-detector
- Up to 10km transmission on SMF
- Duplex LC connector
- Metal enclosure, for lower EMI
- Single 3.3V power supply, Power dissipation < 1.2w
- Electrical interface compliant to SFF-8431
- Specifications for 10 Gigabit “SFP+”
- 2-wire interface for management
- Specifications compliant with SFF 8472
- Case operating temperature range: -5°C to 85°C

APPLICATIONS

- 10GBASE-LR/LW 10G Ethernet

STANDARD

- Compliant with SFF-8472 SFP+ MSA.
- Compliant to SFP+ SFF-8431 and SFF-8432.
- Compliant to 802.3ae 10GBASE-LR.
- RoHS Compliant.

Ⅰ Absolute Maximum Ratings

Parameter	Symbol	Min	Typ	Max	Unit	Ref.
Maximum Supply Voltage	Vcc	-0.5		4.7	V	
Storage Temperature	TS	-40		85	°C	
Case Operating Temperature	Tcase	-5		85	°C	

Ⅱ Electrical Characteristics (Tcase = -5 to 85 ℃, VCC = 3.14 to 3.46 Volts)

Parameter	Symbol	Min	Typ	Max	Unit	Ref.
Supply Voltage	Vcc	3.14	3.3	3.46	V	
Supply Current	Icc			360	mA	
Transmitter						
Input differential impedance	Rin		100		Ω	1
Single ended data input swing	Vin,pp	180		700	mV	
Transmit Disable Voltage	VD	Vcc–1.3		Vcc	V	
Transmit Enable Voltage	VEN	Vee		Vee+ 0.8	V	2
Transmit Disable Assert Time				10	us	
Receiver						
Differential data output swing	Vout,pp	300		850	mV	3
Data output rise time	tr	30			ps	4
Data output fall time	tf	30			ps	4
LOS Fault	VLOS fault	Vcc–1.3		VccHOST	V	5
LOS Normal	VLOS norm	Vee		Vee+0.8	V	5
Power Supply Rejection	PSR	100			mVpp	6

Notes:

1. Connected directly to TX data input pins. AC coupled thereafter.
2. Or open circuit.
3. Into 100 ohms differential termination.
4. 20 – 80 %.
5. Loss Of Signal is LVTTL. Logic 0 indicates normal operation; logic 1 indicates no signal detected.
6. Receiver sensitivity is compliant with power supply sinusoidal modulation of 20 Hz to 1.5 MHz up to specified value applied through the recommended power supply filtering network.

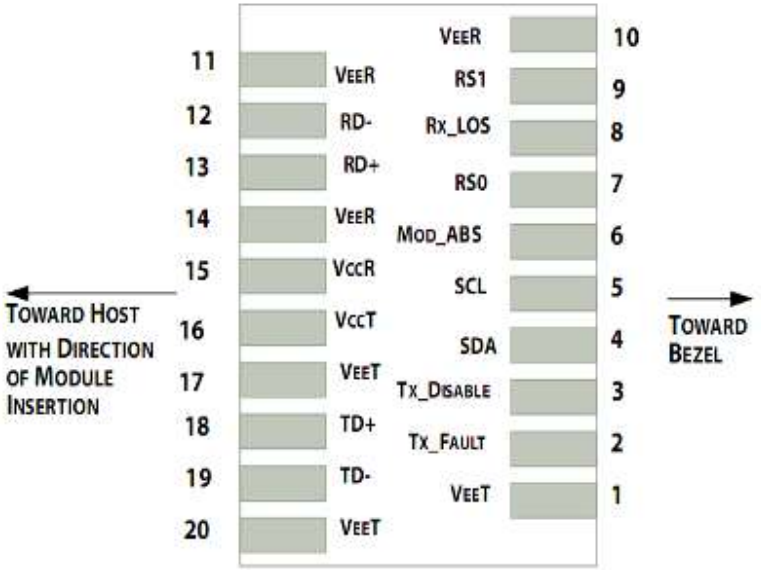
Ⅲ Optical Characteristics (T_{OP} = -5 to 85℃, VCC = 3.14 to 3.46 Volts)

Parameter	Symbol	Min	Typ	Max	Unit	Ref.
Transmitter						
Output Opt. Pwr	POUT	-6		-1	dBm	1
Optical Wavelength	λ	1260	1310	1355	nm	
Spectral Width (-20dB)	σ			1	nm	
Optical Extinction Ratio	ER	3.5			dB	
Transmitter Dispersion Penalty	TDP			3.2	dB	
Side mode Supression ratio	SMSR	30			dB	
RIN	RIN			-128	dB/Hz	
Output Eye Mask	Compliant with IEEE 0802.3ae					
Receiver						
Rx Sensitivity	RSENS			-15	dBm	2
Input Saturation Power (Overload)	Psat	0.5			dBm	
Wavelength Range	λ_C	1270		1610	nm	
LOS De -Assert	LOSD			-17	dBm	
LOS Assert	LOSA	-30			dBm	
LOS Hysteresis		0.5	1.0		dB	

Notes:

1. Class 1 Laser Safety per FDA/CDRH and IEC-825-1 regulations.
2. With worst-case extinction ratio. Measured with a PRBS 2^{31} -1 test pattern, @10.325Gb/s, BER<10^{-12}.

Ⅳ Pin Assignment



Pin out of Connector Block on Host Board

Pin	Symbol	Name/Description	Ref.
1	V_{EET}	Transmitter Ground (Common with Receiver Ground)	1
2	T_{FAULT}	Transmitter Fault.	2
3	T_{DIS}	Transmitter Disable. Laser output disabled on high or open.	3
4	SDA	2-wire Serial Interface Data Line	4
5	SCL	2-wire Serial Interface Clock Line	4
6	MOD_ABS	Module Absent. Grounded within the module	4
7	RS0	Rate Select 0	5
8	LOS	Loss of Signal indication. Logic 0 indicates normal operation.	6
9	RS1	No connection required	1
10	V_{EER}	Receiver Ground (Common with Transmitter Ground)	1
11	V_{EER}	Receiver Ground (Common with Transmitter Ground)	1
12	RD-	Receiver Inverted DATA out. AC Coupled	
13	RD+	Receiver Non-inverted DATA out. AC Coupled	
14	V_{EER}	Receiver Ground (Common with Transmitter Ground)	1
15	V_{CCR}	Receiver Power Supply	
16	V_{CCT}	Transmitter Power Supply	
17	V_{EET}	Transmitter Ground (Common with Receiver Ground)	1
18	TD+	Transmitter Non-Inverted DATA in. AC Coupled.	
19	TD-	Transmitter Inverted DATA in. AC Coupled.	
20	V_{EET}	Transmitter Ground (Common with Receiver Ground)	1

Notes:

1. Circuit ground is internally isolated from chassis ground.
2. T_{FAULT} is an open collector/drain output, which should be pulled up with a 4.7k – 10k Ohms resistor on the host board if intended for use. Pull up voltage should be between 2.0V to Vcc + 0.3V.A high output indicates a transmitter fault caused by either the TX bias current or the TX output power exceeding the preset alarm thresholds. A low output indicates normal operation. In the low state, the output is pulled to <0.8V.
3. Laser output disabled on T_{DIS} >2.0V or open, enabled on T_{DIS} <0.8V.
4. Should be pulled up with 4.7kΩ- 10kΩ host board to a voltage between 2.0V and 3.6V. MOD_ABS pulls line low to indicate module is plugged in.
5. Internally pulled down per SFF-8431 Rev 4.1.
6. LOS is open collector output. It should be pulled up with 4.7kΩ – 10kΩ on host board to a voltage between 2.0V and 3.6V. Logic 0 indicates normal operation; logic 1 indicates loss of signal.

V. Digital Diagnostic Functions

Those transceivers support the 2-wire serial communication protocol as defined in the SFP MSA1.

The standard SFP serial ID provides access to identification information that describes the transceiver's capabilities, standard interfaces, manufacturer, and other information.

Additionally, those S FP+ transceivers provide a unique enhanced digital diagnostic monitoring interface, which allows real-time access to device operating parameters such as transceiver temperature, laser bias current, transmitted optical power, received optical power and transceiver supply voltage. It also defines a sophisticated system of alarm and warning flags, which alerts end-users when particular operating parameters are outside of a factory set normal range.

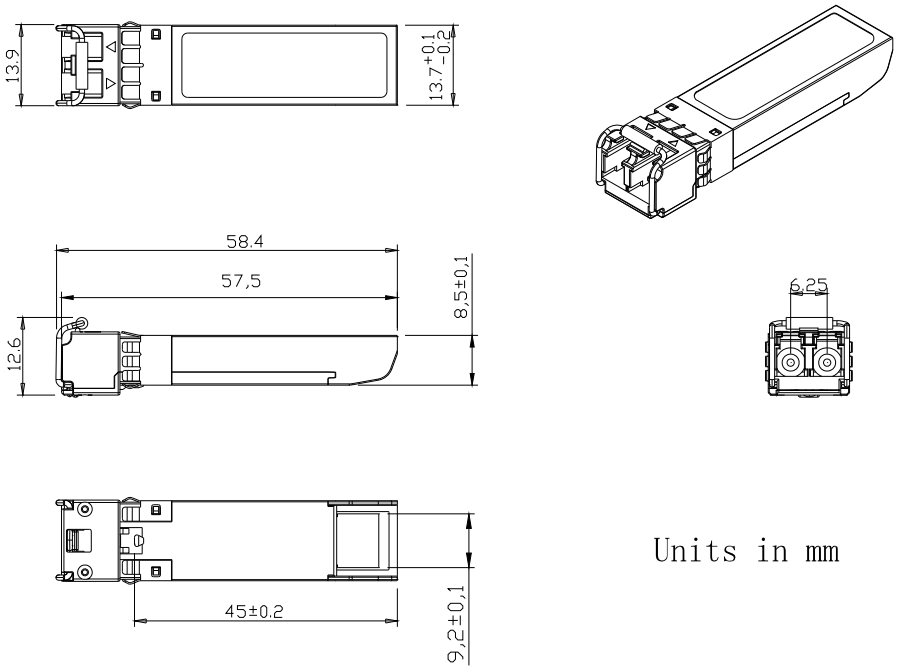
The SFP MSA defines a 256-byte memory map in EEPROM that is accessible over a 2-wire serial interface at the 8 bit address 1010000X (A0h). The digital diagnostic monitoring interface makes use of the 8 bit address 1010001X (A2h), so the originally defined serial ID memory map remains unchanged.

The operating and diagnostics information is monitored and reported by a Digital Diagnostics Transceiver Controller (DDTC) inside the transceiver, which is accessed through a 2-wire serial interface. When the serial protocol is activated, the serial clock signal (SCL, Mod Def 1) is generated by the host. The positive edge clocks data into the SFP transceiver into those segments of the E2PROM that are not write-protected. The negative edge clocks data from the SFP transceiver. The serial data signal (SDA, Mod Def 2) is bi-directional for serial data transfer. The host uses SDA in conjunction with SCL to mark the start and end of serial protocol activation. The memories are organized as a series of 8-bit data words that can be addressed individually or sequentially.

VI. Host - Transceiver Interface Block Diagram

VII. Outline Dimensions

Comply to SFF-8432 rev5.0, the improved Pluggable form factor specification.



VIII. Regulatory Compliance

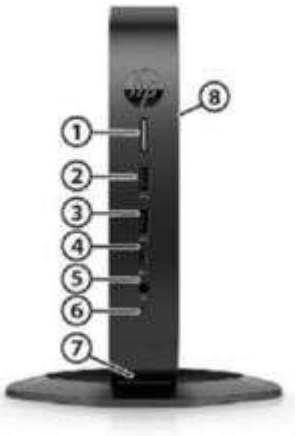
Feature	Reference	Performance
Electrostatic Discharge (ESD) to the Electrical pin（HBM）	MIL-STD-883E Method 3015.7 EIA-JESD22-A114	Class 1
(ESD) to the Simplex Receptacle	IEC/EN 61000-4-2	Compatible with standards
Electromagnetic Interference (EMI)	FCC Part 15 Class B EN 55022 Class B (CISPR 22A)	Compatible with standards
Laser Eye Safety	FDA 21CFR 1040.10, 1040.11 IEC/EN 60825-1，2	Class 1 laser product
Component Recognition	IEC/EN 60950 ，UL	Compatible with standards
ROHS	2002/95/EC	Compatible with standards

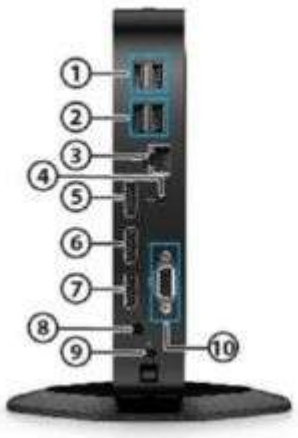
Appendix A. Document Revision

Version No.	Date	Description
1.0	2010-09-01	Preliminary datasheet
2.0	2011-09-10	Update format and company’s logo

Overview

HP Elite t655 Thin Client





FRONT

1	Power button (with integrated power indicator)
2	USB-A 3.1 Gen 1 port
3	USB-A 3.1 Gen 1 port
4	USB-C® 3.2 Gen 2 port
5	3.5mm combo headset/audio jack
6	Flash memory activity indicator
7	Stand
8	100 x 100mm VESA mounting holes

BACK

1	(2) USB-A 2.0 port
2	(2) USB-A 3.2 Gen 1 port
3	RJ-45 (network)jack- option with or without DASH Manageability support
4	Audio-out (headphone)
5	DisplayPort™ 1.2 connector
6	DisplayPort™ 1.2 connector
7	DisplayPort™ 1.2 connector
8	Power connector
9	Security cable slot
10	Configurable Option Port supporting one of the following:

- Blank; no optional configured port
- DP/USB-C® with PD
- DisplayPort™ over USB-C® with USB Power Delivery
- (2) USB-A 3.2 Gen 1 port
- HDMI 2.1 digital video output
- VGA analog video output
- External Wi-Fi® antenna connector
- Fiber Optic NIC connectors (SC or LC)
- Serial port with configurable power
- Dual serial ports (includes cable adapter)

Overview

AT A GLANCE

- AMD Ryzen R2314 System-on-Chip; 2.1 - 3.5 GHz; 4 cores, 4 threads[3]
- DASH Manageability support - Option to order with DASH or non-DASH
- AMD Memory Guard Secure Run technology that encrypts data in main memory
- DDR4 dual-channel SDRAM system memory; up to 2667 MT/s transfer rate; two SODIMM slots
- (3) DisplayPort™ 1.2 video outputs supporting up to UHD/4K (3840 x 2160 @ 60 Hz) resolutions
  NOTE: DisplayPort™ cables and displays sold separately.

- Solid-state flash memory storage; M.2 form factor modules; one slot
- Gigabit Ethernet (GbE) network connection; support for DASH out-of-band remote management
- Optional Allied Telesis M.2 Fiber Optic NICs; Fast Ethernet (100 Mb/s) or Gigabit (1,000 Mb/s)
- Optional Realtek WLAN 8852AE Wi-Fi 6 + Bluetooth® 5.2 WW adapter including antennas integrated internally in the chassis. (Antenna is internally integrated in the chassis with the Wi-fi® SKU)
  NOTE: Fiber optic and Wi-Fi® NIC options cannot be supported together[1]

  NOTE: Wireless features, performance and support may vary depending on environmental variables such placement, settings and firmware of your access points. Please contact your wireless vendor for support of your wireless environment

- Optional remote external Wi-Fi® antenna system
- Option Port with a selection of available factory options (see detailed listing later in this document)
- Integrated PC speaker for basic audio playback; 3.5 mm combo headset/audio port on front and 3.5 mm audio port on rear that can be configured as line in or line out supporting headphones, external speaker systems, or microphone
- 45W non-PFC external power adapter
- Security features include a TCG certified TPM version 2.0 and a system UEFI (BIOS) designed to address NIST SP 800-147 BIOS protection guidelines and NIST SP800-155 BIOS integrity measurement guidelines. A cable lock slot is provided for use with a cable lock to enable the system's physical security
- Passive thermal design (no cooling fans) and active thermal management technology that monitors the system operating temperatures, throttles SOC operation if appropriate and prevents unit thermal shutdown.
- Rated for a maximum ambient operating temperature of 40 degree C
- ENERGY STAR® certified configurations available and EPEAT® Silver registered in the United States. See http://www.epeat.net for registration status in other countries
- Post-consumer recycled plastics content greater than 50% total unit plastics (by weight)
- Low halogen[2] material content
- All models TAA compliant in North America

NOTE: See important legal disclosures for all listed specs in their respective features sections.

Warranty

HP one-year hardware limited warranty in most regions; HP Care Packs* are extended service contracts that go beyond your standard limited warranties; for more details visit http://www.hp.com/go/cpc

*HP Care Packs are sold separately. Service levels and response times for HP Care Packs may vary depending on your geographic location. Service starts on date of hardware purchase. Restrictions and limitations apply. For details, visit www.hp.com/go/cpc. HP Services are governed by the applicable HP terms and conditions of service provided or indicated to Customer at the time of purchase. Customer may have additional statutory rights according to applicable local laws, and such rights are not in any way affected by the HP terms and conditions of service or the HP Limited Warranty provided with your HP Product.

Technical Specifications

OPERATING SYSTEMS

- HP ThinPro, including HP /Smart Zero Core, HP Cloud Endpoint Manager[2]
- Windows 10 IoT Enterprise LTSC 2021[1]
- IGEL
- No OS

1. Not all features are available in all editions or versions of Windows. Systems may require upgraded and/or separately purchased hardware, drivers, software or BIOS update to take full advantage of Windows functionality. ISP fees may apply, and additional requirements may apply over time for updates. See http://www.windows.com.

2. Not all features are available with HP ThinPro, Smart Zero Core and FreeDOS.

PROCESSOR[3,4,5]

Model	CPU Frequency Max/Base	Cores/Threads	GPU Type	GPU Frequency
AMD Ryzen™ Embedded R2314 with Radeon™ Graphics	3.5/2.1 GHz	4/4	Radeon™ Graphics	1,200 MHz

3. Multicore is designed to improve performance of certain software products. Not all customers or software applications will necessarily benefit from use of this technology. Performance and clock frequency will vary depending on application workload and your hardware and software configurations. Intel's numbering, branding and/or naming is not a measurement of higher performance.

4. Processor speed denotes maximum performance mode; processors will run at lower speeds in battery optimization mode.

5. In accordance with Microsoft's support policy, HP does not support the Windows 8 or Windows 7 operating system on products configured with Intel and AMD 7th generation and forward processors or provide any Windows 8 or Windows 7 drivers on http://www.support.hp.com.

DISPLAY SUPPORT

Number of displays supported	**A maximum of 3 displays are supported.** Combination:
	3 x DisplayPort™ (onboard)
	2 x DisplayPort™ (onboard) + 1 x DisplayPort over USB-C (optional)
	2 x DisplayPort™ (onboard) + 1 x HDMI 2.1 (optional)
	2 x DisplayPort™ (onboard) + 1 x VGA (optional)

Video Resolution Support Matrix*

Windows 10 IoT 21H2* *Min of 8GB required	=3 x FHD 1920 X 1080 @ 60Hz	1 x UHD/4K 3840 x 2160 @ 60Hz	2 x UHD/4K 3840 x 2160 @ 60Hz	3 x UHD/4K 3840 x 2160 @ 60Hz (2 x 4G RAM or above)
Static screen (no video)	✓	✓	✓	✓
1080 60fps(or below) video	✓	✓	✓	✓
4K 30fps video	✓	✓	✓	✓
4K 60fps video	✓	✓	✓	✓

Technical Specifications

ThinPro 8	=3 x FHD 1920 X 1080 @ 60Hz	1 x UHD/4K 3840 x 2160 @ 60Hz	2 x UHD/4K 3840 x 2160 @ 60Hz	3 x UHD/4K 3840 x 2160 @ 60Hz (2 x 4G RAM or above)
Static screen (no video)	✓	✓	✓	✓

*For best 4K experience dual channel memory is required.

HP USB-C® DOCK G5 (support requires system to be configured with the Type-C® option Port)

Maximum resolution support for monitors connected to the HP USB-C® Dock G5:
Single Monitor - Max of 2560 x 1440 @60Hz refresh rate
Dual Monitor - Max of two monitors at 1920x1080 @60Hz refresh rate

For an expanded list of supported monitor combinations, refer to the HP Dock Quickspecs document. For the HP Elite t655 Thin Client, refer to the "DP 1.2 MF"? column.

GRAPHICS

Number of displays supported:	3	
Video outputs:	Standard:	(3) DisplayPort™ 1.2
	Optional:	(1) DisplayPort™ over USB-C® with USB Power Delivery (1) VGA analog output (1) HDMI 2.1 digital output
	NOTE: adding an optional output does not increase the number of displays supported.	
Max. screen resolution:	3840 x 2160 @ 60 Hz	

NOTE: HP recommends dual channel memory (two SODIMMs) configurations for optimal display resolution performance

Technical Specifications

MEMORY

Type:	DDR4 dual channel SDRAM
Data Transfer Rate:	Up to 2,667 MT/s
Peak Transfer Rate:	Up to 19,200 MB/s
Number of Slots	2 x SODIMM
Capacities:	4*, 8, 16 and 32 GB

NOTES:

- *** 4GB not configurable with Windows 10 IoT Enterprise LTSC 2021**
- The actual transfer rates will be dependent upon the specification of the SODIMM modules used
- The Graphics Processing Unit (GPU) uses part of the total system memory. System memory dedicated to graphics performance is not available for use by other programs
- HP recommends dual channel memory (two SODIMMs) configurations for optimal system performance

UEFI	
UEFI Specification Revision	2.8
TPM 2.0	Meets requirements for Common Criteria, an independent third-party certification of trustworthiness
	Meets requirements for FIPS 140-2, a standard for cryptographic integrity
Security features	System UEFI designed to address NIST SP 800-147 BIOS protection guidelines and NIST SP800-155 BIOS integrity measurement guidelines

STORAGE*

Type:	NAND flash memory; non-volatile
Number of Sockets:	(1) M.2
Capacities:	
	32 GB M.2 eMMC flash module (32GB not configurable with Windows 10 IoT Enterprise LTSC 2021)
	64 GB M.2 eMMC flash module
	256 GB M.2 PCIe NVMe flash module
	512 GB M.2 PCIe NVMe flash module

*For storage drives, GB = 1 billion bytes. TB = 1 trillion bytes. Actual formatted capacity is less. Up to 30 GB (for Windows 10 and 11) is reserved for system recovery software.

Technical Specifications

Input/Output

USB:
- Front access: (2) USB-A 3.2 Gen 1 port; (1) USB-C® 3.2 Gen 2 port
- Rear access: (2) USB-A 2.0 port (optional WOSK (designated for Power-on from Keyboard if equipped)); (2) USB-A 3.2 port; (2) USB-A 3.2 Gen 1 ports (optional); (1) USB-C® DisplayPort™ (optional)

Video Outputs:
- Standard: (3) DisplayPort™ 1.2 digital outputs
- Optional: (1) VGA analog output; (1) HDMI 2.1 digital output; (1) DisplayPort™ over USB-C® with USB Power Delivery

NOTE: adding an optional output does not increase the number of displays supported.

I/O Interfaces:
- Standard: (1) RJ45 network connector; (1) 3.5 mm combo headset/audio jack (front); (1) 4.5 mm DC audio jack (rear)
- Optional: (1) Dual serial port with configurable power; (2) serial ports enabled with an included cable adapter

Option Port: The rear I/O panel includes an Option Port that can be configured with one of the following factory options:

- Blank; no optional configured port
- 2 x USB-A 3.0 Gen 1 ports
- DisplayPort™ over USB-C® with USB Power Delivery
- HDMI 2.1 digital video output
- VGA analog video output
- External Wi-Fi® antenna connector (requires Wi-Fi® adapter option)
- Fiber Optic NIC connectors; SC or LC connector (requires Fiber Optic NIC option)
- Serial port with configurable power
- Dual serial ports enabled with an included cable adapter

AUDIO/VIDEO

Audio Subsystem
- Internal amplified speaker system for basic audio playback
- 3.5 mm combo headset/headphone/analog microphone audio jack (front access)
- 3.5 mm /4.5mm combo line-out/ line-in socket (rear access)

Audio CODECs
- MP3
- AAC Stereo
- HE AAC
- Includes hardware acceleration support

Video CODECs
- MPEG-4 part 2 (DivX, Xvid)
- MPEG-4 part 10 (H.264, AVC), Advanced Video Coding (AVC) (H.264 encode & decode)
- MPEG-H part 2, High Efficiency Video Coding (HEVC, available with Windows 10 IoT Only) (H.265 (8-bit / 10-bit) decode and (8-bit) encode
- WMV 7/8/9 VC-1 & ASF Demuxer
- Includes hardware acceleration support

Technical Specifications

NETWORKING

Local Area Networking Realtek RTL8111EPH-CG Gigabit Ethernet (GbE) Controller with support for DASH out-of-band remote management and Wake-On-Lan. Realtek RTL8111HSH-CG Gigabit Ethernet (GbE) Controller with support non-DASH and Wake-On-Lan[2]

Wi-Fi® Networking Realtek 8852AE Wi-Fi 6 +BT5.2 WLAN[1]
Realtek 8852BE Wi-Fi 6 +BT5.2 WLAN[1]

1. Wireless access point and Internet service required and sold separately. Availability of public wireless access points limited. Wi-Fi 6 (802.11ax) is backwards compatible with prior 802.11 specs. Wi-Fi 6 is designed to support gigabit data rate when transferring files between two devices connected to the same router. Requires a wireless router, sold separately, that supports 80MHz and higher channels
2. The term "10/100/1000"? or "Gigabit"? Ethernet indicates compatibility with IEEE standard 802.3ab for Gigabit Ethernet, and does not connote actual operating speed of 1 Gb/s. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.

FIBER OPTIC NETWORKING

Adapter Options:
- Allied Telesis AT-27M2/SC Fiber Fast Ethernet M.2 Adapter
- Allied Telesis AT-29M2/SC or LC Fiber Gigabit M.2 Adapter

Features:
- IEEE 802.1p priority encoding/tagging (QoS, CoS)
- IEEE 802.1q VLAN tagging
- IEEE 802.3x flow control
- Buffer/FIFO: 22K transmit and 40K receive
- Loopback mode
- Descriptor-Based Buffer Management
- Wake-on-LAN from S3 (Sleep) and S4 (Hibernate) not supported (AT-29M2)
- Link Detection and PHY interface power; the PHY interface, Link detection and Link LED should be enabled by default at power-up

Performance:

AT-27M2
- >= 85 Mbit/s receive, <= 30% CPU utilization
- >= 85 Mbit/s transmit, <= 30% CPU utilization
- >= 170 Mbit/s total bi-directional, <= 30% CPU utilization

NOTE: The minimum transfer size at 100 Mbit/s is 1 Gbps

AT-29M2
- >= 800 Mbit/s receive, <= 30% CPU utilization
- >= 800 Mbit/s transmit, <= 30% CPU utilization
- >= 1500 Mbit/s total bi-directional, <= 30% CPU utilization

NOTE: The minimum transfer size at 1000 Mbit/s is 1500 Gbps

External Interface: Complies with IEEE 802.3 100BASE-FX operation (AT-27M2)
Complies with IEEE 802.3 1000BASE-X operation (AT-29M2)

Power:
- Uses less than 1800 mW of power at full performance (AT-27M2)
- Uses less than 2100 mW of power at full performance (AT-29M2)
- Supports all PCI Express bus states L0, L0s, L1 and L2

Non-volatile Storage: The MAC address is unique for each system; assigned from the board assembly manufacturer's IEEE registered allocation.

The PCI subsystem ID is unique to HP and unique to each design to allow Windows Update to be finely controlled.

Technical Specifications

SOFTWARE SUPPORT

Host Environment	Protocol	HP ThinPro	Microsoft Windows 10 IoT Enterprise 2021
Remote Desktop Services	Remote FX (RFX), RDP	✓	✓
Citrix®	ICA, HDX	✓	✓
Vmware® Horizon	RDP, PcoIP, Blast Extreme	✓	✓

Protocol Clients	HP ThinPro	Microsoft Windows 10 IoT Enterprise 2021
Citrix® Workspace app	✓	✓
Microsoft Remote Desktop Client		✓
Free Remote Desktop Client	✓	
Vmware™ Horizon View™ Client	✓	✓
HP Remote Graphics Software (RGS)	via add-on	✓
Turbosoft Tterm for Linux® Terminal emulation Software	via add-on	
Turbosoft TTWin Terminal emulation software		via add-on
AVD/Win365	via add-on	via add-on
AWS Amazon Workspace	via add-on	via add-on
HP Anyware[1]	via add-on	via add-on

[1]HP Anyware supports Windows®, Linux® and MacOS® host environments and Window, Linux, MacOS, iOS®, Android®, and Chrome OS® end-user devices. For more on the system requirements for installing HP Anyware, refer to the Admin Guides at: https://docs.teradici.com/find/product/cloud-access-software

Browser Support	HP ThinPro	Microsoft Windows 10 IoT Enterprise 2021
Mozilla Firefox	✓	
Internet Explorer		✓
Microsoft Edge		✓

Security	HP ThinPro	Microsoft Windows 10 IoT Enterprise 2021
Smart Card	✓	✓
Log-on Manager	✓	✓
Read only Operating System	✓	✓
802.1x	✓	✓
Microsoft Firewall		✓
HP Write Manager		✓
Microsoft Unified Write Filter		✓

NOTE: the HP Write Manager is the default active write filter. The Microsoft Unified Write Filter is disabled by default but can be enabled by the user if required.

Technical Specifications

Management Tools	HP ThinPro	Microsoft Windows 10 IoT Enterprise 2021
HP Cloud Endpoint Manager	✓	✓
HP Device Manager	✓	✓
HP ThinUpdate		✓
HP Easy Tools	✓	
HP Smart Zero Client Services	✓	
Microsoft SCCM/EDM agent		✓
HP USB Port Manager		✓
HP User State Tool		Add-on only

Additional Windows Components	HP ThinPro	Microsoft Windows 10 IoT Enterprise 2021
HP Easy Shell		✓
Windows Media Player		✓
Microsoft Direct Access		✓
Microsoft BranchCache		✓
Microsoft AppLocker		✓
Microsoft Sideloading		✓
CyberLink Media Player		✓

NOTE: Other add-on software available (see: http://www.hp.com/support for latest list of available add-ons). Software performance and support may vary depending on customer environment and backend.

Audio/Video CODECs	HP ThinPro	Microsoft Windows 10 IoT Enterprise 2021
MP3	✓	✓
WMA stereo	✓	✓
AAC stereo & HE AAC	✓	
Microsoft AC3 encoder		✓
MPEG-1	✓	
MPEG-4 part 2 (DivX, Xvid, H.263)	✓	✓
MPEG-4 part 10 (H.264, AVC)	✓	✓
h.365/HEVC	✓	✓
WMV 7/8/9/ VC-1 & ASF Demuxer	✓	✓

Recommended TC config for Microsoft Teams media optimization

	TC CPU	Vmware Teams Optimization	Citrix Teams Optimization
t655	2.10 GHz 4 Core	✓	✓

- Not recommended, ✓ recommended

Technical Specifications

WEIGHTS & DIMENSIONS

W x D x H: (vertical orientation)	35 x 200 x 200 mm
Volume:	1.4 liter
System Weight (unit with stand)	1174g Lowest weight. Weight will vary by configuration.
Shipping Weight (System+Stand+AC+PC)	2106g
Shipping Box Dimensions	495 x 132 x 260 mm

NOTE: All measurements are approximate; the addition of optional modules will increase the weight

EXTERNAL POWER SUPPLY

45W non-PFC Smart external power adapter

Worldwide auto-sensing 100 - 240 VAC; nominal voltage is 120 VAC; 50 - 60 Hz

Energy saving automatic power-down; surge tolerant

1.8m output cable

External power adapters are sourced from several suppliers in order to ensure adequate supply and availability is maintained. The actual dimensions of the power brick will vary by supplier.

HP P/N	Vendor	Dimensions:
L25296-001	Lite-On	94 x 40 x 26.5 mm
L25296-002	Chicony	95 x 40 x 26.5 mm
L25296-003	Delta	94 x 39 x 26.5 mm
L25296-004	AcBel	91.4 x 44 x 26.8 mm

COMPLIANCE/CERTIFICATIONS

Accessibility:	Section 508 Accessibility; VPAT report available.
Environmental Stewardship:	Worldwide (ENERGY STAR® configurations available, EPEAT 2.0, RoHS2, ERP, TCO Certified, CECP& SEPA, HP GSE, WEEE, Low Halogen, etc.)
Product Safety:	Worldwide (UL, CB, GS, CCC, BSMI, etc.)
Electromagnetic Compliance (EMC):	Worldwide (FCC/CISPR/EN/VCCI/ICES/AS/NZS/CNS/KCC) "Class B"? EMI regulations
International Medical Safety Standard:	EN60601-1-2 (Medical Equipment EMC) passed

ENVIRONMENTAL

Operating Temperature Range:	50° to 104° F 10° to 40° C	
Non-operating Temperature Range:	-22° to 140° F -30° to 65° C	
Humidity:	Condensing: Non-condensing:	20% to 80% 10% to 90%

NOTE: Specifications are at sea level with altitude derating of 1° C/300m (1.8° F/1000ft) to a maximum of 3 Km (10,000 ft), with no direct, sustained sunlight. Upper limit may be limited by the type and number of options installed.

Technical Specifications

Environmental Data	**Eco-Label Certifications & declarations**	This product has received or is in the process of being certified to the following approvals and may be labeled with one or more of these marks: • IT ECO declaration • US ENERGY STAR® • US Federal Energy Management Program (FEMP) • EPEAT? Gold registered in the United States. See http://www.epeat.net for registration status in your country. • TCO Certified • China Energy Conservation Program (CECP) • China State Environmental Protection Administration (SEPA) • Taiwan Green Mark • Korea Eco-label • Japan PC Green label*		
	Sustainable Impact Specifications	• Ocean-bound plastic in stand • 55% post-consumer recycled plastic • Low halogen • Outside Box and corrugated cushions are 100% sustainably sourced and recyclable • Molded Paper Pulp Cushion inside box is 100% sustainably sourced and recyclable		
	System Configuration	The configuration used for the Energy Consumption and Declared Noise Emissions data for the Notebook model is based on a "Typically Configured Notebook"?.		
	Energy Consumption (in accordance with US ENERGY STAR® test method)	**115VAC, 60Hz**	**230VAC, 50Hz**	**100VAC, 50Hz**
	Normal Operation (Sort idle)	6.44 W	6.94 W	6.96 W
	Normal Operation (Long idle)	5.34 W	5.49 W	6.12 W
	Sleep	0.85 W	0.94 W	0.86 W
	Off	0.8 W	0.86 W	0.76 W
		NOTE: Energy efficiency data listed is for an ENERGY STAR® compliant product if offered within the model family. HP computers marked with the ENERGY STAR® Logo are compliant with the applicable U.S. Environmental Protection Agency (EPA) ENERGY STAR® specifications for computers. If a model family does not offer ENERGY STAR® compliant configurations, then energy efficiency data listed is for a typically configured PC featuring a hard disk drive, a high efficiency power supply, and a Microsoft Windows® operating system.		
	Heat Dissipation*	**115VAC, 60Hz**	**230VAC, 50Hz**	**100VAC, 50Hz**
	Normal Operation (Short idle)	22.0 BTU/hr	23.7 BTU/hr	23.8 BTU/hr
	Normal Operation (Long idle)	18.3 BTU/hr	18.8 BTU/hr	20.9 BTU/hr
	Sleep	2.9 BTU/hr	3.2 BTU/hr	2.9 BTU/hr
	Off	2.7 BTU/hr	2.9 BTU/hr	2.6 BTU/hr
		***NOTE:** Heat dissipation is calculated based on the measured watts, assuming the service level is attained for one hour.		

Technical Specifications

Declared Noise Emissions (in accordance with ISO 7779 and ISO 9296)	Sound Power (L_{Wad}, bels)	Sound Pressure (L_{pAm}, decibels)	
Typically Configured - Idle	2.5	13.7	
Fixed Disk - Random writes	2.5	13.7	
Optical Drive - Sequential reads	2.5	13.9	
Longevity and Upgrading	This product can be upgraded, possibly extending its useful life by several years. Upgradeable features and/or components contained in the Spare parts are available throughout the warranty period and or for up to "5"? years after the end of production.		
Additional Information	• This product is in compliance with the Restrictions of Hazardous Substances (RoHS) directive - 2011/65/EC. • This HP product is designed to comply with the Waste Electrical and Electronic Equipment (WEEE) Directive - 2002/96/EC. • This product is in compliance with California Proposition 65 (State of California; Safe Drinking Water and Toxic Enforcement Act of 1986). • This product is in compliance with the IEEE 1680 (EPEAT) standard at the Gold level, see www.epeat.net • Plastics parts weighing over 25 grams used in the product are marked per ISO11469 and ISO1043. • This product is 90.3% recycle-able when properly disposed of at end of life.		
Packaging Materials	**External:**	PAPER/Corrugated	366 g
		PAPER/Molded Pulp	51 g
		PAPER/Corrugated	80 g
		PAPER/Corrugated	115 g
	Internal:	PLASTIC/Polyethylene low density - LDPE	6 g
	The plastic packaging material contains at least 30.0% recycled content.		
	The corrugated paper packaging materials contains at least 87.0% recycled content.		
RoHS Compliance	HP Inc. complies fully with materials regulations. We were among the first companies to extend the restrictions in the European Union (EU) Restriction of Hazardous Substances (RoHS) Directive to our products worldwide through the HP GSE. HP has contributed to the development of related legislation in Europe, as well as China, India, and Vietnam. We believe the RoHS directive and similar laws play an important role in promoting industry-wide elimination of substances of concern. We have supported the inclusion of additional substances-including PVC, BFRs, and certain phthalates-in future RoHS legislation that pertains to electrical and electronics products. We met our voluntary objective to achieve worldwide compliance with the new EU RoHS requirements for virtually all relevant products by July 2013, and we will continue to extend the scope of the commitment to include further restricted substances as regulations continue to evolve. To obtain a copy of the HP RoHS Compliance Statement, see HP RoHS position statement.		

Material Usage	This product does not contain any of the following substances in excess of regulatory limits (refer to the HP General Specification for the Environment at http://www.hp.com/hpinfo/globalcitizenship/environment/supplychain/gen_specifications.html): • Asbestos • Certain Azo Colorants • Certain Brominated Flame Retardants - may not be used as flame retardants in plastics • Cadmium • Chlorinated Hydrocarbons • Chlorinated Paraffins • Bis(2-Ethylhexyl) phthalate (DEHP) • Benzyl butyl phthalate (BBP) • Dibutyl phthalate (DBP) • Diisobutyl phthalate (DIBP) • Formaldehyde • Halogenated Diphenyl Methanes • Lead carbonates and sulfates • Lead and Lead compounds • Mercuric Oxide Batteries • Nickel - finishes must not be used on the external surface designed to be frequently handled or carried by the user. • Ozone Depleting Substances • Polybrominated Biphenyls (PBBs) • Polybrominated Biphenyl Ethers (PBBEs) • Polybrominated Biphenyl Oxides (PBBOs) • Polychlorinated Biphenyl (PCB) • Polychlorinated Terphenyls (PCT) • Polyvinyl Chloride (PVC) - except for wires and cables, and certain retail packaging has been voluntarily removed from most applications. • Radioactive Substances • Tributyl Tin (TBT), Triphenyl Tin (TPT), Tributyl Tin Oxide (TBTO)
Packaging Usage	HP follows these guidelines to decrease the environmental impact of product packaging: • Eliminate the use of heavy metals such as lead, chromium, mercury and cadmium in packaging materials. • Eliminate the use of ozone-depleting substances (ODS) in packaging materials. • Design packaging materials for ease of disassembly. • Maximize the use of post-consumer recycled content materials in packaging materials. • Use readily recyclable packaging materials such as paper and corrugated materials. • Reduce size and weight of packages to improve transportation fuel efficiency. • Plastic packaging materials are marked according to ISO 11469 and DIN 6120 standards.
End-of-life Management and Recycling	HP offers end-of-life HP product return and recycling programs in many geographic areas. To recycle your product, please go to: http://www.hp.com/go/reuse-recycle or contact your nearest HP sales office. Products returned to HP will be recycled, recovered or disposed of in a responsible manner. The EU WEEE directive (2002/95/EC) requires manufacturers to provide treatment information for each product type for use by treatment facilities. This information (product disassembly instructions) is posted on the Hewlett Packard web site at: http://www.hp.com/go/recyclers. These instructions may be used by recyclers and other WEEE treatment facilities as well as HP OEM customers who integrate and re-sell HP equipment.

Technical Specifications

	HP, Inc. Corporate Environmental Information	For more information about HP's commitment to the environment: Global Citizenship Report http://www.hp.com/hpinfo/globalcitizenship/gcreport/index.html Eco-label certifications http://www8.hp.com/us/en/hp-information/environment/ecolabels.html ISO 14001 certificates: http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=c04755842 and http://www.hp.com/hpinfo/globalcitizenship/environment/pdf/cert.pdf
	footnotes	• Percentage of ocean-bound plastic contained in each component varies by product • Recycled plastic content percentage is based on the definition set in the IEEE 1680.1-2018 standard. • External power supplies, WWAN modules, power cords, cables and peripherals excluded. • 100% outer box packaging and corrugated cushions made from sustainably sourced certified and recycled fibers. • Fiber cushions made from 100% recycled wood fiber and organic materials. • Plastic cushions are made from >90% recycled plastic.

Acoustic Noise Emission Declaration						
Measurement and declaration standards:	ISO 7779: Acoustics - Measurement of airborne noise emitted by information technology and telecommunications equipment					
	ISO 9296: Acoustics - Declared noise emission values of computer and business equipment					
Declared Noise Emission Values in accordance with ISO 9296						
Product Configuration or Operating Mode # (see section below for description)	Declared Sound Power Level, L_{Wad} Bels			Declared Sound Pressure Level, L_{pAm} dBA Tested on ISO Table		
	Idle	Operating		Idle	Operating	
	ECMA-74 C.15.3.2 Idle Mode	ECMA-74 C.9.3.2 Drive Random Seek	ECMA-74 C.15.3.3 g Active Mode	ECMA-74 C.15.3.2 Idle Mode	ECMA-74 C.9.3.2 Drive Random Seek	ECMA-74 C.15.3.3 g Active Mode
Vertical	2.6	2.6	2.6	13.7	13.7	13.7
Horizontal	2.6	2.6	2.6	N/A	N/A	N/A
The Product meets the acoustic noise limits of these voluntary Eco labels:				TCO Certified Blue Angel Nordic Swan EU Flower		
NOTE: Measured under ISO 7779 and ISO 9296 measurement and declaration standards.						

Summary of Changes

Date of change:	Version History:	Type of change	Description of change:
October 7, 2022	From v1 to v2	Changed	PROCESSOR and MEMORY sections
October 13, 2022	From v2 to v3	Changed	PROCESSOR section
October 27, 2022	From v3 to v4	Changed	WEIGHTS & DIMENSIONS, Recommended TC config for Microsoft Teams media optimization sections
February 6, 2023	From v4 to v5	Changed	Format page 1, Storage and COMPLIANCE/CERTIFICATIONS sections
March 17, 2023	From v5 to v6	Changed	SOFTWARE SUPPORT section
April 17, 2023	From v6 to v7	Changed	UEFI section
April 25, 2023	From v7 to v8	Changed	Input/Output section
May 2, 2023	From v8 to v9	Changed	Environmental Data section
May 6, 2023	From v9 to v10	Changed	NETWORKING section
June 12, 2023	From v10 to v11	Changed	HP USB-C DOCK G5 section
October 13, 2023	From v11 to v12	Changed	WEIGHTS & DIMENSIONS section
October 27, 2023	From v12 to v13	Changed	Input/Output section

title

© 2023 HP Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express limited warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein. AMD and Radeon are trademarks of Advanced Micro Devices, Inc. DisplayPort™ and the DisplayPort™ logo are trademarks owned by the Video Electronics Standards Association (VESA®) in the United States and other countries. Amazon Web Services, the "Powered by Amazon Web Services"? logo, and Amazon WorkSpaces are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries. Bluetooth® is a trademark owned by its proprietor and used by HP Inc. under license. ENERGY STAR is a registered trademark owned by the U.S. Environmental Protection Agency. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. VMware Horizon and VMware Horizon View are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Citrix and Citrix Workspace are trademarks of Citrix Systems, Inc. and/or one more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries.

Veeam Backup & Replication 12

User Guide for VMware vSphere

Hardened Repository

https://bp.veeam.com/vbr/Security/hardening_backup_repository_linux.html

https://helpcenter.veeam.com/docs/backup/vsphere/hardened_repository.html?ver=120

https://www.veeam.com/blog/hardened-linux-repository-best-practices.html

Kompletní dokumentace

https://www.veeam.com/veeam_backup_12_user_guide_vsphere_pg.pdf#Hardened%20Repository

Comparison of Windows Server editions

Article • 10/25/2024 •

Applies to: ✅ Windows Server 2025

Use this article to compare the Standard, Datacenter, and Datacenter: Azure Edition of Windows Server to see which is most appropriate for you.

> **Tip**
>
> If you're looking for information about the locks and limits in Windows Server, see Comparison of locks and limits in Windows Server.

Available roles and features

Expand table

Feature	Sub-feature	Standard edition	Datacenter edition	Datacenter: Azure Edition
.NET Framework 3.5 Features		✓	✓	✓
.NET Framework 4.8 Features		✓	✓	✓
Activation		✓	✓	✓
	Automatic Virtual Machine Activation	! [1]	✓	✓ [2]
	Key Management Service (KMS)	✓	✓	✗ [3]
Active Directory Certificate Services		✓	✓	✓
	Certificate Enrollment Policy Web Service	✓	✓	✓

Feature	Sub-feature	Standard edition	Datacenter edition	Datacenter: Azure Edition
	Certificate Enrollment Web Service	✓	✓	✓
	Certification Authority	✓	✓	✓
	Certification Authority Web Enrollment	✓	✓	✓
	Network Device Enrollment Service	✓	✓	✓
	Online Responder	✓	✓	✓
Active Directory Domain Services		✓	✓	✓
Active Directory Federation Services		✓	✓	✓
Active Directory Lightweight Directory Services		✓	✓	✓
Active Directory Rights Management Services		✓	✓	✓
Azure Extended Networking		ⓧ	ⓧ	✓
Background Intelligent Transfer Service (BITS)		✓	✓	✓
BitLocker Drive Encryption		✓	✓	✓
BitLocker Network Unlock		✓ [4]	✓ [4]	✓ [4]

Feature	Sub-feature	Standard edition	Datacenter edition	Datacenter: Azure Edition
BranchCache		✓	✓	✓
Client for NFS		✓	✓	✓
Data Center Bridging		✓	✓	✓
Device Health Attestation		✓	✓	✓
DHCP Server		✓	✓	✓
Direct Play		✓ [3]	✓ [4]	✓ [4]
DLNA codecs and web media streaming		✓ [3]	✓ [4]	✓ [4]
DNS Server		✓	✓	✓
Enhanced Storage		✓	✓	✓
Failover Clustering		✓	✓	✓
Fax Server		✓	✓	✓
File and Storage Services		✓	✓	✓
	BranchCache for Network Files	✓	✓	✓
	Data Deduplication	✓	✓	✓
	DFS Namespaces	✓	✓	✓
	DFS Replication	✓	✓	✓
	File Server	✓	✓	✓
	File Server Resource Manager	✓	✓	✓

Feature	Sub-feature	Standard edition	Datacenter edition	Datacenter: Azure Edition
	File Server VSS Agent Service	✓	✓	✓
	iSCSI Target Server	✓	✓	✓
	iSCSI Target Storage Provider (VDS and VSS hardware providers)	✓	✓	✓
	Server for NFS	✓	✓	✓
	SMB 1.0/CIFS File Sharing Support	✓	✓	✓
	SMB Bandwidth Limit	✓	✓	✓
	SMB over QUIC	✓	✓	✓
	Work Folders	✓	✓	✓
	Storage Migration Service	✓	✓	✓
	Storage Migration Service Proxy	✓	✓	✓
	Storage Spaces	✓	✓	✓
	Storage Spaces Direct	✗	✓	✓
	Storage Replica	✓	✓	✓
Group Policy Management		✓	✓	✓
Host Guardian Hyper-V Support		✗	✓	✓
Host Guardian Service		✓	✓	✓
Hotpatching		✓ [5]	✓ [5]	✓

Feature	Sub-feature	Standard edition	Datacenter edition	Datacenter: Azure Edition
I/O Quality of Service		✓	✓	✓
IIS Hostable Web Core		✓	✓	✓
IP Address Management (IPAM) Server		✓	✓	✓
Management OData IIS Extension		✓	✓	✓
Media Foundation		✓	✓	✓
Message Queuing		✓	✓	✓
	Message Queuing DCOM Proxy	✓	✓	✓
	Message Queuing Services	✓	✓	✓
Microsoft Defender Antivirus		✓	✓	✓
Multipath I/O		✓	✓	✓
MultiPoint Connector		✓	✓	✓
Network ATC		✓	✓	✓
Network Controller		ⓧ	✓	✓
Network Load Balancing		✓	✓	✓
Network Policy and Access Services		✓ [4]	✓ [4]	✓ [4]

Feature	Sub-feature	Standard edition	Datacenter edition	Datacenter: Azure Edition
Network Virtualization		✓	✓	✓
Print and Document Services		✓	✓	✓
	Internet Printing	✓ [4]	✓ [4]	✓ [4]
	Line Printer Daemon (LPD) Service	✓ [4]	✓ [4]	✓ [4]
	Print Server	✓ [4]	✓ [4]	✓ [4]
Quality Windows Audio Video Experience		✓	✓	✓
RAS Connection Manager Administration Kit (CMAK)		✓	✓	✓
Remote Access		✓	✓	✓
	DirectAccess and VPN (RAS)	✓	✓	✓
	Routing	✓	✓	✓
	Web Application Proxy	✓	✓	✓
Remote Assistance		✓ [4]	✓ [4]	✓ [4]
Remote Desktop Services		✓ [4]	✓ [4]	✓ [4]
Remote Differential Compression		✓	✓	✓
Remote Server Administration Tools		✓	✓	✓

Feature	Sub-feature	Standard edition	Datacenter edition	Datacenter: Azure Edition
RPC over HTTP Proxy		✓	✓	✓
Setup and Boot Event Collection		✓	✓	✓
Simple TCP/IP Services		✓ [4]	✓ [4]	✓ [4]
SNMP Service		✓	✓	✓
Software Load Balancer		✓	✓	✓
System Data Archiver		✓	✓	✓
System Insights		✓	✓	✓
Telnet Client		✓	✓	✓
TFTP Client		✓ [4]	✓ [4]	✓ [4]
Virtualization		✓	✓	✓
	Containers	✓	✓	✗
	Hyper-V	✓	✓	✓
	GPU partitioning	✓ [6]	✓	✓ [6]
	VM Shielding Tools for Fabric Management	✓	✓	✓
Volume Activation Services		✓	✓	✓
Web Server (IIS)		✓	✓	✓
	FTP Server	✓	✓	✓
	Web Server	✓	✓	✓
WebDAV Redirector		✓	✓	✓

Feature	Sub-feature	Standard edition	Datacenter edition	Datacenter: Azure Edition
Windows Biometric Framework		✓ [4]	✓ [4]	✓ [4]
Windows Deployment Services		✓	✓	✓
Windows Identity Foundation 3.5		✓ [4]	✓ [4]	✓ [4]
Windows Internal Database		✓	✓	✓
Windows PowerShell		✓	✓	✓
	Windows PowerShell 2.0 Engine	✓	✓	✓
	Windows PowerShell 5.1	✓	✓	✓
	Windows PowerShell Desired State Configuration Service	✓	✓	✓
	Windows PowerShell Web Access	✓	✓	✓
Windows Process Activation Service		✓	✓	✓
Windows Search Service		✓ [4]	✓ [4]	✓ [4]
Windows Server Backup		✓	✓	✓
Windows Server Migration Tools		✓	✓	✓

Feature	Sub-feature	Standard edition	Datacenter edition	Datacenter: Azure Edition
Windows Server Update Services		✓	✓	✓
Windows Standards-Based Storage Management		✓	✓	✓
Windows Subsystem for Linux		✓	✓	✓
Windows TIFF IFilter		✓ [4]	✓ [4]	✓ [4]
WinRM IIS Extension		✓	✓	✓
WINS Server		✓	✓	✓
Wireless LAN Service		✓	✓	✓
WoW64 Support		✓	✓	✓
XPS Viewer		✓ [4]	✓ [4]	✓ [4]

1. As a guest if hosted on a virtualization host activated with Datacenter edition
2. Datacenter: Azure Edition can be used as a nested host or guest
3. Activated by Azure and can't be configured as a KMS host
4. When installed as Server with Desktop Experience
5. Available as an Azure Arc-enabled service. To learn more about Azure Arc pricing, see the Azure Pricing Calculator
6. Available with Windows Server 2025 Standard, and is designed for standalone servers. Live migrate VMs between standalone nodes for planned downtime. If clustering is required for unplanned downtime, Windows Server 2025 Datacenter must be used instead.

LC Optical Patch Cord

Description

Fiber optic patch cord are used for linking the equipment and components in the fiber optic network. They are with various kinds of fiber optic connector types. The fiber optic patch cord types are classfied by the fiber optic connector types. For example, we name LC fiber optic patch cord because this cable is with LC fiber optic connector. There are PC, UPC, APC type fiber patch cord. They are different because of the polish of fiber connector

.



Fiber optic patch cords are with various kinds of cables, PVC, Riser, Plenum, OFNR, OFNP, LSZH (LSOH), 10 Gig Aqua. There are simplex, duplex and multi fiber cable assemblies, and there are ribbon fan out fiber cable assemblies and bundle fiber optic cable assemblies. Typical fiber optic patch cord cable diameter can be 0.9mm, 2mm, or 3mm.

Application



- Telecommunication Networks Local Area Networks (LANs)
- Data processing networks
- CATV
- Video
- Active Device Termination
- Premise Installations
- Wide Area Networks (WANs)
- Industrial, medica & military



Features

- Simplex or duplex
- PC or APC end-face finishing
- Singlemode or Multimode
- 3.0mm, 2.0mm or , 0.9mm
- Compliant with Telcordia GR-326-CORE specification

Specifications

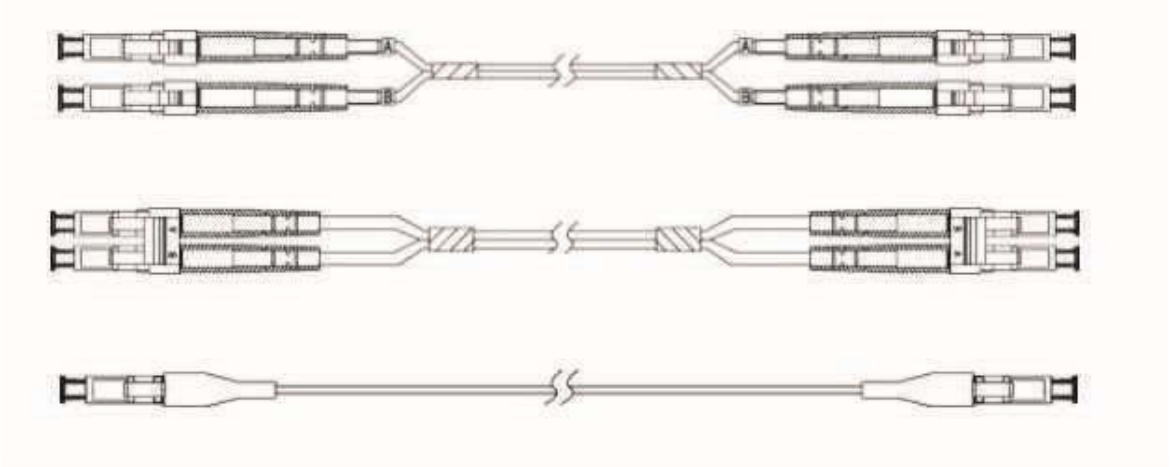
PARAMETERS	TEST CONDITION	VALUE(dB)	
		MULTIMODE	SINGLEMODE
Insertion Loss	850/1300nm or 1310/1550nm	≤0.3	≤0.3
Return Loss	850/1300nm or 1310/1550nm	≥35	≥50 (APC≥60)
Interchangeability	Connect Randomly	<0.2	<0.2
Vibration	5~50Hz, 1.5mm Amplitude	<0.1	<0.1
Tensile Strength	0~15Kgf (0.9mm Exception)	<0.1	<0.1
High Temperature	+85□, During 100 Hours	<0.2	<0.2
Low Temperature	-40°C, During 100 Hours	<0.2	<0.2
Temperature Cycle	-40°C~+85°C, during 5 Cycles	<0.2	<0.2
Humidity	+25°C~+65°C, 93%R. H, 100 hours	<0.2	<0.2

Drawing:

DX 2.0mm

DX 3.0mm

SX 0.9mm



safetica

Intelligent Data Security



2024-11-14

www.safetica.com

safetica

Safetica detects and secures sensitive data and helps teams work safely with information.



Safetica provides Intelligent Data Security, a smarter way to protect data that goes beyond content to understand the context of how data is being used.

Safetica covers today's most critical data security scenarios

Gain data visibility and discover sensitive data

Safetica helps you to discover and classify valuable data using its unique Safetica Unified Classification which combines analysis of file content, file origin and file properties. It offers complete visibility and continuous monitoring, never missing a beat to instantly identify, classify, and track sensitive data.

Protect sensitive and business-critical data

With Safetica, you can protect sensitive business- or customer-related data, source codes, or blueprints from accidental or intentional leakage. It audits all sensitive data activities no matter where the data moves, so you can report and investigate where there's a risk of leakage or theft. These findings are key for data protection.

Prevent insider risks and promote security awareness

Anyone can make a mistake which could put your business at risk. With Safetica, you can analyze insider risks, detect threats, and mitigate them swiftly. Notifications about how to treat sensitive data can help raise awareness around data security and educate your users.

Maintain data security for remote work

Safetica provides the full range of capabilities, including data protection, complete contextual visibility, and incident-driven training, regardless of location or network status. Get your hybrid digital work environment under control, discover unwanted software and services, analyze behavior to detect and audit high-risk users.

Detect and mitigate regulatory compliance violations

Safetica helps you to detect, prevent, and mitigate regulatory violations. Its audit capabilities support incident investigation to comply with regulations and data protection standards like GDPR, HIPAA, SOX, PCI-DSS, GLBA, ISO/IEC 27001, or CCPA.

© Copyright All rights reserved, 2024

safetica

What Makes Us Different

Contextual Defense

Traditional DLP relies heavily on predefined rules and signature-based detection, often resulting in high false positive rates and missed threats. Contextual Defense leverages real-time behavioral analysis to understand the full context of data usage. This includes who is accessing the data, what is being done with it, and where it is being shared.



Smart Classification

Smart classification automatically catalogues all structured and unstructured data in real-time. Sensitive data is precisely detected by inspecting context signals - such as file properties, metadata, and third-party classifications – together with the content of the data.

Contextual factors such as where a file originated, if it has ever been in a specific location, has ever been modified by certain apps, or previously classified by a third-party, are combined to deliver remarkable data classification accuracy.

Say goodbye to false positives.

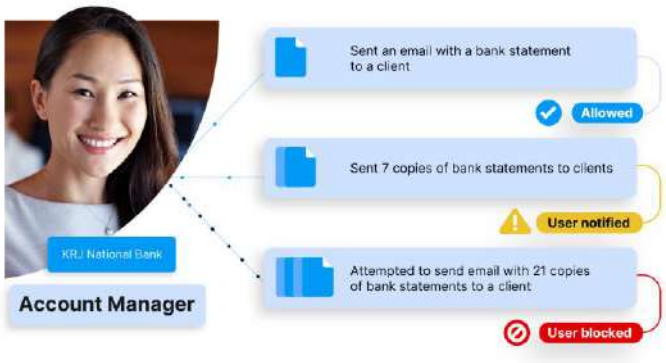


Contextual Risk Analysis

Every data operation your team carries out – such as copying a file or uploading content to a website - is analyzed and scored.

Scoring examines rich context such as time of day, destination, transfer method, and data classification, enabling high risk actions to be immediately spotted, without hindering your team's normal work.

Banish alert fatigue.



Adaptive Security

With contextual defense your security automatically adapts, applying rigorous protection to block dangerous activity without getting in the way of your teams' normal work.

Machine learning tracks how people and your business usually work with data. If contextual risk thresholds are exceeded, dynamic policies allow the user to receive a notification, with override. Continued unusual activity or more serious deviations from normal patterns are automatically blocked.

Eliminate unnecessary interruptions.

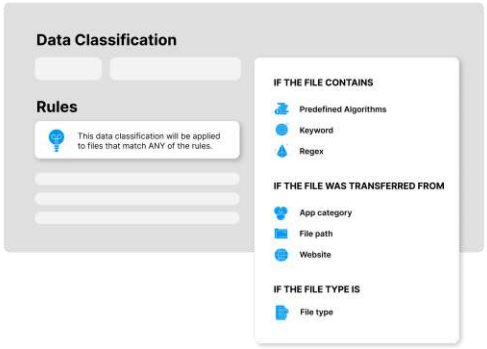
© Copyright All rights reserved, 2024

safetica

How It Works

1 Map and Classify Your Data

Identify, map, and categorize all sensitive information across your data landscape, whether it's in structured databases or unstructured formats like emails or documents. This process spans environments—on-premises, in the cloud, or in hybrid settings—ensuring that data is organized for better management and security.



2 Track user and data activity

Continuously track how data is accessed and used. The Contextual Defense engine observes user actions, application usage, and data movements, building a comprehensive picture of normal behavior.

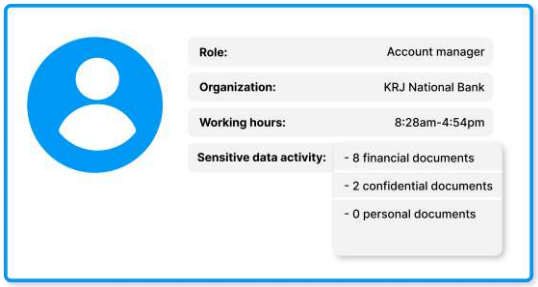
Safetica tracks every data operation.

- Attach file to email
- Send via WhatsApp
- USB upload
- Copies text to GenAI app
- Screen captures technical diagrams and more...



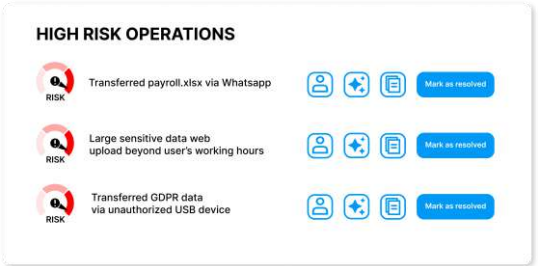
3 Establish baselines

Define what constitutes typical activity for your organization. The engine establishes these baselines to distinguish between usual behavior and anomalies that may signal insider threats or unauthorized data activities.



4 Evaluate Risk in Real-Time

Assess the risk of data actions continuously. For example, if someone attempts to transfer sensitive data externally, the system analyzes the action against historical behavior, the user's role, and the business context to determine if it's a legitimate or suspicious event.



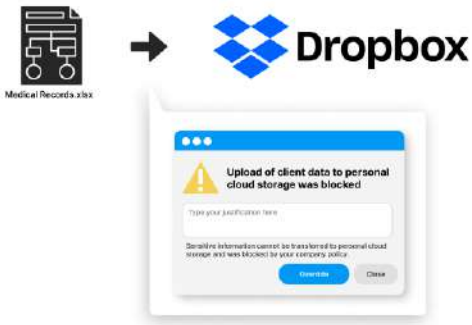
© Copyright All rights reserved, 2024

safetica

How It Works CONTINUED

5 Trigger Adaptive Responses

Respond instantly when a threat is detected. Safetica can enforce tiered responses depending on the risk level of an action, the data involved, and the company's security culture.

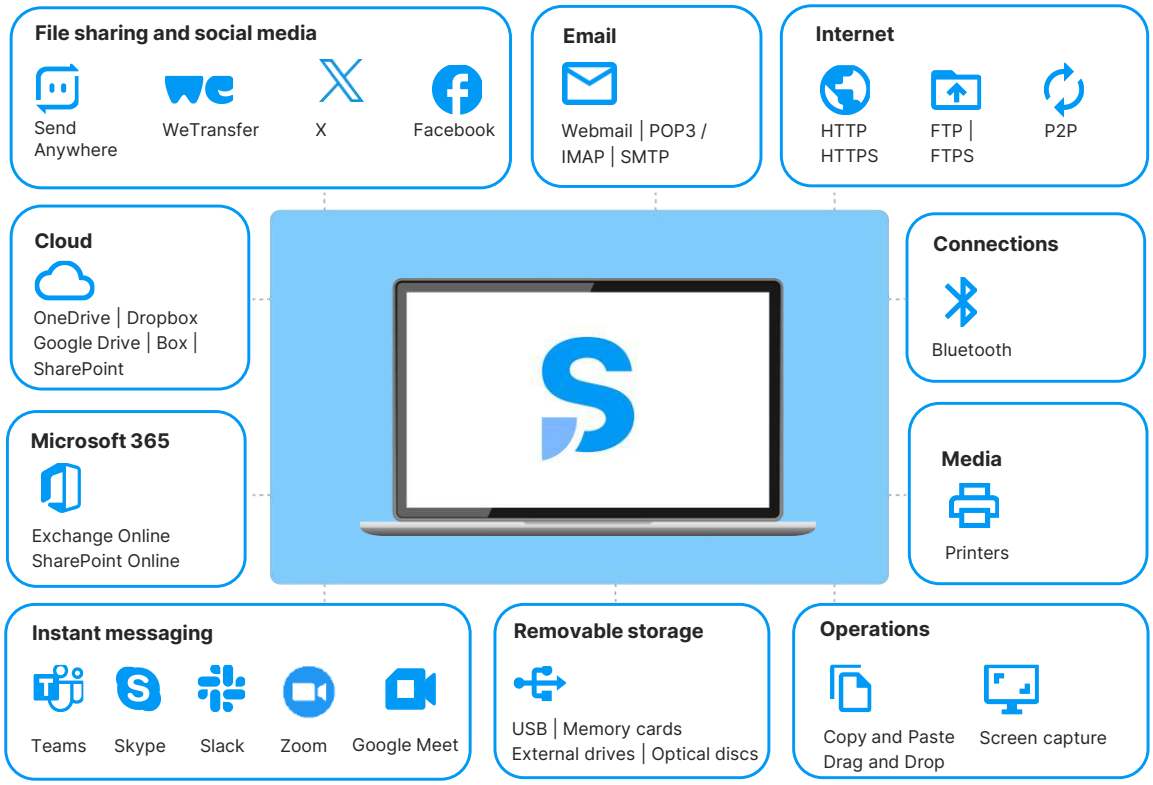


Allow override with a business justification

Safetica can enforce policies such as blocking data transfers to unapproved destinations while allowing employees to override these restrictions for legitimate business reasons. This way, security stays strong without getting in the way of their work.

Data Channels Covered

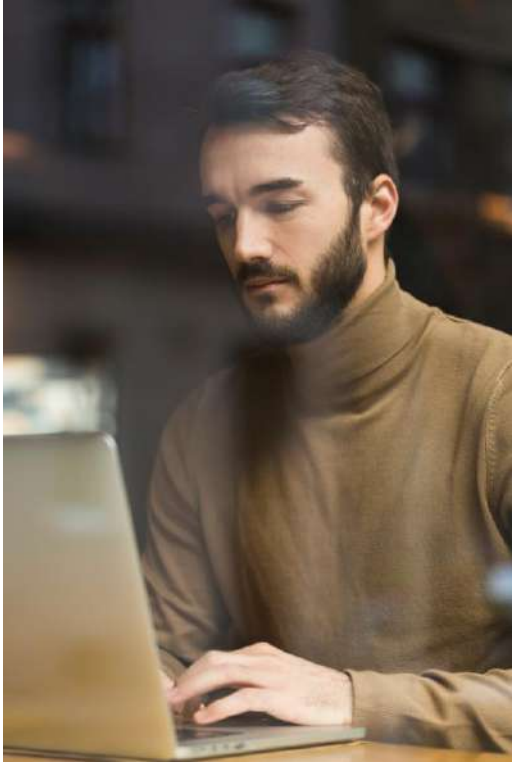
Safetica keeps data protected across a multitude of channels and platforms, ensuring your data is secure wherever it flows.



© Copyright All rights reserved, 2024

safetica

Safetica Essentials
Key Benefits



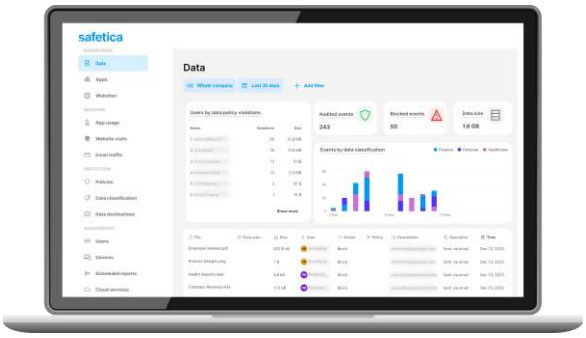
Safetica Essentials audits and classifies all data flows in your organization. It identifies sensitive information and data security risks using content inspection and context awareness. Get a quick overview of what is happening in your work environment in real-time. Better understand all internal activities, processes and data risks to enhance your data security and internal efficiency.

Get insight into data security incidents and regulatory compliance violations to be able to respond and mitigate their impacts

Discover and classify your sensitive data flows in any channel or activity to find out where your data is at risk of loss or theft

Get instant actionable notifications with easy-to-read risk evaluation and incident overview

Get general information about user activities to discover shadow IT – unwanted or unnecessary software, cloud services, devices and peripherals



Safetica Console

Web-based central management console that offers optimized features and seamless user experience:

- Data flow visibility, user activity visibility across email, applications, websites and external devices
- Protected device maintenance, interactive reporting, and data analysis
- Real-time alerts and incident investigation

Data protection starts with data visibility

Identify how company data is used and where the data flow.

- Full Windows and macOS support
- Content inspection and context-aware classification with out-of-box templates
- Data security incident detection
- Easy to upgrade to the full-featured data security platform

© Copyright All rights reserved, 2024

safetica

Safetica Pro
Key Benefits

Safetica Pro identifies risks, prevents human errors and malicious acts, and educates users to protect your data. Adding an intelligent layer to data classification, data loss prevention (DLP) and insider risk management creates a secure environment and supports efficient business operations.



Have full control over sensitive data flows and internal risks based on user behavior and complete content and context analysis

Protect data seamlessly across diverse data storage platforms including clouds, network shares, and Windows & macOS endpoints

Reduce risk and ensure compliance with real-time incident-based user training

Get your data under control, online and offline

Unmatched data visibility across all endpoints, networks, and cloud environments for complete data protection and threat prevention. Safetica uses advanced content classification and OCR for sensitive data detection in image files and scanned documents.

Detect potential threats and analyze internal risks

Respond to threats even before a major incident happens thanks to early discovery of behavior anomalies and data flow risks in your organization.

Set clear policies for all users and data channels

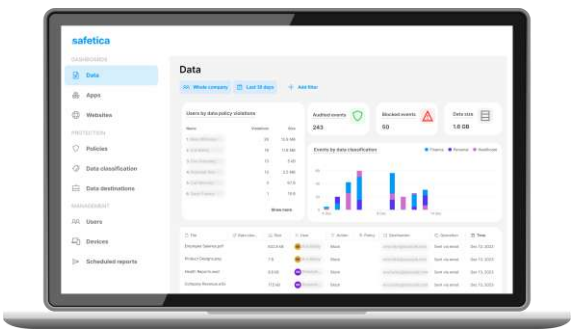
Set up policies for specific groups or individuals. Select the desired workflow with configurable actions from silent auditing, through user notifications to strict blocking.

Empower users to work with sensitive data

Display educational notifications to users when there's a risk of policy violation to let them know or decide. Enforce specific processes to protect the most valuable data.

Built to address the risks of today's work environment

- Safetica smart classification enables content- and context-based approaches
- Pre-defined templates and custom policies
- Flexible DLP actions: record only, notify, justify, or block
- Real-time email notifications
- Multiple hosting options



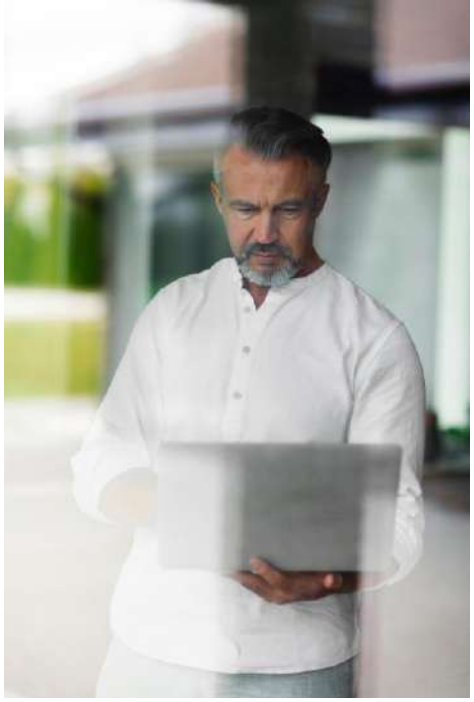
The Safetica Console enables detailed but easy configuration of data classifications, policies, or reports.

© Copyright All rights reserved, 2024

safetica

UEBA
Key Benefits

Knowledge is the first and most important step in understanding your company's workflow, users work habits and productivity. Enrich Safetica product with User and Entity Behavior Analytics add-on to see user activities in detail and uncover their behavior anomalies. Ensure smooth business operations, even when working remotely.



Recognize undesirable user activities
with work activity audit and automated categories for apps used and websites visited by specific users

Get deeper insights into email communication
with records of all incoming and outgoing emails with respect to user privacy

Track changes in user behavior
with detailed overview of user behavior in your organization over time

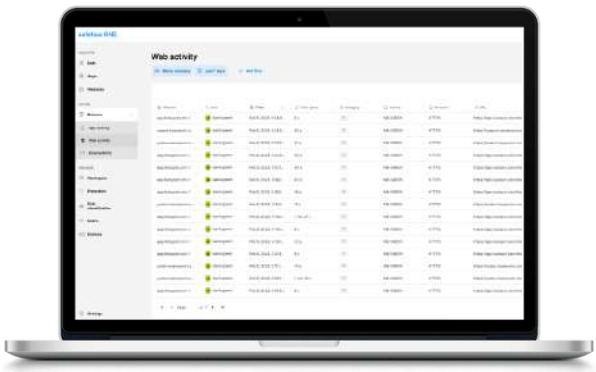
Get comprehensive overview and real-time alerts
about individual user activities, even when working remote, such as via remote desktop etc.

Detect underlying reasons for anomalies

Dig deeper and pinpoint troublesome elements in your environment to address security or business efficiency concerns. Analyze work-related activities of individual users with detailed information. Find out if anyone visits dangerous websites or uses undesirable applications and services.

Maintain transparency in work, even when working remotely

Let top management and department leaders see how their individual reports work. Stay on top of things even when your users work from home or on the go. Prevent security risks and manage users' efficiency by identifying idle workers, and suspicious behavior patterns.



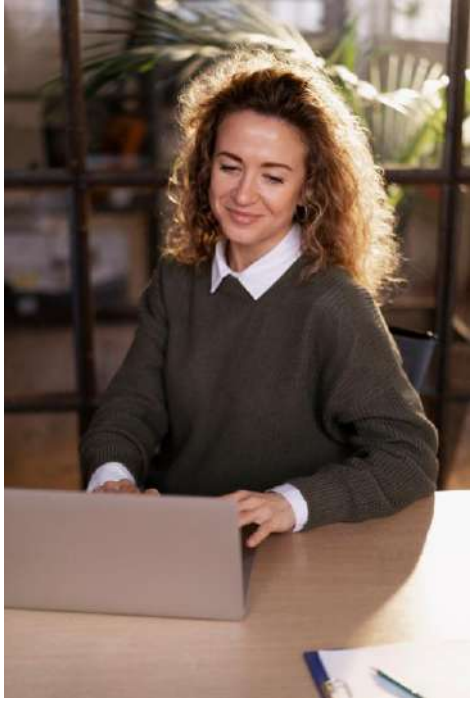
The Safetica Console provides easy-to-understand overview of all possible insider risks. Set up your own overview with customizable views.

© Copyright All rights reserved, 2024

safetica

Cloud Data Protection Key Benefits

Protect your data in the cloud and prevent any unauthorized access or leaks to cloud storage. Secure emailing and file-sharing, safeguard data, and maximize the value of your cloud investment with a dedicated security solution for Microsoft 365.

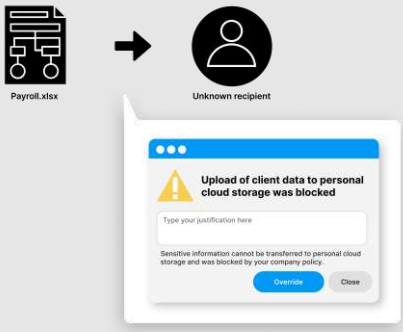


Safetica extends data security to the cloud



Monitor and conduct audits of file uploads and downloads to cloud storage platforms

Safetica can monitor and classify files directly during user operations, such as exports, uploads and downloads, opening files, copying files to a different path, uploading files via web browsers, sending files via email or IM apps, and others.



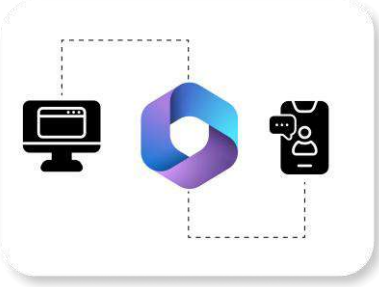
Take real-time action to protect data and educate users

Safetica's flexible protection policies:

- Prevent sensitive data from being stolen or lost
- Inform and instruct users to improve security awareness
- Enable override with a valid justification

Safeguard Remote and Mobile Work with Microsoft 365 Anywhere, Anytime

- Protect, audit, and control access to any data file as you collaborate in Microsoft 365, no matter where the document is stored or who it's shared with.
- Fully leverage the potential of Microsoft 365 cloud applications (OneDrive, Outlook, SharePoint and Teams) on mobile devices
- Automatically apply controls, even when user devices are not within the organization's network.



© Copyright All rights reserved, 2024

safetica

1M+ protected devices

120+ countries

90+ security evangelists

who we are

Safetica is a global software company that provides Data Loss Prevention and Insider Risk Management solutions to organizations of all shapes and sizes. Here at Safetica, we believe everyone deserves to know that their data is safe.

Technology alliances

Awards & achievements

INTELLIGENT DATA SECURITY

@safetica

Book your product demo and sign up for free
www.safetica.com/demo

safetica

What is Active Directory Certificate Services?

- Article
- 03/21/2023
- Applies to:

✅ Windows Server 2025, ✅ Windows Server 2022, ✅ Windows Server 2019, ✅ Windows Server 2016

In this article

1. Issue and manage certificates
2. Key features
3. Benefits
4. Next steps

Active Directory Certificate Services (AD CS) is a Windows Server role for issuing and managing public key infrastructure (PKI) certificates used in secure communication and authentication protocols.

Issue and manage certificates

Digital certificates can be used to encrypt and digitally sign electronic documents and messages as well as for authentication of computer, user, or device accounts on a network. For example, digital certificates are used to provide:

- Confidentiality through encryption.
- Integrity through digital signatures.
- Authentication by associating certificate keys with the computer, user, or device accounts on a computer network.

Key features

AD CS provides the following important features:

- **Certification authorities:** Root and subordinate Certificate Authorities (CAs) are used to issue certificates to users, computers, and services, and to manage certificate validity.
- **Web enrollment:** Web enrollment allows users to connect to a CA with a Web browser in order to request certificates and retrieve certificate revocation lists (CRLs).
- **Online Responder:** The Online Responder service decodes revocation status requests for specific certificates, evaluates the status of these certificates, and sends back a signed response containing the requested certificate status information.
- **Network Device Enrollment Service:** The Network Device Enrollment Service allows routers and other network devices that don't have domain accounts to obtain certificates.

- **TPM key attestation:** Lets the certification authority verify the private key is protected by a hardware-based TPM and that the TPM is one that the CA trusts. TPM key attestation prevents the certificate from being exported to an unauthorized device and can bind the user identity to the device.
- **Certificate Enrollment Policy Web Service:** The Certificate Enrollment Policy Web Service enables users and computers to obtain certificate enrollment policy information.
- **Certificate Enrollment Web Service:** Certificate Enrollment Web Service enables users and computers to perform certificate enrollment through a web service. Together with the Certificate Enrollment Policy Web Service, this enables policy-based certificate enrollment when the client computer isn't a member of a domain or when a domain member isn't connected to the domain.

Benefits

You can use AD CS to enhance security by binding the identity of a person, computer, or service to a corresponding private key. AD CS gives you a cost-effective, efficient, and secure way to manage the distribution and use of certificates. In addition to binding of identities and private keys, AD CS also includes features that allow you to manage certificate enrollment and revocation.

You can use existing endpoint identity information in Active Directory to register certificates, meaning you can have information automatically inserted into certificates. AD CS can also be used to configure Active Directory group policies to designate which users and machines are allowed which types of certificates. Group policy configuration enables role-based or attribute-based access control.

Applications supported by AD CS include Secure/Multipurpose Internet Mail Extensions (S/MIME), secure wireless networks, virtual private network (VPN), Internet Protocol security (IPsec), Encrypting File System (EFS), smart card sign in, Secure Socket Layer/Transport Layer Security (SSL/TLS), and digital signatures.

Next steps

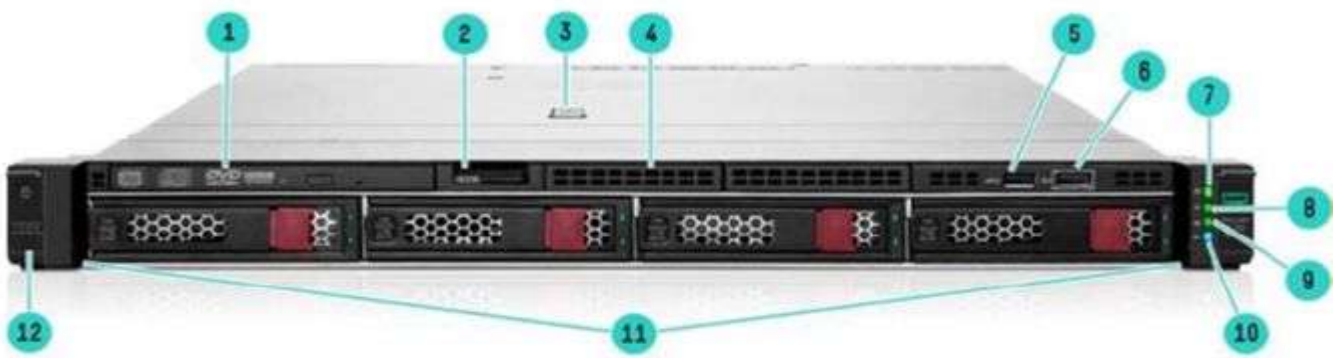
- What is the Certification Authority Role Service?
- Implement and manage Active Directory Certificate Services
- All AD CS role services run on any version
- All AD CS role services can be run on Server Core
- Windows PowerShell Reference for Certificate Services

Overview

HPE ProLiant DL320 Gen11

Are you looking to run edge AI workloads like Computer Vision that need GPU accelerators or distributed data solutions that require expandable storage? The HPE ProLiant DL320 Gen11 is a 1U 1P server with a unique compact and workload-driven modular design that is purpose-built for edge computing delivering exceptional performance at 1P economics and an excellent choice for both virtualized and containerized workloads.

Powered by 4th and 5th Gen Intel® Xeon® Scalable Processors with up to 60 cores and 270W, increased memory capability (up to 2 TB 5200 MT/s), and high-speed PCIe Gen5 supporting up to four single-wide GPUs (or two double-wide), the HPE ProLiant DL320 Gen11 server is a perfect low-cost, 1U 1P, performance solution. The HPE ProLiant Gen11 servers are engineered to optimize IT at the edge with a cloud operating experience, built-in security, and optimized performance for workloads to drive your business forward.

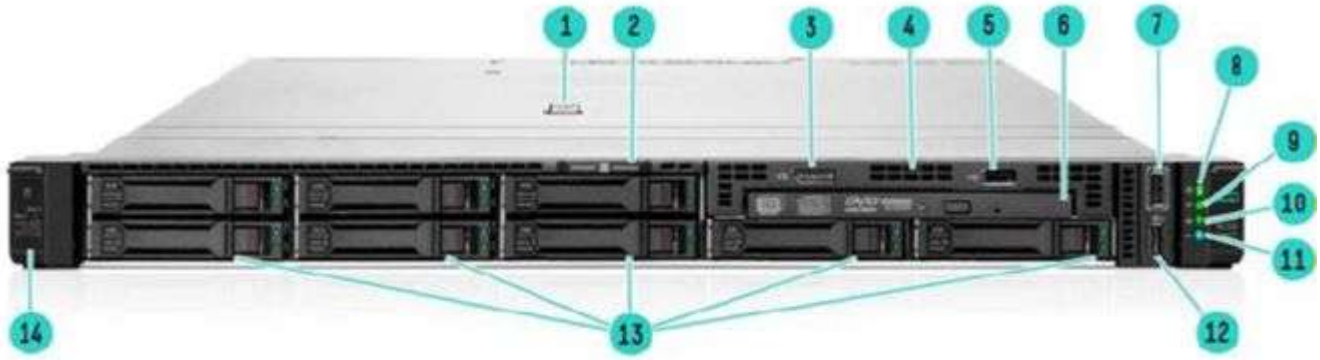


4 LFF Front View - 4 LFF + optional Optical Drive and SAS drives shown

1. Optical Drive (optional - shown)
2. Serial number/iLO information pull tab
3. Quick removal access panel
4. Option Display Port & USB 2.0 port bundle kit (blank shown)
5. USB 3.2 Gen1 port
6. iLO Service Port
7. Power On / Standby button and system power LED
8. Health LED
9. NIC Status[1]
10. Unit ID Button/LED
11. SAS/SATA drive bays
12 Drive support label

Notes: [1] Front NIC LED display doesn't support NIC LED ACT/LINK indication from ALOM/PCIE/FLOM NIC's

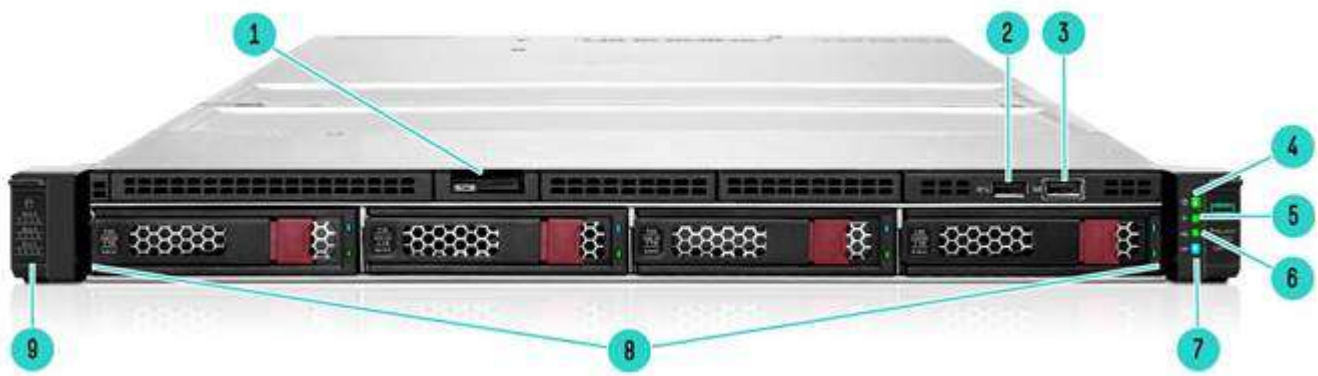
Overview



8SFF Front View - 8 SFF + optional Universal Media Bay, optical Drive, Display Port, USB2.0, and SATA Drive shown

1. Quick removal access panel
2. Serial number/iLO information pull tab
3. Display Port (optional - shown)
4. Universal Media Bay (optional):
   - Option1: Optical drive bay + Display port & USB 2.0 port kit (shown)
   - Option2: 2 SFF x4 Tri-Mode 24G U.3 BC Drive Cage
5. USB 3.2 Gen1 port (optional - shown)
6. Optical Drive (optional- shown)
7. iLO Service Port
8. Power On / Standby button and system power LED
9. Health LED
10. NIC Status [1]
11. Unit ID Button/LED
12. USB 3.2 Gen1 port
13. Drive bays; backplanes options
    - Option1: 8SFF x1 Tri-Mode 24G U.3 BC Backplane
    - Option2: 8SFF x4 Tri-Mode 24G U.3 BC Backplane
14. Drive support label

Notes: [1] Front NIC LED display doesn't support NIC LED ACT/LINK indication from ALOM/PCIE/FLOM NIC's

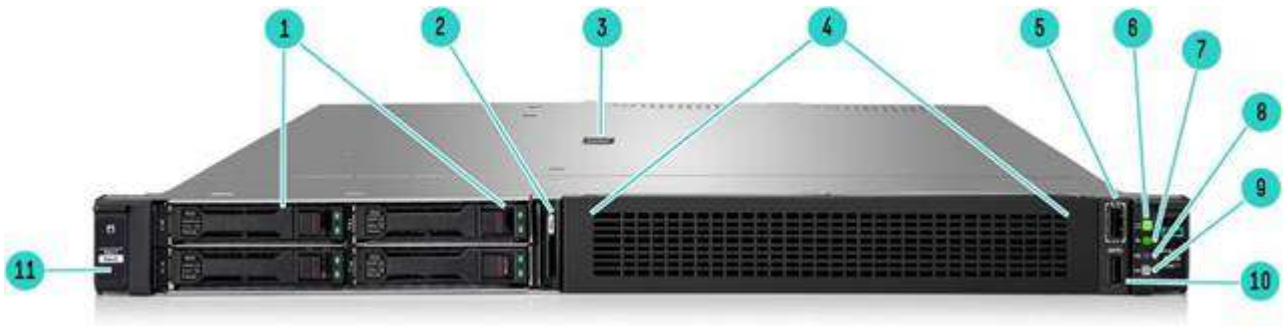


12 LFF Front View - 12 LFF + SAS drives shown

1. Serial number/iLO information pull tab
2. USB 3.2 Gen1 Port
3. iLO Service Port
4. Power On/Standby button and system power LED
5. Health LED
6. NIC Status [1]
7. Unit ID Button/LED
8. SAS/SATA drive bays
9. Drive support label

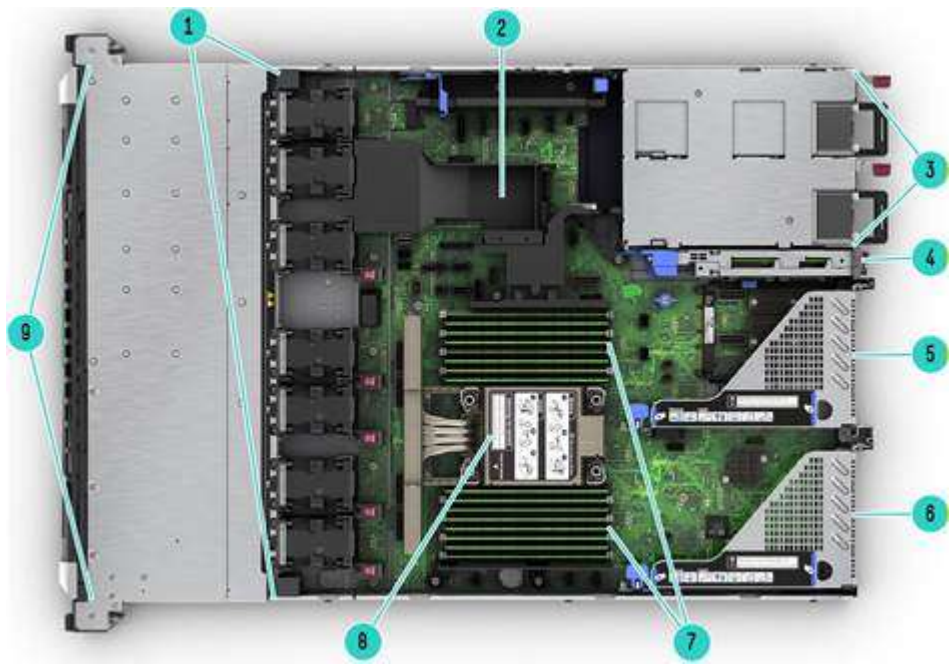
Overview

Notes: [1] Front NIC LED display doesn't support NIC LED ACT/LINK indication from ALOM/PCIE/FLOM NIC's



Front View - GPU CTO Server

1. Drive Bays
   - Option1: 4SFF SAS/SATA/NVMe (Shown)
   - Option2: 8EDSFF E3.s 1T
2. Serial number/iLO information pull tab
3. Quick removal access panel
4. GPU Front Cage
5. iLO Service Port
6. Power On/Standby button and system power LED
7. Health LED
8. NIC Status
9. Unit ID Button/LED
10. USB 3.2 Gen1 Port
11. Drive support label



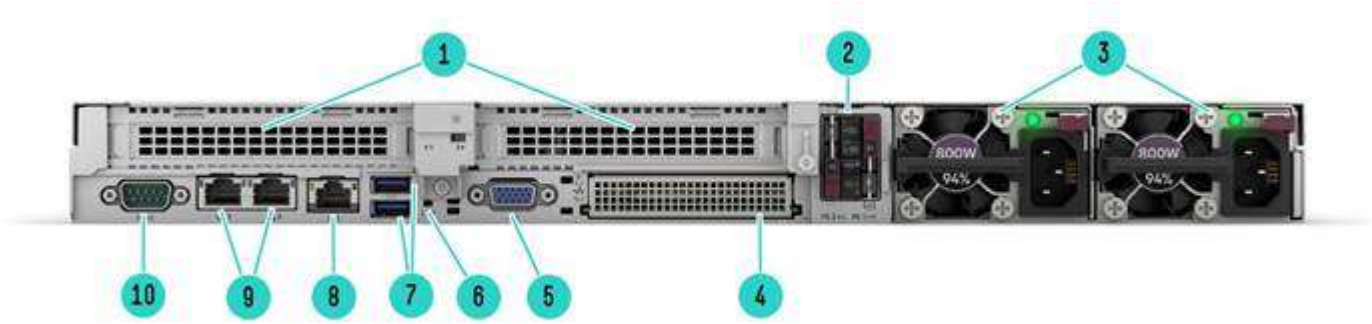
Internal View - Standard for all DL320 Gen11

1. Up to 7 Hot Plug Fans
2. Two (2) M.2 module connectors (under the air baffle- not visible)
3. Redundant Power Supply (1 & 2 as shown)
5. Slot-2 PCIe5 x16 FHHL (Secondary Riser)
6. Slot-1 PCIe5 x16 FHHL (Primary Riser)
7. 16 DDR5 DIMM slots (shown)
8. Heat Sink

Overview

4.	HPE NS204i-u NVMe Hot Plug Boot Optimized Storage Device (optional - shown)	9.	Drive Cage

Notes: 1x Internal USB 3.0 port is not marked.



Rear View - Standard for all DL320 Gen11

1.	Slot-1 & 2 PCIe5 x16 FHHL (Primary and Secondary Riser)	6.	UID Indicator LED
2.	HPE NS204i-u NVMe Hot Plug Boot Optimized Storage Device (optional - shown)	7.	Two (2) USB 3.2 Gen1 ports
3.	Hot-plug Power Supply 1 and 2 (1 +1 Redundancy)	8.	Dedicated iLO management port
4,	OCP 3.0 Slot: x16 PCIe 5.0	9.	Embedded 1GbE x2 network port
5.	Video (VGA) port	10.	Serial port (optional)

Platform Information

Form Factor

- 1U rack

Chassis Types

- 4 LFF (P52765-B21) with optional Optical Disk Drive (use cable P54641-B21), optional Display Port/USB Kit (P48928-B21)
- 8+2 SFF SAS/SATA/NVMe (P52766-B21)
  - Option1: HPE ProLiant DL320 Gen11 8SFF x1 Tri-Mode 24G U.3 BC Backplane Kit (P52743-B21)
  - Option2: HPE ProLiant DL320 Gen11 8SFF x4 Tri-Mode 24G U.3 BC Backplane Kit (P52747-B21)
- 12 LFF (P52767-B21)
- GPU Dense (P61218-B21)
  - Option1: HPE ProLiant DL320 Gen11 8EDSFF x4 NVMe Drive Cage Kit (P61219-B21)
  - Option2: HPE ProLiant DL320 Gen11 4SFF x4 Tri-Mode 24G U.3 BC Drive Cage Kit (P61221-B21)

System Fans

- HPE ProLiant DL3X0 Gen11 1U Standard Fan Kit (P48907-B21), includes five fans.
- HPE ProLiant DL3X0 Gen11 1U High Performance Fan Kit (P48908-B21), includes seven fans.

Overview

- HPE ProLiant DL3X0 Gen11 1U 2P Standard Fan Kit (P54697-B21), includes two fans.

Notes: Generally, the High-Performance fan kit is required when >185W Processors SKUs, High-Performance NVMe drives, or certain backplanes are populated. See notes under each option category or each individual option for specifics.

Standard Features

Processors - One of the following.

The 2nd digit of the processor model number "x4xx" is used to denote the processor generation (i.e., 4=4th generation Intel Scalable Series Processors)

For more information regarding Intel Xeon processors, please see the following **http://www.intel.com/xeon.**

This table covers the public Intel offering only.

Processor Suffix	Description	Offering
H	DB and Analytics	Highest core counts. Database and Analytics usages benefit from DSA and IAA accelerators.
M	Media Transcode	Optimized around AVX frequencies to deliver better performance/watt around Media, AI, and HPC workloads.
N	Network/5G/Edge (High TPT / Low Latency)	Designed for NFV and networking workloads, such as L3 forwarding, 5G UPF, OVS DPDK, VPP FIB router, VPP IPsec, web server/NGINX, vEPC, vBNG, and vCMTS.
S	Storage and HCI	Optimized for Storage UMA use cases with increased UPI Bandwidth for vs. Mainline SKUs.
P	Cloud - IAAS	Designed for cloud IaaS environments to deliver higher frequencies at constrained TDPs.
Q	Liquid Cooling	Liquid-cooled processors with higher frequency and performance at the same TDP.
U	1 Socket Optimized	Optimized for targeted platforms adequately served by the cores, memory bandwidth, and IO capacity available from a single processor.
V	Cloud - SAAS	Optimized for orchestration efficiency that delivers higher core counts and VMs per rack.
Y	Speed Select	Intel® SST-PP increases the base frequency when fewer cores are enabled. Allows greater flexibility, deployment options, and platform longevity.

Notes:

- Covers the Intel public offering only.
- New Built-in Accelerators.
- 1 to 8 socket support
- Intel® Data Streaming Accelerator (DSA)
- Intel® Dynamic Load Balancer (DLB)
- Intel® Quick Assist Technology (QAT)
- Intel® In-Memory Analytics Accelerator (IAA)
- Increased memory bandwidth with 8 channels DDR5, up to 4800 MT/s, 4.0TB maximum RAM per socket.
- Increased I/O bandwidth up to 80 PCIe 5.0 lanes per socket, and new Compute Express Link (CXL).
- Built-in AI Acceleration: Intel® Advanced Matrix Extension (AMX)
- Hardware-enhanced Security: Enhanced Intel® Software Guard Extensions (SGX) - with new cryptographic memory integrity
- Increaed Multi-Socket Bandwidth with new UPI2.0 (up to 16GT/s) with maximum 4 UPI Links

Standard Features

- New FlexBus I/O Interface PCIe5.0 + CXL
- [1] The 4th Generation Intel® Xeon® Scalable Processors are featured with Intel Speed Select Technology (SST) for Infrastructure as as Service, Networking and Virtualized environments worklaods. The SST includes,
- SST- Performace Profile
- SST- Base Frequency
- SST- Core Power
- SST- Turo Frequency
- Default setting in ROM-Based Setp Utility (RBSU) as shown.

Intel® SST Features	RBSU Options	Granular Control over CPU Performance	Default Setting
SST- Performance Profile	Dynamic Intel® Speed Select Technology - Performance Profile	Alllows the CPU to run in one of three performance profiles	CPU hardware-based. Enabled by default
SST-Base Frequency	Intel® Speed Select Technology - Base Frequency	Enables some CPU cores to run at a higher base frequency in return for other cores running at a lower base frequency	Disabled by default
SST-Core Power	Intel® Speed Select Technology - Core Power	Allows software to prioritize with cores will receive excess power after satisfying minimum requirements	Disabled by default
Intel SST Turbo Frequency	Intel® Turbo Boost Technology	Allows software-selected cores to achieve a higher max turbo frequency by reducing other cores' max turbo frequency	Enabled by default

5th Generation Intel® Xeon® Scalable Processor Family (Platinum)								
Intel Xeon Models	**Base Speed**	**Cores**	**L3 Cache**	**Power**	**UPI Links**	**DDR5**	**SGX Enclave Size**	**Die**
Platinum 8581V Processor	2.0 GHz	60	300 MB	270 W	0	4800 (MT/s)	128 GB	XCC

Notes:

- Processors with TDP greater than 185W require High Performance Heatsink (P52756-B21) and High-Performance Fan Kit (P48908-B21).
- Intel® Speed Select enabled processors: 8581V.

Standard Features

5th Generation Intel® Xeon® Scalable Processor Family (Gold 6)								
Intel Xeon Models	**Base Speed**	**Cores**	**L3 Cache**	**Power**	**UPI Links**	**DDR5**	**SGX Enclave Size**	**Die**
Gold 6526Y Processor	2.8 GHz	16	37.5 MB	195W	3	5200 MT/s	128 GB	MCC
Gold 6530 Processor	2.1 GHz	32	160 MB	270W	3	4800 MT/s	128 GB	XCC
Gold 6534 Processor	3.9 GHz	8	22.5 MB	195W	3	4800 MT/s	128 GB	MCC
Gold 6538N Processor	2.1 GHz	32	60.0 MB	205W	3	5200 MT/s	128 GB	MCC
Gold 6538Y+ Processor	2.2 GHz	32	60.0 MB	225W	3	5200 MT/s	128 GB	MCC
Gold 6542Y Processor	2.9 GHz	24	60.0 MB	250W	3	5200 MT/s	128 GB	MCC
Gold 6544Y Processor	3.6 GHz	16	45.0 MB	270W	3	5200 MT/s	128 GB	MCC
Gold 6548N Processor	2.8 GHz	32	60.0 MB	250W	3	5200 MT/s	128 GB	MCC
Gold 6548Y+ Processor	2.5 GHz	32	60.0 MB	250W	3	5200 MT/s	128 GB	MCC
Gold 6554S Processor	2.2 GHz	36	180 MB	270W	4	5200 MT/s	128 GB	XCC

Notes:

- Processors with a TDP greater than 185W require High Performance Heatsink (P52756-B21) and High-Performance Fan Kit (P48908-B21).
- Intel® Speed Select enabled processors: 6526Y, 6538N, 6538Y+, 6542Y, 6544Y, 6548N, 6548Y+ and 6554S.

5th Generation Intel® Xeon® Scalable Processor Family (Gold 5)								
Intel Xeon Models	**Base Speed**	**Cores**	**L3 Cache**	**Power**	**UPI Links**	**DDR5**	**SGX Enclave Size**	**Die**
Gold 5512U Processor	2.1 GHz	28	52.5 MB	185 W	0	4800 MT/s	128 GB	MCC
Gold 5515+ Processor	3.2 GHz	8	22.5 MB	165 W	3	4800 MT/s	128 GB	MCC
Gold 5520+ Processor	2.2 GHz	28	52. 5 MB	205 W	3	4800 MT/s	128 GB	MCC

Notes:

- Processors with TDP greater than 185W require High Performance Heatsink (P52756-B21) and High-Performance Fan Kit (P48908-B21).
- Intel® Speed Select enabled processors: NA.

5th Generation Intel® Xeon® Scalable Processor Family (Silver)								
Intel Xeon Models	**Base Speed**	**Cores**	**L3 Cache**	**Power**	**UPI Links**	**DDR5**	**SGX Enclave Size**	**Die**
Silver 4509Y Processor	2.6 GHz	8	22.5 MB	125W	2	4400 MT/s	64 GB	EE LCC
Silver 4510 Processor	2.4 GHz	12	30.0 MB	150W	2	4000 MT/s	64 GB	EE LCC
Silver 4514Y Processor	2.0 GHz	16	30.0 MB	150W	2	4400 MT/s	64 GB	MCC
Silver 4516Y+ Processor	2.2 GHz	24	45.0 MB	185W	2	4400 MT/s	64 GB	MCC

Notes:

Standard Features

– Intel® Speed Select enabled processors: 4509Y, 4514Y and 4516Y+.

5th Generation Intel® Xeon® Scalable Processor Family (Bronze)								
Intel Xeon Models	**Base Speed**	**Cores**	**L3 Cache**	**Power**	**UPI Links**	**DDR5**	**SGX Enclave Size**	**Die**
Bronze 3508U Processor[1]	2.1 GHz	8	22.5 MB	125W	N/A	4400 MT/s	64 GB	EE LCC

Notes:

– Intel® Speed Select enabled processors: N.A.
– PCIe4.0 only

4th Generation Intel® Xeon® Scalable Processor Family (Gold)								
Intel Xeon Models	**Base Speed**	**Cores**	**L3 Cache**	**Power**	**UPI Links**	**DDR5**	**SGX Enclave Size**	**Die**
Gold 6426Y Processor	2.5 GHz	16	37.5 MB	185 W	3	4800 MT/s	128 GB	MCC
Gold 6442Y Processor	2.6 GHz	24	60 MB	225 W	3	4800 MT/s	128 GB	MCC
Gold 6448Y Processor	2.1 GHz	32	60 MB	225 W	3	4800 MT/s	128 GB	MCC
Gold 6434 Processor	3.7 GHz	8	22.5 MB	195 W	3	4800 MT/s	128 GB	MCC
Gold 6444Y Processor	3.6 GHz	16	45 MB	270 W	3	4800 MT/s	128 GB	MCC
Gold 6430 Processor	2.1 GHz	32	60 MB	270 W	3	4800 MT/s	128 GB	XCC
Gold 6438Y+ Processor	2.0 GHz	32	60 MB	205 W	3	4800 MT/s	128 GB	MCC
Gold 6414U Processor	2.0 GHz	32	60 MB	250 W	0	4800 MT/s	128 GB	XCC
Gold 6438N Processor	2.0 GHz	32	60 MB	205 W	3	4800 MT/s	128 GB	MCC
Gold 6421N Processor	1.8 GHz	32	60 MB	185 W	0	4400 MT/s	128 GB	MCC
Gold 6454S Processor	2.2 GHz	32	60 MB	270 W	4	4800 MT/s	128 GB	XCC
Gold 5415+ Processor	2.9 GHz	8	22.5 MB	150 W	3	4400 MT/s	128 GB	MCC
Gold 5418Y Processor	2.0 GHz	24	45 MB	185 W	3	4400 MT/s	128 GB	MCC
Gold 5420+ Processor	2.0 GHz	28	52.5 MB	205 W	3	4400 MT/s	128 GB	MCC
Gold 5412U Processor	2.1 GHz	24	45 MB	185 W	0	4400 MT/s	128 GB	MCC
Gold 5411N Processor	1.9 GHz	24	45 MB	165 W	0	4400 MT/s	128 GB	MCC
Gold 5418N Processor	1.8 GHz	24	45 MB	165 W	3	4000 MT/s	128 GB	MCC
Gold 5416S Processor	2.0 GHz	16	30 MB	150 W	3	4400 MT/s	128 GB	MCC

Notes:

– Processors with TDP greater than 185W require High Performance Heatsink (P52756-B21) and High-Performance Fan Kit (P48908-B21).
– 96 GB 4800 MT/s Memory cannot be selected if HBM or MCC die

Standard Features

4th Generation Intel® Xeon® Scalable Processor Family (Silver)								
Intel Xeon Models	**Base Speed**	**Cores**	**L3 Cache**	**Power**	**UPI Links**	**DDR5**	**SGX Enclave Size**	**Die**
Silver 4410Y Processor	2.0 GHz	12	30 MB	150 W	2	4000 MT/s	128 GB	MCC
Silver 4416+ Processor	2.0 GHz	20	37.5 MB	165 W	2	4000 MT/s	128 GB	MCC

Notes: 96 GB 4800 MT/s Memory cannot be selected if HBM or MCC die

4th Generation Intel® Xeon® Scalable Processor Family (Bronze)								
Intel Xeon Models	**Base Speed**	**Cores**	**L3 Cache**	**Power**	**UPI Links**	**DDR5**	**SGX Enclave Size**	**Die**
Bronze 3408U Processor	1.8 GHz	8	22.5 MB	125 W	0	4000 MT/s	64 GB	MCC

Notes: 96 GB 4800 MT/s Memory cannot be selected if HBM or MCC die

Chipset

Intel C741 Chipset

Notes: For more information regarding Intel® chipsets, please see the following URL:
https://www.intel.com/content/www/us/en/products/chipsets/server-chipsets.html

System Management Chipset

HPE iLO 6 ASIC

Notes: Read and learn more in the **iLO QuickSpecs**.

Memory

Type	HPE DDR5 Smart Memory Registered (RDIMM)
DIMM Slots Available	16 DIMM Slots, 8 channels, 2 DIMMs per channel
Maximum capacity (RDIMM)	2.04 TB 16 x 128 GB RDIMM @ 4800 MT/s

Notes:

All processors support up to 2TB memory.
- The maximum memory speed is limited by the processor selection.
- To realize the performance memory capabilities listed in this document, HPE DDR5 Smart Memory is required.
- For additional information, please visit the **HPE Memory QuickSpecs and Technical White Papers** or

Standard Features

HPE DDR5 Smart Memory QuickSpecs.

Memory Protection

Advanced ECC

Advanced ECC uses -device data correction to detect and correct single and all multibit errors within a single DRAM chip.

Online Spare

Memory online spare mode detects a degrading rank and switches operation to the spare rank.

Notes: For more information, see our **Memory RAS feature technical whitepaper**.

Network Controller

The HPE ProLiant DL320 Gen11 server offers the customer Embedded Broadcom BCM5720 Ethernet 1Gb 2-port BASE-T Adapter standard with the option to upgrade with a variety of networking options.

Notes: Support document and downloads including firmware and drivers for the Broadcom BCM5720 Ethernet 1Gb 2-port BASET LOM Adapter can be downloaded from the supplier's support and services **webpage**.

PCIe Expansion Slots

Expansion Slots #	Technology	Bus Width	Connector Width	Slot Form Factor
1 (Primary Riser)	PCIe 5.0	x16	x16	FHHL
2 (Secondary Riser)	PCIe 5.0	x16	x16	FHHL

Notes: Slot-1, by default, is with the riser.

OCP Expansion Slots

Expansion Slots #	Technology	Bus Width	Connector Width
1 (OCP 3.0) with Wake on LAN (WoL) and iLO sideband (shared NIC) capability	PCIe 5.0	x16	x16

Internal Storage Devices

- **Optical Drive**
  Available on 8 SFF and 4 LFF CTO Servers as an option (DVD-ROM or DVD-RW)

- **Hard Drives**
  None shipped standard

Graphics

Integrated Video Standard

- Video modes up to 1920 x 1200@60Hz (32 bpp)
- 16MB Video Memory

Maximum Storage

Storage	Capacity	Configuration
Hot Plug LFF SAS HDD	80.0 TB (4 x 20 TB)	HPE ProLiant DL320 Gen11 4 LFF CTO Server (P52765-B21)
Hot Plug LFF SATA HDD	80.0 TB (4 x 20 TB)	
Hot Plug LFF SAS SSD	30.72 TB (4 x 7.68 TB)	
Hot Plug LFF SATA SSD	30.72 TB (4 x 7.68 GB)	
Hot Plug SFF SAS HDD	24.0 TB (8+2 x 2.4 TB)	HPE ProLiant DL320 Gen11 8 SFF CTO Server with optional 2 SFF cage on UMB (P52766-B21)
Hot Plug SFF SATA HDD	20.0 TB (8+2 x 2.0 TB)	
Hot Plug SFF SAS SSD	153.0 TB (8+2 x 15.3 TB)	
Hot Plug SFF SATA SSD	76.8 TB (8+2 x 7.68 TB)	
Hot Plug LFF SAS HDD	240.0 TB (12 x 20 TB)	HPE ProLiant DL320 Gen11 12 LFF CTO Server (P52767-B21)
Hot Plug LFF SATA HDD	240.0 TB (12 x 20 TB)	
Hot Plug LFF SAS SSD	92.16 TB (12 x 7.68 TB)	
Hot Plug LFF SATA SSD	92.16 TB (12 x 7.68 GB)	
EDSFF NVMe SSD	61.44 TB (4 x 15.36)	HPE DL320 Gen11 GPU CTO Server (P61218-B21) with 8EDSFF Drive Cage (P61219-B21)
Hot Plug SFF SAS HDD	9.6 TB (4 x 2.4 TB)	HPE DL320 Gen11 GPU CTO Server (P61218-B21) with 4SFF Drive Cage (P61221-B21)
Hot Plug SFF SATA HDD	8.0 TB (4 x 2.0 TB)	
Hot Plug SFF SAS SSD	61.2 TB (4 x 15.3 TB)	
Hot Plug SFF SATA SSD	30.72 TB (4 x 7.68 TB)	
Hot Plug SFF NVMe PCIe SSD	153.6 TB (8+2 x 15.36 TB)	
M.2 NVMe SSD	3.84 GB (2 x 1.92 GB)	On board with any of the available configurations
M.2 SATA SSD	960 GB (2x 480 GB)	

Storage Controllers

NVMe Boot Devices

- HPE NS204i-u NVMe Hot Plug Boot Optimized Storage Device (P48183-B21)
- HPE ProLiant DL320 Gen11 NS204i-u NVMe Boot Kit (P52786-B21)

Notes: N204i-u Boot device includes 2x 480GB M.2 NVMe SSDs, with preconfigured hardware RAID1. For additional information, please see the **HPE OS Boot Device QuickSpecs**

Standard Features

Software RAID Controllers

The available Gen11 controllers are depicted below.

Intel® VROC SATA for HPE ProLiant Gen11

Notes:

Embedded Intel® VROC SATA for HPE ProLiant Gen11, with 14 SATA ports (10-ports accessible),
Intel® VROC for HPE ProLiant Gen11 is an enterprise, hybrid Software RAID solution specifically designed for
SSDs.
Intel® VROC is a software-based solution utilizing Intel® CPU to RAID or HBA direct connected drives.
RAID Support- 0/1/5/10.
Windows and Linux OS support.
Host Tools- Windows GUI/CLI, Linux CLI.
UEFI Support- HII Utility, OBSE.
iLO Support- IML, Alert, SNMP, AHS.
iLO Redfish- Redfish Read.
Intel® VROC SATA for HPE ProLiant Gen11 will operate in UEFI mode only. For legacy support additional storage controller will be needed.
Intel® VROC SATA is off by default and must be enabled.
8SFF: Separates 8SFF to two drive groups (BOX1 Bay1-4 and Bay5-8) with Intel VROC SATA RAID configuration.
2SFF: 2SFF (BOX2 Bay1 and Bay-2) is one drive group with Intel VROC SATA RAID configuration.
4LFF: 4LFF (BOX1 Bay1-4) is a one drive group with Intel VROC SATA RAID configuration.
10LFF: Each LFF BOX cage is separated drive groups (BOX1, BOX3 & BOX5) with Intel VROC SATA RAID configuration.
Intel® VROC NVMe for HPE ProLiant Gen11

Intel® Virtual RAID on CPU (Intel® VROC) Premium FIO Software for HPE

Intel® Virtual RAID on CPU (Intel® VROC) Standard Software FIO for HPE

In HPE ProLiant Gen11 servers, when secure boot is enabled, Intel Virtual RAID on CPU (VROC) 8.0 Out-of-Band (OOB) management does not function with Linux kernel version 5.4 (or later). VROC OOB will not respond to any PLDM (over-MCTP-over-PCIe) requests from iLO (BMC). VROC Redfish resources will not function (e.g., Redfish actions); therefore, VROC over Redfish management is not available. This is due to a new security feature in Linux kernel version 5.4 (or later).

For more information, please visit **Customer Advisory Document ID: a00128934en_us** at HPE Support Center.

Notes:

All models feature 4 x8 PCIe 5.0 connectors per socket for NVMe connectivity, provides support for up to 8 direct attach x4 NVMe bays.
Only supported on SFF models.
Intel® VROC for HPE ProLiant Gen11 is an enterprise, hybrid Software RAID solution specifically.
designed for NVMe SSDs connected directly to the CPU. Intel® VROC is a software-based solution.
utilizing Intel® CPU to RAID or HBA direct connected drives.
Intel® Virtual RAID on CPU Standard for RAID 0/1/10 (S0E37A/S0E38AAE) or Premium SKU for RAID 0/1/5/10 (R7J57A/R7J59AAE) must be ordered to enable RAID support.
Windows, Linux, VMware OS support.
Host Tools- Windows GUI/CLI, Linux CLI.
UEFI Support- HII Utility, OBSE.
Active health monitoring of NVMe M.2 drives requires use of SMART tools.
Intel® VROC NVMe for HPE ProLiant Gen11 will operate in UEFI mode only. For legacy support an additional Tri-Mode controller will be needed.

Standard Features

For NVMe SSDs only, no PCIe card support.
Intel VROC for HPE ProLiant QuickSpecs

Essential RAID Controllers

- HPE Smart Array E208e-p SR Gen10 Controller

Performance RAID Controllers

- HPE MR216i-p Gen11 x16 Lanes without Cache PCI SPDM Plug-in Storage Controller
- HPE MR216i-o Gen11 x16 Lanes without Cache OCP SPDM Storage Controller
- HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller
- HPE MR416i-o Gen11 x16 Lanes 8GB Cache OCP SPDM Storage Controller
- HPE MR416i-p Gen11 x16 Lanes 8GB Cache PCI SPDM Plug-in Storage Controller
- HPE SR932i-p Gen11 x32 Lanes 8GB Wide Cache PCI SPDM Plug-in Storage Controller[1,2]

Notes:
[1]Requires x16 physical and electrical riser slot.
[2]If second controller is required, must select secondary FH riser. For additional details, please see:
HPE Compute MR Gen11 Controllers QuickSpecs
HPE Compute SR Gen11 Controllers QuickSpecs

Power Supply

- HPE 500W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
  Notes: Available with 94% efficiency.
- HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
  Notes: Available with 94% efficiency.
- HPE 1000W Flex Slot Titanium Hot Plug Low Halogen Power Supply Kit
  Notes: Available with 96% efficiency.
- HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
  Notes:

  Available with 94% efficiency.

  1600W Power supplies only support high line voltage (200 VAC to 240 VAC).
- HPE 1800-2000W Flex Slot Titanium Hot Plug Low Halogen Power Supply Kit

  Notes: Available with 96% efficiency.

HPE Flexible Slot (Flex Slot) Power Supplies share a common electrical and physical design that allows for hot plug, tool-less installation into HPE ProLiant Gen11 Performance Servers. Flex Slot power supplies are certified for high-efficiency operation and offer multiple power output options, allowing users to "right-size" a power supply for specific server configurations. This flexibility helps to reduce power waste, lower overall energy costs, and avoid "trapped" power capacity in the data center.

Standard Features

All pre-configured servers ship with a standard 6-foot IEC C-13/C-14 jumper cord (A0K02A). This jumper cord is also included with each standard AC power supply option kit. Please check the **ProLiant Power Cables** web page if a different power cord is required.

To review the power requirements for your selected system, please use the **HPE Power Advisor Tool.**

For information on power specifications and technical content, visit **HPE Server power supplies.**

European Union Erp Lot 9 Regulation

Beginning on January 1st, 2024, units sold into the European Union (EU), European Economic Area (EEA), the United Kingdom, or Switzerland must include more efficient AC power supplies: 94% for multi-output and 96% for single-output. HPE Flexible Slot power supplies are single-output, and part numbers 865438-B21, P03178-B21, and P44712-B21 are 96% efficient, thus meeting requirements.
HPE is on target to fulfil compliant systems ahead of time and will begin enforcing these requirements in advance to satisfy requests with the current power supplies by the set deadline.

Interfaces

Serial	1 port - Optional
Video	1 Front - Display port (optional) 1 Rear - VGA port (standard on all models) **Notes:** If both ports are used simultaneously, they will display the same image.
Network Ports	None. Choice of OCP or standup card, supporting a wide arrange of NIC adapters
HPE iLO Remote Mgmt Port	1 GbE Dedicated
Front iLO Service Port	1 standard
USB	4 standard on all models: 1 front(USB 3.2 Gen1), 2 rear (USB 3.2 Gen1), 1 internal (USB 2.0) +1 optional USB 2.0 front in 4LFF & 8SFF **Notes:** USB 3.2 Gen1 is referred to as USB 3.0 in some marketing documents.

Operating Systems and Virtualization Software Support for HPE Servers

HPE servers are designed for seamless integration with partner Operating Systems and Virtualization Software. By collaborating closely with our partners, we ensure that their products are optimized, certified, and fully supported within your HPE server environment.

Access the certified and supported servers for each of the OS and Virtualization software: **HPE Servers Support & Certification Matrices**

HPE Server UEFI/Legacy ROM

Unified Extensible Firmware Interface (UEFI) is an industry-standard that provides better manageability and more secured configuration than the legacy ROM while interacting with your server at boot time. HPE ProLiant Gen11 servers have a UEFI Class 3 implementation to support UEFI Mode.

Notes: The UEFI System Utilities tool is analogous to the HPE ROM-Based Setup Utility (RBSU) of legacy BIOS. For more information, please visit **http://www.hpe.com/servers/uefi.**

UEFI enables numerous new capabilities specific to HPE ProLiant servers, such as:

- Secure Boot and Secure Start enabled for enhanced security.
- Embedded UEFI Shell
- Operating system-specific functionality
- Mass Configuration Deployment Tool using iLO RESTful API that is Redfish API Conformant
- Support for > 2.2 TB (using GPT) boot drives.
- PXE boot support for IPv6 networks
- USB 3.0 Stack
- Workload Profiles for simple performance optimization

UEFI Boot Mode only:

- TPM 2.0 Support
- iSCSI Software Initiator Support.
- NVMe Boot Support
- HTTP/HTTPs Boot support as a PXE alternative.
- Platform Trust Technology (PTT) can be enabled.
- Boot support for option cards that only support a UEFI option ROM.

Notes:

Enabling TPM 2.0 no longer requires TPM module option kit for Gen11. It is an embedded feature.
For UEFI Boot Mode, the boot environment and OS image installations should be configured properly to support UEFI.

Industry Standard Compliance

- ACPI 6.3 Compliant
- PCIe 5.0 Compliant
- Wake on LAN (WoL) Support
- Microsoft® Logo certifications
- PXE Support
- VGA
- Display Port

Notes: This support is on the optional Universal Media Bay.

- USB 3.2 Gen1 Compliant
- USB 2.0 Compliant (via Universal Media Bay)

Notes: This support is on the optional Universal Media Bay.

Standard Features

- Energy Star 4.0
- SMBIOS 3.2
- Redfish API
- IPMI 2.0
- Secure Digital 4.0
- Embedded TPM 2.0 Support

Notes: Enabling TPM 2.0 no longer requires TPM module option kit for Gen11 since it is an embedded feature. TPM is disabled on shipments to China.

- Advanced Encryption Standard (AES)
- Triple Data Encryption Standard (3DES)
- SNMP v3
- TLS 1.2
- DMTF Systems Management Architecture for Server Hardware Command Line Protocol (SMASH CLP)
- Active Directory v1.0
- ASHRAE A3/A4

Notes: For additional technical and thermal details regarding ambient temperature, humidity, and feature support, please visit **https://www.hpe.com/support/ASHRAEGen11**

- UEFI (Unified Extensible Firmware Interface Forum) 2.7

Embedded Management

HPE Integrated Lights-Out (HPE iLO)

Monitor your servers for ongoing management, service alerting, reporting, and remote management with HPE iLO.

Learn more at **http://www.hpe.com/info/ilo**.

UEFI

Configure and boot your servers securely with industry-standard Unified Extensible Firmware Interface (UEFI).

Learn more at **http://www.hpe.com/servers/uefi**.

Intelligent Provisioning

Hassle-free server and OS provisioning for one or more servers with Intelligent Provisioning.

Learn more at **http://www.hpe.com/servers/intelligentprovisioning**.

Standard Features

iLO RESTful API

iLO RESTful API is Redfish API conformance and offers simplified server management automation, such as configuration and maintenance tasks based on modern industry standards. Learn more at **http://www.hpe.com/info/restfulapi**.

OpenBMC Support
OpenBMC Capable through iLO6 Transfer of Ownership Process.
Learn more at **OpenBMC Support**

Server Utilities

Active Health System

The HPE Active Health System (AHS) is an essential component of the iLO management portfolio that provides continuous, proactive health monitoring of HPE servers. Learn more at **http://www.hpe.com/servers/ahs**.

Smart Update

Keep your servers up to date with the HPE Smart Update solution by using Smart Update Manager (SUM) to optimize the firmware and driver updates of the Service Pack for ProLiant (SPP).

Learn more at **https://www.hpe.com/us/en/servers/smart-update.html**.

iLO Amplifier Pack

Designed for large enterprise and service provider environments with hundreds of HPE servers, the iLO Amplifier Pack is a free, downloadable open virtual application (OVA) that delivers the power to discover, inventory, and update Gen8, Gen9, Gen10, and Gen10 Plus HPE servers at unmatched speed and scale. Use an iLO Advanced License to unlock full capabilities.

Learn more at **http://www.hpe.com/servers/iLOamplifierpack**.

RESTful Interface Tool

RESTful Interface tool (iLOREST) is a single scripting tool to provision using iLO RESTful API to discover and deploy servers at scale. Learn more at **http://www.hpe.com/info/resttool**.

Scripting Tools

Provision one too many servers, using your scripts to discover and deploy with Scripting Tool (STK) for Windows and Linux or Scripting Tools for Windows PowerShell. Learn more at **http://www.hpe.com/servers/powershell**.

HPE OneView Standard

HPE OneView is an on-premise, multi-generational server monitoring and management solution. HPE OneView Standard can be used for inventory, health monitoring, alerting, and reporting without additional fees. Customers can upgrade their management experience with an HPE OneView Advanced license, all

provided by the same tool. Learn more at **http://www.hpe.com/info/oneview**.

HPE GreenLake for Compute Ops Management

HPE is intelligently transforming compute management with an intuitive cloud operating experience through HPE GreenLake cloud platform to streamline and secure operations from edge-to-cloud. Automated key lifecycle tasks, for onboarding, updating, managing, and monitoring HPE servers, brings agility and greater efficiencies to wherever compute devices reside via a unified single browser-based interface. Manage single locations or multiple, distributed sites. Keep tens to thousands of servers secure with batch policy controls and automated updates.
Compute Ops Management is cloud-native software that is continually updated with new services, features, patches, and fixes. The management application resides in the HPE GreenLake cloud platform (access via **https://console.greenlake.hpe.com**) and leverages the HPE GreenLake architecture, security, and unified operations.
A 3-year subscription to HPE GreenLake for Compute Ops Management is added by default when ordering an HPE ProLiant Gen11 rack, tower, or micro server.

For more information, visit the HPE GreenLake for Compute Ops Management QuickSpecs:
https://www.hpe.com/psnow/doc/a50004263enw

Security

- UEFI Secure Boot and Secure Start support.
- Tamper-free updates - components digitally signed and verified.
- Immutable Silicon Root of Trust
- Ability to rollback firmware
- FIPS 140-2 validation
- Secure erase of NAND/User data
- Common Criteria certification
- Configurable for PCI DSS compliance
- Embedded TPM (Trusted Platform Module) 2.0. Disabled on shipments to China.
- Advanced Encryption Standard (AES) and Triple Data Encryption Standard (3DES) on browser
- Bezel Locking Kit option.
- Support for Commercial National Security Algorithms (CNSA)
- Chassis Intrusion detection option
- Secure Recovery - recover critical firmware to a known good state on detection of compromised firmware

Warranty

This product is covered by a global limited warranty and supported by HPE Services and a worldwide network of Hewlett Packard Enterprise Authorized Channel Partners resellers. Hardware diagnostic support and repair are available for three years from the date of purchase. Support for software and initial setup is available for 90 days from the date of purchase. Enhancements to warranty services are available through HPE Services operational services or customized service agreements. Hard drives have either a one-year or three-year warranty; refer to the specific hard drive QuickSpecs for details.

Notes: Server Warranty includes 3-Year Parts, 3-Year Labor, and 3-Year Onsite support with next business

Standard Features

day response. Warranty repairs may be accomplished through the use of Customer Self Repair (CSR) parts. These parts fall into two categories: 1) Mandatory CSR parts are designed for easy replacement. A travel and labor charge will result when customers decline to replace a Mandatory CSR part; 2) Optional CSR parts are also designed for easy replacement but may involve added complexity. Customers may choose to have Hewlett Packard Enterprise replace Optional CSR parts at no charge. Additional information regarding worldwide limited warranty and technical support is available at:

https://www.hpe.com/support/ProLiantServers-Warranties

https://www.hpe.com/us/en/search-results.html?page=1&q=servers warranty&autocomplete=0

Optional Features

Server Management

HPE iLO Advanced

HPE iLO Advanced licenses offer smart remote functionality for all HPE ProLiant servers without compromise. The license includes the full integrated remote console, virtual keyboard, video, and mouse (KVM), multi-user collaboration, console record and replay, and GUI-based and scripted virtual media and virtual folders. You can also activate the enhanced security and power management functionality.

HPE OneView Advanced

HPE OneView Advanced offers a sophisticated level of automation to infrastructure management by taking a template-driven approach to provisioning, updating, and integrating compute, storage, and networking infrastructure. It builds upon the base features of HPE OpenView Standard and provides full-featured licenses which can be purchased for managing multiple HPE server generations.

To learn more, visit **http://www.hpe.com/info/oneview**.

HPE Insight Cluster Management Utility (CMU)

HPE Insight Cluster Management Utility is a Hyperscale management framework that includes software for the centralized provisioning, management, and monitoring of nodes and infrastructure. Learn more at **http://www.hpe.com/info/cmu**

Accelerator and GPU Information

Hewlett Packard Enterprise supports various accelerators on select HPE ProLiant servers to support different workloads. The accelerators enable seamless integration of GPU computing with HPE ProLiant servers for high-performance computing, large. data center graphics, deep learning and virtual desktop deployments. These accelerators deliver all of the standard benefits of GPU computing while enabling maximum reliability and tight integration with system monitoring and management tools such as HPE Insight Cluster Management Utility.

One Config Simple (OCS/SCE)

OCS/SCE is a guided self-service tool to help sales and non-technical people provide customers with initial configurations in 3 to 5 minutes. You may then send the configuration on for configuration help or use it in your existing ordering processes. If you require "custom" rack configuration or configuration for products not available in SCE, please contact Hewlett Packard Enterprise Customer Business Center or an Authorized Partner for assistance. **https://h22174.www2.hpe.com/SimplifiedConfig/Welcome#**

Rack and Power Infrastructure

The story may end with servers, but it starts with the foundation that makes compute go - and business grow. ' 'We've reinvented our entire portfolio of rack and power products to make IT infrastructure more secure, practical, and efficient. In other words, ' 'we've created a more robust, smarter, and simpler infrastructure to help you get the most out of your IT equipment. As an industry leader, Hewlett Packard Enterprise is uniquely positioned to address the key concerns of power, cooling, cable management, and system access.

HPE G2 Advanced and Enterprise Racks are perfect for the server room or ' 's modern data center, with

Optional Features

enhanced airflow and thermal management, flexible cable management, and a 10-year Warranty to support higher-density computing.

HPE G2 PDUs offer reliable power in flexible form factors that operate at temperatures up to 60°C, include color-coded outlets and load segments, and a low-profile design for optimal access to the rack and support for dense rack environments.

HPE Uninterruptible Power Systems are cost-effective power protection for any workload. Some UPSs include options for remote management and extended runtime modules so ' 'your critical dense data center is covered in power outages.

HPE KVM Solutions include a console and switches designed to work reliably with your server and IT equipment. ' 'We've got a cost-effective KVM switch for your first rack and multiple connection IP switches with remote management and security capabilities to keep your data center rack up and running.

Learn more about HPE Racks, KVM, PDUs, and UPSs at **HPE Rack and Power Infrastructure**.

Service and Support

HPE Services

No matter where you are in your digital transformation journey, you can count on HPE Services to deliver the expertise you need when, where and how you need it. From planning to deployment, ongoing operations and beyond, our experts can help you realize your digital ambitions.
https://www.hpe.com/services

Consulting Services

No matter where you are in your journey to hybrid cloud, experts can help you map out your next steps. From determining what workloads should live where, to handling governance and compliance, to managing costs, our experts can help you optimize your operations.
https://www.hpe.com/services/consulting

HPE Managed Services

HPE runs your IT operations, providing services that monitor, operate, and optimize your infrastructure and applications, delivered consistently and globally to give you unified control and let you focus on innovation.
HPE Managed Services | HPE

Operational services

Optimize your entire IT environment and drive innovation. Manage day-to-day IT operational tasks while freeing up valuable time and resources. Meet service-level targets and business objectives with features designed to drive better business outcomes.
https://www.hpe.com/services/operational

HPE Complete Care Service

HPE Complete Care Service is a modular, edge-to-cloud IT environment service designed to help optimize your entire IT environment and achieve agreed upon IT outcomes and business goals through a personalized experience. All delivered by an assigned team of HPE Services experts. HPE Complete Care Service provides:

- A complete coverage approach -- edge to cloud
- An assigned HPE team
- Modular and fully personalized engagement
- Enhanced Incident Management experience with priority access
- Digitally enabled and AI driven customer experience

https://www.hpe.com/services/completecare

HPE Tech Care Service

HPE Tech Care Service is the operational support service experience for HPE products. The service goes beyond traditional support by providing access to product specific experts, an AI driven digital experience, and general technical guidance to not only reduce risk but constantly search for ways to do things better. HPE Tech Care Service delivers a customer-centric, AI driven, and digitally enabled customer experience to move your business forward. HPE Tech Care Service is available in three response levels. Basic, which provides 9x5 business hour availability and a 2-hour response time. Essential which provides a 15-minute response time 24x7 for most enterprise level customers, and Critical which includes a 6-hour repair commitment where available and outage management response for severity 1 incidents.
https://www.hpe.com/services/techcare

HPE Lifecycle Services

HPE Lifecycle Services provide a variety of options to help maintain your HPE systems and solutions at all stages of the product lifecycle. A few popular examples include:

- Lifecycle Install and Startup Services: Various levels for physical installation and power on, remote access setup, installation and startup, and enhanced installation services with the operating system.
- HPE Firmware Update Analysis Service: Recommendations for firmware revision levels for selected HPE products, taking into account the relevant revision dependencies within your IT environment.
- HPE Firmware Update Implementation Service: Implementation of firmware updates for selected HPE server, storage, and solution products, taking into account the relevant revision dependencies within your IT environment.
- Implementation assistance services: Highly trained technical service specialists to assist you with a variety of activities, ranging from design, implementation, and platform deployment to consolidation, migration, project management, and onsite technical forums.
- HPE Service Credits: Access to prepaid services for flexibility to choose from a variety of specialized service activities, including assessments, performance maintenance reviews, firmware management, professional services, and operational best practices.

Notes: To review the list of Lifecycle Services available for your product go to:
https://www.hpe.com/services/lifecycle
For a list of the most frequently purchased services using service credits, see the **HPE Service Credits Menu**

Other Related Services from HPE Services:

HPE Education Services

Training and certification designed for IT and business professionals across all industries. Broad catalogue of course offerings to expand skills and proficiencies in topics ranging from cloud and cybersecurity to AI and DevOps. Create learning paths to expand proficiency in a specific subject. Schedule training in a way that works best for your business with flexible continuous learning options.
https://www.hpe.com/services/training

Defective Media Retention

An option available with HPE Complete Care Service and HPE Tech Care Service and applies only to Disk or eligible SSD/Flash Drives replaced by HPE due to malfunction.
Consult your HPE Sales Representative or Authorized Channel Partner of choice for any additional questions and services options.

Parts and Materials

HPE will provide HPE-supported replacement parts and materials necessary to maintain the covered hardware product in operating condition, including parts and materials for available and recommended engineering improvements.
Parts and components that have reached their maximum supported lifetime and/or the maximum usage limitations as set forth in the manufacturer's operating manual, product quick-specs, or the technical product data sheet will not be provided, repaired, or replaced as part of these services.

How to Purchase Services

Services are sold by Hewlett Packard Enterprise and Hewlett Packard Enterprise Authorized Service Partners:

- Services for customers purchasing from HPE or an enterprise reseller are quoted using HPE order configuration tools.
- Customers purchasing from a commercial reseller can find services at **https://ssc.hpe.com/portal/site/ssc/**

Service and Support

AI Powered and Digitally Enabled Support Experience
Achieve faster time to resolution with access to product-specific resources and expertise through a digital and data driven customer experience.
Sign into the HPE Support Center experience, featuring streamlined self-serve case creation and management capabilities with inline knowledge recommendations. You will also find personalized task alerts and powerful troubleshooting support through an intelligent virtual agent with seamless transition when needed to a live support agent.
https://support.hpe.com/hpesc/public/home/signin

Consume IT On Your Terms

HPE GreenLake edge-to-cloud platform brings the cloud experience directly to your apps and data wherever they are-the edge, colocations, or your data center. It delivers cloud services for on-premises IT infrastructure specifically tailored to your most demanding workloads. With a pay-per-use, scalable, point-and-click self-service experience that is managed for you, HPE GreenLake edge-to-cloud platform accelerates digital transformation in a distributed, edge-to-cloud world.

- Get faster time to market
- Save on TCO, align costs to business
- Scale quickly, meet unpredictable demand
- Simplify IT operations across your data centers and clouds

To learn more about HPE Services, please contact your Hewlett Packard Enterprise sales representative or Hewlett Packard Enterprise Authorized Channel Partner. Contact information for a representative in your area can be found at "Contact HPE" **https://www.hpe.com/us/en/contact-hpe.html**
For more information
http://www.hpe.com/services

Pre-Configured Models

Pre-Configured models ship with the configurations below.

- Pre-Configured models ship with the configurations below. Options can be selected from the Core or Additional options section of this QuickSpecs.
- Hewlett Packard Enterprise does not allow factory integration of options into pre-configured models. Any additional options purchased will not be shipped inside the server.
- Network Choice models do not include embedded LOM.

SKU Number	P57685-B21 P57685-291 P57685-421	P57686-B21 P57686-291 P57686-421 P57686-AA1
Model Name	P57685-B21 & P57685-291: HPE ProLiant DL320 Gen11 3408U 1.8GHz 8-core 1P 16GB-R 4LFF 500W PS Server	P57686-B21, P57686-291 & P57686-AA1: HPE ProLiant DL320 Gen11 3408U 1.8GHz 8-core 1P 16GB-R 8SFF 500W PS Server
	P57685-421: HPE ProLiant DL320 Gen11 3408U 1.8GHz 8-core 1P 16GB-R 4LFF 1000W PS Server	P57686-421: HPE ProLiant DL320 Gen11 3408U 1.8GHz 8-core 1P 16GB-R 8SFF 1000W PS Server
Chassis	HPE ProLiant DL320 Gen11 4LFF Configure-to-order Server	HPE ProLiant DL320 Gen11 8SFF Configure-to-order Server
Processor	3408U (8 core, 1.8 GHz, 125W)	
Number of Processors	One with standard heatsink	
Memory	16 GB (1x16 GB , 4800 MT/s) **Notes:** Runs at 4000 MT/s due to processor limitation.	
Network Controller	HPE embedded 2-port 1 GbE Ethernet adapter	
Storage Controller	P57685-B21 & P57685-421: Embedded SATA controller (AHCI or Intel SATA software RAID controller)	P57686-B21, P57686-421 & P57686-AA1: Embedded SATA controller (AHCI or Intel SATA software RAID controller)
	P57685-291:	P57686-291:
Included Hard Drives	None ship standard, 4 LFF supported	None ship standard, 8 SFF supported
Optical Drive	P57685-B21 & P57685-421: None included (Optical Drive options available)	P57686-B21, P57686-421 & P57686-AA1: None included (Optical Drive options available)
	P57685-291: Optional HPE 9.5mm SATA DVD-RW Optical Drive, HPE Mobile USB DVD-RW Drive	P57686-291: Optional HPE 9.5mm SATA DVD-RW Optical Drive, HPE Mobile USB DVD-RW Drive
Expansion Slots	2 x16 FHHL PCIe Gen5; 1 OCP 3.0 Slot	

Pre-Configured Models

Power Supply	P57685-B21 & P57685-291: 1x HPE 500W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P57686-B21, P57686-291 & P57686-AA1: 1x HPE 500W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
	P57685-421: 1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit	P57686-421: 1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit
Fans	5x Standard Fans	
Management	HPE iLO 6	
Rail Kit	HPE Easy Install Rail 2 Kit	HPE Easy Install Rail 1 Kit
Security	TPM (Trusted Platform Module)	P57686-B21, P57686-291 & P57686-421: TPM (Trusted Platform Module)
		P57686-AA1: TPM (Trusted Platform Module)
Form Factor	1U Rack	
Warranty	Server warranty includes 3-year parts, 3-year labor, 3-year onsite support with next business day response.	

Pre-Configured Models

SKU Number	P57687-B21 P57687-291 P57687-421	P57688-B21 P57688-291 P57688-421
Model Name	P57687-B21 & P57687-291: HPE ProLiant DL320 Gen11 4410Y 2.0GHz 12-core 1P 16GB-R MR408i-o 8SFF 500W PS Server	P57688-B21 & P57688-291: HPE ProLiant DL320 Gen11 5416S 2.0GHz 16-core 1P 32GB-R MR408i-o 8SFF 500W PS Server
	P57687-421: HPE ProLiant DL320 Gen11 4410Y 2.0GHz 12-core 1P 16GB-R MR408i-o 8SFF 1000W PS Server	P57688-421: HPE ProLiant DL320 Gen11 5416S 2.0GHz 16-core 1P 32GB-R MR408i-o 8SFF 1000W PS Server
Chassis	HPE ProLiant DL320 Gen11 8SFF Configure-to-order Server	
Processor	4410Y (12 core, 2.0 GHz, 150W)	5416S (16 core, 2.0 GHz, 150W)
Number of Processors	One with standard heatsink	
Memory	16 GB (1x16 GB , 4800 MT/s) **Notes:** Runs at 4000 MT/s due to processor limitation.	32 GB (1x32 GB , 4800 MT/s) **Notes:** Runs at 4400 MT/s due to processor limitation.
Network Controller	HPE embedded 2-port 1 GbE Ethernet adapter	
Storage Controller	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller **Notes:** Smart Storage battery included	
Included Hard Drives	None ship standard, 8 SFF supported	
Optical Drive	P57687-B21 & P57687-421: None included (Optical Drive options available)	P57688-B21 & P57688-421: None included (Optical Drive options available)
	P57687-291: Optional HPE 9.5mm SATA DVD-RW Optical Drive, HPE Mobile USB DVD-RW Drive	P57688-291: Optional HPE 9.5mm SATA DVD-RW Optical Drive, HPE Mobile USB DVD-RW Drive
Expansion Slots	2 x16 FHHL PCIe Gen5; 1 OCP 3.0 Slot	
Power Supply	P57687-B21 & P57687-291: 1x HPE 500W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P57688-B21 & P57688-291: 1x HPE 500W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
	P57687-421: 1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit	P57688-421: 1x HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit
Fans	7x Standard Fans	
Management	HPE iLO 6	
Rail Kit	HPE Easy Install Rail 1 Kit	

Pre-Configured Models

Security	TPM (Trusted Platform Module)
Form Factor	1U Rack
Warranty	Server warranty includes 3-year parts, 3-year labor, 3-year onsite support with next business day response.

Pre-Configured Models

Country Code Key

- -B21 = Worldwide
- -291 = Japan
- -421 = Europe, the Middle East and Africa
- -AA1 = China

HPE Smart Choice purchase program

The HPE Smart Choice purchase program features popular fully configured products that can be quoted in minutes and shipped quickly through HPE Authorized Partners. Products are configured and tested in an HPE factory and stocked at HPE Authorized Distributors and Partners. The products arrive in a single box, making onsite integration easier and more efficient for partners and customers. Additionally, there are aggressively priced HPE Tech Care Services available only through the HPE Smart Choice program when you purchase an HPE Smart Choice product.

For additional information on the HPE Smart Choice purchase program, please visit:
https://www.hpe.com/psnow/doc/a50009219enw

Notes: Enabling TPM 2.0 no longer requires TPM module option kit for Gen11 since It is an embedded feature.

Configuration Information

Smart Templates from HPE

HPE is releasing new Smart Template technology in the One Config Advanced (OCA) configurator. These Templates represent the CTO equivalents of the top-selling BTO configurations. They are intended to provide simple starting points to assist you in easily creating and customizing your desired Server solutions. HPE Servers that have Platform Templates developed by HPE Product Managers will have a separate tab in the HPE OCA configurator.

Workload Solutions Templates from HPE

The Workload Solutions Templates build on the Smart Templates technology to easily develop working configurations of the most compelling Workload Solutions. The templates complement the reference Builds developed by HPE. Workload Solutions templates preconfigure some of the key architecture decisions and make it easier for Sellers to get started and complete a differentiated server solution for your ' 'customer's specific workload.

Mainstream SKUs

HPE launched the Mainstream SKU initiative as a market-driven approach to Demand Steering. It is a simplified portfolio of our top-selling options that meet the current and future market trends. HPE has committed to providing a more predictable and faster experience for these options. Mainstream SKUs enjoy higher safety stock levels and have higher fulfillment service levels than non-Mainstream SKUs. Mainstream orders are fulfilled +30% faster than non-Mainstream orders, have fewer shortages, and have better recovery dates. This platform has Mainstream SKUs in the options portfolio and is eligible for an improved Mainstream experience. Mainstream SKUs are designated with a Mainstream symbol in our configurators.

Mainstream Configurations

HPE is using the new Smart Templates technology to present Mainstream configurations. All the options in a Mainstream configuration are pre-selected Mainstream SKUs to optimize the performance, predictability, and fulfillment experience. Check the Template section in our configurators for eligible Mainstream configurations.

This section lists some of the steps required to configure a Factory Integrated Model. To ensure valid configurations are ordered, Hewlett Packard Enterprise recommends using an HPE-approved configurator. Contact your local sales representative for information on configurable product offerings and requirements.

- Factory Integrated Models must start with a CTO Server.
- FIO indicates that this option is only available as a factory installable option.
- Some options may not be integrated at the factory. Contact your local sales representative for additional information

Step 1: Base Configuration (choose one (1) of the following (3) configurable models from the tables below)

Configuration Information

CTO Server	HPE PrLiant DL320 Gen11 4 LFF CTO Server	HPE ProLiant DL320 Gen11 8 SFF CTO Server	HPE ProLiant DL320 Gen11 12 LFF CTO Server	HPE ProLiant DL320 Gen11 GPU CTO Server
SKU Number	P52765-B21	P52766-B21	P52767-B21	P61218-B21
TAA SKU[1]	P52765-B21#GTA	P52766-B21#GTA	P52767-B21#GTA	P61218-B21#GTA
Processor	Not included as standard			
DIMM Slots	16-DIMM slots			
Storage Controller	8 SATA ports, choice of HPE modular Smart Array and PCIe plug-in Controllers, and Intel® VROC Software RAID.			
PCIe	PCIe 5.0: 2 slots (x 16 FHHL) 1x OCP 3.0 in PCIe 5.0 (x16)			
Drive Cage - included	4 SAS/SATA LFF Direct attach for 4 LFF SATA	(8+2) SAS/SATA SFF (8+2) x2 NVMe SFF or, (8+2) x4 NVMe SFF	12 SAS/SATA LFF Direct attach support for 10 LFF SATA	4 SFF SAS/SATA/NVMe or 8 EDSFF E3.s 1T
Network Controller	Two (2) Embedded Broadcom BCM5720 1 GbE networking port Choice of either OCP 3.0 or select standup network adapters for primary networking selection plus additional/optional standup network adapters			
Fans	Support for (7) High-Performance Fans			
Management	HPE iLO with Intelligent Provisioning (standard), iLO Advanced (optional), and OneView (optional)			
USB	Front: One (1) USB 3.2 Gen1 + iLO service port Rear: Two (2) USB 3.2 Gen1 ports Internal: One (1) USB 2.0 port			
	Option: One (1) Front USB 2.0 port	Option: One (1) Front USB 2.0 port		
Security	Embedded TPM 2.0 (Trusted Platform Module) **Notes:** Enabling TPM 2.0 no longer requires TPM module option kit for Gen11 since It is an embedded feature.			
Rail Kit	Optional Easy Install rails and CMA			
Form Factor	1U Rack[2]			
Warranty	3-year parts, 3-year labor, 3-year onsite support with next business day response.			

Notes:

- All DL320 Gen11 CTO Server models require the selection of Processor, Memory and Power Supply. Backplane to be further selected in 8SFF and GPU CTO Server.
- [1]HPE offers multiple Trade Agreement Act (TAA) compliant configurations to meet the needs of US Federal Government customers. These products are either manufactured or substantially transformed in a designated country. TAA compliance is only provided when HPE options are included as part of factory integrated orders (CTO).

Configuration Information

- All CTO servers are Energy Star 3.0 compliant. After January 11, 2024, Energy Star 3.0 compliance is no longer valid. Energy Star 4.0 certification will be valid upon publication. Supported Rail kit to be defaulted for CTO Model in the configurator. But customer can deselect the Rail kit if the CTO Model is selected without Rack (Standalone server).
- If GPU CTO Model is selected, then Performance Heatsink must be selected.
- [2]12 LFF CTO Server requires 1200mm depth rack for proper mounting.

Step 2: Choose Options

Step 2a: Choose Processor

Processor Option Kits (Required Processor)

Notes:

- All SKUs below ship with one processor only. Adequate fans and heatsinks must be selected.
- Processor with TDP greater than 185W require High Performance Heatsink (P52756-B21) and High-Performance Fan Kit (P48908-B21).
- 5th Generation Intel Xeon supports "HPE DDR5 Smart Memory - Registered (RDIMM), up to 5600MT/s".
- 4th Generation Intel Xeon supports "HPE DDR5 Smart Memory - Registered (RDIMM), up to 4800MT/s".
- 96 GB 4800 MT/s Memory cannot be selected if HBM or MCC die.
- 3508U supports PCIe4.0 Only

5th Generation Intel Xeon-Platinum	
Intel Xeon-Platinum 8581V 2.0GHz 60-core 270W Processor for HPE	P67109-B21
5th Generation Intel Xeon-Gold 6	
Intel Xeon-Gold 6526Y 2.8GHz 16-core 195W Processor for HPE	P67080-B21
Intel Xeon-Gold 6542Y 2.9GHz 24-core 250W Processor for HPE	P67081-B21
Intel Xeon-Gold 6548Y+ 2.5GHz 32-core 250W Processor for HPE	P67082-B21
Intel Xeon-Gold 6534 3.9GHz 8-core 195W Processor for HPE	P67083-B21
Intel Xeon-Gold 6544Y 3.6GHz 16-core 270W Processor for HPE	P67084-B21
Intel Xeon-Gold 6530 2.1GHz 32-core 270W Processor for HPE	P67095-B21
Intel Xeon-Gold 6538Y+ 2.2GHz 32-core 225W Processor for HPE	P67096-B21
Intel Xeon-Gold 6538N 2.1GHz 32-core 205W Processor for HPE	P67104-B21
Intel Xeon-Gold 6548N 2.8GHz 32-core 250W Processor for HPE	P67105-B21
Intel Xeon-Gold 6554S 2.2GHz 36-core 270W Processor for HPE	P67110-B21
5th Generation Intel Xeon-Gold 5	
Intel Xeon-Gold 5515+ 3.2GHz 8-core 165W Processor for HPE	P67079-B21
Intel Xeon-Gold 5520+ 2.2GHz 28-core 205W Processor for HPE	P67094-B21
Intel Xeon-Gold 5512U 2.1GHz 28-core 185W Processor for HPE	P67101-B21
5th Generation Intel Xeon-Silver	
Intel Xeon-Silver 4509Y 2.6GHz 8-core 125W Processor for HPE	P67090-B21
Intel Xeon-Silver 4510 2.4GHz 12-core 150W Processor for HPE	P67091-B21

Configuration Information

Intel Xeon-Silver 4514Y 2.0GHz 16-core 150W Processor for HPE	P67092-B21
Intel Xeon-Silver 4516Y+ 2.2GHz 24-core 185W Processor for HPE	P67093-B21
5th Generation Intel Xeon-Bronze	
Intel Xeon-Bronze 3508U 2.1GHz 8-core 125W Processor for HPE	P67100-B21

Notes: Intel Bronze CPU Models (3508U) is capable of supporting up to PCIe Gen4 speed. This will have impact in high speed Networking & InfiniBand cards when selected with Bronze CPU Model.

4th Generation Intel Xeon-Gold 6

Intel Xeon-Gold 6426Y 2.5GHz 16-core 185W Processor for HPE	P49598-B21
Intel Xeon-Gold 6442Y 2.6GHz 24-core 225W Processor for HPE	P49599-B21
Intel Xeon-Gold 6448Y 2.1GHz 32-core 225W Processor for HPE	P49600-B21
Intel Xeon-Gold 6434 3.7GHz 8-core 195W Processor for HPE	P49601-B21
Intel Xeon-Gold 6444Y 3.6GHz 16-core 270W Processor for HPE	P49602-B21
Intel Xeon-Gold 6430 2.1GHz 32-core 270W Processor for HPE	P49614-B21
Intel Xeon-Gold 6438Y+ 2.0GHz 32-core 205W Processor for HPE	P49615-B21
Intel Xeon-Gold 6414U 2.0GHz 32-core 250W Processor for HPE	P49619-B21
Intel Xeon-Gold 6438N 2.0GHz 32-core 205W Processor for HPE	P49638-B21
Intel Xeon-Gold 6421N 1.8GHz 32-core 185W Processor for HPE	P49641-B21
Intel Xeon-Gold 6454S 2.2GHz 32-core 270W Processor for HPE	P49654-B21
4th Generation Intel Xeon-Gold 5	
Intel Xeon-Gold 5415+ 2.9GHz 8-core 150W Processor for HPE	P49597-B21
Intel Xeon-Gold 5418Y 2.0GHz 24-core 185W Processor for HPE	P49612-B21
Intel Xeon-Gold 5420+ 2.0GHz 28-core 205W Processor for HPE	P49613-B21
Intel Xeon-Gold 5412U 2.1GHz 24-core 185W Processor for HPE	P49618-B21
Intel Xeon-Gold 5411N 1.9GHz 24-core 165W Processor for HPE	P49639-B21
Intel Xeon-Gold 5418N 1.8GHz 24-core 165W Processor for HPE	P49640-B21
Intel Xeon-Gold 5416S 2.0GHz 16-core 150W Processor for HPE	P49653-B21
4th Generation Intel Xeon-Silver	
Intel Xeon-Silver 4410Y 2.0GHz 12-core 150W Processor for HPE	P49610-B21
Intel Xeon-Silver 4416+ 2.0GHz 20-core 165W Processor for HPE	P49611-B21
4th Generation Intel Xeon-Bronze	
Intel Xeon-Bronze 3408U 1.8GHz 8-core 125W Processor for HPE	P49617-B21

Notes: Intel Bronze CPU Models (3408U) is capable of supporting up to PCIe Gen4 speed. This will have impact in high speed Networking & InfiniBand cards when selected with Bronze CPU Model.

Step 2b: Choose Memory Options

Please select one or more memory DIMMs from below.

For the new Gen11 memory population rule whitepaper and optimal memory performance guidelines and memory speed table, please go to:

http://www.hpe.com/docs/memory-population-rules

Configuration Information

For memory Reliability, Accessibility, and Serviceability (RAS) features whitepaper like Gen11 Fast Fault Tolerance and legacy mirrored memory feature, etc., please go to: **http://www.hpe.com/docs/memory-ras-feature.**

Notes:

- HPE Server Memory compatibility for a specific server platform may vary or be limited within a server platform depending upon the specific configuration being requested. Because each server environment and requirements can vary, memory compatibility is based not only upon the server family but may also be affected by the amount and type of additional hardware options installed within a specific server configuration. For this reason, some HPE memory DIMMs may be qualified for an HPE server model or family and yet occasionally not be supported with some configurations within that server family.
- Memory DIMMs must be selected in quantities 1, 2, 4, 6, 8, 12, or 16.
- The maximum memory speed and capacity are a function of the memory type, memory configuration, and processor model.

For additional information, please see the **HPE DDR5 Smart Memory QuickSpecs**

Registered DIMMs DDR5 (RDIMMs)

HPE 16GB (1x16GB) Single Rank x8 DDR5-5600 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64705-B21
HPE 32GB (1x32GB) Dual Rank x8 DDR5-5600 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64706-B21
HPE 64GB (1x64GB) Dual Rank x4 DDR5-5600 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64707-B21
HPE 96GB (1x96GB) Dual Rank x4 DDR5-5600 CAS-46-45-45 EC8 Registered Smart Memory Kit	P64708-B21
HPE 128GB (1x128GB) Quad Rank x4 DDR5-5600 CAS-52-45-45 EC8 Registered 3DS Smart Memory Kit	P64709-B21
HPE 16GB (1x16GB) Single Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P43322-B21
HPE 32GB (1x32GB) Dual Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P43328-B21
HPE 64GB (1x64GB) Dual Rank x4 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P43331-B21
HPE 96GB (1x96GB) Dual Rank x4 DDR5-4800 CAS-46-45-45 EC8 Registered Smart Memory Kit	P66675-B21
HPE 128GB (1x128GB) Quad Rank x4 DDR5-4800 CAS-46-39-39 EC8 Registered 3DS Smart Memory Kit	P43334-B21

Notes:

- Only HPE DDR5 5600 MT/s SKUs can be selected with 5th Generation Intel processors.
- Only HPE DDR5 4800 MT/s SKUs can be selected with 4th Generation Intel processors.
- Total memory selection must be in quantities of 1, 2, 4, 6, 8, 12 or 16. For more detailed information regarding memory population rules, please refer to - **https://www.hpe.com/docs/server-memory**.
- If Memory Fault Tolerance is selected, then only x4 memory options can be selected.
- Mixing x4, x8 or 3DS memory is not supported.
- Mixing 4800 Memory and 5600 Memory is not allowed.
- 96 GB Memory cannot be mixed with any other Memory and is supported with only select processors.
- For the GPU CTO Server, 96 GB and higher memory DIMMs are supported only with 25C and lower ambient temperature.
- 96 and 128 GB memory DIMMs require high performance fan kit.
- 5600 MT/s memory SKUs offer a transfer rate of 5600 MT/s at 1 DIMM per channel and 4400 MT/s at 2 DIMMs per channel.
- 4800 MT/s memory SKUs offer a transfer rate of 4800 MT/s at 1 DIMM per channel and 4400 MT/s at 2 DIMMs

Configuration Information

per channel.
– 96 GB 4800 MT/s Memory cannot be selected if HBM or MCC die.

Memory Blank Kit

HPE DDR4 DIMM Blank Kit P07818-B21

Notes: Memory DIMM blanks must be installed in each empty slot.

Step 2c: Choose Power Supplies

Please select one or two power supplies from below.

Notes: Mixing of 2 different power supplies is NOT supported.

HPE Flex Slot Power Supplies

HPE 500W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit 865408-B21

Notes:

- If a 500W power supply is selected with HPE ProLiant DL320 Gen11 8SFF x4 Tri-Mode 24G U.3 BC Backplane Kit (P52747-B21), then the internal controller must be selected. However, this will only support SAS/SATA drives.

– If a 500W power supply is selected with 8SFF CTO Server, NVMe drives cannot be selected, and vice versa.

HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit P03178-B21

HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit P38995-B21

HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit P38997-B21

Notes: Only supports high line voltage (200 VAC to 240 VAC).

HPE 1800W-2200W Flex Slot Titanium Hot Plug Power Supply Kit P44712-B21

Step 2d: Choose backplane /drive cage

No optional backplanes are available for the existing LFF CTO configurations.

8 SFF customers can select the back planes from the following options.

HPE ProLiant DL320 Gen11 8SFF x1 Tri-Mode 24G U.3 BC Backplane Kit P52743-B21

Notes:

– If the 8SFF Drive cage is selected with PCIe Internal controller, then '8SFF TM Cable Kit (P52775-B21) ' must be selected.
– If the 8SFF Drive cage is selected with OROC controller, then '8SFF OROC Cable (P52780-B21)' must be selected.
– 8SFF x1 U.3 Drive cage with Direct attach supports 8 SATA Drives only.
– If MR416i-p/MR216i-p/SR932i-p/MR416i-o/MR216i-o controller is selected along with 8SFF x1 U.3 and 2SFF x4 U.3 drive cage, then one controller is enough to support both the drive cages with the selection of '8SFF cable kit' and '2SFF TM cable kit.'

HPE ProLiant DL320 Gen11 8SFF x4 Tri-Mode 24G U.3 BC Backplane Kit P52747-B21

Notes:

– If the 8SFF Drive cage is selected with PCIe Internal controller, then '8SFF TM Cable Kit (P52775-B21)" must be selected.
– If the 8SFF Drive cage is selected with OROC controller, then '8SFF OROC Cable (P52780-B21)' must be selected.
– 8SFF x4 U.3 Drive cage with Direct attach supports 8 NVMe U.3 Drives only.

Configuration Information

- If 8SFF x4 U.3 Drive cage is selected with 2SFF U.3 drive cage and no Internal controller or **NVMe adapter (P25527-B21)** is selected, then 2SFF U.3 drive cage can support SATA drives only.
- If 8SFF x4 Drive cage and 2SFF drive cage is selected with Qty=1 of Internal controller, then '8SFF cable kit' or '2SFF TM cable kit (P52776-B21)' must be selected
- If 8SFF x4 Drive cage and 2SFF drive cage is selected with Qty=2 of Internal controller, then both '8SFF cable kit' and '2SFF TM cable kit (P52776-B21)' must be selected
- If 8SFF x4 Drive cage and 2SFF Drive cage are selected along with MR408i-o controller, then '2SFF TM cable kit (P52776-B21)' must be selected.
- '2SFF Drive cage' or '1U DP/USB/ODD Blank Kit' requires the selection of an 8SFF Drive cage.
- If 8SFF x4 Drive cage and 2SFF Drive cage is selected along with SR932i-p controller, then '8SFF TM cable kit (P52775-B21)' must be selected.

HPE ProLiant DL320 Gen11 2SFF x4 Tri-Mode 24G U.3 BC Drive Cage Kit P52751-B21

Notes:

- If the 8SFF Drive cage is selected with PCIe Internal controller, then '8SFF TM Cable Kit(P52775-B21)'' must be selected.
- If the 2SFF Drive cage is selected with PCIe Internal controller or OROC controller, then the '2SFF TM Cable Kit (P52776-B21)' must be selected.
- If the 2SFF Drive cage is selected with NVMe 2p Adapter, then '2SFF Retimer Cable(P52777-B21)' must be selected.
- 2SFF U.3 Drive cage with direct attach supports either 2 NVMe U.3 drives or SATA drives.
- If MR416i-p/MR216i-p/SR932i-p/MR416i-o/MR216i-o controller is selected along with 8SFF x1 U.3 and 2SFF x4 U.3 drive cage, then one controller is enough to support both the drive cages with the selection of '8SFF cable kit' and '2SFF TM cable kit (P52776-B21).'
- If the 8SFF x4 Drive cage and 2SFF drive cage are selected with Qty=1 of the Internal controller, then '8SFF cable kit' or '2SFF TM cable kit (P52776-B21)' must be selected.
- If the 8SFF x4 Drive cage and 2SFF drive cage are selected with Qty=2 of the Internal controller, then both '8SFF cable kit' and '2SFF TM cable kit (P52776-B21)' must be selected.
- If 8SFF x4 Drive cage and 2SFF Drive cage are selected along with MR408i-o controller, then '2SFF TM cable kit (P52776-B21)' must be selected.
- If the 8SFF x4 Drive cage and 2SFF Drive cage is selected along with the SR932i-p controller, then the '8SFF TM cable kit (P52775-B21)' must be selected.

HPE ProLiant DL3X0 Gen11 1U 8SFF Display Port/USB/Optical Drive Blank Kit P48926-B21

Notes:

- Either '2SFF drive cage' or '1U DP/USB/ODD Blank' can be selected.
- '2SFF Drive cage' or '1U DP/USB/ODD Blank Kit' requires an 8SFF Drive cage configuration.

HPE ProLiant DL320 Gen11 8EDSFF x4 NVMe Drive Cage Kit P61219-B21

Notes:

- Supported only with HPE ProLiant DL320 Gen11 GPU CTO Server (P61218-B21)
- 8EDSFF Drive cage is to be selected with HPE ProLiant DL320 Gen11 GPU 4x Direct Connected NVMe FIO Enablement Kit (P62181-B21) and drive cage will be connected to Direct attach and can support maximum of 4 EDSFF drives only.

HPE ProLiant DL320 Gen11 4SFF x4 Tri-Mode 24G U.3 BC Drive Cage Kit P61221-B21

Notes:

- Supported only with GPU CTO Server (P61218-B21).

Configuration Information

- If 4SFF U.3 drive cage is selected with MR416i-p controller, then HPE ProLiant DL320 Gen11 4SFF Tri-Mode Cable Kit (P62189-B21) must be selected.
- If 4SFF Drive cage is selected and MR416i-p is not selected, then HPE ProLiant DL320 Gen11 GPU 4x Direct Connected NVMe FIO Enablement Kit () must be selected.
- If 4SFF U.3 Drive cage is selected with HPE ProLiant DL320 Gen11 GPU 4x Direct Connected NVMe FIO Enablement Kit (P62181-B21) then drive cage will be connected to direct attach and support a maximum of 4 NVMe U.3 drives only.

Step 3: Choose Additional Factory Integratable Options (FIO)

One of the following from each list may be selected if desired at the time of factory integration

HPE iLO Common Password FIO Setting P08040-B21

Notes:

- Replaces iLO default randomized password with an HPE-defined common password. HPE highly recommends changing this password immediately after the initial onboarding process.
- Customers who want to choose their own custom iLO default password should use the HPE Factory Express Integration Services.

HPE Bezel Lock Kit 875519-B21

Notes: Bezel lock kit requires selection of Bezel kit.

HPE ProLiant DL320 Gen11 GPU 4x Direct Connected NVMe FIO Enablement Kit P62181-B21

Notes:

- Supported only with GPU CTO Server.
- If 8EDSFF or 4SFF Drive cage is selected and MR416i-p is not selected, then "HPE DL320 G11 GPU 4xDC NVMe FIO Kit" must be selected.
- If 8EDSFF Drive cage is selected with "HPE DL320 G11 GPU 4xDC NVMe FIO Kit" then drive cage will be connected to Direct attach and can support Max=4 EDSFF drives only.
- If 4SFF U.3 Drive cage is selected with "HPE DL320 G11 GPU 4xDC NVMe FIO Kit" then drive cage will be connected to direct attach and support Max=4 NVMe U.3 drives only.
- Cannot be selected if MR416i-p controller is selected.
- Max=1 of 'GPU 4xDC NVMe FIO Kit' or '4SFF TM Cable Kit' or '2xType-p NVMe Cable Kit' can be selected
- If Qty=1 of GPU is selected with GPU CTO Server, then 'GPU 4xDC NVMe FIO Kit' or '2 GPU FIO Enable Kit' must be selected. In GPU CTO Server if ‘GPU 4xDC NVMe FIO Kit' then the first GPU card selected will be installed in the front of the server. This should not be considered for total PCIe selection.

Configuration Information

HPE ProLiant DL320 Gen11 2SFF MLB FIO Power Cable Kit P66961-B21

Notes:

- Max=1
- If 2SFF Drive cage is selected and 8SFF drive cage is not selected then 2SFF MLB FIO Power Cable Kit must be selected.
- 2SFF MLB FIO Power Cable Kit cannot be selected if 8SFF drive cage is selected.
- Not Supported with GPU CTO Server.
- 2SFF MLB FIO Power Cable Kit requires selection of 2SFF drive cage.

HPE Converged Infrastructure Management Software

HPE OneView including 3yr 24x7 Support Flexible Quantity E-LTU E5Y35AAE

Step 4: Choose Additional Options for Factory Integration from the Core and Additional Options sections below

Core Options

HPE Processors (Please select one processor).

Notes:

- All SKUs below ship with a processor only. Adequate fans and heatsinks (standard or high performance) must be selected.
- 4800 MT/S maximum memory speed unless otherwise noted.
- 128GB SGX Enclave unless otherwise noted.
- Processors with TDP greater than 185W require High-Performance Heatsink (P52756-B21).
- Processors with TDP greater than 185W require High-Performance Fan Kit (P48908-B21).

4th Generation Intel Xeon-Gold 6

Intel Xeon-Gold 6426Y 2.5GHz 16-core 185W Processor for HPE	P49598-B21
Intel Xeon-Gold 6442Y 2.6GHz 24-core 225W Processor for HPE	P49599-B21
Intel Xeon-Gold 6448Y 2.1GHz 32-core 225W Processor for HPE	P49600-B21
Intel Xeon-Gold 6434 3.7GHz 8-core 195W Processor for HPE	P49601-B21
Intel Xeon-Gold 6444Y 3.6GHz 16-core 270W Processor for HPE	P49602-B21
Intel Xeon-Gold 6430 2.1GHz 32-core 270W Processor for HPE	P49614-B21
Intel Xeon-Gold 6438Y+ 2.0GHz 32-core 205W Processor for HPE	P49615-B21
Intel Xeon-Gold 6414U 2.0GHz 32-core 250W Processor for HPE	P49619-B21
Intel Xeon-Gold 6438N 2.0GHz 32-core 205W Processor for HPE	P49638-B21
Intel Xeon-Gold 6421N 1.8GHz 32-core 185W Processor for HPE	P49641-B21
Intel Xeon-Gold 6454S 2.2GHz 32-core 270W Processor for HPE	P49654-B21

4th Generation Intel Xeon-Gold 5

Intel Xeon-Gold 5415+ 2.9GHz 8-core 150W Processor for HPE	P49597-B21
Intel Xeon-Gold 5418Y 2.0GHz 24-core 185W Processor for HPE	P49612-B21
Intel Xeon-Gold 5420+ 2.0GHz 28-core 205W Processor for HPE	P49613-B21
Intel Xeon-Gold 5412U 2.1GHz 24-core 185W Processor for HPE	P49618-B21
Intel Xeon-Gold 5411N 1.9GHz 24-core 165W Processor for HPE	P49639-B21
Intel Xeon-Gold 5418N 1.8GHz 24-core 165W Processor for HPE	P49640-B21
Intel Xeon-Gold 5416S 2.0GHz 16-core 150W Processor for HPE	P49653-B21

4th Generation Intel Xeon-Silver

Intel Xeon-Silver 4410Y 2.0GHz 12-core 150W Processor for HPE	P49610-B21

Notes: 64GB SGX Enclave

Intel Xeon-Silver 4416+ 2.0GHz 20-core 165W Processor for HPE	P49611-B21

Notes: 64GB SGX Enclave

4th Generation Intel Xeon-Bronze

Intel Xeon-Bronze 3408U 1.8GHz 8-core 125W Processor for HPE	P49617-B21

HPE Memory

To streamline the configuration process for HPE ProLiant Gen11 servers and to provide the best product availability, HPE recommends memory from the list here: **http://www.hpe.com/products/recommend.**

Core Options

Best product availability is limited to the US, Canada, and Latin America at this time.

Notes: HPE Server Memory compatibility for a specific server platform may vary or be limited within a server platform depending upon the specific configuration being requested. Because each server environment and requirements can vary, memory compatibility is based not only upon the server family but may also be affected by the amount and type of additional hardware options installed within a specific server configuration. For this reason, some HPE memory DIMMs may be qualified for an HPE server model or family and yet occasionally not be supported with some configurations within that server family.

Maximum memory capacity and speed per processor depend on processor model selection or limitation.

HPE DDR5 Memory

Notes: DIMMs must be selected in quantities 1, 2, 4, 6, 8, 12, or 16 per socket.

Registered DIMMs (RDIMMs)

HPE 16GB (1x16GB) Single Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P43322-B21
HPE 32GB (1x32GB) Dual Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P43328-B21
HPE 64GB (1x64GB) Dual Rank x4 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit	P43331-B21
HPE 96GB (1x96GB) Dual Rank x4 DDR5-4800 CAS-46-45-45 EC8 Registered Smart Memory Kit	P66675-B21
HPE 128GB (1x128GB) Quad Rank x4 DDR5-4800 CAS-46-39-39 EC8 Registered 3DS Smart Memory Kit	P43334-B21

Notes:

- Mixing of x4 and x8 DIMMs is not allowed.
- Mixing of 3DS and non-3DS DIMMs is not allowed.
- Mixing of Single Rank and Dual Rank DIMM is allowed only when the total qty of DIMM selected is 16
- If Single Rank and Dual Rank DIMMs are mixed, then the Qty of Single Rank DIMM selected must be equal to the quantity of Dual Rank DIMM.
- 128GB DIMMs require High-Performance Fan Kit (P48908-B21).

HPE DIMM blanks

HPE DDR4 DIMM Blank Kit	P07818-B21

Optical Drives

HPE Mobile USB DVD-RW Optical Drive	701498-B21
HPE 9.5mm SATA DVD-ROM Optical Drive	726536-B21
HPE 9.5mm SATA DVD-RW Optical Drive	726537-B21

Notes:

- If Optical Drive is selected with 8SFF CTO Server, then '1U DP/USB/ODD Blank Kit' must be selected.
- If Optical Drive is selected with 4LFF CTO Server, then '4LFF ODD cable kit (P54641-B21)' must be selected.
- Not Supported with 12LFF or the GPU CTO Server.

Core Options

HPE ProLiant DL3X0 Gen11 1U LFF Display Port/USB Kit	P48928-B21

Notes: Supported only with 4LFF CTO Server

HPE Hard Disk Drives

Mission Critical - 12G SAS - SFF Drives

HPE 2.4TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty 512e Self-encrypting FIPS HDD	P28618-B21
HPE 1.2TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Self-encrypting FIPS HDD	P28622-B21

Enterprise - 12G SAS - SFF Drives

HPE 2.4TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty 512e Multi Vendor HDD	P28352-B21
HPE 1.8TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty 512e Multi Vendor HDD	P53562-B21
HPE 1.2TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Multi Vendor HDD	P28586-B21
HPE 300GB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Multi Vendor HDD	P40430-B21

Midline - 6G SATA - SFF Drives

HPE 2TB SATA 6G Business Critical 7.2K SFF BC 1-year Warranty 512e HDD	P28500-B21

Midline - 12G SAS - LFF Drives

HPE 20TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e ISE Multi Vendor HDD	P53553-B21
HPE 16TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e ISE Multi Vendor HDD	P23608-B21
HPE 14TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e Multi Vendor HDD	P09155-B21
HPE 12TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e Multi Vendor HDD	881781-B21
HPE 8TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty 512e Multi Vendor HDD	834031-B21
HPE 6TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty 512e Multi Vendor HDD	861746-B21
HPE 4TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	833928-B21
HPE 2TB SAS 12G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	833926-B21

Midline - 6G SATA - LFF Drives

HPE 20TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e ISE Multi Vendor HDD	P53554-B21
HPE 16TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e ISE Multi Vendor HDD	P23449-B21
HPE 12TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Helium 512e Multi Vendor HDD	881787-B21
HPE 8TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty 512e Multi Vendor HDD	834028-B21
HPE 6TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty 512e Multi Vendor HDD	861742-B21
HPE 4TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	861683-B21
HPE 2TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	861681-B21
HPE 1TB SATA 6G Business Critical 7.2K LFF LP 1-year Warranty Multi Vendor HDD	861686-B21

Core Options

SSD Selection

For SSD selection guidance, please visit **https://ssd.hpe.com/**

Read Intensive - 12G SAS - SFF - Solid State Drives

HPE 7.68TB SAS 12G Read Intensive SFF BC Value SAS Multi Vendor SSD	P40509-B21
HPE 3.84TB SAS 12G Read Intensive SFF BC Value SAS Multi Vendor SSD	P40508-B21
HPE 1.92TB SAS 12G Read Intensive SFF BC Value SAS Multi Vendor SSD	P40507-B21
HPE 960GB SAS 12G Read Intensive SFF BC Value SAS Multi Vendor SSD	P40506-B21

Read Intensive -24G SAS - SFF - Solid State Drives

HPE 15.36TB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49045-B21
HPE 7.68TB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49041-B21
HPE 3.84TB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49035-B21
HPE 1.92TB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49031-B21
HPE 960GB SAS 24G Read Intensive SFF BC Multi Vendor SSD	P49029-B21
HPE 3.84TB SAS Read Intensive SFF BC Self-encrypting FIPS 140-2 PM7 SSD	P63875-B21

Notes: Require High-Performance Fan Kit (P48908-B21).

Mixed Use - 12 SAS - SFF - Solid State Drives

HPE 3.84TB SAS 12G Mixed Use SFF BC Value SAS Multi Vendor SSD	P40512-B21
HPE 1.92TB SAS 12G Mixed Use SFF BC Value SAS Multi Vendor SSD	P40511-B21
HPE 960GB SAS 12G Mixed Use SFF BC Value SAS Multi Vendor SSD	P40510-B21

Mixed Use -24G SAS - SFF - Solid State Drives

HPE 6.4TB SAS 24G Mixed Use SFF BC Multi Vendor SSD	P49057-B21
HPE 3.2TB SAS 24G Mixed Use SFF BC Multi Vendor SSD	P49053-B21
HPE 1.6TB SAS 24G Mixed Use SFF BC Multi Vendor SSD	P49049-B21
HPE 800GB SAS 24G Mixed Use SFF BC Multi Vendor SSD	P49047-B21
HPE 1.6TB SAS Mixed Use SFF BC Self-encrypting FIPS 140-2 PM7 SSD	P63871-B21

Notes: Require High-Performance Fan Kit (P48908-B21).

Very Read Optimized - 6G SATA - SFF - Solid State Drives

Read Intensive - 6G SATA - SFF - Solid State Drives

HPE 7.68TB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40501-B21
HPE 3.84TB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40500-B21
HPE 3.84TB SATA 6G Read Intensive SFF BC PM893a SSD	P63910-B21
HPE 1.92TB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40499-B21
HPE 480GB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40497-B21
HPE 960GB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40498-B21
HPE 480GB SATA 6G Read Intensive SFF BC Self-encrypting 5400P SSD	P58236-B21
HPE 480GB SATA 6G Read Intensive SFF BC PM893a SSD	P63886-B21
HPE 240GB SATA 6G Read Intensive SFF BC Multi Vendor SSD	P40496-B21

Mixed Use - 6G SATA - SFF - Solid State Drives

HPE 3.84TB SATA 6G Mixed Use SFF BC Multi Vendor SSD	P40505-B21
HPE 1.92TB SATA 6G Mixed Use SFF BC Multi Vendor SSD	P40504-B21
HPE 960GB SATA 6G Mixed Use SFF BC Multi Vendor SSD	P40503-B21
HPE 960GB SATA 6G Mixed Use SFF BC Self-encrypting 5400M SSD	P58244-B21
HPE 480GB SATA 6G Mixed Use SFF BC Multi Vendor SSD	P40502-B21

Mixed Use - 12G SAS - LFF -Solid State Drives

Core Options

HPE 960GB SAS 12G Mixed Use LFF LPC Value SAS Multi Vendor SSD	P37009-B21
Read Intensive - 6G SATA - LFF - Solid State Drives	
HPE 960GB SATA 6G Read Intensive LFF LPC Multi Vendor SSD	P47808-B21
NVMe Gen5 - Read Intensive	
Read Intensive - NVMe - EDSFF - Solid State Drives	
HPE 3.2TB NVMe Gen5 High Performance Mixed Use E3S EC1 EDSFF SPDM CM7 SSD	P61191-B21
HPE 3.84TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM CM7 SSD	P61179-B21
HPE 6.4TB NVMe Gen5 High Performance Mixed Use E3S EC1 EDSFF SPDM CM7 SSD	P61195-B21
HPE 7.68TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM CM7 SSD	P61183-B21
HPE 15.36TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM CM7 SSD	P61187-B21
HPE 3.84TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM PM1743 SSD	P57799-B21
HPE 7.68TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM PM1743 SSD	P57803-B21
HPE 15.36TB NVMe Gen5 High Performance Read Intensive E3S EC1 EDSFF SPDM PM1743 SSD	P57807-B21
Read Intensive - NVMe - SFF - Solid State Drives	
HPE 15.36TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PM1733a SSD	P50224-B21
HPE 15.36TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 CM7 SSD	P63841-B21
HPE 15.36TB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static SPDM Multi Vendor SSD	P69255-B21
HPE 7.68TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PM1733a SSD	P50222-B21
HPE 7.68TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 CM7 SSD	P63837-B21
HPE 3.84TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PM1733a SSD	P50219-B21
HPE 3.84TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 CM7 SSD	P63833-B21
HPE 1.92TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 PM1733a SSD	P50216-B21
HPE 1.92TB NVMe Gen4 High Performance Read Intensive SFF BC U.3 CM7 SSD	P63829-B21
Mixed Use - NVMe - SFF - Solid State Drives	
HPE 6.4TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD	P50233-B21
HPE 6.4TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 CM7 SSD	P63853-B21
HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD	P50230-B21
HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 CM7 SSD	P63849-B21
HPE 1.6TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD	P50227-B21
HPE 1.6TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 CM7 SSD	P63845-B21
HPE 800GB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static V2 Multi Vendor SSD	P64999-B21
HPE 960GB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static V2 Multi Vendor SSD	P64842-B21
HPE 1.6TB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static V2 Multi Vendor SSD	P65007-B21
HPE 1.92TB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static V2 Multi Vendor SSD	P64844-B21
HPE 3.2TB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static V2 Multi Vendor SSD	P65015-B21
HPE 7.68TB NVMe Gen4 Mainstream Performance Read Intensive SFF BC U.3 Static V2 Multi Vendor SSD	P64848-B21

Notes: Require High-Performance Fan Kit (P48908-B21).

Core Options

VROC - Solid State Drives

HPE 3.84TB NVMe Gen4 Mainstream Performance Very Read Optimized E3S EC1 EDSFF P5430 SSD	P63930-B21
HPE 7.68TB NVMe Gen4 Mainstream Performance Very Read Optimized E3S EC1 EDSFF P5430 SSD	P63934-B21
HPE 15.36TB NVMe Gen4 Mainstream Performance Very Read Optimized E3S EC1 EDSFF P5430 SSD	P63938-B21

Notes: Requires selection of 8EDSFF Drive cage (P61219-B21).

Read Intensive - M.2 - Solid State Drives

HPE 480GB NVMe Gen3 Mainstream Performance Read Intensive M.2 Multi Vendor SSD	P40513-B21
HPE 480GB NVMe Gen4 Mainstream Performance Read Intensive M.2 PM9A3 SSD	P69543-B21
HPE 960GB NVMe Gen3 Mainstream Performance Read Intensive M.2 Multi Vendor SSD	P40514-B21
HPE 1.92TB NVMe Gen3 Mainstream Performance Read Intensive M.2 Multi Vendor SSD	P40515-B21
HPE 480GB SATA 6G Read Intensive M.2 Multi Vendor SSD	P47818-B21

Notes:

- Mixing of SATA M.2 and NVMe M.2 SSDs is not allowed.
- Need x7 fans (standard or performance).
- If M.2 SSD is selected then '2P Std Fan Kit' must be selected along with Standard Fan kit or High Performance Fan kit must be selected. This rule is not applicable for GPU CTO server
- All the M.2 NVMe drive will only run maximum PCIe Gen3 speed in DL320 Gen11

Hard Drive Blank Kits

HPE Gen9 LFF HDD Spade Blank Kit	807878-B21
HPE Small Form Factor Hard Drive Blank Kit	666987-B21

HPE Networking

1 Gigabit Ethernet adapters

Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T Adapter for HPE	P51178-B21
Intel I350-T4 Ethernet 1Gb 4-port BASE-T Adapter for HPE	P21106-B21

10 Gigabit Ethernet adapters

Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T Adapter for HPE	P26253-B21
Broadcom BCM57412 Ethernet 10Gb 2-port SFP+ Adapter for HPE	P26259-B21

25 Gigabit Ethernet adapters

Broadcom BCM57414 Ethernet 10/25Gb 2-port SFP28 Adapter for HPE	P26262-B21
Broadcom BCM57504 Ethernet 10/25Gb 4-port SFP28 Adapter for HPE	P26264-B21
Intel E810-XXVDA2 Ethernet 10/25Gb 2-port SFP28 Adapter for HPE	P08443-B21
Intel E810-XXVDA4 Ethernet 10/25Gb 4-port SFP28 Adapter for HPE	P08458-B21
Mellanox MCX631102AS-ADAT Ethernet 10/25Gb 2-port SFP28 Adapter for HPE	P42044-B21
HPE Ethernet 10/25Gb 2-port Secure Network Adapter	S2A69A

Core Options

Notes:

– Recommended ambient temperature with High-Performance Fan Kit (P48908-B21)
– *Conditional support, with DAC (direct attach copper) cable only.

	P26264-B21	P08443-B21	P08458-B21	P42044-B21	S2A69A
4 LFF (P52765-B21)	25C	35C	25C	25C	25C
8 SFF (P52766-B21)	25C	35C	25C	25C	25C
12 LFF (P52767-B21)	25C	35C	25C	25C	25C
GPU Dense (P61218-B21) with 1 Single Wide or 1 Double Wide GPU in front	25C	35C	25C	25C	25C
GPU Dense (P61218-B21) with up to 2 Single Wide GPU in front	25C	35C	25C	25C	25C
GPU Dense (P61218-B21) with 2 Double Wide or 4 Single Wide GPUs in the front	25C *	25C	25C *	25C *	25C *

100 Gigabit Ethernet Adapters

Mellanox MCX623106AS-CDAT Ethernet 100Gb 2-port QSFP56 Adapter for HPE P25960-B21
HPE Ethernet 100Gb 1-port QSFP28 PCIe3 x16 MCX515A-CCAT Adapter P31246-B21
Intel E810-CQDA2 Ethernet 100Gb 2-port QSFP28 Adapter for HPE P21112-B21

Notes: Recommended ambient temperature with High-Performance Fan Kit (P48908-B21)

	P25960-B21	P31246-B21	P21112-B21
4 LFF (P52765-B21)	25C	25C	25C
8 SFF (P52766-B21)	25C	25C	25C
12 LFF (P52767-B21)	Not supported	Not supported	Not supported
GPU Dense (P61218-B21) with 1 Single Wide or 1 Double Wide GPU in front	25C	25C	25C
GPU Dense (P61218-B21) with up to 2 Single Wide GPU in front	25C	25C	25C
GPU Dense (P61218-B21) with 2 Double Wide or 4 Single Wide GPUs in front	Not supported	Not supported	Not supported

200 Gigabit Ethernet Adapters

Mellanox MCX623105AS-VDAT Ethernet 200Gb 1-port QSFP56 Adapter for HPE P10180-B21

Notes: Recommended ambient temperature with High-Performance Fan Kit (P48908-B21)

Core Options

	P10180-B21
4 LFF (P52765-B21)	25C
8 SFF (P52766-B21)	25C
12 LFF (P52767-B21)	Not supported
GPU Dense (P61218-B21) with 1 Single Wide or 1 Double Wide GPU in front	25C
GPU Dense (P61218-B21) with up to 2 Single Wide GPU in front	25C
GPU Dense (P61218-B21) with 2 Double Wide or 4 Single Wide GPUs in front	Not supported

OCP 3.0 Adapters

Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE	P51181-B21
Intel I350-T4 Ethernet 1Gb 4-port BASE-T OCP3 Adapter for HPE	P08449-B21
Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T OCP3 Adapter for HPE	P10097-B21
Broadcom BCM57412 Ethernet 10Gb 2-port SFP+ OCP3 Adapter for HPE	P26256-B21
Broadcom BCM57504 Ethernet 10/25Gb 4-port SFP28 OCP3 Adapter for HPE	P26269-B21
Broadcom BCM57414 Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE	P10115-B21
Intel E810-XXVDA2 Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE	P10106-B21
Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE	P42041-B21
Intel E810-CQDA2 Ethernet 100Gb 2-port QSFP28 OCP3 Adapter for HPE	P22767-B21

Notes: Recommended ambient temperature with High-Performance Fan Kit (P48908-B21)

	P26269-B21	P10106-B21	P42041-B21	P22767-B21
4 LFF (P52765-B21)	35C	35C	35C	35C
8 SFF (P52766-B21)	35C	35C	35C	35C
12 LFF (P52767-B21)	35C	35C	35C	35C
GPU Dense (P61218-B21) with 1 Single Wide or 1 Double Wide GPU in front	25C	35C	35C	25C
GPU Dense (P61218-B21) with up to 2 Single Wide or 1 Double Wide GPU in front	25C	35C	35C	25C
GPU Dense (P61218-B21) with 2 Double Wide or 4 Single Wide GPU in the front	25C	35C	35C	25C

Fibre Channel HBA

HPE SN1610E 32Gb 1-port Fibre Channel Host Bus Adapter	R2J62A
HPE SN1610E 32Gb 2-port Fibre Channel Host Bus Adapter	R2J63A
HPE SN1610Q 32Gb 1-port Fibre Channel Host Bus Adapter	R2E08A
HPE SN1610Q 32Gb 2-port Fibre Channel Host Bus Adapter	R2E09A

Notes: PCIe 3.0 x8// FH or LP for SN1610Q/1610E

HPE SN1700E 64Gb 1-port Fibre Channel Host Bus Adapter	R7N77A

Core Options

HPE SN1700E 64Gb 2-port Fibre Channel Host Bus Adapter	R7N78A
HPE SN1700Q 64Gb 1-port Fibre Channel Host Bus Adapter	R7N86A
HPE SN1700Q 64Gb 2-port Fibre Channel Host Bus Adapter	R7N87A

Notes: PCIe 4.0 x8// FH or LP for SN1700QE 1-port and 2-port

HPE InfiniBand

HPE 100Gb 1-port OP101 QSFP28 x16 PCIe Gen3 with Intel Omni-Path Architecture Adapter	829335-B21
HPE InfiniBand HDR100/Ethernet 100Gb 1-port QSFP56 PCIe4 x16 MCX653105A-ECAT Adapter	P23665-B21
HPE InfiniBand HDR100/Ethernet 100Gb 2-port QSFP56 PCIe4 x16 MCX653106A-ECAT Adapter	P23666-B21

Notes:

- Recommended ambient temperature with High-Performance Fan Kit (P48908-B21)
- *Conditional support, with DAC (direct attach copper) cable only.

	829335-B21	P23665-B21	P23666-B21
4 LFF (P52765-B21)	35C	35C	25C
8 SFF (P52766-B21)	35C	35C	25C
12 LFF (P52767-B21)	35C	35C	25C
GPU Dense (P61218-B21) with 1 Single Wide or 1 Double Wide GPU in front	25C		
GPU Dense (P61218-B21) with up to 2 Single Wide GPUs in front	35C	35C	25C
GPU Dense (P61218-B21) with 2 Double Wide or 4 Single Wide GPU in the front	25C	25C	25C*

HPE I/O Expansion Options

Notes:

- Slot 1: The Primary Riser shipping default in the CTO server is a x16 FHHL.
- Slot 2: The Secondary Riser in the CTO server is optional x16 FHHL.

Risers

Part number	Description	Bus width (Gen5 lanes)			
		Primary	Secondary	Slot 1	Slot 2
P52753-B21	HPE DL320 Gen11 x16 FHHL Riser Kit	D	O	x16	x16

Notes:

Core Options

– D = Default on chassis; O = Optional; N = not supported or slot/connector not present.
– If the Secondary riser is not selected, then the Max PCIe card selection is limited to Qty=1. This is inclusive of all PCIe categories and sub-categories.

HPE Power Supplies

HPE Flexible Slot (Flex Slot) Power Supplies share a common electrical and physical design that allows for hot plug, tool-less installation into HPE ProLiant Servers. Flex Slot power supplies are certified for high-efficiency operation and offer multiple power output options, allowing users to "right-size" a power supply for specific server configurations. This flexibility helps to reduce power waste, lower overall energy costs, and avoid "trapped" power capacity in the data center.

Notes:

– Before making a power supply selection, it is highly recommended that the HPE Power Advisor is run to determine the right size power supply for your server configuration. The HPE Power Advisor is located at:

https://poweradvisorext.it.hpe.com/?Page=Indexv

– HPE ProLiant servers ship with an IEC-IEC power cord for rack mounting with Power Distribution Units (PDUs). Visit

HPE Power Cords and Cables for a full list of optional power cords.

– Mixing of Power supplies is NOT allowed.

– Please refer to "Factory Configuration Setting" section regarding "HPE CE Mark Removal FIO Enablement Kit (P35876-B21)" for non- EU Erp Lot 9 configuration.

– Flex Slot Platinum Plus power supplies support up to 94% power efficiency and include a C-14 power inlet connector that can support HPE Power Discovery Services (blue connector).

– European Union Erp Lot 9 Regulation. Beginning on January 1st, 2024, units sold into the European Union (EU), European Economic Area (EEA), the United Kingdom, or Switzerland must include more efficient AC power supplies: 94% for multi-output and 96% for single output.

– HPE Flexible Slot power supplies are single-output, and part numbers 865438-B21, P03178-B21, and P44712-B21 are 96% efficient, thus meeting requirements.

– HPE is on target to fulfil compliant systems ahead of time and will begin enforcing these requirements in advance to satisfy requests with the current power supplies by the set deadline.

HPE 500W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	865408-B21
HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P38995-B21
HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit	P03178-B21
HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	P38997-B21

Notes: Only supports high line voltage (200VAC to 240VAC).

HPE 1800W-2200W Flex Slot Titanium Hot Plug Power Supply Kit	P44712-B21
HPE 1600W Flex Slot -48VDC Hot Plug Power Supply Kit	P17023-B21

Notes: Requires selection of HPE 1600W DC PSU power lug option kit or HPE 1600W DC PSU Power Cable Kit.

Core Options

HPE 1600W -48VDC Power Cable Lug Kit P36877-B21

Cooling Options

HPE ProLiant DL3X0 Gen11 1U Standard Fan Kit P48907-B21

Notes: Standard Fan Kit Includes 5 fans.

HPE ProLiant DL3X0 Gen11 1U High Performance Fan Kit P48908-B21

Notes:

- High Performance Fan Kit Includes 7 fans.
- High Performance Fans required for processors with a TDP to or greater than 185W

HPE ProLiant DL3X0 Gen11 1U 2P Standard Fan Kit P54697-B21

Notes: 2P Standard Fan Kit Includes 2 fans.

HPE ProLiant DL3X0 Gen11 1U Standard Heat Sink Kit P48904-B21

HPE ProLiant DL320 Gen11 Performance Heat Sink Kit P52756-B21

Notes: Performance Heat Sink Required for processors with a TDP to or greater than 185W.

HPE Computation and Graphics Accelerators

NVIDIA L4 24GB PCIe Accelerator for HPE S0K89C

Notes:

- For best performance across common workloads, HPE recommends system main memory at least twice the memory of all GPU.
- Requires selection of High Performance Fan kit. This rule is not applicable for GPU CTO Server
- If GPU is selected with GPU CTO server then it does not occupy PCIe slot. This must be considered for total PCIe selection
- Mixing of GPU is not allowed
- If Qty=1 Single Wide GPU is selected with GPU CTO Server then 'GPU 4xDC NVMe FIO Kit' or '2 GPU FIO Enable Kit' or '4SW GPU FIO Enable Kit' must be selected
- If Qty=2 Single Wide GPU is selected with GPU CTO server then '2 GPU FIO Enable Kit' or '4SW GPU FIO Enable Kit' must be selected
- If Qty=3 or more GPU is selected then '4SW GPU FIO Enable Kit' must be selected
- 4SW GPU FIO Enable Kit requires selection of Single Wide GPU

NVIDIA A16 64GB PCIe Non-CEC Accelerator for HPE R8T26C

Notes:

- For best performance across common workloads, HPE recommends system main memory at least twice the memory of all GPU.
- If Qty=2 Double Wide GPU is selected with GPU CTO Server then '2 GPU FIO Enable Kit' must be selected
- If Qty=1 Double Wide GPU is selected with GPU CTO Server then 'GPU 4xDC NVMe FIO Kit' or '2 GPU FIO Enable Kit' must be selected
- Supported only with GPU CTO Server
- Max=2
- If Qty=1 of A16 GPU is selected, then Primary GPU 8p Power Cable Kit (P62201-B21) must be

Core Options

selected.
- If Qty=2 of A16 GPU is selected, then Primary GPU 8p Power Cable Kit (P62201-B21) and Sec GPU 8p Power Cable Kit (P62203-B21) must be selected.
- Mixing of GPU is not allowed
- If 4SW GPU FIO Enable Kit or Qty=2 of Double wide GPU is selected then this option cannot be selected

NVIDIA L40S 48GB PCIe Accelerator S2L70C

Notes:

- If Qty=2 Double Wide GPU is selected with GPU CTO Server then '2 GPU FIO Enable Kit' must be selected
- If Qty=1 Double Wide GPU is selected with GPU CTO Server then 'GPU 4xDC NVMe FIO Kit' or '2 GPU FIO Enable Kit' must be selected
- **Warning:** "For best performance across common workloads, HPE recommends system main memory at least twice the memory of all GPU."
- Supported only with GPU CTO Server
- If GPU is selected with GPU CTO server then it does not occupy PCIe slot on Primary or Secondary riser. This must be considered for total PCIe selection
- Max=2
- Mixing of GPU is not allowed
- If 4SW GPU FIO Enable Kit or Qty=2 of Double wide GPU is selected then this option cannot be selected
- If 4516Y+ processor is selected with 4SW GPU FIO Enable Kit or Qty=2 of Double wide GPU, then High performance heat sink must be selected and defaulted
- If Qty=1 of L40s GPU is selected then Prim GPU 16p Adv Pwr Kit must be selected
- If Qty=2 of L40s GPU is selected then Prim GPU 16p Adv Pwr Kit and Sec GPU 16p Adv Pwr Kit must be selected
- These processors (6530,8581V,6554S) and L40s GPU cannot be selected together
- 128GB DIMM and L40s GPU cannot be selected together

NVIDIA L40 48GB PCIe Accelerator for HPE S0K90C

Notes:

- For best performance across common workloads, HPE recommends system main memory at least twice the memory of all GPU.
- If Qty=2 Double Wide GPU is selected with GPU CTO Server then '2 GPU FIO Enable Kit' must be selected
- If Qty=1 Double Wide GPU is selected with GPU CTO Server then 'GPU 4xDC NVMe FIO Kit' or '2 GPU FIO Enable Kit' must be selected
- Supported only with GPU CTO Server
- If GPU is selected with GPU CTO server then it does not occupy PCIe slot. This must be considered for total PCIe selection
- Max=2
- If Qty=1 of L40 GPU is selected, then Prim GPU 16p Power Cable Kit (P62205-B21) must be selected.
- If Qty=2 of L40 GPU is selected, then Primary GPU 16p Power Cable Kit (P62205-B21) and Sec GPU 16p Power Cable Kit (P62207-B21) must be selected.
- Mixing of GPU is not allowed
- If 4SW GPU FIO Enable Kit or Qty=2 of Double wide GPU is selected then this option cannot be selected

Core Options

GPU Support	Ambient Temperature	Fan Kit
up to 2 L4 or 1 L40/A16 - Front	Up to 35C	High Performance Fan kit
up to 2 L4 - Rear	Up to 25C	High Performance Fan kit
2 L40/A16 or 4 L4 - Front	Up to 25C	High Performance Fan kit

Software as a Service Management

HPE Compute Ops Management

Base SKU

HPE Compute Ops Management Enhanced 3-year Upfront ProLiant SaaS R7A11AAE

Upgrade SKUS

HPE Compute Ops Management Enhanced 1-year Upfront ProLiant SaaS R7A10AAE

HPE Compute Ops Management Enhanced 5-year Upfront ProLiant SaaS R7A12AAE

HPE OneView

HPE OneView including 3yr 24x7 Support Flexible Quantity E-LTU E5Y35AAE

Notes: For customers purchasing HPE Compute Ops Management, without a hardware purchase or a BTO purchase, use this base SKU within ASQ order:

HPE Compute Ops Management Base SaaS R6Z73AAE

For more information, visit the HPE GreenLake for Compute Ops Management QuickSpecs:
https://www.hpe.com/psnow/doc/a50004263enw
Supported Servers - CTO only. No OEM. - Complete list can be found here: Latest Supported Server List: **https://www.hpe.com/info/com-supported-servers**

Additional Options

OS Boot Device

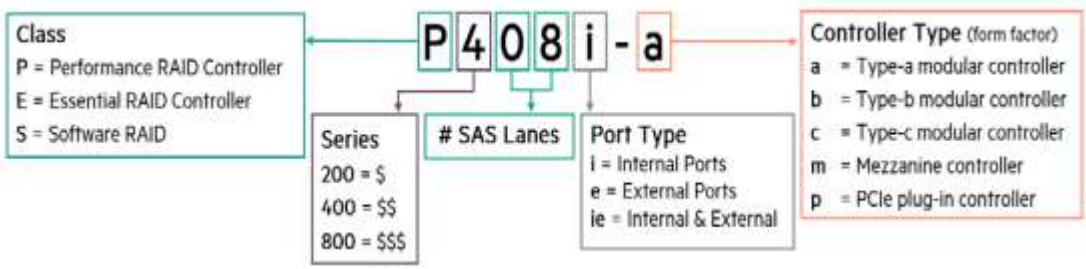
HPE NS204i-u Gen11 NVMe Hot Plug Boot Optimized Storage Device	P48183-B21
HPE ProLiant DL320 Gen11 NS204i-u NVMe OS Boot Device Cable Kit	P52786-B21
HPE ProLiant DL320 Gen11 NS204i-u Security Lock Removal FIO Trigger System Setting	P61855-B21

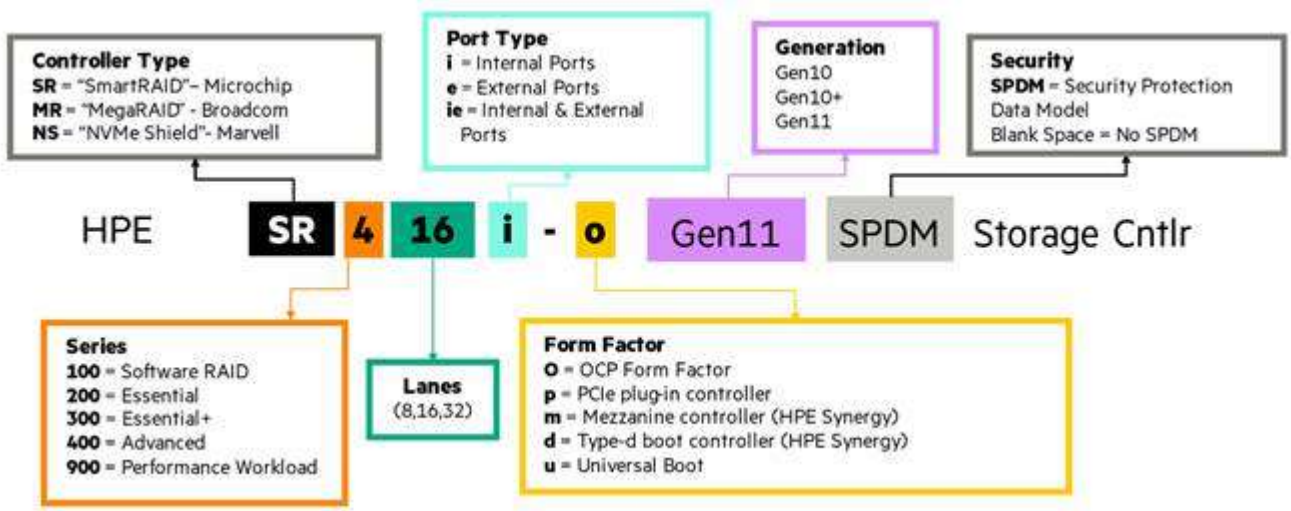
Notes:

- 'NS204i-u Gen11 Hot Plug Boot Opt Dev' requires the selection of 'NS204i-u cable kit' and vice versa.
- Require High-Performance Fan Kit (P48908-B21).
- HPE NS204i-u Gen11 support is limited at 25C ambient while using GPU CTO server.
- Server must use high performance fan kit when supporting NS204i-u for all CTO Configs.
- Ambient temperature for NS204-u is 25C for GPU CTO server.

HPE Storage Controllers

Storage Controllers





Notes:

- When selecting SR RAID controllers for external storage and MR RAID controllers for internal storage, please be aware these two products use different RAID configuration tools. Therefore, there will be a RAID configuration tool for the SR external controller and another for the MR internal controller.
- Mixing MR (Mega RAID) series internal controllers and SR (Smart RAID) series internal Controllers is not supported.
- For more information on the HPE Gen11 Storage Controller, please refer to the following:

HPE Compute MR Gen11 Controllers QuickSpecs

Additional Options

HPE Compute SR Gen11 Controllers QuickSpecs

Essential RAID Controllers

HPE Smart Array E208e-p SR Gen10 (8 External Lanes/No Cache) 12G SAS PCIe Plug-in Controller	804398-B21

Notes: This controller supports up to 8 SAS/SATA Drives (external)
For more information on the HPE Smart Array E208i-p SR Gen10 Controller, please refer to the **QuickSpecs**

Tri-Mode Controllers

HPE MR416i-p Gen11 x16 Lanes 8GB Cache PCI SPDM Plug-in Storage Controller	P47777-B21
HPE MR416i-o Gen11 x16 Lanes 8GB Cache OCP SPDM Storage Controller	P47781-B21
HPE MR216i-p Gen11 x16 Lanes without Cache PCI SPDM Plug-in Storage Controller	P47785-B21
HPE MR216i-o Gen11 x16 Lanes without Cache OCP SPDM Storage Controller	P47789-B21
HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller	P58335-B21
HPE SR932i-p Gen11 x32 Lanes 8GB Wide Cache PCI SPDM Plug-in Storage Controller	P47184-B21

Notes:

- Max 1 of PCIe internal controller can be selected per server.
- Max 2 of PCIe and OROC internal controllers can be selected per server.
- PCIe Internal controller is not supported with 4LFF CTO Server
- OROC controllers not supported with 12LFF CTO Server
- SR932i-p controller is the only controller supported with 8SFF x4 U.2 Drive cage.
- SR932i-p controller can be selected with GPU CTO server only if 8EDSFF drive cage is selected.
- If any controllers are selected, then "HPE 96W Smart Stg Li-ion Batt 145mm Kit" or "HPE Smart Hybrid Capacitor w/ 145mm Cable" must be selected. Vice versa.
- Internal MR-series and Internal SR-series controllers cannot be selected together. Mega RAID Tools cannot be used to script and configure Smart RAID (formerly known as Smart Array) controllers used on HPE Gen9/10/10 Plus/11 servers.
- If MR416i-p/MR216i-p/SR932i-p/MR416i-o/MR216i-o controller is selected along with 8SFF x1 U.3 and 2SFF x4 U.3 drive cage, then one controller is enough to support both the drive cages with the selection of '8SFF cable kit' and '2SFF TM cable kit.'
- If the 8SFF Drive cage is selected with PCIe Internal controller, then the '8SFF TM Cable Kit' must be selected.
- If the 8SFF Drive cage is selected with OROC controller, then '8SFF OROC Cable' must be selected.
- If 8SFF x1 U.3 Drive cage is selected and no Internal controller is selected, then the drive cage will be connected to Direct attach and support SATA Drives only.
- If 8SFF x1 U.3 Drive cage and 2SFF x4 U.3 Drive cage is selected along with Qty=1 of MR216i/MR416i/SR416i/SR932i controller, then 8SFF cable kit OR/AND 2SFF TM cable kit must be selected
- If 8SFF x4 U.3 Drive cage kit is selected and no Internal controller is selected, then the drive cage will be connected to Direct attach and support NVMe U.3 Drives only.
- If the 2SFF Drive cage is selected with PCIe Internal controller or OROC controller, then the '2SFF TM Cable Kit' must be selected.
- If 2SFF U.3 Drive cage is selected and no Internal controller is selected, then the drive cage will be connected to direct attach and supports either Qty=2 of NVMe U.3 drive or SATA drives.
- If 8SFF x4 U.3 Drive cage is selected with 2SFF U.3 drive cage and no Internal controller is selected, then 2SFF U.3 drive cage can support SATA drives only.

Additional Options

- If 8SFF x4 Drive cage and 2SFF drive cage is selected with Qty=1 of Internal controller, then '8SFF cable kit' or '2SFF TM cable kit' must be selected
- If 8SFF x4 Drive cage and 2SFF drive cage is selected with Qty=2 of Internal controller, then both '8SFF cable kit' and '2SFF TM cable kit' must be selected
- If 8SFF x4 Drive cage and 2SFF Drive cage are selected along with MR408i-o controller, then '2SFF TM cable kit' must be selected.
- If SR932i-p is selected with the 8SFF x4 drive cage, the additional 2SFF drive cage will need an extra controller or direct attach.
- MR416i-p is the only controller supported with the GPU Dense CTO server with 4SFF drive cage.
- If SR932i-p/MR416i-p/Gen5 Retimer card is selected along with GPU CTO server then secondary riser must be selected.
- SR932i-p controller can be selected with GPU CTO server only if 8EDSFF drive cage is selected.
- CD7 EDSFF drives requires selection of SR932i-p controller.
- If 8EDSFF drive cage is selected with SR932i-p controller, then 'SR932i-p NVMe Cable Kit' must be selected and defaulted.
- In GPU CTO Server, Mixing of MR416i-p/ Gen5 Retimer card/ SR932i-p controllers are not allowed.

Software RAID

Intel Virtual RAID on CPU Premium E-RTU for HPE	R7J59AAE

Notes:

- Requires UEFI, not supported on Legacy Mode.
- For NVMe SSDs only, no PCIe card support.
- Supports RAID 0/1/5/10.
- Similar to Intel Virtual RAID on CPU Premium FIO Software for HPE (R7J57A) but intended for field deployments (BTO).

NVMe Adapter

HPE DL385 Gen10 Plus 12Gb NVMe 1-port Adapter	P25526-B21
HPE DL385 Gen10 Plus 12Gb NVMe 2-port Adapter	P25527-B21

Notes: The naming is based on HPE ProLiant DL385 Gen10 Plus. However, this option incompatible with HPE ProLiant DL320 Gen11 as well.

HPE Cable Kits

Notes: Some options may not be integrated at the factory. To ensure only valid configurations are ordered, Hewlett Packard Enterprise recommends using an HPE-approved configurator. Contact your local sales representative for additional information.

HPE ProLiant DL320 Gen11 8SFF Tri-Mode Cable Kit	P52775-B21
HPE ProLiant DL320 Gen11 2SFF Tri-Mode Cable Kit	P52776-B21
HPE ProLiant DL320 Gen11 Retimer Card Cable Kit	P52777-B21
HPE ProLiant DL320 Gen11 8SFF OROC Cable Kit	P52780-B21
HPE ProLiant DL320 Gen11 10LFF Direct Attach Cable Kit	P59459-B21
HPE ProLiant DL320 Gen11 12LFF Tri-Mode Cable Kit	P60892-B21

Notes: Customers can self-install the 12LFF tri-mode cable kit. However, due to the complexity involved, HPE recommends customers to call HPE service.

HPE ProLiant DL320 Gen11 2x MR416i-p 8SFF Controller Cable Kit	P66963-B21

Additional Options

Notes:
- Max=1 of 8SFF Cable kit can be selected.
- Not Supported with GPU CTO Server.
- If 8SFF x4 U.3 drive cage is selected along with Qty=2 of MR416i-p then 8SFF 2x MR416i-p Controller Cable Kit must be selected.
- 8SFF 2x MR416i-p Controller Cable Kit Requires selection of 8SFF x4 U.3 drive cage and Qty=2 of MR416i-p controller.
- If 8SFF x4 U.3 drive cage and 2SFF U.3 drive cage is selected along with Qty=2 of MR416i-p controller and 8SFF 2x MR416i-p Controller Cable Kit and no other internal controller is selected then 2SFF x4 U.3 drive cage will be connected to Direct attach and supports either Qty=2 of NVMe U.3 drives or SATA drives.
- Max=1 is not applicable to MR416i-p controller if selected along with 8SFF x4 U.3 drive cage and 8SFF 2x MR416i-p Controller Cable Kit . Max=2 of MR416i-p must be selected for this combination.

HPE ProLiant DL320 Gen11 LFF Optical Cable Kit P54641-B21

Notes: Required if Optical Drive is selected with 4LFF CTO Server.

HPE ProLiant DL320/DL380 Gen11 PCIe Gen5 Retimer Card P48833-B21

Notes:
- Supported only with GPU CTO Server.
- If 4SFF U.3 drive cage is selected with Gen5 Retimer card, then 'HPE DL320 Gen11 4SFF TM Cable Kit' must be selected and defaulted.
- In GPU CTO Server, Mixing of MR416i-p/ Gen5 Retimer card/ SR932i-p controllers are not allowed.
- RAID is allowed on EDSFF drives only if Premium VROC or Standard VROC or Gen5 Retimer card is selected as Primary controller.
- If 8EDSFF Drive cage is selected with Gen5 Retimer card then HPE DL320 Gen11 2xType-p NVMe Cable Kit must be selected and defaulted.
- If 8EDSFF Drive cage is selected with Gen5 Retimer card, then Qty=2 Gen5 Retimer card must be selected and defaulted.
- CD7 EDSFF drives requires selection of Gen5 Retimer card or SR932i-p controller.
- If SR932i-p/MR416i-p/Gen5 Retimer card is selected along with GPU CTO server then secondary riser must be selected.

HPE ProLiant DL320 Gen11 GPU 4x Direct Connected NVMe FIO Enablement Kit P62181-B21

Notes:
- Supported only with GPU CTO Server
- If 8EDSFF or 4SFF Drive cage is selected and MR416i-p is not selected, then "HPE DL320 G11 GPU 4xDC NVMe FIO Kit" must be selected.
- If 8EDSFF Drive cage is selected with "HPE DL320 G11 GPU 4xDC NVMe FIO Kit" then drive cage will be connected to Direct attach and can support Max=4 EDSFF drives only
- If 4SFF U.3 Drive cage is selected with "HPE DL320 G11 GPU 4xDC NVMe FIO Kit" then drive cage will be connected to direct attach and support Max=4 NVMe U.3 drives only
- Cannot be selected if MR416i-p controller is selected.
- Max=1 of 'GPU 4xDC NVMe FIO Kit' or '4SFF TM Cable Kit' or '2xType-p NVMe Cable Kit' can be selected
- If Qty=1 of GPU is selected with GPU CTO Server, then 'GPU 4xDC NVMe FIO Kit' or '2 GPU FIO Enable Kit' must be selected
- In GPU CTO Server if ‘GPU 4xDC NVMe FIO Kit' then the first GPU card selected will be installed in the front of the server. This should not be considered for total PCIe selection.

Additional Options

HPE ProLiant DL320 Gen11 2 GPU FIO Enablement Kit	P62197-B21

Notes:
- If Qty=1 Double Wide GPU is selected with GPU CTO Server, then 'GPU 4xDC NVMe FIO Kit' or '2 GPU FIO Enable Kit' must be selected.
- If Qty=2 Double Wide GPU is selected with GPU CTO Server, then '2 GPU FIO Enable Kit' must be selected
- Cannot be selected with any other GPU enablement kit.
- If '2 GPU FIO Enable Kit' is selected along with 4SFF Drive cage, then MR416i-p controller or Gen5 Retimer card and 4SFF TM Cable Kit must be selected.
- If Qty=1 Single Wide GPU is selected with GPU CTO Server, then 'GPU 4xDC NVMe FIO Kit' or '2 GPU FIO Enable Kit' or '4SW GPU FIO Enable Kit' must be selected.

If Qty=2 Single Wide GPU is selected with GPU CTO server then '2 GPU FIO Enable Kit' or '4SW GPU FIO Enable Kit' must be selected.

HPE ProLiant DL320 Gen11 Serial Port Enablement Kit	P58829-B21

Notes: Max = 1

HPE ProLiant DL320 Gen11 Primary GPU 8-pin Power Cable Kit	P62201-B21
HPE ProLiant DL320 Gen11 Secondary GPU 8-pin Power Cable Kit	P62203-B21
HPE ProLiant DL320 Gen11 Primary GPU 16-pin Power Cable Kit	P62205-B21
HPE ProLiant DL320 Gen11 Secondary GPU 16-pin Power Cable Kit	P62207-B21

Notes:
- P62205-B21 & P62207-B21 require selection of L40 GPU.
- If Qty=1 of L40 GPU is selected, then Primary GPU 16p Power Cable Kit (P62205-B21) must be selected.
- If Qty=2 of L40 GPU is selected, then Primary GPU 16p Power Cable Kit (P62205-B21) and Secondary GPU 8p Power Cable Kit (P62207-B21) must be selected.
- P62201-B21 & P62203-B21 require selection of A16 GPU.
- If Qty=1 of A16 GPU is selected, then Primary GPU 8p Power Cable Kit (P62201-B21) must be selected.
- If Qty=2 of A16 GPU is selected, then Primary GPU 8p Power Cable Kit (P62201-B21) and Secondary GPU 8p Power Cable Kit (P62203-B21) must be selected.

HPE ProLiant DL320 Gen11 Primary GPU 16-pin Advanced Power Cable Kit	P69104-B21
HPE ProLiant DL320 Gen11 Secondary GPU 16-pin Advanced Power Cable Kit	P69106-B21

Notes:
- If Qty=1 of L40s GPU is selected then Prim GPU 16p Adv Power Kit must be selected.
- If Qty=2 of L40s GPU is selected then Prim GPU 16p Adv Power Kit and Sec GPU 16p Advanced Power Kit must be selected.
- Requires selection of L40s GPU

HPE ProLiant DL320 Gen11 4SFF Tri-Mode Cable Kit	P62189-B21

Notes:
- Supported only with GPU CTO Server.
- If 4SFF U.3 drive cage is selected with MR416i-p controller, then 'HPE DL320 Gen11 4SFF TM Cable Kit.'
- Requires selection 4SFF U.3 Drive cage.
- Max=1 of 'GPU 4xDC NVMe FIO Kit' or '4SFF TM Cable Kit' can be selected.

HPE ProLiant DL320 Gen11 2x Type-p NVMe Cable Kit	P62183-B21

Additional Options

Notes:

- Supported only with GPU CTO Server.
- Requires selection of Gen5 Retimer card.
- Requires selection of 8EDSFF Drive cage.
- Max=1 of these kits is allowed for selection GPU 4xDC NVMe FIO Kit or 4SFF TM Cable Ki or Gen5 Retimer cable
- If 8EDSFF Drive cage is selected with Gen5 Retimer card then HPE DL320 Gen11 2xType-p NVMe Cable Kit must be selected and defaulted

HPE ProLiant DL320 Gen11 SR932i-p NVMe Cable Kit P62185-B21

Notes:

- Supported only with GPU CTO Server.
- Requires selection of 8EDSFF Drive cage.
- Max=1 of these kits is allowed for selection GPU 4xDC NVMe FIO Kit or 4SFF TM Cable Ki or Gen5 Retimer cable
- If 8EDSFF drive cage is selected with SR932i-p controller, then 'SR932i-p NVMe Cable Kit' must be selected and defaulted.
- Requires selection of SR932i-p controller

Embedded Management

HPE iLO Advanced

HPE iLO Advanced Electronic License with 1yr Support on iLO Licensed Features E6U59ABE
HPE iLO Advanced 1-server License with 1yr Support on iLO Licensed Features 512485-B21
HPE iLO Advanced AKA Tracking License with 1yr Support on iLO Licensed Features 512487-B21
HPE iLO Advanced Electronic License with 3yr Support on iLO Licensed Features E6U64ABE
HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features BD505A
HPE iLO Advanced AKA Tracking License with 3yr Support on iLO Licensed Features BD507A

HPE Converged Infrastructure Management Software

HPE OneView including 3yr 24x7 Support Flexible Quantity E-LTU E5Y35AAE

Notes: Licenses ship without media. The HPE OneView Media Kit can be ordered separately or **downloaded.**

HPE Security

HPE iLO Common Password FIO Setting P08040-B21

Notes:

- Replaces iLO default randomized password with an HPE-defined common password. HPE highly recommends changing this password immediately after the initial onboarding process.
- Customers who want to choose their own custom iLO default password should use the HPE Factory Express Integration Services.

HPE Bezel Lock Kit 875519-B21

Additional Options

Notes: Requires the bezel kit (P50450-B21).

HPE ProLiant Gen11 1U Common Bezel Kit	P50450-B21
HPE ProLiant DL320 Gen11 Intrusion Detection Cable Kit	P55417-B21

Optional Upgrades

HPE 96W Smart Storage Lithium-ion Battery with 145mm Cable Kit	P01366-B21
HPE Smart Storage Hybrid Capacitor with 145mm Cable Kit	P02377-B21

Notes: Provides backup power for multiple HPE storage controllers or other devices.

HPE Tape Backup

For the complete range of tape drives, autoloaders, libraries, and media, see:

https://www.hpe.com/us/en/storage/storeever-tape-storage.html For hardware and software compatibility of Hewlett Packard Enterprise tape backup products **http://www.hpe.com/storage/BURAcompatibility**

HPE Racks

- Please see the HPE Advanced Series Racks QuickSpecs for information on additional racks options and rack specifications. **HPE G2 Advanced Series Racks**
- Please see the HPE Enterprise Series Racks QuickSpecs for information on additional racks options and rack specifications. **HPE G2 Enterprise Series Racks**

HPE Power Distribution Units (PDUs)

- Please see the **HPE Basic Power Distribution Units (PDU) QuickSpecs** for information on these products and their specifications.
- Please see the **HPE Metered Power Distribution Units (PDU) QuickSpecs** for information on these products and their specifications. Please see the **HPE Intelligent Power Distribution Unit (PDU)** QuickSpecs for information on these products and their specifications.
- Please see the **HPE Metered and Switched Power Distribution Units (PDU)** QuickSpecs for information on these products and their specifications.

Uninterruptible Power Systems (UPS)

- To learn more, please visit the **HPE Uninterruptible Power Systems (UPS)** web page.
- Please see the **HPE Direct Flow Three Phase Uninterruptible Power System QuickSpecs** for information on these products and their specifications.
- Please see the **HPE Line Interactive Single-Phase UPS QuickSpecs** for information on these products and their specifications.

HPE Rack Options

Please see the **HPE KVM Switches web page** for information on these products and their specifications.

Easy Install Rail Kits

Easy Install rail kits contain telescoping rails, which allow for in-rack serviceability.

To assist in installing the server into the rack, an optional installation tool is available by contacting your local services representative.

Notes:

- Hewlett Packard Enterprise recommends that a minimum of two people are required for all Rack Server installations. Please refer to your installation instructions for proper tools and the number of people to use for any installation.
- HPE rail kits are designed to work with HPE racks in compliance with industry standard EIA-310-E. In the event a customer elects to purchase a third-party rack for use with an HPE rail kit, any such use is at customer's own risk. HPE makes no express or implied warranties with respect to such third-party racks and specifically disclaims any implied warranties of merchantability and fitness for a particular purpose. Furthermore, HPE has no obligation and assumes no liability for the materials, design, specifications, installation, safety, and compatibility of any such third-party racks with any rail kits, including HPE rail kits.

HPE USB and SD Options

Notes: In vSphere 7.0, VMware made changes that impact the use of an SD Card/USB media as a standalone boot device and will be removing support for them after version 7.x.
SD Card/USB media can still be used as a standalone boot option through all 7.x releases via published Customer Advisory **Usage of SD Card/USB Devices As Standalone Boot Devices Has Changed Due to System Storage Changes For VMware ESXi 7.0 (Or Later)**.
For any major release beyond VMware ESXi 7.x, VMware will require M.2 or another local persistent device as the standalone boot option.

HPE Support Service

Installation & Startup Services

HPE ProLiant DL/ML Install Service	U4554E
HPE ProLiant DL/ML Startup Service	U4555E

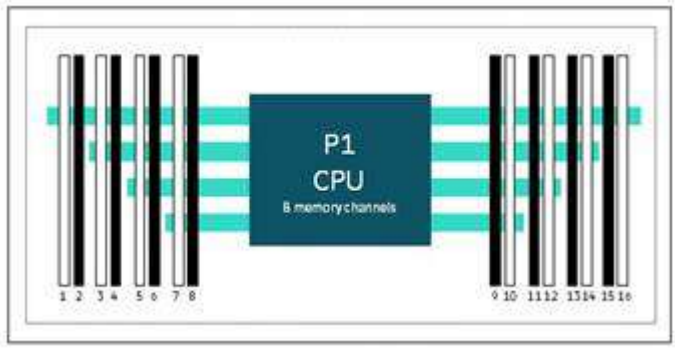
Tech Care

HPE 3 Year Tech Care Essential DL320 Gen11 HW Service	H93M2E
HPE 3 Year Tech Care Essential wDMR DL320 Gen11 HW Service	H93M3E
HPE 5 Year Tech Care Essential DL320 Gen11 HW Service	H93P6E
HPE 5 Year Tech Care Essential wDMR DL320 Gen11 HW Service	H93P7E

Additional Options

Memory

Memory Population guidelines



HPE ProLiant DL320 Gen10 Plus

HPE ProLiant Gen10 Plus 16 slot per CPU DIMM population order																
DIMM population order																
DIMM slot	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**
1 DIMM										10						
2 DIMMs[1]			3							10						
4 DIMMs[1]			3				7			10				14		
6 DIMMs			3		5		7			10				14		16
8 DIMMs[1,2]	1		3		5		7			10		12		14		16
12 DIMMs	1	2	3		5	6	7			10	11	12		14	15	16
16 DIMMs[1,2]	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16

Notes:

- Cells without entries represent configurations not supported, and if populated, the server may result in non-optimal memory performance or other unexpected behavior.
- [1] Support Hemi (hemisphere mode).
- [2] Support Software Guard Extensions (SGX).

General Memory Population Rules and Guidelines

- Install DIMMs only if the corresponding processor is installed.
- If only one processor is installed in a two-processor system, only half of the DIMM slots are available.
- To maximize performance, it is recommended to balance the total memory capacity between all installed processors.
- When two processors are installed, balance the DIMMs across the two processors.
- White DIMM slots denote the first slot to be populated in a channel.
- Mixing of DIMM types (UDIMM, RDIMM, and LRDIMM) is not supported.
- The maximum memory speed is a function of the memory type, memory configuration, and processor model.
- The maximum memory capacity is a function of the number of DIMM slots on the platform, the largest DIMM capacity qualified on the platform, and the number and model of installed processors qualified on the platform.

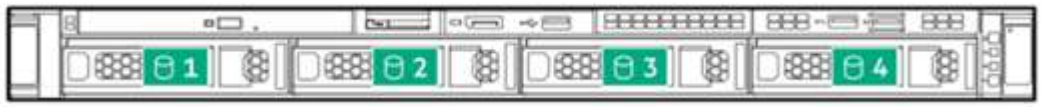
Memory

- For details on the HPE Server Memory Options Population Rules, visit: **http://www.hpe.com/docs/memory-population-rules**
- To realize the performance memory capabilities listed in this document, HPE DDR5 Smart Memory is required. For additional information, please see the **HPE DDR5 Smart Memory QuickSpecs.**

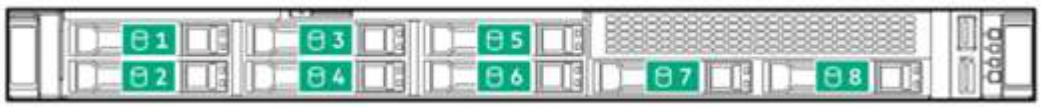
Notes: The maximum memory speed is a function of the memory type, memory configuration, and processor model.

For details on the HPE Server Memory speed, visit: **https://www.hpe.com/docs/server-memory**

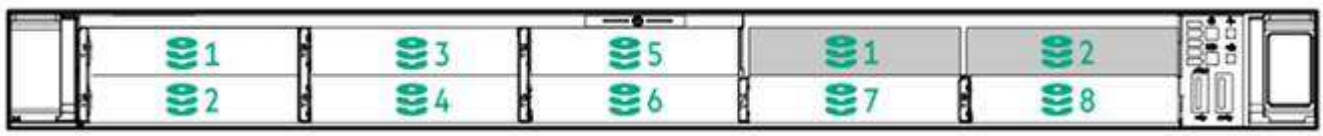
Storage



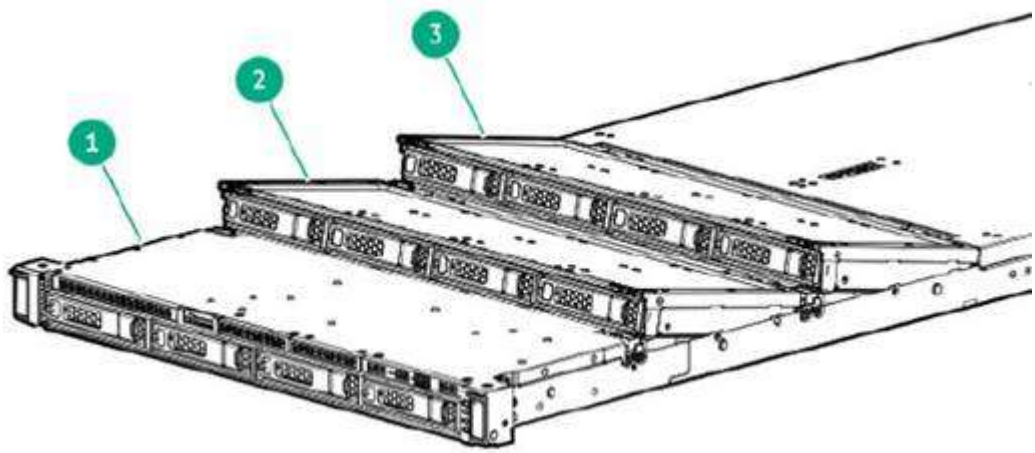
4 LFF device bay numbering



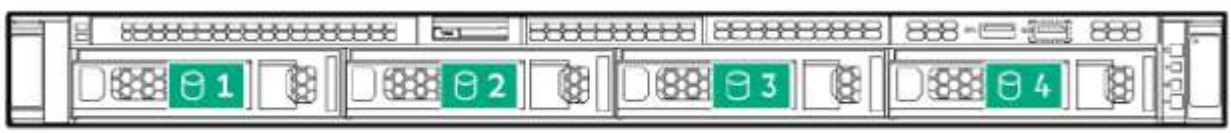
8 SFF drive bay numbering



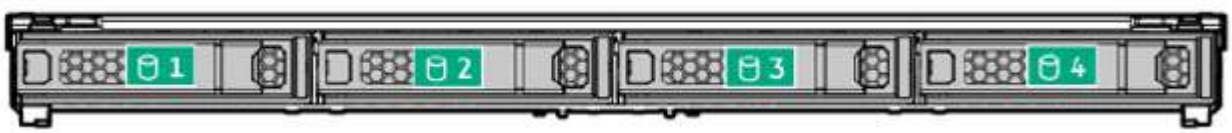
8 SFF + 2 SFF drive bay numbering



12 LFF drive bay positioning



12 LFF drive bay - Box1 numbering



12 LFF drive bay - Box3/Box5 numbering

System Unit

Dimensions (Height x Width x Depth)

8 SFF

- 4.28 x 43.46 x 60.51 cm
  1.69 x 17.11 x 23.82 in

4 LFF CTO Server

- 4.28 x 43.46 x 66.47 cm
  1.69 x 17.11 x 26.17 in

12 LFF CTO Server[1]

- 4.28 x 43.46 x 99.51 cm
  1.69 x 17.11 x 39.18 in

GPU Dense CTO

- 4.28 x 43.46 x 77.42cm
  1.69 x 17.11 x 30.48 in

Notes:

- [1]12 LFF CTO Server requires 1200mm depth rack for proper mounting.

Weight (approximate)

- **8+2 SFF minimum**
  - **12.00 kg (26.45 lb)**: One Drive, one DIMM, one PCIe card, one processor, one heatsink, one power supply, five fans, and two fan blanks.
- **8+2 SFF maximum**
  - **16.71 kg (36.81 lb)**: Ten drives, sixteen DIMMs, two PCIe cards. one processor, one heatsink, two power supplies, and seven fans.
- **4 LFF minimum**
  - **13.20 kg (29.10 lb)**: One drive, one DIMM, one PCIe card, one processor, one heatsink, one power supply, five fans, and two fan blanks.
- **4 LFF maximum**
  - **17.70 kg (39.02 lb)**: Four drives, sixteen DIMMs, and two PCIe standup cards. One processor, one heatsink, two power supplies, and seven fans.
- **12 LFF minimum**
  - **18.70 kg (41.00 lb**): One Drive, one DIMM, one PCIe card, one processor, one heatsink, one power supply, five fans, and two fan blanks.
- **12 LFF maximum**
  - **29.60 kg (65.23 lb**): Twelve drives, sixteen DIMMs, two PCIe cards. one processor, one heatsink, two power supplies, and seven fans.
- **GPU Dense minimum**

  - **14.70 kg (32.41 lb)**: One Drive, one DIMM, one PCIe card, one processor, one heatsink, one power supply, five fans, and two fan blanks.

- **GPU Dense maximum**
  - **20.90 kg (46.08 lb)**: Eight drives, sixteen DIMMs, two PCIe cards. one processor, one heatsink, two power supplies, and seven fans.

Input Requirements (per power supply)

Rated Input Voltage

- 1800-2200W (Titanium):
- 1600W (-48 VDC):
- 1600W (Platinum):
- 1000W (Titanium):
- 800W (Platinum):
- 500W (Platinum): 100-240 VAC

BTU Rating

Maximum

- For 1600W (Platinum)Power Supply: 5918 BTU/hr (at 200 VAC), 5891 BTU/hr (at 230 VAC)
- For 1600W (48VDC) Power Supply: 6026 BTU/hr (at -40 VDC), 6000 BTU/hr (at -48 VDC), 5989 BTU/hr (at -72 VDC)
- For 1000W (Titanium) Power Supply: 3764 BTU/hr (at 100 VAC), 3629BTU/hr (at 200 VAC)
- For 800W (Platinum) Power Supply: 3067 BTU/hr (at 100 VAC), 2958 BTU/hr (at 200 VAC), 2949 BTU/hr (at 240 VAC) China only
- For 500W (Platinum) Power Supply: 1999 BTU/hr (at 100 VAC), 1912 BTU/hr (at 200 VAC), 1904 BTU/hr (at 240 VAC) for China only

---

Power Supply Output (per power supply)

Rated Steady-State Power

- For 1600W (Platinum)Power Supply: 1600 W (at 200 VAC to 240 VAC input), 1600 W (at 240 VDC input)
- For 1600W (48VDC) Power Supply: 1600 W (at -40 VDC to -72 VDC)
- For 1000W (Titanium) Power Supply: 1000 W (at 100 VAC to 127 VAC), 1000 W (at 200 VAC to 240 VAC input)
- For 800W (Platinum) Power Supply: 800 W (at 100 VAC to 127 VAC input), 800 W (at 100 VAC to 240 VAC input), 800 W (at 240 VDC input for China only)
- 3067 BTU/hr (at 100 VAC), 2958 BTU/hr (at 200 VAC), and 2949 BTU/hr (at 240 VAC), China only
- For 500W (Platinum) Power Supply: 500 W (at 100 - 127 VAC input), 500 W (at 100 - 240 VAC input), 500 W (100 VAC - 127 VAC input)

Maximum Peak Power

Technical Specifications

- For 1600W (Platinum)Power Supply: 1600 W (at 200 VAC to 240 VAC input), 1600 W (at 240 VDC input)
- For 1600W (48VDC) Power Supply: 1600 W (at -40 VDC to -72 VDC)
- For 1000W (Titanium) Power Supply: 1000 W (at 100 VAC to 127 VAC), 1000 W (at 200 VAC to 240 VAC input)
- For 800W (Platinum) Power Supply: 800 W (at 100 VAC to 127 VAC input), 800 W (at 100 VAC to 240 VAC input), 800 W (at 240 VDC input for China only)
- For 500W (Platinum) Power Supply: 500 W (at 100 - 127 VAC input), 500 W (at 100 - 240 VAC input), 500 W (100 VAC - 127 VAC input)

System Inlet Temperature

- **Standard Operating Support**

10° to 35°C (50° to 95°F) at sea level with an altitude derating of 1.0°C per every 305 m (1.8°F per every 1000 ft) above sea level to a maximum of 3050 m (10,000 ft), no direct sustained sunlight. The maximum rate of change is 20°C/hr (36°F/hr). The upper limit and rate of change may be limited by the type and number of options installed.

System performance during standard operating support may be reduced if operating with a fan fault or above 30°C (86°F).

- **Extended Ambient Operating Support**

For approved hardware configurations, the supported system inlet range is extended to be: 5° to 10°C (41° to 50°F) and 35° to 40°C (95° to 104°F) at sea level with an altitude derating of 1.0°C per every 175 m (1.8°F per every 574 ft) above 900 m (2953 ft) to a maximum of 3050 m (10,000 ft). The approved hardware configurations for this system are listed at the URL: **https://www.hpe.com/support/ASHRAEGen11**

40°C to 45°C (104°F to 113°F) at sea level with an altitude derating of 1.0°C per every 125 m (1.8°F per every 410 ft) above 900 m (2953 ft) to a maximum of 3050 m (10,000 ft). The approved hardware configurations for this system are listed at the URL: **https://www.hpe.com/support/ASHRAEGen11**

System performance may be reduced if operating in the extended ambient operating range or with a fan fault.

- **Non-operating**

-30° to 60°C (-22° to 140°F). The maximum rate of change is 20°C/hr (36°F/hr).

Relative humidity (non-condensing)

- **Operating**

8% to 90% - Relative humidity (Rh), 28°C (82.4°F) maximum wet bulb temperature, non-condensing.

- **Non-operating**

5 to 95% relative humidity (Rh), 38.7°C (101.7°F) maximum wet bulb temperature, non-condensing.

Altitude

- **Operating**

3050 m (10,000 ft). This value may be limited by the type and number of options installed. The maximum allowable altitude change rate is 457 m/min (1500 ft/min).

- **Non-operating**

9144 m (30,000 ft). The maximum allowable altitude change rate is 457 m/min (1500 ft/min).

Emissions Classification (EMC)

To view the regulatory information for your product, view the Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products, available at the Hewlett Packard Enterprise Support Center:

https://support.hpe.com/hpesc/public/docDisplay?docLocale=en_US&docId=c03471072

Acoustic Noise

Listed are the declared mean A-Weighted sound power levels (LWA,m), declared average bystander position A-Weighted sound pressure levels (LpAm), and the statistical adder for verification, Kv, is a quantity to be added to the declared mean A-weighted sound power level, LWA,m when the product is operating in a 23°C ambient environment. Noise emissions were measured in accordance with ISO 7779 (ECMA 74) and declared in accordance with ISO 9296 (ECMA 109). The listed sound levels apply to standard shipping configurations. Additional options may result in increased sound levels. Please have your HPE representative provide information from the HPE EMESC website for further technical details regarding the configurations listed below.

Test case	Entry 1 (BTO) 4LFF	Entry 2 (BTO) SFF	Mid Range (BTO) SFF	Performance (BTO) SFF
Idle				
LWA,m	4.7 B	4.7 B	4.9 B	4.9 B
LpAm	36 dBA	35 dBA	37 dBA	37 dBA
Kv	0.4 B	0.4 B	0.4 B	0.4 B
Operating				
LWA,m	5.3 B	4.9 B	5.1 B	5.2 B
LpAm	39 dBA	36 dBA	38 dBA	40 dBA
Kv	0.4 B	0.4 B	0.4 B	0.4 B

Notes:

- The Acoustics levels presented here are generated by the test configuration only. Acoustics levels will vary depending on system configuration. Values are subject to change without notification and are for reference only.
- The declared mean A-weighted sound power level, LWA,m, is computed as the arithmetic average of the measured.
- A-weighted sound power levels for a randomly selected sample rounded to the nearest 0,1 B.
- The declared mean A-weighted emission sound pressure level, LpA,m, is computed as the arithmetic average of the measured A-weighted emission sound pressure levels at the bystander positions for a randomly selected sample, rounded to the nearest 1 dB.
- The statistical adder for verification, Kv, is a quantity to be added to the declared mean A-weighted sound power level, LWA,m, such that there will be a 95 % probability of acceptance when using the verification procedures of ISO 9296 if no more than 6,5 % of the batch of new equipment, has A-weighted sound

power levels greater than (LWA,m + Kv).

- The quantity, LWA,c (formerly called LWAd), can be computed from the sum of LWA,m, and Kv.
- All measurements were made to conform to ISO 7779 / ECMA-74 and declared to conform to ISO 9296 / ECMA-109.
- B, dB, abbreviations for bels and decibels, respectively, where 1 B = 10 dB.
- The results in this declaration apply only to the model numbers listed above when operating and tested according to the indicated modes and standards. A system with additional configuration components or increased operating functionality may increase the noise emission values.
- The system under abnormal conditions may increase the noise level, and persons in the vicinity of the product [cabinet] for extended periods should consider wearing hearing protection or using other means to reduce noise exposure.

Environment-friendly Products and Approach End-of-life Management and Recycling

Hewlett Packard Enterprise offers **end-of-life product return, trade-in, and recycling programs**, in many geographic areas, for our products. Products returned to Hewlett Packard Enterprise will be recycled, recovered, or disposed of in a responsible manner.

The EU WEEE directive (2002/95/EC) requires manufacturers to provide treatment information for each product type used by treatment facilities. This information (product disassembly instructions) is posted on the Hewlett Packard Enterprise website. These instructions may be used by recyclers and other WEEE treatment facilities, as well as Hewlett Packard Enterprise OEM customers who integrate and re-sell Hewlett Packard Enterprise equipment.

Summary of Changes

Date	Version History	Action	Description of Change
06-Jan-2025	Version 26	Changed	Optional Features and Core Options sections were updated.
02-Dec-2024	Version 25	Changed	Core Options section was updated.
07-Oct-2024	Version 24	Changed	Core Options section was updated.
03-Sep-2024	Version 23	Changed	Standard Features (Operating Systems and Virtualization Software Support for HPE Servers) and Pre-Configured Models sections were updated.
05-Aug-2024	Version 22	Changed	Pre-Configured Models (TPM China), Core Options and Additional Options sections were updated.
01-Jul-2024	Version 21	Changed	Pre-Configured Models and Configuration Information sections were updated.
03-Jun-2024	Version 20	Changed	Core Options section was updated.
06-May-2024	Version 19	Changed	Core Options section was updated.
01-Apr-2024	Version 18	Changed	Standard Features, Configuration Information, Core Options and Additional Options sections were updated
04-Mar-2024	Version 17	Changed	Standard Features, Configuration Information, Core Options and Additional Options sections were updated
05-Feb-2024	Version 16	Changed	Configuration Information and Core Options sections were updated
14-Dec-2023	Version 15	Changed	Overview, Standard Features, Configuration Information, Core Options and Additional Options sections were updated
13-Nov-2023	Version 14	Changed	Pre-configured Models Section was updated.
06-Nov-2023	Version 13	Changed	Service and Support, Configuration Information and Additional Options sections were updated
02-Oct-2023	Version 12	Changed	Overview and Core Options sections were updated.
05-Sep-2023	Version 11	Changed	Standard Features, Configuration Information, Core Options and Additional Options sections were updated
07-Aug-2023	Version 10	Changed	Standard Features,Core Options and Additional Options sections were updated
17-Jul-2023	Version 9	Changed	Overview section was updated.
10-Jul-2023	Version 8	Changed	Standard Features, Pre-Configured , Configuration Information and Additional Options sections were updated
05-Jun-2023	Version 7	Changed	Standard Features, Pre-Configured , Configuration Information and Additional Options sections were updated.
01-May-2023	Version 6	Changed	Standard Features, and Core Options sections were updated.
17-Apr-2023	Version 5	Changed	Overview, Standard Features, and Optional Features sections were updated.
03-Apr-2023	Version 4	Changed	Pre-Configured and Core Options section were updated.
06-Mar-2023	Version 3	Changed	Overview, Standard Features, Configuration Information, Additional Options and Technical Specifications sections were updated.
06-Feb-2023	Version 2	Changed	Overview, Core Options, and Storage Options sections were updated.
10-Jan-2023	Version 1	New	New QuickSpecs.

Copyright

Make the right purchase decision. Contact our presales specialists.





© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

a50004315enw - 16919 - Worldwide - V26 - 06-January-2025

Gemalto GemPC Twin/TR (ID Bridge CT30)



The IDBridge CT30 (formerly the GemPC Twin/TR) USB desktop reader from Gemalto integrates high quality and functionality with a contemporary design. This reader is fully PC/SC and EMV Level 1 compliant and supports all ISO 7816 microprocessor-based smartcards as well as many popular memory cards. The reader is compatible with previous Gemalto models such as the GCR410.

Simple plug-and-play operation on all versions of Windows (WHQL approved Windows drivers included). The ideal smart card reader for Mac / iMac with USB-A ports with downloadable drivers for MacOS. Linux drivers also available for download from the Thales (Gemalto / SafeNet) support portal.

The IDBridge CT30 is compact and boasts performance to match many other leading brand readers. This unique transparent reader even allows you to see the card when inserted. The latest TR version has a fixed USB cable, but is otherwise identical to the previous Twin model.

There is a software and hardware proof of concept kit for PKI-based logon, which includes this card reader, listed on the Related tab.

Gemalto part number 962-000009-004

©2022 DOT ORIGIN LTD

This unique transparent USB contact reader is highly compact and boasts high performance to match.

To buy, visit:
https://www.smartcardfocus.com/shop/ilp/id~463/p/index.shtml

This Product Briefing has been produced by Dot Origin Ltd, the smart card experts behind SmartcardFocus.com. If you have a query email sales@smartcardfocus.com or call us on +44 (0)1428 685250.

©2022 DOT ORIGIN LTD

Technical Specifications

Product:

Manufacturer: Gemalto

Credential type(s): Contact

Form factor: Portable

Connector: USB CCID (Standard 2.0 Type A connector)

OS support: Up to Windows 10 (including 8.1, 7); Linux; Mac OS

Contact Smartcard Interface:

Contact interface standard: ISO 7816

Contact card types: Processor cards (T=0 and T=1)

Card Voltage 5V, 3V & 1.8V

Data transmission protocol: T=0 / T=1

Card contacts: Sliding

Contact read/write speed: 2kbps - 826kbps

Clock frequency: 4.8MHz

Card insertion cycles: Min 100 000

Host Interface:

Host interface standards: PC/SC

Host data transmission speed (wired): 12 Mbps (USB full speed)

Other Features:

Card detection: Card insertion/extraction detection, Short circuit protection

Status indicator: Green LED with dual state (waiting and operation)

An additional desktop stand is available as an accessory.

Certifications:

Manufacturer's warranty: 2 year

Type approvals: CE, FCC part 15, RoHS 2, REACH, WEEE, IEC 60950, c-Tick

Application approvals: EMV Level 1 (Contact), c-Tick (Australia), VCCI (Japan), BSMI (Taiwan)

Country of origin: China

Physical:

Dimensions: 74 x 63 x13mm (2.9 x 2.5 x 0.5in)

Weight (±5%): 55g (1.9oz)

Cable length supplied: 1.5m (59.1in)

Power:

Supply voltage: 5V DC

Operating current: Less than 200µA (in suspend mode)

Battery power option: No

Environment:

Operating temperature: 0 to 55degC (32 to 131degF)

Storage temperature: 0 to 70degC (32 to 158degF)

Operating humidity: 10-95% RH (non-condensing)

Manufacturer's part number: 962-000009-004

Manufacturer: Thales Gemalto

Connection Method: USB

©2022 DOT ORIGIN LTD

Windows identity and privacy

Article • 12/10/2022 • 2 minutes to read

Malicious actors launch millions of password attacks every day. Weak passwords, password spraying, and phishing are the entry point for many attacks. Knowing that the right user is accessing the right device and the right data is critical to keeping your business, family, and self, safe and secure. Windows Hello, Windows Hello for Business, and Credential Guard enable customers to move to passwordless multifactor authentication (MFA). MFA can reduce the risk of compromise in organizations.

Security capabilities	**Description**
Securing user identity with Windows Hello	Windows Hello and Windows Hello for Business replace password-based authentication with a stronger authentication model to sign into your device using a passcode (PIN) or other biometric based authentication. This PIN or biometric based authentication is only valid on the device that you registered it for and cannot be used on another deviceLearn more: Windows Hello for Business
Windows Defender Credential Guard and Remote Credential Guard	Windows Defender Credential Guard helps protects your systems from credential theft attack techniques (pass-the-hash or pass-the-ticket) as well as helping prevent malware from accessing system secrets even if the process is running with admin privileges. Windows Defender Remote Credential Guard helps you protect your credentials over a Remote Desktop connection by redirecting Kerberos requests back to the device that's requesting the connection. It also provides single sign-on experiences for Remote Desktop sessions. Learn more: Protect derived domain credentials with Windows Defender Credential Guard and Protect Remote Desktop credentials with Windows Defender Remote Credential Guard
FIDO Alliance	Fast Identity Online (FIDO) defined protocols are becoming the open standard for providing strong authentication that helps prevent phishing and are user-friendly and privacy-respecting. Windows 11 supports the use of device sign-in with FIDO 2 security keys, and with Microsoft Edge or other modern browsers, supports the use of secure FIDO-backed credentials to keep user accounts protected. Learn more about the FIDO Alliance.
Microsoft Authenticator	The Microsoft Authenticator app is a perfect companion to help keep secure with Windows 11. It allows easy, secure sign-ins for all your online accounts using multi-factor authentication, passwordless phone sign-in, or password autofill. You also have additional account management options for your Microsoft personal, work, or school accounts. Microsoft Authenticator can be used to set up multi-factor authentication for your users. Learn more: Enable passwordless sign-in with the Microsoft Authenticator app.

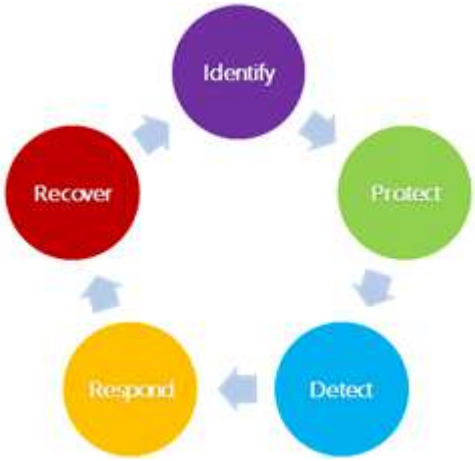
Security capabilities	Description
Smart Cards	Smart cards are tamper-resistant portable storage devices that can enhance the security of tasks in Windows, such as authenticating clients, signing code, securing e-mail, and signing in with Windows domain accounts. Learn more about Smart Cards.
Access Control	Access control is the process of authorizing users, groups, and computers to access objects and assets on a network or computer. Computers can control the use of system and network resources through the interrelated mechanisms of authentication and authorization. Learn more: Access Control.

Windows Credential Theft Mitigation Guide Abstract

Article • 12/09/2022 • 2 minutes to read • Applies to: ✅ Windows 10, ✅ Windows 11

This topic provides a summary of the Windows credential theft mitigation guide, which can be downloaded from the Microsoft Download Center. This guide explains how credential theft attacks occur and the strategies and countermeasures you can implement to mitigate them, following these security stages:

- Identify high-value assets
- Protect against known and unknown threats
- Detect pass-the-hash and related attacks
- Respond to suspicious activity
- Recover from a breach



Attacks that steal credentials

Learn about the different types of attacks that are used to steal credentials, and the factors that can place your organization at risk. The types of attacks that are covered include:

- Pass the hash
- Kerberos pass the ticket
- Kerberos golden ticket and silver ticket
- Key loggers
- Shoulder surfing

Credential protection strategies

This part of the guide helps you consider the mindset of the attacker, with prescriptive guidance about how to prioritize high-value accounts and computers. You'll learn how to architect a defense against credential theft:

- Establish a containment model for account privileges
- Harden and restrict administrative hosts
- Ensure that security configurations and best practices are implemented

Technical countermeasures for credential theft

Objectives and expected outcomes are covered for each of these countermeasures:

- Use Windows 10 with Credential Guard
- Restrict and protect high-privilege domain accounts
- Restrict and protect local accounts with administrative privileges
- Restrict inbound network traffic

Many other countermeasures are also covered, such as using Microsoft Passport and Windows Hello, or multifactor authentication.

Detecting credential attacks

This sections covers how to detect the use of stolen credentials and how to collect computer events to help you detect credential theft.

Responding to suspicious activity

Learn Microsoft's recommendations for responding to incidents, including how to recover control of compromised accounts, how to investigate attacks, and how to recover from a breach.

Windows Hello for Business documentation

Learn how to manage and deploy Windows Hello for Business.

About Windows Hello For Business

OVERVIEW

Windows Hello for Business Overview

CONCEPT

Passwordless Strategy

Why a PIN is better than a password

Windows Hello biometrics in the enterprise

How Windows Hello for Business works

TRAINING

Technical Deep Dive - Device Registration

Technical Deep Dive - Provisioning

Technical Deep Dive - Authentication

Technology and Terminology

Frequently Asked Questions (FAQ)

Configure and manage Windows Hello for Business

CONCEPT

Windows Hello for Business Deployment Overview

Planning a Windows Hello for Business Deployment

Deployment Prerequisite Overview

HOW-TO GUIDE

Hybrid Cloud Kerberos Trust Deployment

Hybrid Azure AD Joined Key Trust Deployment

Hybrid Azure AD Joined Certificate Trust Deployment

On-premises SSO for Azure AD Joined Devices

On-premises Key Trust Deployment

On-premises Certificate Trust Deployment

TRAINING

Manage Windows Hello for Business in your organization

Windows Hello and password changes

Prepare people to use Windows Hello

Windows Hello for Business Features

HOW-TO GUIDE

Conditional Access

PIN Reset

Dual Enrollment

Dynamic Lock

Multi-factor Unlock

Remote Desktop

Windows Hello for Business Troubleshooting

HOW-TO GUIDE

Known Deployment Issues

Errors During PIN Creation

Enable passwordless security key sign-in

Article • 03/10/2023 • 4 minutes to read

For enterprises that use passwords today and have a shared PC environment, security keys provide a seamless way for workers to authenticate without entering a username or password. Security keys provide improved productivity for workers, and have better security.

This document focuses on enabling security key based passwordless authentication. At the end of this article, you'll be able to sign in to web-based applications with your Azure AD account using a FIDO2 security key.

Requirements

- Azure AD Multi-Factor Authentication
- Enable Combined security information registration
- Compatible FIDO2 security keys
- WebAuthN requires Windows 10 version 1903 or higher

To use security keys for logging in to web apps and services, you must have a browser that supports the WebAuthN protocol. These include Microsoft Edge, Chrome, Firefox, and Safari.

Prepare devices

For Azure AD joined devices, the best experience is on Windows 10 version 1903 or higher.

Hybrid Azure AD joined devices must run Windows 10 version 2004 or higher.

Enable passwordless authentication method

Enable the combined registration experience

Registration features for passwordless authentication methods rely on the combined registration feature. Follow the steps in the article Enable combined security information registration, to enable combined registration.

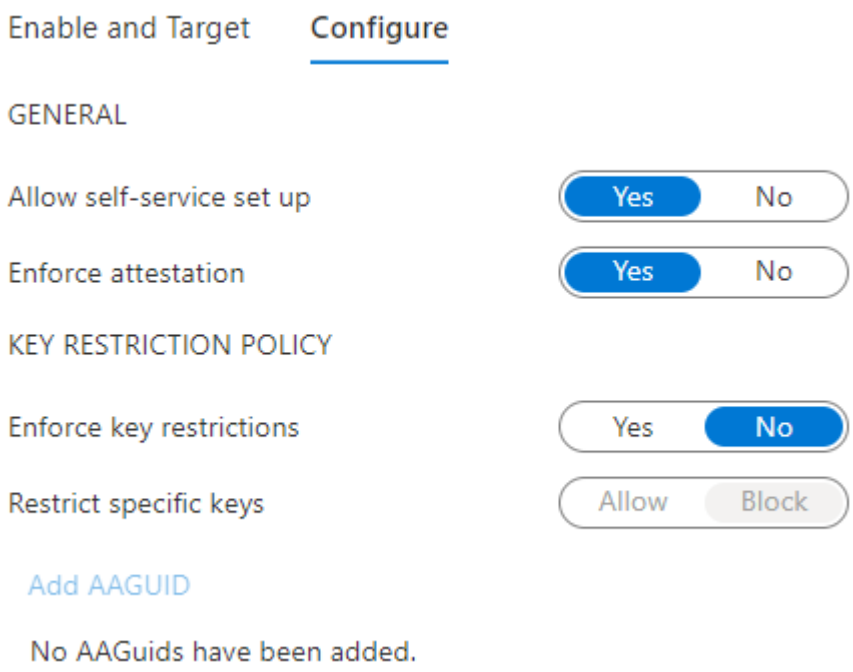
Enable FIDO2 security key method

1. Sign in to the Azure portal.
2. Browse to **Azure Active Directory** > **Security** > **Authentication methods** > **Authentication method policy**.
3. Under the method **FIDO2 Security Key**, click **All users**, or click **Add groups** to select specific groups.
4. **Save** the configuration.

> **Note**
>
> If you see an error when you try to save, the cause might be due to the number of users or groups being added. As a workaround, replace the users and groups you are trying to add with a single group, in the same operation, and then click **Save** again.

FIDO Security Key optional settings

There are some optional settings on the **Configure** tab to help manage how security keys can be used for sign-in.



- **Allow self-service set up** should remain set to **Yes**. If set to no, your users won't be able to register a FIDO key through the MySecurityInfo portal, even if enabled by Authentication Methods policy.
- **Enforce attestation** setting to **Yes** requires the FIDO security key metadata to be published and verified with the FIDO Alliance Metadata Service, and also pass

Microsoft's additional set of validation testing. For more information, see What is a Microsoft-compatible security key?

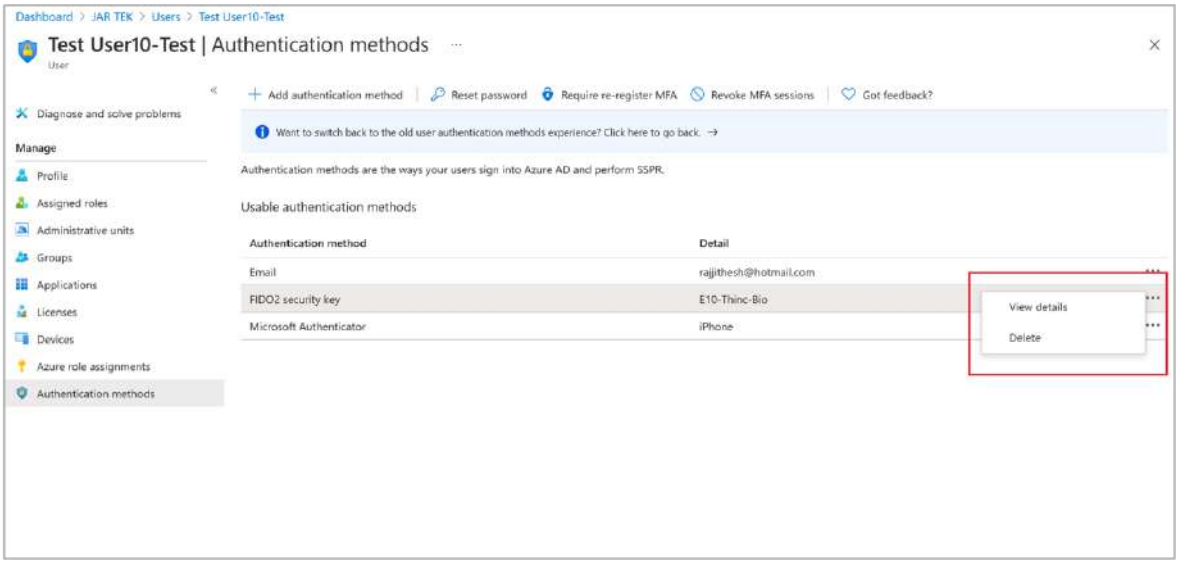
Key Restriction Policy

- **Enforce key restrictions** should be set to **Yes** only if your organization wants to only allow or disallow certain FIDO security keys, which are identified by their AAGuids. You can work with your security key provider to determine the AAGuids of their devices. If the key is already registered, AAGUID can also be found by viewing the authentication method details of the key per user.

Disable a key

To remove a FIDO2 key associated with a user account, delete the key from the user's authentication method.

1. Sign in to the Azure portal and search for the user account from which the FIDO key is to be removed.
2. Select **Authentication methods** > right-click **FIDO2 security key** and click **Delete**.



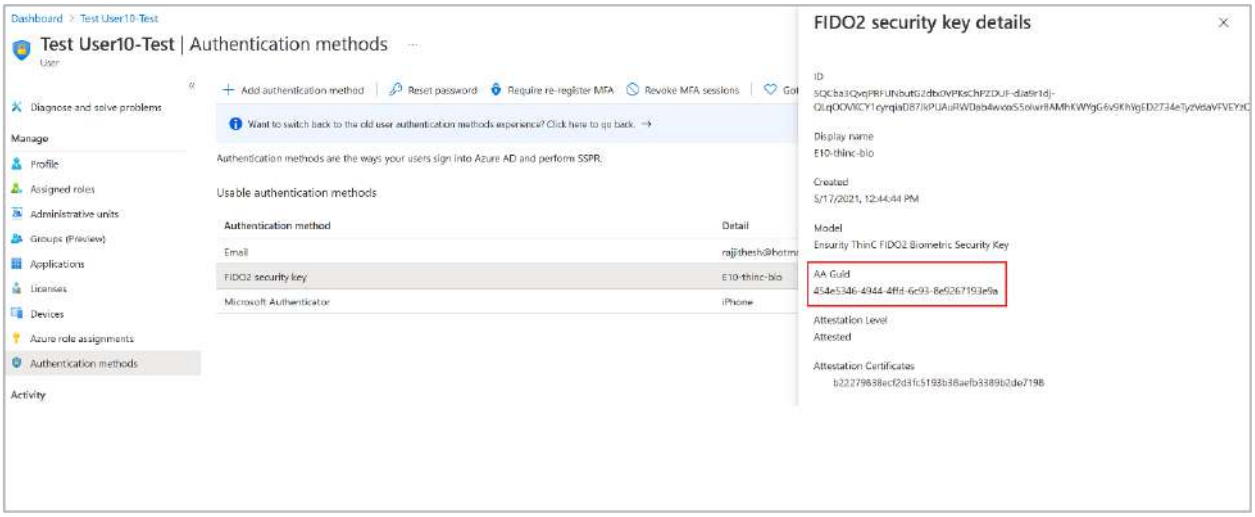
Security key Authenticator Attestation GUID (AAGUID)

The FIDO2 specification requires each security key provider to provide an Authenticator Attestation GUID (AAGUID) during attestation. An AAGUID is a 128-bit identifier indicating the key type, such as the make and model.

> ⓘ **Note**

The manufacturer must ensure that the AAGUID is identical across all substantially identical keys made by that manufacturer, and different (with high probability) from the AAGUIDs of all other types of keys. To ensure, the AAGUID for a given type of security key should be randomly generated. For more information, see **Web Authentication: An API for accessing Public Key Credentials - Level 2 (w3.org)**.

There are two ways to get your AAGUID. You can either ask your security key provider or view the authentication method details of the key per user.

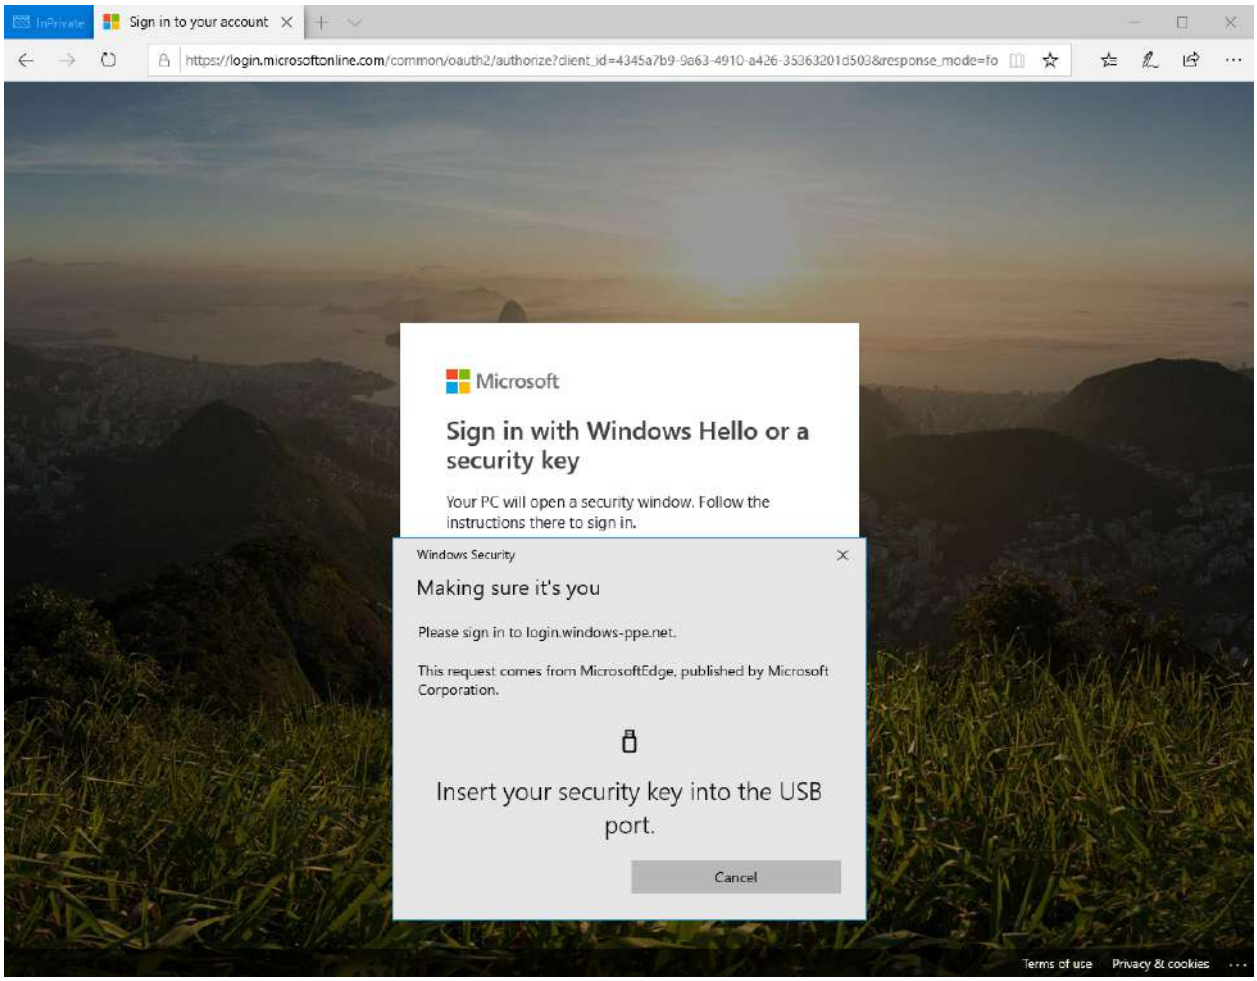


User registration and management of FIDO2 security keys

1. Browse to https://myprofile.microsoft.com .
2. Sign in if not already.
3. Click **Security Info**.
   a. If the user already has at least one Azure AD Multi-Factor Authentication method registered, they can immediately register a FIDO2 security key.
   b. If they don't have at least one Azure AD Multi-Factor Authentication method registered, they must add one.
   c. An Administrator can issue a Temporary Access Pass to allow the user to register a Passwordless authentication method.
4. Add a FIDO2 Security key by clicking **Add method** and choosing **Security key**.
5. Choose **USB device** or **NFC device**.
6. Have your key ready and choose **Next**.
7. A box will appear and ask the user to create/enter a PIN for your security key, then perform the required gesture for the key, either biometric or touch.
8. The user will be returned to the combined registration experience and asked to provide a meaningful name for the key to identify it easily. Click **Next**.
9. Click **Done** to complete the process.

Sign in with passwordless credential

In the example below a user has already provisioned their FIDO2 security key. The user can choose to sign in on the web with their FIDO2 security key inside of a supported browser on Windows 10 version 1903 or higher.



Troubleshooting and feedback

If you'd like to share feedback or encounter issues with this feature, share via the Windows Feedback Hub app using the following steps:

1. Launch **Feedback Hub** and make sure you're signed in.
2. Submit feedback under the following categorization:

    - Category: Security and Privacy
    - Subcategory: FIDO

3. To capture logs, use the option to **Recreate my Problem**.

Known issues

Security key provisioning

Administrator provisioning and de-provisioning of security keys isn't available.

UPN changes

If a user's UPN changes, you can no longer modify FIDO2 security keys to account for the change. The solution for a user with a FIDO2 security key is to sign in to MySecurityInfo, delete the old key, and add a new one.

Next steps

FIDO2 security key Windows 10 sign in

Enable FIDO2 authentication to on-premises resources

Learn more about device registration

Learn more about Azure AD Multi-Factor Authentication

Additional resources

Documentation

Browser support of FIDO2 passwordless authentication - Microsoft Entra

Browsers and operating system combinations support FIDO2 passwordless authentication for apps using Azure Active Directory

Plan a passwordless authentication deployment in Azure Active Directory - Microsoft Entra

Directions for deploying passwordless authentication

Azure Active Directory passwordless sign-in - Microsoft Entra

Learn about options for passwordless sign-in to Azure Active Directory using FIDO2 security keys or Microsoft Authenticator

Passwordless security key sign-in Windows - Azure Active Directory - Microsoft Entra

Learn how to enable passwordless security key sign-in to Azure Active Directory using FIDO2 security keys

Passwordless sign-in with Microsoft Authenticator - Azure Active Directory - Microsoft Entra

Enable passwordless sign-in to Azure AD using Microsoft Authenticator

Overview of Azure Active Directory authentication strength (preview) - Microsoft Entra

Learn how admins can use Azure AD Conditional Access to distinguish which authentication methods can be used based on relevant security factors.

FAQs for hybrid FIDO2 security key deployment - Azure Active Directory - Microsoft Entra

Learn about some frequently asked questions for passwordless hybrid FIDO2 security key sign-in using Azure Active Directory

How to migrate to the Authentication methods policy - Azure Active Directory (preview) - Microsoft Entra

Learn about how to centrally manage multifactor authentication (MFA) and self-service password reset (SSPR) settings in the Authentication methods policy.

Show 5 more

What is Windows LAPS?

Article • 02/16/2023 • 4 minutes to read

Windows Local Administrator Password Solution (Windows LAPS) is a Windows feature that automatically manages and backs up the password of a local administrator account on your Azure Active Directory-joined or Windows Server Active Directory-joined devices. You also can use Windows LAPS to automatically manage and back up the Directory Services Repair Mode (DSRM) account password on your Windows Server Active Directory domain controllers. An authorized administrator can retrieve the DSRM password and use it.

> **ⓘ Important**
>
> Windows LAPS currently is available only in **Windows 11 Insider Preview Build 25145 and later** and the Azure Active Directory LAPS scenario is in private preview. For more information see **Windows LAPS availability and Azure AD LAPS public preview status**.

Windows LAPS supported platforms and Azure AD LAPS preview status

Windows LAPS currently is available only in Windows 11 Insider Preview Build 25145 and later.

The Azure Active Directory LAPS scenario is currently in private preview and is closed to new customers. The Azure Active Directory LAPS scenario will enter public preview in Q2 2023 once all platforms listed below have been updated.

Windows LAPS will be backported to the following OS platforms no later than Q2 2023:

- Windows 10 20H2 and later
- Windows 11 21H2 and later
- Windows Server 2019 and later

Server Core editions will be updated at the same time as their respective desktop Server editions.

This documentation will be updated once a more precise date for the Windows backports and the Azure Active Directory LAPS scenario is available.

Benefits of using Windows LAPS

Use Windows LAPS to regularly rotate and manage local administrator account passwords and get these benefits:

- Protection against pass-the-hash and lateral-traversal attacks
- Improved security for remote help desk scenarios
- Ability to sign in to and recover devices that are otherwise inaccessible
- A fine-grained security model (access control lists and optional password encryption) for securing passwords that are stored in Windows Server Active Directory
- Support for the Azure role-based access control model for securing passwords that are stored in Azure Active Directory

Watch this video to learn about Windows LAPS.
https://www.youtube-nocookie.com/embed/jdEDIXm4JgU

Key Windows LAPS scenarios

You can use Windows LAPS for several primary scenarios:

- Back up local administrator account passwords to Azure Active Directory (for Azure Active Directory-joined devices)
- Back up local administrator account passwords to Windows Server Active Directory (for Windows Server Active Directory-joined clients and servers)
- Back up DSRM account passwords to Windows Server Active Directory (for Windows Server Active Directory domain controllers)
- Back up local administrator account passwords to Windows Server Active Directory by using legacy Microsoft LAPS

In each scenario, you can apply different policy settings.

Windows LAPS doesn't support Azure Active Directory workplace-joined clients.

Understand device join state restrictions

Whether a device is joined to Azure Active Directory or Windows Server Active Directory determines how you can use Windows LAPS.

Devices that are joined only to Azure Active Directory can back up passwords only to Azure Active Directory.

Devices that are joined only to Windows Server Active Directory can back up passwords only to Windows Server Active Directory.

Devices that are hybrid-joined (joined to both Azure Active Directory and Windows Server Active Directory) can back up their passwords either to Azure Active Directory or to Windows Server Active Directory. You can't back up passwords to both Azure Active Directory and Windows Server Active Directory.

Set Windows LAPS policy

To set up and manage policy for your Windows LAPS deployment, you have multiple options:

- Windows LAPS configuration service provider (CSP)
- Windows LAPS Group Policy
- Legacy Microsoft LAPS

Manage and monitor Windows LAPS

You also have various options to manage and monitor Windows LAPS.

Options for Windows include:

- The Windows Server Active Directory Users and Computers properties dialog
- A dedicated event log channel
- A Windows PowerShell module that's specific to Windows LAPS

Azure-based monitoring and reporting solutions are available when you back up passwords to Azure Active Directory.

Windows LAPS vs. legacy Microsoft LAPS

You can still download an earlier version of Local Administrator Password Solution, legacy Microsoft LAPS.

Windows LAPS inherits many design concepts from legacy Microsoft LAPS. If you're familiar with legacy Microsoft LAPS, many Windows LAPS features will be familiar. A key difference is that Windows LAPS is an entirely separate implementation that's native to Windows. Windows LAPS also adds many features that aren't available in legacy

Microsoft LAPS. You can use Windows LAPS to back up passwords to Azure Active Directory, encrypt passwords in Windows Server Active Directory, and store your password history.

> **Important**
>
> Windows LAPS doesn't require you to install legacy Microsoft LAPS. You can fully deploy and use all Windows LAPS features without installing or referring to legacy Microsoft LAPS. But to help migrate an existing legacy Microsoft LAPS deployment, Windows LAPS offers **legacy Microsoft LAPS emulation mode**.

See also

Legacy Microsoft LAPS

Next steps

- Key concepts in Windows LAPS
- Get started with Windows LAPS for Windows Server Active Directory
- Get started with Windows LAPS for Azure Active Directory
- Get started with Windows LAPS in legacy Microsoft LAPS emulation mode

Enterprise Certificate Pinning

Article • 12/09/2022 • 12 minutes to read • Applies to: ✅ Windows 10, ✅ Windows 11

Enterprise certificate pinning is a Windows feature for remembering, or pinning a root issuing certificate authority or end entity certificate to a given domain name. Enterprise certificate pinning helps reduce man-in-the-middle attacks by enabling you to protect your internal domain names from chaining to unwanted certificates or to fraudulently issued certificates.

> ⓘ **Note**
>
> External domain names, where the certificate issued to these domains is issued by a public certificate authority, are not ideal for enterprise certificate pinning.

Windows Certificate APIs (CertVerifyCertificateChainPolicy and WinVerifyTrust) are updated to check if the site's chain that authenticates servers matches a restricted set of certificates. These restrictions are encapsulated in a Pin Rules Certificate Trust List (CTL) that is configured and deployed to Windows 10 computers. Any site certificate that triggers a name mismatch causes Windows to write an event to the CAPI2 event log and prevents the user from navigating to the web site using Microsoft Edge or Internet Explorer.

> ⓘ **Note**
>
> Enterprise Certificate Pinning feature triggering doesn't cause clients other than Microsoft Edge or Internet Explorer to block the connection.

Deployment

To deploy enterprise certificate pinning, you need to:

- Create a well-formatted certificate pinning rule XML file
- Create a pin rules certificate trust list file from the XML file
- Apply the pin rules certificate trust list file to a reference administrative computer
- Deploy the registry configuration on the reference computer using Group Policy Management Console (GPMC), which is included in the Remote Server Administration Tools (RSAT).

Create a Pin Rules XML file

The XML-based pin rules file consists of a sequence of PinRule elements. Each PinRule element contains a sequence of one or more Site elements and a sequence of zero or more Certificate elements.

```
code
```

```
<PinRules ListIdentifier="PinRulesExample" Duration="P28D">

  <PinRule Name="AllCertificateAttributes" Error="None" Log="true">
    <Certificate File="Single.cer"/>
    <Certificate File="Multiple.p7b"/>
    <Certificate File="Multiple.sst"/>
    <Certificate Directory="Multiple"/>
    <Certificate Base64="MIIBy … QFzuM"/>
    <Certificate File="WillExpire.cer" EndDate="2015-05-12T00:00:00Z"/>
    <Site Domain="xyz.com"/>
  </PinRule>

  <PinRule Name="MultipleSites" Log="false">
    <Certificate File="Root.cer"/>
    <Site Domain="xyz.com"/>
    <Site Domain=".xyz.com"/>
    <Site Domain="*.abc.xyz.com" AllSubdomains="true"/>
    <Site Domain="WillNormalize.com"/>
  </PinRule>

</PinRules>
```

PinRules Element

The PinRules element can have the following attributes. For help with formatting Pin Rules, see Representing a Date in XML or Representing a Duration in XML.

Attribute	Description	Required
Duration or **NextUpdate**	Specifies when the Pin Rules will expire. Either is required. **NextUpdate** takes precedence if both are specified. **Duration**, represented as an XML TimeSpan data type, doesn't allow years and months. You represent the **NextUpdate** attribute as an XML DateTime data type in UTC.	**Required?** Yes. At least one is required.
LogDuration or **LogEndDate**	Configures auditing only to extend beyond the expiration of enforcing the Pin Rules. **LogEndDate**, represented as an XML DateTime data type in UTC, takes precedence if both are specified. You represent **LogDuration** as an XML TimeSpan data type, which doesn't allow years and months. If `none of the attributes are specified, auditing expiration uses **Duration** or **NextUpdate** attributes.	No.
ListIdentifier	Provides a friendly name for the list of pin rules. Windows doesn't use this attribute for certificate pinning enforcement; however, it's included when the pin rules are converted to a certificate trust list (CTL).	No.

PinRule Element

The **PinRule** element can have the following attributes.

Attribute	Description	Required
Name	Uniquely identifies the **PinRule**. Windows uses this attribute to identify the element for a parsing error or for verbose output. The attribute isn't included in the generated certificate trust list (CTL).	Yes.
Error	Describes the action Windows performs when it encounters a PIN mismatch. You can choose from the following string values: - **Revoked** - Windows reports the certificate protecting the site as if it was revoked. This typically prevents the user from accessing the site. - **InvalidName** - Windows reports the certificate protecting the site as if the name on the certificate doesn't match the name of the site. This typically results in prompting the user before accessing the site. - **None** - The default value. No error is returned. You can use this setting to audit the pin rules without introducing any user friction.	No.
Log	A Boolean value represents a string that equals **true** or **false**. By default, logging is enabled (**true**).	No.

Certificate element

The **Certificate** element can have the following attributes.

Attribute	Description	Required
File	Path to a file containing one or more certificates. Where the certificate(s) can be encoded as: - single certificate - p7b - sst These files can also be Base64 formatted. All **Site** elements included in the same **PinRule** element can match any of these certificates.	Yes (File, Directory, or Base64 must be present).
Directory	Path to a directory containing one or more of the above certificate files. Skips any files not containing any certificates.	Yes (File, Directory, or Base64 must be present).
Base64	Base64 encoded certificate(s). Where the certificate(s) can be encoded as: - single certificate - p7b - sst This allows the certificates to be included in the XML file without a file directory dependency. Note: You can use **certutil -encode** to convert a .cer file into base64. You can then use Notepad to copy and paste the base64 encoded certificate into the pin rule.	Yes (File, Directory, or Base64 must be present).

Attribute	Description	Required
EndDate	Enables you to configure an expiration date for when the certificate is no longer valid in the pin rule. If you are in the process of switching to a new root or CA, you can set the **EndDate** to allow matching of this element's certificates. If the current time is past the **EndDate**, then, when creating the certificate trust list (CTL), the parser outputs a warning message and excludes the certificate(s) from the Pin Rule in the generated CTL. For help with formatting Pin Rules, see Representing a Date in XML.	No.

Site element

The **Site** element can have the following attributes.

Attribute	Description	Required
Domain	Contains the DNS name to be matched for this pin rule. When creating the certificate trust list, the parser normalizes the input name string value as follows: - If the DNS name has a leading "*", it's removed. - Non-ASCII DNS name is converted to ASCII Puny Code. - Upper case ASCII characters are converted to lower case. If the normalized name has a leading ".", then wildcard left-hand label matching is enabled. For example, ".xyz.com" would match "abc.xyz.com".	Yes.
AllSubdomains	By default, wildcard left-hand label matching is restricted to a single left-hand label. This attribute can be set to "true" to enable wildcard matching of all of the left-hand labels. For example, setting this attribute would also match "123.abc.xyz.com" for the ".xyz.com" domain value.	No.

Create a Pin Rules Certificate Trust List

The command line utility, **Certutil.exe**, includes the **generatePinRulesCTL** argument to parse the XML file and generate the encoded certificate trust list (CTL) that you add to your reference Windows 10 version 1703 computer and subsequently deploy. The usage syntax is:

code

```
CertUtil [Options] -generatePinRulesCTL XMLFile CTLFile [SSTFile]
  Generate Pin Rules CTL
    XMLFile -- input XML file to be parsed.
    CTLFile -- output CTL file to be generated.
    SSTFile -- optional .sst file to be created.
         The .sst file contains all of the certificates
         used for pinning.

Options:
  -f                -- Force overwrite
  -v                -- Verbose operation
```

The same certificate(s) can occur in multiple **PinRule** elements. The same domain can occur in multiple **PinRule** elements. Certutil coalesces these in the resultant pin rules certificate trust list.

Certutil.exe doesn't strictly enforce the XML schema definition. It does perform the following to enable other tools to add/consume their own specific elements and attributes:

- Skips elements before and after the **PinRules** element.
- Skips any element not matching **Certificate** or **Site** within the **PinRules** element.
- Skips any attributes not matching the above names for each element type.

Use the **certutil** command with the **generatePinRulesCTL** argument along with your XML file that contains your certificate pinning rules. Lastly, provide the name of an output file that will include your certificate pinning rules in the form of a certificate trust list.

```
code
```

```
certutil -generatePinRulesCTL certPinRules.xml pinrules.stl
```

Applying Certificate Pinning Rules to a Reference Computer

Now that your certificate pinning rules are in the certificate trust list format, you need to apply the settings to a reference computer as a prerequisite to deploying the setting to your enterprise. To simplify the deployment configuration, it's best to apply your certificate pinning rules to a computer that has the Group Policy Management Console (GPMC) included in the Remote Server Administration Tools (RSAT).

Use **certutil.exe** to apply your certificate pinning rules to your reference computer using the **setreg** argument. The **setreg** argument takes a secondary argument that determines the location of where certutil writes the certificate pining rules. This secondary argument is **chain\PinRules**. The last argument you provide is the name of file that contains your certificate pinning rules in certificate trust list format (.stl). You'll pass the name of the file as the last argument; however, you need to prefix the file name with the '@' symbol as shown in the following example. You need to perform this command from an elevated command prompt.

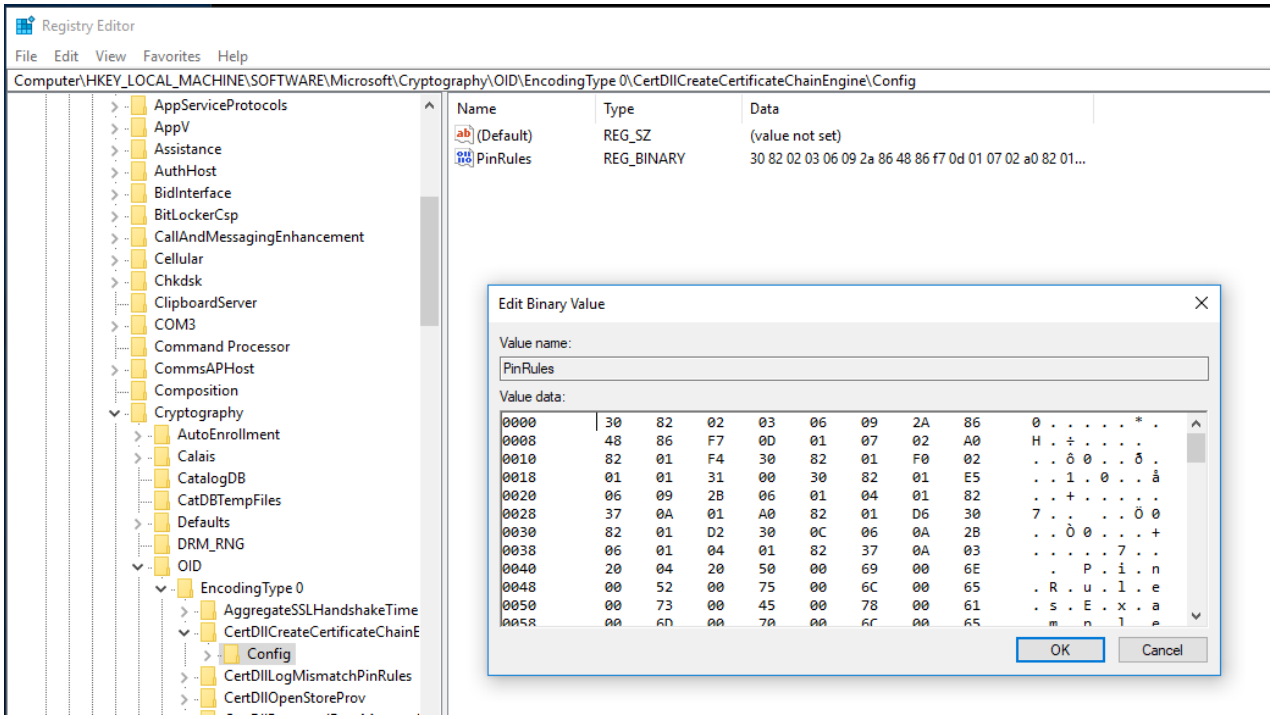
```
code
```

```
Certutil -setreg chain\PinRules @pinrules.stl
```

Certutil writes the binary information to the following registration location:

Name	Value
Key	HKLM\SOFTWARE\Microsoft\Cryptography\OID\EncodingType0\CertDllCreateCertificateChainEngine\Config
Name	PinRules
Value	Binary contents from the certificate pin rules certificate trust list file

Name	Value
Data type	REG_BINARY



Deploying Enterprise Pin Rule Settings using Group Policy

You've successfully created a certificate pinning rules XML file. From the XML file you've created a certificate pinning trust list file, and you've applied the contents of that file to your reference computer from which you can run the Group Policy Management Console. Now you need to configure a Group Policy object to include the applied certificate pin rule settings and deploy it to your environment.

Sign-in to the reference computer using domain administrator equivalent credentials.

1. Start the **Group Policy Management Console** (gpmc.msc)
2. In the navigation pane, expand the forest node and then expand the domain node.
3. Expand the node that contains your Active Directory's domain name
4. Select the **Group Policy objects** node. Right-click the **Group Policy objects** node and click **New**.
5. In the **New GPO** dialog box, type *Enterprise Certificate Pinning Rules* in the **Name** text box and click **OK**.
6. In the content pane, right-click the **Enterprise Certificate Pinning Rules** Group Policy object and click **Edit**.

7. In the **Group Policy Management Editor**, in the navigation pane, expand the **Preferences** node under **Computer Configuration**. Expand **Windows Settings**.

8. Right-click the **Registry** node and click **New**.

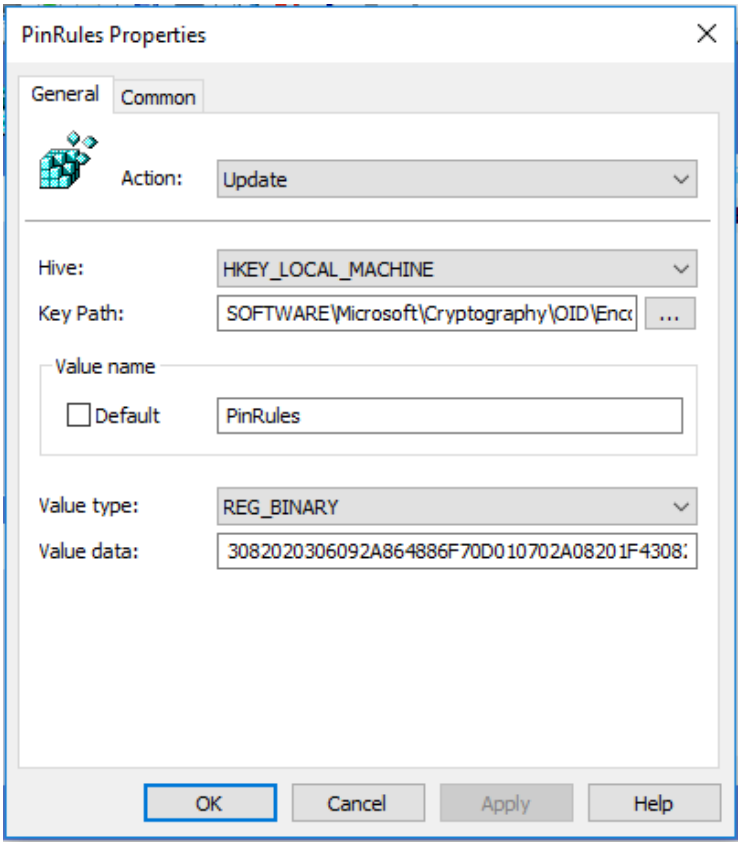
9. In the **New Registry Properties** dialog box, select **Update** from the **Action** list. Select **HKEY_LOCAL_MACHINE** from the **Hive** list.

10. For the **Key Path**, click **...** to launch the **Registry Item Browser**. Navigate to the following registry key and select the **PinRules** registry value name:

    HKLM\SOFTWARE\Microsoft\Cryptography\OID\EncodingType0\CertDllCreateCertificateChainEngine\Config

    Click **Select** to close the **Registry Item Browser**.

11. The **Key Path** should contain the selected registry key. The **Value name** configuration should contain the registry value name ***PinRules***. **Value type** should read ***REG_BINARY*** and **Value data** should contain a long series of numbers from 0-9 and letters ranging from A-F (hexadecimal). Click **OK** to save your settings and close the dialog box.



12. Close the **Group Policy Management Editor** to save your settings.

13. Link the **Enterprise Certificate Pinning Rules** Group Policy object to apply to computers that run Windows 10, version 1703 in your enterprise. When these domain-joined computers apply Group Policy, the registry information configured in the Group Policy object is applied to the computer.

Additional Pin Rules Logging

To assist in constructing certificate pinning rules, you can configure the **PinRulesLogDir** setting under the certificate chain configuration registry key to include a parent directory to log pin rules.

Name	Value
Key	HKLM\SOFTWARE\Microsoft\Cryptography\OID\EncodingType0\CertDllCreateCertificateChainEngine\Config
Name	PinRulesLogDir
Value	The Parent directory where Windows should write the additional pin rule logs
Data type	REG_SZ

Permission for the Pin Rule Log Folder

The folder in which Windows writes the additional pin rule logs must have permissions so that all users and applications have full access. You can run the following commands from an elevated command prompt to achieve the proper permissions.

code

```
set PinRulesLogDir=c:\PinRulesLog
mkdir %PinRulesLogDir%
icacls %PinRulesLogDir% /grant *S-1-15-2-1:(OI)(CI)(F)
icacls %PinRulesLogDir% /grant *S-1-1-0:(OI)(CI)(F)
icacls %PinRulesLogDir% /grant *S-1-5-12:(OI)(CI)(F)
icacls %PinRulesLogDir% /inheritance:e /setintegritylevel (OI)(CI)L
```

Whenever an application verifies a TLS/SSL certificate chain that contains a server name matching a DNS name in the server certificate, Windows writes a .p7b file consisting of all the certificates in the server's chain to one of three child folders:

- AdminPinRules Matched a site in the enterprise certificate pinning rules.
- AutoUpdatePinRules Matched a site in the certificate pinning rules managed by Microsoft.
- NoPinRules Didn't match any site in the certificate pin rules.

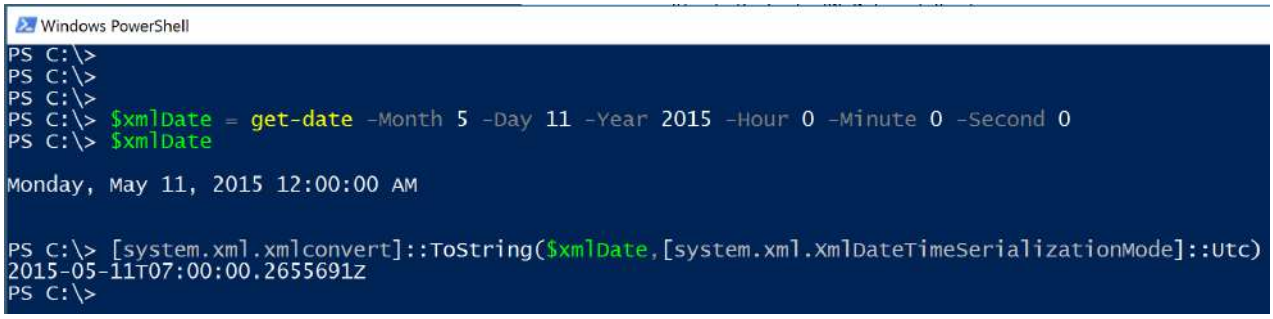
The output file name consists of the leading eight ASCII hex digits of the root's SHA1 thumbprint followed by the server name. For example:

- `D4DE20D0_xsi.outlook.com.p7b`
- `DE28F4A4_www.yammer.com.p7b`

If there's either an enterprise certificate pin rule or a Microsoft certificate pin rule mismatch, then Windows writes the .p7b file to the **MismatchPinRules** child folder. If the pin rules have expired, then Windows writes the .p7b to the **ExpiredPinRules** child folder.

Representing a Date in XML

Many attributes within the pin rules xml file are dates.
These dates must be properly formatted and represented in UTC.
You can use Windows PowerShell to format these dates.
You can then copy and paste the output of the cmdlet into the XML file.



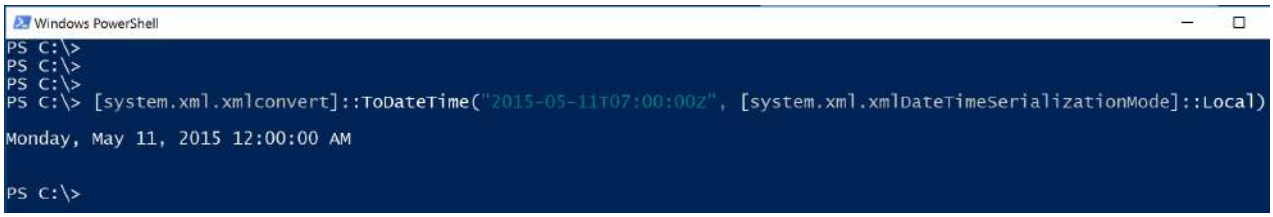
For simplicity, you can truncate decimal point (.) and the numbers after it. However, be certain to append the uppercase “Z” to the end of the XML date string.

code

```
2015-05-11T07:00:00.2655691Z
2015-05-11T07:00:00Z
```

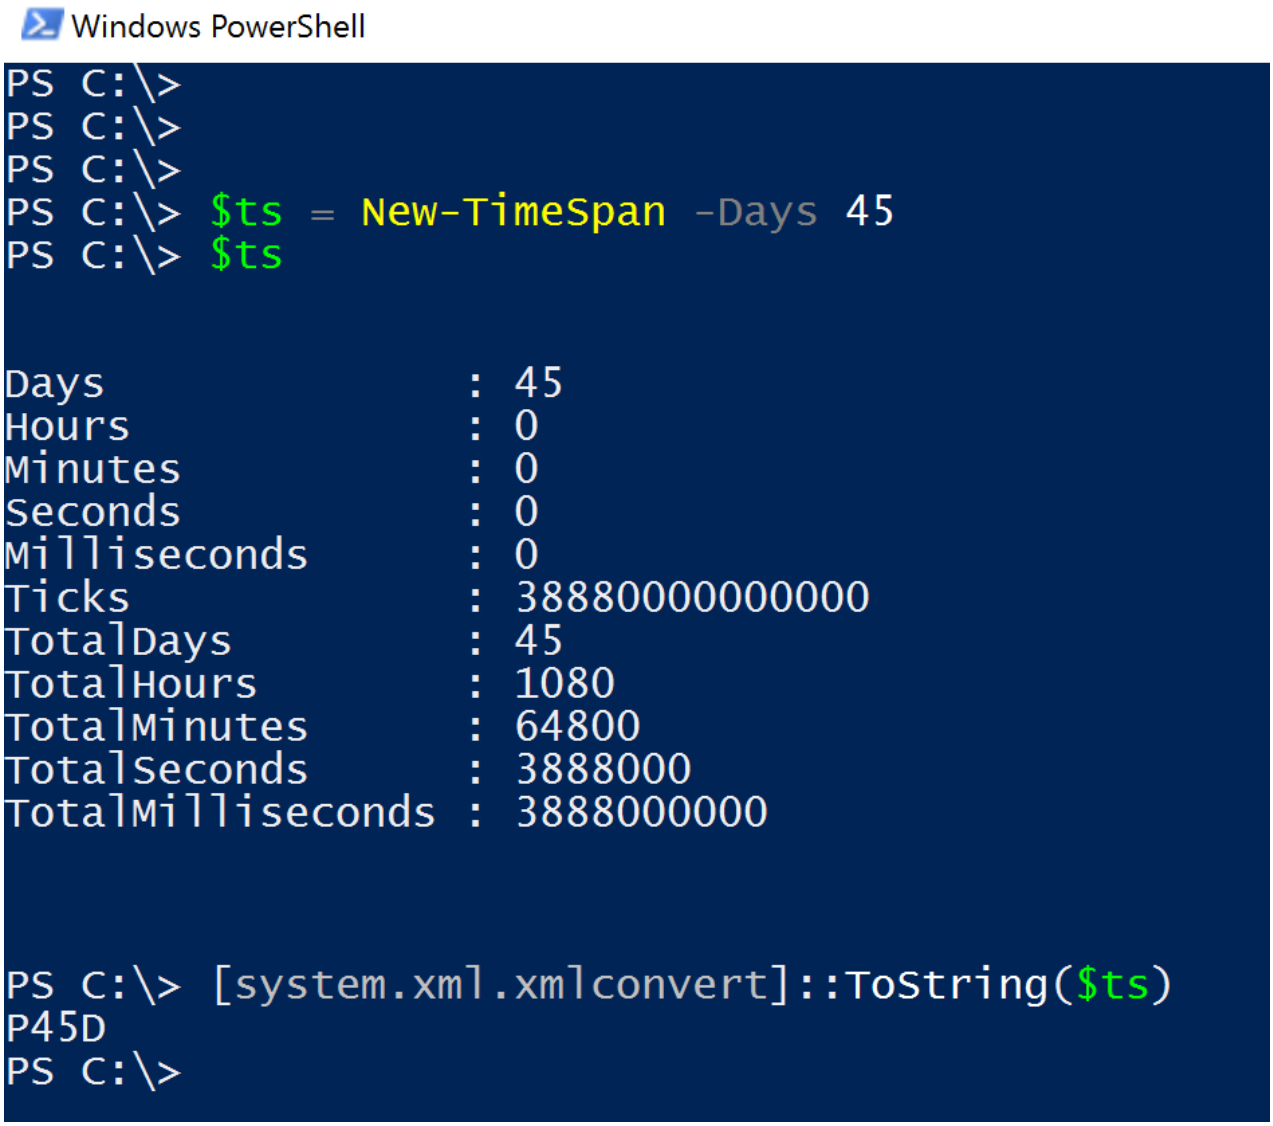
Converting an XML Date

You can also use Windows PowerShell to validate and convert an XML date into a human readable date to validate it’s the correct date.



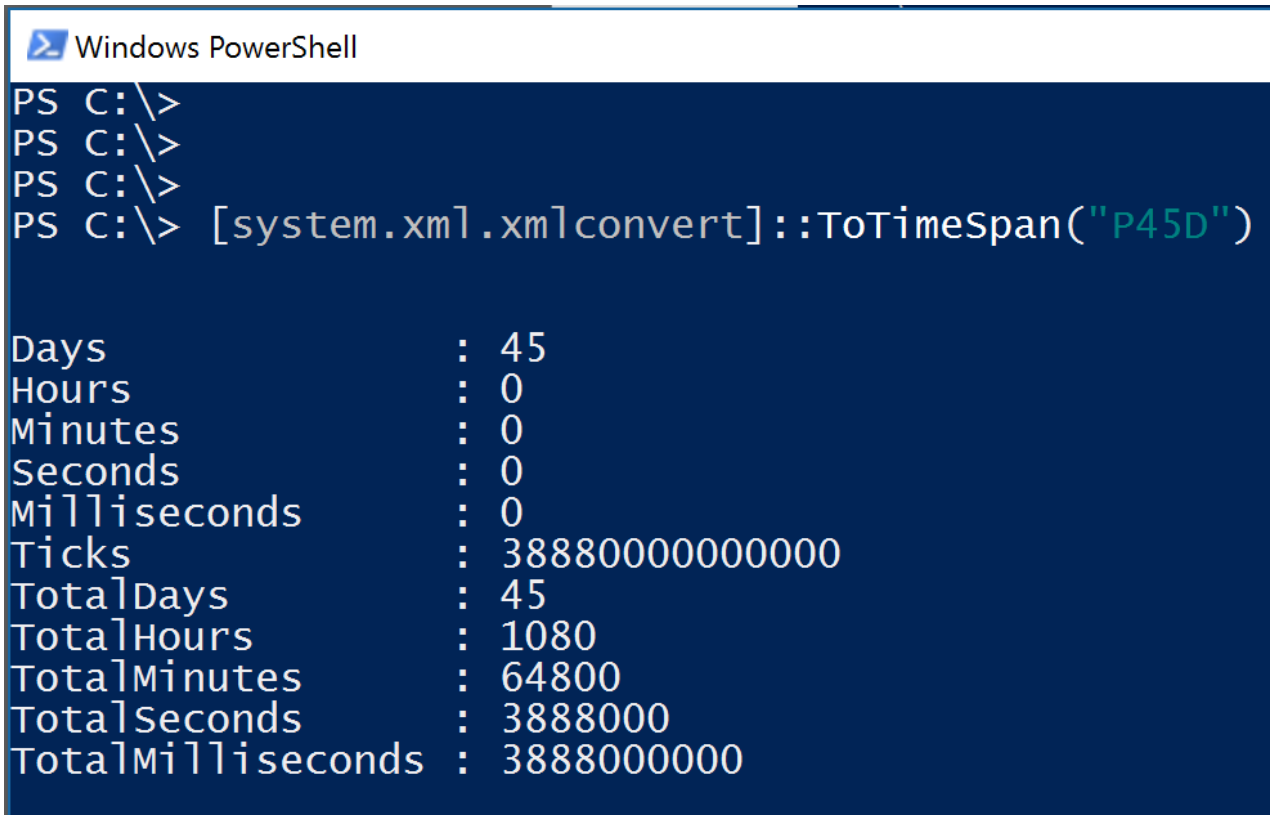
Representing a Duration in XML

Some elements may be configured to use a duration rather than a date. You must represent the duration as an XML timespan data type. You can use Windows PowerShell to properly format and validate durations (timespans) and copy and paste them into your XML file.



Converting an XML Duration

You can convert an XML formatted timespan into a timespan variable that you can read.



Certificate Trust List XML Schema Definition (XSD)

code

```
<xs:schema attributeFormDefault="unqualified" elementFormDefault="qualified"
xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="PinRules">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="PinRule" maxOccurs="unbounded" minOccurs="1">
          <xs:complexType>
            <xs:sequence>
              <xs:element name="Certificate" maxOccurs="unbounded" minOccurs="0">
                <xs:complexType>
                  <xs:simpleContent>
                    <xs:extension base="xs:string">
                      <xs:attribute type="xs:dateTime" name="EndDate"
use="optional"/>
                      <xs:attribute type="xs:string" name="File" use="optional"/>
                      <xs:attribute type="xs:string" name="Directory"
use="optional"/>
                      <xs:attribute type="xs:base64Binary" name="Base64"
use="optional"/>
                    </xs:extension>
                  </xs:simpleContent>
                </xs:complexType>
              </xs:element>
              <xs:element name="Site" maxOccurs="unbounded" minOccurs="1">
                <xs:complexType>
                  <xs:simpleContent>
                    <xs:extension base="xs:string">
                      <xs:attribute type="xs:string" name="Domain"/>
                      <xs:attribute type="xs:boolean" name="AllSubdomains"
```

```
use="optional" default="false"/>
                    </xs:extension>
                  </xs:simpleContent>
                </xs:complexType>
              </xs:element>
            </xs:sequence>
            <xs:attribute type="xs:string" name="Name"/>
            <xs:attribute name="Error" use="optional" default="None">
              <xs:simpleType>
                <xs:restriction base="xs:string">
                  <xs:enumeration value ="Revoked"/>
                  <xs:enumeration value ="InvalidName"/>
                  <xs:enumeration value ="None"/>
                </xs:restriction>
              </xs:simpleType>
            </xs:attribute>
            <xs:attribute type="xs:boolean" name="Log" use="optional"
default="true"/>
          </xs:complexType>
        </xs:element>
      </xs:sequence>
      <xs:attribute type="xs:duration" name="Duration" use="optional"/>
      <xs:attribute type="xs:duration" name="LogDuration" use="optional"/>
      <xs:attribute type="xs:dateTime" name="NextUpdate" use="optional"/>
      <xs:attribute type="xs:dateTime" name="LogEndDate" use="optional"/>
      <xs:attribute type="xs:string" name="ListIdentifier" use="optional"/>
    </xs:complexType>
  </xs:element>
</xs:schema>
```

Protect derived domain credentials with Windows Defender Credential Guard

Article • 02/17/2023 • 2 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

Windows Defender Credential Guard uses virtualization-based security to isolate secrets so that only privileged system software can access them. Unauthorized access to these secrets can lead to credential theft attacks, such as Pass-the-Hash or Pass-The-Ticket. Windows Defender Credential Guard prevents these attacks by protecting NTLM password hashes, Kerberos Ticket Granting Tickets, and credentials stored by applications as domain credentials.

By enabling Windows Defender Credential Guard, the following features and solutions are provided:

- **Hardware security** NTLM, Kerberos, and Credential Manager take advantage of platform security features, including Secure Boot and virtualization, to protect credentials.
- **Virtualization-based security** Windows NTLM and Kerberos derived credentials and other secrets run in a protected environment that is isolated from the running operating system.
- **Better protection against advanced persistent threats** When Credential Manager domain credentials, NTLM, and Kerberos derived credentials are protected using virtualization-based security, the credential theft attack techniques and tools used in many targeted attacks are blocked. Malware running in the operating system with administrative privileges cannot extract secrets that are protected by virtualization-based security. While Windows Defender Credential Guard is a powerful mitigation, persistent threat attacks will likely shift to new attack techniques and you should also incorporate other security strategies and architectures.

> ⓘ **Note**
>
> As of Windows 11, version 22H2, Windows Defender Credential Guard has been enabled by default on all devices which meet the minimum requirements as specified in the **Default Enablement** section. For information about known issues related to default enablement, see **Credential Guard: Known Issues**.

Related topics

- Protecting network passwords with Windows Defender Credential Guard
- Enabling Strict KDC Validation in Windows Kerberos
- What's New in Kerberos Authentication for Windows Server 2012
- Authentication Mechanism Assurance for AD DS in Windows Server 2008 R2 Step-by-Step Guide
- Trusted Platform Module
- Mitigating Credential Theft using the Windows 10 Isolated User Mode
- Isolated User Mode Processes and Features in Windows 10 with Logan Gabriel
- More on Processes and Features in Windows 10 Isolated User Mode with Dave Probert
- Isolated User Mode in Windows 10 with Dave Probert
- Windows 10 Virtual Secure Mode with David Hepkin

How Windows Defender Credential Guard works

Article • 12/15/2022 • 2 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

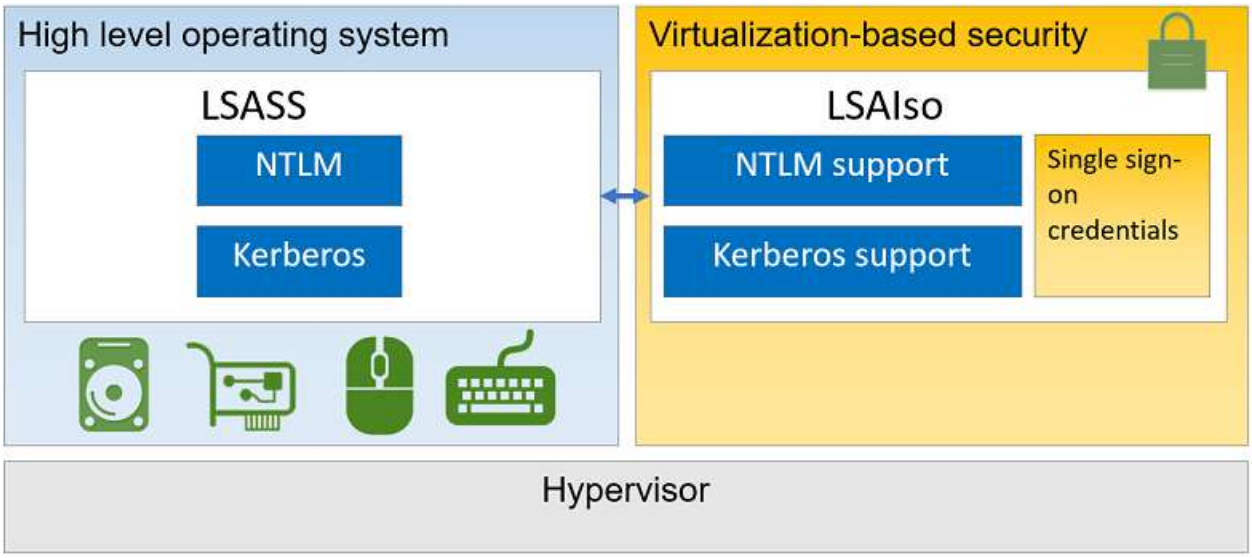
Kerberos, NTLM, and Credential manager isolate secrets by using virtualization-based security. Previous versions of Windows stored secrets in the Local Security Authority (LSA). Prior to Windows 10, the LSA stored secrets used by the operating system in its process memory. With Windows Defender Credential Guard enabled, the LSA process in the operating system talks to a new component called the isolated LSA process that stores and protects those secrets. Data stored by the isolated LSA process is protected using Virtualization-based security and isn't accessible to the rest of the operating system. LSA uses remote procedure calls to communicate with the isolated LSA process.

For security reasons, the isolated LSA process doesn't host any device drivers. Instead, it only hosts a small subset of operating system binaries that are needed for security and nothing else. All of these binaries are signed with a certificate that is trusted by virtualization-based security and these signatures are validated before launching the file in the protected environment.

When Windows Defender Credential Guard is enabled, NTLMv1, MS-CHAPv2, Digest, and CredSSP can't use the signed-in credentials. Thus, single sign-on doesn't work with these protocols. However, applications can prompt for credentials or use credentials stored in the Windows Vault, which aren't protected by Windows Defender Credential Guard with any of these protocols. It is recommended that valuable credentials, such as the sign-in credentials, aren't to be used with any of these protocols. If these protocols must be used by domain or Azure AD users, secondary credentials should be provisioned for these use cases.

When Windows Defender Credential Guard is enabled, Kerberos doesn't allow unconstrained Kerberos delegation or DES encryption, not only for signed-in credentials, but also prompted or saved credentials.

Here's a high-level overview on how the LSA is isolated by using Virtualization-based security:



See also

Related videos

What is Virtualization-based security?

Windows Defender Credential Guard requirements

Article • 12/15/2022 • 8 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

For Windows Defender Credential Guard to provide protection, the computers you are protecting must meet certain baseline hardware, firmware, and software requirements, which we will refer to as Hardware and software requirements. Additionally, Windows Defender Credential Guard blocks specific authentication capabilities, so applications that require such capabilities will break. We will refer to these requirements as Application requirements. Beyond these requirements, computers can meet additional hardware and firmware qualifications, and receive additional protections. Those computers will be more hardened against certain threats. For detailed information on baseline protections, plus protections for improved security that are associated with hardware and firmware options available in 2015, 2016, and 2017, refer to the tables in Security Considerations.

Hardware and software requirements

To provide basic protections against OS level attempts to read Credential Manager domain credentials, NTLM and Kerberos derived credentials, Windows Defender Credential Guard uses:

- Support for Virtualization-based security (required)
- Secure boot (required)
- Trusted Platform Module (TPM, preferred - provides binding to hardware) versions 1.2 and 2.0 are supported, either discrete or firmware
- UEFI lock (preferred - prevents attacker from disabling with a simple registry key change)

The Virtualization-based security requires:

- 64-bit CPU
- CPU virtualization extensions plus extended page tables
- Windows hypervisor (does not require Hyper-V Windows Feature to be installed)

Windows Defender Credential Guard deployment in virtual machines

Credential Guard can protect secrets in a Hyper-V virtual machine, just as it would on a physical machine. When Credential Guard is deployed on a VM, secrets are protected from

attacks inside the VM. Credential Guard does not provide additional protection from privileged system attacks originating from the host.

Requirements for running Windows Defender Credential Guard in Hyper-V virtual machines

- The Hyper-V host must have an IOMMU, and run at least Windows Server 2016 or Windows 10 version 1607.
- The Hyper-V virtual machine must be Generation 2, have an enabled virtual TPM, and be running at least Windows Server 2016 or Windows 10.
  - TPM is not a requirement, but we recommend that you implement TPM.

For information about other host platforms, see Enabling Windows Server 2016 and Hyper-V virtualization based security features on other platforms.

For information about Windows Defender Remote Credential Guard hardware and software requirements, see Windows Defender Remote Credential Guard requirements.

Application requirements

When Windows Defender Credential Guard is enabled, specific authentication capabilities are blocked, so applications that require such capabilities will break. Applications should be tested prior to deployment to ensure compatibility with the reduced functionality.

> ⚠ **Warning**
>
> Enabling Windows Defender Credential Guard on domain controllers is not recommended at this time. Windows Defender Credential Guard does not provide any added security to domain controllers, and can cause application compatibility issues on domain controllers.

> ⓘ **Note**
>
> Windows Defender Credential Guard does not provide protections for the Active Directory database or the Security Accounts Manager (SAM). The credentials protected by Kerberos and NTLM when Windows Defender Credential Guard is enabled are also in the Active Directory database (on domain controllers) and the SAM (for local accounts).

Applications will break if they require:

- Kerberos DES encryption support

- Kerberos unconstrained delegation
- Extracting the Kerberos TGT
- NTLMv1

Applications will prompt and expose credentials to risk if they require:

- Digest authentication
- Credential delegation
- MS-CHAPv2

Applications may cause performance issues when they attempt to hook the isolated Windows Defender Credential Guard process.

Services or protocols that rely on Kerberos, such as file shares, remote desktop, or BranchCache, continue to work and are not affected by Windows Defender Credential Guard.

Security considerations

All computers that meet baseline protections for hardware, firmware, and software can use Windows Defender Credential Guard. Computers that meet additional qualifications can provide additional protections to further reduce the attack surface. The following tables describe baseline protections, plus protections for improved security that are associated with hardware and firmware options available in 2015, 2016, and 2017.

> ⓘ **Note**
>
> Beginning with Windows 10, version 1607, Trusted Platform Module (TPM 2.0) must be enabled by default on new shipping computers.
>
> If you are an OEM, see **PC OEM requirements for Windows Defender Credential Guard**.

Baseline protections

Baseline Protections	Description	Security benefits
Hardware: **64-bit CPU**	A 64-bit computer is required for the Windows hypervisor to provide VBS.	

Baseline Protections	Description	Security benefits
Hardware: **CPU virtualization extensions**, plus **extended page tables**	**Requirements**: - These hardware features are required for VBS: One of the following virtualization extensions: - VT-x (Intel) or - AMD-V And: - Extended page tables, also called Second Level Address Translation (SLAT).	VBS provides isolation of secure kernel from normal operating system. Vulnerabilities and Day 0s in normal operating system cannot be exploited because of this isolation.
Hardware: **Trusted Platform Module (TPM)**	**Requirement**: - TPM 1.2 or TPM 2.0, either discrete or firmware. TPM recommendations	A TPM provides protection for VBS encryption keys that are stored in the firmware. TPM helps protect against attacks involving a physically present user with BIOS access.
Firmware: **UEFI firmware version 2.3.1.c or higher with UEFI Secure Boot**	**Requirements**: - See the following Windows Hardware Compatibility Program requirement: System.Fundamentals.Firmware.UEFISecureBoot	UEFI Secure Boot helps ensure that the device boots only authorized code, and can prevent boot kits and root kits from installing and persisting across reboots.
Firmware: **Secure firmware update process**	**Requirements**: - UEFI firmware must support secure firmware update found under the following Windows Hardware Compatibility Program requirement: System.Fundamentals.Firmware.UEFISecureBoot.	UEFI firmware just like software can have security vulnerabilities that, when found, need to be patched through firmware updates. Patching helps prevent root kits from getting installed.
Software: Qualified **Windows operating system**	**Requirement**: - At least Windows 10 Enterprise, Windows 10 Education, or Windows Server 2016.	Support for VBS and for management features that simplify configuration of Windows Defender Credential Guard.

> **ⓘ Important**
>
> The following tables list additional qualifications for improved security. We strongly recommend meeting the additional qualifications to significantly strengthen the level of security that Windows Defender Credential Guard can provide.

2015 Additional security qualifications starting with Windows 10, version 1507, and Windows Server 2016 Technical Preview 4

Protections for Improved Security	Description
Hardware: **IOMMU** (input/output memory management unit)	**Requirement**: - VT-D or AMD Vi IOMMU **Security benefits**: - An IOMMU can enhance system resiliency against memory attacks. For more information, see Advanced Configuration and Power Interface (ACPI) description tables
Firmware: **Securing Boot Configuration and Management**	**Requirements**: - BIOS password or stronger authentication must be supported. - In the BIOS configuration, BIOS authentication must be set. - There must be support for protected BIOS option to configure list of permitted boot devices (for example, "Boot only from internal hard drive") and boot device order, overriding BOOTORDER modification made by operating system. - In the BIOS configuration, BIOS options related to security and boot options (list of permitted boot devices, boot order) must be secured to prevent other operating systems from starting and to prevent changes to the BIOS settings.
Firmware: **Secure MOR, revision 2 implementation**	**Requirement**: - Secure MOR, revision 2 implementation

2016 Additional security qualifications starting with Windows 10, version 1607, and Windows Server 2016

> **Important**
>
> The following tables list additional qualifications for improved security. Systems that meet these additional qualifications can provide more protections.

Protections for Improved Security	Description	Security Benefits

Protections for Improved Security	Description	Security Benefits
Firmware: **Hardware Rooted Trust Platform Secure Boot**	**Requirements**: - Boot Integrity (Platform Secure Boot) must be supported. See the Windows Hardware Compatibility Program requirements under System.Fundamentals.Firmware.CS.UEFISecureBoot.ConnectedStandby - The Hardware Security Test Interface (HSTI) must be implemented. See Hardware Security Testability Specification.	Boot Integrity (Platform Secure Boot) from Power-On provides protections against physically present attackers, and defense-in-depth against malware. - HSTI provides additional security assurance for correctly secured silicon and platform.
Firmware: **Firmware Update through Windows Update**	**Requirements**: - Firmware must support field updates through Windows Update and UEFI encapsulation update.	Helps ensure that firmware updates are fast, secure, and reliable.
Firmware: **Securing Boot Configuration and Management**	**Requirements**: - Required BIOS capabilities: Ability of OEM to add ISV, OEM, or Enterprise Certificate in Secure Boot DB at manufacturing time. - Required configurations: Microsoft UEFI CA must be removed from Secure Boot DB. Support for 3rd-party UEFI modules is permitted but should leverage ISV-provided certificates or OEM certificate for the specific UEFI software.	- Enterprises can choose to allow proprietary EFI drivers/applications to run. - Removing Microsoft UEFI CA from Secure Boot DB provides full control to enterprises over software that runs before the operating system boots.

2017 Additional security qualifications starting with Windows 10, version 1703

The following table lists qualifications for Windows 10, version 1703, which are in addition to all preceding qualifications.

Protections for Improved Security	Description	Security Benefits
Firmware: **VBS enablement of No-Execute (NX) protection for UEFI runtime services**	**Requirements**: - VBS will enable NX protection on UEFI runtime service code and data memory regions. UEFI runtime service code must support read-only page protections, and UEFI runtime service data must not be executable. UEFI runtime service must meet these requirements: - Implement UEFI 2.6 EFI_MEMORY_ATTRIBUTES_TABLE. All UEFI runtime service memory (code and data) must be described by this table. - PE sections must be page-aligned in memory (not required for in non-volatile storage). - The Memory Attributes Table needs to correctly mark code and data as RO/NX for configuration by the OS: - All entries must include attributes EFI_MEMORY_RO, EFI_MEMORY_XP, or both. - No entries may be left with neither of the above attributes, indicating memory that is both executable and writable. Memory must be either readable and executable or writable and non-executable. (**SEE IMPORTANT INFORMATION AFTER THIS TABLE**)	Vulnerabilities in UEFI runtime, if any, will be blocked from compromising VBS (such as in functions like UpdateCapsule and SetVariable) - Reduces the attack surface to VBS from system firmware.
Firmware: **Firmware support for SMM protection**	**Requirements**: - The Windows SMM Security Mitigations Table (WSMT) specification contains details of an ACPI table that was created for use with Windows operating systems that support Windows virtualization-based security (VBS) features.	- Protects against potential vulnerabilities in UEFI runtime services, if any, will be blocked from compromising VBS (such as in functions like UpdateCapsule and SetVariable) - Reduces the attack surface to VBS from system firmware. - Blocks additional security attacks against SMM.

ⓘ Important

Regarding **VBS enablement of NX protection for UEFI runtime services**:

- This only applies to UEFI runtime service memory, and not UEFI boot service memory.

- This protection is applied by VBS on OS page tables.

  Please also note the following:

  - Do not use sections that are both writable and executable
  - Do not attempt to directly modify executable system memory
  - Do not use dynamic code

Manage Windows Defender Credential Guard

Article • 02/22/2023 • 11 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

Default Enablement

Starting in **Windows 11 Enterprise, version 22H2** and **Windows 11 Education, version 22H2**, compatible systems have Windows Defender Credential Guard turned on by default. This feature changes the default state of the feature in Windows, though system administrators can still modify this enablement state. Windows Defender Credential Guard can still be manually enabled or disabled via the methods documented below.

Known issues arising from default enablement are documented in Windows Defender Credential Guard: Known issues.

Requirements for automatic enablement

Windows Defender Credential Guard will be enabled by default when a PC meets the following minimum requirements:

Component	Requirement
Operating System	**Windows 11 Enterprise, version 22H2** or **Windows 11 Education, version 22H2**
Existing Windows Defender Credential Guard Requirements	Only devices that meet the existing hardware and software requirements to run Windows Defender Credential Guard will have it enabled by default.
Virtualization-based Security (VBS) Requirements	VBS must be enabled in order to run Windows Defender Credential Guard. Starting with Windows 11 Enterprise 22H2 and Windows 11 Education 22H2, devices that meet the requirements to run Windows Defender Credential Guard as well as the minimum requirements to enable VBS will have both Windows Defender Credential Guard and VBS enabled by default.

> ⓘ Note

If Windows Defender Credential Guard or VBS has previously been explicitly disabled, default enablement will not overwrite this setting.

> **Note**
>
> Devices running Windows 11 Pro 22H2 may have Virtualization-Based Security (VBS) and/or Windows Defender Credential Guard automaticaly enabled if they meet the other requirements for default enablement listed above and have previously run Windows Defender Credential Guard (for example if Windows Defender Credential Guard was running on an Enterprise device that later downgraded to Pro).
>
> To determine whether the Pro device is in this state, check if the registry key `IsolatedCredentialsRootSecret` is present in `Computer\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\MSV1_0`. In this scenario, if you wish to disable VBS and Windows Defender Credential Guard, follow the instructions for **disabling Virtualization-Based Security**. If you wish to disable only Windows Defender Credential Guard without disabling Virtualization-Based Security, use the procedures for **disabling Windows Defender Credential Guard**.

Enable Windows Defender Credential Guard

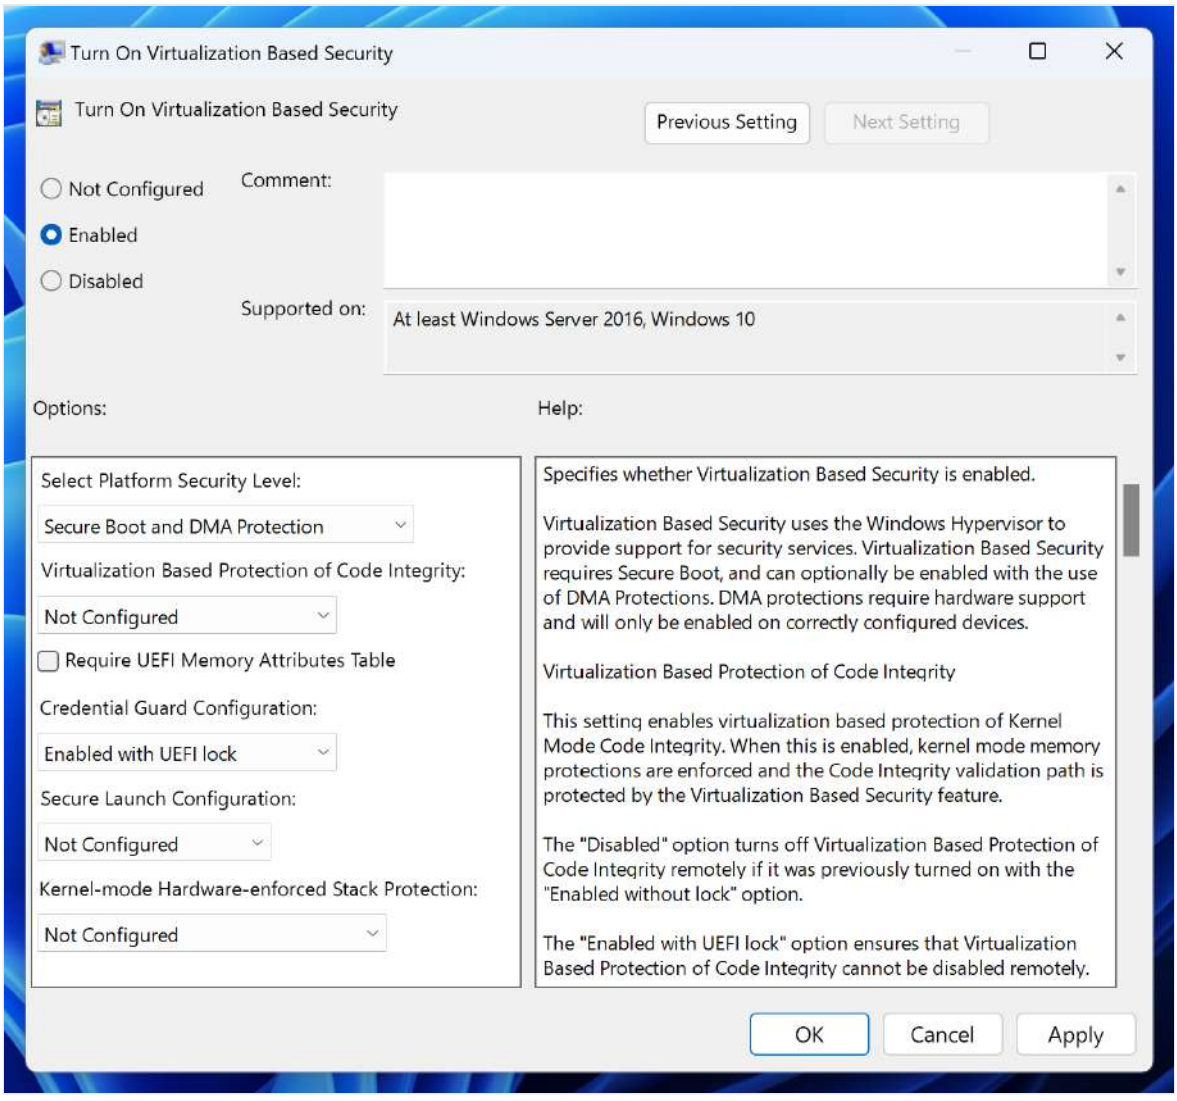
Windows Defender Credential Guard can be enabled either by using Group Policy or the registry. Windows Defender Credential Guard can also protect secrets in a Hyper-V virtual machine, just as it would on a physical machine. The same set of procedures used to enable Windows Defender Credential Guard on physical machines applies also to virtual machines.

> **Note**
>
> Credential Guard and Device Guard are not supported when using Azure Gen 1 VMs. These options are available with Gen 2 VMs only.

Enable Windows Defender Credential Guard by using Group Policy

You can use Group Policy to enable Windows Defender Credential Guard. When enabled, it will add and enable the virtualization-based security features for you if needed.

1. From the Group Policy Management Console, go to **Computer Configuration** > **Administrative Templates** > **System** > **Device Guard**.
2. Select **Turn On Virtualization Based Security**, and then select the **Enabled** option.
3. In the **Select Platform Security Level** box, choose **Secure Boot** or **Secure Boot and DMA Protection**.
4. In the **Credential Guard Configuration** box, select **Enabled with UEFI lock**. If you want to be able to turn off Windows Defender Credential Guard remotely, choose **Enabled without lock**.
5. In the **Secure Launch Configuration** box, choose **Not Configured**, **Enabled** or **Disabled**. For more information, see System Guard Secure Launch and SMM protection.



6. Select **OK**, and then close the Group Policy Management Console.

To enforce processing of the group policy, you can run `gpupdate /force`.

Enable Windows Defender Credential Guard by using Microsoft Intune

1. In the Intune admin center, select **Devices**.
2. Select **Configuration Profiles**.
3. Select **Create Profile** > **Windows 10 and later** > **Settings catalog** > **Create**.
   a. Configuration settings: In the settings picker, select **Device Guard** as category and add the needed settings.

> **Note**
>
> Enable VBS and Secure Boot and you can do it with or without UEFI Lock. If you will need to disable Credential Guard remotely, enable it without UEFI lock.

> **Tip**
>
> You can also configure Credential Guard by using an account protection profile in endpoint security. For more information, see **Account protection policy settings for endpoint security in Microsoft Intune**.

Enable Windows Defender Credential Guard by using the registry

If you don't use Group Policy, you can enable Windows Defender Credential Guard by using the registry. Windows Defender Credential Guard uses virtualization-based security features that have to be enabled first on some operating systems.

Add the virtualization-based security features

Starting with Windows 10, version 1607 and Windows Server 2016, enabling Windows features to use virtualization-based security isn't necessary and this step can be skipped.

If you're using Windows 10, version 1507 (RTM) or Windows 10, version 1511, Windows features have to be enabled to use virtualization-based security. To enable, use the Control Panel or the Deployment Image Servicing and Management tool (DISM).

> **Note**
>
> If you enable Windows Defender Credential Guard by using Group Policy, the steps to enable Windows features through Control Panel or DISM are not required. Group Policy will install Windows features for you.

Add the virtualization-based security features by using Programs and Features

1. Open the Programs and Features control panel.
2. Select **Turn Windows feature on or off**.
3. Go to **Hyper-V** > **Hyper-V Platform**, and then select the **Hyper-V Hypervisor** check box.
4. Select the **Isolated User Mode** check box at the top level of the feature selection.
5. Select **OK**.

Add the virtualization-based security features to an offline image by using DISM

1. Open an elevated command prompt.
2. Add the Hyper-V Hypervisor by running the following command:

```
Windows Command Prompt

dism /image:<WIM file name> /Enable-Feature /FeatureName:Microsoft-
Hyper-V-Hypervisor /all
```

3. Add the Isolated User Mode feature by running the following command:

```
Windows Command Prompt

dism /image:<WIM file name> /Enable-Feature
/FeatureName:IsolatedUserMode
```

> **Note**
>
> In Windows 10, version 1607 and later, the Isolated User Mode feature has been integrated into the core operating system. Running the command in

step 3 above is therefore no longer required.

> **Tip**
>
> You can also add these features to an online image by using either DISM or Configuration Manager.

Enable virtualization-based security and Windows Defender Credential Guard

1. Open Registry Editor.
2. Enable virtualization-based security:
   a. Go to `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\DeviceGuard`.
   b. Add a new DWORD value named **EnableVirtualizationBasedSecurity**. Set the value of this registry setting to 1 to enable virtualization-based security and set it to 0 to disable it.
   c. Add a new DWORD value named **RequirePlatformSecurityFeatures**. Set the value of this registry setting to 1 to use **Secure Boot** only or set it to 3 to use **Secure Boot and DMA protection**.
3. Enable Windows Defender Credential Guard:
   a. Go to `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa`.
   b. Add a new DWORD value named **LsaCfgFlags**. Set the value of this registry setting to 1 to enable Windows Defender Credential Guard with UEFI lock, set it to 2 to enable Windows Defender Credential Guard without lock, and set it to 0 to disable it.
4. Close Registry Editor.

> **Note**
>
> You can also enable Windows Defender Credential Guard by setting the registry entries in the **FirstLogonCommands** unattend setting.

Review Windows Defender Credential Guard performance

Is Windows Defender Credential Guard running?

You can view System Information to check that Windows Defender Credential Guard is running on a PC.

1. Select **Start**, type **msinfo32.exe**, and then select **System Information**.
2. Select **System Summary**.
3. Confirm that **Credential Guard** is shown next to **Virtualization-based security Services Running**.



Virtualization-based security	Running
Virtualization-based security Required Security Properties	Base Virtualization Support, Secure Boot, DMA Protection
Virtualization-based security Available Security Properties	Base Virtualization Support, Secure Boot, DMA Protection, Mode Based ...
Virtualization-based security Services Configured	Credential Guard, Hypervisor enforced Code Integrity
Virtualization-based security Services Running	Credential Guard, Hypervisor enforced Code Integrity

> ⓘ **Note**
>
> For client machines that are running Windows 10 1703, LsaIso.exe is running whenever virtualization-based security is enabled for other features.

- We recommend enabling Windows Defender Credential Guard before a device is joined to a domain. If Windows Defender Credential Guard is enabled after domain join, the user and device secrets may already be compromised. In other words, enabling Credential Guard won't help to secure a device or identity that has already been compromised. So, we recommend turning on Credential Guard as early as possible.
- You should perform regular reviews of the PCs that have Windows Defender Credential Guard enabled. You can use security audit policies or WMI queries. Here's a list of WinInit event IDs to look for:
  - **Event ID 13** Windows Defender Credential Guard (LsaIso.exe) was started and will protect LSA credentials.
  - **Event ID 14** Windows Defender Credential Guard (LsaIso.exe) configuration: [**0x0** | **0x1** | **0x2**], **0**
    - The first variable: **0x1** or **0x2** means that Windows Defender Credential Guard is configured to run. **0x0** means that it's not configured to run.

    - The second variable: **0** means that it's configured to run in protect mode. **1** means that it's configured to run in test mode. This variable should always be **0**.

  - **Event ID 15** Windows Defender Credential Guard (LsaIso.exe) is configured but the secure kernel isn't running; continuing without Windows Defender Credential Guard.

  - **Event ID 16** Windows Defender Credential Guard (LsaIso.exe) failed to launch: [error code]

  - **Event ID 17** Error reading Windows Defender Credential Guard (LsaIso.exe) UEFI configuration: [error code]

- You can also verify that TPM is being used for key protection by checking **Event ID 51** in *Applications and Services logs* > *Microsoft* > *Windows* > *Kernel-Boot* event log. The full event text will read like this: `VSM Master Encryption Key Provisioning. Using cached copy status: 0x0. Unsealing cached copy status: 0x1. New key generation status: 0x1. Sealing status: 0x1. TPM PCR mask: 0x0.` If you're running with a TPM, the TPM PCR mask value will be something other than 0.

- You can use Windows PowerShell to determine whether credential guard is running on a client computer. On the computer in question, open an elevated PowerShell window and run the following command:

  PowerShell

  ```
  (Get-CimInstance -ClassName Win32_DeviceGuard -Namespace root\Microsoft\Windows\DeviceGuard).SecurityServicesRunning
  ```

  This command generates the following output:

  - **0**: Windows Defender Credential Guard is disabled (not running)

  - **1**: Windows Defender Credential Guard is enabled (running)

    > ⓘ **Note**
    >
    > Checking the task list or Task Manager to see if LSAISO.exe is running is not a recommended method for determining whether Windows Defender Credential Guard is running.

Disable Windows Defender Credential Guard

Windows Defender Credential Guard can be disabled via several methods explained below, depending on how the feature was enabled. For devices that had Windows Defender Credential Guard automatically enabled in the 22H2 update and didn't have it enabled prior to the update, it's sufficient to disable via Group Policy.

If Windows Defender Credential Guard was enabled with UEFI Lock, the procedure described in Disabling Windows Defender Credential Guard with UEFI Lock must be followed. The default enablement change in eligible 22H2 devices does **not** use a UEFI Lock.

If Windows Defender Credential Guard was enabled via Group Policy without UEFI Lock, Windows Defender Credential Guard should be disabled via Group Policy.

Otherwise, Windows Defender Credential Guard can be disabled by changing registry keys.

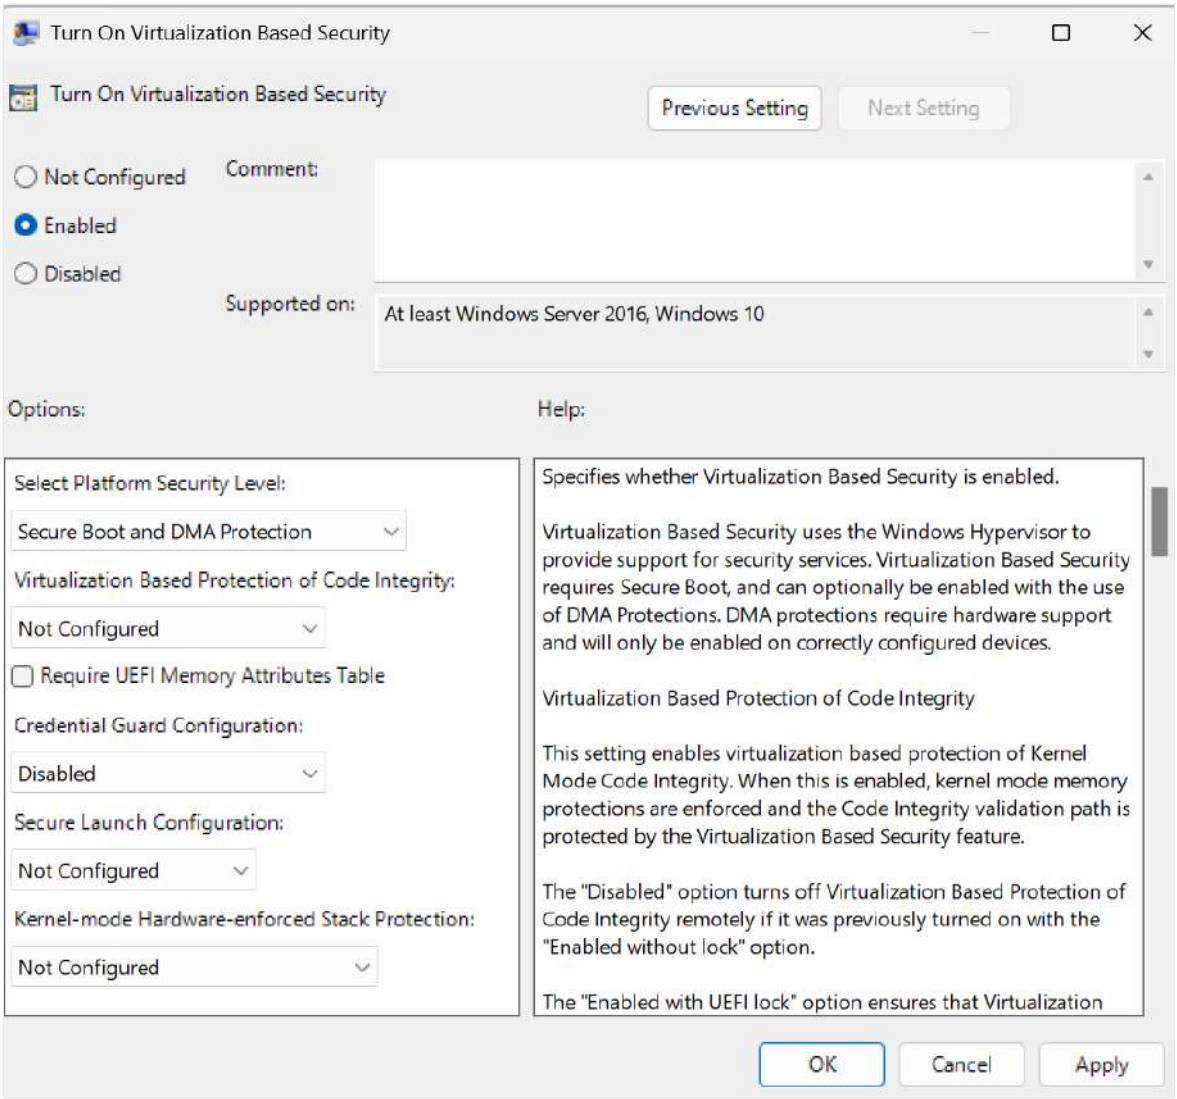
Windows Defender Credential Guard running in a virtual machine can be disabled by the host.

For information on disabling Virtualization-Based Security (VBS), see Disabling Virtualization-Based Security.

Disabling Windows Defender Credential Guard using Group Policy

If Windows Defender Credential Guard was enabled via Group Policy and without UEFI Lock, disabling the same Group Policy setting will disable Windows Defender Credential Guard.

1. Disable the Group Policy setting that governs Windows Defender Credential Guard. Navigate to **Computer Configuration** > **Administrative Templates** > **System** > **Device Guard** > **Turn on Virtualization Based Security**. In the "Credential Guard Configuration" section, set the dropdown value to "Disabled":



2. Restart the machine.

Disabling Windows Defender Credential Guard using Registry Keys

If Windows Defender Credential Guard was enabled without UEFI Lock and without Group Policy, it's sufficient to edit the registry keys as described below to disable Windows Defender Credential Guard.

1. Change the following registry settings to 0:

   - `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\LsaCfgFlags`
   - `HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\DeviceGuard\LsaCfgFlags`

   > ⓘ Note

> Deleting these registry settings may not disable Windows Defender Credential Guard. They must be set to a value of 0.

2. Restart the machine.

Disabling Windows Defender Credential Guard with UEFI Lock

If Windows Defender Credential Guard was enabled with UEFI Lock enabled, then the following procedure must be followed since the settings are persisted in EFI (firmware) variables. This scenario will require physical presence at the machine to press a function key to accept the change.

1. If Group Policy was used to enable Windows Defender Credential Guard, disable the relevant Group Policy setting. Navigate to **Computer Configuration** > **Administrative Templates** > **System** > **Device Guard** > **Turn on Virtualization Based Security**. In the "Credential Guard Configuration" section, set the dropdown value to "Disabled".

2. Change the following registry settings to 0:

   - `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\LsaCfgFlags`
   - `HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\DeviceGuard\LsaCfgFlags`

3. Delete the Windows Defender Credential Guard EFI variables by using bcdedit. From an elevated command prompt, type the following commands:

   Windows Command Prompt

   ```
   mountvol X: /s
   copy %WINDIR%\System32\SecConfig.efi
   X:\EFI\Microsoft\Boot\SecConfig.efi /Y
   bcdedit /create {0cb3b571-2f2e-4343-a879-d86a476d7215} /d "DebugTool"
   /application osloader
   bcdedit /set {0cb3b571-2f2e-4343-a879-d86a476d7215} path
   "\EFI\Microsoft\Boot\SecConfig.efi"
   bcdedit /set {bootmgr} bootsequence {0cb3b571-2f2e-4343-a879-
   d86a476d7215}
   bcdedit /set {0cb3b571-2f2e-4343-a879-d86a476d7215} loadoptions
   DISABLE-LSA-ISO
   bcdedit /set {0cb3b571-2f2e-4343-a879-d86a476d7215} device partition=X:
   mountvol X: /d
   ```

4. Restart the PC. Before the OS boots, a prompt will appear notifying that UEFI was modified, and asking for confirmation. This prompt must be confirmed for the changes to persist. This step requires physical access to the machine.

Disable Windows Defender Credential Guard for a virtual machine

From the host, you can disable Windows Defender Credential Guard for a virtual machine:

```
PowerShell
Set-VMSecurity -VMName <VMName> -VirtualizationBasedSecurityOptOut $true
```

Disabling Virtualization-Based Security

Instructions are given below for how to disable Virtualization-Based Security (VBS) entirely, rather than just Windows Defender Credential Guard. Disabling Virtualization-Based Security will automatically disable Windows Defender Credential Guard and other features that rely on VBS.

> **ⓘ Important**
>
> Other security features in addition to Windows Defender Credential Guard rely on Virtualization-Based Security in order to run. Disabling Virtualization-Based Security may have unintended side effects.

1. If Group Policy was used to enable Virtualization-Based Security, set the Group Policy setting that was used to enable it (**Computer Configuration** > **Administrative Templates** > **System** > **Device Guard** > **Turn on Virtualization Based Security**) to "Disabled".

2. Delete the following registry settings:

   - `HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\DeviceGuard\EnableVirtualizationBasedSecurity`

   - `HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\DeviceGuard\RequirePlatformSecurityFeatures`

> **ⓘ Important**
>
> If you manually remove these registry settings, make sure to delete them all. If you don't remove them all, the device might go into BitLocker recovery.

3. If Windows Defender Credential Guard is running when disabling Virtualization-Based Security and either feature was enabled with UEFI Lock, the EFI (firmware) variables must be cleared using bcdedit. From an elevated command prompt, run the following bcdedit commands after turning off all Virtualization-Based Security Group Policy and registry settings as described in steps 1 and 2 above:

   Windows Command Prompt

   ```
   bcdedit /set {0cb3b571-2f2e-4343-a879-d86a476d7215} loadoptions
   DISABLE-LSA-ISO,DISABLE-VBS
   bcdedit /set vsmlaunchtype off
   ```

4. Restart the PC.

Windows Defender Credential Guard protection limits

Article • 12/15/2022 • 2 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

Some ways to store credentials are not protected by Windows Defender Credential Guard, including:

- Software that manages credentials outside of Windows feature protection
- Local accounts and Microsoft Accounts
- Windows Defender Credential Guard doesn't protect the Active Directory database running on Windows Server domain controllers. It also doesn't protect credential input pipelines, such as Windows Server running Remote Desktop Gateway. If you're using a Windows Server OS as a client PC, it will get the same protection as it would when running a Windows client OS.
- Key loggers
- Physical attacks
- Doesn't prevent an attacker with malware on the PC from using the privileges associated with any credential. We recommend using dedicated PCs for high value accounts, such as IT Pros and users with access to high value assets in your organization.
- Third-party security packages
- Digest and CredSSP credentials
    - When Windows Defender Credential Guard is enabled, neither Digest nor CredSSP have access to users' logon credentials. This implies no Single Sign-On use for these protocols.
- Supplied credentials for NTLM authentication aren't protected. If a user is prompted for and enters credentials for NTLM authentication, these credentials are vulnerable to be read from LSASS memory. These same credentials are vulnerable to key loggers as well.-
- Kerberos service tickets aren't protected by Credential Guard, but the Kerberos Ticket Granting Ticket (TGT) is.
- When Windows Defender Credential Guard is deployed on a VM, Windows Defender Credential Guard protects secrets from attacks inside the VM. However, it doesn't provide additional protection from privileged system attacks originating from the host.
- Windows logon cached password verifiers (commonly called "cached credentials") don't qualify as credentials because they can't be presented to another computer for authentication, and can only be used locally to verify credentials. They're stored

in the registry on the local computer and provide validation for credentials when a domain-joined computer can't connect to AD DS during user logon. These *cached logons*, or more specifically, *cached domain account information*, can be managed using the security policy setting **Interactive logon: Number of previous logons to cache** if a domain controller isn't available.

See also

Deep Dive into Windows Defender Credential Guard: Related videos

Microsoft Cybersecurity Stack: Advanced Identity and Endpoint Protection: Manage Credential Guard

> **Note**
>
> - Note: Requires **LinkedIn Learning subscription** to view the full video

Considerations when using Windows Defender Credential Guard

Article • 01/27/2023 • 6 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

It's recommended that in addition to deploying Windows Defender Credential Guard, organizations move away from passwords to other authentication methods, such as Windows Hello for Business, FIDO 2 security keys or smart cards.

Wi-fi and VPN considerations

When you enable Windows Defender Credential Guard, you can no longer use NTLM classic authentication for single sign-on. You'll be forced to enter your credentials to use these protocols and can't save the credentials for future use.
If you're using WiFi and VPN endpoints that are based on MS-CHAPv2, they're subject to similar attacks as for NTLMv1.

For WiFi and VPN connections, it's recommended to move from MSCHAPv2-based connections (such as PEAP-MSCHAPv2 and EAP-MSCHAPv2), to certificate-based authentication (such as PEAP-TLS or EAP-TLS).

Kerberos considerations

When you enable Windows Defender Credential Guard, you can no longer use Kerberos unconstrained delegation or DES encryption. Unconstrained delegation could allow attackers to extract Kerberos keys from the isolated LSA process.
Use constrained or resource-based Kerberos delegation instead.

Third party Security Support Providers considerations

Some third party Security Support Providers (SSPs and APs) might not be compatible with Windows Defender Credential Guard because it doesn't allow third-party SSPs to ask for password hashes from LSA. However, SSPs and APs still get notified of the password when a user logs on and/or changes their password. Any use of undocumented APIs within custom SSPs and APs aren't supported.
It's recommended that custom implementations of SSPs/APs are tested with Windows Defender Credential Guard. SSPs and APs that depend on any undocumented or

unsupported behaviors fail. For example, using the KerbQuerySupplementalCredentialsMessage API isn't supported. Replacing the NTLM or Kerberos SSPs with custom SSPs and APs.

For more information, see Restrictions around Registering and Installing a Security Package.

Upgrade considerations

As the depth and breadth of protections provided by Windows Defender Credential Guard are increased, new releases of Windows with Windows Defender Credential Guard running may affect scenarios that were working in the past. For example, Windows Defender Credential Guard may block the use of a particular type of credential or a particular component to prevent malware from taking advantage of vulnerabilities.

Test scenarios required for operations in an organization before upgrading a device using Windows Defender Credential Guard.

Saved Windows credentials protected

Domain credentials that are stored in *Credential Manager* are protected with Windows Defender Credential Guard. Credential Manager allows you to store three types of credentials:

- Windows credentials
- Certificate-based credentials
- Generic credentials

Generic credentials, such as user names and passwords that you use to sign in websites, aren't protected since the applications require your clear-text password. If the application doesn't need a copy of the password, they can save domain credentials as Windows credentials that are protected. Windows credentials are used to connect to other computers on a network.

The following considerations apply to the Windows Defender Credential Guard protections for Credential Manager:

- Windows credentials saved by the Remote Desktop client can't be sent to a remote host. Attempts to use saved Windows credentials fail, displaying the error message *Logon attempt failed*.
- Applications that extract Windows credentials fail

- When credentials are backed up from a PC that has Windows Defender Credential Guard enabled, the Windows credentials can't be restored. If you need to back up your credentials, you must do so before you enable Windows Defender Credential Guard. Otherwise, you can't restore those credentials

Clearing TPM considerations

Virtualization-based Security (VBS) uses the TPM to protect its key. When the TPM is cleared, the TPM protected key used to encrypt VBS secrets is lost.

> ⚠ **Warning**
>
> Clearing the TPM results in loss of protected data for all features that use VBS to protect data.
>
> When a TPM is cleared, **all** features that use VBS to protect data can no longer decrypt their protected data.

As a result, Credential Guard can no longer decrypt protected data. VBS creates a new TPM protected key for Credential Guard. Credential Guard uses the new key to protect new data. However, the previously protected data is lost forever.

> ⓘ **Note**
>
> Credential Guard obtains the key during initialization. The data loss will only impact persistent data and occur after the next system startup.

Windows credentials saved to Credential Manager

Since Credential Manager can't decrypt saved Windows Credentials, they're deleted. Applications should prompt for credentials that were previously saved. If saved again, then Windows credentials are protected Credential Guard.

Domain-joined device's automatically provisioned public key

Active Directory domain-joined devices automatically provision a bound public key, for more information about automatic public key provisioning, see Domain-joined Device Public Key Authentication.

Since Credential Guard can't decrypt the protected private key, Windows uses the domain-joined computer's password for authentication to the domain. Unless other policies are deployed, there shouldn't be a loss of functionality. If a device is configured to only use public key, then it can't authenticate with password until that policy is disabled. For more information on Configuring devices to only use public key, see Domain-joined Device Public Key Authentication.

Also if any access control checks including authentication policies require devices to have either the KEY TRUST IDENTITY (S-1-18-4) or FRESH PUBLIC KEY IDENTITY (S-1-18-3) well-known SIDs, then those access checks fail. For more information about authentication policies, see Authentication Policies and Authentication Policy Silos. For more information about well-known SIDs, see [MS-DTYP] Section 2.4.2.4 Well-known SID Structures.

Breaking DPAPI on domain-joined devices

On domain-joined devices, DPAPI can recover user keys using a domain controller from the user's domain. If a domain-joined device has no connectivity to a domain controller, then recovery isn't possible.

> **ⓘ Important**
>
> Best practice when clearing a TPM on a domain-joined device is to be on a network with connectivity to domain controllers. This ensures DPAPI functions and the user does not experience strange behavior.

Auto VPN configuration is protected with user DPAPI. User may not be able to use VPN to connect to domain controllers since the VPN configurations are lost. If you must clear the TPM on a domain-joined device without connectivity to domain controllers, then you should consider the following.

Domain user sign-in on a domain-joined device after clearing a TPM for as long as there's no connectivity to a domain controller:

Credential Type	Behavior
Certificate (smart card or Windows Hello for Business)	All data protected with user DPAPI is unusable and user DPAPI doesn't work at all.
Password	If the user signed in with a certificate or password prior to clearing the TPM, then they can sign-in with password and user DPAPI is unaffected.

Once the device has connectivity to the domain controllers, DPAPI recovers the user's key and data protected prior to clearing the TPM can be decrypted.

Impact of DPAPI failures on Windows Information Protection

When data protected with user DPAPI is unusable, then the user loses access to all work data protected by Windows Information Protection. The impact includes: Outlook is unable to start and work protected documents can't be opened. If DPAPI is working, then newly created work data is protected and can be accessed.

Workaround: Users can resolve the problem by connecting their device to the domain and rebooting or using their Encrypting File System Data Recovery Agent certificate. For more information about Encrypting File System Data Recovery Agent certificate, see Create and verify an Encrypting File System (EFS) Data Recovery Agent (DRA) certificate.

Additional mitigations

Article • 12/15/2022 • 18 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

Windows Defender Credential Guard can provide mitigation against attacks on derived credentials and prevent the use of stolen credentials elsewhere. However, PCs can still be vulnerable to certain attacks, even if the derived credentials are protected by Windows Defender Credential Guard. These attacks can include abusing privileges and use of derived credentials directly from a compromised device, re-using previously stolen credentials prior to Windows Defender Credential Guard, and abuse of management tools and weak application configurations. Because of this, additional mitigation also must be deployed to make the domain environment more robust.

Restricting domain users to specific domain-joined devices

Credential theft attacks allow the attacker to steal secrets from one device and use them from another device. If a user can sign on to multiple devices then any device could be used to steal credentials. How do you ensure that users only sign on with devices that have Windows Defender Credential Guard enabled? By deploying authentication policies that restrict them to specific domain-joined devices that have been configured with Windows Defender Credential Guard. For the domain controller to know what device a user is signing on from, Kerberos armoring must be used.

Kerberos armoring

Kerberos armoring is part of RFC 6113. When a device supports Kerberos armoring, its TGT is used to protect the user's proof of possession which can mitigate offline dictionary attacks. Kerberos armoring also provides the additional benefit of signed KDC errors this mitigates tampering which can result in things such as downgrade attacks.

To enable Kerberos armoring for restricting domain users to specific domain-joined devices

- Users need to be in domains that are running Windows Server 2012 R2 or higher
- All the domain controllers in these domains must be configured to support Kerberos armoring. Set the **KDC support for claims, compound authentication, and Kerberos armoring** Group Policy setting to either **Supported** or **Always provide claims**.

- All the devices with Windows Defender Credential Guard that the users will be restricted to must be configured to support Kerberos armoring. Enable the **Kerberos client support for claims, compound authentication and Kerberos armoring** Group Policy settings under **Computer Configuration** -> **Administrative Templates** -> **System** -> **Kerberos**.

Protecting domain-joined device secrets

Since domain-joined devices also use shared secrets for authentication, attackers can steal those secrets as well. By deploying device certificates with Windows Defender Credential Guard, the private key can be protected. Then authentication policies can require that users sign on to devices that authenticate using those certificates. This prevents shared secrets stolen from the device to be used with stolen user credentials to sign on as the user.

Domain-joined device certificate authentication has the following requirements:

- Devices' accounts are in Windows Server 2012 domain functional level or higher.
- All domain controllers in those domains have KDC certificates which satisfy strict KDC validation certificate requirements:
  - KDC EKU present
  - DNS domain name matches the DNSName field of the SubjectAltName (SAN) extension
- Windows devices have the CA issuing the domain controller certificates in the enterprise store.
- A process is established to ensure the identity and trustworthiness of the device in a similar manner as you would establish the identity and trustworthiness of a user before issuing them a smartcard.

Deploying domain-joined device certificates

To guarantee that certificates with the required issuance policy are only installed on the devices these users must use, they must be deployed manually on each device. The same security procedures used for issuing smart cards to users should be applied to device certificates.

For example, let's say you wanted to use the High Assurance policy only on these devices. Using a Windows Server Enterprise certificate authority, you would create a new template.

Creating a new certificate template

1. From the Certificate Manager console, right-click **Certificate Templates**, and then click **Manage.**
2. Right-click **Workstation Authentication**, and then click **Duplicate Template**.
3. Right-click the new template, and then click **Properties**.
4. On the **Extensions** tab, click **Application Policies**, and then click **Edit**.
5. Click **Client Authentication**, and then click **Remove**.
6. Add the ID-PKInit-KPClientAuth EKU. Click **Add**, click **New**, and then specify the following values:

    - Name: Kerberos Client Auth
    - Object Identifier: 1.3.6.1.5.2.3.4

7. On the **Extensions** tab, click **Issuance Policies**, and then click **Edit**.
8. Under **Issuance Policies**, click**High Assurance**.
9. On the **Subject name** tab, clear the **DNS name** check box, and then select the **User Principal Name (UPN)** check box.

Then on the devices that are running Windows Defender Credential Guard, enroll the devices using the certificate you just created.

Enrolling devices in a certificate

Run the following command:

```PowerShell
CertReq -EnrollCredGuardCert MachineAuthentication
```

> ⓘ **Note**
>
> You must restart the device after enrolling the machine authentication certificate.

How a certificate issuance policy can be used for access control

Beginning with the Windows Server 2008 R2 domain functional level, domain controllers support for authentication mechanism assurance provides a way to map certificate issuance policy OIDs to universal security groups. Windows Server 2012 domain controllers with claim support can map them to claims. To learn more about authentication mechanism assurance, see Authentication Mechanism Assurance for AD DS in Windows Server 2008 R2 Step-by-Step Guide on TechNet.

To see the issuance policies available

- The get-IssuancePolicy.ps1 shows all of the issuance policies that are available on the certificate authority.
  From a Windows PowerShell command prompt, run the following command:

```
PowerShell

.\get-IssuancePolicy.ps1 -LinkedToGroup:All
```

To link an issuance policy to a universal security group

- The set-IssuancePolicyToGroupLink.ps1 creates a Universal security group, creates an organizational unit, and links the issuance policy to that Universal security group.
  From a Windows PowerShell command prompt, run the following command:

```
PowerShell

.\set-IssuancePolicyToGroupLink.ps1 -IssuancePolicyName:"<name of issuance policy>" -groupOU:"<Name of OU to create>" -groupName:"<name of Universal security group to create>"
```

Restricting user sign-on

So we now have completed the following:

- Created a special certificate issuance policy to identify devices that meet the deployment criteria required for the user to be able to sign on
- Mapped that policy to a universal security group or claim
- Provided a way for domain controllers to get the device authorization data during user sign-on using Kerberos armoring. Now what is left to do is to configure the access check on the domain controllers. This is done using authentication policies.

Authentication policies have the following requirements:

- User accounts are in a Windows Server 2012 domain functional level or higher domain.

Creating an authentication policy restricting users to the specific universal security group

1. Open Active Directory Administrative Center.
2. Click **Authentication**, click **New**, and then click **Authentication Policy**.
3. In the **Display name** box, enter a name for this authentication policy.
4. Under the **Accounts** heading, click **Add**.

5. In the **Select Users, Computers, or Service Accounts** dialog box, type the name of the user account you wish to restrict, and then click **OK**.
6. Under the **User Sign On** heading, click the **Edit** button.
7. Click **Add a condition**.
8. In the **Edit Access Control Conditions** box, ensure that it reads **User** > **Group** > **Member of each** > **Value**, and then click **Add items**.
9. In the **Select Users, Computers, or Service Accounts** dialog box, type the name of the universal security group that you created with the set-IssuancePolicyToGroupLink script, and then click **OK**.
10. Click **OK** to close the **Edit Access Control Conditions** box.
11. Click **OK** to create the authentication policy.
12. Close Active Directory Administrative Center.

> **Note**
>
> When the authentication policy enforces policy restrictions, users will not be able to sign on using devices that do not have a certificate with the appropriate issuance policy deployed. This applies to both local and remote sign on scenarios. Therefore, it is strongly recommended to first only audit policy restrictions to ensure you don't have unexpected failures.

Discovering authentication failures due to authentication policies

To make tracking authentication failures due to authentication policies easier, an operational log exists with just those events. To enable the logs on the domain controllers, in Event Viewer, navigate to **Applications and Services Logs\Microsoft\Windows\Authentication, right-click AuthenticationPolicyFailures-DomainController**, and then click **Enable Log**.

To learn more about authentication policy events, see Authentication Policies and Authentication Policy Silos.

Appendix: Scripts

Here is a list of scripts mentioned in this topic.

Get the available issuance policies on the certificate authority

Save this script file as get-IssuancePolicy.ps1.

PowerShell

```
########################################
##     Parameters to be defined       ##
##     by the user                    ##
########################################
Param (
$Identity,
$LinkedToGroup
)
########################################
##     Strings definitions            ##
########################################
Data getIP_strings {
# culture="en-US"
ConvertFrom-StringData -stringdata @'
help1 = This command can be used to retrieve all available Issuance Policies
in a forest. The forest of the currently logged on user is targeted.
help2 = Usage:
help3 = The following parameter is mandatory:
help4 = -LinkedToGroup:<yes|no|all>
help5 = "yes" will return only Issuance Policies that are linked to groups.
Checks that the linked Issuance Policies are linked to valid groups.
help6 = "no" will return only Issuance Policies that are not currently
linked to any group.
help7 = "all" will return all Issuance Policies defined in the forest.
Checks that the linked Issuance policies are linked to valid groups.
help8 = The following parameter is optional:
help9 = -Identity:<Name, Distinguished Name or Display Name of the Issuance
Policy that you want to retrieve>. If you specify an identity, the option
specified in the "-LinkedToGroup" parameter is ignored.
help10 = Output: This script returns the Issuance Policy objects meeting the
criteria defined by the above parameters.
help11 = Examples:
errorIPNotFound = Error: no Issuance Policy could be found with Identity "
{0}"
ErrorNotSecurity = Error: Issuance Policy "{0}" is linked to group "{1}"
which is not of type "Security".
ErrorNotUniversal = Error: Issuance Policy "{0}" is linked to group "{1}"
whose scope is not "Universal".
ErrorHasMembers = Error: Issuance Policy "{0}" is linked to group "{1}"
which has a non-empty membership. The group has the following members:
LinkedIPs = The following Issuance Policies are linked to groups:
displayName = displayName : {0}
Name = Name : {0}
dn = distinguishedName : {0}
        InfoName = Linked Group Name: {0}
        InfoDN = Linked Group DN: {0}
NonLinkedIPs = The following Issuance Policies are NOT linked to groups:
'@
}
##Import-LocalizedData getIP_strings
import-module ActiveDirectory
########################################
```

```
##           Help                    ##
#######################################
function Display-Help {
    ""
    $getIP_strings.help1
    ""
$getIP_strings.help2
""
$getIP_strings.help3
"     " + $getIP_strings.help4
"             " + $getIP_strings.help5
    "             " + $getIP_strings.help6
    "             " + $getIP_strings.help7
""
$getIP_strings.help8
    "     " + $getIP_strings.help9
    ""
    $getIP_strings.help10
""
""
$getIP_strings.help11
    "     " + '$' + "myIPs = .\get-IssuancePolicy.ps1 -LinkedToGroup:All"
    "     " + '$' + "myLinkedIPs = .\get-IssuancePolicy.ps1 -
LinkedToGroup:yes"
    "     " + '$' + "myIP = .\get-IssuancePolicy.ps1 -Identity:""Medium
Assurance"""
""
}
$root = get-adrootdse
$domain = get-addomain -current loggedonuser
$configNCDN = [String]$root.configurationNamingContext
if ( !($Identity) -and !($LinkedToGroup) ) {
display-Help
break
}
if ($Identity) {
    $OIDs = get-adobject -Filter {(objectclass -eq "msPKI-Enterprise-Oid") -
and ((name -eq $Identity) -or (displayname -eq $Identity) -or
(distinguishedName -like $Identity)) } -searchBase $configNCDN -properties *
    if ($OIDs -eq $null) {
$errormsg = $getIP_strings.ErrorIPNotFound -f $Identity
write-host $errormsg -ForegroundColor Red
    }
    foreach ($OID in $OIDs) {
        if ($OID."msDS-OIDToGroupLink") {
# In case the Issuance Policy is linked to a group, it is good to check
whether there is any problem with the mapping.
            $groupDN = $OID."msDS-OIDToGroupLink"
            $group = get-adgroup -Identity $groupDN
    $groupName = $group.Name
# Analyze the group
            if ($group.groupCategory -ne "Security") {
$errormsg = $getIP_strings.ErrorNotSecurity -f $Identity, $groupName
                write-host $errormsg -ForegroundColor Red
            }
```

```
            if ($group.groupScope -ne "Universal") {
                $errormsg = $getIP_strings.ErrorNotUniversal -f $Identity,
$groupName
write-host $errormsg -ForegroundColor Red
            }
            $members = Get-ADGroupMember -Identity $group
            if ($members) {
                $errormsg = $getIP_strings.ErrorHasMembers -f $Identity,
$groupName
write-host $errormsg -ForegroundColor Red
                foreach ($member in $members) {
                    write-host "          "  $member -ForeGroundColor Red
                }
            }
        }
    }
    return $OIDs
    break
}
if (($LinkedToGroup -eq "yes") -or ($LinkedToGroup -eq "all")) {
    $LDAPFilter = "(&(objectClass=msPKI-Enterprise-Oid)(msDS-
OIDToGroupLink=*)(flags=2))"
    $LinkedOIDs = get-adobject -searchBase $configNCDN -LDAPFilter
$LDAPFilter -properties *
    write-host ""
    write-host "**************************************************"
    write-host $getIP_strings.LinkedIPs
    write-host "**************************************************"
    write-host ""
    if ($LinkedOIDs -ne $null){
      foreach ($OID in $LinkedOIDs) {
# Display basic information about the Issuance Policies
          ""
  $getIP_strings.displayName -f $OID.displayName
  $getIP_strings.Name -f $OID.Name
  $getIP_strings.dn -f $OID.distinguishedName
# Get the linked group.
          $groupDN = $OID."msDS-OIDToGroupLink"
          $group = get-adgroup -Identity $groupDN
          $getIP_strings.InfoName -f $group.Name
          $getIP_strings.InfoDN -f $groupDN
# Analyze the group
          $OIDName = $OID.displayName
    $groupName = $group.Name
          if ($group.groupCategory -ne "Security") {
          $errormsg = $getIP_strings.ErrorNotSecurity -f $OIDName,
$groupName
          write-host $errormsg -ForegroundColor Red
          }
          if ($group.groupScope -ne "Universal") {
          $errormsg = $getIP_strings.ErrorNotUniversal -f $OIDName,
$groupName
          write-host $errormsg -ForegroundColor Red
          }
          $members = Get-ADGroupMember -Identity $group
```

```
            if ($members) {
            $errormsg = $getIP_strings.ErrorHasMembers -f $OIDName, $groupName
            write-host $errormsg -ForegroundColor Red
                foreach ($member in $members) {
                    write-host "          "  $member -ForeGroundColor Red
                }
            }
            write-host ""
        }
    }else{
write-host "There are no issuance policies that are mapped to a group"
    }
    if ($LinkedToGroup -eq "yes") {
        return $LinkedOIDs
        break
    }
}
if (($LinkedToGroup -eq "no") -or ($LinkedToGroup -eq "all")) {
    $LDAPFilter = "(&(objectClass=msPKI-Enterprise-Oid)(!(msDS-
OIDToGroupLink=*))(flags=2))"
    $NonLinkedOIDs = get-adobject -searchBase $configNCDN -LDAPFilter
$LDAPFilter -properties *
    write-host ""
    write-host "************************************************************"
    write-host $getIP_strings.NonLinkedIPs
    write-host "************************************************************"
    write-host ""
    if ($NonLinkedOIDs -ne $null) {
      foreach ($OID in $NonLinkedOIDs) {
# Display basic information about the Issuance Policies
write-host ""
$getIP_strings.displayName -f $OID.displayName
$getIP_strings.Name -f $OID.Name
$getIP_strings.dn -f $OID.distinguishedName
write-host ""
      }
    }else{
write-host "There are no issuance policies which are not mapped to groups"
    }
    if ($LinkedToGroup -eq "no") {
        return $NonLinkedOIDs
        break
    }
}
```

> **Note**
>
> If you're having trouble running this script, try replacing the single quote after the ConvertFrom-StringData parameter.

Link an issuance policy to a group

Save the script file as set-IssuancePolicyToGroupLink.ps1.

PowerShell

```
########################################
##     Parameters to be defined       ##
##     by the user                    ##
########################################
Param (
$IssuancePolicyName,
$groupOU,
$groupName
)
########################################
##     Strings definitions            ##
########################################
Data ErrorMsg {
# culture="en-US"
ConvertFrom-StringData -stringdata @'
help1 = This command can be used to set the link between a certificate
issuance policy and a universal security group.
help2 = Usage:
help3 = The following parameters are required:
help4 = -IssuancePolicyName:<name or display name of the issuance policy
that you want to link to a group>
help5 = -groupName:<name of the group you want to link the issuance policy
to>. If no name is specified, any existing link to a group is removed from
the Issuance Policy.
help6 = The following parameter is optional:
help7 = -groupOU:<Name of the Organizational Unit dedicated to the groups
which are linked to issuance policies>. If this parameter is not specified,
the group is looked for or created in the Users container.
help8 = Examples:
help9 = This command will link the issuance policy whose display name is
"High Assurance" to the group "HighAssuranceGroup" in the Organizational
Unit "OU_FOR_IPol_linked_groups". If the group or the Organizational Unit do
not exist, you will be prompted to create them.
help10 = This command will unlink the issuance policy whose name is
"402.164959C40F4A5C12C6302E31D5476062" from any group.
MultipleIPs = Error: Multiple Issuance Policies with name or display name "
{0}" were found in the subtree of "{1}"
NoIP = Error: no issuance policy with name or display name "{0}" could be
found in the subtree of "{1}".
IPFound = An Issuance Policy with name or display name "{0}" was
successfully found: {1}
MultipleOUs = Error: more than 1 Organizational Unit with name "{0}" could
be found in the subtree of "{1}".
confirmOUcreation = Warning: The Organizational Unit that you specified does
not exist. Do you want to create it?
OUCreationSuccess = Organizational Unit "{0}" successfully created.
OUcreationError = Error: Organizational Unit "{0}" could not be created.
```

```
OUFoundSuccess = Organizational Unit "{0}" was successfully found.
multipleGroups = Error: More than one group with name "{0}" was found in
Organizational Unit "{1}".
confirmGroupCreation = Warning: The group that you specified does not exist.
Do you want to create it?
groupCreationSuccess = Univeral Security group "{0}" successfully created.
groupCreationError = Error: Univeral Security group "{0}" could not be
created.
GroupFound = Group "{0}" was successfully found.
confirmLinkDeletion = Warning: The Issuance Policy "{0}" is currently linked
to group "{1}". Do you really want to remove the link?
UnlinkSuccess = Certificate issuance policy successfully unlinked from any
group.
UnlinkError = Removing the link failed.
UnlinkExit = Exiting without removing the link from the issuance policy to
the group.
IPNotLinked = The Certificate issuance policy is not currently linked to any
group. If you want to link it to a group, you should specify the -groupName
option when starting this script.
ErrorNotSecurity = Error: You cannot link issuance Policy "{0}" to group "
{1}" because this group is not of type "Security".
ErrorNotUniversal = Error: You cannot link issuance Policy "{0}" to group "
{1}" because the scope of this group is not "Universal".
ErrorHasMembers = Error: You cannot link issuance Policy "{0}" to group "
{1}" because it has a non-empty membership. The group has the following
members:
ConfirmLinkReplacement = Warning: The Issuance Policy "{0}" is currently
linked to group "{1}". Do you really want to update the link to point to
group "{2}"?
LinkSuccess = The certificate issuance policy was successfully linked to the
specified group.
LinkError = The certificate issuance policy could not be linked to the
specified group.
ExitNoLinkReplacement = Exiting without setting the new link.
'@
}
# import-localizeddata ErrorMsg
function Display-Help {
""
write-host $ErrorMsg.help1
""
write-host $ErrorMsg.help2
""
write-host $ErrorMsg.help3
write-host "`t" $ErrorMsg.help4
write-host "`t" $ErrorMsg.help5
""
write-host $ErrorMsg.help6
write-host "`t" $ErrorMsg.help7
""
""
write-host $ErrorMsg.help8
""
write-host $ErrorMsg.help9
".\Set-IssuancePolicyToGroupMapping.ps1 -IssuancePolicyName ""High
```

```
Assurance"" -groupOU ""OU_FOR_IPol_linked_groups"" -groupName
""HighAssuranceGroup"" "
""
write-host $ErrorMsg.help10
'.\Set-IssuancePolicyToGroupMapping.ps1 -IssuancePolicyName
"402.164959C40F4A5C12C6302E31D5476062" -groupName $null '
""
}
# Assumption:  The group to which the Issuance Policy is going
#              to be linked is (or is going to be created) in
#              the domain the user running this script is a member of.
import-module ActiveDirectory
$root = get-adrootdse
$domain = get-addomain -current loggedonuser
if ( !($IssuancePolicyName) ) {
display-Help
break
}
########################################
##     Find the OID object            ##
##     (aka Issuance Policy)          ##
########################################
$searchBase = [String]$root.configurationnamingcontext
$OID = get-adobject -searchBase $searchBase -Filter { ((displayname -eq
$IssuancePolicyName) -or (name -eq $IssuancePolicyName)) -and (objectClass -
eq "msPKI-Enterprise-Oid")} -properties *
if ($OID -eq $null) {
$tmp = $ErrorMsg.NoIP -f $IssuancePolicyName, $searchBase
write-host $tmp -ForeGroundColor Red
break;
}
elseif ($OID.GetType().IsArray) {
$tmp = $ErrorMsg.MultipleIPs -f $IssuancePolicyName, $searchBase
write-host $tmp -ForeGroundColor Red
break;
}
else {
$tmp = $ErrorMsg.IPFound -f $IssuancePolicyName, $OID.distinguishedName
write-host $tmp -ForeGroundColor Green
}
########################################
##  Find the container of the group   ##
########################################
if ($groupOU -eq $null) {
# default to the Users container
$groupContainer = $domain.UsersContainer
}
else {
$searchBase = [string]$domain.DistinguishedName
$groupContainer = get-adobject -searchBase $searchBase -Filter { (Name -eq
$groupOU) -and (objectClass -eq "organizationalUnit")}
if ($groupContainer.count -gt 1) {
$tmp = $ErrorMsg.MultipleOUs -f $groupOU, $searchBase
write-host $tmp -ForegroundColor Red
break;
```

```
}
elseif ($groupContainer -eq $null) {
$tmp = $ErrorMsg.confirmOUcreation
write-host $tmp " ( (y)es / (n)o )" -ForegroundColor Yellow -nonewline
$userChoice = read-host
if ( ($userChoice -eq "y") -or ($userChoice -eq "yes") ) {
new-adobject -Name $groupOU -displayName $groupOU -Type "organizationalUnit"
-ProtectedFromAccidentalDeletion $true -path $domain.distinguishedName
if ($?){
$tmp = $ErrorMsg.OUCreationSuccess -f $groupOU
write-host $tmp -ForegroundColor Green
}
else{
$tmp = $ErrorMsg.OUCreationError -f $groupOU
write-host $tmp -ForeGroundColor Red
break;
}
$groupContainer = get-adobject -searchBase $searchBase -Filter { (Name -eq
$groupOU) -and (objectClass -eq "organizationalUnit")}
}
else {
break;
}
}
else {
$tmp = $ErrorMsg.OUFoundSuccess -f $groupContainer.name
write-host $tmp -ForegroundColor Green
}
}
########################################
##  Find the group                ##
########################################
if (($groupName -ne $null) -and ($groupName -ne "")){
##$searchBase = [String]$groupContainer.DistinguishedName
$searchBase = $groupContainer
$group = get-adgroup -Filter { (Name -eq $groupName) -and (objectClass -eq
"group") } -searchBase $searchBase
if ($group -ne $null -and $group.gettype().isarray) {
$tmp = $ErrorMsg.multipleGroups -f $groupName, $searchBase
write-host $tmp -ForeGroundColor Red
break;
}
elseif ($group -eq $null) {
$tmp = $ErrorMsg.confirmGroupCreation
write-host $tmp " ( (y)es / (n)o )" -ForegroundColor Yellow -nonewline
$userChoice = read-host
if ( ($userChoice -eq "y") -or ($userChoice -eq "yes") ) {
new-adgroup -samAccountName $groupName -path
$groupContainer.distinguishedName -GroupScope "Universal" -GroupCategory
"Security"
if ($?){
$tmp = $ErrorMsg.GroupCreationSuccess -f $groupName
write-host $tmp -ForegroundColor Green
}else{
$tmp = $ErrorMsg.groupCreationError -f $groupName
```

```
write-host $tmp -ForeGroundColor Red
break
}
$group = get-adgroup -Filter { (Name -eq $groupName) -and (objectClass -eq
"group") } -searchBase $searchBase
}
else {
break;
}
}
else {
$tmp = $ErrorMsg.GroupFound -f $group.Name
write-host $tmp -ForegroundColor Green
}
}
else {
#####
## If the group is not specified, we should remove the link if any exists
#####
if ($OID."msDS-OIDToGroupLink" -ne $null) {
$tmp = $ErrorMsg.confirmLinkDeletion -f $IssuancePolicyName, $OID."msDS-
OIDToGroupLink"
write-host $tmp " ( (y)es / (n)o )" -ForegroundColor Yellow -nonewline
$userChoice = read-host
if ( ($userChoice -eq "y") -or ($userChoice -eq "yes") ) {
set-adobject -Identity $OID -Clear "msDS-OIDToGroupLink"
if ($?) {
$tmp = $ErrorMsg.UnlinkSuccess
write-host $tmp -ForeGroundColor Green
}else{
$tmp = $ErrorMsg.UnlinkError
write-host $tmp -ForeGroundColor Red
}
}
else {
$tmp = $ErrorMsg.UnlinkExit
write-host $tmp
break
}
}
else {
$tmp = $ErrorMsg.IPNotLinked
write-host $tmp -ForeGroundColor Yellow
}
break;
}
########################################
##  Verify that the group is          ##
##  Universal, Security, and          ##
##  has no members                    ##
########################################
if ($group.GroupScope -ne "Universal") {
$tmp = $ErrorMsg.ErrorNotUniversal -f $IssuancePolicyName, $groupName
write-host $tmp -ForeGroundColor Red
break;
```

```
}
if ($group.GroupCategory -ne "Security") {
$tmp = $ErrorMsg.ErrorNotSecurity -f $IssuancePolicyName, $groupName
write-host $tmp -ForeGroundColor Red
break;
}
$members = Get-ADGroupMember -Identity $group
if ($members -ne $null) {
$tmp = $ErrorMsg.ErrorHasMembers -f $IssuancePolicyName, $groupName
write-host $tmp -ForeGroundColor Red
foreach ($member in $members) {write-host "   $member.name" -ForeGroundColor
Red}
break;
}
########################################
##  We have verified everything. We  ##
##  can create the link from the     ##
##  Issuance Policy to the group.    ##
########################################
if ($OID."msDS-OIDToGroupLink" -ne $null) {
$tmp = $ErrorMsg.ConfirmLinkReplacement -f $IssuancePolicyName, $OID."msDS-
OIDToGroupLink", $group.distinguishedName
write-host $tmp  "( (y)es / (n)o )" -ForegroundColor Yellow -nonewline
$userChoice = read-host
if ( ($userChoice -eq "y") -or ($userChoice -eq "yes") ) {
$tmp = @{'msDS-OIDToGroupLink'= $group.DistinguishedName}
set-adobject -Identity $OID -Replace $tmp
if ($?) {
$tmp = $Errormsg.LinkSuccess
write-host $tmp -Foreground Green
}else{
$tmp = $ErrorMsg.LinkError
write-host $tmp -Foreground Red
}
} else {
$tmp = $Errormsg.ExitNoLinkReplacement
write-host $tmp
break
}
}
else {
$tmp = @{'msDS-OIDToGroupLink'= $group.DistinguishedName}
set-adobject -Identity $OID -Add $tmp
if ($?) {
$tmp = $Errormsg.LinkSuccess
write-host $tmp -Foreground Green
}else{
$tmp = $ErrorMsg.LinkError
write-host $tmp -Foreground Red
}
}
```

ⓘ Note

If you're having trouble running this script, try replacing the single quote after the ConvertFrom-StringData parameter.

Windows Defender Credential Guard: Known issues

Article • 12/15/2022 • 7 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

Windows Defender Credential Guard has certain application requirements. Windows Defender Credential Guard blocks specific authentication capabilities. So applications that require such capabilities won't function when it's enabled. For more information, see Application requirements.

Known Issue: Single Sign-On (SSO) for Network services breaks after upgrading to Windows 11, version 22H2

Symptoms of the issue:

Devices that use 802.1x wireless or wired network, RDP, or VPN connections that rely on insecure protocols with password-based authentication will be unable to use SSO to log in and will be forced to manually re-authenticate in every new Windows session when Windows Defender Credential Guard is running.

Affected devices:

Any device that enables Windows Defender Credential Guard may encounter this issue. As part of the Windows 11, version 22H2 update, eligible devices which had not previously explicitly disabled Windows Defender Credential Guard had it enabled by default. This affected all devices on Enterprise (E3 and E5) and Education licenses, as well as some Pro licenses*, as long as they met the minimum hardware requirements.

* All Pro devices which previously ran Windows Defender Credential Guard on an eligible license and later downgraded to Pro, and which still meet the minimum hardware requirements, will receive default enablement.

> **Tip**
>
> To determine if your Pro device will receive default enablement when upgraded to **Windows 11, version 22H2**, do the following **before** upgrading:

Check if the registry key `IsolatedCredentialsRootSecret` is present in `Computer\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\MSV1_0`. If it is present, the device will have Windows Defender Credential Guard enabled after upgrading. Note that Windows Defender Credential Guard can be disabled after upgrade by following the **disablement instructions**.

Why this is happening:

Applications and services are affected by this issue when they rely on insecure protocols that use password-based authentication. Windows Defender Credential Guard blocks the use of these insecure protocols by design. These protocols are considered insecure because they can lead to password disclosure on the client and the server, which is in direct contradiction to the goals of Windows Defender Credential Guard. Affected procols include:

- Kerberos unconstrained delegation (both SSO and supplied credentials are blocked)
- Kerberos when PKINIT uses RSA encryption instead of Diffie-Hellman (both SSO and supplied credentials are blocked)
- MS-CHAP (only SSO is blocked)
- WDigest (only SSO is blocked)
- NTLM v1 (only SSO is blocked)

Since only SSO is blocked for MS-CHAP, WDigest, and NTLM v1, these protocols can still be used by prompting the user to supply credentials.

> ⓘ **Note**
>
> MS-CHAP and NTLMv1 are particularly relevant to the observed SSO breakage after the Windows 11, version 22H2 update. To confirm whether Windows Defender Credential Guard is blocking either of these protocols, check the NTLM event logs in Event Viewer at `Application and Services Logs\Microsoft\Windows\NTLM\Operational` for the following warning and/or error:
>
> **Event ID 4013** (Warning)

```
<string
  id="NTLMv1BlockedByCredGuard"
  value="Attempt to use NTLMv1 failed.
  Target server: %1%nSupplied user: %2%nSupplied domain: %3%nPID of
```

```
client process: %4%nName of client process: %5%nLUID of client process:
%6%nUser identity of client process: %7%nDomain name of user identity
of client process: %8%nMechanism OID: %9%n%nThis device does not
support NTLMv1. For more information, see
https://go.microsoft.com/fwlink/?linkid=856826."
/>
```

Event ID 4014 (Error)

```
<string
   id="NTLMGetCredentialKeyBlockedByCredGuard"
   value="Attempt to get credential key by call package blocked by
Credential Guard.%n%nCalling Process Name: %1%nService Host Tag: %2"
/>
```

Options to fix the issue:

Microsoft recommends that organizations move away from MSCHAPv2-based connections such as PEAP-MSCHAPv2 and EAP-MSCHAPv2, to certificate-based authentication such as PEAP-TLS or EAP-TLS. Windows Defender Credential Guard will not block certificate-based authentication.

For a more immediate but less secure fix, disable Windows Defender Credential Guard. Note that Windows Defender Credential Guard does not have per-protocol or per-application policies, and must either be completely on or off. Disabling Windows Defender Credential Guard will leave some stored domain credentials vulnerable to theft. Windows Defender Credential Guard can be disabled after it has already been enabled, or it can be explicitly disabled prior to updating to Windows 11, version 22H2, which will prevent default enablement from occurring.

Tip

To *prevent* default enablement, **use Group Policy to explicitly disable Windows Defender Credential Guard** before updating to Windows 11, version 22H2. If the GPO value is not configured (which is the default state), the device will receive default enablement after updating, if eligible. If the GPO value is set to "disabled", it will not be enabled after updating. This process can also be done via Mobile Device Management (MDM) policy rather than Group Policy if the devices are currently being managed by MDM.

Known issues involving third-party applications

The following issue affects MSCHAPv2:

- Credential guard doesn't work with MSCHAPv2 configurations, of which Cisco ISE is a very popular enterprise implementation.

The following issue affects the Java GSS API. See the following Oracle bug database article:

- JDK-8161921: Windows Defender Credential Guard doesn't allow sharing of TGT with Java

When Windows Defender Credential Guard is enabled on Windows, the Java GSS API won't authenticate. This is expected behavior because Windows Defender Credential Guard blocks specific application authentication capabilities and won't provide the TGT session key to applications regardless of registry key settings. For more information, see Application requirements.

The following issue affects Cisco AnyConnect Secure Mobility Client:

- Blue screen on Windows computers running Hypervisor-Protected Code Integrity and Windows Defender Credential Guard with Cisco Anyconnect 4.3.04027

The following issue affects McAfee Application and Change Control (MACC):

- KB88869 Windows machines exhibit high CPU usage with McAfee Application and Change Control (MACC) installed when Windows Defender Credential Guard is enabled Note 1

The following issue affects Citrix applications:

- Windows machines exhibit high CPU usage with Citrix applications installed when Windows Defender Credential Guard is enabled. Note 1

> **Note**
>
> **Note 1**: Products that connect to Virtualization Based Security (VBS) protected processes can cause Windows Defender Credential Guard-enabled Windows 10, Windows 11, Windows Server 2016, or Windows Server 2019 machines to exhibit high CPU usage. For technical and troubleshooting information, see **KB4032786 High CPU usage in the LSAISO process on Windows**.
>
> For more technical information on LSAISO.exe, see **Isolated User Mode (IUM) Processes**.

Vendor support

For more information on Citrix support for Secure Boot, see Citrix Support for Secure Boot

Windows Defender Credential Guard isn't supported by the following products, products versions, computer systems, or Windows 10 versions:

- Support for Hypervisor-Protected Code Integrity and Windows Defender Credential Guard on Windows with McAfee encryption products
- Check Point Endpoint Security Client support for Microsoft Windows Defender Credential Guard and Hypervisor-Protected Code Integrity features
- "VMware Workstation and Device/Credential Guard are not compatible" error in VMware Workstation on Windows 10 host (2146361)
- ThinkPad support for Hypervisor-Protected Code Integrity and Windows Defender Credential Guard in Microsoft Windows
- Windows devices with Windows Defender Credential Guard and Symantec Endpoint Protection 12.1

This list isn't comprehensive. Check whether your product vendor, product version, or computer system supports Windows Defender Credential Guard on systems that run Windows or specific versions of Windows. Specific computer system models may be incompatible with Windows Defender Credential Guard.

Microsoft encourages third-party vendors to contribute to this page by providing relevant product support information and by adding links to their own product support statements.

Previous known issues that have been fixed

The following known issues have been fixed in the Cumulative Security Update for November 2017:

- Scheduled tasks with domain user-stored credentials fail to run when Credential Guard is enabled. The task fails and reports Event ID 104 with the following message:

```
Console
```

```
Task Scheduler failed to log on '\Test'.
Failure occurred in 'LogonUserExEx'.
User Action: Ensure the credentials for the task are correctly
specified.
Additional Data: Error Value: 2147943726. 2147943726:
ERROR\_LOGON\_FAILURE (The user name or password is incorrect).
```

- When you enable NTLM audit on the domain controller, an Event ID 8004 with an indecipherable username format is logged. You also get a similar user name in a user logon failure event 4625 with error 0xC0000064 on the machine itself. For example:

  Console

  ```
  Log Name: Microsoft-Windows-NTLM/Operational
  Source: Microsoft-Windows-Security-Netlogon
  Event ID: 8004
  Task Category: Auditing NTLM
  Level: Information
  Description:
  Domain Controller Blocked Audit: Audit NTLM authentication to this
  domain controller.
  Secure Channel name: <Secure Channel Name>
  User name:
  @@CyBAAAAUBQYAMHArBwUAMGAoBQZAQGA1BAbAUGAyBgOAQFAhBwcAsGA6AweAgDA2AQQAM
  EAwAANAgDA1AQLAIEADBQRAADAtAANAYEA1AwQA0CA5AAOAMEAyAQLAYDAxAwQAEDAEBwMA
  MEAwAgMAMDACBgRA0HA
  Domain name: NULL
  ```

  - This event stems from a scheduled task running under local user context with the Cumulative Security Update for November 2017 or later and happens when Credential Guard is enabled.
  - The username appears in an unusual format because local accounts aren't protected by Credential Guard. The task also fails to execute.
  - As a workaround, run the scheduled task under a domain user or the computer's SYSTEM account.

The following known issues have been fixed by servicing releases made available in the Cumulative Security Updates for April 2017:

- KB4015217 Windows Defender Credential Guard generates double bad password count on Active Directory domain-joined Windows machines

  This issue can potentially lead to unexpected account lockouts. For more information, see the following support articles:
  - KB4015219
  - KB4015221

Protect Remote Desktop credentials with Windows Defender Remote Credential Guard

Article • 02/17/2023 • 8 minutes to read • Applies to: ✅ **Windows 10**, ✅ **Windows Server 2016**

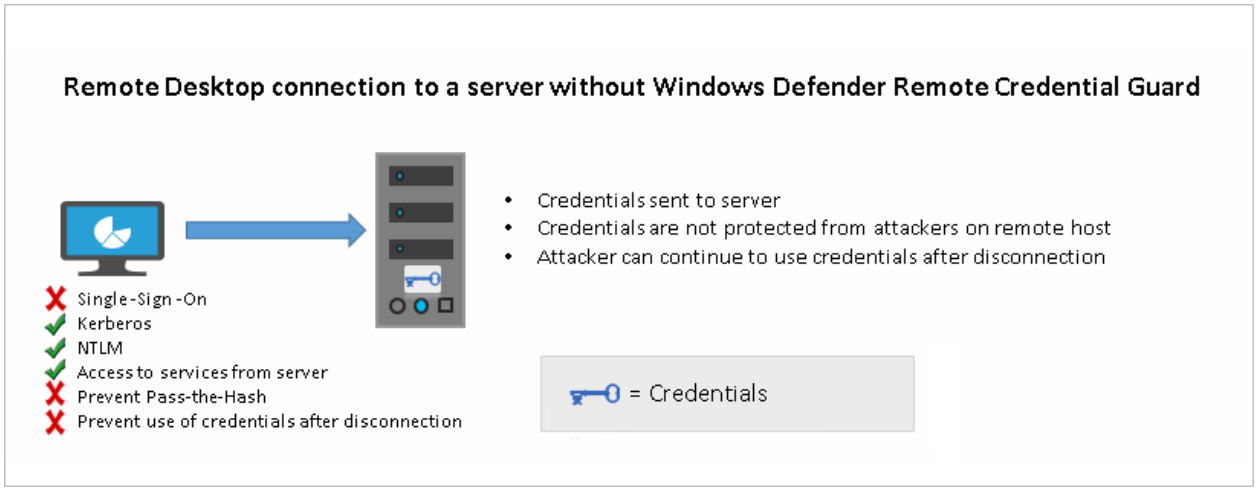
Introduced in Windows 10, version 1607, Windows Defender Remote Credential Guard helps you protect your credentials over a Remote Desktop connection by redirecting Kerberos requests back to the device that's requesting the connection. It also provides single sign-on experiences for Remote Desktop sessions.

Administrator credentials are highly privileged and must be protected. By using Windows Defender Remote Credential Guard to connect during Remote Desktop sessions, if the target device is compromised, your credentials are not exposed because both credential and credential derivatives are never passed over the network to the target device.

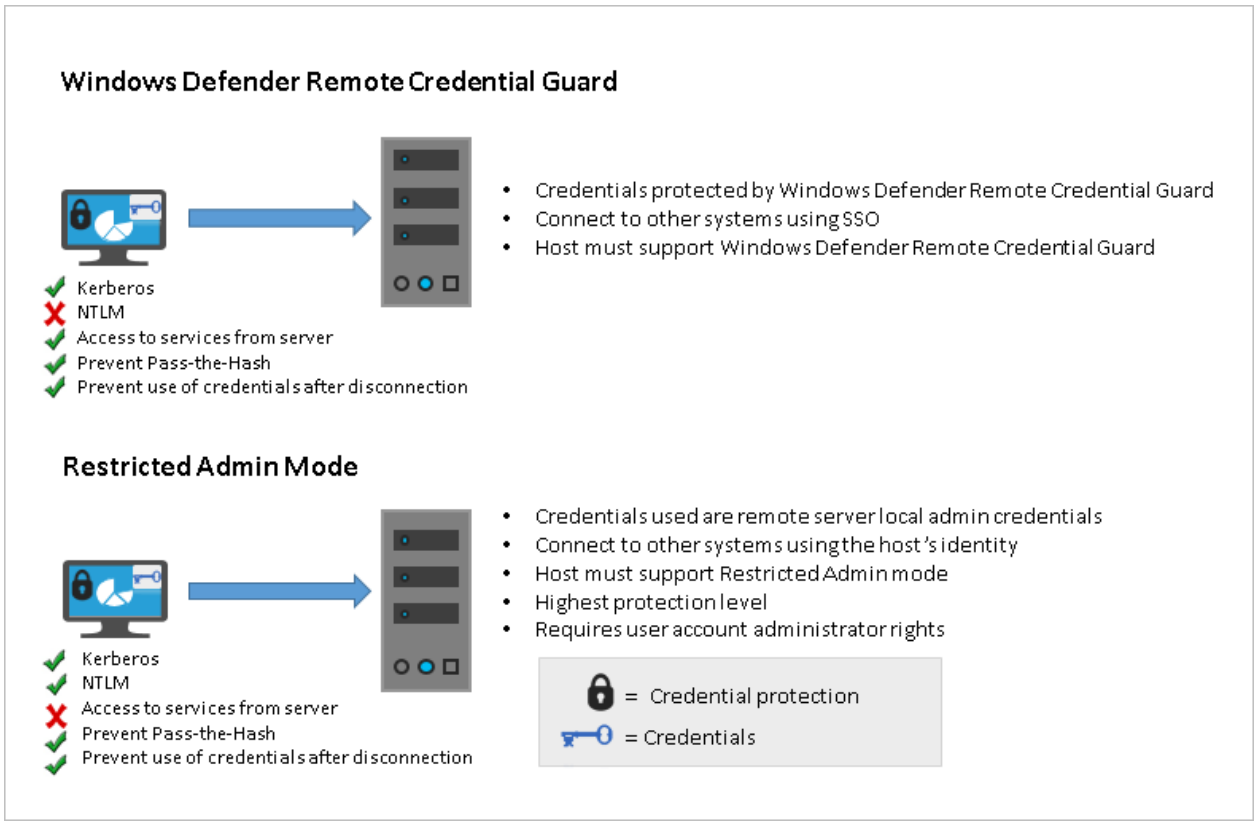
> ⓘ **Important**
>
> For information on Remote Desktop connection scenarios involving helpdesk support, see **Remote Desktop connections and helpdesk support scenarios** in this article.

Comparing Windows Defender Remote Credential Guard with other Remote Desktop connection options

The following diagram helps you to understand how a standard Remote Desktop session to a server without Windows Defender Remote Credential Guard works:



The following diagram helps you to understand how Windows Defender Remote Credential Guard works, what it helps to protect against, and compares it with the Restricted Admin mode option:



As illustrated, Windows Defender Remote Credential Guard blocks NTLM (allowing only Kerberos), prevents Pass-the-Hash (PtH) attacks, and also prevents use of credentials after disconnection.

Use the following table to compare different Remote Desktop connection security options:

Feature	Remote Desktop	Windows Defender Remote Credential Guard	Restricted Admin mode
Protection benefits	Credentials on the server are not protected from Pass-the-Hash attacks.	User credentials remain on the client. An attacker can act on behalf of the user *only* when the session is ongoing	User logs on to the server as local administrator, so an attacker cannot act on behalf of the "domain user". Any attack is local to the server
Version support	The remote computer can run any Windows operating system	Both the client and the remote computer must be running **at least Windows 10, version 1607, or Windows Server 2016**.	The remote computer must be running **at least patched Windows 7 or patched Windows Server 2008 R2**. For more information about patches (software updates) related to Restricted Admin mode, see Microsoft Security Advisory 2871997.
Helps prevent	N/A	• Pass-the-Hash • Use of a credential after disconnection	• Pass-the-Hash • Use of domain identity during connection
Credentials supported from the remote desktop client device	• **Signed on** credentials • **Supplied** credentials • **Saved** credentials	• **Signed on** credentials only	• **Signed on** credentials • **Supplied** credentials • **Saved** credentials
Access	**Users allowed**, that is, members of Remote Desktop Users group of remote host.	**Users allowed**, that is, members of Remote Desktop Users of remote host.	**Administrators only**, that is, only members of Administrators group of remote host.
Network identity	Remote Desktop session **connects to other resources as signed-in user**.	Remote Desktop session **connects to other resources as signed-in user**.	Remote Desktop session **connects to other resources as remote host's identity**.

Feature	Remote Desktop	Windows Defender Remote Credential Guard	Restricted Admin mode
Multi-hop	From the remote desktop, **you can connect through Remote Desktop to another computer**	From the remote desktop, you **can connect through Remote Desktop to another computer**.	Not allowed for user as the session is running as a local host account
Supported authentication	Any negotiable protocol.	Kerberos only.	Any negotiable protocol

For further technical information, see Remote Desktop Protocol and How Kerberos works.

Remote Desktop connections and helpdesk support scenarios

For helpdesk support scenarios in which personnel require administrative access to provide remote assistance to computer users via Remote Desktop sessions, Microsoft recommends that Windows Defender Remote Credential Guard should not be used in that context. This is because if an RDP session is initiated to a compromised client that an attacker already controls, the attacker could use that open channel to create sessions on the user's behalf (without compromising credentials) to access any of the user's resources for a limited time (a few hours) after the session disconnects.

Therefore, we recommend instead that you use the Restricted Admin mode option. For helpdesk support scenarios, RDP connections should only be initiated using the /RestrictedAdmin switch. This helps ensure that credentials and other user resources are not exposed to compromised remote hosts. For more information, see Mitigating Pass-the-Hash and Other Credential Theft v2.

To further harden security, we also recommend that you implement Local Administrator Password Solution (LAPS), a Group Policy client-side extension (CSE) introduced in Windows 8.1 that automates local administrator password management. LAPS mitigates the risk of lateral escalation and other cyberattacks facilitated when customers use the same administrative local account and password combination on all their computers. You can download and install LAPS here.

For further information on LAPS, see Microsoft Security Advisory 3062591.

Remote Credential Guard requirements

To use Windows Defender Remote Credential Guard, the Remote Desktop client and remote host must meet the following requirements:

The Remote Desktop client device:

- Must be running at least Windows 10, version 1703 to be able to supply credentials, which is sent to the remote device. This allows users to run as different users without having to send credentials to the remote machine.
- Must be running at least Windows 10, version 1607 or Windows Server 2016 to use the user's signed-in credentials. This requires the user's account be able to sign in to both the client device and the remote host.
- Must be running the Remote Desktop Classic Windows application. The Remote Desktop Universal Windows Platform application doesn't support Windows Defender Remote Credential Guard.
- Must use Kerberos authentication to connect to the remote host. If the client cannot connect to a domain controller, then RDP attempts to fall back to NTLM. Windows Defender Remote Credential Guard does not allow NTLM fallback because this would expose credentials to risk.

The Remote Desktop remote host:

- Must be running at least Windows 10, version 1607 or Windows Server 2016.
- Must allow Restricted Admin connections.
- Must allow the client's domain user to access Remote Desktop connections.
- Must allow delegation of non-exportable credentials.

There are no hardware requirements for Windows Defender Remote Credential Guard.

> **Note**
>
> Remote Desktop client devices running earlier versions, at minimum Windows 10 version 1607, only support signed-in credentials, so the client device must also be joined to an Active Directory domain. Both Remote Desktop client and server must either be joined to the same domain, or the Remote Desktop server can be joined to a domain that has a trust relationship to the client device's domain.

GPO **Remote host allows delegation of non-exportable credentials** should be enabled for delegation of non-exportable credentials.

- For Windows Defender Remote Credential Guard to be supported, the user must authenticate to the remote host using Kerberos authentication.
- The remote host must be running at least Windows 10 version 1607, or Windows Server 2016.
- The Remote Desktop classic Windows app is required. The Remote Desktop Universal Windows Platform app doesn't support Windows Defender Remote Credential Guard.

Enable Windows Defender Remote Credential Guard

You must enable Restricted Admin or Windows Defender Remote Credential Guard on the remote host by using the Registry.

1. Open Registry Editor on the remote host.
2. Enable Restricted Admin and Windows Defender Remote Credential Guard:
    - Go to HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa.
    - Add a new DWORD value named **DisableRestrictedAdmin**.
    - To turn on Restricted Admin and Windows Defender Remote Credential Guard, set the value of this registry setting to 0.
3. Close Registry Editor.

You can add this by running the following command from an elevated command prompt:

```
Console

reg add HKLM\SYSTEM\CurrentControlSet\Control\Lsa /v DisableRestrictedAdmin
/d 0 /t REG_DWORD
```

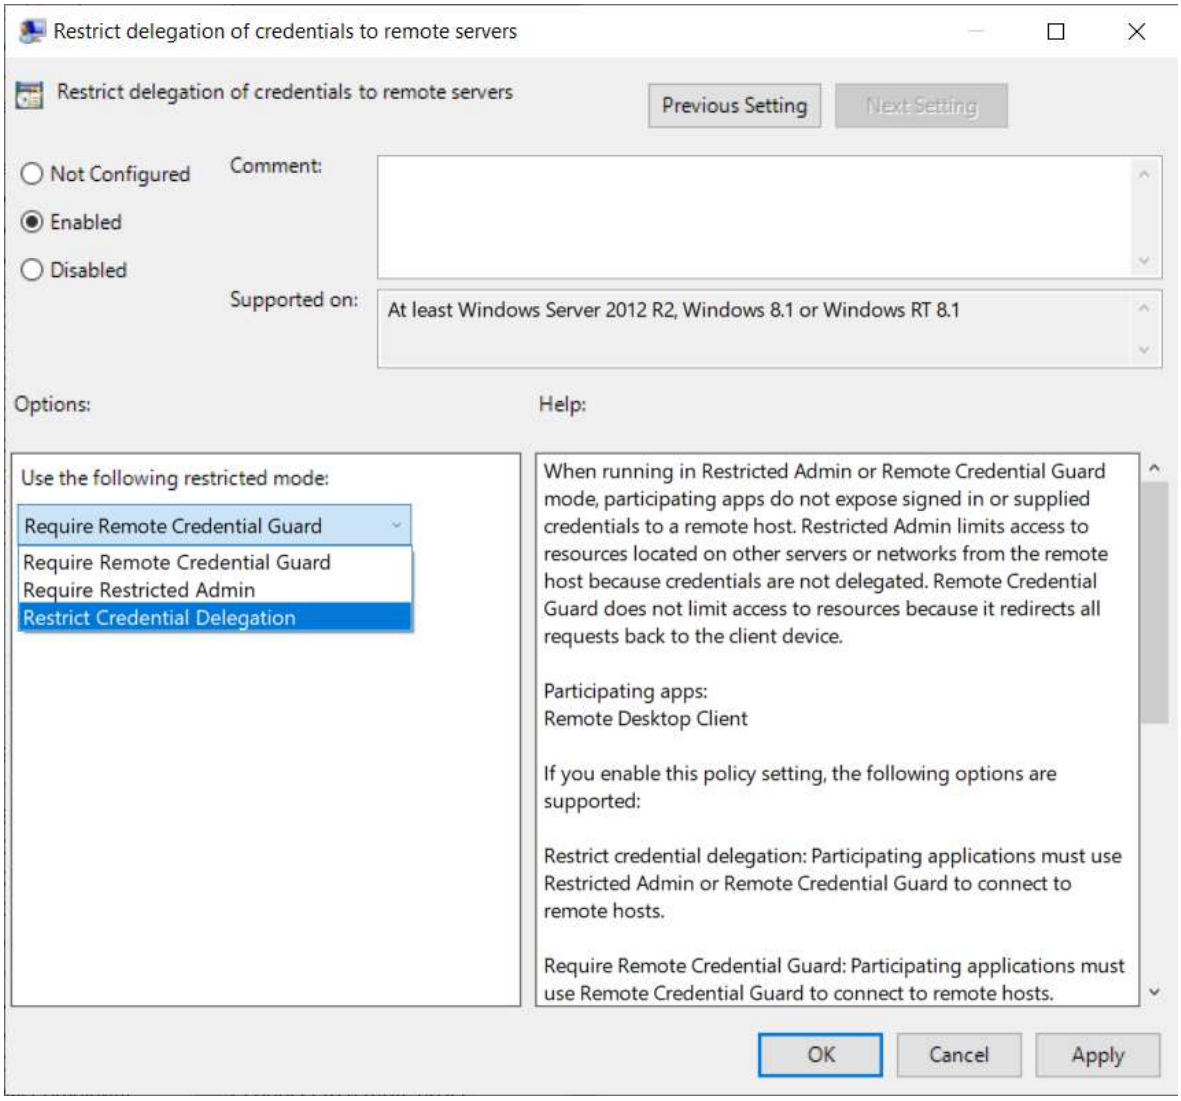
Using Windows Defender Remote Credential Guard

Beginning with Windows 10 version 1703, you can enable Windows Defender Remote Credential Guard on the client device either by using Group Policy or by using a parameter with the Remote Desktop Connection.

Turn on Windows Defender Remote Credential Guard by using Group Policy

1. From the Group Policy Management Console, go to **Computer Configuration** -> **Administrative Templates** -> **System** -> **Credentials Delegation**.

2. Double-click **Restrict delegation of credentials to remote servers**.



3. Under **Use the following restricted mode**:

   - If you want to require either Restricted Admin mode or Windows Defender Remote Credential Guard, choose **Restrict Credential Delegation**. In this configuration, Windows Defender Remote Credential Guard is preferred, but it will use Restricted Admin mode (if supported) when Windows Defender Remote Credential Guard cannot be used.

> **Note**
>
> Neither Windows Defender Remote Credential Guard nor Restricted Admin mode will send credentials in clear text to the Remote Desktop server. When **Restrict Credential Delegation** is enabled, the /restrictedAdmin switch will be ignored. Windows will enforce the policy configuration instead and will use Windows Defender Remote Credential Guard.

- If you want to require Windows Defender Remote Credential Guard, choose **Require Remote Credential Guard**. With this setting, a Remote Desktop connection will succeed only if the remote computer meets the requirements listed earlier in this topic.
- If you want to require Restricted Admin mode, choose **Require Restricted Admin**. For information about Restricted Admin mode, see the table in Comparing Windows Defender Remote Credential Guard with other Remote Desktop connection options, earlier in this topic.

4. Click **OK**.
5. Close the Group Policy Management Console.
6. From a command prompt, run **gpupdate.exe /force** to ensure that the Group Policy object is applied.

Use Windows Defender Remote Credential Guard with a parameter to Remote Desktop Connection

If you don't use Group Policy in your organization, or if not all your remote hosts support Remote Credential Guard, you can add the remoteGuard parameter when you start Remote Desktop Connection to turn on Windows Defender Remote Credential Guard for that connection.

Console

```
mstsc.exe /remoteGuard
```

> **Note**

The user must be authorized to connect to the remote server using Remote Desktop Protocol, for example by being a member of the Remote Desktop Users local group on the remote computer.

Considerations when using Windows Defender Remote Credential Guard

- Windows Defender Remote Credential Guard does not support compound authentication. For example, if you're trying to access a file server from a remote host that requires a device claim, access will be denied.
- Windows Defender Remote Credential Guard can be used only when connecting to a device that is joined to a Windows Server Active Directory domain, including AD domain-joined servers that run as Azure virtual machines (VMs). Windows Defender Remote Credential Guard cannot be used when connecting to remote devices joined to Azure Active Directory.
- Remote Desktop Credential Guard only works with the RDP protocol.
- No credentials are sent to the target device, but the target device still acquires Kerberos Service Tickets on its own.
- The server and client must authenticate using Kerberos.

Configuring Additional LSA Protection

Article • 09/30/2022 • 11 minutes to read

> Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016

This article for the IT professional explains how to configure additional protection for the Local Security Authority (LSA) process to prevent code injection that could compromise credentials.

The LSA, which includes the Local Security Authority Server Service (LSASS) process, validates users for local and remote sign-ins and enforces local security policies. The Windows 8.1 operating system and later provides additional protection for the LSA to prevent reading memory and code injection by non-protected processes. This feature provides added security for the credentials that LSA stores and manages. The protected process setting for LSA can be configured in Windows 8.1 and later. When this setting is used with UEFI lock and Secure Boot, additional protection is achieved because disabling the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa registry key has no effect.

Protected process requirements for plug-ins or drivers

For an LSA plug-in or driver to successfully load as a protected process, it must meet the following criteria:

1. Signature verification

   Protected mode requires that any plug-in that is loaded into the LSA is digitally signed with a Microsoft signature. Therefore, any plug-ins that are unsigned or aren't signed with a Microsoft signature will fail to load in LSA. Examples of these plug-ins are smart card drivers, cryptographic plug-ins, and password filters.

   LSA plug-ins that are drivers, such as smart card drivers, need to be signed by using the WHQL Certification. For more information, see WHQL Release Signature.

   LSA plug-ins that don't have a WHQL Certification process, must be signed by using the file signing service for LSA.

2. Adherence to the Microsoft Security Development Lifecycle (SDL) process guidance

All of the plug-ins must conform to the applicable SDL process guidance. For more information, see the Microsoft Security Development Lifecycle (SDL) Appendix.

Even if the plug-ins are properly signed with a Microsoft signature, non-compliance with the SDL process can result in failure to load a plug-in.

Recommended practices

Use the following list to thoroughly test that LSA protection is enabled before you broadly deploy the feature:

- Identify all of the LSA plug-ins and drivers that are in use within your organization. Include non-Microsoft drivers or plug-ins such as smart card drivers and cryptographic plug-ins, and any internally developed software that is used to enforce password filters or password change notifications.
- Ensure that all of the LSA plug-ins are digitally signed with a Microsoft certificate so that the plug-in won't fail to load.
- Ensure that all of the correctly signed plug-ins can successfully load into LSA and that they perform as expected.
- Use the audit logs to identify LSA plug-ins and drivers that fail to run as a protected process.

Limitations introduced with enabled LSA protection

If additional LSA protection is enabled, you can't debug a custom LSA plugin. You can't attach a debugger to LSASS when it's a protected process. In general, there's no supported way to debug a running protected process.

Auditing to identify LSA plug-ins and drivers that fail to run as a protected process

You can use the audit mode to identify LSA plug-ins and drivers that will fail to load in LSA Protection mode. While in the audit mode, the system will generate event logs, identifying all of the plug-ins and drivers that will fail to load under LSA if LSA Protection is enabled. The messages are logged without blocking the plug-ins or drivers.

The events described in this section are located in the Operational log under Applications and Services Logs\Microsoft\Windows\CodeIntegrity. They can help you identify LSA plug-ins and drivers that are failing to load due to signing reasons. To manage these events, you can use the **wevtutil** command-line tool. For information about this tool, see Wevtutil.

> **Important**
>
> Audit events will not be generated if **Smart App Control** is enabled on a device. To check or change the enablement state of Smart App Control, open the Windows Security Application and go to the App & browser control page. Select the Smart App Control settings link to check the enablement state and change the configuration to Off if you are trying to audit additional LSA protection.

Automatic enablement of Audit mode

Audit mode for additional LSA protection is enabled by default on devices running **Windows 11, 22H2**. If your device is running this build, no additional actions are needed to audit additional LSA protection.

Enable the audit mode for Lsass.exe on a single computer by editing the Registry

1. Open the Registry Editor (RegEdit.exe), and navigate to the registry key that is located at: HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Image File Execution Options\LSASS.exe.
2. Set the value of the registry key to **AuditLevel=dword:00000008**.
3. Restart the computer.

Analyze the results of event 3065 and event 3066.

After these steps, you may see these events in Event Viewer in Applications and Services Logs/Microsoft/Windows/CodeIntegrity:

- **Event 3065**: This event records that a code integrity check determined that a process (usually lsass.exe) attempted to load a particular driver that did not meet the security requirements for Shared Sections. However, due to the system policy that is set, the image was allowed to load.
- **Event 3066**: This event records that a code integrity check determined that a process (usually lsass.exe) attempted to load a particular driver that did not meet the Microsoft signing level requirements. However, due to the system policy that is set, the image was allowed to load.

> **Important**

These operational events are not generated when a kernel debugger is attached and enabled on a system.

If a plug-in or driver contains Shared Sections, Event 3066 is logged with Event 3065. Removing the Shared Sections should prevent both the events from occurring unless the plug-in does not meet the Microsoft signing level requirements.

To enable audit mode for multiple computers in a domain, you can use the Registry Client-Side Extension for Group Policy to deploy the Lsass.exe audit-level registry value. You need to modify HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Image File Execution Options\LSASS.exe registry key.

How to create the AuditLevel value setting in a GPO

1. Open the Group Policy Management Console (GPMC).
2. Create a new Group Policy Object (GPO) that is linked at the domain level or that is linked to the organizational unit that contains your computer accounts. Or you can select a GPO that is already deployed.
3. Right-click the GPO, and then select **Edit** to open the Group Policy Management Editor.
4. Expand **Computer Configuration**, expand **Preferences**, and then expand **Windows Settings**.
5. Right-click **Registry**, point to **New**, and then select **Registry Item**. The **New Registry Properties** dialog box appears.
6. In the **Hive** list, select **HKEY_LOCAL_MACHINE.**
7. In the **Key Path** list, browse to **SOFTWARE\Microsoft\Windows NT\CurrentVersion\Image File Execution Options\LSASS.exe**.
8. In the **Value name** box, type **AuditLevel**.
9. In the **Value type** box, select **REG_DWORD**.
10. In the **Value data** box, type **00000008**.
11. Select **OK**.

Note

For the GPO take effect, the GPO change must be replicated to all domain controllers in the domain.

To opt in for additional LSA protection on multiple computers, you can use the Registry Client-Side Extension for Group Policy by modifying HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa. For instructions, see How to configure additional LSA protection of credentials in this article.

After opting in: How to identify plug-ins and drivers loaded by the lsass.exe

You can use the event log to identify LSA plug-ins and drivers that failed to load in LSA protection mode. When the LSA protected process is enabled, the system generates event logs that identify all of the plug-ins and drivers that failed to load under LSA.

You may see these events in Event Viewer: Microsoft-Windows-Codeintegrity/Operational:

- **Event 3033**: This event records that a code integrity check determined that a process (usually lsass.exe) attempted to load a driver that did not meet the Microsoft signing level requirements.
- **Event 3063**: This event records that a code integrity check determined that a process (usually lsass.exe) attempted to load a driver that did not meet the security requirements for Shared Sections.

Shared Sections are typically the result of programming techniques that allow instance data to interact with other processes that use the same security context. This can create security vulnerabilities.

How to configure additional LSA protection of credentials

On devices running Windows 8.1 or later, configuration is possible by performing the procedures described in this section.

On x86-based or x64-based devices using Secure Boot and UEFI or not

On x86-based or x64-based devices that use Secure Boot or UEFI, a UEFI variable can be set in the UEFI firmware when LSA protection is enabled by using registry key or policy. When the setting is stored in the firmware, the UEFI variable can't be deleted or changed by modifying the registry or by policy used to enable additional LSA protection. The UEFI variable must be reset using the instructions below on how to disable LSA protection.

x86-based or x64-based devices that don't support UEFI or where Secure Boot are disabled can't store the configuration for LSA protection in the firmware. These devices rely solely on the presence of the registry key. In this scenario, it's possible to disable LSA protection by using remote access to the device. Disablement of LSA protection won't take effect until the device reboots.

Automatic Enablement

For devices running Windows RT 8.1, additional LSA protection is always enabled, and it can't be turned off.

For client devices running Windows 11, 22H2, additional LSA protection will be enabled by default if the following criteria are met:

- The device is a new install of Windows 11, 22H2 (not upgraded from previous release).
- The device is enterprise joined (Active Directory domain joined, Azure AD domain joined, or hybrid Azure AD domain joined).
- The device is capable of Hypervisor-protected code integrity (HVCI)

Automatic enablement of additional LSA protection on Windows 11, 22H2 doesn't set a UEFI variable for the feature. If you want to set a UEFI variable, you can use a registry configuration or policy.

How to enable LSA protection on a single computer

Using the Registry

1. Open the Registry Editor (RegEdit.exe), and navigate to the registry key that is located at: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa.

2. Set the value of the registry key to:
   a. "RunAsPPL"=dword:00000001 to configure the feature with a UEFI variable.
   b. "RunAsPPL"=dword:00000002 to configure the feature without a UEFI variable (only on Windows 11, 22H2).

3. Restart the computer.

Using Local Group Policy on Windows 11, 22H2

1. Open Local Group Policy Editor (gpedit.msc)
2. Expand **Computer Configuration**, expand **Administrative Templates**, expand **System**, and then expand **Local Security Authority**.
3. Open the **Configure LSASS to run as a protected process** policy.
4. Set the policy to **Enabled**.
5. Under **Options**, set Configure LSA to run as a protected process to:
   a. "Enabled with UEFI Lock" to configure the feature with a UEFI variable.
   b. "Enabled without UEFI Lock" to configure the feature without a UEFI variable.
6. Restart the computer.

How to enable LSA protection using Group Policy

1. Open the Group Policy Management Console (GPMC).
2. Create a new GPO that is linked at the domain level or that is linked to the organizational unit that contains your computer accounts. Or you can select a GPO that is already deployed.
3. Right-click the GPO, and then select **Edit** to open the Group Policy Management Editor.
4. Expand **Computer Configuration**, expand **Preferences**, and then expand **Windows Settings**.
5. Right-click **Registry**, point to **New**, and then select **Registry Item**. The **New Registry Properties** dialog box appears.
6. In the **Hive** list, select **HKEY_LOCAL_MACHINE**.
7. In the **Key Path** list, browse to **SYSTEM\CurrentControlSet\Control\Lsa**.
8. In the **Value name** box, type **RunAsPPL**.
9. In the **Value type** box, select **REG_DWORD**.
10. In the **Value data** box, type:
    a. **00000001** to enable LSA protection with a UEFI variable.

b. **00000002** to enable LSA protection without a UEFI variable (only enforced on Windows 11, 22H2).

11. Select **OK**.

How to disable LSA protection

How to disable using the Registry

1. Open the Registry Editor (RegEdit.exe), and navigate to the registry key that is located at: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa.
2. Set "RunAsPPL"=dword:00000000 or delete the DWORD.
3. If PPL was enabled with a UEFI variable, use the Local Security Authority Protected Process Opt-out tool to remove the UEFI variable.
4. Restart the computer.

How to disable using local policy on Windows 11, 22H2

Open Local Group Policy Editor (gpedit.msc)

1. Expand **Computer Configuration**, expand **Administrative Templates**, expand **System**, and then expand **Local Security Authority**.
2. Open the **Configure LSASS to run as a protected process** policy.
3. Set the policy to **Enabled**.
4. Under **Options**, set Configure LSA to "Disabled"
5. Restart the computer.

If you set this policy to **Not Configured** and the policy was previously enabled, the previous setting doesn't get cleaned up and will continue to be enforced. You must set the policy to **Disabled** under the **Configure LSA to run as a protected process** drop-down to disable the feature.

How to remove the LSA protection UEFI variable

Use the Local Security Authority (LSA) Protected Process Opt-out tool to delete the UEFI variable if the device is using Secure Boot.

For more information about the opt-out tool, see Download Local Security Authority (LSA) Protected Process Opt-out from Official Microsoft Download Center.

For more information about managing Secure Boot, see UEFI Firmware.

> ⚠ **Warning**
>
> When Secure Boot is turned off, all the Secure Boot and UEFI-related configurations are reset. You should turn off Secure Boot only when all other means to disable LSA protection have failed.

Verifying LSA protection

To discover if LSA was started in protected mode when Windows started, search for the following WinInit event in the **System** log under **Windows Logs**:

- 12: LSASS.exe was started as a protected process with level: 4

Additional resources

Credentials Protection and Management

File signing service for LSA

Technical support policy for lost or forgotten passwords

Article • 11/08/2022 • 2 minutes to read

Microsoft takes security seriously. This is for your protection. Microsoft accounts, the Windows operating system, and other Microsoft products include passwords to help secure your information. This article provides some options that you can use to reset or recover your password if you forget it. If these options don't work, Microsoft support engineers can't help you retrieve or circumvent a lost or forgotten password.

If you lose or forget a password, you can use the links in this article to find published support information that will help you reset the password.

How to reset a password for a domain account

If you lose or forget the password for a domain account, contact your IT administrator or Helpdesk. For more information, see Change or reset your Windows password.

How to reset a password for a Microsoft account

If you lose or forget the password for your Microsoft Account, use the Recover your account wizard.

This wizard requests your security proofs. If you've forgotten your security proofs, or no longer have access to them, select **I no longer have these anymore**. After you select this option, fill out a form for the Microsoft Account team. Provide as much information as you can on this form. The Microsoft Account team reviews the information that you provide to determine whether you're the account holder. This decision is final. Microsoft doesn't influence the team's choice of action.

How to reset a password for a local account on a Windows device

Local accounts on a device include the device's Administrator account.

Windows 10

If you lose or forget the password for a local account on a device that runs Windows 10, see Reset your Windows 10 local account password.

Windows 8.1 or Windows 7

If you lose or forget the password for a local account on a device that runs Windows 8.1 or Windows 7, see Change or reset your Windows password. In that article, you can select your operating system version from the **Select Product Version** menu.

How to reset a hardware BIOS password

If you lose or forget the password for the hardware BIOS of a device, contact the device manufacturer for help and support. If you do contact the manufacturer online, make sure that you visit the manufacturer website and not the website of some third party.

How to reset a password for an individual file

Some applications let you password-protect individual files. If you lose or forget such a password, you can rely on that application only to reset or recover it. Microsoft support engineers can't help you reset, retrieve, or circumvent such passwords.

Using third-party password tools

Some third-party companies claim to be able to circumvent passwords that have been applied to files and features that Microsoft programs use. For legal reasons, we can't recommend or endorse any one of these companies. If you want help to circumvent or reset a password, you can locate and contact a third party for this help. However, you use such third-party products and services at your own risk.

Access Control Overview

Article • 02/21/2023 • 6 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

This topic for the IT professional describes access control in Windows, which is the process of authorizing users, groups, and computers to access objects on the network or computer. Key concepts that make up access control are permissions, ownership of objects, inheritance of permissions, user rights, and object auditing.

Feature description

Computers that are running a supported version of Windows can control the use of system and network resources through the interrelated mechanisms of authentication and authorization. After a user is authenticated, the Windows operating system uses built-in authorization and access control technologies to implement the second phase of protecting resources: determining if an authenticated user has the correct permissions to access a resource.

Shared resources are available to users and groups other than the resource's owner, and they need to be protected from unauthorized use. In the access control model, users and groups (also referred to as security principals) are represented by unique security identifiers (SIDs). They are assigned rights and permissions that inform the operating system what each user and group can do. Each resource has an owner who grants permissions to security principals. During the access control check, these permissions are examined to determine which security principals can access the resource and how they can access it.

Security principals perform actions (which include Read, Write, Modify, or Full control) on objects. Objects include files, folders, printers, registry keys, and Active Directory Domain Services (AD DS) objects. Shared resources use access control lists (ACLs) to assign permissions. This enables resource managers to enforce access control in the following ways:

- Deny access to unauthorized users and groups
- Set well-defined limits on the access that is provided to authorized users and groups

Object owners generally grant permissions to security groups rather than to individual users. Users and computers that are added to existing groups assume the permissions of that group. If an object (such as a folder) can hold other objects (such as subfolders and files), it is called a container. In a hierarchy of objects, the relationship between a

container and its content is expressed by referring to the container as the parent. An object in the container is referred to as the child, and the child inherits the access control settings of the parent. Object owners often define permissions for container objects, rather than individual child objects, to ease access control management.

This content set contains:

- Dynamic Access Control Overview
- Security identifiers
- Security Principals
  - Local Accounts
  - Active Directory Accounts
  - Microsoft Accounts
  - Service Accounts
  - Active Directory Security Groups

Practical applications

Administrators who use the supported version of Windows can refine the application and management of access control to objects and subjects to provide the following security:

- Protect a greater number and variety of network resources from misuse.
- Provision users to access resources in a manner that is consistent with organizational policies and the requirements of their jobs.
- Enable users to access resources from a variety of devices in numerous locations.
- Update users' ability to access resources on a regular basis as an organization's policies change or as users' jobs change.
- Account for a growing number of use scenarios (such as access from remote locations or from a rapidly expanding variety of devices, such as tablet computers and mobile phones).
- Identify and resolve access issues when legitimate users are unable to access resources that they need to perform their jobs.

Permissions

Permissions define the type of access that is granted to a user or group for an object or object property. For example, the Finance group can be granted Read and Write permissions for a file named Payroll.dat.

By using the access control user interface, you can set NTFS permissions for objects such as files, Active Directory objects, registry objects, or system objects such as processes. Permissions can be granted to any user, group, or computer. It is a good practice to assign permissions to groups because it improves system performance when verifying access to an object.

For any object, you can grant permissions to:

- Groups, users, and other objects with security identifiers in the domain.
- Groups and users in that domain and any trusted domains.
- Local groups and users on the computer where the object resides.

The permissions attached to an object depend on the type of object. For example, the permissions that can be attached to a file are different from those that can be attached to a registry key. Some permissions, however, are common to most types of objects. These common permissions are:

- Read
- Modify
- Change owner
- Delete

When you set permissions, you specify the level of access for groups and users. For example, you can let one user read the contents of a file, let another user make changes to the file, and prevent all other users from accessing the file. You can set similar permissions on printers so that certain users can configure the printer and other users can only print.

When you need to change the permissions on a file, you can run Windows Explorer, right-click the file name, and click **Properties**. On the **Security** tab, you can change permissions on the file. For more information, see Managing Permissions.

> ⓘ **Note**
>
> Another kind of permissions, called share permissions, is set on the Sharing tab of a folder's **Properties** page or by using the Shared Folder Wizard. For more information see **Share and NTFS Permissions on a File Server**.

Ownership of objects

An owner is assigned to an object when that object is created. By default, the owner is the creator of the object. No matter what permissions are set on an object, the owner of

the object can always change the permissions. For more information, see Manage Object Ownership.

Inheritance of permissions

Inheritance allows administrators to easily assign and manage permissions. This feature automatically causes objects within a container to inherit all the inheritable permissions of that container. For example, the files within a folder inherit the permissions of the folder. Only permissions marked to be inherited will be inherited.

User rights

User rights grant specific privileges and sign-in rights to users and groups in your computing environment. Administrators can assign specific rights to group accounts or to individual user accounts. These rights authorize users to perform specific actions, such as signing in to a system interactively or backing up files and directories.

User rights are different from permissions because user rights apply to user accounts, and permissions are associated with objects. Although user rights can apply to individual user accounts, user rights are best administered on a group account basis. There is no support in the access control user interface to grant user rights. However, user rights assignment can be administered through **Local Security Settings**.

For more information about user rights, see User Rights Assignment.

Object auditing

With administrator's rights, you can audit users' successful or failed access to objects. You can select which object access to audit by using the access control user interface, but first you must enable the audit policy by selecting **Audit object access** under **Local Policies** in **Local Security Settings**. You can then view these security-related events in the Security log in Event Viewer.

For more information about auditing, see Security Auditing Overview.

See also

- For more information about access control and authorization, see Access Control and Authorization Overview.

Local Accounts

Article • 02/17/2023 • 18 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

This article describes the default local user accounts for Windows operating systems, and how to manage the built-in accounts.

About local user accounts

Local user accounts are stored locally on the device. These accounts can be assigned rights and permissions on a particular device, but on that device only. Local user accounts are security principals that are used to secure and manage access to the resources on a device, for services or users.

Default local user accounts

The *default local user accounts* are built-in accounts that are created automatically when the operating system is installed. The default local user accounts can't be removed or deleted and don't provide access to network resources.

Default local user accounts are used to manage access to the local device's resources based on the rights and permissions that are assigned to the account. The default local user accounts, and the local user accounts that you create, are located in the *Users* folder. The Users folder is located in the Local Users and Groups folder in the local *Computer Management* Microsoft Management Console (MMC). *Computer Management* is a collection of administrative tools that you can use to manage a local or remote device.

Default local user accounts are described in the following sections. Expand each section for more information.

▼ **Administrator**

The default local Administrator account is a user account for system administration. Every computer has an Administrator account (SID S-1-5-*domain*-500, display name Administrator). The Administrator account is the first account that is created during the Windows installation.

The Administrator account has full control of the files, directories, services, and other resources on the local device. The Administrator account can create other local users,

assign user rights, and assign permissions. The Administrator account can take control of local resources at any time by changing the user rights and permissions.

The default Administrator account can't be deleted or locked out, but it can be renamed or disabled.

Windows setup disables the built-in Administrator account and creates another local account that is a member of the Administrators group.

Members of the Administrators groups can run apps with elevated permissions without using the *Run as Administrator* option. Fast User Switching is more secure than using `runas` or different-user elevation.

Account group membership

By default, the Administrator account is a member of the Administrators group. It's a best practice to limit the number of users in the Administrators group because members of the Administrators group have Full Control permissions on the device.

The Administrator account can't be removed from the Administrators group.

Security considerations

Because the Administrator account is known to exist on many versions of the Windows operating system, it's a best practice to disable the Administrator account when possible to make it more difficult for malicious users to gain access to the server or client computer.

You can rename the Administrator account. However, a renamed Administrator account continues to use the same automatically assigned security identifier (SID), which can be discovered by malicious users. For more information about how to rename or disable a user account, see Disable or activate a local user account and Rename a local user account.

As a security best practice, use your local (non-Administrator) account to sign in and then use **Run as administrator** to accomplish tasks that require a higher level of rights than a standard user account. Don't use the Administrator account to sign in to your computer unless it's entirely necessary. For more information, see Run a program with administrative credentials.

Group Policy can be used to control the use of the local Administrators group automatically. For more information about Group Policy, see Group Policy Overview.

> **ⓘ Important**

- Blank passwords are not allowed in the versions designated in the **Applies To** list at the beginning of this topic.
- Even when the Administrator account has been disabled, it can still be used to gain access to a computer by using safe mode. In the Recovery Console or in safe mode, the Administrator account is automatically enabled. When normal operations are resumed, it is disabled.

▼ Guest

The Guest account lets occasional or one-time users, who don't have an account on the computer, temporarily sign in to the local server or client computer with limited user rights. By default, the Guest account is disabled and has a blank password. Since the Guest account can provide anonymous access, it's considered a security risk. For this reason, it's a best practice to leave the Guest account disabled, unless its use is necessary.

Account group membership

By default, the Guest account is the only member of the default Guests group (SID S-1-5-32-546), which lets a user sign in to a device.

Security considerations

When enabling the Guest account, only grant limited rights and permissions. For security reasons, the Guest account shouldn't be used over the network and made accessible to other computers.

In addition, the guest user in the Guest account shouldn't be able to view the event logs. After the Guest account is enabled, it's a best practice to monitor the Guest account frequently to ensure that other users can't use services and other resources. This includes resources that were unintentionally left available by a previous user.

▼ HelpAssistant

The HelpAssistant account is a default local account that is enabled when a Remote Assistance session is run. This account is automatically disabled when no Remote Assistance requests are pending.

HelpAssistant is the primary account that is used to establish a Remote Assistance session. The Remote Assistance session is used to connect to another computer running the Windows operating system, and it's initiated by invitation. For solicited remote

assistance, a user sends an invitation from their computer, through e-mail or as a file, to a person who can provide assistance. After the user's invitation for a Remote Assistance session is accepted, the default HelpAssistant account is automatically created to give the person who provides assistance limited access to the computer. The HelpAssistant account is managed by the Remote Desktop Help Session Manager service.

Security considerations

The SIDs that pertain to the default HelpAssistant account include:

- SID: `S-1-5-<domain>-13`, display name Terminal Server User. This group includes all users who sign in to a server with Remote Desktop Services enabled. Note: In Windows Server 2008, Remote Desktop Services is called Terminal Services.
- SID: `S-1-5-<domain>-14`, display name Remote Interactive Logon. This group includes all users who connect to the computer by using a remote desktop connection. This group is a subset of the Interactive group. Access tokens that contain the Remote Interactive Logon SID also contain the Interactive SID.

For the Windows Server operating system, Remote Assistance is an optional component that isn't installed by default. You must install Remote Assistance before it can be used.

For details about the HelpAssistant account attributes, see the following table.

HelpAssistant account attributes

Attribute	Value
Well-Known SID/RID	`S-1-5-<domain>-13 (Terminal Server User), S-1-5-<domain>-14 (Remote Interactive Logon)`
Type	User
Default container	`CN=Users, DC=<domain>`
Default members	None
Default member of	Domain Guests Guests
Protected by ADMINSDHOLDER?	No
Safe to move out of default container?	Can be moved out, but we don't recommend it.
Safe to delegate management of this group to non-Service admins?	No

▼ **DefaultAccount**

The DefaultAccount account, also known as the Default System Managed Account (DSMA), is a well-known user account type. DefaultAccount can be used to run processes that are either multi-user aware or user-agnostic.

The DSMA is disabled by default on the desktop SKUs and on the Server operating systems with the desktop experience.

The DSMA has a well-known RID of `503`. The security identifier (SID) of the DSMA will thus have a well-known SID in the following format: `S-1-5-21-\<ComputerIdentifier>-503`.

The DSMA is a member of the well-known group **System Managed Accounts Group**, which has a well-known SID of `S-1-5-32-581`.

The DSMA alias can be granted access to resources during offline staging even before the account itself has been created. The account and the group are created during first boot of the machine within the Security Accounts Manager (SAM).

How Windows uses the DefaultAccount

From a permission perspective, the DefaultAccount is a standard user account. The DefaultAccount is needed to run multi-user-manifested-apps (MUMA apps). MUMA apps run all the time and react to users signing in and signing out of the devices. Unlike Windows Desktop where apps run in context of the user and get terminated when the user signs off, MUMA apps run by using the DSMA.

MUMA apps are functional in shared session SKUs such as Xbox. For example, Xbox shell is a MUMA app. Today, Xbox automatically signs in as Guest account and all apps run in this context. All the apps are multi-user-aware and respond to events fired by user manager. The apps run as the Guest account.

Similarly, Phone auto logs in as a *DefApps* account, which is akin to the standard user account in Windows but with a few extra privileges. Brokers, some services and apps run as this account.

In the converged user model, the multi-user-aware apps and multi-user-aware brokers will need to run in a context different from that of the users. For this purpose, the system creates DSMA.

How the DefaultAccount gets created on domain controllers

If the domain was created with domain controllers running Windows Server 2016, the DefaultAccount will exist on all domain controllers in the domain. If the domain was created with domain controllers running an earlier version of Windows Server, the DefaultAccount will be created after the PDC Emulator role is transferred to a domain controller that runs Windows Server 2016. The DefaultAccount will then be replicated to all other domain controllers in the domain.

Recommendations for managing the Default Account (DSMA)

Microsoft doesn't recommend changing the default configuration, where the account is disabled. There's no security risk with having the account in the disabled state. Changing the default configuration could hinder future scenarios that rely on this account.

Default local system accounts

▼ **SYSTEM**

The *SYSTEM* account is used by the operating system and by services running under Windows. There are many services and processes in the Windows operating system that need the capability to sign in internally, such as during a Windows installation. The SYSTEM account was designed for that purpose, and Windows manages the SYSTEM account's user rights. It's an internal account that doesn't show up in User Manager, and it can't be added to any groups.

On the other hand, the SYSTEM account does appear on an NTFS file system volume in File Manager in the **Permissions** portion of the **Security** menu. By default, the SYSTEM account is granted Full Control permissions to all files on an NTFS volume. Here the SYSTEM account has the same functional rights and permissions as the Administrator account.

> ⓘ **Note**
>
> To grant the account Administrators group file permissions does not implicitly give permission to the SYSTEM account. The SYSTEM account's permissions can be removed from a file, but we do not recommend removing them.

▼ **NETWORK SERVICE**

The NETWORK SERVICE account is a predefined local account used by the service control manager (SCM). A service that runs in the context of the NETWORK SERVICE

account presents the computer's credentials to remote servers. For more information, see NetworkService Account.

▼ LOCAL SERVICE

The LOCAL SERVICE account is a predefined local account used by the service control manager. It has minimum privileges on the local computer and presents anonymous credentials on the network. For more information, see LocalService Account.

How to manage local user accounts

The default local user accounts, and the local user accounts you create, are located in the Users folder. The Users folder is located in Local Users and Groups. For more information about creating and managing local user accounts, see Manage Local Users.

You can use Local Users and Groups to assign rights and permissions on only the local server to limit the ability of local users and groups to perform certain actions. A right authorizes a user to perform certain actions on a server, such as backing up files and folders or shutting down a server. An access permission is a rule that is associated with an object, usually a file, folder, or printer. It regulates which users can have access to an object on the server and in what manner.

You can't use Local Users and Groups on a domain controller. However, you can use Local Users and Groups on a domain controller to target remote computers that aren't domain controllers on the network.

> ⓘ **Note**
>
> You use Active Directory Users and Computers to manage users and groups in Active Directory.

You can also manage local users by using NET.EXE USER and manage local groups by using NET.EXE LOCALGROUP, or by using various PowerShell cmdlets and other scripting technologies.

Restrict and protect local accounts with administrative rights

An administrator can use many approaches to prevent malicious users from using stolen credentials such as a stolen password or password hash, for a local account on one

computer from being used to authenticate on another computer with administrative rights. This is also called "lateral movement".

The simplest approach is to sign in to your computer with a standard user account, instead of using the Administrator account for tasks. For example, use a standard account to browse the Internet, send email, or use a word processor. When you want to perform administrative tasks such as installing a new program or changing a setting that affects other users, you don't have to switch to an Administrator account. You can use User Account Control (UAC) to prompt you for permission or an administrator password before performing the task, as described in the next section.

The other approaches that can be used to restrict and protect user accounts with administrative rights include:

- Enforce local account restrictions for remote access.
- Deny network logon to all local Administrator accounts.
- Create unique passwords for local accounts with administrative rights.

Each of these approaches is described in the following sections.

> ⓘ **Note**
>
> These approaches do not apply if all administrative local accounts are disabled.

Enforce local account restrictions for remote access

User Account Control (UAC) is a security feature that informs you when a program makes a change that requires administrative permissions. UAC works by adjusting the permission level of your user account. By default, UAC is set to notify you when applications try to make changes to your computer, but you can change when UAC notifies you.

UAC makes it possible for an account with administrative rights to be treated as a standard user non-administrator account until full rights, also called elevation, is requested and approved. For example, UAC lets an administrator enter credentials during a non-administrator's user session to perform occasional administrative tasks without having to switch users, sign out, or use the **Run as** command.

In addition, UAC can require administrators to specifically approve applications that make system-wide changes before those applications are granted permission to run, even in the administrator's user session.

For example, a default feature of UAC is shown when a local account signs in from a remote computer by using Network logon (for example, by using NET.EXE USE). In this instance, it's issued a standard user token with no administrative rights, but without the ability to request or receive elevation. Consequently, local accounts that sign in by using Network logon can't access administrative shares such as C$, or ADMIN$, or perform any remote administration.

For more information about UAC, see User Account Control.

The following table shows the Group Policy and registry settings that are used to enforce local account restrictions for remote access.

No.	Setting	Detailed Description
	Policy location	Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options
1	Policy name	User Account Control: Admin Approval Mode for the Built-in Administrator account
	Policy setting	Enabled
2	Policy location	Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options
	Policy name	User Account Control: Run all administrators in Admin Approval Mode
	Policy setting	Enabled
3	Registry key	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System
	Registry value name	LocalAccountTokenFilterPolicy
	Registry value type	DWORD
	Registry value data	0

Note

You can also enforce the default for LocalAccountTokenFilterPolicy by using the custom ADMX in Security Templates.

To enforce local account restrictions for remote access

1. Start the **Group Policy Management** Console (GPMC)
2. In the console tree, expand <*Forest*>\Domains\<*Domain*>, and then **Group Policy Objects** where *forest* is the name of the forest, and *domain* is the name of the domain where you want to set the Group Policy Object (GPO)
3. In the console tree, right-click **Group Policy Objects** > **New**
4. In the **New GPO** dialog box, type <**gpo_name**>, and > **OK** where *gpo_name* is the name of the new GPO. The GPO name indicates that the GPO is used to restrict local administrator rights from being carried over to another computer
5. In the details pane, right-click <**gpo_name**>, and > **Edit**
6. Ensure that UAC is enabled and that UAC restrictions apply to the default Administrator account by following these steps:

- Navigate to the Computer Configuration\Windows Settings\Security Settings\Local Policies\, and > **Security Options**
- Double-click **User Account Control: Run all administrators in Admin Approval Mode** > **Enabled** > **OK**
- Double-click **User Account Control: Admin Approval Mode for the Built-in Administrator account** > **Enabled** > **OK**

1. Ensure that the local account restrictions are applied to network interfaces by following these steps:

- Navigate to *Computer Configuration**Preferences and Windows Settings*, and > **Registry**
- Right-click **Registry**, and > **New** > **Registry Item**
- In the **New Registry Properties** dialog box, on the **General** tab, change the setting in the **Action** box to **Replace**
- Ensure that the **Hive** box is set to **HKEY_LOCAL_MACHINE**
- Select (**...**), browse to the following location for **Key Path** > **Select** for: `SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System`
- In the **Value name** area, type `LocalAccountTokenFilterPolicy`
- In the **Value type** box, from the drop-down list, select **REG_DWORD** to change the value
- In the **Value data** box, ensure that the value is set to **0**
- Verify this configuration, and > **OK**

1. Link the GPO to the first **Workstations** organizational unit (OU) by doing the following:

- Navigate to the `*Forest*\<Domains>\*Domain*\*OU*` path
- Right-click the **Workstations** > **Link an existing GPO**
- Select the GPO that you created, and > **OK**

1. Test the functionality of enterprise applications on the workstations in that first OU and resolve any issues caused by the new policy
2. Create links to all other OUs that contain workstations
3. Create links to all other OUs that contain servers

Deny network logon to all local Administrator accounts

Denying local accounts the ability to perform network logons can help prevent a local account password hash from being reused in a malicious attack. This procedure helps to prevent lateral movement by ensuring that stolen credentials for local accounts from a compromised operating system can't be used to compromise other computers that use the same credentials.

> ⓘ **Note**
>
> To perform this procedure, you must first identify the name of the local, default Administrator account, which might not be the default user name "Administrator", and any other accounts that are members of the local Administrators group.

The following table shows the Group Policy settings that are used to deny network logon for all local Administrator accounts.

No.	Setting	Detailed Description
	Policy location	Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment
1	Policy name	Deny access to this computer from the network
	Policy setting	Local account and member of Administrators group
2	Policy location	Computer Configuration\Windows Settings\Security Settings\Local Policies\User Rights Assignment

No.	Setting	Detailed Description
	Policy name	Deny log on through Remote Desktop Services
	Policy setting	Local account and member of Administrators group

To deny network logon to all local administrator accounts

1. Start the **Group Policy Management** Console (GPMC)
2. In the console tree, expand <*Forest*>\Domains\<*Domain*>, and then **Group Policy Objects**, where *forest* is the name of the forest, and *domain* is the name of the domain where you want to set the Group Policy Object (GPO)
3. In the console tree, right-click **Group Policy Objects**, and > **New**
4. In the **New GPO** dialog box, type <**gpo_name**>, and then > **OK** where *gpo_name* is the name of the new GPO indicates that it's being used to restrict the local administrative accounts from interactively signing in to the computer
5. In the details pane, right-click <**gpo_name**>, and > **Edit**
6. Configure the user rights to deny network logons for administrative local accounts as follows:
7. Navigate to the Computer Configuration\Windows Settings\Security Settings\, and > **User Rights Assignment**
8. Double-click **Deny access to this computer from the network**
9. Select **Add User or Group**, type **Local account and member of Administrators group**, and > **OK**
10. Configure the user rights to deny Remote Desktop (Remote Interactive) logons for administrative local accounts as follows:
11. Navigate to Computer Configuration\Policies\Windows Settings and Local Policies, and then select **User Rights Assignment**
12. Double-click **Deny log on through Remote Desktop Services**
13. Select **Add User or Group**, type **Local account and member of Administrators group**, and > **OK**
14. Link the GPO to the first **Workstations** OU as follows:

- Navigate to the <*Forest*>\Domains\<*Domain*>\OU path
- Right-click the **Workstations** OU, and > **Link an existing GPO**
- Select the GPO that you created, and > **OK**

1. Test the functionality of enterprise applications on the workstations in that first OU and resolve any issues caused by the new policy
2. Create links to all other OUs that contain workstations

3. Create links to all other OUs that contain servers

> **Note**
>
> You might have to create a separate GPO if the user name of the default Administrator account is different on workstations and servers.

Create unique passwords for local accounts with administrative rights

Passwords should be unique per individual account. While it's true for individual user accounts, many enterprises have identical passwords for common local accounts, such as the default Administrator account. This also occurs when the same passwords are used for local accounts during operating system deployments.

Passwords that are left unchanged or changed synchronously to keep them identical add a significant risk for organizations. Randomizing the passwords mitigates "pass-the-hash" attacks by using different passwords for local accounts, which hamper the ability of malicious users to use password hashes of those accounts to compromise other computers.

Passwords can be randomized by:

- Purchasing and implementing an enterprise tool to accomplish this task. These tools are commonly referred to as "privileged password management" tools
- Configuring Local Administrator Password Solution (LAPS) to accomplish this task
- Creating and implementing a custom script or solution to randomize local account passwords

User Account Control

Article • 02/17/2023 • 2 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

User Account Control (UAC) helps prevent malware from damaging a PC and helps organizations deploy a better-managed desktop. With UAC, apps and tasks always run in the security context of a non-administrator account, unless an administrator specifically authorizes administrator-level access to the system. UAC can block the automatic installation of unauthorized apps and prevent inadvertent changes to system settings.

UAC allows all users to log on to their computers using a standard user account. Processes launched using a standard user token may perform tasks using access rights granted to a standard user. For instance, Windows Explorer automatically inherits standard user level permissions. Additionally, any apps that are started using Windows Explorer (for example, by double-clicking a shortcut) also run with the standard set of user permissions. Many apps, including those that are included with the operating system itself, are designed to work properly in this way.

Other apps, especially those that were not specifically designed with security settings in mind, often require additional permissions to run successfully. These types of apps are referred to as legacy apps. Additionally, actions such as installing new software and making configuration changes to the Windows Firewall, require more permissions than what is available to a standard user account.

When an app needs to run with more than standard user rights, UAC allows users to run apps with their administrator token (with administrative groups and privileges) instead of their default, standard user access token. Users continue to operate in the standard user security context, while enabling certain apps to run with elevated privileges, if needed.

Practical applications

Admin Approval Mode in UAC helps prevent malware from silently installing without an administrator's knowledge. It also helps protect from inadvertent system-wide changes. Lastly, it can be used to enforce a higher level of compliance where administrators must actively consent or provide credentials for each administrative process.

In this section

Topic	Description
How User Account Control works	User Account Control (UAC) is a fundamental component of Microsoft's overall security vision. UAC helps mitigate the impact of malware.
User Account Control security policy settings	You can use security policies to configure how User Account Control works in your organization. They can be configured locally by using the Local Security Policy snap-in (secpol.msc) or configured for the domain, OU, or specific groups by Group Policy.
User Account Control Group Policy and registry key settings	Here's a list of UAC Group Policy and registry key settings that your organization can use to manage UAC.

How User Account Control works

Article • 02/17/2023 • 14 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

User Account Control (UAC) is a fundamental component of Microsoft's overall security vision. UAC helps mitigate the impact of malware.

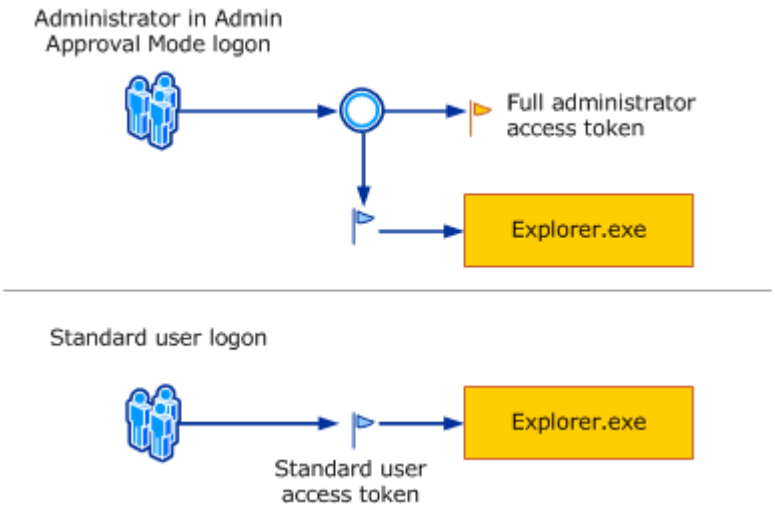
UAC process and interactions

Each app that requires the administrator access token must prompt for consent. The one exception is the relationship that exists between parent and child processes. Child processes inherit the user's access token from the parent process. Both the parent and child processes, however, must have the same integrity level. Windows protects processes by marking their integrity levels. Integrity levels are measurements of trust. A "high" integrity application is one that performs tasks that modify system data, such as a disk partitioning application, while a "low" integrity application is one that performs tasks that could potentially compromise the operating system, such as a Web browser. Apps with lower integrity levels cannot modify data in applications with higher integrity levels. When a standard user attempts to run an app that requires an administrator access token, UAC requires that the user provide valid administrator credentials.

To better understand how this process happens, let's look at the Windows logon process.

Logon process

The following shows how the logon process for an administrator differs from the logon process for a standard user.



By default, standard users and administrators access resources and run apps in the security context of standard users. When a user logs on to a computer, the system creates an access token for that user. The access token contains information about the level of access that the user is granted, including specific security identifiers (SIDs) and Windows privileges.

When an administrator logs on, two separate access tokens are created for the user: a standard user access token and an administrator access token. The standard user access token contains the same user-specific information as the administrator access token, but the administrative Windows privileges and SIDs are removed. The standard user access token is used to start apps that do not perform administrative tasks (standard user apps). The standard user access token is then used to display the desktop (explorer.exe). Explorer.exe is the parent process from which all other user-initiated processes inherit their access token. As a result, all apps run as a standard user unless a user provides consent or credentials to approve an app to use a full administrative access token.

A user that is a member of the Administrators group can log on, browse the Web, and read e-mail while using a standard user access token. When the administrator needs to perform a task that requires the administrator access token, Windows automatically prompts the user for approval. This prompt is called an elevation prompt, and its behavior can be configured by using the Local Security Policy snap-in (Secpol.msc) or Group Policy. For more info, see User Account Control security policy settings.

The UAC User Experience

When UAC is enabled, the user experience for standard users is different from that of administrators in Admin Approval Mode. The recommended and more secure method of running Windows, is to make your primary user account a standard user account. Running as a standard user helps to maximize security for a managed environment. With the built-in UAC elevation component, standard users can easily perform an administrative task by entering valid credentials for a local administrator account. The default, built-in UAC elevation component for standard users is the credential prompt.

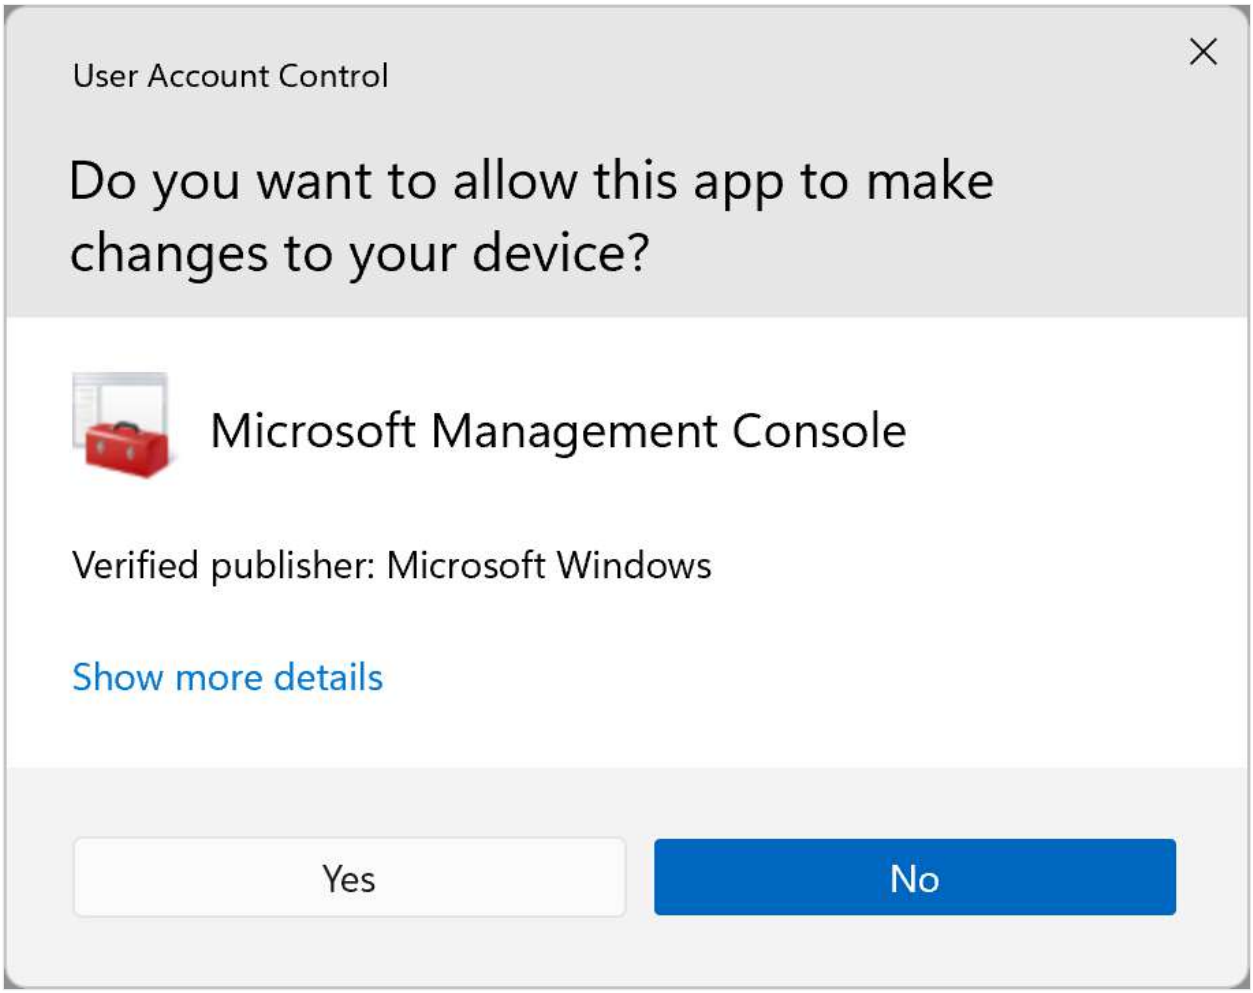
The alternative to running as a standard user is to run as an administrator in Admin Approval Mode. With the built-in UAC elevation component, members of the local Administrators group can easily perform an administrative task by providing approval. The default, built-in UAC elevation component for an administrator account in Admin Approval Mode is called the consent prompt.

The consent and credential prompts

With UAC enabled, Windows prompts for consent or prompts for credentials of a valid local administrator account before starting a program or task that requires a full administrator access token. This prompt ensures that no malicious software can be silently installed.

The consent prompt

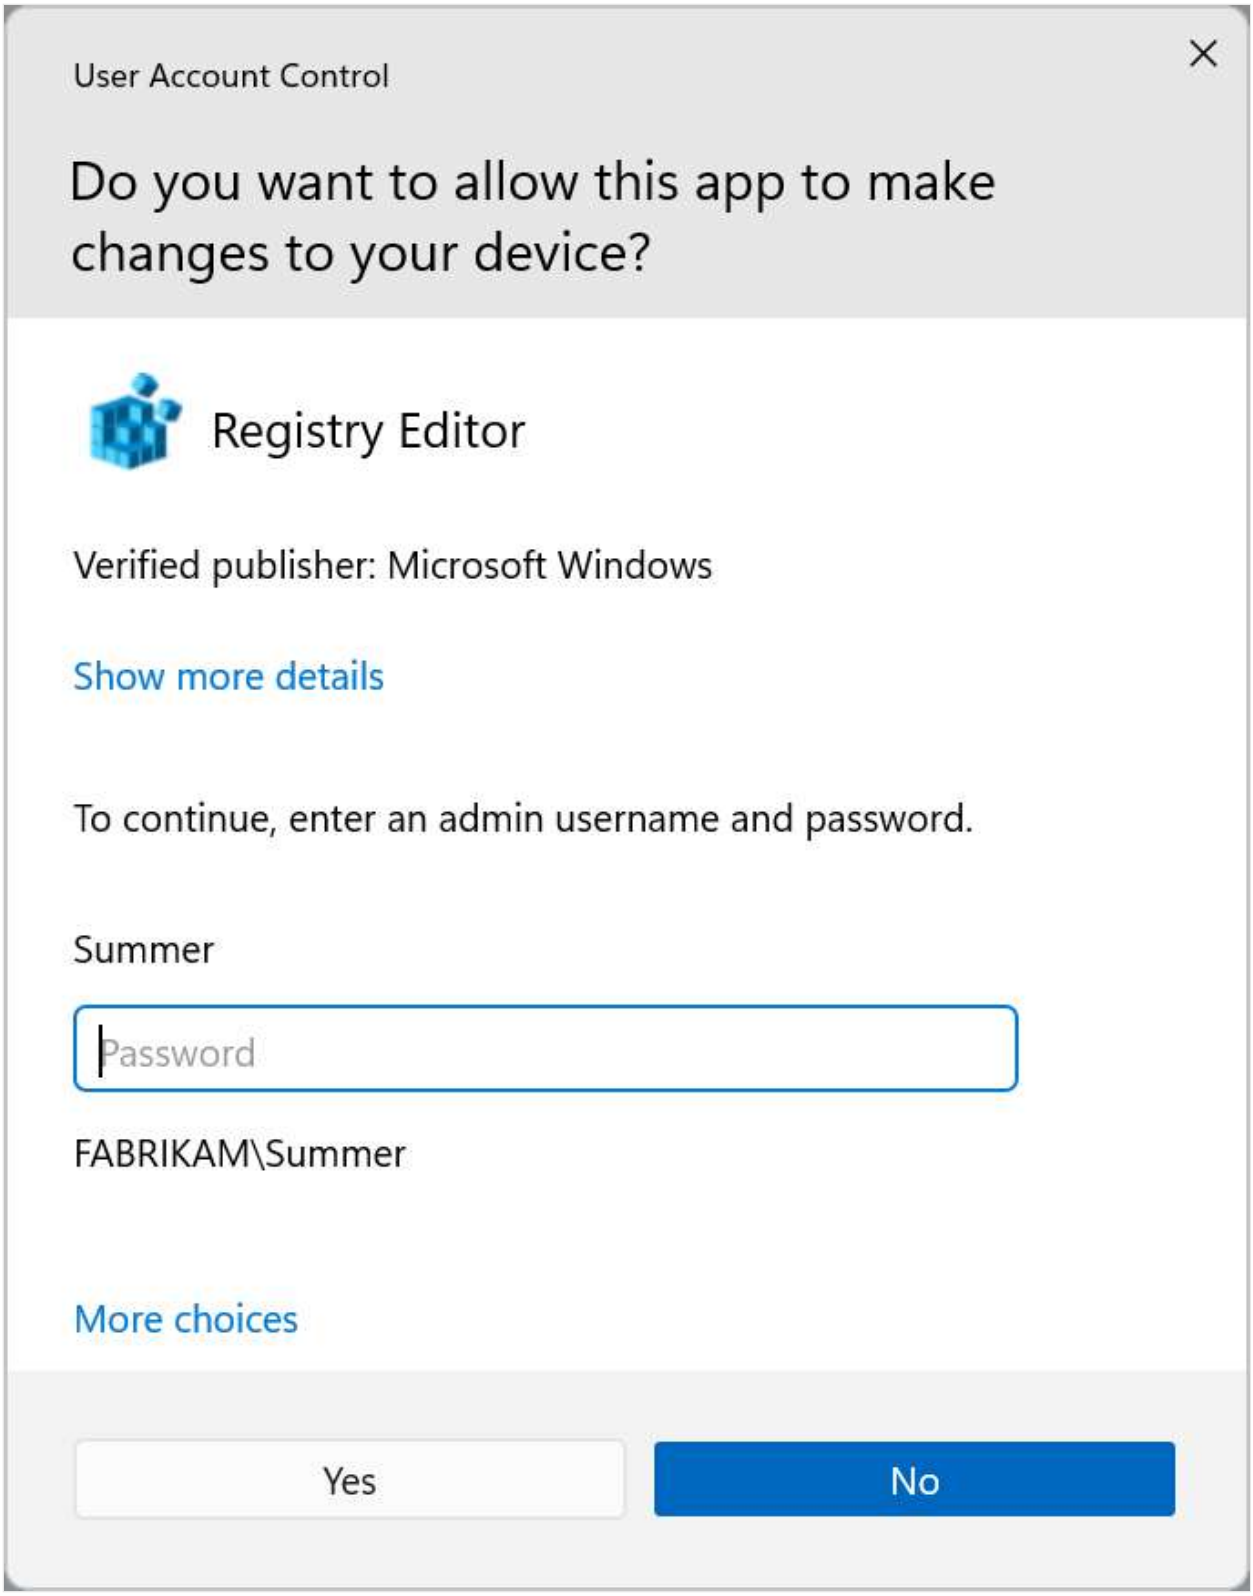
The consent prompt is presented when a user attempts to perform a task that requires a user's administrative access token. The following is an example of the UAC consent prompt.



The credential prompt

The credential prompt is presented when a standard user attempts to perform a task that requires a user's administrative access token. Administrators can also be required to provide their credentials by setting the **User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode** policy setting value to **Prompt for credentials**.

The following is an example of the UAC credential prompt.



UAC elevation prompts

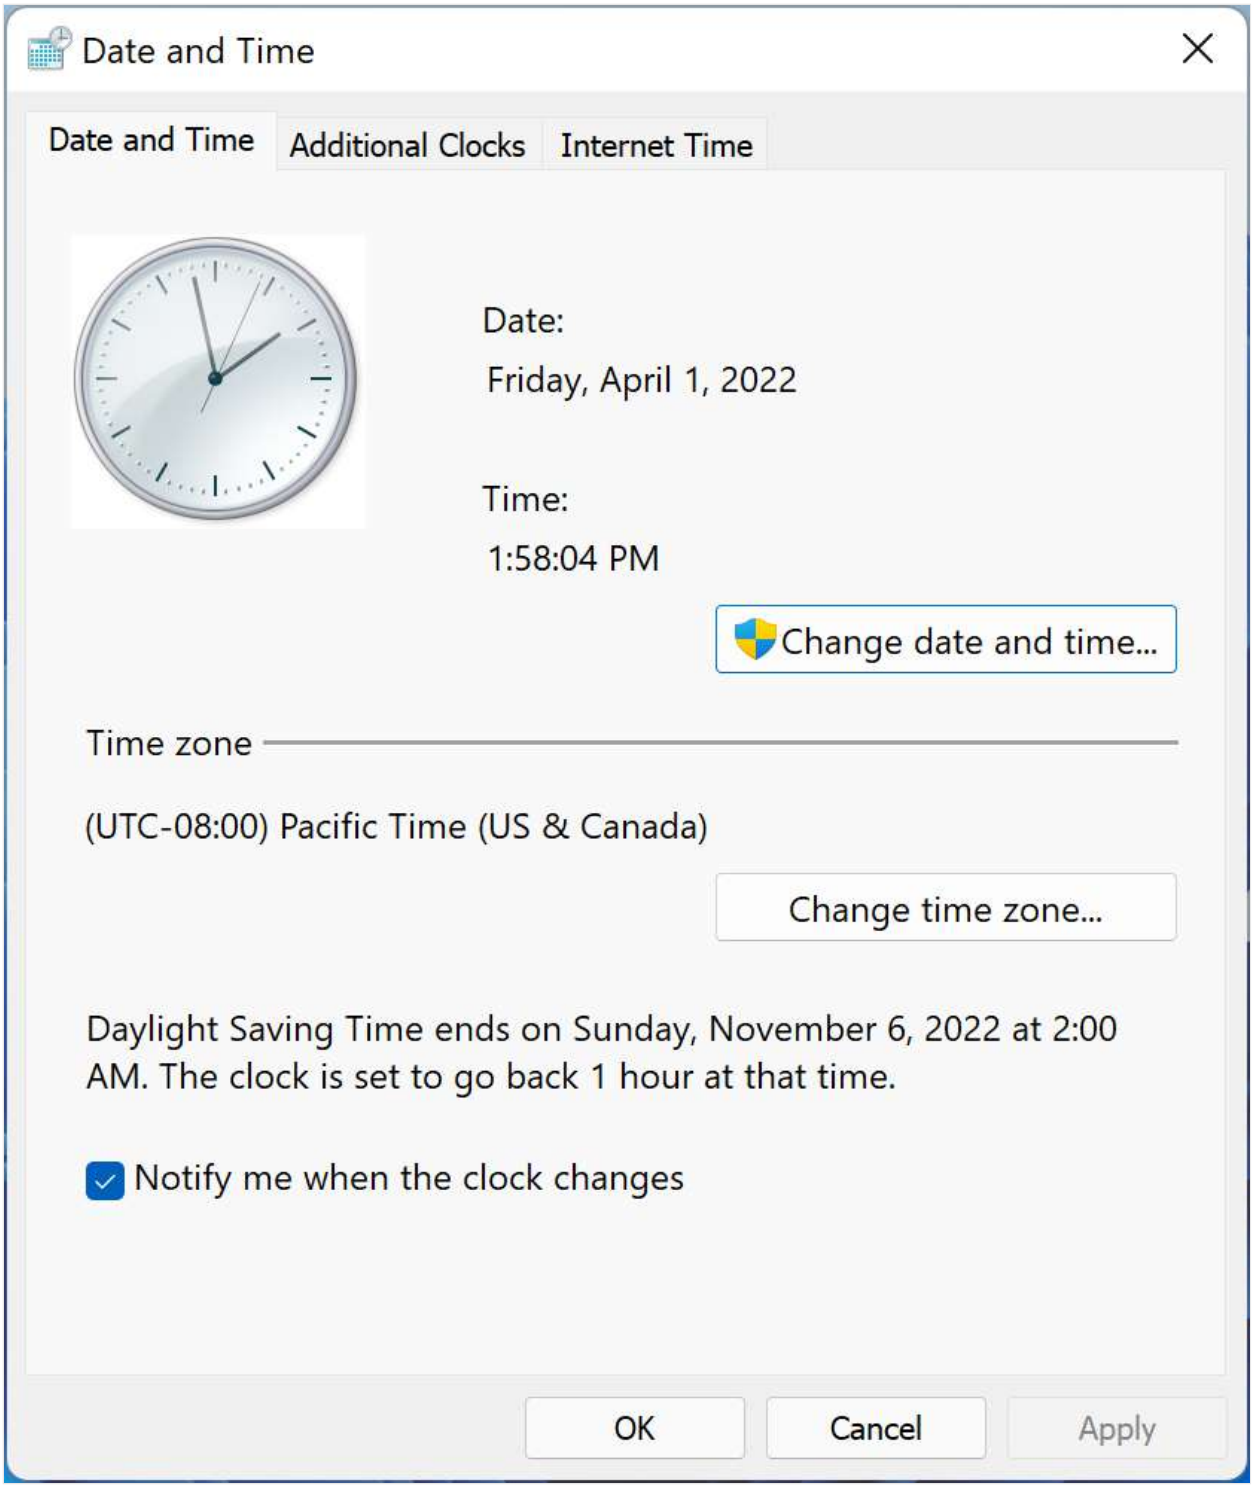
The UAC elevation prompts are color-coded to be app-specific, enabling for immediate identification of an application's potential security risk. When an app attempts to run with an administrator's full access token, Windows first analyzes the executable file to determine its publisher. Apps are first separated into three categories based on the file's publisher: Windows 10 or Windows 11, publisher verified (signed), and publisher not verified (unsigned). The following diagram illustrates how Windows determines which color elevation prompt to present to the user.

The elevation prompt color-coding is as follows:

- Red background with a red shield icon: The app is blocked by Group Policy or is from a publisher that is blocked.
- Blue background with a blue and gold shield icon: The application is a Windows 10 and Windows 11 administrative app, such as a Control Panel item.
- Blue background with a blue shield icon: The application is signed by using Authenticode and is trusted by the local computer.
- Yellow background with a yellow shield icon: The application is unsigned or signed but is not yet trusted by the local computer.

Shield icon

Some Control Panel items, such as **Date and Time Properties**, contain a combination of administrator and standard user operations. Standard users can view the clock and change the time zone, but a full administrator access token is required to change the local system time. The following is a screenshot of the **Date and Time Properties** Control Panel item.



The shield icon on the **Change date and time** button indicates that the process requires a full administrator access token and will display a UAC elevation prompt.

Securing the elevation prompt

The elevation process is further secured by directing the prompt to the secure desktop. The consent and credential prompts are displayed on the secure desktop by default in Windows 10 and Windows 11. Only Windows processes can access the secure desktop. For higher levels of security, we recommend keeping the **User Account Control: Switch to the secure desktop when prompting for elevation** policy setting enabled.

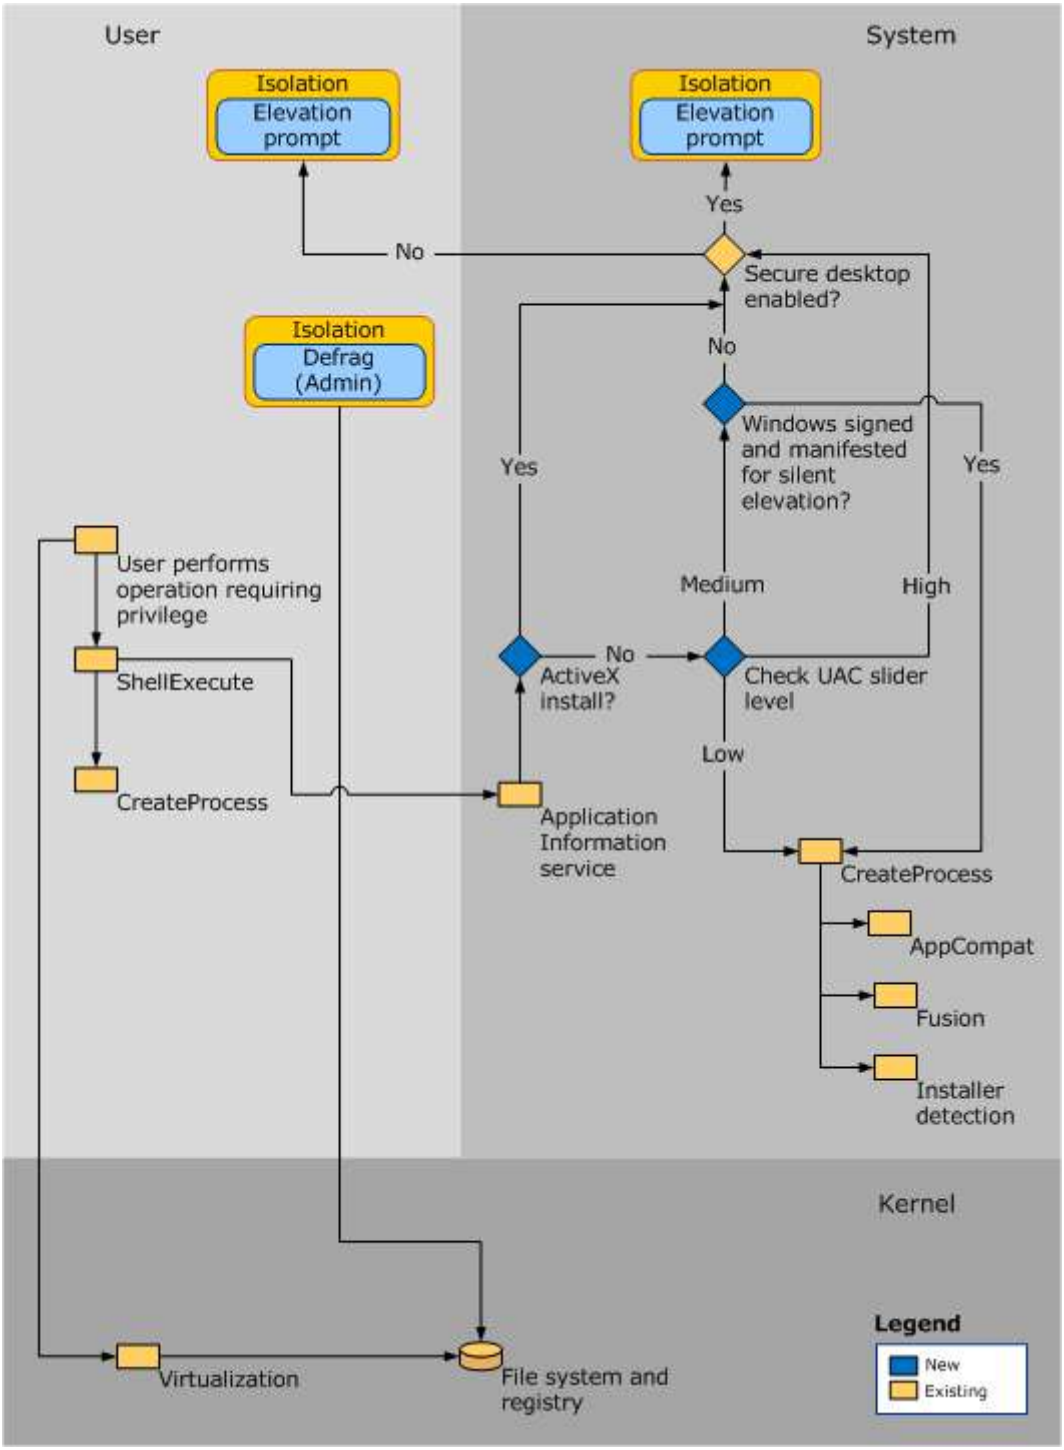
When an executable file requests elevation, the interactive desktop, also called the user desktop, is switched to the secure desktop. The secure desktop dims the user desktop and displays an elevation prompt that must be responded to before continuing. When the user clicks **Yes** or **No**, the desktop switches back to the user desktop.

Malware can present an imitation of the secure desktop, but when the **User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode** policy setting is set to **Prompt for consent**, the malware does not gain elevation if the user clicks **Yes** on the imitation. If the policy setting is set to **Prompt for credentials**, malware imitating the credential prompt may be able to gather the credentials from the user. However, the malware does not gain elevated privilege and the system has other protections that mitigate malware from taking control of the user interface even with a harvested password.

While malware could present an imitation of the secure desktop, this issue cannot occur unless a user previously installed the malware on the PC. Because processes requiring an administrator access token cannot silently install when UAC is enabled, the user must explicitly provide consent by clicking **Yes** or by providing administrator credentials. The specific behavior of the UAC elevation prompt is dependent upon Group Policy.

UAC Architecture

The following diagram details the UAC architecture.



To better understand each component, review the table below:

User

Component	Description
User performs operation requiring privilege	If the operation changes the file system or registry, Virtualization is called. All other operations call ShellExecute.

Component	Description
ShellExecute	ShellExecute calls CreateProcess. ShellExecute looks for the ERROR_ELEVATION_REQUIRED error from CreateProcess. If it receives the error, ShellExecute calls the Application Information service to attempt to perform the requested task with the elevated prompt.
CreateProcess	If the application requires elevation, CreateProcess rejects the call with ERROR_ELEVATION_REQUIRED.

System

Component	Description
Application Information service	A system service that helps start apps that require one or more elevated privileges or user rights to run, such as local administrative tasks, and apps that require higher integrity levels. The Application Information service helps start such apps by creating a new process for the application with an administrative user's full access token when elevation is required and (depending on Group Policy) consent is given by the user to do so.
Elevating an ActiveX install	If ActiveX is not installed, the system checks the UAC slider level. If ActiveX is installed, the **User Account Control: Switch to the secure desktop when prompting for elevation** Group Policy setting is checked.

Component	Description
Check UAC slider level	UAC has a slider to select from four levels of notification. • **Always notify** will: ◦ Notify you when programs try to install software or make changes to your computer. ◦ Notify you when you make changes to Windows settings. ◦ Freeze other tasks until you respond. Recommended if you often install new software or visit unfamiliar websites. • **Notify me only when programs try to make changes to my computer** will: ◦ Notify you when programs try to install software or make changes to your computer. ◦ Not notify you when you make changes to Windows settings. ◦ Freeze other tasks until you respond. Recommended if you do not often install apps or visit unfamiliar websites. • **Notify me only when programs try to make changes to my computer (do not dim my desktop)** will: ◦ Notify you when programs try to install software or make changes to your computer. ◦ Not notify you when you make changes to Windows settings. ◦ Not freeze other tasks until you respond. Not recommended. Choose this only if it takes a long time to dim the desktop on your computer. • **Never notify (Disable UAC prompts)** will: ◦ Not notify you when programs try to install software or make changes to your computer. ◦ Not notify you when you make changes to Windows settings. ◦ Not freeze other tasks until you respond. Not recommended due to security concerns.
Secure desktop enabled	The **User Account Control: Switch to the secure desktop when prompting for elevation** policy setting is checked: • If the secure desktop is enabled, all elevation requests go to the secure desktop regardless of prompt behavior policy settings for administrators and standard users. • If the secure desktop is not enabled, all elevation requests go to the interactive user's desktop, and the per-user settings for administrators and standard users are used.

Component	Description
CreateProcess	CreateProcess calls AppCompat, Fusion, and Installer detection to assess if the app requires elevation. The file is then inspected to determine its requested execution level, which is stored in the application manifest for the file. CreateProcess fails if the requested execution level specified in the manifest does not match the access token and returns an error (ERROR_ELEVATION_REQUIRED) to ShellExecute.
AppCompat	The AppCompat database stores information in the application compatibility fix entries for an application.
Fusion	The Fusion database stores information from application manifests that describe the applications. The manifest schema is updated to add a new requested execution level field.
Installer detection	Installer detection detects setup files, which helps prevent installations from being run without the user's knowledge and consent.

Kernel

Component	Description
Virtualization	Virtualization technology ensures that non-compliant apps do not silently fail to run or fail in a way that the cause cannot be determined. UAC also provides file and registry virtualization and logging for applications that write to protected areas.
File system and registry	The per-user file and registry virtualization redirects per-computer registry and file write requests to equivalent per-user locations. Read requests are redirected to the virtualized per-user location first and to the per-computer location second.

The slider will never turn UAC completely off. If you set it to **Never notify**, it will:

- Keep the UAC service running.
- Cause all elevation request initiated by administrators to be auto-approved without showing a UAC prompt.
- Automatically deny all elevation requests for standard users.

> ⓘ **Important**
>
> In order to fully disable UAC you must disable the policy **User Account Control: Run all administrators in Admin Approval Mode**.

> ⚠ **Warning**

Some Universal Windows Platform apps may not work when UAC is disabled.

Virtualization

Because system administrators in enterprise environments attempt to secure systems, many line-of-business (LOB) applications are designed to use only a standard user access token. As a result, you do not need to replace the majority of apps when UAC is turned on.

Windows 10 and Windows 11 include file and registry virtualization technology for apps that are not UAC-compliant and that require an administrator's access token to run correctly. When an administrative app that is not UAC-compliant attempts to write to a protected folder, such as Program Files, UAC gives the app its own virtualized view of the resource it is attempting to change. The virtualized copy is maintained in the user's profile. This strategy creates a separate copy of the virtualized file for each user that runs the non-compliant app.

Most app tasks operate properly by using virtualization features. Although virtualization allows a majority of applications to run, it is a short-term fix and not a long-term solution. App developers should modify their apps to be compliant as soon as possible, rather than relying on file, folder, and registry virtualization.

Virtualization is not an option in the following scenarios:

- Virtualization does not apply to apps that are elevated and run with a full administrative access token.
- Virtualization supports only 32-bit apps. Non-elevated 64-bit apps simply receive an access denied message when they attempt to acquire a handle (a unique identifier) to a Windows object. Native Windows 64-bit apps are required to be compatible with UAC and to write data into the correct locations.
- Virtualization is disabled if the app includes an app manifest with a requested execution level attribute.

Request execution levels

An app manifest is an XML file that describes and identifies the shared and private side-by-side assemblies that an app should bind to at run time. The app manifest includes entries for UAC app compatibility purposes. Administrative apps that include an entry in the app manifest prompt the user for permission to access the user's access token. Although they lack an entry in the app manifest, most administrative app can run

without modification by using app compatibility fixes. App compatibility fixes are database entries that enable applications that are not UAC-compliant to work properly.

All UAC-compliant apps should have a requested execution level added to the application manifest. If the application requires administrative access to the system, then marking the app with a requested execution level of "require administrator" ensures that the system identifies this program as an administrative app and performs the necessary elevation steps. Requested execution levels specify the privileges required for an app.

Installer detection technology

Installation programs are apps designed to deploy software. Most installation programs write to system directories and registry keys. These protected system locations are typically writeable only by an administrator in Installer detection technology, which means that standard users do not have sufficient access to install programs. Windows 10 and Windows 11 heuristically detect installation programs and requests administrator credentials or approval from the administrator user in order to run with access privileges. Windows 10 and Windows 11 also heuristically detect updates and programs that uninstall applications. One of the design goals of UAC is to prevent installations from being run without the user's knowledge and consent because installation programs write to protected areas of the file system and registry.

Installer detection only applies to:

- 32-bit executable files.
- Applications without a requested execution level attribute.
- Interactive processes running as a standard user with UAC enabled.

Before a 32-bit process is created, the following attributes are checked to determine whether it is an installer:

- The file name includes keywords such as "install," "setup," or "update."
- Versioning Resource fields contain the following keywords: Vendor, Company Name, Product Name, File Description, Original Filename, Internal Name, and Export Name.
- Keywords in the side-by-side manifest are embedded in the executable file.
- Keywords in specific StringTable entries are linked in the executable file.
- Key attributes in the resource script data are linked in the executable file.
- There are targeted sequences of bytes within the executable file.

> ⓘ **Note**

The keywords and sequences of bytes were derived from common characteristics observed from various installer technologies.

> **Note**
>
> The User Account Control: Detect application installations and prompt for elevation policy setting must be enabled for installer detection to detect installation programs. For more info, see **User Account Control security policy settings**.

User Account Control security policy settings

Article • 12/15/2022 • 6 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

You can use security policies to configure how User Account Control works in your organization. They can be configured locally by using the Local Security Policy snap-in (secpol.msc) or configured for the domain, OU, or specific groups by Group Policy.

User Account Control: Admin Approval Mode for the Built-in Administrator account

This policy setting controls the behavior of Admin Approval Mode for the built-in Administrator account.

- **Enabled** The built-in Administrator account uses Admin Approval Mode. By default, any operation that requires elevation of privilege will prompt the user to approve the operation.
- **Disabled** (Default) The built-in Administrator account runs all applications with full administrative privilege.

User Account Control: Allow UIAccess application to prompt for elevation without using the secure desktop

This policy setting controls whether User Interface Accessibility (UIAccess or UIA) programs can automatically disable the secure desktop for elevation prompts used by a standard user.

- **Enabled** UIA programs, including Windows Remote Assistance, automatically disable the secure desktop for elevation prompts. If you don't disable the "User Account Control: Switch to the secure desktop when prompting for elevation" policy setting, the prompts appear on the interactive user's desktop instead of the secure desktop.
- **Disabled** (Default) The secure desktop can be disabled only by the user of the interactive desktop or by disabling the "User Account Control: Switch to the secure desktop when prompting for elevation" policy setting.

User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode

This policy setting controls the behavior of the elevation prompt for administrators.

- **Elevate without prompting** Allows privileged accounts to perform an operation that requires elevation without requiring consent or credentials.

  > **Note:** Use this option only in the most constrained environments.

- **Prompt for credentials on the secure desktop** When an operation requires elevation of privilege, the user is prompted on the secure desktop to enter a privileged user name and password. If the user enters valid credentials, the operation continues with the user's highest available privilege.
- **Prompt for consent on the secure desktop** When an operation requires elevation of privilege, the user is prompted on the secure desktop to select either Permit or Deny. If the user selects Permit, the operation continues with the user's highest available privilege.
- **Prompt for credentials** When an operation requires elevation of privilege, the user is prompted to enter an administrative user name and password. If the user enters valid credentials, the operation continues with the applicable privilege.
- **Prompt for consent** When an operation requires elevation of privilege, the user is prompted to select either Permit or Deny. If the user selects Permit, the operation continues with the user's highest available privilege.
- **Prompt for consent for non-Windows binaries** (Default) When an operation for a non-Microsoft application requires elevation of privilege, the user is prompted on the secure desktop to select either Permit or Deny. If the user selects Permit, the operation continues with the user's highest available privilege.

User Account Control: Behavior of the elevation prompt for standard users

This policy setting controls the behavior of the elevation prompt for standard users.

- **Prompt for credentials** (Default) When an operation requires elevation of privilege, the user is prompted to enter an administrative user name and password. If the

user enters valid credentials, the operation continues with the applicable privilege.

- **Automatically deny elevation requests** When an operation requires elevation of privilege, a configurable access denied error message is displayed. An enterprise that is running desktops as standard user may choose this setting to reduce help desk calls.
- **Prompt for credentials on the secure desktop** When an operation requires elevation of privilege, the user is prompted on the secure desktop to enter a different user name and password. If the user enters valid credentials, the operation continues with the applicable privilege.

User Account Control: Detect application installations and prompt for elevation

This policy setting controls the behavior of application installation detection for the computer.

- **Enabled** (Default) When an app installation package is detected that requires elevation of privilege, the user is prompted to enter an administrative user name and password. If the user enters valid credentials, the operation continues with the applicable privilege.
- **Disabled** App installation packages aren't detected and prompted for elevation. Enterprises that are running standard user desktops and use delegated installation technologies, such as Group Policy or Microsoft Intune should disable this policy setting. In this case, installer detection is unnecessary.

User Account Control: Only elevate executable files that are signed and validated

This policy setting enforces public key infrastructure (PKI) signature checks for any interactive applications that request elevation of privilege. Enterprise administrators can control which applications are allowed to run by adding certificates to the Trusted Publishers certificate store on local computers.

- **Enabled** Enforces the certificate certification path validation for a given executable file before it's permitted to run.
- **Disabled** (Default) Doesn't enforce the certificate certification path validation before a given executable file is permitted to run.

User Account Control: Only elevate UIAccess applications that are installed in secure locations

This policy setting controls whether applications that request to run with a User Interface Accessibility (UIAccess) integrity level must reside in a secure location in the file system. Secure locations are limited to the following folders:

- ...\Program Files\, including subfolders
- ...\Windows\system32\
- ...\Program Files (x86)\, including subfolders for 64-bit versions of Windows

> **Note:** Windows enforces a digital signature check on any interactive app that requests to run with a UIAccess integrity level regardless of the state of this security setting.

- **Enabled** (Default) If an app resides in a secure location in the file system, it runs only with UIAccess integrity.
- **Disabled** An app runs with UIAccess integrity even if it doesn't reside in a secure location in the file system.

User Account Control: Turn on Admin Approval Mode

This policy setting controls the behavior of all User Account Control (UAC) policy settings for the computer. If you change this policy setting, you must restart your computer.

- **Enabled** (Default) Admin Approval Mode is enabled. This policy must be enabled and related UAC policy settings must also be set appropriately. They'll allow the built-in Administrator account and all other users who are members of the Administrators group to run in Admin Approval Mode.
- **Disabled** Admin Approval Mode and all related UAC policy settings are disabled. Note: If this policy setting is disabled, the Windows Security app notifies you that the overall security of the operating system has been reduced.

User Account Control: Switch to the secure desktop when prompting for elevation

This policy setting controls whether the elevation request prompt is displayed on the interactive user's desktop or the secure desktop.

- **Enabled** (Default) All elevation requests go to the secure desktop regardless of prompt behavior policy settings for administrators and standard users.
- **Disabled** All elevation requests go to the interactive user's desktop. Prompt behavior policy settings for administrators and standard users are used.

User Account Control: Virtualize file and registry write failures to per-user locations

This policy setting controls whether application write failures are redirected to defined registry and file system locations. This policy setting mitigates applications that run as administrator and write run-time application data to %ProgramFiles%, %Windir%, %Windir%\system32, or HKLM\Software.

- **Enabled** (Default) App write failures are redirected at run time to defined user locations for both the file system and registry.
- **Disabled** Apps that write data to protected locations fail.

User Account Control Group Policy and registry key settings

Article • 02/17/2023 • 12 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

Group Policy settings

There are 10 Group Policy settings that can be configured for User Account Control (UAC). The table lists the default for each of the policy settings, and the following sections explain the different UAC policy settings and provide recommendations. These policy settings are located in **Security Settings\Local Policies\Security Options** in the Local Security Policy snap-in. For more information about each of the Group Policy settings, see the Group Policy description. For information about the registry key settings, see Registry key settings.

Group Policy setting	Registry key	Default
User Account Control: Admin Approval Mode for the built-in Administrator account	FilterAdministratorToken	Disabled
User Account Control: Allow UIAccess applications to prompt for elevation without using the secure desktop	EnableUIADesktopToggle	Disabled
User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode	ConsentPromptBehaviorAdmin	Prompt for consent for non-Windows binaries
User Account Control: Behavior of the elevation prompt for standard users	ConsentPromptBehaviorUser	Prompt for credentials
User Account Control: Detect application installations and prompt for elevation	EnableInstallerDetection	Enabled (default for home) Disabled (default for enterprise)

Group Policy setting	Registry key	Default
User Account Control: Only elevate executables that are signed and validated	ValidateAdminCodeSignatures	Disabled
User Account Control: Only elevate UIAccess applications that are installed in secure locations	EnableSecureUIAPaths	Enabled
User Account Control: Run all administrators in Admin Approval Mode	EnableLUA	Enabled
User Account Control: Switch to the secure desktop when prompting for elevation	PromptOnSecureDesktop	Enabled
User Account Control: Virtualize file and registry write failures to per-user locations	EnableVirtualization	Enabled

User Account Control: Admin Approval Mode for the built-in Administrator account

The **User Account Control: Admin Approval Mode for the built-in Administrator account** policy setting controls the behavior of Admin Approval Mode for the built-in Administrator account.

The options are:

- **Enabled.** The built-in Administrator account uses Admin Approval Mode. By default, any operation that requires elevation of privilege will prompt the user to approve the operation.
- **Disabled.** (Default) The built-in Administrator account runs all applications with full administrative privilege.

User Account Control: Allow UIAccess applications to prompt for elevation without using the secure desktop

The **User Account Control: Allow UIAccess applications to prompt for elevation without using the secure desktop** policy setting controls whether User Interface Accessibility (UIAccess or UIA) programs can automatically disable the secure desktop for elevation prompts used by a standard user.

The options are:

- **Enabled.** UIA programs, including Windows Remote Assistance, automatically disable the secure desktop for elevation prompts. If you do not disable the **User**

Account Control: Switch to the secure desktop when prompting for elevation policy setting, the prompts appear on the interactive user's desktop instead of the secure desktop.

- **Disabled.** (Default) The secure desktop can be disabled only by the user of the interactive desktop or by disabling the **User Account Control: Switch to the secure desktop when prompting for elevation** policy setting.

UIA programs are designed to interact with Windows and application programs on behalf of a user. This policy setting allows UIA programs to bypass the secure desktop to increase usability in certain cases; however, allowing elevation requests to appear on the interactive desktop instead of the secure desktop can increase your security risk.

UIA programs must be digitally signed because they must be able to respond to prompts regarding security issues, such as the UAC elevation prompt. By default, UIA programs are run only from the following protected paths:

- ...\Program Files, including subfolders
- ...\Program Files (x86), including subfolders for 64-bit versions of Windows
- ...\Windows\System32

The **User Account Control: Only elevate UIAccess applications that are installed in secure locations** policy setting disables the requirement to be run from a protected path.

While this policy setting applies to any UIA program, it is primarily used in certain remote assistance scenarios, including the Windows Remote Assistance program in Windows 7.

If a user requests remote assistance from an administrator and the remote assistance session is established, any elevation prompts appear on the interactive user's secure desktop and the administrator's remote session is paused. To avoid pausing the remote administrator's session during elevation requests, the user may select the **Allow IT Expert to respond to User Account Control prompts** check box when setting up the remote assistance session. However, selecting this check box requires that the interactive user respond to an elevation prompt on the secure desktop. If the interactive user is a standard user, the user does not have the required credentials to allow elevation.

If you enable this policy setting, requests for elevation are automatically sent to the interactive desktop (not the secure desktop) and also appear on the remote administrator's view of the desktop during a remote assistance session. This allows the remote administrator to provide the appropriate credentials for elevation.

This policy setting does not change the behavior of the UAC elevation prompt for administrators.

If you plan to enable this policy setting, you should also review the effect of the **User Account Control: Behavior of the elevation prompt for standard users** policy setting. If it is configured as **Automatically deny elevation requests**, elevation requests are not presented to the user.

User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode

The **User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode** policy setting controls the behavior of the elevation prompt for administrators.

The options are:

- **Elevate without prompting.** Allows privileged accounts to perform an operation that requires elevation without requiring consent or credentials.

  Note Use this option only in the most constrained environments.

- **Prompt for credentials on the secure desktop.** When an operation requires elevation of privilege, the user is prompted on the secure desktop to enter a privileged user name and password. If the user enters valid credentials, the operation continues with the user's highest available privilege.

- **Prompt for consent on the secure desktop.** When an operation requires elevation of privilege, the user is prompted on the secure desktop to select either **Permit** or **Deny**. If the user selects **Permit**, the operation continues with the user's highest available privilege.

- **Prompt for credentials.** When an operation requires elevation of privilege, the user is prompted to enter an administrative user name and password. If the user enters valid credentials, the operation continues with the applicable privilege.

- **Prompt for consent.** When an operation requires elevation of privilege, the user is prompted to select either **Permit** or **Deny**. If the user selects **Permit**, the operation continues with the user's highest available privilege.

- **Prompt for consent for non-Windows binaries.** (Default) When an operation for a non-Microsoft application requires elevation of privilege, the user is prompted on the secure desktop to select either **Permit** or **Deny**. If the user selects **Permit**, the operation continues with the user's highest available privilege.

User Account Control: Behavior of the elevation prompt for standard users

The **User Account Control: Behavior of the elevation prompt for standard users** policy setting controls the behavior of the elevation prompt for standard users.

The options are:

- **Automatically deny elevation requests.** When an operation requires elevation of privilege, a configurable access denied error message is displayed. An enterprise that is running desktops as standard user may choose this setting to reduce help desk calls.
- **Prompt for credentials on the secure desktop.** When an operation requires elevation of privilege, the user is prompted on the secure desktop to enter a different user name and password. If the user enters valid credentials, the operation continues with the applicable privilege.
- **Prompt for credentials.** (Default) When an operation requires elevation of privilege, the user is prompted to enter an administrative user name and password. If the user enters valid credentials, the operation continues with the applicable privilege.

User Account Control: Detect application installations and prompt for elevation

The **User Account Control: Detect application installations and prompt for elevation** policy setting controls the behavior of application installation detection for the computer.

The options are:

- **Enabled.** (Default for home) When an application installation package is detected that requires elevation of privilege, the user is prompted to enter an administrative user name and password. If the user enters valid credentials, the operation continues with the applicable privilege.
- **Disabled.** (Default for enterprise) Application installation packages are not detected and prompted for elevation. Enterprises that are running standard user desktops and use delegated installation technologies such as Group Policy Software Installation or Systems Management Server (SMS) should disable this policy setting. In this case, installer detection is unnecessary.

User Account Control: Only elevate executables that are signed and validated

The **User Account Control: Only elevate executables that are signed and validated** policy setting enforces public key infrastructure (PKI) signature checks for any interactive applications that request elevation of privilege. Enterprise administrators can control which applications are allowed to run by adding certificates to the Trusted Publishers certificate store on local computers.

The options are:

- **Enabled.** Enforces the PKI certification path validation for a given executable file before it is permitted to run.
- **Disabled.** (Default) Does not enforce PKI certification path validation before a given executable file is permitted to run.

User Account Control: Only elevate UIAccess applications that are installed in secure locations

The **User Account Control: Only elevate UIAccess applications that are installed in secure locations** policy setting controls whether applications that request to run with a User Interface Accessibility (UIAccess) integrity level must reside in a secure location in the file system. Secure locations are limited to the following:

- ...\Program Files, including subfolders
- ...\Windows\system32
- ...\Program Files (x86), including subfolders for 64-bit versions of Windows

Note Windows enforces a PKI signature check on any interactive application that requests to run with a UIAccess integrity level regardless of the state of this security setting.

The options are:

- **Enabled.** (Default) If an application resides in a secure location in the file system, it runs only with UIAccess integrity.
- **Disabled.** An application runs with UIAccess integrity even if it does not reside in a secure location in the file system.

User Account Control: Run all administrators in Admin Approval Mode

The **User Account Control: Run all administrators Admin Approval Mode** policy setting controls the behavior of all UAC policy settings for the computer. If you change this policy setting, you must restart your computer.

The options are:

- **Enabled.** (Default) Admin Approval Mode is enabled. This policy must be enabled and related UAC policy settings must also be set appropriately to allow the built-in Administrator account and all other users who are members of the **Administrators** group to run in Admin Approval Mode.
- **Disabled.** Admin Approval Mode and all related UAC policy settings are disabled.

Note If this policy setting is disabled, the Windows Security app notifies you that the overall security of the operating system has been reduced.

User Account Control: Switch to the secure desktop when prompting for elevation

The **User Account Control: Switch to the secure desktop when prompting for elevation** policy setting controls whether the elevation request prompt is displayed on the interactive user's desktop or the secure desktop.

The options are:

- **Enabled.** (Default) All elevation requests go to the secure desktop regardless of prompt behavior policy settings for administrators and standard users.
- **Disabled.** All elevation requests go to the interactive user's desktop. Prompt behavior policy settings for administrators and standard users are used.

When this policy setting is enabled, it overrides the **User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode** policy setting. The following table describes the behavior of the elevation prompt for each of the administrator policy settings when the **User Account Control: Switch to the secure desktop when prompting for elevation** policy setting is enabled or disabled.

Administrator policy setting	Enabled	Disabled
Prompt for credentials on the secure desktop	The prompt appears on the secure desktop.	The prompt appears on the secure desktop.
Prompt for consent on the secure desktop	The prompt appears on the secure desktop.	The prompt appears on the secure desktop.

Administrator policy setting	Enabled	Disabled
Prompt for credentials	The prompt appears on the secure desktop.	The prompt appears on the interactive user's desktop.
Prompt for consent	The prompt appears on the secure desktop.	The prompt appears on the interactive user's desktop.
Prompt for consent for non-Windows binaries	The prompt appears on the secure desktop.	The prompt appears on the interactive user's desktop.

When this policy setting is enabled, it overrides the **User Account Control: Behavior of the elevation prompt for standard users** policy setting. The following table describes the behavior of the elevation prompt for each of the standard user policy settings when the **User Account Control: Switch to the secure desktop when prompting for elevation** policy setting is enabled or disabled.

Standard policy setting	Enabled	Disabled
Automatically deny elevation requests	No prompt. The request is automatically denied.	No prompt. The request is automatically denied.
Prompt for credentials on the secure desktop	The prompt appears on the secure desktop.	The prompt appears on the secure desktop.
Prompt for credentials	The prompt appears on the secure desktop.	The prompt appears on the interactive user's desktop.

User Account Control: Virtualize file and registry write failures to per-user locations

The **User Account Control: Virtualize file and registry write failures to per-user locations** policy setting controls whether application write failures are redirected to defined registry and file system locations. This policy setting mitigates applications that run as administrator and write run-time application data to %ProgramFiles%, %Windir%, %Windir%\system32, or HKLM\Software.

The options are:

- **Enabled.** (Default) Application write failures are redirected at run time to defined user locations for both the file system and registry.
- **Disabled.** Applications that write data to protected locations fail.

Registry key settings

The registry keys are found in **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System**. For information about each of the registry keys, see the associated Group Policy description.

Registry key	Group Policy setting	Registry setting
FilterAdministratorToken	User Account Control: Admin Approval Mode for the built-in Administrator account	0 (Default) = Disabled 1 = Enabled
EnableUIADesktopToggle	User Account Control: Allow UIAccess applications to prompt for elevation without using the secure desktop	0 (Default) = Disabled 1 = Enabled
ConsentPromptBehaviorAdmin	User Account Control: Behavior of the elevation prompt for administrators in Admin Approval Mode	0 = Elevate without prompting 1 = Prompt for credentials on the secure desktop 2 = Prompt for consent on the secure desktop 3 = Prompt for credentials 4 = Prompt for consent 5 (Default) = Prompt for consent for non-Windows binaries
ConsentPromptBehaviorUser	User Account Control: Behavior of the elevation prompt for standard users	0 = Automatically deny elevation requests 1 = Prompt for credentials on the secure desktop 3 (Default) = Prompt for credentials
EnableInstallerDetection	User Account Control: Detect application installations and prompt for elevation	1 = Enabled (default for home) 0 = Disabled (default for enterprise)
ValidateAdminCodeSignatures	User Account Control: Only elevate executables that are signed and validated	0 (Default) = Disabled 1 = Enabled

Registry key	Group Policy setting	Registry setting
EnableSecureUIAPaths	User Account Control: Only elevate UIAccess applications that are installed in secure locations	0 = Disabled 1 (Default) = Enabled
EnableLUA	User Account Control: Run all administrators in Admin Approval Mode	0 = Disabled 1 (Default) = Enabled
PromptOnSecureDesktop	User Account Control: Switch to the secure desktop when prompting for elevation	0 = Disabled 1 (Default) = Enabled
EnableVirtualization	User Account Control: Virtualize file and registry write failures to per-user locations	0 = Disabled 1 (Default) = Enabled

Smart Card Technical Reference

Article • 02/27/2023 • 2 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

The Smart Card Technical Reference describes the Windows smart card infrastructure for physical smart cards and how smart card-related components work in Windows. This document also contains information about tools that information technology (IT) developers and administrators can use to troubleshoot, debug, and deploy smart card-based strong authentication in the enterprise.

Audience

This document explains how the Windows smart card infrastructure works. To understand this information, you should have basic knowledge of public key infrastructure (PKI) and smart card concepts. This document is intended for:

- Enterprise IT developers, managers, and staff who are planning to deploy or are using smart cards in their organization.
- Smart card vendors who write smart card minidrivers or credential providers.

What are smart cards?

Smart cards are tamper-resistant portable storage devices that can enhance the security of tasks such as authenticating clients, signing code, securing e-mail, and signing in with a Windows domain account.

Smart cards provide:

- Tamper-resistant storage for protecting private keys and other forms of personal information.
- Isolation of security-critical computations that involve authentication, digital signatures, and key exchange from other parts of the computer. These computations are performed on the smart card.
- Portability of credentials and other private information between computers at work, home, or on the road.

Smart cards can be used to sign in to domain accounts only, not local accounts. When you use a password to sign in interactively to a domain account, Windows uses the

Kerberos version 5 (v5) protocol for authentication. If you use a smart card, the operating system uses Kerberos v5 authentication with X.509 v3 certificates.

Virtual smart cards were introduced in Windows Server 2012 and Windows 8 to alleviate the need for a physical smart card, the smart card reader, and the associated administration of that hardware.

> ⚠ **Warning**
>
> **Windows Hello for Business** is the modern, two-factor authentication for Windows. Microsoft will deprecate virtual smart cards in the near future. Customers using virtual smart cards are strongly encouraged to move to Windows Hello for Business. Microsoft will publish the deprecation date to ensure customers have adequate lead time to move to Windows Hello for Business. We recommend that new Windows deployments use Windows Hello for Business.

In this technical reference

This reference contains the following topics.

- How Smart Card Sign-in Works in Windows
  - Smart Card Architecture
  - Certificate Requirements and Enumeration
  - Smart Card and Remote Desktop Services
  - Smart Cards for Windows Service
  - Certificate Propagation Service
  - Smart Card Removal Policy Service
- Smart Card Tools and Settings
  - Smart Cards Debugging Information
  - Smart Card Group Policy and Registry Settings
  - Smart Card Events

How Smart Card Sign-in Works in Windows

Article • 12/09/2022 • 2 minutes to read •

Applies to: ✅ Windows 10, ✅ Windows 11, ✅ Windows Server 2016, ✅ Windows Server 2019, ✅ Windows Server 2022

This topic for IT professional provides links to resources about the implementation of smart card technologies in the Windows operating system. It includes the following resources about the architecture, certificate management, and services that are related to smart card use:

- Smart Card Architecture: Learn about enabling communications with smart cards and smart card readers, which can be different according to the vendor that supplies them.
- Certificate Requirements and Enumeration: Learn about requirements for smart card certificates based on the operating system, and about the operations that are performed by the operating system when a smart card is inserted into the computer.
- Smart Card and Remote Desktop Services: Learn about using smart cards for remote desktop connections.
- Smart Cards for Windows Service: Learn about how the Smart Cards for Windows service is implemented.
- Certificate Propagation Service: Learn about how the certificate propagation service works when a smart card is inserted into a computer.
- Smart Card Removal Policy Service: Learn about using Group Policy to control what happens when a user removes a smart card.

Smart Card Architecture

Article • 12/09/2022 • 19 minutes to read •

Applies to: ✅ Windows 10, ✅ Windows 11, ✅ Windows Server 2016, ✅ Windows Server 2019, ✅ Windows Server 2022

This topic for the IT professional describes the system architecture that supports smart cards in the Windows operating system, including credential provider architecture and the smart card subsystem architecture.

Authentication is a process for verifying the identity of an object or person. When you authenticate an object, such as a smart card, the goal is to verify that the object is genuine. When you authenticate a person, the goal is to verify that you are not dealing with an imposter.

In a networking context, authentication is the act of proving identity to a network application or resource. Typically, identity is proven by a cryptographic operation that uses a key only the user knows (such as with public key cryptography), or a shared key. The server side of the authentication exchange compares the signed data with a known cryptographic key to validate the authentication attempt. Storing the cryptographic keys in a secure central location makes the authentication process scalable and maintainable.

For smart cards, Windows supports a provider architecture that meets the secure authentication requirements and is extensible so that you can include custom credential providers. This topic includes information about:

- Credential provider architecture
- Smart card subsystem architecture

Credential provider architecture

The following table lists the components that are included in the interactive sign-in architecture of the Windows Server and Windows operating systems.

Component	Description
Winlogon	Provides an interactive sign-in infrastructure.
Logon UI	Provides interactive UI rendering.
Credential providers (password and smart card)	Describes credential information and serializing credentials.

Component	Description
Local Security Authority (LSA)	Processes sign-in credentials.
Authentication packages	Includes NTLM and the Kerberos protocol. Communicates with server authentication packages to authenticate users.

Interactive sign-in in Windows begins when the user presses CTRL+ALT+DEL. The CTRL+ALT+DEL key combination is called a secure attention sequence (SAS). To keep other programs and processes from using it, Winlogon registers this sequence during the boot process.

After receiving the SAS, the UI then generates the sign-in tile from the information received from the registered credential providers. The following graphic shows the architecture for credential providers in the Windows operating system.

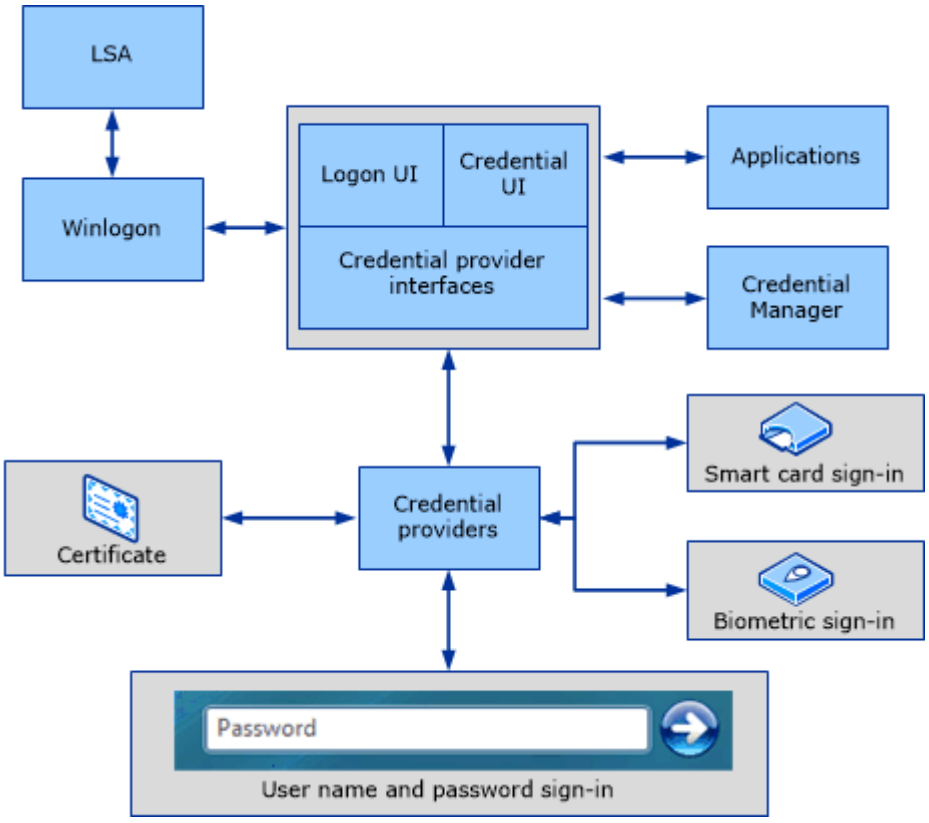


Figure 1 Credential provider architecture

Typically, a user who signs in to a computer by using a local account or a domain account must enter a user name and password. These credentials are used to verify the user's identity. For smart card sign-in, a user's credentials are contained on the smart card's security chip. A smart card reader lets the computer interact with the security chip on the smart card. When users sign in with a smart card, they enter a personal identification number (PIN) instead of a user name and password.

Credential providers are in-process COM objects that run on the local system and are used to collect credentials. The Logon UI provides interactive UI rendering, Winlogon provides interactive sign-in infrastructure, and credential providers work with both of these components to help gather and process credentials.

Winlogon instructs the Logon UI to display credential provider tiles after it receives an SAS event. The Logon UI queries each credential provider for the number of credentials it wants to enumerate. Credential providers have the option of specifying one of these tiles as the default. After all providers have enumerated their tiles, the Logon UI displays them to the user. The user interacts with a tile to supply the proper credentials. The Logon UI submits these credentials for authentication.

Combined with supporting hardware, credential providers can extend the Windows operating system to enable users to sign in by using biometrics (for example, fingerprint, retinal, or voice recognition), password, PIN, smart card certificate, or any custom authentication package. Enterprises and IT professionals can develop and deploy custom authentication mechanisms for all domain users, and they may explicitly require users to use this custom sign-in mechanism.

> **Note** Credential providers are not enforcement mechanisms. They are used to gather and serialize credentials. The LSA and authentication packages enforce security.

Credential providers can be designed to support single sign-in (SSO). In this process, they authenticate users to a secure network access point (by using RADIUS and other technologies) for signing in to the computer. Credential providers are also designed to support application-specific credential gathering, and they can be used for authentication to network resources, joining computers to a domain, or to provide administrator consent for User Account Control (UAC).

Multiple credential providers can coexist on a computer.

Credential providers must be registered on a computer running Windows, and they are responsible for:

- Describing the credential information that is required for authentication.
- Handling communication and logic with external authentication authorities.
- Packaging credentials for interactive and network sign-in.

> **Note** The Credential Provider API does not render the UI. It describes what needs to be rendered.

> Only the password credential provider is available in safe mode.
> The smart card credential provider is available in safe mode during networking.

Smart card subsystem architecture

Vendors provide smart cards and smart card readers, and in many cases the vendors are different for the smart card and the smart card reader. Drivers for smart card readers are written to the Personal Computer/Smart Card (PC/SC) standard. Each smart card must have a Cryptographic Service Provider (CSP) that uses the CryptoAPI interfaces to enable cryptographic operations, and the WinSCard APIs to enable communications with smart card hardware.

Base CSP and smart card minidriver architecture

Figure 2 illustrates the relationship between the CryptoAPI, CSPs, the Smart Card Base Cryptographic Service Provider (Base CSP), and smart card minidrivers.

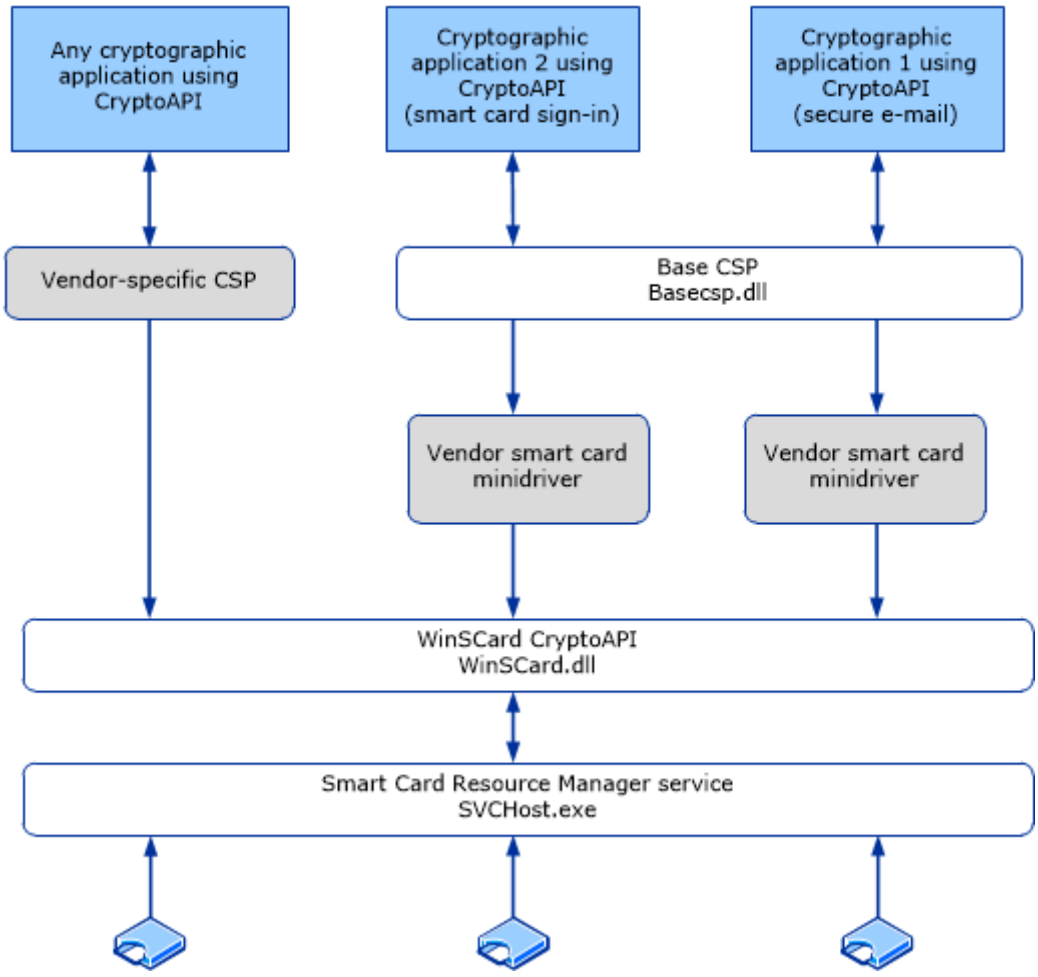


Figure 2 Base CSP and smart card minidriver architecture

Caching with Base CSP and smart card KSP

Smart card architecture uses caching mechanisms to assist in streamlining operations and to improve a user's access to a PIN.

- Data caching: The data cache provides for a single process to minimize smart card I/O operations.
- PIN caching: The PIN cache helps the user from having to reenter a PIN each time the smart card is unauthenticated.

Data caching

Each CSP implements the current smart card data cache separately. The Base CSP implements a robust caching mechanism that allows a single process to minimize smart card I/O operations.

The existing global cache works as follows:

1. The application requests a cryptographic operation. For example, a user certificate is to be read from the smart card.
2. The CSP checks its cache for the item.
3. If the item is not found in the cache, or if the item is cached but is not up-to-date, the item is read from the smart card.
4. After any item has been read from the smart card, it is added to the cache. Any existing out-of-date copy of that item is replaced.

Three types of objects or data are cached by the CSP: pins (for more information, see PIN caching), certificates, and files. If any of the cached data changes, the corresponding object is read from the smart card in successive operations. For example, if a file is written to the smart card, the CSP cache becomes out-of-date for the files, and other processes read the smart card at least once to refresh their CSP cache.

The global data cache is hosted in the Smart Cards for Windows service. Windows includes two public smart card API calls, SCardWriteCache and SCardReadCache. These API calls make global data caching functionality available to applications. Every smart card that conforms to the smart card minidriver specification has a 16-byte card identifier. This value is used to uniquely identify cached data that pertains to a given smart card. The standard Windows GUID type is used. These APIs allow an application to add data to and read data from the global cache.

PIN caching

The PIN cache protects the user from entering a PIN every time the smart card is unauthenticated. After a smart card is authenticated, it will not differentiate among host-side applications—any application can access private data on the smart card.

To mitigate this, the smart card enters an exclusive state when an application authenticates to the smart card. However, this means that other applications cannot communicate with the smart card and will be blocked. Therefore, such exclusive connections are minimized. The issue is that a protocol (such as the Kerberos protocol) requires multiple signing operations. Therefore, the protocol requires exclusive access to the smart card over an extended period, or it requires multiple authentication operations. This is where the PIN cache is used to minimize exclusive use of the smart card without forcing the user to enter a PIN multiple times.

The following example illustrates how this works. In this scenario, there are two applications: Outlook and Internet Explorer. The applications use smart cards for different purposes.

1. The user starts Outlook and tries to send a signed e-mail. The private key is on the smart card.
2. Outlook prompts the user for the smart card PIN. The user enters the correct PIN.
3. E-mail data is sent to the smart card for the signature operation. The Outlook client formats the response and sends the e-mail.
4. The user opens Internet Explorer and tries to access a protected site that requires Transport Layer Security (TLS) authentication for the client.
5. Internet Explorer prompts the user for the smart card PIN. The user enters the correct PIN.
6. The TLS-related private key operation occurs on the smart card, and the user is authenticated and signed in.
7. The user returns to Outlook to send another signed e-mail. This time, the user is not prompted for a PIN because the PIN is cached from the previous operation. Similarly, if the user uses Internet Explorer again for another operation, Internet Explorer will not prompt the user for a PIN.

The Base CSP internally maintains a per-process cache of the PIN. The PIN is encrypted and stored in memory. The functions that are used to secure the PIN are RtlEncryptMemory, RtlDecryptMemory, and RtlSecureZeroMemory, which will empty buffers that contained the PIN.

Smart card selection

The following sections in this topic describe how Windows leverages the smart card architecture to select the correct smart card reader software, provider, and credentials for a successful smart card sign-in:

- Container specification levels
- Container operations
- Context flags
- Create a new container in silent context
- Smart card selection behavior
- Make a smart card reader match
- Make a smart card match
- Open an existing default container (no reader specified)
- Open an existing GUID-named container (no reader specified)
- Create a new container (no reader specified)
- Delete a container

Container specification levels

In response to a CryptAcquireContext call in CryptoAPI, the Base CSP tries to match the container that the caller specifies to a specific smart card and reader. The caller can provide a container name with varying levels of specificity, as shown in the following table, and sorted from most-specific to least-specific requests.

Similarly, in response to a NCryptOpenKey call in CNG, the smart card KSP tries to match the container the same way, and it takes the same container format, as shown in the following table.

> **Note** Before opening a key by using the smart card KSP, a call to NCryptOpenStorageProvider (MS_SMART_CARD_KEY_STORAGE_PROVIDER) must be made.

Type	Name	Format

Type	Name	Format
I	Reader Name and Container Name	\\.\<Reader Name>\<Container Name>
II	Reader Name and Container Name (NULL)	\\.\<Reader Name>
III	Container Name Only	<Container Name>
IV	Default Container (NULL) Only	NULL

The Base CSP and smart card KSP cache smart card handle information about the calling process and about the smart cards the process has accessed. When searching for a smart card container, the Base CSP or smart card KSP first checks its cache for the process. If the cached handle is invalid or no match is found, the SCardUIDlg API is called to get the card handle.

Container operations

The following three container operations can be requested by using CryptAcquireContext:

1. Create a new container. (The CNG equivalent of CryptAcquireContext with dwFlags set to CRYPT_NEWKEYSET is NCryptCreatePersistedKey.)
2. Open an existing container. (The CNG equivalent of CryptAcquireContext to open the container is NCryptOpenKey.)
3. Delete a container. (The CNG equivalent of CryptAcquireContext with dwFlags set to CRYPT_DELETEKEYSET is NCryptDeleteKey.)

The heuristics that are used to associate a cryptographic handle with a particular smart card and reader are based on the container operation requested and the level of container specification used.

The following table shows the restrictions for the container creation operation.

Specification	Restriction
No silent context	Key container creation must always be able to show UI, such as the PIN prompt.
No overwriting existing containers	If the specified container already exists on the chosen smart card, choose another smart card or cancel the operation.

Context flags

The following table shows the context flags used as restrictions for the container creation operation.

Flag	Description
CRYPT_SILENT	No UI can be displayed during this operation.
CRYPT_MACHINE_KEYSET	No cached data should be used during this operation.
CRYPT_VERIFYCONTEXT	Only public data can be accessed on the smart card.

In addition to container operations and container specifications, you must consider other user options, such as the CryptAcquireContext flags, during smart card selection.

> **Important** The CRYPT_SILENT flag cannot be used to create a new container.

Create a new container in silent context

Applications can call the Base CSP with CRYPT_DEFAULT_CONTAINER_OPTIONAL, set the PIN in silent context, and then create a new container in silent context. This operation occurs as follows:

1. Call CryptAcquireContext by passing the smart card reader name in as a type II container specification level, and specifying the CRYPT_DEFAULT_CONTAINER_OPTIONAL flag.
2. Call CryptSetProvParam by specifying PP_KEYEXCHANGE_PIN or PP_SIGNATURE_PIN and a null-terminated ASCII PIN.
3. Release the context acquired in Step 1.
4. Call CryptAcquireContext with CRYPT_NEWKEYSET, and specify the type I container specification level.
5. Call CryptGenKey to create the key.

Smart card selection behavior

In some of the following scenarios, the user can be prompted to insert a smart card. If the user context is silent, this operation fails and no UI is displayed. Otherwise, in response to the UI, the user can insert a smart card or click **Cancel**. If the user cancels the operation, the operation fails. The flow chart in Figure 3 shows the selection steps performed by the Windows operating system.

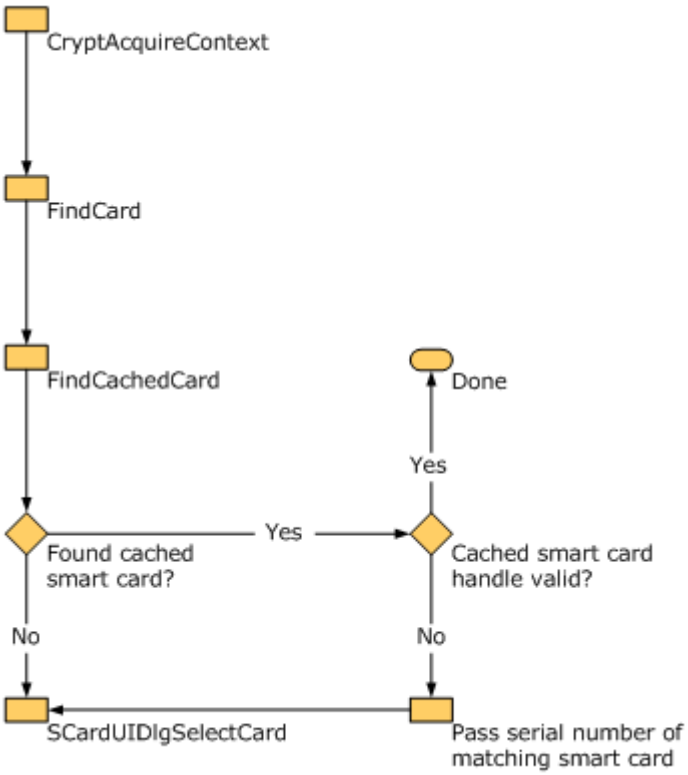


Figure 3 Smart card selection behavior

In general, smart card selection behavior is handled by the SCardUIDlgSelectCard API. The Base CSP interacts with this API by calling it directly. The Base CSP also sends callback functions that have the purpose of filtering and matching candidate smart cards. Callers of CryptAcquireContext provide smart card matching information. Internally, the Base CSP uses a combination of smart card serial numbers, reader names, and container names to find specific smart cards.

Each call to SCardUI * may result in additional information read from a candidate smart card. The Base CSP smart card selection callbacks cache this information.

Make a smart card reader match

For type I and type II container specification levels, the smart card selection process is less complex because only the smart card in the named reader can be considered a match. The process for matching a smart card with a smart card reader is:

1. Find the requested smart card reader. If it cannot be found, the process fails. (This requires a cache search by reader name.)

2. If no smart card is in the reader, the user is prompted to insert a smart card. (This is only in non-silent mode; if the call is made in silent mode, it will fail.)

3. For container specification level II only, the name of the default container on the chosen smart card is determined.

4. To open an existing container or delete an existing container, find the specified container. If the specified container cannot be found on this smart card, the user is prompted to insert a smart card.

5. If the system attempts to create a new container, if the specified container already exists on this smart card, the process fails.

Make a smart card match

For container specification levels III and IV, a broader method is used to match an appropriate smart card with a user context, because multiple cached smart cards might meet the criteria provided.

Open an existing default container (no reader specified)

> **Note** This operation requires that you use the smart card with the Base CSP.

1. For each smart card that has been accessed by the Base CSP and the handle and container information are cached, the Base CSP looks for a valid default container. An operation is attempted on the cached SCARDHANDLE to verify its validity. If the smart card handle is not valid, the Base CSP continues to search for a new smart card.

2. If a matching smart card is not found in the Base CSP cache, the Base CSP calls to the smart card subsystem. SCardUIDlgSelectCard() is used with an appropriate callback filter to find a matching smart card with a valid default container.

Open an existing GUID-named container (no reader specified)

> **Note** This operation requires that you use the smart card with the Base CSP.

1. For each smart card that is already registered with the Base CSP, search for the requested container. Attempt an operation on the cached SCARDHANDLE to verify its validity. If the smart card handle is not valid, the smart card's serial number is passed to the SCardUI * API to continue searching for this specific smart card (rather than only a general match for the container name).

2. If a matching smart card is not found in the Base CSP cache, a call is made to the smart card subsystem. SCardUIDlgSelectCard() is used with an appropriate callback filter to find a matching smart card with the requested container. Or, if a smart

card serial number resulted from the search in Step 1, the callback filter attempts to match the serial number, not the container name.

Create a new container (no reader specified)

> **Note** This operation requires that you use the smart card with the Base CSP.

If the PIN is not cached, no CRYPT_SILENT is allowed for the container creation because the user must be prompted for a PIN, at a minimum.

For other operations, the caller may be able to acquire a "verify" context against the default container (CRYPT_DEFAULT_CONTAINER_OPTIONAL) and then make a call with CryptSetProvParam to cache the user PIN for subsequent operations.

1. For each smart card already known by the CSP, refresh the stored SCARDHANDLE and make the following checks:

   a. If the smart card has been removed, continue the search.

   b. If the smart card is present, but it already has the named container, continue the search.

   c. If the smart card is available, but a call to CardQueryFreeSpace indicates that the smart card has insufficient storage for an additional key container, continue the search.

   d. Otherwise, use the first available smart card that meets the above criteria for the container creation.

2. If a matching smart card is not found in the CSP cache, make a call to the smart card subsystem. The callback that is used to filter enumerated smart cards verifies that a candidate smart card does not already have the named container, and that CardQueryFreeSpace indicates the smart card has sufficient space for an additional container. If no suitable smart card is found, the user is prompted to insert a smart card.

Delete a container

1. If the specified container name is NULL, the default container is deleted. Deleting the default container causes a new default container to be selected arbitrarily. For this reason, this operation is not recommended.

2. For each smart card already known by the CSP, refresh the stored SCARDHANDLE and make the following checks:

    a. If the smart card does not have the named container, continue the search.

    b. If the smart card has the named container, but the smart card handle is no longer valid, store the serial number of the matching smart card and pass it to SCardUI *.

3. If a matching smart card is not found in the CSP cache, make a call to the smart card subsystem. The callback that is used to filter enumerated smart cards should verify that a candidate smart card has the named container. If a serial number was provided as a result of the previous cache search, the callback should filter enumerated smart cards on serial number rather than on container matches. If the context is non-silent and no suitable smart card is found, display UI that prompts the user to insert a smart card.

Base CSP and KSP-based architecture in Windows

Figure 4 shows the Cryptography architecture that is used by the Windows operating system.

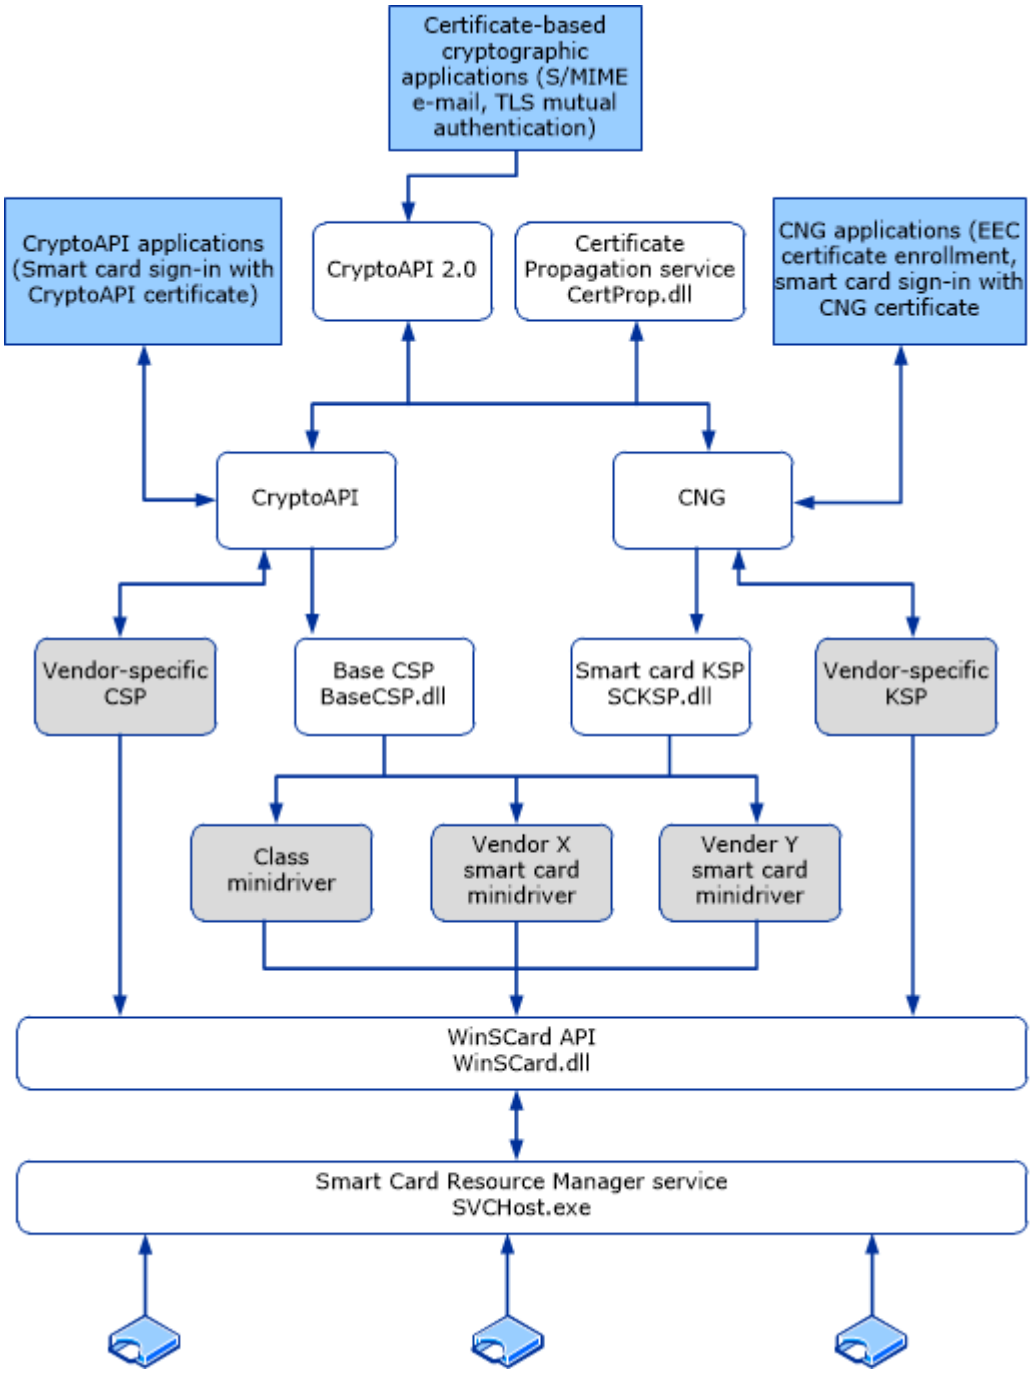


Figure 4 Cryptography architecture

Base CSP and smart card KSP properties in Windows

The following properties are supported in versions of Windows designated in the **Applies To** list at the beginning of this topic.

> **Note** The API definitions are located in WinCrypt.h and WinSCard.h.

Property	Description

Property	Description
PP_USER_CERTSTORE	- Used to return an HCERTSTORE that contains all user certificates on the smart card - Read-only (used only by CryptGetProvParam) - Caller responsible for closing the certificate store - Certificate encoded using PKCS_7_ASN_ENCODING or X509_ASN_ENCODING - CSP should set KEY_PROV_INFO on certificates - Certificate store should be assumed to be an in-memory store - Certificates should have a valid CRYPT_KEY_PROV_INFO as a property
PP_ROOT_CERTSTORE	- Read and Write (used by CryptGetProvParam and CryptSetProvParam) - Used to write a collection of root certificates to the smart card or return HCERTSTORE, which contains root certificates from the smart card - Used primarily for joining a domain by using a smart card - Caller responsible for closing the certificate store
PP_SMARTCARD_READER	- Read-only (used only by CryptGetProvParam) - Returns the smart card reader name as an ANSI string that is used to construct a fully qualified container name (that is, a smart card reader plus a container)
PP_SMARTCARD_GUID	- Return smart card GUID (also known as a serial number), which should be unique for each smart card - Used by the certificate propagation service to track the source of a root certificate
PP_UI_PROMPT	- Used to set the search string for the SCardUIDlgSelectCard card insertion dialog box - Persistent for the entire process when it is set - Write-only (used only by CryptSetProvParam)

Implications for CSPs in Windows

Cryptographic Service Providers (CSPs), including custom smart card CSPs, continue to be supported but this approach is not recommended. Using the existing Base CSP and smart card KSP with the smart card minidriver model for smart cards provides significant benefits in terms of performance, and PIN and data caching. One minidriver can be configured to work under CryptoAPI and CNG layers. This provides benefits from enhanced cryptographic support, including elliptic curve cryptography and AES.

If a smart card is registered by a CSP and a smart card minidriver, the one that was installed most recently will be used to communicate with the smart card.

Write a smart card minidriver, CSP, or KSP

CSPs and KSPs are meant to be written only if specific functionality is not available in the current smart card minidriver architecture. For example, the smart card minidriver architecture supports hardware security modules, so a minidriver could be written for a hardware security module, and a CSP or KSP may not be required unless it is needed to support algorithms that are not implemented in the Base CSP or smart card KSP.

For more information about how to write a smart card minidriver, CSP, or KSP, see Smart Card Minidrivers.

Certificate Requirements and Enumeration

Article • 01/24/2023 • 19 minutes to read •

Applies to: ✅ Windows 10, ✅ Windows 11, ✅ Windows Server 2016, ✅ Windows Server 2019, ✅ Windows Server 2022

This topic for the IT professional and smart card developers describes how certificates are managed and used for smart card sign-in.

When a smart card is inserted, the following steps are performed.

> **Note** Unless otherwise mentioned, all operations are performed silently (CRYPT_SILENT is passed to CryptAcquireContext).

1. The smart card resource manager database searches for the smart card's cryptographic service provider (CSP).
2. A qualified container name is constructed by using the smart card reader name, and it is passed to the CSP. The format is \\.\<*Reader name*>\
3. CryptAcquireContext is called to retrieve a context to the default container. If a failure occurs, the smart card will be unusable for smart card sign-in.
4. The name of the container is retrieved by using the PP_CONTAINER parameter with CryptGetProvParam.
5. Using the context acquired in Step 3, the CSP is queried for the PP_USER_CERTSTORE parameter (added in Windows Vista). For more information, see Smart Card Architecture. If the operation is successful, the name of a certificate store is returned, and the program flow skips to Step 8.
6. If the operation in Step 5 fails, the default container context from Step 3 is queried for the AT_KEYEXCHANGE key.
7. The certificate is then queried from the key context by using KP_CERTIFICATE. The certificate is added to an in-memory certificate store.
8. For each certificate in the certificate store from Step 5 or Step 7, the following checks are performed:
    a. The certificate must be valid, based on the computer system clock (not expired or valid with a future date).

b. The certificate must not be in the AT_SIGNATURE part of a container.

c. The certificate must have a valid user principal name (UPN).

d. The certificate must have the digital signature key usage.

e. The certificate must have the smart card logon EKU.

Any certificate that meets these requirements is displayed to the user with the certificate's UPN (or e-mail address or subject, depending on the presence of the certificate extensions).

> **Note** These requirements are the same as those in Windows Server 2003, but they are performed before the user enters the PIN. You can override many of them by using Group Policy settings.

9. The process then chooses a certificate, and the PIN is entered.
10. LogonUI.exe packages the information and sends it to Lsass.exe to process the sign-in attempt.
11. If successful, LogonUI.exe closes. This causes the context acquired in Step 3 to be released.

About Certificate support for compatibility

Although versions of Windows earlier than Windows Vista include support for smart cards, the types of certificates that smart cards can contain are limited. The limitations are:

- Each certificate must have a user principal name (UPN) and the smart card sign-in object identifier (also known as OID) in the extended key usage (EKU) attribute field. There is a Group Policy setting, Allow ECC certificates to be used for logon and authentication, to make the EKU optional.
- Each certificate must be stored in the AT_KEYEXCHANGE portion of the default CryptoAPI container, and non-default CryptoAPI containers are not supported.

The following table lists the certificate support in older Windows operating system versions.

Operating system	Certificate support

Operating system	Certificate support
Windows Server 2008 R2 and Windows 7	Support for smart card sign-in with ECC-based certificates. ECC smart card sign-in is enabled through Group Policy. ECDH_P256 ECDH Curve P-256 from FIPS 186-2 ECDSA_P256 ECDSA Curve P-256 from FIPS 186-2 ECDH_P384 ECDH Curve P-384 from FIPS 186-2 ECDH_P521 ECDH Curve P-521 from FIPS 186-2 ECDSA_P256 ECDH Curve P-256 from FIPS 186-2 ECDSA_P384 ECDSA Curve P-384 from FIPS 186-2 ECDSA_P521 ECDSA Curve P-384 from FIPS 186-2
Windows Server 2008 and Windows Vista	Valid certificates are enumerated and displayed from all smart cards and presented to the user. Keys are no longer restricted to the default container, and certificates in different containers can be chosen. Elliptic curve cryptography (ECC)-based certificates are not supported for smart card sign-in

Smart card sign-in flow in Windows

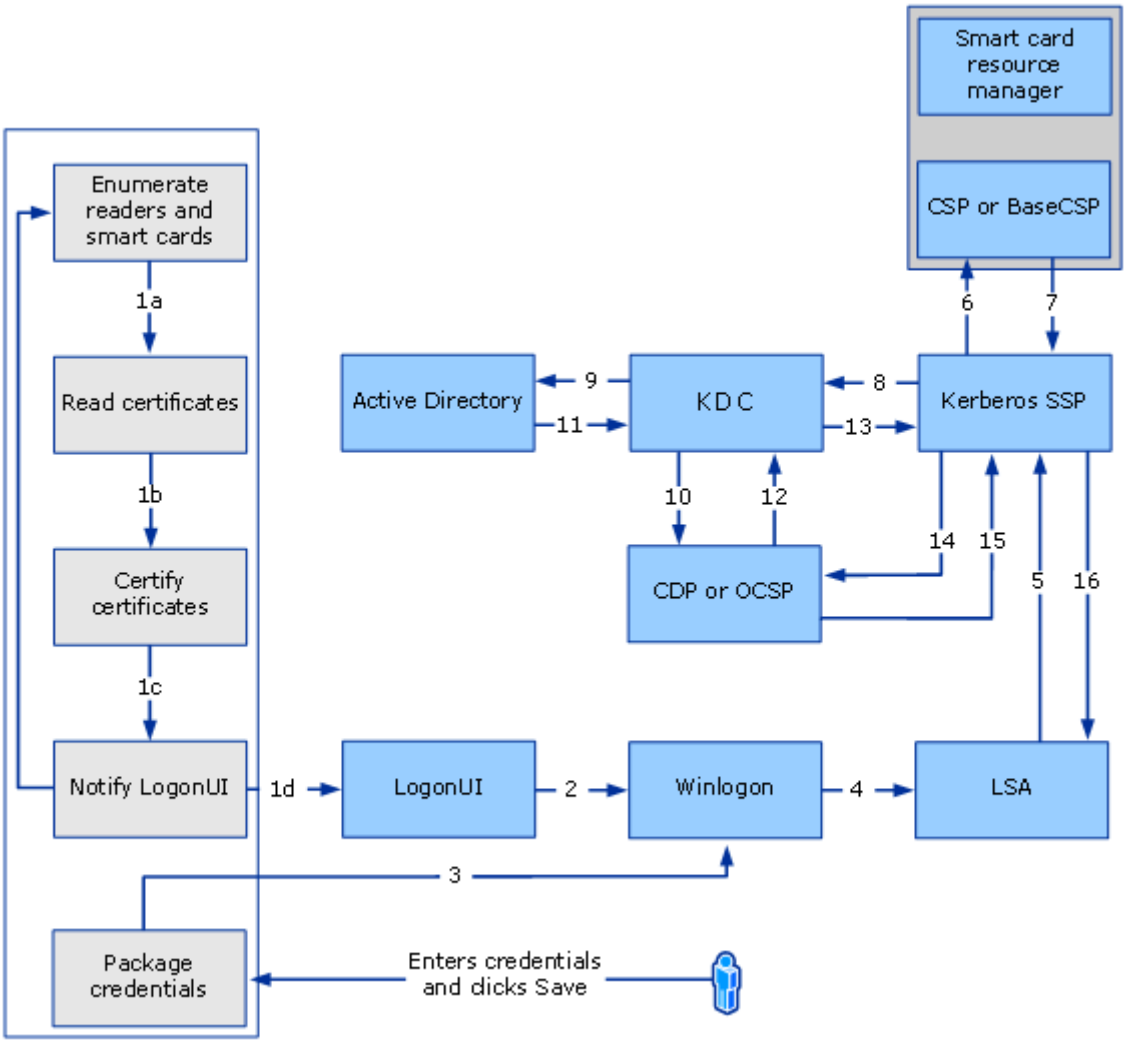
Most issues during authentication occur because of session behavior changes. When changes occur, the Local Security Authority (LSA) does not reacquire the session context; it relies instead on the Cryptographic Service Provider to handle the session change.

In the supported versions of Windows designated in the **Applies To** list at the beginning of this topic, client certificates that do not contain a UPN in the **subjectAltName** (SAN) field of the certificate can be enabled for sign-in, which supports a wider variety of certificates and supports multiple sign-in certificates on the same card.

Support for multiple certificates on the same card is enabled by default. New certificate types must be enabled through Group Policy.

If you enable the **Allow signature keys valid for Logon** credential provider policy, any certificates that are available on the smart card with a signature-only key are listed on the sign-in screen. This allows users to select their sign-in experience. If the policy is disabled or not configured, smart card signature-key-based certificates are not listed on the sign-in screen.

The following diagram illustrates how smart card sign-in works in the supported versions of Windows.



Smart card sign-in flow

Following are the steps that are performed during a smart card sign-in:

1. Winlogon requests the sign-in UI credential information.
2. Asynchronously, smart card resource manager starts, and the smart card credential provider does the following:
    a. Gets credential information (a list of known credentials, or if no credentials exist, the smart card reader information that Windows detected).
    b. Gets a list of smart card readers (by using the WinSCard API) and the list of smart cards inserted in each of them.
    c. Enumerates each card to verify that a sign-in certificate that is controlled by Group Policy is present. If the certificate is present, the smart card credential provider copies it into a temporary, secure cache on the computer or terminal.

    > **Note** Smartcard cache entries are created for certificates with a subject name or with a subject key identifier. If the certificate has a subject name, it is stored with an index that is based on the subject name and certificate issuer. If another certificate with the same subject name and certificate issuer is used, it will replace the existing cached entry. A change in this behavior after Windows Vista, allows for the condition when the certificate does not have a subject name, the cache is created with an index that is based on the subject key identifier and certificate issuer. If another certificate has the same the subject key identifier and certificate issuer, the cache entry is replaced. When certificates have neither a subject name nor subject key identifier, a cached entry is not created.

    d. Notifies the sign-in UI that it has new credentials.
3. The sign-in UI requests the new credentials from the smart card credential provider. As a response, the smart card credential provider provides each sign-in certificate to the sign-in UI, and corresponding sign-in tiles are displayed. The user selects a smart card-based sign-in certificate tile, and Windows displays a PIN dialog box.
4. The user enters the PIN, and then presses ENTER. The smart card credential provider encrypts the PIN.
5. The credential provider that resides in the LogonUI system collects the PIN. As part of packaging credentials in the smart card credential provider, the data is packaged in a KERB_CERTIFICATE_LOGON structure. The main contents of the KERB_CERTIFICATE_LOGON structure are the smart card PIN, CSP data (such as reader name and container name), user name, and domain name. User name is

required if the sign-in domain is not in the same forest because it enables a certificate to be mapped to multiple user accounts.

6. The credential provider wraps the data (such as the encrypted PIN, container name, reader name, and card key specification) and sends it back to LogonUI.

7. Winlogon presents the data from LogonUI to the LSA with the user information in LSALogonUser.

8. LSA calls the Kerberos authentication package (Kerberos SSP) to create a Kerberos authentication service request (KRB_AS_REQ), which containing a preauthenticator (as specified in RFC 4556: Public Key Cryptography for Initial Authentication in Kerberos (PKINIT)).

   If the authentication is performed by using a certificate that uses a digital signature, the preauthentication data consists of the user's public certificate and the certificate that is digitally signed with the corresponding private key.
   If the authentication is performed by using a certificate that uses key encipherment, the preauthentication data consists of the user's public certificate and the certificate that is encrypted with the corresponding private key.

9. To sign the request digitally (as per RFC 4556), a call is made to the corresponding CSP for a private key operation. Because the private key in this case is stored in a smart card, the smart card subsystem is called, and the necessary operation is completed. The result is sent back to the Kerberos security support provider (SSP).

10. The Kerberos SSP sends an authentication request for a ticket-granting-ticket (TGT) (per RFC 4556) to the Key Distribution Center (KDC) service that runs on a domain controller.

11. The KDC finds the user's account object in Active Directory Domain Services (AD DS), as detailed in Client certificate requirements and mappings, and uses the user's certificate to verify the signature.

12. The KDC validates the user's certificate (time, path, and revocation status) to ensure that the certificate is from a trusted source. The KDC uses CryptoAPI to build a certification path from the user's certificate to a root certification authority (CA) certificate that resides in the root store on the domain controller. The KDC then uses CryptoAPI to verify the digital signature on the signed authenticator that was included in the preauthentication data fields. The domain controller verifies the signature and uses the public key from the user's certificate to prove that the request originated from the owner of the private key that corresponds to the

public key. The KDC also verifies that the issuer is trusted and appears in the NTAUTH certificate store.

13. The KDC service retrieves user account information from AD DS. The KDC constructs a TGT, which is based on the user account information that it retrieves from AD DS. The TGT's authorization data fields include the user's security identifier (SID), the SIDs for universal and global domain groups to which the user belongs, and (in a multidomain environment) the SIDs for any universal groups of which the user is a member.

14. The domain controller returns the TGT to the client as part of the KRB_AS_REP response.

    > **Note** The KRB_AS_REP packet consists of:
    >
    > - Privilege attribute certificate (PAC)
    > - User's SID
    > - SIDs of any groups of which the user is a member
    > - A request for ticket-granting service (TGS)
    > - Preauthentication data

    TGT is encrypted with the master key of the KDC, and the session key is encrypted with a temporary key. This temporary key is derived based on RFC 4556. Using CryptoAPI, the temporary key is decrypted. As part of the decryption process, if the private key is on a smart card, a call is made to the smart card subsystem by using the specified CSP to extract the certificate corresponding to the user's public key. (Programmatic calls for the certificate include CryptAcquireContext, CryptSetProvParam with the PIN, CryptgetUserKey, and CryptGetKeyParam.) After the temporary key is obtained, the Kerberos SSP decrypts the session key.

15. The client validates the reply from the KDC (time, path, and revocation status). It first verifies the KDC's signature by the construction of a certification path from the KDC's certificate to a trusted root CA, and then it uses the KDC's public key to verify the reply signature.

16. Now that a TGT has been obtained, the client obtains a service ticket, which is used to sign in to the local computer.

17. With success, LSA stores the tickets and returns a success message to LSALogonUser. After this success message is issued, user profile for the device is selected and set, Group Policy refresh is instantiated, and other actions are performed.

18. After the user profile is loaded, the Certification Propagation Service (CertPropSvc) detects this event, reads the certificates from the smart card (including the root certificates), and then populates them into the user's certificate store (MYSTORE).

19. CSP to smart card resource manager communication happens on the LRPC Channel.

20. On successful authentication, certificates are propagated to the user's store asynchronously by the Certificate Propagation Service (CertPropSvc).

21. When the card is removed, certificates in the temporary secure cache store are removed. The Certificates are no longer available for sign-in, but they remain in the user's certificate store.

> **Note** A SID is created for each user or group at the time a user account or a group account is created within the local security accounts database or within AD DS. The SID never changes, even if the user or group account is renamed.

For more information about the Kerberos protocol, see Microsoft Kerberos.

By default, the KDC verifies that the client's certificate contains the smart card client authentication EKU szOID_KP_SMARTCARD_LOGON. However, if enabled, the **Allow certificates with no extended key usage certificate attribute** Group Policy setting allows the KDC to not require the SC-LOGON EKU. SC-LOGON EKU is not required for account mappings that are based on the public key.

KDC certificate

Active Directory Certificate Services provides three kinds of certificate templates:

- Domain controller
- Domain controller authentication
- Kerberos authentication

Depending on the configuration of the domain controller, one of these types of certificates is sent as a part of the AS_REP packet.

Client certificate requirements and mappings

Certificate requirements are listed by versions of the Windows operating system. Certificate mapping describes how information from the certificate is mapped to the

user account.

Certificate requirements

The smart card certificate has specific format requirements when it is used with Windows XP and earlier operating systems. You can enable any certificate to be visible for the smart card credential provider.

Component	Requirements for Windows 8.1, Windows 8, Windows 7, Windows Vista, Windows 10, and Windows 11	Requirements for Windows XP
CRL distribution point location	Not required	The location must be specified, online, and available, for example: [1]CRL Distribution Point Distribution Point Name: Full Name: URL=`<https://server1.contoso.com/CertEnroll/caname.crl>`
Key usage	Digital signature	Digital signature
Basic constraints	Not required	[Subject Type=End Entity, Path Length Constraint=None] (Optional)
extended key usage (EKU)	The smart card sign-in object identifier is not required. **Note** If an EKU is present, it must contain the smart card sign-in EKU. Certificates with no EKU can be used for sign-in.	- Client Authentication (1.3.6.1.5.5.7.3.2) The client authentication object identifier is required only if a certificate is used for SSL authentication. - Smart Card Sign-in (1.3.6.1.4.1.311.20.2.2)

Component	Requirements for Windows 8.1, Windows 8, Windows 7, Windows Vista, Windows 10, and Windows 11	Requirements for Windows XP
Subject alternative name	E-mail ID is not required for smart card sign-in.	Other Name: Principal Name=(UPN), for example: UPN=user1@contoso.com The UPN OtherName object identifier is 1.3.6.1.4.1.311.20.2.3. The UPN OtherName value must be an ASN1-encoded UTF8 string.
Subject	Not required	Distinguished name of user. This field is a mandatory extension, but the population of this field is optional.
Key exchange (AT_KEYEXCHANGE field)	Not required for smart card sign-in certificates if a Group Policy setting is enabled. (By default, Group Policy settings are not enabled.)	Not required
CRL	Not required	Not required
UPN	Not required	Not required
Notes	You can enable any certificate to be visible for the smart card credential provider.	There are two predefined types of private keys. These keys are Signature Only (AT_SIGNATURE) and Key Exchange (AT_KEYEXCHANGE). Smart card sign-in certificates must have a Key Exchange (AT_KEYEXCHANGE) private key type.

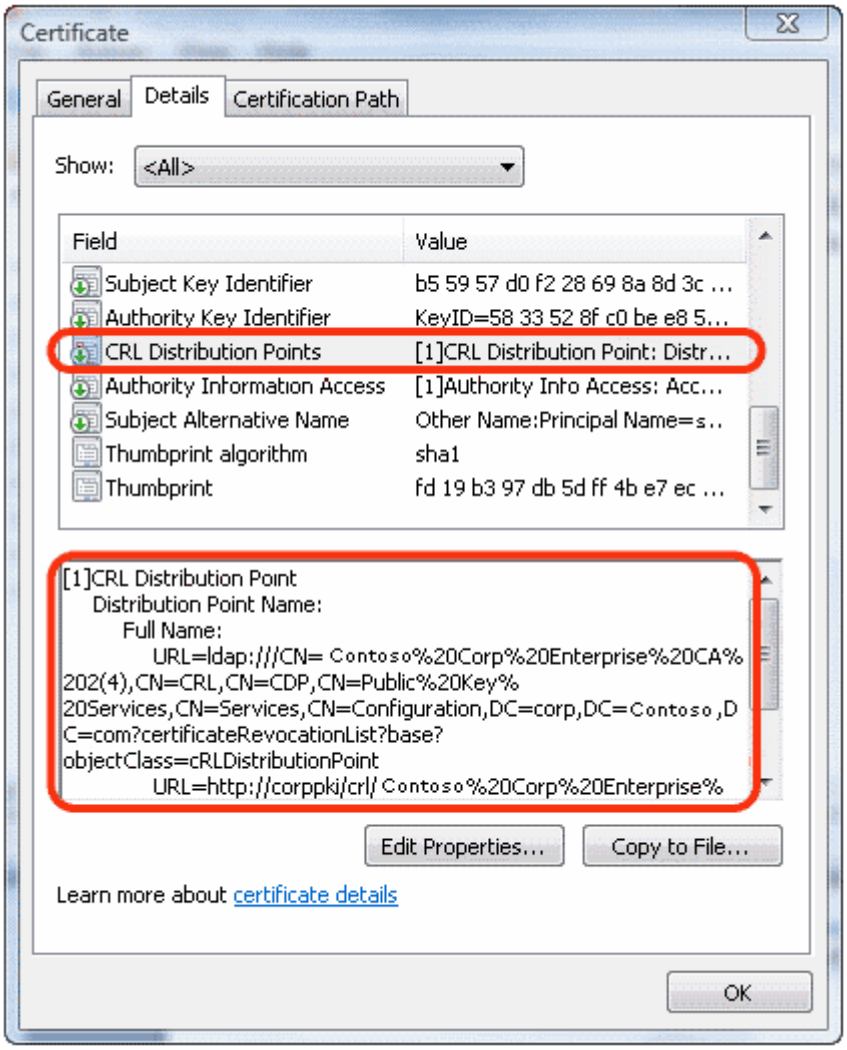
Client certificate mappings

Certificate mapping is based on the UPN that is contained in the subjectAltName (SAN) field of the certificate. Client certificates that do not contain information in the SAN field

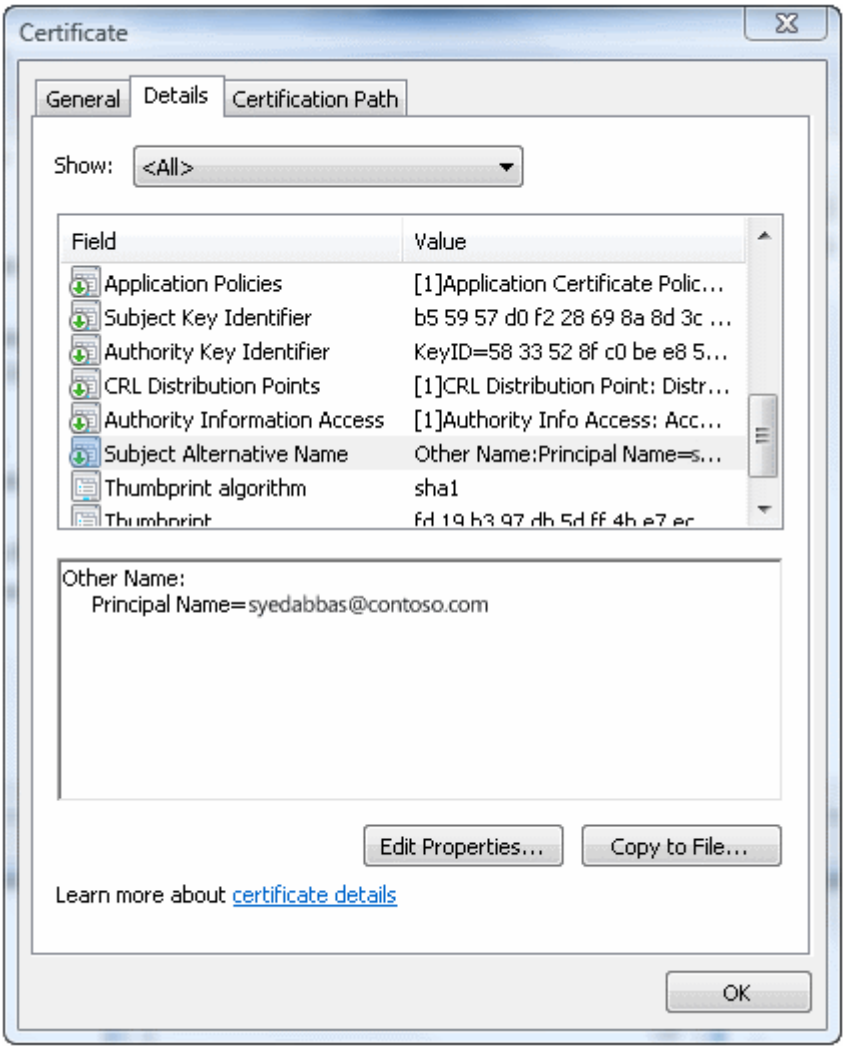
are also supported.

SSL/TLS can map certificates that do not have SAN, and the mapping is done by using the AltSecID attributes on client accounts. The X509 AltSecID, which is used by SSL/TLS client authentication is of the form "X509: <I>"<*Issuer Name*>"<S>"<*Subject Name*>. The <*Issuer Name*> and <*Subject Name*> are taken from the client certificate, with '\r' and '\n' replaced with ','.

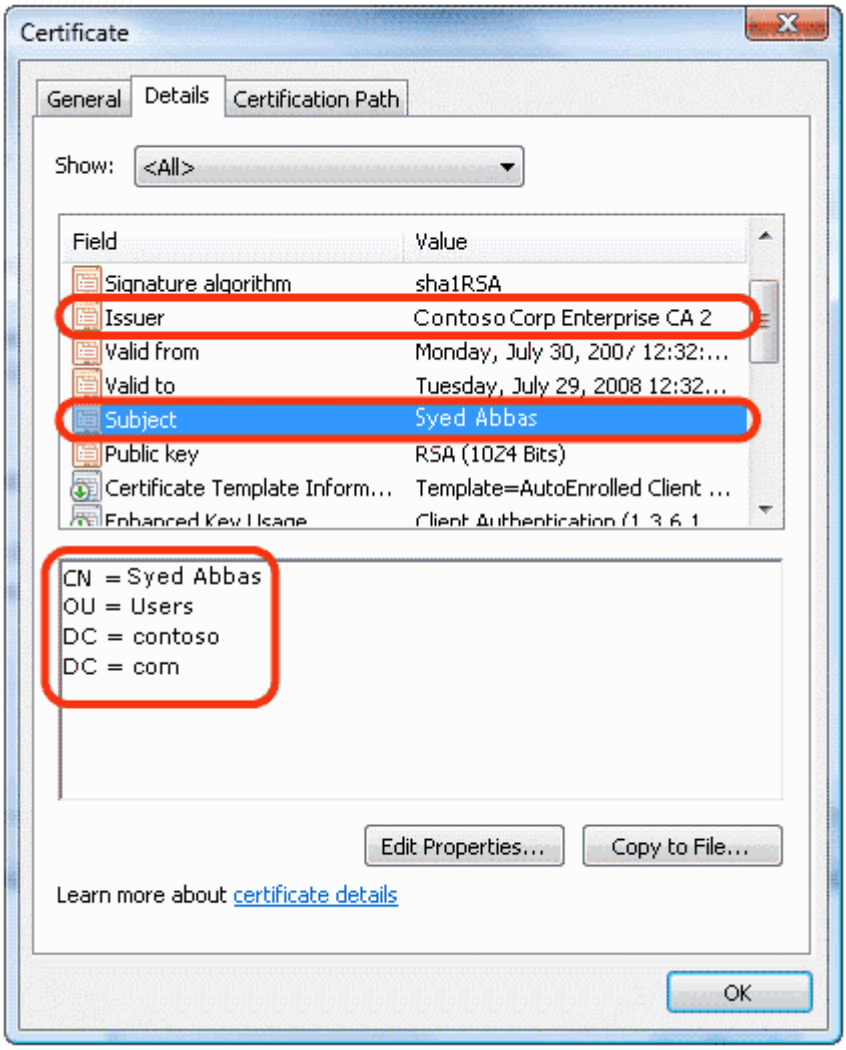
Certificate revocation list distribution points



UPN in Subject Alternative Name field

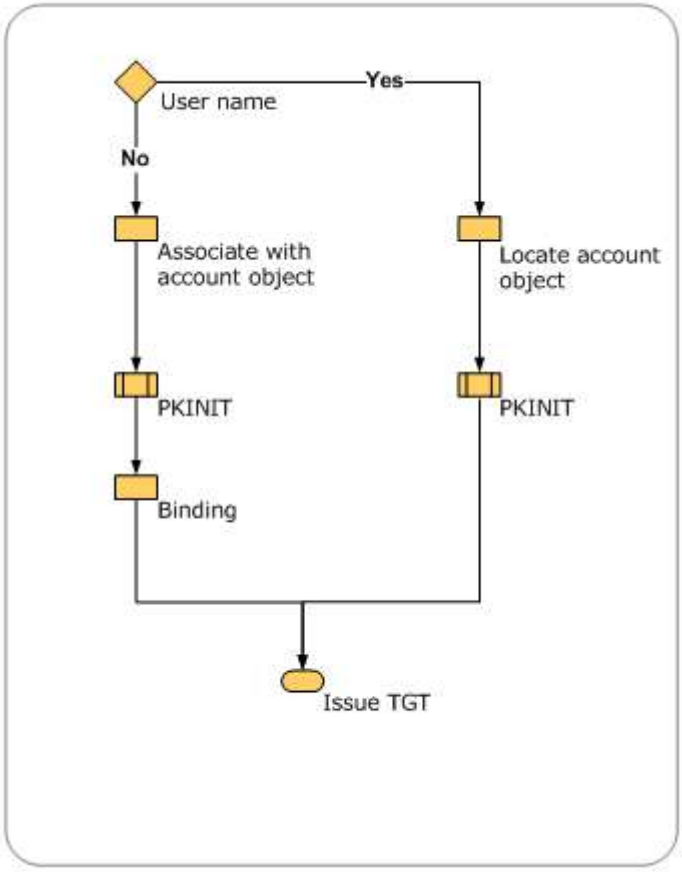


Subject and Issuer fields



This account mapping is supported by the KDC in addition to six other mapping methods. The following figure demonstrates a flow of user account mapping logic that is used by the KDC.

High-level flow of certificate processing for sign-in



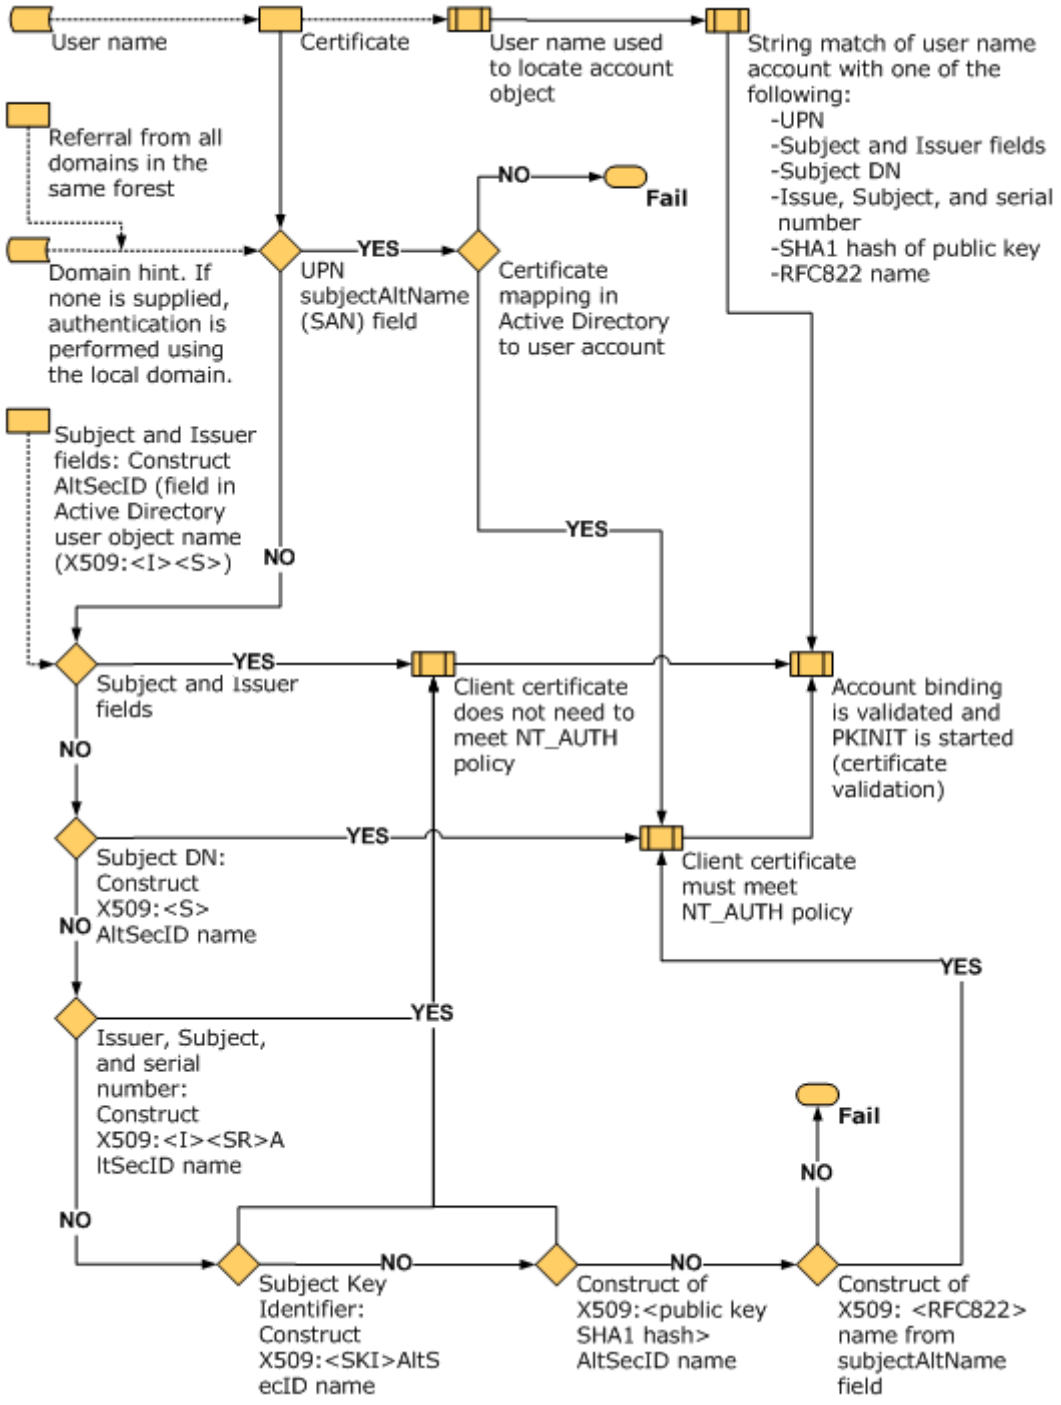
The certificate object is parsed to look for content to perform user account mapping.

- When a user name is provided with the certificate, the user name is used to locate the account object. This operation is the fastest, because string matching occurs.
- When only the certificate object is provided, a series of operations are performed to locate the user name to map the user name to an account object.
- When no domain information is available for authentication, the local domain is used by default. If any other domain is to be used for lookup, a domain name hint should be provided to perform the mapping and binding.

Mapping based on generic attributes is not possible because there is no generic API to retrieve attributes from a certificate. Currently, the first method that locates an account successfully stops the search. But a configuration error occurs if two methods map the same certificate to different user accounts when the client does not supply the client name through the mapping hints.

The following figure illustrates the process of mapping user accounts for sign-in in the directory by viewing various entries in the certificate.

Certificate processing logic



NT_AUTH policy is best described in the CERT_CHAIN_POLICY_NT_AUTH parameter section of the CertVerifyCertificateChainPolicy function. For more information, see CertVerifyCertificateChainPolicy.

Smart card sign-in for a single user with one certificate into multiple accounts

A single user certificate can be mapped to multiple accounts. For example, a user might be able to sign in to a user account and also to sign in as a domain administrator. The mapping is done by using the constructed AltSecID based on attributes from client

accounts. For information about how this mapping is evaluated, see Client certificate requirements and mappings.

> **Note** Because each account has a different user name, we recommend that you enable the **Allow user name hint** Group Policy setting (**X509HintsNeeded** registry key) to provide the optional fields that allow users to enter their user names and domain information to sign in.

Based on the information that is available in the certificate, the sign-in conditions are:

1. If no UPN is present in the certificate:
    a. Sign-in can occur in the local forest or in another forest if a single user with one certificate needs to sign in to different accounts.
    b. A hint must be supplied if mapping is not unique (for example, if multiple users are mapped to the same certificate).
2. If a UPN is present in the certificate:
    a. The certificate cannot be mapped to multiple users in the same forest.
    b. The certificate can be mapped to multiple users in different forests. For a user to sign in to other forests, an X509 hint must be supplied to the user.

Smart card sign-in for multiple users into a single account

A group of users might sign in to a single account (for example, an administrator account). For that account, user certificates are mapped so that they are enabled for sign-in.

Several distinct certificates can be mapped to a single account. For this to work properly, the certificate cannot have UPNs.

For example, if Certificate1 has CN=CNName1, Certificate2 has CN=User1, and Certificate3 has CN=User2, the AltSecID of these certificates can be mapped to a single account by using the Active Directory Users and Computers name mapping.

Smart card sign-in across forests

For account mapping to work across forests, particularly in cases where there is not enough information available on the certificate, the user might enter a hint in the form of a user name, such as *domain\user*, or a fully qualified UPN such as *user@contoso.com*.

> **Note** For the hint field to appear during smart card sign-in, the **Allow user name hint** Group Policy setting (**X509HintsNeeded** registry key) must be enabled on the client.

OCSP support for PKINIT

Online Certificate Status Protocol (OCSP), which is defined in RFC 2560, enables applications to obtain timely information about the revocation status of a certificate. Because OCSP responses are small and well bound, constrained clients might want to use OCSP to check the validity of the certificates for Kerberos on the KDC, to avoid transmission of large CRLs, and to save bandwidth on constrained networks. For information about CRL registry keys, see Smart Card Group Policy and Registry Settings.

The KDCs in Windows attempt to get OCSP responses and use them when available. This behavior cannot be disabled. CryptoAPI for OCSP caches OCSP responses and the status of the responses. The KDC supports only OCSP responses for the signer certificate.

Windows client computers attempt to request the OCSP responses and use them in the reply when they are available. This behavior cannot be disabled.

Smart card root certificate requirements for use with domain sign-in

For sign-in to work in a smart card-based domain, the smart card certificate must meet the following conditions:

- The KDC root certificate on the smart card must have an HTTP CRL distribution point listed in its certificate.
- The smart card sign-in certificate must have the HTTP CRL distribution point listed in its certificate.
- The CRL distribution point must have a valid CRL published and a delta CRL, if applicable, even if the CRL distribution point is empty.
- The smart card certificate must contain one of the following:

- A subject field that contains the DNS domain name in the distinguished name. If it does not, resolution to an appropriate domain fails, so Remote Desktop Services and the domain sign-in with the smart card fail.
- A UPN where the domain name resolves to the actual domain. For example, if the domain name is Engineering.Corp.Contoso, the UPN is username@engineering.corp.contoso.com. If any part of the domain name is omitted, the Kerberos client cannot find the appropriate domain.

Although the HTTP CRL distribution points are on by default in Windows Server 2008, subsequent versions of the Windows Server operating system do not include HTTP CRL distribution points. To allow smart card sign-in to a domain in these versions, do the following:

1. Enable HTTP CRL distribution points on the CA.
2. Restart the CA.
3. Reissue the KDC certificate.
4. Issue or reissue the smart card sign-in certificate.
5. Propagate the updated root certificate to the smart card that you want to use for the domain sign-in.

The workaround is to enable the **Allow user name hint** Group Policy setting (**X509HintsNeeded** registry key), which allows the user to supply a hint in the credentials user interface for domain sign-in.

If the client computer is not joined to the domain or if it is joined to a different domain, the client computer can resolve the server domain only by looking at the distinguished name on the certificate, not the UPN. For this scenario to work, the certificate requires a full subject, including DC=<*DomainControllerName*>, for domain name resolution.

To deploy root certificates on a smart card for the currently joined domain, you can use the following command:

certutil -scroots update

For more information about this option for the command-line tool, see -SCRoots.

See also

How Smart Card Sign-in Works in Windows

Smart Card and Remote Desktop Services

Article • 12/09/2022 • 5 minutes to read •

Applies to: ✅ Windows 10, ✅ Windows 11, ✅ Windows Server 2016, ✅ Windows Server 2019, ✅ Windows Server 2022

This topic for the IT professional describes the behavior of Remote Desktop Services when you implement smart card sign-in.

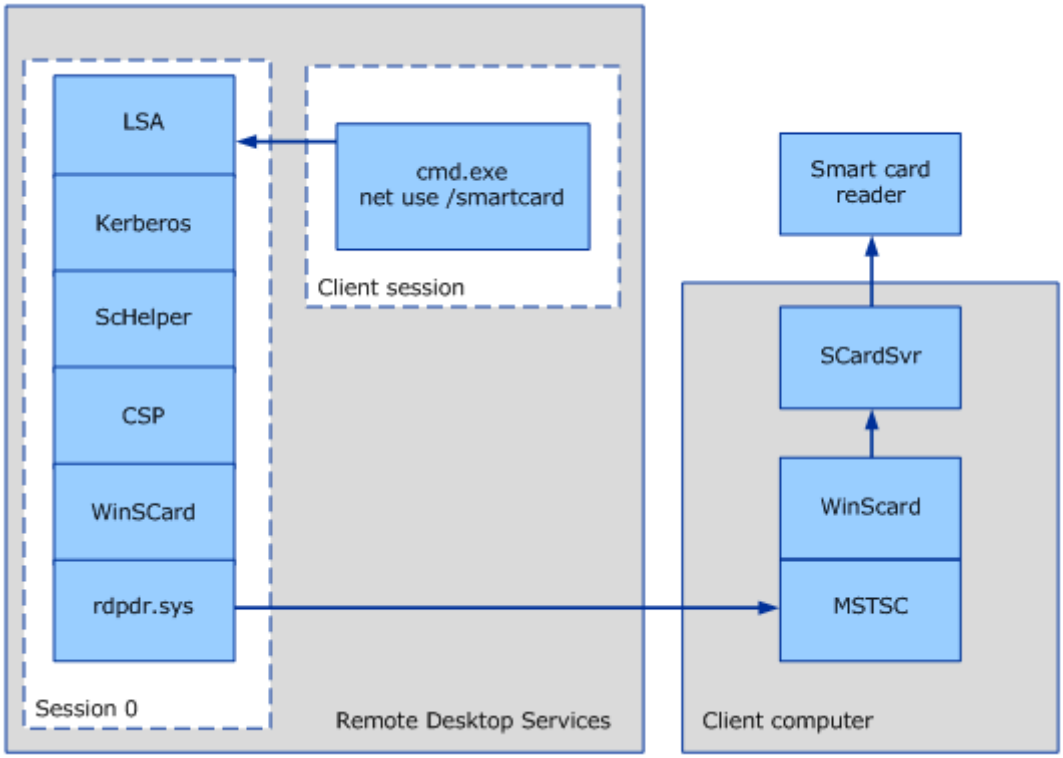
The content in this topic applies to the versions of Windows that are designated in the **Applies To** list at the beginning of this topic. In these versions, smart card redirection logic and **WinSCard** API are combined to support multiple redirected sessions into a single process.

Smart card support is required to enable many Remote Desktop Services scenarios. These include:

- Using Fast User Switching or Remote Desktop Services. A user is not able to establish a redirected smart card-based remote desktop connection. That is, the connect attempt is not successful in Fast User Switching or from a Remote Desktop Services session.
- Enabling Encrypting File System (EFS) to locate the user's smart card reader from the Local Security Authority (LSA) process in Fast User Switching or in a Remote Desktop Services session. If EFS is not able to locate the smart card reader or certificate, EFS cannot decrypt user files.

Remote Desktop Services redirection

In a Remote Desktop scenario, a user is using a remote server for running services, and the smart card is local to the computer that the user is using. In a smart card sign-in scenario, the smart card service on the remote server redirects to the smart card reader that is connected to the local computer where the user is trying to sign in.



Remote Desktop redirection

Notes about the redirection model:

1. This scenario is a remote sign-in session on a computer with Remote Desktop Services. In the remote session (labeled as "Client session"), the user runs **net use /smartcard**.

2. Arrows represent the flow of the PIN after the user types the PIN at the command prompt until it reaches the user's smart card in a smart card reader that is connected to the Remote Desktop Connection (RDC) client computer.

3. The authentication is performed by the LSA in session 0.

4. The CryptoAPI processing is performed in the LSA (Lsass.exe). This is possible because RDP redirector (rdpdr.sys) allows per-session, rather than per-process, context.

5. The WinScard and SCRedir components, which were separate modules in operating systems earlier than Windows Vista, are now included in one module. The ScHelper library is a CryptoAPI wrapper that is specific to the Kerberos protocol.

6. The redirection decision is made on a per smart card context basis, based on the session of the thread that performs the SCardEstablishContext call.

7. Changes to WinSCard.dll implementation were made in Windows Vista to improve smart card redirection.

RD Session Host server single sign-in experience

As a part of the Common Criteria compliance, the RDC client must be configurable to use Credential Manager to acquire and save the user's password or smart card PIN. Common Criteria compliance requires that applications not have direct access to the user's password or PIN.

Common Criteria compliance requires specifically that the password or PIN never leave the LSA unencrypted. A distributed scenario should allow the password or PIN to travel between one trusted LSA and another, and it cannot be unencrypted during transit.

When smart card-enabled single sign-in (SSO) is used for Remote Desktop Services sessions, users still need to sign in for every new Remote Desktop Services session. However, the user is not prompted for a PIN more than once to establish a Remote Desktop Services session. For example, after the user double-clicks a Microsoft Word document icon that resides on a remote computer, the user is prompted to enter a PIN. This PIN is sent by using a secure channel that the credential SSP has established. The PIN is routed back to the RDC client over the secure channel and sent to Winlogon. The user does not receive any additional prompts for the PIN, unless the PIN is incorrect or there are smart card-related failures.

Remote Desktop Services and smart card sign-in

Remote Desktop Services enables users to sign in with a smart card by entering a PIN on the RDC client computer and sending it to the RD Session Host server in a manner similar to authentication that is based on user name and password.

In addition, Group Policy settings that are specific to Remote Desktop Services need to be enabled for smart card-based sign-in.

To enable smart card sign-in to a Remote Desktop Session Host (RD Session Host) server, the Key Distribution Center (KDC) certificate must be present on the RDC client computer. If the computer is not in the same domain or workgroup, the following command can be used to deploy the certificate:

certutil -dspublish NTAuthCA "*DSCDPContainer*"

The *DSCDPContainer* Common Name (CN) is usually the name of the certification authority.

Example:

certutil -dspublish NTAuthCA <*CertFile*> **"CN=NTAuthCertificates,CN=Public Key Services,CN=Services,CN=Configuration,DC=engineering,DC=contoso,DC=com"**

For information about this option for the command-line tool, see -dsPublish.

Remote Desktop Services and smart card sign-in across domains

To enable remote access to resources in an enterprise, the root certificate for the domain must be provisioned on the smart card. From a computer that is joined to a domain, run the following command at the command line:

certutil -scroots update

For information about this option for the command-line tool, see -SCRoots.

For Remote Desktop Services across domains, the KDC certificate of the RD Session Host server must also be present in the client computer's NTAUTH store. To add the store, run the following command at the command line:

certutil -addstore -enterprise NTAUTH <*CertFile*>

Where <*CertFile*> is the root certificate of the KDC certificate issuer.

For information about this option for the command-line tool, see -addstore.

> **Note** If you use the credential SSP on computers running the supported versions of the operating system that are designated in the **Applies To** list at the beginning of this topic: To sign in with a smart card from a computer that is not joined to a domain, the smart card must contain the root certification of the domain controller. A public key infrastructure (PKI) secure channel cannot be established without the root certification of the domain controller.

Sign-in to Remote Desktop Services across a domain works only if the UPN in the certificate uses the following form: <*ClientName*>@<*DomainDNSName*>

The UPN in the certificate must include a domain that can be resolved. Otherwise, the Kerberos protocol cannot determine which domain to contact. You can resolve this issue

by enabling GPO X509 domain hints. For more information about this setting, see Smart Card Group Policy and Registry Settings.

See also

How Smart Card Sign-in Works in Windows

Smart Cards for Windows Service

Article • 12/09/2022 • 2 minutes to read •

Applies to: ✅ Windows 10, ✅ Windows 11, ✅ Windows Server 2016, ✅ Windows Server 2019, ✅ Windows Server 2022

This topic for the IT professional and smart card developers describes how the Smart Cards for Windows service (formerly called Smart Card Resource Manager) manages readers and application interactions.

The Smart Cards for Windows service provides the basic infrastructure for all other smart card components as it manages smart card readers and application interactions on the computer. It is fully compliant with the specifications set by the PC/SC Workgroup. For information about these specifications, see the PC/SC Workgroup Specifications website.

The Smart Cards for Windows service runs in the context of a local service, and it is implemented as a shared service of the services host (svchost) process. The Smart Cards for Windows service, Scardsvr, has the following service description:

PowerShell

```
<serviceData
    dependOnService="PlugPlay"
    description="@%SystemRoot%\System32\SCardSvr.dll,-5"
    displayName="@%SystemRoot%\System32\SCardSvr.dll,-1"
    errorControl="normal"
    group="SmartCardGroup"
    imagePath="%SystemRoot%\system32\svchost.exe -k
LocalServiceAndNoImpersonation"
    name="SCardSvr"
    objectName="NT AUTHORITY\LocalService"
    requiredPrivileges="SeCreateGlobalPrivilege,SeChangeNotifyPrivilege"
    sidType="unrestricted"
    start="demand"
    type="win32ShareProcess"
    >
  <failureActions resetPeriod="900">
      <actions>
         <action
             delay="120000"
             type="restartService"
         />
         <action
             delay="300000"
             type="restartService"
         />
         <action
              delay="0"
```

```
                type="none"
            />
        </actions>
    </failureActions>
    <securityDescriptor name="ServiceXSecurity"/>
</serviceData>

    <registryKeys buildFilter="">
        <registryKey
keyName="HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\SCardSvr\Param
eters">
        <registryValue
            name="ServiceDll"
            value="%SystemRoot%\System32\SCardSvr.dll"
            valueType="REG_EXPAND_SZ"
            />
        <registryValue
            name="ServiceMain"
            value="CalaisMain"
            valueType="REG_SZ"
            />
        <registryValue
            name="ServiceDllUnloadOnStop"
            value="1"
            valueType="REG_DWORD"
            />
        </registryKey>
    </registryKeys>
```

> **Note** For winscard.dll to be invoked as the proper class installer, the INF file for a smart card reader must specify the following for **Class** and **ClassGUID**:
>
> `Class=SmartCardReader`
>
> `ClassGuid={50DD5230-BA8A-11D1-BF5D-0000F805F530}`

By default, the service is configured for manual mode. Creators of smart card reader drivers must configure their INFs so that they start the service automatically and winscard.dll files call a predefined entry point to start the service during installation. The entry point is defined as part of the **SmartCardReader** class, and it is not called directly. If a device advertises itself as part of this class, the entry point is automatically invoked to start the service when the device is inserted. Using this method ensures that the service is enabled when it is needed, but it is also disabled for users who do not use smart cards.

When the service is started, it performs several functions:

1. It registers itself for service notifications.

2. It registers itself for Plug and Play (PnP) notifications related to device removal and additions.
3. It initializes its data cache and a global event that signals that the service has started.

> **Note** For smart card implementations, consider sending all communications in Windows operating systems with smart card readers through the Smart Cards for Windows service. This provides an interface to track, select, and communicate with all drivers that declare themselves members of the smart card reader device group.

The Smart Cards for Windows service categorizes each smart card reader slot as a unique reader, and each slot is also managed separately, regardless of the device's physical characteristics. The Smart Cards for Windows service handles the following high-level actions:

- Device introduction
- Reader initialization
- Notifying clients of new readers
- Serializing access to readers
- Smart card access
- Tunneling of reader-specific commands

See also

How Smart Card Sign-in Works in Windows

Certificate Propagation Service

Article • 12/09/2022 • 2 minutes to read •

Applies to: ✅ Windows 10, ✅ Windows 11, ✅ Windows Server 2016, ✅ Windows Server 2019, ✅ Windows Server 2022

This topic for the IT professional describes the certificate propagation service (CertPropSvc), which is used in smart card implementation.

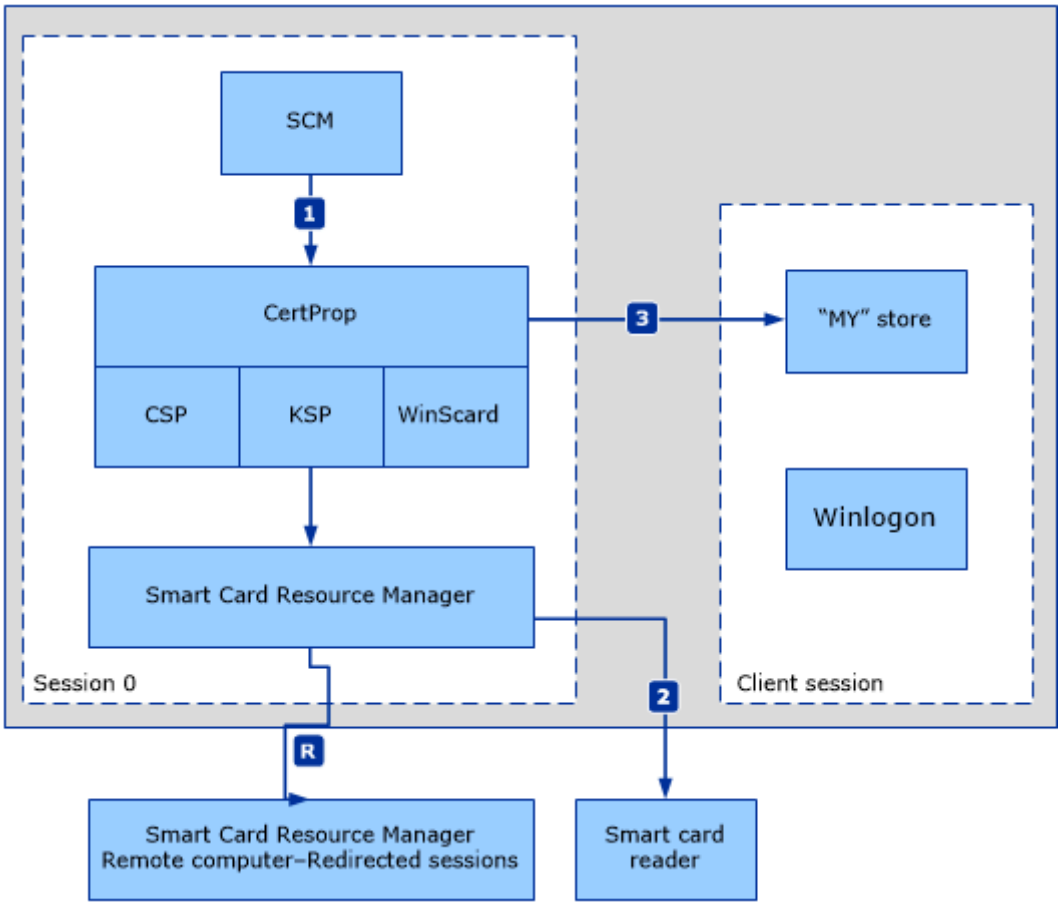
The certificate propagation service activates when a signed-in user inserts a smart card in a reader that is attached to the computer. This action causes the certificate to be read from the smart card. The certificates are then added to the user's Personal store. Certificate propagation service actions are controlled by using Group Policy. For more information, see Smart Card Group Policy and Registry Settings.

> **Note** The certificate propagation service must be running for smart card Plug and Play to work.

The following figure shows the flow of the certificate propagation service. The action begins when a signed-in user inserts a smart card.

1. The arrow labeled **1** indicates that the Service Control Manager (SCM) notifies the certificate propagation service (CertPropSvc) when a user signs in, and CertPropSvc begins to monitor the smart cards in the user session.
2. The arrow labeled **R** represents the possibility of a remote session and the use of smart card redirection.
3. The arrow labeled **2** indicates the certification to the reader.
4. The arrow labeled **3** indicates the access to the certificate store during the client session.

Certificate propagation service



1. A signed-in user inserts a smart card.
2. CertPropSvc is notified that a smart card was inserted.
3. CertPropSvc reads all certificates from all inserted smart cards. The certificates are written to the user's personal certificate store.

> **Note** The certificate propagation service is started as a Remote Desktop Services dependency.

Properties of the certificate propagation service include:

- CERT_STORE_ADD_REPLACE_EXISTING_INHERIT_PROPERTIES adds certificates to a user's Personal store.
- If the certificate has the CERT_ENROLLMENT_PROP_ID property (as defined by wincrypt.h), it filters empty requests and places them in the current user's request store, but it does not propagate them to the user's Personal store.
- The service does not propagate any computer certificates to a user's Personal store or propagate user certificates to a computer store.
- The service propagates certificates according to Group Policy options that are set, which may include:

- **Turn on certificate propagation from the smart card** specifies whether a user's certificate should be propagated.
- **Turn on root certificate propagation from smart card** specifies whether root certificates should be propagated.
- **Configure root certificate cleanup** specifies how root certificates are removed.

Root certificate propagation service

Root certificate propagation is responsible for the following smart card deployment scenarios when public key infrastructure (PKI) trust has not yet been established:

- Joining the domain
- Accessing a network remotely

In both cases, the computer is not joined to a domain, and therefore, trust is not being managed by Group Policy. However, the objective is to authenticate to a remote server, such as the domain controller. Root certificate propagation provides the ability to use the smart card to include the missing trust chain.

When the smart card is inserted, the certificate propagation service propagates any root certificates on the card to the trusted smart card root computer certificate stores. This process establishes a trust relationship with the enterprise resources. You may also use a subsequent cleanup action when the user's smart card is removed from the reader, or when the user signs out. This is configurable with Group Policy. For more information, see Smart Card Group Policy and Registry Settings.

For more information about root certificate requirements, see Smart card root certificate requirements for use with domain sign-in.

See also

How Smart Card Sign-in Works in Windows

Smart Card Removal Policy Service

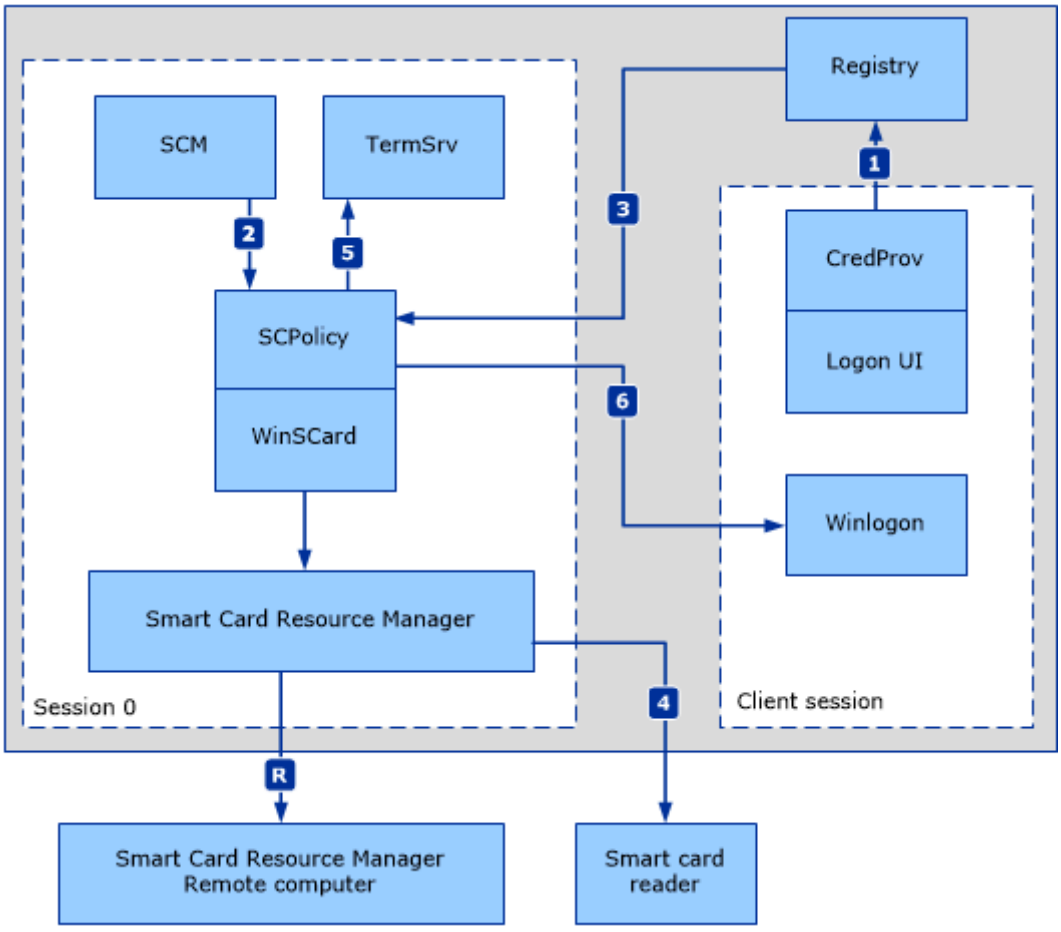
Article • 12/09/2022 • 2 minutes to read •

Applies to: ✅ Windows 10, ✅ Windows 11, ✅ Windows Server 2016, ✅ Windows Server 2019, ✅ Windows Server 2022

This topic for the IT professional describes the role of the removal policy service (ScPolicySvc) in smart card implementation.

The smart card removal policy service is applicable when a user has signed in with a smart card and then removes that smart card from the reader. The action that is performed when the smart card is removed is controlled by Group Policy settings. For more information, see Smart Card Group Policy and Registry Settings.

Smart card removal policy service



The numbers in the previous figure represent the following actions:

1. Winlogon is not directly involved in monitoring for smart card removal events. The sequence of steps that are involved when a smart card is removed begins with the smart card credential provider in the sign-in UI process. When a user successfully signs in with a smart card, the smart card credential provider captures the reader

name. This information is then stored in the registry with the session identifier where the sign-in was initiated.

2. The smart card resource manager service notifies the smart card removal policy service that a sign-in has occurred.

3. ScPolicySvc retrieves the smart card information that the smart card credential provider stored in the registry. This call is redirected if the user is in a remote session. If the smart card is removed, ScPolicySvc is notified.

4. ScPolicySvc calls Remote Desktop Services to take the appropriate action if the request is to sign out the user or to disconnect the user's session, which might result in data loss. If the setting is configured to lock the computer when the smart card is removed, ScPolicySvc sends a message to Winlogon to lock the computer.

See also

How Smart Card Sign-in Works in Windows

Smart Card Tools and Settings

Article • 12/09/2022 • 2 minutes to read •

Applies to: ✅ Windows 10, ✅ Windows 11, ✅ Windows Server 2016, ✅ Windows Server 2019, ✅ Windows Server 2022

This topic for the IT professional and smart card developer links to information about smart card debugging, settings, and events.

This section of the Smart Card Technical Reference contains information about the following:

- Smart Cards Debugging Information: Learn about tools and services in supported versions of Windows to help identify certificate issues.
- Smart Card Group Policy and Registry Settings: Learn about smart card-related Group Policy settings and registry keys that can be set on a per-computer basis, including how to edit and apply Group Policy settings to local or domain computers.
- Smart Card Events: Learn about events that can be used to manage smart cards in an organization, including how to monitor installation, use, and errors.

See also

Smart Card Technical Reference

Smart Card Troubleshooting

Article • 02/17/2023 • 5 minutes to read •

Applies to: ✅ Windows 10, ✅ Windows 11, ✅ Windows Server 2016, ✅ Windows Server 2019, ✅ Windows Server 2022

This article explains tools and services that smart card developers can use to help identify certificate issues with the smart card deployment.

Debugging and tracing smart card issues requires a variety of tools and approaches. The following sections provide guidance about tools and approaches you can use.

- Certutil
- Debugging and tracing using Windows software trace preprocessor (WPP)
- Kerberos protocol, Key Distribution Center (KDC), and NTLM debugging and tracing
- Smart Card service
- Smart card readers
- CryptoAPI 2.0 Diagnostics

Certutil

For a complete description of Certutil including examples that show how to use it, see Certutil [W2012].

List certificates available on the smart card

To list certificates that are available on the smart card, type `certutil -scinfo`.

> ⓘ **Note**
>
> Entering a PIN is not required for this operation. You can press ESC if you are prompted for a PIN.

Delete certificates on the smart card

Each certificate is enclosed in a container. When you delete a certificate on the smart card, you're deleting the container for the certificate.

To find the container value, type `certutil -scinfo`.

To delete a container, type **certutil -delkey -csp "Microsoft Base Smart Card Crypto Provider"** "<*ContainerValue*>".

Debugging and tracing using WPP

WPP simplifies tracing the operation of the trace provider. It provides a mechanism for the trace provider to log real-time binary messages. Logged messages can be converted to a human-readable trace of the operation. For more information, see Diagnostics with WPP - The NDIS blog.

Enable the trace

Using WPP, use one of the following commands to enable tracing:

- **tracelog.exe -kd -rt -start** <*FriendlyName*> **-guid #**<*GUID*> **-f .** <*LogFileName*>**.etl -flags** <*flags*> **-ft 1**
- **logman start** <*FriendlyName*> **-ets -p {**<*GUID*>**} -**<*Flags*> **-ft 1 -rt -o .** <*LogFileName*>***.etl -mode 0x00080000***

You can use the parameters in the following table.

Friendly name	GUID	Flags
`scardsvr`	13038e47-ffec-425d-bc69-5707708075fe	0xffff
`winscard`	3fce7c5f-fb3b-4bce-a9d8-55cc0ce1cf01	0xffff
`basecsp`	133a980d-035d-4e2d-b250-94577ad8fced	0x7
`scksp`	133a980d-035d-4e2d-b250-94577ad8fced	0x7
`msclmd`	fb36caf4-582b-4604-8841-9263574c4f2c	0x7
`credprov`	dba0e0e0-505a-4ab6-aa3f-22f6f743b480	0xffff
`certprop`	30eae751-411f-414c-988b-a8bfa8913f49	0xffff
`scfilter`	eed7f3c9-62ba-400e-a001-658869df9a91	0xffff
`wudfusbccid`	a3c09ba3-2f62-4be5-a50f-8278a646ac9d	0xffff

Examples

To enable tracing for the SCardSvr service:

- **tracelog.exe -kd -rt -start scardsvr -guid #13038e47-ffec-425d-bc69-5707708075fe -f .\scardsvr.etl -flags 0xffff -ft 1**
- **logman start scardsvr -ets -p {13038e47-ffec-425d-bc69-5707708075fe} 0xffff -ft 1 -rt -o .\scardsvr.etl -mode 0x00080000**

To enable tracing for scfilter.sys:

- **tracelog.exe -kd -rt -start scfilter -guid #eed7f3c9-62ba-400e-a001-658869df9a91 -f .\scfilter.etl -flags 0xffff -ft 1**

Stop the trace

Using WPP, use one of the following commands to stop the tracing:

- **tracelog.exe -stop** <*FriendlyName*>
- **logman -stop** <*FriendlyName*> **-ets**

Examples

To stop a trace:

- **tracelog.exe -stop scardsvr**
- **logman -stop scardsvr -ets**

Kerberos protocol, KDC, and NTLM debugging and tracing

You can use these resources to troubleshoot these protocols and the KDC:

- Kerberos and LDAP Troubleshooting Tips.
- Windows Driver Kit (WDK) and Debugging Tools for Windows (WinDbg). You can use the trace log tool in this SDK to debug Kerberos authentication failures.

To begin tracing, you can use `Tracelog`. Different components use different control GUIDs as explained in these examples. For more information, see Tracelog.

NTLM

To enable tracing for NTLM authentication, run the following command on the command line:

- **tracelog.exe -kd -rt -start ntlm -guid #5BBB6C18-AA45-49b1-A15F-085F7ED0AA90 -f .\ntlm.etl -flags 0x15003 -ft 1**

To stop tracing for NTLM authentication, run this command:

- **tracelog -stop ntlm**

Kerberos authentication

To enable tracing for Kerberos authentication, run this command:

- **tracelog.exe -kd -rt -start kerb -guid #6B510852-3583-4e2d-AFFE-A67F9F223438 -f .\kerb.etl -flags 0x43 -ft 1**

To stop tracing for Kerberos authentication, run this command:

- **tracelog.exe -stop kerb**

KDC

To enable tracing for the KDC, run the following command on the command line:

- **tracelog.exe -kd -rt -start kdc -guid #1BBA8B19-7F31-43c0-9643-6E911F79A06B -f .\kdc.etl -flags 0x803 -ft 1**

To stop tracing for the KDC, run the following command on the command line:

- **tracelog.exe -stop kdc**

To stop tracing from a remote computer, run this command: logman.exe -s <*ComputerName*>.

> ⓘ **Note**
>
> The default location for logman.exe is %systemroot%system32\. Use the **-s** option to supply a computer name.

Configure tracing with the registry

You can also configure tracing by editing the Kerberos registry values shown in the following table.

Element	Registry Key Setting
NTLM	HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\MSV1_0 Value name: NtLmInfoLevel Value type: DWORD Value data: c0015003
Kerberos	HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos Value name: LogToFile Value type: DWORD Value data: 00000001 HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters Value name: KerbDebugLevel Value type: DWORD Value data: c0000043 HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\Kerberos\Parameters Value name: LogToFile Value type: DWORD Value data: 00000001
KDC	HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Kdc Value name: KdcDebugLevel Value type: DWORD Value data: c0000803

If you used `Tracelog`, look for the following log file in your current directory: kerb.etl/kdc.etl/ntlm.etl.

If you used the registry key settings shown in the previous table, look for the trace log files in the following locations:

- NTLM: %systemroot%\tracing\msv1_0
- Kerberos: %systemroot%\tracing\kerberos
- KDC: %systemroot%\tracing\kdcsvc

To decode event trace files, you can use `Tracefmt` (tracefmt.exe). `Tracefmt` is a command-line tool that formats and displays trace messages from an event trace log file (.etl) or a real-time trace session. `Tracefmt` can display the messages in the Command Prompt window or save them in a text file. It is located in the \tools\tracing subdirectory of the Windows Driver Kit (WDK). For more information, see Tracefmt.

Smart Card service

The smart card resource manager service runs in the context of a local service. It's implemented as a shared service of the services host (svchost) process.

To check if Smart Card service is running

1. Press CTRL+ALT+DEL, and then select **Start Task Manager**.
2. In the **Windows Task Manager** dialog box, select the **Services** tab.
3. Select the **Name** column to sort the list alphabetically, and then type **s**.
4. In the **Name** column, look for **SCardSvr**, and then look under the **Status** column to see if the service is running or stopped.

To restart Smart Card service

1. Run as administrator at the command prompt.
2. If the **User Account Control** dialog box appears, confirm that the action it displays is what you want, and then select **Yes**.
3. At the command prompt, type `net stop SCardSvr`.
4. At the command prompt, type `net start SCardSvr`.

You can use the following command at the command prompt to check whether the service is running: `sc queryex scardsvr`.

The following code sample is an example output from this command:

Console

```
SERVICE_NAME: scardsvr
    TYPE        : 20 WIN32_SHARE_PROCESS
    STATE       : 4 RUNNING
                  (STOPPABLE, NOT_PAUSABLE, ACCEPTS_SHUTDOWN)
    WIN32_EXIT_CODE  : 0 (0x0)
    SERVICE_EXIT_CODE : 0 (0x0)
    CHECKPOINT     : 0x0
    WAIT_HINT     : 0x0
    PID        : 1320
    FLAGS       :
C:\>
```

Smart card readers

As with any device connected to a computer, Device Manager can be used to view properties and begin the debug process.

To check if smart card reader is working

1. Navigate to **Computer**.
2. Right-click **Computer**, and then select **Properties**.
3. Under **Tasks**, select **Device Manager**.
4. In Device Manager, expand **Smart card readers**, select the name of the smart card reader you want to check, and then select **Properties**.

> **Note**
>
> If the smart card reader is not listed in Device Manager, in the **Action** menu, select **Scan for hardware changes**.

CryptoAPI 2.0 Diagnostics

CryptoAPI 2.0 Diagnostics is available in Windows versions that support CryptoAPI 2.0 and can help you troubleshoot public key infrastructure (PKI) issues.

CryptoAPI 2.0 Diagnostics logs events in the Windows event log. The logs contain detailed information about certificate chain validation, certificate store operations, and signature verification. This information makes it easier to identify the causes of issues and reduces the time required for diagnosis.

For more information about CryptoAPI 2.0 Diagnostics, see Troubleshooting an Enterprise PKI.

See also

Smart Card Technical Reference

Smart Card Group Policy and Registry Settings

Article • 01/24/2023 • 20 minutes to read •

Applies to: ✅ Windows 10, ✅ Windows 11, ✅ Windows Server 2016, ✅ Windows Server 2019, ✅ Windows Server 2022

This article for IT professionals and smart card developers describes the Group Policy settings, registry key settings, local security policy settings, and credential delegation policy settings that are available for configuring smart cards.

The following sections and tables list the smart card-related Group Policy settings and registry keys that can be set on a per-computer basis. If you use domain Group Policy Objects (GPOs), you can edit and apply Group Policy settings to local or domain computers.

- Primary Group Policy settings for smart cards
  - Allow certificates with no extended key usage certificate attribute
  - Allow ECC certificates to be used for logon and authentication
  - Allow Integrated Unblock screen to be displayed at the time of logon
  - Allow signature keys valid for Logon
  - Allow time invalid certificates
  - Allow user name hint
  - Configure root certificate clean up
  - Display string when smart card is blocked
  - Filter duplicate logon certificates
  - Force the reading of all certificates from the smart card
  - Notify user of successful smart card driver installation
  - Prevent plaintext PINs from being returned by Credential Manager
  - Reverse the subject name stored in a certificate when displaying
  - Turn on certificate propagation from smart card
  - Turn on root certificate propagation from smart card
  - Turn on Smart Card Plug and Play service
- Base CSP and Smart Card KSP registry keys
- CRL checking registry keys
- Additional smart card Group Policy settings and registry keys

Primary Group Policy settings for smart cards

The following smart card Group Policy settings are in Computer Configuration\Administrative Templates\Windows Components\Smart Card.

The registry keys are in the following locations:

- **HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\ScPnP\EnableScPnP**
- **HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\SmartCardCredentialProvider**

- **HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\CertProp**

> ⓘ **Note**
>
> Smart card reader registry information is in **HKEY_LOCAL_MACHINE\Software\Microsoft\Cryptography\Calais\Readers**.
> Smart card registry information is in **HKEY_LOCAL_MACHINE\Software\Microsoft\Cryptography\Calais\SmartCards**.

The following table lists the default values for these GPO settings. Variations are documented under the policy descriptions in this article.

Server type or GPO	Default value
Default Domain Policy	Not configured
Default Domain Controller Policy	Not configured
Stand-Alone Server Default Settings	Not configured
Domain Controller Effective Default Settings	Disabled
Member Server Effective Default Settings	Disabled
Client Computer Effective Default Settings	Disabled

Allow certificates with no extended key usage certificate attribute

You can use this policy setting to allow certificates without an extended key usage (EKU) set to be used for sign-in.

> ⓘ **Note**
>
> extended key usage certificate attribute is also known as extended key usage.
>
> In versions of Windows before Windows Vista, smart card certificates that are used to sign in require an EKU extension with a smart card logon object identifier. This policy setting can be used to modify that restriction.

When this policy setting is turned on, certificates with the following attributes can also be used to sign in with a smart card:

- Certificates with no EKU
- Certificates with an All Purpose EKU
- Certificates with a Client Authentication EKU

When this policy setting isn't turned on, only certificates that contain the smart card logon object identifier can be used to sign in with a smart card.

Item	Description
Registry key	AllowCertificatesWithNoEKU
Default values	No changes per operating system versions Disabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None
Notes and resources	

Allow ECC certificates to be used for logon and authentication

You can use this policy setting to control whether elliptic curve cryptography (ECC) certificates on a smart card can be used to sign in to a domain.

When this setting is turned on, ECC certificates on a smart card can be used to sign in to a domain.

When this setting isn't turned on, ECC certificates on a smart card can't be used to sign in to a domain.

Item	Description
Registry key	**EnumerateECCCerts**
Default values	No changes per operating system versions Disabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None
Notes and resources	This policy setting only affects a user's ability to sign in to a domain. ECC certificates on a smart card that are used for other applications, such as document signing, aren't affected by this policy setting. If you use an ECDSA key to sign in, you must also have an associated ECDH key to permit sign in when you're not connected to the network.

Allow Integrated Unblock screen to be displayed at the time of logon

You can use this policy setting to determine whether the integrated unblock feature is available in the sign-in user interface (UI). The feature was introduced as a standard feature in the Credential Security Support Provider in Windows Vista.

When this setting is turned on, the integrated unblock feature is available.

When this setting isn't turned on, the feature is not available.

Item	Description
Registry key	**AllowIntegratedUnblock**
Default values	No changes per operating system versions Disabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None
Notes and resources	To use the integrated unblock feature, the smart card must support it. Check with the hardware manufacturer to verify that the smart card supports this feature. You can create a custom message that the user sees when the smart card is blocked by configuring the policy setting Display string when smart card is blocked.

Allow signature keys valid for Logon

You can use this policy setting to allow signature key–based certificates to be enumerated and available for sign-in.

When this setting is turned on, any certificates that are available on the smart card with a signature-only key are listed on the sign-in screen.

When this setting isn't turned on, certificates available on the smart card with a signature-only key aren't listed on the sign-in screen.

Item	Description

Item	Description
Registry key	**AllowSignatureOnlyKeys**
Default values	No changes per operating system versions Disabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None
Notes and resources	

Allow time invalid certificates

You can use this policy setting to permit certificates that are expired or not yet valid to be displayed for sign-in.

> **Note**
>
> Before Windows Vista, certificates were required to contain a valid time and to not expire. For a certificate to be used, it must be accepted by the domain controller. This policy setting only controls which certificates are displayed on the client computer.

When this setting is turned on, certificates are listed on the sign-in screen whether they have an invalid time, or their time validity has expired.

When this policy setting isn't turned on, certificates that are expired or not yet valid aren't listed on the sign-in screen.

Item	Description
Registry key	**AllowTimeInvalidCertificates**
Default values	No changes per operating system versions Disabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None
Notes and resources	

Allow user name hint

You can use this policy setting to determine whether an optional field appears during sign-in and provides a subsequent elevation process where users can enter their username or username and domain, which associates a certificate with the user.

When this policy setting is turned on, users see an optional field where they can enter their username or username and domain.

When this policy setting isn't turned on, users don't see this optional field.

Item	Description
Registry key	**X509HintsNeeded**
Default values	No changes per operating system versions Disabled and not configured are equivalent

Item	Description
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None
Notes and resources	

Configure root certificate clean-up

You can use this policy setting to manage the cleanup behavior of root certificates. Certificates are verified by using a trust chain, and the trust anchor for the digital certificate is the Root Certification Authority (CA). A CA can issue multiple certificates with the root certificate as the top certificate of the tree structure. A private key is used to sign other certificates. This creates an inherited trustworthiness for all certificates immediately under the root certificate.

When this policy setting is turned on, you can set the following cleanup options:

- **No cleanup**. When the user signs out or removes the smart card, the root certificates used during their session persist on the computer.
- **Clean up certificates on smart card removal**. When the smart card is removed, the root certificates are removed.
- **Clean up certificates on log off**. When the user signs out of Windows, the root certificates are removed.

When this policy setting isn't turned on, root certificates are automatically removed when the user signs out of Windows.

Item	Description
Registry key	**RootCertificateCleanupOption**
Default values	No changes per operating system versions Disabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None
Notes and resources	

Display string when smart card is blocked

You can use this policy setting to change the default message that a user sees if their smart card is blocked.

When this policy setting is turned on, you can create and manage the displayed message that the user sees when a smart card is blocked.

When this policy setting isn't turned on (and the integrated unblock feature is also enabled), the user sees the system's default message when the smart card is blocked.

Item	Description
Registry key	**IntegratedUnblockPromptString**
Default values	No changes per operating system versions Disabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: This policy setting is only effective when the Allow Integrated Unblock screen to be displayed at the time of logon policy is enabled.
Notes and resources	

Filter duplicate logon certificates

You can use this policy setting to configure which valid sign-in certificates are displayed.

> ⓘ **Note**
>
> During the certificate renewal period, a user's smart card can have multiple valid sign-in certificates issued from the same certificate template, which can cause confusion about which certificate to select. This behavior can occur when a certificate is renewed and the old certificate has not expired yet.
>
> If two certificates are issued from the same template with the same major version and they are for the same user (this is determined by their UPN), they are determined to be the same.

When this policy setting is turned on, filtering occurs so that the user can select from only the most current valid certificates.

If this policy setting isn't turned on, all the certificates are displayed to the user.

This policy setting is applied to the computer after the Allow time invalid certificates policy setting is applied.

Item	Description
Registry key	**FilterDuplicateCerts**
Default values	No changes per operating system versions Disabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None
Notes and resources	If there are two or more of the same certificates on a smart card and this policy setting is enabled, the certificate that is used to sign in to computers running Windows 2000, Windows XP, or Windows Server 2003 will be displayed. Otherwise, the certificate with the most distant expiration time will be displayed.

Force the reading of all certificates from the smart card

You can use this policy setting to manage how Windows reads all certificates from the smart card for sign-in. During sign-in, Windows reads only the default certificate from the smart card unless it supports retrieval of all certificates in a single call. This policy setting forces Windows to read all the certificates from the smart card.

When this policy setting is turned on, Windows attempts to read all certificates from the smart card, regardless of the CSP feature set.

When this policy isn't turned on, Windows attempts to read only the default certificate from smart cards that don't support retrieval of all certificates in a single call. Certificates other than the default aren't available for sign-in.

Item	Description
Registry key	**ForceReadingAllCertificates**
Default values	No changes per operating system versions Disabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None **Important**: Enabling this policy setting can adversely impact performance during the sign-in process in certain situations.

Item	Description
Notes and resources	Contact the smart card vendor to determine if your smart card and associated CSP support the required behavior.

Notify user of successful smart card driver installation

You can use this policy setting to control whether the user sees a confirmation message when a smart card device driver is installed.

When this policy setting is turned on, the user sees a confirmation message when a smart card device driver is installed.

When this setting isn't turned on, the user doesn't see a smart card device driver installation message.

Item	Description
Registry key	**ScPnPNotification**
Default values	No changes per operating system versions Disabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None
Notes and resources	This policy setting applies only to smart card drivers that have passed the Windows Hardware Quality Labs (WHQL) testing process.

Prevent plaintext PINs from being returned by Credential Manager

You can use this policy setting to prevent Credential Manager from returning plaintext PINs.

> **Note**
>
> Credential Manager is controlled by the user on the local computer, and it stores credentials from supported browsers and Windows applications. Credentials are saved in special encrypted folders on the computer under the user's profile.

When this policy setting is turned on, Credential Manager doesn't return a plaintext PIN.

When this setting isn't turned on, Credential Manager can return plaintext PINs.

Item	Description
Registry key	**DisallowPlaintextPin**
Default values	No changes per operating system versions Disabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None
Notes and resources	If this policy setting is enabled, some smart cards might not work in computers running Windows. Consult the smart card manufacturer to determine whether this policy setting should be enabled.

Reverse the subject name stored in a certificate when displaying

You can use this policy setting to control the way the subject name appears during sign-in.

> **Note**
>
> To help users distinguish one certificate from another, the user principal name (UPN) and the common name are displayed by default. For example, when this setting is enabled, if the certificate subject is CN=User1, OU=Users, DN=example, DN=com and the UPN is user1@example.com, "User1" is displayed with "user1@example.com." If the UPN is not present, the entire subject name is displayed. This setting controls the appearance of that subject name, and it might need to be adjusted for your organization.

When this policy setting is turned on, the subject name during sign-in appears reversed from the way that it's stored in the certificate.

When this policy setting isn't turned on, the subject name appears the same as it's stored in the certificate.

Item	Description
Registry key	**ReverseSubject**
Default values	No changes per operating system versions Disabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None
Notes and resources	

Turn on certificate propagation from smart card

You can use this policy setting to manage the certificate propagation that occurs when a smart card is inserted.

> **Note**
>
> The certificate propagation service applies when a signed-in user inserts a smart card in a reader that is attached to the computer. This action causes the certificate to be read from the smart card. The certificates are then added to the user's Personal store.

When this policy setting is turned on, certificate propagation occurs when the user inserts the smart card.

When this policy setting is turned off, certificate propagation doesn't occur, and the certificates aren't available to applications, like Outlook.

Item	Description
Registry key	**CertPropEnabled**
Default values	No changes per operating system versions Enabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: This policy setting must be enabled to allow the Turn on root certificate propagation from smart card setting to work when it is enabled.
Notes and resources	

Turn on root certificate propagation from smart card

You can use this policy setting to manage the root certificate propagation that occurs when a smart card is inserted.

> **Note**
>
> The certificate propagation service applies when a signed-in user inserts a smart card in a reader that is attached to the computer. This action causes the certificate to be read from the smart card. The certificates are then added to the user's Personal store.

When this policy setting is turned on, root certificate propagation occurs when the user inserts the smart card.

When this policy setting isn't turned on, root certificate propagation doesn't occur when the user inserts the smart card.

Item	Description
Registry key	**EnableRootCertificate Propagation**
Default values	No changes per operating system versions Enabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: For this policy setting to work, the Turn on certificate propagation from smart card policy setting must also be enabled.
Notes and resources	

Turn on Smart Card Plug and Play service

You can use this policy setting to control whether Smart Card Plug and Play is enabled.

> **Note**
>
> Your users can use smart cards from vendors who have published their drivers through Windows Update without needing special middleware. These drivers will be downloaded in the same way as drivers for other devices in Windows. If an appropriate driver isn't available from Windows Update, a PIV-compliant mini driver that's included with any of the supported versions of Windows is used for these cards.

When this policy setting is turned on, the system attempts to install a smart card device driver the first time a smart card is inserted in a smart card reader.

When this policy setting isn't turned on, a device driver isn't installed when a smart card is inserted in a smart card reader.

Item	Description
Registry key	**EnableScPnP**
Default values	No changes per operating system versions Enabled and not configured are equivalent
Policy management	Restart requirement: None Sign off requirement: None Policy conflicts: None
Notes and resources	This policy setting applies only to smart card drivers that have passed the Windows Hardware Quality Labs (WHQL) testing process.

Base CSP and Smart Card KSP registry keys

The following registry keys can be configured for the base cryptography service provider (CSP) and the smart card key storage provider (KSP). The following tables list the keys. All keys use the DWORD type.

The registry keys for the Base CSP are in the registry in **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Cryptography\Defaults\Provider\Microsoft Base Smart Card Crypto Provider**.

The registry keys for the smart card KSP are in **HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Control\Cryptography\Providers\Microsoft Smart Card Key Storage Provider**.

Registry keys for the base CSP and smart card KSP

Registry Key	Description
AllowPrivateExchangeKeyImport	A non-zero value allows RSA exchange (for example, encryption) private keys to be imported for use in key archival scenarios. Default value: 00000000
AllowPrivateSignatureKeyImport	A non-zero value allows RSA signature private keys to be imported for use in key archival scenarios. Default value: 00000000
DefaultPrivateKeyLenBits	Defines the default length for private keys, if desired. Default value: 00000400 Default key generation parameter: 1024-bit keys
RequireOnCardPrivateKeyGen	This key sets the flag that requires on-card private key generation (default). If this value is set, a key generated on a host can be imported into the smart card. This is used for smart cards that don't support on-card key generation or where key escrow is required. Default value: 00000000
TransactionTimeoutMilliseconds	Default timeout values allow you to specify whether transactions that take an excessive amount of time will fail. Default value: 000005dc The default timeout for holding transactions to the smart card is 1.5 seconds.

Additional registry keys for the smart card KSP

Registry Key	Description
AllowPrivateECDHEKeyImport	This value allows Ephemeral Elliptic Curve Diffie-Hellman (ECDHE) private keys to be imported for use in key archival scenarios. Default value: 00000000
AllowPrivateECDSAKeyImport	This value allows Elliptic Curve Digital Signature Algorithm (ECDSA) private keys to be imported for use in key archival scenarios. Default value: 00000000

CRL checking registry keys

The following table lists the keys and the corresponding values to turn off certificate revocation list (CRL) checking at the Key Distribution Center (KDC) or client. To manage CRL checking, you must configure settings for both the KDC and the client.

CRL checking registry keys

Registry Key	Details
HKEY_LOCAL_MACHINE\SYSTEM\CCS\Services\Kdc\UseCachedCRLOnlyAndIgnoreRevocationUnknownErrors	Type = DWORD Value = 1

Registry Key	Details
HKEY_LOCAL_MACHINE\SYSTEM\CCS\Control\LSA\Kerberos\Parameters\UseCachedCRLOnlyAndIgnoreRevocationUnknownErrors	Type = DWORD Value = 1

Additional smart card Group Policy settings and registry keys

In a smart card deployment, additional Group Policy settings can be used to enhance ease-of-use or security. Two of these policy settings that can complement a smart card deployment are:

- Turning off delegation for computers
- Interactive logon: Do not require CTRL+ALT+DEL (not recommended)

The following smart card-related Group Policy settings are in Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options.

Local security policy settings

Group Policy setting and registry key	**Default**	**Description**
Interactive logon: Require smart card **scforceoption**	Disabled	This security policy setting requires users to sign in to a computer by using a smart card. **Enabled** Users can sign in to the computer only by using a smart card. **Disabled** Users can sign in to the computer by using any method.
Interactive logon: Smart card removal behavior **scremoveoption**	This policy setting isn't defined, which means that the system treats it as **No Action**.	This setting determines what happens when the smart card for a signed-in user is removed from the smart card reader. The options are: **No Action** **Lock Workstation**: The workstation is locked when the smart card is removed, so users can leave the area, take their smart card with them, and still maintain a protected session. **Force Logoff**: The user is automatically signed out when the smart card is removed. **Disconnect if a Remote Desktop Services session**: Removal of the smart card disconnects the session without signing out the user. The user can reinsert the smart card and resume the session later, or at another computer that's equipped with a smart card reader, without having to sign in again. If the session is local, this policy setting functions identically to the **Lock Workstation** option. **Note**: In earlier versions of Windows Server, Remote Desktop Services was called Terminal Services.

From the Local Security Policy Editor (secpol.msc), you can edit and apply system policies to manage credential delegation for local or domain computers.

The following smart card-related Group Policy settings are in Computer Configuration\Administrative Templates\System\Credentials Delegation.

Registry keys are in **HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Control\Lsa\Credssp\PolicyDefaults**.

> ⓘ Note
>
> In the following table, fresh credentials are those that you are prompted for when running an application.

Credential delegation policy settings

Group Policy setting and registry key	**Default**	**Description**

Group Policy setting and registry key	Default	Description
Allow Delegating Fresh Credentials **AllowFreshCredentials**	Not configured	This policy setting applies: When server authentication was achieved through a trusted X509 certificate or Kerberos protocol. To applications that use the CredSSP component (for example, Remote Desktop Services). **Enabled**: You can specify the servers where the user's fresh credentials can be delegated. **Not configured**: After proper mutual authentication, delegation of fresh credentials is permitted to Remote Desktop Services running on any computer. **Disabled**: Delegation of fresh credentials to any computer isn't permitted. **Note**: This policy setting can be set to one or more service principal names (SPNs). The SPN represents the target server where the user credentials can be delegated. A single wildcard character is permitted when specifying the SPN, for example: Use *TERMSRV/** for Remote Desktop Session Host (RD Session Host) running on any computer. Use *TERMSRV/host.humanresources.fabrikam.com* for RD Session Host running on the host.humanresources.fabrikam.com computer. Use *TERMSRV/*.humanresources.fabrikam.com* for RD Session Host running on all computers in .humanresources.fabrikam.com
Allow Delegating Fresh Credentials with NTLM-only Server Authentication **AllowFreshCredentialsWhenNTLMOnly**	Not configured	This policy setting applies: When server authentication was achieved by using NTLM. To applications that use the CredSSP component (for example, Remote Desktop). **Enabled**: You can specify the servers where the user's fresh credentials can be delegated. **Not configured**: After proper mutual authentication, delegation of fresh credentials is permitted to RD Session Host running on any computer (TERMSRV/*). **Disabled**: Delegation of fresh credentials isn't permitted to any computer. **Note**: This policy setting can be set to one or more SPNs. The SPN represents the target server where the user credentials can be delegated. A single wildcard character (*) is permitted when specifying the SPN. See the **Allow Delegating Fresh Credentials** policy setting description for examples.
Deny Delegating Fresh Credentials **DenyFreshCredentials**	Not configured	This policy setting applies to applications that use the CredSSP component (for example, Remote Desktop). **Enabled**: You can specify the servers where the user's fresh credentials can't be delegated. **Disabled** or **Not configured**: A server is not specified. **Note**: This policy setting can be set to one or more SPNs. The SPN represents the target server where the user credentials can't be delegated. A single wildcard character (*) is permitted when specifying the SPN. For examples, see the "Allow delegating fresh credentials" policy setting.

If you're using Remote Desktop Services with smart card logon, you can't delegate default and saved credentials. The registry keys in the following table, which are at **HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Control\Lsa\Credssp\PolicyDefaults**, and the corresponding Group Policy settings are ignored.

Registry key	Corresponding Group Policy setting
AllowDefaultCredentials	Allow Delegating Default Credentials
AllowDefaultCredentialsWhenNTLMOnly	Allow Delegating Default Credentials with NTLM-only Server Authentication

Registry key	Corresponding Group Policy setting
AllowSavedCredentials	Allow Delegating Saved Credentials
AllowSavedCredentialsWhenNTLMOnly	Allow Delegating Saved Credentials with NTLM-only Server Authentication

See also

Smart Card Technical Reference

Smart Card Events

Article • 12/09/2022 • 13 minutes to read •

Applies to: ✅ Windows 10, ✅ Windows 11, ✅ Windows Server 2016, ✅ Windows Server 2019, ✅ Windows Server 2022

This topic for the IT professional and smart card developer describes events that are related to smart card deployment and development.

A number of events can be used to monitor smart card activities on a computer, including installation, use, and errors. The following sections describe the events and information that can be used to manage smart cards in an organization.

- Smart card reader name
- Smart card warning events
- Smart card error events
- Smart card Plug and Play events

Smart card reader name

The Smart Card resource manager doesn't use the device name from Device Manager to describe a smart card reader. Instead, the name is constructed from three device attributes that are queried directly from the smart card reader driver.

The following three attributes are used to construct the smart card reader name:

- Vendor name
- Interface device type
- Device unit

The smart card reader device name is constructed in the form *<VendorName> <Type> <DeviceUnit>*. For example 'Contoso Smart Card Reader 0' is constructed from the following information:

- Vendor name: Contoso
- Interface device type: Smart Card Reader
- Device unit: 0

Smart card warning events

> **Note** IOCTL in the following table refers to input and output control.

Event ID	Warning Message	Description
620	Smart Card Resource Manager was unable to cancel IOCTL %3 for reader '%2': %1. The reader may no longer be responding. If this error persists, your smart card or reader may not be functioning correctly. %n%nCommand Header: %4	This occurs if the resource manager attempts to cancel a command to the smart card reader when the smart card service is shutting down or after a smart card is removed from the smart card reader and the command could not be canceled. This can leave the smart card reader in an unusable state until it's removed from the computer or the computer is restarted. %1 = Windows error code %2 = Smart card reader name %3 = IOCTL being canceled %4 = First 4 bytes of the command that was sent to the smart card
619	Smart Card Reader '%2' hasn't responded to IOCTL %3 in %1 seconds. If this error persists, your smart card or reader may not be functioning correctly. %n%nCommand Header: %4	This occurs when a reader hasn't responded to an IOCTL after an unusually long period of time. Currently, this error is sent after a reader doesn't respond for 150 seconds. This can leave the smart card reader in an unusable state until it's removed from the computer or the computer is restarted. %1 = Number of seconds the IOCTL has been waiting %2 = Smart card reader name %3 = IOCTL sent %4 = First 4 bytes of the command that was sent to the smart card

Smart card error events

Event ID	Error Message	Description
202	Failed to initialize Server Application	An error occurred, and the service cannot initialize properly. Restarting the computer may resolve the issue.
203	Server Control has no memory for reader reference object.	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue.

Event ID	Error Message	Description
204	Server Control failed to create shutdown event: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Windows error code
205	Reader object has duplicate name: %1	There are two smart card readers that have the same name. Remove the smart card reader that is causing this error message. %1 = Name of the smart card reader that is duplicated
206	Failed to create global reader change event.	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue.
401	Reader shutdown exception from eject smart card command	A smart card reader could not eject a smart card while the smart card reader was shutting down.
406	Reader object cannot Identify Device	A smart card reader did not properly respond to a request for information about the device, which is required for constructing the smart card reader name. The smart card reader will not be recognized by the service until it's removed from the computer and reinserted or until the computer is restarted.
502	Initialization of Service Status Critical Section failed	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue.
504	Resource Manager cannot create shutdown event flag: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Windows error code
506	Smart Card Resource Manager failed to register service: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Windows error code

Event ID	Error Message	Description
506	Smart Card Resource Manager received unexpected exception from PnP event %1	An attempt to add a Plug and Play reader failed. The device may already be in use or may be defective. To resolve this error message, try to add the device again or restart the computer. %1 = The affected handle name
507	No memory available for Service Status Critical Section	There is not enough system memory available. This prevents the service from managing the status. Restarting the computer may resolve the issue.
508	Smart Card Resource Manager received unexpected exception from PnP event %1	An attempt to add a Plug and Play reader failed. The device may already be in use or may be defective. To resolve this error message, try to add the device again or restart the computer. %1 = The affected handle name
509	Smart Card Resource Manager received unexpected exception from PnP event %1	An attempt to add a Plug and Play reader failed. The device may already be in use or may be defective. To resolve this error message, try to add the device again or restart the computer. %1 = The affected handle name
510	Smart Card Resource Manager received NULL handle from PnP event %1	An attempt to add a Plug and Play smart card reader failed. The device may already be in use or may be defective. To resolve this error message, try to add the device again or restart the computer. %1 = The affected handle name
511	Smart Card Resource Manager received unexpected exception from PnP event %1	An attempt to add a Plug and Play reader failed. The device may already be in use or may be defective. To resolve this error message, try to add the device again or restart the computer. %1 = The affected handle name
512	Smart Card Resource Manager received NULL handle from PnP event %1	An attempt to add a Plug and Play smart card reader failed. The device may already be in use or may be defective. To resolve this error message, try to add the device again or restart the computer. %1 = The affected handle name
513	Smart Card Resource Manager received unexpected exception from PnP event %1	An attempt to add a Plug and Play reader failed. The device may already be in use or may be defective. To resolve this error message, try to add the device again or restart the computer. %1 = The affected handle name

Event ID	Error Message	Description
514	Smart Card Resource Manager failed to add reader %2: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Windows error code %2 = Smart card reader name
515	Smart Card Resource Manager failed to declare state: %1	This is an internal unrecoverable error that indicates a failure in the smart card service. The smart card service may not operate properly. Restarting the service or computer may resolve this issue. %1 = Windows error code
516	Smart Card Resource Manager Failed to declare shutdown: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The smart card service may not be able to stop. Restarting the computer may resolve this issue. %1 = Windows error code
517	Smart Card Resource Manager received unexpected exception attempting to add reader %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Smart card reader name
521	Smart Card Resource Manager received NULL handle from PnP event %1	An attempt to add a Plug and Play smart card reader failed. The device may already be in use or may be defective. To resolve this error message, try to add the device again or restart the computer. %1 = The affected handle name
523	Smart Card Resource Manager received NULL handle from PnP event %1	An attempt to add a Plug and Play smart card reader failed. The device may already be in use or may be defective. To resolve this error message, try to add the device again or restart the computer. %1 = The affected handle name
602	WDM Reader driver initialization cannot open reader device: %1	The service cannot open a communication channel with the smart card reader. You cannot use the smart card reader until the issue is resolved. %1 = Windows error code
603	WDM Reader driver initialization has no memory available to control device %1	There is not enough system memory available. This prevents the service from managing the smart card reader that was added. Restarting the computer may resolve the issue. %1 = Name of affected reader

Event ID	Error Message	Description
604	Server control cannot set reader removal event: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Windows error code
605	Reader object failed to create overlapped event: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Windows error code
606	Reader object failed to create removal event: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Windows error code
607	Reader object failed to start monitor thread: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Windows error code
608	Reader monitor failed to create power down timer: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Windows error code
609	Reader monitor failed to create overlapped event: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Windows error code

Event ID	Error Message	Description
610	Smart Card Reader '%2' rejected IOCTL %3: %1 If this error persists, your smart card or reader may not be functioning correctly.%n%nCommand Header: %4	The reader cannot successfully transmit the indicated IOCTL to the smart card. This can indicate hardware failure, but this error can also occur if a smart card or smart card reader is removed from the system while an operation is in progress. %1 = Windows error code %2 = Name of the smart card reader %3 = IOCTL that was sent %4 = First 4 bytes of the command sent to the smart card These events are caused by legacy functionality in the smart card stack. It can be ignored if there is no noticeable failure in the smart card usage scenarios. You might also see this error if your eSIM is recognized as a smartcard controller.
611	Smart Card Reader initialization failed	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve this issue.
612	Reader insertion monitor error retry threshold reached: %1	This occurs when a smart card reader fails several times to respond properly to the IOCTL, which indicates whether a smart card is present in the reader. The smart card reader is marked as defective, and it is not recognized by the service until it's removed from the computer and reinserted or until the computer is restarted. %1 = Windows error code
615	Reader removal monitor error retry threshold reached: %1	This occurs when a smart card reader fails several times to respond properly to the IOCTL, which indicates whether a smart card is present in the reader. The smart card reader is marked as defective, and it is not recognized by the service until it's removed from the computer and reinserted or until the computer is restarted. %1 = Windows error code
616	Reader monitor '%2' received uncaught error code: %1	This occurs when a smart card reader fails several times to respond properly to the IOCTL, which indicates whether a smart card is present in the reader. The smart card reader is marked as defective, and it is not recognized by the service until it's removed from the computer and reinserted or until the computer is restarted. %1 = Windows error code %2 = Reader name

Event ID	Error Message	Description
617	Reader monitor '%1' exception -- exiting thread	An unknown error occurred while monitoring a smart card reader for smart card insertions and removals. The smart card reader is marked as defective, and it is not recognized by the service until it's removed from the computer and reinserted or until the computer is restarted. %1 = Smart card reader name
618	Smart Card Resource Manager encountered an unrecoverable internal error.	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue.
621	Server Control failed to access start event: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Windows error code These events are caused by legacy functionality in the smart card stack. It can be ignored if there is no noticeable failure in the smart card usage scenarios.
622	Server Control failed to access stop event: %1	This is an internal, unrecoverable error that indicates a failure in the smart card service. The most common cause is limited computer resources. Restarting the computer may resolve the issue. %1 = Windows error code

Smart card Plug and Play events

Event ID	Event type	Event Message	Description
1000	Error	Could not get device ID for smart card in reader %1. The return code is %2.	Smart card Plug and Play could not obtain the device ID for the smart card. This information is required to determine the correct driver. The smart card may be defective. %1 = Smart card reader name %2 = Windows error code
1001	Information	Software successfully installed for smart card in reader %1. The smart card name is %2.	Smart card Plug and Play successfully installed a minidriver for the inserted card. %1 = Smart card reader name %2 = Name of new smart card device

See also

Smart Card Technical Reference

Virtual Smart Card Overview

Article • 02/27/2023 • 5 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

> ⚠ **Warning**
>
> **Windows Hello for Business** is the modern, two-factor authentication for Windows. Microsoft will deprecate virtual smart cards in the near future. Customers using virtual smart cards are strongly encouraged to move to Windows Hello for Business. Microsoft will publish the deprecation date to ensure customers have adequate lead time to move to Windows Hello for Business. We recommend that new Windows deployments use Windows Hello for Business.

This article provides an overview of the virtual smart card technology.

Feature description

Virtual smart card technology offers comparable security benefits to physical smart cards by using two-factor authentication. Virtual smart cards emulate the functionality of physical smart cards, but they use the Trusted Platform Module (TPM) chip that is available on devices. Virtual smart cards don't require the use of a separate physical smart card and reader. You create virtual smart cards in the TPM, where the keys used for authentication are stored in cryptographically-secured hardware.

By utilizing TPM devices that provide the same cryptographic capabilities as physical smart cards, virtual smart cards accomplish the three key properties that are desired for smart cards: non-exportability, isolated cryptography, and anti-hammering.

Practical applications

Virtual smart cards are functionally similar to physical smart cards, appearing in Windows as smart cards that are always-inserted. Virtual smart cards can be used for authentication to external resources, protection of data by encryption, and integrity through signing. You can deploy virtual smart cards by using in-house methods or a purchased solution, and they can be a replacement for other methods of strong authentication in a corporate setting of any scale.

Authentication use cases

Two-factor authentication–based remote access

After a user has a fully functional TPM virtual smart card, provisioned with a sign-in certificate, the certificate is used to gain authenticated access to corporate resources. When the proper certificate is provisioned to the virtual card, the user need only provide the PIN for the virtual smart card, as if it was a physical smart card, to sign in to the domain.

In practice, this is as easy as entering a password to access the system. Technically, it's far more secure. Using the virtual smart card to access the system proves to the domain that the user who is requesting authentication has possession of the personal computer upon which the card has been provisioned and knows the virtual smart card PIN. Because this request couldn't have possibly originated from a system other than the system certified by the domain for this user's access, and the user couldn't have initiated the request without knowing the PIN, a strong two-factor authentication is established.

Client authentication

Virtual smart cards can also be used for client authentication by using TLS/SSL or a similar technology. Similar to domain access with a virtual smart card, an authentication certificate can be provisioned for the virtual smart card, provided to a remote service, as requested in the client authentication process. This adheres to the principles of two-factor authentication because the certificate is only accessible from the computer that hosts the virtual smart card, and the user is required to enter the PIN for initial access to the card.

Virtual smart card redirection for remote desktop connections

The concept of two-factor authentication associated with virtual smart cards relies on the proximity of users to the devices that they use to access domain. When you connect to a device that is hosting virtual smart cards, you can't use the virtual smart cards located on the remote device during the remote session. However, you can access the virtual smart cards on the connecting device (which is under your physical control), which are loaded onto the remote device. You can use the virtual smart cards as if they were installed by using the remote devices' TPM, extending your privileges to the remote device, while maintaining the principles of two-factor authentication.

Confidentiality use cases

S/MIME email encryption

Physical smart cards are designed to hold private keys. You can use the private keys for email encryption and decryption. The same functionality exists in virtual smart cards. By

using S/MIME with a user's public key to encrypt email, the sender of an email is assured that only the person with the corresponding private key can decrypt the email. This assurance is a result of the non-exportability of the private key. It never exists within reach of malicious software, and it remains protected by the TPM—even during decryption.

BitLocker for data volumes

BitLocker Drive Encryption technology makes use of symmetric-key encryption to protect the content of a user's hard drive. BitLocker ensures that if the physical ownership of a hard drive is compromised, an adversary won't be able to read data off the drive. The key used to encrypt the drive can be stored in a virtual smart card, which necessitates knowledge of the virtual smart card PIN to access the drive, and possession of the device that is hosting the TPM virtual smart card. If the drive is obtained without access to the TPM that hosts the virtual smart card, any brute force attack will be difficult.

You can use BitLocker to encrypt portable drives, storing keys in virtual smart cards. In this scenario, unlike using BitLocker with a physical smart card, the encrypted drive can be used only when it's connected to device for the virtual smart card that is used to encrypt the drive, because the BitLocker key is only accessible from the device. This method can be useful to ensure the security of backup drives and personal storage uses outside the main hard drive, too.

Data integrity use case

Signing data

To verify authorship of data, a user can sign it by using a private key stored in the virtual smart card. Digital signatures confirm the integrity and origin of the data.

- Storing the key in an operating system that is accessible, malicious users could access it and use it to modify already signed data or to spoof the key owner's identity
- Storing the key in a virtual smart card, means that you can only use it to sign data on the host device. You can't export the key to other systems (intentionally or unintentionally, such as with malware theft), making digital signatures more secure than other methods for private key storage

Hardware requirements

To use the virtual smart card technology, TPM 1.2 is the minimum required for devices running a supported operating system.

Understand and Evaluate Virtual Smart Cards

Article • 02/27/2023 • 13 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

> ⚠ **Warning**
>
> **Windows Hello for Business** is the modern, two-factor authentication for Windows. Microsoft will deprecate virtual smart cards in the near future. Customers using virtual smart cards are strongly encouraged to move to Windows Hello for Business. Microsoft will publish the deprecation date to ensure customers have adequate lead time to move to Windows Hello for Business. We recommend that new Windows deployments use Windows Hello for Business.

This article describes the virtual smart card technology and how it can fit into your authentication design.

Virtual smart card technology uses cryptographic keys that are stored on computers that have the Trusted Platform Module (TPM) installed. Virtual smart cards offer comparable security benefits to conventional smart cards by using two-factor authentication. The technology also offers more convenience for users and has a lower cost to deploy. By utilizing TPM devices that provide the same cryptographic capabilities as conventional smart cards, virtual smart cards accomplish the three key properties that are desired for smart cards: non-exportability, isolated cryptography, and anti-hammering.

Virtual smart cards are functionally similar to physical smart cards. They appear as always-inserted smart cards, and they can be used for authentication to external resources, protection of data by secure encryption, and integrity through reliable signing. Because TPM-enabled hardware is readily available and virtual smart cards can be easily deployed by using existing certificate enrollment methods, virtual smart cards can become a full replacement for other methods of strong authentication in a corporate setting of any scale.

This topic contains the following sections:

- Comparing virtual smart cards with physical smart cards: Compares properties, functional aspects, security, and cost.
- Authentication design options: Describes how passwords, smart cards, and virtual smart cards can be used to reach authentication goals in your organization.

Comparing virtual smart cards with physical smart cards

Virtual smart cards function much like physical smart cards, but they differ in that they protect private keys by using the TPM of the computer instead of smart card media.

A virtual smart card appears to applications as a conventional smart card. Private keys in the virtual smart card are protected, not by isolation of physical memory, but by the cryptographic capabilities of the TPM. All sensitive information is encrypted by using the TPM and then stored on the hard drive in its encrypted form.

All cryptographic operations occur in the secure, isolated environment of the TPM, and the unencrypted private keys are never used outside this environment. So like physical smart cards, virtual smart cards remain secure from any malware on the host. Additionally, if the hard drive is compromised in some way, a malicious user won't be able to access keys that are stored in the virtual smart card because they're securely encrypted by using the TPM. Keys can also be protected by BitLocker Drive Encryption.

Virtual smart cards maintain the three key properties of physical smart cards:

- **Non-exportability**: Because all private information on the virtual smart card is encrypted by using the TPM on the host computer, it can't be used on a different computer with a different TPM. Additionally, TPMs are designed to be tamper-resistant and non-exportable, so a malicious user can't reverse engineer an identical TPM or install the same TPM on a different computer. For more information, see Evaluate Virtual Smart Card Security.
- **Isolated cryptography**: TPMs provide the same properties of isolated cryptography that are offered by physical smart cards, and this is utilized by virtual smart cards. Unencrypted copies of private keys are loaded only within the TPM and never into memory that is accessible by the operating system. All cryptographic operations with these private keys occur inside the TPM.
- **Anti-hammering**: If a user enters a PIN incorrectly, the virtual smart card responds by using the anti-hammering logic of the TPM, which rejects further attempts for a period of time instead of blocking the card. This is also known as lockout. For more information, see Evaluate Virtual Smart Card Security.

The following subsections compare the functionality, security, and cost of virtual smart cards and physical smart cards.

Functionality

The virtual smart card system that was designed by Microsoft closely mimics the functionality of conventional smart cards. The most striking difference to the end user is that the virtual smart card is essentially a smart card that is always inserted into the computer. There's no method to export the user's virtual smart card for use on other computers, which adds to the security of virtual smart cards. If a user requires access to network resources on multiple computers, multiple virtual smart cards can be issued for that user. Additionally, a computer that is shared among multiple users can host multiple virtual smart cards for different users.

The basic user experience for a virtual smart card is as simple as using a password to access a network. Because the smart card is loaded by default, the user must simply enter the PIN that is tied to the card to gain access. Users are no longer required to carry cards and readers or to take physical action to use the card.

Additionally, although the anti-hammering functionality of the virtual smart card is equally secure to that of a physical smart card, virtual smart card users are never required to contact an administrator to unblock the card. Instead, they simply wait a period of time (depending on the TPM specifications) before they reattempt to enter the PIN. Alternatively, the administrator can reset the lockout by providing owner authentication data to the TPM on the host computer.

Security

Physical smart cards and virtual smart cards offer comparable levels of security. They both implement two-factor authentication for using network resources. However, they differ in certain aspects, including physical security and the practicality of an attack. Due to their compact and portable design, conventional smart cards are most frequently kept close to their intended user. They offer little opportunity for acquisition by a potential adversary, so any sort of interaction with the card is difficult without committing some variety of theft.

TPM virtual smart cards, however, reside on a user's computer that may frequently be left unattended, which provides an opportunity for a malicious user to hammer the TPM. Although virtual smart cards are fully protected from hammering (as are physical smart cards), this accessibility makes the logistics of an attack somewhat simpler. Additionally, the anti-hammering behavior of a TPM smart card differs in that it only presents a time delay in response to repeated PIN failures, as opposed to fully blocking the user.

However, there are several advantages provided by virtual smart cards to mitigate these slight security deficits. Most importantly, a virtual smart card is much less likely to be lost. Virtual smart cards are integrated into computers and devices that the user already owns for other purposes and has incentive to keep safe. If the computer or device that hosts the virtual smart card is lost or stolen, a user will more immediately notice its loss

than the loss of a physical smart card. When a computer or device is identified as lost, the user can notify the administrator of the system, who can revoke the certificate that is associated with the virtual smart card on that device. This precludes any future unauthorized access on that computer or device if the PIN for the virtual smart card is compromised.

Cost

If a company wants to deploy physical smart cards, they need to purchase smart cards and smart card readers for all employees. Although relatively inexpensive options can be found, options that ensure the three key properties of smart card security (most notably, non-exportability) are more expensive. If employees have computers with a built-in TPM, virtual smart cards can be deployed with no additional material costs. These computers and devices are relatively common in the market.

The maintenance cost of virtual smart cards is less than that for physical smart cards, which are easily lost, stolen, or broken from normal wear. TPM virtual smart cards are only lost or broken if the host computer or device is lost or broken, which in most cases is much less frequently.

Comparison summary

Physical Smart Cards	TPM virtual smart cards
Protects private keys by using the built-in cryptographic functionality of the card.	Protects private keys by using the cryptographic functionality of the TPM.
Stores private keys in isolated non-volatile memory on the card, which means that access to private keys is only from the card, and access is never allowed to the operating system.	Stores encrypted private keys on the hard drive. The encryption ensures that these keys can only be decrypted and used in the TPM, not in the accessible memory of the operating system.
Guarantees non-exportability through the card manufacturer, which includes isolating private information from operating system access.	Guarantees non-exportability through the TPM manufacturer, which includes the inability of an adversary to replicate or remove the TPM.
Performs and isolates cryptographic operations within the built-in capabilities of the card.	Performs and isolates cryptographic operations in the TPM of the user's computer or device.
Provides anti-hammering through the card. After a certain number of failed PIN entry attempts, the card blocks further access until administrative action is taken.	Provides anti-hammering through the TPM. Successive failed attempts increase the device lockout time (the time the user has to wait before trying again). This can be reset by an administrator.

Physical Smart Cards	TPM virtual smart cards
Requires that users carry their smart card and smart card reader with them to access network resources.	Allows users to access their TPM-enabled computers or devices, and potentially access the network, without other equipment.
Enables credential portability by inserting the smart card into smart card readers that are attached to other computers.	Prevents exporting credentials from a given computer or device. However, virtual smart cards can be issued for the same user on multiple computers or devices by using additional certificates.
Enables multiple users to access network resources through the same computer by inserting their personal smart cards.	Enables multiple users to access network resources through the same computer or device by issuing a virtual smart card for each user on that computer or device.
Requires the user to carry the card, making it more difficult for an attacker to access the device and launch a hammering attempt.	Stores virtual smart card on the user's computer, which may be left unattended and allow a greater risk window for hammering attempts.
Provides a generally single-purpose device that is carried explicitly for the purpose of authentication. The smart card can be easily misplaced or forgotten.	Installs the virtual smart card on a device that has other purposes for the user, so the user has greater incentive to be responsible for the computer or device.
Alerts users that their card is lost or stolen only when they need to sign in and notice it's missing.	Installs the virtual smart card on a device that the user likely needs for other purposes, so users will notice its loss much more quickly. This reduces the associated risk window.
Requires companies to invest in smart cards and smart card readers for all employees.	Requires that companies ensure all employees have TPM-enabled computers, which are relatively common.
Enables using a smart card removal policy to affect system behavior when the smart card is removed. For example, the policy can dictate if the user's sign-in session is locked or terminated when the user removes the card.	Eliminates the necessity for a smart card removal policy because a TPM virtual smart card is always present and can't be removed from the computer.

Authentication design options

The following section presents several commonly used options and their respective strengths and weaknesses, which organizations can consider for authentication.

Passwords

A password is a secret string of characters that is tied to the identification credentials for a user's account. This establishes the user's identity. Although passwords are the most commonly used form of authentication, they're also the weakest. In a system where passwords are used as the sole method of user authentication, only individuals who know their passwords are considered valid users.

Password authentication places a great deal of responsibility on the user. Passwords must be sufficiently complex so they can't be easily guessed, but they must be simple enough to be committed to memory and not stored in a physical location. Even if this balance is successfully achieved, a wide variety of attacks exist (such as brute force attacks, eavesdropping, and social engineering tactics) where a malicious user can acquire a user's password and impersonate that person's identity. A user often won't realize that the password is compromised, which makes it's easy for a malicious user to maintain access to a system if a valid password has been obtained.

One-time passwords

A one-time password (OTP) is similar to a traditional password, but it's more secure in that it can be used only once to authenticate a user. The method for determining each new password varies by implementation. Assuming a secure deployment of each new password, OTPs have several advantages over the classic password model of authentication. Most importantly, if a given OTP token is intercepted in transmission between the user and the system, the interceptor can't use it for any future transactions. Similarly, if a malicious user obtains a valid user's OTP, the interceptor will have limited access to the system (only one session).

Smart cards

Smart cards are physical authentication devices, which improve on the concept of a password by requiring that users actually have their smart card device with them to access the system, in addition to knowing the PIN that provides access to the smart card. Smart cards have three key properties that help maintain their security:

- **Non-exportability**: Information stored on the card, such as the user's private keys, can't be extracted from one device and used in another medium
- **Isolated cryptography**: Any cryptographic operations that are related to the card (such as secure encryption and decryption of data) occur in a cryptographic processor on the card, so malicious software on the host computer can't observe the transactions
- **Anti-hammering**: To prevent access to the card by a brute-force attack, a set number of consecutive unsuccessful PIN entry attempts blocks the card until

administrative action is taken

Smart cards provide greatly enhanced security over passwords alone, because it's much more difficult for a malicious user to gain and maintain access to a system. Most importantly, access to a smart card system requires that users have a valid card and that they know the PIN that provides access to that card. It's difficult for a thief to acquire the card and the PIN.

Additional security is achieved by the singular nature of the card because only one copy of the card exists, only one individual can use the sign-in credentials, and users will quickly notice if the card has been lost or stolen. This greatly reduces the risk window of credential theft when compared to using a password alone.

The additional security comes with added material and support costs. Traditional smart cards are expensive to purchase (cards and card readers must be supplied to employees), and users can misplace or lose them.

Virtual smart cards

Virtual smart cards emulate the functionality of traditional smart cards. Instead of requiring the purchase of additional hardware, virtual smart cards utilize technology that users already own and are more likely to always have with them. Theoretically, any device that can provide the three key properties of smart cards (non-exportability, isolated cryptography, and anti-hammering) can be commissioned as a virtual smart card. The virtual smart card platform is limited to the use of the Trusted Platform Module (TPM) chip, which is on most modern devices.

Virtual smart cards that utilize a TPM provide the three main security principles of traditional smart cards: non-exportability, isolated cryptography, and anti-hammering. Virtual smart cards are less expensive to implement and more convenient for users. Since many corporate computers already have a built-in TPM, there's no cost associated with purchasing new hardware. The user's possession of a computer or device is equivalent to the possession of a smart card, and a user's identity can't be assumed from any other computer or device without administrative provisioning of further credentials. Thus, two-factor authentication is achieved because the user must have a computer that is set up with a virtual smart card and know the PIN to use the virtual smart card.

Get Started with Virtual Smart Cards: Walkthrough Guide

Article • 02/27/2023 • 5 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

> ⚠ **Warning**
>
> **Windows Hello for Business** is the modern, two-factor authentication for Windows. Microsoft will deprecate virtual smart cards in the near future. Customers using virtual smart cards are strongly encouraged to move to Windows Hello for Business. Microsoft will publish the deprecation date to ensure customers have adequate lead time to move to Windows Hello for Business. We recommend that new Windows deployments use Windows Hello for Business.

This topic for the IT professional describes how to set up a basic test environment for using TPM virtual smart cards.

Virtual smart cards are a technology from Microsoft that offer comparable security benefits in two-factor authentication to physical smart cards. They also offer more convenience for users and lower cost for organizations to deploy. By utilizing Trusted Platform Module (TPM) devices that provide the same cryptographic capabilities as physical smart cards, virtual smart cards accomplish the three key properties that are desired by smart cards: non-exportability, isolated cryptography, and anti-hammering.

This step-by-step walkthrough shows you how to set up a basic test environment for using TPM virtual smart cards. After you complete this walkthrough, you will have a functional virtual smart card installed on the Windows computer.

Time requirements

You should be able to complete this walkthrough in less than one hour, excluding installing software and setting up the test domain.

Walkthrough steps

- Prerequisites
- Step 1: Create the certificate template
- Step 2: Create the TPM virtual smart card
- Step 3: Enroll for the certificate on the TPM Virtual Smart Card

> **Important** This basic configuration is for test purposes only. It is not intended for use in a production environment.

Prerequisites

You will need:

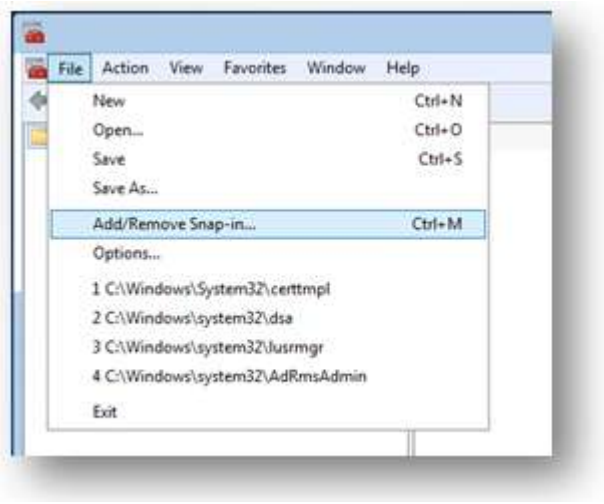
- A computer running Windows 10 with an installed and fully functional TPM (version 1.2 or version 2.0).
- A test domain to which the computer listed above can be joined.
- Access to a server in that domain with a fully installed and running certification authority (CA).

Step 1: Create the certificate template

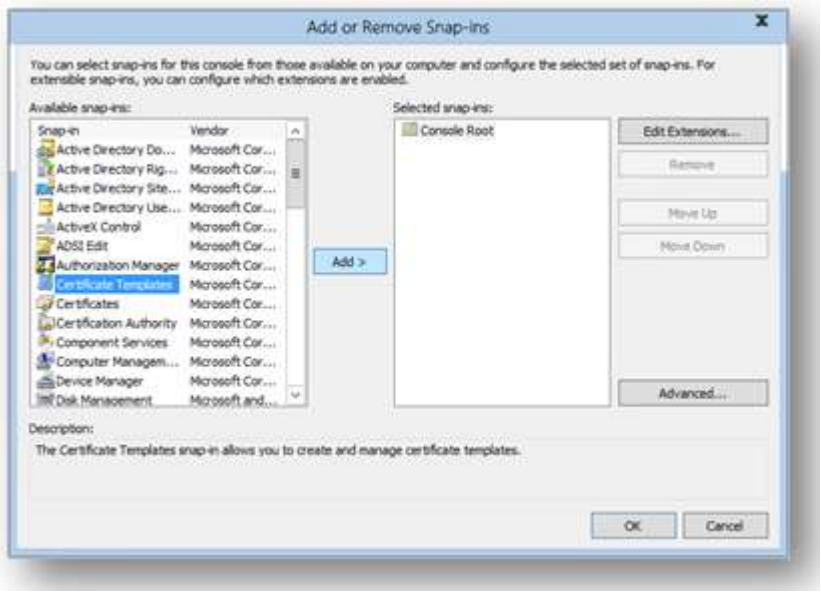
On your domain server, you need to create a template for the certificate that you will request for the virtual smart card.

To create the certificate template

1. On your server, open the Microsoft Management Console (MMC). One way to do this is to type **mmc.exe** from the **Start** menu, right-click **mmc.exe**, and click **Run as administrator**.
2. Click **File**, and then click **Add/Remove Snap-in**.

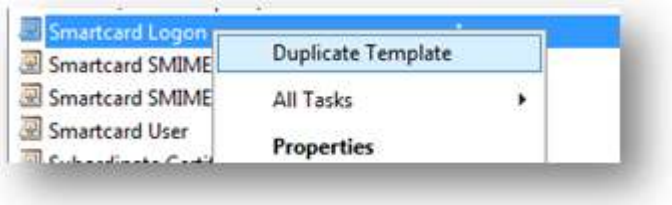


3. In the available snap-ins list, click **Certificate Templates**, and then click **Add**.

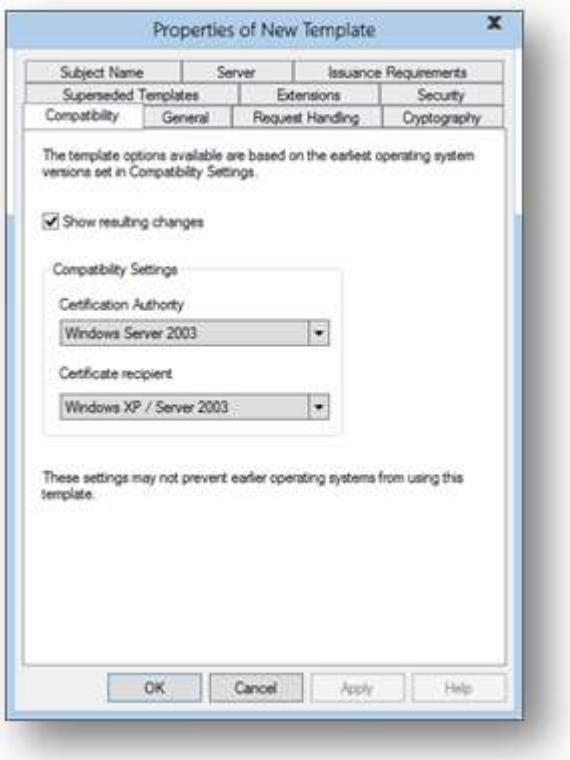


4. Certificate Templates is now located under **Console Root** in the MMC. Double-click it to view all the available certificate templates.

5. Right-click the **Smartcard Logon** template, and click **Duplicate Template**.



6. On the **Compatibility** tab, under **Certification Authority**, review the selection, and change it if needed.



7. On the **General** tab:

a. Specify a name, such as **TPM Virtual Smart Card Logon**.

b. Set the validity period to the desired value.

8. On the **Request Handling** tab:

   a. Set the **Purpose** to **Signature and smartcard logon**.

   b. Click **Prompt the user during enrollment**.

9. On the **Cryptography** tab:

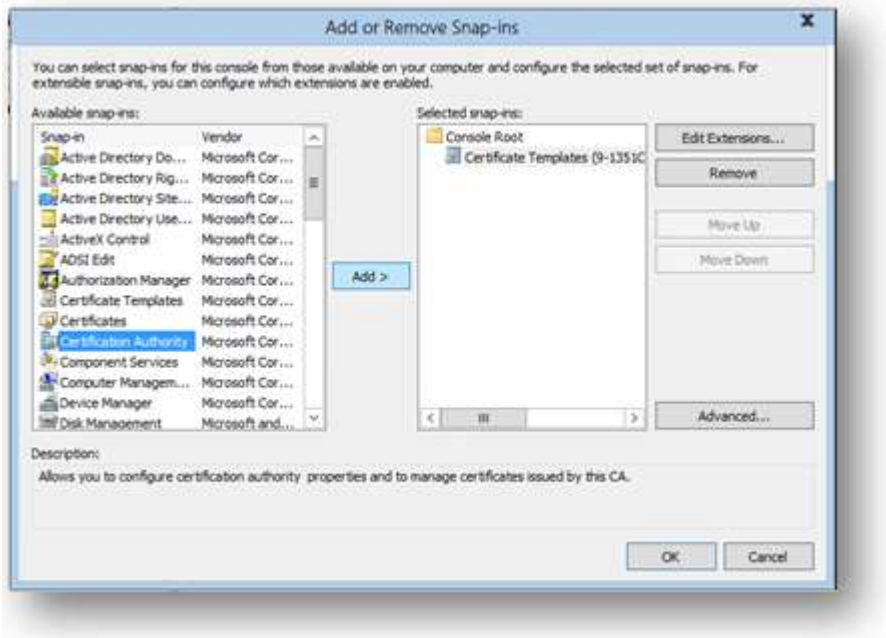
   a. Set the minimum key size to 2048.

   b. Click **Requests must use one of the following providers**, and then select **Microsoft Base Smart Card Crypto Provider**.

10. On the **Security** tab, add the security group that you want to give **Enroll** access to. For example, if you want to give access to all users, select the **Authenticated users** group, and then select **Enroll** permissions for them.

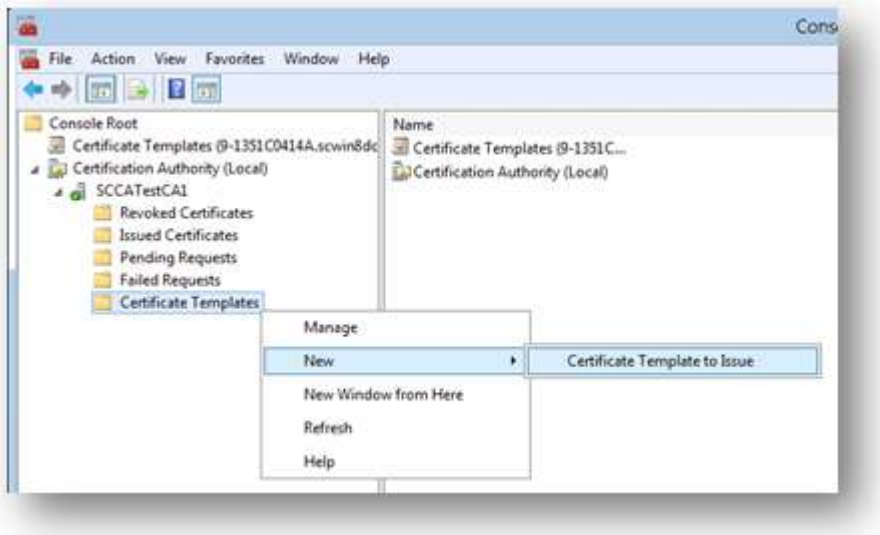
11. Click **OK** to finalize your changes and create the new template. Your new template should now appear in the list of Certificate Templates.

12. Select **File**, then click **Add/Remove Snap-in** to add the Certification Authority snap-in to your MMC console. When asked which computer you want to manage, select the computer on which the CA is located, probably **Local Computer**.



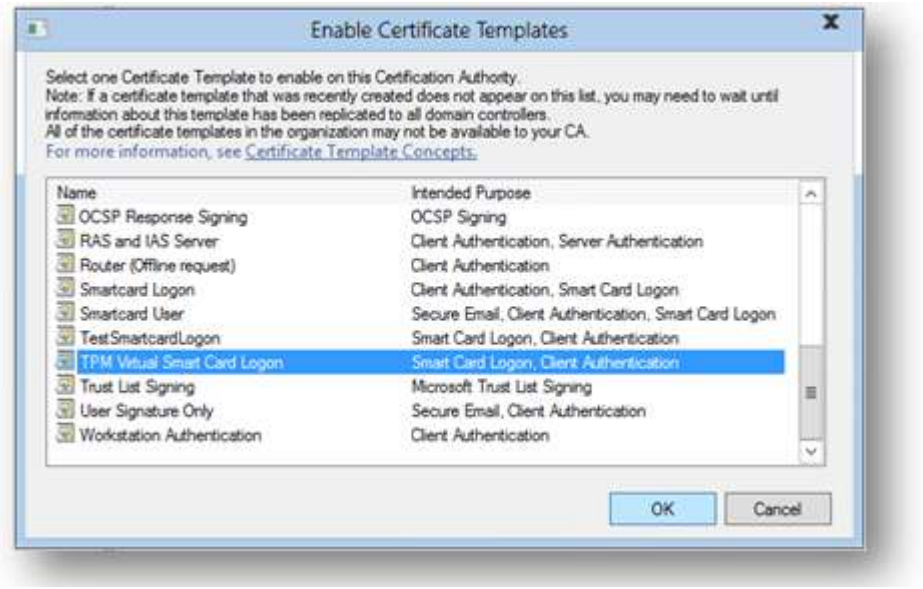
13. In the left pane of the MMC, expand **Certification Authority (Local)**, and then expand your CA within the Certification Authority list.

14. Right-click **Certificate Templates**, click **New**, and then click **Certificate Template to Issue**.

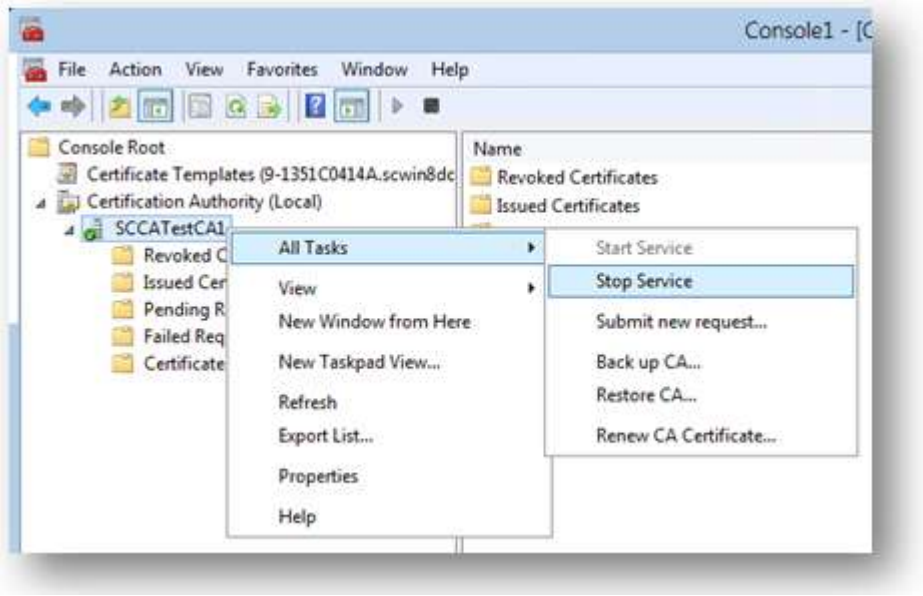


15. From the list, select the new template that you just created (**TPM Virtual Smart Card Logon**), and then click **OK**.

> **Note** It can take some time for your template to replicate to all servers and become available in this list.



16. After the template replicates, in the MMC, right-click in the Certification Authority list, click **All Tasks**, and then click **Stop Service**. Then, right-click the name of the CA again, click **All Tasks**, and then click **Start Service**.

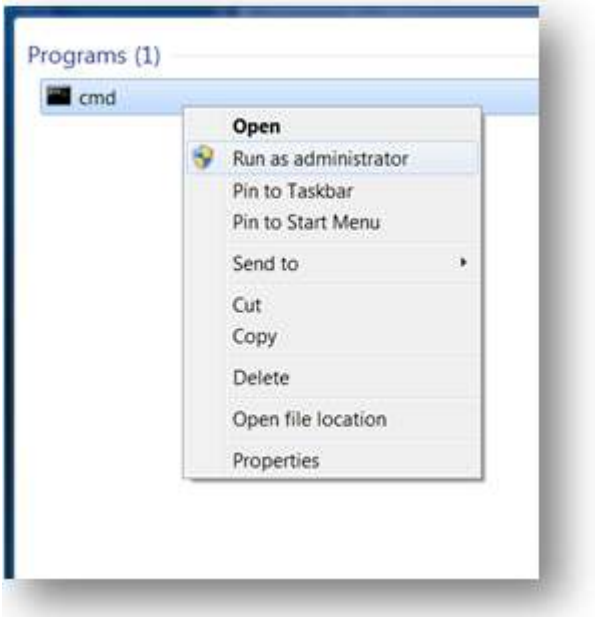


Step 2: Create the TPM virtual smart card

In this step, you will create the virtual smart card on the client computer by using the command-line tool, Tpmvscmgr.exe.

To create the TPM virtual smart card

1. On a domain-joined computer, open a Command Prompt window with Administrative credentials.



2. At the command prompt, type the following, and then press ENTER:

```
tpmvscmgr.exe create /name TestVSC /pin default /adminkey random /generate
```

This will create a virtual smart card with the name **TestVSC**, omit the unlock key, and generate the file system on the card. The PIN will be set to the default, 12345678. To be prompted for a PIN, instead of **/pin default** you can type **/pin prompt**.

For more information about the Tpmvscmgr command-line tool, see Use Virtual Smart Cards and Tpmvscmgr.

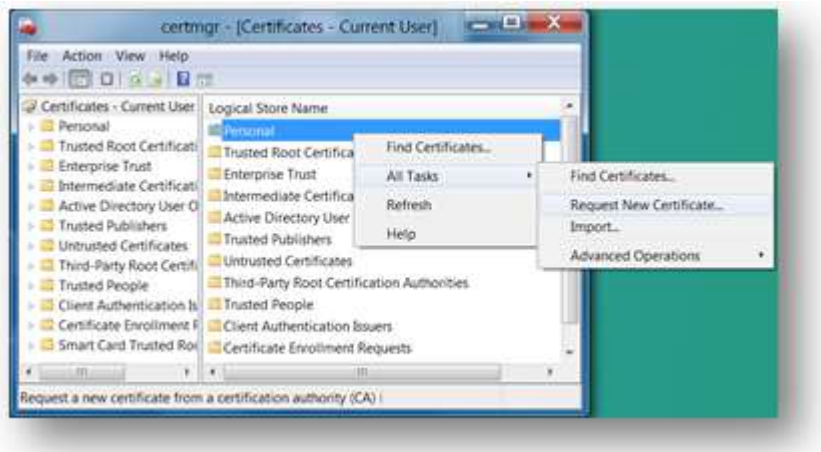
3. Wait several seconds for the process to finish. Upon completion, Tpmvscmgr.exe will provide you with the device instance ID for the TPM Virtual Smart Card. Store this ID for later reference because you will need it to manage or remove the virtual smart card.

Step 3: Enroll for the certificate on the TPM Virtual Smart Card

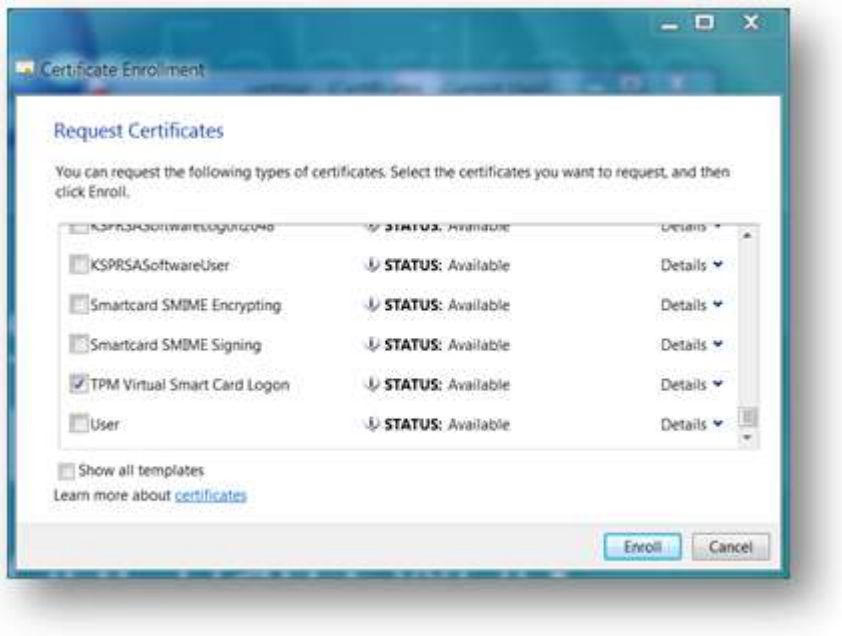
The virtual smart card must be provisioned with a sign-in certificate for it to be fully functional.

To enroll the certificate

1. Open the Certificates console by typing **certmgr.msc** on the **Start** menu.
2. Right-click **Personal**, click **All Tasks**, and then click **Request New Certificate**.



3. Follow the prompts and when offered a list of templates, select the **TPM Virtual Smart Card Logon** check box (or whatever you named the template in Step 1).



4. If prompted for a device, select the Microsoft virtual smart card that corresponds to the one you created in the previous section. It displays as **Identity Device (Microsoft Profile)**.

5. Enter the PIN that was established when you created the TPM virtual smart card, and then click **OK**.

6. Wait for the enrollment to finish, and then click **Finish**.

The virtual smart card can now be used as an alternative credential to sign in to your domain. To verify that your virtual smart card configuration and certificate enrollment were successful, sign out of your current session, and then sign in. When you sign in, you will see the icon for the new TPM virtual smart card on the Secure Desktop (sign in) screen or you will be automatically directed to the TPM smart card sign-in dialog box. Click the icon, enter your PIN (if necessary), and then click **OK**. You should be signed in to your domain account.

See also

- Understanding and Evaluating Virtual Smart Cards
- Use Virtual Smart Cards
- Deploy Virtual Smart Cards

Use Virtual Smart Cards

Article • 02/27/2023 • 5 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

> ⚠ **Warning**
>
> **Windows Hello for Business** is the modern, two-factor authentication for Windows. Microsoft will deprecate virtual smart cards in the near future. Customers using virtual smart cards are strongly encouraged to move to Windows Hello for Business. Microsoft will publish the deprecation date to ensure customers have adequate lead time to move to Windows Hello for Business. We recommend that new Windows deployments use Windows Hello for Business.

Learn about the requirements for virtual smart cards, how to use and manage them.

Requirements, restrictions, and limitations

Area	Requirements and details
Supported operating systems	Windows Server 2016 Windows Server 2012 R2 Windows Server 2012 Windows 10 Windows 8.1 Windows 8
Supported Trusted Platform Module (TPM)	Any TPM that adheres to the TPM main specifications for version 1.2 or version 2.0 (as set by the Trusted Computing Group) is supported for use as a virtual smart card. For more information, see the TPM Main Specification.
Supported virtual smart cards per computer	Ten smart cards can be connected to a computer or device at one time. This includes physical and virtual smart cards combined. **Note** You can create more than one virtual smart card; however, after creating more than four virtual smart cards, you may start to notice performance degradation. Because all smart cards appear as if they're always inserted, if more than one person shares a computer or device, each person can see all the virtual smart cards that are created on that computer or device. If the user knows the PIN values for all the virtual smart cards, the user will also be able to use them.

Area	Requirements and details
Supported number of certificates on a virtual smart card	A single TPM virtual smart card can contain 30 distinct certificates with the corresponding private keys. Users can continue to renew certificates on the card until the total number of certificates on a card exceeds 90. The reason that the total number of certificates is different from the total number of private keys is that sometimes the renewal can be done with the same private key—in which case a new private key isn't generated.
PIN, PIN Unlock Key (PUK), and Administrative key requirements	The PIN and the PUK must be a minimum of eight characters that can include numerals, alphabetic characters, and special characters. The Administrative key must be entered as 48 hexadecimal characters. It's a 3-key triple DES with ISO/IEC 9797 padding method 2 in CBC chaining mode.

Using Tpmvscmgr.exe

To create and delete TPM virtual smart cards for end users, the Tpmvscmgr command-line tool is included as a command-line tool with the operating system. You can use the **Create** and **Delete** parameters to manage virtual smart cards on local or remote computers. For information about using this tool, see Tpmvscmgr.

Create and delete virtual smart cards programmatically

Virtual smart cards can also be created and deleted by using APIs. For more information, see the following classes and interfaces:

- TpmVirtualSmartCardManager
- RemoteTpmVirtualSmartCardManager
- ITpmVirtualSmartCardManager
- ITPMVirtualSmartCardManagerStatusCallBack

You can use APIs that were introduced in the Windows.Device.SmartCards namespace in Windows Server 2012 R2 and Windows 8.1 to build Microsoft Store apps to manage the full lifecycle of virtual smart cards. For information about how to build an app to do this, see Strong Authentication: Building Apps That Leverage Virtual Smart Cards in Enterprise, BYOD, and Consumer Environments | Build 2013 | Channel 9.

The following table describes the features that can be developed in a Microsoft Store app:

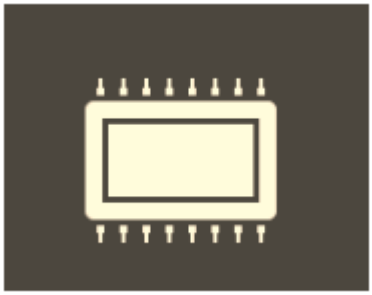
Feature	Physical Smart Card	Virtual Smart Card
Query and monitor smart card readers	Yes	Yes
List available smart cards in a reader, and retrieve the card name and card ID	Yes	Yes
Verify if the administrative key of a card is correct	Yes	Yes
Provision (or reformat) a card with a given card ID	Yes	Yes
Change the PIN by entering the old PIN and specifying a new PIN	Yes	Yes
Change the administrative key, reset the PIN, or unblock the smart card by using a challenge/response method	Yes	Yes
Create a virtual smart card	Not applicable	Yes
Delete a virtual smart card	Not applicable	Yes
Set PIN policies	No	Yes

For more information about these Windows APIs, see:

- Windows.Devices.SmartCards namespace (Windows)
- Windows.Security.Cryptography.Certificates namespace (Windows)

Distinguishing TPM-based virtual smart cards from physical smart cards

To help users visually distinguish a Trusted Platform Module (TPM)-based virtual smart card from physical smart cards, the virtual smart card has a different icon. The following icon is displayed during sign-in, and on other screens that require the user to enter the PIN for a virtual smart card.



A TPM-based virtual smart card is labeled **Security Device** in the user interface.

Changing the PIN

The PIN for a virtual smart card can be changed by following these steps:

- Sign in with the old PIN or password.
- Press Ctrl+Alt+Del and choose **Change a password**.
- Select **Sign-in Options**.
- Select the virtual smart card icon.
- Enter and confirm the new PIN.

Resolving issues

TPM not provisioned

For a TPM-based virtual smart card to function properly, a provisioned TPM must be available on the computer. If the TPM is disabled in the BIOS, or it isn't provisioned with full ownership and the storage root key, the TPM virtual smart card creation will fail.

If the TPM is initialized after creating a virtual smart card, the card will no longer function, and it will need to be re-created.

If the TPM ownership was established on a Windows Vista installation, the TPM won't be ready to use virtual smart cards. The system administrator needs to clear and initialize the TPM for it to be suitable for creating TPM virtual smart cards.

If the operating system is reinstalled, prior TPM virtual smart cards are no longer available and need to be re-created. If the operating system is upgraded, prior TPM virtual smart cards will be available to use in the upgraded operating system.

TPM in lockout state

Sometimes, due to frequent incorrect PIN attempts from a user, the TPM may enter the lockout state. To resume using the TPM virtual smart card, it's necessary to reset the lockout on the TPM by using the owner's password or to wait for the lockout to expire. Unblocking the user PIN doesn't reset the lockout in the TPM. When the TPM is in lockout, the TPM virtual smart card appears as if it's blocked. When the TPM enters the lockout state because the user entered an incorrect PIN too many times, it may be necessary to reset the user PIN by using the virtual smart card management tools, such as Tpmvscmgr command-line tool.

See also

For information about authentication, confidentiality, and data integrity use cases, see Virtual Smart Card Overview.

Deploy Virtual Smart Cards

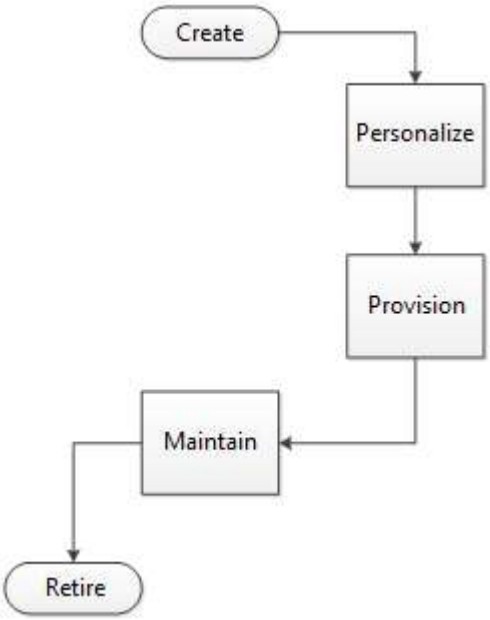
Article • 02/27/2023 • 24 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

> ⚠ **Warning**
>
> **Windows Hello for Business** is the modern, two-factor authentication for Windows. Microsoft will deprecate virtual smart cards in the near future. Customers using virtual smart cards are strongly encouraged to move to Windows Hello for Business. Microsoft will publish the deprecation date to ensure customers have adequate lead time to move to Windows Hello for Business. We recommend that new Windows deployments use Windows Hello for Business.

This article discusses the factors to consider when you deploy a virtual smart card authentication solution.

Traditional identity devices, such as physical smart cards, follow a predictable lifecycle in any deployment, as shown in the following diagram.



A device manufacturer creates physical devices, and then an organization purchase and deploy them. The device passes through the personalization stage, where its unique properties are set. In smart cards, these properties are the *administrator key*, *Personal Identification Number (PIN)*, *PIN Unlock Key (PUK)*, and its physical appearance. During the device provisioning phase, the required certificates are installed, such as a sign-in certificate. After you provision the device, it's ready for use. You'll maintain the device, for example you may replace cards when they're lost or stolen, or reset PINs when users

forget them. Finally, you'll retire devices when they exceed their intended lifetime or when employees leave the company.

This topic contains information about the following phases in a virtual smart card lifecycle:

- Create and personalize virtual smart cards
- Provision virtual smart cards
- Maintain virtual smart cards

Create and personalize virtual smart cards

A corporation purchases the devices to deploy then. The device passes through the personalization stage, where its unique properties are set. In smart cards, these properties are the administrator key, Personal Identification Number (PIN), PIN Unlock Key (PUK), and its physical appearance. The security that is provided for a TPM virtual smart card is fully provisioned in the host TPM.

Trusted Platform Module readiness

The TPM Provisioning Wizard, which is launched from the **TPM Management Console**, takes the user through all the steps to prepare the TPM for use.

When you create virtual smart cards, consider the following actions in the TPM:

- **Enable and Activate**: TPMs are built into many devices. In some cases, the TPM must be enabled and activated through the BIOS
- **Take ownership**: When you provision the TPM, you set an owner password for managing the TPM in the future, and you establish the *storage root key*. To provide anti-hammering protection for virtual smart cards, the user or a domain administrator must be able to reset the TPM owner password. For corporate use of TPM virtual smart cards, the domain administrator should restrict access to the TPM owner password by storing it in Active Directory, and not in the local registry. When TPM ownership is set, you must clear and reinitialize the TPM
- **Manage**: You can manage ownership of a virtual smart card by changing the owner password, and you can manage anti-hammering logic by resetting the lockout time

A TPM might operate in reduced functionality mode, which may occur if the operating system can't determine if the owner password is available to the user. During reduce functionality mode, you can use the TPM to create a virtual smart card, but it's

preferable to bring the TPM to a fully ready state so that any unexpected circumstances won't leave the user blocked from using the device.

Those smart card deployment management tools that require a status check of a TPM before attempting to create a TPM virtual smart card can do so using the TPM WMI interface.

Depending on the setup of the device designated for installing TPM virtual smart cards, it may be necessary to provision the TPM before continuing with the virtual smart card deployment. For more information about provisioning, see Use Virtual Smart Cards.

For more information about managing TPMs by using built-in tools, see Trusted Platform Module Services Group Policy Settings.

Creation

A TPM virtual smart card simulates a physical smart card, using the TPM to provide the same functionality as physical smart card hardware.
A virtual smart card appears within the operating system as a physical smart card that is always inserted. Windows presents a *virtual smart card reader* and a *virtual smart card* to applications using the same interface as physical smart cards. The messages to and from the virtual smart card are translated to TPM commands, ensuring the integrity of the virtual smart card through the three properties of smart card security:

- **Non-exportability**: Because all private information on the virtual smart card is encrypted by using the TPM on the host computer, it can't be used on a different computer with a different TPM. Additionally, TPMs are designed to be tamper-resistant and non-exportable, so a malicious user can't reverse engineer an identical TPM or install the same TPM on a different computer. For more information, see Evaluate Virtual Smart Card Security.

- **Isolated cryptography**: TPMs provide the same properties of isolated cryptography that is offered by physical smart cards, which is utilized by virtual smart cards. Unencrypted copies of private keys are loaded only within the TPM and never into memory that is accessible by the operating system. All cryptographic operations with these private keys occur inside the TPM.

- **Anti-hammering**: If a user enters a PIN incorrectly, the virtual smart card responds by using the anti-hammering logic of the TPM, which rejects further attempts for some time instead of blocking the card. This is also known as lockout. For more information, see Blocked virtual smart card and Evaluate Virtual Smart Card Security.

There are several options for creating virtual smart cards, depending on the size of the deployment and budget of the organization. The lowest cost option is using `tpmvscmgr.exe` to create cards individually on users' computers. Alternatively, a virtual smart card management solution can be purchased to more easily accomplish virtual smart card creation on a larger scale and aid in further phases of deployment. Virtual smart cards can be created on computers that are to be provisioned for an employee or on those that are already in an employee's possession. In either approach, there should be some central control over personalization and provisioning. If a computer is intended for use by multiple employees, multiple virtual smart cards can be created on a computer.

For information about the TPM Virtual Smart Card command-line tool, see Tpmvscmgr.

Personalization

During virtual smart card personalization, the values for the administrator key, PIN, and PUK are assigned. As with a physical card, knowing the administrator key is important for resetting the PIN or for deleting the card in the future. (If you set a PUK, you can't use the administrator key to reset the PIN.)

Because the administrator key is critical to the security of the card, it's important to consider the deployment environment and decide on the proper administrator key setting strategy. Options for these strategies include:

- **Uniform**: Administrator keys for all the virtual smart cards deployed in the organization are the same. Although using the same key makes the maintenance infrastructure easy (only one key needs to be stored), it's highly insecure. This strategy might be sufficient for small organizations, but if the administrator key is compromised, all virtual smart cards that use the key must be reissued.

- **Random, not stored**: Administrator keys are assigned randomly for all virtual smart cards, and they aren't recorded. This is a valid option if the deployment administrators don't require the ability to reset PINs, and instead prefer to delete and reissue virtual smart cards. This is a viable strategy if the administrator prefers to set PUK values for the virtual smart cards and then use this value to reset PINs, if necessary.

- **Random, stored**: you assign the administrator keys randomly, storing them in a central location. Each card's security is independent of the others. This is a secure strategy on a large scale, unless the administrator key database is compromised.

- **Deterministic**: Administrator keys are the result of some function or known information. For example, the user ID could be used to randomly generate data

that can be further processed through a symmetric encryption algorithm by using a secret. This administrator key can be similarly regenerated when needed, and it doesn't need to be stored. The security of this method relies on the security of the secret used.

Although the PUK and the administrator key methodologies provide unlocking and resetting functionality, they do so in different ways. The PUK is a PIN that is entered on the computer to enable a user PIN reset.

The administrator key methodology takes a challenge-response approach. The card provides a set of random data after users verify their identity to the deployment administrator. The administrator then encrypts the data with the administrator key and gives the encrypted data back to the user. If the encrypted data matches that produced by the card during verification, the card will allow PIN reset. Because the administrator key is never accessible by anyone other than the deployment administrator, it can't be intercepted or recorded by any other party (including employees). This provides significant security benefits beyond using a PUK, an important consideration during the personalization process.

TPM virtual smart cards can be personalized on an individual basis when they're created with the Tpmvscmgr command-line tool. Or organizations can purchase a management solution that can incorporate personalization into an automated routine. Another advantage of such a solution is the automated creation of administrator keys. Tpmvscmgr.exe allows users to create their own administrator keys, which can be detrimental to the security of the virtual smart cards.

Provision virtual smart cards

Provisioning is the process of loading specific credentials onto a TPM virtual smart card. These credentials consist of certificates that are created to give users access to a specific service, such as domain sign-in. A maximum of 30 certificates is allowed on each virtual smart card. As with physical smart cards, several decisions must be made regarding the provisioning strategy, based on the environment of the deployment and the desired level of security.

A high-assurance level of secure provisioning requires absolute certainty about the identity of the individual who is receiving the certificate. Therefore, one method of high-assurance provisioning is utilizing previously provisioned strong credentials, such as a physical smart card, to validate identity during provisioning. In-person proofing at enrollment stations is another option, because an individual can easily and securely prove his or her identity with a passport or driver's license, although this can become infeasible on a larger scale. To achieve a similar level of assurance, a large organization

can implement an "enroll-on-behalf-of" strategy, in which employees are enrolled with their credentials by a superior who can personally verify their identities. This creates a chain of trust that ensures individuals are checked in person against their proposed identities, but without the administrative strain of provisioning all virtual smart cards from a single central enrollment station.

For deployments in which a high-assurance level isn't a primary concern, you can use self-service solutions. These can include using an online portal to obtain credentials or simply enrolling for certificates by using Certificate Manager, depending on the deployment. Consider that virtual smart card authentication is only as strong as the method of provisioning. For example, if weak domain credentials (such as a password alone) are used to request the authentication certificate, virtual smart card authentication will be equivalent to using only the password, and the benefits of two-factor authentication are lost.

For information about using Certificate Manager to configure virtual smart cards, see Get Started with Virtual Smart Cards: Walkthrough Guide.

High-assurance and self-service solutions approach virtual smart card provisioning by assuming that the user's computer has been issued prior to the virtual smart card deployment, but this isn't always the case. If virtual smart cards are being deployed with new computers, they can be created, personalized, and provisioned on the computer before the user has contact with that computer.

In this situation, provisioning becomes relatively simple, but identity checks must be put in place to ensure that the recipient of the computer is the individual who was expected during provisioning. This can be accomplished by requiring the employee to set the initial PIN under the supervision of the deployment administrator or manager.

When you're provisioning your computers, you should also consider the longevity of credentials that are supplied for virtual smart cards. This choice must be based on the risk threshold of the organization. Although longer lived credentials are more convenient, they're also more likely to become compromised during their lifetime. To decide on the appropriate lifetime for credentials, the deployment strategy must take into account the vulnerability of their cryptography (how long it could take to crack the credentials), and the likelihood of attack.

For compromised virtual smart cards, administrators should be able to revoke the associated credentials, like they would with a lost or stolen laptop. Revoking credentials requires a record of which credentials match which user and device, but the functionality doesn't natively exist in Windows. Deployment administrators might want to consider add-on solutions to maintain a record.

Virtual smart cards on consumer devices used for corporate access

There are techniques that allow employees to provision virtual smart cards and enroll for certificates that can be used to authenticate the users. This is useful when employees attempt to access corporate resources from devices that aren't joined to the corporate domain. Those devices can be further defined to not allow users to download and run applications from sources other than the Microsoft Store.

You can use APIs to build Microsoft Store apps that you can use to manage the full lifecycle of virtual smart cards. For more information, see Create and delete virtual smart cards programmatically.

TPM ownerAuth in the registry

When a device or computer isn't joined to a domain, the TPM ownerAuth is stored in the registry under HKEY_LOCAL_MACHINE. This exposes some threats. Most of the threat vectors are protected by BitLocker, but threats that aren't protected include:

- A malicious user possesses a device that has an active local sign-in session before the device locks. The malicious user could attempt a brute-force attack on the virtual smart card PIN, and then access the corporate secrets.
- A malicious user possesses a device that has an active virtual private network (VPN) session. The device is then compromised.

The proposed mitigation for the previous scenarios is to use Exchange ActiveSync (EAS) policies to reduce the automatic lockout time from five minutes to 30 seconds of inactivity. You can set policies for automatic lockout while provisioning virtual smart cards. If an organization wants more security, they can also configure a setting to remove the ownerAuth from the local device.

For configuration information about the TPM ownerAuth registry key, see the Group Policy setting **Configure the level of TPM owner authorization information** available to the operating system.

For information about EAS policies, see Exchange ActiveSync Policy Engine Overview.

Managed and unmanaged cards

The following table describes the important differences between managed and unmanaged virtual smart cards that exist on consumer devices:

Operation	Managed and unmanaged cards	Unmanaged cards
Reset PIN when the user forgets the PIN	Yes	No. Delete and recreate the card.
Allow user to change the PIN	Yes	No. Delete and recreate the card.

Managed cards

A managed virtual smart card can be serviced by the IT administrator or another person in that designated role. It allows the IT administrator to have influence or complete control over specific aspects of the virtual smart card from its creation to deletion. To manage these cards, a virtual smart card deployment management tool is often required.

Managed card creation

A user can create blank virtual smart card by using the *Tpmvscmgr* command-line tool, which is a built-in tool executed with administrative credentials through an elevated command prompt. The virtual smart card must be created with well-known parameters (such as default values), and it should be left unformatted (specifically, the **/generate** option shouldn't be specified).

The following command creates a virtual smart card that can later be managed by a smart card management tool launched from another computer (as explained in the next section):

```
tpmvscmgr.exe create /name "VirtualSmartCardForCorpAccess" /AdminKey DEFAULT /PIN PROMPT
```

Alternatively, instead of using a default administrator key, a user can enter an administrator key at the command line:

```
tpmvscmgr.exe create /name "VirtualSmartCardForCorpAccess" /AdminKey PROMPT /PIN PROMPT
```

In either case, the card management system needs to be aware of the initial administrator key. The requirement is so that the card management system can take ownership of the virtual smart card and change the administrator key to a value that is only accessible through the card management tool operated by the IT administrator. For example, when you use the default, the administrator key is set to:

`1020304050607080102030405060708010203040506070 8`

For information about using this command-line tool, see Tpmvscmgr.

Managed card management

After the virtual smart card is created, the user needs to open a remote desktop connection to an enrollment station, for example, in a computer that is joined to the domain. Virtual smart cards that are associated with a client computer are available for use in the remote desktop connection. The user can open a card management tool inside the remote session that can take ownership of the card and provision it for use by the user. This requires that a user is allowed to establish a remote desktop connection from a non-domain-joined computer to a domain-joined computer. This might require a specific network configuration, such as through IPsec policies.

When users need to reset or change a PIN, they need to use the remote desktop connection to complete these operations. They can use the built-in tools for PIN unlock and PIN change or the smart card management tool.

Certificate management for managed cards

Similar to physical smart cards, virtual smart cards require certificate enrollment.

Certificate issuance

Users can enroll for certificates from within a remote desktop session that is established to provision the card. This process can also be managed by the smart card management tool that the user runs through the remote desktop connection. This model works for deployments that require the user to sign a request for enrollment by using a physical smart card. The driver for the physical smart card doesn't need to be installed on the client computer if it's installed on the remote computer. This is made possible by smart card redirection functionality that was introduced in Windows Server 2003, which ensures that smart cards that are connected to the client computer are available for use during a remote session.

Alternatively, without establishing a remote desktop connection, users can enroll for certificates from the Certificate Management console (certmgr.msc) on a client computer. Users can also create a request and submit it to a server from within a custom certificate enrollment application (for example, a registration authority) that has controlled access to the certification authority (CA). This requires specific enterprise

configuration and deployments for Certificate Enrollment Policies (CEP) and Certificate Enrollment Services (CES).

Certificate lifecycle management

You can renew certificates through remote desktop connections, certificate enrollment policies, or certificate enrollment services. Renewal requirements could be different from the initial issuance requirements, based on the renewal policy.

Certificate revocation requires careful planning. When information about the certificate to be revoked is reliably available, the specific certificate can be easily revoked. When information about the certificate to be revoked isn't easy to determine, all certificates that are issued to the user under the policy that was used to issue the certificate might need to be revoked. For example, this could occur if an employee reports a lost or compromised device, and information that associates the device with a certificate isn't available.

Unmanaged cards

Unmanaged virtual smart cards aren't serviceable by an IT administrator. Unmanaged cards might be suitable if an organization doesn't have an elaborate smart card deployment management tool and using remote desktop connections to manage the card isn't desirable. Because unmanaged cards aren't serviceable by the IT administrator, when a user needs help with a virtual smart card (for example, resetting or unlocking a PIN), the only option available to the user is to delete the card and create it again. This results in loss of the user's credentials and he or she must re-enroll.

Unmanaged card creation

A user can create a virtual smart card by using the Tpmvscmgr command-line tool, which is run with administrative credentials through an elevated command prompt. The following command creates an unmanaged card that can be used to enroll certificates:

```
tpmvscmgr.exe create /name "VirtualSmartCardForCorpAccess" /AdminKey RANDOM /PIN PROMPT /generate
```

This command creates a card with a randomized administrator key. The key is automatically discarded after the creation of the card. If users forget or want to change their PIN, they need to delete the card and create it again. To delete the card, a user can run the following command:

```
tpmvscmgr.exe destroy /instance <instance ID>
```

where <instance ID> is the value that is printed on the screen when the user creates the card. Specifically, for the first card created, the instance ID is ROOT\SMARTCARDREADER\0000).

Certificate management for unmanaged cards

Depending on the security requirements that are unique to an organization, users can initially enroll for certificates from the certificate management console (certmgr.msc) or from within custom certificate enrollment applications. The latter method can create a request and submit it to a server that has access to the Certification Authority. This requires specific organizational configurations and deployments for certificate enrollment policies and certificate enrollment services. Windows has built-in tools, specifically Certreq.exe and Certutil.exe, which can be used by scripts to perform the enrollment from the command line.

Requesting the certificate by providing domain credentials only

The simplest way for users to request certificates is to provide their domain credentials through a script that can perform the enrollment through built-in components you have in place for certificate requests.

Alternatively, an application (such as a line-of-business app) can be installed on the computer to perform enrollment by generating a request on the client. The request is submitted to an HTTP server, which can forward it to a registration authority.

Another option is to have the user access an enrollment portal that is available through Internet Explorer. The webpage can use the scripting APIs to perform certificate enrollment.

Signing the request with another certificate

You can provide users with a short-term certificate through a Personal Information Exchange (.pfx) file. You can generate the .pfx file by initiating a request from a domain-joined computer. You can enforce other policy constraints on the .pfx file to assert the identity of the user.

The user can import the certificate into the **MY** store (which is the user's certificate store). And your organization can present the user with a script that can be used to sign the request for the short-term certificate and to request a virtual smart card.

For deployments that require users to use a physical smart card to sign the certificate request, you can use the procedure:

1. Users initiate a request on a domain-joined computer.
2. Users complete the request by using a physical smart card to sign the request.
3. Users download the request to the virtual smart card on their client computer.

Using one-time password for enrollment

Another option to ensure that users are strongly authenticated before virtual smart card certificates are issued, is to send a user a one-time password through SMS, email, or phone. The user then types the one-time password during the certificate enrollment from an application or a script on a desktop that invokes built-in command-line tools.

Certificate lifecycle management

Certificate renewal can be done from the same tools that are used for the initial certificate enrollment. Certificate enrollment policies and certificate enrollment services can also be used to perform automatic renewal.

Certificate revocation requires careful planning. When information about the certificate to be revoked is reliably available, the specific certificate can be easily revoked. When information about the certificate to be revoked isn't easy to determine, all certificates issued to the user under the policy that was used to issue the certificate might need to be revoked. For example, if an employee reports a lost or compromised device, and information that associates the device with a certificate isn't available.

Maintain virtual smart cards

Maintenance is a significant portion of the virtual smart card lifecycle and one of the most important considerations from a management perspective. After virtual smart cards are created, personalized, and provisioned, they can be used for convenient two-factor authentication. Deployment administrators must be aware of several common administrative scenarios, which can be approached by using a purchased virtual smart card solution or on a case-by-case basis with in-house methods.

Renewal: Renewing virtual smart card credentials is a regular task that is necessary to preserve the security of a virtual smart card deployment. Renewal is the result of a signed request from a user who specifies the key pair desired for the new credentials. Depending on user's choice or deployment specification, the user can request

credentials with the same key pair as previously used, or choose a newly generated key pair.

When renewing with a previously used key, no extra steps are required because a strong certificate with this key was issued during the initial provisioning. However, when the user requests a new key pair, you must take the same steps that were used during provisioning to assure the strength of the credentials. Renewal with new keys should occur periodically to counter sophisticated long-term attempts by malicious users to infiltrate the system. When new keys are assigned, you must ensure that the new keys are being used by the expected individuals on the same virtual smart cards.

Resetting PINs: Resetting virtual smart card PINs is also a frequent necessity, because employees forget their PINs. There are two ways to accomplish this, depending on choices made earlier in the deployment: Use a PUK (if the PUK is set), or use a challenge-response approach with the administration key. Before resetting the PIN, the user's identity must be verified by using some means other than the card—most likely the verification method that you used during initial provisioning (for example, in-person proofing). This is necessary in user-error scenarios when users forget their PINs. However, you should never reset a PIN if it has been compromised because the level of vulnerability after the PIN is exposed is difficult to identify. The entire card should be reissued.

Lockout reset: A frequent precursor to resetting a PIN is the necessity of resetting the TPM lockout time because the TPM anti-hammering logic will be engaged with multiple PIN entry failures for a virtual smart card. This is currently device specific.

Retiring cards: The final aspect of virtual smart card management is retiring cards when they're no longer needed. When an employee leaves the company, it's desirable to revoke domain access. Revoking sign-in credentials from the certification authority (CA) accomplishes this goal.

The card should be reissued if the same computer is used by other employees without reinstalling the operating system. Reusing the former card can allow the former employee to change the PIN after leaving the organization, and then hijack certificates that belong to the new user to obtain unauthorized domain access. However, if the employee takes the virtual smart card-enabled computer, it's only necessary to revoke the certificates that are stored on the virtual smart card.

Emergency preparedness

Card reissuance

The most common scenario in an organization is reissuing virtual smart cards, which can be necessary if the operating system is reinstalled or if the virtual smart card is compromised in some manner. Reissuance is essentially the recreation of the card, which involves establishing a new PIN and administrator key and provisioning a new set of associated certificates. This is an immediate necessity when a card is compromised, for example, if the virtual smart card-protected computer is exposed to an adversary who might have access to the correct PIN. Reissuance is the most secure response to an unknown exposure of a card's privacy. Additionally, reissuance is necessary after an operating system is reinstalled because the virtual smart card device profile is removed with all other user data when the operating system is reinstalled.

Blocked virtual smart card

The anti-hammering behavior of a TPM virtual smart card is different from that of a physical smart card. A physical smart card blocks itself after the user enters the wrong PIN a few times. A TPM virtual smart card enters a timed delay after the user enters the wrong PIN a few times. If the TPM is in the timed-delay mode, when the user attempts to use the TPM virtual smart card, the user is notified that the card is blocked. Furthermore, if you enable the integrated unlock functionality, the user can see the user interface to unlock the virtual smart card and change the PIN. Unlocking the virtual smart card doesn't reset the TPM lockout. The user needs to perform an extra step to reset the TPM lockout or wait for the timed delay to expire.

For more information about setting the Allow Integrated Unblock policy, see Allow Integrated Unblock screen to be displayed at the time of logon.

Evaluate Virtual Smart Card Security

Article • 02/27/2023 • 3 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

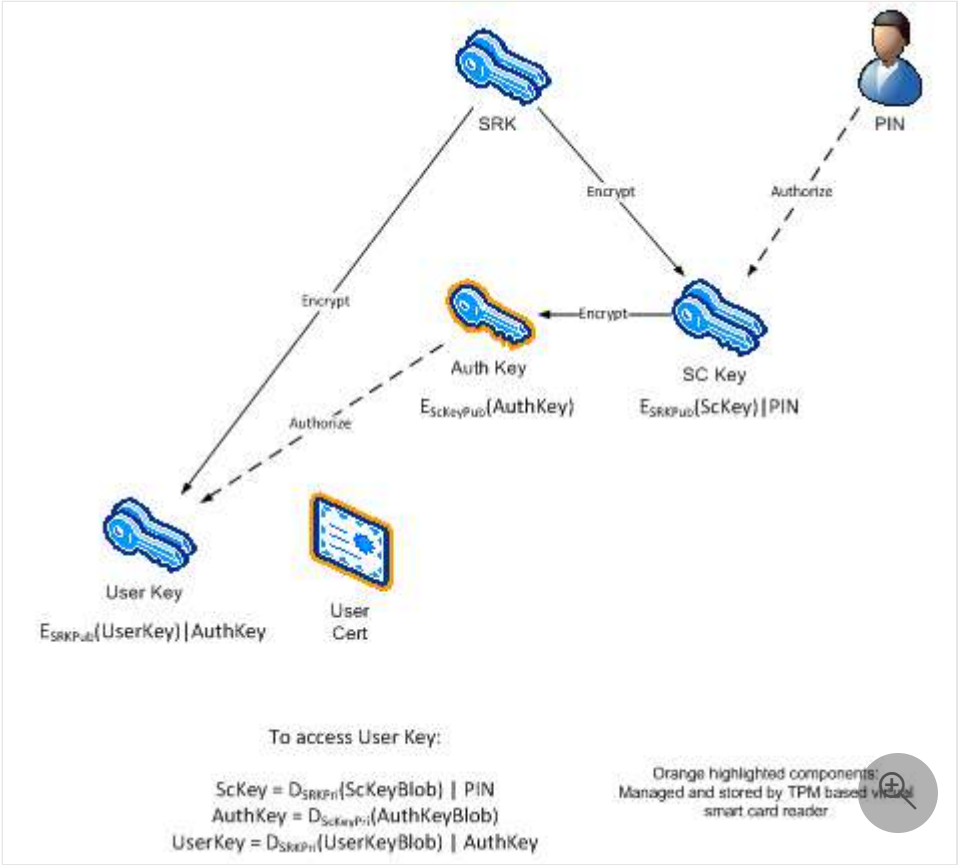
> ⚠ **Warning**
>
> **Windows Hello for Business** is the modern, two-factor authentication for Windows. Microsoft will deprecate virtual smart cards in the near future. Customers using virtual smart cards are strongly encouraged to move to Windows Hello for Business. Microsoft will publish the deprecation date to ensure customers have adequate lead time to move to Windows Hello for Business. We recommend that new Windows deployments use Windows Hello for Business.

In this article, you'll learn about security characteristics and considerations when deploying TPM virtual smart cards.

Virtual smart card non-exportability details

A crucial aspect of TPM virtual smart cards is their ability to securely store and use secret data. Specifically, that the secured data is non-exportable.
Data can be accessed and used within the virtual smart card system, but it's meaningless outside of its intended environment. In TPM virtual smart cards, security is ensured with a secure key hierarchy, which includes several chains of encryption. The chain originates with the TPM storage root key, which is generated and stored within the TPM and never exposed outside the chip. The TPM key hierarchy is designed to allow encryption of user data with the storage root key, but it authorizes decryption with the user PIN so that changing the PIN doesn't require re-encryption of the data.

The following diagram illustrates the secure key hierarchy and the process of accessing the user key.



The following keys are stored on the hard disk:

- User key
- Smart card key, which is encrypted by the storage root key
- Authorization key for the user key decryption, which is encrypted by the public portion of the smart card key

When the user enters a PIN, the use of the decrypted smart card key is authorized with this PIN. If this authorization succeeds, the decrypted smart card key is used to decrypt the auth key. The auth key is then provided to the TPM to authorize the decryption and use of the user's key that is stored on the virtual smart card.

The auth key is the only sensitive data used as plaintext outside the TPM, but its presence in memory is protected by Microsoft Data Protection API (DPAPI), such that before being stored in any way, it's encrypted. All data other than the auth key is processed only as plaintext within the TPM, which is isolated from external access.

Virtual smart card anti-hammering details

The anti-hammering functionality of virtual smart cards relies on the anti-hammering functionality of the TPM that is enabling the virtual smart card. However, the TPM version 1.2 and subsequent specifications (as designed by the Trusted Computing Group) provide flexible guidelines for responding to hammering. The spec requires only

that the TPM implement protection against trial-and-error attacks on the user PIN, PUK, and challenge/response mechanism.

The Trusted Computing Group specifies that if the response to attacks involves suspending proper function of the TPM for some period of time, or until administrative action is taken, the TPM must prevent running the authorized TPM commands. The TPM can prevent running any TPM commands until the termination of the attack response. Beyond using a time delay or requiring administrative action, a TPM could also force a reboot when an attack is detected. The Trusted Computing Group allows manufacturers a level of creativity in their choice of implementation. The methodology used by TPM manufacturers determines the anti-hammering response of TPM virtual smart cards. Some typical aspects of protection from attacks include:

1. Allow only a limited number of wrong PIN attempts before enabling a lockout that enforces a time delay before any further commands are accepted by the TPM.

   > ⓘ **Note**
   >
   > If the user enters the wrong PIN five consecutive times for a virtual smart card (which works in conjunction with the TPM), the card is blocked. When the card is blocked, it must be unblocked by using the administrative key or the PUK.

2. Increase the time delay exponentially as the user enters the wrong PIN so that an excessive number of wrong PIN attempts quickly trigger long delays in accepting commands.

3. Have a failure leakage mechanism to allow the TPM to reset the timed delays over a period of time. This is useful in cases where a valid user has entered the wrong PIN occasionally, for example, due to complexity of the PIN.

For example, it will take 14 years to guess an eight character PIN for a TPM that implements the following protection:

1. Number of wrong PINs allowed before entering lockout (threshold): 9
2. Time the TPM is in lockout after the threshold is reached: 10 seconds
3. Timed delay doubles for each wrong PIN after the threshold is reached

Tpmvscmgr

Article • 02/27/2023 • 5 minutes to read •

Applies to: ✅ Windows 10 and later, ✅ Windows Server 2016 and later

> ⚠ **Warning**
>
> **Windows Hello for Business** is the modern, two-factor authentication for Windows. Microsoft will deprecate virtual smart cards in the near future. Customers using virtual smart cards are strongly encouraged to move to Windows Hello for Business. Microsoft will publish the deprecation date to ensure customers have adequate lead time to move to Windows Hello for Business. We recommend that new Windows deployments use Windows Hello for Business.

The Tpmvscmgr command-line tool allows users with Administrative credentials to create and delete TPM virtual smart cards on a computer. For examples of how this command can be used, see Examples.

Syntax

```
Tpmvscmgr create [/quiet] /name <name> /AdminKey {DEFAULT | PROMPT | RANDOM} [/PIN {DEFAULT | PROMPT}] [/PUK {DEFAULT | PROMPT}] [/generate] [/machine <machine name>] [/pinpolicy [policy options]] [/attestation {AIK_AND_CERT | AIK_ONLY}] [/?]
```

```
Tpmvscmgr destroy [/quiet] [/instance <device instance ID>] [/machine <machine name>] [/?]
```

Parameters for Create command

The Create command sets up new virtual smart cards on the user's system. It returns the instance ID of the newly created card for later reference if deletion is required. The instance ID is in the format `ROOT\SMARTCARDREADER\000n` where n starts from 0 and is increased by 1 each time you create a new virtual smart card.

Parameter	Description
/name	Required. Indicates the name of the new virtual smart card.

Parameter	Description
/AdminKey	Indicates the desired administrator key that can be used to reset the PIN of the card if the user forgets the PIN. **DEFAULT** Specifies the default value of 010203040506070801020304050607080102030405060708. **PROMPT** Prompts the user to enter a value for the administrator key. **RANDOM** Results in a random setting for the administrator key for a card that is not returned to the user. This creates a card that might not be manageable by using smart card management tools. When generated with RANDOM, the administrator key is set as 48 hexadecimal characters.
/PIN	Indicates desired user PIN value. **DEFAULT** Specifies the default PIN of 12345678. **PROMPT** Prompts the user to enter a PIN at the command line. The PIN must be a minimum of eight characters, and it can contain numerals, characters, and special characters.
/PUK	Indicates the desired PIN Unlock Key (PUK) value. The PUK value must be a minimum of eight characters, and it can contain numerals, characters, and special characters. If the parameter is omitted, the card is created without a PUK. **DEFAULT** Specifies the default PUK of 12345678. **PROMPT** Prompts the user to enter a PUK at the command line.
/generate	Generates the files in storage that are necessary for the virtual smart card to function. If the /generate parameter is omitted, it's equivalent to creating a card without this file system. A card without a file system can be managed only by a smart card management system such as Microsoft Configuration Manager.
/machine	Allows you to specify the name of a remote computer on which the virtual smart card can be created. This can be used in a domain environment only, and it relies on DCOM. For the command to succeed in creating a virtual smart card on a different computer, the user running this command must be a member in the local administrators group on the remote computer.
/pinpolicy	If **/pin prompt** is used, **/pinpolicy** allows you to specify the following PIN policy options: **minlen** <minimum PIN length> If not specified, defaults to 8. The lower bound is 4. **maxlen** <maximum PIN length> If not specified, defaults to 127. The upper bound is 127. **uppercase** Can be **ALLOWED**, **DISALLOWED**, or **REQUIRED.** Default is **ALLOWED.** **lowercase** Can be **ALLOWED**, **DISALLOWED**, or **REQUIRED.** Default is **ALLOWED.** **digits** Can be **ALLOWED**, **DISALLOWED**, or **REQUIRED.** Default is **ALLOWED.** **specialchars** Can be **ALLOWED**, **DISALLOWED**, or **REQUIRED.** Default is **ALLOWED.** When using **/pinpolicy**, PIN characters must be printable ASCII characters.

Parameter	Description
/attestation	Configures attestation (subject only). This attestation uses an Attestation Identity Key (AIK) certificate as a trust anchor to vouch that the virtual smart card keys and certificates are truly hardware bound. The attestation methods are: **AIK_AND_CERT** Creates an AIK and obtains an AIK certificate from the Microsoft cloud certification authority (CA). This requires the device to have a TPM with an EK certificate. If this option is specified and there's no network connectivity, it's possible that creation of the virtual smart card will fail. **AIK_ONLY** Creates an AIK but doesn't obtain an AIK certificate.
/?	Displays Help for this command.

Parameters for Destroy command

The Destroy command securely deletes a virtual smart card from a computer.

> ⚠ **Warning**
>
> When a virtual smart card is deleted, it cannot be recovered.

Parameter	Description
/instance	Specifies the instance ID of the virtual smart card to be removed. The instanceID was generated as output by Tpmvscmgr.exe when the card was created. The **/instance** parameter is a required field for the Destroy command.
/machine	Allows you to specify the name of a remote computer on which the virtual smart card will be deleted. This can be used in a domain environment only, and it relies on DCOM. For the command to succeed in deleting a virtual smart card on a different computer, the user running this command must be a member in the local administrators group on the remote computer.
/?	Displays Help for this command.

Remarks

Membership in the Administrators group (or equivalent) on the target computer is the minimum required to run all the parameters of this command.

For alphanumeric inputs, the full 127 character ASCII set is allowed.

Examples

The following command shows how to create a virtual smart card that can be later managed by a smart card management tool launched from another computer.

```
Console

tpmvscmgr.exe create /name "VirtualSmartCardForCorpAccess" /AdminKey DEFAULT
/PIN PROMPT
```

Alternatively, instead of using a default administrator key, you can create an administrator key at the command line. The following command shows how to create an administrator key.

```
Console

tpmvscmgr.exe create /name "VirtualSmartCardForCorpAccess" /AdminKey PROMPT
/PIN PROMPT
```

The following command will create the unmanaged virtual smart card that can be used to enroll certificates.

```
Console

tpmvscmgr.exe create /name "VirtualSmartCardForCorpAccess" /AdminKey RANDOM
/PIN PROMPT /generate
```

The preceding command will create a virtual smart card with a randomized administrator key. The key is automatically discarded after the card is created. This means that if the user forgets the PIN or wants to the change the PIN, the user needs to delete the card and create it again. To delete the card, the user can run the following command.

```
Console

tpmvscmgr.exe destroy /instance <instance ID>
```

where <instance ID> is the value printed on the screen when the user created the card. Specifically, for the first card created, the instance ID is ROOT\SMARTCARDREADER\0000.

The following command will create a TPM virtual smart card with the default value for the administrator key and a specified PIN policy and attestation method:

```
Console
```

```
tpmvscmgr.exe create /name "VirtualSmartCardForCorpAccess" /PIN PROMPT
/pinpolicy minlen 4 maxlen 8 /AdminKey DEFAULT /attestation AIK_AND_CERT
/generate
```

GREYCORTEX Mendel Data Sheet

NDR

Mendel od společnosti GREYCORTEX, NDR řešení pro firmy, státní správu a kritickou infrastrukturu nabízí hlubokou viditelnost v síti, pokročilou detekci a reakci na hrozby.

ZDROJE DAT

- Zrcadlený síťový provoz
- NetFlow/IPFIX
- Síťové a aplikační protokoly
- Zaznamenaný síťový provoz (PCAP)

DETAILNÍ VIDITELNOST SÍTĚ

- Všechny podsítě, hosté, služby a toky obohacené o podrobné informace
- Metadata poskytují bohaté informace o chování sítě pro účely forenzního vyšetřování, zjištění souladu s předpisy atd.
- Tunelovaný provoz
- Dešifrování šifrované komunikace pomocí příslušného klíče
- Automatická identifikace kritických zařízení v síti, jako je Active Directory, e-mailový server, SMB server atd.
- Historická data jsou indexována a tedy rychle dostupná pro zobrazení
- Široké možnosti filtrování nad uloženými daty, přitom uživatelsky přívětivé

VÝKONNÁ DETEKCE

na základě procesů strojového učení, matematických modelů, s eliminací falešně pozitivních výsledků

PREDIKTIVNÍ ANALÝZA

- Volumetrické anomálie
- DDoS útoky

Využívá pokročilé algoritmy umělé inteligence a strojového učení k předvídání a identifikaci potenciálních bezpečnostních hrozeb dříve, než se projeví jako plnohodnotný útok.

ANALÝZA ZMĚN NA SÍTI

- Nová zařízení, služby a komunikační vektory
- Nové administrátorské přístupy

Proces systematického zkoumání a mapování síťového prostředí s cílem identifikovat aktiva, zranitelná místa a potenciální místa kompromitace.

ANALÝZA TOKŮ

- Skeny
- Útoky hrubou silou
- Enumerace

Investigace a interpretace toků síťového provozu s cílem zajistit přehled o komunikačních vzorcích, chování a potenciálních bezpečnostních hrozbách.

REPETITIVNÍ ANALÝZA

- Command & Conquer
- Útoky hrubou silou

Identifikace a analýza opakujících se vzorců nebo chování v síťovém provozu s cílem odhalit potenciální bezpečnostní hrozby nebo anomálie.

ANALÝZA NA ZÁKLADĚ PRAVIDEL

- Detekce známého malwaru a exploitů
- Porušení bezpečnostních zásad
- Security policies violations

Metoda identifikace a vyhodnocování bezpečnostních hrozeb na základě předem definovaných pravidel nebo kritérií pro detekci a reakci v reálném čase.

ANALÝZA VÝKONU

- Problémy s výkonem sítě: vysoká doba odezvy, slow round trip time

Vyhodnocování a optimalizace provozní efektivity, spolehlivosti a škálovatelnosti bezpečnostních systémů a síťové infrastruktury prostřednictvím průběžného sledování výkonnostních ukazatelů.

ANALÝZA LOGŮ

- Nový firmware, noví uživatelé
- Změny v registrech a nastaveních

Shromažďování, analýza a rozbor dat protokolů generovaných různými síťovými zařízeními, systémy a aplikacemi za účelem identifikace bezpečnostních incidentů, anomálií a provozních problémů.

REAKCE

- Centralizovaná platforma pro sledování, dokumentování a hlášení bezpečnostních incidentů.
- Systém pluginů pro univerzální způsob interakce s firewally, NAC atd.
- Rozsáhlé možnosti exportu dat pro sdílení dat s jinými bezpečnostními nástroji

KLÍČOVÉ SCHOPNOSTI

Analýza chování sítě

Analýza síťového provozu založená na tocích prostřednictvím „unsupervised“ strojového učení a několika dalších detekčních metod (viz výše).

Detekční schopnosti:

- Aktivita malwaru – šíření, stahování, spamování atd.
- Aktivita útočníka – skenování, brute-force, exploitace atd.
- Aktivita C&C – RAT, APT, AVT, boti, červi, rootkity atd.
- Exfiltrace dat

Hloubková kontrola packetů

- Monitoruje jakoukoli interakci s interní sítí nebo uvnitř ní
- Umožňuje kontrolovat provoz až do rychlosti 100 Gbit/s
- Detekční signatury pro malware, porušení politik, útoky a další závadné aktivity
- Detekce škodlivých souborů pomocí hashování
- Přehled komunikace s hosty na blacklistu
- Možnost přidávat signatury (pravidla) vytvořené uživateli

Network Inventory

- Sloučení vrstev Viditelnost a Detekce do jednoho společného pohledu
- Zobrazení síťové infrastruktury z pohledu podsítí a zařízení, doplněná o rizika a další bezpečnostní pohledy
- Přehledná interpretace dat formou srozumitelných tabulek a grafických přehledů

Monitorování výkonu

Analýza výkonu sítě a aplikací založená na tocích (NPM, APM):

- Povědomí o aplikacích
- Monitorování aktuální a průměrné šířky pásma
- Monitorování metrik výkonu jako jsou doba odezvy aplikace, round-trip, user-experience
- Detekce založená na pravidlech (například SLA)
- Automatická detekce založená na anomáliích

Historická metadata a forenzní analýza

Proprietární protokol Advanced Security Network Metrics (ASNM) je zaměřený na bezpečnost a výkon a slouží pro širší popis síťového provozu.

Zahrnuje:

- Obousměrný záznam toku (jeden tok obsahuje request i response)
- Metadata aplikačních protokolů pro FTP, SSH, Telnet, SMTP, DNS, DHCP, HTTP/S, NTP, SMB, SNMP, LDAP, NFS, MS-SQL, SIP, SSL/TLS, Kerberos atd.
- Retence dat v rozsahu měsíců až několika let (v závislosti na kapacitě úložiště)

Kategorizace zařízení

- Rozšířená klasifikace zařízení a jejich rolí prostřednictvím systémových či vlastních tagů
- Dynamická viditelnost díky sledování nových aktivit nebo změn způsobených zařízeními komunikujícími ve vaší síti
- Ruční nebo automatizovaný způsob označování hostitelů nebo podsítí podle pravidel vytvořených uživatelem pomocí snadno pochopitelného průvodce

Záznam a přehrávání provozu

- Zachycení paketů na základě definované podmínky (události) nebo podle dalších pravidel na základě zdrojové a cílové IP, MAC, protokolu, portu atd.
- Možnost přehrávat a analyzovat PCAPy zaznamenané samotným nástrojem Mendel nebo zpracovat přenosy zachycené jinými nástroji

OT schopnosti

- Detekce, parsování a hlášení událostí včetně bohatých metadat pro průmyslové protokoly BACnet, CC-link, COAP, DLMS/COSEM, DNP3, ENIP, EtherCAT, GE-STRP, IEC-104, IEC61850 (GOOSE, SV, MMS), MODBUS, MQTT, OMRON FINS, OPC UA, Profinet IO DCE/RPC, PROFINET-DCP, Siemens S7, SNMP
- Specifické detekční OT signatury definované GREYCORTEX
- Metriky OT aplikací
- Detekce průmyslových zařízení (Asset discovery)
- Společná prezentace dat podle Purdue a MITRE ATT&CK

VSTUPY

Komplexní přístup k analytickému zpracování dat z mnoha různých zdrojů pro zajištění síťové bezpečnosti, a to včetně různých vrstev modelu OSI, využití znalostní báze hrozeb, funkcí dohledu chování jednotlivých sítí a jejich uživatelů pro účinné prosazování nastavených interních politik a zmírňování či eliminaci hrozeb.

Síťová data

- Zrcadlený provoz (TAP, SPAN, RSPAN, ERSPAN nebo jiný typ zrcadleného provozu)
- Podpora linkové vrstvy
- Podpora síťové vrstvy včetně protokolů IPv6
- Dekapsulace tunelovaného provozu
- Podpora transportní vrstvy
- Podpora aplikační vrstvy
- Protokoly založené na Flow (NetFlow family, IPFIX, sFLOW, JFlow, NetStream a protokoly toku VPC)
- Z jiných zařízení Mendel (senzor, mikrosenzor nebo kolektor)

Znalostní báze hrozeb

- Soubor detekčních pravidel Proofpoint Emerging Threats a interní výzkum GREYCORTEX
- Další databáze informací o hrozbách IP adres, doménových reputací a škodlivých souborů
- Možnost integrace vlastních informačních kanálů TI (včetně národních platforem MISP)
- Definice události podle MITRE ATT&CK Enterprise a Mitre ATT&CK ICS

Znalost sítě

- Definice politik podle segmentů/podsítí, které sdílejí stejné vzorce chování, například management, obchod, servery, WiFi, VoIP, tiskárny, DMZ atd.
- Spojení IP s doménou (pomocí DNS záznamů)

Znalost uživatelů

- Spojení IP s názvem hosta (prostřednictvím logů z doménových řadičů, AD, LDAP, CISCO ISE a všechny poskytnuté protokoly s těmito informacemi v rámci sledované sítě od různých služeb)

VÝSTUPY

Systém poskytuje komplexní sadu výstupů a integrací pro efektivní monitorování a správu zabezpečení sítě a zajišťuje účinné monitorování, analýzu a reakci na bezpečnostní incidenty.

Grafické uživatelské rozhraní (GUI)

- Webové uživatelské rozhraní (Firefox, Chrome, Opera, Edge)
- Hlavní interaktivní pohled založený na znalostech GREYCORTEX a frameworku MITRE ATT&CK®.
- Rozsáhlé snadno přizpůsobitelné dashboardy
- Manažerské a bezpečnostní dashboardy pro zjednodušený přehled
- Operativní a široké možnosti filtrování
- Průvodce tvorbou vlastních IDS pravidel
- Dva motivy designu (světlý a tmavý)
- Kontextová nápověda a obsáhlá uživatelská dokumentace

Reporty a notifikace

- Reporty generované na základě definovaných podmínek
- Bohatý a přizpůsobitelný výstupní formát s možností přidání přímého odkazu do GUI
- Formáty vhodné pro koncového uživatele: e-mail (HTML) a PDF

Integrace

- SIEM – formáty CEF, LEEF, Syslog nebo API
- SOAR – přizpůsobitelná integrace založená na exportu událostí a API s integračním balíčkem pro Palo Alto XSOAR
- XDR – přizpůsobitelná integrace s platformami EDR na základě rozhraní API nebo podle specifik dodavatele
- IPFIX – export toků ve formátu IPFIX
- IDENTITA – Active Directory, Cisco ISE a běžné externí protokoly pro doplnění identity uživatele
- FIREWALL/NAC – MikroTik, Juniper, FortiGate, Palo Alto, Checkpoint atd.
- OUTPUT – Přizpůsobitelné datové výstupy
- API – Obecné rozhraní RESTful API pro integraci s další infrastrukturou

NASAZENÍ

Škálovatelnost nasazení se liší podle konkrétních podmínek a kombinací v dané infrastruktuře.

Senzor

- Monitorování sítě s propustností až 100 Gbps
- Až 20 monitorovacích rozhraní na každém HW zařízení s libovolnou kombinací 1GbE, 10GbE, 25GbE nebo s HW akcelerovanými kartami FPGA až 8× 25/10GbE nebo 4×100/40GbE
- V režimu virtualizace nebo Cloud nasazení zpracování až 10 Gbps
- Volitelně malý HW mikrosenzor s monitorovacími porty až 3× 1GbE

Kolektor

- Až 100 senzorů pod jedním kolektorem (na základě celkové zpracovávané kapacity)
- Až 150 000 monitorovaných síťových uzlů na jeden kolektor
- Uchovávání historických dat je omezeno pouze velikostí dostupného (navrženého) datového úložiště
- Nasazení ve virtualizovaném prostředí (včetně Cloudu) se zpracováním až 20 připojených senzorů
- Úložiště s více diskovými oddíly s podporou rychlých disků (NVMe, SSD, SAS)
- Možnost online a offline aktualizace samotného systému nebo připojených senzorů

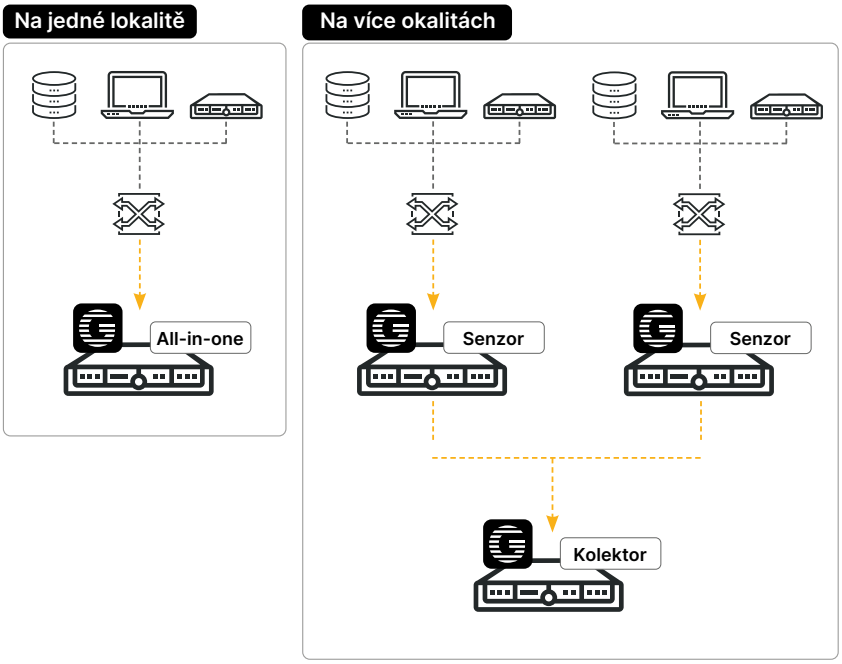
All-in-One

- Jedno zařízení obsahující zároveň senzor i kolektor
- Monitorování sítě s propustností až 50 Gbps
- Až 20 monitorovacích rozhraní na každém HW zařízení s libovolnou kombinací 1GbE, 10GbE, 25GbE nebo s HW akcelerovanými kartami FPGA až 8× 25/10GbE nebo 4×100/40GbE
- Až 20 připojených dalších senzorů na jedno All-in-One zařízení
- Až 50 000 monitorovaných uzlů na jedno All-in-One zařízení
- Úložiště s více diskovými oddíly s podporou rychlých disků (NVMe, SSD, SAS)
- Možnost online a offline aktualizace samotného systému nebo připojených senzorů

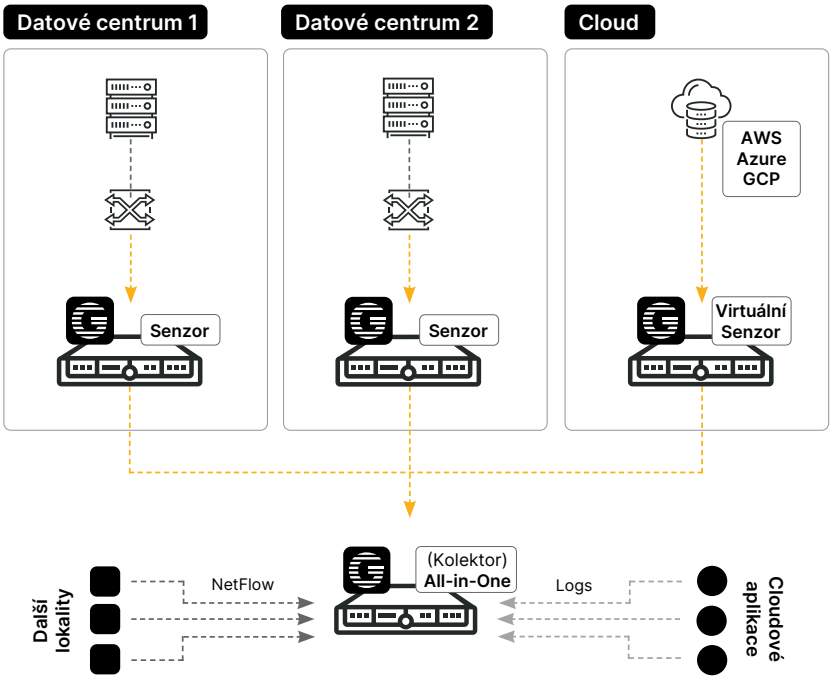
Centrální správa událostí

- Spojení až 20 kolektorů do jednoho centrálního místa
- Přehled a správa všech událostí na jednom místě z celé připojené infrastruktury

Nasazení



Nasazení na více lokalitách



Vyzkoušejte GREYCORTEX Mendel

Proof of Concept (PoC)
Pro vyzkoušení GREYCORTEX Mendel vám nabízíme měsíční bezpečnostní audit sítě.

Copyright © 2024 GREYCORTEX s.r.o. All rights reserved.
Purkyňova 127, 612 00 Brno, Czech Republic, +420 511 205 216, info@greycortex.com

www.greycortex.com

Documentation | GREYCORTEX Mendel

4.3.0

GCX Team

Copyright © 2024 Greycortex

Table of contents

1. User Documentation - GREYCORTEX MENDEL 4.3.0 4
2. Getting started 5
2.1 Introduction 5
2.2 Pre-Deployment 6
2.3 Deployment 7
2.4 Post-Deployment 26
3. User Interface 43
3.1 Overview page 43
3.2 Applications and filters 45
3.3 Dashboard 48
3.4 Inventory 49
3.5 Network 51
3.6 Events 64
3.7 Incidents 75
3.8 Settings 82
4. Setup 148
4.1 Automated Threat Response 148
4.2 Active Survey 150
4.3 Central Event Management 152
4.4 Dashboard Manager 161
4.5 Exports & Integrations 162
4.6 Ipfix Configuration 203
4.7 Multitenant Configuration (SOC) 206
4.8 Network Analysis 208
4.9 Remote Shell 209
4.10 Samba Backup 211
4.11 Samba Recovery 213
4.12 Sensor Collector installation 215
4.13 Sensor Collector limitations 221
4.14 SNMP Monitoring 222
4.15 SNMP Trap 223
4.16 Storage Wipe 224
4.17 Universal Tagging Engine 241
4.18 Cloud Deployment 246
4.19 Operational Technologies 247

Copyright © 2024 Greycortex

4.20 NetFlow Configuration 255
4.21 Napatech Configuration 260
5. How to 262
5.1 End User / Entry level Training 262
6. About 298
6.1 EULA 298
6.2 License 305
6.3 Support 314
6.4 Helpdesk 315
6.5 Release notes 317

Copyright © 2024 Greycortex

1. User Documentation - GREYCORTEX MENDEL 4.3.0

GREYCORTEX MENDEL is NTA/NDR solution that combines advanced detection techniques with full network visibility and integrations with security infrastructure tools to detect threats, visualize communications and respond quickly and effectively. This includes advanced and unknown threats which are able to bypass existing security tools. GREYCORTEX MENDEL helps secure your network, protecting data, reducing risk and maintaining safety.

Let´s start | **How it looks** | **How it works** | **How to ...** | **API**

For GREYCORTEX website visit greycortex.com.

Copyright © 2024 Greycortex

2. Getting started

2.1 Introduction

GREYCORTEX MENDEL is a network traffic analysis tool designed for the operational and security monitoring of computer network traffic. The solution has a few different deployment architectures to move closer to the customers' needs. Find out more detailed information about basic deployment architectures and discover the best option for you.

ALL-IN-ONE

This is the default deployment architecture of MENDEL. The system can be set up easily and everything is ready within the arm's reach. One appliance collects and processes all the data from the network and provides an overview for security analysts about potential risks and anomalies in the network.

SENSOR-COLLECTOR

This is the deployment architecture where two types of appliances are used - sensors and a collector. Multiple sensors are usually connected to one collector. Data received from sensors are processed and displayed on the collector. This solution is most appreciated by companies that have branches all over the world but just one department for analyzing the data from multiple networks.

MULTITENANT

SOC means Security Operation Center and for this kind of organization, there is a special deployment architecture. The main benefit of this deployment is that any user can see only his data based on user security permissions. On the other hand, Analysts can have access to all relevant users and see all the data in one place.

CENTRAL EVENT MANAGEMENT

For customers or partners with larger deployments, MENDEL provides the possibility to connect multiple appliances using a multi-level structure consisting of sensors, collectors, and a Central Event Management console. This gives the administrator of CEM more comprehensive overview of the network.

The sensor collects and analyzes network traffic based on detection signatures and generates records of the flows it stores in the database.

The collector stores all the data from the sensor, analyzes the behavior of individual network data streams and provides all service and user services. The GUI provides reporting and system setup.

The hardware and software part of the solution can be expanded to meet the needs of the network in case of its topology change or increased network traffic.

Copyright © 2024 Greycortex

2.2 Pre-Deployment

2.2.1 Prerequisites

In order to do pre-deployment configuration of MENDEL there is some information which must be received from the end-user, and there are also certain steps which the end-user must perform to ensure proper functionality of MENDEL upon installation.

2.2.2 Network

Network parameters for the management interface of the appliance

Network parameters for iDrac interface of a server

- The information has to be provided for each of the sensors/collectors that will be deployed.
- Network parametres for iDrac interface are not needed if the installation will be done on a virtual machine.

2.2.3 Network - subnets

IP address ranges of internal subnets which will be monitored and their short descriptions (up to 64 characters). These may be sent in a comma separated file. Example: 172.16.10.0/24, Public Wi-Fi

Attribute	Description
	Name must be from 2 to 63 characters long. Valid characters are ASCII letters from a to z, A to Z, digits from 0 to 9, and the hyphen (-).
	If the sensor will access an update server through web proxy, add IP, port number, and user credentials.
	An email to which the license will be bound.

2.2.4 Settings in the monitored infrastructure

Please check/configure the following at the point of installation:

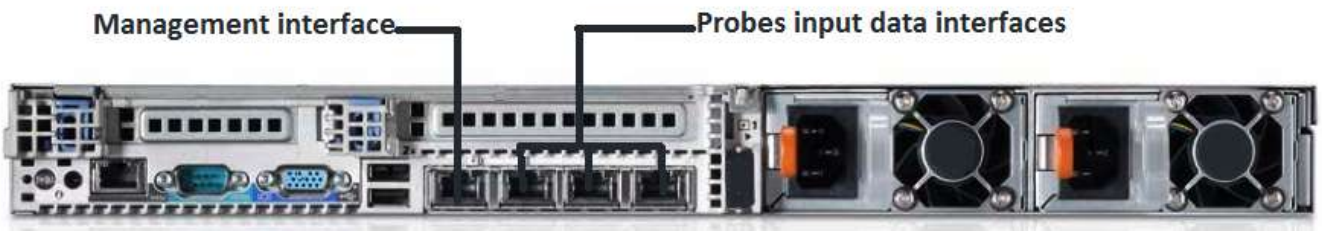
- In case a sensor-collector model will be deployed, then all the sensors must have connectivity in their collector to port 5432/tcp and 80/tcp.
- A mirror of the traffic from a core switch or multiple switches using port mirroring, SPAN or network TAP. SPAN/mirror ports should be prepared before installation of a sensor/AiO.
- Is there a global NTP server used at all of the locations?
- The sensor/collector must have an access to GREYCORTEX update server (which has the DNS name: update.greycortex.com and archive.greycortex.com) at port 443/tcp. The IP address of the server is not static and might change from time to time, so we recommend to check whether your firewall supports DNS names.
- The type of SPAN port medium – Metallic or Optic.
- If an integration with Active Directory will be implemented, then there must be a special account created at the Active Directory server - see Active Directory setup.
- All necessary network cables are available at the installation (at least 3 per sensor or AiO and 2 per collector).

Copyright © 2024 Greycortex

2.3 Deployment

2.3.1 MENDEL Installation

GREYCORTEX MENDEL is distributed as an appliance or installation ISO image file. For a description of how to deploy the appliance to the network, refer to the manual that came with the hardware.



Example of network interface

WIRING REQUIREMENTS

The product mirrors network traffic to monitor the network segment. This data can be obtained from backbone routers or switches using mirror ports or TAP devices duplicating network traffic on the lines. Depending on how you deploy on the network, the server is shipped with 2, 4, or 8 network interfaces to monitor the desired number of subnets.

The solution can be installed in an internal cloud configuration that allows one to monitor geographically remote networks.

Installation from Image File

Before installing the software, prepare a USB/DVD media with a copy of the installation image and the license code assigned to you based on your e-mail address, as well as, the IP address for the network interface of the sensor and the network management interface of the application.

HOW TO CREATE THE INSTALLATION USB DRIVE

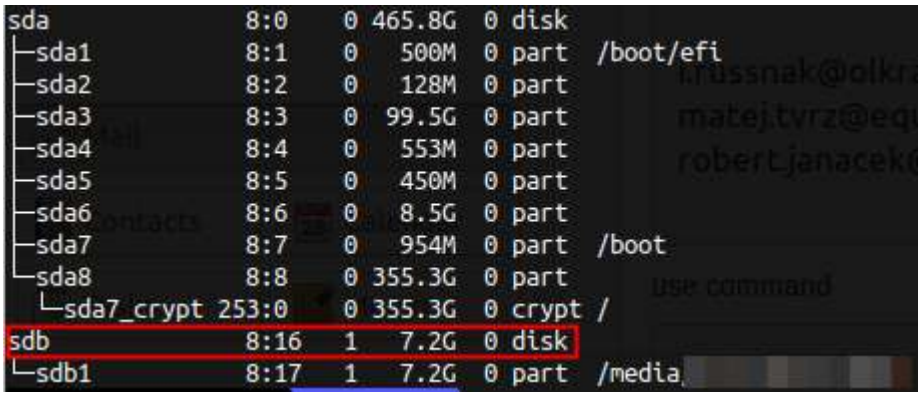
Prerequisites:

- Downloaded MENDEL ISO file.

All files on the USB drive will be deleted during the process! This is our recommended way of creating the installation USB drive.

1. Insert the USB.
2. Open terminal of your choosing.
3. Use lsblk command to obtain the USB drive information.

Copyright © 2024 Greycortex



1. In this example, it is a /dev/sdb.
2. Unmount the flash drive.

1. Use dd command to copy the ISO file onto the USB drive.

1. Wait until the dd command finishes.
2. The USB drive is now ready to be used as an installation drive for MENDEL.

HARDWARE REQUIREMENTS

- In order to pass the hardware requirements test at the beginning of the installation, MENDEL requires a processor with at least 8 CPU cores, 32 GB RAM, a 400 GB hard drive (or two 400 GB hard drives in case of multidisk installation), and two network interfaces. These specifications enable the Mendel appliance to work as a sensor and collector (all-in-one) for a maximum throughput of 500Mbps.
- When installing MENDEL as a virtual server solution, it is also necessary to meet the minimum hardware requirements above. Please remember that all network traffic from the SPAN port must be via a mirrored network bridge from the physical network interface used in the virtual server.
- To verify correct operation of our software solution, your hardware or virtual server has to be compatible with CentOS7 distribution. Please confirm that information online with Red Hat Enterprise Linux 7.
- Disk encryption with the strong AES-128 encryption is provided to ensure that all stored network information is protected. Changing or disabling encryption during the installation process requires the presence of an authorized person.

Installation Process

- The installation consists of two parts. The first part runs on Linux in user shell. If you are installing on a network that is not connected to the Internet, use the offline installation.
- The second part takes place via a web interface.

SHELL INSTALLATION

Before installing MENDEL, you will need the following: the USB/DVD with a copy of the MENDEL installation .iso image file and the IP address for the management network interface of the application. The license code will be assigned to you by GREYCORTEX based on your email address.

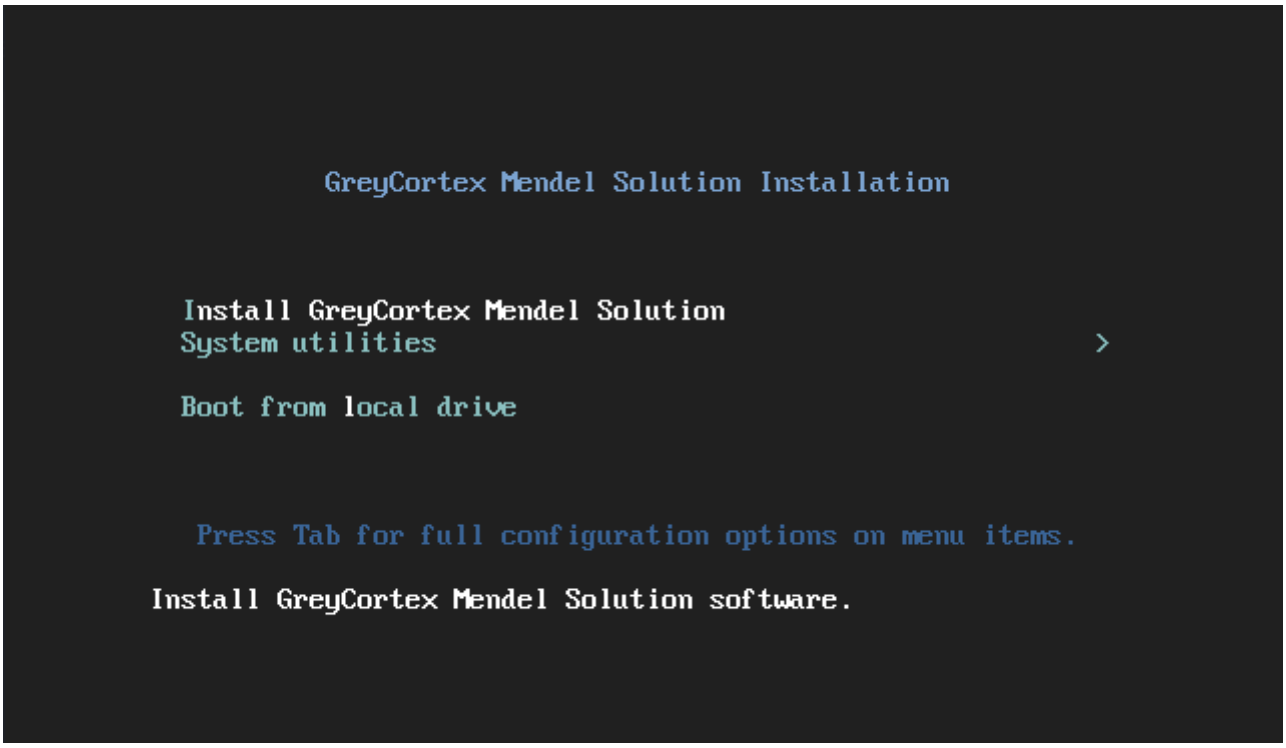
Installation Process

1. The installation consists of two parts. The first part runs via a Linux user shell.
2. The second part of the installation process takes place via a web interface.

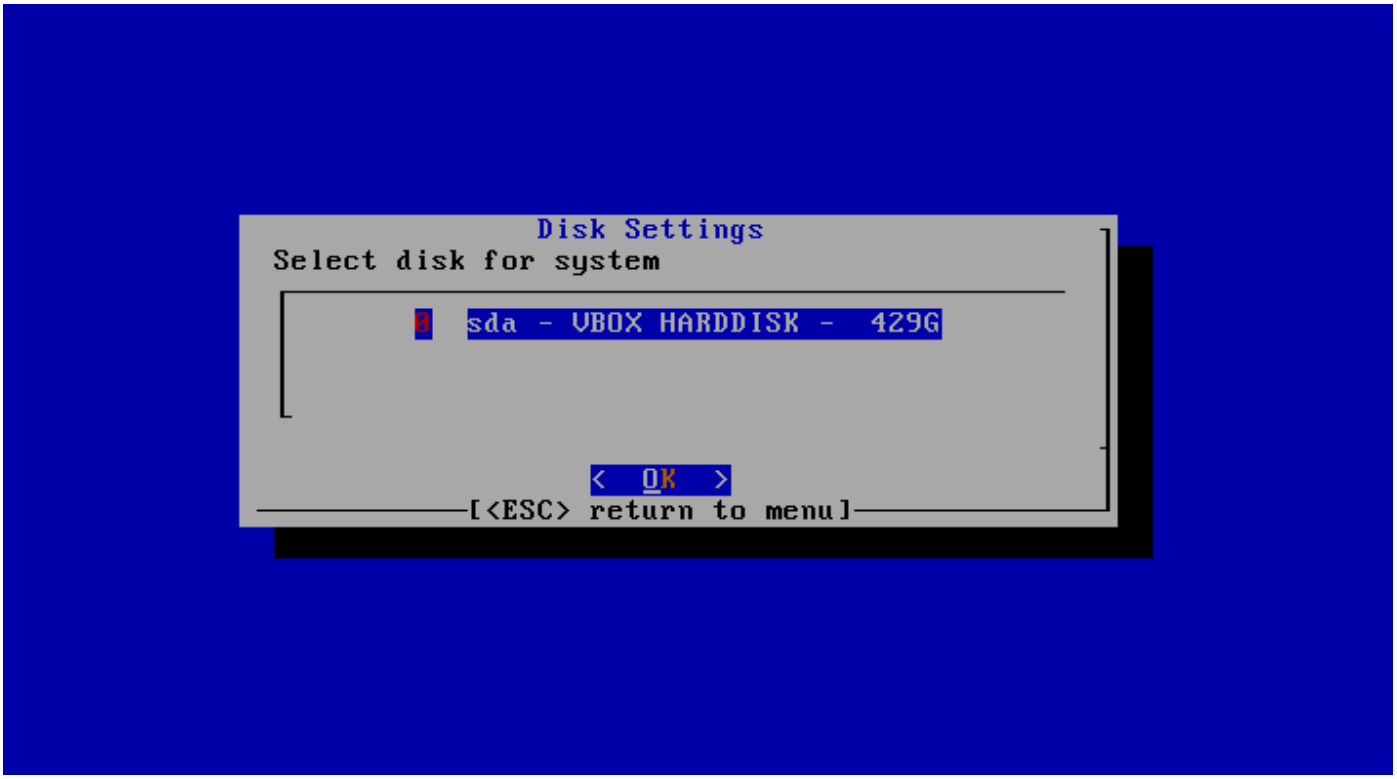
Copyright © 2024 Greycortex

Installation Part I: Installation Within the Linux Shell

1. Insert an SD card or USB/DVD media with a copy of the .iso image file, and boot from it. Select the boot option \"Install GREYCORTEX MENDEL Solution\".

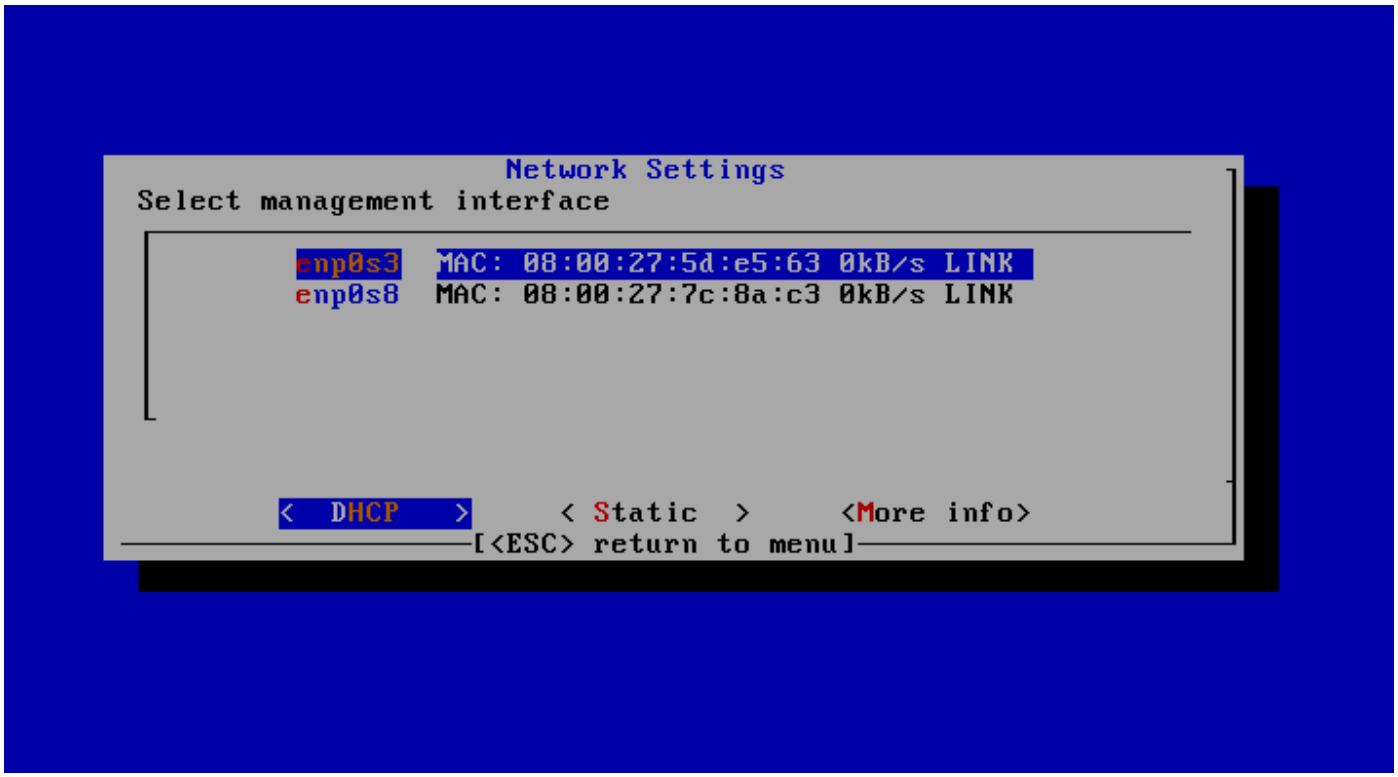


2. Select the hard drive where MENDEL will be installed - this can be any disk larger than 400 GB. MENDEL can be installed on one or two disks. In case of multidisk installation, you will be prompted firstly for the system disk and then for the data disk. Minimum size requirement for both disks is 400 GB.

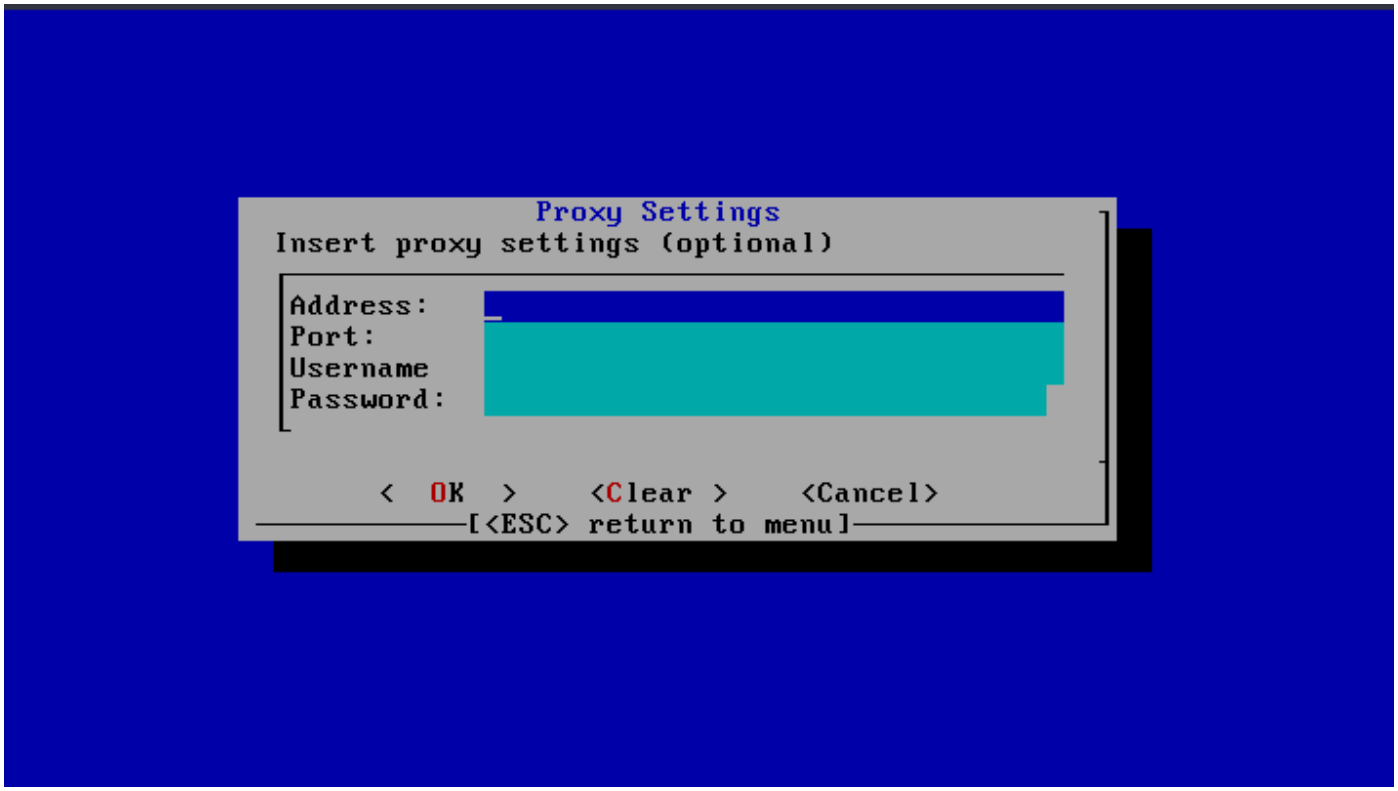


3. After selecting a disk, select the network interface with which to connect to the internet. This will be used during the installation process to download the file with the installation key as well software updates and to access the remote console. Choose between obtaining an IP address from DHCP server or setting a static IP address.

Copyright © 2024 Greycortex

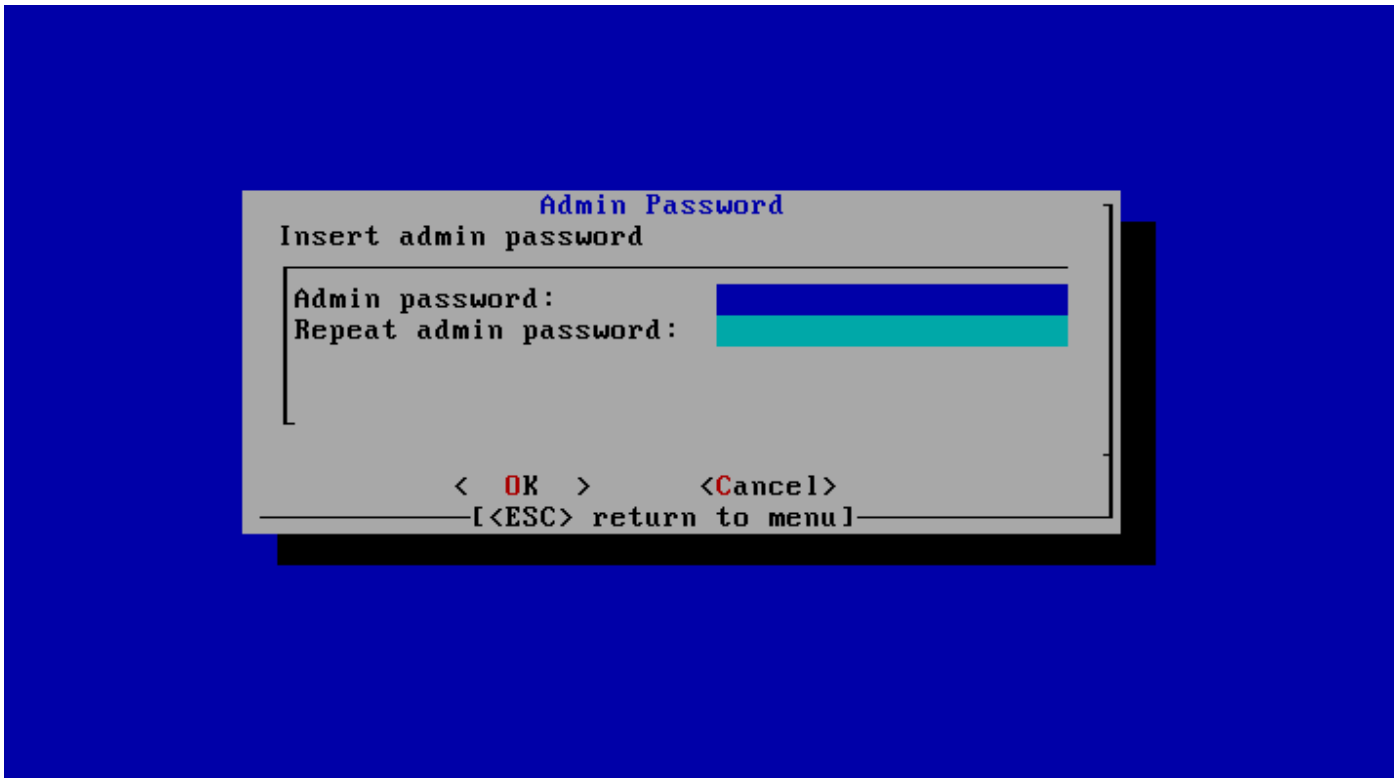


4. Next, enter the proxy server address if your network uses one, or leave it blank and continue by pressing OK.

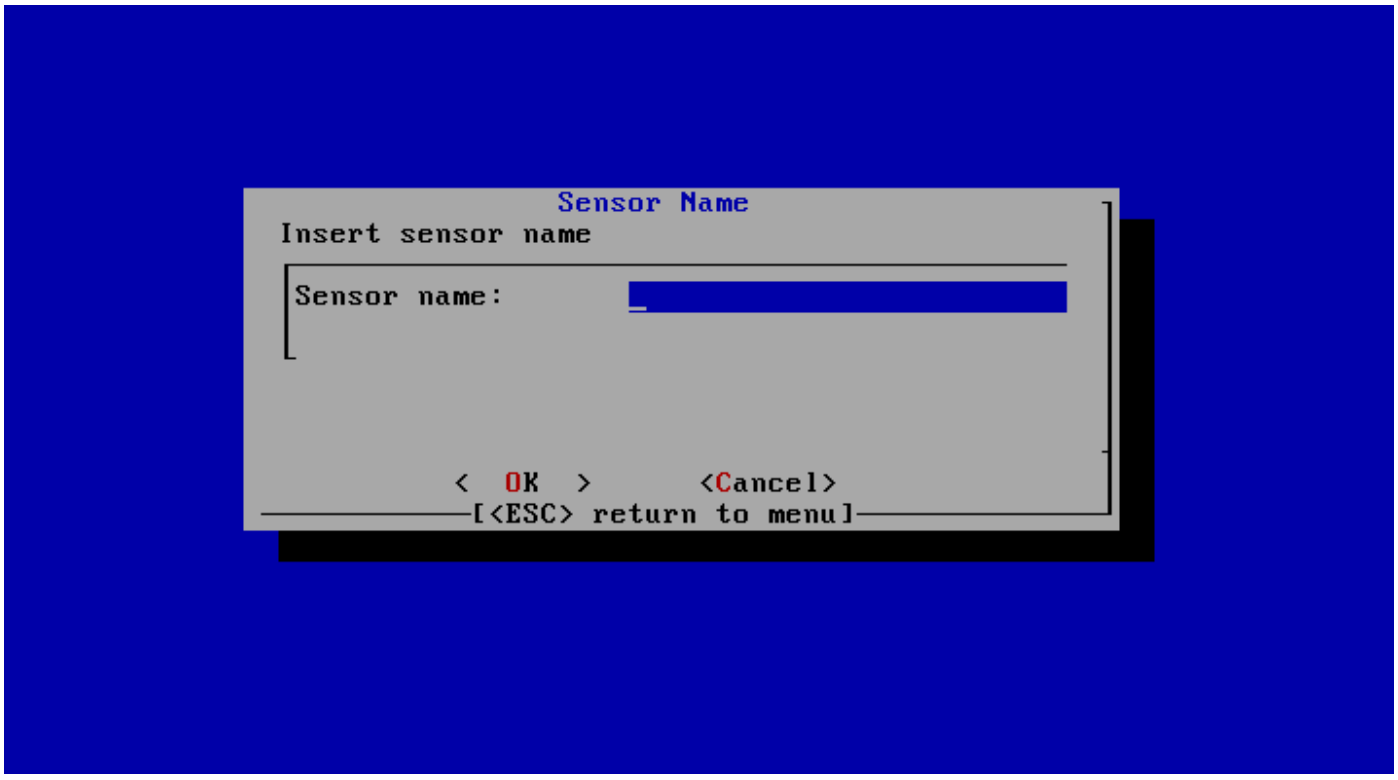


5. Create and enter the administrator password you will use to login to MENDEL.

Copyright © 2024 Greycortex

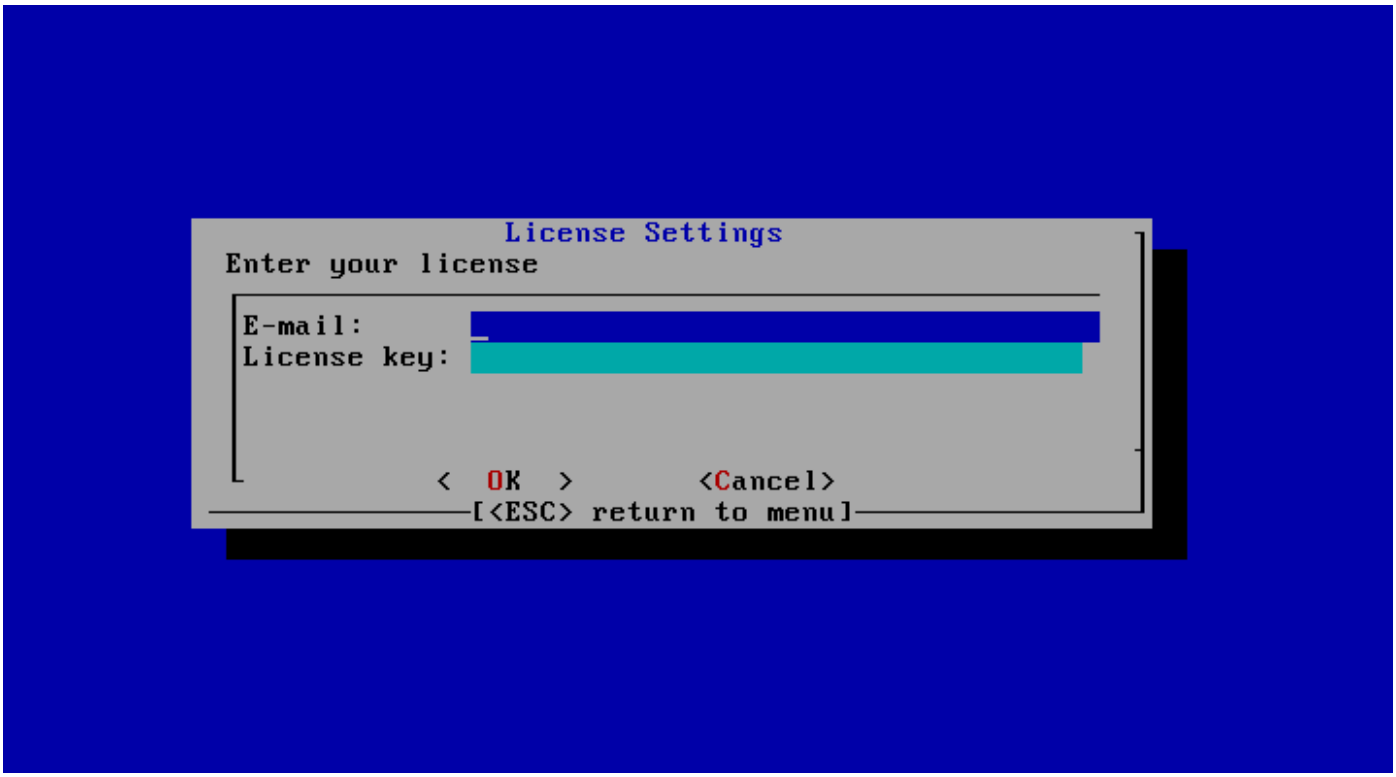


6. Next, enter the hostname of the MENDEL sensor you are installing.



7. Enter email and the activation code that you have received from GREYCORTEX.

Copyright © 2024 Greycortex

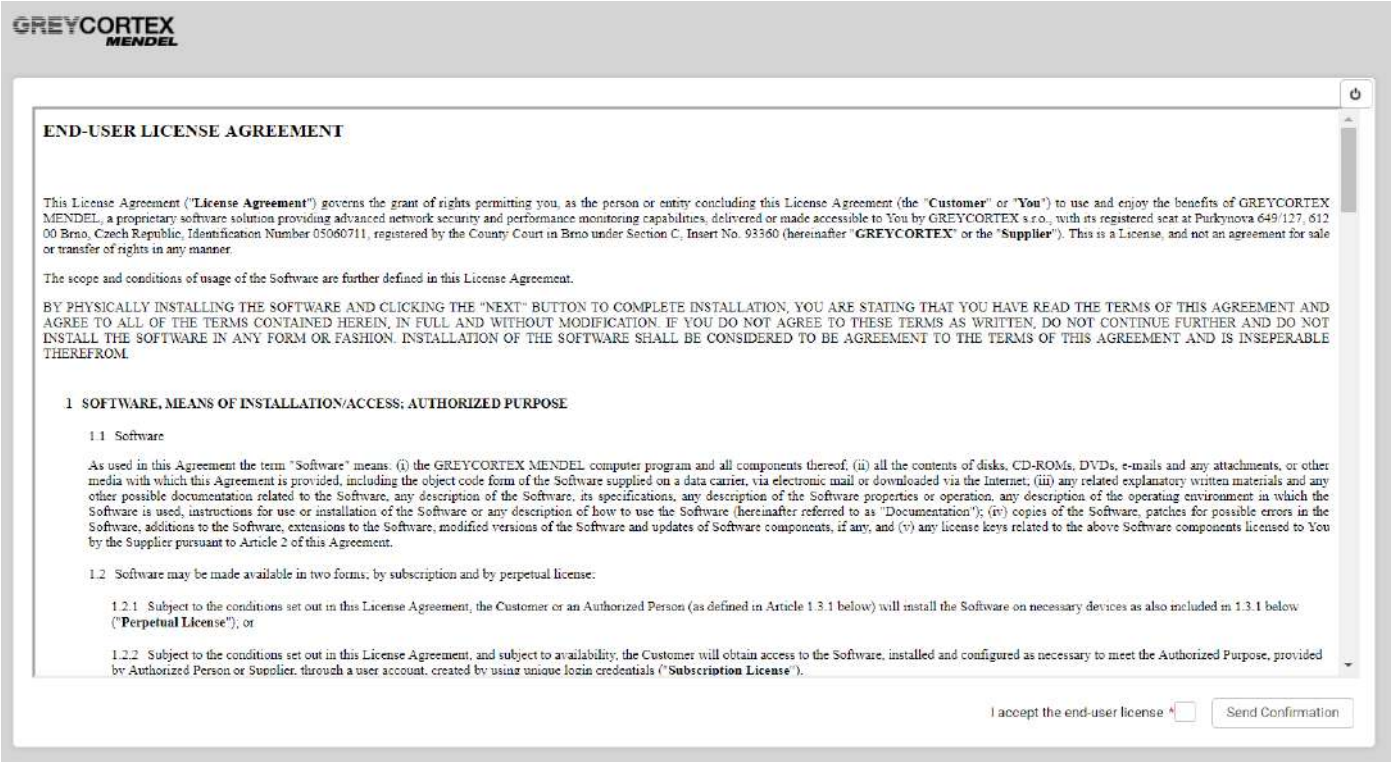


8. Then the summary of the installation is shown. If you want to change it, press Change. If everything is filled correctly, you can start the installation by pressing on Install MENDEL.
9. Once the installation starts, MENDEL automatically detects the hardware configuration of the server and compares it with hardware requirements contained with software. In the case where hardware fails to meet these hardware requirements, a message reflecting this status will be displayed. If the hardware requirements are met, installation can be continued.
10. Wait for the installation to finish. This may take several minutes, and machine will reboot. Then, please continue with the installation part II in your browser

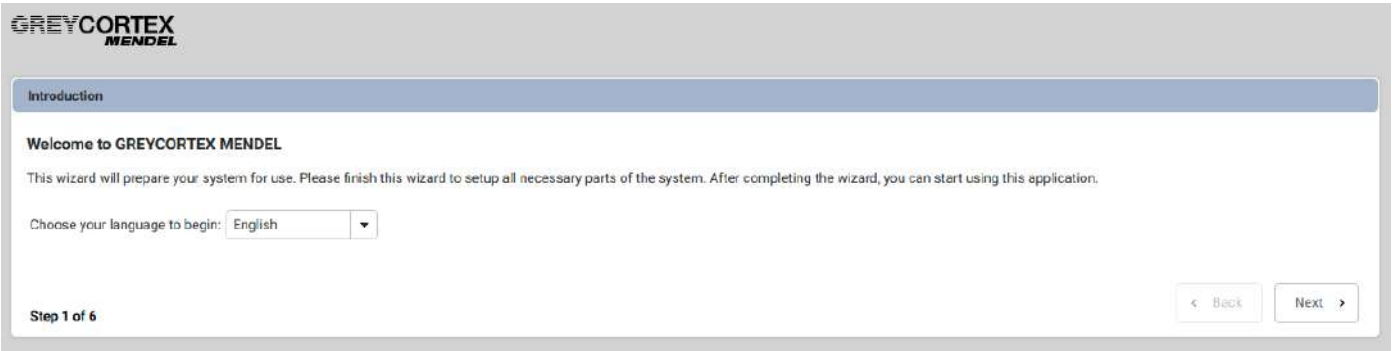
Installation Part II: Installation via Web Interface

1. First, enter the address where you have installed MENDEL into your browser. This will be the IP address that you entered in the first part of the installation, or the address via a DNS record that has been assigned to your IP address.
2. This will now display the MENDEL login screen. Enter the **administrator** username and password you have created.
3. Following the installation wizard screen, you will see the GREYCORTEX MENDEL license agreement. Confirm it by checking the checkbox and click the Next button to agree to the license agreement and proceed with installation.

Copyright © 2024 Greycortex



4. The next screen of this wizard welcomes you and allows you to choose the language.

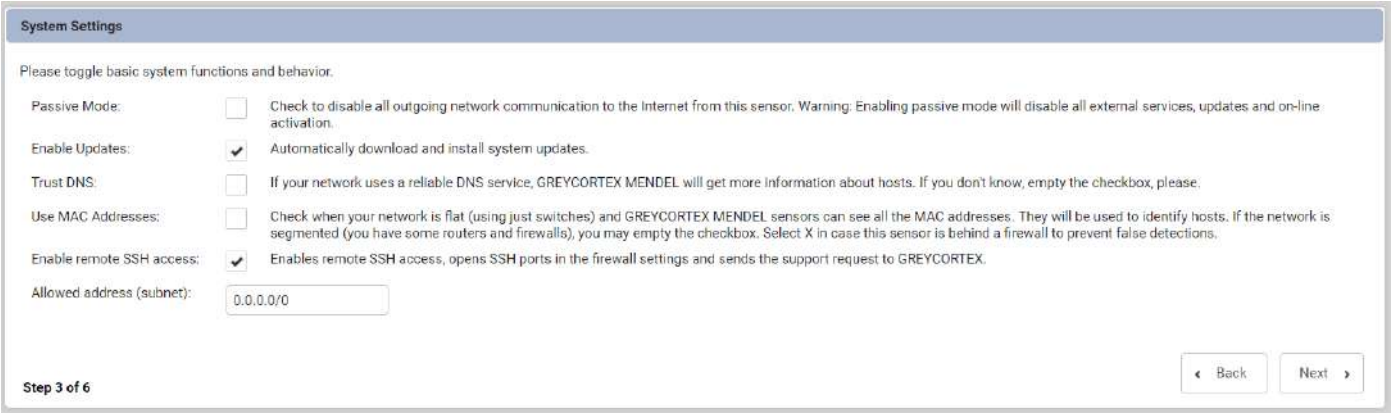


5. The next screen allows you to load settings from your recovery file if you created one from a previous installation.

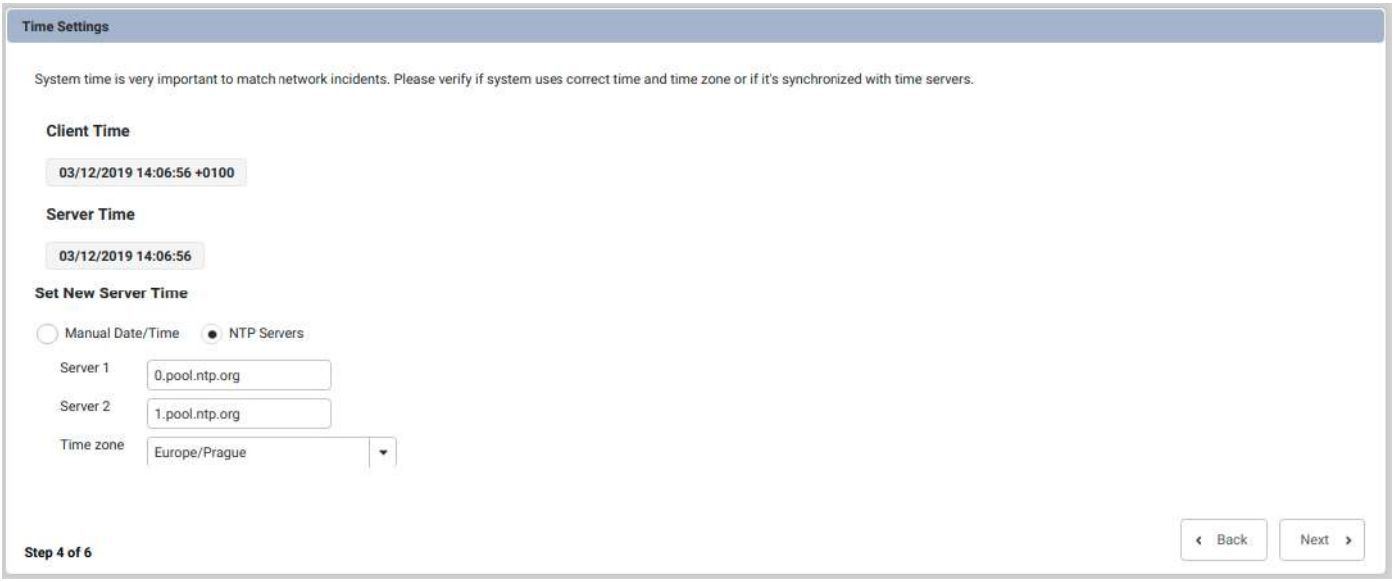


6. System Settings: Enables/disables outgoing communications from the server. When disabled, it works in \'off-line\' mode. Otherwise, this option allows you to automatically download and install operating system updates. In addition, it allows you to search for IP and MAC addresses of computers on the network and using the DNS service and SSH protocol setup. If you check the Enable remote SSH access check box, you will be able to control the application as a user by the administrator from a network other than the application. If you leave 0.0.0.0/0 in the Allowed Address (subnet) box, you can access the application from any network. These options can be changed after the completion of your installation, if necessary.

Copyright © 2024 Greycortex

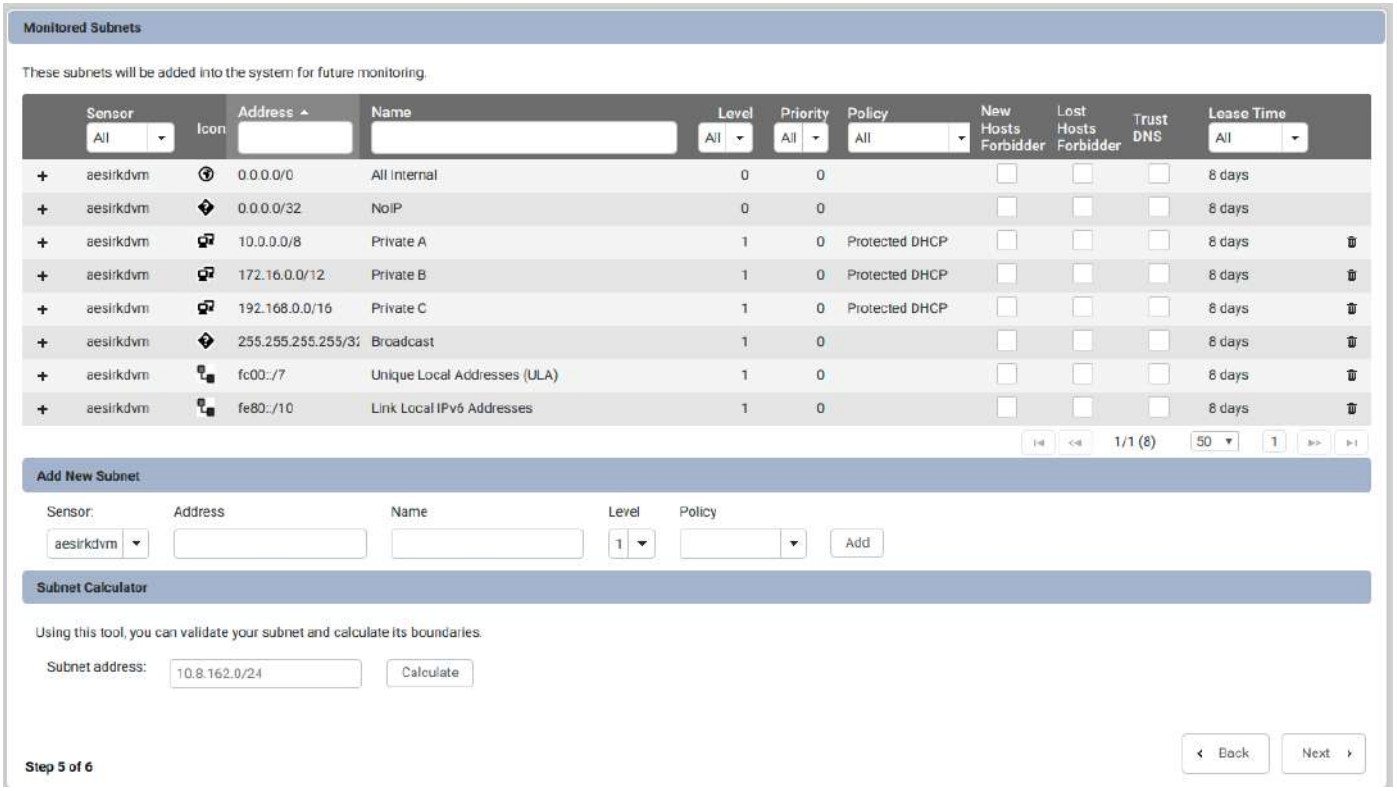


7. Time Settings: On the next screen, check the NTP server. In the boxes Server 1 and Server 2, please enter your local NTP servers or publicly available NTP servers, e.g. tik.cesnet.cz.

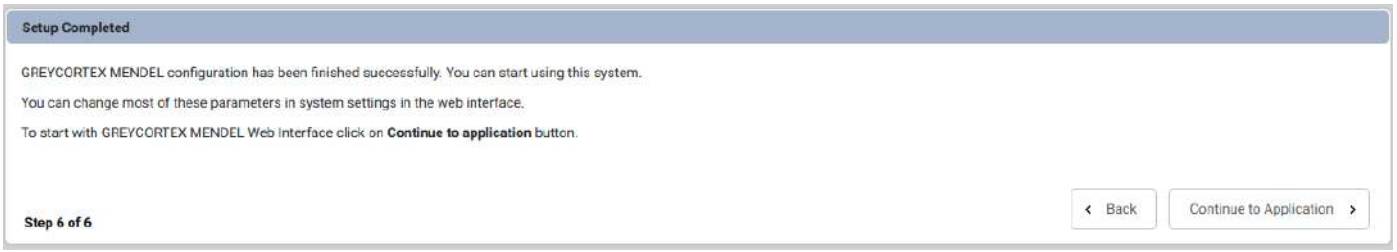


8. Monitored subnets: You can choose what subnets you want to monitor. Only the default subnets are shown, and you can manually edit them, add new subnets or delete the ones you don't need. These can be changed later in the application settings if needed.

Copyright © 2024 Greycortex



9. Completion of the installation process: The next step starts MENDEL's GUI application.



Deployment approaches

MENDEL can be set up for different functionality other than an all-in-one approach depending on the user's infrastructure needs.

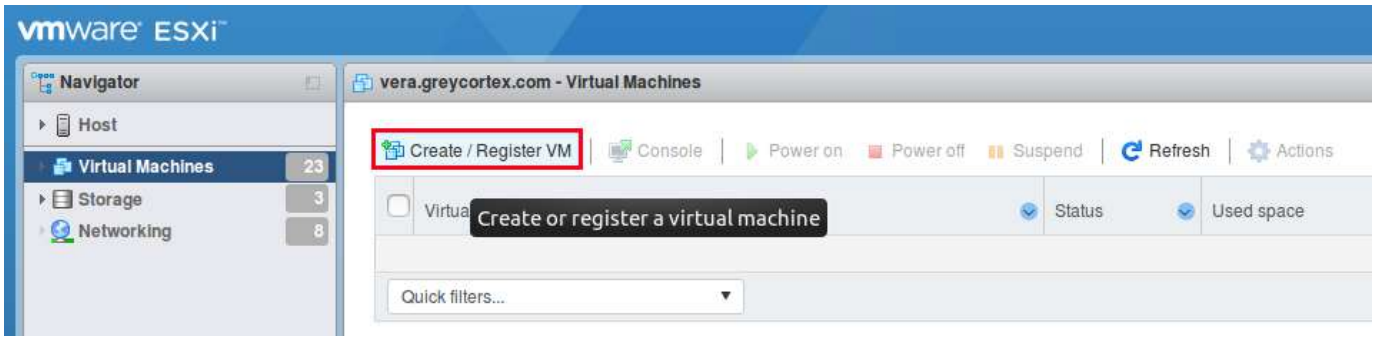
- **Sensor** - In this type of setup, MENDEL monitors the network and sends data to the associated collector. Sent data contains flows and events generated by IDS.
- **Collector** - In this type of deployment all services are running, and everything is being processed on this machine.
- **Sensor-Collector** - In the sensor collector architecture, a sensor or sensors are monitoring the network with IDS and sending data to the collector, where everything else is being processed. One collector is capable of gathering data from up to 30 sensors.
- **Central Event Management** - This type of MENDEL can monitor the network and receives data from other collectors including all events and data about hosts.
- **Multitenant configuration for SOC** - Several Sensors for different customers are connected to one Collector, data are protected so that customers can’t see the data of the others.

Copyright © 2024 Greycortex

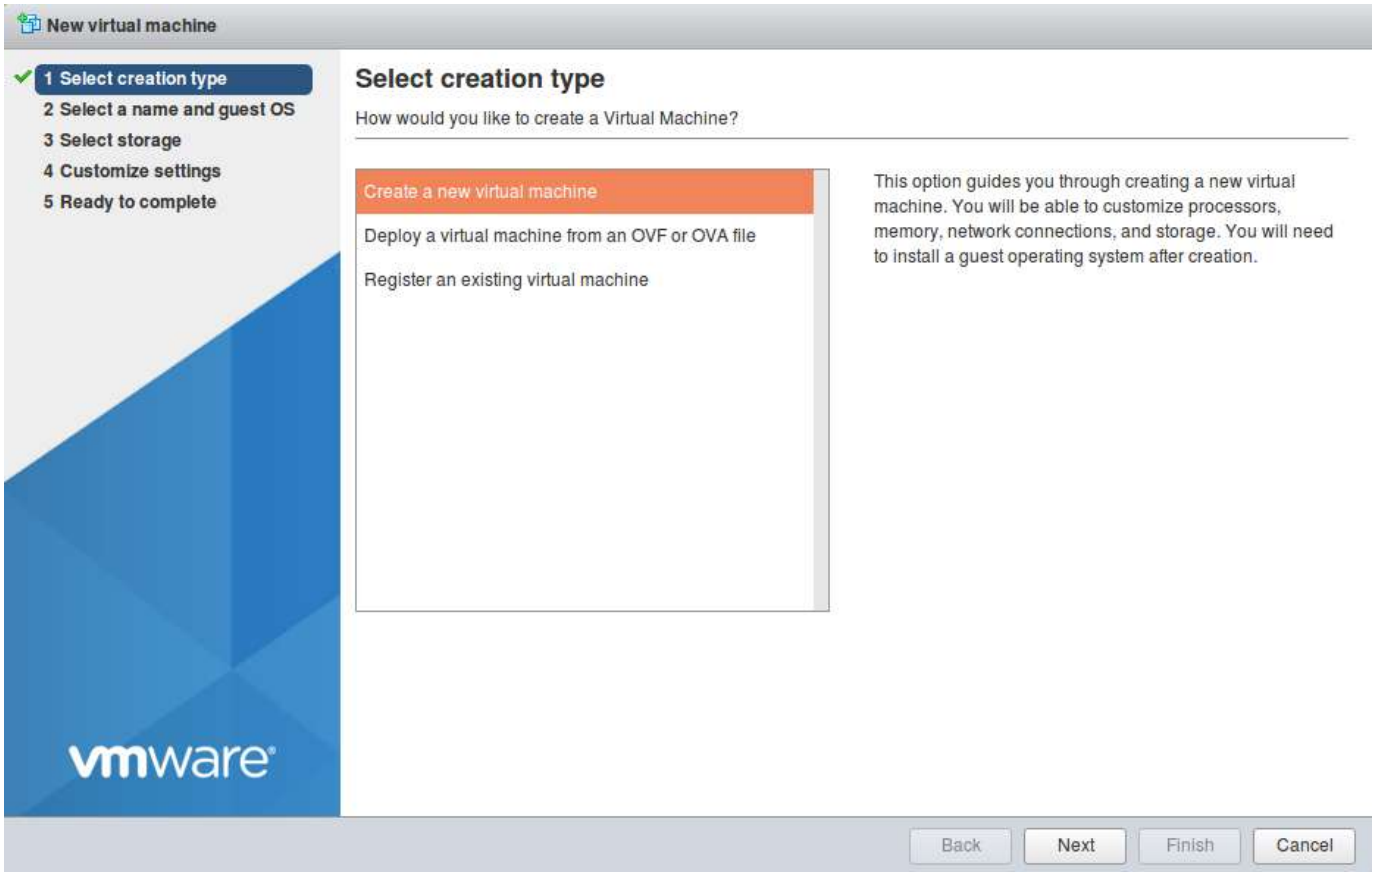
2.3.2 VMware Esxi Setup

VMware ESXi Setup Guide

To create a new virtual machine using a VMware host client, click “Virtual Machines” on the left panel. Then on the first menu tab, click “Create/Register VM .”

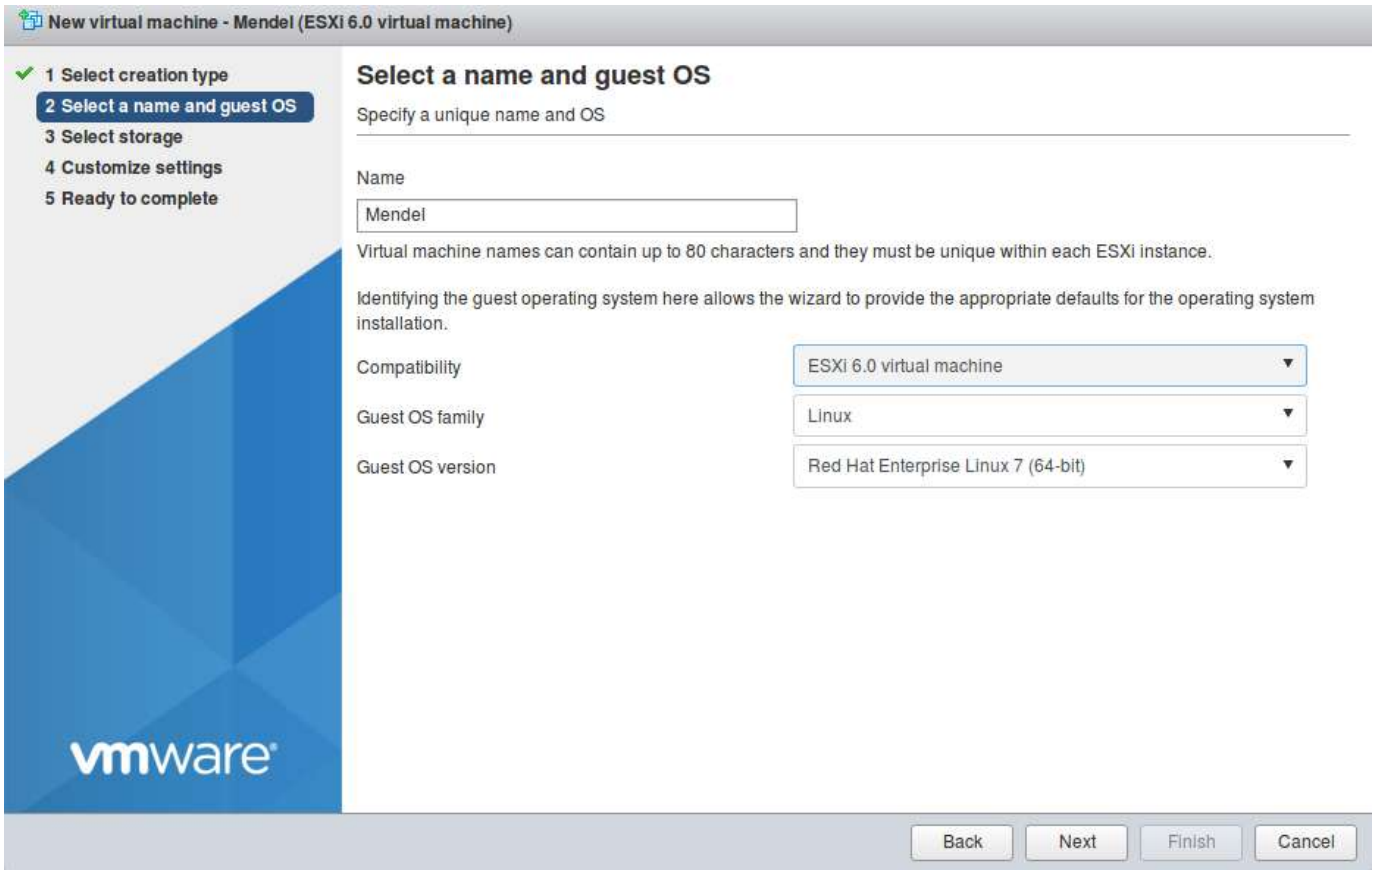


In the pop-up window, choose the first step; “ 1. Select creation type ,” select “Create a new virtual machine ,” then click “ Next .”



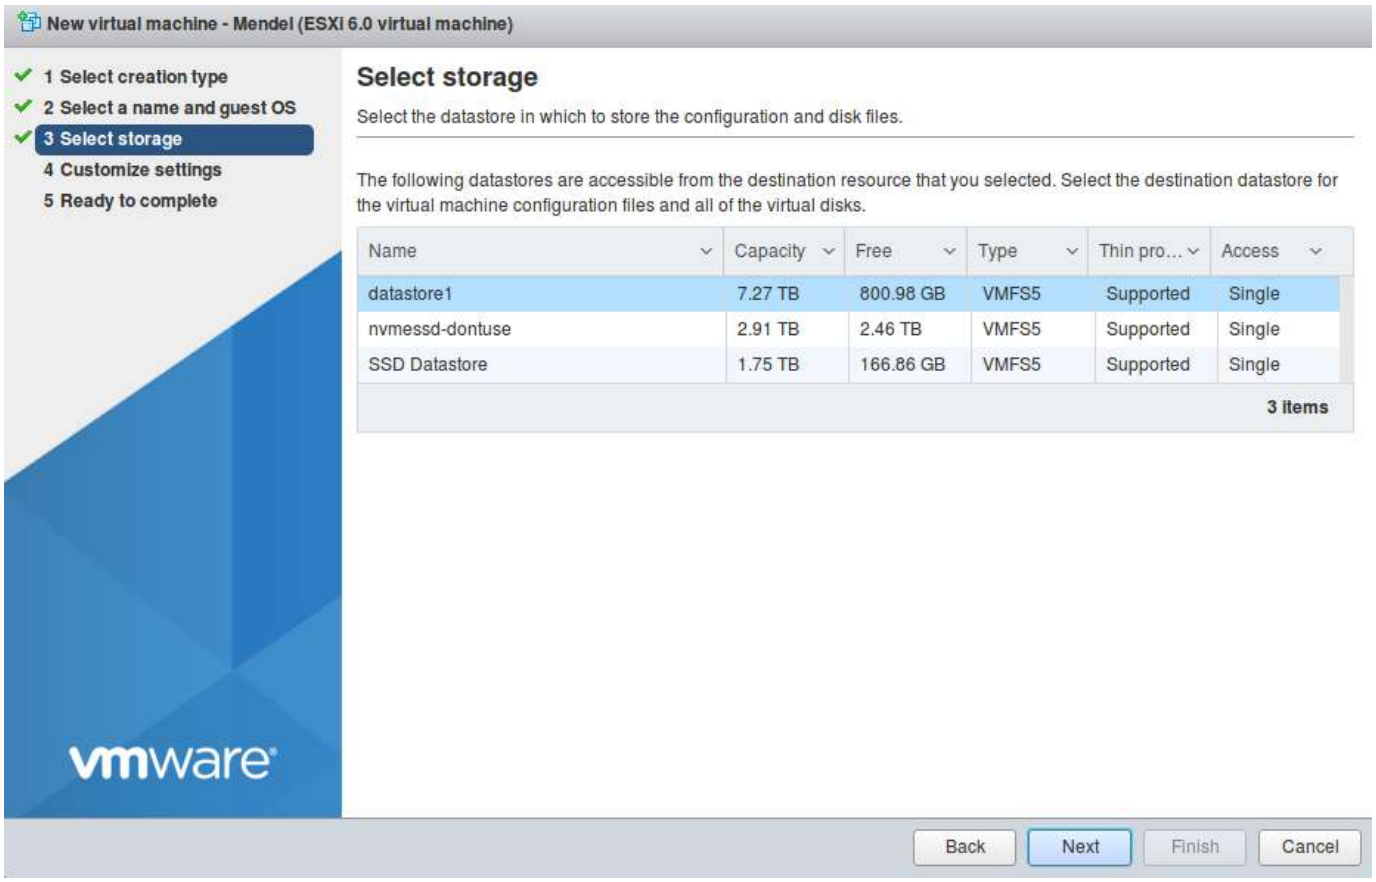
In the second step, type the name of the virtual machine, e.g. ”MENDEL.” Under Compatibility, choose “ESXi 6.0 virtual machine ” and under “Guest OS family ,” choose Linux as the operating system. For the version, select “ Red Hat Enterprise Linux 7 (64-bit)” if available. If it is not available, select an older version (e.g. version 6) and click “ Next .”

Copyright © 2024 Greycortex



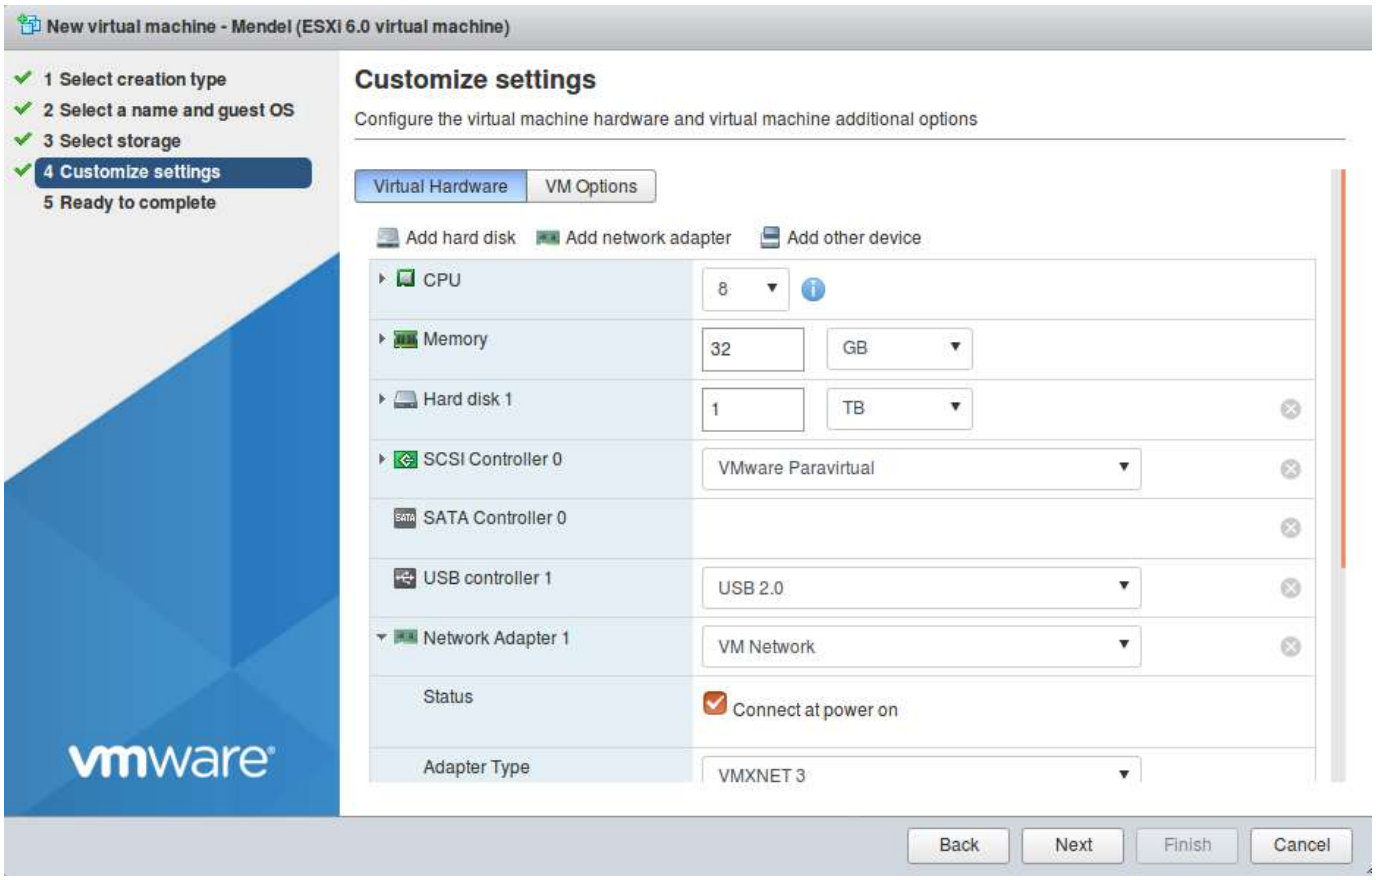
In the third step, choose the destination for data storage, i.e. where MENDEL will be installed, then click “ Next .” Installing MENDEL on fast storage (SAS or SSD disks in RAID configuration) is recommended.

Copyright © 2024 Greycortex



In the fourth step, set the number of CPUs to 8 and enter the correct memory size. 32 GB of memory is recommended. Next, select 1 TB of disk size. This configuration is recommended for 500 Mbit/s traffic. If you have higher volumes of network traffic, we recommend setting these parameters to higher values. For the SCSI Controller, choose “ VMware Paravirtual .” Scroll further in the window to find the remaining input parameters.

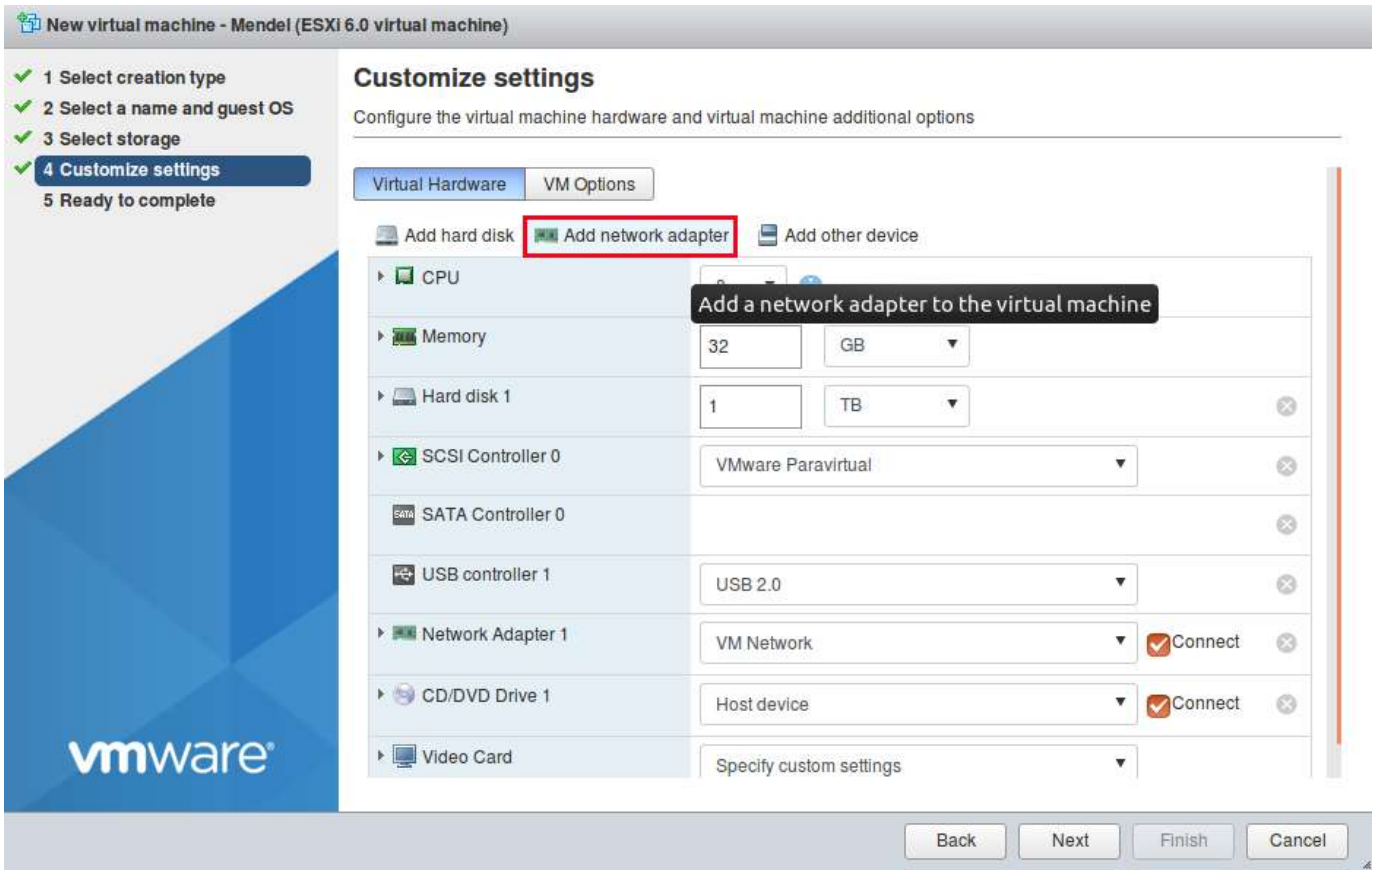
Copyright © 2024 Greycortex

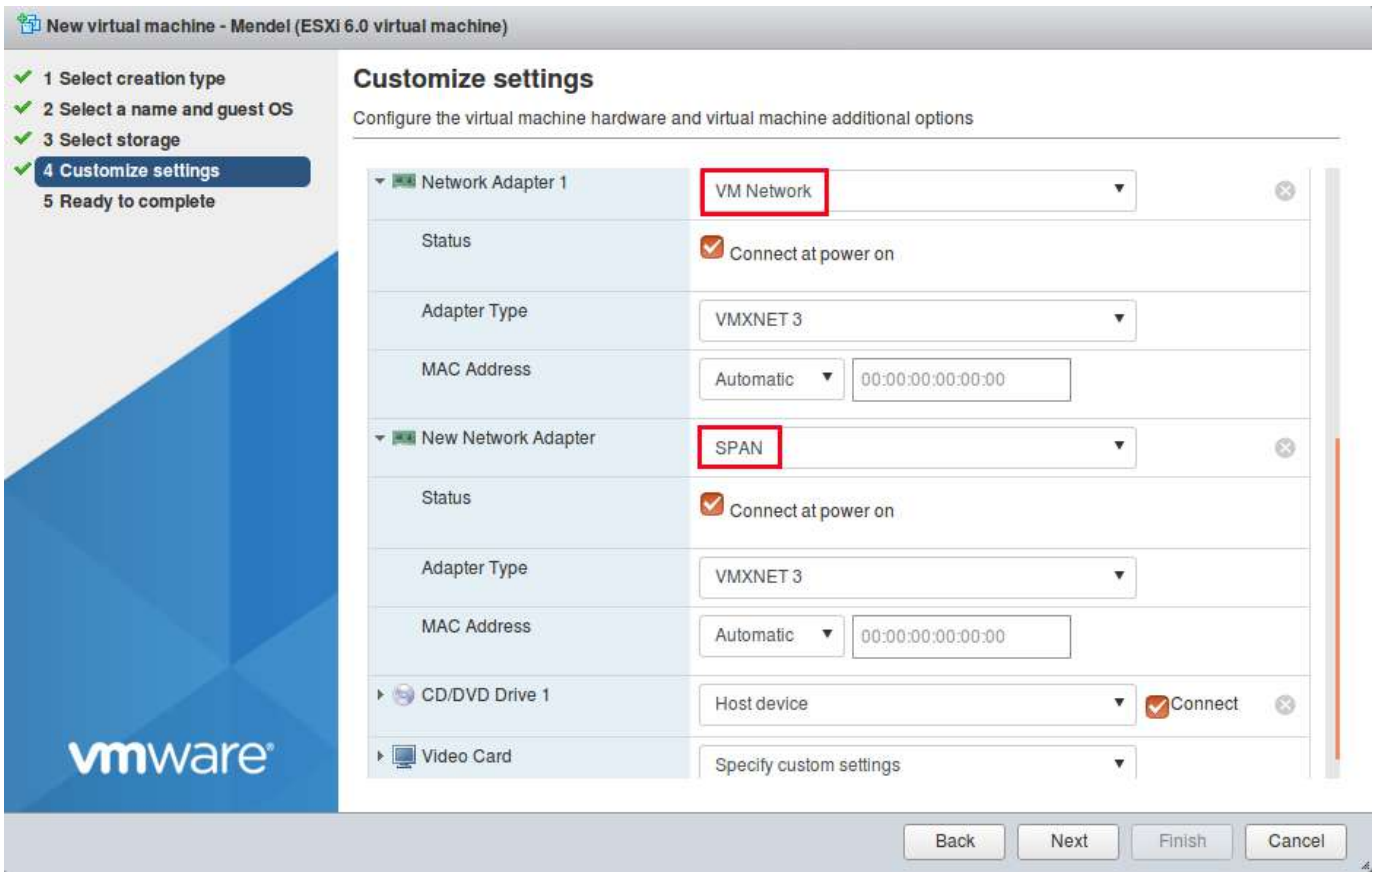


GREYCORTEX MENDEL needs two network adapters to function properly. One is for the management port, and one is for the SPAN port (mirrored traffic). Your physical machine must have at least two network interfaces. To configure this, select “Add network adapter” from the top of the menu in this window, then select “VM Network” for one interface and “SPAN” for the other one interface. Then click “Next.”

Please note that promiscuous mode must be configured either on the interface for the SPAN/Mirror port or on the virtual switch.

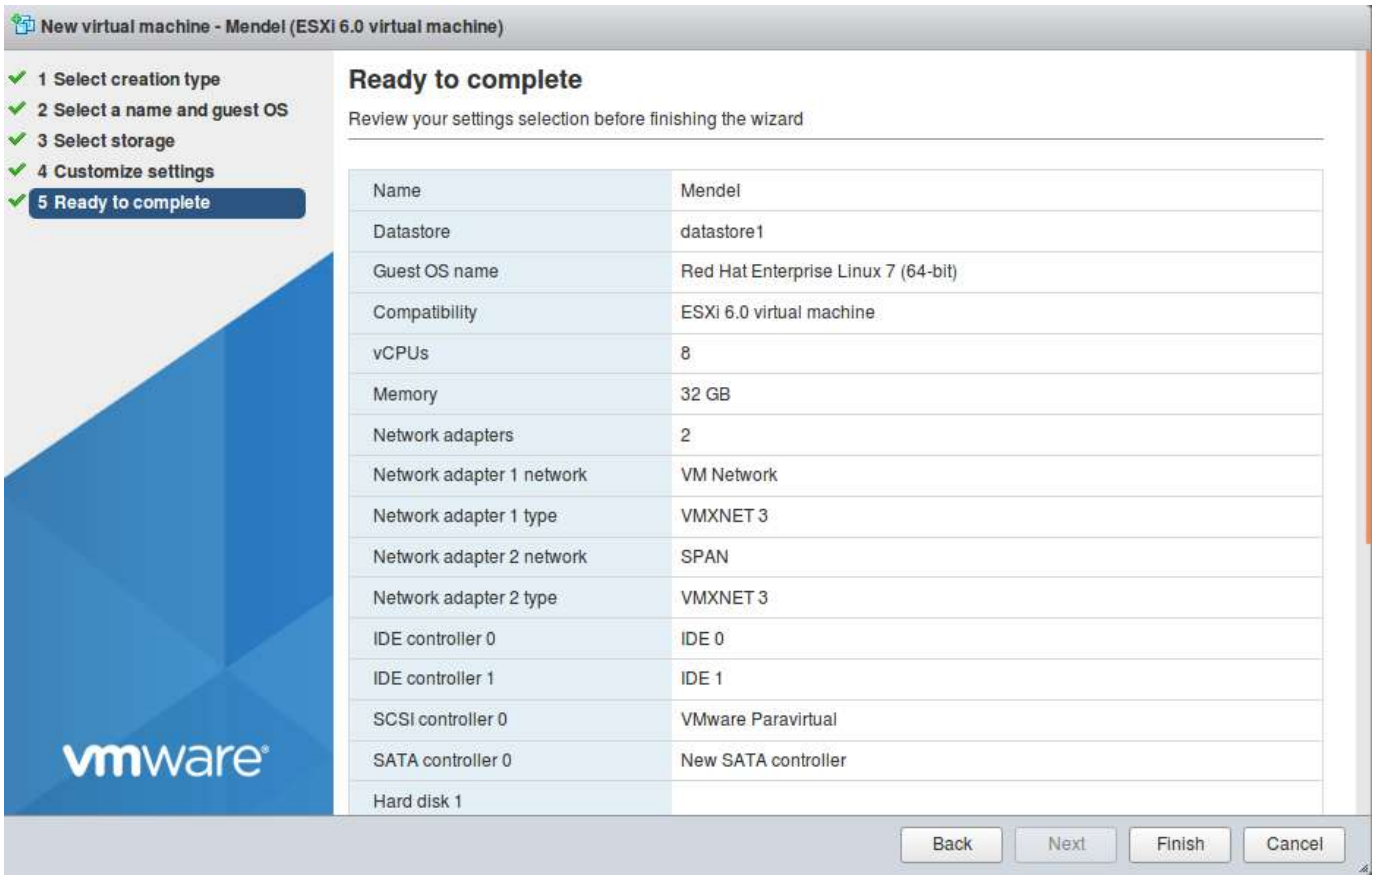
Copyright © 2024 Greycortex



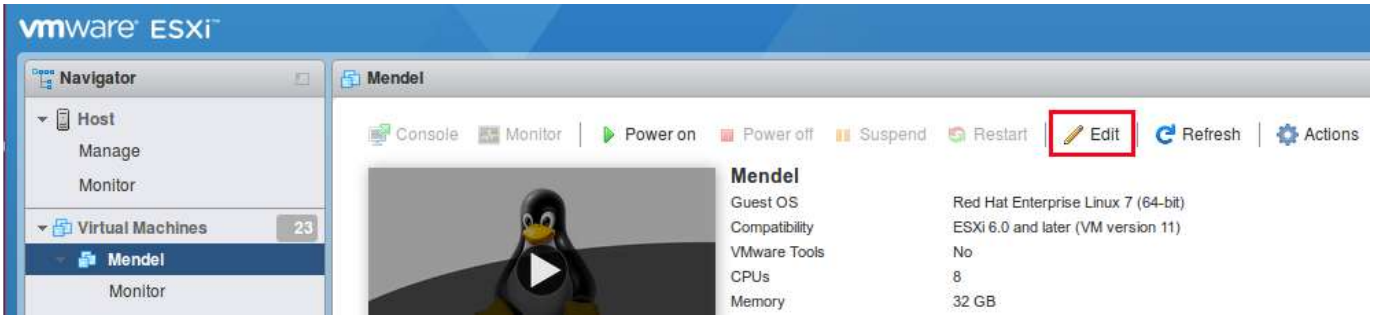


Review the selected settings in Step 5, and then click “Finish” to complete the wizard.

Copyright © 2024 Greycortex

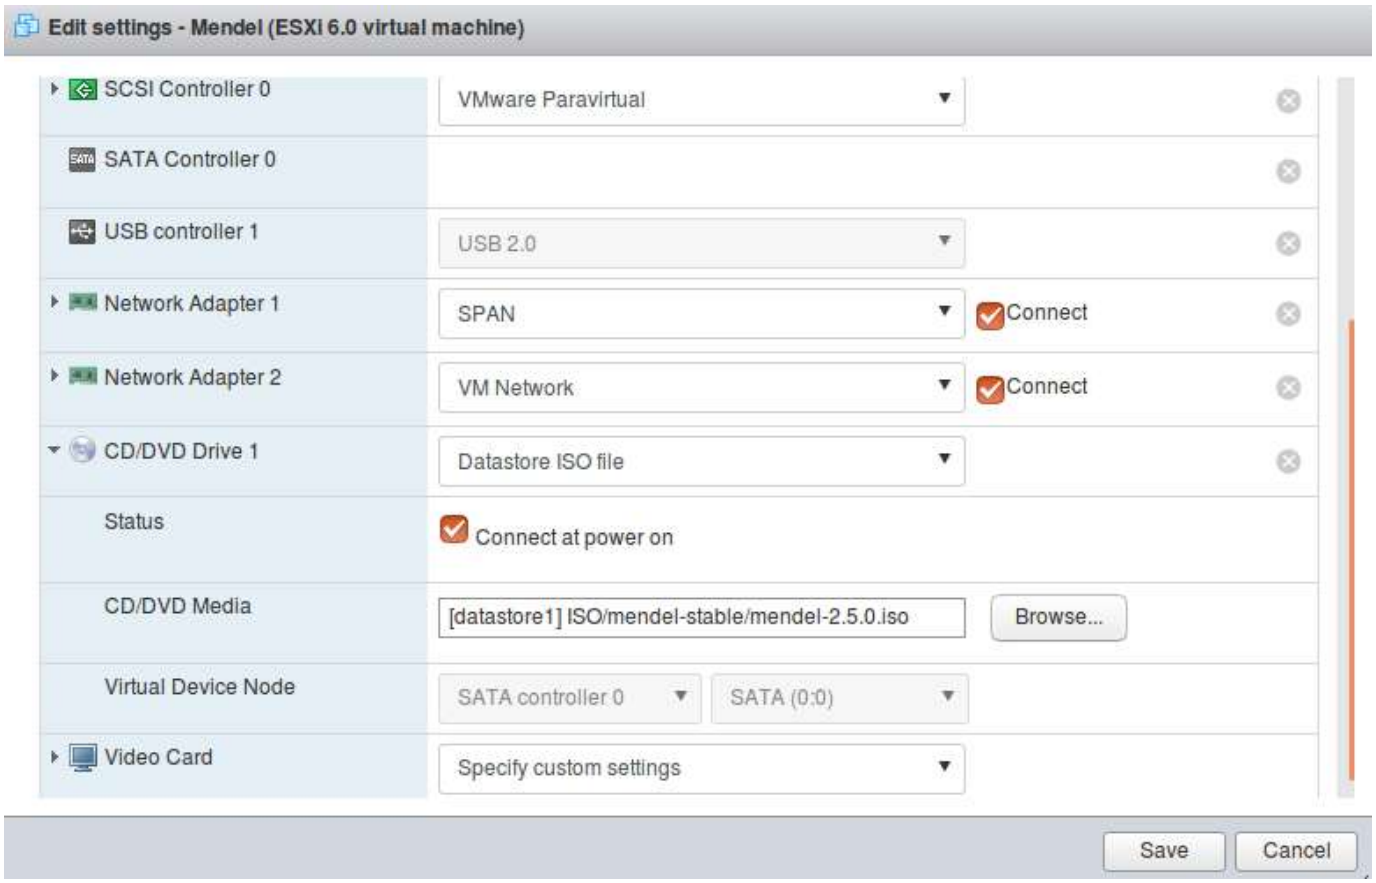


Select MENDEL, the virtual machine you just installed, from the list of installed virtual machines, and click “Edit” in the top menu.



In the pop-up window, click on “CD/DVD Drive 1” and in the first box, choose “Datastore ISO file.” In the “CD/DVD Media” field, click “Browse,” then select the MENDEL installation .iso file from your data storage. Check the “Connect at power on” box. Optionally, you can insert the .iso installation file using the “Host device” option.

Copyright © 2024 Greycortex



Now that the virtual machine setup is complete, power on this virtual machine and begin the process of installing GREYCORTEX MENDEL.

You can find instructions for installing GREYCORTEX MENDEL in the MENDEL installation guide.

Copyright © 2024 Greycortex

2.3.3 Virtual Machine Templates

Cloning MENDEL from VM templates

This documentation provides a step-by-step guide on creating a VMware template for installing the Mendel. The process includes initial setup, installation, and template creation. Thin provisioning is recommended, especially for testing purposes.

Prerequisites

Before starting, ensure you have the following:

- VMware environment set up and accessible.
- Mendel ISO file available on a datastore.
- Administrator access to the VMware environment.

STEP 1: INITIAL VM SETUP

Create a New Virtual Machine:

- Open VMware vSphere Client.
- Right-click on the desired host or cluster and select .
- Follow the wizard steps to create a new VM. Select configuration.

Configure VM Settings:

- Provide a name for the VM.
- Select the appropriate ESXi version.
- Choose .
- Select the appropriate version that matches the Mendel installation requirements.

Network Interfaces:

- Add the same number of network interfaces as required. For example, if using one VM Network and one VM SPAN, add two network interfaces.
- Ensure and options are checked.
- Set to .

Storage:

- Allocate at least 80GB HDD using .
- The allocated storage value is set during instantiation and must be the same or larger than the template.

CD/DVD Drive:

- Choose .
- Select the Mendel ISO from the datastore.

STEP 2: INSTALL GREYCORTEX MENDEL

Power On the VM:

- Start the virtual machine.
- The VM will boot from the Mendel ISO.

Copyright © 2024 Greycortex

Install Mendel:

- Follow the usual installation steps for GreyCortex Mendel.
- Complete the installation process as per the software's documentation.

STEP 3: PREPARE FOR TEMPLATE CREATION

Power Off the VM:

- Once the installation is complete, power off the VM.
- **Remove Installation Media:**
- Unmount or remove the GreyCortex Mendel installation media from the CD/DVD drive.

STEP 4: CREATE THE TEMPLATE

Convert VM to Template:

- Right-click on the VM.
- Navigate to .
- This action will create a template with the same name as the VM.

Rename the Template:

- Locate the newly created template.
- Rename it to a desired, descriptive name for easy identification.

Template Storage:

- The template will be stored on the ESXi datastore.

STEP 5: EXPORT THE OVF TEMPLATE

Export Template:

- Right-click on the template.
- Navigate to .

Exported Files:

- The export will generate several files, typically:
- (virtual disk file)
- (ISO file if included)
- (NVRAM settings)
- (manifest file)
- (Open Virtualization Format descriptor file)

Notes

- **Compatibility:** When importing the OVF template to another ESXi, be aware that there might be compatibility issues if importing from a newer ESXi version to an older one.
- **Storage Allocation:** Ensure the target environment has sufficient storage to accommodate the template size.
- **Network Configuration:** After deploying from the template, review and configure network settings as required for your specific environment.

Copyright © 2024 Greycortex

Deployment of created VM Template

Following steps are recommended for proper template deployment and may slightly vary depending on your specific VM environment.

1. Deploy OVF template
   - Select all 4 files (VMDK, OVF, NVRAM and MF) and import them
2. Select the mapping of networks
   - VM Network for management, SPAN for SPAN port
3. Do NOT set power on automatically
4. Connect ISO file via VM Settings (edit)
   - CD/DVD drive 1 -> Datastore ISO File and select the virtual-template-icense ISO image created above, make sure to check Connect At Power On
5. Power On VM and wait until setup and reboots are done
6. Now deployed Mendel should be an instance with new UUID and new connection strings:
   - IP address (use of DHCP is recommended, static IP Address configuration in template may lead to conflict of IP Adresses)
   - New license from the ISO image or relicensed same license.
7. Log file is available in:

Copyright © 2024 Greycortex

2.4 Post-Deployment

2.4.1 Post-Deployment Guide

This brief Post-Deployment Guide addresses important steps following installation of the GREYCORTEX MENDEL appliance on the network.

NETWORK SETTINGS

It is necessary to confirm that the appliance is connected to the Internet and a mirror port is set up.

MANAGEMENT PORT

To configure the appliance’s network settings, the use of its management port is necessary.

SPAN PORTS

Mirrored traffic which should be analyzed by GREYCOREX MENDEL must flow to the SPAN ports of the appliance.

CORRECT NETWORK SETTINGS

If the Management Port is correctly set up, an IP address assigned to the sensor will be displayed. This IP address is used to access the sensor over SSH or a web browser.

SOLVING ISSUES WITH NETWORK

To solve issues with network interfaces, use these commands:

- network show
- network DHCP
- network IP netmask gateway DNS DNS2
- ping

WIZARD INSTALLATION

Type the IP address or hostname assigned to GREYCORTEX MENDEL into an Internet browser.

Login as Administrator with the password set during installation.

Copyright © 2024 Greycortex

(Preliminary step) - Accept the terms and conditions.

- **Step 1** – **Welcome message**
- **Step 2** – **Restore settings**: Used to restore previous settings or import subnets in .csv format
- **Step 3** – **System settings**: Various functionalities
- **Passive mode**: Check to disable all outgoing network communication to the Internet from this sensor. This will disable all external services, updates, and online activation.
- **Enable updates**: Automatically download and install system updates
- **Trust DNS**: An option to trust DNS records to identify hosts on the network.
- **Use MAC Addresses**: If the monitored network is flat and there are no routers or firewalls; use this setting. MAC addresses will be used to identify hosts.
- **Enable remote SSH access**: Enables remote SSH access, opens SSH ports on the firewall settings, and sends a support request to GREYCORTEX.
- **Allowed addresses(subnet)** - Specify allowed address range for SSH access.
- **Step 4 - Time settings**: Set up the system time. On the screen, NTP servers or a Manual Date/Time setting may be selected. An NTP server that can't be reached when setting system time can't be used. This setting can be adjusted later.
- **Step 5 - Subnet configuration**:
- Press + to add subnet from detected subnets.
- It is possible to add subnets within the **Add New Subnets** section as well.
- **Levels** are set up automatically to reflect the infrastructure.
- **Priority** - adjusts severity value of generated events. Use this setting to reflect importance of critical subnets.
- **Policy setting** - specify policy applied for the subnets.
- **New Host Forbidden** - When a new host appears in subnet, there will be an event generated.
- **Lost Hosts Forbidden** - When a host disappears from network, an event will be generated.
- **Trust DNS**
- **Lease Time** – The length of time for each created model on the monitored subnet. Lease time should be equal to lease time of DHCP, which is 8 days by default.
- **Step 6 - Setup Completed**: Installation is finished and now GREYCORTEX MENDEL application is ready to use.

REMOTE SUPPORT

Click on the button below to send a request for remote access from the GREYCORTEX support team. Access will be activated when the request is confirmed. Your system will be unlocked for SSH access by our support team members who will help you maintain the software. A request should be sent if the customer has purchased a support subscription or requires the L2 support.

TRAFFIC

To check the traffic flow on the monitored network, go to the Status Monitor section. There, we can see the Traffic Overview graph which shows how much traffic is received on the network interface. To specify on which sensor you want to see this traffic, you can use the advanced filter.

Copyright © 2024 Greycortex

This manual describes now to configure a clean Mendel installation for Energy Distribution networks. It is written for an engineering audience, engineers involved in deploying and or operating an Intrusion Detection System, like Mendel. It does not replace any existing documentation.

2.4.2 Deployment Scenario Terminology

All-in-one

This is the standard deployment architecture of IT-MENDEL. This single-node platform system is easy to deploy in a small, non-distributed, environment. One appliance collects and processes all the data from the network and provides an overview for security analysts about potential risks and anomalies in the network. This scenario is not in scope for this document.

Sensor-Collector

This is the deployment architecture where two types of appliances are used - sensors and a collector. This architecture is discussed in this document. Multiple sensors are usually connected to one collector. The sensors can be dedicated IT or OT-sensors. Data received from sensors are processed and displayed on the collector. This solution is most appreciated by organizations with multiple geographical locations. The dedicated OT-sensor deployment scenario is discussed in this document.

There are also Multitenant & Central Event Management deployment and multiple hardware/virtualization options. Please contact your account manager for further instructions.

Configurations & Best practices

After the OT-sensor has been deployed and connected to a collector-node the following configuration & best-practices are advised.

NETWORKS

For each network that is covered by Mendel, decide to what Purdue level it belongs to. In the configuration located here: UI --> Settings --> Subnets, please apply the following for each of the networks, or better, create three Policies (UI --> Settings --> Policies) with the following settings. Recommended settings for each Purdue Level:

Copyright © 2024 Greycortex

• Purdue Level 1 (PLC, RTU, Control)

• Edit Name

• Edit Description

• Policy = Protected Static

• Priority = 0

• Trust DNS = false

• Lease time = 30 days

• Tag =

• New hosts Forbidden = true

• Lost Hosts Forbidden = true

• Allowed/Denied Ports:

• Outgoing traffic to external network: Allow: None + Deny: All

• Purdue Level 2 (EWS, HMI, Protocol Clients)

• Edit Name

• Edit Description

• Policy = Protected Static

• Priority = 1

• Trust DNS = false

• Lease time = 30 days

• Tag =

• New hosts Forbidden = true

• Lost Hosts Forbidden = true

• Allowed/Denied Ports:

• Outgoing traffic to external network: Allow: None + Deny: All

• Purdue Level 3 (SCADA, Historians)

• Edit Name

• Edit Description

• Policy = Protected Static

• Priority = 2

• Trust DNS = false

• Lease time = 30 days

• Tag =

• New hosts Forbidden = true

• Lost Hosts Forbidden = true

• Allowed/Denied Ports:

• Outgoing traffic to external network: Allow: None + Deny: All

Copyright © 2024 Greycortex

Hosts

For each host that is covered by Mendel, decide to what Purdue level it belongs to. In the configuration located here: UI --> Settings --> Hosts, please apply the following for each of the critical hosts. Mendel has it's own internal tagging system (UnTE) that will apply tags automatically over time the system is operational.

- Name
- Description
- Hostname
- Policy (See previous chapter)
- Add relevant tags: You can create your own tree structure, but for convenience the PURDUE-model is pre-loaded in the sensor.

Detection

In an Energy Distribution networks the commonly most used protocols are:

• MMS • GOOSE • SV • IEC104 • DNP3 • DLMS/COSEM • SNMP

!! NOTE: For all these protocols Mendel provides full protocol dissectors (parsers). This allows the user to define it's own IDS-detection signatures.

IDS

For each ICS-protocol that is used in the Purdue Levels 1, 2 & 3, check that it is enabled in UI --> Settings --> Services. You can search the name of the protocol by typing the name in the "service" search box. If the protocol is running on a non standard port (e.g. Modbus on 5020 instead of the standardized 502 port) please add the ports used. Ensure that the profile is set to "PARSE"

Go to UI => Settings => User Integration => Sensor Groups: Add group "OT-sensors"

Go to UI --> Settings --> System --> Sensors & Collectors --> Other

Press configure and change the parameters as stated in the list below. - IT-profile => IT-Custom - OT-profile => OT-expert - Optionally give the sensor a description, Latitude & Longitude.

Advanced users can tweak the Database/Detection & Netflow performance profiles. Please ask your support for details.

!! NOTE !! For creating custom detection rules a short tutorial is listed in Appendix D.

Energy Distribution specific Detection signatures

- MMS: See Appendix A
- IEC104: See Appendix B
- DLMS/COSEM: See Appendix C
- DNP3: Shipped with product (verify if enabled)

NBA

Disable the following NBA-rules for all sensors connected to

ID	Description
-2600	Periodic: Repetitive Connections (>> 1 every minute in an hour)
-2605	Periodic: Repetitive Connections (same size >> 1 in 5 minutes)
-2700	Periodic: Repetitive Connections (>> 1 every 2 minutes in 4 hours)
-2800	Periodic: Repetitive Connections (every 30 minutes in 6 hours)

Copyright © 2024 Greycortex

UNTE

No customizations required. There is a known limitation in tagging several OT-client devices.

User Accounts

Set up the user-accounts and groups that are specified by the customer.

Administrators	Administrators
Default	Default User Group
Report manager	Reporting Privileges
Support	Elevated access
System manager	System Manager
User manager	User Manager

API

For adding API_KEYS to user-accounts please follow the documentation available in the help-section of your collector deployment.

Sensor Profile

(See IDS-section)

False Positives (Via API)

After configuring the IT-custom profile in the sensor profile, no further initial false-positives need to be configured.

Assumptions

COMMON PROTOCOL IN ENERGY DISTRIBUTION NETWORKS

- MMS
- GOOSE
- SV
- DNP3
- OPCUA
- DLMS/COSEM

If other protocols are present, please consult the product datasheet for support.

Strategic considerations

Based on the organization structure, staffing and knowledge level available at the partner and the end-customer it can be a strategy to use a Mendel OT-sensor loaded with Threat Intelligence made available by GCX or to run a system without any detection intelligence.

- Mendel with minimum Threat Intelligence
- Fully loaded with GCX-Signatures
- User profile
- Sensor profile

Copyright © 2024 Greycortex

Appendix A (Energy Distribution Signatures – MMS/GOOSE/SV)

Copyright © 2024 Greycortex

```lua alert tcp any any -> any 102 (msg:"MMS takeControl";output.request; luascript:function init (args) local needs = {} needs["buffer"] = tostring(true) return needs end function match(args) buf = args["buffer"] if string.find(buf, "request") and string.find(buf, "\"takeControl\"") then return 1 end end return 1; priority:1; classtype:information; sid:1000010007; rev:1;)
```

```lua alert tcp any any -> any 102 (msg:"MMS deleteEventAction";output.request; luascript:function init (args) local needs = {} needs["buffer"] = tostring(true) return needs end function match(args) buf = args["buffer"] if string.find(buf, "request") and string.find(buf, "\"deleteEventAction\"") then return 1 end end return 1; priority:1; classtype:information; sid:1000010011; rev:1;)
```

Appendix B (Energy Distribution Signatures - IEC104)

Copyright © 2024 Greycortex

Appendix C (Energy Distribution Signatures - DLMS/COSEM)

Copyright © 2024 Greycortex

Copyright © 2024 Greycortex

Appendix D: Creating LUA Rules

For protocols that are currently supported for full parsing by GreyCortex Mendel, version 3.7 of the Mendel software supports a LUA-scripting language that is enabling the operator to extract JSON-output from the flows. With this information the operator can compose his own anomaly-detectors, leading to events. This appendix describes different use-cases.

Integer example rule

Observe the suricata detection signature for the IEC 60870-5 -104 protocol below.

Understanding the LUA-part of the signature requires protocol knowledge

String example rule

Copyright © 2024 Greycortex

MMS examples

Reset of ICD detected.

Unauthorized write to ICDevice

This example uses the new host variable feature and will deny each write command from devices that are not labeled as Engineering Workstation.

Blacklisted MMS Method

for example method request.MMSpdu.confirmed-RequestPDU.confirmedServiceRequest getNameList is blacklsted

MMS Write method detected

Copyright © 2024 Greycortex

3rd party rules (examples)

Copyright © 2024 Greycortex

2.4.3 Post-Deployment OT Industry

This brief Post-Deployment Guide addresses important steps following installation of the GREYCORTEX MENDEL appliance on the network.

NETWORK SETTINGS

It is necessary to confirm that the appliance is connected to the Internet and a mirror port is set up.

MANAGEMENT PORT

To configure the appliance’s network settings, the use of its management port is necessary.

SPAN PORTS

Mirrored traffic which should be analyzed by GREYCOREX MENDEL must flow to the SPAN ports of the appliance.

CORRECT NETWORK SETTINGS

If the Management Port is correctly set up, an IP address assigned to the sensor will be displayed. This IP address is used to access the sensor over SSH or a web browser.

SOLVING ISSUES WITH NETWORK

To solve issues with network interfaces, use these commands:

- network show
- network DHCP
- network IP netmask gateway DNS DNS2
- ping

WIZARD INSTALLATION

Type the IP address or hostname assigned to GREYCORTEX MENDEL into an Internet browser.

Login as Administrator with the password set during installation.

Copyright © 2024 Greycortex

(Preliminary step) - Accept the terms and conditions.

- **Step 1** – **Welcome message**
- **Step 2** – **Restore settings**: Used to restore previous settings or import subnets in .csv format
- **Step 3** – **System settings**: Various functionalities
- **Passive mode**: Check to disable all outgoing network communication to the Internet from this sensor. This will disable all external services, updates, and online activation.
- **Enable updates**: Automatically download and install system updates
- **Trust DNS**: An option to trust DNS records to identify hosts on the network.
- **Use MAC Addresses**: If the monitored network is flat and there are no routers or firewalls; use this setting. MAC addresses will be used to identify hosts.
- **Enable remote SSH access**: Enables remote SSH access, opens SSH ports on the firewall settings, and sends a support request to GREYCORTEX.
- **Allowed addresses(subnet)** - Specify allowed address range for SSH access.
- **Step 4 - Time settings**: Set up the system time. On the screen, NTP servers or a Manual Date/Time setting may be selected. An NTP server that can't be reached when setting system time can't be used. This setting can be adjusted later.
- **Step 5 - Subnet configuration**:
- Press + to add subnet from detected subnets.
- It is possible to add subnets within the **Add New Subnets** section as well.
- **Levels** are set up automatically to reflect the infrastructure.
- **Priority** - adjusts severity value of generated events. Use this setting to reflect importance of critical subnets.
- **Policy setting** - specify policy applied for the subnets.
- **New Host Forbidden** - When a new host appears in subnet, there will be an event generated.
- **Lost Hosts Forbidden** - When a host disappears from network, an event will be generated.
- **Trust DNS**
- **Lease Time** – The length of time for each created model on the monitored subnet. Lease time should be equal to lease time of DHCP, which is 8 days by default.
- **Step 6 - Setup Completed**: Installation is finished and now GREYCORTEX MENDEL application is ready to use.

REMOTE SUPPORT

Click on the button below to send a request for remote access from the GREYCORTEX support team. Access will be activated when the request is confirmed. Your system will be unlocked for SSH access by our support team members who will help you maintain the software. A request should be sent if the customer has purchased a support subscription or requires the L2 support.

TRAFFIC

To check the traffic flow on the monitored network, go to the Status Monitor section. There, we can see the Traffic Overview graph which shows how much traffic is received on the network interface. To specify on which sensor you want to see this traffic, you can use the advanced filter.

Copyright © 2024 Greycortex

2.4.4 Post Deployment Checklist

Before going through this checklist, you should already have installed MENDEL on a physical or virtual machine.

Please open the MENDEL post-install setup wizard in web browser. Let’s go through the steps!

- This is a welcome message. You can click next.
- You will now be offered the option to upload previously exported settings from another MENDEL installation, for example Policies, Subnets, Global Settings, False Positives, or Users. You can upload a whole .tar file.
- Next, you will be asked to choose some basic settings. We suggest allowing SSH.
- You will then be able to choose your time settings, and whether it will be a manual date/time combination, or from an NTP Server.
- The fifth step is really important, and helps us A LOT with our security consultations. In this step, you will be defining subnets. To be more specific, you will choose the sensor, which icon that subnet should have, the IP range, and the name. The options that follow are related to Events. It is necessary to define which subnets are in your network in order to avoid false positive events.
- If you see this screen, setup was a success!

However, we’re not finished with deployment. We still need to configure the Email Server, Incidents, False Positives, and more.

EMAIL SERVER

The Email Server will allow you to set up notifications which will be sent from MENDEL to desired email addresses. When someone reports an incident (described below), an email is generated. Those emails are then sent to those who are monitoring the incident. If a valid filter is set up, the user can also be notified about events in real time. You can configure the Email Server in settings under the System tab, selecting the Email option. Click on it. You should see a new window called Email System Settings. Global settings are mandatory.

Set up:

- **Server name** - address of mail server
- **Port** - port under which mails are sent, usually 465 (SMTPS)
- **Security** - preferably SSL/TLS
- **Authentication** - password or no password
- **User Name** - used to log into mail server (it has to be a valid user on that mail server)
- **Password**
- **From Address**

INCIDENTS

Incident can be a certain event or group of multiple events in a certain time. It is defined by summary, status, priority, description, assignee, reported, and commentaries. Now that our mail configuration is done, we need to allow MENDEL to send emails.

Go into Settings, under the System Tab, select the System Components row. Click on it. You should see new a window called System Components. Choose the sensor for which you want to allow Incidents. Incident Management is in the last row. Turn it on, verify that email notification is checked as well.

FALSE POSITIVES

After deployment, this is the most important step. What is a False Positive? It is a type of event or communication that is expected on the network but might be wrongly marked as dangerous or anomalous by MENDEL. In order to optimize the monitoring activity of the network, is necessary to eliminate false positives. After successful initial selection of False Positives, new anomalies on the network will be more obvious and easily noticeable.

You can identify events as False Positives in the Events tab. Click on a desired event, it should show you basic information, like hosts, subnets, and services. Choose the combination that you want to mark as a False Positive event, and then click on False Positive. This will open new window called False positive manager. There you can specify the False Positive rule in more depth. If you set up False Positives properly, work in Events tab, it will be easier to identify dangerous events.

Copyright © 2024 Greycortex

How does one spot a False Positive? It is necessary to be familiar with client's infrastructure to correctly eliminate False Positives. The following is a list of some common False Positives typically present in each network:

- Devices that are exposed to the Internet are under constant attacks from the outside. The most common type of event generated by these attack attempts are Inbound Scans. You can adjust the severity of Inbound Scans to 1 and Horizontal Scans to 2.
- Before subnets are specified, every new device in the network will trigger the "Discovery: New Subnet" event. It is advisable to decrease the severity of this event to 1 until the subnet range is properly specified.
- "Info:" events should also have decreased severity.
- Events identified as "Suspicious User-Agent" should have their severity decreased to 3. Many programs are using their own User Agents to act on behalf of a user.
- Internal DNS servers generally create many False Positive events. It's advisable to mark this communication to remove unnecessary events. We use False Positives to filter out events like:
- **DNS**: DNS Query … , DNS: Query …
- **Brute-force**: Potential DNS Brute-Force attempt
- **Periodic**: Repetitive Connections
- **Policy**: Request to an external DNS server (Only internal DNS servers should be allowed to communicate with external DNS servers. Other hosts on the network should always ask internal DNS servers.)
- **Flow**: DNS High Number of Connections
- Other similar events
- Email servers are similar to DNS servers. They generate unnecessary events that can be marked as False Positives. Some of those events are:
- **Periodic**: Repetitive Connections, **scan**: Rapid POP3 Connections - Possible Brute Force Attack, **scan**: Rapid IMAPS Connections - Possible Brute Force Attack, ...
- Mark Tor events on port 123 as False Positives. Some NTP nodes might be the same as Tor nodes.
- Inbound communication with blacklisted IP addresses is also a False Positive. Set its severity to 1. This behavior is common around the internet. Not every blacklisted device is part of a malicious group which will scan various devices or try to use an exploit on them.
- Outgoing web communication can be marked as False Positive, if it leads to legitimate web servers like Microsoft, ESET, Avast, Skype, Facebook, ...
- P2P events in the internal network can be marked as False Positives or their severity can be decreased.

SUBNETS

Subnets are a portion of the network which shares a common address component, usually the same prefix.

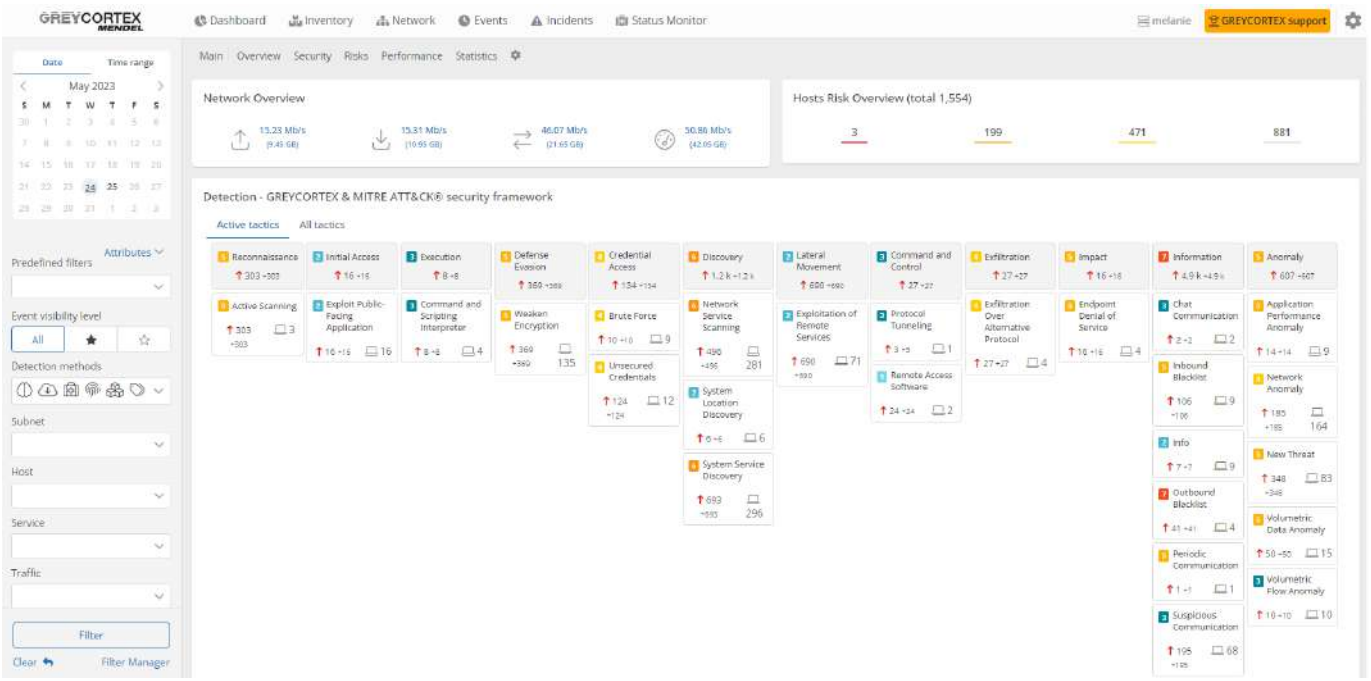
Setting these values is the same as the one in the wizard. If you have to define new subnets, you can do so in Settings, under the Network tab, in the row called Subnets. Every subnet that's in a customer's network infrastructure needs to be defined, for the reasons stated above.

Copyright © 2024 Greycortex

3. User Interface

3.1 Overview page

3.1.1 Basic User Interface



3.1.2 Web Interface

The application's user interface is implemented through a web application based on JAVA technology. The app supports viewing in the following web browsers:

3.1.3 Authentication

App authentication depends on deployment configuration. A basic user authentication is a form for entering a username and password for the authentication of locally created accounts. Other supported authentication options are:

3.1.4 Basic Controls

The application is controlled by user-defined filters (data, time) and visualization tabs, DASHBOARD, NETWORK, and EVENTS. The data visualization and data splicing principles are based on the use of single or double mouse clicks on the desired data row.

- A single-click displays basic detail, basic information, an information window, or applies the assigned function (display, filter, switch, etc.)
- A double-click always drills-down into more detailed display details

Outside the graph, only the left mouse button is used for control. In the charts, the right mouse button is used to indicate the desired interval.

Copyright © 2024 Greycortex

Moveable boxes

Every new window or dialog box (light box) which is opened as overlay is moveable.

3.1.5 Data Visualization

The application consists of two main areas:

- **The configuration area** - selecting display data
- **The display area** - visualization of selected data

3.1.6 Top Bar Menu



The display area is used to visualize the information you have acquired. The basic navigation is accomplished with the help of the upper tabs.

Attribute	Description
	one source or destination of the network communication defined by its IP address and sensor.
	User-defined information about the behavior of the monitored network, subnets, and users.
	A description of network communication with the ability to trace individual flows.
	Captured security events.
	Reported events.
	Monitoring local system resources and events.

3.1.7 Options menu

This menu is located in the top right corner of the application. The options are:

Attribute	Description
	System settings.
	Editing a user account, changing language, theme, show login history, setup level of displayed data.
	The URL generates a link to the current page as HASH or a set of parameters defining a preview.
	Used for reporting an incident. This is only visible in the options when you are on the pages in the Event Tab.
	Used for log out from account.
	Monitoring local system resources and events.
	Manuals for every released version and Guides for specific actions.

Copyright © 2024 Greycortex

3.2 Applications and filters

3.2.1 Introduction

The configuration area is used to select the date and time interval to display. You must use the button to apply the desired change or selection.

The configuration area contains the following setup options:

- Select time or time interval (calendar)
- Predefined filters selection
- Event visibility level
- Detection method
- Fast filters
- button
- button
- access

Side Bar Menu

SELECTING DATE AND TIME INTERVAL

Date

The Calendar field is used to select a specific date. Only days for which records are available can be selected. Clicking on the date will directly filter data for the selected date.

Time range

The time range is used to select a time interval: -.
The shift the time interval for the selected value.
Button sets the time interval from the midnight of the current day till midnight of the next day.
ensures that the new data in visual components, e.g. graphs, are loaded and displayed.

EVENT VISIBILITY LEVEL

ALL	All events without applied event visibility level filter
Top events	The most relevant and important detected events
Other events	Less important detected events

Important

This Event visibility level filter is superior to all other filters. It´s applied immediately after selection and it´s NOT cleared by button.

Copyright © 2024 Greycortex

DETECTION METHODS FILTER

The filter allows the selection of the detection method.

- Network Behavior Analysis (NBA)
- Intrusion Detection System (IDS)
- Event Correlation (EC)
- Log Processing (EM)
- Universal Tagging Engine (UnTE)
- Status Monitor (SM)

CONTENT FILTERING

Use the quick filter to select data for visualization. You can filter based on arbitrarily defined values like subnet, IP address, service, data flow direction, event name, event severity, or other attributes.

BASIC FILTERS

Filters are divided into categories -- fast filters, saved filters, and advanced filters. Individual filters can be combined with each other.

The Fast filter allows filtering by direct selection. It includes options for setting, extending, and subsequent storage, thus creating custom saved filters.

CONFIRMING BASIC FILTERS

The filter function is applied after entering the required parameter into the field or fields and then confirming it with the button.

The meaning of buttons for advanced filtering options:

- Advanced filter manager
- Clear all filter forms
- Undo the last step in filter
- Attributes allow customizing the filter

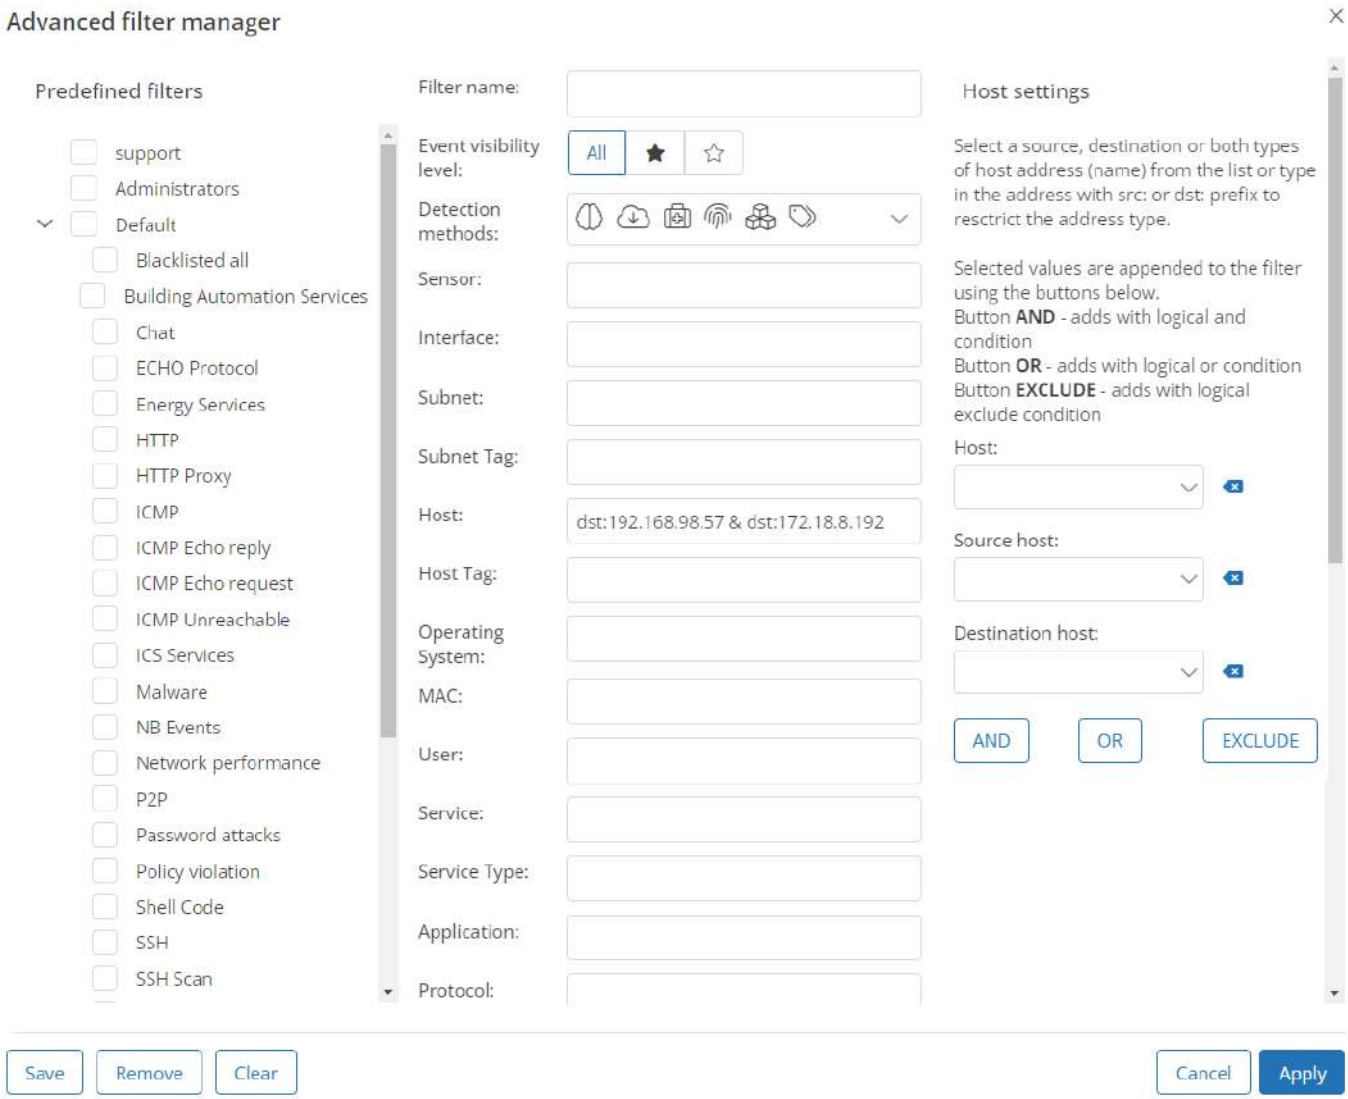
SETUP OF ADVANCED FILTERS

Advanced filters can be set up using the button.
Advanced filter manager allows you to save, create, and edit custom filters according to any criteria.
The Advanced filter manager dialog consists of three columns:

- List of saved filters -- **Left part**.
- List of search attributes -- **Middle part**.
- Input Wizard/Requirements Importer -- **Right part**. (appears for the given filter attribute)

Click on the desired .
The required importer will be activated.
In the importer, you need to enter/select one or more desired values separated by a comma in the appropriate field,
for example: source address, destination address, address.
Using the buttons below the parameters, the desired parameter field is inserted into the appropriate field.
Attributes can be inserted manually to direct editing.

Copyright © 2024 Greycortex



Example:

Edit buttons for Advanced filter manager:

- Save filter.
- Delete filter.
- Clear filter.



Copyright © 2024 Greycortex

3.3 Dashboard

Main Dashboard

GREYCORTEX & MITRE ATT&CK® security framework

An interactive dashboard based on GREYCORTEX's and MITRE ATT&CK®'s knowledge in the field of cyber security. A clear and comprehensive view of specific threat models and methodologies in accordance with detection in supervised networks. For both IT and OT environments.

All events are sorted into relevant tactics and techniques with added information about:

Severity	based on the event highest severity
Trend	rising or falling number of detecte events
Hosts	number of involved hosts

User-defined Views

The Dashboard provides a user-defined basic overview of overall network behavior and displays the risks and basic network and security events statistics.

Predefined Dashboards	Description
	Helps to visualize overall network state.
	Shows the overview of the monitored hosts divided into categories. These categories are described later in the signatures section. By clicking on the chart, you are redirected to the Hosts by risk page and the filter is applied.
	Shows hosts that are the most threatening in the monitored network. There is comprehensive information about security of devices in network traffic.
	Helps to visualize the overall performance. By clicking on the dots in the charts, you are redirected to the Hosts page.
	Shows top statistics through the whole network. It illustrates situations in the network and gives the chance to see all the needed information about network traffic in one place.
	Monitoring local system resources and events.

Dashboards can be added, modified or deleted via the Dashboard manager. For detailed information see Dashboard Manager Guide.

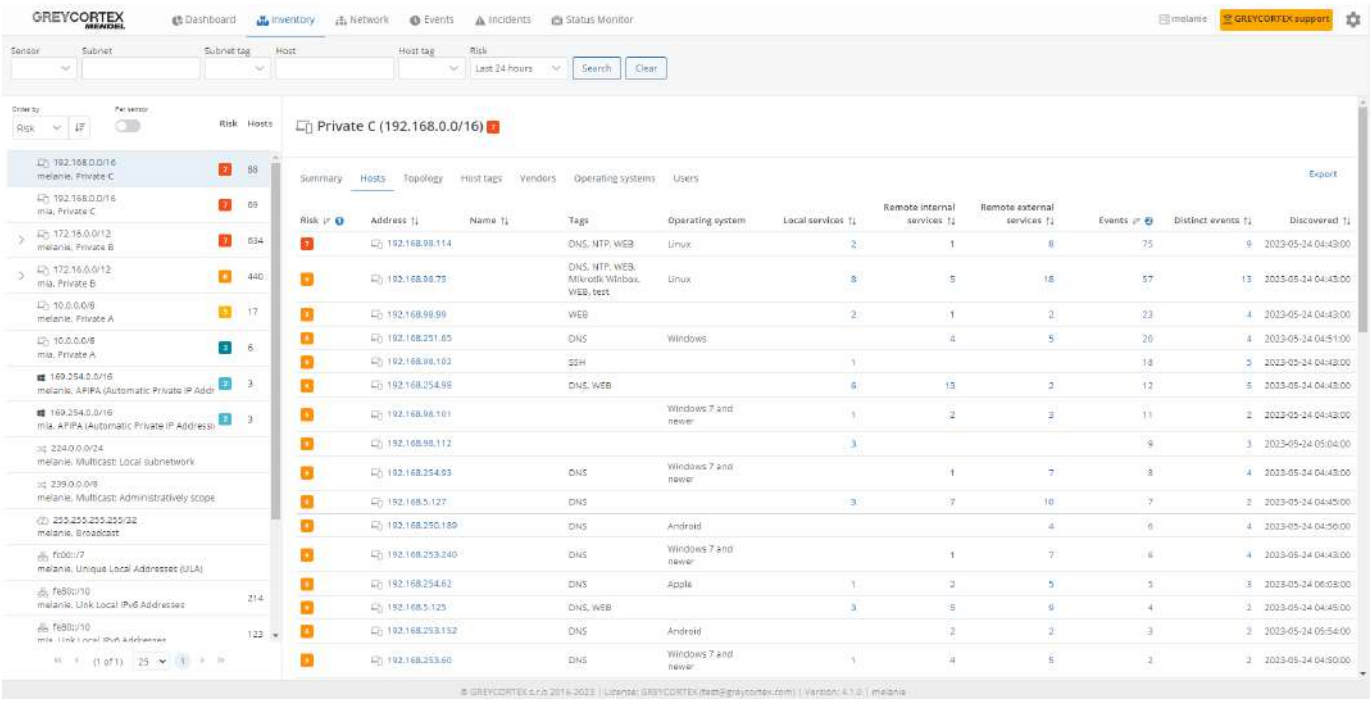
Copyright © 2024 Greycortex

3.4 Inventory

3.4.1 Inventory

Your network topology based on network discovery with calculated risk, broken down by subnets and devices.

Layout



HORIZONTAL FILTER

Filter option by , , , , , .

is calculated from chosen time frame.

SUBNET TREE

Network topology tree based on network discovery (or user defined) with option to broke it down to the lowest subnet level.

Each subnet contains information about: , , , , .

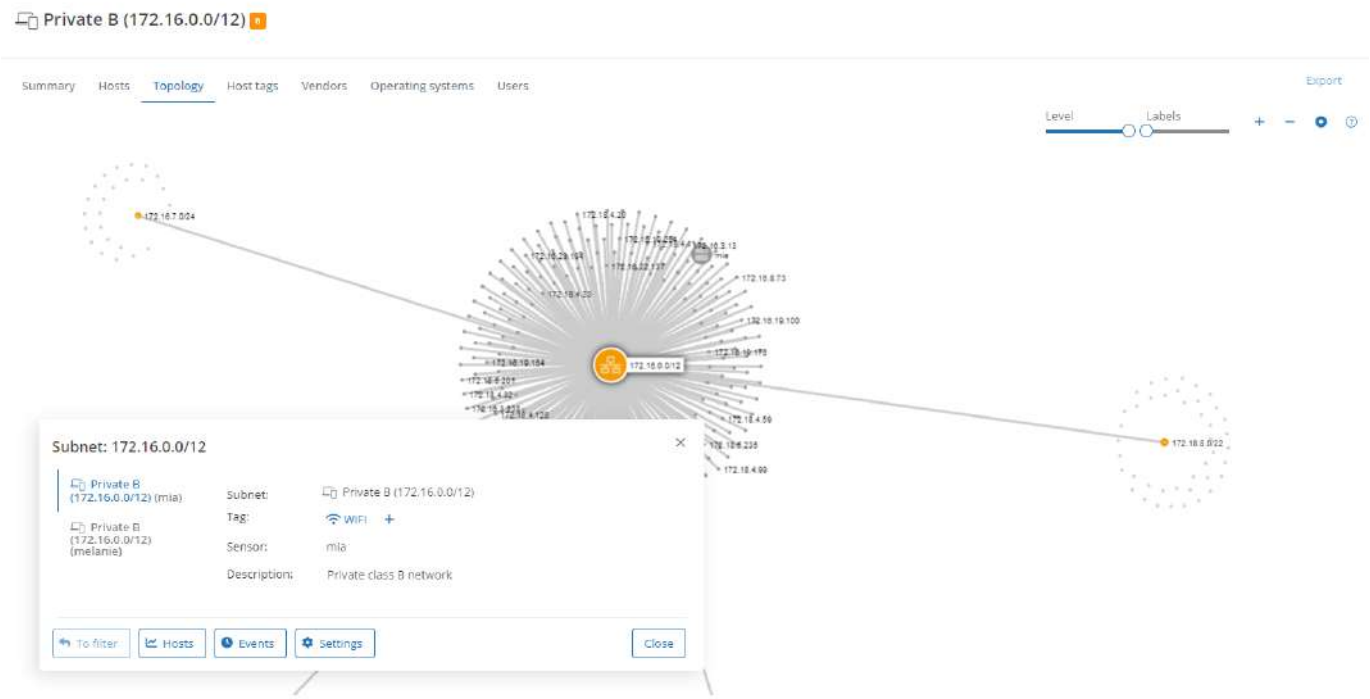
Order option by , , and .

button to change view from cumulated subnet view to the one divided per sensors.

Copyright © 2024 Greycortex

SUBNET INFORMATION BOX

Attribute	Description
	Subnet icon, name, CIDR and calculated risk
	All discovered or added information about specific subnet
	List of all hosts from specific subnet
	Rendered graphical representation of specific subnet with hosts
	List of all host tags added (with link to respective hosts and events)
	List of device vendor/manufacturer
	List of operting systems used on hosts in specific subnet
	List of identified users (based on Identity integration)



MISCELLANEOUS

button with possibility to edit subnet name and export generated view into the csv file.

Copyright © 2024 Greycortex

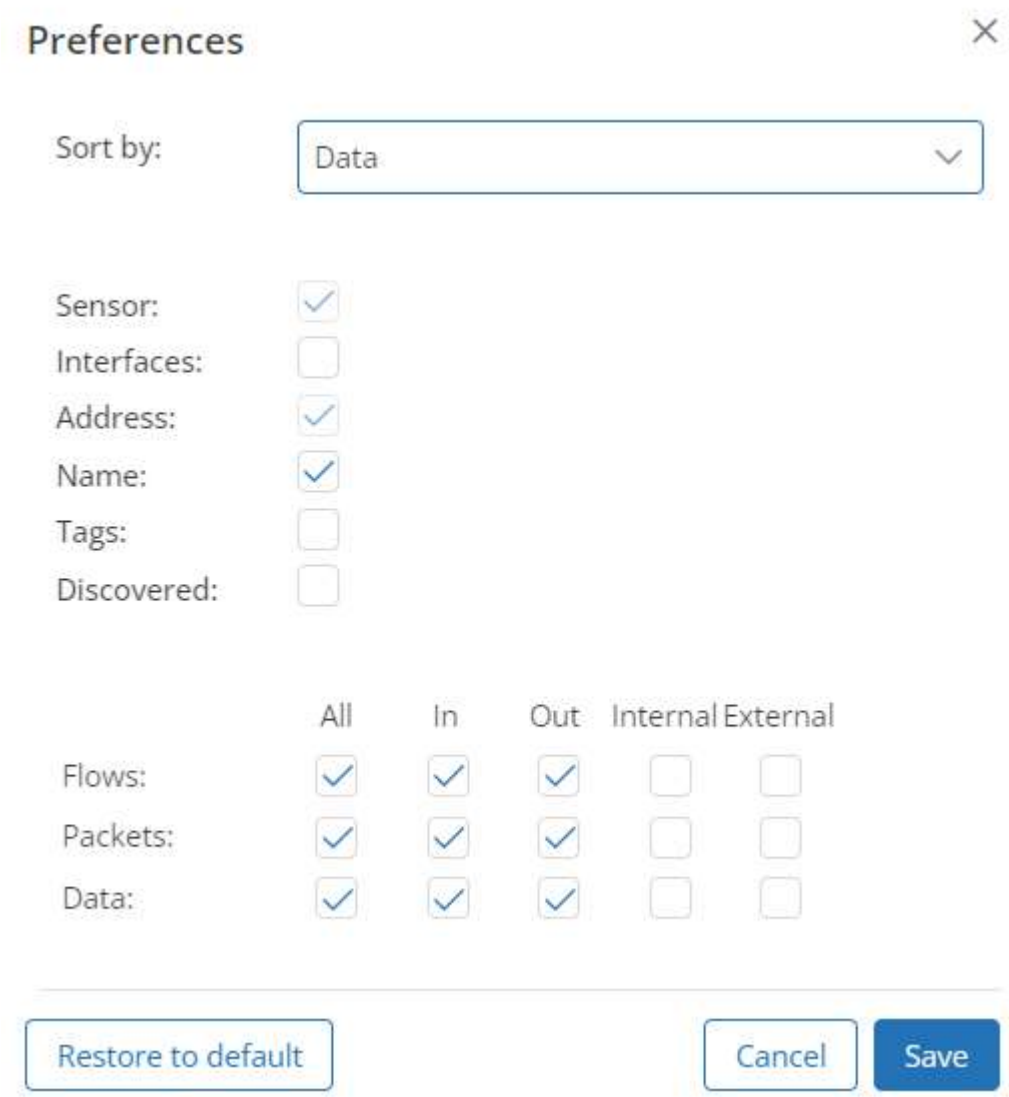
3.5 Network

3.5.1 Introduction

Network

Network Tab

Use the Network Menu to track individual data streams. This Menu shows basic overviews of the entire network like subnets, hosts, access to individual services, and sub-network streams. The data is mostly visualized in the graph and table. Visible parameters can be changed for each view in the Preferences.



NETWORK ACTIVITY GRAPH - SUBNETS

This graph is used to visualize flows in the monitored network over time (x-axis). The y-axis represents the quantity of the parameter. By default, the axis is depicted on a logarithmic scale. Clicking on the button can switch views from logarithmic to linear.

The graph always defines the contents of the data in the table below, based on time. You can move the graph (left mouse clicking and dragging), zoom in/out (ctrl + mouse wheel), or select the desired area (right mouse click and dragging). To control this zooming, use the mouse or buttons located in the top right corner of the chart.

Copyright © 2024 Greycortex

The individual curves displayed in the graph can be activated using the color checkboxes at the top of the graph.

NETWORK ACTIVITY TABLE

The table provides an overview of information about individual networks, hosts, services, and flows. The parameters displayed in both the table and the chart can be set when you click the button.

EXPORT DATA

This option is located in the right corner above the tables next to Preferences. The visible columns can be exported to .csv file from Subnets, Hosts, Services, Flows, Analysis, Events by Severity, Host by Risk, and Status Monitor. The limit for exporting to one file is 10 000 rows.

TABLE SETTINGS

Tables can be set as shown:

Attribute	Description
	Default sorting of the table.
	The number of records displayed in the table.

DIRECTIONAL PARAMETERS

Most parameters can also be combined in the following variants:

Attribute	Description
	All data.
	Traffic to the monitored network.
	Traffic from the monitored network.
	Traffic with both source and destination within the monitored network range.
	Traffic where either the source, the target or both are outside the monitored network range.

TABLE TABS

Included table tabs:

Attribute	Description
	Individual columns contain sums of all values (Packets, Flows, Data In, Data Out) shown in the graph.
	The value at a specific time.
	Network model (average + possible deviation).

All columns in the table can be sorted according to your needs. Sorting takes place over the full database, which may be more time-consuming in some cases.

Copyright © 2024 Greycortex

BASIC CONTROLS

Basic control of the entire data visualization mechanism is based on the **single click** and **double click** method:

Attribute	Description
	Visualize the relevant data in the graph.
	Device information and other analytics features.
	Drill down to a deeper display level: • To subnets on hosts. • To hosts on services. • To services on flows or peers. • The Analysis tab is used to visualize selected statistics.

Further advanced controls are possible by filtering the appropriate parameters in the filter form.

USING FILTERS

Data can be filtered using fast or saved filters. At each level, you can filter only using active (non-gray) quick filter fields.

The filters are applied sequentially during browsing of all pages.

Copyright © 2024 Greycortex

3.5.2 Subnets

Subnets

The Network is logically divided into subnets. Each subnet has own role and needs another level of security. Subnet tab is showing all subnets in the system. There is all information that users need to know about Flows, Packets and Data including advanced details about the direction (In/Out).

Clicking on the row of the table, the chart will change and shows just the data of that specific subnet. Clicking on the value in the Flow column will redirect to the Flows tab and display only flows that are specific for this subnet.



Sum Current Export data Preferences

Sensor	Address	Name	Flows	Flows In	Flows Out	Packets	Packets In	Packets Out	Data [B]	Data In [B]	Data Out [B]
melanie	0.0.0.0/0	All Internal	2.93 M	2.57 M	2.91 M	62.44 M	53.44 M	50.67 M	49.31 G	39.81 G	35.32 G
melanie	172.16.0.0/12	Private B	2.85 M	2.54 M	2.85 M	36.32 M	31.44 M	27.5 M	26.8 G	21.27 G	17.05 G
melanie	192.168.0.0/16	Private C	49.24 k	16.54 k	41.15 k	24.96 M	21.29 M	22.08 M	21.48 G	18.11 G	17.27 G
mia	0.0.0.0/0	All Internal	44.95 k	19.74 k	41.03 k	2.59 M	1.12 M	1.73 M	2.43 G	856.77 M	1.68 G
mia	172.16.0.0/12	Private B	29.13 k	12.94 k	27.4 k	2.24 M	930.8 k	1.5 M	2.12 G	688.58 M	1.47 G
melanie	10.0.0.0/8	Private A	19.65 k	1.8 k	19.64 k	921.9 k	552.16 k	921.9 k	862.88 M	356.52 M	862.88 M
mia	192.168.0.0/16	Private C	9.8 k	2.91 k	7.76 k	168.41 k	106.69 k	98.46 k	146.21 M	99.73 M	74.52 M
mia	172.18.0.0/22	Automatically created subnet	499	360	475	113.1 k	22.19 k	92.07 k	120.74 M	2.04 M	119.11 M
melanie	172.18.0.0/22	Automatically created subnet	379	285	362	112.89 k	22.17 k	91.92 k	120.48 M	2.04 M	118.86 M

DISPLAYED PARAMETERS

The following data can be displayed in columns:

Attribute	Description
	Source sensor.
	Interface of the particular subnet.
	The corresponding icon assigned to each subnet and host.
	Range of IP addresses of the logical subnet of the monitored network defined by the user.
	The name of the logical subnet of the monitored network, defined by the user.
	Number of flows.
	The number of transmitted packets.
	The amount of data transferred.

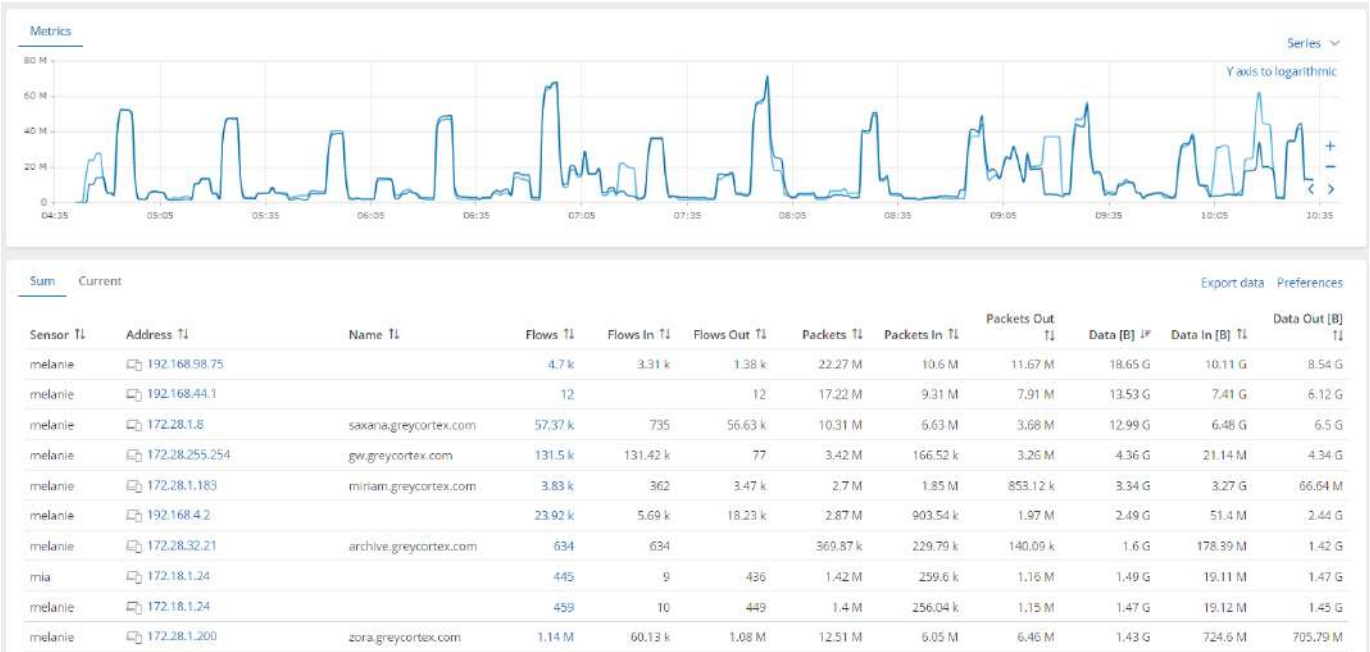
Copyright © 2024 Greycortex

3.5.3 Hosts

Hosts

Devices of the selected subnet -> hosts are defined as unique MAC address + IP addresses.
Clicking on any line displays the flow of the selected device in the graph.
Selecting a device by clicking its IP address will open
the "Host Information" dialog box with the information that are also visible in the:

for the corresponding host. Clicking on the in the menu will redirect to the services of the selected device.



Sensor	Address	Name	Flows	Flows In	Flows Out	Packets	Packets In	Packets Out	Data [B]	Data In [B]	Data Out [B]
melanie	192.168.98.75		4.7 k	3.31 k	1.38 k	22.27 M	10.6 M	11.67 M	18.65 G	10.11 G	8.54 G
melanie	192.168.44.1		12		12	17.22 M	9.31 M	7.91 M	13.53 G	7.41 G	6.12 G
melanie	172.28.1.8	saxana.greycortex.com	57.37 k	735	56.63 k	10.31 M	6.63 M	3.68 M	12.99 G	6.48 G	6.5 G
melanie	172.28.255.254	gw.greycortex.com	131.5 k	131.42 k	77	3.42 M	166.52 k	3.26 M	4.36 G	21.14 M	4.34 G
melanie	172.28.1.183	miriam.greycortex.com	3.83 k	362	3.47 k	2.7 M	1.85 M	853.12 k	3.34 G	3.27 G	66.64 M
melanie	192.168.4.2		23.92 k	5.69 k	18.23 k	2.87 M	903.54 k	1.97 M	2.49 G	51.4 M	2.44 G
melanie	172.28.32.21	archive.greycortex.com	634	634		369.87 k	229.79 k	140.09 k	1.6 G	178.39 M	1.42 G
mia	172.18.1.24		445	9	436	1.42 M	259.6 k	1.16 M	1.49 G	19.11 M	1.47 G
melanie	172.18.1.24		459	10	449	1.4 M	256.04 k	1.15 M	1.47 G	19.12 M	1.45 G
melanie	172.28.1.200	zora.greycortex.com	1.14 M	60.13 k	1.08 M	12.51 M	6.05 M	6.46 M	1.43 G	724.6 M	705.79 M

DISPLAYED PARAMETERS

The following data can be displayed in columns:

Attribute	Description
	Source sensor.
	Interface of the particular subnet.
	The corresponding icon assigned to each host.
	IP adress of the host.
	The name of the host.
	The amount of names that has specific host.
	Name of the detected operation system.
	Tags identify what the host is used for, and how the detectors behave.
	The time that the host was firstly discovered in the network. It is displayed only if the host is from the internal subnet. • Flows -> Number of flows. • Packets -> The number of transmitted packets. • Data -> The amount of data transferred.

Copyright © 2024 Greycortex

HOST LIFETIME

Host Lifetime is a complex cycle and depends on various conditions.
In MENDEL, a host is represented by a sensor and an IP address and
is detected by the existence of an unknown IP address in the monitored subnets.
When an unassigned IP address appears, a new host is created.
When a new destination IP address is found, a new host is created only
if the communication was evaluated as a valid one (a valid answer is visible).
Then a set of detectors evaluate if the new IP address means that the host has moved from an existing IP address
and it’s only the change of existing host, or whether the host is really new. Every subnet has a lease time. It represents how long the record will be valid without any communication.
The default value is eight days. It means that after that time,
a host without communication for that time will be removed and a lost host event will be reported.

HOSTNAME ASSIGNMENT

MENDEL uses various sources how to determine a hostname.
It parses DHCP and DNS records and store IP addresses and hostnames into DNS cache structure for further use.
MENDEL populates source and destination domain name columns in the events and flows records.

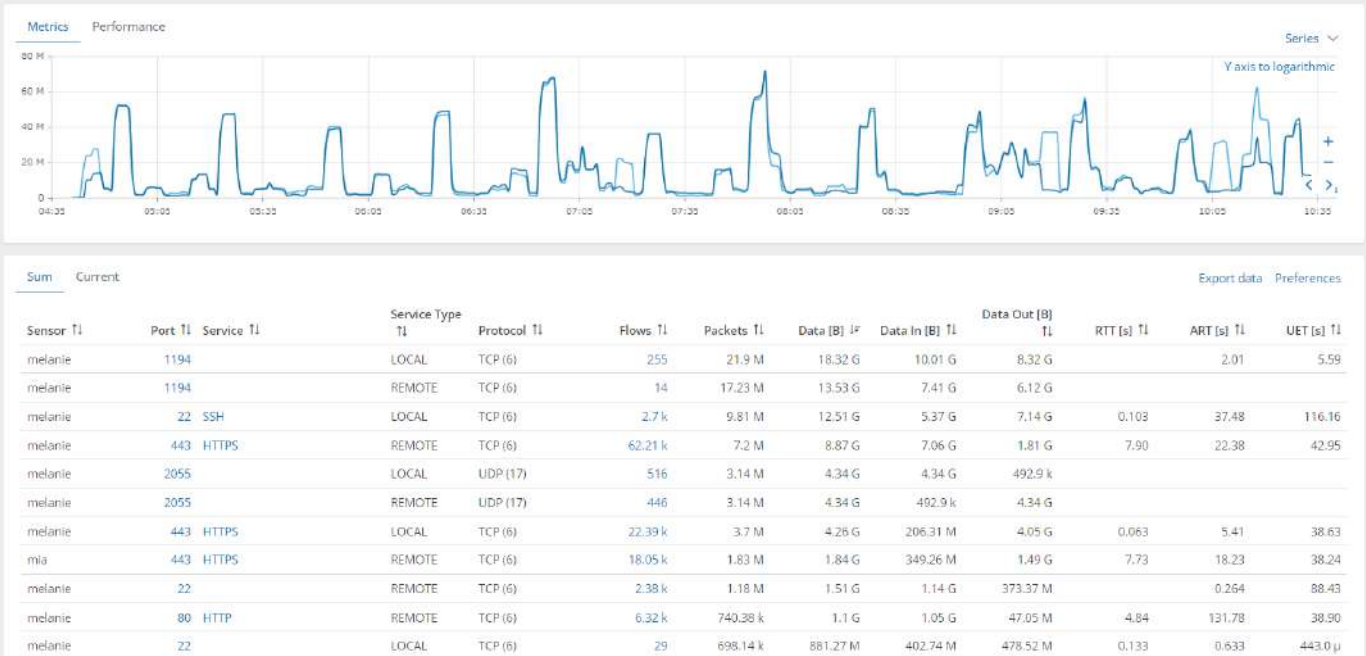
- When an event/flow is created, the names from the DNS Cache are stored in the object
  if they can be resolved at that time. If the name is resolved, then it is not changed during the object lifetime.
  If the name is not resolved, attempts to resolve are made each time the events record or flow update record is about to be created.
- If protocol is based on TLS and the "Client Hello" request contains a "server name" extension,
  then the name from this extension is used as the destination domain name. The source domain name is always taken from DNS Cache.
- If the protocol is HTTP and host request header contains a name, then the name from the host request header is used as a destination domain.
- If the protocol is HTTP and a connection is established through proxy, then we have host-to-proxy and proxy-to-server connections.
- For a proxy-to-server connection, MENDEL tries to find a name for the IP address from x-forwarded-for request header. If a name is found, then this IP address and name are used as the source IP and domain.

The DNS cache structure is periodically maintained to contain only valid, unexpired, and up to date records.

Copyright © 2024 Greycortex

3.5.4 Services

Services



Sensor	Port	Service	Service Type	Protocol	Flows	Packets	Data [B]	Data In [B]	Data Out [B]	RTT [s]	ART [s]	UET [s]
melanie	1194		LOCAL	TCP (6)	255	21.9 M	18.32 G	10.01 G	8.32 G		2.01	5.59
melanie	1194		REMOTE	TCP (6)	14	17.23 M	13.53 G	7.41 G	6.12 G			
melanie	22	SSH	LOCAL	TCP (6)	2.7 k	9.81 M	12.51 G	5.37 G	7.14 G	0.103	37.48	116.16
melanie	443	HTTPS	REMOTE	TCP (6)	62.21 k	7.2 M	8.87 G	7.06 G	1.81 G	7.90	22.38	42.95
melanie	2055		LOCAL	UDP (17)	516	3.14 M	4.34 G	4.34 G	492.9 k			
melanie	2055		REMOTE	UDP (17)	446	3.14 M	4.34 G	492.9 k	4.34 G			
melanie	443	HTTPS	LOCAL	TCP (6)	22.39 k	3.7 M	4.26 G	206.31 M	4.05 G	0.063	5.41	38.63
mia	443	HTTPS	REMOTE	TCP (6)	18.05 k	1.83 M	1.84 G	349.26 M	1.49 G	7.73	18.23	38.24
melanie	22		REMOTE	TCP (6)	2.38 k	1.18 M	1.51 G	1.14 G	373.37 M		0.264	88.43
melanie	80	HTTP	REMOTE	TCP (6)	6.32 k	740.38 k	1.1 G	1.05 G	47.05 M	4.84	131.78	38.90
melanie	22		LOCAL	TCP (6)	29	698.14 k	881.27 M	402.74 M	478.52 M	0.133	0.633	443.0 µ

DISPLAYED PARAMETERS

The following data can be displayed in columns:

Attribute	Description
	Source sensor.
	Interface of the particular subnet.
	Predefined service identifier (name).
	List that defines the ports used by the service.
	Represents the type of the service. • Displays device communication on the selected port: • LOCAL represents the local service to which other devices are connecting. • REMOTE is a remote service to which the analyzed host is connecting to.
	The Internet protocol defines the TCP/UDP protocols used by the service.
	Number of flows.
	The number of transmitted packets.
	The amount of data transferred in Bytes.
	The duration that it takes for a network request to go from a starting point to a destination and back again to the starting point.
	It describes a number of neural network models which use supervised and unsupervised learning methods, and address problems such as pattern recognition and prediction.
	User Experience Time.

Copyright © 2024 Greycortex

SERVICE DETAIL

Double-click on the used by the service opens a dialog box with Service information details.

Attribute	Description
	Predefined service identifier (name).
	The Internet protocol defines the TCP/UDP protocols used by the service.
	Suggested network protocols based on a service/port.
	Timestamp represents first discovery of the Service in supervised network.
	Source sensor.
	Option to find additional information about Service/Port through search engines and knowledge bases.
	Description of the detected service, protocol, etc.
	Suggestion of the commonly used service if the service is not recognized.

Copyright © 2024 Greycortex

3.5.5 Peers

Peers

An aggregated list of all communication partners of that device or the result of the filter used. The Peers parameter (the number of communication partners) listed in the tables is partly estimated due to computational difficulty and may not be quite accurate.

The application allows three different views on the connection of the communication partners: **chart, hosts, subnets**.

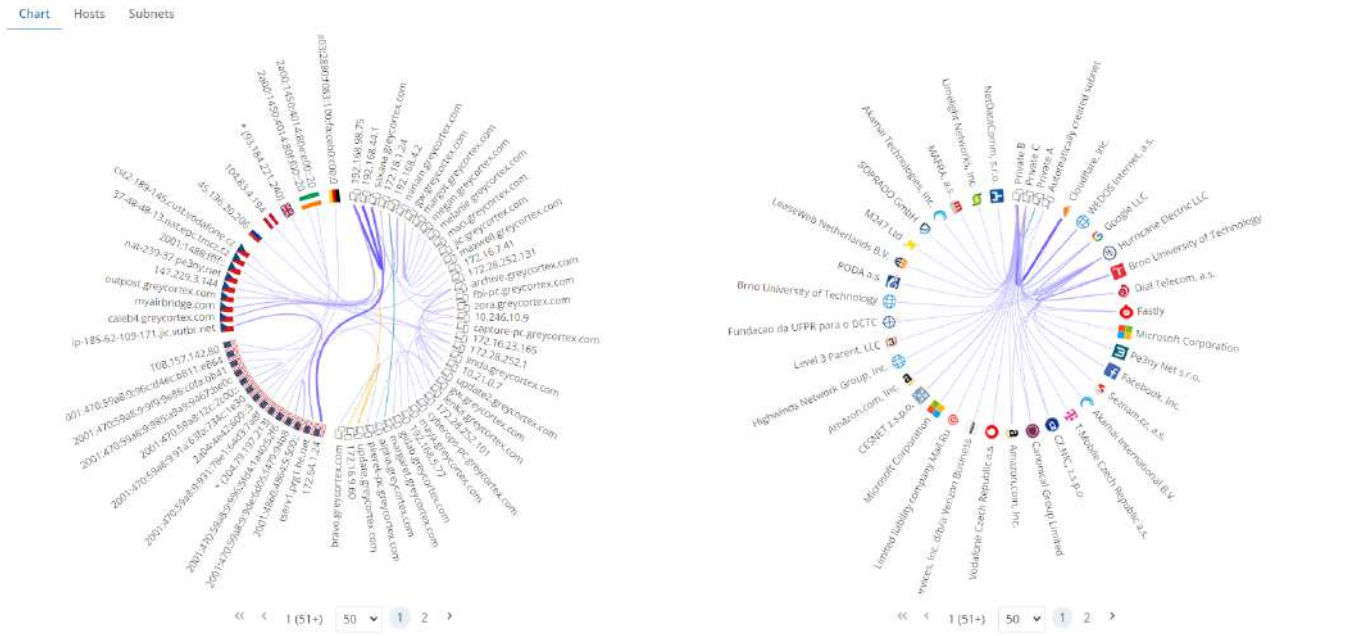


Chart Hosts Subnets

Export data Preferences

Peer			Peer		Summary	
Host ↑↓	Packet Count ↑↓	Packet Size ↑↓	Host ↑↓	Packet Count ↑↓	Packet Size ↑↓	Protocol ↑↓	Flows ↑↓
192.168.98.75	9.3 M	7.4 G	192.168.44.1	7.9 M	6.1 G	TCP	12
172.64.1.24	6.1 M	5.4 G	172.28.1.8	3.3 M	6.5 G	TCP	260
216.66.86.122	2.5 M	2.7 G	185.62.109.171	6.7 M	2.8 G	IPV6	23
192.168.4.2	1.6 M	2.2 G	172.28.1.183	623.4 k	25.0 M	TCP	70
172.18.1.24	1.2 M	1.5 G	88.208.76.141	257.9 k	15.6 M	TCP	66
172.18.1.24	1.2 M	1.5 G	88.208.76.141	254.1 k	15.4 M	TCP	68
172.28.255.254	519.8 k	720.0 M	172.28.1.178	0	0	UDP	11
172.28.255.254	519.8 k	720.0 M	172.28.1.173	19.0 k	10.9 M	UDP	28
172.28.255.254	519.8 k	720.0 M	172.28.1.170	0	0	UDP	11
172.28.255.254	519.8 k	720.0 M	172.28.32.10	0	0	UDP	11
172.28.255.254	519.8 k	720.0 M	172.28.1.171	543	310.9 k	UDP	13
172.28.255.254	519.8 k	720.0 M	172.28.32.14	0	0	UDP	11
172.28.32.21	60.7 k	476.0 M	172.28.1.8	143.3 k	8.0 M	TCP	1
192.168.98.75	1.2 M	336.0 M	89.221.218.75	1.7 M	3.1 G	TCP	24
192.168.98.75	235.2 k	288.3 M	94.142.239.37	137.0 k	37.1 M	TCP	4
192.168.98.75	568.7 k	269.0 M	46.28.107.158	529.2 k	688.8 M	TCP	14

Copyright © 2024 Greycortex

Chart Hosts Subnets

Export data Preferences

Peer			Peer			Summary	
Subnet	Packet Count	Packet Size	Subnet	Packet Count	Packet Size	Protocol	Flows
Private B (172.16.0.0/12)	3.3 M	6.5 G	Cloudflare, Inc.	6.1 M	5.4 G	TCP	1.5 k
Private C (192.168.0.0/16)	2.2 M	2.5 G	Private B (172.16.0.0/12)	1.1 M	342.2 M	TCP	28.4 k
Private B (172.16.0.0/12)	1.1 M	1.3 G	Dial Telecom, a.s.	235.7 k	14.3 M	TCP	59
Private B (172.16.0.0/12)	1.1 M	1.3 G	Dial Telecom, a.s.	232.0 k	14.1 M	TCP	83
Google LLC	1.4 M	801.5 M	Hurricane Electric LLC	6.1 M	2.7 G	UDP	111.1 k
Private C (192.168.0.0/16)	1.8 M	605.0 M	WEDOS Internet, a.s.	2.3 M	3.8 G	TCP	128
Fastly	232.6 k	378.9 M	Hurricane Electric LLC	136.2 k	11.0 M	TCP	310
Private B (172.16.0.0/12)	553.1 k	356.7 M	Private A (10.0.0.0/8)	921.8 k	859.7 M	TCP	20.6 k
Private C (192.168.0.0/16)	235.2 k	288.3 M	Pe3ny Net s.r.o.	137.0 k	37.1 M	TCP	4
Hurricane Electric LLC	206.2 k	270.0 M	Hurricane Electric LLC	244.9 k	63.3 M	TCP	8.0 k
Seznam.cz, a.s.	157.3 k	234.2 M	Hurricane Electric LLC	119.0 k	14.9 M	TCP	1.2 k
Google LLC	273.1 k	224.0 M	Hurricane Electric LLC	198.0 k	59.3 M	TCP	8.8 k
Private C (192.168.0.0/16)	210.1 k	219.7 M	T-Mobile Czech Republic a.s.	141.2 k	22.7 M	TCP	100
CZ.NIC, z.s.p.o.	124.6 k	205.4 M	Hurricane Electric LLC	74.6 k	5.7 M	TCP	8

Copyright © 2024 Greycortex

3.5.6 Flows

Flows

Flow is a bidirectional sequence of packets that all share the following values: Protocol, Source IP address, Destination IP address, Source port, Destination port, Tunnel. The flow which is not finished (is running longer than a minute) is reported every 5 minutes.

Export data Preferences

Src Host	Dst Host	Protocol	Dst Port	Service	Applications	Src Packet Count	Src Packet Size	Dst Packet Count	Dst Packet Size	Src Flags	Dst Flags
192.168.209.10	172.18.0.10	TCP	445	SMB		22	6.7 k	19	3.4 k	CE.AP.S.	.E.AP.S.
192.168.209.10	172.18.0.10	TCP	445	SMB		11	4.6 k	8	1.3 k	CE.AP.S.	.E.AP.S.
192.168.209.10	172.18.0.10	TCP	445	SMB		13	4.6 k	9	1.9 k	CE.AP.S.	.E.AP.S.
192.168.209.10	172.18.0.10	TCP	445	SMB		13	4.5 k	9	1.7 k	CE.AP.S.	CE.AP.S.
172.18.1.103	172.18.0.31	TCP	445	SMB		2	206	1	182	...AP...	...AP...
192.168.209.10	172.18.0.10	TCP	445	SMB		1	60	0	0	...A....	.....
192.168.209.10	172.18.0.10	TCP	445	SMB		1	126	1	126	...AP...	...AP...
192.168.209.10	172.18.0.10	TCP	445	SMB		1	126	1	126	C..AP...	...AP...
192.168.209.10	172.18.0.10	TCP	445	SMB		1	60	0	0	...A.R..	.....
192.168.209.10	172.18.0.10	TCP	445	SMB		1	60	0	0	...A.R..	.....
192.168.209.10	172.18.0.10	TCP	445	SMB		1	60	0	0	...A.R..	.....

This page lists all the flows based on the quick filter -- some value must be filtered. After clicking on the (symbol) the details of the flow are displayed. The details consist of the information about the flow and the data gathered according to the layers:

FLOW

Information about flow like Reported Timestamp, Total Duration, Flow Type etc. are displayed here.

LINK LAYER

Packets, MAC addresses, and interface information

NETWORK LAYER

Information about IP addresses, Hostname and IP Family

- Values **Min TTL**, **Max TTL** and **L3/L4 Feature Vector** serve for the operating system detection.
- Values **Min DSCP** and **Max DSCP** show the minimal and maximal Differentiated Services Code Point. This value from the IP header classifies the packet to provide a quality of service.

Copyright © 2024 Greycortex

TRANSPORT LAYER

Protocol, Flags and Port is listed here. Additional records will be reported for TCP flows only:

Attribute	Description
	Flow which was blocked on a Firewall and the Firewall responded with some ICMP answer (Reset or ICMP Reject, Prohibited, Unreachable).
	Starting with SYN + ACK.
	Timeout or other unknown reason.
	0 reply packets including no ICMP.
	Retransmission is the resending of packets which have been either damaged or lost. Retransmission is one of the basic mechanisms used by protocols operating over a packet switched computer network to provide reliable communication.
	Packets which were in a different order from which they were sent.
	TCP Zero Window is when the Window size in a machine remains at zero for a specified amount of time.
	Packets with every single option set for every used protocol.
	Total time of separate request, first with handshake, others as they come. The value shown in the column is the average.
	Is the length time it takes for a data packet to be sent to a destination plus the time it takes for an acknowledgment of that packet to be received back at the origin.
	Time between the first server acknowledgement and the start of the data transfer.

APPLICATION LAYER

If parsed, service information and application data are displayed here. It is possible to full-text filter the data if the data is available. In particular protocols, it is possible to see additional values:

- In the HTTP and HTTPS Requests and Responses, values NTT, and ART are added for separate transactions. Values ART and NTT are in milliseconds. They are visible only in the flows that are not encrypted. You can see the statistics according to ART and NTT on the Network -> Analysis page.
- **Network Transfer Time NTT** -- In the Request, this metric is the time from when the request message from the client is sent until the acknowledgement from the server is received. In the response, it is the time that it took for the data to be transferred from the server to the client and the last acknowledgement from the client is received.
- Protocols using SSL/TLS:
- **ja3Hash** -- JA3 fingerprint of SSL/TLS client connection. It is the MD5 hash of the ja3Str.
- **ja3Str** -- String containing several fields from the TLS Client Hello message sent during initiation of the TLS session (SSL/TLS handshake), that are concatenated together.

If you want to decrypt HTTPS traffic, you can upload a certificate and a private key of the device that you want the traffic to be decrypted for to the MENDEL. You can upload them in the:

After that, if it's not using **Diffie-Hellman**, the traffic should be decrypted.

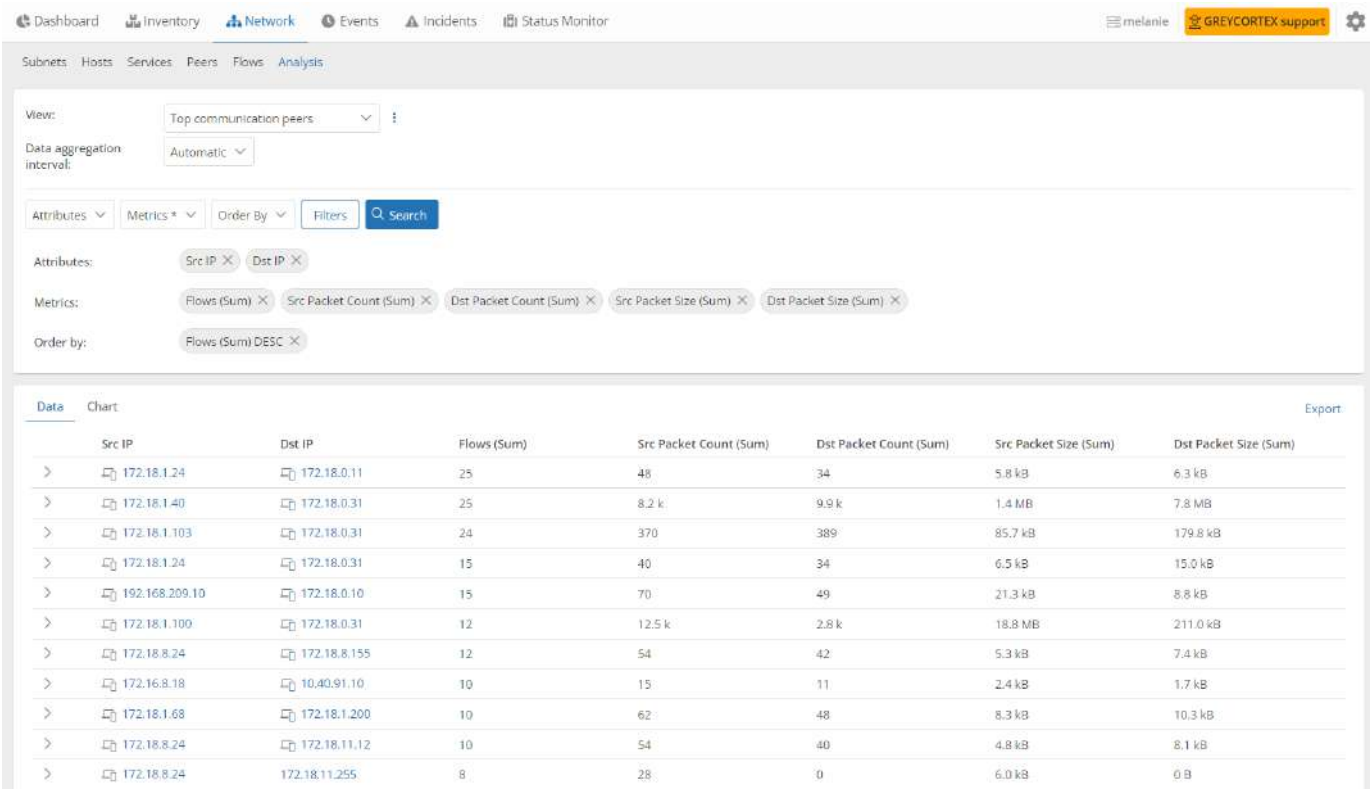
Copyright © 2024 Greycortex

3.5.7 Analysis

Analysis

Provides statistical visualization of selected parameters according to the user selection. Visualized data is based on non-computed views and you can define any view over the processed and stored data using attributes, metrics and other variables.

Generating views can be computationally advanced, and thus time-consuming.



Important

The data are selected from an indexed database. Individual queries can take more time due to database size (100 GB -- 100 TB). In order to optimize DB machines, each query is performed for a maximum of five minutes.

If the query does not retrieve the requested data within this time frame, the data returned will be blank. If you have a blank query returned, please adjust the time interval (limit the size of the searched data) and re-select.

For detailed information how to define your own query see Network Analysis Guide.

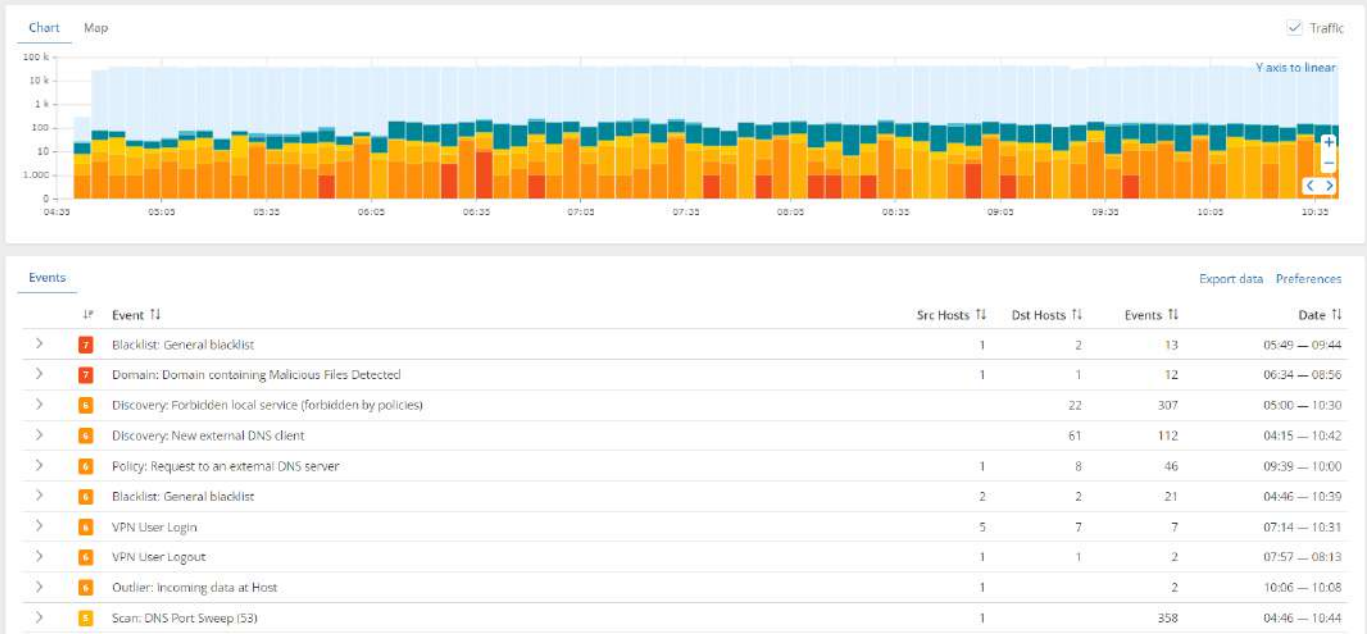
Copyright © 2024 Greycortex

3.6 Events

3.6.1 Events by Severity

Events by Severity

The Events Menu shows detected events in the appropriate time span.



CHART

The graph visualizes the severity of events and flows in the monitored
network over time (the X axis). The Y axis represents the amount of data.
The axis is depicted on a logarithmic scale. Click the button
if you want to switch the scale from logarithmic to linear.

The graph visualizes the number of identified events shown in the table below.
The severity of the event is identified by its color.
The graph is broken down into columns representing:

- 1 minute
- 5 minutes
- 30 minutes
- 4 hours
- 1 day

where one column represents the number of flows that have passed during a given time interval
(light blue color -- passage traffic), and the number of security events recorded.

You can drag the graph by left-clicking and dragging it, zoom in/out
(mouse wheel or Ctrl + mouse wheel when there is a scroller bar in the page)
or select the desired area (right mouse click). To control this, use the mouse or buttons located in the top right corner of the chart.
The "Traffic" item on the graph serves to turn off/on the number of displayed streams.

Copyright © 2024 Greycortex

TABLES

Identified events are displayed in the overview table below the graph. The default view is sorted by the severity of the events and contains the following columns:

Attribute	Description
	The severity of the event is shown as 1 (lowest severity) to 10 (highest severity). The severity of the events is generated contextually, i.e. for the same event, it can display different values in different dependencies (according to user settings).
	The basic description of the identified event.
	The number of unique source IP addresses in the event.
	The number of unique target IP addresses in the event.
	The total number of events of that type.
	The timing of the event determining the moment of occurrence or the time interval in which the event occurred.

You can change the table display options by using the Preferences configuration button.

SETTING TABLE CONTENT

Clicking on the button will bring up the menu to view more display options. You can see the following items in this step:

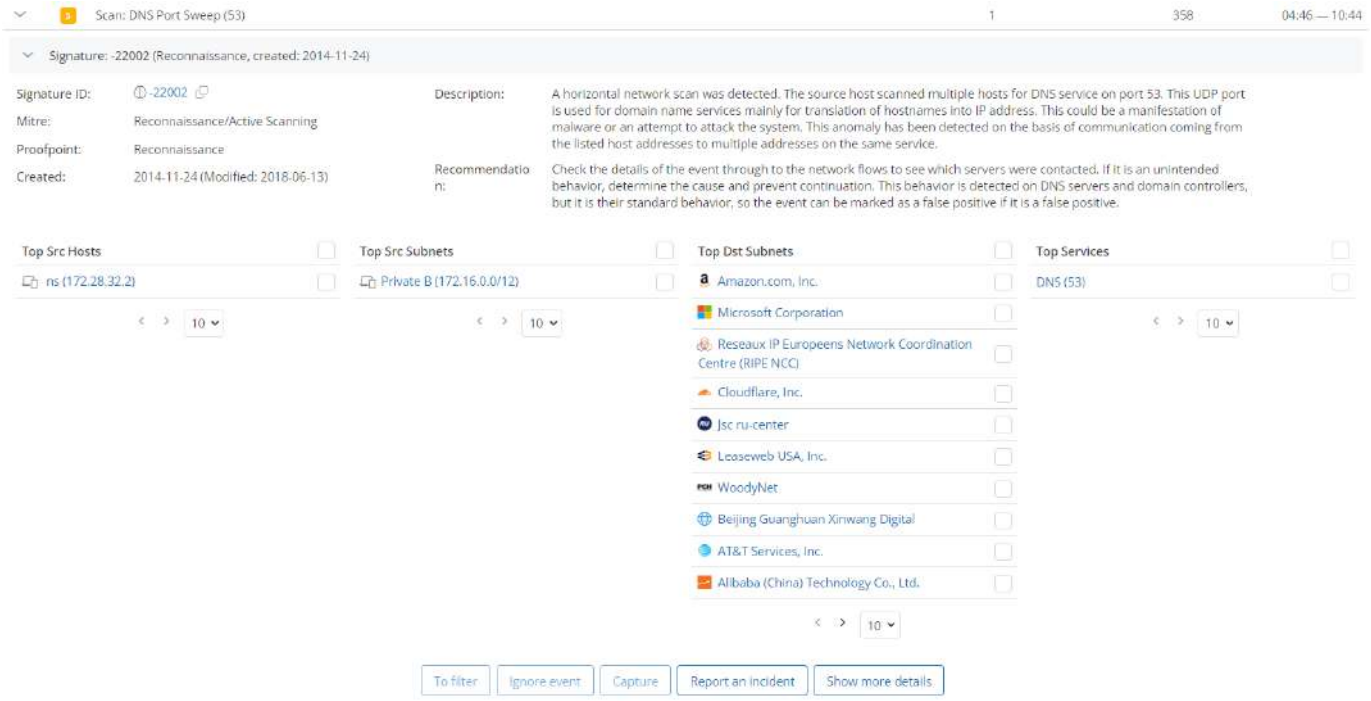
Attribute	Description
	The severity of the detected threat.
	Display the name of the sensor on which the event was detected.
	Basic description of the detected event.
	Source user of communication (LDAP).
	Destination user of the communication (LDAP).
	The network service on which the event was detected. (SSH, http, ...)
	Determines whether it is part of the local (LOCAL) or remote (REMOTE).
	Protocol family. (IPv4, IPv6, TCP, UDP, ICMP, ...)
	The number of identified events of that type.
	The occurrence of events determining the moment of occurrence or the time interval in which the event occurred -> cannot be turned off.
	Source host name/IP address.
	Destination hostname/IP address.
	Source subnet.
	Destination subnet.
	Viewing individual source and target IP addresses in a given event.
	View aggregated information about the event.
	Default sorting of the events.
	The number of displayed events.

Copyright © 2024 Greycortex

BASIC EVENT DETAIL

Each row in the report table can be expanded to show the basic detail with one mouse click on the line of the event. This displays a basic description of the event, and an overview of addresses, subnets, and services covered by the event.

The source and destination IP addresses, subnets, and services are sorted into columns according to their frequency in the event. The specific event and its details can be displayed by double-clicking on the appropriate value field, or by simply clicking the , , or .



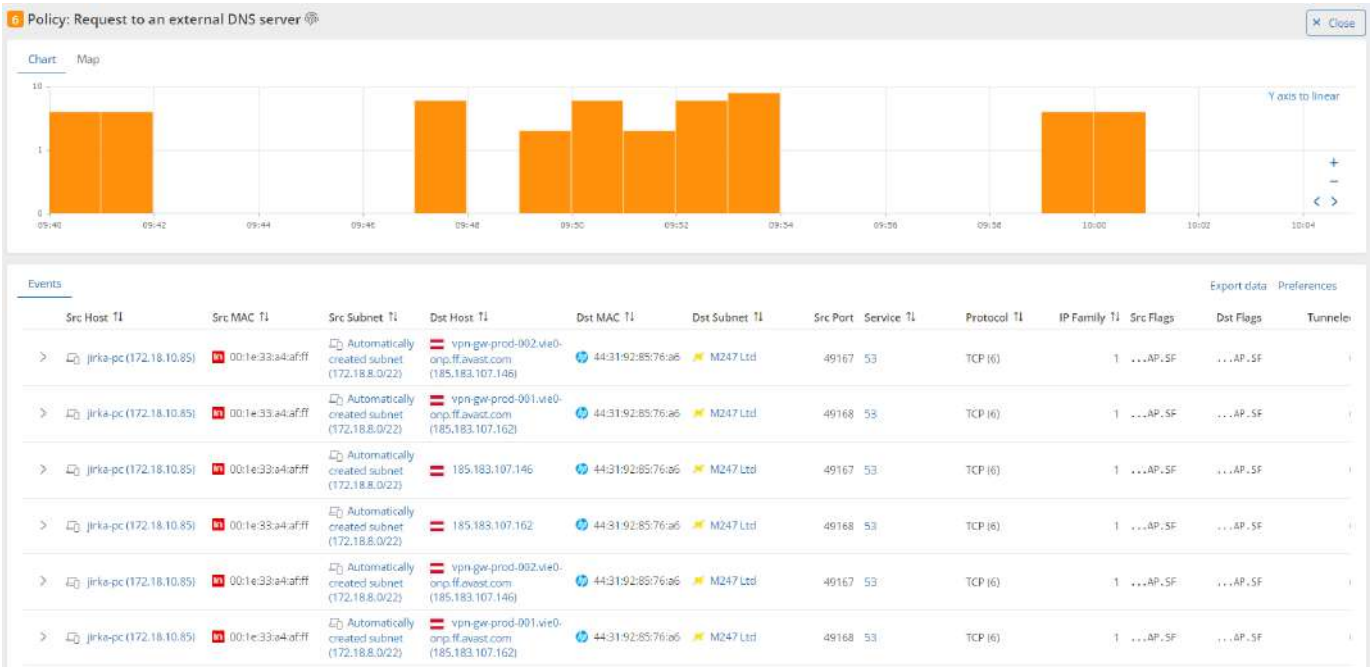
Functions of the buttons underneath the event detail:

Attribute	Description
	Adds an IP address to the filter, where it can then show all events related to this address or analyze its behavior in other tabs -- Dashboard, Network, etc.
	Adds an event to the list of false positive rules.
	Creates a new rule to record all network traffic of that event to the .pcap file.
	Allows to perform an action with chosen values using the selected plugin, e.g. blocking an IP address in your firewall.
	Reports an event to the Incidents Tab. This button is visible only when the Incident Management option is activated in the System Components under Settings.
	Displays a detailed description of all events. The event detail is always displayed after double clicking on the event line in the original view, the source or destination IP address, the source or destination subnet/autonomous network, the services, or the user.

ADVANCED EVENT DETAIL

The full details of each event are displayed in the "Lightbox" component. At the top is a chart, that has the same properties as the chart in the Events tab, and can be used as a time filter.

Copyright © 2024 Greycortex



The table below shows the details of separate events of the same type that make up the whole event.

Attribute	Description
	A link to more detailed information about the event.
	Source IP address, hostname, and username (if available).
	Source subnet.
	Destination IP address.
	Destination subnet.
	The service to which the flow relates.
	Number of flows.
	The amount of data transferred.
	Timestamp of the event.

You can change displayed columns through Preferences in the top right of the table.

Copyright © 2024 Greycortex

LIST OF EVENTS

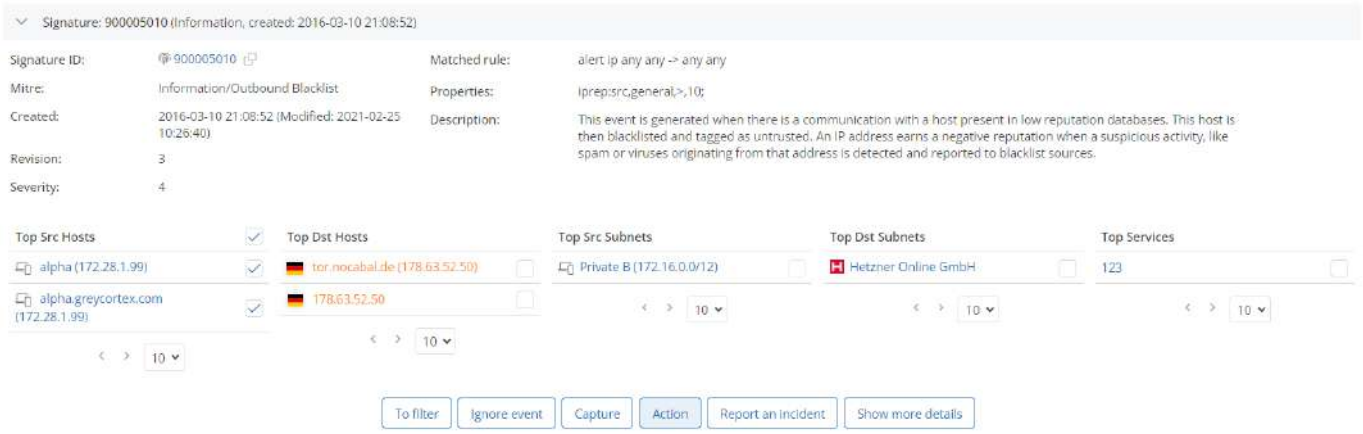
Simply click on the line or the button to expand and display specific details. Details vary depending on the detection module.

Attribute	Description
	Events are displayed statistically in relation to the given IP, subnet, and service. The values in the Current Model table show the information on which the event was detected. • Flows - Number of flows. • Packets - The number of transmitted packets. • Data - The amount of data transferred. • Peers - Number of communication partners. • Ports - Number of ports. • Duration -- Session duration. (network flow)
	Shows the captured packet and the rule (regular expression) that captured it.
	Displays the label of the logged-in event.
	Shows the list of events which triggered the correlation detection.

Double-clicking on the event line will show specific flows from which the event is generated.

Action

The Action option allows users to run scripts previously added to MENDEL and perform tasks described by the script e.g. adding a firewall rule, immediately, without leaving the MENDEL user interface. Executing the script can be done by clicking on the button in the Event Detail and then clicking on the button in the lightbox.



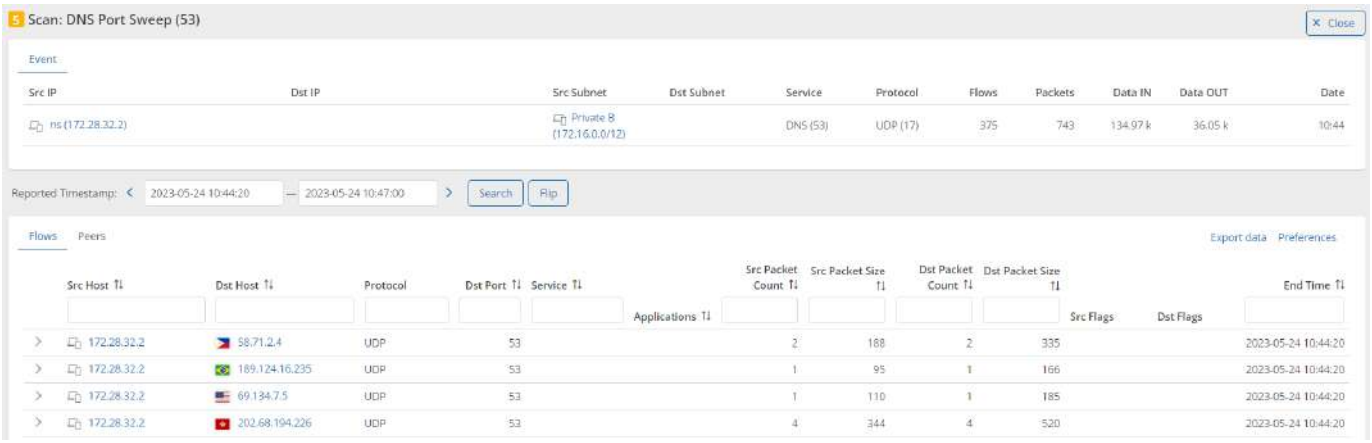
After clicking, the dialog presented below will appear with the following menu:

Attribute	Description
	Specific profile for the script used for the action the user wants to perform.
	Source or destination IP addresses of the subnet the user wants to execute the action with.
	Source or destination IP address of the host the user wants to execute the action with.
	Specific protocol the user wants to execute the action with.
	Network service the user wants to execute the action with.
	Source and destination network port the user wants to execute the action with.

Copyright © 2024 Greycortex

VIEWING THE EVENT FLOWS

The Flows tab shows event flows, the Peers tab displays an aggregated view of the flow from source or target IPs. To track all flows, set the start and end times. The Flip button is used to rotate the search condition so that the source is shifted to the target.

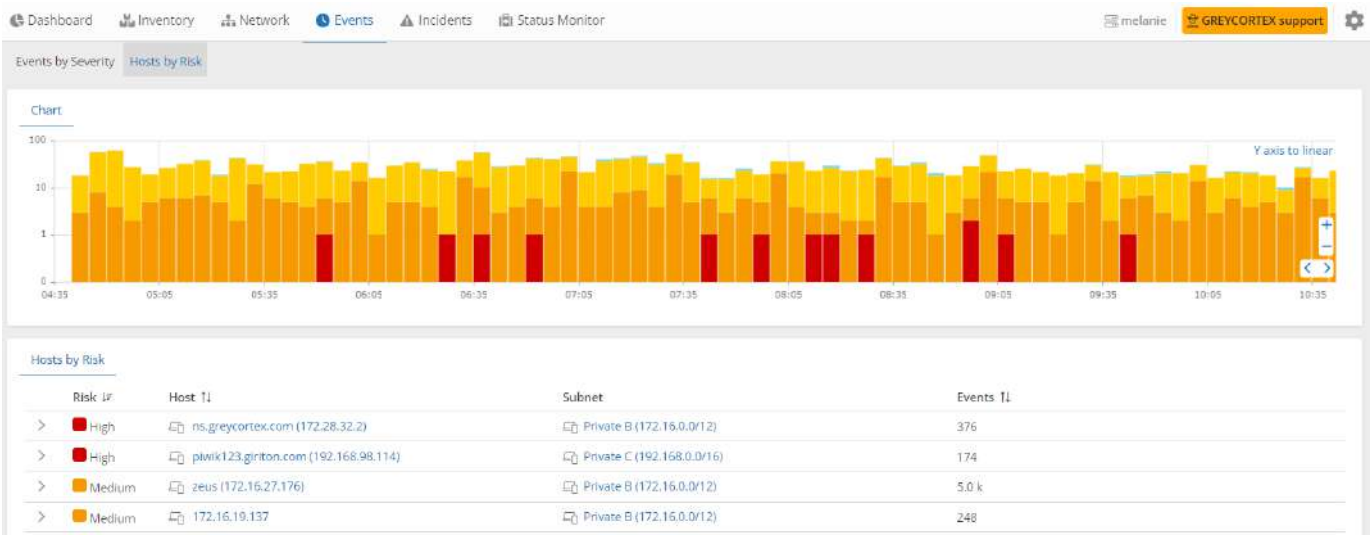


Copyright © 2024 Greycortex

3.6.2 Hosts by Risk

Hosts by Risk

This page shows detected events according to the host that was involved. The risk of the host is derived from the highest severity of the event.



Risk categories:

Attribute	Description
	Risk 9 to 10.
	Risk 7 to 8.
	Risk 5 to 6.
	Risk 2 to 4.
	Risk 1.

The bar graph shows the number of hosts (IP addresses) with a specific risk over time.
The hosts without risk or in false positives are not shown.
If the table underneath the chart has the column Sensor, the same IP address is counted 2x.

The table contains a list of all internal hosts ranked according to the risk who communicated within the timeframe defined in the calendar.
When you click on or in the guest table, information dialogs are displayed.
Within the expanded row of the table, a list of events associated with that guest is displayed.
This inner table has the same behavior as the table on the main event page, so it is possible to:

- create false positive,
- create capture,
- report an incident,
- to run an external script action and view detailed information in lightbox.

In the table, the Sensor column is displayed only if it is active on the Event page.
Columns can be sorted by the Risk, Host, Subnet, and Events.

Copyright © 2024 Greycortex

Filtering can be used as expected. The only special filtering is according to traffic:

Attribute	Description
	Only destination hosts.
	Only source hosts.
	Only destination hosts on the internal network.
	Only source hosts on the internal network.
	Cannot be used.
	Cannot be used.

Copyright © 2024 Greycortex

3.6.3 False Positives

Eliminating False Positive Detection

In some cases, the detection engine may mark legitimate behavior as harmful or anomalous. There are several ways to eliminate these phenomena and limit them:

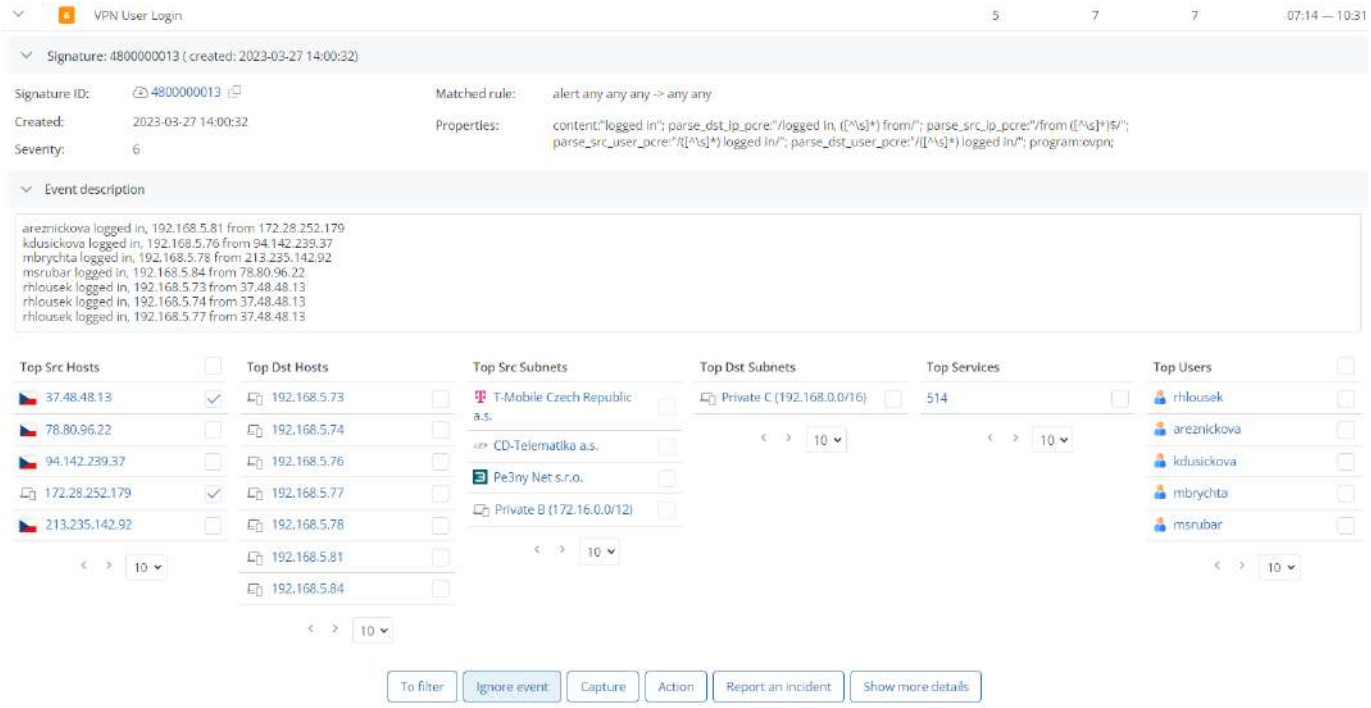
- **Turning off a rule or signature** - Prevents the event from occurring again. See NBA and IDS settings.
- **Event filtering** - Creating a filter rule that eliminates unwanted events in the appropriate outputs view.
- **Creating a false positive rule** - Creating a filter rule that eliminates unwanted events in the appropriate outputs view.

The goal is to eliminate any unwanted events so that the network appears to be clean. Each new event will be instantly identified and will not be "overlooked."

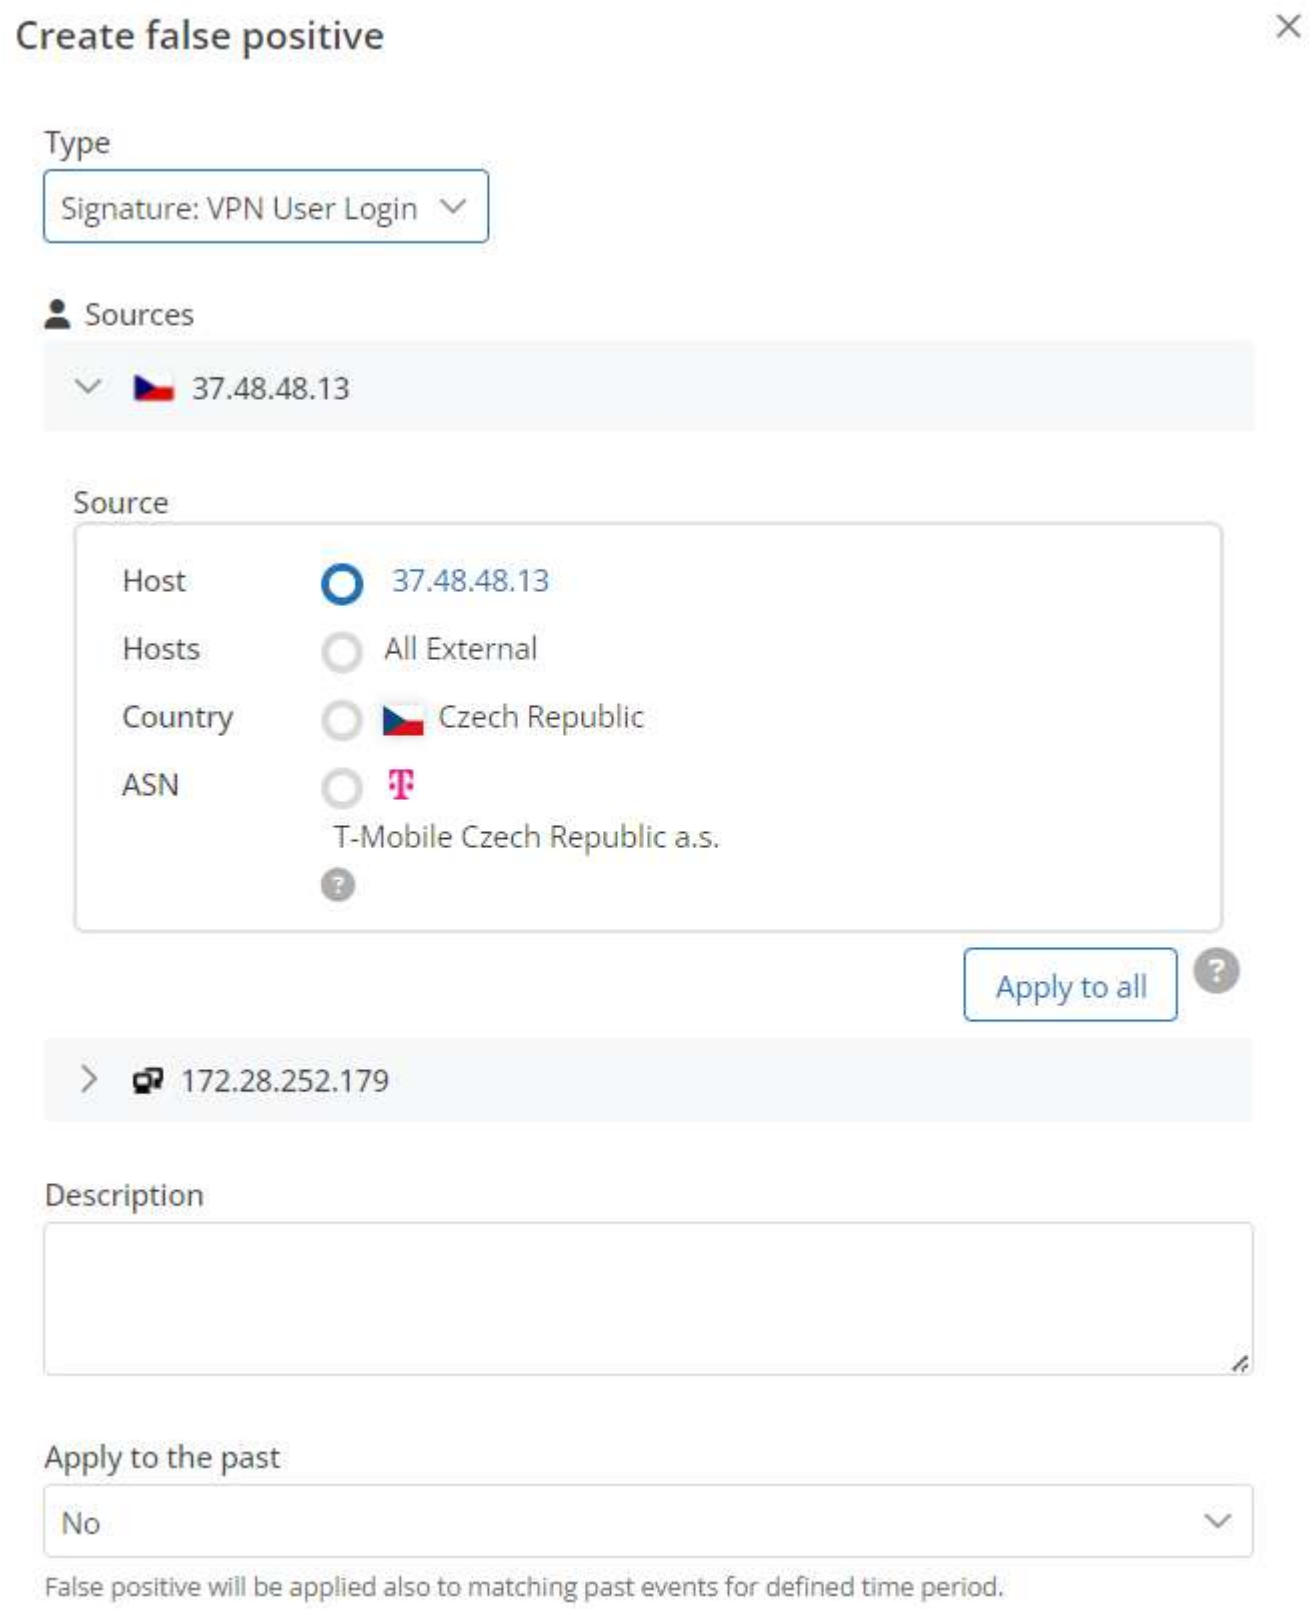
There are two kinds of false positives: system and user. System false positives are defined by GREYCORTEX and distributed from the update server on the update time. These false positives can't be removed or edited but can be enabled or disabled by a user. User false positives are created by a user.

DEFINING A FALSE POSITIVE/IGNORED EVENT

The creation of a "false positive" (i.e., misquoted reports) in detection can be applied to the basic and advanced details of the event. In the basic details of the event, you can identify the sources of the false positive (IP, subnet, service) and create a mass false positive corresponding to one of these sources.



Copyright © 2024 Greycortex



In the advanced detail view, a false positive can be created for a specific event only. The false positive rule, once created, always refers to a specific selection of events (or attributes) and will have a unique identification number. The false positive form contains the following parameters that can specify the behavior of the entered a false positive rule:

Copyright © 2024 Greycortex

Type Type of event. It can be a specific event, component, class, or source. Source and Destination Source and destination of communications, which can be defined as:

Attribute	Description
	Type of event. It can be a specific event, component, class, or source.
	Source and destination of communications, which can be defined as: • The IP address of a host • Hostname • MAC address • Host tag • Port number • Country • Autonomous system number • Subnet • Internal or external network • Buttons: • - chosen attributes will replace that attribute for all the hosts in that specific direction (source or destination) There is a feature that you can change the hostname to a specific address. So it will be applied to all hosts in a specific direction. E.g. *.google.com • - It will delete a specific host from the false positive dialog.
	Commentary describing false positive.
	False positives can be applied to match past events for the defined time period as well.
	The definition of the sensor for which false positive is valid. The false positive will be applied only to data from the respective sensor.
	The VLAN number for which false positive is valid.
	The type of protocol.
	The IP protocol family like IPv4, IPv6, IP, etc.
	Local or remote service.
	Decreasing or increasing the priority, i.e. changing the value of severity by the appropriate value. E.g. the current event has a severity of 7, when the priority - 1 is set, the false positive is detected by the weighting of the severity at 6.

Copyright © 2024 Greycortex

3.7 Incidents

3.7.1 Incident Management



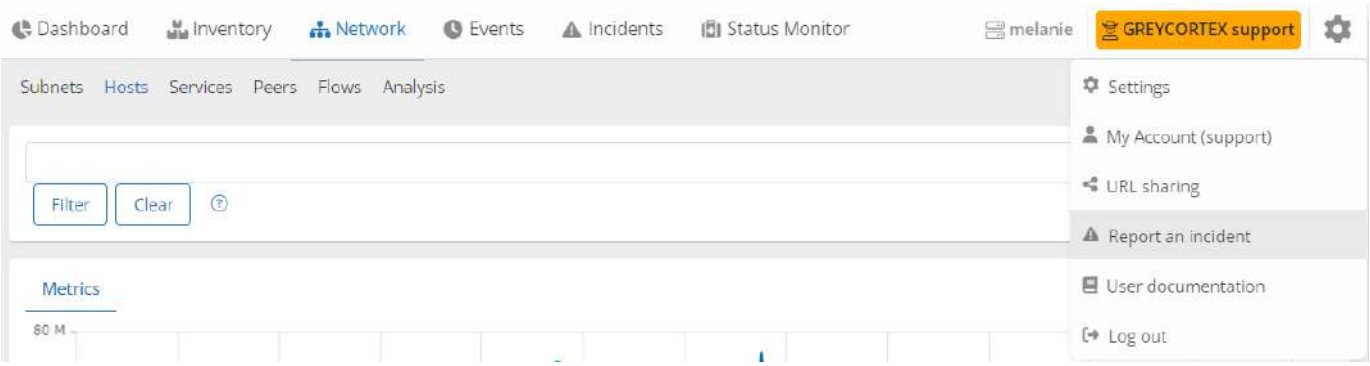
The **Incident Management** functionality allows MENDEL users **to report** events and other network anomalies for further analysis and gather them in the **Incidents**.
Reported incidents can be filtered, sorted, discussed and exported.

Reporting incidents

Reporting is possible over all items available under the **Network** and **Events** tabs.

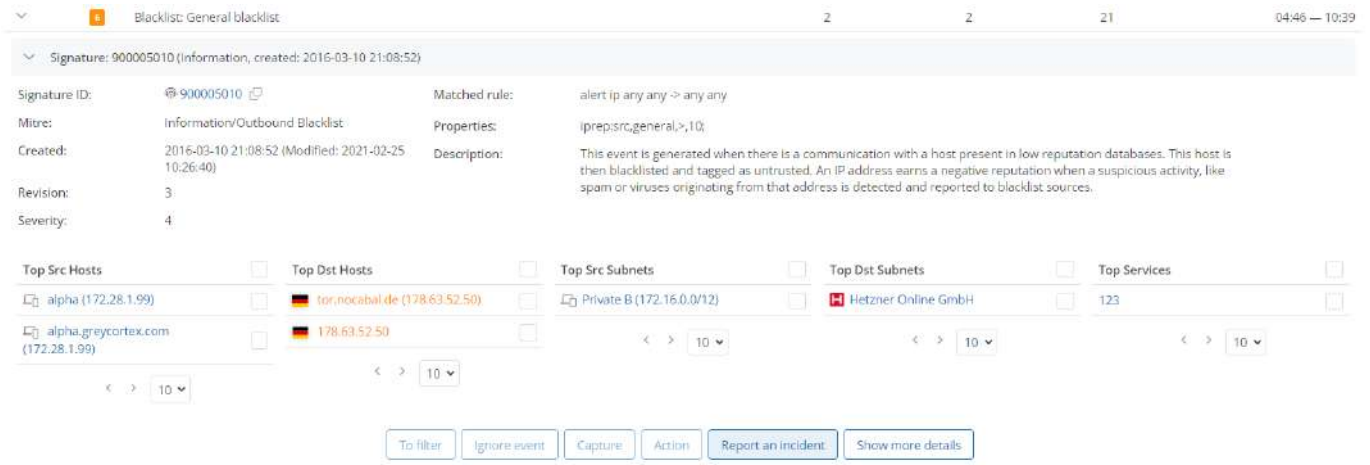
Specific event or anomaly can be reported as Incident in two ways: - by clicking button (available both in the Network and in the Event tab) - Report Incident button (available only in the event tab)

OPTION 1. REPORT INCIDENT IN SETTINGS MENU



OPTION 2. REPORT INCIDENT BUTTON

Use the button in Event details after opening the Event row in the Event Tab.



Copyright © 2024 Greycortex

Creating Incident Report

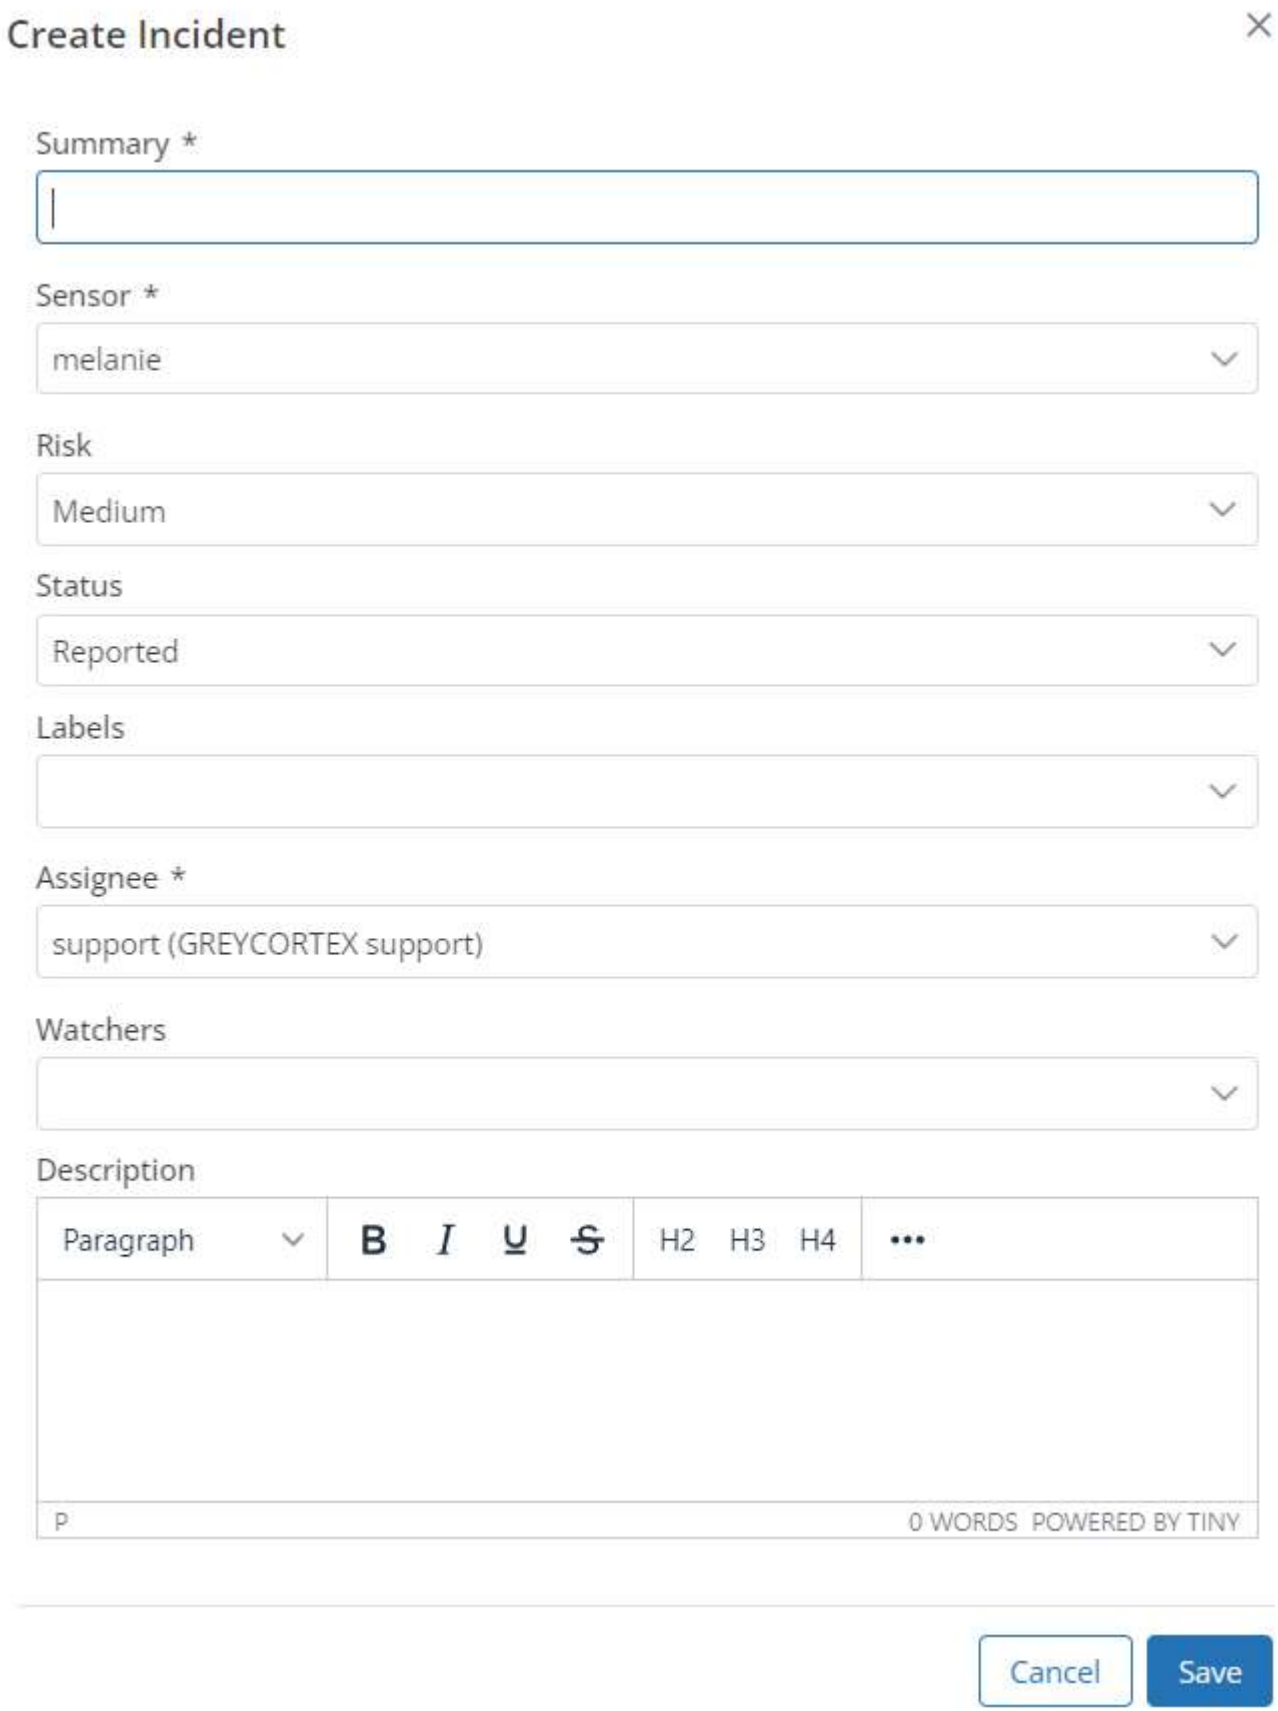
The next step in the Incident reporting process is to complete and submit an **Incident report form**.

Attribute	Description
	Text describing the incident.
	The ID of sensor where incident was reported.
	Incidents are ranked according to seriousness. The range goes from Info (no importance) to Critical (highest importance).
	State of the specific incident. It can be: Reported, Analyze or Resolved.
	Categorize incidents with tags.
	A person who is responsible for resolving the incident.
	A person who is interested in how the incident is handled.
	Description of the incident by attaching any important and related information.

Hyperlink reffering to reported Incident is automatically attached to the generated report.

***Important**: If you label reported incident as a False positive, it will be excluded from the Incident inventory list by default. If you want to show them, you have to set Label filter - False Positive.*

Copyright © 2024 Greycortex



By submiting the completed information, the **Incident report** will be **sent** and **stored** in the **Incidents**.

Copyright © 2024 Greycortex

3.7.2 Incidents



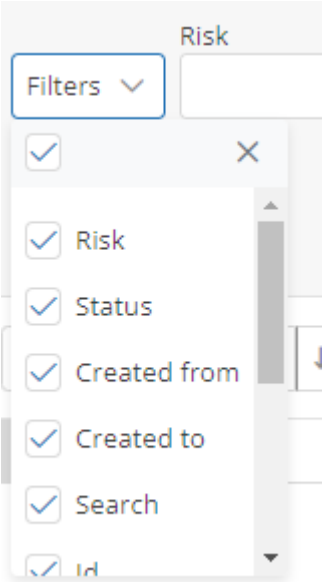
Used to manage workflow over reported incidents with filtering and editing tools.

Filtering Options

You can filter incidents by the options below.



Managing what filtering options are displayed.



Incidents - Inventory and Details

List of all reported (filtred) incidents, with details displayed when you click one of them.

Incidents list
An overview of reported (filtered) incidents with a graphical resolution of risk level and sorting options.

Incident detail
Detailed information about the selected Incident with the possibility of editing individual items and adding comments.

Incidents - Editing and other actions

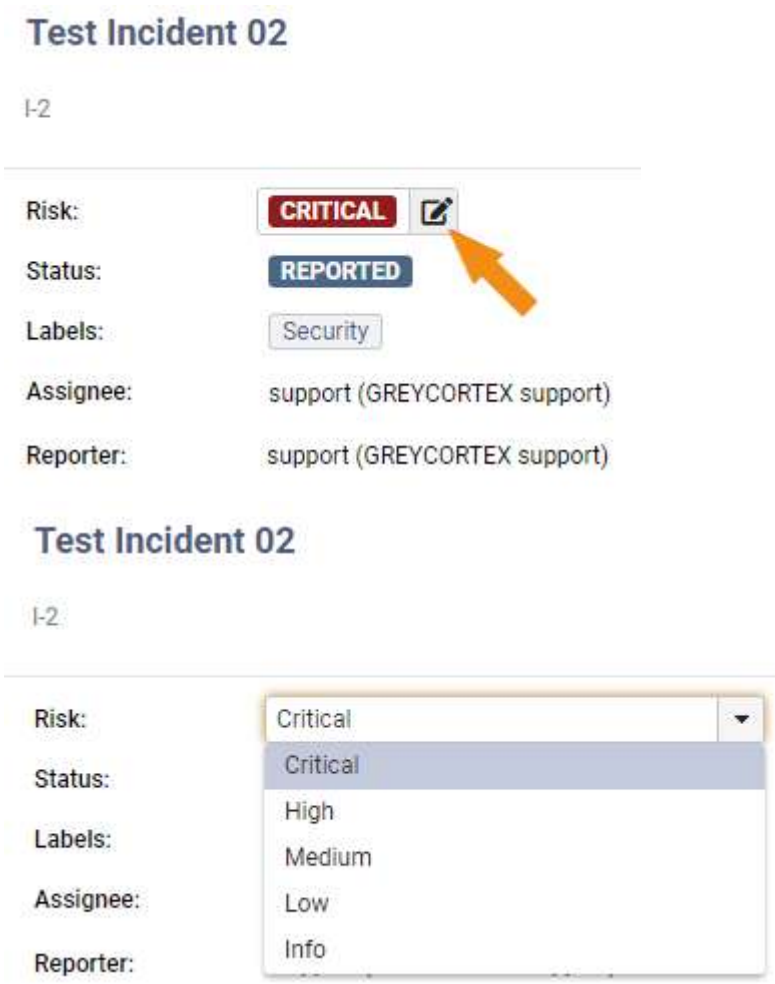
User with sufficient privileges, can: - edit Name, Risk, Status, Labels, Assignee, Watchers, Description - add a comment and start watching this incident - remove the incident *(available only in "options" dropodown menu)* - sharing URL *(available only in "options" dropodown menu)*

This can be achieved in two ways:

OPTION 1. QUICK EDIT (INPLACE)

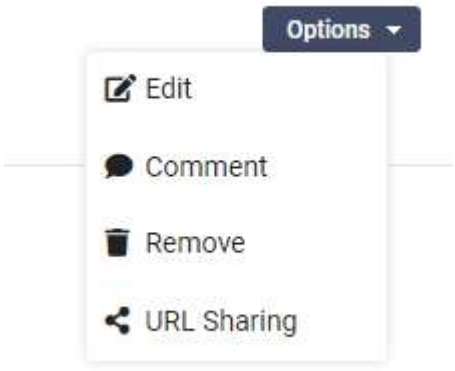
First method is by clicking **directly** on the detailed information about selected Incident.

Copyright © 2024 Greycortex



OPTION 2. FORM EDIT

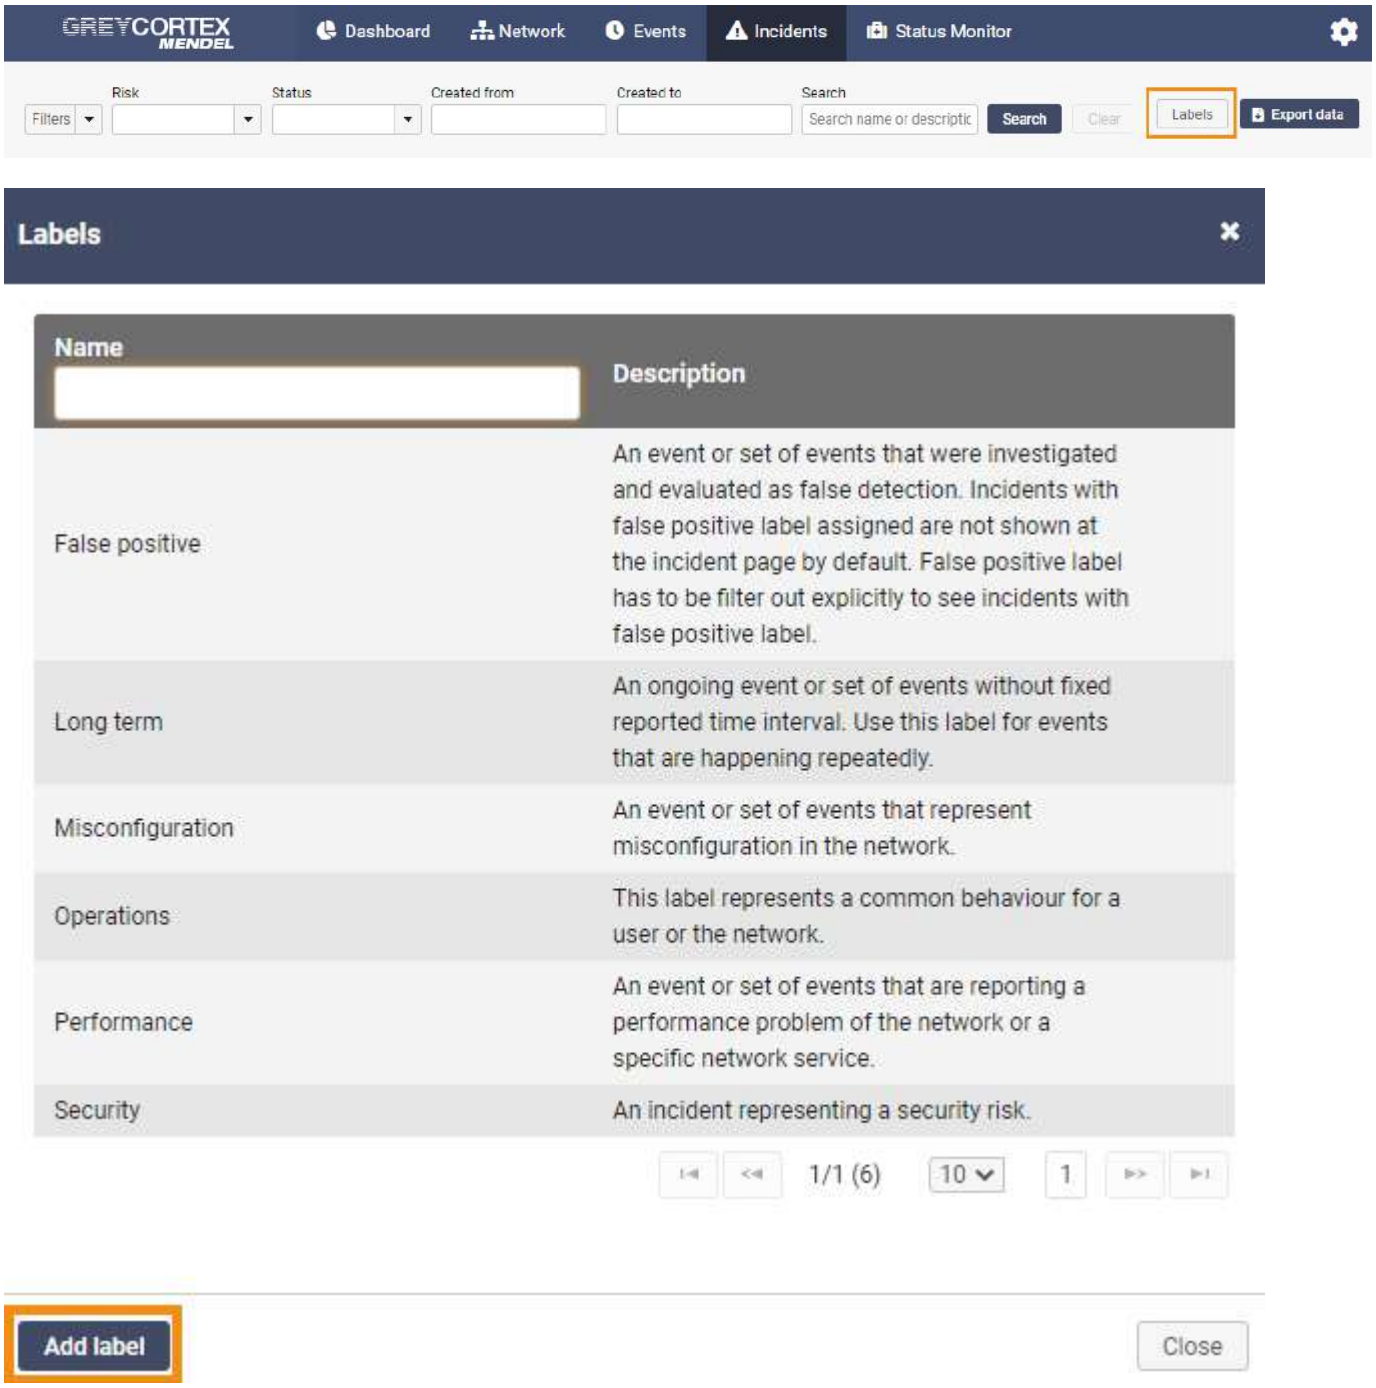
Second method is by clicking dropdown menu and: - editing Incident details in new window through the form - leaving new comment - removing Incident - sharing hyperlink to the Incident



Label management

With this funcionality you can add and managed your own labels.
Note: *Some labels are default and system predefined (False positive, Long term, Misconfiguration, Operations, Performance, Security).*

Copyright © 2024 Greycortex



Incidents - Description and Comments

Essential tools for more efficient incident management.

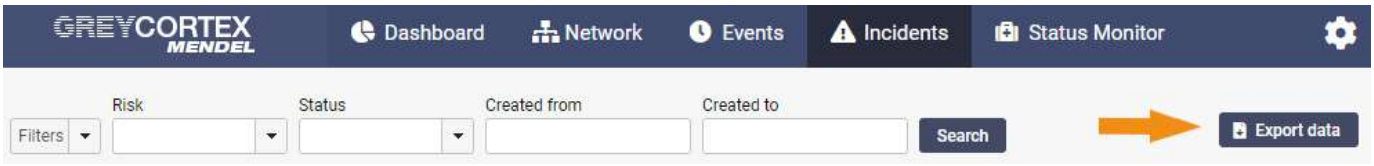
Description
The incident description field is used to enter all important information that could have an impact on its resolution and also serves as an overview of all this related information in one place.

Comments
Are used to insert additional information, mainly related to the incident resolution process. In addition, individual comments contain information about the author, date and time of insertion.

Copyright © 2024 Greycortex

Export

This feature allows you to create data export from the displayed (filtered) incidents.



INCIDENTS CAN BE EXPORTED INTO CSV AND PDF FORMAT.



CSV export
Exported data in a format suitable for further processing with sorting into columns.
ID, Sensor, Summary, Risk, Status, Labels, Assignee, Reporter, Created, Updated, Resolved, Link

PDF export
Understandable and organized output, suitable for visual presentation. Processed into one file, which contains an overview of all incidents in the report, as well as detailed information and description for each of them separately.

Incidents Email Notification

The Incidents Email Notification functionality provides email notification for the Assignee and the Watchers when one of the following actions occur:

- Creation of a new incident
- Deleting an incident
- Commenting on an incident
- Editing a comment
- Deleting a comment
- Editing the name of an incident
- Editing the risk of an incident
- Editing an assignee of an incident
- Editing the state of an incident

Incident Email notification is enabled by default along with enabling the Incidents component in the system components settings. Any user commenting on an incident will automatically become a Watcher for the specific incident.

For correct functionality, **correct Email settings are required**. If email settings are not done correctly, then the status monitor event **"Incident Email Error"** will appear in the Events tab when the notification email should be sent.

Copyright © 2024 Greycortex

3.8 Settings

3.8.1 Introduction

Introduction

Settings are used to set individual application components. In this section, you will use the second icon from the left, located in the upper-right corner of the application. The first screen displays the basic application information. The setting consists of five parts corresponding to the basic tabs:

Attribute	Description
	Setting up monitored networks and policies.
	Setting detection mechanisms.
	Set basic application exports.
	Setting basic system properties.
	Setting authentication mechanisms, users, and their rights.

BASIC APPLICATION INFORMATION

An overview of the following fields is displayed here:

- License in use
- License expiration date
- Current version and its build number
- Data version (last data update from update server)
- Sensor name and its IP address

Copyright © 2024 Greycortex

3.8.2 Network

Policies

POLICIES

Policies are used to set up communication rules for monitored networks. This tracks the presence and downtime of devices and network services and communicates with remote services.

You can use a predefined policy or create your own set of policies. The policy can then be applied to the selected subnet or device on the network.

Policy Definition

A policy is defined by a set of rules. Each policy contains:

Attribute	Description
	The icon displayed on a network with a user-defined policy. The icon appears on the subnet and all the devices on that subnet.
	The user-defined policy name.
	The parameter defines the priority of events recorded in the application policy environment. The priority value is added to or subtracted from the common severity of the identified event. E.g. Priority +2 turns an event with a severity of 7 into an event with a severity of 9. It is useful to use the functionality for critical or non-critical monitored segments.
	A policy creating a rule for reporting events about newly connected hosts and services. If a new device or service is discovered on this subnet, the event will be identified.
	A rule defining a policy for reporting events about the unavailability of existing devices or services on the network, i.e. If a device or service is unavailable on this subnet, the event will be reported. With this rule, it is also necessary to set the period of inactivity of the device -- the Lease Time parameter -- after which time the event should occur.
	A parameter defining the time of inactivity of the device after which an event is reported. An interval between 1 minute and 8 days can be set.
	A description of the displayed policy.
	Defines allowed or disabled communication ports. Detection of unauthorized activities will be reported by reporting tools. Ports are defined by a port number and are separated by a comma (e.g. 445,80,67,443).
	Defines value limits for specific metrics to fine tune the policy settings

Definition of Packet Flow Communication Paths

Attribute	Description
	Definition of communication services coming from internal networks.
	Definition of communication services coming from external networks.
	Definition of communication services leaving for internal networks.
	Definition of communication services leaving for external networks.

Definition of Ports Allowance

Attribute	Description
	Definition of allowed communication services further defined by Packet Flow Communication Paths.
	Definition of prohibited communication services further defined by Packet Flow Communication Paths.

Copyright © 2024 Greycortex

Definition of Communication Directions

Attribute	Description
	The direction of network communication defined by the policy.
	The direction of network communication defined by the policy.
	The monitored part of the network solution.
	The unsupervised part of the network solution.

Pre-defined Policies

We prepared basic policies for subnets. Network should be separated to the subnets that have a different level. In MENDEL you can set up another policy to each subnets. You can describe it, set up ports for communication or limits. We pre-defined policies based on usual network.

Attribute	Description
	It separates LAN from Internet. It should be used for external-facing servers, resources and services (Proxy, DNS, FTP, VoIP, ...). They can be accessible from the Internet but the rest of LAN is unreachable.
	Devices from the subnet with this policy are not accessible from the Internet. Machines in the subnet get an address from private address space. Devices from that subnet can communicate just with each others.
	Devices from the subnet with this policy are not accessible from the Internet. Admin has to allocate a specific IP address from private address space for the machines. Devices from that subnet can communicate just with each others.
	Devices from this subnet can access the Internet via NAT. Machines get IP address by DHCP
	Devices from this subnet can access the Internet via NAT. Admin has to allocate a specific IP address for the machines.

Copyright © 2024 Greycortex

Subnets

In this part of the ME settings, the monitored subnets are defined and managed. The system compute behavioral models for each of the defined subnets. For this reason, it is appropriate to specify each functional sub-network individually depending on its functionality.

Subnets can be partitionate to:

users | printers | classrooms | VoIP | Wi-Fi | Servers | DMZ | Etc.

Subdivision of the subnet brings advantages for each defined subnet:

- Filter according to each partition
- A model for detecting behavioral anomalies is created
- Separate detection policies can be set

Adding a Subnet

A. Automatic Subnet Discovery

System automatically recognizes unnamed subnets and creates generic entry in the subnet list.
with tag.

This fucntion is activated/deactivated in

B. Import from a file

It is possible to add multiple number of subnets through file import (csv) by clicking on the button and following the wizard.

Import function supports both types of separators and CSV file needs to follow this format:

Subnets can be inserted with existing **tags** with reference ID. New column with single or multiple tags can be added to csv file. Tags in the file has to be in format as Zone/administrators describing whole path to the tag.

C. Add individually per subnet

Click on the button and follow the wizard.

Attribute	Description
Sensor	List of sensors for which it will be added
Address	Subnet address
Name	Subnet name.
Policy	Application of a defined policy
Description	Further description of the subnet
Subnet Calculator	Subnet validator and calculation its boundaries

Copyright © 2024 Greycortex

A subnet is defined by 3 parameters:

- Subnet address in 192.168.1.0/24 format
- Network name or short description
- Subnet network level (LEVEL), value from 1 to 9 *(to determine, please see article below)*

How to determine subnet level

This is very important to set it up right. When the network is more specific (divided to more levels), the detectors will work better and precisely. User can assign a value 1 to 9 as a level of the subnet. The default and most precise level is 1, and it can contain a maximum of 2000 subnets but no other nested subnets. When a user creates a new subnet, the level should be higher than 1; by that application knows there is another network on the smaller level that is nested. Therefore, subnets from the higher level covers a smaller levels; so there is more space in the first level.

The best way to split subnets to different levels is starting with the most specific one. There are 2 subnets - 10.1.0.0/16 and 10.0.0.0/8. You can see that the second one is not so specific as the first one. That means 10.1.0.0/16 can be nested to the 10.0.0.0/8. So 10.1.0.0/16 is level 2 and 10.0.0.0/8 is level 1.

Tip

Setting the subnet level is an important internal function for running the application. By assigning values 1-9 to each subnet, the administrator defines an application that the subnet protects.

The default value is 1 and defines a partial subnet as a simple one. A simple subnet of level 1 can contain any number of devices, but no other (nested) subnets. By setting a level above 1, the application automatically says that the subnet contains additional (nested) subnets. The application will always assign a device to the lowest-level subnet. For this reason, the configurator will not allow overlapping subnets of the same level to be set. In network, there can be added maximum 1000 subnets on level 1. If there are more subnets it's necessary to choose the right level for nesting.

Example:

- If you have a 'nested' subnet in the network, e.g. 10.1.0.0/16 in a simple 10.0.0.0/8 (Private A) subnet, you need to set the subnet level 10.0.0.0/8 to 2. Then add Subnet 10.1.0.0/16 with level 1.
- If you have another 'nested' subnet, e.g. 10.1.2.0/24, you have to set the subnet level of 10.0.0.0/8 to at least 3 and the subnet 10.1.0.0/16 to level 2 before adding the subnet 10.1.2.0/24. You can also set a gap between levels, e.g. 1, 3, 7.
- If you do not want to protect the entire subnet 10.0.0.0/8 on the selected sensor, you can delete it. The default top-level subnet is 0.0.0.0.0, which also matches IPv6's:/0, and represents the entire protected network. (This also involves connecting multiple sensors/collectors in a distributed network). If you do not have other "nested" subnets, but you want to receive a report on the existence of other subnetting in the network (or IP 'spoofing'), you can set the level to 10 or more.

Note

Changing subnet settings or levels will result in the loss of learned data in changed subnets. Therefore, it is best to prepare all the subnet settings together after installation.

Important

Create subnets is essential for proper detection and network visibility. So please do this really carefully. All subnets can be defined for specific sensor. You can set up the pre-defined policy or create a new one. If you want to have better protection, there is need to define subnets policies. In policies you can define ports for communication (outbound, inbound) or limits per hosts.

Each protected network can be configured in the application by specifying a pre-defined and stored policy, or directly configuring settings similar to policy settings -- see Policies.

Copyright © 2024 Greycortex

Hosts

HOSTS

The Host tab is used to manage devices in monitored subnets. MENDEL models the behavior of each device on the monitored network and each of its services. Devices cannot be added or removed. This is information obtained from automatically analyzed network traffic.

Definition of Host

Attribute	Description
	The sensor on which the device was found.
	The subnet to which the device belongs.
	Icon that is inherited from the settings of the network in which the device is located (The icon may be changed if necessary).
	IP address.
	Can be obtained from DNS or DHCP and in the case of the too long hostname, it is cropped so only the last 64 characters are left. Can be edited by the user.
	The time that the host was firstly discovered in the network. It is displayed only if the host is from the internal subnet.
	Applied policy.
	Monitoring the emergence of new services.
	Defines device priority as well as subnet setting.
	Automatically identified operating system.

In detail of each device is listed below:

Attribute	Description
	Name of the host, can be derived from the Hostname or edited by the user.
	Entered by the user.
	Allowed/denied ports in a particular direction.
	List of MAC addresses with which this IP address was detected. It is possible to add new MAC addresses to optimize the Discovery detector.
	Tags of the host and their state. Here it is possible to add new, modify the state of existing or delete unwanted tags. They identify what the host is used for, and how the detectors behave. For example, you can set up Proxy, Mail, DNS, DC, WEB, SSH, etc.
	List of users who use or used this IP address.
	Limits that, if exceeded, trigger an event.
	Button to shows services on the host.
	Defines device priority as well as subnet setting.
	Automatically identified operating system.

Setting a proxy visibility can be found in the Proxy Visibility.

Copyright © 2024 Greycortex

Services

SERVICES

The tab is for service management (service name / protocol displayed on the network port of all devices in the network) including the appropriate ports. Each service group is represented by a unique name and contains one or more numbers indicating a port (80,443,8080) and a string specifying the communication protocol (TCP, UDP, etc.).

Each service contains the following information:

Attribute	Description
	Optionally user-replaceable or user-modifiable.
	Predefined service identifier (name).
	List that defines the ports used by the service. You can edit this list and add a new item to it.
	Defines if it is a Port or ICMP Code.
	With reference to an information resource (e.g. wikipedia.org).
	The Internet protocol defines the TCP / UDP protocols used by the service.
	Profile settings can setup either parsing or detecting Layer 7 Protocols. Selecting "PARSE" shows the services which are set up to parse Layer 7 application layer metadata by the service in the selected row. For "DETECT", the service which can be detected by DPI is shown. If you don’t want to use this service, you can set the Profile to “Disabled”.
	A list of selected most commonly used services displayed at the beginning of the filtering lists.

Copyright © 2024 Greycortex

3.8.3 Detection

Intrusion Detection System

SIGNATURE CLASS

Reffered to as category in IDS. Signatures are divided into six categories/classes.

Attribute	Description
	A suspicious, miscellaneous, unknown or potentially bad traffic.
	Attempted or successful attacks.
	Generic informative signatures like ICMP event, blacklists, etc..
	Trojans, malicious files, exploits, and malware.
	Inappropriate content detection, information leaks, privacy violations and sensitive data transmissions.
	Network scan, potential spoofing or decode of an RPC query.

IDS

The main task of the IDS engine is to analyze packets in the network, their reconstructed content and related metadata to look for patterns that match their database similarly to an antivirus engine. In contrast to the NB analysis, it doesn't use any intelligent or adaptive rules, but a strictly defined set of signatures. In this way, Mendel's IDS system can detect a wide range of known malware, attacks or corporate policy violations, which behave "normally" in the network.

For a better understanding of the network, you can define IDS signatures to reveal what is normal in your network such as usage of servers, response time or SLA limits. In MENDEL there are pre-defined IDS detection rules, you can find them in . The administrator can activate/deactivate, modify existing, or add new ones.

The rule applies to network data corresponding to the protocol, source, and target IP addresses and ports, and further depends on other parameters contained in the network flow (defined in Properties). The system responds according to the action defined by the rule. Individual rules are contained in rows that contain the attributes of sorting and filtering:

Attribute	Description
	The name of the rule, the type of rule is given as the name prefix (see Source). The rule name follows the colon.
	Type of rule. (files/dos/web server/exploit/etc.)
	The rule is included in one of the signature classes. These classes are described in the Detection section.
	Action to be taken if the rule is applied: • alert - An event is created (event) • allow - No action will be taken
	Rule set based on the license.
	Event severity -- the number on scale 1-10, where 10 is the most serious event.
	The severity of the event if it is an Inbound Network Flow.
	The severity of the event if it is an Outbound Network Flow.
	Rule is *active/inactive* by default.

Copyright © 2024 Greycortex

Details of the IDS signature:

Attribute	Description
	Event definition, protocol, source IP addresses, ports, and destination IPs, ports.
	Definition of the rule with the direction crossing the boundary between private and public network.
	Links to information sources on a given rule.
	For the System IDS signatures, there is a possibility to enable or disable them for all or just some sensors through Edit option. First, the sensor must be chosen, then the possibility to enable is displayed. This information then can be only seen in the details of the rule and not in the Enabled column. Other modifications for the predefined rules are not possible.

Group operations are used for enabling or disabling multiple rules at once. First, the table has to be filtered according to the rules we want the group operations to be applied to. Then the sensor must be chosen. The operation can be , or . The last option is used when resetting to the default state is required.

Execute button is used for confirmation of the group operations.

Create new custom signature

You can create your own IDS signature for accurate detection or label some specific communication as an application. User IDS Signatures can be added in the form of the Full Rule String (Raw Sytax) or step-by-step with the help of the IDS Rule Wizard. These rules can be fully modified. It can be created in in this format:

Raw Syntax

Attribute	Description
	Defines what happens when the rule is matched.
	Considered protocol.
	Defines the source and destination that the rule should match.
	Text of the rule that will show as title of the event.
	Flow keyword that consists of the action taken and the name of flowbit. • flowbits:noalert -> Does not generate event if the rule is matched. • flowbits:set -> Generate the event when the rule is matched. • a_ -> Name of the application, it will be visible as the parameter of the flows. Filter can be applied on this value.
	Mandatory numeric value.
	Classification of the rule.
	Unique signature ID of the rule.
	Revision, represents the version of the rule.
	The first part of the content is the part of the rule that you would like to be matched. The content modifier that follows says what buffer should match only the requested URI buffer.

Rule Wizard

Step-by-step way how to easily create a rule.

Copyright © 2024 Greycortex

First step is to define major content as definition of the (action, protocol, addresses, ports and direction) and (ID, priority, revision number, rule message and protocol specific options).

Second step is to change manually the rule syntax if needed with the option to verify all your changes.

Third step gives the liberty to add , define and apply .

> **Info**
>
> GREYCORTEX Support, can assist you with creation of these rules.

Edit Signature

You can edit your signature or created by others based on your permissions. If you want to edit the rule, go to and find that rule in the list. Then click on the pencil. It will open the dialog where all of your values are paired and you can comfortably edit all the parameters.

Import/Export Custom Signature(s)

You can easily import or export more rules in a one bulk operation.

Applications

If you need better performance statistics or choice of filtering based on applications, you can create one. You can label basic communication in your network. You can label communication with your servers to have better statistics of usage. When you define the application, you can easily use the filter.

To define your own application, go to and . Syntax is the same as in case of IDS signature, you just have to define the application with prefix **a_** after the a_ there should be a name of that specific application. See the example below:



Services

You have an option to mark specific communication per service. It is really useful in case that you communicate through different sensor. You can define your own service that you can label the flows. You can use defined services or define your own service that is really useful in IoT networks. You can filter based on your new service. For mark the communication, you need to do signature with prefix **p_** after the p_ there should be name of the specific service that you want to see in flows. You will see that the communication in flows is marked as that new service.

Example of defining Service



APM

Application performance monitoring (APM) allows a user to see detailed metrics like **NTT** and **ART** of individual transactions in the specific flow. It can be very useful for monitoring performance of your servers or different problems in the network. You can see the whole picture of your network in **Analysis** section and find the performance there. You can see information about APM in . Expand the row of the table by the plus sign. All information about APM is in requests and responses on the flows.

Copyright © 2024 Greycortex

Request	Response
uri: / **host**: www.google.com **method**: HEAD **protocol**: HTTP/1.1 **userAgent**: python-requests/2.12.1 **ntt**: 136	**status**: 302 **contentType**: text/html **protocol**: HTTP/1.1 **location**: http://www.google.cz/?gfe_rd=cr&dcr=0&ei=O4hxWrOdHYeO8Qe1kYGQDg **ntt**: 85 **art**: 19
uri: / **host**: www.google.com **method**: HEAD **protocol**: HTTP/1.1 **userAgent**: python-requests/2.12.1 **ntt**: 104	**status**: 302 **contentType**: text/html **protocol**: HTTP/1.1 **location**: http://www.google.cz/?gfe_rd=cr&dcr=0&ei=O4hxWrz4H5SO8QfejoGgBw **ntt**: 232 **art**: 10

Measure the performance

You can add metrics to the APM. So you can easily know if your server is overloaded or your SLA limits are not right. You can find out anomalies of your appliance in the network. You can specify it in the IDS signature in the format:

Example for ART

Example for NTT

Analysis

This module is great when you want to know the whole picture of the defined application. Here you see application, request count, average, maximum and minimum time that the request takes. You can select APM metrics as *Server NTT*, *Client NTT* or *ART*. So here you can see if your SLA is as it should be and the other side is not breaking the agreement. For more information see Network Analysis Module.

Copyright © 2024 Greycortex

Network Behavior Analysis

NBA

The tab is used to set NBA detection mechanisms. Here is a list of all basic NBA detection rules. Predefined rules are based on complex algorithms. These rules cannot be deleted or edited; the administrator can only influence their functionality by modifying some parameters labeled as Key ID.

The list of NBA rules can be filtered and/or sorted by any column (parameter).

Each NBA rule contains:

Attribute	Description
	Opens a rule description, a list of subordinate rules, and modifiable parameters of the algorithm. (described in more detail below)
	Unambiguous numeric code identifier • 0 - 999: reserved for internal system reports, displayed only to users with support privileges. • 1000 or higher: NBA rules for event evaluation in network traffic. • Each detection mechanism has the same first two digits. The following numbers designate \"stained\" methods based on a given set of detection algorithms. (rules)
	Rule name.
	Derived rule identifiers that form a subset of the rule.
	Event category.
	Setting the severity of an event in the range 1-10, where 1 is the minimum severity and 10 is the most serious.
	Serves to deactivate the rule, the rule is active/inactive.

Details of the NBA rule:

Attribute	Description
	Date of creation or modification of the rule.
	Detailed description of the rule.
	Recommendations on how to proceed in the event of an event based on this rule (text description).
	List of rules.
	ID key of a rule.

Table Color Rules

In some cases, it is desirable to distinguish the events detected in more detail according to one rule, the network service, and/or traffic direction. For these purposes serves colors. The ID of derived rules usually begins with the same number.

Parameters of Detection Algorithms

The administrator can modify the value of some of the key parameters of the algorithm and thereby change the sensitivity of the rule. Optional parameters are specific to each rule.

Copyright © 2024 Greycortex

Signature Class

Signatures are divided into six categories/classes. These classes are used for IDS and EC signatures as well.

Attribute	Description
	A suspicious, miscellaneous, unknown or potentially bad traffic.
	Attempted or successful attacks.
	Generic informative signatures like ICMP event, blacklists, etc.
	Trojans, malicious files, exploits, and malware.
	Inappropriate content detection, information leaks, privacy violations and sensitive data transmissions.
	Network scan, potential spoofing or decode of an RPC query.

Copyright © 2024 Greycortex

Universal Tagging Engine

UNIVERSAL TAGGING ENGINE (UNTE)

Dynamic visibility by tracking new activities or changes caused by devices communicating in your network. A completely new engine that brings a manual or automated way of tagging hosts or subnets through a system of user-defined rules with easy-to-understand syntax.

Tag is unique combination of network entity identifier (host_id, subnet_id), name and description object.

Examples of simple tags

host_id	Name	Meaning
melanie, 10.0.0.1	DNS Server	machine 10.0.0.1 on sensor melanie acts as DNS server
melanie, 10.0.0.2	Domain controller	machine 10.0.0.1 on sensor melanie acts as Domain controller
melanie, 10.0.0.3	My pink computer	manually assigned tag my pink computer to machine 10.0.0.3 on sensor melanie
melanie, 10.10.0.1	DHCP server	machine 10.10.0.1 on sensor melanie acts as Domain controller

UnTE signatures System tagging signatures

Pre-prepared rules by GREYCORTEX.

Every system signature can be affected by user in a way of:

- Enable/disable the signature processing over all data (all sensors)
- Enable/disable the signature processing for specific data (sensor)

User defined signatures

For more information how to add your own signature see Universal Tagging Engine Guide.

Tag Management

Copyright © 2024 Greycortex

Tag Management

Here you can create and manage the Tags use mainly by the Universal Tagging and Correlation Engine (UnTE). Tags are defined by a tree structure.

TAG CREATION

Simply use the button to create a new tag. At least the name of the tag is obligatory, but if you intend to use the tag in UnTE rules (signatures) you need to define as well.

Attribute	Description
	Name of the tag.
	Unique identifier used in UnTE rules (signatures).
	Affects the severity of an event. The default (0) value has no affect.
	Definition of the tag tree structure. Choose to create a parent.
	The sensors for which the tag is valid. If empty, there is no restriction.
	Priority configuration for specific sensor only.

For the enhanced tag usage check:

- Universal Tagging Engine
- Universal Tagging Engine Guide

Copyright © 2024 Greycortex

Log Processing

LOG PROCESSING

Feature with ability to process received logs and generate security events from them by semi-passive approach. Logs are received by Mendel on specific port (514, Syslog).

Feature activation

Log processing is disabled by default. You can activate/deactivate this feature in

Upon activation, the Mendel is listening on port 514 (Syslog).

Signature Management System signatures

Attribute	Description
	Identification number of the rule.
	The name of the rule, the type of rule is given as the name prefix (see Source). The rule name follows the colon.
	Type of rule. (files/dos/web server/exploit/etc.)
	The rule is included in one of the signature classes. These classes are described in the Detection section.
	Action to be taken if the rule is applied: • alert - An event is created (event) • allow - No action will be taken
	Rule set based on the license.
	Event severity -- the number on scale 1-10, where 10 is the most serious event.
	The severity of the event if it is an Inbound Network Flow.
	The severity of the event if it is an Outbound Network Flow.
	Rule is *active/inactive* by default.

Custom signatures

You can create your own Log processing signature for accurate detection. Log processing signatures can be added in the form of the Full Rule String. These rules can be fully modified. It can be created in in this format:

Copyright © 2024 Greycortex

Attribute	Description
	This informs Mendel how to flag the event. Valid options are , or . When using the pass option and the signatures conditions are met, no other signatures are processed.
	Valid options for this field are , , or . In most cases, you will likely want to specify any. The protocal is determined by the , or liblognorm rule keywords.
	This informs Mendel where the source IP address or addresses must be coming from in order to trigger. By default the variable is used. This filed is populated by whatever the source IP address within the log might be. For example, if the signature lacks or , then the syslog source is adopted. If or rule option is used, then data (if any) that is extracted from the log is used.
	The next is the source port. If the or rule option is used, it will be applied here. This can be useful in filtering out certain subnets or syslog clients.
	his would be the direction. From the $EXTERNAL_NET $HOME_NETWORK.
	This works similarly to how functions. Rather than being the source of the traffic, this is the destination of the traffic. Also like the , network CIDR notation can be used (ie - 192.168.1.0). Data from this is populated by the and rule options.
	The final rule option is the destination port. If the normalize or default_dst_port rule option is used, it will be applied here. This can be useful in filtering out events from certain subnets.

Our Log processing engine has similar syntax as SAGAN (rule syntax, rule keywords).

GREYCORTEX Support can assist you with creation of these rules.

Edit Signature

You can edit your signature or created by others based on your permissions. If you want to edit the rule, go to and find that rule in the list. Then click on the pencil. It will open the dialog where all of your values are paired and you can comfortably edit all the parameters.

Import/Export Custom Signature(s)

You can easily import or export more rules in a one bulk operation.

Copyright © 2024 Greycortex

Event Corellation

EVENT CORRELATION SETTINGS

The Event correlation settings allow you to enable or disable existing event correlation signatures. They consist of specific IDS or NBA rules. The Event correlation engine is currently in beta development, which means users cannot edit existing event correlation signatures or add new ones at this time.

Attribute	Description
	Unique signature identifier.
	Signature name always starts with the correlation prefix.
	Describes types of events the correlation is focused on. The rule is included in one of the signature classes. These classes are described in the NBA section of the manual.
	Event correlation signature description.
	Recommendations when the situation occurs.
	A condition which looks for specific filter results in the time interval that is a multiplication of the signature interval.
	Sets the severity of the correlation event using a number on 1-10. (1 -- OK, 10 -- Critical)
	Sets the severity of the correlation event if it is in the inbound network traffic.
	Sets the severity of the correlation event if it is in the outbound network traffic.
	A time setting governing how often the event correlation signature checks for the specific NBA or IDS signatures.
	Turns signatures on and off.

Copyright © 2024 Greycortex

False Positives

FALSE POSITIVES

The system allows the detection of false detection as false positive and its management. The tab is used to display a list of defined System 'false positive' event detection events with the option of enabling or disabling them.

Each record for the system false positive contains:

Attribute	Description
	The identification number of the false positive.
	Specify if the false positive is based on the component, signature, class or source.
	The name consists of the number and the name of the signature for which the false positive was created.
	The description of the false positive.
	The source IP address and/or Subnet to which the false positive relates.
	The target IP address or Subnet to which the false positive relates.
	The service to which the false positive relates.
	The date of the last time the false positive matched with the occurred event.
	False positive creation date.
	If checked, the false positive is active.
	The name of the component to which the false positive belongs. (NB -- NBA, ID -- IDS, SM -- Status Monitor)

For the User false positives, there are these details in addition to the mentioned above:

Attribute	Description
	Login authorization.
	Sensor, for which the false positive applies. If it is not filled, it relates to all Sensors.
	Editing false settings.
	The ability to delete a created rule.

Copyright © 2024 Greycortex

Threat Intelligence

Use built-in Threat Innteligence which comes with the product and is divided into three categories (, ,) or enhance the detection with the Threat Intelligence you will insert by yourself.

BLACKLISTS

System Blacklists lists all blacklisted IP addresses and ranges. It is updated several times a day along with the IDS rules. On average, it contains 30,000 -- 100,000 entries.

In User Blacklists, a user can add their own blacklist containing the IP address or a range of IP addresses and reputations. Importing blacklist from the .csv file is also available.

Blacklist Contains:

Attribute	Description
	an Internet Protocol address is a numerical label assigned to device connected to a computer network that uses the Internet Protocol for communication.
	IPv4 or IPv6.
	Ranging from: • 1 - low criticality • to values 3 - high criticality i.e.: If an event is recorded with an appropriate IP address from the blacklist, the reputation will increase its severity and will be marked in the respective color in the GUI.
	description of the source from which the information is based.

Checking blacklisted IP addresses is done at the IDS level. If a user wants to add their own address to the blacklist, it is necessary to create a custom rule to include that address.

MALICIOUS FILES

A malware definition based on calculated hash differentiator.

MALICIOUS DOMAINS

Malicious domains refer to website addresses that are created and used for nefarious purposes, such as distributing malware, phishing attacks, or other forms of cybercrime.

Copyright © 2024 Greycortex

Data Capture

The Data Captures part is used to capture, store and download full communication based on a defined filter in .pcap files. The filter allows you to specify the communication based on the source and destination IP addresses, MAC addresses, and ports.

The bookmark lists the created rules with a filter for storing data. Each row contains:

Attribute	Description
	Filter name.
	User that created the data capture.
	Text description of the data capture.
	Applied sensor for capturing network data.
	Enable/disable logging of data.
	Full modification except for the name.
	Delete the rule.

Details of the data capture:

Attribute	Description
	Date of creation.
	Maximum size of the file in MB.
	Protocol type.
	Detection port.
	Source IP address.
	Source port.
	Source MAC address.
	Destination IP address.
	Destination port.
	Destination MAC address.

Copyright © 2024 Greycortex

PCAP Analyser

PCAP playback and analysis is an extension to Mendel's core capabilities not only to analyse mirrored traffic in real-time but offers the same power to do it over earlier recorded data with the built-in PCAP recorder or captured by any other tool you have. Replayed PCAPs are processed and stored separately to any other data in Mendel to ensure the highest user experience during retrospective analysis with a maximum focus on data consistency and security.

PCAP UPLOAD

Hit the button to import a PCAP into the Mendels internal storage. Choose the specific file and upload it. You can find all the uploaded PCAPs in the tab.

PCAP REPLAY

Whenever you have the PCAP files listed in the tab you can start the analysis by replaying them. You can find the button under the menu on each row.

Two atributes need to be specified to replay the PCAP files.

Attribute	Description
	A unique name to distinguish the analysed data.
	Basic virtual sensor configuration will based on a chosen physical sensor.

You can find all replayed PCAPs in the tab.

PCAP DATA

Analyzed (replayed) data from the PCAP can be found by filtering the virtual sensor name in the main filter. You can also use the hyperlink in tab for each reaplayed PCAP.

Limitations

- Not all detection capabilities are available for PCAP playback and analysis, as some detection methods are based on a longer time period than is the time range of your PCAP file.
- Stored PCAP files are currently NOT covered by any backup scenario. You need to backup them on your own.

Copyright © 2024 Greycortex

Proxy Visibility

Proxy Visibility is a feature which allows MENDEL to detect traffic passing through a proxy server. This proxy visibility feature also provides the ability to pair flows from a source device to a proxy server, with flows from a proxy server to a destination device into a single flow.

- In the collector UI, select .
- Click the button to mark a network host as a Proxy server.

Attribute	Parameter
	User defined name
	Choose one of the available sensors
	IP address of the Proxy server
	Port number of the Proxy server

Copyright © 2024 Greycortex

3.8.4 Exports

Exports

EXPORT

EXPORTING

The information and notifications received from MENDEL can be sent to other devices and accounts using system logs and emails. The reports contain data and can be ordered arbitrarily. Each report uses a single data format and time period, but many data formats and time intervals are available.

Use the Reports tab to set up the reports and the time of reports. Use the Data tab to report and create alerts, or to continuously send logs to your system.

Copyright © 2024 Greycortex

Data

EXPORT DATA TO EXTERNAL SYSTEM

Detected security events and network flows can be sent either to the email of the users or sent as external log security systems like Security Incident and Event Management (SIEM) format, Common Event Format (CEF), IDEA system format, or IBM Qradar's Log Event Extended Format (LEEF). Data exporting allows users to export security events or network flows.

Export rules are defined by these items:

Attribute	Description
	The name of the collector from which the events are being sent
	Output type - email, email with the attachment, log
	Name of the export
	User name of the rule maker
	Enables or disables the export
	Export expiration. Empty means there is no expiration
	How often the export is being sent

For event-based reports, it is necessary to set it further in these fields:

Attribute	Description
	The filter applied to the output.
	The maximum number of reported events.
	Email addresses of message recipients.
	The time zone in which the report is created.
	IP address of the receiving server's mail server.
	Output file format. (.xls, .xml)
	Transmission type. (UDP, TCP, TCP with SSL)
	Destination port for logging.
	Specification of the syslog protocol.
	Check of the server certificate.
	Provide your certificate key if the server requires it.

For a network flows-based report, it is necessary to set these further fields:

Attribute	Description
	The filter applied to the output.
	Destination port for logging.
	IP address of the receiving server.
	the type of transport protocol. (UDP or TCP)
	number specifying the sender by IPFIX protocol standards.

Copyright © 2024 Greycortex

Reported Information

The reported information can be modified in the **Exported columns** table, which defines the parameters to be reported.

There are these options:

- Active columns are checked with a checkbox.
- The sum symbol (Σ) enables the count of items in the column for the export.

Copyright © 2024 Greycortex

Description of reported atributes

Copyright © 2024 Greycortex

CEF standard format:

Copyright © 2024 Greycortex

• class_id --> Category of event.

• cnt --> Number of detected events with same parameters in the same minute.

• dst_ip_addr --> Destination host IP address.

• dst_ip_icon --> Destination host icon.

• dst_ip_name --> Destination host hostname.

• dst_ip_rep --> Destination host reputation based on blacklisted IP addresses.

• dst_mac_addrs --> List of destination host MAC addresses.

• dst_port --> Destination port of the communication.

• dst_sn_addr --> Destination host subnet.

• dst_sn_icon --> Destination host subnet icon.

• dst_sn_name --> Destination host subnet name.

• dst_us_icons --> List of destination host user icons.

• dst_us_ids --> List of destination host user identifiers.

• dst_us_names --> List of destination host user names.

• end_time --> Time when was the last event detected.

• interfaces --> List of network card interfaces on which was event detected.

• ip_family --> Link to IANA.

• ip_proto --> Link to wikipedia.

• link --> url address to the reported event in MENDEL.

• min_time --> Time when was the first flow with event started.

• name --> Name of the event.

• s_type --> Type of provided service.
LOCAL - Communication is received. Services is provided.
REMOTE - Communication with remote machine.
ANY - cannot be determined.

• sensor --> Name of the MENDEL sensor.

• service --> Name of detected engine.

• service_icon --> Icon for service.

• severity --> Risk of the event in min = 0 to max = 10.

• sid --> ID of the event.

• src_ip_addr --> Source host IP address.

• src_ip_icon --> Source host icon.

• src_ip_name --> Source host hostname.

• src_ip_rep --> Source host reputation based on blacklisted IP addresses.

• src_mac_addrs --> List of source host MAC addresses.

• src_sn_addr --> Source host subnet.

• src_sn_icon --> Source host subnet icon.

• src_sn_name --> Source host subnet name.

• src_us_icons --> List of source host user icons.

• src_us_ids --> List of source host user identifiers.

• src_us_names --> List of source host user names.

• start_time --> Time when was the first event detected.

• sys --> type of detection: NB: network behaviour methods, ID: intrusion detection methods, SM: internal Status Monitor.

• timestamp --> Time when event was reported.

• traffic --> Indicates direction of communication. Possible values: ALL, INBOUND, OUTBOUND, INTERNAL, EXTERNAL.

• tunneled --> Boolean value which means that the traffic was tunneled.

Copyright © 2024 Greycortex

• vlan_ids --> List of VLAN tags on which was event detected.

Copyright © 2024 Greycortex

LEEF standard format:

Copyright © 2024 Greycortex

• class_id --> Category of event.

• dst_ip_addr --> Destination host IP address.

• dst_mac_addr --> List of destination host MAC addresses.

• dst_port --> Destination port of the communication.

• ip_proto --> Link to wikipedia.

• severity --> Risk of the event in min = 0 to max = 10.

• sid --> ID of the event.

• src_ip_addr --> Source host IP address.

• src_mac_addrs --> List of source host MAC addresses.

• timestamp --> Time when event was reported.

• sensor --> Name of the Mendel sensor.

• sysComponent --> type of detection: NB: network behavior methods, ID: intrusion detection methods, SM: internal Status Monitor.

• name --> Name of the event.

• traffic --> Indicates direction of communication. Possible values: ALL, INBOUND, OUTBOUND, INTERNAL, EXTERNAL.

• src_ip_name --> Source host hostname.

• src_ip_icon --> Source host icon.

• src_ip_rep --> Source host reputation based on blacklisted IP addresses.

• dst_ip_name --> Destination host hostname.

• dst_ip_icon --> Destination host icon.

• dst_ip_rep --> Destination host reputation based on blacklisted IP addresses.

• src_sn_addr --> Source host subnet.

• src_sn_name --> Source host subnet name.

• src_sn_icon --> Source host subnet icon.

• dst_sn_addr --> Destination host subnet.

• dst_sn_name --> Destination host subnet name.

• dst_sn_icon --> Destination host subnet icon.

• src_us_ids --> List of source host user identifiers.

• src_us_names --> List of source host user names.

• src_us_icons --> List of source host user icons.

• dst_us_ids --> List of destination host user identifiers.

• dst_us_names --> List of destination host user names.

• dst_us_icons --> List of destination host user icons50.

• service --> Name of detected engine.

• service_icon --> Icon for service.

• s_type --> provisioning of the service. LOCAL, REMOTE, ANY = can not be determined.

• ip_family --> Link to IANA.

• vlan_ids --> List of VLAN tags on which was event detected.

• interfaces --> List of network card interfaces on which was event detected.

• tunneled --> Boolean value which meansif the traffic was in tunnel.

• cnt --> Number of detected events with same parameters in same minute.

• start_time --> Time when was the first event detected.

• min_time --> Time when was the first flow with event started.

• end_time --> Time when was the last event detected.

• link --> url address to the reported event in Mendel.

Copyright © 2024 Greycortex

Reports

REPORTS

Reports are able to create and send selected information at defined (regular) intervals.

Adding a New Report

A new report is created with the button and filling the attributes below.

Attribute	Description
	Name of report.
	Description.
	Email addresses of the recipients of the report.
	Disables or enables sending the report.
	The time interval after which the report is to be sent.
	Setting the validity of the report. (empty means infinite validity)
	The time interval after which the report is to be sent.
	The time zone definition.

Creating Report Content

Creating a report content is done using the table. These sections are pre-defined components based on the Dashboards components with the ability to apply a saved filters. Simply put a new section into your report by clicking the button and select a sections from the dropdown menu and apply desired filters if you want.

For each section added into the report you can change the order, edit the applied filter or for the type section you can define number of rows.

Generating a report

To generate a one time (test) report based on defined content, just hit the button in table. To start sendind a report automatically and repeatedly, you need to it. In the same table you can check your reports history as well.

Copyright © 2024 Greycortex

Managerial & Security

MANAGERIAL & SECURITY

On this page you can generate two new types of reports. There are two preset reports. To enable them, click on the . You will be prompted to fill the email and after saving the configuration the report is enabled.

- **Managerial Report -** Primarily intended for managers and non-technical persons. These reports provide high level overview of the security situation in the network.
- **Security Report** - More detailed reports primarily intended for the administrator or security analyst. These reports summarize incidents and more technical details.

To create a new report, click on the button. In the Company dropdown, you can choose the company that the report is for. Then the logo of the company will be shown in the first page of the report. However, this is not compulsory.

Then you choose the type of report you want to generate. You need to fill the name of the report and the email addresses that the report should be delivered to. Choose the time interval for which the report should be generated. If you select the interval other than day, choose the day of the week that the report will be delivered on.

In Managerial reports, you can select up to 5 critical systems that will be included in the report. If you need more, you have to create the second report.

In the advanced settings you can set the language for the report, the time of the delivery, history size (the number of saved previously generated reports, up to 30), sensor and subnets.

Copyright © 2024 Greycortex

Companies

COMPANIES

- Company is important for Reports.
- You can set it up by clicking on button.
- Expand company configuration by clicking on symbol.

Attribute	Description
	Name of the company, will be visible only in list of companies
	Logo which is displayed on the first page of the report. • Remove - remove picture • Upload - upload picture • Picture can be only in .png and the best resolution is 400 x 250 px. If image has different resolution, logo in report will be changed to the recommended resolution. • Maximum size is 100 KB.

Copyright © 2024 Greycortex

3.8.5 User Integration

Sensor Groups

Sensor groups

Ability to create logical sensor based on a physical sensors.

Copyright © 2024 Greycortex

Identity provider configuration

Identity Provider Configuration

Integration to obtain user identity from various identity providers, such as , , , etc.

You can fing more information in guides below:

- Active Directory Configuration Guide
- Identity Service Engine Configuration Guide
- Identity Logs Configuration Guide

Copyright © 2024 Greycortex

Users

USERS

When using an integration to obtain user identity from various sources it lists all this users and their information.

Copyright © 2024 Greycortex

3.8.6 System

Global Settings

Httpd Settings

- Certificate for the running web server on MENDEL.

Update Settings

Set up an automatic updates and upgrades.

Attribute	Description
	The IP address of the server with updates and upgrades. If the box is empty, the default GREYCORTEX update server update.greycortex.com is used.
	Application updates are released in Alpha/Beta/Release phases. By default, the automatic download is set. If you have an administrator account, you can choose Beta or Alpha download (Integrator/Support assistance required).
	The frequency of how often automatic updates are triggered. In the case of starting updates once a day, you can also set a specific start time. • Note: The start of the update may be postponed for up to 1 hour due to system load.
	Time of the last data update checkup.
	Time of next data update checkup. You can use in order to update manually.

Important

Automatic updates don't work when the Passive mode is set or when the application doesn't not have access to an external network. Alternatives are Offline Updates (see or).

Maintenance Settings

- The deadline for system and application upgrades.

Server Time Settings

For the correct evaluation of MENDEL data, the same time is required for all sensors. The following section is for checking and correcting the synchronization of time between the collector and the client computer from which you are making the settings. The collector then sets the time on the sensors.

Attribute	Description
	Time on the client computer with a web browser accessing the application\'s user interface.
	Time on the collector.
	To set time on the collector either manually or using the NTP server. Changes to the time must be saved by clicking on the button in the bottom right corner.

MISP Configuration

The Threat Sharing project consists of multiple initiatives, from software to facilitate threat analysis and sharing to freely usable structured Cyber Threat Information and Taxonomies.

Copyright © 2024 Greycortex

To connect to specific MISP platform you need to fill in:

Attribute	Description
	IP address or domain where the MISP platform can be reached
	Access token
	How often the data will be exchanged
	Distinguisher to filter MISP records (optional)

Discovery Settings

Attribute	Description
	The addresses of the monitored stations are converted to domain names. **Important**: In some cases, this may lead to delays in data processing. In this case, it is recommended to disable DNS usage.
	Interval for updating internal DNS records from external DNS servers, defaulted to 1 hour.
	A time setting indicating for how long the DNS record will be in the DNS cache.
	Interval for letting the system use expired DNS records.
	System automatically recognizes unnamed subnets and creates generic entry in the subnet list.
	Hostnames are taken from DHCP data.
	It checks if the network is flat (uses only switches) and sensors can "see" all MAC addresses. They will then be used to identify devices on the network. For a segmented network (routers and firewalls are included in the network), leave the checkbox blank. Check this box if you are sure you are monitoring the network behind the firewall. This prevents false detections.

Storage Settings

- Fast storage provides faster response and faster read speed from used data. This option can be set only if MENDEL is installed on two disk configuration with fast (SSD) disks on primary storage.

Primary storage

- If enabled, storage compression will compress old files for better data retention.
- For primary storage you can set how many days data will be stored in uncompressed format for better performance.

Secondary storage

- This option can be set only if MENDEL is installed on two disk configuration.
- If enabled, storage compression will compress old files for better data retention.

Copyright © 2024 Greycortex

System Components

Activation/Deactivation of main system components.

Intrusion Detection System

The main task of the Intrusion Detection System is to analyze packets in the network, their reconstructed content and related metadata to look for patterns that match their database similarly to an antivirus engine. In contrast to the NB analysis, it doesn't use any intelligent or adaptive rules, but a strictly defined set of signatures. In this way, Mendel's ID system can detect a wide range of known malware, attacks or corporate policy violations, which behave "normally" in the network.

Network Behavior Analysis

It is an analysis of input data with the Artificial Intelligence/machine learning engine. NBA can detect security threat vectors even when there are no signatures available. It can detect various exploration and brute-force techniques as well as data leaks and other anomalous activities within the network not limited to the perimeter. MENDEL's NB analyzer does not focus on a method of attack nor a piece of malicious code, but on symptoms of the attack. By detecting atomic symptoms of malicious behavior and anomalies in network traffic, it identifies emerging security incidents in their early stages. Thereby, it decreases incident response time, prevents further damage and helps to decrease overall cyber security risks.

Subcomponent	Description
	NBA signatures denoted with keyword Discovery.
	NBA signatures denoted with keywords Flow, Scan, or Periodic.
	NBA signatures denoted with keywords Limit or Outlier.
	MENDEL can detect the most widespread OS (including distributions) like Linux, Windows, macOS, and Android. The system will not detect OS on hosts which are marked by proxy tag.

Network Statistics

Network Statistics enable the computing of advanced network statistics from detected data flows. These further populate the Subnets, Hosts, and Services tabs.

Status Monitor

The component does several essential tasks to monitor system running and system services, monitor hardware resource usage (CPU usage, memory usage, disk space usage, or current traffic on network interfaces).

Event Correlation

The Event Correlation (EC) engine correlates events from other detection methods for more precise detection and highlights the most critical threats.

Data Updates

This enables downloading of regular updates for IDS rules, blacklists, CVEs, GeoIPs, and other details which are later used in the process of event discovery and result presentation. The frequency of these updates can be set using the Auto Updates Time value in the Global Settings section.

Central Collector

This indicates whether or not this machine is set to function as the Central Collector. If so, it receives processed events from subordinate Collectors (which can further receive unprocessed data from their subordinate Sensors).

Incident Management

Allows further escalation of events, situations, and states of your choice by marking them as an Incident using the Report Incident button on an event details. Single or multiple events, hosts, discovered subnets, and more can form an incident. Reported incidents can later be viewed and managed from the Incidents tab.

OT Metrics

Metrics for OT / ICS systems.

Copyright © 2024 Greycortex

Asset Survey

Active survey that can obtain more detailed asset information. This function uses active interrogation against selected IP address.

Log processing

Processing logs from 3rd party sources.

Copyright © 2024 Greycortex

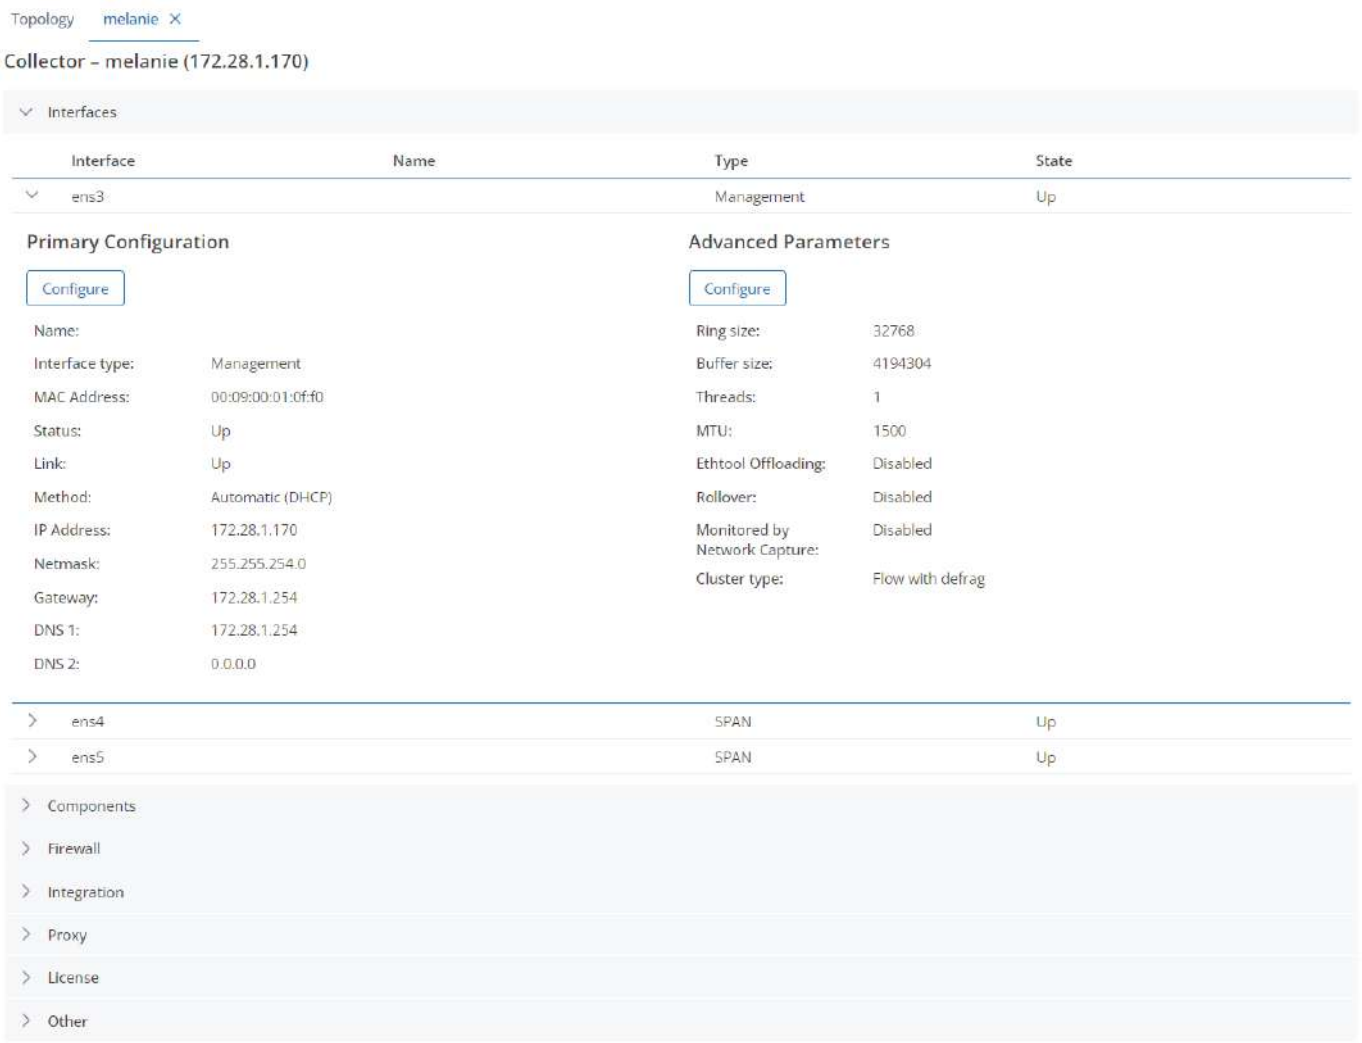
Sensors & Collectors

SENSORS AND COLLECTORS

The application consists of two basic components -> the sensor and the collector/analyzer. The sensor receives the network data stream and is implemented using a network interface and software extensions.

The **collector** processes and saves the data into the database system and ensures their visualization.

The **Sensors** option displays all the sensors and collectors connected to the system.



Interfaces

List of physical interfaces in the system with the following parameters:

Attribute	Description
	Interface identification/name.
	Indicative interface type designation. (eth/WLAN/span/etc.)
	Current status. (Up/Restarting)
	Allows you to name the interface.

Copyright © 2024 Greycortex

Some parameters can be changed by clicking on the button. In the configure dialog you can change:

- Name of the interface
- Method of getting an IP address -> Automatically by DHCP or manual settings
- State of the interface -> restart interface

Interface bonding

To bond interfaces into the virtual interface, use the button .

You need to define a , and "slave interfaces". You can further set bonding options in detail.

Components

Depending on the component, the user can restart it, turn it on/off or configure it.

Components in the system:

- **Application Server** -> this component is responsible for the GUI. Users can just restart this component; it can't be shut down.
- **Database** -> component is responsible for the database, it is a really important part of the system, so the user can just restart this component.
- **Network Capture** -> This component is responsible for the overall reception. User can turn it on/off, restart or configure these parameters:
- *Rmem Max* -> Maximum memory for receive buffer.
- *Netdev Max Backlog* -> Maximum size of receive queue to netdev_max_backlog frames.
- *Flow Inspection Depth* -> how much of the content in the request body is inspected.
- *Stream Reassembly Memcap* -> Max memory usage (in bytes) for TCP session tracking.
- *Stream Memcap* -> maximal amount of memory allowed for TCP stream objects.
- *Flow Memcap* -> maximal amount of RAM allowed for Flow objects.
- *Deep Packet Inspection Enabled* -> When enabled, Deep Packet Inspection ensures, that the application layer protocol has data sent for parsing per flow.
- *Rules Enabled* -> When enabled, the event detection can be based on the IDS rules.
- *Translate non-IP flow MAC addresses into IPv6 addresses* -> when this option is on, it distinguishes IP flows according to MAC address.
  Recommendation: Turn on when using SCADA.
- *Match* -> You can choose if the events should be matched according to the Network definition of the rule or Perimeter definition of the rule.
- *BPF filter* -> In some cases there are reasons to ignore certain traffic. Maybe a trusted host or network, or a site. Through BPFs the kernel can be told what to send to a Network Capture.
  For example: a simple filter 'tcp' will only send TCP packets, filter src net 192.168.5.0/24 and 'dst' host 192.168.20.41 will send only this source subnet and destination host.
- **Detection Engine** -> This component detects anomalies in the system. Users can turn it on/off, restart or in configuration dialog turn on Layer 7 analysis parsing.
- **Netflow** -> Responsible for processing NetFlow data. User can turn it on/off or restart.
- **Status Monitor** -> Monitors the state of the system. User can turn it on/off or restart.
- **System Upgrades** -> Enables upgrades for the MENDEL sensor. User can turn it on/off or restart.

Firewall

The firewall is used to control access to the appliance. By default, all 80/TCP and 443/TCP ports are disabled. It is also possible to have an open port 22/TCP for SSH for support purposes.

Copyright © 2024 Greycortex

By clicking on the button the "Add Exception" dialog will pop up with the following items:

Attribute	Description
	Action you want to perform with the specific port: ACCEPT, REJECT, or DENY.
	Number of the port of the service.
	The subnet for which rule will apply.
	Protocol of the service TCP or UDP.
	Description of the rule.

Rules for Internet access are managed on your network firewall. Please allow license.greycortex.com, update.greycortex.com, and archive.greycortex.com on your network firewall so that system and data updates, as well as license verification, are working.

Integration

In this tab, you can connect appliances one to another. There are two options:

- **Connect some other appliance to the appliance that the user is currently logged** -> for this you need to generate Connection token and login to another appliance to establish the connection
- **Connect appliance that the user is currently logged to another appliance** -> User can choose the type of connection. It's possible appliance to be Sensor, Collector or CEM based on many conditions. For a better understanding of this topic, please read the guide.

Proxy

This component allows set the Proxy server (if any) for that appliance.

Others

Set up for the description of the appliance, Latitude, and Longitude.

Copyright © 2024 Greycortex

NetFlow

NETFLOW

The NetFlow collector is used to record data at the NetFlow level that is generated in the appropriate network devices that are not covered by GREYCORTEX MENDEL (routers or switches).

Processed versions of the NetFlow protocol

The current product version processes the following NetFlow formats:

- **NetFlow v1**
- **NetFlow v5**
- **NetFlow v9**
- **IPFIX**

The application does not support the use of encrypted communications at the NetFlow protocol level.

Application Settings

The NetFlow collector can be configured in the System section of Settings.

Attribute	Description
	The sensor name is selected from the list of connected sensors.
	Unambiguous designation of a network source. The user can enter any number that is not occupied in the list of connected sensors. For reasons of clarity, we recommend using up to three-digit numbers.
	The port number on which the sensor listens to incoming traffic in the NetFlow format. For this port, NetFlow must be sent using a remote sensor/generator. The recommended value is the standard NetFlow port 4739. If this port is already being used, it is recommended to use ports in the range of 2000-3000 to prevent other applications from being corrupted. A firewall rule for a selected port is added automatically.
	Allows you to select the transport log type from the menu. Currently, only the UDP protocol is supported.

Example of a NetFlow Generating Sensor

These are third-party facilities. After connecting the sensor to the network, you need to set the parameters for the collection:

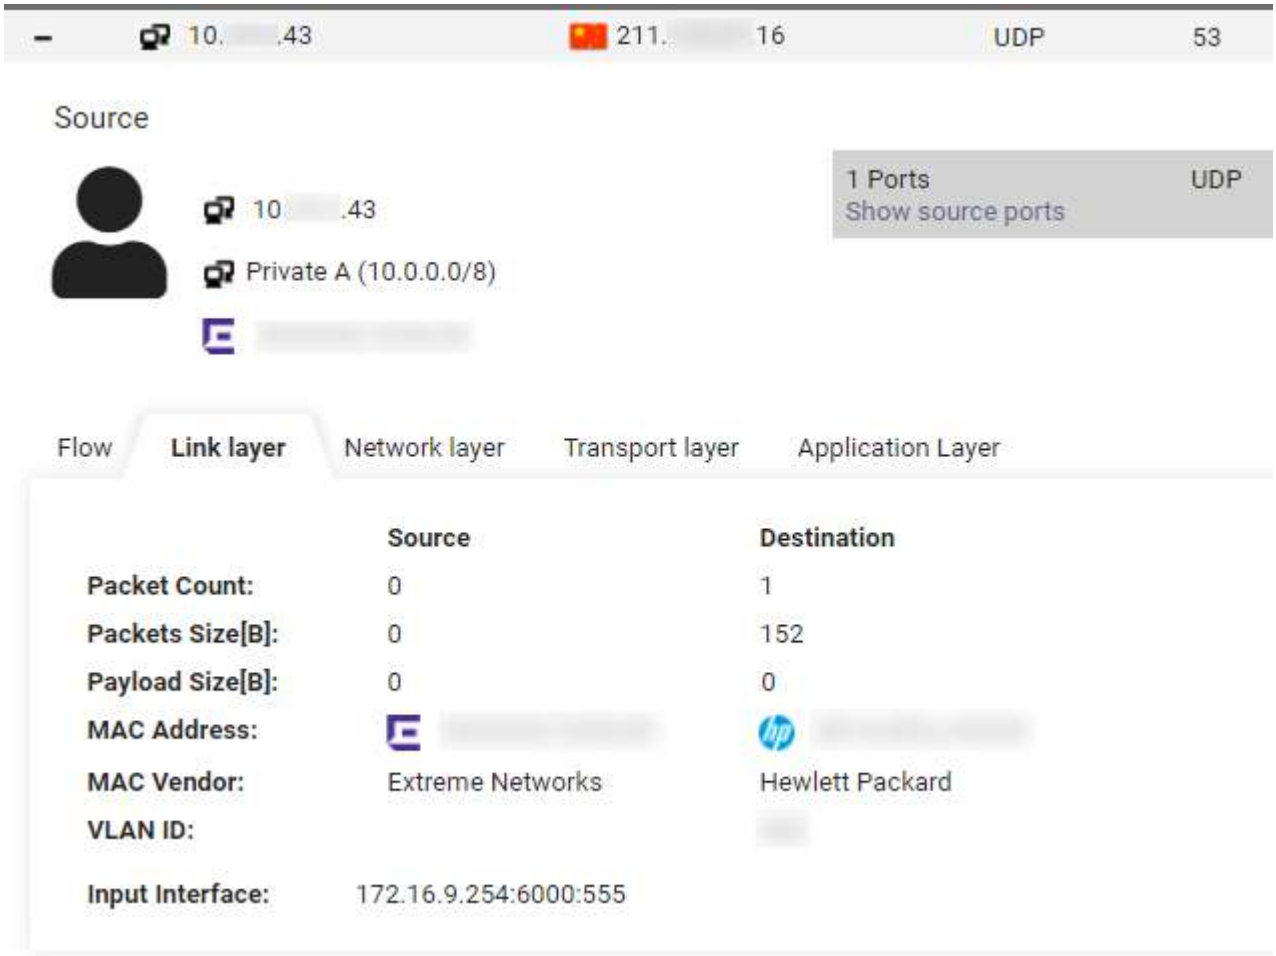
Attribute	Description
	The source address of the subnet from which NetFlow will be logged.
	Target IP address of the application.
	The port number to which NetFlow will be sent.
	Version of the transport protocol. Typically, these are TCP or UDP.
	Set the time interval in which data is to be sent. The recommended value is 600 seconds, i.e. 10 minutes.
	The buffer timeout during which the evaluated data is stored in the database and deletes the unnecessary data. The recommended value is 1800 seconds.

Copyright © 2024 Greycortex



Recognize NetFlow in Flows

Flows from a NetFlow exporter can be recognized by the input interface on the link layer. In the image below, we can see a flow from an exporter. In the link layer, we can see that the input interface has an IP address 172.16.9.254, which is the address of the NetFlow exporter which sends data to the sensor on port 6000. The number 555 is the NetFlow ID number setup in settings.



Copyright © 2024 Greycortex

Online KBs

ONLINE KBS

Set up external sources of information. The IP addresses/domains/ports displayed in the application can be accessed by connecting to external information sources. Use this page to set up these resources.

An external information source can be any web service that allows IP addresses/domains/ports to be searched and provides more detailed information. In addition to publicly available sites (e.g. google.com), you can also use your own service, such as an internal Wikipedia server.

When you click on the IP address/domain/port, an information dialog with a list of page names is displayed. The selected page opens in the new browser panel at the appropriate address.

External resources are defined by the following items:

Attribute	Description
	The URL where the asterisk (*) is replaced by an IP address, domain, or port.The URL where the asterisk (*) is replaced by an IP address, domain, or port. • Examples: https://www.google.com/search?= or https://www.google.com/search?=port+
	Specifies the search subject: IP address/domain/port.
	Allows you to delete a page from the list.

Copyright © 2024 Greycortex

Email

EMAIL SYSTEM SETTINGS

To send emails from the application you need to configure the SMTP server. Messages can be encrypted and signed. GnuPG is used to encrypt emails.

Mail server set up / Global settings

Basic Settings	
	Name of the mail server.
	Security protocol usage: None, StartTLS, or SSL/TLS.
	The port number for communication with the mail server.
	Authentication method for access to the mail server.
	The user name used to log in to the mail server.
	The user password used to log in to the mail server.
	Mail address on the mail server from which emails will be sent.

Signature & encryption	
	The email signature with a Private PGP key.
	The email signature with a Private PGP key password.
	Public PGP key for mail encryption.
	Type of encryption algorithm.
	Type of message digest function.
	Verifies that the outgoing mail server is set up correctly.

GnuPG Available Keys

To encrypt/sign email from an application, PGP keys are used from the Available Keys list. You can import keys from a file, select Import, or create a new key pair directly by filling out the form and clicking the button.

Attribute	Description
	Type of PGP key.
	Length of the PGP key in bits.
	Name of the PGP Key.
	Email of the owner of the PGP Key.
	Time interval for how long the PGP Key is valid.
	Password for the PGP Key.

The appropriate key is then used for its respective encryption. Email signing is defined by assigning a key in Set Signing and Encryption.

GnuPG Key Import

Imports existing Public or Private GnuPG keys.

Copyright © 2024 Greycortex

Email Notification

Enables sending email notification for sepcific situations.

- **Incident notifications** - to users when significant events occur (e.g. creation of an incident; completion of an incident).

Copyright © 2024 Greycortex

Certificate

CERTIFICATE

The certificate bookmark is used to import trusted certificates for the https server connection. Creating and importing certificates is based on the commonName standard.

+ (symbol) -> Shows detailed information about a certificate.

Upload Certificate

Option to input an already existing certificate to the collector:

Attribute	Description
	Name for the certificate of your choice.
	The collector you want the certificate to be uploaded to.
	Type of the file you are uploading: Certificate, Certificate+Key, Certificate Bundle.
	Source of the upload; file Upload or Plain Pem certificate format.
	The button for a certificate upload if "File Upload" option is selected or an input area for "Plain Pem" option.

Create Certification Request

Generates the file necessary to request a new certificate from the certification authority:

Attribute	Description
	Name for the certificate of your choice.
	The sensor you want the certificate to be uploaded to.
	The name of the server for which the certificate is issued.
	Organization, Unit, Location, State and Country.
	Length of certification key in bytes.
	Type of the hash algorithm.

Copyright © 2024 Greycortex

Maintenance

MAINTENANCE

System Maintenance Actions

Attribute	Description
	You can download logs from the Sensor from particular Components.
	Restarts the collector. During the restart, the graphical user interface will not be available.
	To use in a situation where remote support from the manufacturer is required. After the agreement with the manufacturer's technical department, by clicking this option, you enable **remote access** to the manufacturer's technical support for your installation. Don’t forget to enable the backup by switching the toggle.

Data Backup to Samba Server

Allows you to set up access to the Samba server for automatic backup of system settings, data and events. In configuration you need to fill all the required values. Server must be external Samba server. NFS system is not supported. The recovery process must be performed by the Support.

Backup Settings

The application allows you to backup:

- **Policies**
- **Subnets**
- **Global Settings**
- **False Positives**
- **Users**

Backups are provided in a .csv format for later renewal purposes, or for setting another sensor. The backup is stored on the client computer disk as a *.tar.gz archive.

Restore/Import Settings Allows you to restore policies, configuration, etc. from backup files, or you can use properly formatted .csv files to set up your system.

Copyright © 2024 Greycortex

Data Import

DATA IMPORT

This tool allows for manually updating Mendel's data, such as IDS rules, signatures, GeoIP or ASNs. This operation can be based on a provided update file or upon given checkpoint date, which will then proceed automatically.

Copyright © 2024 Greycortex

Changelog

CHANGELOG

On the Changelog page, you can find the listed versions with their descriptions. To see available versions, you must have the update server configured. The default GREYCORTEX update server is update.greycortex.com and the certificate is already in MENDEL after installation. The backup server for data and older system updates is archive.greycortex.com Make sure that both update servers are allowed through your network firewall using TCP port 443.

The information about the new available version of the system - besides changelog - is displayed when the user - the system manager group with at least read rights - logs in. The information is displayed as the message on the top of the page.

Copyright © 2024 Greycortex

System Upgrade

SYSTEM UPGRADE

System upgrade page allows you to update MENDEL for the newest available version. Older sensors and collectors are not capable of working together if they differ by more than one version.

Online Upgrade

If you have the newest available version, you will get notification on the top of the page:

You are able to see the version you are using and have the option to upgrade all sensors and collectors or click to which will open a new tab with the released versions and information to them.

Attribute	Description
	Consists of collector and all available sensors. • Collector processes and saves the data into the database system and ensures their visualization. Collector can be only one. • Sensor receives the network data stream and is implemented using a network interface and software extensions.
	The name of the collector or sensor.
	The version of the collector or sensor. You can click on the version and see Changelog page.
	Update server column consists of these information: 1) The status of a collector or sensor can be available or unavailable. 2) An option for checking the availability of the update server for all sensors. 3) An information of the last check.
	Can only be seen if you do not have the latest version. You can click on button and start updating specific sensor or collector.

Offline Upgrade

In situations where the appliance have no access to the update server (e.g., secure network zone with no internet access - DMZ) it has the ability to upgrade itself (collector, AiO) and connected sensors through mounted local device (e.g., USB drive) or remote SMB folder.

For this scenario simply choose the mode on the main page.

After the proper file for the upgrade is mounted the new version availability is rendered and you are able to initiate the common upgrade process.

Local device

Whenever the local storage device is connected to the server it appears in a list.

Samba Share

To mount a remote storage you need to specify the path, user name and password.

Important

If you update your collector your interface will stop working, new tab will open in your browser where you can see the progression of the update.

Copyright © 2024 Greycortex

EULA

EULA

The End User License Agreement is listed here for the reference. It is also listed in this manual.

Copyright © 2024 Greycortex

Plugin

PLUGIN

The Plugin page lists all the plugins/scripts that are available in the system. These plugins can interact with the network devices, e.g. blocking the communication in the firewall. It is necessary to create a plugin profile for the plugin. Then the action from the event can be executed.

An administrator can add a new plugin, but the script (Bash or Python) must be approved and signed by the GREYCORTEX. For the approval, please send the plugin file to the GREYCORTEX Support reachable at e-mail: support@greycortex.com.

For the deletion of the user added plugin, it is necessary to delete all the created profiles that are based on this plugin first.

Copyright © 2024 Greycortex

Plugin Profile

PLUGIN PROFILE

For every plugin, you can create plugin profiles, where each profile serves as a separate configuration of a plugin for the particular network device. Therefore, if you want to configure several network devices using e.g. Fortinet plugin, you need to create a plugin profile for each one of them.

The profile is added by clicking on the button at the top of the page.
Select the plugin for which the profile should be created.
All the compulsory fields marked with the star symbol must be filled, other fields are optional.
For more information please follow Automated Threat Response, Plugin Development and Firewall Integration guides.

When the profiles are created, they are listed in the table.
The details of the profile are shown after clicking on the symbol at the beginning of the row.
Attributes shown are based on the plugin they are created for and are editable.
Besides the attributes, the basic information about the plugin is displayed.
Plugin profiles can be disabled or removed.
After the plugin is created, you can take actions by clicking on the button in Events. If there are no plugin profiles, this button is hidden.

Copyright © 2024 Greycortex

Authentication

AUTHENTICATION

Setting authentication methods allows you to specify the security level and authentication type.

Default Built-in Password Policy

Designed to define a password security policy, customizable parameters are predefined:

Attribute	Description
	Policy settings to check repeat passwords used by the user.
	Password setting policy.
	Turn on/off forced password strength.
	Policy to force password change over time.

Default Built-in Authentication Policy

Setting the default policies for user authentication, the policy always applies to a combination of username and IP address.

Attribute	Description
	Sets the maximum number of attempts to log in. It is necessary to wait for the defined time for the next attempt.
	Sets the invalid duration after invalid attempts before an attempt can be made again.
	Idle timeout for automatic logout.

Active Directory authentication method

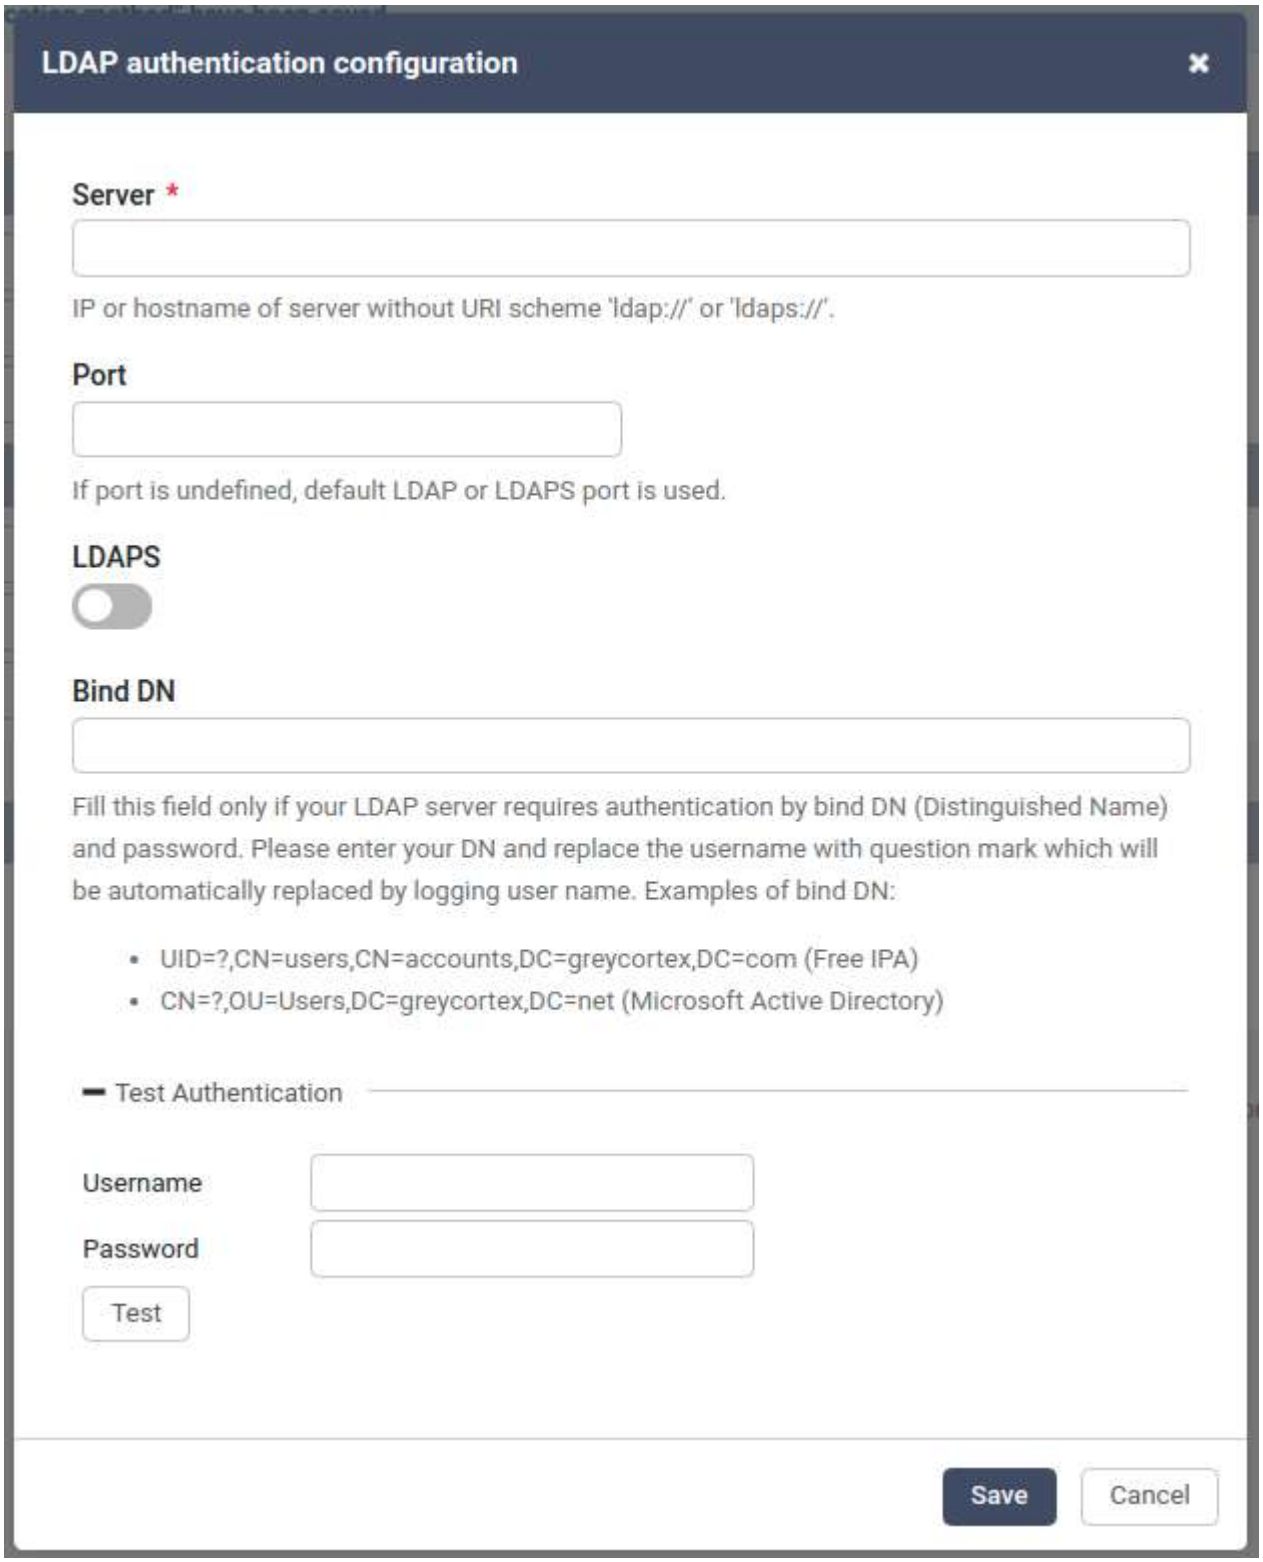
Native/default authentication is HTTP authentication through HTTPS. It allows you to authenticate users who have a valid password account with the right to access the application.

Active Directory

This enables authentication of users through Domain Controller and Active Directory. It is necessary to create a user account with the same user name as in the LDAP structure or as the Kerberos structure. Once created the user needs to be configured as an EXTERNAL USER. This can be done on the Persons page in the Accounts section of the settings.

LDAP

Copyright © 2024 Greycortex



Allows users to authenticate through LDAP or LDAPs. To use this service, you must enter an IP address in the LDAP server settings. You can use authentication with a certificate.

To specify that a request is sent to the LDAP server, you must fill out the Bind DN. It can include parameters like UID, CN, DC, etc. depending on the LDAP server configuration.

Copyright © 2024 Greycortex

Kerberos

Allows users to authenticate through Kerberos. To use this service, you must enter the IP address of the Kerberos Server and Realm into the Kerberos settings.



Copyright © 2024 Greycortex

Status Monitor

STATUS MONITOR

It is used to set individual signatures allowing system monitoring and supervision. Signatures cannot be added or deleted. They are set by default.

Attribute	Description
	It serves to display the details of the signature.
	Unique signature identifier.
	The signature name always starts with the status monitor prefix.
	Signature description.
	Recommendations when the situation occurs.
	Allocates the signatures to the class.
	Sets the severity of the event using a number on 1-10. (1 -- OK, 10 -- Critical)
	Turns signatures on and off.

Copyright © 2024 Greycortex

User Activity log

Audit every step of a Mendel user or API queries to comply with legislative requirements or internal policies.

AUDIT LOG

Narrow down your search by using the filter attributes.

Attribute	Description
	Start time of the record.
	End time of the record.
	User login name.
	IP address of the user client.
	Specific Mendel service affected by the listed action.
	Database method performed by the listed action.
	Hyperlink to the site which trigger the listed action.
	Formated and raw changes caused by the listed action.

Copyright © 2024 Greycortex

3.8.7 Accounts

Persons

PERSONS

User administration in the application allows full management -- creation, editing, and deletion.

The basic operation is to create a new user; for this purpose, you must use the button to fill in the user information, the star (*) attributes are required:

Attribute	Description
	The characters [a-zA-Z][\ -_ a-zA-Z0-9 @.] are accepted. The first the letter may be uppercase or lowercase, and the length is limited to 100 characters.
	An optional parameter may contain any characters.
	Mandatory parameter may contain characters _A-Za-z0-9. * @
	The default language that will be used in the application when the user logs in. User may change it later in the “My account”.
	Account is active, and it is possible to use it.

When the user exists, you can edit the user’s settings, except User Login. You can set him as an External User, this setting must be active when the Active Directory is used for login into the application. Here it is possible to manually reset user’s password. You can also add him to the system or custom group and set individual restrictions that are not inherited from the groups.

If the individual user has the shell access, you can see the information like when the shell account was Created, changed and last used for the access, the method used for authentication, used IP address and the number of failed access.

Copyright © 2024 Greycortex

Groups

GROUPS

On this page you can manage manually created groups. It is possible to create new groups and restrict them as necessary. The name of already created group cannot be changed.

Underneath is the list of all System groups. These groups cannot be changed or deleted.

To assign a user to a group, please go to the page (url:).

Members of the same group without further individual restrictions can see (or manage with Write permissions) objects created by other members. This includes Incidents, Reports (new), Companies, and Data Captures.

Copyright © 2024 Greycortex

Permissions

PERMISSIONS

General permissions for groups:

Attribute	Description
	User can view the content.
	The user can view the content, edit it and delete it. User, that is assigned to group with Write permissions, can also manage filters in this group.
	The user can view the content, edit it, delete it and delete the group itself if it is a group created by a user. We do not recommend using this permission.

System groups

Attribute	Description
	Network administrator can manage Network and Detection sections in Settings.
	System administrator can manage System section in Settings. Can activate the shell account on My Account page.
	Can manage Exports section in Settings. If restricted by sensor, only new type of Reports and Companies are available.
	Can manage user accounts.
	User can see the list of pre-defined filters. It is recommended to use Read permission.

Custom groups

Permissions in custom groups depends purely on restrictions that you add to the group.

Restrictions

The access to network information from selected sources can be restricted. This relates to any data shown on the Dashboards, Network Tab or Events Tab. If there are no restrictions, the user will see all the data.

If there is set at least one restriction, the user can see only data from this source. If you want the user to see e.g. data from all components but one, you must set restrictions to all components that you want to see. Please use only Read permissions when setting the restrictions.

Available restrictions:

- **Component** (detection methods)
- **Sensor** -- Restricted sections in Settings:
- Network -> Subnets, Hosts, Users and Detection -> User False Positives.
- **Subnet**
- **Host**
- **MAC address**
- **User**

By combining the groups memberships with permissions and restrictions you can set what the user should be able to see and manage in the application. If you have any trouble or a specific use case, please contact the partner or support.

Copyright © 2024 Greycortex

4. Setup

4.1 Automated Threat Response

It is an extension of the classical capability to integrate MENDEL´s network traffic analysis solution with existing customer **firewalls** through separate output module for exporting data to external systems outside MENDEL, with the possibility of **AUTOMATIC EXECUTION of PLUGINS**.

Funcionality

Based on the defined **rules**, the selected **plugin** (notification, blocking, etc.) will be **executed automatically** when the conditions are met.

Benefits

- Connection to customer security solutions (firewall, etc.).
- Possibility of operational and unattended use of data obtained and processed in MENDEL.
- Proactive solution of detected security incidents according to defined rules.

4.1.1 Creating rules for Automated Plugin Execution

For proper operation, the following procedure must be followed:

- Create script plugin (Plugin Development).
- Set firewall integration (Firewall Integration).
- Create a new record for data transfer to an external system.
- Set details before execution.
- Execution (enabling) of Automatic Plugin.

4.1.2 Creating a new record for data transfer to an external system

1. To create a new record or edit (delete) existing ones, you need to enter the:

2. Clicking the button will open a wizard with a form for creating a new export item.
3. By filling in the form and then confirming with the button a new record will be created for the automatic plugin execution.

Attribute	Description
	Type of the Export naming
	select "Events"
	Select "Plugin Execution"
	fixed value set for 1 minute (immutable value)

Copyright © 2024 Greycortex

4. Detailed settings before execution:

Atribute	Description
	the date by which (inclusive) the requested automatic action will be performed.
	fixed value set for 1 minute (immutable value).
	specifies which events should be processed by selecting default predefined filters (leave blank for no filter restriction).
	select the plugin profile that matches the intended request.
	the number of iterations to execute the request with identical arguments.
	define which parameters will be sent to the external system after plugin execution.

5. Execution (enable) of Automatic Plugin:

automatic execution is not performed
automatic execution is in operation

Important

Incorrect or inappropriate setting of parameters for automatic execution of scripts in the plugin can significantly limit or compromise the target device to which this data is passed. Please be extra careful when setting it up.

Copyright © 2024 Greycortex

4.2 Active Survey

4.2.1 Host Survey Guide

When the module Asset Survey is enabled, you are able to click on the button in the host detail dialog.

Pressing the survey button triggers an active scan on the selected host, which will result in additional information about hosts.

For example if a host, in this case a PLC is running the services **OPC UA Server** (OT) & **S7 Server** (OT) and flows have been detected by Mendel for these services, the survey button will trigger an reconnaissance packet to these service, leading to the identification-data to be stored in the Mendel system and displayed in the host detail dialog.

Additional benefits

- Mendel is using the detecte hardware & firmware versions to detect if the host is subject to security vulnerabilities. If your host is vulnerable an additional tab will appear with the listing of the CVE's. In case this it is advised to study the CVE-description and consider patching the host.
- Upon re-surveying the host any changes in the firmware-version are detected. In the Events tab, a message appears informing you that a firmware-upgrade (or downgrade) was detected. For this feature to work, please ask your supplier to activate it, or contact GreyCortex support for assistance.

Copyright © 2024 Greycortex

Enable Host Survey

- Go to the:

  Find and click on the Asset Survey .

Usage summary

- Click on the in the topbar.
- Afterwards click on .
- Select an and click on it.
- A new tab will pop up with Host Information.
- In **Asset information** click on .
- After a brief moment, when the scan has finished, additional information will show up.

Important

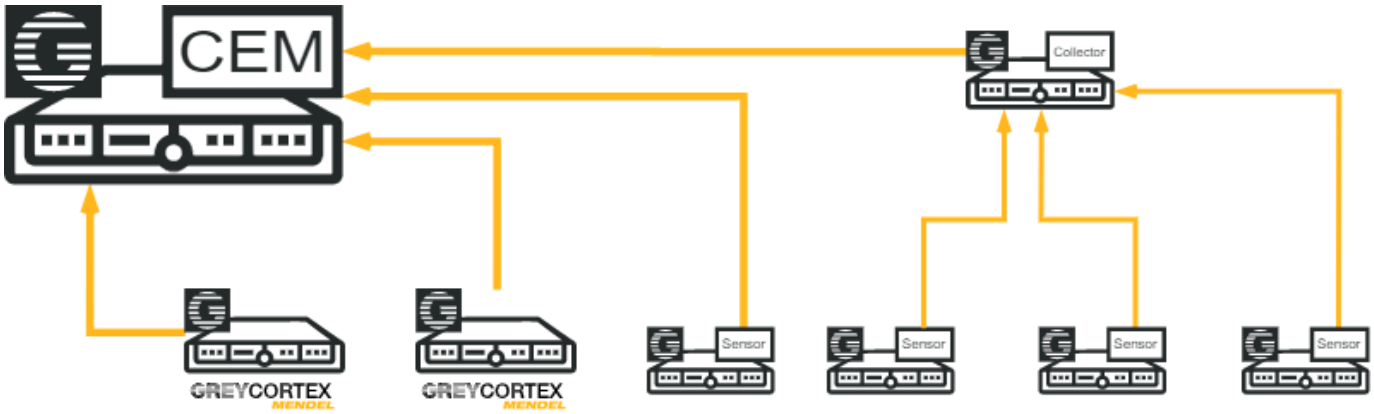
Actively surveying a mature OT-network has the potential to cause (temporary) interruptions to the OT-process. It is advised to test this feature in a test-environment, prior to using this in production. Surveying during a maintenance window is another option

Copyright © 2024 Greycortex

4.3 Central Event Management

4.3.1 Central Event Management

Central **E**vent **M**anager (**CEM**) is the highest possible element for overview and management over the entire deployed MENDEL architecture.



Designed for customers or partners with larger deployments, MENDEL provides the possibility to connect multiple appliances using a multi-level structure with **CEM** integrated.

CEM Architecture components

Sensors, Collectors and All-in-One(s), every of these MENDEL appliances you can connect and manage from one place = CEM.

- **Sensor** - Collects and analyzes network traffic based on detection signatures and generates records of the flows it stores in the database.
- **Collector** - Stores all the data from the sensor(s), analyzes the behavior of individual network data streams and provides all service and user services.
- **All-in-One (AiO)** - Solution with Sensor and Collector located on the same device. Procedurally it acts as a Collector.
- **CEM** - Offers the ability to connect appliances using a multi-level structure.

4.3.2 Benefits of using CEM

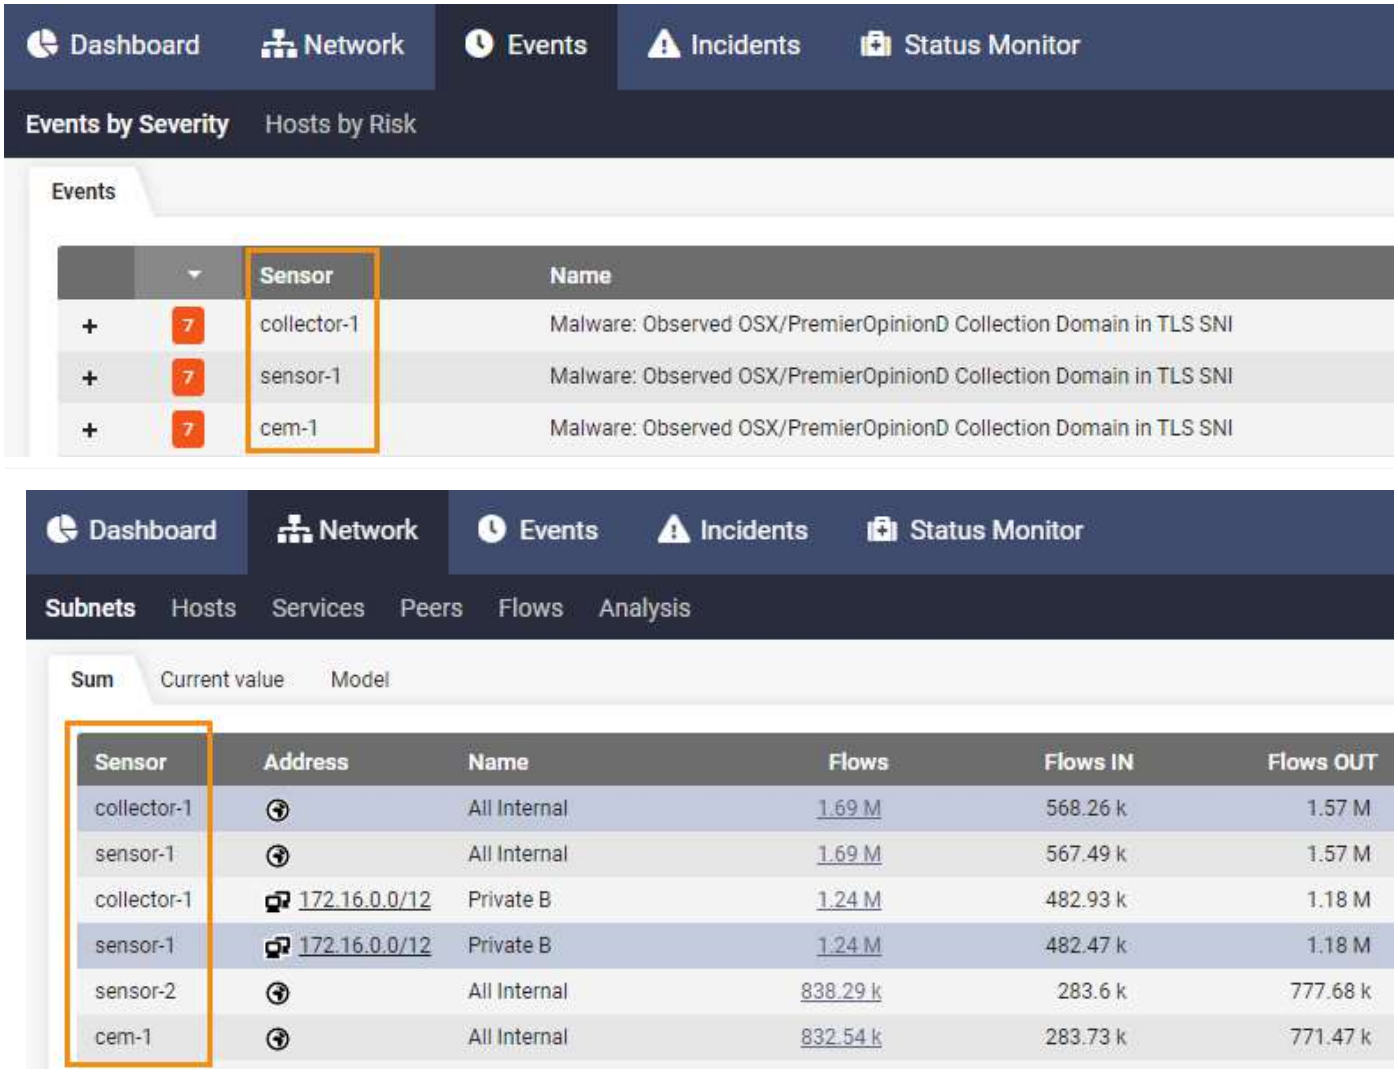
CEM displays a summary view of the network environment within the monitored infrastructure, including detected events from all devices in this network. It allows you to quickly identify potential problems and provides redirection to the target machine where the incident can be resolved.

Supervision of the overall infrastructure

Events and network traffic on subnets are displayed as an general overview.

Copyright © 2024 Greycortex

4.3.3

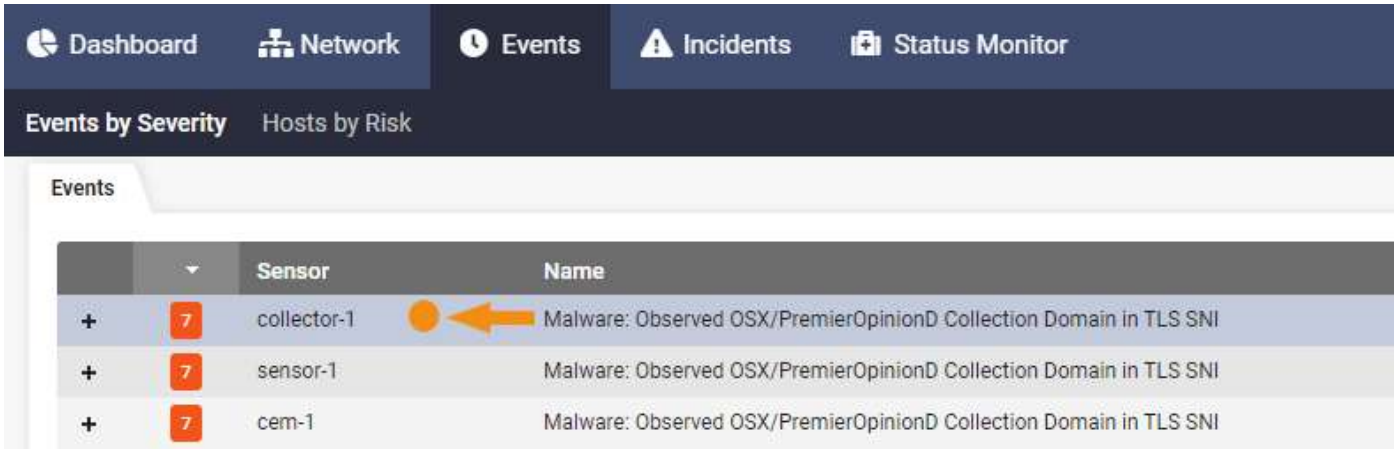


Quick way to detailed information

For more detailed information, just double-click on the relevant line with Events (Events by Severity) or with a subnet (Network Subnets). You will be redirected to the corresponding appliance, where you can continue to work with these data on a regular basis.

Copyright © 2024 Greycortex

4.3.4





Important information related to the deployment of CEM

INTERFACE CAPACITY USAGE

Deployment of Sensors and Collectors to target infrastructure has an impact on interface capacity usage.

Connection from **Sensor to Collector**
Each sensor generate 1-2% load from current traffic on collector management interface.

Connection from **Collector to Central Event Manager** (CEM)
Each sensor on each collector gernerate 0.1% load from current traffic on CEM management interface.

Example:
Deployment of 4x1000Gb sensor requires 40 - 80Mbit link capacity to collector and 4Mbit from collector to CEM.

4.3.5 Central Event Management - Commissioning

Description of individual steps that are used to create a CEM and connect the collector to it.

How to

... connect a standalone sensor or AiO directly to CEM, use the Sensor-Collector Guide.

In this scenario, the CEM acts as a collector for directly connected sensors or AiOs.

Copyright © 2024 Greycortex

4.3.6 1. Connect sensor-collector to the CEM

Before connecting, please check if you have allowed these ports on your firewall as well as in Mendel:

- For connection to CEM allow **port 80 and 443**
- Database is working on **port 5432**
- Logs can be downloaded on **port 22**

MENDEL database uses **OpenSSL** to encrypt connection. It uses certificate and strong encryption with **key longer than 128 bits**. Detailed cyber suite is **ECDHE-RSA-AES256-GCM-SHA384**

For demonstration in this tutorial, we use these components:

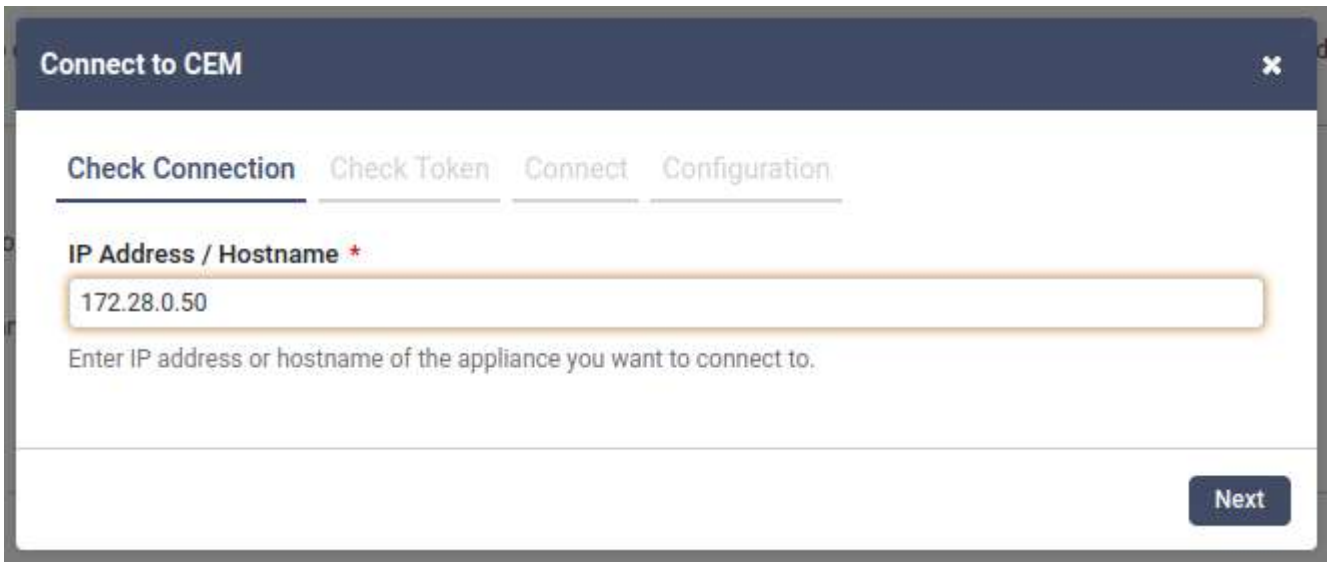
Collector with a connected sensor	All-in-one as future CEM
	Hostname: **mendel01**
	IP address: **172.28.0.50**

Step 1. Open settings on Collector (e.g. mendel02)

Go to and click on tab. There are few options for connecting more appliances together. You need connect to CEM - it doesn't matter if it is just all-in-one architecture or existing CEM. Click on button.

Step 2. Add IP address of the CEM

Insert the IP address of the future CEM and click to the button . You will see input field for connection token, it must be generate from future CEM.



Do not close this window, you will need it after generating token.

Step 3. Open settings on CEM (e.g. mendel01)

Go to the UI of the second appliance (in our case mendel01) and navigate to

and then click on the button.

Copyright © 2024 Greycortex

Step 4. Add IP address of the colector

If the appliance knows the IP address of the future collector, it will be automatically filled out. In case of empty box for IP address, you can click to the button.

This will show you what IP address is trying to connect. In case that this option is not working either (e.g. you have NAT in your network) you need to fill out the IP address manually.

Token will be generated by clicking on the button. You can copy it by the link on the right side .

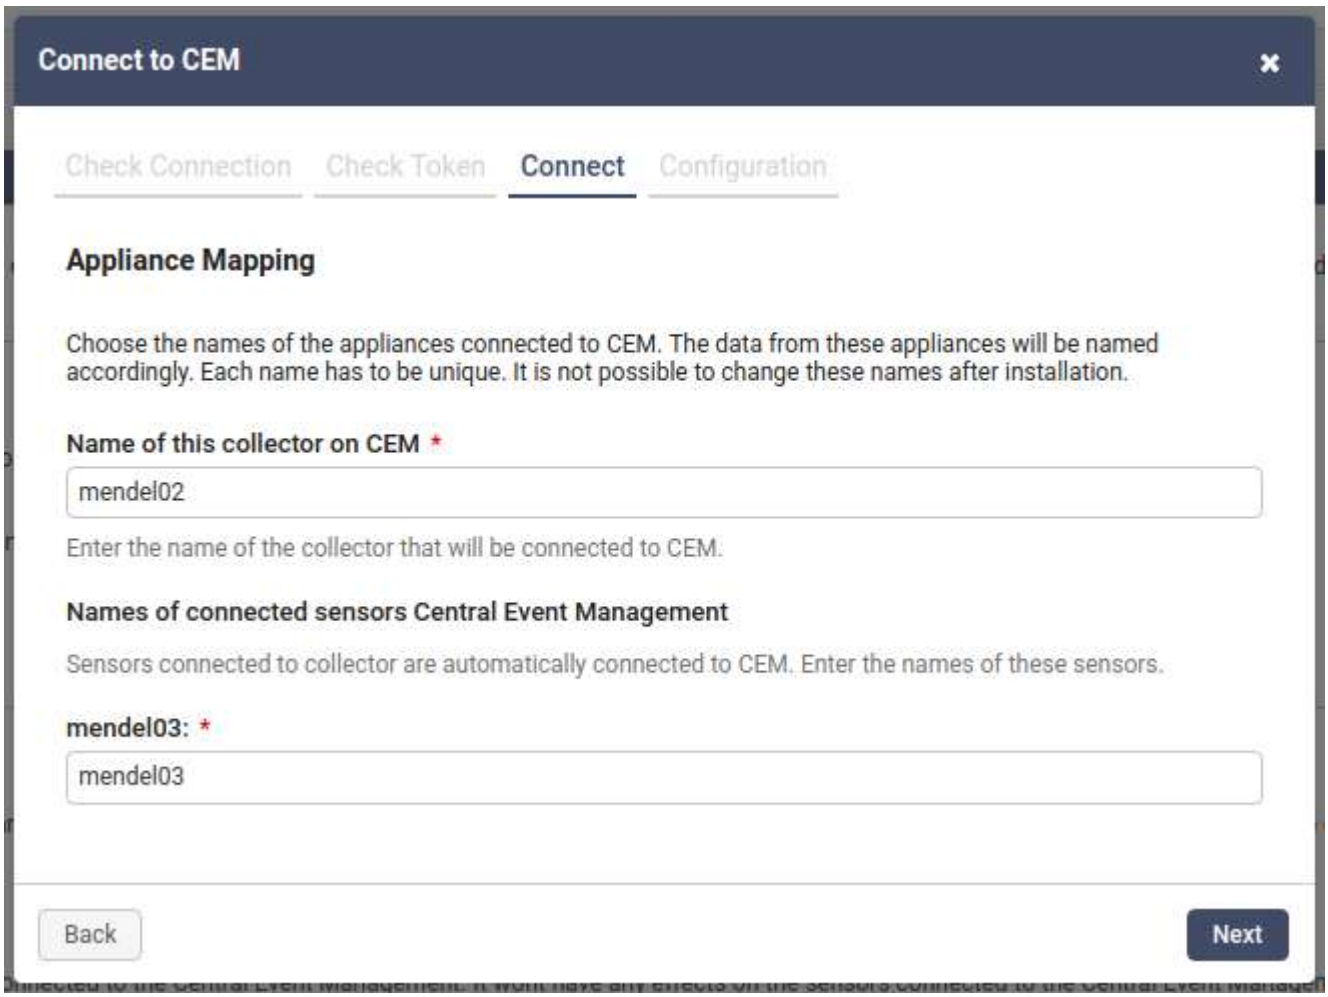
Step 5. Insert the token to the collector (e.g. mendel02)

Now when you have connection token, you need to go back to the UI of the other appliance (future Collector). And add the token and click on button.

Step 6. Appliance Mapping

In this step, you need to set up the mapping of the devices. That means the names of the appliances which you will see in the graphical interface of the CEM. When you are connecting architectures that have more devices as Sensor-Collector or Multitenant you have to name all the appliances here.

It can have the same name that has independently on the CEM cascade. That is the default value. So in our case, it is mendel02 (collector has the same name) and mendel03 (sensor has the same name).



Mapping of the appliances is really important mainly in deployment architectures as Sensor-collector or Multitenant where you have more appliances connected together.

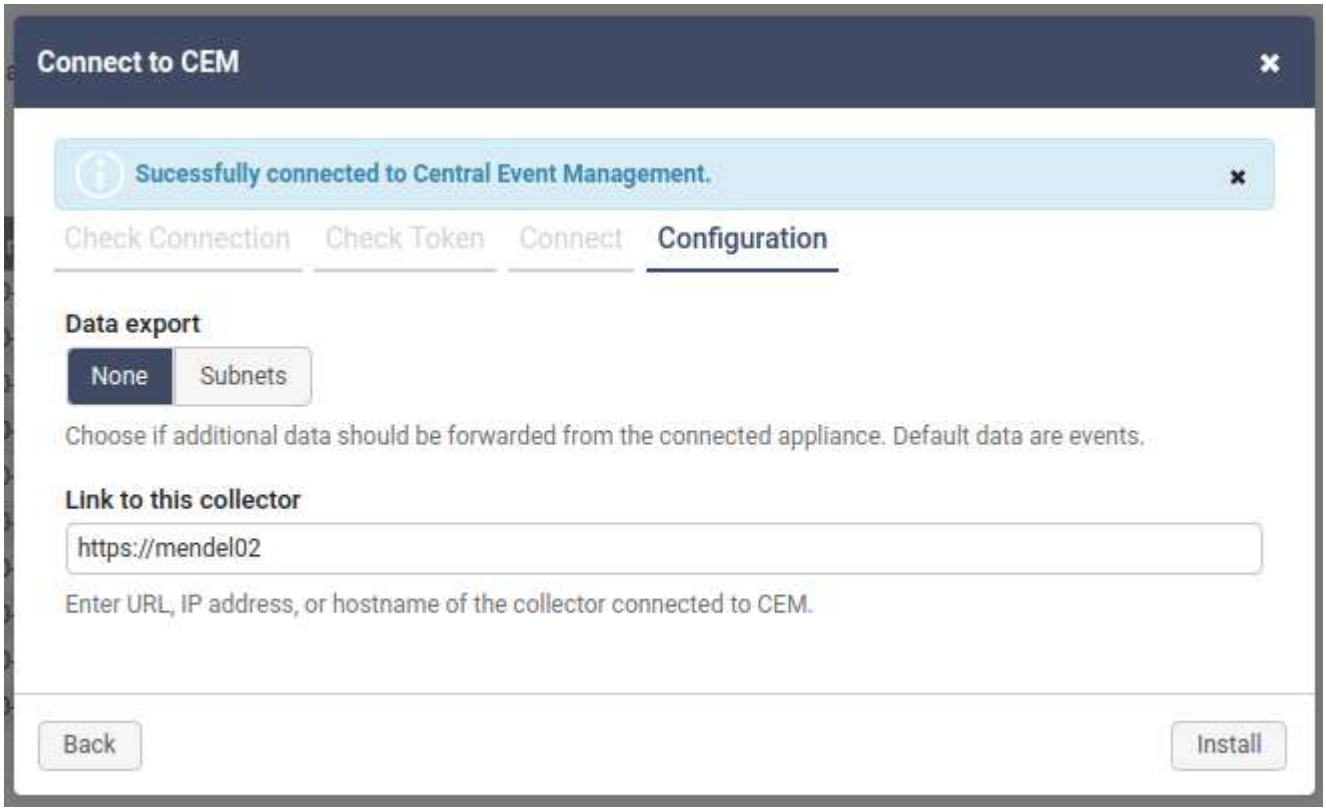
Copyright © 2024 Greycortex

The problem could come when appliances in your network will have common names as in this tutorial (mendel02 and mendel03). When you have more appliances named the same and want them to see it together in CEM, you need to map (rename) them.

Step 7. Additional configuration

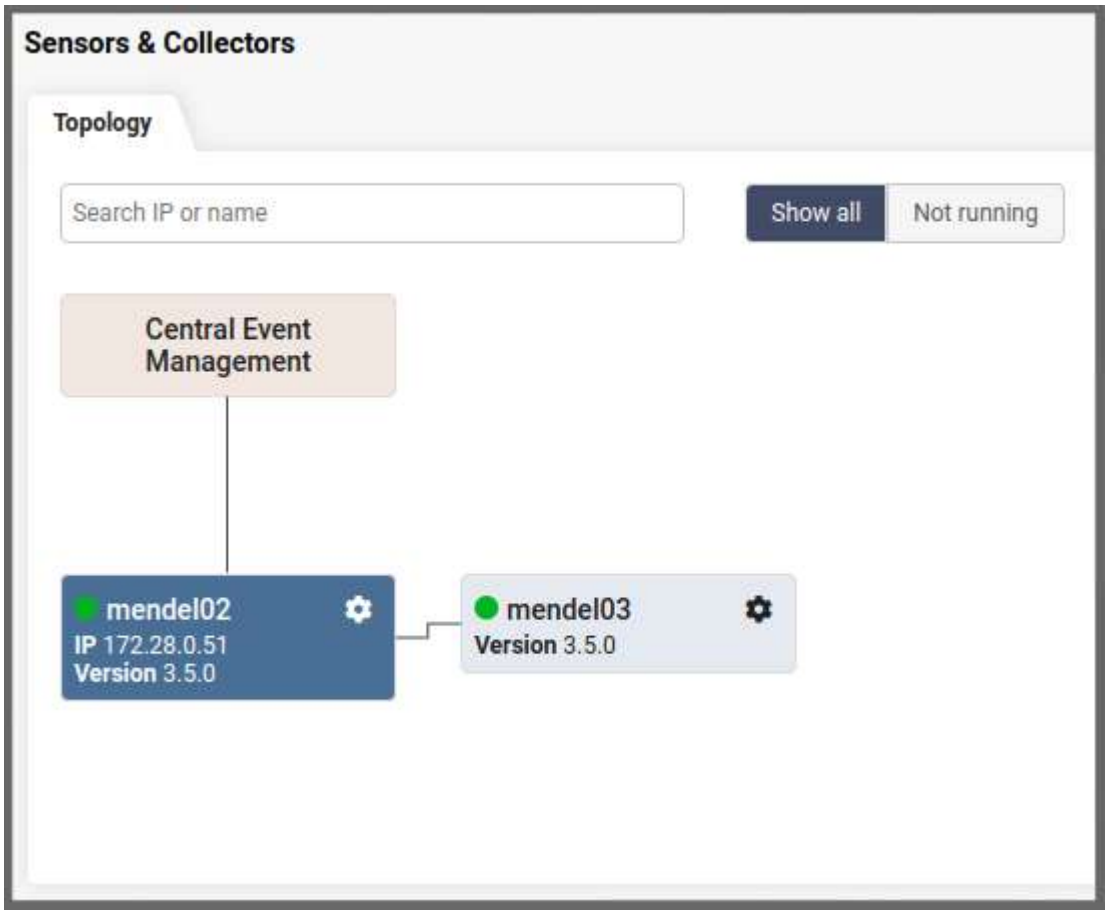
The last step of the CEM connection is and additional configuration.

- You can choose if you want to send **subnets** to the CEM or none of the additional data export. By default, you can see just **events** of the connected appliances.
- You can define here a **link to the collector** so you can be redirected to this address from the user interface of the CEM.



You can check the success of the installation on the collector in . It opens the topology tab.

Copyright © 2024 Greycortex

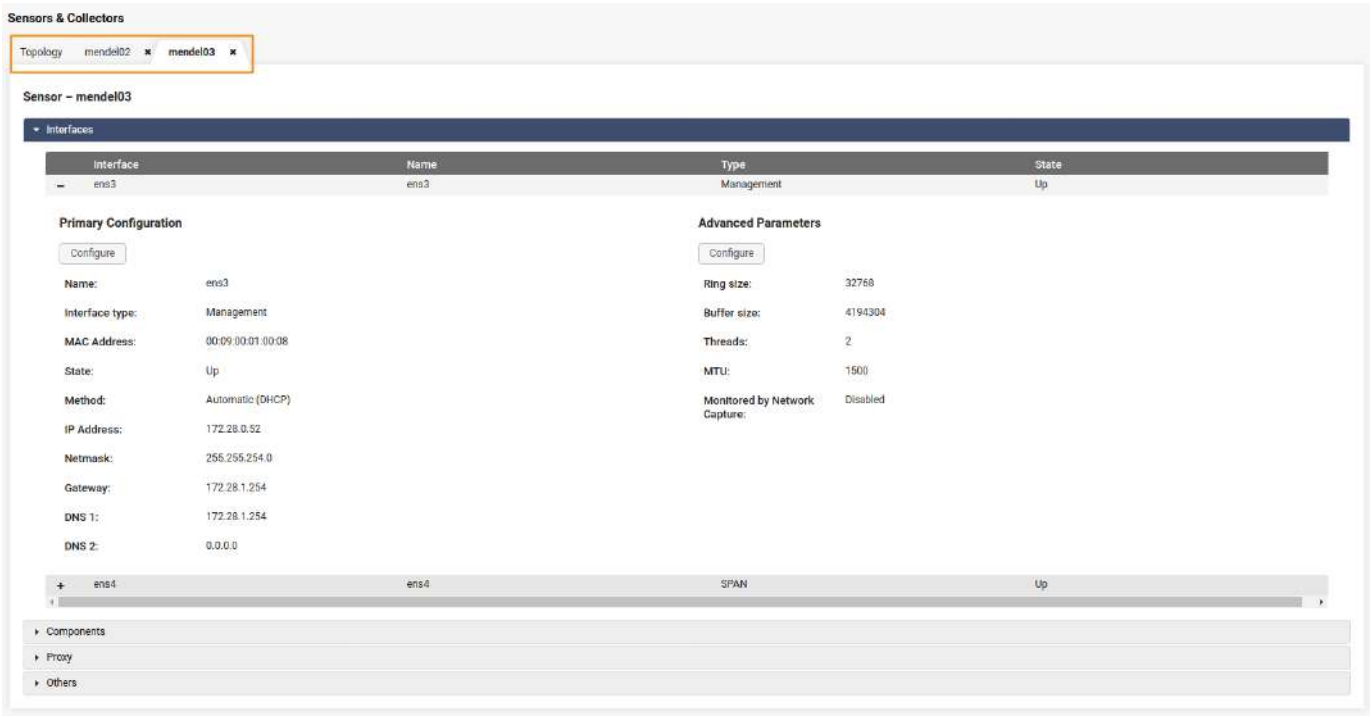


4.3.7 2. Edit settings of the appliances

If you want to edit the settings of the appliances, you have to go to the collector and change it from there. In the CEM there is not an option to change any different settings than itself.

All settings are now available after clicking on the in the right corner of the appliance in the topology tab. It will open a new tab next to the Topology. So you now can control the sensors and the collector in one place.

Copyright © 2024 Greycortex



4.3.8 3. Disconnect the collector from the CEM

Open graphical interface of the collector that is connected to the CEM.

Go to .

Check the part and there is a button.



After clicking on it, there will be confirmation dialog with options. Hit and collector is no more in CEM.



You can check the success of the disconnection on the collector in .

Copyright © 2024 Greycortex



Copyright © 2024 Greycortex

4.4 Dashboard Manager

Dashboards can be added, modified or deleted via the Dashboard manager. The dashboards are divided into different sections that allow for a different breakdown of columns in ratios.

Dashboard Manager

The Dashboard manager allows the following functions:

- - Allows you to add, modify, delete and change the order of the whole dashboards.
- - Add a new section to current dashboard .
- For each section to:
- Add component -- Add Charts or Top-N Reports to the selected section.
- Change the section layout.
- Remove the components or whole sections.
- Reorder components or whole sections in case there is more than one. You cannot move a component to another section.
- - Ability to export chosen dashboards and than import them.
- - Revert all changes to factory state.

Modifying Components

Individual components show specific data about security events or network statistics from multiple sensors. Each component allows the following functionality:

Action	Description
	Using quick filters allows components of the dashboard to be filtered according to chosen parameters, modified by time using the calendar, scrolled using a mouse wheel while hovering over the components, or modified by clicking and dragging.
	For the most components can be permanently applied a predefined or user-defined filter. These filters are then indicated in the title of the component.
	Most views allow a direct click on the detail of the selected data.
	Reordering components within the section. Components cannot be moved to different section nor dashboard.
	By hovering above the title of each component, the basic explanation of what the component displays is shown.
	Optional menu for more functions related to each component: • Delete component -> Removes the selected component from the section. • Settings -> Several options for components can be set in the settings like User title, Font size, Height of the component, Data series etc. The options in settings are dependable on the type of the component. • Export data -> Only table components have this option. It is possible to export from the table to the .csv file. It will export rows and columns that are visible in the component.

Copyright © 2024 Greycortex

4.5 Exports & Integrations

4.5.1 Email Alerts

You can set MENDEL to notify you with an email if an event of a particular type or severity has occurred.

Setting up email alerts

> **Important**
>
> At first, you have to set parameters for the external server (SMTP) through which email notifications will be sent.
> This procedure is described in Email settings.

Email settings

1. Go to the main page by clicking on the **GREYCORTEX MENDEL** logo in your top left corner.
   - Click on in your bottom left corner.
   - Enter a **Filter name**.
   - Select an **Event severity** you want to be informed of.
   - Click on button in left bottom corner and afterwards click on again.
   - Click on button.
2. Go to:
   - Click on .
   - Enter a **Name**.
   - Select an **Email** in **Type**.
   - And choose which **Interval** you want.
   - Click on .
3. Click on (plus) icon.
   - Add a new email data export with desired sending interval.
   - Select the predefined filter.
   - Select your time zone.
   - In **Exported Columns** you can choose which attributes will appear.
   - Click on checkbox under section in order to enable an export.
4. From this point forward, you will start receiving an email when the event is triggered.

Copyright © 2024 Greycortex

4.5.2 Firewall Integration

Firewall Integration

GREYCORTEX currently offers the capability to integrate its MENDEL network traffic analysis solution with existing customer firewalls. This integration allows users to block (or use other supported firewall command) communication vector using subnets, IPs, services, protocols, and/or ports via the firewall, directly from the MENDEL UI.

How to set up the Integration

The firewall integration is done by adding a firewall plugin profile in:

1. Click on the button and afterwards select one of the available firewall plugins.
2. Fill in all the necessary information

Attribute	Description
	Plugin profile name
	Custom description
	Defines maximum number of execution of the script per minute
	Adjustment specific for individual plugins

3. Click on the button to finalize plugin profile.
4. After you add at least one plugin profile you are able to start using the button in the Events Tab.

Hint

- To create plugin profile you have to enrich MENDEL with particular plugin/script.
- To perform this task please follow Plugin Development

Copyright © 2024 Greycortex

System plugins (supported by default)

Practical Example (How to take an action)

- If an analyst determines that a specific communication vector should be blocked on the firewall, he can take an action. This can be done directly from an event with the button.
- This action can create a specified firewall rule on the specified firewall and thus can block the nefarious communication. Once the rule is defined and the user clicks on the button, a new rule will be created on the specified firewall with these parameters.

Copyright © 2024 Greycortex

Integration with other firewall solutions

- Because we do not integrate with all firewall solutions out of the box, MENDEL has the capability to use any plugin script, if the API of the firewall solution in question supports it. This means that you can write your own custom integration script.
- However, before you can add your custom integration plugin into MENDEL, you have to send it to us for revision. If there are no compatibility issues, we will sign it and return it to you, so you can upload it to MENDEL and start using it. MENDEL will not allow you to use custom unsigned plugins.

Copyright © 2024 Greycortex

4.5.3 Plugin Development

GREYCORTEX currently offers the ability to integrate its network traffic analysis solution, MENDEL, with existing customer firewalls. This integration, for instance, allows users to block communication coming from and going to subnets, IPs, services, protocols, and/or ports via the firewall, directly from the MENDEL UI. Apart from blocking, other standard firewall commands are also supported.

The system includes some default scripts for common firewalls, on the other hand, it also allows system integrators to create their own firewall plugins. Custom plugin scripts can allow GREYCORTEX MENDEL to interact with external devices based on their access credentials and selected attributes from the event.

A custom plugin is a simple file written in supported scripting languages (Bash, Python or Expect). The script has an interface that defines the way of communication with MENDEL. The script accepts command line parameters and returns exit code. When the script is ready for production, it has to be verified and signed by GREYCORTEX, and only then it can be uploaded into the selected MENDEL appliance and properly configured.

PLUGIN INTERFACE

An essential part of every plugin is a plugin interface. The interface itself is defined in the JSON format and has to be written in standard output when running a plugin script with the "--interface" argument:

```
```

The interface output format is the following:

```
```

The attributes version, name, and description define the whole plugin and its purpose. The attribute params include all custom parameters which will be defined in the user interface when the plugin will be configured. It should be used to provide access credentials for the plugin - IP address, username, and password. Each parameter has to contain these attributes:

Attribute	Description
	The name of the parameter, typically "username" or "password".
	The type of parameter, "text" for plain text and "password" for password fields are currently implemented.
	Description to be shown when setting up Plugin profile in Mendel UI.
	Makes the parameter required and cannot be empty when setting up Plugin profile in Mendel UI.

PARAMETERS

The plugin will be executed with the following parameters:

```
```


Copyright © 2024 Greycortex

Parametr	Description
	IPv4 or IPv6 address.
	IPv4 or IPv6 address.
	IPv4 or IPv6 subnet address using CIDR format.
	IPv4 or IPv6 subnet address using CIDR format.
	source port in the range from 1 to 65535.
	destination port in the range from 1 to 65535.
	service abbreviation (HTTP, SSH).
	IP protocol number in the range from 0 to 255.
	a list of all parameters defined in plugin interface.

All parameters are separated by space and empty parameters have zero length.

Return value

Plugin will return value of 0 for success and 1 for error. In case of an error, the system will read the plugin standard error output and will display an error message in MENDEL UI or in the description of the Status monitor event related to the plugin execution.

SAMPLE PLUGIN

Below, you can see an example of a firewall plugin. This plugin only disallows IP protocol 0 (HOPOPT) with an error message but passes if IP protocol is not set to 0:

The script can be tested by running this command:

Copyright © 2024 Greycortex

Parametr	Description	Example
	Source IP detected on event.	172.0.0.1
	Protocol from detected event.	6 for TCP
	Router IP.	192.168.1.1
	Router password.	User
	Router password.	Pass

Copyright © 2024 Greycortex

4.5.4 API

EXTERNAL API

Higher **integration** potential with a new ability to connect with external information sources or provide processed data to recipients (SIEMs, etc.) for further processing.

External (public) API in Mendel is versatile way of integration based on a API architectural style.

API COVERAGE

Currently Mendel is able to provide these typ of data through API:

- Events
- IDS and log processing signatures, IDS variables
- Malicious domains
- Data captures (a direct connection into the database where all captured network data is stored)
- False positive management
- Blacklists based on IP addresses (including MISP)
- Malicious Files

API TECHNICAL DESCRIPTION

- Mendel External API is an RPC API for getting information from without using Web UI directly. It can be used by external systems and scripts.
- JSON format is used for transferring data structures. Design follows REST service conventions.
- Authorization is HTTP Basic with **ExtApi Key** used as and **ExtApi Secret** used as . ExtAPI Key and Secret may be generated in the Mendel web application. It is always generated under some user and inherits user's permissions.
- Clients must accept at least application/json content type for all calls. Even if the requested method returns some other content, errors may be returned in JSON.

HOW TO ACCESS API

Preconditions

To be able to access Mendel External API you need to have credentials.

They are represented by

- ExtAPI Key (username)
- ExtAPI Secret (password)

and you can obtain them in UI of the appliance in

> **Note**
>
> You can generate more than one API key for every user (account) you need to have an access to the API.

> **Important**
>
> Make sure to copy and save this API key and secret. You will not be able to access secret after closing this dialog

Copyright © 2024 Greycortex

API Access

API can be accessed with the URL address of the destination appliance with ME running

API Documentation (raw)

There is an OpenApi v3 description of ExtAPI on URL

API Documentation (structured)

There is also visualized interpretation in an interactive UI generated by Swagger on URL

REQUEST EXAMPLE

Curl

```
curl -X 'GET' \
  'https://melanie.greycortex.com/extapi/v2/detection/captures/1' \
  -H 'accept: application/json'
```

Request URL

```
https://melanie.greycortex.com/extapi/v2/detection/captures/1
```

Server response

Code	Details
200	Response body


```
{
  "id": 1,
  "name": "172.16.27.176",
  "groupId": "support",
  "sizeLimitBytes": 20000000,
  "sensorName": "melanie",
  "srcHostAddr": "172.16.27.176",
  "enabled": true,
  "created": "2020-09-24T16:12:00.086+02:00"
}
```


Download

Copyright © 2024 Greycortex

4.5.5 Active Directory Configuration

In order to configure Active Directory (AD) integration, the first and most essential step is to properly set up the AD user which will be able to read security logs from your Domain Controller (DC). That way, MENDEL will be able to pair the log-in and log-out events of users with the respective hosts and events. This can be done with the help of our guide Active Directory Setup Guide.

Active Directory Configuration Guide

In order to configure Active Directory (AD) integration, the first and most essential step is to properly set up the AD user which will be able to read security logs from your Domain Controller (DC). That way, MENDEL will be able to pair the log-in and log-out events of users with the respective hosts and events. This can be done with a help of our guide **gcx_active_directory_setup.pdf**. You can find this document on our partners section.

INTEGRATION SET UP

In order to set up the integration, go to:

1. Click on the **Add** button to enter the wizard 2. Fill in all the necessary/optional information

Copyright © 2024 Greycortex

Attribute	Description
	Custom name for your AD integration profile.
	One or more IP address(es) or hostname(s), you have to confirm each item with enter key. Server priority is defined by order of entered servers from left to right.
	Username used for the integration.
	Password used for the integration.
Advanced	*Optional*
	Choose between NTLM v1 and NTLM v2.
	Users will be assigned to the hosts from these sensor groups.
	How often will be the server used for getting events (in minutes).
	Processing of logout events can increase accuracy of assigning users to hosts. Recommended option: disabled. For more details, please see the in-depth guide.

1. If the entered server is reachable and the provided credentials valid, you will be presented with a confirmation of validity of the provided information and with the option to save the configuration.
2. If you want to configure the LDAP integration (detailed information about detected users) toggle the 'Configure LDAP' slider. - This can be also done after the initial wizard. *(Optional)*
3. Fill in all the necessary/optional information. *(Optional)*

Attribute	Description
	Either LDAP (Toggled off) or LDAPS (Toggled on) will be used.
	Base distinguished name.
	Search filter for users.
Advanced	*(Optional)*
	If the username differs from the Active directory user.
	If the password differs from the Active directory user.
	If port is undefined, default LDAP or LDAPS port is used.

4. Click .
5. You can verify proper functionality of this feature by clicking on the hamburger menu button and selecting 'LDAP search'. If you try to search for a known user in the domain, information about this user should be returned to you.
6. You can customize the additional provided information about AD users in the 'LDAP Fields mapping' option.
7. You can enable/disable the querying of your DC by toggling the slider right next to the hamburger menu button of the desired AD profile.

AD User information

If your settings for both MENDEL and the Windows Server are correct, you will be able to see recently connected user accounts in , and some additional host information in corresponding hosts and events.

Copyright © 2024 Greycortex

4.5.6 Active Directory in Depth

Active Directory In-depth Guide

In this portion of the guide, you can find detailed information about how MENDEL obtains and work with the data from Active Directory (AD).

How MENDEL obtains data from the AD

When the integration with AD is set up and user with permissions to read security logs is set, MENDEL will start downloading the log-in and log-out events. This is done via following wmic query each minute (or other set interval in the Active Directory Settings) where the last interval is selected:

```
```

The following eventcodes are downloaded:
4624 - An account was successfully logged on
4634 - An account was logged off
4647 - User initiated logoff

How MENDEL obtains information about AD users

MENDEL will obtain and retain information about all detected users for 7 days. This information is obtained via ldapsearch query. If the information about the active user is older than 7 days, it will be re-downloaded.

The ldapsearch query:

```
```

What ports must be open in order for the AD integration to work

Source	Src_port	Direction	Dst_port	Destination	Description
Collector	any/tcp	--->	135/tcp and 49152-65535/tcp	The AD server (DC)	AD Integration. Port 135 is used for DCE RPC Endpoint Manager DCOM Service Control and range 49152-65535 is used by DCOM
Collector	any/tcp	--->	389/tcp or 636/tcp	The AD server (DC)	LDAP/LDAPS - to be able to extract information about users

What happens when new user is detected

When a new user is detected, new record is added to the ti.users table and ldapsearch query is issued to obtain information about the user from AD. This user is then paired to the respective host in the ti.hosts table. From the next minute, this user will be tagged in all events which will occur on this host until this user either logs out, or new user rewrites this user in the ti.hosts table. This depends on the used AD mode.

There are 2 modes in which MENDEL can operate.

1. Processing log-off events
2. Not processing log-off events

This mode can be switched in the Settings -> System -> Active Directory -> Active Directory Settings -> Process logout events checkbox.

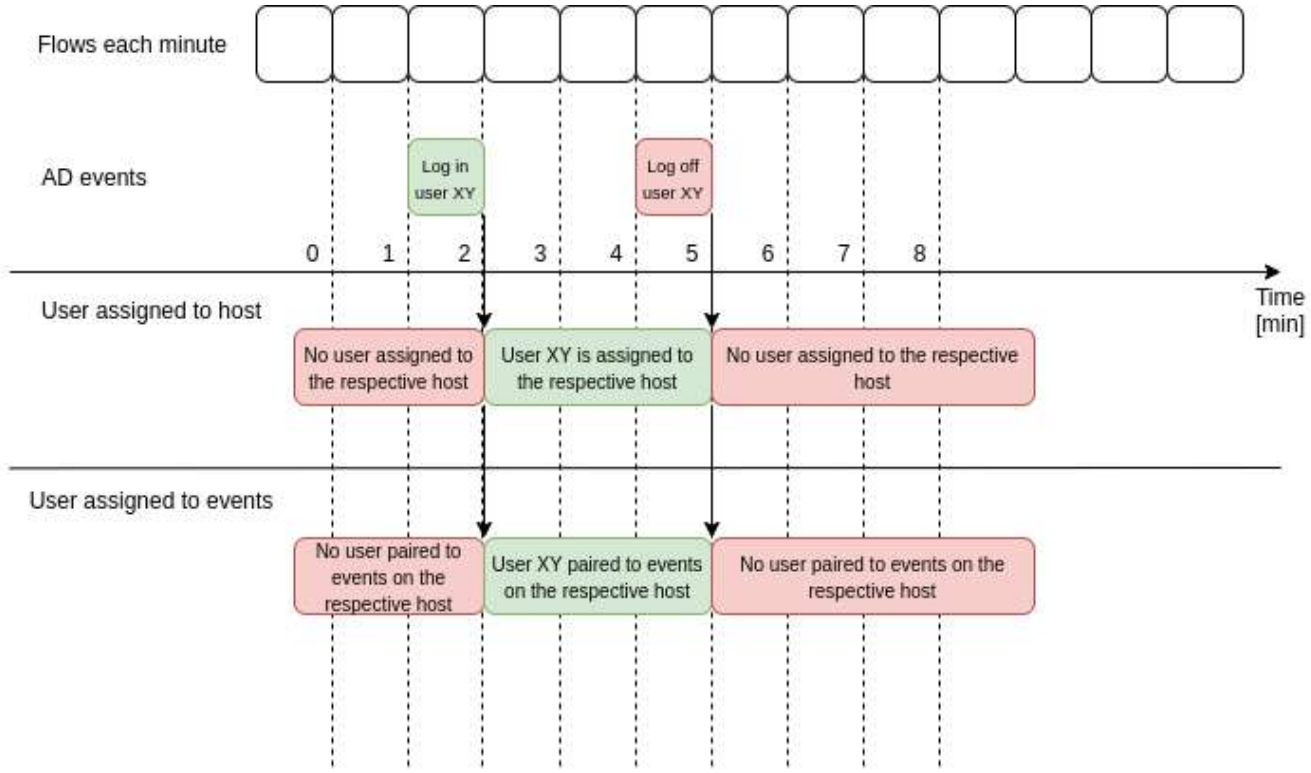
Copyright © 2024 Greycortex



1. Processing log-off events

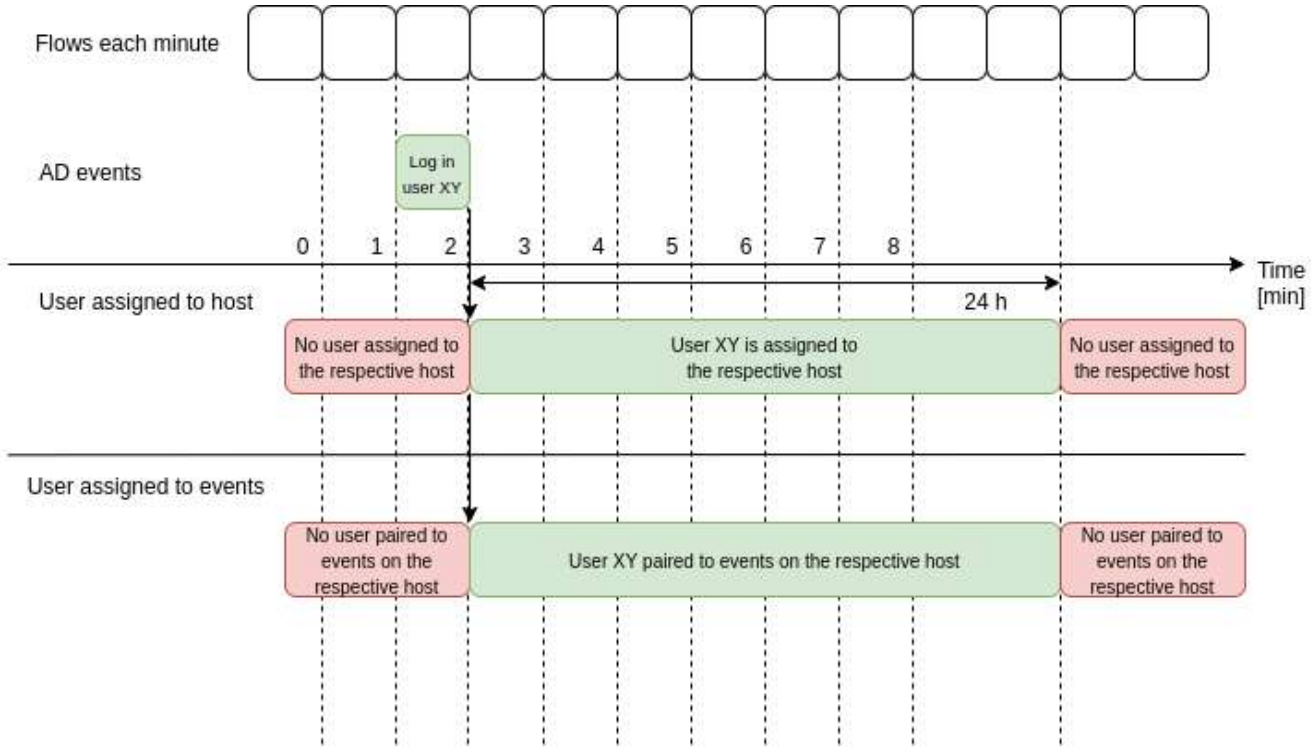
This method will create sessions for each log-in event and will assign user to host until this session expires. The expiration period is 24 hours (You can imagine this as an user lease time). If a log-off event occurs during the user active session period, the session is terminated.

User logs in and logs out:



User logs in but won't log out:

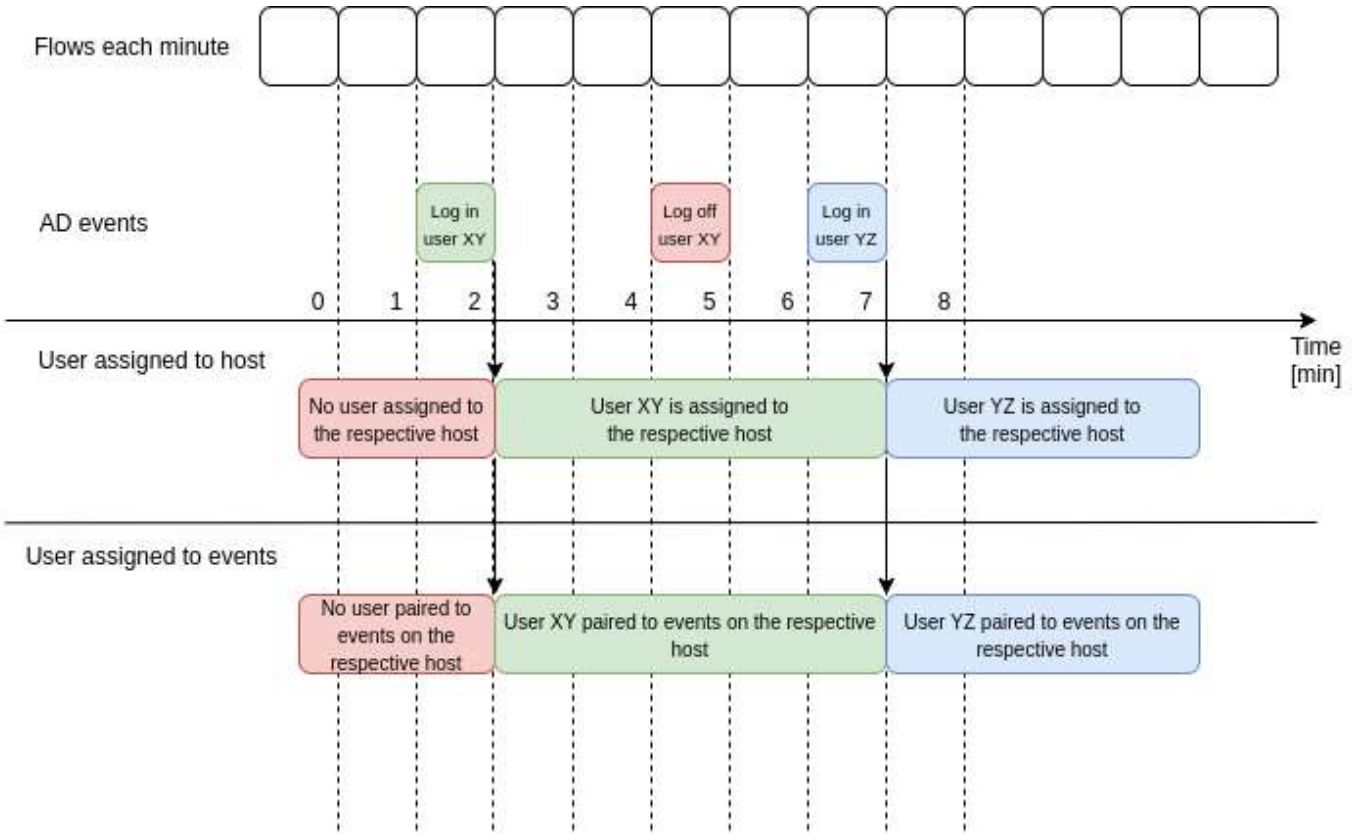
Copyright © 2024 Greycortex



2. Not processing log-off events

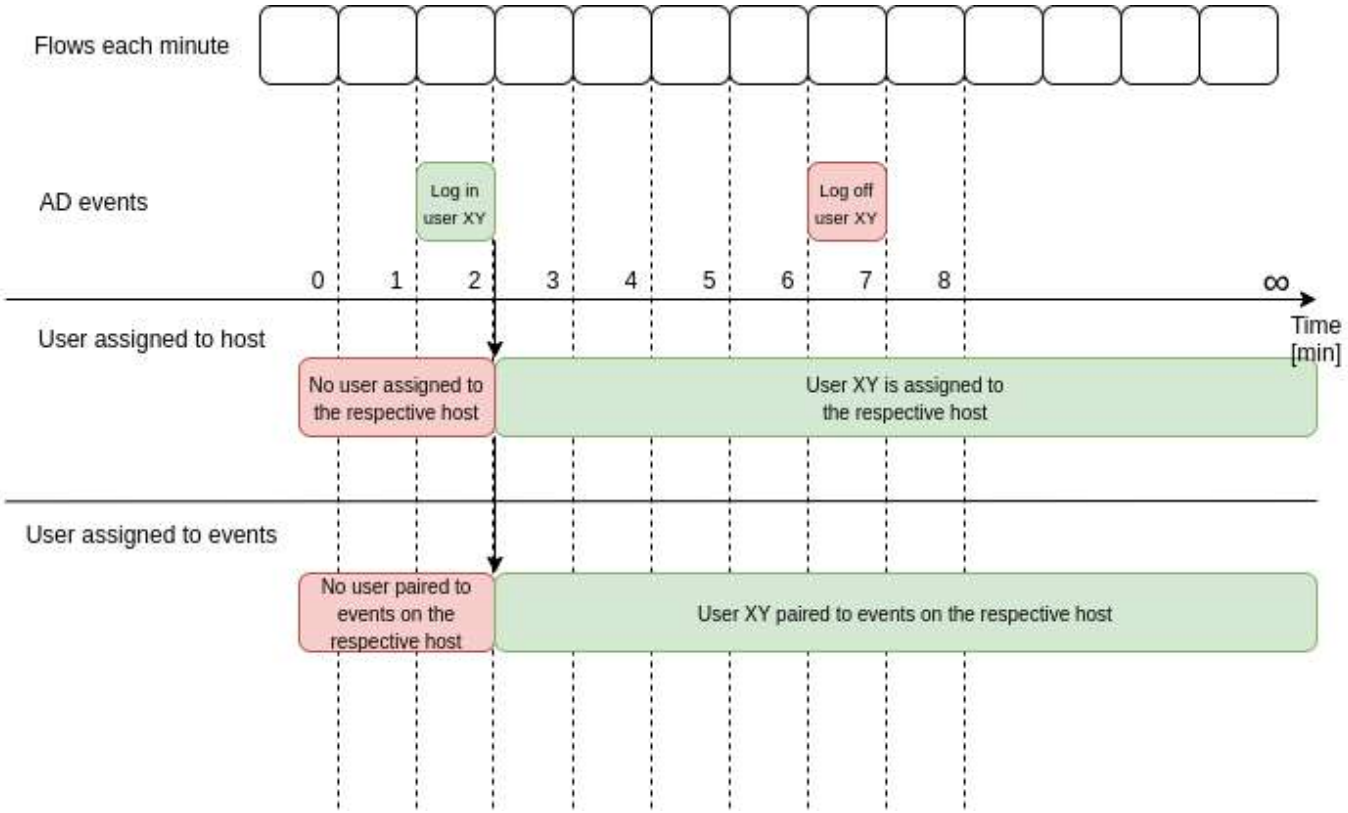
This simplified method assumes that log-in of new user will rewrite the previous user and thus no user active sessions are created.

User logs in, logs out and another user logs in afterwards:



User logs in but won't log out:

Copyright © 2024 Greycortex



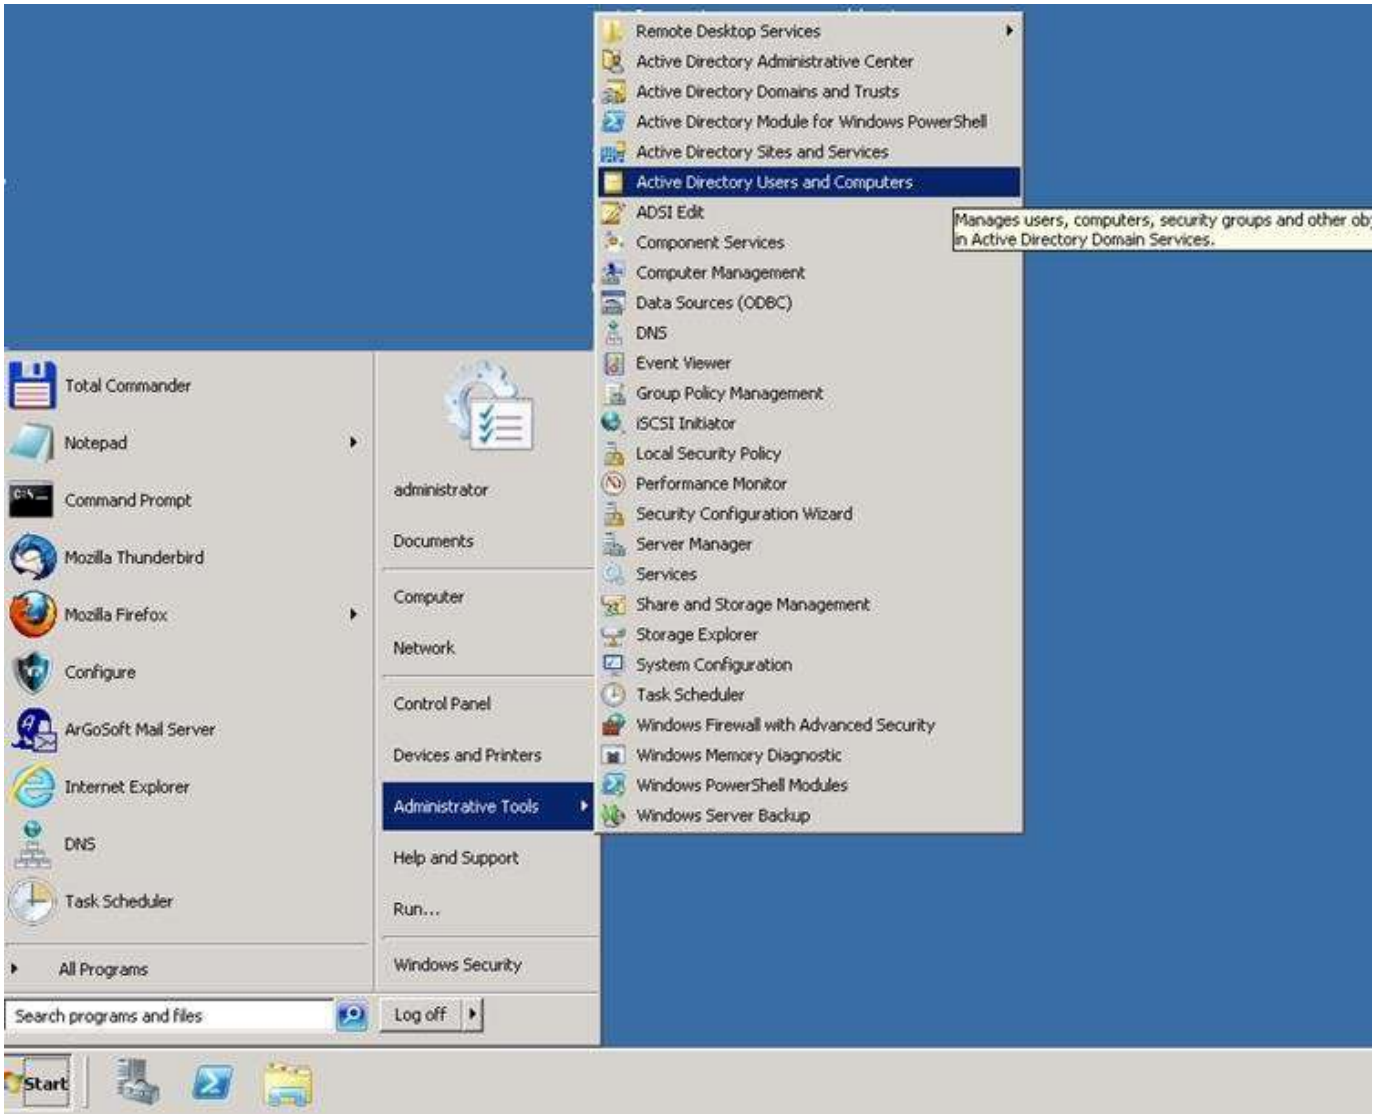
Copyright © 2024 Greycortex

4.5.7 Active Directory Setup

Active Directory Setup Guide

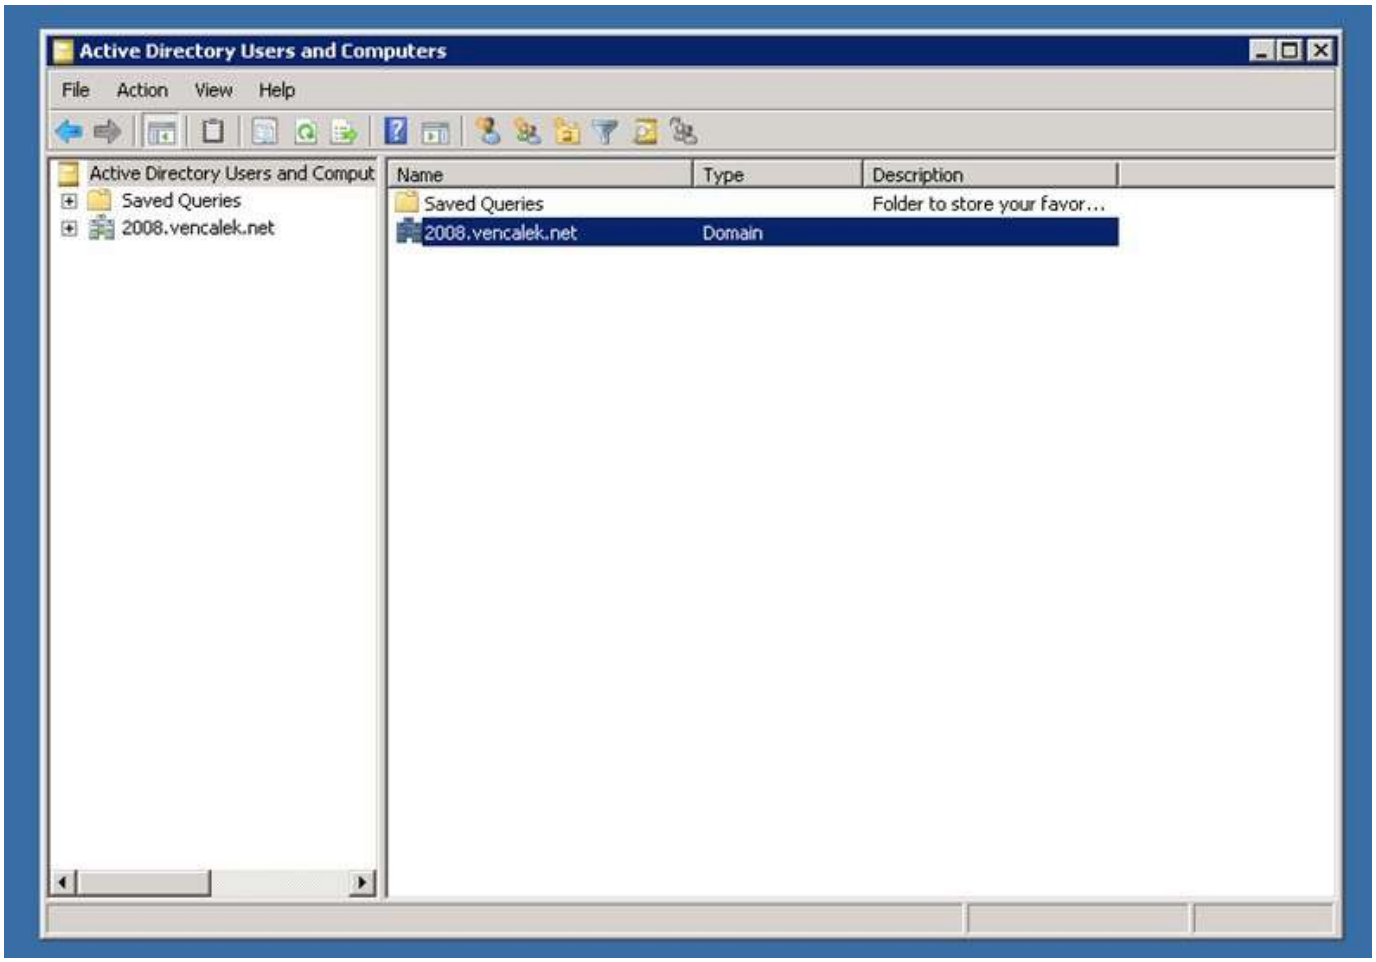
USER CREATION AND ADDITION TO NECESSARY GROUPS

First, create a New User in Start -> Administrative Tools -> Active Directory Users and Computers.



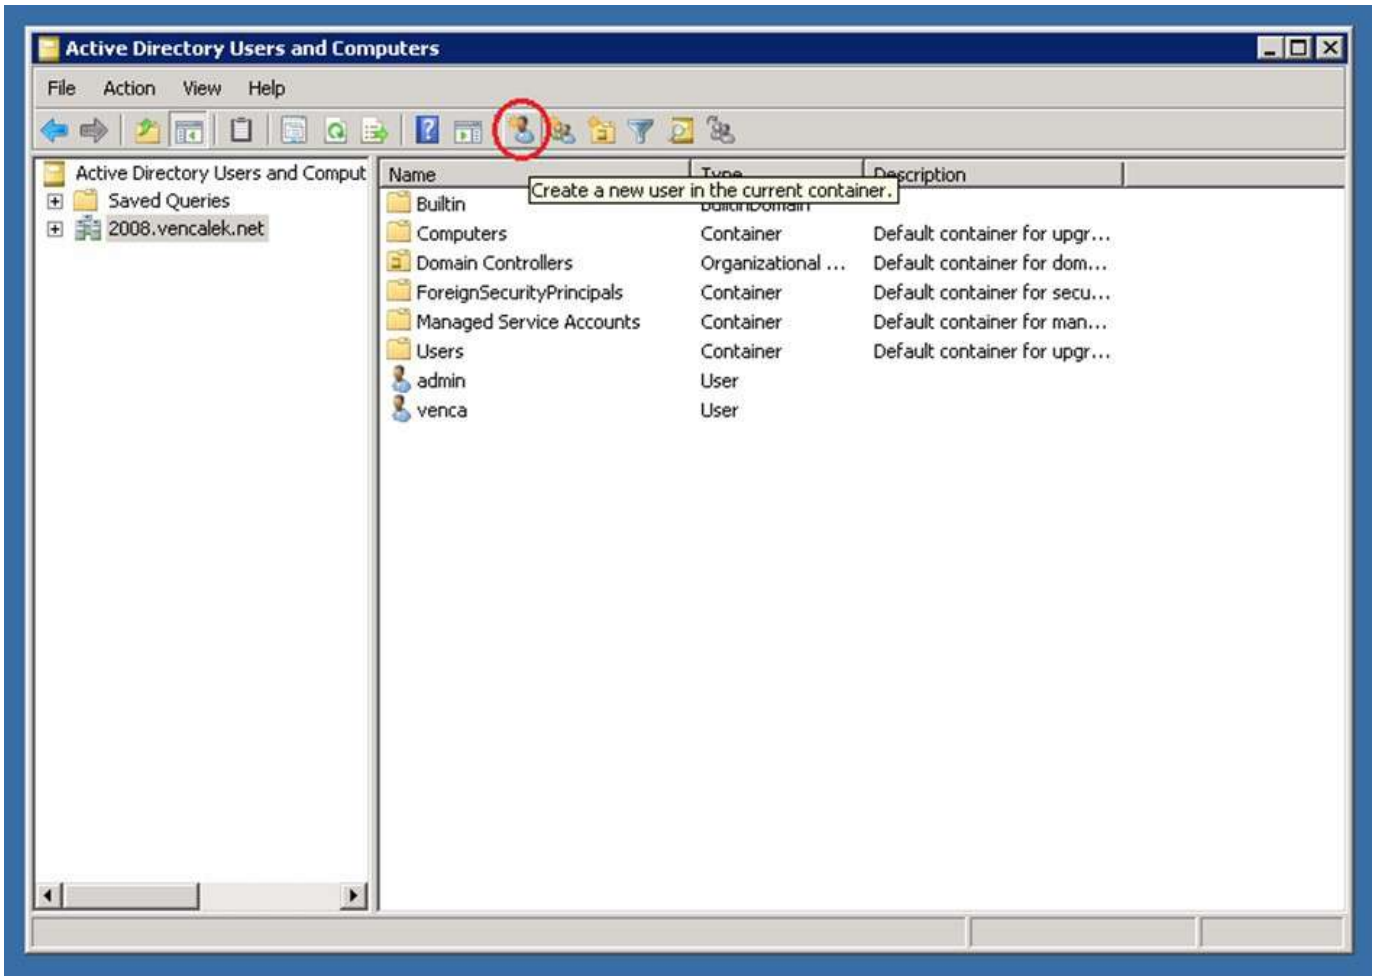
In those settings you will need to select the domain in which the new user will be created.

Copyright © 2024 Greycortex



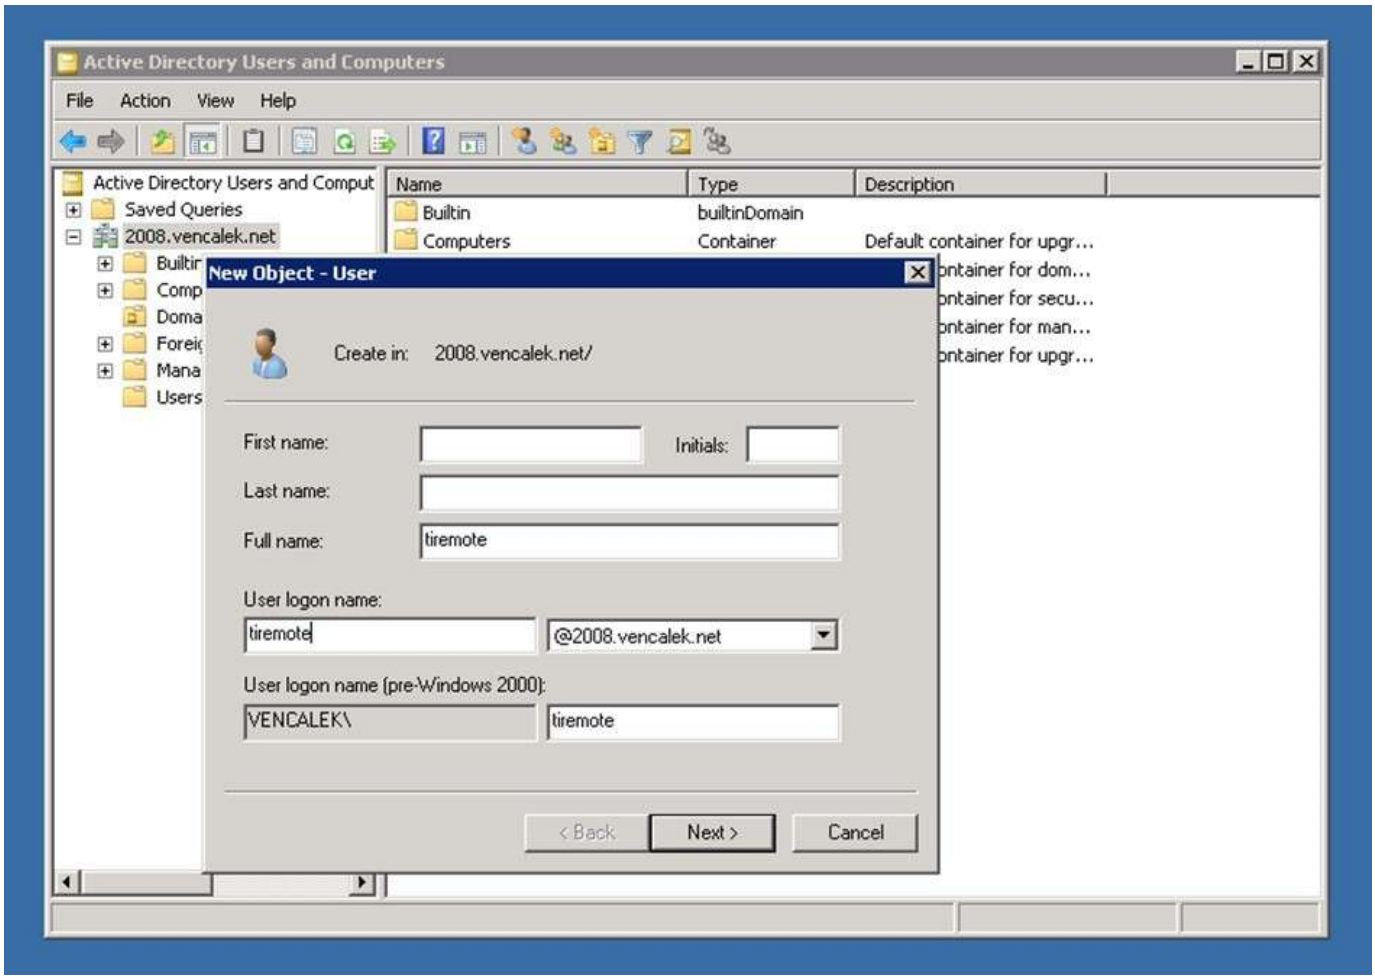
By clicking on the highlighted icon, you will be prompted to click through the dialogue window.

Copyright © 2024 Greycortex



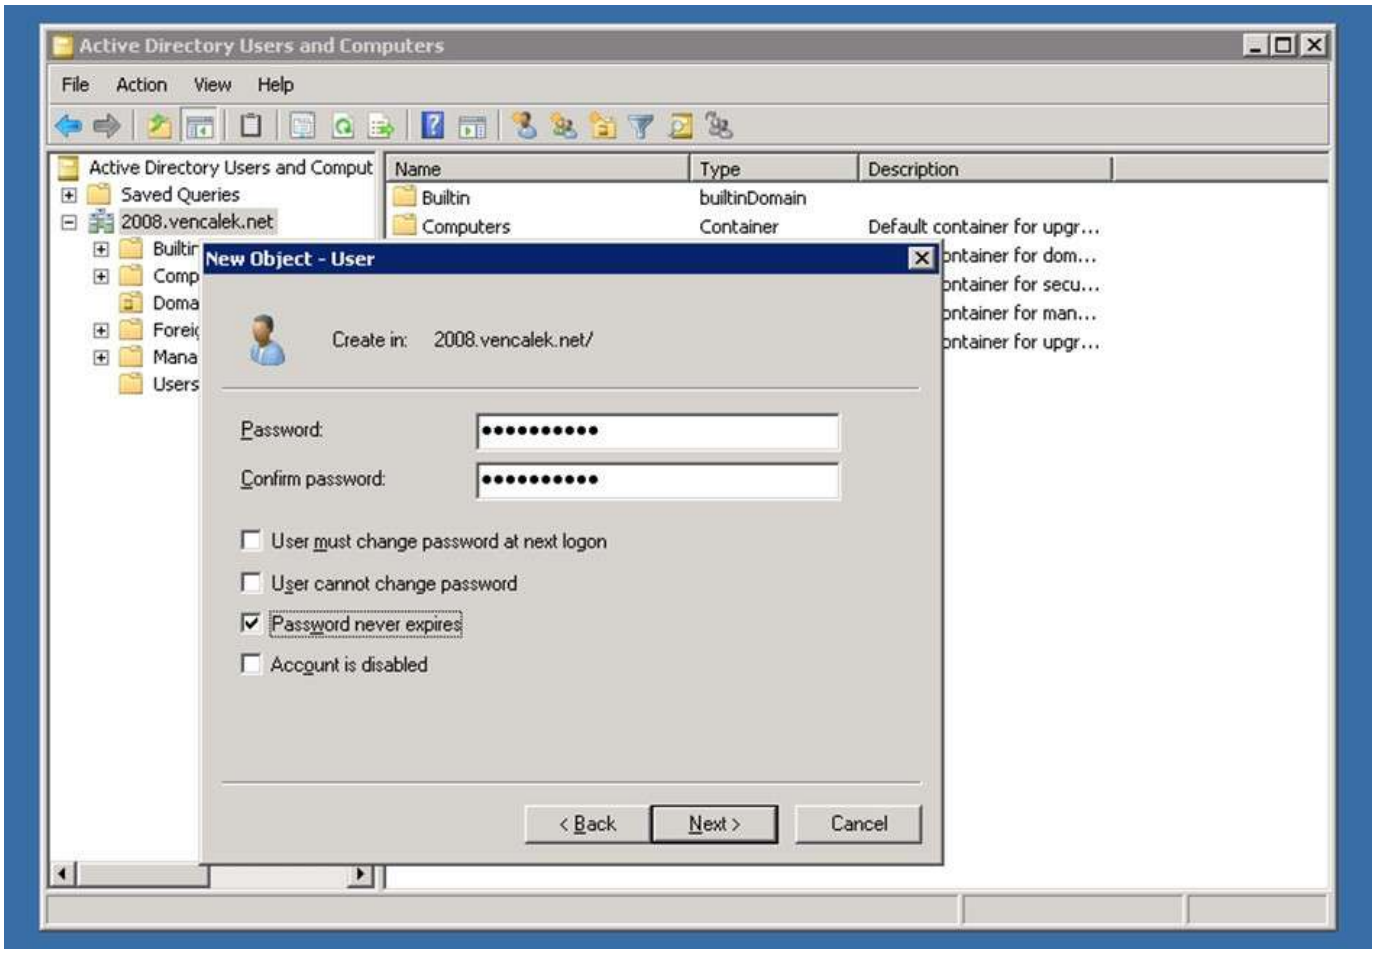
You need to fill in the login credentials.

Copyright © 2024 Greycortex



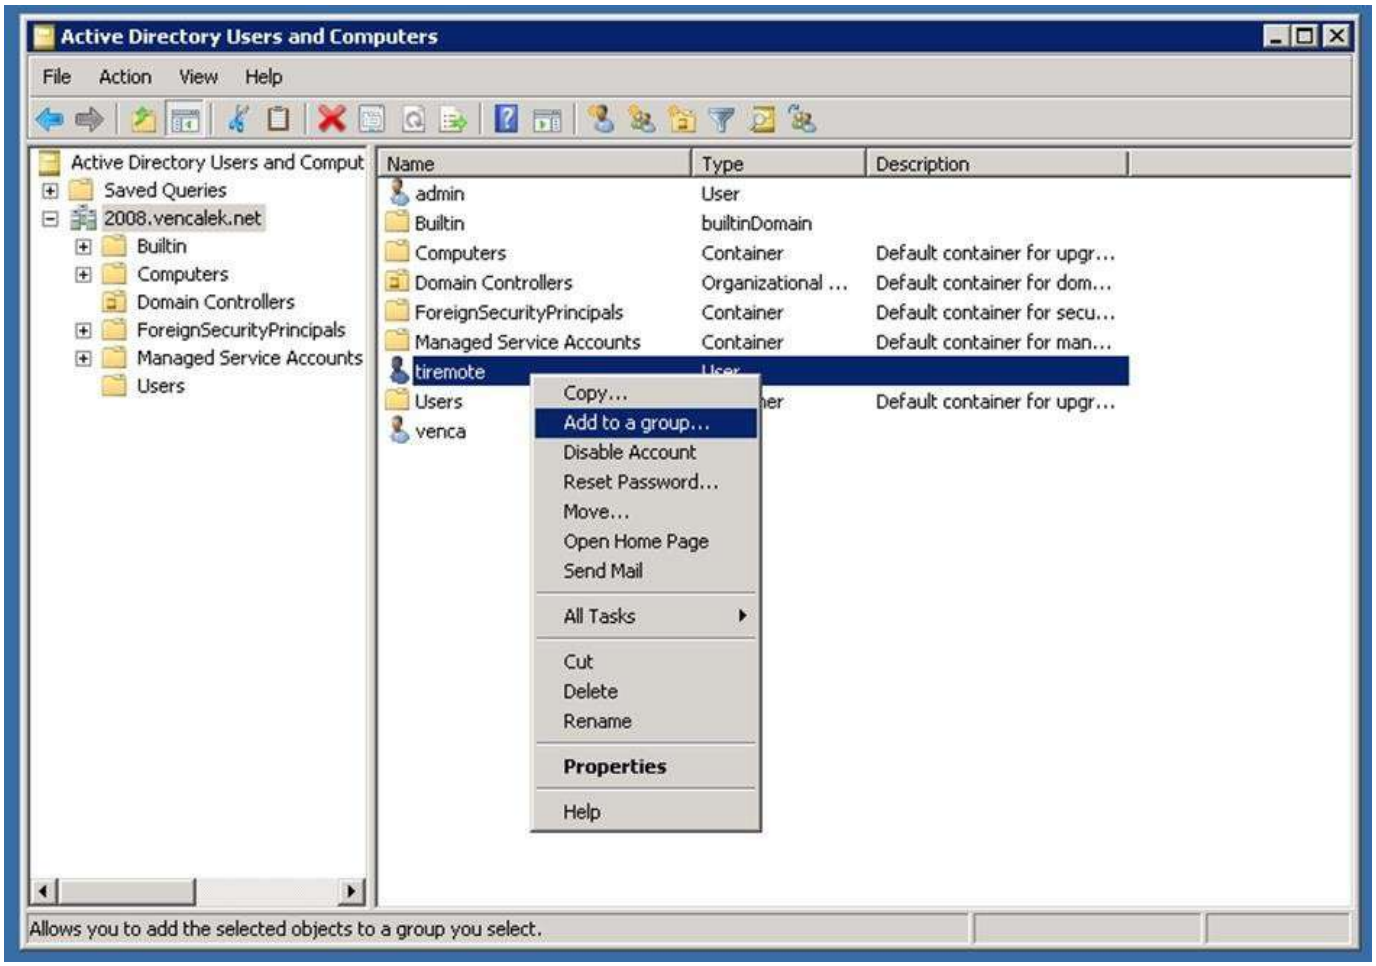
We advise you to select the option “Password never expires” and to unselect “User must change password on next logon.”

Copyright © 2024 Greycortex



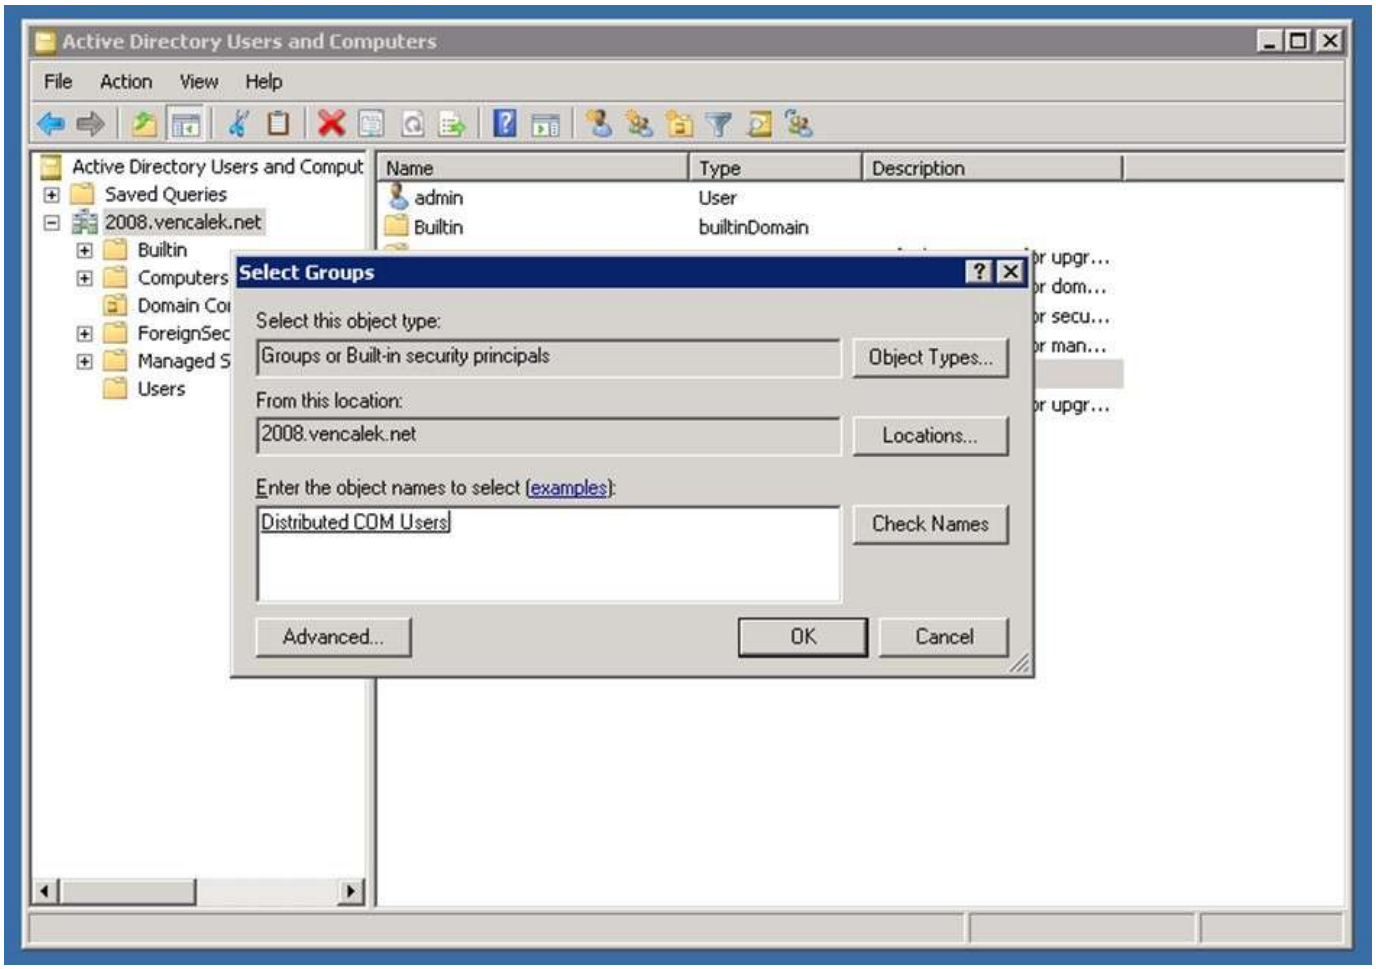
Then add the new user into groups:

Copyright © 2024 Greycortex



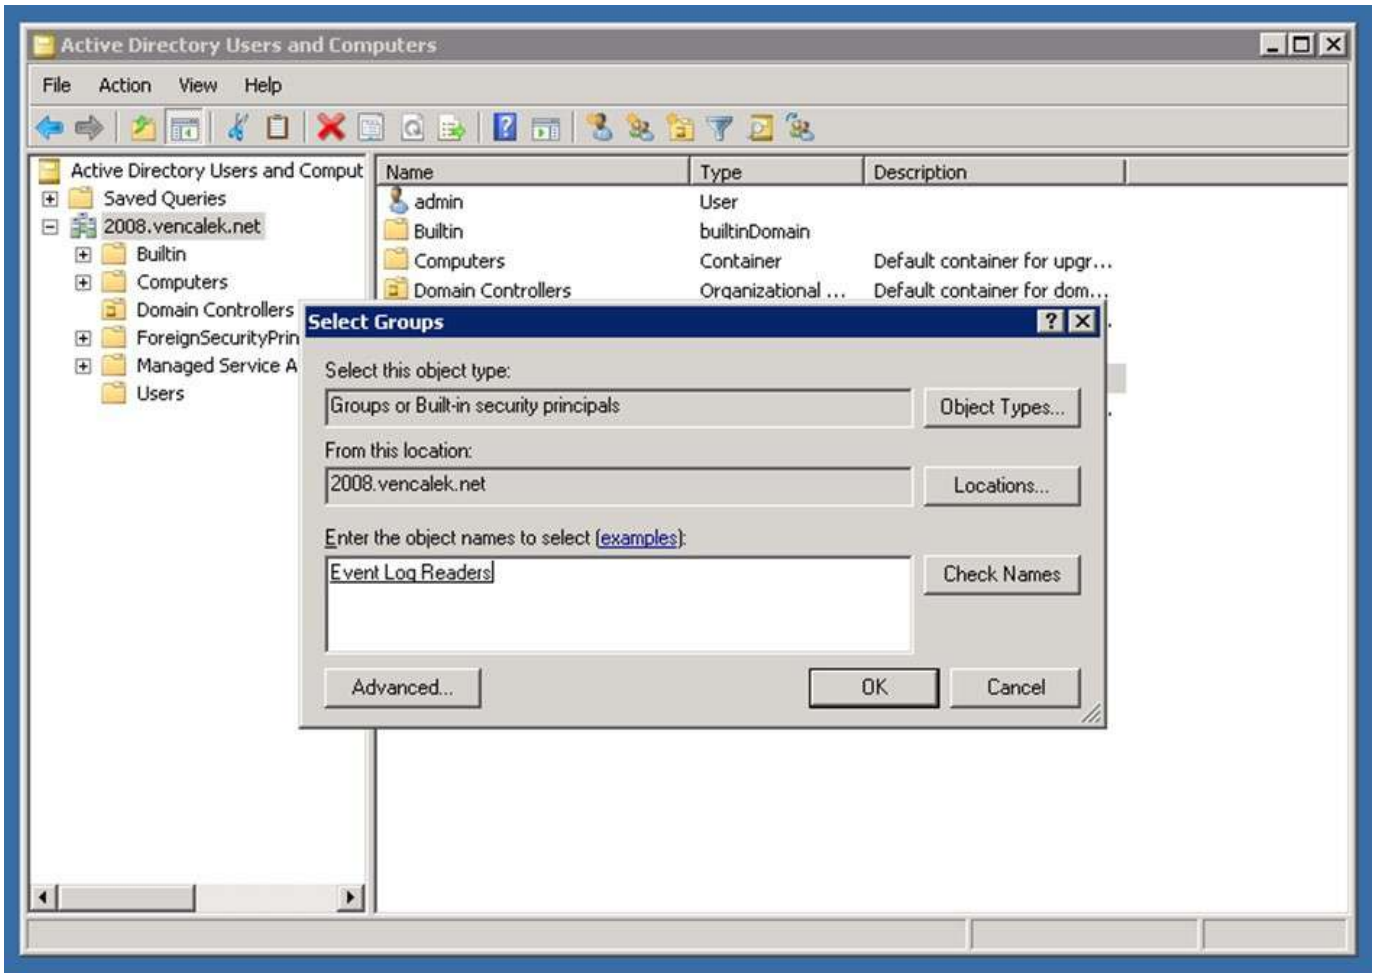
Distributed COM Users

Copyright © 2024 Greycortex



And Event Log Readers.

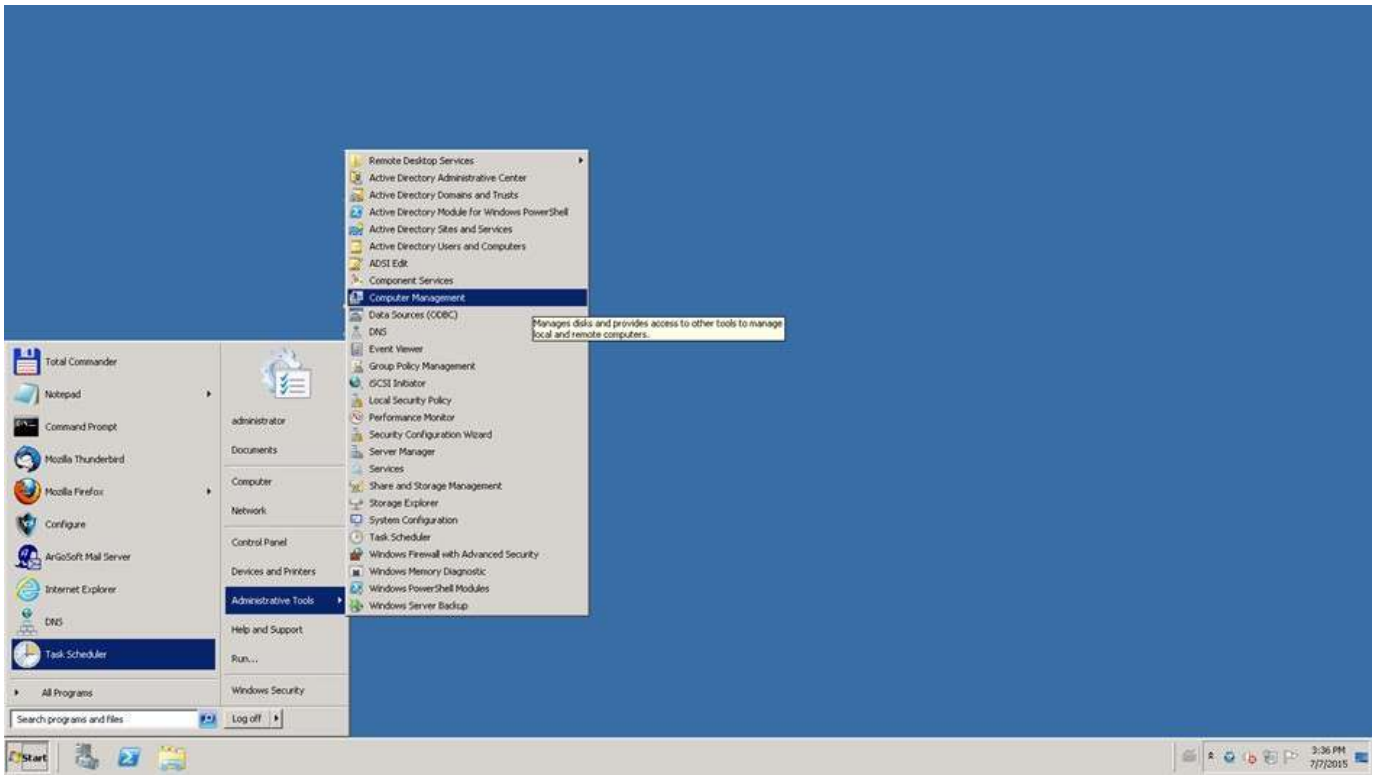
Copyright © 2024 Greycortex



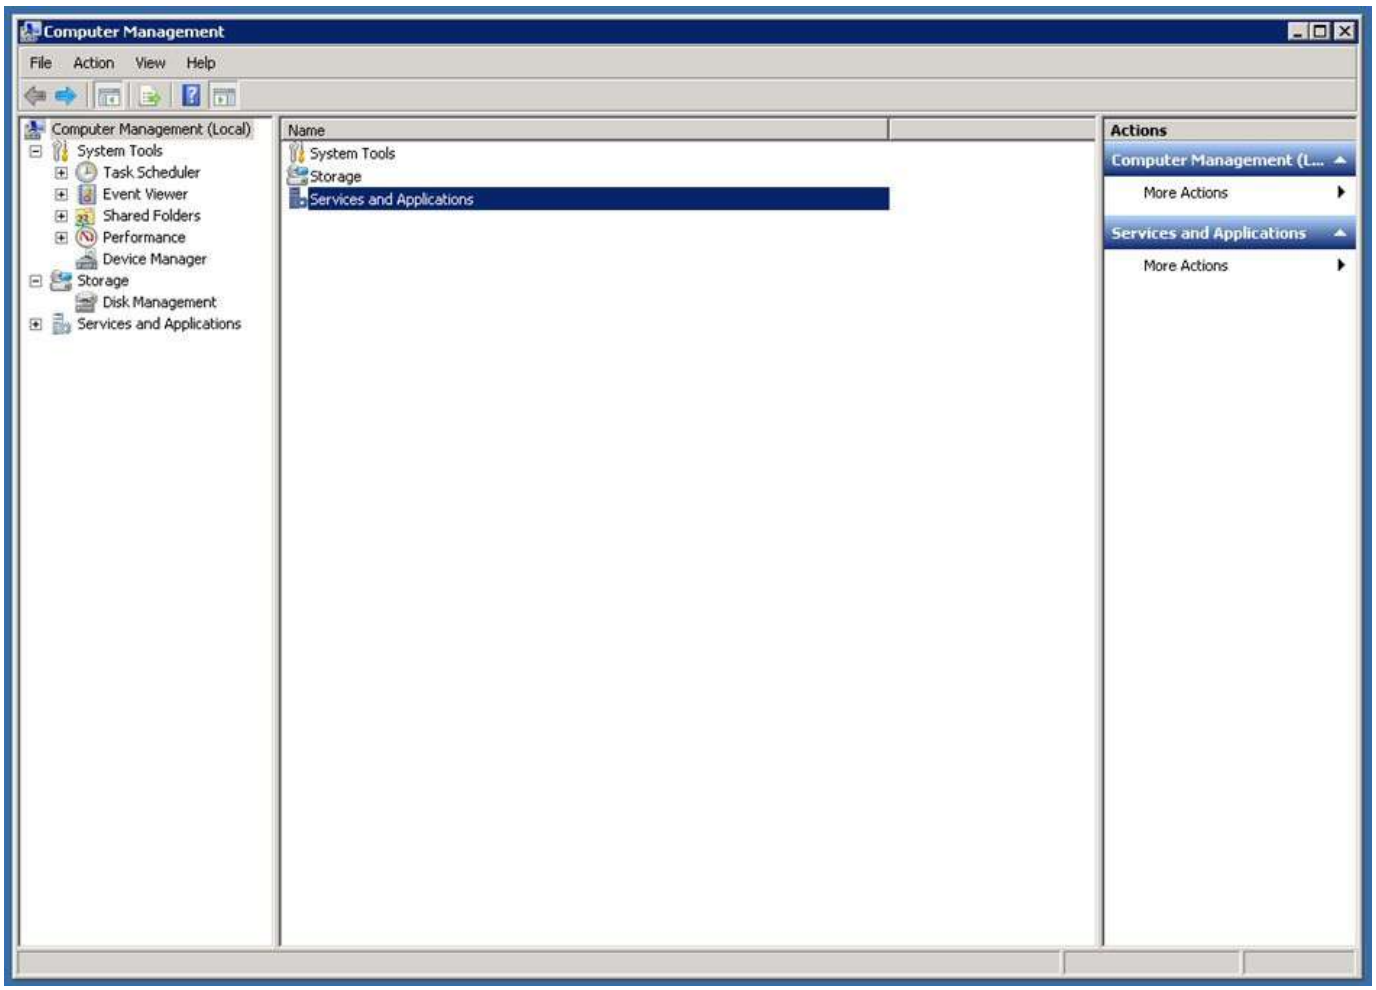
WMI CONTROL

You can set it up in Start -> Administrative Tools -> Computer Management.

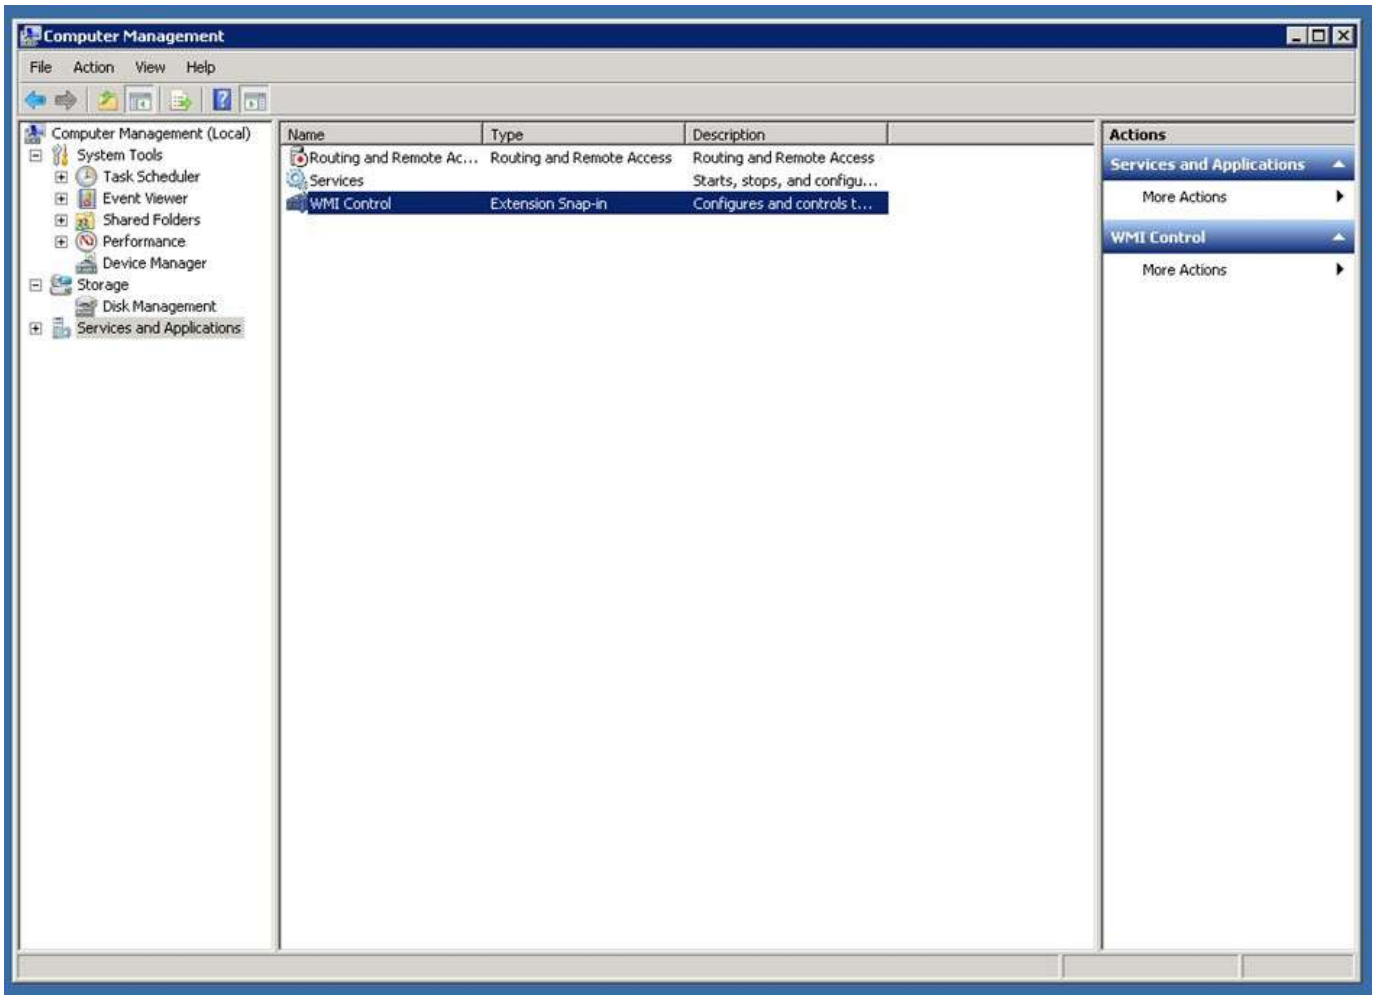
Copyright © 2024 Greycortex



You need to select Services and Remote Access -> WMI Control.

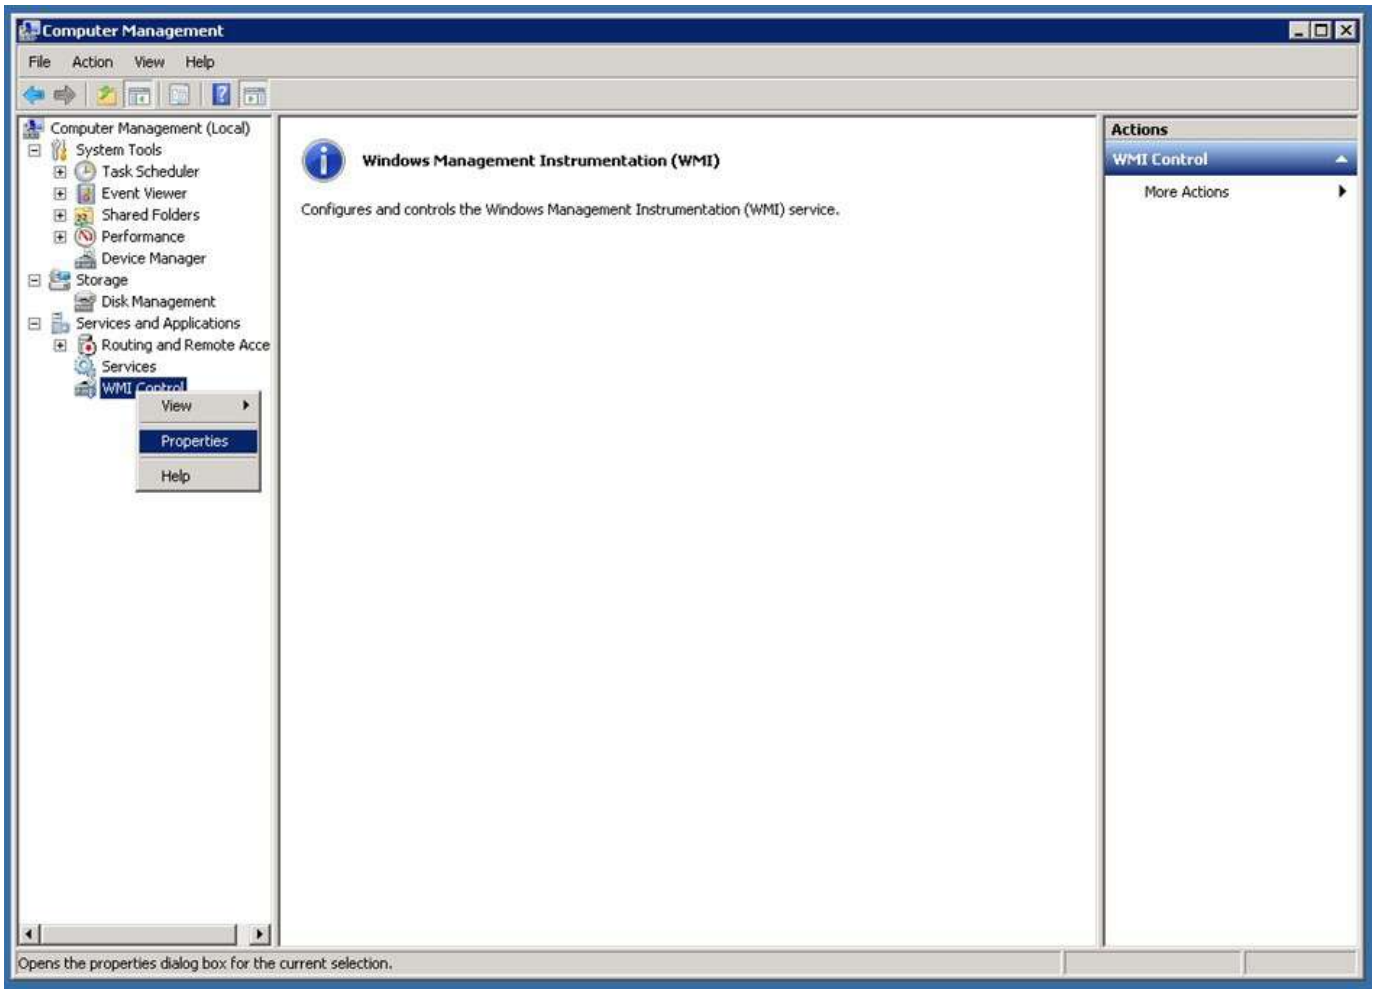


Copyright © 2024 Greycortex



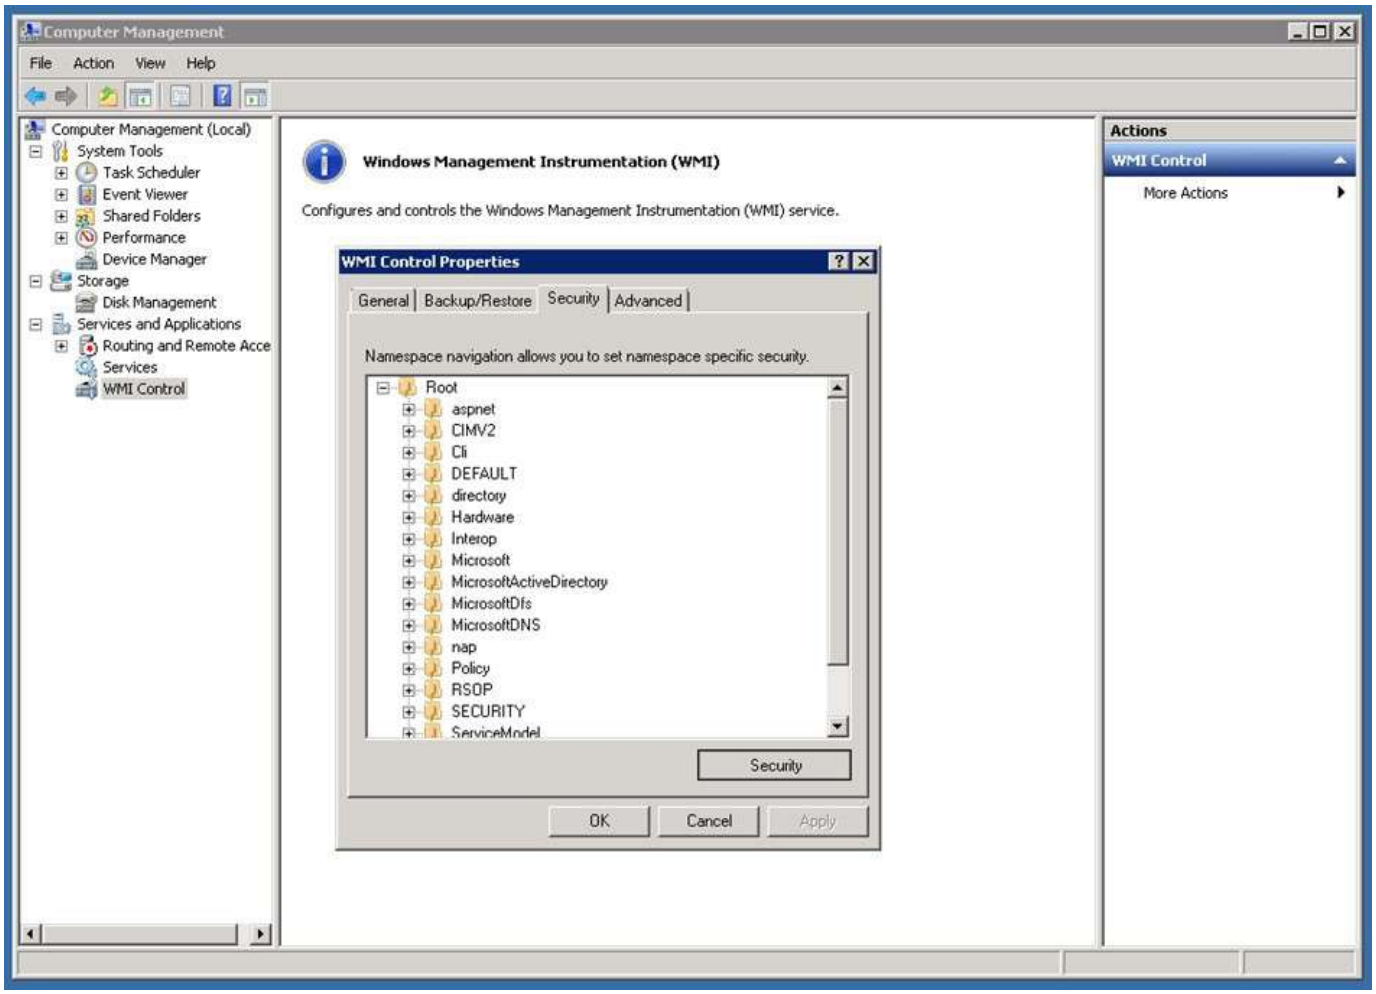
In the menu, you need to right-click on “WMI control” and choose properties.

Copyright © 2024 Greycortex



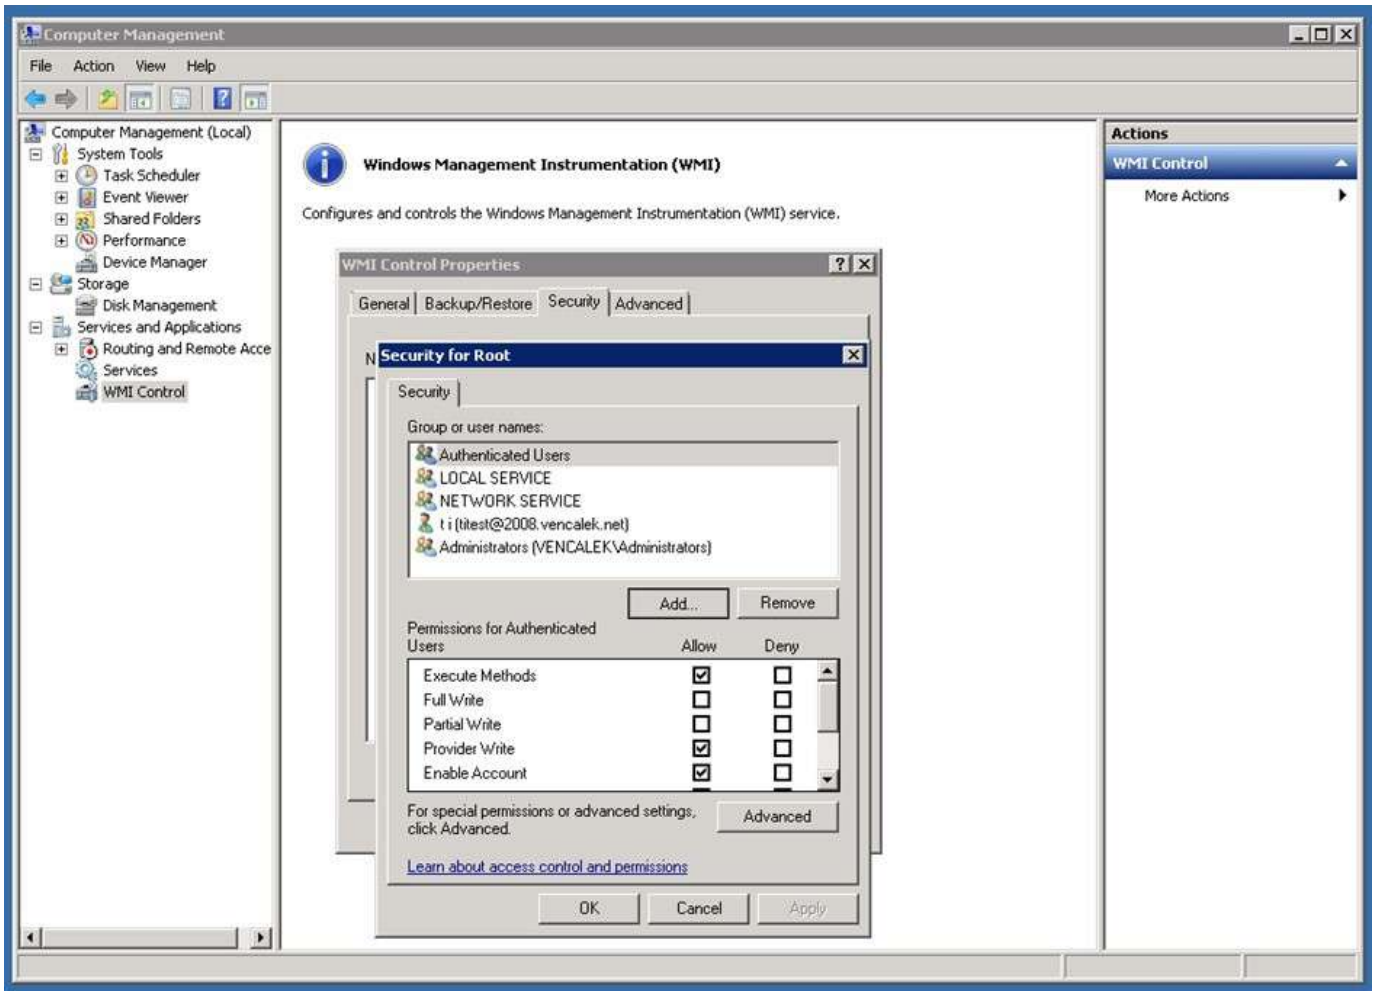
Move to the Security tab, select Root, and click on the button “Security.”

Copyright © 2024 Greycortex



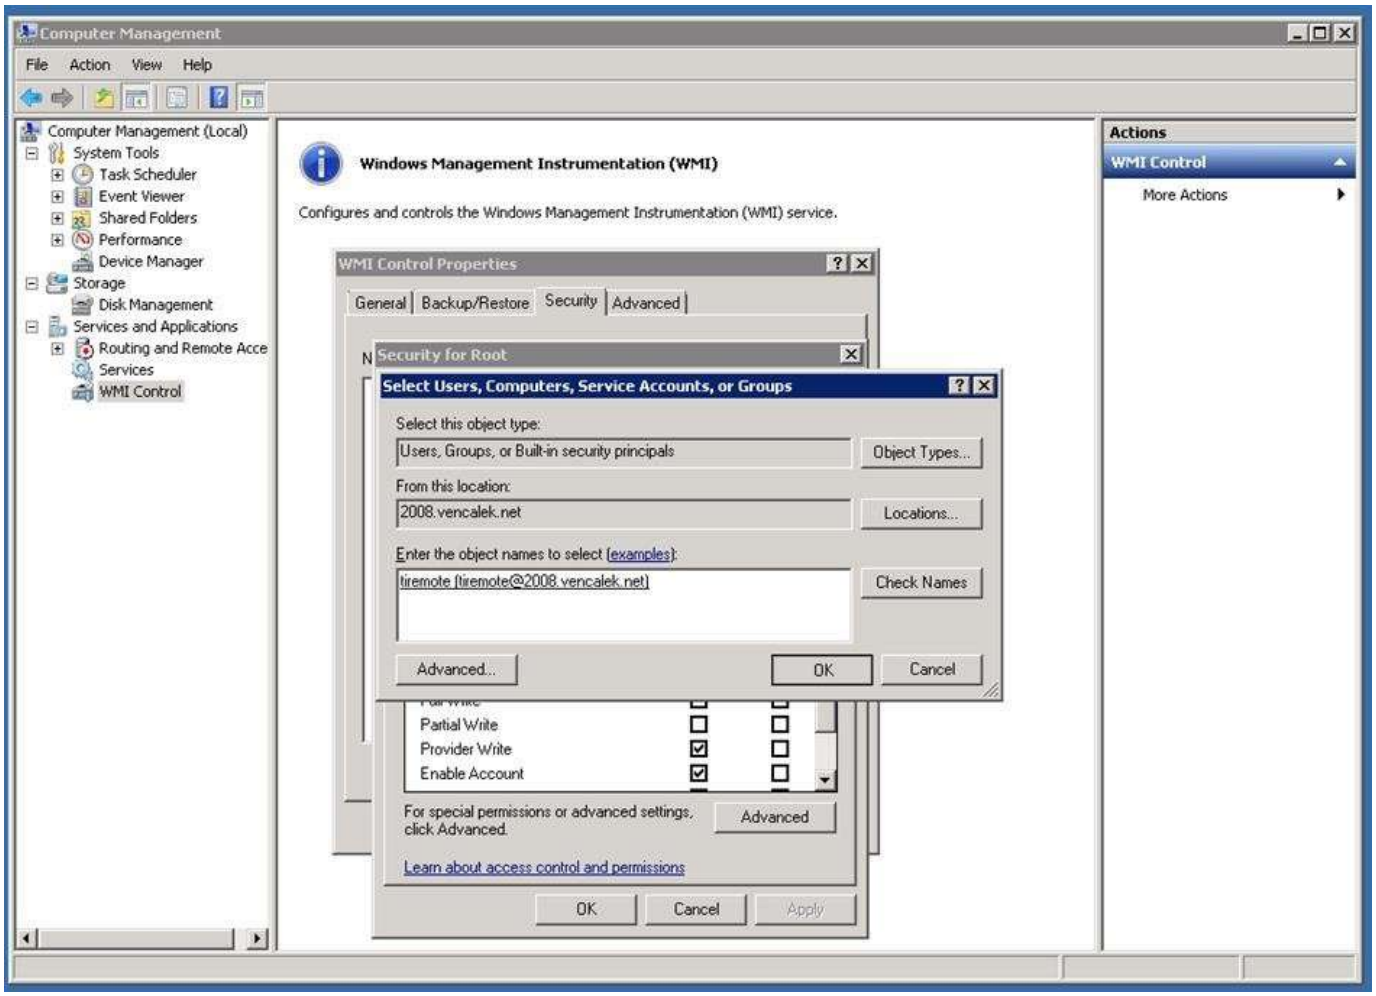
Select your new user record and add it using the “Add” button.

Copyright © 2024 Greycortex



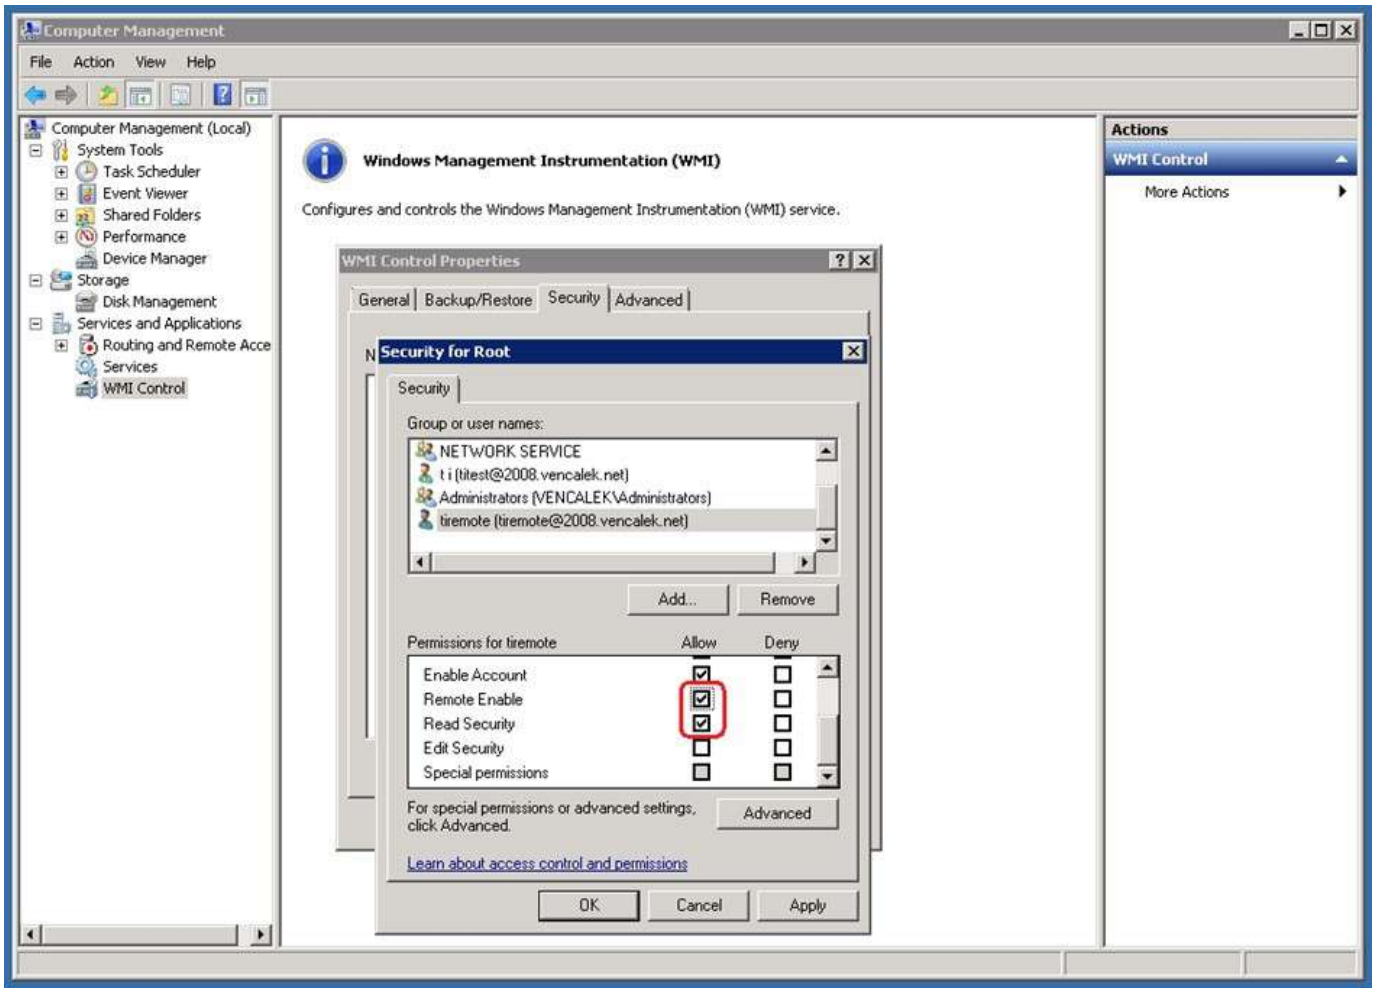
To add a name from already created user it’s advised to use the “Check names” function.

Copyright © 2024 Greycortex



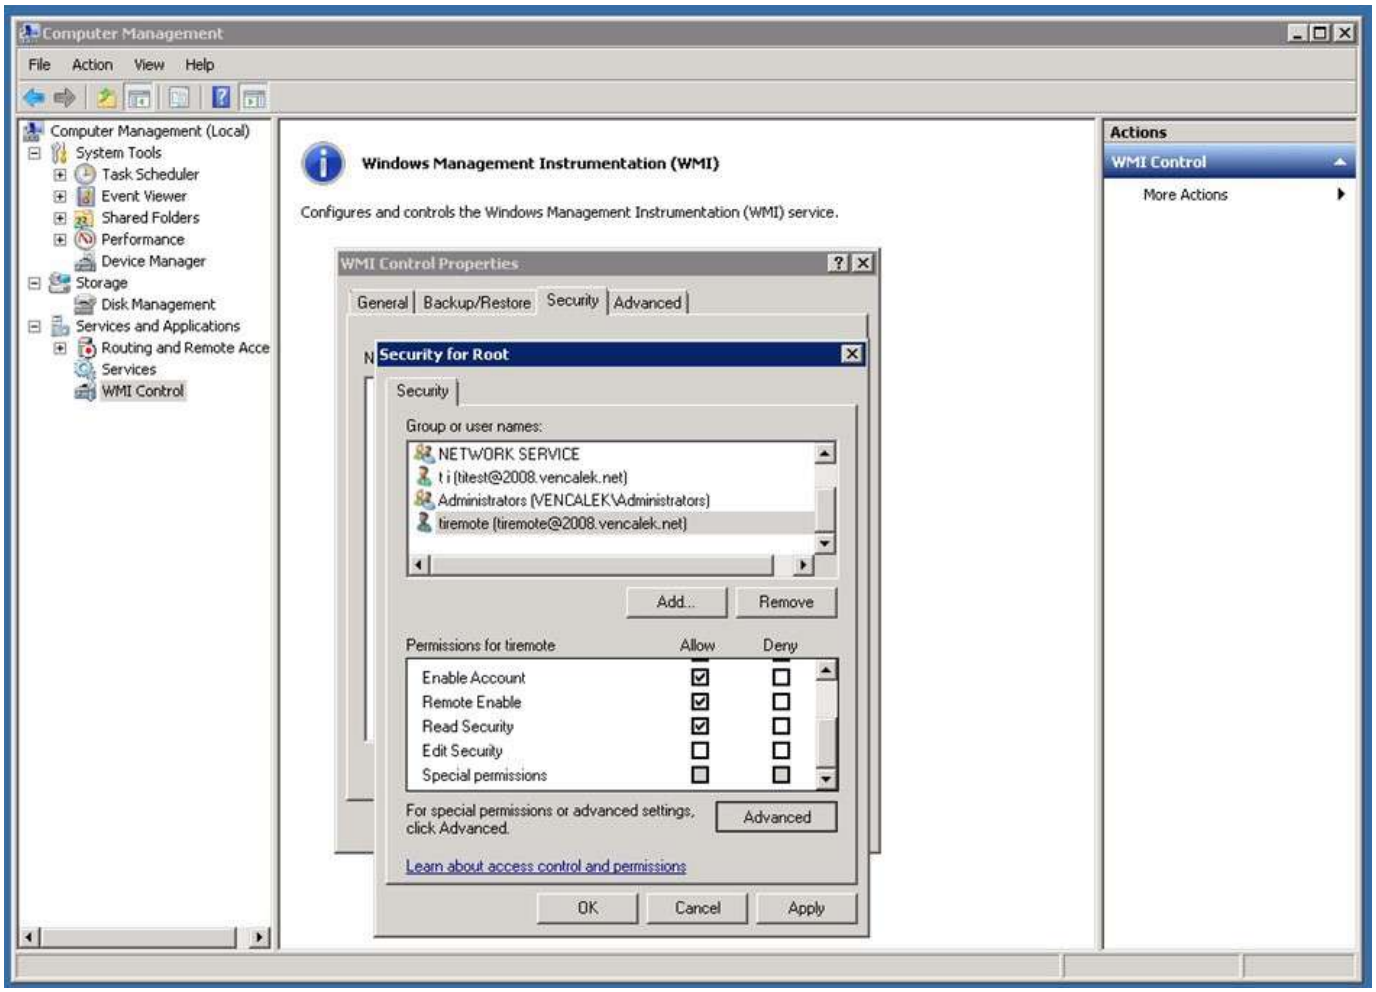
You also need to check “Remote Enable” and “Read Security.”

Copyright © 2024 Greycortex

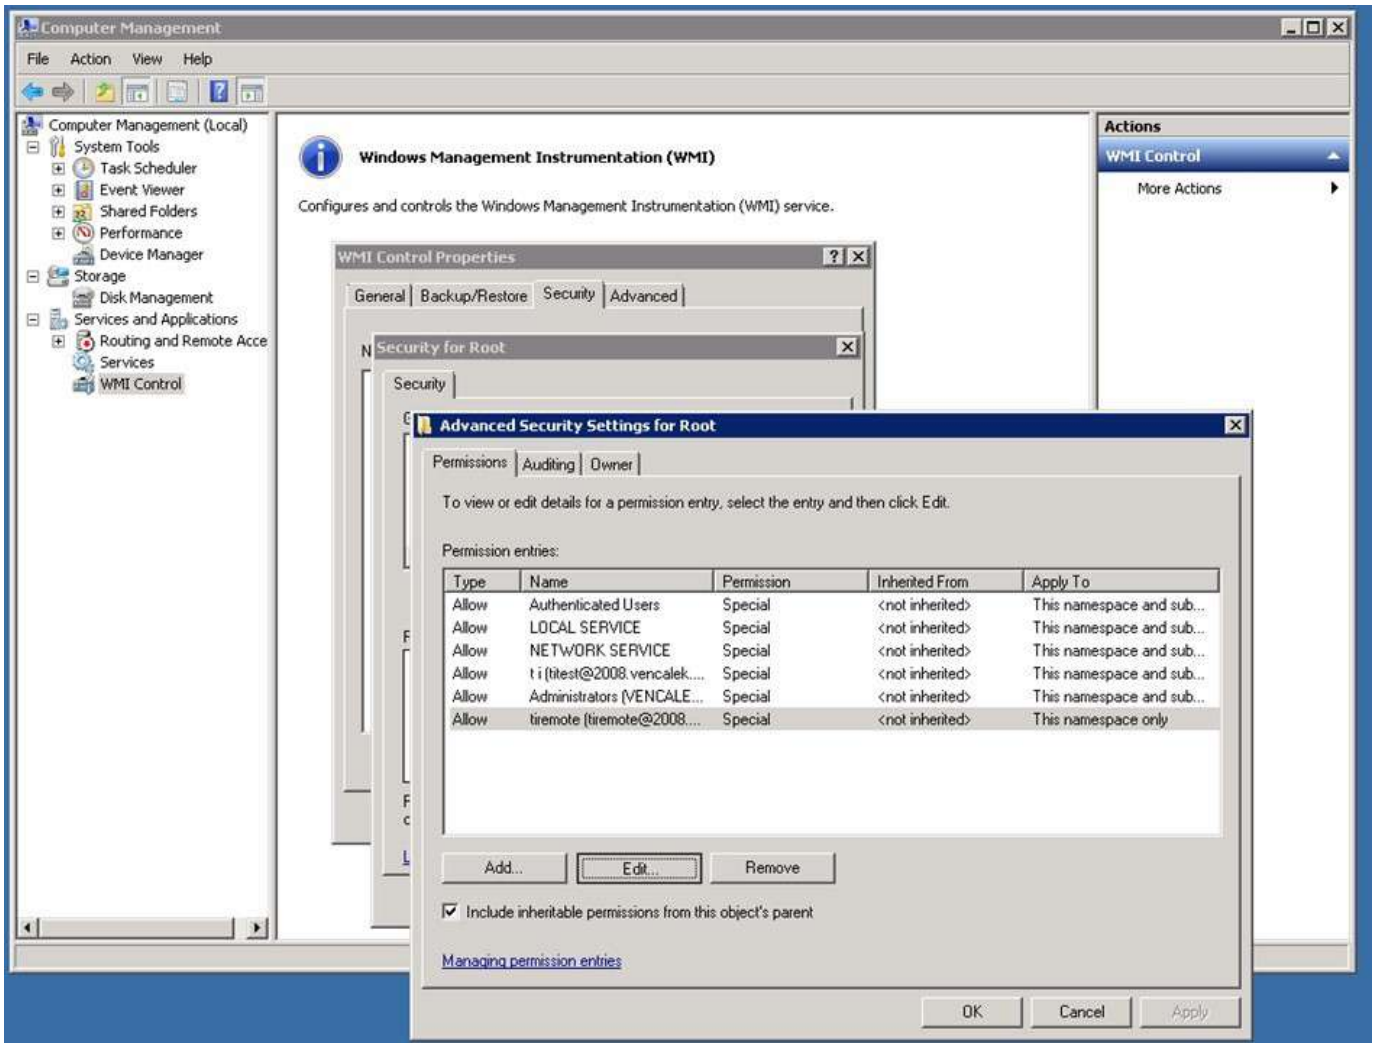


The final step is to select your user and to select “Edit” in the “Advanced” tab.

Copyright © 2024 Greycortex

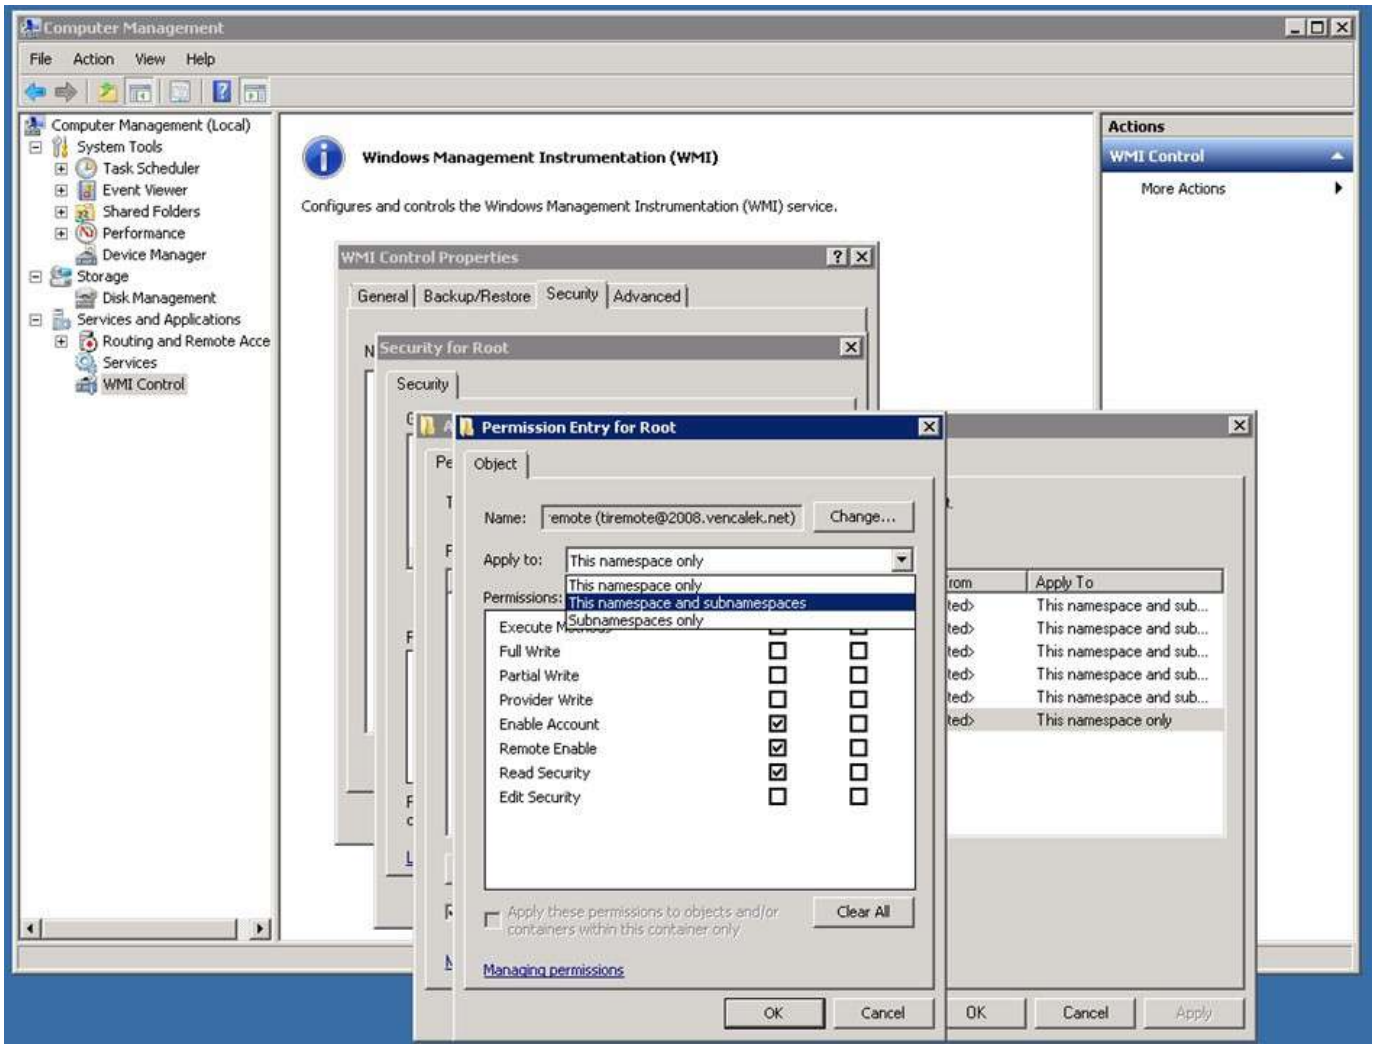


Copyright © 2024 Greycortex



Select “This namespace and subnamespaces” from the dropdown next to “Apply to:”

Copyright © 2024 Greycortex



Copyright © 2024 Greycortex

4.5.8 CISCO ISE Configuration

ISE Configuration Guide

In order to configure Identity Services Engine (ISE) integration in Mendel, the first step is to obtain necessary certificates from the Cisco ISE GUI and put them into Mendel's internal trust store. This can be done with a help of our guide ISE setup guide. You can find this document on our partners section.

Integration set up

In order to set up the integration, go to:

1. Click on the **Add** button to enter the wizard 2. Fill in all the necessary/optional information

Attribute	Description
	Custom name for your ISE integration profile.
	One or more IP address(es) or hostname(s), you have to confirm each item with enter key. Server priority is defined by order of entered servers from left to right.
	Username used for the integration. You can either pick already approved user, or select random one and approve it later in the Cisco ISE GUI.
Advanced	*Optional*
	Users will be assigned to the hosts from these sensor groups.
	How often will be the server used for getting events (in minutes).

1. If the entered server is reachable and the provided certificates are valid, you will be presented with a confirmation of validity of the provided information and with the option to save the configuration. If there is a problem with the authentication, you need to approve the newly created user in your Cisco ISE GUI at
2. (Optional) If you want to configure the LDAP integration (detailed information about detected users) toggle the 'Configure LDAP' slider. - This can be also done after the initial wizard.
3. (Optional) Fill in all the necessary/optional information.

Attribute	Description
	Either LDAP (Toggled off) or LDAPS (Toggled on) will be used.
	Base distinguished name.
	Search filter for users.
Advanced	*Optional*
	Valid domain username.
	Valid password for the domain username.
	If port is undefined, default LDAP or LDAPS port is used.

1. Click Save.
2. You can verify proper functionality of this feature by clicking on the hamburger menu button and selecting 'LDAP search'. If you try to search for a known user in the domain, information about this user should be returned to you.
3. You can customize the additional provided information about users in the 'LDAP Fields mapping' option.
4. You can enable/disable the querying of your ISE by toggling the slider right next to the hamburger menu button of the desired ISE profile.

Copyright © 2024 Greycortex

If your settings for both Mendel and the ISE are correct, you will be able to see recently connected user accounts in , and some additional host information in corresponding hosts and events.

Copyright © 2024 Greycortex

4.5.9 CISCO ISE Setup

Identity Services Engine Setup Guide

This guide will help you to generate and upload all the necessary certificates and a key from the Cisco Identity Services Engine (ISE) GUI to the Mendel's GUI.

GENERATING THE NECESSARY FILES

1. Log in to your Cisco ISE graphical user interface (GUI).
2. Go to .
3. Select from the drop-down window.
4. Fill in all the necessary information:

Attribute	Description
	Hostname of your Cisco ISE.
	Description for he new certificate.
	Insert the IP address of your Cisco ISE.
	Select the PEM format.
	Choose a passphrase for the certificate (it will be removed later).
	Confirm the password above.

5. Click on the button.
6. Download the cert package (zip).

PREPARING AND UPLOADING THE FILES

1. Extract all files from the cert .zip package.
2. Remove the passphrase from the key e.g. using linux command line:
3. Upload the Certificate+Key without a passphrase in the Mendel's internal certificate trust store select Type: and choose and upload both files.
4. Upload the root certificate chain (RootCA, NodeCA, and EndpointSubCA) selecting the Type: and repeat the process for each individual file.
5. You can proceed with the Identity services engine integration.

Copyright © 2024 Greycortex

4.5.10 CISCO ACI

CISCO ACI INTEGRATION

Generic solution to inherit important information about objects in the software defined network from SDN management layer to acquire extended visibility with applied SDN definitions (politics). Current approach is based on universal data exchange through APIs and marking the network objects with respective information = tags. This solution is capable to be used for any SDN vendor, pilot project is demonstrated on integration with CISCO ACI platform.

CISCO ACI INTEGRATION

- automatically created tag tree structure based on EPG/ESG definitions (Endpoint Groups / Endpoint Security Groups)
- automatically labeled objects (host or subnets)

On the basis of this extended information about hosts and subnets, it is also possible to perform advanced filtering and create detection and response rules.

Please contact us for more information.

Copyright © 2024 Greycortex

4.5.11 Identity logs

Identity Logs Configuration Guide

In case you want to enhanced the Visibility by getting user identity from your current tool´s logs you can use this feature.

INTEGRATION SET UP

In order to set up the integration, go to:

1. Click on the **Add** button to enter the wizard 2. Fill in all the necessary/optional information

Attribute	Description
	Custom name for your Identity logs integration profile.
Advanced	*Optional*
	Users will be assigned to the hosts from these sensor groups.
	How often will be the server used for getting events (in minutes).
	Processing of logout events can increase accuracy of assigning users to hosts. Recommended option: disabled. For more details, please see the in-depth guide.

3. If the entered server is reachable and the provided credentials valid, you will be presented with a confirmation of validity of the provided information and with the option to save the configuration. 4. If you want to configure the LDAP integration (detailed information about detected users) toggle the 'Configure LDAP' slider. - This can be also done after the initial wizard. *(Optional)* 5. Fill in all the necessary/optional information. *(Optional)*

, , , , , , , , , ,

6. Click .

Copyright © 2024 Greycortex

4.5.12 EDR Integration

EDR INTEGRATION

Generic (vendor plugin based) solution to cooperate with different brands of endpoint clients (via APIs). Current coverage within the first stage is focused on generic data acquisition from the endpoints to enrich data in Mendel.

Tip

The sample integration is prepared as a Proof of Concept based on the example of connection with the SentinelOne© platform. To find out more see SentinelOne Integration Guide.

Copyright © 2024 Greycortex

4.5.13 SentinelOne Integration

The integration covers the scenario for the transfer of security data from the SentinelOne platform and its further use in the GREYCORTEX Mendel product.

Main abilities

- obtaining detailed information about the end client (host)
- obtaining Incidents generated by the SentinelOne platform
- execution of an action based on the fulfillment of the conditions for the response (for example: blocking/quarantine of the host appliance performed by the EDR client)

HOST DETAILED INFORMATION

Granular information about specific endpoint client (host) acquired from SentinelOne platform.

Example of retrieved details:

, , , , , , , , , , , , .

INCIDENTS

By using the native processing of logs in the product GREYCORTEX Mendel, the extracted information is processed and the subsequent creation of events reporting security Incidents that were reported by the SentinelOne platform.

RESPONSE

The action is based on our Firewall pluging response system.

We have created a script that can be imported into the GREYCORTEX Mendel product as a plugin and by subsequently creating rules and defining the conditions for its execution, you will achieve an automated response.

How to use plugin response system, see Automated Threat Response Guide.

> **Warning**
>
> To use this integration, cooperation of a partner or manufacturer (GREYCORTEX) is necessary. Feel free to contact your partner or directly support@greycortex.com.

Copyright © 2024 Greycortex

4.5.14 XSOAR Integration

Copyright © 2024 Greycortex

4.6 Ipfix Configuration

IPFIX (IP Flow Information Export) is a protocol intended for data collection from routers or probes, which observe communications in the network. MENDEL has IPFIX export implementation built in. This means you can send all flows to an IPFIX collector in the IPFIX format.

4.6.1 How to set up MENDEL to start exporting IPFIX:

There are three approaches of IPFIX setup in MENDEL

1. **Collector** - One Collector captures network traffic data. From this data, the IPFIX exporter generates exports.
2. **Central Event Manager** - The instance of MENDEL acting as the Central Collector (Central Event Manager) receives only events from other, connected collectors. Because of this, IPFIX exports must be configured on every collector separately. This is because flow data stays only in the collector instance that captures it, or which receives it from a sensor.
3. **Sensor-Collector** - In this type of deployment, the collector gets flows from all sensors in the same architecture. This means that a collector is able to export IPFIX from all sensors. If you want a separated export for every sensor, you must to configure the filter for that output.

4.6.2 Setup an IPFIX Export

1. Go to:

2. Press the button
3. Complete the following fields in the dialog:

Attribute	Description
	Name of the export
	Specification of data being exported (For IPFIX exports, select “Flows.”)
	Format of export (For IPFIX exports, select “IPFIX.”)
	How often data will be send (In IPFIX case only option is 1min)

4. Click the button.

Copyright © 2024 Greycortex

4.6.3 IPFIX configuration

1. Press the button in the export row which you want to edit.

Important options	Description
	Port based on the IPFIX collector
	IP address of the IPFIX collector
	Transport protocol is based on IPFIX collector configuration

Advanced options	Description
	An identifier of the Observation Point (Probe) that is locally unique to an Exporting Process.
	Exports IPFIX of specific traffic based on the selected filter.
	In future you will be able to edit which data will be exported.
	L7 data will be added to the export as a json file in proprietary fields 902 and 903.

Copyright © 2024 Greycortex

2. After configuration is done, you can click the checkbox. This will start the configured IPFIX exporter. Then you should see exports in your IPFIX collector.

These columns are currently implemented in MENDEL's IPFIX exporter	Description
	IPv4 source. address
	IPv4 destination. address
	IPv6 source. address
	IPv6 destination. address
	Source MAC. address
	Destination MAC. address
	Source. port
	Destination. port
	Incoming counter with length N x 8 bits for number of bytes associated with an IP Flow.
	Incoming counter with length N x 8 bits for the number of packets associated with an IP Flow.
	Outgoing counter with length N x 8 bits for number of bytes associated with an IP Flow.
	Outgoing counter with length N x 8 bits for the number of packets associated with an IP Flow.
	These are control bits that indicate different connection states or information about how a packet should be handled.
	The IP version field in the IP packet header.
	System uptime at which the first packet of this flow was switched.
	System uptime at which the last packet of this flow was switched.
	The relative timestamp of the first packet of this Flow. Counted from 1.1.1970.
	The relative timestamp of the last packet of this Flow. Counted from 1.1.1970.

4 different IPFIX templates are used:	Templates contain these proprietary fields:
256 - ipv4 request	900 - sensor
257 - ipv4 response	901 - service
258 - ipv6 request	902 - json
259 - ipv6 response	903 - json

Copyright © 2024 Greycortex

4.7 Multitenant Configuration (SOC)

MENDEL can be installed in several deployment architectures. One of them is SOC (Security Operation Center) architecture. In this architecture, every customer has their own Sensor, and all of them are processed in one Collector that is managed by the Managed Security Service Provider (MSSP) partner.

4.7.1 Types of accounts

- **Super Support** – The support account for GREYCORTEX to be able to help the partner or the customer.
- **Super Administrator** – Account dedicated for MSSP partner, so he can see and edit everything (Collector, Sensors, etc.).
- **Analyst** – Without assigned group, analyst can see everything depending on their membership in a system group.
- **Common User** – Account dedicated for the customer, must be restricted by assigned Sensor.

4.7.2 Permissions

There are five System groups – Administrators, Default, Report manager, System manager, and User manager. With the combination of user groups and system groups, you can manage what each user should or shouldn't be able to see and do.

Restrictions are used to limit the access to information that are visible on Dashboard Tab, Network Tab and Events Tab. Use only Read permissions when configuring the restrictions. When you configure the restrictions, only data from the selected source will be available to the user belonging to this group. If you want to deny the source, you must configure all the others.

4.7.3 Creating the Super Administrator and Super Support

Super Administrator and Super Support users are created during the installation process.

Setting up the customer

For each customer, you need to create two new user groups, which will belong to the customer:

***Company Name* Users** – This group will be used for the customer's administrators and users. Users will be able to see other members and able to see Incidents, Reports, Filters, etc. created by other members within this group. Analyst user will not see it. Add restrictions for this group:

- Customer's Sensor with Read permissions
- Components: NBA, IDS and EC with Read permissions -> The user will not be able to see the Status Monitor

***Company Name* Shared** – In this group there will be customer's users and Analyst user, so that Analyst user is not restricted by the Sensor and users can assign Incidents to the Analyst users. Do not apply any restrictions for this group.

Create a new user (this will be the Common User) with the name of the customer. Assign the user to the groups:

- ***Company Name* Users** with Read permissions
- ***Company Name* Incidents** with Read permissions
- **Report manager** - When the user is also in the group restricted to Sensor, only new types of Reports and Companies will be visible in the Reports section of settings. When not restricted, Data exports and old type of Reports are also visible. Assign the user into this group with Read permissions if you want him to be able to download existing reports and request the new generation of the already created Report and see existing companies. Use Write permissions if you want him to fully manage these sections.
- **Administrators** with Write permissions if you want him to be able to edit Network and Detection settings or Read permissions if you want him to only see them.

  (Administrators mean network administrators; appliance administrators are in the System manager group).

This user is now able to see the data and create incidents only on customer's Sensor, can see only users in the Company group, can see only the Network and Detection settings and cannot see users from other groups or their data.

Copyright © 2024 Greycortex

Creating the Analyst user

Create a new user and use the name “Analyst”.

Add the user to the groups:

- **Administrators** with Write permissions if you want him to be able to edit settings or Read permissions if you want him to only see them.
- All ***Company Name*** **Shared** groups for which you want the Analyst to see the incidents, with Read permissions.

4.7.4 Management

Customer will be able to:

- See data on Dashboards, Network Tab and Events Tab restricted by Sensor.
- Create incidents and assign them to other members of the same group, including Analyst
- See or manage Network and Detection settings (visibility restricted by the Sensor).
- See or manage Exports section – New type Reports and Companies
- See or manage False Positives
- See Network policies
- See and use Filters saved in Administrators group
- Take actions based on all already created Plugin profiles
- See or manage (according to permissions) objects created by other members of the same *Company Name* Users group:
- Data Captures
- New type Reports
- Companies
- Incidents

MSSP partner who owns the Collector will be able to:

- Manage the whole appliance
- Manage users for the customer
- Manage authentication policy for the customer
- Manage Network policies for the customers (Do not use the name of the customer as the name of the policy)
- Create Reports and configure them to be sent regularly to the customer via email
- Create and configure Data Exports (Events or Flows) to the external system on customer’s request – SIEM alerts, IDEA alerts or alerts via email
- Manage Plugins and Plugin profiles

Copyright © 2024 Greycortex

4.8 Network Analysis

Stay tunned - will be added soon.



You can currently be inspired by predefined views to define your own queries.

Back to Network Analysis User Interface.

Copyright © 2024 Greycortex

4.9 Remote Shell

4.9.1 Remote Shell

MENDEL can also be managed remotely through an SSH connection through an administrator account. To connect, use the administrator's username and password, and the IP address of the MENDEL device.

Notice: SSH connection must be enabled on TCP port 22 in MENDEL's firewall (see firewall settings in User Manual).

Remote Shell (MShell) is divided into four sections:

System information:

- **check**– hardware requirements and settings
- **drives**– show system hard-drives
- **status**– show system service status
- **info**– show system & service information
- **version-check**– show system & packages version
- **uuid**– shows system's ID

System management tasks:

- **reboot**– reboot this system
- **shutdown**– shutdown this system
- **hostname**– show system hostname
- **passwd**– change user password
- **gencert**– generate a new self-signed certificate
- **restart**– restart system services
- **services**– run service manager
- **support-me**– generates "enable support account" request
- **license**– change mendel license

Support tasks:

- **help**– screen with list of commands
- **lookup**– lookup for domain/server using DNS
- **fw-man**– change firewall rules for web interface access
- **check-db**– check whether new flow and event data are being added
- **failures**– show GUI login failures and optionally reset counter
- **ping [-v] \<ip>**– ping \<ip>, optionally with verbose option (-v)
- **ifstats [iface]**– show network statistics
- **rescue**– create a rescue request file

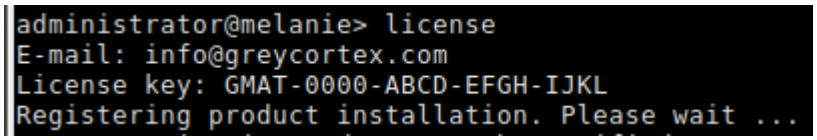
Copyright © 2024 Greycortex

Miscellaneous tasks:

- **join**– join a collector and act as a sensor
- **leave**– change role from sensor to collector
- **network**– show or change network interface settings
- **update**– run system update directly from shell
- **pkg-update**– run package update directly form shell
- **exit**– logout from shell
- **quit**– logout from shell

MENDEL license change

Step 1. - use "license" command
Step 2. - enter email and license key



Copyright © 2024 Greycortex

4.10 Samba Backup

4.10.1 Samba Backup Guide

GREYCORTEX MENDEL offers mechanism for continual backup of data via SMB protocol. Those data can be later restored and observed for additional forensic analysis. Customer can store historical data from multiple months or years depending on volume of hardware provided for SMB share.

Samba Data Backup contains:

- System Settings
- Data (Flows)
- Events

Configuration in Settings

Set up of backup mechanism can be configured in the

- Click button.
- . Fill the path to the shared folder on the external server.
- . Fill the name of the folder that will be used for backup from MENDEL. This folder will be created automatically in the shared folder that you have set in the previous step, please do not create it manually.
- is optional parameter. It is the domain that will be used for authentication.
- . Choose the time when the backup should run. Please check when is your Maintenance time in the and set the time in to Samba configuration at least two hours before or after this time, otherwise it is possible the backup will not start due to database upgrade in progress.
- Fill the and for the authentication.

Technical information

SMB backup folder holds 2 types of backup: and .

1. It’s necessary to maintain 4 newer backups by timestamp . The backup contains all important database tables in the system.
2. Backup consist of historical data. Deleting older data will result in loss of historical data by timestamp. The backup contains backups of all partitions that were finalized in the last 24 hours.
3. backup will be stored for last 7 days. If backup fails for some reason it will be backed up next day.
4. Hash file with same name as backup is present for integrity checking.

Maintenance

Periodically delete backups while maintaining 4 newest backups by timestamp.

backups should be deleted to free space on SMB share. Data should be deleted in chronological order from oldest to the newest. Deleted historical data with timestamp won’t be recoverable.

In short, to maintain your SMB share, delete all data that are older (by timestamp) than desired time period of your backup. Remember that backup will keep continuity of 7 days.

Copyright © 2024 Greycortex

Restoration process

- Restoration process is described in Samba Recovery Guide.

Copyright © 2024 Greycortex

4.11 Samba Recovery

4.11.1 Database recovery from automatic Mendel Samba backup

To create Mendel Samba backup follow Samba Backup Guide.

Preconditions
Database recovery can be used with backups of version 3.4.0 and higher. However, the recovery function is available for versions 3.4.1 and higher.
The data is restored from the newest backup to the oldest and automatically stops when the disk reaches 65% of its capacity. Any older data beyond this point will not be stored.
For recovery you will need an existing machine, which will be reinstalled or a new hardware with a new license.
Data recovery time depends on the amount of data to be recovered as well as the performance of the machine and the speed of the connection between Samba storage and Mendel. It may take up to couple of hours or even days for the recovery to finish.

Step by step Guide

1. SETUP SMB BACKUP ON RUNNING MENDEL.

Version should be the same as the one you are going to restore.

Expected result:

2. LOGIN AS ADMINISTRATOR VIA SSH.

Enter command:

Input "" and confirm with ENTER.

Info

Mshell will notice that your connection is made on SSH and it could drop connection and break process of restoration. Note that if you login via physical terminal, notice will not be there. Performing restore via physical terminal is considered to be safer.

3. MSHELL WILL ASK FOR UNC PATH TO THE BACKUP AS IS IT STORED ON SAMBA SHARE.

Enter correct path to your backups e.g.

- Mshell will ask for username that has access to your backups.
- Enter correct username.
- Mshell will ask for password for user that has access to your backups.
- Enter correct password.

Important

If path or username or password are wrong, mounting samba share will fail and process of restoration will end without any changes. Otherwise stopping of services will start.

Copyright © 2024 Greycortex

4. MSHELL WILL ASK:

- Pressing " " without ENTER will begin process of backup restoration.
- Pressing " " without ENTER will cancel process of backup restoration. Services and settings are restored to the original state.

Opportunity to change management interface will be still offered.

5. PROCESS OF BACKUP RESTORATION BEGINS.

- No errors connected to restoration should appear.
- There may be record like: . This may happen and it is harmless if suricata eventually (after restore is completed) starts and works as it should.

6. MSHELL WILL SHOW NOTICES REMINDING USER TO VERIFY LDAP, AD, EMAIL, ... SETTINGS AFTER RESTORE.

Info

Mshell will print system network interfaces and offer the opportunity to set management interface. This is only needed in situation when interfaces after restoration are not according to user's expectation. Changing management interface while on SSH may cause session drop resulting in premature restoration end.

Expected results:

- Restoration is finished.
- Services are in state as before restoration start.
- Data from backup are restored.

7. VERIFY THAT DATA (FLOWS) WERE RESTORED SUCCESSFULLY.

Explore events tab in UI and navigate to interval of restored data.

Expected result:

- Flows and events within interval are the same as on original sensor.

False-positives

If false positives were applied to the past, recovered data will not be affected. These false positives will have to be applied to the past again for them to take effect.

- Beware, that the false positives should be applied to the past one-by-one because this operation causes a heavy load on the database.

Copyright © 2024 Greycortex

4.12 Sensor Collector installation

4.12.1 Sensor-Collector Installation Guide

The Sensor-Collector approach allows users to visualize data from up to 30 sensors in one collector.

Warning: **MENDEL does not support the Sensor-Collector approach on instances with the same hostname!**

This guide describes how to connects two appliances through the Graphical User Interface (GUI). There is an option to connect it through mshell as well. All you need is log as administrator into mshell of the future sensor, write "join" as a command and continue by following the suggestions in mshell.

Important Settings

Before starting, please check if you have allowed these ports on your firewall as well as in Mendel:

- For connection Sensor-Collector allow **port 443**.
- Database is working on **port 5432**.
- Logs can be downloaded on **port 22**.

Our database uses OpenSSL to encrypt connection. It uses certificate and strong encryption with **key longer than 128 bits**. Detailed cyber suite is **ECDHE-RSA-AES256-GCM-SHA384**.

Interface capacity usage

Deployment of Sensors and Collectors to target infrastructure has an impact on interface capacity usage.

Connection from **Sensor to Collector**
Each sensor generate 1-2% load from current traffic on collector management interface.

Connection from **Collector to Central Event Manager** (CEM)
Each sensor on each collector gernerate 0.1% load from current traffic on CEM management interface.

Example: Deployment of 4x1000Gb sensor requires 40 - 80Mbit link capacity to collector and 4Mbit from collector to CEM.

4.12.2 Connect sensor to the collector via GUI

Data deletion

Changing the destination of the machine from All-in-One (or Collector) to Sensor has the effect of deleting all previously saved data, which was created by analyzing network traffic.

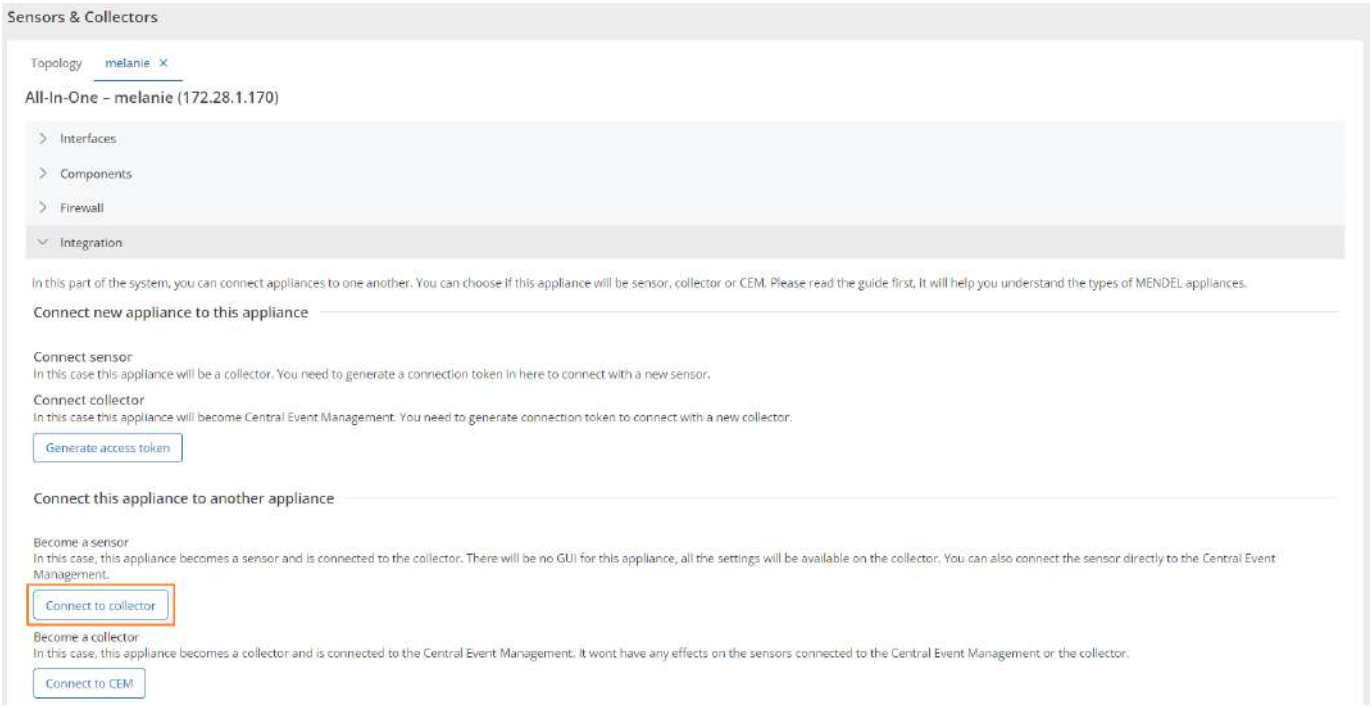
In next steps are **two all-in-one appliances** used to connect together, to get Sensor-Collector scenario.

For **sensor** is used Hostname: **mendel03** and IP address: 172.28.0.52
For **collector** is used Hostname: **mendel02** and IP address: 172.28.0.51

Step 1. Open settings on sensor (e.g. mendel03)

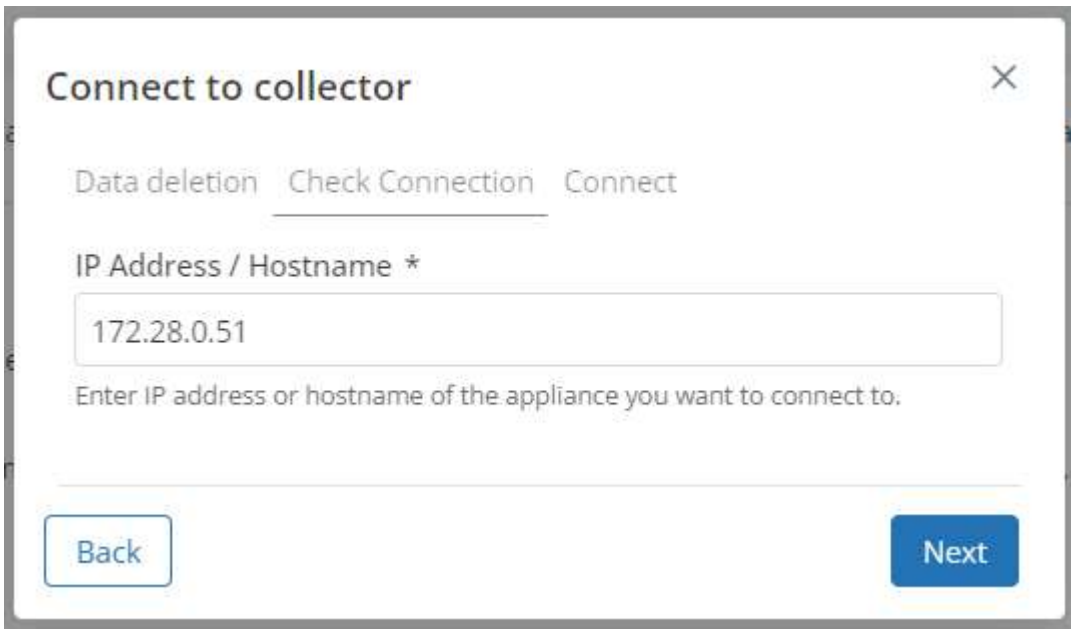
Go to and click on tab. There are few options for connecting more appliances together. You need that appliance be a sensor, that means it has to be connected to the collector. Click on button.

Copyright © 2024 Greycortex



Step 2. Add IP address of the collector

Insert the IP address of the future collector and click to the button . You will see input field for connection token, it must be generate from future collector appliance.

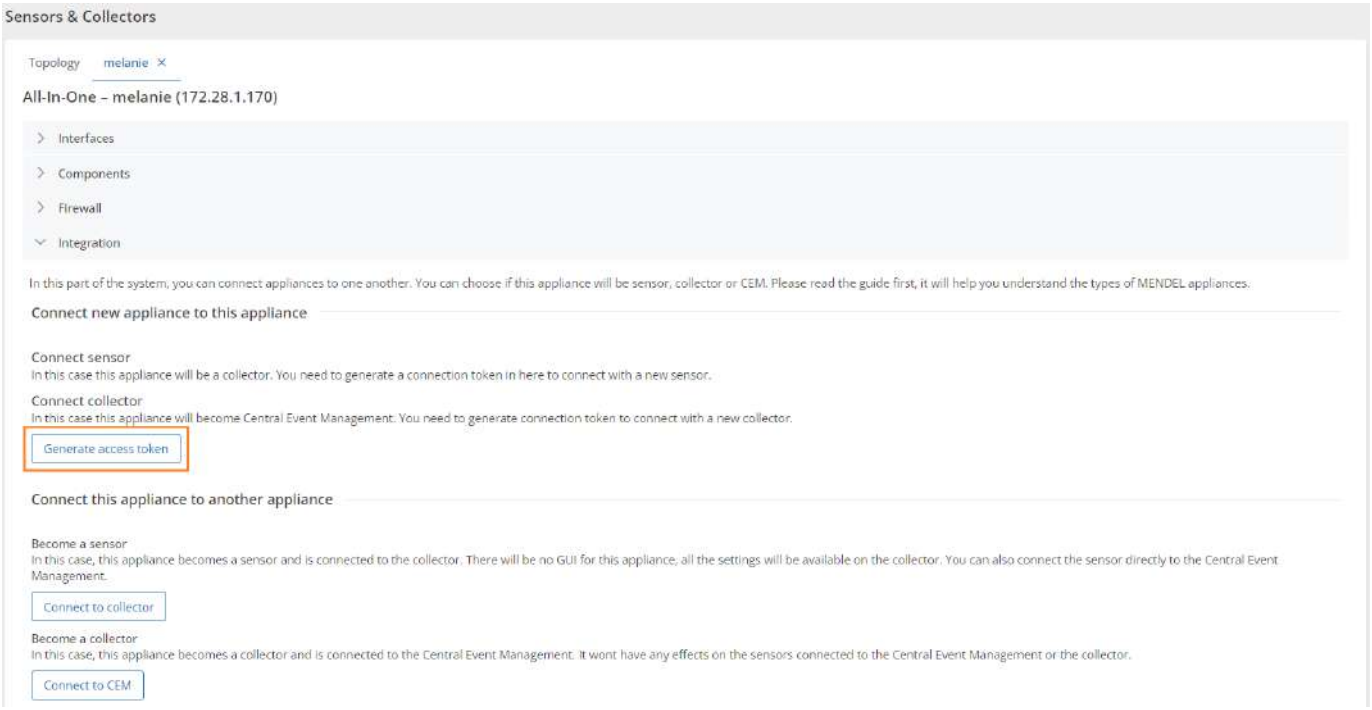


Please, do not close this window, you will need it after generating token.

Step 3. Open settings on collector (e.g. mendel02)

Go to the UI of the second appliance (in our case mendel02) and click to . Click on the button.

Copyright © 2024 Greycortex



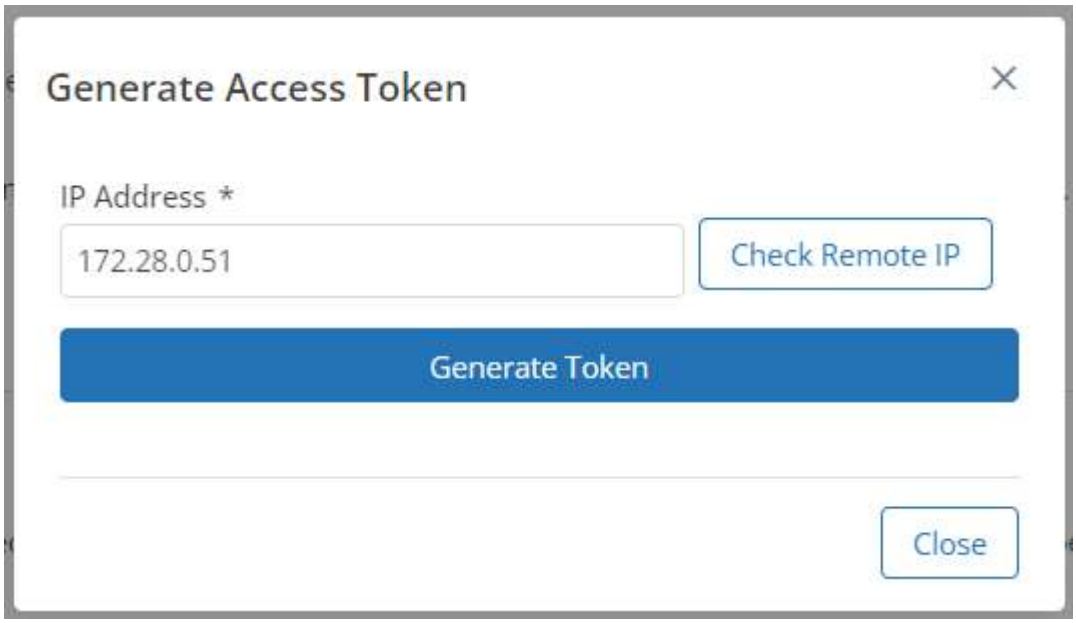
Step 4. Add IP address of the sensor

If the appliance knows the IP address of the future sensor, it will be automatically filled out. In case of empty box for IP address, you can click to the button.

This will show you what IP address is trying to connect. In case that this option is not working either (e.g. you have NAT in your network) you need to fill out the IP address manually.

Step 5. Generate the token on collector (e.g. mendel02)

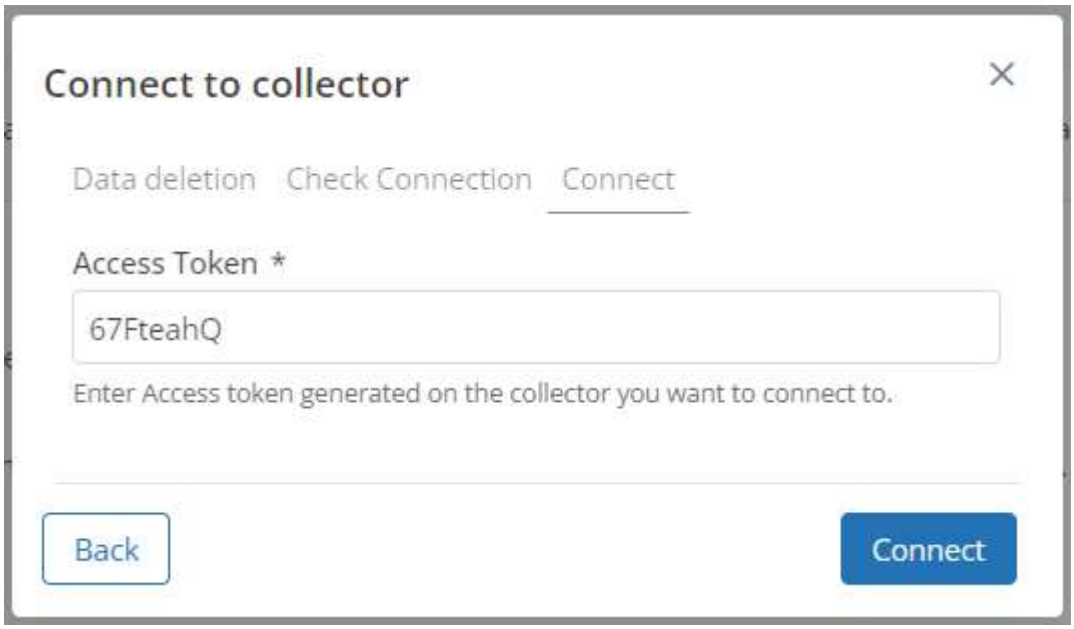
Token will be generated by clicking on the button. You can copy it by the link on the right side .



Copyright © 2024 Greycortex

Step 6. Insert the token to the sensor (e.g. mendel03)

Now when you have connection token, you need to go back to the UI of the other appliance (future sensor). And add the token and click on button.



After clicking on the button, there will loader on the whole UI. That means that two devices are now successfully connected. You are not able to go to the UI of the sensor (e.g. mendel03) all settings are now in the collector (e.g. mende02).

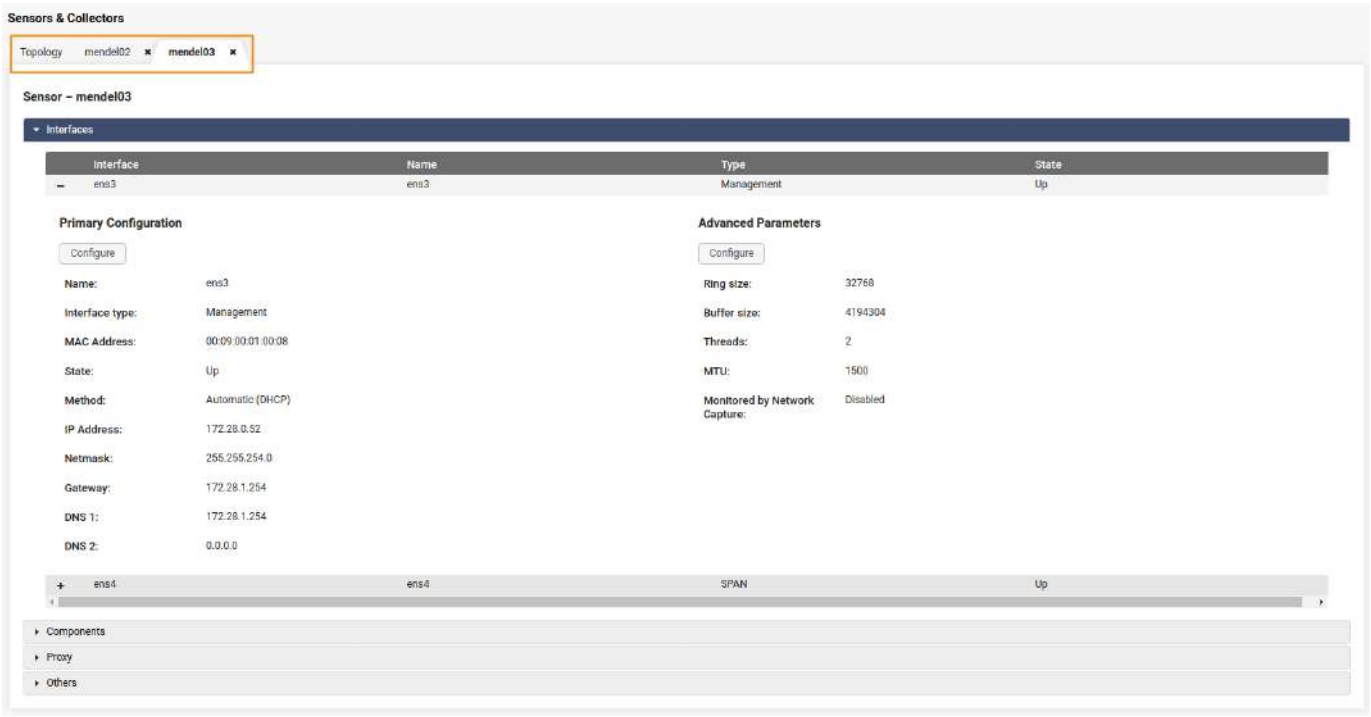
You can check the success of the installation on the collector in . It opens the topology tab.



4.12.3 Edit settings of the appliances

All settings are now available after clicking on the in the right corner of the appliance in the topology tab. It will open a new tab next to the Topology. So you now can control the sensor and the collector in one place.

Copyright © 2024 Greycortex

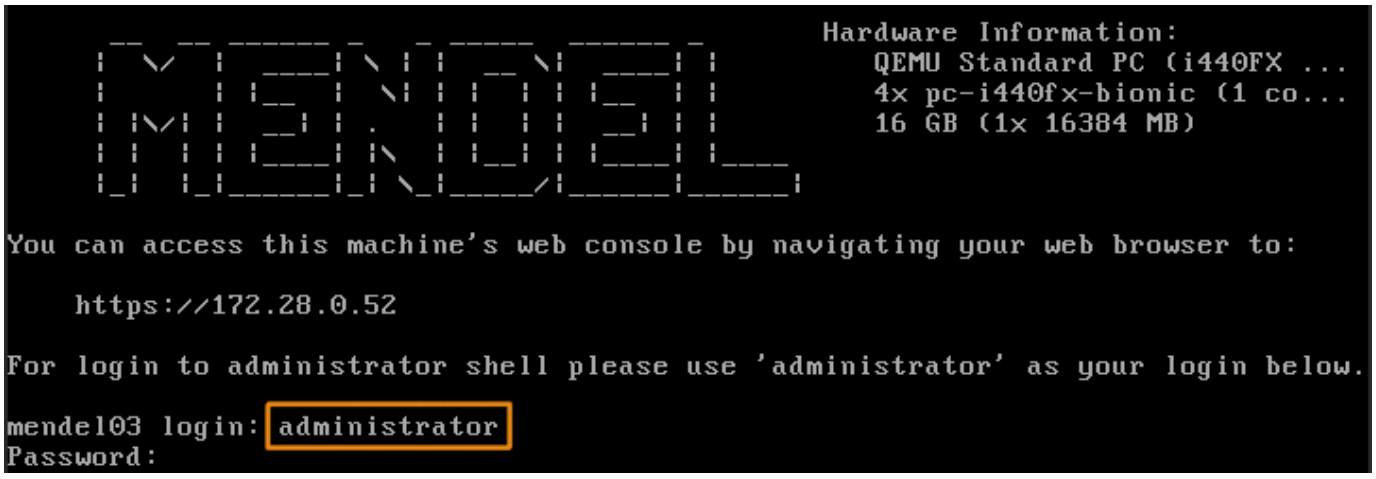


4.12.4 Disconnecting the sensor from the collector

Disconnecting is possible only through mshell.

Step 1. Log into mshell on the sensor (e.g. mendel03)

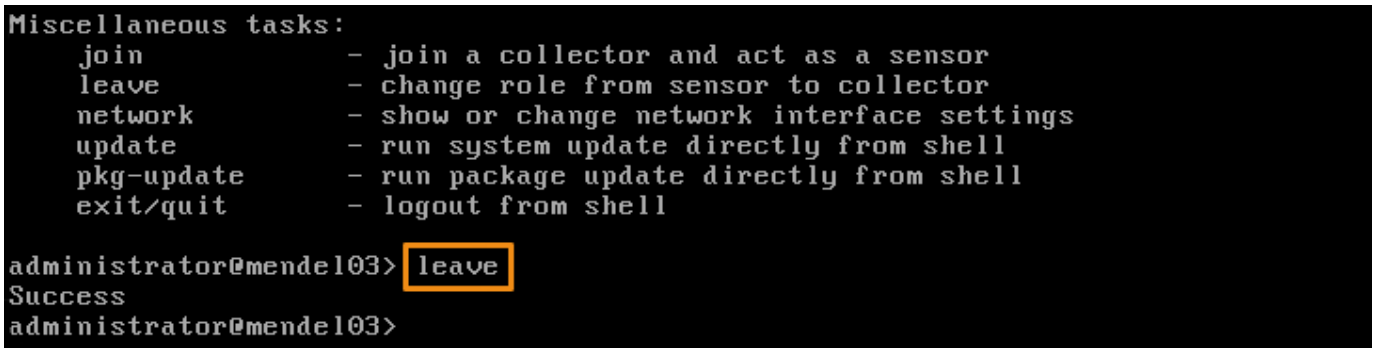
Using your administrator login, log into mshell on the machine you want to use as a sensor.



Step 2. Disconnect the collector (e.g. mendel02)

When you are logged in mshell, write command **leave** for disconnecting the sensor from the collector. Wait until there is success notice in the mshell.

Copyright © 2024 Greycortex



Step 3. Check the UI on former sensor (e.g. mendel03)

You can now check if you can see GUI of the appliance - that means disconnection was really successful. In you see just one appliance.

Copyright © 2024 Greycortex

4.13 Sensor Collector limitations

4.13.1 Limitations

Sensor-Collector approach brings a lightweight sensor to save resources. This setup does not provide all features and it have some limitations compared to collector-only deployments. In this section are described differences and limitations (possible issues) from standard deployment.

Network connection

Use of static IP addresses is highly recommended, because change of IP Address may casuse break connection between sensor and collector. This includes DHCP renewal, changes in firewall configuration, or changes in NAT configuration. It could be partially solved by extension of IP Address range in firewall settings (not recomemended).

All sensors must be disconnected and reconnected when collector changes its IP Address. Download of system logs is not possible when bi-directional connection (NAT) is not available.

System functions

There are some currently unsupported functionalities: * download of PCAP files from sensor * configuration of mshell users on sensor * changes of license including its extending on sensor

The sensor-collector solution require time synchronization using NTP server for its proper operation (change of NTP server is currently not possible on sensor side).

Copyright © 2024 Greycortex

4.14 SNMP Monitoring

4.14.1 SNMP Monitoring Guide

Simple Network Management Protocol (SNMP) is an Internet Standard protocol used for collecting information from managed devices.

SNMP Monitoring is useful for anyone who is responsible for servers and network devices, it lets you keep track of the traffic on the network.

Steps to set up SNMP Monitoring:

1. First of all you have to go to **Settings** and afterwards to **Sensors & Collectors** under the **System** tab.
2. In the **Topology** you have to choose which sensor you want to set up for monitoring.
3. Click on the to modify selected sensor.
4. Go to **Components** and click on SNMP Monitoring's button.
5. Write down username. Only alphanumeric characters are allowed.
6. Choose security level.
7. Depending on security level you've chosen, set up your authentication password and privacy password. Minimum password lenght is 8 characters.

Protocols:

Security levels:

If all previous steps are finished successfully, the state of component **SNMP Monitoring** should change in next minute to active.

Copyright © 2024 Greycortex

4.15 SNMP Trap

4.15.1 SNMP Trap Guide

Simple Network Management Protocol (SNMP) is an Internet Standard protocol used for collecting information from managed devices.
SNMP Trap are alert messages sent from a SNMP **Agent** to a central collector (SNMP **Manager**). It is important to be able to view these information when managing your data, as they provide device information on metrics that could impact your applications.

Steps to set up SNMP Traps:

1. The first step is go to the **Settings** and visit the **Data** page under the **Exports** tab.
2. Click on a button.
3. Write down a **Name** for the SNMP Trap.
4. Select an in the **Subject**.
5. Select the in the **Type**.
6. Select an in the **Interval**.
   Trap can only send data as often as it is set to, even if the values are exceeded sooner.)
7. Click . Afterwards click on (plus sign).
8. It is going to show you these information:

- Name --> Is set from your previous configuration however you can change it at anytime.
- Description --> You can write basic information about your SNMP Trap.
- Valid Until --> Means the time until the SNMP Trap occurs to end the period of validity.
- Interval --> Is set from your previous configuration however you can change it at anytime.
- Filter Name --> In this option you can set to filter what you want to keep an eye on.
- Record Count --> In this option you can set how many alert messages you can get during the interval you set.
- Export Time Zone --> Choose your time zone.
- Servers --> In this option you have to assign an IP address or a hostname for SNMP Manager.
- Username --> In this option you have to insert an username for the SNMP Manager.
- Authentication Protocol --> If it is set to none, communication is without authentication.
  The protocols for Authentication are **MD5** (message-digest) and **SHA** (Secure Hash Algorithm).
- Authentication Password --> Minimum password lenght is 8 characters.
- Privacy Protocol --> If it is set to none, communication is without privacy.
  The protocols used for Privacy are **DES** (Data Encryption Standart) and **AES** (Advanced Encryption Standard) protocols.
- Privacy Password --> Minimum password lenght is 8 characters.

Protocols:

If all previous steps are finished successfully, click on Enable checkbox.
You are all set now!

Copyright © 2024 Greycortex

4.16 Storage Wipe

4.16.1 Storage wipe guide

In this guide, you will learn how to wipe out disks after a PoC, to erase all data from the customer’s network.

Please note: During this process, any virtual disks will be deleted and you need to recreate them for the next deployment.

Furthermore; if the disk does not support cryptographic erase (instant), you will have to use the clear operation, which can take approximately 1.5 hours for the 1.2TB SAS drive.

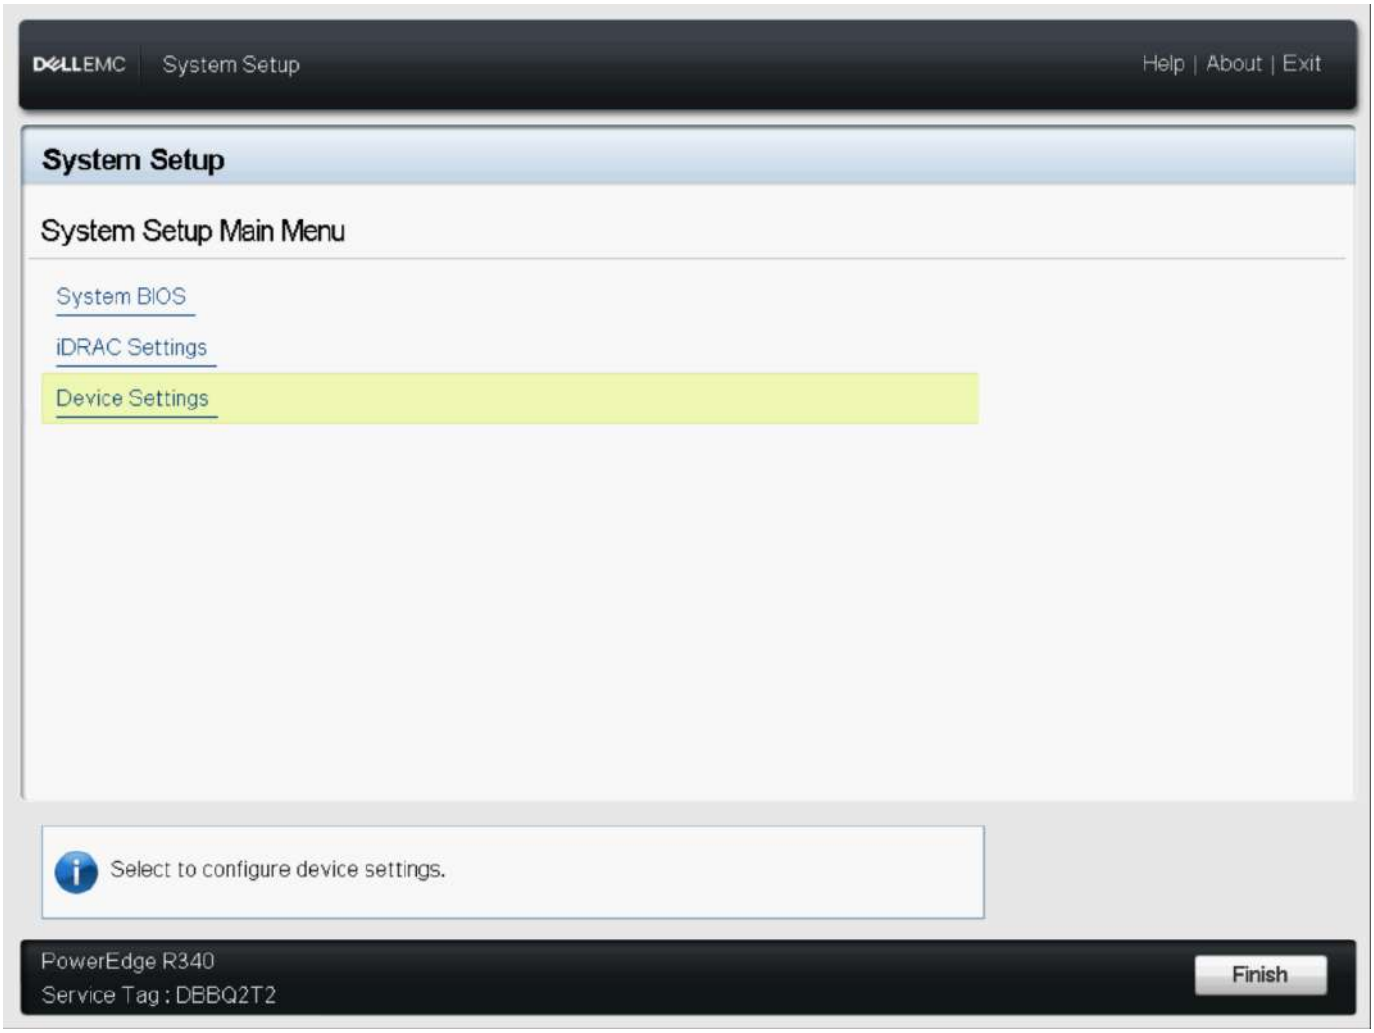
The process of storage wipeout is as follows:

- Boot into the iDRAC System Setup menu.



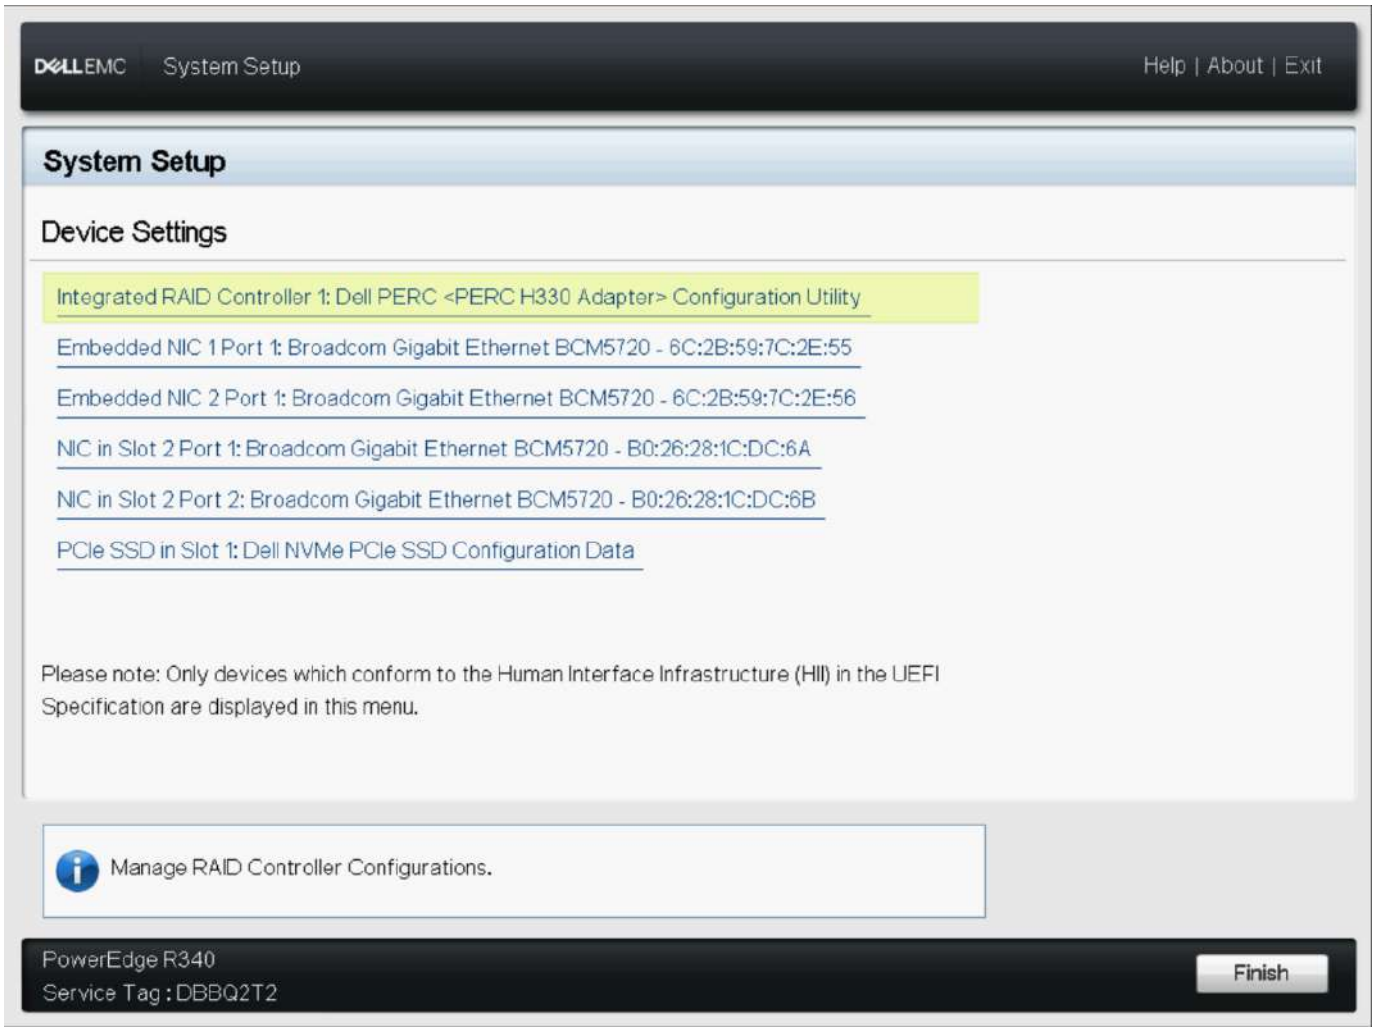
- Select Device Settings .

Copyright © 2024 Greycortex



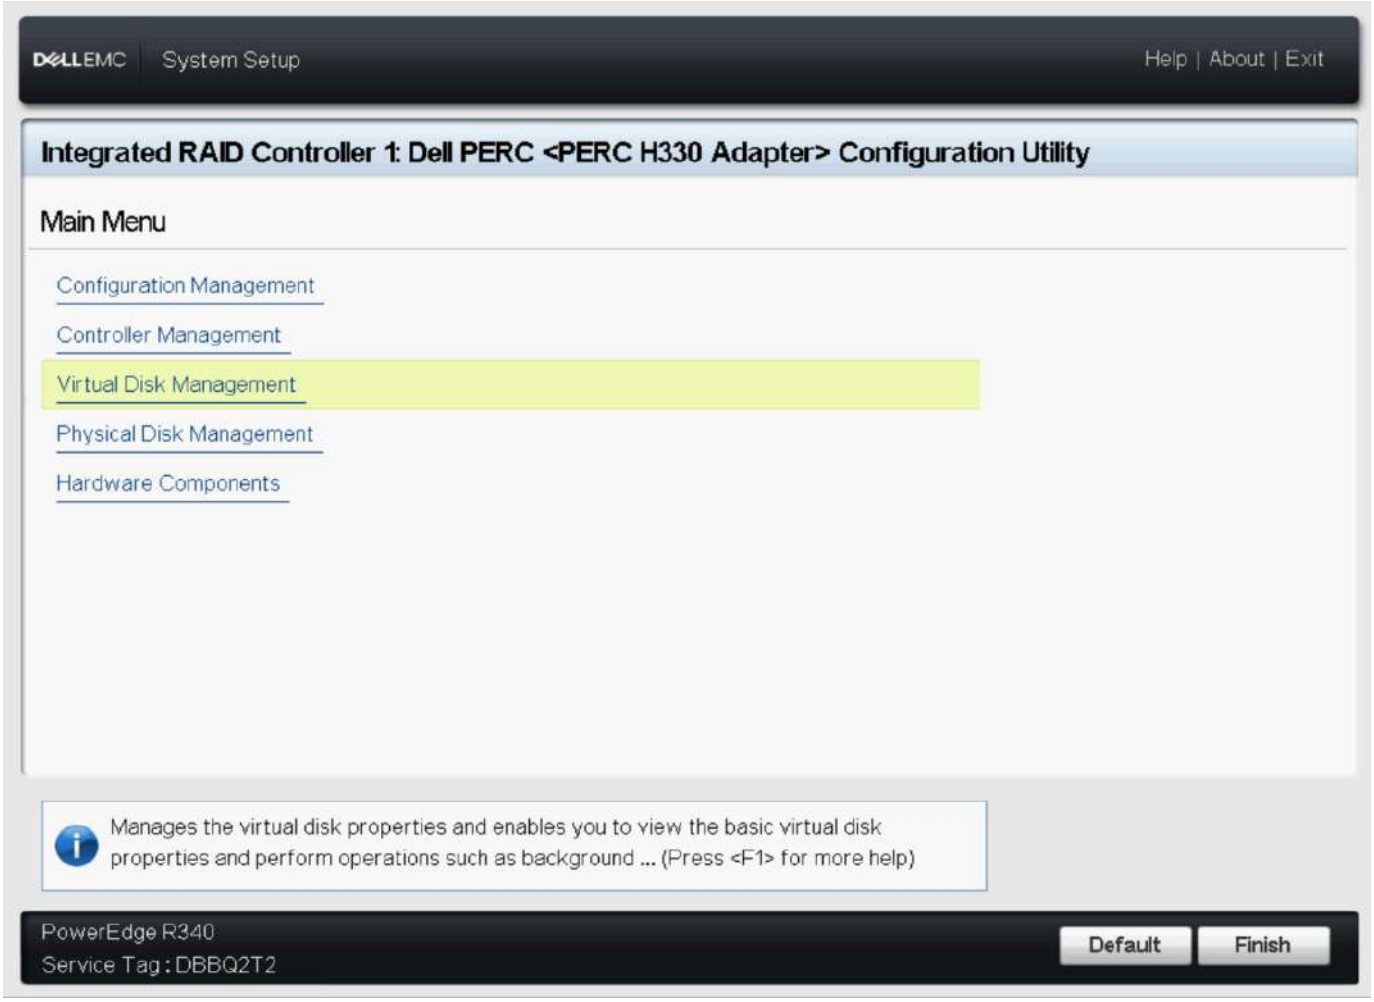
• Select your RAID Controller.

Copyright © 2024 Greycortex



• Select Virtual Disk Management.

Copyright © 2024 Greycortex



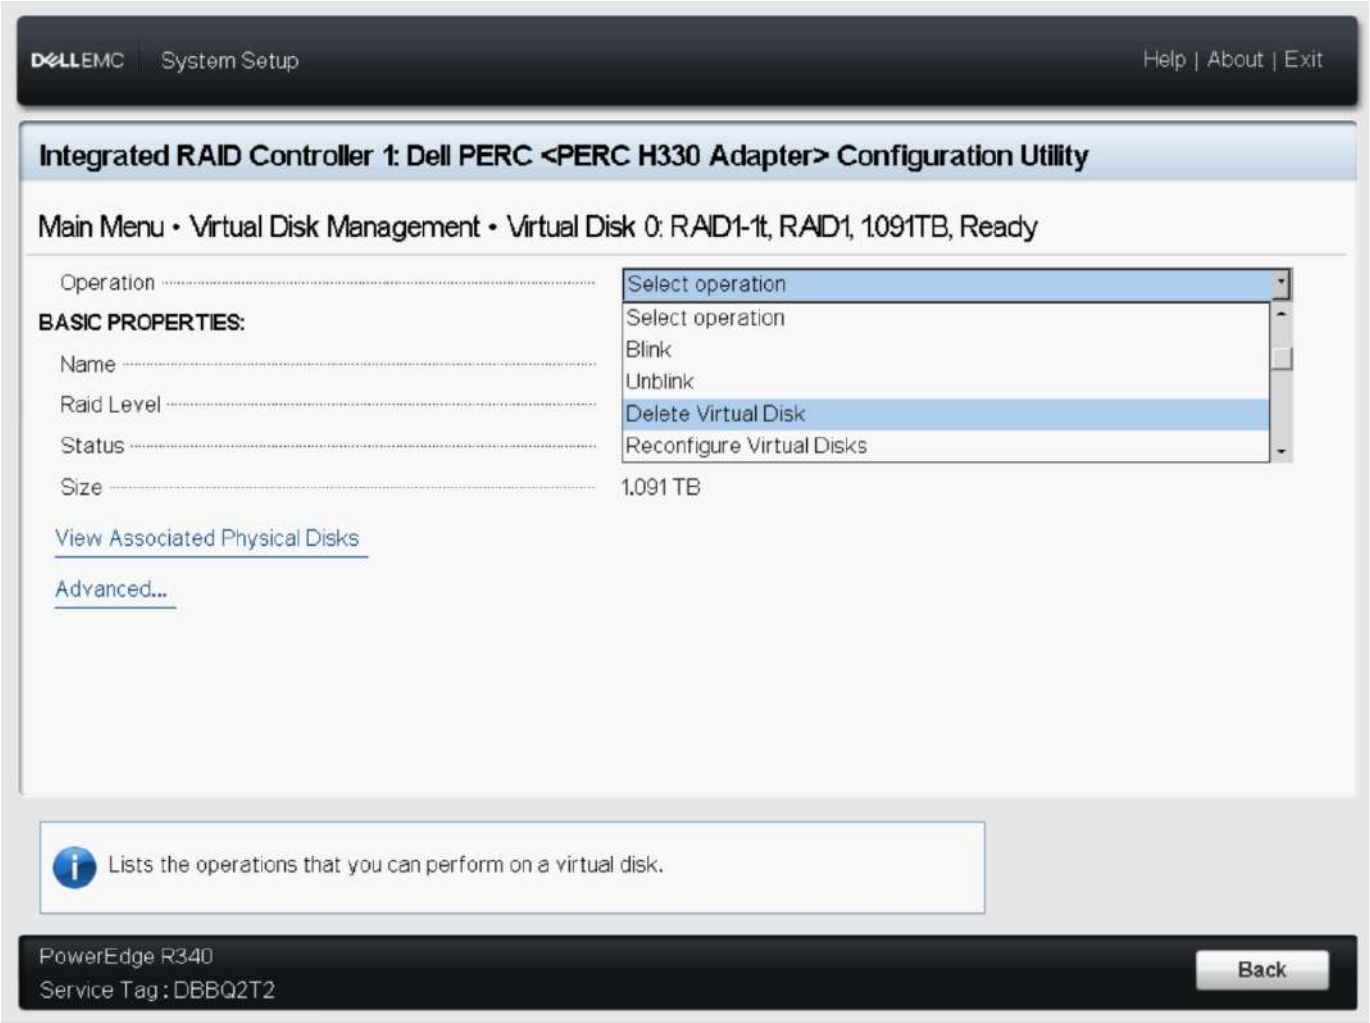
• Select your Virtual Disk/Disks.

Copyright © 2024 Greycortex



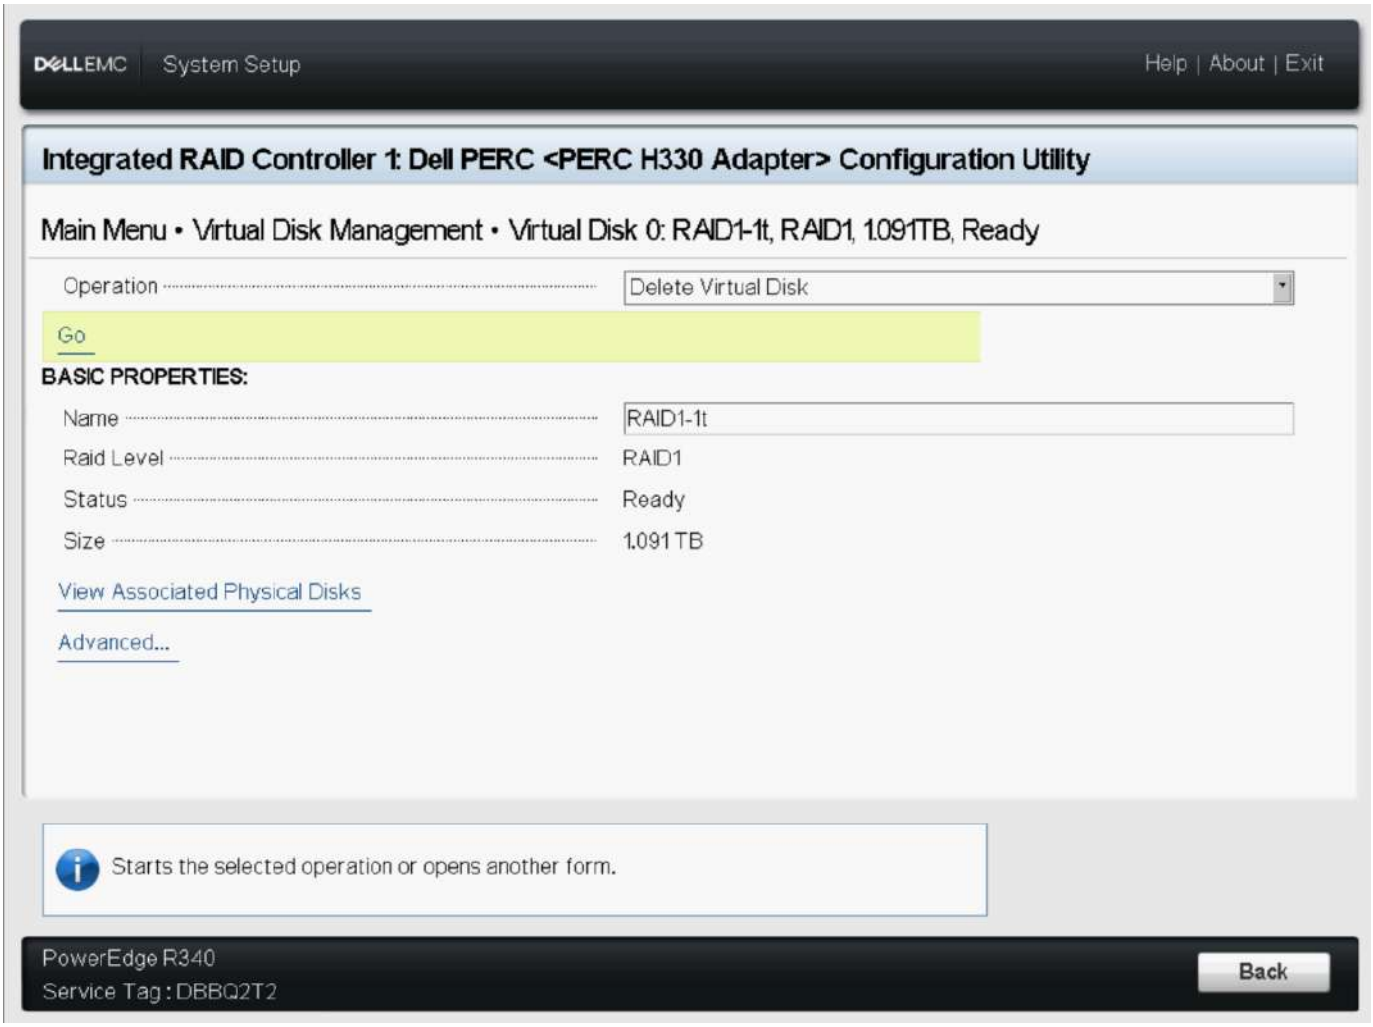
• Select “ Delete Virtual Disk.”

Copyright © 2024 Greycortex



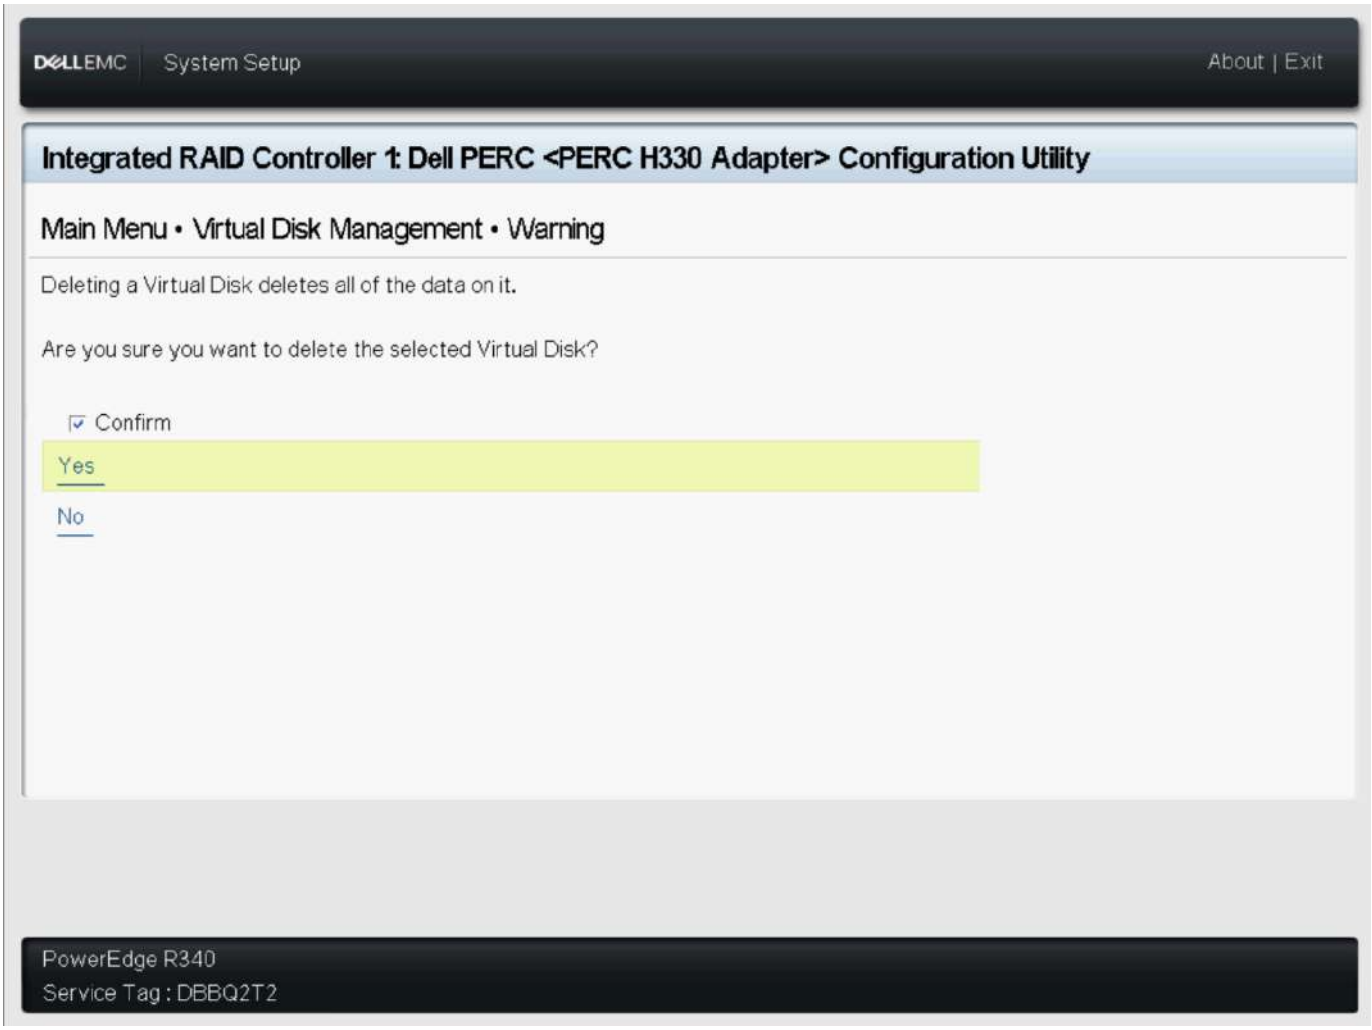
• Start the operation.

Copyright © 2024 Greycortex



• Click to confirm the operation, then select “ Yes .”

Copyright © 2024 Greycortex



• Wait for the success prompt, then click .

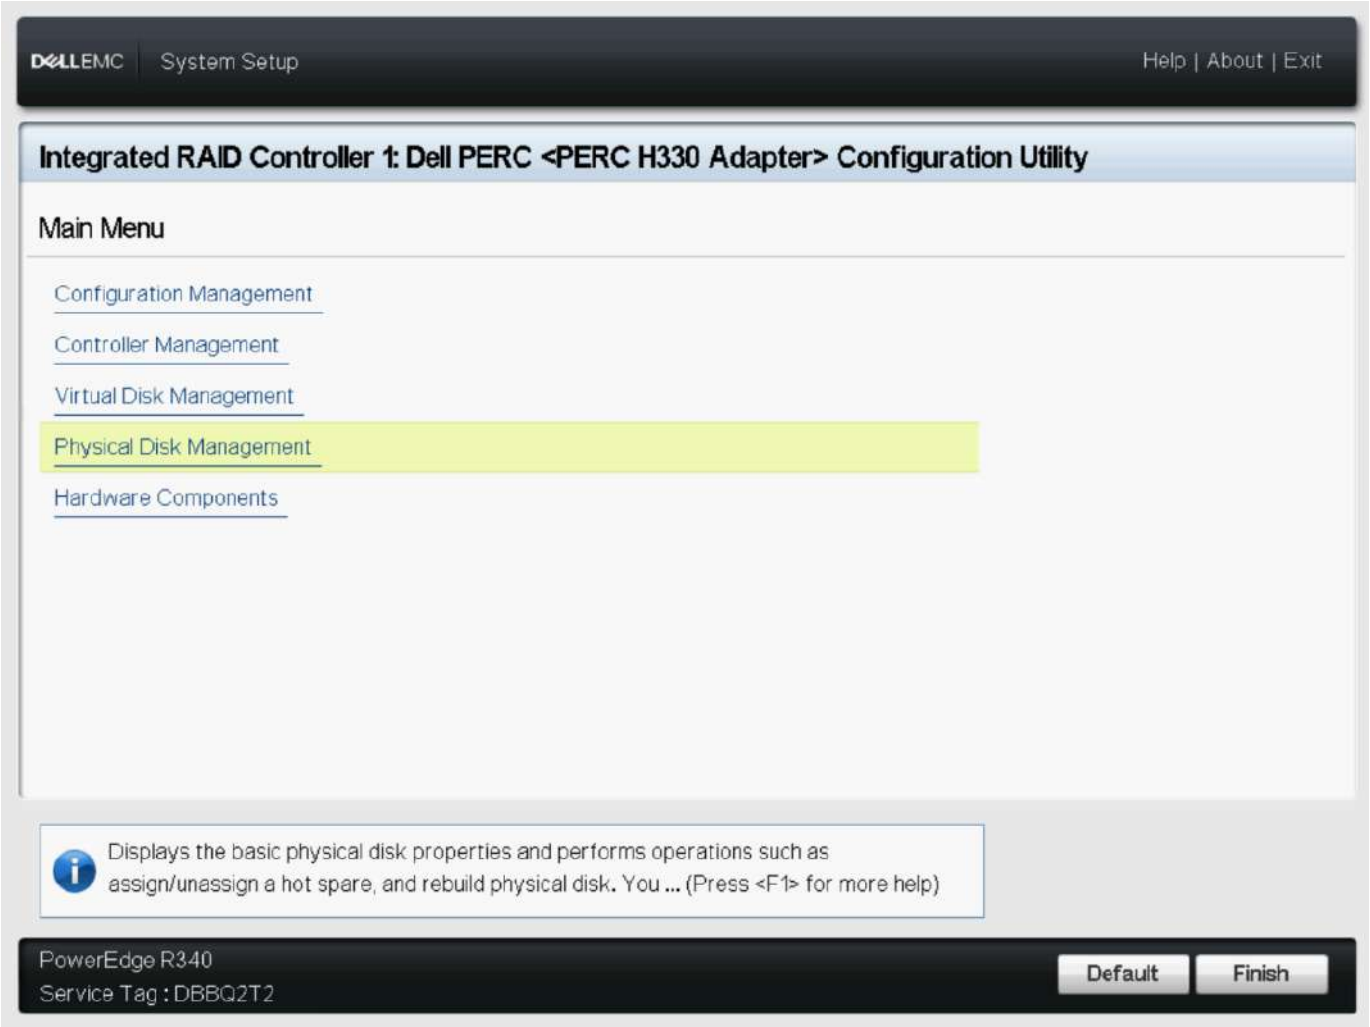
• Repeat these steps for other Virtual Disks, if present.

Copyright © 2024 Greycortex



• Select “ Physical Disk Management .”

Copyright © 2024 Greycortex



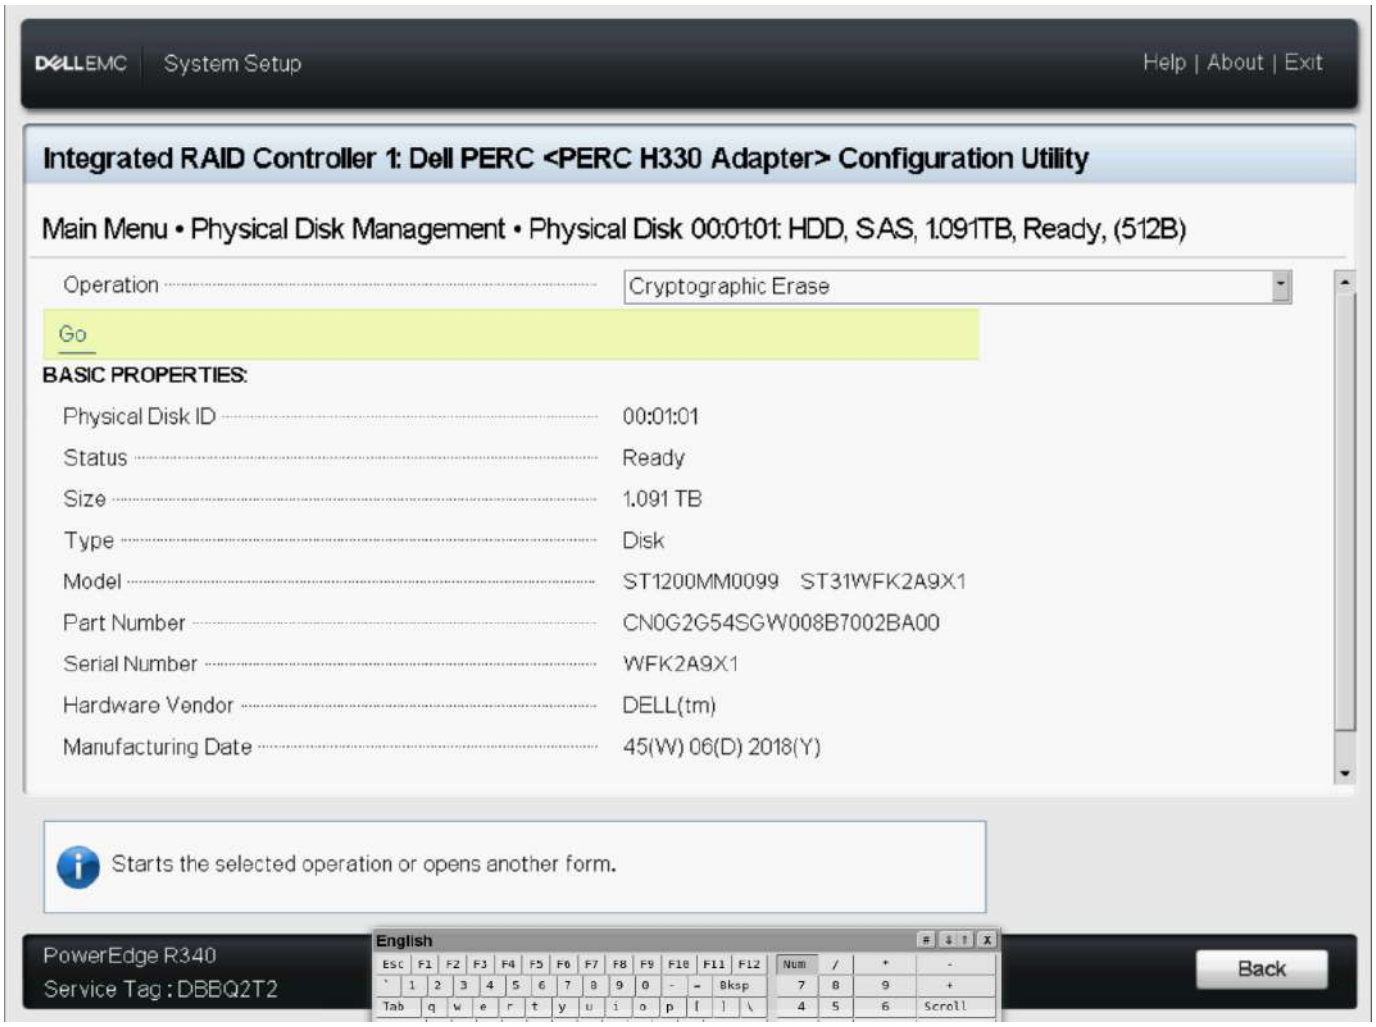
• Select the first Physical Disk.

Copyright © 2024 Greycortex



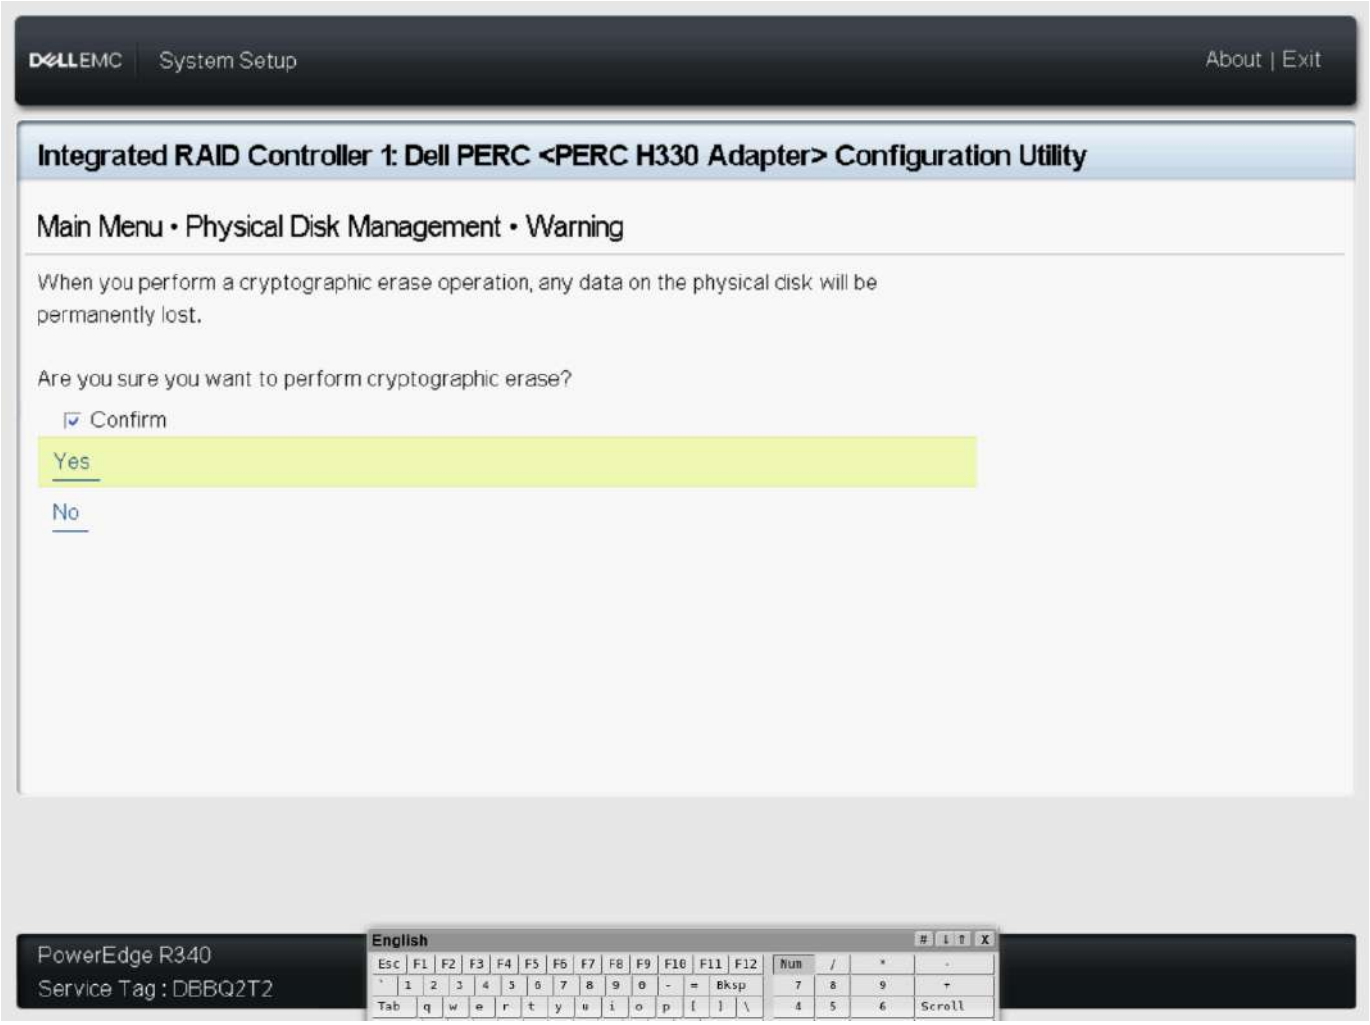
• Choose “ Cryptographic Erase ” if the disk supports, it and select “ Go .” (if not, skip to Step 17.) **This operation can be run simultaneously on all disks.**

Copyright © 2024 Greycortex



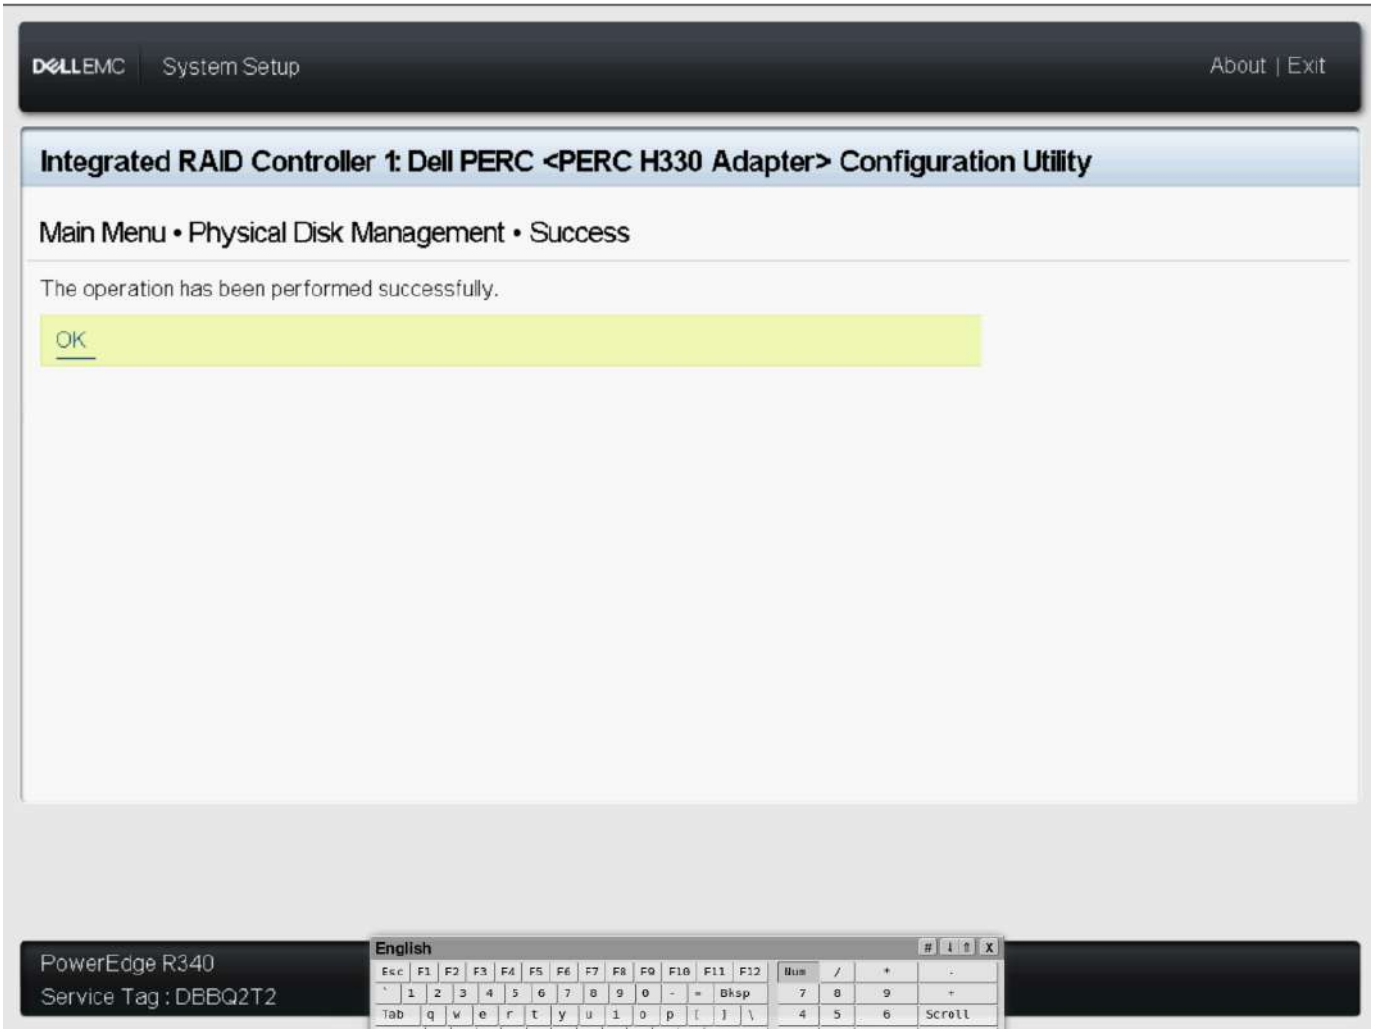
• Click to confirm the operation, then select .

Copyright © 2024 Greycortex



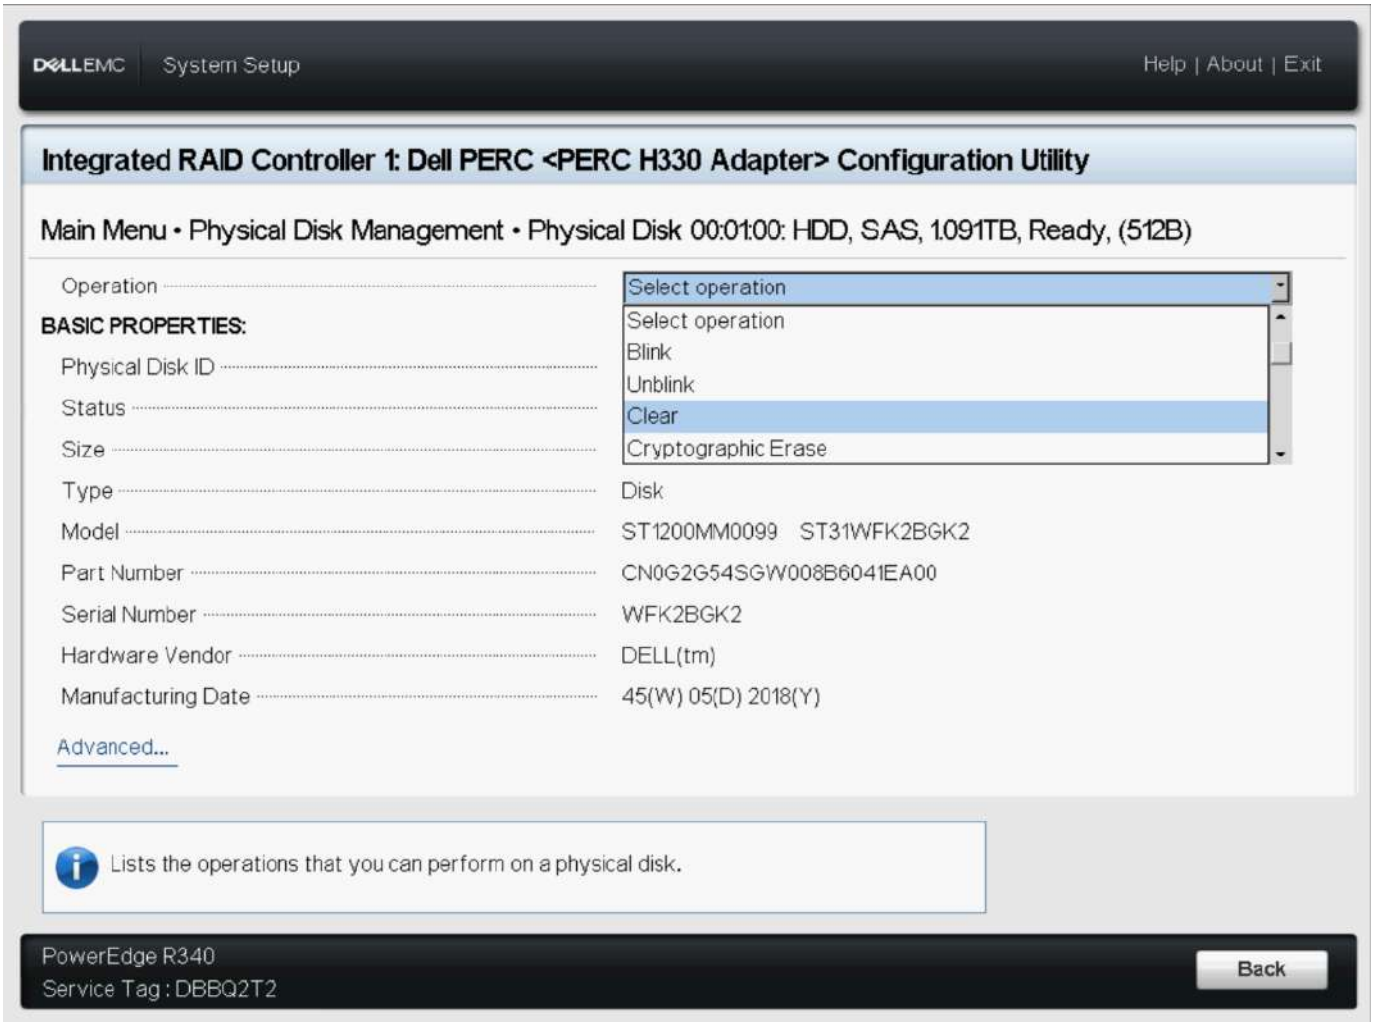
• Wait for the success prompt and click .

Copyright © 2024 Greycortex



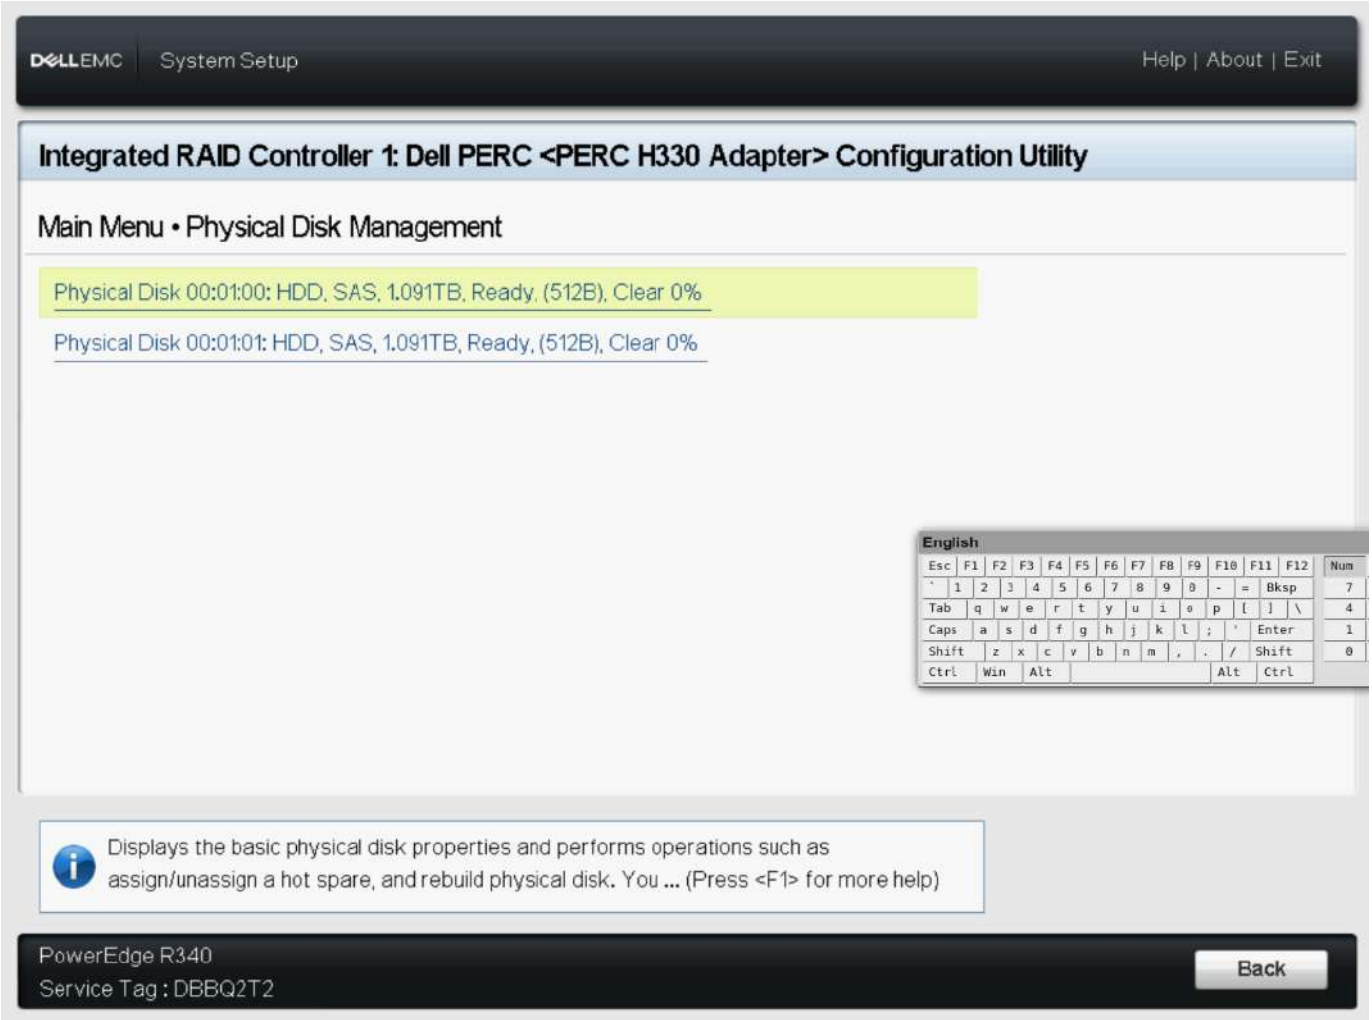
- Repeat for ALL physical disks you want to cryptographically erase.
- Select the Clear operation and click , confirm the operation and select , then wait for the success prompt. This operation can be run simultaneously on all disks.
- Repeat for all physical disks you want to wipe

Copyright © 2024 Greycortex

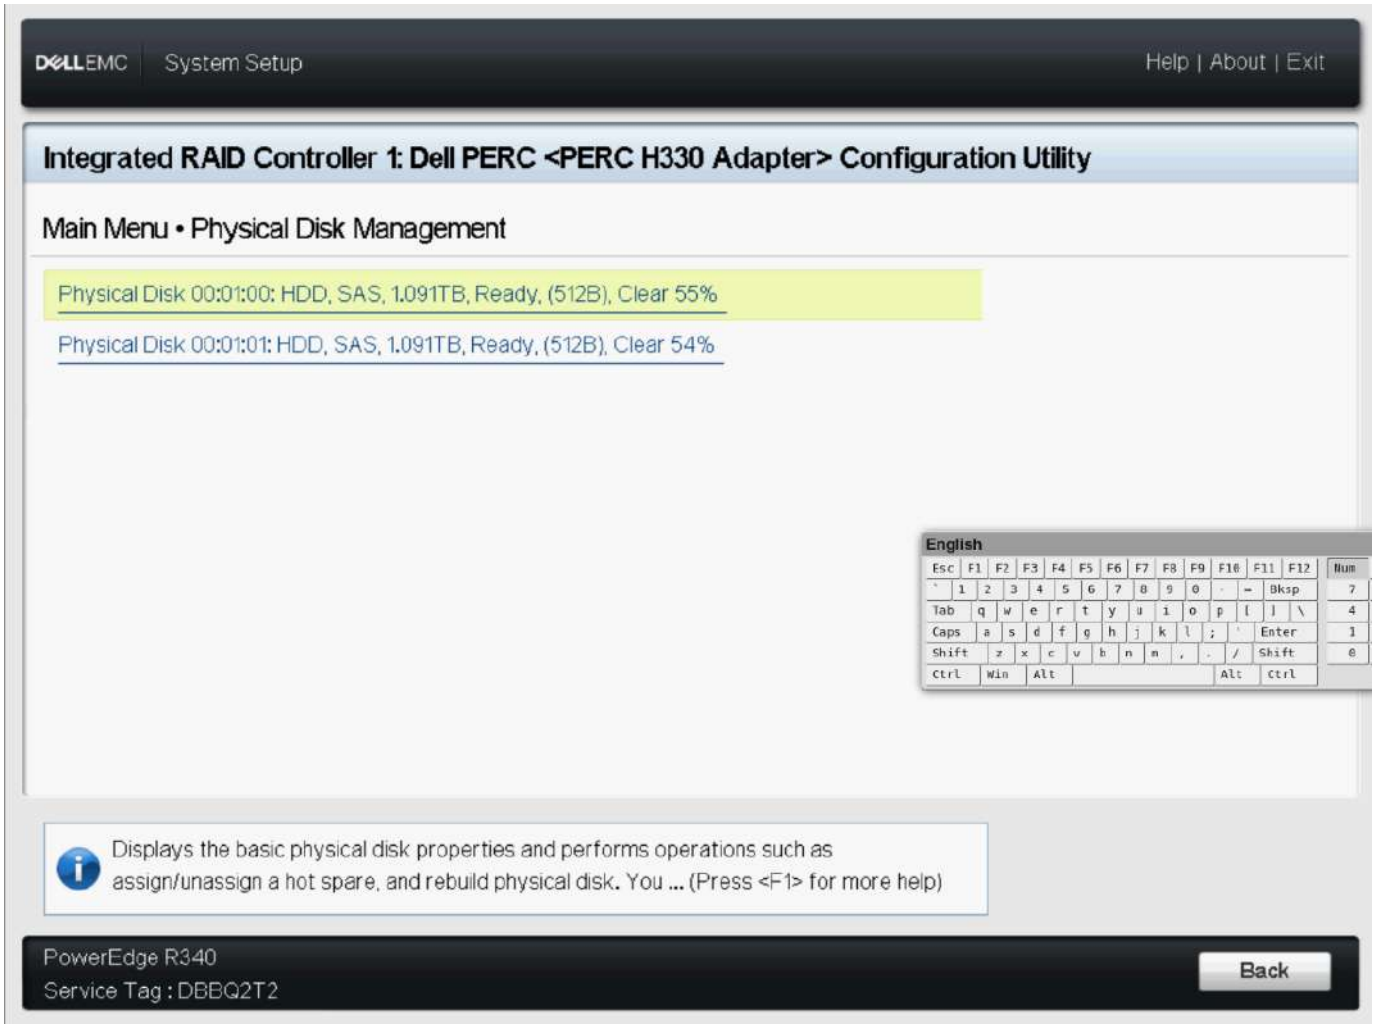


• If you now look at the Physical Disk management , you will see the progress of the clear operation (or operations) which are running.

Copyright © 2024 Greycortex



Copyright © 2024 Greycortex



• Once the progress of the clear operation hits 100%, the disks will be filled with 0s.

• After the completion of the cryptographic erase operation, the disks will be cryptographically erased.

Copyright © 2024 Greycortex

4.17 Universal Tagging Engine

4.17.1 Universal Tagging Engine Guide (UnTE)

Dynamic visibility by tracking new activities or changes caused by devices communicating in your network. A completely new engine that brings a manual or automated way of tagging hosts or subnets through a system of user-defined rules with easy-to-understand syntax.

Tag is unique combination of network entity identifier (host_id, subnet_id), name and description object.

EXAMPLES OF SIMPLE TAGS

host_id	Name	Meaning
melanie, 10.0.0.1	DNS Server	machine 10.0.0.1 on sensor melanie acts as DNS server
melanie, 10.0.0.2	Domain controller	machine 10.0.0.1 on sensor melanie acts as Domain controller
melanie, 10.0.0.3	My pink computer	manually assigned tag my pink computer to machine 10.0.0.3 on sensor melanie
melanie, 10.10.0.1	DHCP server	machine 10.10.0.1 on sensor melanie acts as Domain controller

EXAMPLES OF TAGS WITH ADDITIONAL DESCRIPTION

FLAGS

Every tag can contain various flags with defined meaning.

Flag	Meaning
	Tag was assigned automatically by UnTE
	Tag was assigned manually by mendel user
	Don't show this tag in UI. This is usable for tag detected by UnTE, which user doesn't want to see anymore for some reason.

Copyright © 2024 Greycortex

HOW TO WRITE THE RULE AND HOW IT WORKS

Each rule consists of the following mandatory parts:

Part	Description
	Name of the rule, used for human identification, logging.
	What kind of data are processed by this rule.
	Predicate matching the data from entry point. Rule may contain multiple steps.
	If data matches all predicates, all provided actions are triggered in consecutive order. Rule may contain multiple actions. Difference between action and step is step is side-effect-free while action causes persistent changes.

EXAMPLE OF THE RULE

UnTE rule evaluation

is using rules defined in UnTE signatures and processes various input objects. High level view on evaluation is described by following pseudo code algorithm:

UnTE rule execution

UnTE engine is executed every minute. Execution can be described by following pseudocode.

RULE REQUISITES

Root

Type of the object supported by the rule.

Object	Description
hosts	Host existing in ti.hosts_base table
events	Event reported in current minute (ui.tab01)
flows	Flow processed in current minute (nb.flows01)

Copyright © 2024 Greycortex

Step(s)

Steps are dependent on the root. Only steps supporting certain root value can be used in the rule. This is enforced by grammar. When the argument contains +, multiple values separated by semicolon are accepted.

Step name	Arguments	Description
	Service_type, String+	Check whether there exists row in ti.host_services_base having given service type and one of given detected services. Save information about service into data log.
	Service_type, Trans_proto, Int+	Check whether there exists row in ti.host_services_base having given service type, transport protocol and one of given ports. Save information about service into data log.
	String	Check whether host has assigned tag with given name.
	String+	Check whether host belongs to given subnets.

Step name	Example
	host_has_service(LOCAL, "HTTP", "HTTPS")
	host_has_service(LOCAL, TCP, 80, 443)
	host_has_tag("sample_tag")
	host_filter_subnet("10.1.0.0/16", "10.2.0.0/16")

Step can be used multiple times on single line with || operator:

Copyright © 2024 Greycortex

Steps for flows root

Action name	Arguments	Description	Example
	Direction, String	True if src/dst host has assigned tag. Adds information about tag into data log.	flow_filter_tag(SRC, "sample_tag")
	Direction, ! String	True if src/dst host has NOT assigned tag. Adds information about tag into data log.	flow_filter_tag(SRC, ! "sample_tag")
	Direction, String	True if src/dst host belongs to subnet.	flow_filter_tag(DST, "10.1.0.0/16")
	Direction, ! String	True if src/dst host does NOT belong to subnet.	flow_filter_tag(DST, ! "10.1.0.0/16")
	Direction	True if src/dst host exists in ti.hosts_base.	flow_filter_internal_host()
	String+	True flow contains one of given services.	flow_filter_service("HTTP", "HTTPS")
	Trans_proto, Int+	True if flow contains given transport protocol and one of destination ports	flow_filter_service(TCP, 80, 443)
	String+, Trans_proto, Int+	Overloaded combination of 2 steps above.	flow_filter_service("HTTP", "HTTPS", TCP, 80, 443)
	Flow_field_selector, Comparision_operator, Int Size_unit	True if selected field compared by operator wih given size matches	flow_filter_metric(DST_OCTETS, <=, 50 MB)

Step for events root

Action name	Arguments	Description	Example
	Int+	True if event has one of given sids	event_filter_sid(1001, 1002)
	Int+	True if event has one of given severity	event_filter_serverity(1,2,3,4)
		True if event has severity <=0 or at least one FP was applied on it	event_ignore_fp()
	Traffic+	True if traffic of event matches one of given parameters	event_filter_traffic(INBOUND)
	Event_regex_field, String	True if given field matches regex	event_filter_regex_match(EVENT_NAME, "Discovery:.*")
	Int+	True if src or dst ASN number matcher param	event_filter_asn(18176, 4455)
	Duration	True if event (src_ip_addr, dst_ip_addr, sid) wasn't reported for at least given time	event_appeared_for_first_time('10 hours')

Copyright © 2024 Greycortex

ACTION

Actions is dependent on the step a thus transitively on the root. Therefore it's not possible to mix actions from different roots.

Actions for hosts root

Action name	Arguments	Description	Example
	String	Remove all tags having given tag tag name from host	host_remove_tag("sample_tag")
	String	Create host tag with parameters: name, description, duration	host_create_tag("sample_tag")
	String, Duration	Create host tag with parameters: name, description, duration	host_create_tag("sample tag", 10 hours)
	Frequency, Sid, Service_type	Report event with given Sid and service_type. When the frequency is ONCE, event will be reported only when tag is detected for first time. Other event parameters: sensor, dst ip/subnet; service information	host_report_event(ONCE, -12345, LOCAL)

Actions for flows root

Action name	Arguments	Description	Example
	Direction, String, Duration	Create tag on src/dst host with parameters: name, description, duration	flow_create_host_tag(DST, "sample_tag", 20 minutes)
	Direction, String, Duration	Create tag on src/dst host with parameters: name, description, duration. Tag is created with given flag.	flow_create_host_tag(DST, 20 minutes, "sample_tag", HIDE)
	Direction, String	Create tag on src/dst host with parameters: name, duration, description	flow_create_host_servcie_tag(DST, "sample_tag")
	Frequency, Direction, Sid, Service_type	Report event with given Sid and service_type. When the frequency is ONCE, event will be reported only when tag is detected for first time. Other event parameters: sensor, src/dst ip/subnet, service information, traffic	flow_report_event (ONCE, DST, SID, LOCAL)

Actions for events root

Action name	Arguments	Description	Example
	Direction, String, Duration, Flag	Create tag on src/dst host with parameters: name, description, duration	create_host_tag(DST, "sample_tag", '1 day') create_host_tag(DST, "sample_tag", '1 day', HIDE)
	Sid	Report event with given sid. All other params are taken from processing log.	report_event_once(SID)

UnTE rule evaluation

is using rules defined in UnTE signatures and processes various input objects. High level view on evaluation is described by following pseudo code algorithm:

Copyright © 2024 Greycortex

4.18 Cloud Deployment

4.18.1 Amazon Web Services

Cloud deployment - Amazon Web Services

Deployment of the Mendel network security solution into the AWS cloud services is able to be covered by multiple options:

Option A - With Greycortex partial support

- Deployment steps 1-3 are done by GCX
- Deployment steps 4-5 are done by partner/customer (owner of the target account)

Option B - With Greycortex full support

- All deployment steps are done by GCX (it's needed to share access to target account with GCX)

Option C - Without Greycortex support

- All deployment steps are done by partner/customer

Deployment steps	
1. Create Mendel image	RAW image type is required (common mendel ISO image has to be converted)
2. Upload Mendel image into AWS S3 storage bucket	Upload RAW image to AWS S3 bucket; Import RAW image as EBS snapshot into EC2 cluster; Register EBS snapshot to EC2 to AMIs (validated Amazon Machine Image)
3. Share Amazon AMI image	Copy AMI image to another AWS region (same with destination acc); Source account steps (GCX); Destination account steps (Partner/Customer)
4. AWS EC2 Instance	Launch a new instance from AMIs and configure details; Management interface setup; Secondary management interface setup (optional)
5. AWS Traffic mirroring configuration	Launch an AWS MENDEL instance with 1 Ethernet interface (step 3 above); Create a new interface for mirror traffic (SPAN); Adding secondary interface (optional); Setup mirroring

For detailed step by step procedure see AWS deployment in depth.

Pre-conditions:

- Virtual instance needs to meet below criteria: EC2 NITRO specific instances (Instances built on the Nitro System)
- The AMI image is accessible only in the same zone (region)

Copyright © 2024 Greycortex

4.19 Operational Technologies

4.19.1 OT Metrics

Introduction

Where **IT** refers to , **OT** is refering to . Although these abbreviations are similar the industries behind them are very different. The term **OT** was introduced to visualize the differences between traditional **IT** systems and Industrial Control Systems environment.

Some examples of **OT** assets are:

- Programmable logic Controller
- Supervisory Control And Data Acquisition
- Distributed Control Systems
- Computer Numerical Control systems, including 3d printers, printing presses and other computerized machining tools.
- Building Management and Building Automation systems
- Energy monitoring

Operational Technology is widely used in oil refineries, (nuclear) power plants, HVAC-systems in (public) buildings, (railway) transportation, etc. and as such has become a common, crucial element of critical infrastructure systems.

With the OT-Metrics module of GREYCORTEX MENDEL, you are able to process several common protocols used in these applications. The solution processes both network and application data layer. These protocols include , , , , , , etc.

Info

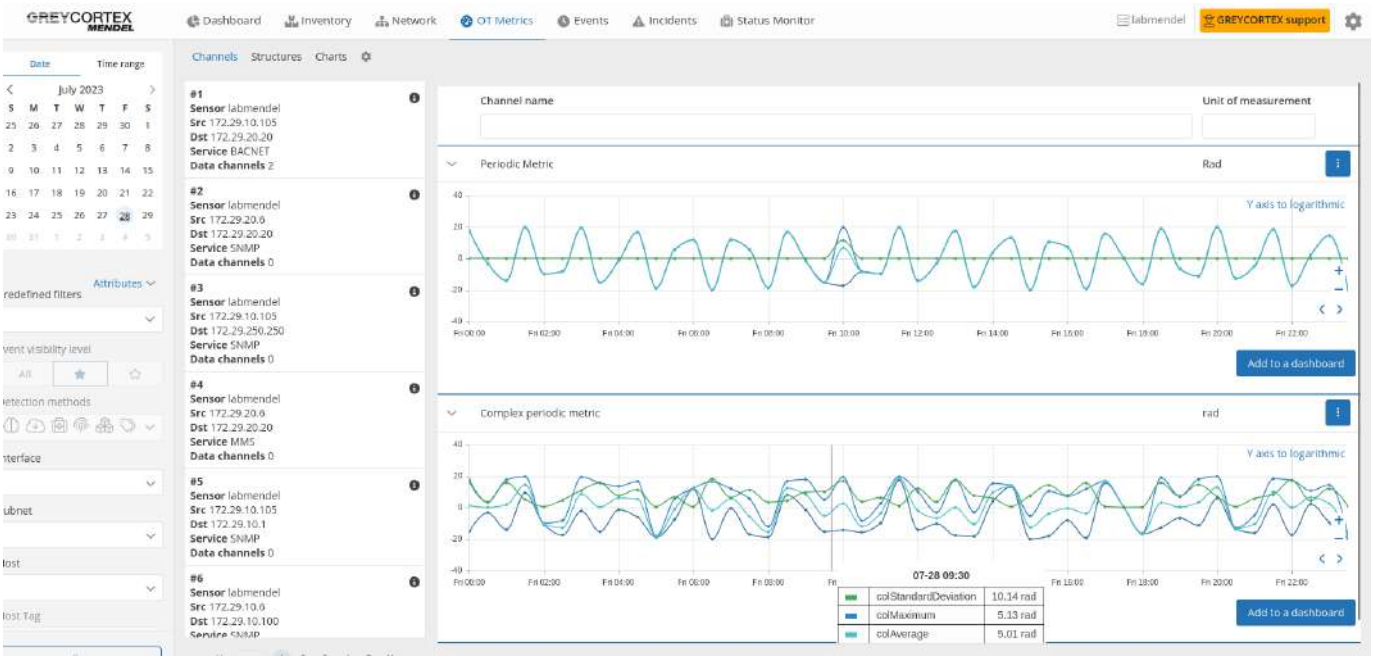
This document guides you through the terms used by GREYCORTEX related to this module, configuration, and description of the module's user interface.

SCOPE

OT-metrics module focuses on OT-protocols used commonly in the following verticals:

- Facility Management or HVAC/FMIS -
- Industrial Manufacturing & Processing -
- Energy Distribution -

Copyright © 2024 Greycortex



Important

When you have a particular usecase for implementing the OT-metrics and your protocol is not in the list please contact your GREYCORTEX sales partner to request an implementation.

GLOSSARY

one source or destination of the network communication defined by its IP address and sensor. a logical connection of two hosts. These to hosts are expected to communicate. an object configured on host on the application layer. Example: logical node in MMS. data values of a structure in time. a storage for configuration files. Contains revisions of the same file in time.

CONFIGURATION

- OT-Metrics functionality is available for those who have it bought within their license. If it´s not avaiable for you, please contact GREYCORTEX support.
- For proper configuration, there should already exist hosts for all devices to be monitored by MENDEL. If the network traffic is being sent to MENDEL, it creates hosts in several minutes.
- Next step is to configure the . There are two ways how to configure them in MENDEL.
- *Automated detection* - When there already are flows existing in MENDEL for a host to be monitored, you can detect the links using an automated mechanism. It can be invoked on page by pressing button.
- *CSV import* - You can also configure links using CSV import. This can be done also on page. The format of the file is supposed to look like the following example:

Modbus CSV-file specification

What columns should Modbus CSV contain:

Copyright © 2024 Greycortex



MENDEL starts to parse application data now. When a new structure appears, it can be seen on dashboard. All structures have to be confirmed on dashboard in order to make MENDEL process application data for them.

(Optional) If MMS is supposed to be processed by MENDEL, the SCD files for all networks to be monitored have to be provided. The configuration can be done using . Each of the new files should be in a separate container. All structures configured by SCD are confirmed by default.

From now on you can see processed application data in dashboard.

User interface layout

CHANNELS

List of links is displayed on the left part of the page. For each link, you can see more detailed information such as

- sensor on which it was detected,
- source address,
- destination address,
- detected service,
- number of detected channels.

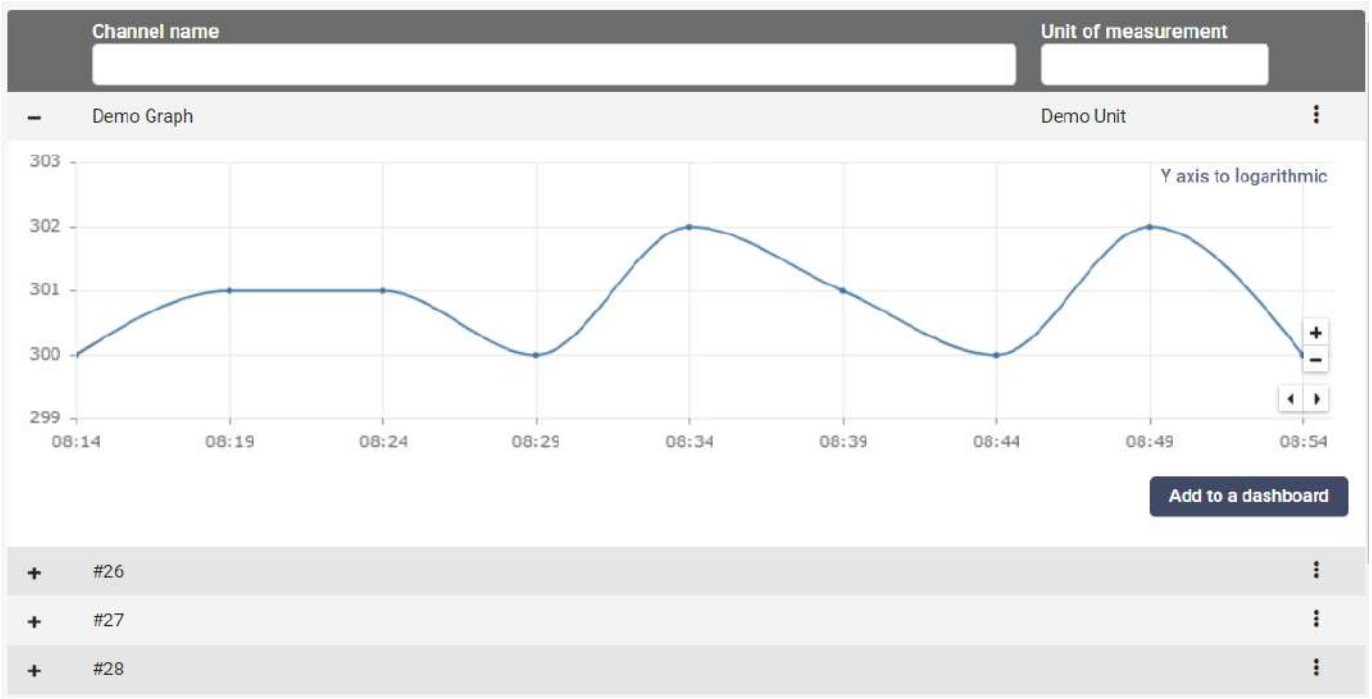
You can display more details by hovering above symbol in top right corner of the link.

Copyright © 2024 Greycortex



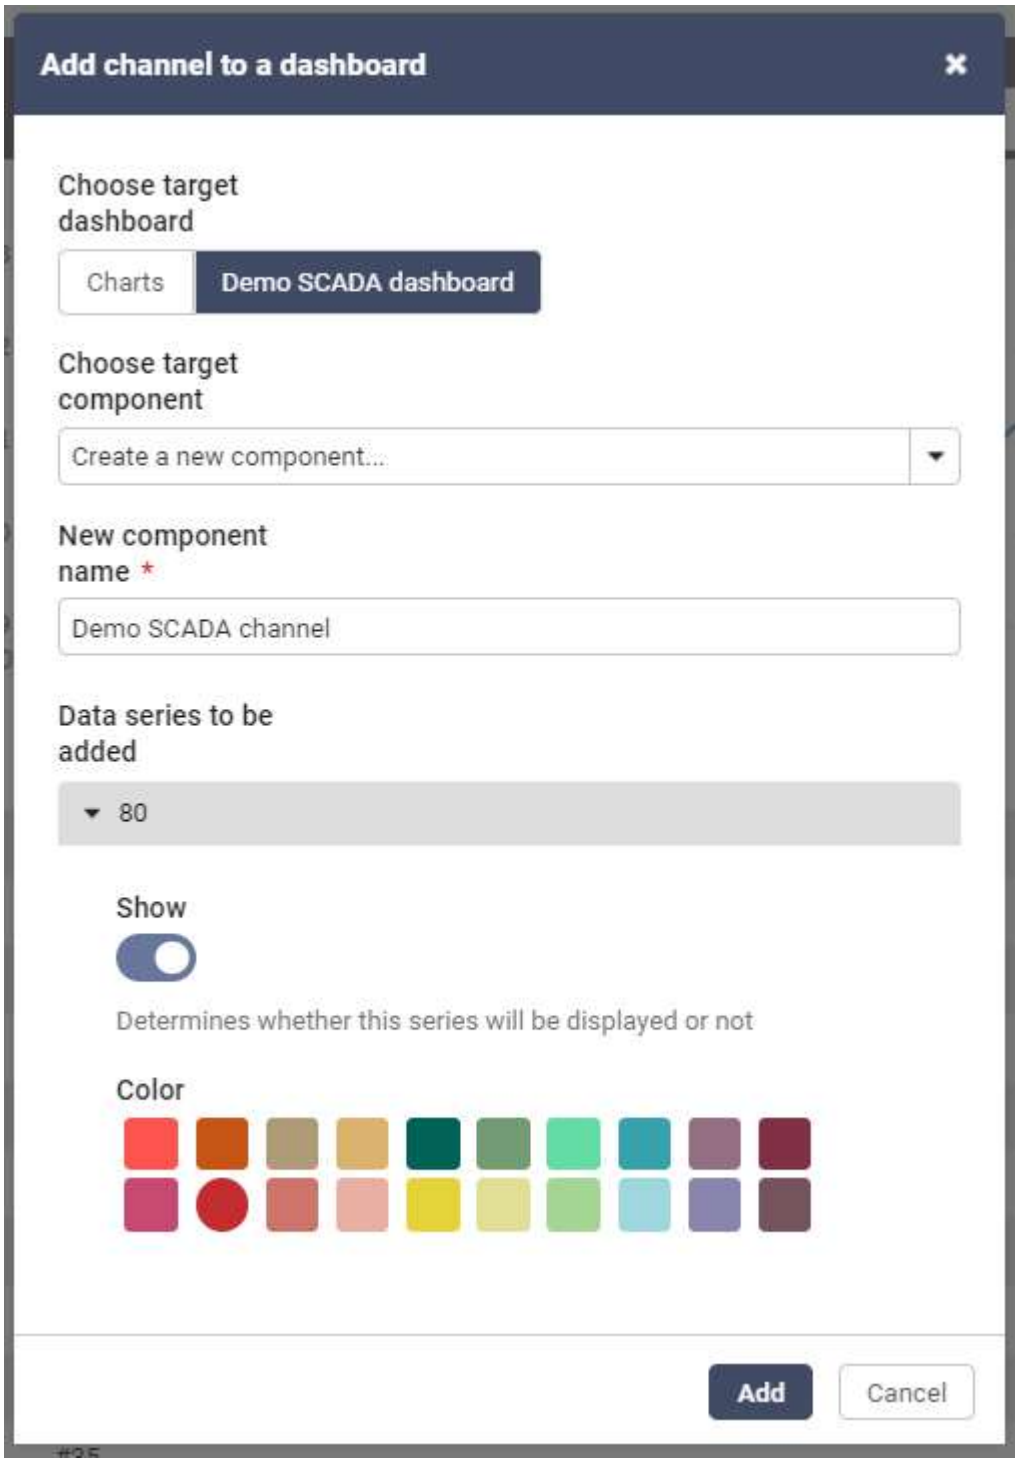
By clicking on the link, channels are displayed on the right part of the page. By expanding the channel row, you can see a graph of the data detected, but firstly, you have to enable the structure for the link in the Structures section.

By default, the channel name is in the format and there are no units, but you can change them manually. Editing menu is located in the three dots menu at the end of the row. Afterwards, it is possible to sort and search the channels table by these parameters.

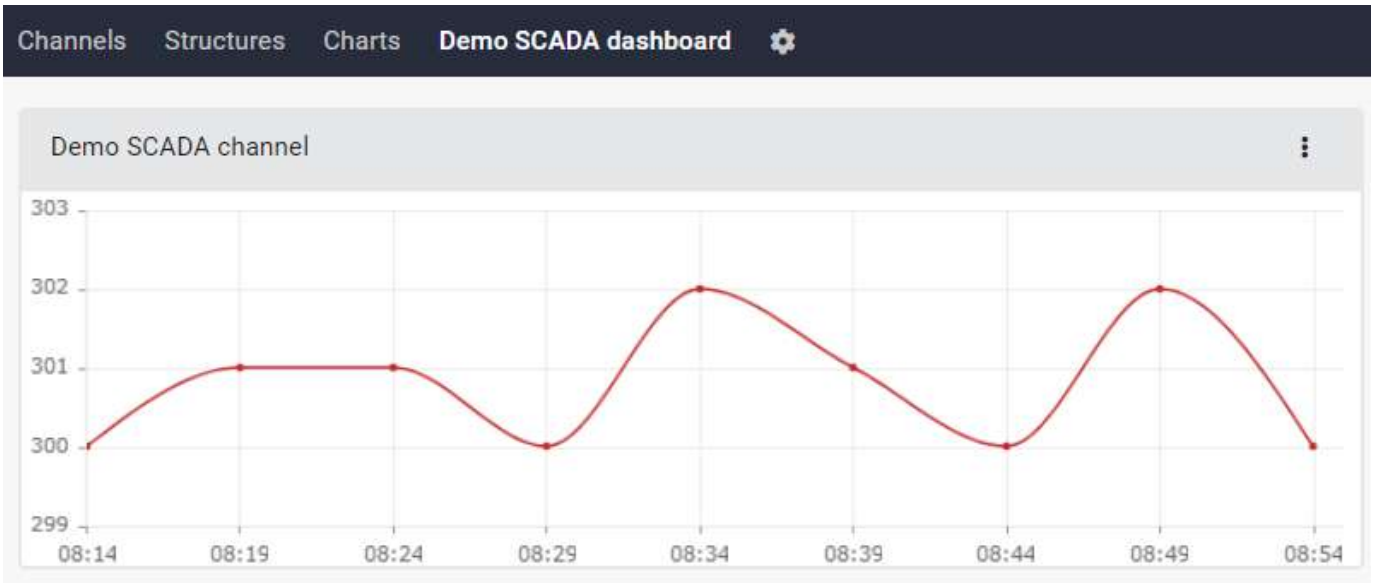


If a particular graph interests you, you can add it to the dashboard by clicking button. You can choose to which dashboard you want to add the graph, choose the component or create a new one, name the component and edit the preferences of the data series such as choose a colour or in case of multiple series, choose which to display.

Copyright © 2024 Greycortex



Copyright © 2024 Greycortex



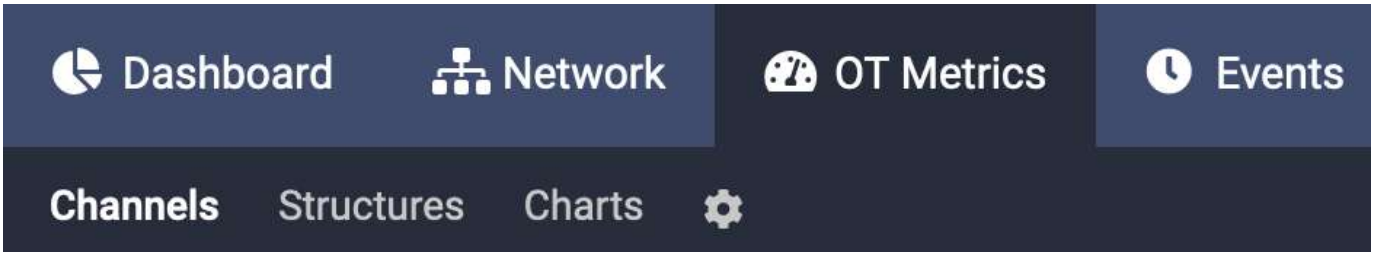
STRUCTURES

Each can have several structures. You can see the list of the by clicking on the link. The graph is a tree representing related data . Each tree needs initial review and confirmation. By confirming the tree, it becomes relevant to the data processing engine.



CHARTS

Default dashboard for the OT-Metrics components. You can modify, add new or remove dashboards using the Dashboard Manager.



To the OT-Metrics dashboards, you can add another OT-Metrics components as well as other components from MENDEL Dashboards (first tab in the menu), but on MENDEL Dashboards, you cannot add OT-Metrics components.

Copyright © 2024 Greycortex

4.19.2 OT use cases

This document elaborates on GreyCortex' target verticals for their OT-security platform and is targeting both end-customers as well as channel-partners & integrators. GreyCortex is focusing their OT-Mendel product on the following verticals:

- Energy & Utilities
- Healthcare
- Energy Distribution
- Industrial Manufacturing & Processing
- Maritime

Each vertical has its own specific use-cases and their own motivations for implementing and operating a security platform. This document describes a few of them, it is not an exhaustive summary.

Maritime

GreyCortex understands the IT/OT convergence in the maritime vertical comes with its challenges. Vessels are becoming more connected and data-driven. Adding more digital components, often required by law, will add complexity and potential vulnerabilities. With the Mendel product the Officer responsible for the vessels IT & OT infrastructure on board the vessel is able to get clear visibility on the current state of the on-board security situation from cyber perspective, while providing the on-shore MSP with the data needed to support the officer in making decisions in time-critical situations.

- The OT-Mendel product supports management of the martitime vertical with the right data to include OT-cybersecurity into the mandatory Maritime **Incident Response Plan**, **functional description documentation** documenting cyber-enabled assets & network diagrams & creating the **Plan of Actions & Milestones**.
- Provides high level & detailed knowledge required for the on-board crew and for the designated OT-officer.
- Support in documenting the high-level and low-level state of the IT & OT-infrastructure on board the vessel.
- what OT-assets are not connected to IT
- what OT-assets are unknowingly connected to IT
- what OT-assets are knowingly connected to IT.
- Provide insights on how OT systems are isolated from IT onboard the vessel.
- Provides real time security events to be sent to shore Managed Service Party and/or to the vessels incident response officer.

Energy & Utilities

Both energy & utilities sectors are typically categorized as being responsible for operating a country's **critical infrastructure**. Continuity & Stability are keywords that come to mind. In these sectors typical use cases are:

- **Asset inventory** GreyCortex OT-Mendel supports the verticals operator by providing knowledge what assets & networks are deployed in what purdue-level in the organisation. This helps the operator to understand his environment and determine priorities to successfully make a security risk assessment.
- **Asset details** Mendel enables energy & utilities operators in understanding the risk of each of their OT-assets & IT-network elements by reporting known security vulnerabilities (**CVEs & CPEs**) and potential misconfigurations of those assets. This enables this vertical to make an informed decision on how to execute patch-management versus the operational risk that involves.
- **Intrusion prevention** Mendel enables the OT-operators to communicate with the IT-operator of their critical network-elements to close and open firewalls for high risk traffic, lowering the overall risk-level of the IT/OT network.
- **Infrastructure weaknesses** Mendel leads the beginning operator in it's journey in identifying weak points in it's infrastructure.

Healthcare

- **Asset inventory** Mendel supports the operator by knowing what assets are deployed in what purdue-level in the organisation. This helps the operator to understand his environment and determining priorities to succesfully make a security and patch-management assesment.

Copyright © 2024 Greycortex

- **Asset details** Mendel enables the operator in knowing the risk of it's assets by reporting know security vulnerabilitie **CVEs & CPEs** and misconfiguration of those assets.
- **Data flow detection** GreyCortex Mendel will identity what patient data is shared with what part of the healthcare-organisation.
- **Intrusion prevention** Mendel enables the operator to close firewalls for high risk traffic, lowering the overall risk-level of the network.

Manufacturing

Industrial manufacturing networks are often vulnerable to both targeted and random cyber attacks. Even if these attacks are not always catastrophic, even the smallest malfunction can negatively impact regular business operations. How to execute OT-asset patch-management varies from business to business, Mendel is here to identify risks and provide insights in data-flows. This combination of informations enables the management of this sector to set cost-efficient priorities.

- **Asset inventory** Mendel supports the manufacturing operator by providing knowledge about what assets are deployed in what purdue-level in the organisation. This helps the operator to understand his environment and determining priorities to succesfully make a security and patch-management assesment.
- **Asset details** Mendel enables the operator in knowing the risk of it's assets by reporting know security vulnerabilitie **CVEs & CPEs** and misconfiguration of those assets.
- **Intrusion prevention** Mendel enables the operator to close firewalls for high risk traffic, lowering the overall risk-level of the network.
- **Asset Tagging** with clever automatic tagging of both OT & IT assets, the manufacturing integration patner can quickly deploy a custom set of risk-assesment signatures in case of an high priority deployment.
- **Mitre ATT@CK** GreyCortex understands that IT & OT-teams are not the same in their way of thinking and communication. Mendel is using communication terminology for it's security related events by implementing the Mitre ATT@CK for ICS framework for OT-operators and at the same time Mitre ATT@CK for Enterprise for IT-operators. We believe this is a common set of terminology for both teams to optimize their communication and lower risks due to miscomunication.

Processing

Like in the manufacturing vertical, in the processing vertical, the smallest malfunction can negatively impact regular business operations. GreyCortex Mendel assists in defining security policies and provides the toolset to create rulesets that detect policy violations. With the adoptation of most popular firewall vendors, into it's plugin API, GreyCortex Mendel can not just monitor policy violations, but also enforce them by instructing firewalls to be closed for certain traffic.

- **Asset inventory** Mendel supports the Processing operator by knowing what assets are deployed in what purdue-level in the organisation. This helps the operator to understand his environment and determining priorities to succesfully make a security and patch-management assesment.
- **Asset details** Mendel enables the operator in knowing the risk of it's assets by reporting know security vulnerabilitie **CVEs & CPEs** and misconfiguration of those assets.
- **Intrusion prevention** Mendel enables the operator to close firewalls for high risk traffic, lowering the overall risk-level of the network.
- **Technical Network Operatios** Mendel identifies which assets on the technical op

Protocol parsers

The PLC manufacturers that are currently supported by GREYCORTEX Mendel & the Microsensor are:

Copyright © 2024 Greycortex

4.20 NetFlow Configuration

4.20.1 Processing NetFlow and IPFIX

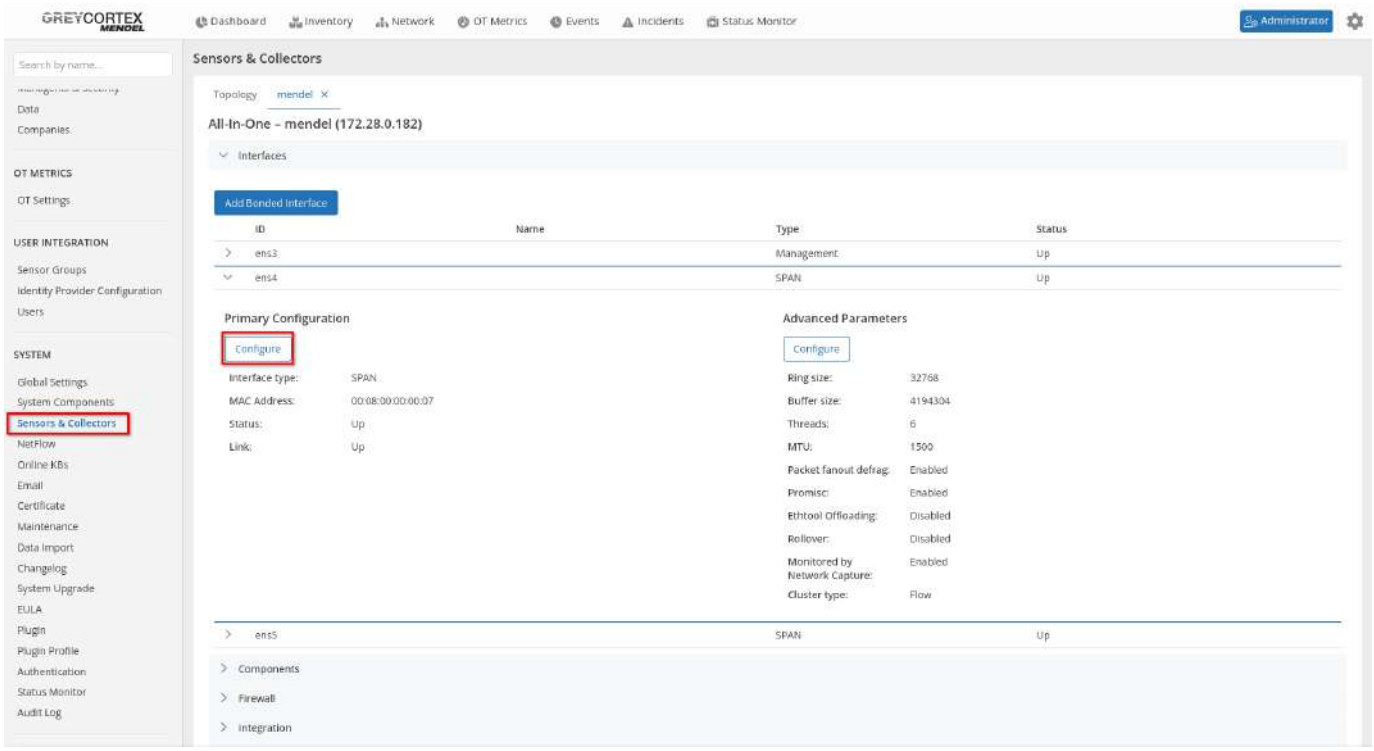
The main purpose of the GREYCORTEX Mendel is to process and analyse network packets usually received from SPAN ports. But there are also situations when getting NetFlow/IPFIX from a specific part of the network can enrich the visibility of your network. Mendel is capable of processing and analysing NetFlows records.

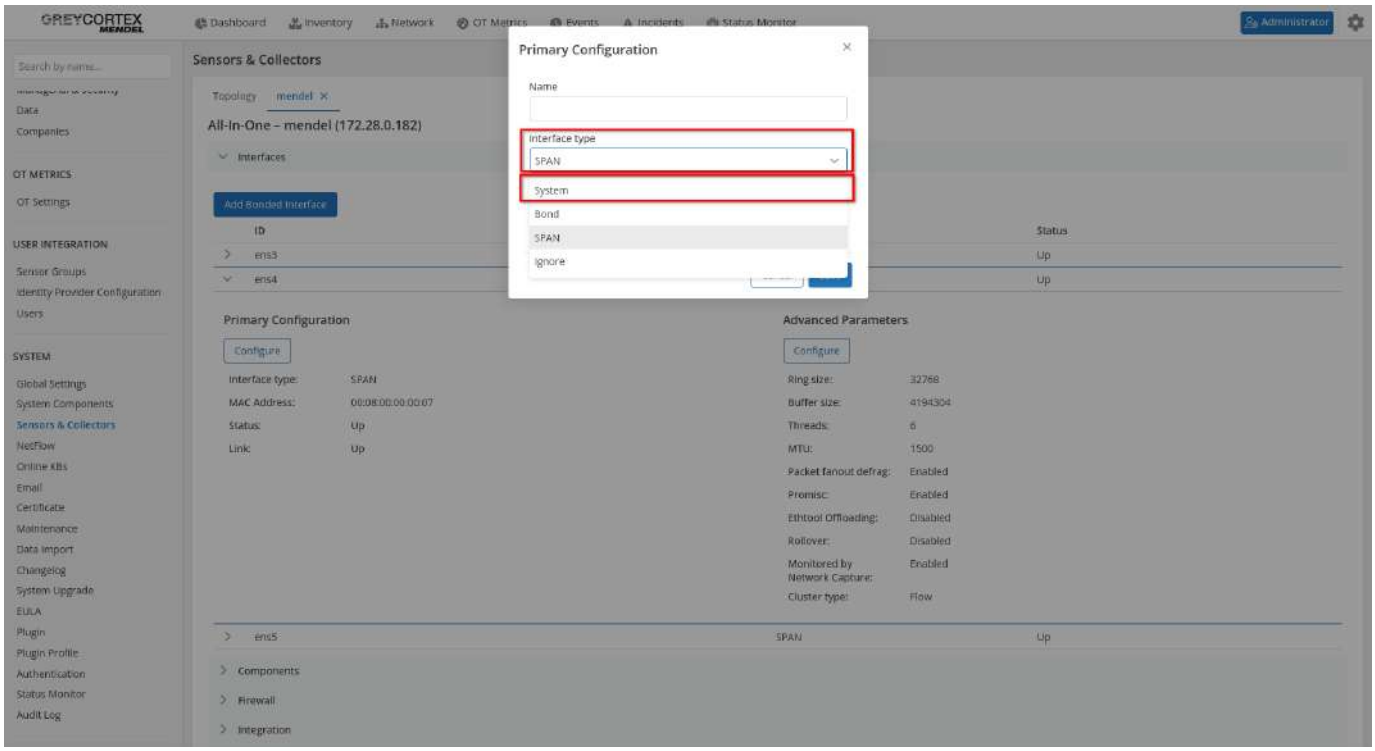
Interface settings

To receive and process NetFlow/IPFIX, you need an interface with an IP address. That can be the Management interface or System interface. We recommend using the System Interface.

To create a System interface, open configuration of one of your SPAN interfaces, click the Configuration button on the left side, and select the System interface type. This type of interface requires an IP address, which has to be configured manually.

Copyright © 2024 Greycortex



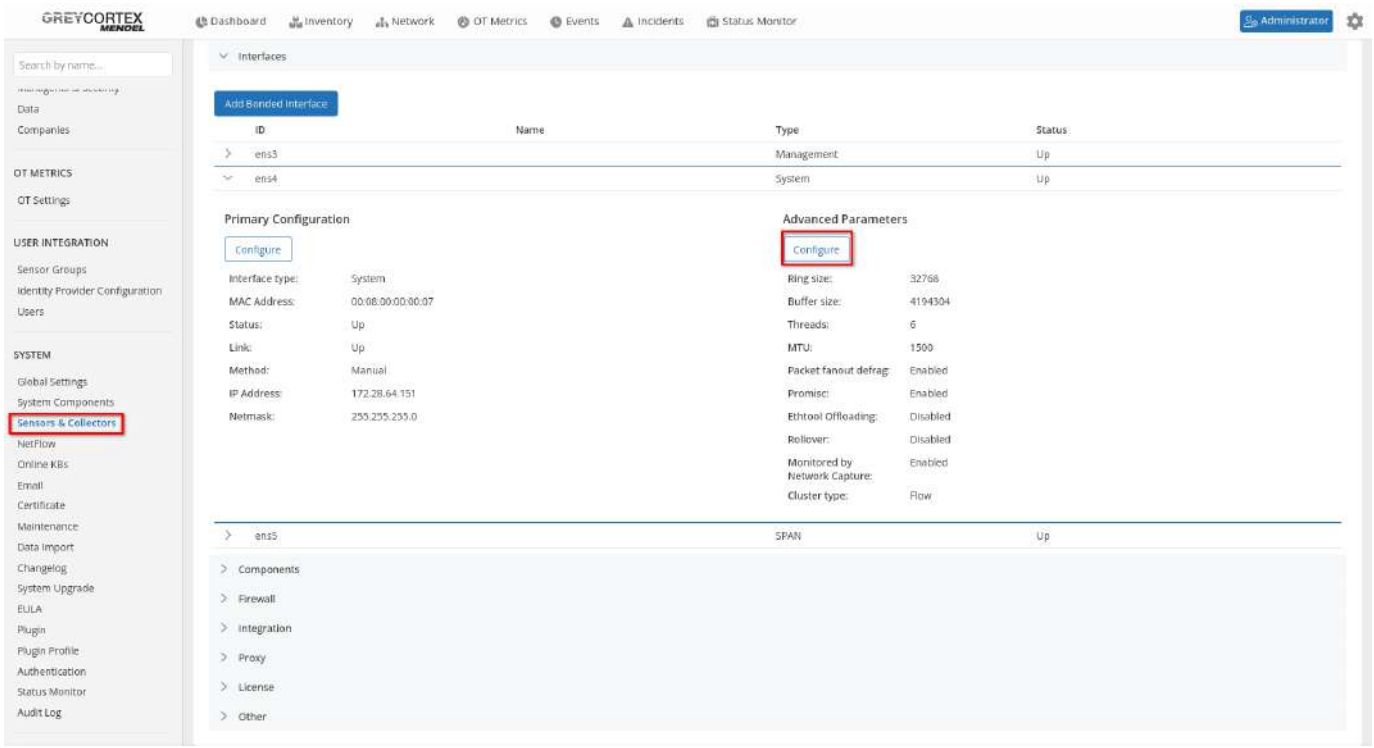


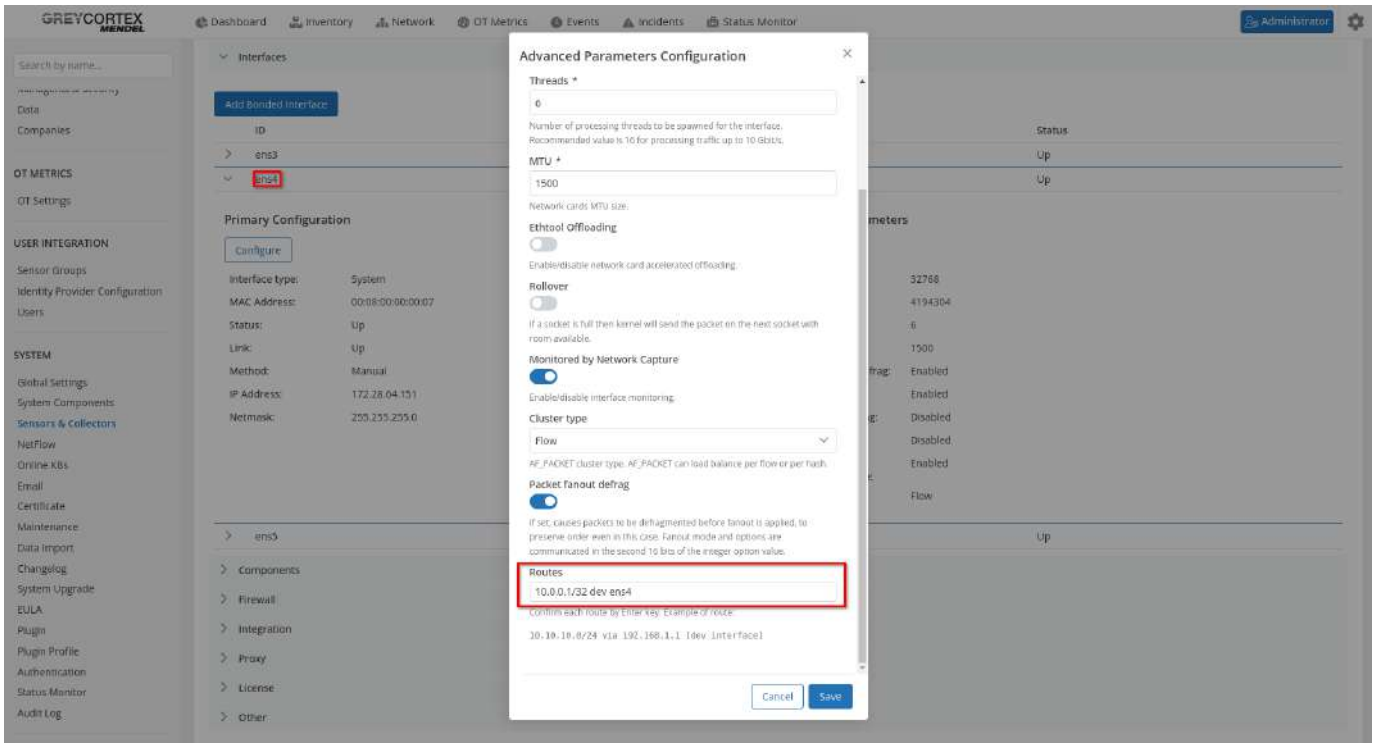
To work correctly, this interface needs a route to the NetFlow producer. New routes can be added via the Advanced Parameters. Click on the Configure button on the right side and scroll down. You should see a text field, which you can use to add routes. The route for NetFlow can be in the following format:

```
```

ProducerIP is the IP address of the NetFlow producer, interface_name is the name of the interface that will be receiving the NetFlow traffic, for example .

Copyright © 2024 Greycortex



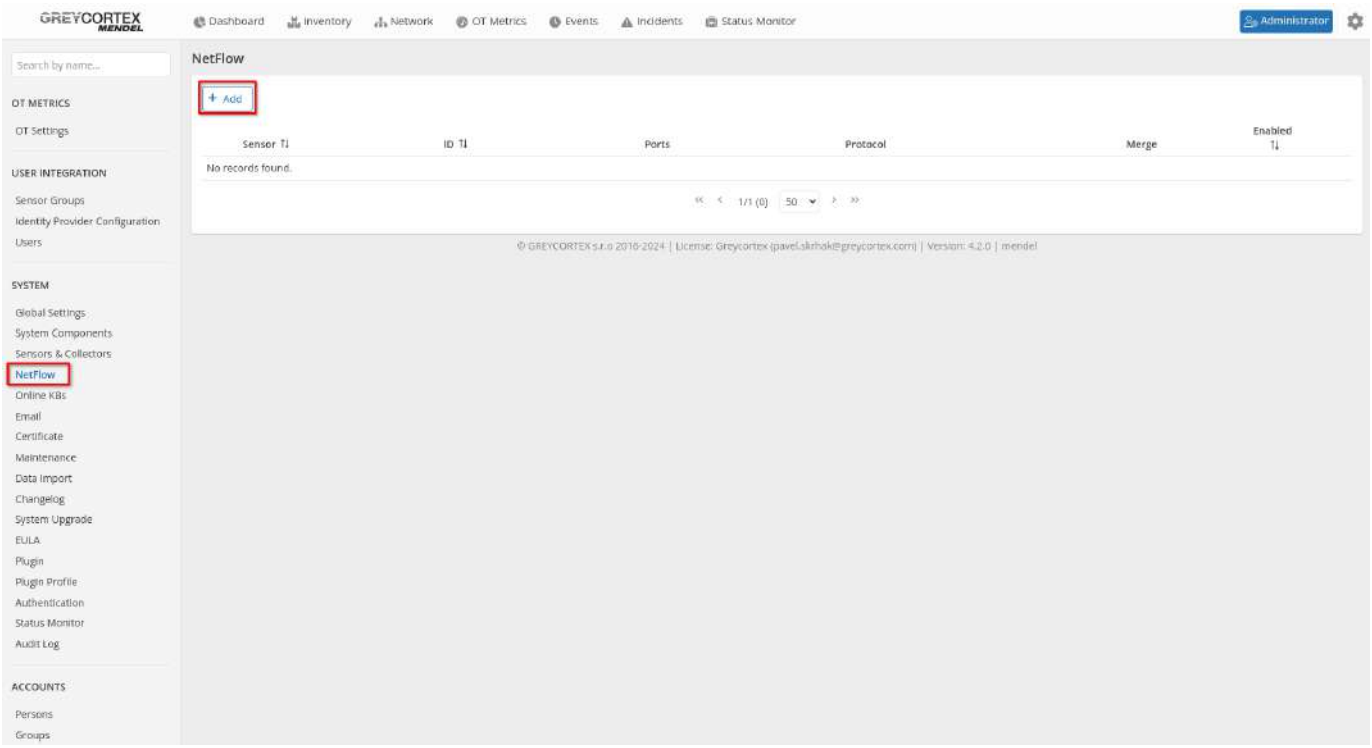


NetFlow settings

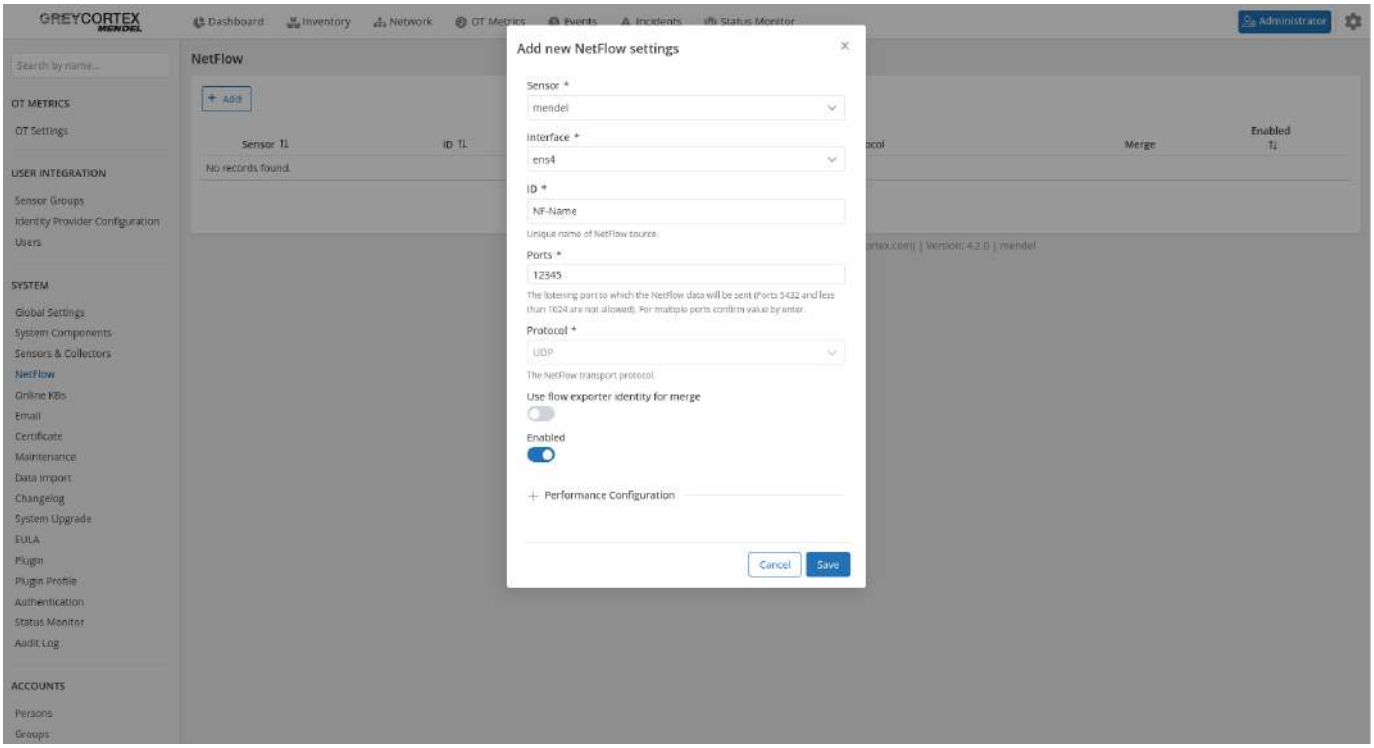
Second step is to create a NetFlow consumer on the Mendel side. Such a setting can be found in System>NetFlow and is used to configure the NetFlow component.

To add a new NetFlow source, click the Add button. This will bring up a configuration window where you can configure all the necessary options.

Copyright © 2024 Greycortex



The Sensor field allows you to select which sensor should receive the NetFlow. If you have only one machine, you can leave it on the default. Interface selects the interface that will be receiving the NetFlow. This should be either the Management interface, or System interface. If you're not sure which is the correct interface, you can check Sensors&Collectors to see the interface names. ID will be used to identify the NetFlow data in Mendel flows. You should use some descriptive name, as you will be using this to find and filter the NetFlow traffic. Ports tell Mendel which network port the NetFlow will come to.



There are some necessities that NetFlow must have. Mendel cannot start processing NetFlow until it gets a template frame that tells it which field means what. A NetFlow record must also contain five core fields. Those are source IP address, destination IP address, source port, destination port and IP protocol.

Copyright © 2024 Greycortex

Troubleshooting

- If you’re receiving the traffic on Management interface, make sure the **Monitored by Network Capture** option is enabled.
- Check in Flows that you're receiving traffic on the correct **port** on the selected interface.
- If you’re receiving the NetFlow traffic on the System interface, make sure the **route** to the **producer** is set up correctly.
- Check that the template is containing the source and **destination IP address, source** and **destination Port** and **IP protocol**.

Copyright © 2024 Greycortex

4.21 Napatech Configuration

4.21.1 Setting Up Napatech Network Interface Card in Mendel UI

This documentation provides a comprehensive guide for setting up the Napatech Network Interface Card (NIC) within the Mendel cyber security software. Follow the steps below to ensure proper configuration and integration.

4.21.2 Accessing Napatech Configuration

Navigate to Napatech Configuration:

- Open the Mendel cyber security software.
- Go to .
- Select .
- Click on next to the component.
- Click on the button.

4.21.3 Download Current Ntservice Configuration

To estimate what needs to be changed, you can download current settings.

- Click on the button.
- This action will download the current configuration file.
- Review this file to understand the current settings for the Napatech running on the appliance.

4.21.4 Modifying Napatech Parameters

Network Capture Configuration

Access Network Capture Configuration:

- In the section, locate the section.
- This section contains parameters relevant to the Network Capture component.

Modify Parameters:

- Adjust the parameters as necessary.
- The section is prefilled with safe default values, but adjustments may be needed based on hardware specifications.
- For proper tuning, contact Greycortex support.

Ntservice Configuration

Access Ntservice Configuration:

- Locate the section within the area.
- This section contains parameters relevant to the .

Modify Parameters:

- Adjust the parameters if necessary.
- This section is also prefilled with safe default values.
- Contact Greycortex support for guidance on proper tuning based on your hardware.

Copyright © 2024 Greycortex

4.21.5 Resetting Napatech Configuration

Restore Default Settings:

- If you need to revert to the default settings, locate the switch.
- Flip the switch to restore the settings to their default values.

4.21.6 Applying Changes

Restart Network Capture Component:

- After making any changes to the settings, the Network Capture component needs to be restarted.
- Ensure all modifications are saved.
- Restart the Network Capture component to apply the changes.

Notes

- **Hardware Dependency:** The values in both the Network Capture Configuration and Ntservice Configuration sections are hardware-dependent. Proper tuning is essential for optimal performance.
- **Support:** For any tuning or configuration support, please contact Greycortex support.
- **Default Values:** Both configuration sections come prefilled with safe default values. Restoring to default settings can be done using the reset switch if needed.

Copyright © 2024 Greycortex

5. How to

5.1 End User / Entry level Training

5.1.1 Introduction

1. Introduction

Computer systems are all around us in today's world. These systems contain large amounts of information that they share with other computer systems that are connected to each other through computer networks. On each device runs software that helps us in our daily work, provide entertainment, accessibility and many other functions that we take for granted today. The software on these devices talks to software on other devices through messages. The messages that these devices exchange are part of the network traffic. By monitoring and storing this network traffic, it is possible to gain insight into what data is being exchanged between devices, to whom they are transmitting that data, and also when.

Monitoring device behaviour and data flow is essential today to ensure security of the devices themselves, as well as the data they store. There are many tools that help address the issue of network traffic monitoring. The Mendel system developed by GREYCORTEX is one of them.

This training material is intended for end users who have no practical experience with Mendel. This material firstly introduces the user to how the Mendel system fits into monitoring today's TCP/IP networks and how it can help him with visibility into his network. The user will learn how the system works with network data, how it stores them, and how it extracts information from this data that is useful for analysis and visualization of the monitored computer network. The user will also learn about the detection capabilities of the system, which can assist in debugging network problems, detecting network anomalies or finding malicious communications. Everything is presented on network data that have been created for this training material. The data are stored in pcap files and are part of the training material. The user will also learn how to build a test environment in which they can practically test their newly acquired knowledge through exercises.

The training material was created by the CyberOps team, based on practical experience with the daily use of the monitoring system, computer networks and dealing with operational and security issues. If you are interested in more advanced user training directly from team members CyberOps, please contact our customer support team at support@greycortex.com.

Copyright © 2024 Greycortex

5.1.2 Monitoring network traffic

2. Monitoring network traffic

2.1 TCP/IP NETWORKS

In today's TCP/IP based networks, there are many types of of devices that communicate using different protocols and are interconnected by different technologies. Today, these devices can include mobile and desktop computers that people use for everyday work. This can also include other supporting devices that help them in their work, such as printers, telephones, televisions and others. In addition, the network may include servers that run applications or store large amounts of data. All of these end devices are connected by various wired and wireless technologies to the network elements, that allow these devices to communicate with each other.

Figure 1 - Example of infrastructure and the elements it contains.

In Figure 1 we can see a diagram of an example infrastructure with network elements. Terminal devices are connected to network devices, called access (or edge) switches. These switches can be further connected to so-called distribution switches and these are usually connected to the core switch. The core switch tends to be the heart of the entire network, as it provides many of the functions for the entire network and also sees most of the important communications on the network.

The individual devices on the network communicate using messages that are referred to as *packets*. A packet can be thought of as a sequence of bytes. These bytes are given the protocols that are used during the communication. By analysing packets, it is possible to to learn a lot of information not only about the network and its architecture, but also about the devices that are connected to the network and how the network devices behave on the network, who they communicate with, and what messages they receive during that communication. A good place to get packets from the network for communication monitoring purposes is the core switch. Every packet sent from end device to/from the Internet goes trough this switch. This switch also usually performs routing between VLANs.

There are multiple ways to get the packets that pass through the core switch. The most commonly used technology is referred to as , i.e. Switched Port ANalyzer. Using this technology it is possible to mirror incoming, outgoing or incoming and outgoing packets from a particular port of the switch to another port. This will result in forwarding of each packet from one port to the so-called SPANport. The SPAN port then sends them to a network packet analysis device, which can be connected to the SPAN port. On some devices it is possible to specify which packets should be forwarded to the SPAN port, e.g. only a specific VLAN ID can be forwarded. With this technology, it is possible to monitor all packets that pass through the core switch and gain a passive overview of what is happening in the network from a network perspective.

2.2 PACKET CAPTURE FILE TYPE

If we were to start saving the data on the other side of the SPAN port to a file, we would create a *packet capture* file type that contains individual packets. Files of this type usually have the extension , and are generally referred to as pcaps. Those files will be referred to as such in this training material. To read such a file we can use the program Wireshark, which is an excellent tool for analyzing files containing network packets. Let us now look at a pcap file download-vlc-merlin-http.pcap. When a skilled analyst quickly goes through the packets, he will see DNS queries at the beginning of the communication and then communication with the web server on port 80/tcp. Wireshark offers many views and features to quickly guide the analyst through a given communication. For example, we can use the *Endpoints* view from the *Statistics* menu, which tells us how many IPv4 addresses are communicating in a given pcap. In Figure 2 we can see that there are only three IP addresses communicating in the pcap, only one of which is in the private range, 192.168.1.2.

Figure 2 - List of IP addresses that communicated in the captured communication.

Next, Wireshark can tell us how the end-to-end devices were communicating with each other. We can see the view for each protocol in the *Conversations* view from the Statistics menu. From the IPv4 view, we can see in Figure 3 that the private IP address 192.168.1.2 communicated with two public IP addresses, 83.240.0.135 and 147.229.176.19.

Figure 3 - Viewing IPv4 protocol level conversations.

From a UDP transport protocol perspective, we can see in Figure 4 that devices from private network communicated only with the public IP address 83.240.0.135.

Figure 4 - Private IP address communication from a UDP transport protocol perspective.

We can also notice that in both cases Wireshark is capable of telling exactly who was communicating with whom, on what transport protocol, how much data was transferred during the communication, and also who initiated the communication. In Figure 5 we can see that Wireshark is also able to view and save files

Copyright © 2024 Greycortex

that have been transferred on some protocols. In the pcap file we analyzed, the file transferred was executable file named , whose size was 43 MB and was downloaded from a server named .

Figure 5 - List of files that were intercepted in the network communication.

The tool offers many other statistics that can significantly facilitate the work of an analyst examining network traffic. Wireshark is virtually unrivalled for analysis of small pcap files today. But what if we need to analyze large amounts of data? Or we need to analyze data on the fly? Or we want to do more advanced analysis with the data? That's where the Mendel system comes in, which is developed by the company GREYCORTEX.

2.3 MENDEL

You can think of the system as Wireshark on steroids. Mendel can not only analyze network traffic, it can also store metadata about monitored traffic, which can be used for retrospective analysis. Data about traffic data is stored in its internal database so we can always go back to data from the previous week, for example. Further analysis is performed over the stored data, which makes it possible to extract a lot of useful information from the network data. In general, the Mendel system specialises in three areas , which are:

1. visibility into the computer network,
2. network anomaly and threat detection, and
3. proactive response to detected events.

2.3.1 Network visibility

By monitoring network traffic, the Mendel system is able to provide a very detailed level of visibility into the infrastructure that it monitors. The system is able to classify individual devices into groups, which can then be worked with not only by the user but also by the detection mechanisms of the system itself. The system is also capable of recording, analysing and displaying communication vectors across the entire network without defining profiles or additional configuration. As can be seen in Figure 6, the system user has a detailed view of who is communicating with whom, how, and what is happening during the communication. This functionality is not only useful for debugging and finding network problems, but also for finding network threats.

Figure 6 - The Mendel system provides detailed insight into a computer network.

2.3.2 Problem and threat detection

There can be a lot of communication hidden in network traffic today that can be undesirable for a given network from both security and operational perspective. The Mendel system has several detection modules that aim to look for these operational and security issues. Thanks to this detection, the system is able to detect operational problems before they become critical. It can detect malware installed on equipment that is involved in the monitored infrastructure or targeted activity of an attacker who has access to the network and is preparing to attack or is already in the final stages of their attack. In Figure 7 we can see many types of of network activity that may be undesirable for elements on the network. These activities include:

- Tunneling traffic through other protocols.
- Violating security policies.
- Exploratory activities within the network.
- Transmission of unwanted software, called *malware.*
- Unwanted communications with unauthorized servers.
- Leakage of sensitive data and many others.

Figure 7 - The Mendel system detects operational and security anomalies in network traffic.

2.3.3 Threat Response

Network insight and detection mechanisms alone can tell you a lot of information about a network. However, in the event of detection of operational or security problems, you usually need to take action. Action may consist of adjusting the configuration of one of the monitored systems, creating a new ACL rule, adding a firewall rule, running an antivirus program, and many others. These actions can be performed by the network administrator himself, in which case however, there is a delay between the detection of the threat itself and the action to eliminate the problem or reduce its risk.

The Mendel system can be interfaced with systems that can execute actions on the network to mitigate risk or block malicious activity. With this integration, the user of the Mendel system gains the ability to take any of the following actions directly when a problem is detected, or configure the system to to perform these

Copyright © 2024 Greycortex

prevention actions completely automatically. In Figure 8 we can see that Mendel can be linked to firewalls to create rules that block incoming or outgoing traffic. The Mendel system can also integrate with products that are able to perform isolation of individual devices on the network.

Figure 8 - The Mendel system can communicate with other network and security features.

This training material will guide you through the different parts of the Mendel system. It will teach you how to use it, explain each section in detail and how to work effectively with the system.

Copyright © 2024 Greycortex

5.1.3 Test environment

3. Test environment

The most effective way to learn how to use the Mendel system is going through the use cases by hand. For each example, the test data was created in the form of pcap files, similar to what we seen in Section 2.2. Each test pcap will be used to demonstrate specific functionality of the Mendel system or to demonstrate some network problem. This data will be played back into the Mendel system and we will examine and learn the system's response. For these purposes, it will be necessary to have the Mendel system installed and also a second system from which we will to play the data into the Mendel system. We will name this system REPLAY. For purposes of this training material, the test environment chosen for this software VMware ESXi. The same goal can be achieved in other ways.

3.1 VMWARE ESXI

3.1.1 Target state description

Our test environment will consist of two virtual switches and two virtual machines:

1. VM REPLAY -- packet replayer.
2. VM Mendel -- the Mendel system.
3. Virtual switch HUB -- packet distribution from VM REPLAY to VM Mendel.
4. Virtual switch vSwitch0 -- management switch.

The first system named REPLAY can be any Linux distribution, on which a tcpreplay tool can be installed. This program can send packets from a file to a selected network interface. The program is distributed as a standard package on most of the Linux distributions in use today. The REPLAY system will send a file over the network interface packets from the selected pcap file, to a virtual switch named *HUB*. The virtual switch will be set up in such a way that all packets received on one port will be forwarded to all other ports, i.e. it will behave like hub.

TODO: It wants a logical diagram of that target state.

These packets will be received by a second system, which will be the Mendel system. The system will process the received packets and this will achieve playback of the packets from the pcap file to the Mendel system. The second virtual switch named vSwitch0 will serve as the management switch. The network interfaces that are connected to this switch will have an IP address, through which we will manage the devices. In the following section, we will look at the creation of individual switches, port groups, and virtual machines for the systems, that we're going to install.

3.1.2 Creating virtual switches and port groups

First, we will create and set up a standard virtual switch (vSwitch0) that we will use to manage our devices. If you have been running VMware ESXi for a long time for other purposes, there is high chance that this switch already exists. It is important that the virtual machines we create in the following sections have access to the Internet and are able to connect to both systems on ports 22/tcp and 443/tcp. If you are working with a freshly installed VMware ESXi instance, then you can create this switch as shown in Figure 9. In the figure, we can see that vSwitch0 has the vmnic0 uplink interface connected. Using this uplink interface and a port group named *Management Network,* we will be able to connect to this switch over the network.

Figure 9 - Creating a virtual switch for device management.

In the next step, we will create a new port group for this switch, to which we will later connect the virtual machines created below. As we can see in Figure 10, we have created this special port group named GREYCORTEX. We did not set any special properties for this group , and for parameters where it was possible, we left the option to inherit the settings from the vSwitch0 virtual switch. This group we will later assign the management network interfaces of the REPLAY and Mendel, which will have an IP address.

Figure 10 - Creating a new port group named GREYCORTEX.

Now we will create a second virtual switch (HUB) that will be used for sending packets from our REPLAY device to the Mendel system. In Figure 11 we can see the setup of this switch. In particular, notice the security policy settings, where we need to set a promiscuous mode, disable the use of the MAC Address table, and allow MAC address changes. If these options are not set to Yes, then the switch will behave like a traditional L2 switch and will drop most packets that we will send from the pcap file to its ports, because the ports on this switch will not have the devices or MAC addresses from the packets being replayed.

Figure 11 - Creating and setting up the HUB switch.

Copyright © 2024 Greycortex

As a final step, we will create a new port group for this switch, which we will name *Mirroring.* In Figure 12, we can see that for this group we set the VLAN ID value to 4095. Using this value, it will be possible to receive and send packets with any VLAN ID in this group. We connect the group to our newly created virtual switch named HUB. Next, note that we inherit the security policy from the virtual switch.

Figure 12 - Creating a new port group called Mirroring on the switch HUB.

In Figure 13 we can see the state of the virtual HUB switch after the creation and addition of a new port group called Mirroring. In the following sections, we will connect the network interfaces of the REPLAY and Mendel system to this group, trough which we will send and recieve selected network traffic.

Figure 13 - State of the virtual HUB switch (packet replay) after creation new port group.

Now we have all the virtual switches and port groups created and configured, and we can move on to installing the virtual machines, which we will then connect to the newly created port groups.

3.1.3 REPLAY System Installation

First, we will create the REPLAY virtual machine. This system will be used to play pcap files to the Mendel system. For this system, we will only need the minimum system requirements, as its only purpose will be to play pcap files using tcpreplay. The final configuration of this virtual machine prior to its actual creation can be seen in Figure 14.

Figure 14 - Final configuration of the REPLAY machine, before it is created.

The Ubuntu distribution chosen for the installation was version 18.04, but other distribution of your choice can be used. The installation of your chosen Linux distribution itself will be performed by you, since it is beyond the scope of this training material. After the installation is completed, you will need to install and set up an SSH server for remote access and also the tcpreplay tool itself for pcap playback. In the case of our distribution, we can use the command:

To verify that the SSH server is correctly installed and configured, move the test pcap file test-icmp-ping.pcap to the newly installed system REPLAY using scp:

In the command, replace the user and the IP address according to the parameters selected during the REPLAY installation. Copied pcap contains approximately 3 minutes of network ICMP communication, which we will use in Section 3.1.5 to verify that the systems in our test environment are correctly installed and connected. Using the following command we can verify that the tcpreplay tool is installed correctly.

On Ubuntu, we also recommend disabling IPv6 connectivity due to unnecessary multicast noise and other protocols. Disabling IPv6 connectivity on every boot can be done using the GRUB system boot loader, where in /etc/default/grub we add the value ipv6.disable=1 to GRUB_CMDLINE_LINUX_DEFAULT and GRUB_CMDLINE_LINUX variables. You can edit the /etc/default/grub file with the editor of your choice, such as vim or nano.

In Figure 15 we can see the contents of these variables for the installed system.

Figure 15 - Setting GRUB system variables to disable IPv6.

In order to apply this change, we also need to generate new settings for GRUB bootloader using the command:

Now we can use the ifconfig utility to view the list of network interfaces. As we can see in Figure 16, our installed distribution Ubuntu has two network interfaces. The interface named ens160 was given the IP address from the DHCP server and will serve as the management interface. The other interface called ens192 will be used to play packets from pcap files via the virtual switch HUB to the Mendel system.

Figure 16 - Network interfaces on an installed REPLAY system.

For pcap playback to be possible, the ens192 interface must be in the *up* state. Since this interface does not have an IP address, our Ubuntu instance will not use it after booting the system and sets it to the down state. So that we don't have to set the interface to the up state every time the system boots, we make a configuration change in the Ubuntu system. The configuration is done in the file /etc/netplan/01-netcfg.yaml, where we add the missing ens192 interface and set

Copyright © 2024 Greycortex

the dhcp option to No. The resulting configuration content can be seen in Figure 17. Adding this configuration will ensure that the interface is in the up state when the system boots.

Figure 17 - Contents of the 01-netcfg.yaml file after adding the ens192 interface.

Now we can reboot the REPLAY system and after the reboot we can verify that the network interface for pcap playback is in the up state and does not have IPv6 or IPv4 address. The verification can be done using the ifconfig or ip programs. The REPLAY system is now ready and we can move on to installing the Mendel system.

3.1.4 Installing MENDEL

3.1.4.1 Creating a virtual machine

The installation of Mendel is very similar to a normal Linux distribution installation. First, we prepare a virtual machine in the VMware ESXi environment and then install the system from an ISO file. To obtain the ISO file and license for the installation, please contact your partner or distributor.

First, we will create the virtual machine again according to the following parameters:

- Guest OS family: Linux
- Guest OS version: CentOS 7 (64-bit)
- CPU: 8 vCPUs
- RAM: 32 GB
- Hard disk 1: 100 GB
- Network adapter 1: VM Network
- Network adapter 2: SPAN Port group

The final configuration of all parameters can be seen in Figure 18. Version of the Mendel system that we will be using in this training material is 3.9.0 released in June 2022.

Figure 18 - Final virtual machine configuration for the Mendel system.

3.1.4.2 Part One: Text Installer

After starting the virtual machine for the first time, we get to the Mendel system installer. In Figure 19 we can see that in addition to the actual installation we can also test the embedded installation media or start booting system from the local disk.

Figure 19 - Initial launch of the Mendel installer.

In the next step, the installer shows us the currently set values to install. In Figure 20 we can see that the system has correctly detected the presence of one disk and two network cards. For the installation itself there are still a few pieces of information to fill in. So we proceed by selecting Next. In the next step in Figure 21 we can see that the system shuts down support for multi-disk installation. Installation on multiple disk system can be used to increase the performance, but will not be used for purposes of this training material.

Figure 20 - The Mendel installer shows the current settings.

Figure 21 - With a single disk, it is only possible to perform a single disk configuration.

For common installations of the smallest sensors in a virtual environment, it is recommended to use a disk space of at least 400 GB. For the purposes of this training material, we will make do with the set disk space of 100 GB. In order for the system to work with this space efficiently, we must select the minimum layout of disk volumes in the expert settings. In the disk selection step, as we see in Figure 22, select Expert, then, as you can see in Figure 23, select option 2, i.e. *GPT developer (minimal) layout*. After saving the minimum layout, the installer returns us to the disk selection step, where we can continue by selecting OK.

Figure 22 - Selecting the disk to install the system.

Figure 23 - Selecting the minimum disk volume layout.

In the next step, the installer will show us the available network interfaces, their MAC addresses, the amount of data currently flowing on the interface, and whether the interface is connected to any active element with which it was connecting to at the link layer level. In Figure 24 we can see that interface named

Copyright © 2024 Greycortex

ens192 was selected as the management port, , and has a MAC address 00:0c:29:a3:76:66. This interface is connected to a port group called GREYCORTEX, which is defined on the virtual switch vSwitch0. For this interface, we select the DHCP option and request that the system attempts to receive IP address from the DHCP server.

Figure 24 - Selecting a network interface for management.

In the next step, the installer will tell us if it was successful in obtaining network configuration. In Figure 25 we can see that in our case the installer received the IP address 172.29.10.102/24. We will later access the system's web interface or SSH to the *mshell* text interface of the Mendel system. Mshell is shell, which is automatically started when connecting to port 22/tcp.

Figure 25 - The installer received the network configuration from the DHCP server.

During installation, the system needs to verify the license key against the GREYCORTEX's server license.greycortex.com. In order for this verification to occur, the installer needs to be able to access the aforementioned server on port 443/tcp. In some infrastructures, it is necessary to use proxy server for the connection to the Internet. Therefore, the installer offers the option to configure proxy server in the next step. In Figure 26 we can see, that the option to specify a proxy is optional. If the proxy server is not necessary for accessing the Internet, you can leave the values for the proxy settings blank and continue on.

Figure 26 - Option to specify a proxy server.

In the next step, the user is prompted to enter a password for an account named administrator. This account has the maximum permissions on the installed system and should only be used for administrative tasks. You can imagine yourself as the root user on Unix systems or as the Administrator in the Windows world. This account can then be used to create other local accounts. In Figure 27 we can see that the installer prompts the user for an Admin password, but remember that the account name is .

Figure 27 - Creating a password for the local user administrator.

In the next step, we need to select the so-called *sensor name*. This name will be used as the hostname for the device. It will also be used as the CN during the certificate generation for the web server and it will be also used within the Mendel system. For example, if we have multiple sensors connected to the Mendel system, then this name is the identifier for each sensor. After selecting and entering the name during installation, the name cannot be can be changed in the future. Therefore, choose a name that clearly describes the device. Since we will be running only one system in our test environment, we choose the name as you can see in Figure 28.

Figure 28 - Choosing a system name to further identify it.

As we can see in Figure 29, the system will request the license email and the associated license key in the next step. The license key starts with the characters GM, where the first character represents the name of the company, i.e. GREYCORTEX, and the second the product name, i.e. Mendel. The license key always consists of capital letters and numbers. In the figure we can see that the license email used was and license key, that was partialy deleted for security purposes. Verification of the specified license is then performed against the server license.greycortex.com. There is also the option of *offline installation* for cases where the installation has to take place in an environment that is completely isolated from Internet access. This type of installation is described in in the installation manual and will not be discussed further in this text.

Image 29 - Inserting the license key and email.

In the next step, the installer will show the user some information that the user has entered and asks the user to check it. If all the information is correct, we can use the option to start the installation of the system itself. In Figure 30 we can see a summary of of the installer in our test environment.

Figure 30 - Summary of information before installation.

In the last step, the installer prompts us for one last confirmation that the disk, that we have selected in the installer can indeed be formatted and used for writing Mendel data. In Figure 31 we can see that in our test environment, data has been written to the disk marked . With this last confirmation we tell the system that the installation itself can begin. So the installer has collected all the necessary information from the user and can begin verifying the license, writing data to disk and configuring the system. Two system reboots will occur during the installation.

Image 31 - Last confirmation before writing data to disk.

Once the Mendel installation has been successfully completed, the system will show the system banner in text text console, which shows basic information about the hardware used, information about the IP address of the web interface and the option login to the Mendel web interface. In Figure 32 we can we can see that the test system was installed in a VMware environment. We will use the option to log in to the administrator account to verify the chosen password, that was entered during installation.

Figure 32 Mendel system banner after successful installation.

Copyright © 2024 Greycortex

After a successful login, the Mendel shell (in short) is launched. Mendel shell is a special Mendel system shell. Mshell can be used for system administration when the web interface is not available or during initial installation. Mshell is also run when you log in to the system with an administrator account using SSH, which we will test once we have enabled SSH access on the system. Other users are not allowed to use mshell unless explicitly set by the system administrator. As we can see in Figure 33, after logging into mshell the command is ran right away, which displays help for all commands that mshell has. This completes the first part of the installation and we can now proceed to the second part, where we perform the initial setup of the system after installation. So now we log into the web interface of the Mendel system.

Figure 33 - The Mendel shell is launched when the user text logs in to system.

3.1.4.3 Part Two: Web Interface Configuration

After entering the IP address into the web browser, the browser will probably display some warning that you are connecting to an insecure system, that your data may be stolen, etc. In our case, the web server sends the browser a certificate that contains the Common Name Mendel and expects that this value will also be used to access the web server. If we use an IP address, then the web browser cannot verify that the system we're connecting to is really the one it's pretending to be. We will ignore this risk for now and discuss it further in other sections. As we can see in Figure 34, after skipping the security browser message, we are taken to the login form of Mendel system web interface. To log in, we use the previously mentioned administrator username and password that was chosen during the installation.

Figure 34 - Connecting to the Mendel web interface.

Once the installation is complete, there are six configuration steps to be performed before the Mendel system is fully operational. First, we need to agree to the end-user license agreement for using the system, which can be seen on Figure 35.

Figure 35 - Confirmation of the Mendel license terms.

In the next step we have to select the language. The Mendel system is currently localized into Czech, English, Polish and Japanese. List of supported languages is still being expanded and we are open to cooperation with our partners to localize into their native language. In Figure 36 we can we can see that we choose English for our instance.

Figure 36 - Choosing the English language for the Mendel system.

In the second step, we can insert a backup of the settings from the previous installation. This option is useful when, for example, you are running PoC (Proof of Concept), after which a backup of the settings is made. If the customer decides to purchase the system and a new installation occurs, the the settings can be restored in this step. It is important to note here that the backup should be restored to the same version of the system from which it was created. In Figure 37 we can see that in our test installation we will not be restoring data from a backup.

Figure 37 - It is possible to insert and restore settings to the system from a backup.

In the third step, you can set some system behavior options. For test environment, we can leave all options at their default settings, except for remote access via SSH. In Figure 38 we can see that we enable remote access for subnet 0.0.0.0/0, which represents any source IP address. The system will create firewall rules in background, that will allow connection to port 22/tcp from any IP address.

Figure 38 - Enabling remote access using SSH.

For the system to run properly and display the results correctly, it needs to be set to an accurate time. Setting the time can be done manually or the NTP protocol can be used to synchronize the time with NTP servers. In Figure 39 we can see that in our case we will use the publicly available NTP servers, as there are no NTP servers in our test environment.

Figure 39 - Selecting public NTP servers for time synchronization.

The fifth step is to create and correctly name the subnets that we will to monitor. In our test environment and throughout this training material, we will only use private ranges A, B, and C. These ranges are already included in the default system installation. In Figure 40 we can see that in addition to these private subnets, the system also contains some other ranges from which we can by default see communication in computer networks. How does the system work with these networks and why is their definition very important will be explained further in Section 4.4.

Figure 40 - Setting up the subnets that the system will monitor and store data from them.

With this step, the system configuration is complete and we can tell the system that we want to go to the actual web application of the Mendel system. As we can see in Figure 41, the sixth step is the last step before entering the Mendel web application itself.

Figure 41 - System configuration is complete.

Copyright © 2024 Greycortex

If you have made it to the screen that can be seen in Figure 42, then you have successfully installed and configured the Mendel system for collection and analysis of network traffic.

Figure 42 - Mendel System Web Interface.

3.1.5 Playing pcap for the first time

Now that the Mendel system is fully functional, we can proceed to verify that the system will process the data that we are going to replay from the REPLAY system. First we perform an initial test using the pcap file test-icmp-ping.pcap. The file contains roughly 3 minutes of ICMP communication. We log into the Mendel system using SSH.

After logging into the Mendel system, we will get into mshell. This shell contains a useful show network command, that will show us the available network interfaces on the system, their configuration and purpose. In Figure 43 we can see that the Mendel system we have installed has two interfaces, one of them, ens192, is the management interface and has the IP address to which we have connected using SSH. The second interface, named ens224, is designated as the monitoring interface. Any packet that flows into this interface will be processed by the Mendel system.

Figure 43 - Output of the network show command in mshell.

We will now use the second useful ifstats command to display the number of incoming data into the Mendel system in real time. When we use this command, we can see in Figure 44 that the ens224 interface is currently not receiving any data.

Figure 44 - The ens224 interface is receiving 0 bits.

We leave the ifstats command running and open a second console in which we use SSH to login to the REPLAY system.

After logging in, verify that the home directory of the user you created is the test-icmp-ping.pcap file and that it has the correct checksum. Next, view the network interfaces and note the name of the interface that does not have set IP address. In Figure 45 you can see that in our case this interface is named ens192. We will use this interface to play packets from the selected pcap files.

Figure 45 - REPLAY login and pcap file verification.

Now we are ready to play packets from pcap to the Mendel system for the first time. The actual replay will be done using the tcpreplay program. Pcap is is about 3 minutes long, so watch the console during playback with Mendel system with ifstats enabled. After the playback starts the values of the incoming bits should be incremented, i.e. the
. The actual playback of the pcap in the REPLAY system is triggered by the following command.

Increase of the RX bits value after starting the pcap playback means, that the system is successfully receiving data. In Figure 46 we can see, that the Mendel system we have installed is receiving data on the ens224 monitoring port at a rate of 1.7 K bits per second. This value will change over time. If you see a non-zero value in that column, then it means that your test environment is fully ready to play another pcaps.

Figure 46 - The Mendel system receives 1.7 Kb of data per second.

When the playback is finished, we can log out of the REPLAY system and in the Mendel system we can disable ifstats using the keyboard shortcut Ctrl+C. Also we can delete the currently captured data using the command , which will delete all captured data. As we can see in Figure 47, the system prompts the user to copy a short string to confirm that he really wants to delete the data, since this command can be very dangerous. Put it in the same category as the in the Unix world.

Figure 47 - Deleting test data collected by Mendel.

3.2 TODO: VMWARE WORKSTATION

Instructions for playing pcap from host directly to guest (Mendel) on vmware Workstation. The replay system is therefore the one where Workstation is is installed.

Copyright © 2024 Greycortex

3.3 TODO: VIRTUALBOX

Instructions for replaying pcap from host directly to guest (Mendel) on Virtualbox. The replay system is therefore the one where Virtualbox is installed.
Basically this: https://confluence.greycortex.com/pages/viewpage.action?pageId=8619317

Copyright © 2024 Greycortex

5.1.4 Basic use of the MENDEL system

4. Basic use of the MENDEL system

The test environment has been successfully commissioned, now we can switch back to the Mendel web interface. We open the web browser and go to the IP address we chose during installation or received from DHCP. In our case, the Mendel system has an IP address of 172.29.10.102.

4.1 WEB INTERFACE PHILOSOPHY

You will perform most of your analysis activities in the web interface, which allows us to access the collected network traffic, allows us to ask analytical questions, change detection algorithm settings, and many other operations. The whole philosophy of the system is as follows. The system processes and stores data in its internal database. The web interface then provides several views that represent different views of these data, always over a selected time interval. The data displayed on a particular view can be further specified in detail by selecting specific filters. Thanks to this philosophy, the system can be asked many analytical queries. Examples of analytical queries include:

- Display the list of devices that communicated on VLAN 20 on Monday.
- Which devices attempted to communicate with CnC server with IP X.X.X.X since the date the phishing email was received?
- I moved the antivirus administration to a new IP address on Monday. Are there any devices that tried to communicate with the old IP address since Monday?
- What is the RTT value between specific servers and how fast is the line between them?
- And many more.

The system is also able to use its detection modules to report large number of different events that occur on individual devices or the entire network. The system also includes mechanisms that can be used to tell the system which events are relevant to the infrastructure being monitored and which are not. Using these mechanisms, the system can be tuned to the needs of a particular infrastructure. With the help of a calendar and time intervals, it is possible to easily search the recorded data to the past, select any time interval or look at new data in real time.

4.2 BASIC WEB INTERFACE ELEMENTS

In Figure 48 we can see a typical view after logging into the Mendel system, unless some manual adjustments have already been made. On the left side we can notice the calendar, which we use to tell the system what time period we are interested in. We can select specific days or exact time periods. The calendar affects all the views on which the calendar is displayed.

Figure 48 - Typical Mendel view after login.

Down below the calendar, we can see a set of filters and icons that can be used to further define which data the system should show on a particular view. If the user does not select any filter, then he tells the system to display all the data that were collected over the selected time period. For a particular user account, it is possible to set which filters the account sees and which ones it does not. This allows the user to select only the ones he uses most often, such as filters for subnets, device names, user names, port numbers, etc. You can select these filters using the attributes, which can be seen in Figure 49. In the figure we can see that the current user has selected filters that can further specify subnet, device name or IP address, device label, MAC address, service, service type, application or protocol. Not all filters are active in all views. If a filter is not active, it is grayed out and the user cannot select any value using it.

Figure 49 - Custom filters are defined using attributes.

In the bottom left part we can then see the most important button , which we use to tell the system that our filter is composed and we want to display the data. Below this button we also find the Clear option, which deletes all the values that are currently entered in the filters. Note that once we clear a filter, we can no longer use back button (left arrow). Next, in Figure 50 we can see the filter Manager, which we will become more familiar with as part of the individual exercises.

Figure 50 - Using the Filter button, we tell the system to find the specified data.

Basically, the whole system is built on the principle of views that show all the data unless we restrict or *filter* it using filters. With individual filters and their advanced use in the form of their composing, storing and using them in different places in the system. We will cover the details of the filters in this training material.

At the top of the system we can see the different view categories that system offers to the user. In Figure 51 we can see that right after the login, the system shows us a view called *Main* from the category *Dashboard.* The entire category of dashboards contains user-defined views that are set up by the manufacturer

Copyright © 2024 Greycortex

in the default installation, but can can be completely changed by the user. At the same time, each user can define their own dashboards that can show exactly what the user is interested in.

The second category is named *Network* because it offers several views that give us insight into the infrastructure that is being monitored. The third category is the *Events* view, where the user can find a list of all events or anomalies that the system has detected while monitoring your infrastructure. Different views show different data, from these views we can build an incident that will describe an operational or security problem. You can list these incidents and work with them in a category called *Incidents.* The last category is *Status Monitor*, which is a special module of the Mendel system whose purpose is to to ensure that the system operates without failure, i.e. 24 hours a day, seven days a week, 365 days a year. All views from each category will be discussed in the following sections.

Figure 51 - The different categories of views can be found at the top.

In the upper right section we can see the name of the logged in user and the entry to the system settings. In Figure 52 we can see that the currently logged in user is the Administrator. Under the gear icon we can then find several important options. The first option is to enter the actual system settings, which we will enter from time to time during the following sections. Next, we can see the settings for the currently logged in user account, where we can change the password for that account or, for example, the graphical display of the system. Another important option is the so-called *URL sharing,* which allows us to create a link to the view we currently have open. This link can then be saved in the incident or sent to a colleague to whom we want to show what we are looking at. The last important option is the link to the user documentation, which can answer many questions.

We have covered the basic elements of the web interface. Now we can move on to the first view where we explain how the system works with packets it receives on its monitoring ports.

Figure 52 - Logged in user and settings.

4.3 NETWORK FLOWS OR FLOWS

Network communication between two devices always has a beginning and an end. During the communication, one device is a client and the other device is the server. During this communication some protocols are used and messages are exchanged.

4.3.1 Pcap - download text file

An example of such a communication is the download of a text file, which is captured in the file http-get-flow.pcap. This file contains communications from a device on the internal network with a private IP address 192.168.1.11 with the merlin.fit.vutbr.cz server, which has a public IP address 147.229.176.19. During this communication, a text file named was downloaded using wget program. The full command to download the file was as follows.

If we open the communication in Wireshark, we can see that the first three packets represent a TCP handshake between the client and server. The operating system of the client device has selected a random port with number 58000/tcp for the communication. The client, when queried for a given file, specified the HTTP protocol, so communication with the server was established on the destination port number 80/tcp. Port 80/tcp is the default port for HTTP protocol. Establishing this connection is the responsibility of the operating systems of the devices in question. This initial communication is done at the TCP protocol level is referred to as a and can be seen in Figure 53.

Figure 53 - TCP handshake between client and server.

Once the TCP connection is successfully established, the application on the client device comes next, i.e. the wget program. The program sends packet number 4, which already contains the application layer in the form of an HTTP message. The content of this message can be seen in Figure 54. The HTTP application protocol is text readable protocol whose format is usually in the form of a header and body. The first HTTP message contains only the header, in which we can see that the client wants to use the GET method to download an object that is on the server merlin.fit.vutbr.cz located on the path mirrors/centos/7.9.2009/isos/x86_64/0_README.txt. We can also notice the attribute in which the used program has written its name and version.

Figure 54 - Content of the first HTTP message sent by the client.

With the fifth packet, the remote server acknowledges to the client that it has received its request and in the sixth and eighth packets sends the requested data to the client. The tenth packet contains the HTTP message that can be seen in Figure 55. The message consists of three segments, namely packets 6, 8 and 10. The server in the header responds to the client using the code which tells the client that it has the requested object on its file system and that the client can receive the file. After the HTTP message header, we can see the actual content of the requested text file 0_README.txt.

Figure 55 - In the sixth packet, the server sends the data to the client.

Copyright © 2024 Greycortex

In packet number eleven, the client tells the server that it has successfully received the data and the remote server knows that it does not need to send the data again. With the twelfth packet the client tells the server that it will not request any more data and the connection can therefore be terminated. In the last two packets, the devices only confirm that they are aware that the communication is ending and that there may be the release of resources on both sides. This reliable transmission is provided by the TCP transport protocol.

4.3.2 How network traffic is generated

We will now play the same communication log into the Mendel system and examine how the system handles this communication and what it is able to tell us about this communication. The system analyzes each packet of network communication and tries to create logical network flows from those packets. Each such flow has a source device that initiated the communication and a destination device it wanted to communicate with. The network flow then contains all the information about that communication. First, we transfer a file http-get-flow.pcap to the home directory of the REPLAY system using the scp program.

Then log on to the REPLAY system using SSH.

After logging in, we verify that the file has been transferred correctly by verifying checksum, which must be 8c219c730874148ec90d7e2c8174c769.

Now we can transfer the pcap to the Mendel system using the program tcpreplay. We need to replay the pcap under user with uid 0, i.e. user root. The interface we use for replay is an interface called ens192, which has no IP address set and is connected to a virtual switch called HUB that we created in Section 3.1.2.

If all goes well, exactly fourteen packets should be replayed. Now we can log into the Mendel system, select the correct day and time interval in the calendar, filter the client's IP address from that communication and let the system display the network flows. In Figure 56 we can see that we first log in with the username and enter the password we chose during installation of Mendel system in section 3.1.4.2. After loggin in, we select the flows view, which can be found in the Network category. When switching to this view we can notice that the system does not show us any data. This is because the network flows contain the most detailed information about network traffic. Without the other parameters, the system would have to display a large amount of data, so it asks us to specify what we're looking for, using filters.

Figure 56 - Logging in and displaying the network stream from the replayed traffic.

We replayed the pcap in our test environment on January 31, 2023 at 12:42 PM. First, we use the calendar to limit the search interval start to the date 31.1 from 12:00. Then we use the filter to tell the system, that we are only interested in network flows that involve an IP 192.168.1.11. This is the IP address in the pcap of the client device that that was downloading the data. After specifying all the filters, we can now use the button to display the data. Now we can see one line that represents the network flow that the Mendel system created from the packets that are captured in the network traffic that is being played back. The created network flow can be seen in Figure 57. We can notice that in the header of the flow, the system shows the direction of communication in which the IP address 192.168.1.11 is the source, and IP address 147.229.176.19 the destination. Thus, from the communication arrow header it is immediately apparent who initiated the communication.

Figure 57 - Network flow created from replayed packets.

4.3.2.1 Application Layer

Copyright © 2024 Greycortex

Under the header, we can find the different layers and information that the system stores at these layers. The first layer we see is the application layer, where we can see the HTTP transaction in progress. Here we can notice that the system stores important information from the HTTP message header such as which object the client is requesting from the server and what is the response of the server to this request. We can also notice that the system has enriched the transaction with additional information that is not part of the packet, such as the checksum of the downloaded file. The Mendel system normally enriches network flows with other useful information that is not part of the packets. We will take a closer look at this useful information in this training material.

4.3.2.2 Transport Layer

In the transport layer in Figure 58 we can see that during the communication the TCP protocol was used along with other information that is related to this protocol, such as port numbers or the number of flags that was transmitted. Furthermore, we can also notice that the performance metrics was calculated from the communication, for example RTT. This metric tells us how fast the TCP handshake was between the devices. From this information, we can infer how fast is the data medium that is between the devices. In our case, we can see that the handshake took place within 8 micro seconds. In your case the value may be different.

Figure 58 - Recorded information from the transport layer.

4.3.2.3 Network layer

In the network layer, which we can see in Figure 59, in addition to the IP addresses we can also find e.g. TTL (Time To Live) values. From these values, the L3/L4 Feature Vector is made, which the system uses e.g. to heuristic detection of operating systems running on individual devices.

Figure 59 - Logged information from the network layer.

4.3.2.4 Link layer

In the link layer in Figure 60 we can find the number of packets that each side of the communication has sent, the MAC addresses of the communicating devices or the interface name of the network adapter that received the packets from a given network communication. The interface name can be useful if the system receives data from more than one SPAN port.

Figure 60 - Recorded link layer information.

4.3.2.5 Network Flow Information

The first tab, called flow, contains information about the network flow itself as the Mendel system has assembled it from the individual packets that it received on its monitoring interface. From the information shown in Figure 61 we can see that the flow type has a value of , which means that the system has seen packets that signal the termination of the TCP connection and considers the flow to be completed. We can also notice that the system reports a signature with an identifier

Figure 61 Information about the network flow created by Mendel.

If the system finds a pattern defined in the network flow signature, it will log the event, which can be seen in the view *Events by Severity* in the *Events* category. If this view is viewed, we can see, as in Figure 62, that the system is reporting an event with the name Description event generation and how to work with them will be discussed in more detail in Chapter 5.

Figure 62 - A generated event describing a captured pattern in a network communication.

Copyright © 2024 Greycortex

4.3.3 Exercise - Installing a Linux Package

On the Mendel system, upload a file named package-install.pcap. The pcap contains the communication of the Linux system on which the package was installed from the official distribution server.

1. What is the IP address of the client that established communication with the server on the Internet?
2. What is the MAC address of the client device that established communication with the server on the Internet?
3. What is the IP address of the destination server?
4. On what target port did the communication take place?
5. What source port was used to download the package?
6. What application protocol was used during the communication?
7. What is the name of the package that the client downloaded from the server?
8. What is the MD5 checksum of the downloaded package?
9. What is the DNS name of the server from which the package was downloaded?
10. What transport protocol was used to download the package?
11. What is the L3/L4 Feature Vector calculated by Mendel?
12. What is the signature ID that was in the download stream package download flow?

4.4 SUBNETS AND VLANS

In today's TCP/IP networks, individual devices are most commonly addressed by IP addresses that fall within reserved private ranges that are:

- Private A (10.0.0.0/8),
- Private B (172.16.0.0/12), and
- Private C (192.168.0.0/16).

Most devices on the internal network communicate from one or more IP addresses from the above ranges. Using segmentation, it is then possible to create separate logical segments for different categories of devices. An example of such segments could be:

- 192.168.100.0/24 Servers,
- 192.168.200.0/24 Users,
- 192.168.300.0/24 Public WiFi network,
- and so on.

In many infrastructures, these ranges are then also given their own VLAN ID.

While we will see the IP address of a device or VLAN ID in packets when communicating, we can't see the name of that subnet or VLAN. Names, however, can add meaning to individual subnets and can be a great help to an analyst, who has to analyse the network being monitored. Let's take a look at how Mendel works with subnets and VLAN IDs.

4.4.1 Pcap - download text file

For this section you need to have played http-get-flow.pcap from section 4.3.1. The file is not necessary to replay if you have already done so.

4.4.2 Subnet Identification

At the point when the Mendel system receives packets about network traffic and creates a network flow from those packets, it also starts to extract other useful information. One of this information is statistics about communications on each subnet. These statistics can be found at *Subnets* view in the *Network* category. In Figure 63 you can see a total of three subnets. The replay file contains a total of 10 packets and consists of one logical flow. As you can also notice the system automatically named the subnet as . The other subnet is named . This subnet represents all user-defined networks together. The meaning of this subnet will become clearer during the following sections.

Figure 63 - A view of the statistics of each subnet.

Copyright © 2024 Greycortex

TODO: Take a new screenshot because we disabled IPv6 in section 3.1.3, so it won't be there anymore.

All networks shown are defined in the system by default after a standard installation. For example, if we click on the network 192.168.0.0/16, then the system will show us that the current name of that network is Private C and its description is These parameters can be seen on the Figure 64 and we can also change them in the settings of that network. So let's move on to the settings of that network using the *Settings* button.

Figure 64 - Network Information 192.168.0.0/16.

Change the following parameters for the 192.168.0.0/16 network:

- Users,
- Windows domain users.

Now we go back to the Subnet view in the Network category, where we can notice that the system now shows the new name of the subnet. As we can can see in Figure 65, the new name can now also be filtered in the filter with name Subnet.

Figure 65 - The new subnet name can now also be filtered.

For the Users subnet, if we now double-click on the number of flows, we will be redirected to the list of network flows that have occurred in that time period in the Users network for the given time period. Note here that with the redirection the value 192.168.0.0/16 was automatically filled into the subnet filter. Also note we can notice that the subnet was renamed in an already created flow. We can see these two facts in Figure 66. The subnet name change is thus applied to both new data and historical data already in the system.

Figure 66 - ythe system added the 192.168.0.0/16 network automatically to the filter.

4.4.3 Define your subnets

The routing tables of a monitored network can contain dozens, hundreds, and thousands of subnets or VLANs. Therefore, it is advisable to include the names of these subnets/VLANs in the system so that the system can classify individual devices into correct subnets. The individual subnets can then be handled as a whole, e.g. for filtering or creating detection rules. Also many internal mechanisms work with these subnets. The correct definition of subnets is therefore `necessary' for the full functionality of the system.

4.5 DEVICES ON THE NETWORK

For the Mendel system, each unique combination of MAC address and IPv4/IPv6 address is a single so-called *host*. The Hosts view in the Network category is able to show a list of all communicating hosts on the network for the selected time period. If we look at this view, we can see that the system has recognized from the pcap that was played http-get-flow.pcap host with IP address 192.168.1.11 and MAC address 00:0c:29:2a:8a:8a. If the IP address is clicked, the system displays information about that host, which we can see in Figure 67. In the details for that host, we can also notice the time marker that tells us when the host was discovered on the network. This marker represents the time when the system first saw the guest communicate. If a guest is going to be inactive on the network for an extended period of time (8 days or more by default), then the system will forget the guest.

Figure 67 - Detail of a guest with IP address 192.168.1.11 in the all view guests.

4.6 PORTS AND THE SERVICES THAT RUN ON THEM

If we move further to the right in the view we can notice view called *Services*. In this view the system shows the services, that are communicating on the network, again for the selected time interval and the selected filters. If no filter is selected, then the system shows all ports and the recognized protocols on them that are on the network for the selected time interval. For the file http-get-flow.pcap, as can be seen in Figure 68, you can see that the system recognized the HTTP service running on port 80/tcp. Next we can notice, that the service type has been recognized as . This type means that devices from the monitored network are connecting to the external network on port 80/tcp and are using the protocol HTTP.

Figure 68 - View of the services that have been detected on the network.

Copyright © 2024 Greycortex

4.6.1 Exercise - Subnets, Guests and Services

Upload the file subnets-hosts-services.pcap to the Mendel system. The pcap contains communication with IP address 10.168.0.198. The device communicated with two servers on the Internet and one on a private network on three different services.

1. How much data was sent from subnet 10.0.0.0/8? Use the view Subnets.
2. How much data came into the 10.0.0.0/8 network? Use the Subnets view.
3. How many hosts are created from the replayed pcap in the Hosts view?
4. How many network flows are created for the host with IP address 10.168.0.198?
5. How many communicating services were detected for the host with the IP address 10.168.0.198 (Use the Services view)?
6. What type of service was detected on port 21/tcp?
7. What type of service was detected on port 22/tcp?
8. What type of service was detected on port 80/tcp?

4.7 WHO IS COMMUNICATING WITH WHOM?

Another useful view offered by the Mendel system is called *Peers*, i.e. a graphical view of the communication partners. This view can be useful when we are interested in dependencies between networks or VLANs.

4.7.1 Pcap - SSH scan

To demonstrate this view we will use the file network-peers-ssh-scan.pcap. If we open the file in Wireshark and do a quick analysis, then we'll see that the pcap only contains communication on port 22/tcp, always to the destination IP 172.16.13.10. First, we transfer the file to the REPLAY machine using scp.

The file captures roughly twenty minutes of communication. We can adjust the playback speed using the mbps parameter of tcpreplay, so we don't have to wait so long for the pcap to play. Replaying pcap should only take a few seconds.

4.7.2 Viewing communication partners

Now we can go to the Mendel system, set the time interval to the time we played the pcap, and on the *Peers* view in the Network category, filter the target IP address 172.16.13.10. In the left part we can see a graphical representation of the communication IP:IP, i.e. the communication of each IP address with each other. In our case we are filtering out a specific IP address, and therefore we can only see communication with IP 172.16.13.10. In Figure 69 we can notice that if we click on a specific IP address, the system will display information about that IP address. Also we can notice that if we hover over IP address 172.16.13.10, the system will display the data transferred to/from that IP address. In the picture we can see that a total of 14 devices from the Internet have sent to our internal address a total of 32.27kB of data. The internal device sent 35.82kB of data back to them.

Figure 69 - Viewing IP address 172.16.13.10 in the communications view partners.

On the right side we can see the communication with the internal IP address from the individual autonomous systems. So we can see that IP addresses from the following systems have tried to communicate with the IP address:

- DIGITALOCEAN-ASN,
- Chang Way Technologies Co. Limited,
- Kanzas LLC,
- China Telecom Group, and
- Chinanet.

Thus, the right-hand side of the view shows the IP:ASN communication if filtering a specific IP address. If we were to map a specific IP address not filtered, then the right side shows the subnet/VLAN:ASN communication, which can be seen in Figure 70 if we take the IP address out of the device filter.

Figure 70 - The right side of the peers shows the subnet/vlan:asn view.

Copyright © 2024 Greycortex

If we double click on the connecting line between any two communication peers, then the system will redirect us to the network flow view between the selected peers. In Figure 71, we can notice that the moment the user has double-clicked on the link line between the IP addresses 172.16.13.10 and 193.106.191.80, he was redirected to the network flow view. At the same time, the system automatically populated the device filter, into which it inserted the value , where represents logical AND conjunction. This filter tells the system to filter only those flows that occurred between the IP addresses mentioned.

In the list of network flows, we can see that the system attempted to establish communication on destination port 22/tcp with the destination IP address 172.16.13.10. If we open this flow, we can see that sixty-two packets were transmitted in this communication. Interesting information may be the data from the application layer. In the data we can notice that the source system from Russia has used a program to make the connection that sends the information that it uses the So it's probably not classic SSH client program, but a special program that is used to to establish communication. On the target side there is SSH server implemented by the library version 6.7p1, which probably runs on a system Raspbian.

Figure 71 Double-click to switch from peers to network flows.

4.7.3 Exercise - Who is my device communicating with?

Download the pcap peers-ipv4.pcap to the Mendel system. The file contains the communications of a device from the internal network with IP address 172.28.252.85, that communicated with servers on the Internet. Answer the following questions using the view Network -> Peers:

1. How many servers on the Internet did the device from the internal network communicate with?
2. What is the hostname of the Czech server with which the device communicated?
3. What is the hostname of the American server that the device communicated with?
4. How much data [B] did the device send from the internal network to the server with IP address 147.229.176.19?
5. How much data [MB] did the server with IP address 151.101.131 send to the device on the internal network?
6. What is the name of the subnet on which the server with the IP address 147.229.176.19 is?
7. What is the ASN number of the server with the IP address 151.101.131.5?

4.8 ANALYSIS OVER STORED DATA

The Mendel system stores all network flow information in an internal database. We have already seen in Section 4.3 that we can view these flows and filter them based on the available filters in the left menu. However, the system offers more advanced analytical queries that the user can also extend or define new ones. These analytical queries can be asked via the *Analysis* view in the Network category. On this view several analytical queries are already defined, those are available in each installation. In Figure 72 we can see that one of the predefined analytic queries is a query called *Top HTTPS Cipher Suites*. This query will display the cryptographic algorithms that have been used during TLS communication from all network flows for the selected time period.

Figure 72 - Selecting a predefined Top HTTPS Cipher Suits query in Analysis view.

4.8.1 Pcap - TLS connection

In order to try out the Analysis module, let's first look at the file tls-communication.pcap. The file contains the communication of the 192.168.1.2 device with twelve web servers using HTTP protocol wrapped in TLS. For example, the first server, with which a connection is established in the captured communication is the Czech server named root.cz. This DNS name was used during communication was resolved to the IP address 91.213.160.188. This translation is not part of the intercepted communication. In Figure 73 we can see that in the fourth packet, the source device sends a TLS Client Hello message in which it offers the server a total of 29 sets of cryptographic functions with which it can to communicate.

Image 73 The client offers the server a set of cryptographic algorithms, using that it can use to communicate.

In the sixth packet, we can then see that the server responds with a TLS message of type Server Hello, in which it selects the TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 set, as we can see in Figure 74. These parameters are handled by the Mendel system and stored to the network flows that the system creates from the packets.

Figure 74 - The server selects a set for communication TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384.

4.8.2 Analysis using the analysis module

Now we can play the file tls-communication.pcap using the system REPLAY to our Mendel instance. Since all communication in the captured file was from the internal IP address 192.168.1.2, we can first filter it on the network flows view to verify that the system has processed the packets correctly. As we can see in Figure 75, we select the correct time interval and then filter the IP address in question. We can see that the system has found communication with a total of ten

Copyright © 2024 Greycortex

web servers. If we take a closer look at the logged data for the communication with the mentioned server root.cz, then we can see the system for this TLS communication, e.g. the name of the server, the TLS protocol version or also JA3 hash. From the server's point of view, we can see e.g. information about the certificate that the server sends, or information about the selected set of cryptographic functions.

Figure 75 - Data filtering for IP 192.168.1.2 in the network flow view.

It is possible to perform analysis over the stored data for each network flow. We move to the analytical view and perform analysis over the used cryptographic functions in the data that we have entered into the Mendel system. As we can see, for the analysis we use the already predefined view called Top HTTPS Cipher Suits, to which we will also add an attribute for the target IP address. By adding this attribute, we can see which cryptographic suites each web server has chosen to communicate with. Once we run the query, the system will show us that seven web servers used the TLS_AES_256_GCM_SHA384 set for communication and the other three used the TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384.

Figure 76 - Analysis over the cryptographic functions used.

If we look at the predefined query in more detail, we can see that the system searches for network flows in which the answers contain the string . The system will only show us two from these flows. Attributes are the cryptographic set itself and the destination IP address. The system also displays a count that indicates how many times has the cryptographic set attribute been intercepted in the communication with a given server. This is sorted in descending order of the number of these communications from largest to smallest.

We can further see in Figure 77 that by using filters, the analyst can access the large number of attributes that the system captures from L7 application log layer. In this particular case, the query is looking for the Cipher Suite attribute from the server side. However, it is possible to create custom queries that work with e.g. information about certificate issuer (Response Issuer dn), with the name of the server (Request Server Name) and many others. Operators can then be applied to these attributes, that you may be familiar with from SQL languages. In this way, we can define our own queries and store these queries or share them with other users of the system.

Figure 77 - Filters allow the analyst to work with data from the L7 layer application logs.

4.8.3 Exercise - Analyzing DNS Communications

Play the network traffic captured in the file dns.pcap. The file contains the communication of a device on the internal network with IP address 172.28.252.109 with a DNS server with IP address 172.28.255.254. After playing the communication, filter the IP address 172.28.252.109 in the Mendel system and analyze the communication using the Analysis module and the predefined views for the DNS protocol.

1. How many domain names did the source device try to translate during the communication (rrnames)?
2. What 5 types of DNS records did the source device attempt to resolve?
3. How many DNS record type A queries did the source device send?
4. Was any DNS query answered with a return code other than 0? (Yes/No)

4.9 TAG ME OR TAGGING

So far we have only shown and explained the basic use of the Mendel system for data analysis, which with more effort could be done even using Wireshark. Now we're going to look at more advanced functionality, which is, for example, the so-called *tagging*. You will know the principle from other systems. An example would be email client, which allows you to add a "Project Lazarus" label to all emails, that are related to that project. At the same time, you usually have the option to define rules that will automatically assign the label e.g. based on the sender of the email. There is very similar functionality in the Mendel system, except that you can label any device or subnets. At the same time, you can write rules that will perform tagging automatically. The system includes several rules that already do the tagging in the default installation.

4.9.1 Pcap - download file from SMB server

For the labeling demonstration we will use the file smb-file-download.pcap. In this file is captured communication with a Windows system that publishes on the local network shared storage that is accessible via the SMB protocol. The system has an IP address 172.28.0.35 and the path to the accessible storage is \\fileshare.greycortex.com\\DATA. If we examine the pcap using Wireshark, we can first see the TCP Handshake that is created from the IP address 172.28.1.170 to the destination server with IP address 172.28.0.35, followed by the creation of an SMB session using the SMBv2 protocol. In Figure 78 we can notice that in the thirteenth packet, authentication takes place using NTLM with the username

Image 78 - NTLM authentication with the username msrubar.

In packet 21 we can then see an attempt to connect to the shared storage named DATA on the fileshare.greycortex.com server. We can see this attempt in Figure 79.

Copyright © 2024 Greycortex

Figure 79 - Attempting to connect to a shared storage fileshare.greycortex.com\\DATA.

Furthermore, we can see in packet 27 that a user who connected to the shared storage went to a directory named and from this directory then downloaded a file named salaries-1999-2022.zip. As can we can see in Figure 80 the download of the file was successful.

Figure 80 Download of the file named salaries-1999-2022.zip.

Now we can upload the file to the Mendel system using our prepared REPLAY computer that we prepared in section 3.1.3.

4.9.2 Stick a label on me

After successfully playing the pcap, we first examine the network flows that the Mendel system has created from the replayed communication. We then log into the Mendel system, set the correct time interval, filter the source IP address of the client and go to the network flows view, which we can see in Figure 81.

Figure 81 - Login and filter the destination server.

From the communication that is captured in the pcap, the Mendel system created one network flow. In Figure 81 we can see that the client device in the flow is device with IP address 172.28.1.170. The connection was established using a transport protocol TCP to the destination port 445/tcp. We can also see that the Mendel system has recognised the SMB (Server Message Block) and the client downloaded 12.5 MB of data from the server. If we further look at the application data, that the system recorded for the flow, we see important information that were recorded about that SMB session. In Figure 82 we can see that system has recorded the individual transactions. For example, in the transaction with ID 7 we can see that the user accessed the shared storage with the name DATA using the server name fileshare.greycortex.com. In the transaction with ID 9 we can see the salaries directory was opened. In recent transactions we can see the actual reading of the file named salaries-1999-2022.zip and also that the download of the file was successful.

Figure 82 - SMB logged application data.

So we can see that the system correctly recognized all the transactions that occurred at the SMB protocol level and has stored the useful information in the created network flow. This is not all, however, because the system inherently also includes functionality to automatically recognize and mark such servers. If we take a closer look at the device that represents the server itself, i.e. IP address 172.28.0.35, then we can see that it has an SMB label. In Figure 83 we can see that when we click on this label, the system tells us that it is an automatically assigned label, that falls into the Role/Server/File Server/SMB hierarchy. This tells us that that the device is acting as an SMB file server on the network. The label has been automatically assigned to the device based on its behavior. Each SMB server, that's on the network gets that label automatically. This label can also be used in a filter, so it is very easy to to see all SMB servers on the network in the Hosts view.

Figure 83 - Viewing the automatically assigned SMB label.

If we switch to the event view, we can see that the discovery of a new SMB server is reported by an event called . The detector will report you all newly discovered SMB servers on the monitored network. In Figure 84 we can see that in the details of the event, we can see the IP address of the SMB server. In further details, the system will show us the network flows based on which the new SMB server was discovered.

Figure 84 - Each new SMB server on the network is reported by the system.

In this way, the Mendel system is able to find and flag several types of servers on the network. In Section 5.5 we will take a closer look at how the rules that mark and report servers on the network as event are defined. We will also look at how to write custom rules that find and tag your devices based on the requirements you define.

Copyright © 2024 Greycortex

4.9.3 Exercise - Labeling New Services

Upload the krb-816.pcap file to the Mendel system. The pcap contains the communication of two devices on the internal network. The system will discover a new service. Detecting the service, tagging it and reporting it as an event can can take up to 10 minutes.

1. Which service was discovered on the network?
2. What is the IP address of the device that is acting as the client?
3. What is the IP address of the device that is acting as the server?
4. What is the name of the event that detected the new service on the server?
5. What is the ID of the event that discovered the new service?
6. What is the name of the tag that is used in the rule definition with the SID -50023?
7. What is the name of the tag that was assigned to the server (Tag Assignment can take up to 10 minutes)?
8. The captured in the communication also contains an IDS signature reporting weak encryption parameters. What is the ID of this event?

4.10 SUMMARY

We have now covered all five views of the Mendel system, which brings visibility to the infrastructure that the system monitors. We learned that everything starts with packets that must arrive at the the monitoring ports. From the packets, the system composes logical network flows that represent network communications. The packets are then used to create hosts that communicate from subnets/VLANs. The hosts use different transport and application protocols. All of this data is stored in an internal database, which allows us to ask the system various analytical queries or simply graphically display who is communicating with whom. All this can be done historically for any time interval or for data processed in current time. Now we can take a look at the detection mechanisms of the system.

Copyright © 2024 Greycortex

5.1.5 Introduction of detection modules

5. Introduction of detection modules

The Mendel system combines several detection mechanisms that look for malicious patterns, anomalies, and other information in network traffic that can be helpful in analyzing the monitored network. The outputs from these detection methods are presented by the system in the form of events. These events can be viewed from the perspective of the individual devices on which they occurred or from the perspective of their `severity'. The list of detected events can be found in the view Events. Like most other views in the Mendel system, this view is influenced by the time period selected in the calendar and the filters selected by the user. The following are newly applied to the Events view filters:

- Event view level and
- detection modules.

With the view level filter, the user defines whether he wants to see all events or only the most relevant ones. Using the detection modules filter it is possible to define from which detection modules the system should receive events to display. As can be seen in Figure 85, by default the system displays events from all detection modules. If we hover the mouse over the icon of one of the detection modules, the system will display a short description of the selected module. The system contains the following detection modules:

1. Intrusion Detection System (IDS) module.
2. Network Behaviour Analysis (NBA) module.
3. Event Correlation Module.
4. Log Processing module.
5. Tagging module (UNTE - Universal Tagging system).
6. Monitoring module (Status Monitor).

Each of these modules is capable of generating events that say something about the Mendel system itself or the data it processes. The individual modules and their events will be further described in the following sections.

Figure 85 - Iconic filters for displaying events from individual modules.

5.1 SEARCHING IN PACKETS (IDS)

The pattern detection module or *Intrusion detection system* (IDS) is generally a system that is capable of detecting patterns in network traffic that are described by rules, which are generally called . The Mendel system includes the well-known IDS module surikata. The module works with more than seventy thousand rules, which are updated from the GREYCORTEX servers every four hours. IDS Module thus captures specific patterns in network communication. Example of such a rule can be seen in Figure 86, where we can see that the signature has an ID number 2034581. The identification number is unique across all signatures or detection methods in the Mendel system. We can also see that the signature falls into the *Mitre* category Information/Suspicious Communication. The categories can be used in filters or when sending events to external systems (e.g. SIEM systems). We can also see in the signature details when the signature was created. However, the most important information can be found in the properties of the signature. In the properties section we can see the syntactic notation of the signature and what is actually looked for in the network traffic. The signature searches the URI for string ".txt" when using the HTTP GET method. It is therefore looking for the download of a text file. A description of the different options that can be used when for creating signatures can be found in the official Documentation.

Figure 86 - Example IDS signature.

We now return to the network communication that was captured in the pcap file http-get-flow.pcap. In this communication, a text file was downloaded.

So, let's reload the mentioned pcap into the Mendel system or go back to the saved data and go to the Events view. In Figure 87 we can see that an event with the name . This event was detected by the IDS module using a signature with ID 2034581, which can be seen in Figure 86. In Figure 87, we further note that the system, when clicking on a given signature, not only displays the properties of that signature, but also in which direction the communication took place and on which device was the signature detected. Thus, we can see that the signature was detected during an http communication established by a device with IP 192.168.1.11 with the merlin.fit.vutbr.cz server. We can also see that the system offers several options for the given event.

Figure 87 - Detail of the IDS event that reports a text file download.

Copyright © 2024 Greycortex

These options include the following:

1. -- moves selected parameters from the communication matrix as such as source and destination devices to the filter.
2. -- tells the system to not report the given event between these devices. This allows the system to be debugged. System debugging will be described further in section 5.4.
3. -- the system will start recording the network communication into the pcap file.
4. -- creates an incident from the event. Work with incidents will be further described in Section 6.4.
5. -- displays additional information or the complete network flow.

These options are available for most incidents that the system reports. For IDS events, the most interesting section is the one that shows part of the packet in which the signature was detected. In Figure 88 we can see that part of the packet can be found in the event detail in the section labeled , where we can see the name of the downloaded text file. This section can be particularly useful to a security analyst because here he can find additional information necessary for the analysis.

Figure 88 - Viewing the detail of the packet in which the signature found the wanted pattern.

In Figure 89 we can see that the signature detail contains a button to display the network flows. This option will show us the network flows in which the signature was detected. We can see that we are looking at the same flow, that we found in the flows view from section 4.3. Thanks to this, we don't have to search for each flow individually, the system will show us the relevant flows for the selected signature. Network flows can be displayed for events generated by most detection modules.

Figure 89 - Network flows can be displayed for most events.

5.1.1 Exercise - Communication to the TOR network

Upload a file tor_bl.pcap to the Mendel system. The pcap contains the communication of a system from the internal network that communicates through the network TOR.

1. What is the IP address of the internal network device that communicated via TOR network?
2. What is the MAC address of the source device?
3. What version of TLS was used during the communication?
4. What is the identification number of the signature that detected the communication?
5. What is the name of the target subnet of the TOR node that the device started communicating with?
6. On what target port did the communication take place?
7. What ja3Hash was detected in the network stream at the source device?

5.2 RECORDING PACKETS

For each event there is an option to start recording the communication to a pcap file, which becomes active after selecting at least one of the items from the communication matrix. Using this option, it is possible to create a rule that will upload a specific network communication when it occurs from the time of creation of the rule itself. It should be emphasized here that the Mendel system does not store a copy of the network traffic and the recording of data only starts after the rule for the recording of traffic is created. It is not possible to download a pcap with captured network traffic for events, that have been already created.

5.2.1 Capturing network traffic

To demonstrate the functionality of traffic capturing we will use the file wget-http-outbound-download.pcap. Let's play it using tcpreplay and wait for an event called Info: Terse Request for .txt - Likely Hostile. In Figure 90 we can can see that if we view the detail of the event and click on the source IP address, then we can use the button. After clicking this button, we can select a name for the capturing of traffic in the dialog box. Next, we can notice that the system itself pre-selected that we want to start capturing traffic from the source device with IP address 192.168.1.11. Last remaining thing is seting the size limit of the file to which the network traffic will be saved and we can save the traffic capture record.

Figure 90 - Creating an entry to record traffic to a pcap file.

Every time the system receives network traffic from an IP address 192.168.1.11 since saving this record, it will then begin capturing it. Now let's play the pcap file wget-http-outbound-download.pcap again and wait for the event named *Info to be reported again: Terse Request for .txt - Likely Hostile*. The captured network communication files can be downloaded in the settings under Detection and Data Captures subsection. In Figure 91 we can see that the expected event

Copyright © 2024 Greycortex

has been reported again, so let's go to the settings and look at the list of rules for capturing traffic. In the figure we can see that the system contains only one rule that we created for capturing traffic. After viewing the details of this rule, we can see that there is a downloadable pcap file with total size of 4.1kB called new traffic capture.pcap.1. If this file is downloaded and examined, we find that it contains the network traffic.

Figure 91 - Downloading the recorded network traffic.

5.3 SEARCHING NETWORK FLOWS (NBA)

Surikata signatures are useful for finding specific patterns in network traffic. However, legitimate programs, attacker or malware can perform activities that cannot be detected by a specific pattern, that would be found in the packets, but a pattern-based behaviour analysis is needed. A simple verbal description of such a pattern might look like the following. If one device connects to more than thirty other devices on different transport ports in a short defined interval, then report this behavior as a possible port scan. Example of such an NBA detection method can be seen in Figure 92. In the figure, we can see that for the source device with IP address 172.28.252.85, a network activity that corresponds in behavior to that of a program or user, that would try to detect open ports on the device with the IP address 172.28.1.170 was detected. We can also notice that, in contrast to the suricata signatures, NBA methods do not contain the rule definition itself and thus it is not possible to look at the exact definition of the rule. The procedure for detecting a given behavior is part of the Mendel system know-how. NBA detection methods therefore only contain the identifier, the categorization, the creation date, a description and a recommendation for the user.

Figure 92 - Example of an event detected by the NBA module.

This and many other patterns are searched for in network communications by the NBA module or *Network Behaviour Analysis*. The module does not examine the content of individual packets, but analyzes the individual flows as a whole. As a result, it is able to detect a large amount of reconnaissance activity, changes on the network such as connection of new devices, communication from an undefined VLAN, anomalies in the network traffic, periodic communications and many others.

5.3.1 Pcap - nmap scan

In this exercise we will work with a file nmap-scan.pcap. This file captures the communication of the program nmap, which was run with the following parameters under the privileged user with the highest privileges.

If we examine this pcap, we find that it contains the communication of the source device with IP address 172.28.252.85 with the system with IP address 172.28.1.170. As we can see in Figure 93, during the port scan there was communication that was using several protocols. As we will see below, nmap discovered several open ports on the target device, it attempted to connect to these ports and save banners, which are provided by the processes listening on those ports.

Figure 93 - List of logs captured during port scanning.

If we examine individual conversations at the level of TCP protocol, we find that nmap attempted to connect to a total of 2152 ports. In Figure 93 we can see that if we sort the destination ports in ascending order, we can see that nmap proceeded from port 1 upwards. The program picks numbers based on the most used ports. In the figure we can also see that the first 9 ports on the target system did not send back any response.

Figure 94 Attempting to connect to the first nine ports on the target device.

Now we play the aforementioned pcap into the Mendel system using the REPLAY system described in section 3.1.3.

5.3.2 Events looking for specific behavior

After playing the mentioned pcap and logging into the system, we first verify, that the system has processed the packets correctly and created network flows from them. At Figure 95 we can see that we first log into the Mendel system using the username administrator. Then, in the calendar, we set the date August 19 (the day we played the pcap). Next, we go to the Hosts view, where we use the right mouse button on the graph to adjust the time interval to more specific values. We can notice that selecting the time interval also adjusts the the time interval in the calendar. In this way it is possible to quickly adjust calendar time interval if you don't want to set the interval manually in the calendar. Next, click on the number of flows for IP address 172.28.252.85. The system will redirect us to the network flows view and also fill in the filter with IP address we are interested in.

In the network flows list, we first sort the flows by destination port, where we can see several flows where nmap has tried to connect to the low ports. However, we can see that in these flows there were only two packets containing the SYN flag sent, to which the target device sent no response. Next we will only display flows, where the target system sent back a response. We can do this by using the filter which we insert into the destination packet count filter. This tells the Mendel system to display only flows in which the target system has sent back one or more packets.

Figure 95 - Checking flows after pcap has been played.

Copyright © 2024 Greycortex

If we take a closer look at flows that contain bidirectional communication, we can infer which services are running on the target device. In Figure 96 we can see that the device is responding to ICMP Echo messages. We can also see that port 22/tcp is open on the device and is running OpenSSH server. If we look at the flow details for port 22 we can see that the OpenSSH server on the target device is running version 7.4. A web server is also running on the device, which responds on ports 80/tcp and 443/tcp. In the flow detail for port 80 we can see a response code 302 for each transaction. The web server therefore tries to make all requests through the encrypted channel on port 443/tcp. On the target system, the ports 514 and 5432 are also open, and nmap has attempted to send several requests to determine what type of service is listening on that port. We can also notice that nmap has attempted to connect to ports 9010/tcp and 9011/tcp for which, however, the target system sent back only packets containing the TCP flags ACK+RST.

Figure 96 - Viewing flows in which the target system replied back.

Thus, the Mendel system correctly pieced together the network traffic from the replayed flows. Now we're going to look at what it found in the network communication. Let's switch to the detected events view. In the list of events we can see several events that were generated by different modules of the Mendel system. We will focus on the event called . In Figure 97 we can see that this event was detected with a severity of 8 out of 10, i.e. very severe. If we view the the detail of this event, we can see that this event informs us that activity has been detected on the network that corresponds to scanning activity that may have been triggered by a program or user. In the detail we can also see that the system informs us between which devices this activity occurred. If we go deeper into into the details of the event, we can get down to the individual network flows, that we've already explored.

Figure 97 - The Mendel system detected the port scanning using the NBA module method.

Thus, the NBA module methods do not look for patterns in the packets themselves, but instead in the network flows. This is a significant difference between IDS methods and NBA methods. Using the link next to the signature number, we can go to the signature settings, which are located in the Detection section and subsection NBA. In Figure 98 we can see that the system we have installed contains 224 detection methods. The methods are divided into a total of six categories. Each category has its own prefix, which can be found in the name of the signature. The categories are as follows:

1. *Scan* -- scanning and exploration activities.
2. *Discovery* -- detection of changes on the network such as adding a new device, a new communicating service on a device, a new server on the network, etc.
3. *Flow* -- anomalies in network flows.
4. *Periodic* -- periodic and recurring communication.
5. *Outlier* -- an anomaly in terms of data transmitted, communicating partners, and other parameters.
6. Limit -- detection of various parameters based on user defined limits.

Note also that all NBA detectors have an ID with a negative value, while IDS signatures have IDs with a positive value.

Figure 98 - A list of all NBA detectors can be found in the settings.

5.3.3 Exercise - Port Scanning

Download the nba-scan.pcap file to the Mendel system. The file contains network communication, during which an nmap port scan occurred.

1. What is the name of the NBA event that was detected in the replayed communication?
2. What is the identification number of this NBA event?
3. What is the IP address of the device that performed the scan on the network?
4. What is the port number that the device scanned?
5. What is the IP address of the device that has an open port for the RDP service?

5.4 DEBUG THE DISCOVERY METHODS

Let's now consider a situation where we have a device on the network that regularly performs a live test of other devices on the network using ICMP. It performs this activity repeatedly and several times during the day. In this case, the the system's detection of this behavior is correct, but for the monitored infrastructure, it is an expected behaviour. We would therefore like to say to the system that if the monitoring system performs this activity, then we don't want the Mendel system to report this activity any further. But at the same time we require that such activity is reported if it is initiated by another device on the network.

Copyright © 2024 Greycortex

It is possible to define rules in the Mendel system that govern the behaviour of individual detection methods or entire detection modules for specific communication directions and their selected parameters. Thanks to these rules, we are able to tell the system to ignore pings from our monitoring device, but still report if this activity occurs from another device.

5.4.1 Pcap - ICMP Scan

We will demonstrate the ability to debug the system using network communication captured in the file nmap-icmp-scan.pcap. This pcap captures the communication of a monitoring system that scans devices in subnet 172.28.1.0/24. The monitoring system has an IP address of 172.28.252.85. Now we can replay this communication into the Mendel system. This network activity should be responded to by the NBA module, which should classify it as a scan.

5.4.2 Event Ignore Rule (false positive)

After logging into the Mendel system, we can go straight to the event view and use the graph to set the correct time interval for the pcap playback time. In Figure 99 we can see that the mentioned exploration activity was detected by an NBA event called , which has ID -22020.

Figure 99 - Logging in and finding the event detecting the ICMP scan.

Let us now assume that this activity is legitimate, and also expected. So we want to modify the Mendel system so that this event is no longer reported for our monitoring system. In Figure 100 we can see that if we select the source IP address 172.28.252.85 (our monitoring system), the button will be visible. With this option it is possible to create rules mentioned above. After pressing it, the system will offer us a dialog for creating rule that will ignore future behavior. In the picture we can see, that the system has pre-populated the event type and the source device. To description, we'll add the information that the source device is a monitoring system and performs this activity on a regular basis. We can leave the other options on the default settings and save the rule. If the monitoring system will be performing the aforementioned ICMP scan after the rule was created, the system will record this fact, but will not display it by default.

We will now simulate the situation by replaying the file nmap-icmp-scan.pcap.

Figure 100 - Creating a rule to ignore a specific event.

Now let's move to the time interval of the second pcap playback and take a look at the list of events. In Figure 101 we can see that we have changed the time interval to the approximate time of the second pcap replay. On the graph we can see the grey bar that represents that the system has processed and recorded some network communication. However, the system does not show any events. In the left part of the filter we can notice a filter named . Each event that the system reports has a severity that can range from 1 to 10, where a value of 10 represents critical risk and a value of 1 is merely informative. We can notice that by default, this filter has only selected values from 1 to 10. If we look at the filter in more detail, we can see that it also contains a severity value of 0. This severity is given to all events that have some rule applied to them. So if we check the zero option in the severtiy filter, then the system will show us the event that it detected but which we chose to ignore. In the event details, we can also notice that the system displays the number of the rule that was applied to the event. In our case, this is the rule with the label FP-U1, which stands for *user false positive with the number 1.*

Figure 101 - Displaying events that are ignored by the system.

A list of all user created rules can be found in the settings in the Detection section. In Figure 102 we can see that the system only contains one user rule that we have just created. In this setting, a rule can be deleted, modified or a new rule can be created.

Figure 102 - List of created rules (false positives) can be found in settings.

In this way, it is possible to tune the individual detection methods of all modules for the monitored network where the Mendel system is deployed. Tuning the system is a daily affair, as each computer network changes over time. The area of security and malware is also very dynamic, every day brings new challenges that the system tries to to respond to by updating signatures and detection rules.

Copyright © 2024 Greycortex

5.4.3 Exercise - Severity Rule

Upload the malware_exe4.pcap file to the Mendel system. The pcap contains communication of device from the internal network to the US server.

1. What is the IP address of the source device?
2. What is the MAC address of the source device?
3. What destination port was used during the communication?
4. The system reports a file download event for the intercepted communication. What is the name of this event?
5. What is the identification number of this event?
6. What is the name of the object (file) that the client queried and on which the event was reported?
7. Assume that the event is reported with severity 4. Create rule that sets the severity value to 4. This will increase the severity for this event. What is the new severity of the reported event after the second play of the pcap?

5.5 UNTE (UNIVERSAL TAGGING ENGINE)

Another module that the Mendel system contains is called *UNTE* or *Universal Tagging Engine*. With this module it is possible to define custom rules that are able to be assigned to individual hosts or subnets to assign a user-defined *tag*. We have already seen example of tagging in section 4.9.2. By using this module we can define rules that work with individual network flows, hosts or services that run on them or connect to them. For example, imagine a situation where we wanted to find and tag all the DHCP servers on the monitored network. We would probably filter ports 67/udp, 68/udp, or the DHCP service that runs locally on on a device. This way we could find the DHCP servers. We want the system to identify such servers automatically every time such a server appears on the monitored network.

5.5.1 UNTE rule analysis

The Mendel system already contains a rule that searches for DHCP servers in network communication. The rule has ID -50009 and the name

. In the system settings in the Detection section we can find the definition of this rule, which can also be seen in Figure 103. Each UNTE rule contains four basic parts, which are:

1. rule name,
2. the data source (root),
3. predicates (step), and
4. actions.

In the figure, we can see that the rule in question is named DHCP server rule. We can also see that the rule is defined over a list containing devices (hosts) that the system already knows about and has stored in its internal database. On the third line of the rule we see a logical predicate that says, that the system should only look for hosts that have local service DHCP. That is, devices that other devices connect to on local port using DHCP. On the last two lines we can we see the definition of the actions to be performed on devices that meet our defined predicates. The first action is to add a DHCP tag to the selected device and the second is to report this fact in the form of an event that will be reported only once for each device that will be on the network acting as a DHCP server. Thus, this rule will tag each DCHP server on on the network and also inform us about it in the form of an event.

Figure 103 - Discovery rule definition: new DHCP service.

5.5.2 Pcap - Receiving an IP address from a DHCP server

To get a closer look at how the rule works, we will use the network communication that is captured in the file dhcp.pcap. This file captures device with MAC address 00:0b:82:01:fc:42, that connects to the network and wants to receive an IP address using the DHCP protocol. You can examine the communication in Wireshark and then play it back into the Mendel system. In Figure 104 we can see that the system has detected two devices that have IP addresses 192.168.0.1 and 192.168.0.10 from the replayed data. If we click on flows of device with IP address 192.168.0.1, then the system will redirect us to the network flows view with the filter for that IP address automatically filled in. If we look at the detail of the flow, we can see that Mendel has created one logical network flow from the four packets that were played. Although at the beginning of this flow, the device communicated with a source IP address of 0.0.0.0 and was sending the initial DHCP message to broadcast IP address 255.255.255.255, the system created a logical flow that correctly stored the source and destination IP address. We can see that the source device was given IP address 192.168.0.10 from the DHCP server 192.168.0.1. We can see the DHCP tag in the detail of the DHCP server. The tag was automatically assigned to the server using the aforementioned Discovery: new DHCP service rule.

Figure 104 - Inspecting the generated flow from DHCP traffic.

Copyright © 2024 Greycortex

If we now switch from the network flow view to the event view, then we can see the UNTE event itself, which tells us that a new DHCP server has been discovered on the network. In Figure 105, we can notice in the graph, that the discovery of the new DHCP server on the network took place slightly later than the actual assignment of the IP address. UNTE detectors are only triggered in certain time intervals. The detection method that looks for DHCP servers is triggered every five minutes, so it can take up to five minutes to report an event. In the event detail itself, we can see that the system reports that it has detected a new local DHCP service on a device with an IP address of 192.168.0.1. we can then see the same network flow that we could see in Figure 104 in the flow details.

Figure 105 - The UNTE module has discovered a new DHCP server on the network.

So the UNTE module allows you to define rules in which you can use many different conditions that work with already stored information about devices, services, or network flows that have occurred in the monitored infrastructure. Based on the conditions, you can assign individual devices or entire subnetworks to assign tags or report events. UNTE can also correlate events from other modules, making it very versatile tool for detecting behavior you define over data that the system has stored in the database. For more detailed information we recommend exploring the page describing the UNTE module in the user documentation which can be found at https://IP/data/docs/guides/unte-guide.html. We also recommend that you review the other UNTE system rules, which are included in every standard Mendel installation. The UNTE rules definitions themselves can be found at https://IP/settings/system_unte_signatures.xhtml.

Copyright © 2024 Greycortex

5.1.6 Infected device analysis

6. Infected device analysis

The individual detection modules of the Mendel system can show possible indicators of compromised devices on the network. In the previous chapters, we have shown which modules the system has and what are their capabilities. We already know that the system is capable of searching for specific patterns in packets that are defined by signatures. We also know that the system creates logical network flows from individual packets and also keeps track of a list of devices, services and subnets that are active on the monitored network. It then searches these records for behavioral patterns that can again serve as an indicator of compromisation. The system offers several views into the details of network communications that can be used to further investigate what happened on the network or to reconstruct security incident. In the following chapters, we will take a closer look at how use the knowledge gained to investigate a device that has been infected by a malicious program. We will also show how to use these indicators to create security incident and how to export this incident to a text form. We will also learn about the incident management module, which can be helpful if more than one user is working with the system. For demonstration of the analysis of an infected device, we will use a communication log of a computer that has been infected with a known Trojan horse named TrickBot.

6.1 PCAP - TRICKBOT MALWARE

For the purposes of this section, we will analyze the file Trickbot-infection-traffic.pcap, which contains network traffic of a device with IP address 10.9.25.101 that was infected with the TrickBot malware. The communication log does not capture the progress of the infection of the device, but rather its network activity after successful compromisation. If we examine the communication in Wireshark, we find that it captures a total of 16731 packets. This communication could be examined on a packet-by-packet basis, but it's more sensible to play it back into the Mendel system. The captured communication is 72 minutes long. Therefore, you can start the playback and take a lunch break.

6.2 ANALYSIS OF RECORDED EVENTS

After the playback of the pcap is complete, we can log into the Mendel system, set the correct time interval, select the display of all events and switch to the event view. In Figure 106 we can see that the system reports several events with different severities at different time intervals.

Figure 106 - List of detected events of a malware infected device TrickBot.

We will now sort the list of events in descending order by the time of occurrence of a particular event so that we can see how the malware was detected over time. We also disable the aggregation of target devices because the source infected device is only one, so we can see which servers were contacted first over the time.

In Figure 107 - we can see that after sorting the events, the first two are the following events:

1. Info: Terse Request for .txt - Likely Hostile (ID 2034581) and
2. Info: Weak cipher suite used in communication (ID 900013025).

The events were logged at 10:34 and 10:35. In the event details we can see that the infected device downloaded a text file from the server www.msftncsi.com and communicated with another Microsoft server. In this case, it is correctly detected behavior, but it is standard for Windows systems. The event is therefore not related to the malware itself, but the Windows operating system on which the malware was installed.

Figure 107 - Sorting events by time of occurrence and disabling aggregation targets.

At the same minute, a file was downloaded from the German server with the IP address 144.91.69.195. The file download was captured by the signature with the name (ID 2021076), which detects access to the server using the IP address in the Host attribute of the HTTP protocols. Most users today use the domain name of the server they wish to connect to. Accessing the IP address directly can be suspicious. In Figure 108, if we look in the event detail at the network flow itself, we can notice that IP 144.91.69.195 was used in the Host attribute and there was a object named downloaded. We can also notice a strange string in the user-agent attribute that was set by the malware to value .

Figure 108 - Downloading a file from an IP address without DNS name resolution.

Another event named (ID 2014520) informs us that an executable EXE file has been downloaded. If we look into the details of the event and in particular the network flows, then we will see that the signature was detected in the same network flow, in which the communication directly to the IP address was detected. In Figure 109 we can see that if we look again at the network flow, we can see that the process is requesting a php script called solar.php, but the server sends back a file named .

Copyright © 2024 Greycortex

Image 109 - Query php script returns an executable EXE file.

If we export this file from the parsed file and use file program on it, then we will find that it is an executable file for MS Windows systems. In Figure 110 we can also see that the downloaded object has an MD5 hash of 367b6a5c0e0e8ec68ea14a085b1d32b3. If we check the hash, for example on VirusTotal server, then we'll find that it's a payload of the TrickBot malware. The malware therefore first found out what device it was on and then downloaded the payload for a specific operating system and its architecture.

Figure 110 - Getting the file type and md5 hash of the downloaded object.

The download of this Windows executable was also detected by a signature called (ID 2018959). After downloading the payload, it was executed, i.e. execution of the TrickBot malware itself.

After the payload is executed, the malware starts communicating with its controller servers (CnC). The control servers instruct the infected device with other steps. One of these communications was intercepted using the signature (ID 2837550). In Figure 111 we can see that the event detail shows that the infected device connected to the server with IP 187.58.56.26 using port 449/tcp. This port is again very non-standard and the TLS protocol was detected during the communication. If we look into the details of the event, we can see that the infected device communicated with this server more than once. If we look on the details of one of these communications, we can see that the system has saved certificate information from application data and calculated JA3 checksums for the encrypted connection. In particular, note the field, which clearly states that this encrypted communication corresponds to more than one malicious program.

Image 111 - Communication with the control server 187.58.56.26.

Other logged events include a signature called (ID 2026743). This signature detects a device that is trying to connect to a DNS server named ident[.]me. This is a public API, using which it is possible to find out the current public IP address. Malware detects the public IP address from which it is communicating via this server and passes the address to its control server. This makes the management server and the author able to detect the approximate geolocation of the infected device, or even the company in which the device is located. In Figure 112 we can see that the infected device has connected to said server using the TLS protocol and the detection was therefore based on the name that appears in the certificate of the target server. The public IP address returned by the server, is not in the stored network stream, nor are we able to find it in the pcap file.

Figure 112 - Detecting the discovery of a public IP address using the service ident[.]me.

At 10:48 we can then see another important event that was captured by a signature called (ID 2032218). The signature intercepted the infected device communicating with server with IP address 170.238.117.187 on port 8082/tcp. Again, this is communication with another control server.

In this case, the device itself notifies the management server and sends the name of the infected computer (hostname) and other information. This data is sent in the form of a POST request in the URI of the HTTP protocol. In this communication, we can again notice that the connection was made directly to the IP address without domain translation name and was also made to the not-so-standard port 8082/tcp.

In Figure 113 we can see that if we look into the details of the event, then in the network flow we see that the device is sending a string . In this string, is the name of the infected device. We can also see the IP address and port used in the Host attribute and also the special string in the user-agent attribute.

This communication is referred to as a *malware checkin* and the device is giving the control server information, that it is active. The control server then takes this information from the web server log and stores the device in the list of active infected systems.

Figure 113 - Malware sends a checkin packet to the management server.

Once the malware knows it is connected to the Internet, it can start sending useful data to the control server. The first thing the malware will do is try to get the username and password of the user under which the malware was run. The malware infected a device where there was only one user who was also part of the local Administrators group. In Figure 114 we see, that the exfiltration of this data was detected by a signature with the name (ID 2035356). The data goes to the same server to which the initial checkin was sent, i.e., the server with the IP address 170.238.117.187 on port 8082/tcp. In the event detail we can see the packet bytes section, where a piece of the packet, in which the signature was captured, is saved. For better formatting and visibility, let's forward the data to the [CyberChef] tool (https://gchq.github.io/CyberChef/), which is integrated directly into the Mendel system. In the formatted data we can see the string randybachman||P@ssw0rd\$. This is how the malware sends to the control server that the username of the infected device is and the password is `P@ssw0rd\$. This way the administrator of the control server gained access to the account in question, which may be further used in subsequent phases of the attack.

Figure 114 - Username and password discovery and exfiltration.

Copyright © 2024 Greycortex

In the next step, the control server instructs the infected device to send more information about itself. The malware then begins to determine whether the infected device is connected to a domain, what network configuration it has, what processes are running on the device, etc. For this purpose, the malware downloads its own libraries from the control server to find out this information. In this way, it is possible to bypass programs that monitor user access to these information. In Figure 115 we can see that the download of this module was logged by an event called (ID 2834883). The download was from a server with IP address 185.98.87.185. As we can see in the network flow detail, the connection to server was made using HTTP protocol to the destination port 80/tcp. Malware queries two objects from the control server named and . The objects look like images based on the extension, but they are Windows libraries. In the application data transaction, we can we can see that the server returns a return code of 200 for both objects, which tells us that the server has the objects on its local storage and can therefore provide it. So the malware has successfully downloaded these libraries and can use them in the next step to find out the system information.

Figure 115 - Downloading additional network discovery libraries information.

Once the malware has detected the necessary system information, it will again prepare it to send to the control server. This transmission was intercepted by the signature called (ID 2027117). In Figure 116 we can see that the system information was sent to the same server that username and password were sent to. In the packet byte section, in event detail, we can again see the information that the infected device is sending to the remote server. Again, we use the CyberChef tool to to transform the sent data into a more readable form. In the sent data we can see, for example, a list of running processes. This information can be used by the server to know if there are any running security products on the device, that need to be bypassed. In the event that it detects such product, it would again download a module that would put the software out of service. In the list we can also see information about the device, i.e. again the name of the computer, the name of the operating system, its version and more. The malware also sends information about network settings such as IP address, domain name and more. This information can later be used in attacks on other devices on the network.

Figure 116 - Exfiltration of process list information, network settings, and other.

In this way we could go on and analyze other events as well, or individual flows. However, I think that from the events described above you have already gotten an idea of the process of investigating an infected device. In the next section, we'll look at how to record this information and create security incident report.

6.3 EXERCISE - TRICKBOT

Let's revisit the list of events recorded during the communication TrickBot malware described in Section 6.1.

1. The malware attempted to translate a TOR domain. What is the ID of the signature that detected this translation?
2. What is the TOR domain name that the infected device attempted to translate.
3. What is the IP address of the only DNS server with which the infected device communicated?
4. The infected device attempted to attack the domain controller on the port 445/tcp. What was the IP address of this controller?
5. The infected device communicated with three control servers on port 449/tcp. What are the IP addresses of (3) these servers?
6. Using the analysis view, determine to which server (IP) did the infected device send the most packets.

6.4 DESCRIPTION OF THE SECURITY INCIDENT

In Section 6.2, we took a closer look at how we would investigate a compromised device. During this analysis, we discovered many facts that clearly indicate that there has been a security incident. For this incident, we could write a simple description from the information we've gathered so far. A summary of this information would might look like this:

1. IP address of the infected device: 10.9.25.101
2. Download of an executable EXE file as solar.php (phn34ycjtghm.exe) from server 144.91.69.195 (VirusTotal).
3. First contact with CnC server at IP 187.58.56.26:449.
4. First checkin to another CnC server 170.238.117.187:8082 and sending the hostname in the URI (/ono19/BACHMANN-BTO-PC_W617601.AC3B679F4A22738281E6D7B0C5946E42/90).
5. Sending username and password to CnC server 170.238.117.187:8082.
6. Malware downloads auxiliary libraries from server 185.98.87.185:80.
7. Malware sends additional device information such as network settings, process list, system information, and more.

Copyright © 2024 Greycortex

We would somehow like to store the information for the incident. In the Figure 117 we can see that the system allows us to create an incident from a view that we currently have open. When we click on the report incident option, the system will open a dialog box to fill out the incident report. In this dialog box we fill in the information we have found so far. The first thing we select is the name of the incident, where we use "Device infected with TrickBot malware". Next, we select the risk that we rate as critical, since it is an active infection. For the status of the incident, we select "BEING ANALYZED", as we can further analyze the intercepted communication. The system also offers the ability to add predefined or user-defined tags. Note that these tags have nothing to do with the tags described in in Section 4.9. The logged in user is automatically filled in as the assigned user. In this case, the user is the administrator. The important part is then the field for the actual incident description. The field offers wysiwyg editor, which can be used to create a formatted description of the incident. Once the description is complete, we can save the incident description. The system will then notify us that the incident has been successfully saved to the incident manager.

Figure 117 - Saving a security incident description to management incident.

The Mendel system has an incident manager that can be used to describe incidents, share incidents with other users and further work with the incidents. The incident management module essentially acts as a ticketing system. You can think of each incident as a ticket that was created by someone, has a status, a description, and can be discussed in the form of comments from other users. In Figure 118 we can see the transition to the *Incidents* view. There is only one incident in the incident list, that describes an infected device. We can notice that the newly created incident has been given the identification number I-1. Each incident has a unique identification number in the system. In the image we can also see the ability to add comments to incidents. So, if multiple users are working with the system, there can be a discussion over individual incidents.

Figure 118 - Switching to the incident management view.

An important attribute of each incident is also a data link, which takes the user to the view that the user was looking at when the the incident was reported. As we can see in Figure 119, using this link, we can go back to the list of events that were recorded at the time of the device infection. From the current view, it is also possible to create a link that can be sent to a colleague to whom we want to show exactly what we're currently looking at.

Figure 119 - Each incident contains a link to the view from which it was the incident was created.

Another useful option is to export selected incidents to a PDF or CSV file. This option is useful when reporting or summarizing incidents for a selected period. The export of an I-1 incident can be seen in Figure 120. The resulting PDF file contains a predefined front page, a table of exported incidents, a description of each incident, and a predefined header, that contains information about the system itself, time when the incident was created and about the user, that reported this incident. If we want to view this incident in more detail, you can find it as an attachment Export_2022-08-23.pdf. Incident Management Module can be used by a team of analysts who use the Mendel system to analyse network communications and needs to share information about individual incidents and discuss them.

Figure 120 - Exporting a selected incident to PDF.

This demonstrated how to proceed with the analysis of detected incidents, how to investigate these events and how to save the investigated facts as an incident directly into the Mendel system.

Copyright © 2024 Greycortex

5.1.7 Exercise answers

7. Exercise answers

4.3.3 EXERCISE - INSTALLING A LINUX PACKAGE

4.6.1 EXERCISE - SUBNETS, GUESTS AND SERVICES

4.7.3 EXERCISE - WHO IS MY DEVICE COMMUNICATING WITH?

4.8.3 EXERCISE - ANALYZING DNS COMMUNICATION

4.9.3 EXERCISE - LABELING NEW SERVICES

5.1.1 EXERCISE - COMMUNICATION TO TOR NETWORK

5.3.3 EXERCISE - PORT SCANNING

Copyright © 2024 Greycortex

5.4.3 EXERCISE - SEVERITY RULE

6.3 EXERCISE - TRICKBOT

Copyright © 2024 Greycortex

5.1.8 Attachments

8. Attachments

8.1 PCAP FILES

- download-vlc-merlin-http.pcap (MD5: f796a9e31aefb5c95784d80b7a160167)
- http-get-flow.pcap (MD5: 34dc7544f04aab32a2d52a039de8843a)
- test-icmp-ping.pcap (MD5: f4ee0fcdc78a7d72e702bfce29d6adbc)
- network-peers-ssh-scan.pcap (MD5: c81493c86be46a4f6e69c743dae29d38)
- tls-communication.pcap (MD5: 79b85873c0edeb08df7ca8c29275812e)
- smb-file-download.pcap (MD5: 31ef42b6f79a51a74ef726065572f668)
- wget-http-outbound-download.pcap (MD5: 7327a5ffee9982d9c8b5ed8257418e4)
- nmap-scan.pcap (MD5: 23c17894e1ab46d4c49b42baa727a647)
- nmap-icmp-scan.pcap (MD5: fa8908b5d919428c137865ea0c4369af)
- dhcp.pcap (MD5: 350a4a71b742023fcfed000c5f664fd6)
- Trickbot-infection-traffic.pcap (MD5: d2674d0c60c5da77008a0ea6d403dd4a)

8.2 PCAP FILES FOR THE EXERCISE

- dns.pcap (MD5: c192a8adfe7e1e5e9f02386cfba708de)
- krb-816.pcap (MD5: 7c66ae741447efcabfa2dbd847afe8b5)
- malware_exe4.pcap (MD5: 5d8de4abb2e31be47bc7b79e1f8f7050)
- nba-scan.pcap (MD5: 7a375a8a0fc7c06b01e0aabb142ee811)
- package-install.pcap (MD5: adabae8d35c586537d7be78d97dff496)
- peers-ipv4.pcap (MD5: af8081e0ac9a8b45b2fee6ea946bdda8)
- subnets-hosts-services.pcap (MD5: 074757251734ad24b1d237851bb19db1)
- tor_bl.pcap (MD5: 5667548ce885ca381c837440348435c6)

8.3 INCIDENT EXPORTS

- Export_2022-08-23.pdf (MD5: 210ad2b44f3b50a997194b19cfdd7336)

Copyright © 2024 Greycortex

6. About

6.1 EULA

6.1.1 End-user license agreement (EULA)

This License Agreement ("**License Agreement**") governs the grant of rights permitting you, as the person or entity concluding this License Agreement (the "**Customer**" or "**You**") to use and enjoy the benefits of GREYCORTEX MENDEL, a proprietary software solution providing advanced network security and performance monitoring capabilities, delivered or made accessible to You by GREYCORTEX s.r.o., with its registered seat at Purkynova 649/127, 612 00 Brno, Czech Republic, Identification Number 05060711, registered by the County Court in Brno under Section C, Insert No. 93360 (hereinafter "**GREYCORTEX**" or the "**Supplier**"). This is a License, and not an agreement for sale or transfer of rights in any manner.

The scope and conditions of usage of the Software are further defined in this License Agreement.

BY PHYSICALLY INSTALLING THE SOFTWARE AND CLICKING THE "NEXT" BUTTON TO COMPLETE INSTALLATION, YOU ARE STATING THAT YOU HAVE READ THE TERMS OF THIS AGREEMENT AND AGREE TO ALL OF THE TERMS CONTAINED HEREIN, IN FULL AND WITHOUT MODIFICATION. IF YOU DO NOT AGREE TO THESE TERMS AS WRITTEN, DO NOT CONTINUE FURTHER AND DO NOT INSTALL

Copyright © 2024 Greycortex

THE SOFTWARE IN ANY FORM OR FASHION. INSTALLATION OF THE SOFTWARE SHALL BE CONSIDERED TO BE AGREEMENT TO THE TERMS OF THIS AGREEMENT AND IS INSEPERABLE THEREFROM.

Copyright © 2024 Greycortex

1. SOFTWARE, MEANS OF INSTALLATION/ACCESS; AUTHORIZED PURPOSE

a. Software

As used in this Agreement the term "Software" means: (i) the GREYCORTEX MENDEL computer program and all components thereof; (ii) all the contents of disks, CD-ROMs, DVDs, e-mails and any attachments, or other media with which this Agreement is provided, including the object code form of the Software supplied on a data carrier, via electronic mail or downloaded via the Internet; (iii) any related explanatory written materials and any other possible documentation related to the Software, any description of the Software, its specifications, any description of the Software properties or operation, any description of the operating environment in which the Software is used, instructions for use or installation of the Software or any description of how to use the Software (hereinafter referred to as "Documentation"); (iv) copies of the Software, patches for possible errors in the Software, additions to the Software, extensions to the Software, modified versions of the Software and updates of Software components, if any, and (v) any license keys related to the above Software components licensed to You by the Supplier pursuant to Article 2 of this Agreement.

b. Software may be made available in two forms; by subscription and by perpetual license:

i. Subject to the conditions set out in this License Agreement, the Customer or an Authorized Person (as defined in Article 1.3.1 below) will install the Software on necessary devices as also included in 1.3.1 below ("**Perpetual License**"); or

ii. Subject to the conditions set out in this License Agreement, and subject to availability, the Customer will obtain access to the Software, installed and configured as necessary to meet the Authorized Purpose, provided by Authorized Person or Supplier, through a user account, created by using unique login credentials ("**Subscription License**").

iii. For the avoidance of doubt – "Software" as used singly herein shall refer to all versions, variants, or other types of the Software, collectively.

c. Installation. Software supplied on a data carrier, sent via electronic mail, downloaded from the Internet, downloaded from the Supplier’s servers or obtained from other sources requires installation. You must install the Software, or ensure that the software is installed correctly by an Authorized Person, and correctly configure the device upon which it is installed or ensure that it is correctly configured, complying at least with requirements set out in the Documentation. The installation methodology is described in the Documentation. No computer programs or hardware which could have an adverse effect on the Software may be installed on the device on which You install the Software.

i. For the avoidance of doubt, you may install and configure the Perpetual License, or cause the same to be installed and configured by Supplier or Supplier’s authorized Reseller or Distributor; their employees and/or management (hereafter "**Authorized Person**"), on as many devices (including sensors, collectors, all-in-one appliances; both virtual and/or physical) as permitted by separate agreement between You and Supplier or Authorized Person.

d. The purpose of use of the license granted herein shall be to use the Software as installed for the purposes for which the Software is designed by its functionality; i.e. network security and performance monitoring of Your computer network, as defined in the agreement for purchase of the license to Software between You and Supplier or You and Authorized Person. (hereinafter referred to as the "**Authorized Purpose**").

e. The Supplier may, at its discretion, release minor Software enhancements and/or bug fixes (related to Software functionality and which shall be updated automatically in Your version of the Software) as well as complex Software upgrades (hereinafter referred to collectively as "**Updates**"). The terms of this License Agreement shall apply to all free Updates that are made available by the Supplier and that the Supplier does not expressly exclude from such licensing regime. The Supplier reserves the right to release major upgrades, especially if providing a particular enhancement, on a paid basis ("**Paid Upgrades**"). Such Paid Upgrades may be subject to separate license terms provided with such major upgrade release. If no separate license terms are provided with the major upgrade, the terms of this License Agreement shall apply (save for the payment of the fee for the Paid Upgrades).

2. LICENSE

Copyright © 2024 Greycortex

The following rights are granted by the Supplier to the Customer (the "License"):

a. License to Install and/or Use (Perpetual License)

The Supplier grants to the Customer a non-exclusive, non-assignable, non-sublicensable (except as permitted herein), and royalty-free license to install and use the Software. The rights granted under this Section 2.1. are granted for the period as described in Article 5 hereof. You are entitled to make a necessary number of non-functional back-up copies of the Software, exclusively for disaster recovery, emergency restart, or backup purposes, provided that these reproductions contain all notes related to the copyrights, trademarks, and/or other intellectual property rights to the Software, by which the Software and/or any part thereof is identified.

i. As referenced in Article 1.2.1. above, this grant of rights is specific to You.

ii. Exercising End User rights. You must exercise End User rights by engaging in the Authorized Purpose in person, via your employees, or via an Authorized Person. You are only entitled to use the Software on network components ("devices" as referenced in Article 1), owned or rightfully used by You, and for Your network (comprising the network connections, data, data transfer within and without these devices and the internet).

b. License to Use (Subscription License)

Subject to availability, Supplier grants to the Customer a non-exclusive, non-assignable, non-sublicensable (except as permitted herein), and royalty-free license to use the Subscription License version of Software for the Authorized Purpose. The rights granted under this Section 2.1. are granted for the period as described in Article 5 hereof.

i. As referenced in Article 1.2.2. above, this grant of rights is specific to You.

ii. Exercising End User rights. You must exercise End User rights by engaging in the Authorized Purpose in person, via your employees, or Authorized Person. You are only entitled to use the Software on network components ("devices" as referenced in Article 1) owned or rightfully used by You, and for Your network and for Your network (comprising the network connections, data, data transfer within and without these devices and the internet).

c. With respect to the Documentation accompanying or related to the Software (the "**Documentation**"), the Customer shall be entitled to make such number of reproductions (in a printed or electronic form) as may be strictly necessary to benefit from the Software provided that these reproductions are used by the Customer in compliance with this License Agreement (and any other applicable documents), especially as regards intellectual property, and the Documentation is not further published or distributed to third parties.

d. Each reproduction of the Software and/or Documentation shall be created only in the manner permitted under this License Agreement and shall be identified by the Customer, using all copyright, trademark, and/or any other legends and notices for the protection of intellectual property rights of the Supplier, its suppliers and/or third parties who are the rightful owners and/or licensees of the Software and/or Documentation.

e. You may sub-license the rights contained in Article 2.1 and 2.2 of this License, to Authorized Person, as defined in Article 1.3.1 herein, as required for the Authorized Purpose, as defined in Article 1.4 herein.

3. LICENSE RESTRICTIONS

Copyright © 2024 Greycortex

Unless expressly permitted by this License Agreement or by the respective legal regulations, the Customer shall not be entitled to:

a. Lend, rent, sell, assign, sublicense (except as permitted herein), provide or otherwise distribute the Software or any rights granted by this License Agreement without explicit written permission of the Supplier;

b. Install the Software, or cause it to be installed on more devices than permitted by agreement between You and Supplier or You and Authorized Person, except as provided in Article 1.

c. Use the Software for purposes other than the Authorized Purpose, including, but not limited to the use of the Software for trainings or counselling services, or use the Software for any commercial services.

d. Change, adapt, or adjust the Software or make any derivative work from the Software;

e. Reverse engineer, decompose, disassemble, or decompile the Software (except as may be permitted by applicable law and only in the scope and for the purpose permitted by the applicable law);

i. You may not reverse engineer the Software, except in such cases as it is absolutely necessary to insure interoperability as contemplated in Section 15 of DIRECTIVE 2009/24/EC OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 23 April 2009 on the legal protection of computer programs. In the event of such permitted reverse engineering, You will send the Supplier a prior written notice to sales@greycortex.com, and enable the Supplier, at its own discretion, and in a reasonable time, to submit an offer and provide the You with information and assistance that is reasonably required to ensure the interoperability of the Software with Your other software and/or hardware. Such information and assistance may be subject to a mutually agreed fee. In the written request, the You shall provide the Supplier with all the necessary technical specifications and plans for the intended reverse engineering along with the purpose of the reverse engineering.

A. In the event You undertake such reverse engineering as permitted under Article 3.5.1, You agree that any such results of this undertaking (including but not limited to source and/or object code, documentation, notes, etc.), whether successful or unsuccessful, immediately becomes property of Supplier, in full, as though such results had been produced solely by Supplier.

f. Use the Software for development of other Software or for creation of any type of a competitive product;

g. Permit third parties access to the Software or its use, except as explicitly permitted in Articles 1.2.1, 2.5, 3.2, and 8 of this License Agreement; and/or

h. Grant the consent to use the Software and/or assign or sublicense the License to third parties, except as permitted by this License Agreement.

i. You may not be, and may not act on behalf of: (a) any person who is a citizen, national, or resident of, or who is controlled by, the government of any country to which the United States has prohibited export transactions; or (b) any person or entity listed on the U.S. Treasury Department list of Specially Designated Nationals and Blocked Persons, or the U.S. Commerce Department Denied Persons List or Entity List. You further state that you will comply with all applicable regulations of the United States Department of Commerce and with the United States Export Administration Act, as amended from time-to-time, and with all applicable laws and regulations of other jurisdictions with respect to the importation and use of the Third-Party software (as further referenced in Article 8 hereof), to assure that such software and confidential information related thereto are not exported, imported, or used in violation of law or applicable regulation.

4. INTELLECTUAL PROPERTY

a. The Software and all rights, title, and interest without limitation including proprietary rights and intellectual property rights thereto are owned by Supplier and/or its licensors in full ("Intellectual Property"). Said rights include, but are not limited to rights generally categorized as those existing in copyright, trademark, patent, and/or trade secret rights. They are protected by international treaty provisions and by all other applicable national laws of the country in which the Software is being used. The structure, organization, and code of the Software are the valuable trade secrets and confidential information of Supplier and/or its licensors. You must not copy the Software, except as set forth in Article 2.1. You are expressly forbidden from reverse engineering the Software except as provided in Article 3.5.1. In no circumstances, may you remove trademarks or any other source-identifiers from the Software or any Documentation there accompanying. Any copies which You are permitted to make pursuant to this Agreement must contain the same copyright, trademark, and other proprietary notices that appear on the original Software and/or Documentation. If You reverse engineer, reverse compile, disassemble or otherwise attempt to discover the source code of the Software, in breach of the provisions of this Agreement, You hereby agree that any information thereby obtained shall automatically and irrevocably be deemed to be transferred to and owned by the Supplier in full, from the moment such information comes into being, notwithstanding the Supplier’s rights in relation to breach of this Agreement.

i. This Agreement is a license, and as such does not confer any ownership rights in the Intellectual Property to You either in whole or in part.

5. TERM AND TERMINATION

Copyright © 2024 Greycortex

This Agreement is effective from the date You agree to the terms of this Agreement by installing the Software, and shall continue for the term provided herein:

a. In the case of a Subscription License, the term of this Agreement shall be one year (12 months), which shall renew upon payment of such fees corresponding to the next one-year period as are agreed upon by You and Supplier or You and Authorized Person.

b. In the case of a Perpetual License, the term of this grant of rights shall be perpetual and Software features shall continue in full following payment of yearly maintenance fees as agreed upon by You and Supplier or You and Authorized person (Maintenance Fee). Failure to pay said Maintenance fee will result in the loss of certain functionalities made available through third-party providers. For the avoidance of doubt, such Maintenance Fee must be paid to Supplier or Authorized Person, and may not be made to any third parties directly.

c. You may terminate this Agreement at any time by permanently uninstalling, destroying and returning, at your own costs, the Software, all back-up copies and all related materials provided by the Supplier or Authorized Person, in which case You shall have no rights to compensation of any type for any reason. Irrespective of the manner of termination of this Agreement, the provisions of Articles 3.5.1.1, 4, 7-11 shall continue to apply to you in perpetuity.

6. RESERVATION OF RIGHTS

The Supplier hereby reserves all rights to the Software, with the exception of rights expressly granted under the terms of this Agreement to You as the End User of the Software.

7. DATA TRANSFER

You acknowledge and agree that the Software is, or may collect, and may transfer to Supplier, data, including, but not limited to data about server condition and/or status (including but not limited to version, license, update status, CPU, ram usage, disk space, performance statistics for each system component, summary traffic throughput, disc space usage, system configuration, system disruption, and to satisfy Supplier that Your license key is authentic), hereafter "Data Collection". Data Collection may be necessary to provide You the relevant Software functionalities and updates. You may be required to set Offline mode to stop further Data Collection that supports these functions, or take steps during configuration to facilitate this Data Collection.

The decision to collect, when to collect, and what data to collect is reserved solely to Supplier and shall be exercised in Suppliers sole discretion, and may be changed from time to time without notice, but at no time shall Supplier collect Your or End User’s personal data, or data which can be used to identify an individual or their activities.

8. THIRD-PARTY RIGHTS

The Software may contain products and/or parts of such products owned by (or licensed to) independent third parties, including software developed by using open source software programs ("**Third-Party Software**"). Unless stipulated otherwise with respect to such Third-Party Software, the use of such Third-Party Software is subject to this License Agreement. Specific licenses made available by such third parties may apply in the event that:

a. It is expressly stipulated so in the documentation relating to the Third-Party Software; or

b. The Third-Party Software is used as a standalone product separate from the Software.

c. For the avoidance of doubt, in the event of a discrepancy between applicable rights reserved by Third Parties and GREYCORTEX, the most restrictive grant of rights to You shall be controlling for Your use of the Software. Relevant Third-Party Software licenses shall be referenced in Documentation accompanying the Software.

9. END USER DECLARATIONS

AS THE END USER, YOU ACKNOWLEDGE THAT THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW. NEITHER THE SUPPLIER, ITS LICENSORS OR AFFILIATES NOR THE COPYRIGHT HOLDERS MAKE ANY REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR THAT THE SOFTWARE WILL NOT INFRINGE ANY THIRD-PARTY PATENTS, COPYRIGHTS, TRADEMARKS, OR OTHER RIGHTS. THERE IS NO WARRANTY BY THE SUPPLIER OR BY ANY OTHER PARTY THAT THE FUNCTIONS CONTAINED IN THE SOFTWARE WILL MEET YOUR REQUIREMENTS OR THAT THE OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED OR ERROR-FREE. YOU ASSUME ALL RESPONSIBILITY AND RISK FOR THE SELECTION OF THE SOFTWARE TO ACHIEVE YOUR INTENDED RESULTS AND FOR THE INSTALLATION, USE AND RESULTS OBTAINED FROM IT.

No other obligations. This Agreement creates no obligations on the part of the Supplier and its licensors other than as specifically set forth herein.

Copyright © 2024 Greycortex

10. LIMITATION OF LIABILITY

a. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, WHATEVER THE LEGAL BASIS FOR THE CLAIM, THE SUPPLIER WILL NOT BE LIABLE FOR ANY DIRECT OR INDIRECT DAMAGES (INCLUDING, WITHOUT LIMITATION, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY KIND WHATSOEVER (INCLUDING, WITHOUT LIMITATION, ATTORNEYS' FEES), DAMAGES FOR LOST PROFITS OR REVENUES, BUSINESS INTERRUPTION, OR LOSS OF BUSINESS INFORMATION) DUE TO, RESULTING FROM, OR ARISING IN CONNECTION WITH THIS LICENSE AGREEMENT, THE SERVICES, THE SOFTWARE, THE DOCUMENTATION, OR THE FAILURE TO PERFORM SUPPLIER’S OBLIGATIONS.

11. MISCELANEOUS

a. Choice of Law and Venue. This License Agreement shall be construed in accordance with the laws of the Czech Republic without giving effect to any contrary principles of conflicts of laws or the United Nations Convention on the International Sale of Goods. The Parties agree that any disputes arising under or relating to this License Agreement shall be resolved by the competent courts in the Czech Republic.

b. Notice. All notices and return of the Software and Documentation must be delivered to: GREYCORTEX s.r.o, Purkynova 127, 612 00 Brno, Czech Republic.

c. Waiver/Modification. This Agreement may only be modified in written form, signed by an authorized representative of the Supplier or a person expressly authorized to act in this capacity under the terms of a power of attorney. Failure by Supplier to enforce a provision of this License does not constitute waiver of that provision, nor does it constitute a waiver of the right to enforce that provision at a later time of Supplier’s choosing or to do so retroactively, nor does it constitute waiver of all of the provisions herein.

d. Integration. This is the entire Agreement between Supplier and You relating to the Software and it supersedes any prior representations, discussions, undertakings, communications or advertising relating to the Software.

e. Severability. Should any of the provisions of this Agreement be invalid or unenforceable, this shall not affect the validity of the other provisions of the Agreement, which shall remain valid and enforceable in accordance with the conditions stipulated therein.

Copyright © 2024 Greycortex

6.2 License

6.2.1 Third-Party End-User Licenses

As referenced in the GREYCORTEX MENDEL End-User License Agreement, MENDEL contains some third-party software. The following Third-Party End-User Licenses are provided as referenced in the End User License Agreement and are in English only.

Suricata

The most recent Suricata Public License can be found at:

http://suricata.readthedocs.io/en/latest/licenses/gnu-gpl-v2.0.html

GeoLite2

This product includes GeoLite2 data created by MaxMind, available from:

https://www.maxmind.com.

Proofpoint

(A full copy of the most recent End User license may be found on

subsequent pages of this document)

PROOFPOINT, INC.

END-USER LICENSE AGREEMENT

This End-User License Agreement (this "Agreement" or "EULA") applies to

all Subscriptions of Proofpoint, Inc., a Delaware corporation

("Proofpoint"), purchased by you, as an "End-User" (or "you"), from

either Proofpoint, or an Authorized Party. This Agreement is effective

on the earliest of (i) the click acceptance by the End-User, (ii) the

execution of this Agreement or (iii) the download, installation or use

of any Subscription, or part thereof, by the End-User.

BY CLICKING "ACCEPT" BUTTON BELOW (AS APPLICABLE), SIGNING THIS

AGREEMENT OR BY DOWNLOADING, INSTALLING, OR USING THE SUBSCRIPTIONS, (A)

EACH PERSON REPRESENTS AND WARRANTS IT HAS THE RIGHT, POWER AND

AUTHORITY TO BIND THE END-USER FOR WHICH SUCH PERSON IS EMPLOYED OR ACTS

AS A CONSULTANT AND (B) ON BEHALF OF SUCH END-USER CONFIRMS IT HAS READ

THIS AGREEMENT AND IS CONSENTING TO BE BOUND BY THIS AGREEMENT. IF YOU

DO NOT AGREE TO THE TERMS AND CONDITIONS OF THIS AGREEMENT, DO NOT SIGN

THIS AGREEMENT, AND DO NOT DOWNLOAD, INSTALL, OR USE ANY PART OF THE

SUBSCRIPTIONS. NEW VERSIONS OF THIS AGREEMENT MAY BE PUBLISHED FROM TIME

TO TIME AND EACH END-USER SHALL ACCEPT SUCH TERMS PRIOR TO THE DOWNLOAD

OR INSTALLATION OF ANY SUBSCRIPTION OR SUBSCRIPTION UPDATE.

Copyright © 2024 Greycortex

1. DEFINITIONS

Terms used in this Agreement, not otherwise defined, shall have the following meaning.

"Authorized Party" means a Person authorized and licensed by Proofpoint as a provider of the Subscriptions (or a specific Subscription) as an authorized reseller, an original equipment manufacturer or otherwise duly authorized as a provider of the Subscriptions (or a specific Subscription).

"Compilation" means a work, which combines a Subscription or portions thereof with any subscriptions, services, programs, code, or other products not governed by the terms of this Agreement.

"Third Party" means any Person other than End-User hereunder or Proofpoint.

"Person" means any individual, entity, corporation, partnership, joint venture, limited liability association, joint-stock trust, unincorporated organization or governmental body.

"Subscription" means all subscriptions, products and proprietary information or services of Proofpoint, which have been purchased by End-User pursuant to this Agreement either directly from Proofpoint or through an Authorized Party. Such Subscriptions include, but are not limited to, Proofpoint's ETPro (IDS/IPS ruleset), IQRisk Rep List, IQRisk Rep Query, IQRisk Suite, PCAP Feed, URL Feed, any other customized data feed, subscription or product from Proofpoint and any modification and any historical data product subsets of any such Subscriptions.

2. LIMITED GRANT

Subject to the terms and conditions of this Agreement, and any separate agreement between End-User and an Authorized Party (if applicable), and in consideration of the license fees paid, Proofpoint hereby grants to End-User a worldwide, non-transferable, non sub-licensable, non-assignable, non-exclusive license for the Subscriptions purchased, subscribed to and paid for by End-User.

The Subscriptions, number of licenses and/or access levels for certain Subscriptions, shall be reflected in End-User's electronic subscription, a quote from Proofpoint and the corresponding purchase order from End-User or such other agreement entered into with an Authorized Party.

End-User may purchase, subscribe to and pay for the Subscriptions

Copyright © 2024 Greycortex

through an Authorized Party or directly with Proofpoint through an electronic means or by purchase order. In no event shall the license granted under this Agreement extend to any Subscription, or for any term, not duly paid for by End-User.

Any Subscription purchase of the IQRisk Reputation List shall only entitle End-User to affect or respond to data or communications captured by, or within the data-path of the device(s) paid for by End-User. Use of the IQRisk Reputation List content to affect, respond, or correlate data aggregated from other devices and data sources requires specific approval to do so. Such approval, or grant of license, may be provided in a purchase order under this Agreement.

Proofpoint may grant a limited temporary license for evaluation purposes to End-User for a maximum of five (5) sensors. Such limited license shall be restricted solely to an End-User evaluating the applicable Subscription internally and such End-User shall not have any other rights as set forth in this Agreement. All additional restrictions on use under this Agreement shall apply. For any such evaluation-limited license, in which the End-User has not paid the full commercial price for a Subscription, Proofpoint reserves its rights to unilaterally terminate this grant and restrict all access to such Subscription. If such End-User does not purchase a Subscription after the evaluation period communicated by Proofpoint, such End-User shall remove the Subscription from all systems.

All rights, not granted under this Agreement, are reserved by Proofpoint.

3. LIMITATIONS ON LICENSE; RESTRICTED ACTIVITIES.

End-User recognizes and agrees that the Subscriptions (i) are the property of Proofpoint, (ii) contain valuable assets, proprietary information and property of Proofpoint, and (iii) are provided to End-User under the terms and conditions of this Agreement and any additional agreement with the Authorized Party (if applicable).

Notwithstanding anything to the contrary in this Agreement, End-User agrees that it shall *not* do any of the following:

(a) download, use, install, deploy, perform, modify, license, display, reproduce, distribute, or disclose the Subscriptions (even if

Copyright © 2024 Greycortex

merged with other materials as a Compilation) other than as

permitted under this Agreement;

(b) sell, license, transfer, rent, loan, reproduce, distribute to a

Third Party or disclose the Subscriptions (in whole or in part and

whether done independently or as part of a Compilation) to a third party;

(c) share any user authentication information and/or password provided

to an End-User with any Third Party to allow such party to access

the End-User's account with Proofpoint or to otherwise access the Subscriptions;

(d) use the IQRisk Reputation List to affect, respond, or correlate data

aggregated from other devices and data sources; or

(e) alter or remove any copyright notice or proprietary legend contained in or on the Subscriptions.

Paragraphs (a) though (e) of this Section 3 are collectively referred to

as the **"Restricted Activities"**

4. REPRODUCTION OBLIGATIONS

Each End-User agrees that any embodiment of the Subscriptions

permitted under this Agreement will contain the notices set forth in below:

The contents of this file are subject to an End-User License Agreement

or a Data License Agreement with Proofpoint, Inc. You may not use this

file except in compliance with the applicable Agreement. You may obtain

a copy of the applicable terms of use at http://www.proofpoint.com/license/.

Subscriptions distributed under the Agreement are distributed on an "AS

IS" basis, WITHOUT WARRANTY OF ANY KIND, either expressed or implied.

See the Agreement for the specific language governing rights and

limitations under the Agreement. The developer of the Subscriptions is

Proofpoint, Inc.

In addition, to the extent End-User makes any copies of the

Subscriptions as permitted under this Agreement, End-User agrees to

ensure that any and all such copies shall contain:

(a) a copy of an appropriate copyright notice and all other applicable

proprietary legends;

(b) a disclaimer of any warranty consistent with this Agreement; and

(c) any and all notices referencing this Agreement and absence of warranties.

5. TERMINATION

Unless earlier terminated as set forth herein or in an agreement with

Copyright © 2024 Greycortex

the Authorized Party (if applicable), this Agreement is effective so long as the Subscriptions are used by End-User. End-User's rights under this Agreement will terminate automatically if End-User fails to comply with any of the limitations or other requirements described herein. Upon any termination or expiration of this Agreement, End-User must cease use of the Subscriptions and must destroy all copies of the Subscriptions. Notwithstanding the foregoing, limitations on the use of the Subscriptions and the liability of Proofpoint shall survive any termination of this Agreement.

6. FEES; TERM

Subscriptions purchased directly from Proofpoint shall be subject to the agreed fee for the applicable term, which shall be reflected in the electronic purchase or a quote by Proofpoint and a corresponding purchase order from End-User. Once End-User has submitted a purchase electronically or submitted a purchase order and Proofpoint has provided the access information for such Subscription, the sale will be considered final and not subject to any refunds. Subscriptions purchased through an Authorized Party are subject to the fee agreement between End-User and such Authorized Party for the applicable term.

7. AUTOMATIC RENEWAL

If End-User has purchased a Subscription through Proofpoint (and not an Authorized Party) End-User's subscription will be automatically renewed at the expiration of the then-current term (subject to local regulation restrictions), excluding promotional and discount pricing. The new term will be for the same duration as the expired term unless otherwise specified at the time of renewal. If End-User wishes to cancel its subscription, End-User must, prior to the expiration of the then-current subscription term, provide written notice to Proofpoint. If End-User has purchased the Subscription through a purchase order or other payment method, Proofpoint will invoice End-User for such renewal term. For the avoidance of doubt, if payment is not duly and timely received or available for any renewal term, Proofpoint reserves the right to terminate End-User's access to the Subscription until such payment has been received. End-user must provide current, complete and accurate

Copyright © 2024 Greycortex

information for its billing account. End-User is responsible for ensuring this information is correct and must promptly update all information. If End-User has purchased a Subscription as part of a product offered by an Authorized Party, End-User's agreement with such Authorized Party will control any renewal.

8. MODIFICATIONS; DERIVATIVE WORKS; DISTRIBUTION

In the event Proofpoint creates a modification or any derivative work of a Subscription, the use, reproduction and distribution of such modifications or derivative works shall be governed by the terms and conditions of this Agreement.

9. INABILITY TO COMPLY DUE TO STATUTE OR REGULATION

If it is impossible for End-User to comply with any of the terms of this Agreement with respect to some or all of the Subscriptions due to statute, judicial order, or regulation then End-User must: (a) comply with the terms of this Agreement to the maximum extent possible; and (b) immediately provide written notice to Proofpoint describing these limitations. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

10. DISCLAIMER OF WARRANTY

THE SUBSCRIPTIONS ARE PROVIDED UNDER THIS AGREEMENT ON AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE SUBSCRIPTIONS ARE FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGING. THE ENTIRE RISK AS TO PERFORMANCE OF THE SUBSCRIPTIONS IS WITH END-USER. SHOULD THE SUBSCRIPTIONS PROVE DEFECTIVE IN ANY RESPECT, END-USER (NOT PROOFPOINT) ASSUMES THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS AGREEMENT. IN NO EVENT WILL STATUTORY WARRANTIES (IF ANY) APPLY UNLESS REQUIRED BY STATUTE NOTWITHSTANDING EXCLUSION BY CONTRACT.

11. LIMITATION OF LIABILITY

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL PROOFPOINT BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR

Copyright © 2024 Greycortex

CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, SECURITY BREACHES OR FAILURES, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATIONS. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY. NOTWITHSTANDING ANY OTHER PROVISION OF THIS AGREEMENT, PROOFPOINT'S LIABILITY UNDER THIS AGREEMENT, WHETHER UNDER CONTRACT LAW, TORT LAW, WARRANTY OR OTHERWISE, SHALL BE LIMITED TO DIRECT DAMAGES NOT TO EXCEED THE AMOUNTS ACTUALLY RECEIVED BY PROOFPOINT UNDER THIS AGREEMENT FOR THE SUBSCRIPTION TERM IMMEDIATELY PRECEDING THE ACTION GIVING RISE TO THE CLAIM.

12. TERMINATION

This Agreement and the rights granted hereunder will terminate automatically if End-User breaches any of the terms herein and fails to cure such breach within thirty (30) days of becoming aware of the breach. Provisions, which, by their nature, must remain in effect, beyond the termination of this Agreement, shall survive.

13. MONITORING

Proofpoint reserves the right to monitor access to its servers, by an End-User or otherwise. Such monitoring by Proofpoint will include, but not be limited to, identification of the IP addresses that access and download Subscriptions from a Proofpoint server. By entering into this Agreement and using the Subscriptions, End-User agrees to the transfer of such information to Proofpoint's offices in the United States.

14. VERSIONS OF THE AGREEMENT

Proofpoint may publish revised and/or new versions of this Agreement from time to time. Any such revised or new version of the Agreement shall be binding and effective upon the earlier of (i) the date that is thirty (30) days after posting of the revised or new version of the Agreement on Proofpoint's website or (ii) the date that Proofpoint provides notice to End-User of the amended Agreement. End-User may immediately terminate the Agreement upon written notice to Proofpoint ifv

Copyright © 2024 Greycortex

a change is unacceptable to End-User. End-User's continued use of the Subscriptions indicates End-User's agreement to such revised or new versions of the Agreement.

15. NOTICES

Notices given under this Agreement shall be in writing and sent by facsimile, email, or by first class mail or equivalent. Proofpoint shall direct notice to End-User at the facsimile number, email address, or physical mailing address provided in the registration process. End-User shall direct notice to Proofpoint at the following address:

Proofpoint, Inc.

892 Ross Drive

Sunnyvale, CA 94089

USA

Fax: (408) 850-4022

Either party may change its notice contact information at any time by giving notice of the new contact information as provided in this section.

16. ASSIGNMENT

End-User may not assign this Agreement without Proofpoint's prior written consent. Proofpoint may assign its rights and obligations under this Agreement without End-User's consent.

17. SEVERABILITY; GOVERNING LAW; VENUE; CONSTRUCTION

This Agreement represents the complete agreement concerning the subject matter hereof. No terms stated in any purchase order shall be incorporated into this Agreement and all such terms shall be void and of no effect. If any provision of this Agreement is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This Agreement shall be governed by California law, excluding its conflict-of-law provisions. Any litigation relating to this Agreement shall be subject to the jurisdiction of the state and federal Courts serving Santa Clara County, California. In the event Proofpoint files suit to enforce the terms of this Agreement, Proofpoint shall be entitled to recover all of its costs, including without limitation, court costs and attorneys' fees and expenses. End-User hereby submits to jurisdiction and venue in such courts. The

Copyright © 2024 Greycortex

application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation, which provides that the language of a contract shall be construed against the drafter, shall not apply to this Agreement. Headings and section references are used for reference only and shall not be used to define, limit or describe such section. The pronouns used herein shall include, where appropriate, either gender or both, singular and plural.

18. COUNTERPARTS

This Data License Agreement may be executed by facsimile, electronic signature and in counterparts.

Copyright © 2024 Greycortex

6.3 Support

Information

Technical support is provided by the manufacturer via HELPDESK.

Business information is available at PARTNER PORTAL or by email at info@greycortex.com.

Maintenance

GREYCORTEX is constantly developing and enhancing MENDEL's solutions by adding new features and bug fixes in software. These changes are continually provided to registered customers as part of Maintenance Support purchased, including updating detection samples from external and internal resources.

Copyright © 2024 Greycortex

6.4 Helpdesk

If a technical issue regarding your Mendel appliance occurs, our Helpdesk is a place for you to open a ticket which will be then resolved by our Support department. You can also use it if you find yourself with a generally or technically oriented question regarding your Mendel appliance, which needs to be answered by our Support or R&D department. There is also an option to define a feature request.

6.4.1 How can I obtain access to the Helpdesk?

Contact one of our sales representatives through email. After you obtain credentials to your personal account you can log in to Helpdesk.

6.4.2 The types of tickets

Technical Issue

Create this type of ticket if you have ongoing or long term technical issue with your appliance.

Example

- No events in the GUI.
- Can not upgrade to newer version.
- The GUI and/or ssh is unreachable.
- We had a hardware failure, how should we proceed?

General Inquiry / Technical Question

Create this type of ticket ONLY if you can´t find the answer to your question in MENDEL User Documentation and have:

- Request for support account activation
- Technical question
- General question

Example

- I've installed Mendel, but I can't see any events in the GUI, why is that?
- Aside from DHCP, what other methods does MENDEL use to get the hostname?

Feature Request

Create this type of ticket if you have an idea for a new feature that will help us to improve the product.

Example

- One click upgrade for the sensor-collector type of deployment.
- Time-limited false positives.

6.4.3 How to open a new ticket

- Visit our Helpdesk and select .
- This will take you to the login screen where you can log in with credentials provided by our sales representative.

Copyright © 2024 Greycortex

- Select one of the available help topics.
- Once you select the topic a form will appear.

Example

Technical issue
General Inquiry / Technical Question
Feature Request

6.4.4 Solution and feedback

We will answer your inquiry as quickly as we can based on the information you have provided. We may need to ask for some follow-up information depending on the nature of the issue. As soon as the inquiry is processed, you will be notified by our helpdesk via email, which will contain direct link to your ticket.

Copyright © 2024 Greycortex

6.5 Release notes

6.5.1 4.3.0

MAIN FEATURES

OT METRICS IN PRODUCTION - STAGE 1

With OT Metrics, you gain real-time insights into key metrics that drive your operational success. Monitor critical parameters with precision and accuracy. Our intuitive visualisation provides a holistic view of your OT environment, allowing you to identify alarming trends, and potential security issues and make informed decisions. All this is for the most used protocols such as Modbus, MMS, SNMP, BACnet, and IEC-104 with the potential to enlarge the OT protocol coverage in future versions upon your interest and request.

BUILT UP OT CAPABILITIES

- New OT parsers for Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol (LLDP), GE-SRTP, and Mitsubishi MELSEC
- Enhanced SNMP parser to version 3
- User configuration of OT protocol ports using variables
- Specific IDS detection rules for protocols in OT environment:
- Energy Distribution
- Industrial Manufacturing and Processing
- Facility management networks
- Up to 100 small sensors managed by one collector
- Our latest product for quick asset inventory is the GREYCORTEX Microsensor. This tool allows our partners to quickly make an overview of a network with all it’s OT-assets, without the deployment complexity of a full GREYCORTEX Mendel-product
- OT POST deployment guides for: Energy Distribution and Industrial Manufacturing & Processing

PCAP RECORDER 2.0 WITH MORE RELIABLE DATA CAPTURE AND BETTER UX

We are introducing enhanced reliability and UX in data capture. Our new PCAP recorder now leverages:

- primary keys from the database for flawless PCAP naming and compatibility across languages and charsets
- innovative LST_ flags, database signalling ensures real-time tracking of PCAP changes for each sensor
- recording all packets with matching filtering options, even across multiple PCAPs, ensures comprehensive data capture
- recording criteria to include Ethernet type, MPLS, tunnels, protocol, and source port
- all PCAP definitions and database interactions are seamlessly managed through specialised database functions, guaranteeing secure and efficient data operations.

UX improvements related to optimised PCAP recording and storage management:

- The file rotation system supports both one-time pass recordings and continuous recording with FIFO-based erasing of the oldest files.
- PCAP rotation across the entire storage partition, giving you the flexibility to protect specific files from deletion and choose between FIFO rotation or halting recording when the volume is full.
- Information about overall disk volume and available space

NETFLOW ENHANCEMENTS

- Maximum processing efficiency by sharing output threads across multiple NetFlow pipelines to ensure optimal resource utilization, reducing overhead and enhancing throughput for faster, more reliable data processing.
- Significant memory savings with optimised ParsedFlow class to reduce memory consumption.
- Dedicated flow enrichment pipeline component, designed to operate independently for enhanced modularity. This allows for more flexible deployment, easier maintenance, and the ability to scale or update components without impacting the entire system.

Copyright © 2024 Greycortex

- Overall improvement accelerates data processing, enhances code readability, and reduces CPU load, providing a smoother and more efficient operation.
- And for all these changes there is a new NetFlow configuration in the UI to manage NetFlow settings across multiple pipelines easily.

OTHER FEATURES

UI CHANGES FOR BETTER UX

- Visual changes in CEM to distinguish it from other types of appliances
- Whisperer in Main filter for Applications data type
- Policies, Subnets, and Host in Settings with reworked tables and filters
- An option to manage users by multiple administrators (user managers)
- Filtering by network interfaces and NetFlow sources
- Improved import function for subnet's settings and added bulk operations
- The incident filter was adjusted to work with False positives and Long Term labels as attributes
- Default (system) subnets in settings reworked to be more comprehensible
- Plugins in Settings with code visualisation
- Group operations in Subnets and Hosts
- PDF reports in the Czech language
- Training forbidden ports on subnets
- Tag rendering

API EXTENSION TO PROVIDE SUPPORT FOR

- Incidents
- Custom categories in Signatures
- Saved views in Network Analysis
- Audit (user activity) logs

MISCELLANEOUS

- Improved Unte signatures
- Matching Mac addresses for vendor recognition
- New predicate when tag does not exist on host egration with Zabbix monitoring platform to enrich host/device information
- Cloning Mendel instances from templates in VMware for automatic multiple locations installation using Ansible

ENHANCEMENTS

- Logs in LEEF and CEF format support the substitution of empty fields
- Added UnTE filtering by subnet tag
- Jitter and Delay metrics + RTP header

FIXED ISSUES

- Fixed issues with parent subnet assignment during subnet import
- Fixed issues with Mendel upgrade on AWS
- Fixed issues with traffic calculation on Napatech cards
- Fixed issues with numbering Napatech cards interfaces
- Fixed issues with greycortex.mib specification
- Fixed issues with adding custom domain to Threat intelligence component

Copyright © 2024 Greycortex

- Fixed issues with DNS traffic on high ports
- Fixed issues with HART-IP parser crashing on input with nonzero length
- Fixed issues with rendering some values of HTTP2 protocol on flows
- Fixed issues with memory when specific decrypted data is processed
- Fixed issues with reporting application data from flows
- Fixed issues with reporting protocols specific events for Modbus
- Fixed issues with SMTP undefined values for some response codes
- Fixed issues with huge number of entries in IDS variables
- Fixed issues with detection of possible TOR SSL traffic on invalid port
- Fixed issues with reporting incorrect interface traffic on Broadcom cards
- Fixed issues with support membership group in API
- Fixed issues with ignore event button on NBA events
- Fixed issues with sending events via syslog protocol
- Fixed issues with adding routes in UI
- Fixed issues with processing offline Pcaps
- Fixed issues with service port validation in Data Capture
- Fixed issues with timing of IDS events
- Fixed issues with removing parental tag
- Fixed issues with adding tag to non-existing host
- Fixed issues with cancelling queries in Network analysis module
- Fixed issues with incident links to images on systems using proxy
- Fixed issues with exporting incidents to PDF
- Fixed issues with event count in Host inventory
- Fixed issues with Top events list for external host
- Fixed issues with Peers Subnet table displaying source subnet
- Fixed issues with Test button on LDAP authentication configuration
- Fixed issues with user identities when LDAP is not configured
- Fixed issues with using multiple LDAP servers
- Fixed issues with password when no authentication is used
- Fixed issues with changing Kerberos settings
- Fixed issues with comma in Certification request
- Fixed issues with reloading existing components on Dashboards

OFFICIAL MENDEL PRODUCT SUPPORT

With the release of version 4.3.0, full-service support is provided for versions 4.3.x and 4.2.x. Limited service support is provided for the previous version, 4.1.x. Versions 4.0.x and older are no longer supported. End-users with valid support and maintenance or an active software subscription are advised to upgrade to a supported version(s).

Copyright © 2024 Greycortex

6.5.2 4.2.0

MAIN FEATURES

PCAP PLAYBACK AND ANALYSIS

An extension to Mendel's core capabilities not only to analyse mirrored traffic in real-time but offers the same power to do it over earlier recorded data with the built-in PCAP recorder or captured by any other tool you have. Replayed PCAPs are processed and stored separately to any other data in Mendel to ensure the highest user experience during retrospective analysis with a maximum focus on data consistency and security.

XSOAR INTEGRATION

An integration pack that connects Mendel and XSOAR and enables native data ingestion. This integration streamlines the flow of data from Mendel directly into XSOAR, the security orchestration, automation, and response platform.

IDS RULE WIZARD

With this new step-by-step tool, you are able to tailor the rules in no time with significantly less knowledge the same way as a skilled professional, and can use the time saved more effectively.

NEW REPORTING MODULE

A redesigned way to create and visually process reports in PDF format. Using the new and more effective engine, you can select added options for every part of the report to make it exactly meet your expectations. Use any predefined dashboard and filter, with the ability to define the number of rows for table type content, comes as standard. This is the foundation for a broad and unified powerful reporting service, streamlining the process of making comprehensive reports from any data available in the product.

USER ACTIVITY LOG

Audit every step of a Mendel user or API queries to comply with legislative requirements or internal policies.

USER EVENT CATEGORIZATION BETA

A brand new option to configure user-accounts as IT or OT users. This categorizes events faster and more efficiently with an emphasis on different needs and the focus of IT and OT security and operational specialists. This allows both user-groups to collaborate via the same Mendel, while still having their own perspective.

OT METRICS PROTOTYPE FOR BACNET EXPERIMENTAL

The ability to visualize parsed values or metrics, derived from your OT protocols network flows within Mendel, offers you an unprecedented window into understanding your industrial equipment’s behaviour. This exclusive experimental feature brings a new dimension to monitoring and comprehension. Contact us to explore this non-public offering and gain a deeper understanding of your industrial systems.

OTHER FEATURES

UI changes for better UX

- Detection part in settings (a merged view for system and customer rules + TI)
- Group operations in the detection part
- LACP settings applicable via the UI
- The MITRE page as a regular dashboard
- Import subnets with tags

API extension to provide support over data related to:

- hosts
- users (persons)
- reports

Copyright © 2024 Greycortex

Calculated delay and jitter from VoIP time metrics File extraction from analyzed traffic for external analysis in the Sandbox BETA Processing AWS FlowLogs protocol in the NetFlow module EXPERIMENTAL

ENHANCEMENTS

Network capture parsers and enhancements

- Parsers (PostgreSQL, Bittorrent , IKE1/IKE2, QUIC and HTTP2)
- VOIP and delay metrics, improved ART/RTT
- MPLS and VLAN tags filtering from the recorded PCAPs

Improved OT protocol parsers

- BACnet
- OPC-UA
- S7

NetFlow export enhancements

- export version 5/9
- application data
- time metrics

Extended UnTE engine to support event correlations BETA

FIXED ISSUES with

- Fixed issue with Automatic Subnet Discovery when feature is disabled
- Fixed issue with filtering ICMP codes in Network Analysis module
- Fixed issue with displaying graph legend in Network Analysis module
- Fixed issue with filtering events from incidents with deprecated categories
- Fixed issue with applying filter in Performance dashboard
- Fixed issue with rendering Identity provider configuration page
- Fixed issue with detection of service in bypassed flows
- Fixed issue with passing packets from bond0 interface to processing
- Fixed issue with filtering by tags in Network Analysis module
- Fixed issue with failsafe mode in unstable environment to eliminate gaps in data
- Fixed issue with generating encrypted PDF attachments in PGP signed emails
- Fixed issue with rights in incidents for user with sensor restriction in multitenant environment
- Fixed issue with with filtering data with asn code
- Fixed issue with Fortigate plugin for Fortigate versions higher than 7.0.1
- Fixed issue with NTP configuration on sensor
- Fixed issue with mounting NTFS storage in Shell
- Fixed issue with displaying data in graphs user with subnet restrictions
- Fixed issue with content-type in export email
- Fixed issue with Purdue tags structure
- Fixed issue with invalid service recognition on port 53
- Fixed issue with corrupted or malformed DNS message results in flow

Copyright © 2024 Greycortex

- Fixed issue with fulltext filter whisperer
- Fixed issue with case sensitivity in filter
- Fixed issue with sorting flows by ART metrics
- Fixed issue with Analysis module (blacklisted IPs not colored)
- Fixed issue with configuration of network stack and DB restart
- Fixed issue with categories for Status monitor events
- Fixed issue with invalid ASN filtering in Analysis module
- Fixed issue with filtering IP addresses from Outlier events
- Fixed issue with exporting Top Tags dashboard
- Fixed issue with reporting MISP connector errors
- Fixed issue with Bacnet parser
- Fixed issue with removal certificate chain which are used
- Fixed issue with rendering Autocomplete menu behind dialog

OFFICIAL MENDEL PRODUCT SUPPORT

With the release of version 4.2.0, full-service support is provided for versions 4.2.x and 4.1.x. Limited service support is provided for the previous version, 4.0.x. Versions 3.9.x and older are no longer supported. End-users with valid support and maintenance or an active software subscription are advised to upgrade to a supported version(s).

Copyright © 2024 Greycortex

6.5.3 4.1.0

MAIN FEATURES

NEW LOOK

The new user interface (UI) is a new redesigned look. It has been created to reduce visual complexity, provide easy access to essential data and progressively disclose complex functionality as needed – resulting in a cleaner look and feel.

EXTENDED INTEGRATION FOR (NOT ONLY) XDR COVERAGE

EDR INTEGRATION / STAGE 1

- Generic (vendor plugin based) solution to cooperate with different brands of endpoint clients (APIs)
- Extended host/device/endpoint information in Mendel UI
- Events from EDR visible in Mendel´s UI with hyperlinks to EDR console
- Response capability triggered (manually/automatically) by Mendel and executed by EDR client

CISCO ACI/APIC INTEGRATION

- Extended visibility with applied definitions of the software defined infrastructure/network (e.g. host/subnet tags corresponding to ESG/EPG information).

API VERSION 2

IMPORTANT: With the new version 4.1.0 the product (Mendel) is using higher version of the API = v2. This version comes with unified language and new set of keywords.

Current API coverage:

- Events
- IDS and log processing signatures, IDS variables
- Data captures (a direct connection into the database where all captured network data is stored)
- External visualization or analysis of pre-processed data (network data stored in ME database)
- False positive management
- UnTE signatures (4.1.0)
- Flows records (4.1.0)
- Blacklists based on IP addresses (including MISP)
- Malicious Files
- Malicious domains

ADVANCED FILTERING WITH POWERFUL QUERIES

Master the data outputs with queries based on rich syntax with the new fulltext filter.

OTHER FEATURES

- New OT parsers for BACnet and Profinet on DCE/RPC
- LACP bonding on Mendel interfaces
- IoC validation in history using API
- Processing ERSPAN communication (decapsulation on Mendel interfaces)

ENHANCEMENTS

- System profiles are now configurable by users with administrator role

Copyright © 2024 Greycortex

- Upgraded Napatech drivers for wider compatibility
- Newly discovered subnets are now reported in events
- Optimization of hosts/services tab
- Add interface to Asset Survey Configuration
- Single page for AD + ISE + Identity logs settings

FIXED ISSUES with

- Fixed issues with Syslog structure in RFC3164/RFC5424
- Fixed issues with automatic rule setup for port 514 (if log processing is enabled)
- Fixed issues with Log processing module stability
- Fixed issues with validation of IP address captured by Log processing module
- Fixed issues with flipped IDS events addresses compared to TCP flags
- Fixed issues with pairing DHCP transactions in DHCP protocol
- Fixed issues with handling NaN values in parsed data
- Fixed issues with host count calculation in Network Inventory
- Fixed issues with choosing proper destination IP address in Proxy pairing module
- Fixed issues with policy inheritance for new discovered subnets
- Fixed issues with evaluation Subnet inclusion in false positives
- Fixed issues with detection source subnet in New Host detector
- Fixed issues with nondeterministic detection of domain controllers
- Fixed issues with various Unte validator errors
- Fixed issues with assigning user defined tags
- Fixed issues with value editing of hosts in settings
- Fixed issues with Disabling/Enabling data capture in UI
- Fixed issues with search False positive filters in settings
- Fixed issues with data capture filter defined by source and destination ip address
- Fixed issues with resetting filter in Incident module
- Fixed issues with filter name saving in Filter manager
- Fixed issues with predefined filters
- Fixed issues with paginator in Network Analysis module
- Fixed issues with data export for selected number of rows and page number
- Fixed issues with event category component filtering by category
- Fixed issues with configuration ssh access in Setup wizard in UI
- Fixed issues with redirection on Upgrade page
- Fixed issues with restore to factory settings
- Fixed issues with detection of traffic on network interfaces in Installer
- Fixed issues with displaying LDAP authentication information
- Fixed issues with detection of SMB base backup
- Fixed issues with user permission on Export-Data interval
- Fixed issues with removing users with existing exports
- Fixed issues of displaying Top Events on Host dialog

Copyright © 2024 Greycortex

• Fixed displaying Napatech interface in flows

OFFICIAL MENDEL PRODUCT SUPPORT

With the release of version 4.1.0, full-service support will be provided for versions 4.1.x and 4.0.x. Limited service support is provided for the previous version 3.9.x. Versions 3.8.x and older are no longer supported. End-users with valid support and maintenance or an active software subscription are advised to upgrade to a supported version(s).

RELEASE NOTES

For speeding up and improving system updates and data updates, we are preparing a new update/upgrade server. To be prepared for further use of the new server, please allow IP address 46.28.107.158 on your firewall. Current IP 89.221.218.75 (update.greycortex.com) needs to be preserved as well.

Copyright © 2024 Greycortex

6.5.4 4.0.0

MAIN FEATURES

NETWORK INVENTORY

The Visibility and Detection layers have been merged into one clear view. Have a complete overview of your network infrastructure with the added value of subnet and host detailed information colored with calculated risk and a security view. See the data as a sortable table or a scalable graphical interpretation.

API/STAGE 3

Even greater integration potential with the new ability to connect Mendel with external information sources or provide processed data to recipients (SIEMs, etc.) for further processing. #restful

Current API coverage:

- Events (4.0.0)
- IDS and log processing signatures, IDS variables (4.0.0)
- Malicious domains (4.0.0)
- Data captures (a direct connection into the database where all captured network data is stored)
- False positive management
- Blacklists based on IP addresses (including MISP)
- Malicious Files

INTEGRATION

Utilize threat intelligence platforms and enrich your security feeds. #MISP

Extract identity information about logged users from logs from firewalls or domain controllers and pair them with hosts. #Radius

Community ID Flow Hashing: Blue-team security specialists often use a variety of monitoring applications to correlate network flows from each of them. It's often desirable to pivot quickly from one dataset to another. In Mendel 4.0, GREYCORTEX introduces the industry standard Community ID in its flows, decreasing the workload of such dataset "joins" for blue-team members working with Mendel.

- Security Tooling Supporting Community ID: Arkime, Elasticsearch, GREYCORTEX Mendel (v4.0+), Security Onion, VAST, Wireshark, Zeek and many others.

ENHANCEMENTS

Scalable NetFlow processing

- Up to 50 Gb of origin traffic
- Up to one thousand NetFlow sources
- Stores HTTP, TLS and DNS fields from IPFix
- Detects blacklisted IPs in traffic
- Extracts performance metrics and parameters (e.g., incoming interface/port)
- Supports performance profiles to optimize deployment and processing
- Receivies and processes NetFlow on multiple network interfaces of a single machine

Estimatated source of data anomalies on outliers with the capability to define false positives and reduce invalid detection.

Custom rules defined (added or imported) by the user for Log processing = Log processing is now in full production stage with UI support.

Failsafe mode with enhanced performance to solve infrastructure connection issues.

Flow identifier string to quickly and simply represent a given network flow. #Community ID

Copyright © 2024 Greycortex

User interface enhancements for better UX

- Top events and services in a (small) host info dialog box
- A copy button for IP and MAC in the host info dialog box
- Common services rendered into flows detailed information
- Handover of the Supervisor role to another account
- User defined data retention

Network capture enhancements * Signatures are capable of testing TCP flags in flows * GE-SRTP OT parser * Support for the latest network card drivers

FIXED ISSUES with

- Fixed issues with processing detection module with 10,000 subnets
- Fixed issues with AWS installation using LVM
- Fixed possible issue CVE-2022-42889
- Fixed issue with unknown protocol value
- Fixed performance issues with host services statistics
- Fixed performance issues with tagging subnets
- Fixed issues with inconsistent computation of network statistics per subnet
- Fixed issues with virtio_net driver
- Fixed issues in inconsistent data in Host by Risk overview dashboard
- Fixed issues with payload limited to 2,400 bytes
- Fixed issues with importing from CSV files
- Fixed issues with detection of new services on flows
- Fixed issue with TLS support in RDP flows
- Fixed issues with data restore from a corrupted backup
- Fixed issues with Checkpoint firewall configuration using TLS
- Fixed issue with ICMP detection
- Fixed issue with computation of event severity based on tags
- Fixed issues with UnTE detection on missing flow MAC addresses
- Fixed issue with filtering substring in service filter on flows
- Fixed issue with displaying IP protocol in data capture
- Fixed issue with threshold limits for Ethernet protocol in signatures
- Fixed issues with accessing flowvar/flowint in LUA signatures
- Fixed issue with validation of plugin parameters
- Fixed issue with storage compression with huge records
- Fixed issues with locks in startup phase of processing module
- Fixed issues with subnet level calculation
- Fixed issues with validation of categories and tags in filter
- Fixed issues with user permissions in tag management
- Fixed issues with filtering by tunnelled and operating system
- Fixed issues with telemetry reporting
- Fixed issues with logging errors from web interface

Copyright © 2024 Greycortex

OFFICIAL MENDEL PRODUCT SUPPORT

With the release of version 4.0.0, full-service support will be provided for versions 4.0.x and 3.9.x. Limited service support is provided for the previous version 3.8.x. Versions 3.7.x and older are no longer supported. End-users with valid support and maintenance or an active software subscription are advised to upgrade to a supported version(s).

Copyright © 2024 Greycortex

6.5.5 3.9.1

ENHANCEMENTS

Event visibility level store its configuration on user level (keep the last state before logout)

Improved performance and reliability for Failsafe mode

Improved subnet filtering by substring search in filter

FIXED ISSUES with

- Performance in network capture module
- Invalid licence during Sensor&Collector upgrade
- Default firewall configuration for asset discovery tool
- Checkpoint firewall rule policies
- Detecting TOR traffic by IDS signatures
- Resizing LVM storage on AWS
- Two or more DNS servers on management interface
- Empty subnet graph for subnets filtered by tag(s)
- User permissions
- SSL configuration for Fortigate firewall plugin
- Invalid CSV header in subnet import
- Malformed input for network parsers

OFFICIAL MENDEL PRODUCT SUPPORT

With the release of version 3.9.1, full-service support will be provided for versions 3.9.x and 3.8.x. Limited service support is provided for the previous versions, 3.7.x. Versions 3.6.x and older are no longer supported. End-users with valid support and maintenance or active SW subscription are advised to upgrade to a supported version(s).

Copyright © 2024 Greycortex

6.5.6 3.9.0

MAIN FEATURES

THREAT VISUALISATION

An interactive dashboard based on GREYCORTEX's and MITRE ATT&CK®'s knowledge in the field of cyber security. A clear and comprehensive view of specific threat models and methodologies in accordance with detection in supervised networks. For both IT and OT environments.

API / STAGE 2

Higher integration potential with a new ability to connect Mendel with external information sources or provide processed data to recipients (SIEMs, etc.) for further processing. #restful

Current API coverage:

- Data captures (a direct connection into the database where all captured network data is stored)
- False positives management
- Blacklists based on IP addresses (including MISP)
- Malicious Files

OFFLINE UPGRADE

Native capability to upgrade itself and the whole cascade (sensor via collector) via the user interface (UI) without the need of a live connection to the update server. Supports local (removable device) or remote (Samba) storage.

USER ACTIVITY LOG

To ensure the higher level of internal security of the monitoring tool itself, there is now the possibility to monitor the activity of users within their work.

ENHANCEMENTS

Network Analysis improvements for a better user experience

- Added the new metric "All" for the sum of the source and destination values
- Added new operators (is empty, is not empty, etc.)
- Added "Save as" and "To filter" buttons
- Import/export functionality for saved predefined views

Added support in the UI for downloading PCAPs from sensors

- Full support in the user interface for creating (conditional recording), storing and downloading raw traffic data from the sensor appliance

UnTE (Tagging) enhancements

- Dashboards based on the new tagging system
- Richer rule-making syntax
- Optimization for higher traffic speed (20 Gbps and higher)

Compressed data on the fast primary storage

- Capability to compress stored data on the fast (SSD) primary storage

  This feature is available only for a multipartition appliance setup.

Netflow processing for higher speeds and from a scaled number of netflow sources

Import/Export of (custom) IDS rules via the Settings page (UI)

Copyright © 2024 Greycortex

User defined time settings for data update distribution

Hardened DCOM remote protocol by authentication level with packet integrity + Windows 2022 support

Multidomain mapping in Suricata to handle proxy pairing in flows

Support for Intel E810 network cards

Support for Broadcom 10/25GbE cards

Changed visualization of event details for better comprehensibility

Geoip detection in IDS rules

FIXED ISSUES with

- Fixed issues with user rights restrictions in Host info dialog
- Fixed issues with flows lookup for new subnet events
- FIxed issues with detection of network interfaces in AWS
- Fixed AD module to work with Windows 2022
- Fixed duplicates in data exports using syslog protocol
- Fixed performance issues with file hash calculation on higher speeds
- Fixed issues with MTU configuration in system
- Fixed issues with parsing ARP protocols
- Fixed issues with saving existing filters
- Fixed issues with generation managerial and security reports on multiple sensors
- Fixed issues with file hash calculation
- Fixed issues with unsupported characters in data capture
- Fixed issues on Sensor & Collector page for CEM deployment
- Fixed issues with calculation appliance total traffic
- Fixed issues with configuration system interface
- Fixed issues with configuration ip address in shell
- Fixed issue with missing default subnets on new sensors
- Fixed issues with proxy configuration with non standard characters
- Fixed issues with time settings and timezones
- Fixed issues with creating false positives with defined time validity
- Fixed issues with rendering fields for application data
- Fixed issues with saving PGP signing and encryption configuration
- Fixed issues with memory consumption on tcp flow reassembly
- Fixed issues with UnTE signature settings dialog
- Fixed issues with importing CSV data from Excel
- Fixed issues with subnet tab on configuration wizard
- Fixed issue with exporting events in email with csv attachment
- Fixed issues with too long IDS variables
- Fixed issues with displaying incidents to users assigned to another sensor
- Fixed issues with detection up/down state of network interface
- Fixed priority configuration for IMAP and other services

Copyright © 2024 Greycortex

- Fixed detection of flow direction on TLS flows
- Fixed parsing on concurrent FTP data stream
- Fixed issues with NTP configuration and time synchronisation
- Fixed memory issues in OPC parser
- Fixed issues with host assignment into imported subnets
- Fixed issue with removing host tag in settings
- Fixed issues with validators limits in configuration for network capture module
- Fixed issues with displaying network data using name sorting
- Fixed issues with missing secondary DNS server in configuration
- FIxed filtering issues for various fields in Analysis module
- Fixed issues with incorrect subnets assigned to hosts
- Fixed issues with switched flow direction on FTP flows
- Fixed issues with long service restarts
- Fixed performance issues on outputs to QRadar using TCP Syslog

OFFICIAL MENDEL PRODUCT SUPPORT

With the release of version 3.9.0, full-service support will be provided for versions 3.9.x and 3.8.x. Limited service support is provided for the previous versions, 3.7.x. Versions 3.6.x and older are no longer supported. End-users with valid support and maintenance or active SW subscription are advised to upgrade to a supported version(s).

New Support SLA

PROFESSIONAL PLUS SUPPORT (SAME AS PROFESSIONAL SUPPORT WITH ADDITIONAL):

- Priority solution times – support tickets are responded to before waiting or open non-priority tickets.
- Priority software fix times – SW issues are fixed before other/non-priority tickets.
- Proactive Care – active system health monitoring by GREYCORTEX systems, including preventive identification of potential issues and problems (only performance statistics are analyzed, no specific network data and no sensitive data at all are collected)
- Priority handling of feature requests – via a special form sent to feature@greycortex.com
- On-demand consulting, training and security services provided by GREYCORTEX – 4 hours per quarter (16 hours per year) by default.
- Available for All-in-ones and Collectors only

Copyright © 2024 Greycortex

6.5.7 3.8.0

MAIN FEATURES

EXTENDED CLASSIFICATION OF DEVICES AND THEIR ROLES (TAGGING)

Dynamic visibility by tracking new activities or changes caused by devices communicating in your network. A completely new engine that brings a manual or automated way of tagging hosts or subnets through a system of user-defined rules with easy-to-understand syntax.

NEW CATEGORIZATION OF DETECTED SECURITY EVENTS

Mendel interprets all events captured in your network with more clarity thanks to the MITRE ATT&CK® tactics and techniques, Proofpoint rules and Top events (most relevant events at the top).

ANALYSIS MODULE

Deeper and more advanced data analysis thanks to the redesigned Analysis module. You can define any view over the processed and stored data using attributes, metrics and other variables.

CLOUD SUPPORT

New ability for Partners or Customers to deploy Mendel into the Cloud environment without support by utilization of our pre-prepared image ISO file. AWS is fully supported (sensor, AiO, collector). MS Azure and Google Cloud ready (support for collector deployment).

SYSTEM PROFILING FOR SPEEDS UP TO 100GB

Mendel now officially supports 4x10Gbit on Intel X710 network cards, 10/25/40/100 Gbit on Napatech network cards and experimentally Nvidia Mellanox network cards

SUBNET DISCOVERY

Simplify the post deployment process by automatically identifying subnets in the monitored infrastructure.

ENHANCED DASHBOARD SETTINGS

Added ability to import/export, clone and move dashboards to a different layout.

Added text field with wysiwyg editor and addition of text-only dashboard component.

LOG PROCESSING - BETA

Ability to process received logs and generate security events from them by semi-passive approach (logs received by Mendel on specified port).

OT/ICS/SCADA FEATURES

- Asset Discovery as a Mendel feature or standalone application BETA
- CVE matching for vulnerable assets BETA
- Implementation of OT/ICS parsers based on Siemens S7 and OPC protocols
- Implementation of KNX and BACnet discovery into Asset Discovery tool
- Added detection of DICOM and KNX protocol
- UI tab displaying data from industrial networks renamed to OT metrics (former Scada)

ENHANCEMENTS

- Optimized database processing for better UI response
- Added description of common service in service info dialog box
- Added interface for IDS variables adjustment in UI
- Added option to export event alerts in separate emails

Copyright © 2024 Greycortex

- Added response plugin to send email about the detected response request
- Added option to generate application data in IPFIX flow export
- Added option to insert subnet to multiple sensors
- Added signature for detection SMB broadcast request
- Added support for new arguments in firewall plugins
- Added ability to import csv files for importing subnets
- Added option to merge multiple netflow sources
- Added capability to process IPv4 over IPv4 tunnel
- Added support for parsing GENEVE Protocol
- Added view over application data for application IDS events
- Added option to merge span traffic from multiple interfaces to process separated RX and TX
- Enhanced validation and error reporting for new IDS signatures
- Enhanced local service detection on flows with errors
- Enhanced Checkpoint firewall plugin
- Enhanced scan detector for ARP scans
- Improved Netflow processing to scale to higher speeds
- Improved filtering in UI (autocomplete, port number for undefined services and more)
- Improved scan detection to ignore common antivirus scans
- Improved service descriptions
- Improved reporting events to external systems on multiple sensors
- All dialog windows are now movable

FIXED ISSUES with

- Activation certificate bundles for UI
- ASN assignment for proxy traffic
- Assigning host IP and name to filter
- COAP parser
- Detection already expired certificates
- Detection of lost hosts on incomplete flows
- Exporting data into csv file
- Failed network configuration
- Filtering custom subnet ranges
- Filtering numbers in flow application data
- Filtering on False Positives
- Hostnames for ICMP flows
- HTTP parser to detect correct HTTP methods
- Changing DNS configuration via shell
- Inserting a big number of custom IDS signatures
- Inserting a big number of new subnets on huge networks
- Invalid category in new user IDS signature
- Invalid characters in SNMP configuration

Copyright © 2024 Greycortex

- License validation in installer
- Multiple sensor version info after upgrade
- NTP validation in wizard
- Pairing requests and responses in TCP flows
- Parsing MODBUS data on reversed flows
- Performance issues with data capture on higher speeds
- Processing appliance telemetry
- Processing broken kerberos flow
- Processing parameters in LDAP query
- Redirection from top countries dashboard
- Search in application data on flow
- Sending appliance log via proxy
- Sending sensor log over collector
- Slow shell on slow machines
- Spaces in hostname on host settings
- Subnet autocomplete in filter
- Updating information on new management interface
- Using invalid dns servers for updating dns cache
- Using restricted TLS on SMTP
- Validation custom IDS signatures

OFFICIAL MENDEL PRODUCT SUPPORT

With the release of version 3.8.0 full-service support will be provided for the versions 3.8.x and 3.7.x. Limited service support is provided for the previous version 3.6.x. Versions 3.5.x and older are no longer supported. End-users with valid support and maintenance or active SW subscription are advised to upgrade to a supported version(s).

IMPORTANT

Please note that upgrading to version 3.8.0 will replace the system kernel and result in a reboot of the appliance. We recommend having direct or remote access to the appliance in order to be able to restart it if necessary.

If you are using the standalone Greycortex upgrade proxy application as a replacement for access to the update server, then you need to upgrade this proxy solution to a newer version first, to ensure that the Mendel product upgrade (3.8.0) is available.

Copyright © 2024 Greycortex

6.5.8 3.7.0

MAIN FEATURES

CISCO ISE USER IDENTITY INTEGRATION AND INCIDENT RESPONSE

When MENDEL has no direct access to AD/LDAP server or with limited permissions then user identity could be provided via integration with CISCO Identity Service Engine (ISE).

CISCO FIREPOWER INCIDENT RESPONSE

For situations where it is necessary to respond to emerging threats, we will ensure appropriate steps through integration with CISCO network elements. If this is unavoidable, you can block endpoint communication, isolate part of the network, etc.

SNMP APPLIANCE MONITORING & SNMP TRAP

With incorporation of SNMP agent and trap functionality you are able to oversee MENDEL appliances with your current infrastructure monitoring solution.

UPGRADE MANAGEMENT OVER APPLIANCES

Upgrade the whole Mendel deployment through a single point = collector´s UI. Choose either “One click” multi upgrade or upgrade each sensor individually. Upgrade is performed by two step method, to keep sensor running for maximum time and shorten the maintenance time.

AWS, MS AZURE AND GOOGLE CLOUD DEPLOYABILITY

Amazon Web Services, Microsoft Azure and Google Cloud are now supported for deployment of Collector or Central Event Management (CEM). Early access - GREYCORTEX support required.

UTILIZATION OF HIGH-SPEED DISKS WITHIN MULTITIER STORAGE AND OPTIMIZED DATABASE QUERIES

Use your fast disks not only for the operation of the system itself, but also for a much faster response of the user interface when displaying the „hot“ data and views of them. If your deployment does not have multi-tier storage with fast disks, we still bring you a faster response in the GUI by optimizing the database queries.

FALSE POSITIVES FOR LIMITED TIME PERIOD

Hide events only for the time that is relevant and related to the maintenance of your infrastructure, tests, etc. Apply false positives with specific time frame and/or recurrence.

CONDITIONAL PCAP RECORDING

Data captures can be triggered on-demand or by specified conditions (user-defined & event-based).

OT/ICS/SCADA FEATURES

- Asset Discovery
- Parsing MQTT, COAP and Profinet protocols
- Detection of LoRaWAN protocol
- Detection anomaly in IDS on data level
- New IDS rules for common OT/ICS protocols

ENHANCEMENTS

- Processing netflow data with NAT information
- Switch flow errors from flags to real calculation
- Connect Mendel sensor to secondary collector (HA)
- Deactivate inactive Sensor on Collector
- User Documentation available via GUI

Copyright © 2024 Greycortex

- Time validity of false positives
- Connect Mendel sensor to secondary collector (HA)
- Deactivate inactive Sensor on Collector

OTHER ENHANCEMENTS

- Process VMware ESXi NSX-T IPFIX format
- Add support for storing Suricata Variables in DB
- Enhance update server update data sources
- Show IDS signatures sources in license UI
- Semi-automated restoration of SMB backup
- IDS signatures using the detected application
- Display the logged-in user name on all pages
- False positive change Priority field Default text
- False positive not applicable into past by default
- Import new JA3 hash codes from ja3er.com
- Add description field into data exports
- Hide user from managerial/security reports and email
- Added assignee, reporter and date of last updated to Incident exports (PDF)
- Reworked Firewall settings with new location in UI
- Better explanation over data transfer between hosts in peers graph
- Evaluate and add IPv6 multicast address into monitored subnets
- System logs in mshell
- CAT tool for ME localization

FIXED ISSUES with

- Events: toplist in toggled row were not always loaded
- Discovery: New DNS Server reports hosts that didn´t respond to DNS requests
- SMB backup didn't work on sensors with capital letters in host name
- Incidents PDF report wasn´t generated due to invalid formatting
- Failed subnet restore
- Email exports (notifications) were not sent when 4h or 24h interval were used
- Failed export incidents to PDF with inherited tables
- Packets and data were doubled when processing NSEL netflow
- Flow detectors detected scans on NSEL firewall events
- Proxy settings without password in installer
- Name column couldn´t be sorted on hosts tab
- Possible memory leak in DNS parser
- New service detector didn´t report forbidden remote services in external network
- Offline installation with two usb disks
- Subnet negative priority didn´t affected event severities
- Filtering by substring in accounts and profiles

Copyright © 2024 Greycortex

- CSV export contained unstructured from/to dates
- ASN didn´t filter traffic on events correctly
- Segfault in DLMS parser
- Invalid icon names for Mac addresses
- Invalid handling of 0.0.0.0/0 subnet in negative filters
- Missing signature name at User False Positives table
- TDS parser memory allocation issues
- Performance profile detection failed on more than 256 GB of RAM
- Installer didn't support NVMe disk
- Table caption converted to paragraph text after saving in Incident description field
- Settings menu was not scrolling vertically in Firefox
- Interface filter did not affect number of flows
- Proxy settings allowed without password
- Severity filter was not applied when redirected from Top Events dashboard component
- SMTP flows did not show Application Layer data
- Invalid host tag removal in multi-sensor computation
- Double subnet in report info
- Mendel Shell: dhcp renew lead to shell error
- Shell login worked even after deactivation
- Tag events were without subnet icon
- Fixed bugs in plugin execution management

OTHER

Greycortex MENDEL helpdesk

OFFICIAL MENDEL PRODUCT SUPPORT

With release of version 3.7.0 full-service support will be provided for the versions 3.7.x and 3.6.x. Limited service support is provided for previous version 3.5.x. Versions 3.4.x and older are no longer supported, end-users with valid support and maintenance or active SW subscription can upgrade to the supported version(s).

Copyright © 2024 Greycortex

6.5.9 3.6.1

Enhancements

Improved Incident management

- Added incident label management (custom labels)
- Excluded false positive incidents by default
- Possibility to add watchers from incident page form
- Ability to change incident state
- Changed PDF report title
- Added time range into PDF report title page
- Added new items into incident overview header in PDF report
- Added ability for no-reinstall recovery after motherboard replacement on DELL server
- Faster representation of event queries in the lightbox

Removed user information from managerial/security reports and emails

Reworked firewall plugins compatibility with PaloAlto

Fixed Issues

- Fixed issues with data exports for 4 hour or 1 day intervals
- Fixed issues with Colorful text not exported to pdf
- Fixed issues with Incident PDF report doesn't export image
- Fixed issues with PDF export from new incidents fails under some conditions
- Fixed issues with inability to show first 50 incidents after you scroll through them
- Fixed issues with applying false positives into past on compressed data
- Fixed issues with ssh keys in Mikrotik plugin
- Fixed issues with using proxy settings in MENDEL installer
- Fixed issues with restarting application server after upgrade to latest version
- Fixed issues with installation of MENDEL in VirtualBox 6.1
- Fixed issues with Mismatched MAC address vendor and its icon
- Fixed issues with Saving icons for subnets not working
- Fixed issues with MMS configuration parser failure on specific data

Official MENDEL Product Support

Full-service support is provided for the versions 3.6.x and 3.5.x. Limited service support is provided for previous version 3.4.x. Versions 3.3.x and older are no longer supported, end-users with valid support and maintenance or active SW subscription can upgrade to the supported version(s).

Copyright © 2024 Greycortex

6.5.10 3.6.0

NOTICE: THIS VERSION BRINGS COMPRESSED STORAGE. THIS REQUIRES SOME DATA TRANSFORMATIONS THAT CAN BE PERFORMED IN THE BACKGROUND, BUT CAN CAUSE HISTORICAL DATA TO BE HIDDEN FOR A FEW HOURS DEPENDING ON DISK SPEED AND AMOUNT OF DATA. WHEN TRANSFORMATION IS FINISHED, THE APPLIANCE WILL DISPLAY ALL DATA.

Main Features

INCIDENT MANAGEMENT

Improves the ability to work with incidents, providing an easier way to define and describe security incidents and their management, and provide tools to export them into a final printable report suitable for presentation and reporting.

AUTOMATED THREAT RESPONSE

MENDEL allows the automatic blocking of traffic on firewalls or sending notifications to external systems through custom scripts. The execution of the script is based on predefined event filter plugins.

MULTITIER STORAGE FOR LONGER DATA RETENTION

A change in the way MENDEL stores data, with the goal of saving disk space and providing more effective storage for machines, with data retention over three months.

APPLICATION DETECTION

Enhanced flow detection, which can now recognize up to 4,000 commonly used applications to improve better visibility into specific services.

PARSING NEW RDP PROTOCOL AND ENHANCE EXISTING PARSERS

Many improvements have been implemented for SMB, SNMP, SIP, SMTP, Modbus(TCP/UDP), and many other protocols.

JA3S SUPPORT

Enhanced support for the JA3 fingerprint method, making it possible to fingerprint the entire cryptographic negotiation between a client and its server by combining JA3 + JA3S (server side fingerprinting).

ENHANCED DETECTION OF MALICIOUS ENCRYPTED COMMUNICATION

MENDEL is able to detect malicious TLS certificates, malicious clients, or servers using JA3 fingerprints.

Enhancements

PROCESSING CISCO ASA NETFLOW

Appliances are now able to process NetFlow data from Cisco ASA solutions in the NetFlow Secure Event Logging format.

CLONING MENDEL INSTANCES FROM VM TEMPLATES

Allows the cloning of existing MENDEL instances in a virtual environment for faster deployment without installation.

DB OPTIMIZATIONS AND MAINTENANCE PROCESSES

Optimization and improvement for faster processing in user interface response for large networks.

EXTENDED PROTOCOL DETECTION

Improvements in flow processing, with better flow direction for greater accuracy as well as support for dealing with asynchronous flows.

SAMBA BACKUP

Enhanced configuration and password handling of data backup to Samba share, including a better connection error handling and consistency check for more reliable backup in the event of a backup error.

Copyright © 2024 Greycortex

Other Enhancements

- Removed deprecated IDS signatures from Proofpoint
- Added a certificate valid date into TLS protocol
- Added the option to cancel data requests to most pages
- Added support for Citrix Xen virtualization platform
- Enhanced processing for false positives for external networks
- Improve the processing of network statistics on a huge number of subnets in the User Interface
- Enhanced reporting for Active Directory errors during processing logs
- Enhanced port filter with full text
- Added IPv4 Link-local subnet (APIPA) to default subnets
- Updated Dell hardware monitoring tools to the latest version
- Upgraded to latest Intel 10Gbits network cards drivers
- Removed deprecated blacklisted sources

Fixed Issues

- Fixed issues with computing negative subnet priorities in events
- Fixed issue with processing UTF-8 and UTF-16 data in MS-SQL parser
- Fixed issues with evaluating expired TLS certificates
- Fixed issues with tagging hosts with invalid service
- Fixed issues with processing UTF-8 data in network capture module
- Fixed issues with displaying the reputation filter after setting and applying
- Fixed issues with unnecessary restarting application server
- Fixed issues with storing application data for IPv6 flows in IPv4 tunnel
- Fixed issue with processing new service detection on multiple sensors
- Fixed issues with processing the source port in NetFlow
- Fixed issues with displaying specific flow applications data for various services
- Fixed issues in access rights while connecting sensor to collector
- Fixed issues with displaying the SCADA channels table
- Fixed issues with generating security reports in specific time intervals
- Fixed time synchronization issues in event graphs
- Fixed issues with calculating top host dashboards on multiple sensors
- Fixed issues with protocol filters using protocol name
- FIxed issues with deadlock in shell commands for service management
- Fixed escaping issues in false positive description in Events
- Fixed issues with saving filter in a shared group
- Fixed issues with invalid GPT header
- Fixed issues with attachments extensions in email reports
- Fixed issues with segfault in shell before proper configuration
- Fixed issues with filtering specific filters in the Peers graph
- Fixed issues with processing Application performance on the client side
- Fixed issues with processing DNS reply codes

Copyright © 2024 Greycortex

- Fixed issues with excluding ASN from filter in flows
- Fixed issues with renaming and enabling data export
- Fixed issues with applying false positives to status monitor events
- Fixed issues with processing specific hostnames in IPv6 flows
- Fixed filtering issues for events not reported in incidents
- Fixed false positive alerts due to lost packets
- Fixed issues with invalid detection of ENIP protocol
- Fixed issues with deleting users
- Fixed issues with calculation statistics in Network Analysis module

Official MENDEL Product Support

- Full-service support is provided for the new released version 3.6.0 and previous version 3.5.x.
- Limited service support is provided for previous version 3.4.x.
- Versions 3.3.x and older are no longer supported, end-users with valid support and maintenance or active SW subscription can upgrade to the supported version(s).

Copyright © 2024 Greycortex

6.5.11 3.5.0

NOTICE: THIS VERSION CONTAINS A NUMBER OF MAJOR CHANGES IN THE SYSTEM. TO ENSURE A SMOOTH UPGRADE PROCESS AND TO PROVIDE SUPPORT TO ALL OUR PARTNERS AND CUSTOMERS, WE WILL BE INTRODUCING THIS RELEASE GRADUALLY OVER THE NEXT 14 DAYS.

We recommend having direct or remote access to the appliance in order to be able to restart it if necessary.

New Features

DATA EXPORT INTO CSV FORMAT

MENDEL users can now export data regarding hosts, network, flows, and even incidents into csv. format for further processing and creation of new network data visualizations.

CENTRAL EVENT MANAGEMENT

For customers or partners with larger deployments, MENDEL offers the ability to connect appliances using a multi-level structure; consisting of sensors, collectors, and a Central Event Management console. This provides a more comprehensive overview of the full network.

VALIDATING SSL AND TLS CERTIFICATES

For encrypted communication, MENDEL detects expired or invalid SSL and TLS certificates and alerts the user.

ARP PROTOCOL PARSER

We have added the ability to parse the communications using the Address Resolution Protocol for even better processing of non-IP data.

Enhancements

OPERATING SYSTEM IDENTIFICATION USING L7 DATA

MENDEL is able to detect the operating system of the host more precisely, using an advanced data model based on Samba, DHCP, HTTP, SSH, and L3/L4 parameters, among others. Data is also presented within a new dashboard showing the top operating systems in the network for the chosen period.

FILTERING DATA BY ADDITIONAL VALUES

We added the option to filter by additional variables, including operating system, interface, application, and port range.

NEW PREDEFINED DASHBOARDS

We have provided two new dashboards: Risks and Statistics; for our users to quickly and easily review the situation in their network.

Additional Enhancements:

- Upgrade of system components
- Printer tagging
- Browser protocol parser
- Sensor-Collector management
- Enhanced TLS 1.3 protocol parser
- Extended host/subnet lease time
- Configurable display level
- Decoding QoS/DSCP
- System improvements
- Network capture module improvements
- GUI improvements
- Localization improvements

Copyright © 2024 Greycortex

- SCADA
- MMS protocol processing
- For the visualization of MMS protocol data and further analysis, we added MMS protocol processing.
- Asset resources management
- We added the ability to name, manage, and add new devices in the network.
- DLMS/COSEM protocol parser
- We added parsing for DLMS/COSEM, one of the most widely accepted international standards for utility meter data exchange.
- OMRON FINS protocol parser
- We added parsing for the OMRON FINS protocol, which can be used by a PLC program to transfer data and perform other services with a remote PLC connected on an ethernet network.

Fixed Issues

- Fixed issues with default subnet policy configuration
- Fixed issues with support for VLAN filtering in BPF filters
- Fixed issues with computing event severity based on multiple priorities
- Fixed issues with displaying NTT and ART in application data
- Fixed issues with a limited number of subnets
- Fixed issues with sending emails without valid configuration
- Fixed issues with broken SMB backup for slower network connection
- Fixed issues when searching for flows for events with ASN
- Fixed issues with pre-filled source/destination field in false positive dialog
- Fixed a number of issues with false positive creation, editing, and filtering
- Fixed issues with connecting a sensor to a collector through proxy
- Fixed issues with parsing TDS (MS-SQL) data
- Fixed issues with displaying the JA3 client hash on the server side
- Fixed issues with reporting flow errors for UDP flows
- Fixed issues with assigning operating systems to known proxy servers
- Fixed issues with generating URL links
- Fixed issues with missing IP addresses in email exports
- Fixed issues with changed custom certificates after upgrade
- Fixed issues with certificates for the update server
- Fixed issues with the validation of email addresses
- Fixed issues with creating incident dialogues
- Fixed issues with the delay in processing Active Directory logs
- Fixed issues with using shortened IPv6 addresses in the filter
- Fixed issues with filtering by sensor combined with additional parameters
- Fixed issues with filtering unreported events
- Fixed issues with country code in the filter
- Fixed issues with disabling system interfaces
- Fixed issues with detected hosts without any detected service
- Fixed issues with disabling system proxy
- Fixed issues with false detection of Telnet protocol

Copyright © 2024 Greycortex

As well as more improvements to user experience, system stability, and performance.

Copyright © 2024 Greycortex

6.5.12 3.4.0

New features

- Added detection of operating system on monitored network hosts using network metadata
- Added new network metrics TTL and QoS
- Added Application Performance Metrics and Monitoring on the transaction level
- Extended detection of non-IP protocols including support for IDS rules to detect Ethernet based protocols
- Added default rules for non-IP protocols including Wake-On-LAN, IEC61850 - GSE, Sercos III, LLDP, and Ethernet Powerlink
- Added support for parsing EtherCAT protocol
- Added plugin for integration with Palo Alto Firewalls
- Added configurable database backup to external Samba storage

Enhancements

- Improved managerial and security reports - consistent calculation of hosts, improved data ordering, fixed formatting and design issues
- Improved report translations for Japanese, Polish, Chinese, Korean, and Czech
- Improved Group Permissions for better support of multi-tenant environments
- Improved processing of non-IP flows
- Improved data representation for GOOSE protocol
- Improved parsing capabilities of Samba protocol
- Improved parsing capabilities of MMS protocol
- Improved LDAP configuration dialog
- Allowed login for usernames starting with numbers

Bug Fixes

- Fixed security issues with HTTPS certificate
- Fixed issues with Failsafe Mode recovery on slow connections
- Fixed issues with invalid Telnet protocol detection
- Fixed issues with assigning hostnames from DHCP
- Fixed displaying flows for Discovery events
- Fixed issues with detection time for long-term events
- Fixed an issue with IDS detection on decrypted data streams
- Fixed an issue with the lost host detector on multiple sensors
- Fixed issues with false positive detection on multiple sensors
- Fixed issues with time settings in SCADA time-series graphs
- Fixed issues with url links to System events
- Fixed issues with invalid characters in sensor proxy parameters
- Fixed timeout issues in the user interface
- Fixed issues with disk space calculation in the administration shell
- Fixed issues with detecting sensor outages
- Fixed issues with sensor settings after connecting to collector
- Fixed issues with changed gateway in system

Copyright © 2024 Greycortex

- Fixed issues with IP configuration during upgrade
- Fixed issue with DNS settings during installation

Copyright © 2024 Greycortex

GREYCORTEX Mendel Hardware Specs

Update November 1st, 2024

	AiO	AiO	AiO	AiO	AiO	AiO	AiO	AiO	AiO	AiO	AiO
	XS-25	**S-25**	**SA-25 (alternative)**	**M-25**	**MA-25**	**L-25**	**LE-25 (economic)**	**XL-25**	**2XL-25**	**3XL-25**	**2XL-24**	**L-24**
DELL Model	**R360**	**R360**	R660xs	**R660xs**	R360	**R6615**	R660xs	**R6615**	**R7615**	**R7625**	R760	R660xs
chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	16x2.5 chassis	16x2.5 chassis	16x2.5 chassis	8x2.5 chassis
CPU	1x Intel Xeon E-2436, 2.9 GHz, 6/12	1x Intel Xeon E-2478, 2.8 GHz, 8/16	1x Intel Xeon Silver 4510, 2.4 GHz, 12/24	1x Intel Xeon Silver 4510, 2.4 GHz, 12/24	1x Intel Xeon E-2478, 2.8 GHz, 8/16	1x AMD EPYC 9174F, 4.1 GHz, 16/32	1x Intel Xeon Silver 4416+, 2.0 GHz, 20/40	1x AMD EPYC 9274F, 4,05 GHz, 24/48	1x AMD EPYC 9374F, 3,85 GHz, 32/64	2x AMD EPYC 9374F, 3,85 GHz, 32/64	2x Intel Xeon Gold 5420+, 2.0GHz, 28/56	1x Intel Xeon Gold 5420+, 2.0GHz, 28/56
RAM	2x 16GB RAM DDR5	2x 32GB RAM DDR5	4x 16GB RAM DDR5	8x 16GB RAM DDR5	2x 32GB RAM DDR5	8x 32GB RAM DDR5	8x 32GB RAM DDR5	8x 64GB RAM DDR5	12x 64GB RAM DDR5	24x 64GB RAM DDR5	16x 64GB RAM DDR5	8x 32GB RAM DDR5
PERC	PERC H355	PERC H355	PERC H355	PERC H355	PERC H355	PERC H755N	PERC H755	PERC H755N	Dual PERC H755N	Dual PERC H755N	Dual PERC H755N	PERC H755N
Primary Storage	2x 960GB SSD SATA read-intensive	2x 1.92TB SSD SATA read-intensive	2x 1.92TB SSD SATA read-intensive	2x 3.84 SSD SATA read-intensive	2x 1.92TB SSD SATA read-intensive	2x 7.68 SSD NVMe Data Center	3x 3.84 SSD SATA read-intensive	4x 7.68 SSD NVMe Data Center	5x 7.68 SSD NVMe Data Center	7x 7.68 SSD NVMe Data Center	6x 7.68 SSD NVMe Datacenter	2x 7.68 SSD NVMe Datacenter
Primary RAID (default)	RAID 1 or 10	RAID 1 or 10	RAID 1 or 10	RAID 1 or 10 (upgrade possible)	RAID 1 or 10	RAID1 or RAID6 recommended for upgrade, any RAID possible	RAID1 or RAID6 recommended for upgrade, any RAID possible	RAID6, any RAID possible	RAID6, any RAID possible	RAID6, any RAID possible	RAID6, any RAID possible	RAID6, any RAID possible
Riser config	Riser config 2, 2 x8 LP	Riser config 2, 2 x8 LP	Riser Config 5, 1 x8 LP + 1 x16 LP	Riser Config 5, 1 x8 LP + 1 x16 LP	Riser config 2, 2 x8 LP	Riser Config 3, 2 x16 FH (Gen5)		Riser Config 3, 2 x16 FH (Gen5)	Riser Config 4, 2 x16 FH (Gen5) + 2 x8 FH + 2 x16 LP, Half Length	Riser Config 4, 2 x16 FH (Gen5) + 2 x8 FH + 2 x16 LP, Half Length	Riser Config 2, 2x8 FH Slots (Gen4), 4x8 FH Slots (Gen5), 2x16 LP Slots (Gen4)	Riser Config 5, 1 x8 LP + 1 x16 LP
OCP 3.0 NIC			Intel E810 Dual Port 10/25GbE SFP SFP28, OCP 3.0	Intel E810 Dual Port 10/25GbE SFP+/SFP28 OCP 3.0		Intel E810 Quad Port 10/25GbE SFP28 OCP 3.0	Intel E810-XXVDA4 Quad Port 10/25GbE SFP28 OCP 3.0	Intel E810 Quad Port 10/25GbE SFP28 OCP 3.0	Intel E810 Quad Port 10/25GbE SFP28 OCP 3.0	Intel E810 Quad Port 10/25GbE SFP28 OCP 3.0	Intel E810 Quad Port 10/25GbE OCP 3.0	Intel E810 Dual Port 10/25GbE OCP 3.0
PCI-E NIC	Broadcom 5719 Quad Port	Broadcom 5719 Quad Port	Broadcom 5719 Quad Port	Broadcom 5719 Quad Port	Broadcom 5719 Quad Port Intel E810 Dual Port 10/25GbE	Broadcom 5719 Quad Port	Broadcom 5719 Quad Port	Broadcom 5719 Quad Port				Broadcom 5719 Quad Port
PSU	600W dual PSU	600W dual PSU	1100W dual PSU	700W dual PSU	600W dual PSU	1100W dual PSU	700W dual PSU	1100W dual PSU	1100W dual PSU	1400W dual PSU	1100W dual PSU	700W dual PSU
iDRAC	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise
Rails	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management
Support	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD
Optional FPGA [1]				Napatech FPGA Options: 2x25/10GbE	Napatech FPGA Options: 2x25/10GbE	Napatech FPGA Options: - 4x10GbE or 8x10GbE - 2x25/10GbE or 4x25/10GbE - 2x40/100GbE	Napatech FPGA Options: 2x25/10GbE	Napatech FPGA Options: - 4x10GbE or 8x10GbE - 2x25/10GbE or 4x25/10GbE - 2x40/100GbE	Napatech FPGA Options: - 4x10GbE or 8x10GbE - 2x25/10GbE or 4x25/10GbE - 2x40/100GbE	Napatech FPGA Options: - 4x10GbE or 8x10GbE - 2x25/10GbE or 4x25/10GbE - 2x40/100GbE	Napatech FPGA Options: - 4x10GbE - 2x25/10GbE or 4x25/10GbE - 2x40/100GbE	Napatech FPGA Options: 2x25/10GbE
Maximum monitoring ports w/o Napatech FPGA	Onboard (2x1GbE metalic) & iDRAC **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP+ only SR SFP modules possible (no LR) - Intel X710 Dual or Quad Port 10GbE Base-T copper	Onboard (2x1GbE metalic) & iDRAC **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP+ only SR SFP modules possible (no LR) - Intel X710 Dual or Quad Port 10GbE Base-T copper	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0**: Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP (SR only, no LR possible) - Intel X710 Dual or Quad Port 10GbE Base-T copper	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0**: Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP (SR only, no LR possible) - Intel X710 Dual or Quad Port 10GbE Base-T copper - or FPGA options	Onboard (2x1GbE metalic) & iDRAC **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP+ only SR SFP modules possible (no LR) - Intel X710 Dual or Quad Port 10GbE Base-T copper - or FPGA options	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0:** Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x full height:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Quad Port 10/25 GbE SFP - Intel X710 Dual or Quad Port 10GbE Base-T copper - or FPGA options	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0:** Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP (SR only, no LR possible) - Intel X710 Dual or Quad Port 10GbE Base-T copper - or FPGA options	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0:** Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x full height:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Quad Port 10/25 GbE SFP - Intel X710 Dual or Quad Port 10GbE Base-T copper - or FPGA options	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0:** Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 4x full height & 2x low profile:** - see FPGA options	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0:** Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 4x full height & 2x low profile:** - see FPGA options		Onboard (2x1GbE metalic) & iDRAC **OCP 3.0**: Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP (SR only, no LR possible) - Intel X710 Dual or Quad Port 10GbE Base-T copper - or FPGA options

	Collector	Collector	Collector	Collector	Collector	Collector
	S	**M**	**L**	**XL**	**2XL**	**3XL**
DELL Model	**R360**	**R660xs**	**R6615**	**R6615**	**R7615**	**R7625**
chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	16x2.5 chassis	16x2.5 chassis
CPU	1x Intel Xeon E-2478, 2.8 GHz, 8/16	1x Intel Xeon Silver 4510, 2.4 GHz, 12/24	1x AMD EPYC 9174F, 4.1 GHz, 16/32	1x AMD EPYC 9274F, 4,05 GHz, 24/48	1x AMD EPYC 9374F, 3,85 GHz, 32/64	2x AMD EPYC 9374F, 3,85 GHz, 32/64
RAM	2x 32GB RAM DDR5	8x 16GB RAM DDR5	8x 32GB RAM DDR5	8x 64GB RAM DDR5	12x 64GB RAM DDR5	24x 64GB RAM DDR5
PERC	PERC H355	PERC H755N	PERC H755N	PERC H755N	Dual PERC H755N	Dual PERC H755N
Primary Storage	2x 3.84 SSD SATA read-intensive	2x 7.68 NVMe Datacenter	4x 7.68 NVMe Datacenter	6x 7.68 NVMe Datacenter	8x 7.68 SSD NVMe Datacenter	6x 15.36 TB NVMe
Primary RAID (default)	RAID1	RAID1	RAID6	RAID6	RAID6	RAID6
Riser config	Riser config 2, 2 x8 LP	Riser Config 5, 1 x8 LP + 1 x16 LP	Riser Config 3, 2 x16 FH (Gen5)	Riser Config 3, 2 x16 FH (Gen5)	Riser Config 4, 2 x16 FH (Gen5) + 2 x8 FH + 2 x16 LP, Half Length	Riser Config 4, 2 x16 FH (Gen5) + 2 x8 FH + 2 x16 LP, Half Length
OCP 3.0 NIC		Broadcom 5720 Quad Port 1GbE Base-T copper Adapter, OCP NIC 3.0	Intel E810-XXV Dual Port 10/25GbE SFP28, OCP 3.0	Intel E810-XXVDA4 Dual Port 10/25GbE SFP28 OCP 3.0	Intel E810-XXVDA4 Dual Port 10/25GbE SFP28 OCP 3.0	Intel E810-XXVDA4 Quad Port 10/25GbE SFP28 OCP 3.0
PCI-E NIC						
PSU	600W dual PSU	700W dual PSU	1100W dual PSU	1100W dual PSU	1100W dual PSU	1800W dual PSU
iDRAC	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise
Rails	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management
Support	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD

	Sensor	Sensor	Sensor	Sensor	Sensor	Sensor	Sensor	Sensor	Sensor	Sensor
	XS-25	**S-25**	**SA-25**	**M25**	**MA25**	**L25**	**LA25**	**XL**	**2XL**	**3XL**
DELL Model	**R360**	**R360**	**R660xs**	**R360**	R660xs	**R660xs**	R6615	**R6615**	**R6615**	**R6615**
chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis	8x2.5 chassis
CPU	1x Intel Xeon E-2434, 3.4 GHz, 4/8	1x Intel Xeon E-2436, 2.9 GHz, 6/12	1x Intel Xeon Silver 4510, 2.4 GHz, 12/24	1x Intel Xeon E-2478, 2.8 GHz, 8/16	1x Intel Xeon Silver 4510, 2.4 GHz, 12/24	1x Intel Xeon Gold 6526Y, 2.8 GHz, 16/32	1x AMD EPYC 9174F, 4.1 GHz, 16/32	1x AMD EPYC 9174F, 4.1 GHz, 16/32	1x AMD EPYC 9274F, 4,05 GHz, 24/48	1x AMD EPYC 9374F, 3,85 GHz, 32/64
RAM	1x 16GB RAM DDR5	2x 16GB RAM DDR5	2x 16GB RAM DDR5	2x 32GB RAM DDR5	4x 16GB RAM DDR5	8x 16GB RAM DDR5	8x 16GB RAM DDR5	8x 32GB RAM DDR5	12x 32GB RAM DDR5	12x 64GB RAM DDR5
PERC	PERC H355	PERC H355	PERC H355	PERC H355	PERC H355	PERC H355	PERC H355	PERC H755	PERC H755	PERC H755
Storage	2x 600GB SAS 2.5" 10k HDD	2x 1.2TB SAS 2.5" 10k HDD	2x 1.2TB SAS 2.5" 10k HDD	2x 960GB SSD read-intensive SATA	2x 960GB SSD read-intensive SATA	2x 960GB SSD read-intensive SATA	2x 960GB SSD read-intensive SATA	2x 1.92GB SSD read-intensive SATA	2x 1.92GB SSD read-intensive SATA	2x 3.84GB SSD read-intensive SATA
RAID	RAID1	RAID1	RAID1	RAID1	RAID1	RAID1	RAID1	RAID1	RAID1	RAID1
Riser config	Riser config 2, 2 x8 LP	Riser config 2, 2 x8 LP	Riser Config 5, 1 x8 LP + 1 x16 LP	Riser config 2, 2 x8 LP	Riser Config 5, 1 x8 LP + 1 x16 LP	Riser Config 5, 1 x8 LP + 1 x16 LP	Riser Config 3, 2 x16 FH (Gen5)	Riser Config 3, 2 x16 FH (Gen5)	Riser Config 3, 2 x16 FH (Gen5)	Riser Config 3, 2 x16 FH (Gen5)
OCP 3.0 NIC					Intel E810 Quad Port 10/25GbE OCP 3.0	Intel E810 Quad Port 10/25GbE OCP 3.0	Intel E810 Quad Port 10/25GbE OCP 3.0	Intel E810 Quad Port 10/25GbE OCP 3.0	Intel E810 Quad Port 10/25GbE OCP 3.0	Intel E810 Quad Port 10/25GbE OCP 3.0
NIC	Broadcom 5719 Quad Port	Broadcom 5719 Quad Port	Broadcom 5719 Quad Port	Broadcom 5719 Quad Port Intel E810 Dual Port 10/25GbE	Broadcom 5719 Quad Port	Broadcom 5719 Quad Port	Broadcom 5719 Quad Port			
PSU	600W dual PSU	600W dual PSU	700W dual PSU	600W dual PSU	700W dual PSU	700W dual PSU	700W dual PSU	700W dual PSU	1100W dual PSU	1100W dual PSU
iDRAC	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise	iDRAC9 Enterprise
Rails	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management	ReadyRails with cable management
Support	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD	5 yrs. Dell ProSupport NBD
Optional FPGA [1]				Napatech FPGA Options: 2x25/10GbE	Napatech FPGA Options: 2x25/10GbE	Napatech FPGA Options: 2x25/10GbE	Napatech FPGA Options: - 4x10GbE - 2x25/10GbE or 4x25/10GbE - 2x40/100GbE	Napatech FPGA Options: - 4x10GbE - 2x25/10GbE or 4x25/10GbE - 2x40/100GbE	Napatech FPGA Options: - 4x10GbE - 2x25/10GbE or 4x25/10GbE - 2x40/100GbE	Napatech FPGA Options: - 4x10GbE - 2x25/10GbE or 4x25/10GbE - 2x40/100GbE
Maximum monitoring ports w/o Napatech FPGA	Onboard (2x1GbE metalic) & iDRAC **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP+ only SR SFP modules possible (no LR) - Intel X710 Dual or Quad Port 10GbE Base-T copper	Onboard (2x1GbE metalic) & iDRAC **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP+ only SR SFP modules possible (no LR) - Intel X710 Dual or Quad Port 10GbE Base-T copper	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0**: Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP (SR only, no LR possible) - Intel X710 Dual or Quad Port 10GbE Base-T copper	Onboard (2x1GbE metalic) & iDRAC **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP+ only SR SFP modules possible (no LR) - Intel X710 Dual or Quad Port 10GbE Base-T copper	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0**: Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP (SR only, no LR possible) - Intel X710 Dual or Quad Port 10GbE Base-T copper - or FPGA options	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0**: Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Quad Port 10/25 GbE SFP - Intel X710 Dual or Quad Port 10GbE Base-T copper - or FPGA options	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0**: Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x low profile:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP (SR only, no LR possible) - Intel X710 Dual or Quad Port 10GbE Base-T copper - or FPGA options	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0**: Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x full height:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Quad Port 10/25 GbE SFP - Intel X710 Dual or Quad Port 10GbE Base-T copper - or FPGA options	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0**: Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x full height:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Dual Port 10/25 GbE SFP - Intel X710 Dual or Quad Port 10GbE Base-T copper - or FPGA options	Onboard (2x1GbE metalic) & iDRAC **OCP 3.0**: Up to Intel E810 QP or DP (SR only, no LR possible), or BCM 5720 QP **PCI-E 2x full height:** - Broadcom 5719 Quad Port 1GbE copper - Intel E810 Quad Port 10/25 GbE SFP - Intel X710 Dual or Quad Port 10GbE Base-T copper - or FPGA options

Notes:

1 Optional FPGAs allow only for NAPATECH network cards to be used as monitoring ports. No combination of Intel/Broadcom and Napatech FPGA cards processing traffic simultanously is possible. Intel or Broadcom cards are kept on the servers because they can be used only when Napatech is not.

Logmanager dokumentace

Kompletní dokumentace pro systém Logmanager.

Úvod

- End-User License Agreement (EULA)
- Zákaznická podpora
- Poznámky k vydání
- Logmanager technický bulletin
- Průvodce pro rychlý start
- Průvodce pro rychlý start virtuálního Logmanageru
- Použitý open-source software

Logmanager Beats Agent

- Automatická migrace
- Logmanager Orchestrator
- Beat agenti
- Beat globální konfigurace
- Filtry Beat agentů
- Microsoft security audit

Webové rozhraní

- Webové rozhraní

CLI rozhraní

- Nastavení CLI

Virtuální Logmanager

- Virtuální Logmanager

Logmanager Forwarder

- Logmanager Forwarder

Zdrojová zařízení a aplikace

- Konfigurace zdrojových zařízení a aplikací

Doporučení a návody

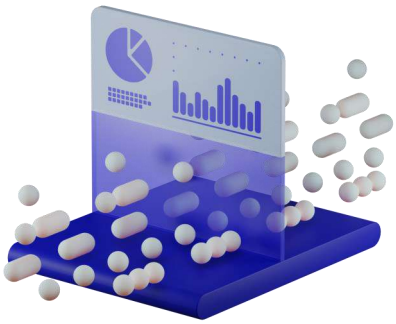
- Obecná doporučení na minimalizaci nebezpečí a možných dopadů cílených útoků
- Návod na vytvoření parseru
- Popis k provádění PoC

- Microsoft security audit
- Microsoft SQL audit
- Dokumentace ke tvorbě šablon
- Jak postupovat v případě nefunkčního/neexistujícího parseru

Další informace

- Komunikace Logmanageru
- Zpracování událostí pomocí blockly
- Vzdálená správa systému/serveru Logmanager
- Popis migrace stávajícího Logmanageru na nový
- Downgrade software Logmanager
- Logmanager s High Availability nasazením
- Standardizované názvy proměnných
- Programová aktualizace Vyhledávací tabulky
- Logmanager log flow

Logmanager



Centrální úložiště logů a dostupný SIEM

V dnešním přetechnizovaném světě jsou informace kritickým zdrojem umožňujícím správné rozhodnutí ve správný čas. V protikladu k tomu konstatování stojí fakt, že důležité informace jsou distribuovány v nejrůznějších zařízeních a aplikacích napříč celou organizací, ne vždy ve snadno pochopitelném formátu a s rozdílnou dostupností. Sjednocení informací z mnoha zdrojů a jejich přeložení do lidsky srozumitelného tvaru, nastavení pevných pravidel pro nakládání s informacemi a jejich nezpochybnitelnost jsou proto klíčové požadavky pro efektivitu bezpečnostních i operativních činností každé organizace. Když se k tomu přidá i přehledná interpretace těchto informací v kompaktním a výkonném nástroji, získá IT organizace nástroj pro realizaci správných rozhodnutí. A tímto nástrojem je český systém Logmanager.

Určení systému Logmanager

Logmanager je HW řešení pro centralizovanou správu logů a jiných strojových dat z libovolných zdrojů. Je založen na výkonné databázi s obrovskou kapacitou, rychlým vyhledáváním ve „velkých datech“ a okamžitou vizualizací vyžádaných dat. Jeho podstatou je sběr, dlouhodobé nezpochybnitelné ukládání a analýza strojových dat organizace. Umožňuje prohledávat agregovaná data v reálném čase, vytvářet analýzy, reporty a upozornění na události korelované z dat více zdrojů. Nedílnou součástí řešení Logmanager je taktéž podpora souladu s požadavky zákonných norem.

Při správné implementaci pomůže organizaci k zajištění shody s ČSN/ISO 27001:2013 o pořizování auditních záznamů, plnění požadavků GDPR či Zákona o kybernetické bezpečnosti.

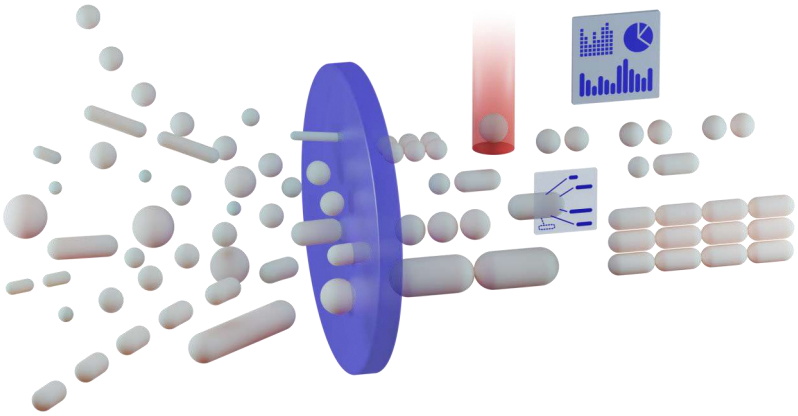
Logmanager však není určen pouze pro oddělení bezpečnosti IT.

Při vývoji Logmanageru je kladen velký důraz na radikální jednoduchost užívání a jeho reálný přínos pro IT obecně. Logmanager na jednom místě shromáždí provozní a diagnostická data ze všech důležitých systémů organizace. Operátor IT tak dostane možnost zjistit během několika sekund informace o provozních stavech a případných závadách, které by jinak musel hodiny komplikovaně vyhledávat v distribuovaných zdrojích. K tomu je i automaticky informován o sledovaných událostech a může tak předcházet kritickým IT nebo bezpečnostním incidentům.

Podporovaná zařízení

Logmanager nativně podporuje více než 125 zdrojů ze všech oblastí IT, od bezpečnostních řešení, přes sítě, virtualizace, operační systémy, databáze, až po cloud aplikace. Seznam zdrojů se každou aktualizací rozšiřuje. Logmanager dále podporuje standardizované strukturované formáty logů jako jsou CEF, LEEF, RFC5424 a JSON. Pro unikátní zdroje dat umožňuje Logmanager rychlé a snadné vytvoření nové integrace.

Vlastnosti produktu



Klíčové vlastnosti

- Centrální úložiště logů, událostí a strojových dat organizace
- Sjednocení formátu zdrojových logů do lidsky srozumitelné formy
- Zpracování a vizualizace přijímaných dat v reálném čase
- Rychlé prohledávání dat bez nutnosti znalosti SQL jazyka
- SIEM funkce. Alerty na základě podmínek s limity a korelacemi
- Windows Agent pro sběr logů ze serverů a stanic
- Unikátní grafické konfigurační a programovací rozhraní
- Radikální jednoduchost a uživatelská přívětivost
- Snadné vytváření reportů a auditních zpráv za běhu
- Bez licenčních omezení na množství zdrojů, výkon a uložená data
- Úspora na licencích při budoucím rozšíření směrem k SIEM/SOC
- Jednoduchá klasifikace pomocí unikátního nástroje Blockly

Umožní snadnější splnění požadavků na shodu s regulacemi pro:

GDPR

Zákon o kybernetické bezpečnosti a návazné vyhlášky

ČSN/ISO 27001:2013 pro pořizování auditních záznamů

PCI DSS 3.2

NIS 2

Konkurenční výhody

- Řešení „vše v jednom“ - Server, OS, Aplikace, DB na jedné platformě
- Varianta diskového úložiště až 160 TB
- Podpora velkého množství zdrojových zařízení, OS a aplikací
- Centrálně řízený klient pro sběr logů z Windows OS
- Možnost vysoké dostupnosti v clusteru
- Snadná integrace se SIEM/UBA systémy třetích stran
- Rychlé nasazení a snadné zaškolení pro běžné operace
- Rozhraní i kompletní dokumentace v českém jazyce
- Rozsáhlá síť spolehlivých a technicky zdatných partnerů
- Přímá technická podpora výrobcem a možnost testování zdarma
- Podpora více než 150 zařízení

Informace o výrobci Logmanager a.s. a reference

Logmanager je vyvíjen od roku 2014 jako nosný produkt firmy Logmanager a.s., která sídlí v Praze. Na stránkách logmanager.com naleznete vybrané reference. Pro podrobnější seznam referencí přímo z oblasti Vaší činnosti nás neváhejte poptat. Příslušné kontakty na stávající zákazníky, kteří souhlasí s uváděním na referenčním listu, rádi předáme.

Specifikace produktu

Logmanager Appliance

Logmanager XL

Server 2U Dell - 4 x Gbit a management, redundantní napájení

Procesor	Paměť	Disky	RAID	Kapacita úložiště	Odhad retence při maximálním vytížení	Trvalý výkon	Špičkový výkon
2 × 16 core 3.2 GHz	128 GB	12 × 12 TB	6	120 TB / 160 TB	220 / 290 dní	10 000 EPS	20 000 EPS

Logmanager L

Server 2U Dell - 4 x Gbit a management, redundantní napájení

Procesor	Paměť	Disky	RAID	Kapacita úložiště	Odhad retence	Trvalý výkon	Špičkový výkon
2 × 16 core 2.8 GHz	128 GB	12 × 4 TB	6	40 TB / 80 TB	140 / 280 dní	5 000 EPS	10 000 EPS

Logmanager M

Server 1U Dell - 4 x Gbit a management, redundantní napájení

Procesor	Paměť	Disky	RAID	Kapacita úložiště	Odhad retence při maximálním vytížení	Trvalý výkon	Špičkový výkon
1 × 16 core 3 GHz	64 GB	4 × 4 TB	5	12 TB	120 dní	2 000 EPS	4 000 EPS

Logmanager Demo

MicroPC 1 x Gbit, pouze jako neprodukční zařízení pro lab, nebo PoC)

Procesor	Paměť	Disky	RAID	Kapacita úložiště	Odhad retence při maximálním vytížení	Trvalý výkon	Špičkový výkon
1 x 8 core 2.6 GHz	32 GB	1 × 500 GB	–	500 GB	30 dní	500 EPS	1 000 EPS

Logmanager Cluster

V rámci clusteru Logmanager lze kombinovat pouze totožné modely. Pro více informací kontaktujte vašeho autorizovaného partnera.

Logmanager Forwarder

Řešení pro bezpečný a spolehlivý sběr logů ze vzdálených poboček, Internetu či DMZ

Logmanager VF

Virtuální Logmanager Forwarder - podpora Hyper-V a VMware

Procesor	Paměť	Disky	RAID	Kapacita úložiště	Odhad retence při maximálním vytížení	Trvalý výkon	Špičkový výkon
2 × vCPU	6 GB vRAM	1 x vDisk	–	8 / 16 / 128 GB	–	9 000	18 000

Logmanager HF

MicroPC - 1 x Gbit

Procesor	Paměť	Disky	RAID	Kapacita úložiště	Odhad retence při maximálním vytížení	Trvalý výkon	Špičkový výkon
1 x 8 core 2.6 GHz	8 GB	1 x 256 GB	–	250 GB	–	9 000	18 000

Rozšíření

Síťové rozšíření

Volitelné rozšíření síťového rozhraní o 2 x nebo 4 x SFP+ porty (10 GbE) pro Logmanager L a XL. Volitelné i vhodné SFP+ transceivery.

DB Rozšíření

Volitelné rozšíření kapacity úložiště o dalších 40 TB pro Logmanager L na 80 TB a XL na 160 TB. (nutno specifikovat s objednávkou Logmanager).

- EPS - Events per secconds, nebo-li události za vteřinu. Hodnota udávající výkon Logmanager při plném zpracování logů.
- Špičkový výkon - udává maximální výkon po dobu 20 minut při dvojnásobném zatížení.
- Logmanager hardware je automaticky dodáván s 5letou NBD zárukou, 1letou software podporou a 1 licencí pro virtuální Logmanager Forwarder.
- Odhad retence EPS - Očekávané množství událostí za sekundu, log mix s RAW velikostí logů průměrně 700 Byte.
- Workload Akcelerátor - navyšuje výkon u modelu L o 1000 EPS, u Logmanager XL je vždy součástí.
- Logmanager využívá systémové zobrazení značení, vycházející ze součtu kapacit disků tak, jak je udávají výrobci, kteří používají tzv. dekadické předpony (1 kB = 1000 B). Logmanager interně používá systémové zobrazení, tzv. binární předpony (1 KiB = 1024 B).

Dell PowerEdge R360

Technical Guide

Regulatory Model: E77S Series
Regulatory Type: E77S002
December 2024
Rev. A02

DELL Technologies

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

© 2024 Dell Inc. or its subsidiaries. All rights reserved. Dell Technologies, Dell, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Contents

Chapter 1: System overview 5
Key workloads 5
New technologies 5

Chapter 2: System features and generational comparison 7

Chapter 3: Chassis views and features 9
Front view of the system 9
Rear view of the system 10
Inside the system 12

Chapter 4: Processor 13
Processor features 13
Supported processors 13

Chapter 5: Memory subsystem 14
Supported memory 14

Chapter 6: Storage 15
Storage controllers 15
Storage controller feature matrix 15
Supported Drives 17
Hard Disk Drives (HDDs) 17
Solid State Drives (SSDs) 19
Internal storage configuration 20
External Storage 21

Chapter 7: Networking 22
Supported network cards 22
Overview 22

Chapter 8: PCIe subsystem 23
PCIe risers 23

Chapter 9: Power, thermal, and acoustics 25
Power 25
Power Supply Units 26
Thermal 27
Thermal design 27
Acoustics 28
Acoustical performance 28

Chapter 10: Rack, rails, and cable management 30

Rack Rails....30

Chapter 11: Operating Systems and Virtualization.... 36
Supported Operating Systems....36

Chapter 12: Dell Systems Management....37
Integrated Dell Remote Access Controller (iDRAC).... 37
Systems Management software support matrix.... 38

Chapter 13: Appendix D: Service and support.... 40
Why attach service contracts....40
ProSupport Infrastructure Suite 40
Specialty Support Services....42
ProDeploy Infrastructure Suite....43
Supplemental Deployment Services 46
Unique Deployment Scenarios....47
DAY 2 – Automation Services with Ansible.... 48
Dell Technologies Consulting Services....49

Chapter 14: Appendix A: Additional specifications....51
Chassis dimensions 51
Chassis weight....52
NIC port specifications.... 52
Video specifications....52
USB Ports.... 53
PSU rating....54
Environmental specifications....55
Particulate and gaseous contamination specifications.... 56
Thermal restriction matrix....57
Thermal air restrictions....58

Chapter 15: Appendix A. Standards compliance.... 59

Chapter 16: Appendix C Additional resources....60

1

System overview

The Dell PowerEdge R360 is a powerful entry-level one-socket rack server designed to meet the evolving computing demands from organizations.

The system features:

- One Intel ® Xeon-E series processor or one Intel ® Pentium processor
- Four DDR5 DIMM slots
- Two redundant AC or DC power supply units
- PCI Express® (PCIe) 4.0 enabled expansion slots
- Up to 8 x 2.5-inch SAS/SATA or up to 4 x 3.5 inch SAS/SATA
- Onboard LAN ports

Topics:

- Key workloads
- New technologies

Key workloads

PowerEdge R360 is versatile enough to address many customer segments and workloads affordably, that includes:

- Collaboration/Sharing: Built-in features to enable collaborative applications between groups of people that share information and processes on-site or remotely.
- ROBO Business/Database: Provide computing performance for ROBO business and database workloads including mailing, data process and analysis.
- Near Edge: Suitable form-factor for near-edge computing with GPU applications, such as video and audio analysis, surveillance, and VDI

New technologies

Table 1. New technologies

Technology	Detailed Description
Intel Raptor lake - E Processor (Socket V0, LGA-1700)	Core count: Up to 8 core processor
	Maximum number of PCIe lanes: Integrated 16 PCIe Gen5 lanes @ 32 GT/s, 4 PCIe Gen4 lanes @ 16GT/s
	Maximum TDP: 95 W
Intel Alder lake Pentium Processor (Socket V0, LGA-1700)	Core count: Up to 2 core processor
	Maximum number of PCIe lanes: Integrated 16 PCIe Gen5 lanes @ 32 GT/s, 4 PCIe Gen4 lanes @ 16GT/s
	Maximum TDP: 46 W
4800 MT/s DDR5 Memory	Maximum 4 DIMM slots per system
	Supports DDR5 ECC UDIMM up to 4400 MT/s
Flex I/O	LOM, 2x 1 GB with BCM5720 LAN controller
	Rear I/O with: ● 1 GB Dedicated Management Network Port ● USB 3.0 x1

Table 1. New technologies (continued)

Technology	Detailed Description
	• USB 2.0 x1 • VGA port • Serial port
	Front I/O with: • USB 2.0 x1 • Dedicated iDRAC Direct micro-USB x 1
CPLD 1-wire	Support payload data of Riser, BOSS-N1, BP, and Rear I/O to BIOS and iDRAC
Dedicated PERC	PERC 11
Power Supplies	60 mm dimension is the new PSU form factor design.
	600 W AC/DC Platinum
	700 W AC/DC Titanium

2

System features and generational comparison

The following table shows the comparison between the PowerEdge R360 with the PowerEdge R350.

Table 2. Features comparison

Features	PowerEdge R360	PowerEdge R350
Processors	One Intel Xeon E-2400 series processors with up to 8 cores or one Intel Pentium processor with up to 2 cores	One Intel Xeon E-2300 series processor with up to 8 cores or one Intel Pentium processor with up to 2 cores
Memory	Up to 4 x DDR5 ECC UDIMMs DIMM Speed: Up to 4400 MT/s	Up to 4 x DDR4 UDIMMs DIMM Speed: Up to 3200 MT/s
Storage Controllers	Internal Controllers: PERC H355, PERC H755, PERC H355f, PERC H755f, HBA355i, HBA355i (front)	Internal Controllers: PERC H355, PERC H355f, PERC H345, PERC H755, PERC H755f, HBA355i, PERC H345f, HBA355f
	External Controllers: HBA355e	External Controllers: HBA355e
	Software RAID: S160	Software RAID: S150
Drive Bays	Front: Up to 4 x 3.5-inch SAS/SATA (HDD/SSD) Up to 8 x 2.5-inch SAS/SATA (HDD/SSD)	Front: Up to 4 x 3.5-inch SAS/SATA (HDD/SSD) Up to 8 x 2.5-inch SAS/SATA (HDD/SSD)
Power Supplies	600 W Mixed Mode Platinum 100 V-240 V AC or 240 V DC. Hot Swap redundant 700 W Mixed Mode Titanium 200 V-240 V AC or 240 V DC. Hot Swap redundant	600 W Mixed Mode Platinum 100 V-240 V AC or 240 V DC. Hot Swap redundant 700 W Mixed Mode Titanium 200 V-240 V AC or 240 V DC. Hot Swap redundant
Cooling Options	Air cooled	Air cooled
Fans	Up to 3 STD fans and 1 HPR fan	Up to 3 STD fans and 1 HPR fan
Dimension	H x W x D: 1U x 482 mm x 563.3 mm	H x W x D: 1U x 482 mm x 534.59 mm
Form Factor	1U	1U
Embedded Management	iDRAC9, iDRAC Direct, iDRAC RESTful API with Redfish, iDRAC Service Manual	iDRAC9 Enterprise, iDRAC Direct, Datacenter license options; iDRAC Service Module, iDRAC RESTful API with Redfish
OpenManage Software	• OpenManage Enterprise • OpenManage Power Manager plugin • OpenManage Service plugin • OpenManage Update Manager plugin • CloudIQ for PowerEdge plug in • OpenManage Enterprise Integration for VMware vCenter • OpenManage Integration for Microsoft System Center	• OpenManage Enterprise • OpenManage Power Manager plugin • OpenManage SupportAssist plugin • OpenManage Update Manager plugin

Table 2. Features comparison (continued)

Features	PowerEdge R360	PowerEdge R350	
	• OpenManage Integration with Windows Admin Center		
Mobility	OpenManage Mobile	OpenManage Mobile	
Integrations and Connections	OpenManage Integrations • BMC Truesight • Microsoft System Center • OpenManage Integration with ServiceNow • Red Hat Ansible Modules • Terraform Providers • VMware vCenter and vRealize Operations Manager	Integrations: • Microsoft® System Center • VMware® vCenter™ and vRealize Operations Manager • BMC Truesight • Red Hat Ansible Modules	Connections: • Nagios Core & Nagios XI • Micro Focus Operations Manager i (OMi) • IBM Tivoli Netcool/ OMNIbus • IBM Tivoli® Network Manager IP Edition
Security	• Cryptographically signed firmware • Data at Rest Encryption (SEDs with local or external key mgmt) • Secure Boot • Secured Component Verification (Hardware integrity check) • Secure Erase • Secured-core server • Silicon Root of Trust • System Lockdown (requires iDRAC9 Enterprise or Datacenter) • TPM 2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ	• TPM 1.2/2.0 FIPS, CC-TCG certified, TCM 2.0 optional • Cryptographically signed firmware • Silicon Root of Trust • Secure Boot • System Lockdown (requires iDRAC9 Enterprise or Datacenter) • Secure Erase	
Embedded NIC	2 x 1 GbE LOM	2 x 1 GbE LOM	
Networking Options	Optional Network Card	Optional Network Card	
GPU Options	Up to 1 x 60 W SW	None	
Ports	Front: 1 x USB 2.0, 1x iDRAC Direct (Micro-USB) Rear: 1 x VGA, 1 x USB2.0 + 1 x USB 3.2 Gen1 Internal : USB 3.2 Gen1	Front: 1 x USB 2.0, 1 managed (micro USB) + front VGA Rear: 1 x Gen2 + 1 x Gen3 USB Optional Internal USB	
PCIe	Up to 2 x PCIe Slots Gen4	Up to 2 x PCIe Slots Gen4	
Operating System and Hypervisors	• Canonical Ubuntu Server LTS • Microsoft Windows Server with Hyper-V • Red Hat Enterprise Linux • SUSE Linux Enterprise Server • VMware ESXi For specifications and interoperability details, see OS support.	• Canonical® Ubuntu® LTS • Citrix® Hypervisor • Microsoft® Windows Server® LTSC with Hyper-V • Red Hat® Enterprise Linux • SUSE® Linux Enterprise Server For specifications and interoperability details, see Dell.com/OSsupport	

3

Chassis views and features

Topics:

- Front view of the system
- Rear view of the system
- Inside the system

Front view of the system



Figure 1. Front view of 4 x 3.5-inch drive system

Table 3. Features are available on the front of the system for a 3.5-inch drive system

Item	Ports, panels, and slots	Icon	Description
1	Left control panel	N/A	Contains the system health, system ID, and the status LED indicator.
2	3.5-inch drive or 2.5-inch drive in 3.5-inch carrier adapter.	NA	Enables you to install drives that are supported on your system. .
3	Right control panel	N/A	Contains the Power button, USB 2.0 port and iDRAC direct Micro USB port



Figure 2. Front view of 8 x 2.5-inch drive system

Table 4. Features are available on the front of the system for a 2.5-inch drive system

Item	Ports, panels, and slots	Icon	Description
1	Left control panel	N/A	Contains the system health, system ID, and the status LED indicator.
2	2.5-inch drive	NA	Enables you to install drives that are supported on your system. .
3	Right control panel	N/A	Contains the Power button, USB 2.0 port and iDRAC direct Micro USB port

Rear view of the system

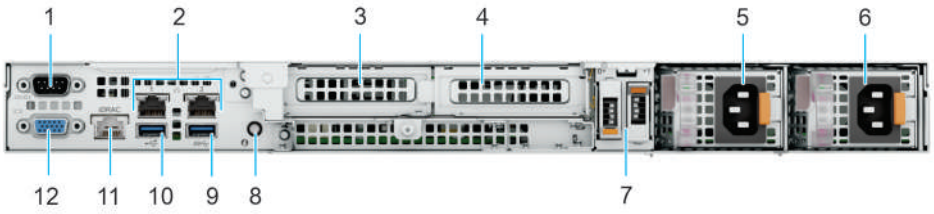


Figure 3. Rear view of the system

Table 5. Features available on the rear of the system

Item	Ports, panels, or slots	Icon	Description
1	Serial port		Enables you to connect a serial device to the system.
2	NIC ports		The NIC ports that are integrated on the system board provide network connectivity.
3	PCIe expansion riser card slot 1	N/A	The expansion card riser enables you to connect PCI Express expansion cards.For more information, see the Expansion card installation guidelines section.
4	PCIe expansion riser card slot 2	N/A	The expansion card riser enables you to connect PCI Express expansion cards.For more information, see the Expansion card installation guidelines section.
5	Power supply unit (PSU 1)		Indicates the PSU.
6	Power supply unit (PSU 2)		Indicates the PSU.
7	BOSS-N1 module	N/A	BOSS-N1 module for internal system boot.

Table 5. Features available on the rear of the system (continued)

Item	Ports, panels, or slots	Icon	Description
8	System Identification (ID) button		The System Identification (ID) button is available on the front and back of the system. Press the button to identify a system in a rack by turning on the system ID button. You can also use the system ID button to reset iDRAC and to access BIOS using the step through mode. When pressed, the system ID LED in the back panel blinks until either the front or rear button is pressed again. Press the button to toggle between on or off mode. **NOTE:** If the server stops responding during POST, press and hold the **System ID** button for more than five seconds to enter the BIOS progress mode **NOTE:** To reset the iDRAC (if not disabled on the iDRAC setup page by pressing F2 during system boot), press and hold the **System ID** button for more than 15 seconds.
9	USB 3.2 Gen1 port		This port is USB 3.2 Gen1-compliant.
10	USB 2.0 port		This port is USB 2.0-compliant.
11	Dedicated iDRAC Ethernet port	iDRAC	Enables you to remotely access iDRAC. For more information, see the Integrated Dell Remote Access Controller User's Guide at PowerEdge Manuals.
12	VGA port		Enables you to connect a display device to the system.

Inside the system

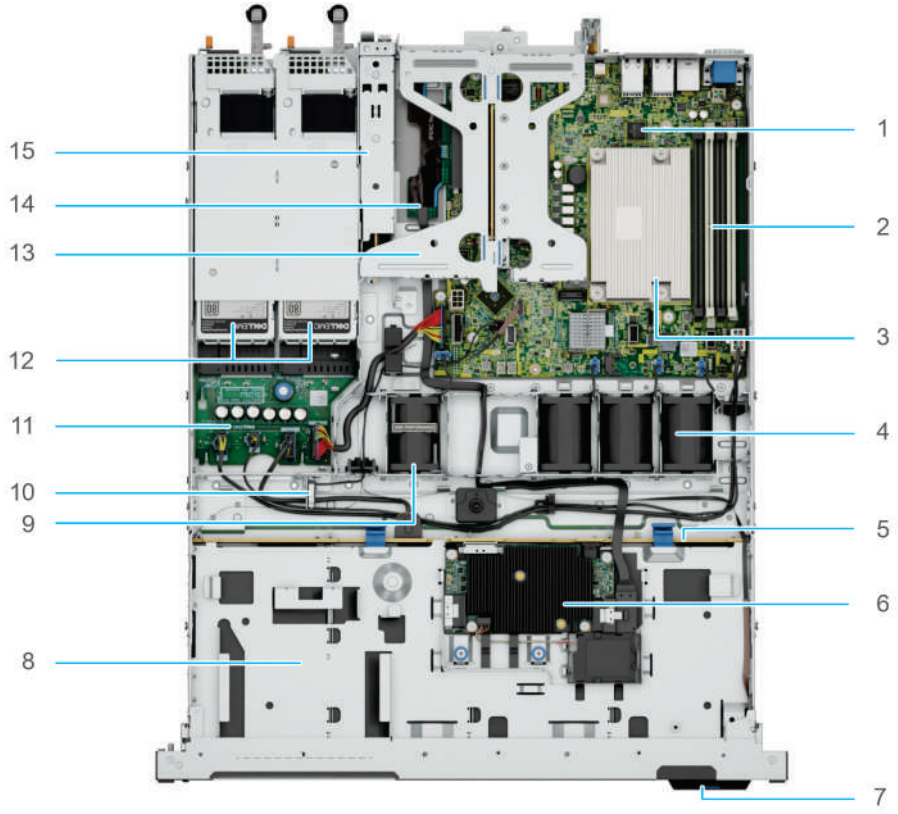


Figure 4. Inside the system

1. System board
2. DIMM slots
3. Processor heatsink
4. Standard (STD) cooling fans
5. Drive backplane
6. Front PERC card
7. Express Service Code Tag
8. Front drives
9. High Performance (HPR) fan - Optional
10. Intrusion switch
11. Power Interposer board (PIB)
12. Power supply units
13. Butterfly Riser with 2x PCIe slots
14. Front PERC card riser on the dedicated internal PCIe slot
15. BOSS N1 module

4

Processor

Topics:

- Processor features

Processor features

The following lists the features and functions that are in the upcoming Intel® Xeon E-2400 series Processor offering:

- For small businesses, Xeon E provides a reliable solution with for business ready servers to support business-critical services and customer data needs.
- For cloud services, Xeon E is the perfect option for Bare Metal Instances and Code/Data Security, with a cost-effective solution for essential performance to support entry-level bare metal services.

Key updates for the Intel® Xeon E-2400 series include increased performance with a new processor core architecture:

- 4, 6, and 8 core option
- Up to 95 W TDP
- Up to 128 GB memory
- DDR5 up to 4800 MT/s

Supported processors

Table 6. Supported processors for PowerEdge R360

Processor	Clock Speed (GHz)	Cache (M)	Cores	Threads	Turbo *	Memory Speed (MT/s)	Memory Capacity	TDP
E-2488	3.2	24	8	16	Turbo	4800	128 GB	95 W
E-2486	3.5	18	6	12	Turbo	4800	128 GB	95 W
E-2478	2.8	24	8	16	Turbo	4800	128 GB	80 W
E-2468	2.6	24	8	16	Turbo	4800	128 GB	65 W
E-2456	3.3	18	6	12	Turbo	4800	128 GB	80 W
E-2436	2.9	18	6	12	Turbo	4800	128 GB	65 W
E-2434	3.4	12	4	8	Turbo	4800	128 GB	55 W
E-2414	2.6	12	4	4	Turbo	4800	128 GB	55 W
G7400	3.7	6	2	4	No Turbo	4800	128 GB	46 W
G7400T	3.1	6	2	4	No Turbo	4800	128 GB	35 W

NOTE: Intel E-2400 Series processors under turbo mode will not be able to meet maximum frequency due to processor TDP thermal limitations.

5

Memory subsystem

Topics:

- Supported memory

Supported memory

Table 7. Memory technology

DIMM type	Rank	Capacity	DIMM rated voltage and speed	Operating Speed	
				1 DIMM per channel (DPC)	2 DIMM per channel (DPC)
ECC UDIMM	1 R	16 GB	DDR5 (1.1 V), 4800 MT/s	4400 MT/s	4000 MT/s
	2 R	32 GB	DDR5 (1.1 V), 4800 MT/s	4400 MT/s	3600 MT/s
	1 R	16 GB	DDR5 (1.1 V), 5600 MT/s	4400 MT/s	4000 MT/s
	2 R	32 GB	DDR5 (1.1 V), 5600 MT/s	4400 MT/s	3600 MT/s

Table 8. Supported DIMMs

Rated DIMM Speed (MT/s)	DIMM Type	DIMM Capacity (GB)	Ranks per DIMM	Data Width	DIMM Volts
4800	UDIMM	16	1	8	1.1
4800	UDIMM	32	2	8	1.1
5600	UDIMM	16	1	8	1.1
5600	UDIMM	32	2	8	1.1

6

Storage

Topics:

- Storage controllers
- Supported Drives
- Internal storage configuration
- External Storage

Storage controllers

NOTE: The size of the RAID 1 drives must be less than that of the second RAID container.

Dell's RAID controller options offer performance improvements, including the Adapter PERC solution. Adapter PERC provides a base RAID HW controller without consuming a PCIe slot by using a small form factor and high-density connector to the base planar.

Table 9. PERC Series controller offerings

Performance Level	Controller and Description
Entry	S160 - Windows software RAID
Value	HBA355i Adapter (Internal), HBA355i Front
Value Performance	H755 Adapter, H355 Adapter, PERC H355 Front, PERC H755 Front
External Controllers	HBA355e

NOTE: For more information about the features of the Dell PowerEdge RAID controllers (PERC), Software RAID controllers, or BOSS card, and on deploying the cards, see the storage controller documentation at Storage Controller Manuals.

NOTE: H355 will replace H345 as the entry raid controller.

Storage controller feature matrix

Table 10. Storage controller feature matrix

Model and Form Factors	Interface Support	PCI Support	SAS Connection	Cache Memory Size	Write Back Cache	RAID Levels	Max Drive Support	RAID Support
PowerEdge Server-Storage Controllers (PERC and SAS HBA) Series 11								
H755 Adapter	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA Gen3 (8 GT/s) NVMe	PCIe Gen 4	16 ports- 2x8 Internal	8 GB NV	Flash Backed Cache	0,1,5,6,10,50,60	16/ controller 50 with SAS Expander	Hardware

Table 10. Storage controller feature matrix (continued)

Model and Form Factors	Interface Support	PCI Support	SAS Connection	Cache Memory Size	Write Back Cache	RAID Levels	Max Drive Support	RAID Support
	Gen4 (16 GT/s) NVMe							
H755 Front	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA Gen3 (8 GT/s) NVMe Gen4 (16 GT/s) NVMe	PCIe Gen 4	16 ports- 2x8 Internal	8 GB NV	Flash Backed Cache	0,1,5,6,10,50,60	16/ controller 50 with SAS Expander	Hardware
HBA355i Adapter	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA	PCIe Gen 4	16 ports- 2x8 Internal	N/A	N/A	N/A	16/ controller 50 with SAS Expander	N/A
HBA355i Front	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA	PCIe Gen 4	16 ports- 2x8 Internal	N/A	N/A	N/A	16/ controller 50 with SAS Expander	N/A
HBA355e Adapter	12Gb/s SAS 6Gb/s SAS/SATA 3Gb/s SAS/SATA	PCIe Gen 4	16 ports- 4x4 external	N/A	N/A	N/A	240	N/A
H355 Adapter	12Gb/s SAS 6Gb/s SAS/SATA	PCIe Gen 4	16 ports- 2x8 Internal	No Cache	No Cache	0,1,10	Up to 32 RAID, or 32 Non- RAID	Hardware
H355 Front	12Gb/s SAS 6Gb/s SAS/SATA	PCIe Gen 4	16 ports- 2x8 Internal	No Cache	No Cache	0,1,10	Up to 32 RAID, or 32 Non- RAID	Hardware
S160 Software RAID	Gen4 (16 GT/s) NVMe	PCIe Gen 4	N/A	No Cache	No Cache	0,1,5,10	Up to 8	Software RAID - Windows only

NOTE:

1. RAID 5/50 removed from entry RAID card
2. SWRAID support for Linus provides a pre-boot configuration utility to configure MDRAID and degraded boot capability.
3. For information, post-RTS, see the Storage controller documentation at Storage Controller Manuals.

This document is updated as changes happen, so for the latest version be sure to bookmark it rather than downloading an offline copy or refer to the Storage Controller Matrix on sales portal.

Internal storage configuration

PowerEdge R360 supports the following internal storage configurations:

- 8 x 2.5 inch SAS/SATA RAID
- 4 x 3.5 inch SAS/SATA RAID

NOTE: The PowerEdge R360 system does not support rear storage.

Supported Drives

The table shown below lists the internal drives supported by the PowerEdge R360. Refer to Agile for the latest SDL.

Table 11. Supported Drives

Form Factor	Type	Speed	Rotational Speed	Capacities
2.5 inches	vSAS	12 Gb	SSD	960 Gb, 1.92 TB, 3.84 TB, 7.68 TB
	SAS	24 Gb	SSD	800 Gb, 960 GB, 1.6 TB, 1.92 TB, 3.84 TB, 7.68 TB
	SATA	6 Gb	SSD	480 Gb, 960 Gb, 1.92 TB, 3.84 TB
3.5 inches	SAS	12 Gb	7.2 K	4 TB, 8 TB, 12 TB, 16 TB
	SATA	6 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16 TB

Hard Disk Drives (HDDs)

HDD feature matrix

Table 12. HDD feature matrix

Type	Interface	Form Factor	RPM	Sector	Security	Capacity
HDD	12 Gbps SAS	2.5	10K	512n	ISE	600 GB
HDD	12 Gbps SAS	2.5	10K	512n	ISE	1.2 TB
HDD	12 Gbps SAS	2.5	10K	512e	ISE	2.4 TB
HDD	12 Gbps SAS	2.5	10K	512e	FIPS-140	2.4 TB
HDD	12 Gbps SAS	3.5	7.2K	512n	ISE	4 TB
HDD	12 Gbps SAS	3.5	7.2K	512e	ISE	8 TB
HDD	12 Gbps SAS	3.5	7.2K	512e	FIPS-140	8 TB
HDD	12 Gbps SAS	3.5	7.2K	512e	ISE	12 TB
HDD	12 Gbps SAS	3.5	7.2K	512e	ISE	16 TB
HDD	12 Gbps SAS	3.5	7.2K	512e	FIPS-140	16TB
HDD	6 Gbps SATA	3.5	7.2K	512n	ISE	2 TB
HDD	6 Gbps SATA	3.5	7.2K	512n	ISE	4 TB
HDD	6 Gbps SATA	3.5	7.2K	512e	ISE	8 TB
HDD	6 Gbps SATA	3.5	7.2K	512e	ISE	12 TB
HDD	6 Gbps SATA	3.5	7.2K	512e	ISE	16 TB

This document is updated as changes happen, so be sure to bookmark it rather than downloading an offline copy to stay with the latest information or see the Drive and Platform Matrix.

HDD Facts

HDD (Hard Disk Drive) is a storage media that are characterized by a set of spinning platters with arms to move recording heads over the surfaces to the right locations for reading and writing designated data.



Figure 5. Hard Disk Drive

The heads read or write the data and transfer it through the interface to the server. That interface for Dell standard enterprise HDDs can be either Serial Attached SCSI (SAS) or Serial ATA (SATA) and affects the speed at which data is transferred. Typically, SATA is 6 gigabits/sec. Where SAS is 12 gigabits/sec so throughput for SAS can be twice that of SATA. Additionally, due to better signal to noise ratios, SAS can have longer cable lengths, allowing it to connect to external data storage. SAS is also considered a more robust protocol.

Enterprise HDDs are typically used with multiple-user servers running enterprise software. Examples are transaction processing databases, internet infrastructure (email, webserver, e-commerce), scientific computing software, and nearline storage management software. Enterprise drives commonly operate continuously ("24/7") in demanding environments while delivering the highest possible performance without sacrificing reliability.

The fastest enterprise HDDs spin at 10,000 RPM and 15,000 RPM, and can achieve sequential media transfer speeds above 290 MB/s. Drives running at 10,000 or 15,000 RPM use smaller platters to mitigate increased power requirements and therefore generally have lower capacity than the highest capacity 7,200 RPM drives. 10 K and 15 K drives are labeled Mission Critical or Performance Optimized, while 7.2 K are called Business Critical or Capacity Optimized. Since 7.2 K drives spin slower, they can have larger platters and space for more platters in an HDD case. That allows for higher capacity drives – 16 TB, 18 TB and so on.

Random read/write task speeds are usually measured in IOPs (Input/Output operations per second) and for 15 K drives can be up to 290. That may sound like a lot but is dwarfed by IOPS (in the hundreds of thousands) available on SSDs. Below is a link to a chart showing HDD performance characteristics:

HDD_Characteristic_and_Metrics

Depending on the workload type, storage device capabilities are prioritized differently. Below is discussed several capabilities and the media that you would select for each.

- Better performance measured in IOPS: Storage performance for random workloads is measured in IOPS. When ordered in terms of general IOPS performance, choose 15 K, and then 10 K, then 7.2 K and within those SAS HDDs, NL-SAS HDDs, and SATA HDDs.
- Better performance is measured in throughput, or gigabytes per second (GB/s): Unless dealing with heavy sequential workloads that would benefit from flash technology, HDDs are a good choice for most sequential workloads, such as media viewing or database logging. NAND caching can further boost the performance of HDD storage as necessary.
- Lower latency: For workloads sensitive to latency, internal storage on the server itself typically has less latency than storage on external arrays where longer fetch times over the network can greatly add to existing storage latency. It is important to note that SSDs present much lower latency than mechanical HDDs.
- Greater capacity: For capacity-driven workloads, such as email archives, disk-based backup, and object storage applications, high IOPS or throughput performance may be less of a priority as compared to capacity. In this case, choose cost-efficient HDDs, which can offer the greatest capacity at the lowest cost.

While HDDs generally provide lower performance and higher latency than SSDs, they are still an excellent option when used as part of a complete storage strategy that balances cost per GB, capacity, application needs, and performance.

Solid State Drives (SSDs)

SSD Feature Matrix

Table 13. SSD feature matrix

Type	Interface	Speed	From Factor	Endurance	Security	Capacity	Drive Description
SSD	SAS	24 Gbps	2.5	MU	ISE	1.6 TB	Agnostic SAS MU 1600GB
SSD	SAS	24 Gbps	2.5	MU	ISE	800 GB	Agnostic SAS MU 800GB
SSD	SAS	24 Gbps	2.5	MU	SED FIPS	1.92 TB	Kioxia PM6 FIPS MU 1920GB
SSD	SAS	24 Gbps	2.5	MU	SED FIPS	3.84 TB	Kioxia PM6 FIPS MU 3840GB
SSD	SAS	24 Gbps	2.5	MU	SED FIPS	960 GB	Kioxia PM6 FIPS MU 960GB
SSD	SAS	24 Gbps	2.5	RI	ISE	1.92 TB	Agnostic SAS RI 1920GB
SSD	SAS	24 Gbps	2.5	RI	ISE	3.84 TB	Agnostic SAS RI 3840GB
SSD	SAS	24 Gbps	2.5	RI	ISE	7.68 TB	Agnostic SAS RI 7680GB
SSD	SAS	24 Gbps	2.5	RI	SED FIPS	1.92 TB	Kioxia PM6 FIPS RI 1920GB
SSD	SAS	24 Gbps	2.5	RI	SED FIPS	3.84 TB	Kioxia PM6 FIPS RI 3840GB
SSD	SAS	24 Gbps	2.5	RI	SED FIPS	7.68 TB	Kioxia PM6 FIPS RI 7680GB
SSD	SATA	6 Gbps	2.5	MU	ISE	1.92 TB	Agnostic SATA MU 1920GB
SSD	SATA	6 Gbps	2.5	MU	ISE	480 GB	Agnostic SATA MU 480GB
SSD	SATA	6 Gbps	2.5	MU	ISE	960 GB	Agnostic SATA MU 960GB
SSD	SATA	6 Gbps	2.5	MU	ISE	3.84 TB	Agnostic SATA MU 3840GB
SSD	SATA	6 Gbps	2.5	RI	ISE	1.92 TB	Agnostic SATA RI 1920GB
SSD	SATA	6 Gbps	2.5	RI	ISE	3.84 TB	Agnostic SATA RI 3840GB
SSD	SATA	6 Gbps	2.5	RI	ISE	480 GB	Agnostic SATA RI 480GB
SSD	SATA	6 Gbps	2.5	RI	ISE	960 GB	Agnostic SATA RI 960GB
SSD	vSAS	12 Gbps	2.5	MU	SED	1.92 TB	Agnostic Value SAS SED MU 1920GB
SSD	vSAS	12 Gbps	2.5	MU	SED	3.84 TB	Agnostic Value SAS SED MU 3840GB
SSD	vSAS	12 Gbps	2.5	MU	SED	960 GB	Agnostic Value SAS SED MU 960GB
SSD	vSAS	12 Gbps	2.5	RI	SED	1.92 TB	Agnostic Value SAS SED RI 1920GB
SSD	vSAS	12 Gbps	2.5	RI	SED	7.68 TB	Agnostic Value SAS SED RI 7680GB
SSD	vSAS	12 Gbps	2.5	RI	SED	960 GB	Agnostic Value SAS SED RI 960GB

This document is updated as changes happen, so be sure to bookmark it rather than downloading an offline copy to stay with the latest information or see the Drive and Platform Matrix.

SSD Facts

Unlike hard disk drives (HDDs) which use a spinning platter to store data, solid state drives (SSDs) use solid state memory NAND flash. HDDs have several different mechanical moving parts which make them susceptible to vibrational and handling interference. Solid state drives, on the other hand have no moving parts and are therefore much less susceptible to vibrational or handling damage even when impacted during use.

SSDs deliver high-performance input/output operations per second (IOPS), and very low latency for transaction - intensive server and storage applications. Properly used in systems with HDDs, they reduce total cost of ownership (TCO) through low power consumption and low operating temperature.

Dell offers different solid-state drive (SSD) solutions to meet different customer needs. Enterprise SSDs, as a class, are unique compared to client or consumer-based SSD in terms of reliability, performance and architecture. While consumer-based SSDs, such as those utilized in notebooks are designed with a focus on consumer-based workloads, rigidity and battery life, enterprise-class SSDs are designed around enterprise application I/O (input/output) requirements with focus points of random I/O performance, reliability, and protection of data during a sudden power-down.

Understanding the basics of enterprise-class SSDs allow customers to make informed decisions when comparing solutions:

- Over-provisioning: The Achilles' heel of SSDs are their write characteristics. To rewrite an area of an SSD that has already been written, the data must be erased and then written. In order to overcome a portion of the write performance penalty, Dell enterprise SSDs found across Dell PowerEdge products, all employ a practice that is known as over-provisioning of Flash. This practice keeps native Flash capacity beyond the user-defined capacity and uses the additional space as a scratch pad of sorts to quickly put down application write data on areas of Flash that are already in an erased state. The SSDs perform cleanup functions of this over-provisioned Flash space during time periods typically not impacting application performance.
- Write Endurance: Write endurance is the number of program/erase (P/E or write cycles) that can be applied to a block of flash memory before the storage media becomes unreliable. Due to different data center workloads and read/write needs, Dell offers different enterprise SSDs with different endurance ratings so customers can design the right solution for their needs.

Below are the different categories (swim lanes) of enterprise SSDs Dell offers:

- Mixed Use (MU, 3 WPD): 70/30 read/write workloads with medium endurance. E-mail/messaging, OLTP, and E-commerce are example workloads.
- Read Intensive (RI, 1 WPD): 90/10 read/write workloads with lower endurance. Database warehousing, media streaming, and VOD solutions are example workloads.

Dell enterprise SSDs support five kinds of host interface options:

- NVMe SSD: NVMe SSDs are a mainstream, high-performance, high reliability solid-state storage device that enables IOPS performance of up to 2000x more than conventional rotating hard drives.
- Datacenter NVMe: datacenter NVMe SSDs share the same value proposition as NVMe SSDs, but with a reduction in cost at only a minor performance tradeoff as compared to NVMe.
- SAS SSD: SAS SSDs are based on the industry standard SAS interface. SAS SSDs combine superior reliability, data integrity, and data fail recovery making them suitable for enterprise applications.
- Value SAS: Value SAS is a new class of SAS SSD that leverages the PowerEdge SAS server infrastructure to deliver SAS like performance at a cost that is competitive with SATA.
- SATA SSD: SATA SSDs are based on the industry standard SATA interface. SATA SSDs provide reasonable performance for enterprise servers.

Dell Enterprise SSDs will support a new form factor in addition to several existing:

- E3.S: Part of the EDSFF family, E3.S is targeted to NVMe SSDs with x4 PCIe link widths though it can mechanically fit an x16 card edge. It supports power profiles up to 25 W and positioned to be a primary form factor for mainstream NVMe server storage subsystems as it can be used across a wide variety of platforms including modular and short depth chassis.

Internal storage configuration

PowerEdge R360 supports the following internal storage configurations:

- 8 x 2.5 inch SAS/SATA RAID
- 4 x 3.5 inch SAS/SATA RAID

NOTE: The PowerEdge R360 system does not support rear storage.

External Storage

The PowerEdge R360 supports the external storage device types that are listed in the table below.

Table 14. Supported External Storage Devices

Device Type	Description
External Tape	Supports connection to external USB tape products
NAS/IDM appliance software	Supports NAS software stack
JBOD	Supports connection to 12 Gb MD-series JBODs

7

Networking

Topics:

- Supported network cards
- Overview

Supported network cards

Table 15. Supported network cards

Vendor	Port type	Port speed	Port count
Broadcom	F1	1 GbE	4
Intel	BT	10 GbE	2
Intel	BT	10 GbE	4
Broadcom	BT	10 GbE	3
Broadcom	BT	10 GbE	4
Intel	F1	1 GbE	4

Overview

PowerEdge offers a wide variety of options to get information moving to and from our servers. Industry best technologies are chosen, and systems management features are added by our partners to firmware to tie in with iDRAC. These adapters are rigorously validated for worry-free, fully supported use in Dell servers.

8

PCIe subsystem

Topics:

- PCIe risers

PCIe risers

The PowerEdge R360 system supports a Gen4 butterfly riser on the system.

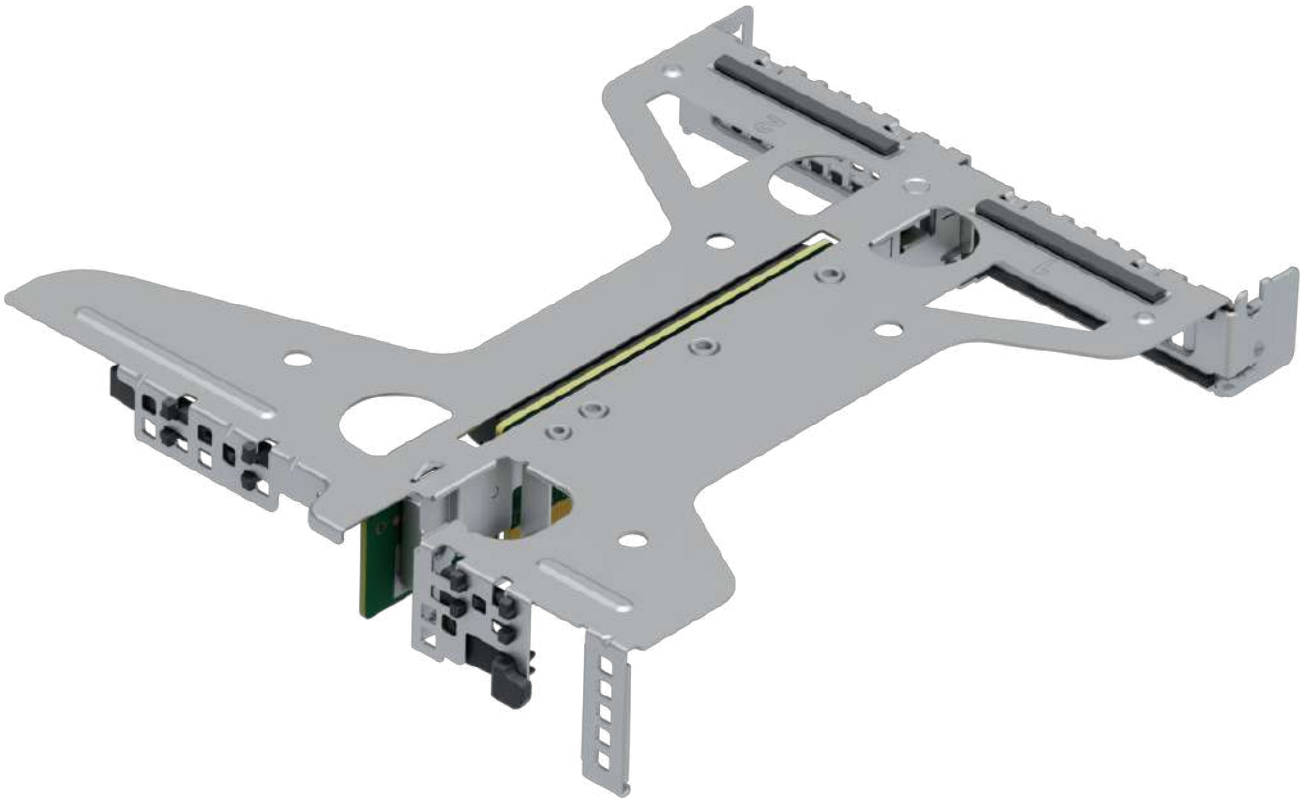


Figure 6. Butterfly Riser

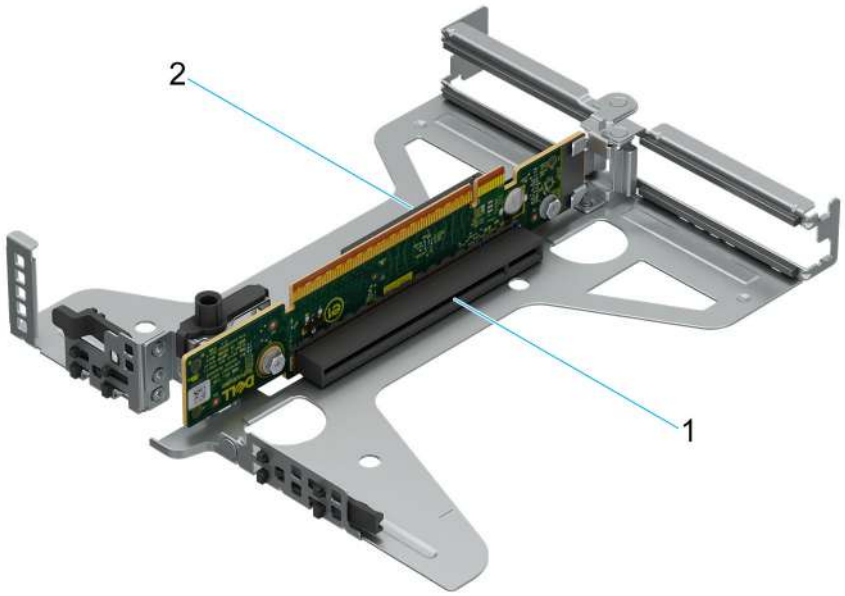


Figure 7. Butterfly Riser

1. PCIe Slot 2
2. PCIe Slot 1

Table 16. PCIe riser configurations

Configuration No.	Riser Configuration	No. of Processors	PERC type supported (internal slot)	Rear Storage Possible
RC0	1x PERC - No Riser	1	aPERC or fPERC	No
RC2	1x PERC + 1x Riser (Gen4)	1	aPERC or fPERC	No

9

Power, thermal, and acoustics

PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps to regulate temperature by reducing server noise and power consumption. The table below lists the tools and technologies Dell offers to lower power consumption and increase energy efficiency.

Topics:

- Power
- Thermal
- Acoustics

Power

Table 17. Power tools and technologies

Feature	Description
Power Supply Units(PSU) portfolio	Dell's PSU portfolio includes intelligent features such as dynamically optimizing efficiency while maintaining availability and redundancy. Find additional information in the Power supply units section.
Tools for right sizing	Enterprise Infrastructure Planning Tool (EIPT) is a tool that can help you determine the most efficient configuration possible. With Dell's EIPT, you can calculate the power consumption of your hardware, power infrastructure, and storage at a given workload. Learn more at Dell EIPT.
Industry Compliance	Dell's servers are compliant with all relevant industry certifications and guide lines, including 80 PLUS, Climate Savers and ENERGY STAR.
Power monitoring accuracy	PSU power monitoring improvements include: ● Dell's power monitoring accuracy is currently 1%, whereas the industry standard is 5% ● More accurate reporting of power ● Better performance under a power cap
Power capping	Use Dell's systems management to set the power cap limit for your systems to limit the output of a PSU and reduce system power consumption. Dell is the first hardware vendor to leverage Intel Node Manager for circuit-breaker fast capping.
Systems Management	iDRAC Enterprise and Datacenter provides server-level management that monitors, reports and controls power consumption at the processor, memory and system level. Dell OpenManage Power Center delivers group power management at the rack, row, and data center level for servers, power distribution units, and uninterruptible power supplies.
Active power management	Intel Node Manager is an embedded technology that provides individual server-level power reporting and power limiting functionality. Dell offers a complete power management solution comprised of Intel Node Manager accessed through Dell iDRAC9 Datacenter and OpenManage Power Center that allows policy-based management of power and thermal at the individual server, rack, and data center level. Hot spare reduces power consumption of redundant power supplies. Thermal control off a speed optimizes the thermal settings for your environment to reduce fan consumption and lower system power consumption. Idle power enables Dell servers to run as efficiently when idle as when at full workload.
Rack infrastructure	Dell offers some of the industry's highest-efficiency power infrastructure solutions, including:

Table 17. Power tools and technologies (continued)

Feature	Description
	• Power distribution units (PDUs) • Uninterruptible power supplies (UPSs) • Energy Smart containment rack enclosures Find additional information at: Power and Cooling.

Power Supply Units

Energy Smart power supplies have intelligent features, such as the ability to dynamically optimize efficiency while maintaining availability and redundancy. Also featured are enhanced power-consumption reduction technologies, such as high-efficiency power conversion and advanced thermal-management techniques, and embedded power-management features, including high-accuracy power monitoring. The table below shows the power supply unit options that are available for the PowerEdge R360.

Table 18. Power Supply Unit Options

Wattage	Frequency	Voltage/Current	Class	Heat dissipation
600 W	50/60 Hz	100-240 Vac/7.1 - 3.6 A	Platinum	2250 BTU/hr
	NA	240 Vdc/2.9 A	NA	
700 W	50/60 Hz	200-240 Vac/4.1 A	Titanium	2625 BTU/hr
	NA	240 Vdc/3.4 A	NA	

NOTE: If a system with AC 1100W or 1400W PSUs operate at low line 100 ~ 120 Vac, then the power rating per PSU is de-rated to 1050W.

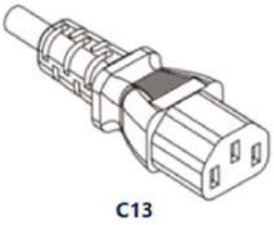


Figure 8. PSU power cord

Table 19. PSU power cords

Form factor	Output	Power cord
Redundant 60 mm	600 W	C13 /C14 (inlet)
	700 W	

Thermal

PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption.

Thermal design

Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.

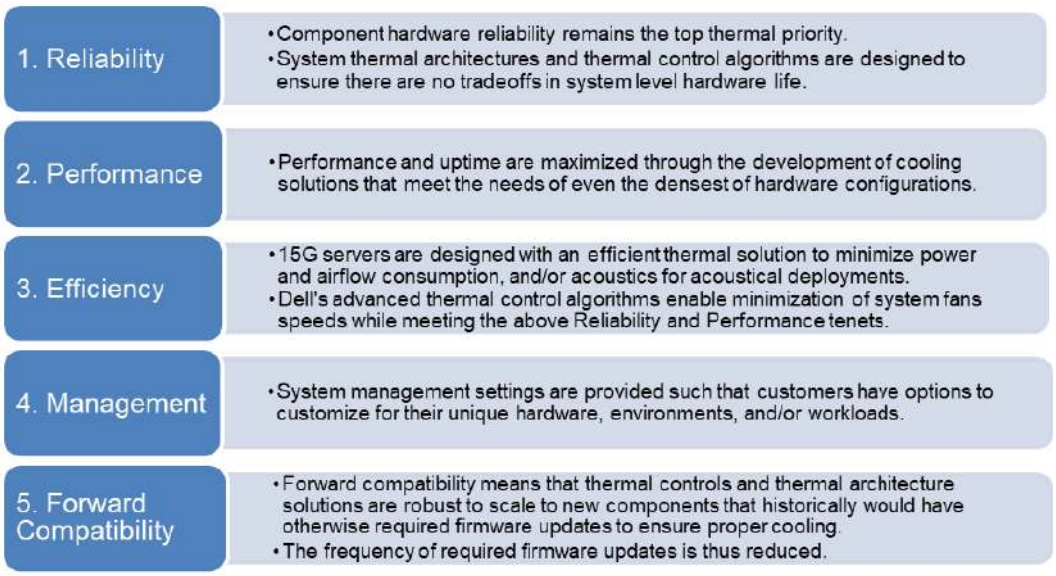


Figure 9. Thermal design characteristics

The thermal design of the PowerEdge R360 reflects the following:

- Optimized thermal design: The system layout is architected for optimum thermal design.
- System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power.
- Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP.
- Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds.
- User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R360 Installation and Service Manual at PowerEdge Manuals and “Advanced Thermal Control: Optimizing across Environments and Power Goals” on Dell.com.
- Cooling redundancy: The PowerEdge R360 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system.
- Environmental Specifications: The optimized thermal management makes the PowerEdge R360 reliable under a wide range of operating environments.

Acoustics

Acoustical performance

Dell PowerEdge R360 is a rack server appropriate for general use space or datacenter environment.

Table 20. Hardware and software configurations for lower acoustical output

Configuration	Quietest	Volume	Volume with BOSS	Feature Rich with GPU	Feature Rich with max storage
Processor Type	Raptor Lake processor	Raptor Lake processor	Raptor Lake processor	Raptor Lake processor	Raptor Lake processor
Processor TDP	65 W	65 W	65 W	95 W	95 W
Processor Quantity	1	1	1	1	1
RDIMM Memory	16 GB DDR5 UDIMM	16 GB DDR5 UDIMM	16 GB DDR5 UDIMM	32 GB DDR5 UDIMM	32 GB DDR5 UDIMM
Memory Quantity	1	2	2	4	4
Backplane Type	4 x 3.5-inch BP	4 x 3.5-inch BP	4 x 3.5-inch BP	8 x 2.5 inch BP	8 x 2.5 inch BP
HDD Type	3.5 inch SATA 2 TB HDD	3.5 inch SATA 2 TB HDD	3.5 inch SATA 2 TB HDD	2.5 inch SAS 600 GB HDD	2.5 inch SAS 600 GB HDD
HDD Quantity	2	2	2	8	8
PSU Type	600 W (60 mm)	600 W (60 mm)	600 W (60 mm)	600 W (60 mm)	600 W (60 mm)
PSU Quantity	2	2	2	2	2
PCIe Card	PERC H355	PERC H355	PERC H355	PERC H755 2x 1 GbE NIC A2 GPU (60 W)	PERC H755 2x 1 GbE NIC
Bezel	N/A	Yes	Yes	Yes	Yes
BOSS	N/A	N/A	Yes	N/A	N/A

Table 21. Acoustical performance of R360 acoustical configurations

Configuration		Quietest	Volume	Volume with BOSS	Feature Rich with GPU	Feature Rich with max storage
Acoustical Performance: Idle/ Operating @ 25 °C Ambient						
$L_{wA,m}$ (B)	Idle	5.1	5.1	5.5	5.7	5.1
	Operating	5.1	5.1	5.5	8.0	5.7
K_v (B)	Idle	0.4	0.4	0.4	0.4	0.4
	Operating	0.4	0.4	0.4	0.4	0.4
$L_{pA,m}$ (dB)	Idle	35	35	39	40	35
	Operating	35	35	39	65	41
Prominent tones		No prominent tones in Idle and Operating				
Acoustical Performance: Idle @ 28 °C Ambient						
$L_{wA,m}$ (B)		5.9	5.9	5.9	6.0	5.9

Table 21. Acoustical performance of R360 acoustical configurations (continued)

Configuration	Quietest	Volume	Volume with BOSS	Feature Rich with GPU	Feature Rich with max storage
K_v (B)	0.4	0.4	0.4	0.4	0.4
$L_{pA,m}$ (dB)	43	43	43	44	44
Acoustical Performance: Max. Loading @ 35 °C Ambient					
$L_{wA,m}$ (B)	7.1	5.9	5.9	8.0	7.7
K_v (B)	0.4	0.4	0.4	0.4	0.4
$L_{pA,m}$ (dB)	53	53	53	65	62

LwA,m: The declared mean A-weighted sound power level (LwA) is calculated per section 5.2 of ISO 9296 (2017) with data collected using the methods described in ISO 7779 (2010). Data presented here may not be fully compliant with ISO 7779.

LpA,m: The declared mean A-weighted emission sound pressure level is at the bystander position per section 5.3 of ISO 9296 (2017) and measured using methods described in ISO 7779 (2010). The system is placed in a 24U rack enclosure, 25 cm above a reflective floor. Data presented here may not be fully compliant with ISO 7779.

Prominent tones: Criteria of D.6 and D.11 of ECMA-74 (17th ed., Dec. 2019) are followed to determine if discrete tones are prominent and to report them, if so.

Idle mode: The steady-state condition in which the server is energized but not operating any intended function.

Operating mode: The maximum of the steady state acoustical output at 50% of CPU TDP or active HDDs per C.9.3.2 in ECMA-74 (17th ed., Dec. 2019)

Power Capping

The PowerEdge R360 supports the NVIDIA A2 GPU, which provides enterprise-level performance, therefore, louder acoustic performance is expected. Power capping solutions provides better acoustic performance by limiting GPU performance by up to 20%.

Table 22. Acoustical performance of R360 on power capping

R360	Without Power Capping	With Power Capping
Acoustic Performance	8.0 bels	6.3 bels
Acoustic Category	Category 6	Category 4

NOTE: PowerEdge R360 with GPU workload is not recommend for an acoustically sensitive environment.

10

Rack, rails, and cable management

Topics:

- Rack Rails

Rack Rails

The rail offerings for the PowerEdge R360 consist of two types — sliding and static. The sliding and static rail systems for the PowerEdge R360 provide tool-less support for four-post racks with tooled installation for square, round or threaded holes including all generations of Dell racks. Both also support tooled mounting in four-post threaded racks, and the static rails support tooled mounting in two-post (Telco) racks as well for added versatility.

The sliding rails for the PowerEdge R360 offer native support for threaded hole racks via the ReadyRails™ II mounting interface. The rails ship in the tool-less mounting configuration, but can be converted to the tooled configuration very quickly and easily as shown in the figure below.

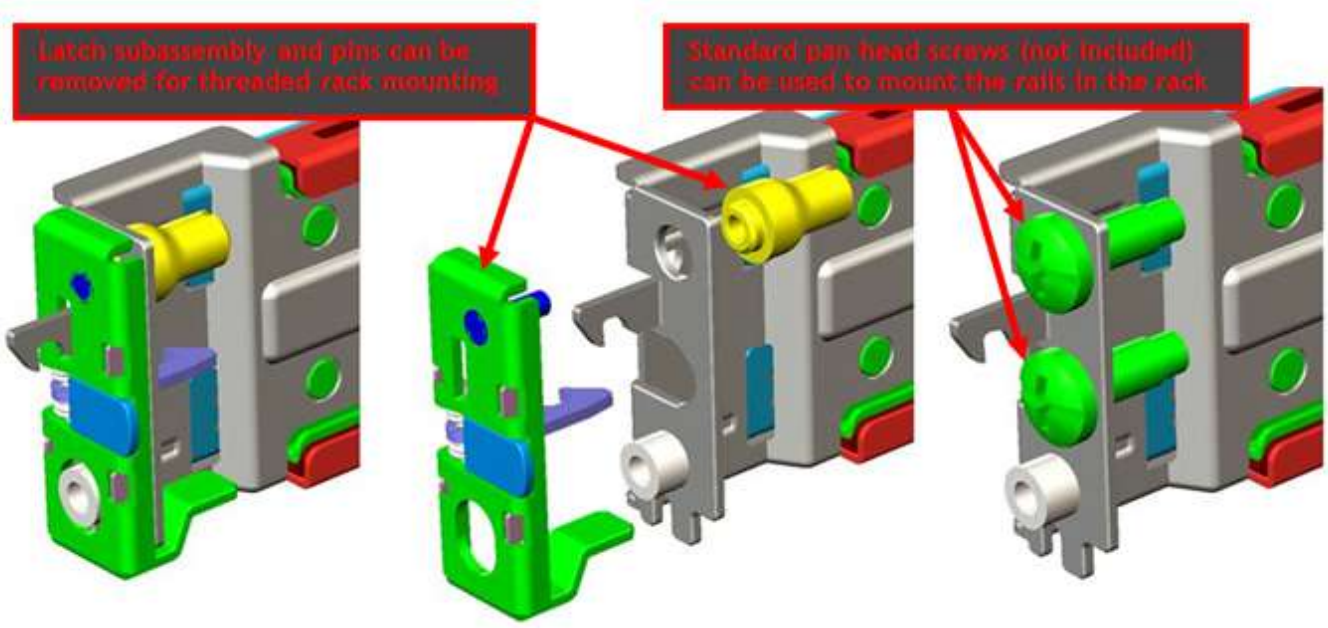


Figure 10. ReadyRails II Mounting Interface

One key factor in selecting the proper rails is identifying the type of rack in which they will be installed. Both the sliding rails and the static rails support tool-less mounting in 19-inch wide, EIA-310-E compliant square hole and unthreaded round hole four-post racks. Both also support tooled mounting in threaded hole four-post racks, but only the static rails, as the more “universal” solution, support mounting in two-post (Telco) racks.

Table 23. Sliding and static rail configurations

Sliding and Static Rail								
Product	Rail Identifier	Mounting Interface	Rail Type	Rack Types Supported				
				4–Post			2–Post	
				Square	Round	Thread	Flush	Center
PowerEdge R360	A12	ReadyRails II	Sliding	√	√	√[1]	X	X
	A8	ReadyRails	Static	√	√	√[1]	√[1]	√[1]

[1] Minor conversion required.

Note that screws are not included in either kit due to the fact that threaded racks are offered with a variety of thread designations. Users must therefore provide their own screws when mounting the rails in threaded racks.

NOTE: Screw head diameter for the sliding rails must be 10 mm or less.

Other key factors governing proper rail selection include the spacing between the front and rear mounting flanges of the rack, the type and location of any equipment mounted in the back of the rack such as power distribution units (PDUs), and the overall depth of the rack. The static rails offer a greater adjustability range and a smaller overall mounting footprint than the sliding rails. This is because of their reduced complexity and lack of need for CMA support.

Table 24. Rail adjustable range and rail depth specifications

Rail Adjustable Range and Rail Depth											
Product	Rail Identifier	Chassis configuration	Rail Type	Rail Adjustable Range (mm)						Rail Depth (mm)	
				Square*		Round*		Threaded*		without CMA	with CMA
				Min	Max	Min	Max	Min	Max		
PowerEdge R360	A12	2.5-inch	Sliding	631	868	617	861	631	883	720	845
		3.5-inch	Sliding	681	868	667	861	681	883	770	895
	A8		Static	608	879	594	872	618	898	622	-

* Values represent the distance between the front and rear mounting flanges on the rack.

NOTE: The adjustment range of the rails is a function of the type of rack in which they are being mounted. The Min/Max values listed above represent the allowable distance between the front and rear mounting flanges in the rack. Rail depth without the CMA represents the minimum depth of the rail with the outer CMA brackets removed (if applicable) as measured from the front mounting flanges of the rack.

NOTE: 3.5 inch HDD configurations utilize a self-adjusting slide feature. To fully install the system, this feature requires extra force during the last 50 mm travel.

Ready Rails II Sliding Rails for 4–Post Racks

- Support for tool-less installation in 19” EIA-310-E compliant square or unthreaded round hole four-post racks including all generations of Dell racks.
- Support for tooled installation in 19-inch EIA-310-E compliant threaded hole four-post racks.
- Support full extension of the system out of the rack to allow serviceability of key internal components.
- Support for optional cable management arm (CMA).
- Minimum rail mounting depth without the CMA (2.5-inch configuration): 720 mm
- Minimum rail mounting depth without the CMA (3.5-inch configuration): 770 mm
- Minimum rail mounting depth with the CMA (2.5-inch configuration): 845 mm
- Minimum rail mounting depth with the CMA (3.5-inch configuration): 895 mm
- Square-hole rack adjustment range (2.5-inch configuration): 631-868 mm
- Square-hole rack adjustment range (3.5-inch configuration): 681-868 mm
- Round-hole rack adjustment range (2.5-inch configuration): 617-861 mm
- Round-hole rack adjustment range (3.5-inch configuration): 667-861 mm

- Threaded-hole rack adjustment range (2.5-inch configuration): 631-883 mm
- Threaded-hole rack adjustment range (3.5-inch configuration): 681-883 mm

The sliding rails allow the system to be fully extended out of the rack for service. They are available with or without the optional cable management arm (CMA).

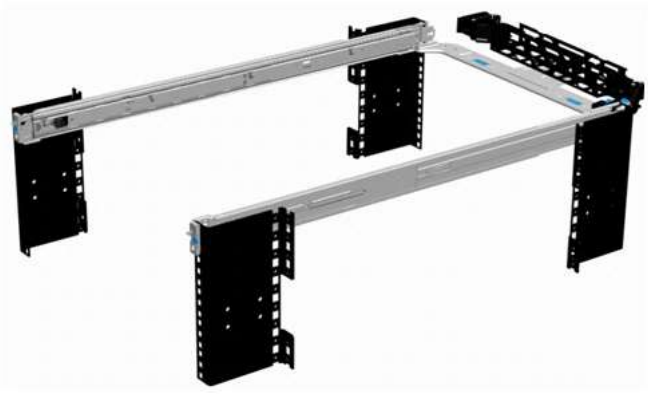


Figure 11. Sliding rails with optional CMA

Ready Rails Static Rails for 4-Post and 2-Post Racks

- Support tool-less installation in 19-inch EIA-310-E compliant square or unthreaded round hole four-post racks including all generations of Dell racks.
- Support tooled installation in 19-inch EIA-310-E compliant threaded hole four-post and two-post racks.
- Minimum rail mounting depth: 622 mm
- Square-hole rack adjustment range: 608-879 mm
- Round-hole rack adjustment range: 594-872 mm
- Threaded-hole rack adjustment range: 618-898 mm

The static rails support a wider variety of racks than the sliding rails. However, they do not support serviceability in the rack and are thus not compatible with the CMA.



Figure 12. Static rails

Cable Management Arm (CMA)

The optional cable management arm (CMA) for the PowerEdge R360 organizes and secures the cords and cables exiting the back of the server. It unfolds to allow the server to extend out of the rack without having to detach the cables. Some key features of the PowerEdge R360 CMA include:

- Large U-shaped baskets to support dense cable loads.
- Open vent pattern for optimal airflow.
- Ability to be mounted on either side by simply swinging the spring-loaded brackets from one side to the other.
- Utilizes hook-and-loop straps rather than plastic tie wraps to eliminate the risk of cable damage during cycling.
- Includes a low-profile fixed tray to both support and retain the CMA in its fully closed position.

- Both the CMA and the tray mount without the use of tools via simple and intuitive snap-in designs.

The CMA can be mounted to either side of the sliding rails without the use of tools or the need for conversion. However, it is recommended that it be mounted on the side opposite to the power supplies to allow easier access to the power supplies and rear hard drives (if applicable) for service or replacement.



Figure 13. System mounted in sliding rails with the CMA

The PowerEdge R360 static rails are a "stab-in" design. This means that the inner (chassis) rail members must first be attached to the sides of the system and then inserted into the outer (cabinet) members installed in the rack.



Figure 14. System mounted in static rails in two-post center-mount configuration

Rack Installation

The PowerEdge R360 sliding rails are a "drop-in" design. This means that the system is installed vertically into the rails by inserting the standoffs on the sides of the system into the "J-slots" in the inner rail members with the rails in the fully extended position. The recommended method of installation is to first insert the rear standoffs on the system into the rear J-slots on the rails in order to free up a hand and then rotate the system down into the remaining J-slots while using the free hand to hold the rail against the side of the system.

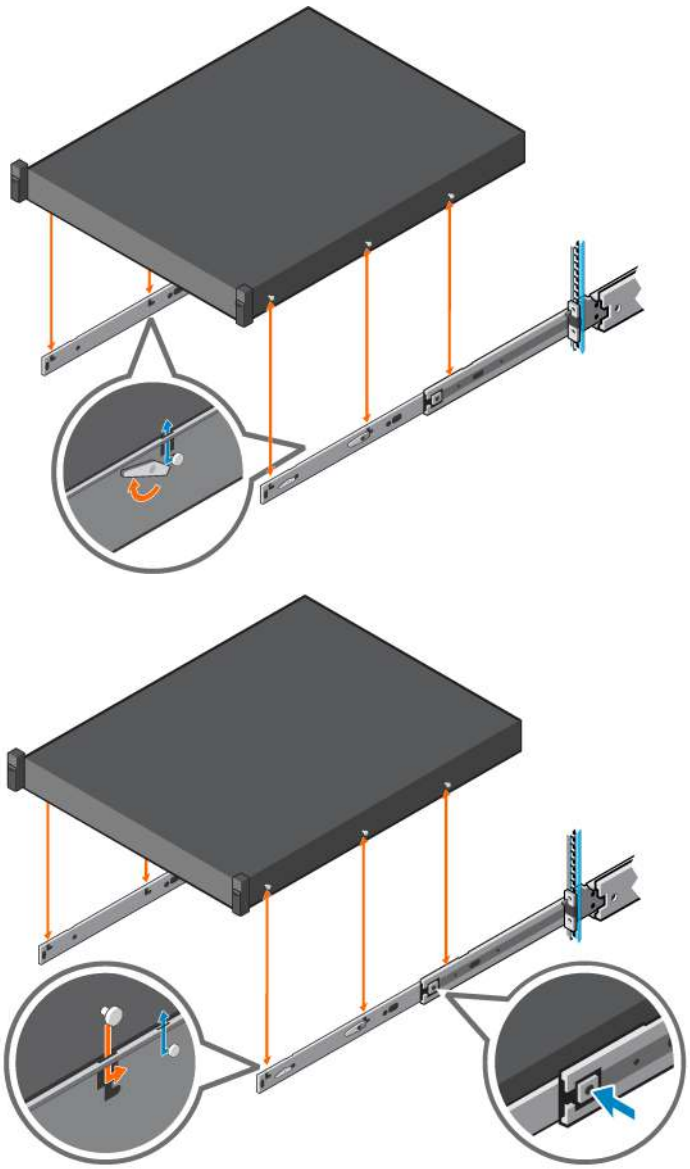


Figure 15. Sliding rails

Installing Systems in Sliding Rails

To install the system in sliding rails:

1. Pull the inner slide rails out of the rack until they lock into place.
   - Hold the system by its front and back ends and tilt the front end upward slightly.
   - Check the window on the sides of both rails to confirm that the heads of the standoffs are visible before adjusting or releasing your grip on the back end of the system.
   - Use your free hand to hold the rail against the side of the chassis if needed as the system is rotated downward into the remaining J-slots.
2. Locate the rear rail standoff on each side of the system and lower them into the rear J-slots of the slide assembly.
3. Rotate the system downward until all the rail standoffs are seated in the J-slots.
4. Push the system inward until the lock levers click into place. Press the slide-release lock buttons on both rails and slide the system into the rack.

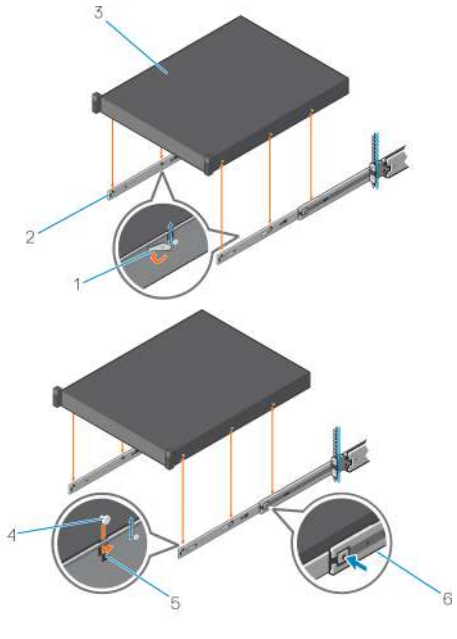


Figure 16. Sliding rails installation

1. Server-lock latch (2)
2. inner slide rail (2)
3. system
4. standoff on the system
5. J-slot on the rail
6. slide-release lock latch (2)

11

Operating Systems and Virtualization

Topics:

- Supported Operating Systems

Supported Operating Systems

The PowerEdge system supports the following operating systems:

- Canonical® Ubuntu® Server LTS
- Microsoft® Windows Server® with Hyper-V
- Red Hat® Enterprise Linux
- SUSE® Linux Enterprise server
- VMware® ESXi®

Links to specific OS versions and editions, certification matrices, Hardware Compatibility Lists (HCL) portal, and Hypervisor support are available at Dell Enterprise Operating Systems.

12

Dell Systems Management

Dell delivers management solutions that help IT administrators effectively deploy, update, monitor, and manage IT assets. Dell solutions and tools enable you to quickly respond to problems by helping them to manage Dell servers efficiently; in physical, virtual, local, and remote environments; all without the need to install an agent in the operating system.

The OpenManage portfolio includes:

- Innovative embedded management tools - integrated Dell Remote Access Controller (iDRAC)
- Consoles - OpenManage Enterprise
- Extensible with plug-ins - OpenManage Power Manager
- Update tools - Repository Manager

Dell has developed comprehensive systems management solutions that are based on open standards and has integrated with management consoles from partners such as Microsoft and VMware, allowing advanced management of Dell servers. Dell management capabilities extend to offerings from the industry's top systems management vendors and frameworks such as Ansible, Splunk, and ServiceNow. OpenManage tools automate the full span of server life cycle management activities along with powerful RESTful APIs to script or integrate with your choice of frameworks.

For more information about the entire OpenManage portfolio, see:

- The latest Dell Systems Management Overview Guide.

Topics:

- Integrated Dell Remote Access Controller (iDRAC)
- Systems Management software support matrix

Integrated Dell Remote Access Controller (iDRAC)

iDRAC9 delivers advanced, agent-free, local and remote server administration. Embedded in every PowerEdge server, iDRAC9 provides a secure means to automate a multitude of common management tasks. Because iDRAC is embedded within every PowerEdge server, there is no additional software to install; just plug in power and network cables, and iDRAC is ready to go. Even before installing an operating system (operating system) or hypervisor, IT administrators have a complete set of server management features at their fingertips.

With iDRAC9 in-place across the Dell PowerEdge portfolio, the same IT administration techniques and tools can be applied throughout. This consistent management platform allows easy scaling of PowerEdge servers as an organization's infrastructure grows. Customers can use the iDRAC RESTful API for the latest in scalable administration methods of PowerEdge servers. With this API, iDRAC enables support for the Redfish standard and enhances it with Dell extensions to optimize at-scale management of PowerEdge servers. By having iDRAC at the core, the entire OpenManage portfolio of Systems Management tools allows every customer to tailor an effective, affordable solution for any size environment.

Zero Touch Provisioning (ZTP) is embedded in iDRAC. ZTP - Zero Touch Provisioning is Intelligent Automation Dell's agent-free management puts IT administrators in control. Once a PowerEdge server is connected to power and networking, that system can be monitored and fully managed, whether you're standing in front of the server or remotely over a network. In fact, with no need for software agents, an IT administrator can: · Monitor · Manage · Update · Troubleshoot and remediate Dell servers With features like zero-touch deployment and provisioning, iDRAC Group Manager, and System Lockdown, iDRAC9 is purpose-built to make server administration quick and easy. For those customers whose existing management platform utilizes in-band management, Dell does provide iDRAC Service Module, a lightweight service that can interact with both iDRAC9 and the host operating system to support legacy management platforms.

When ordered with DHCP enabled from the factory, PowerEdge servers can be automatically configured when they are initially powered up and connected to your network. This process uses profile-based configurations that ensure each server is configured per your specifications. This feature requires an iDRAC Enterprise license.

iDRAC9 offers following license tiers:

Table 25. iDRAC9 license tiers

License	Description
iDRAC9 Basic	• Available only on 100-500 series rack/tower • Basic instrumentation with iDRAC web UI • For cost conscious customers that see limited value in management
iDRAC9 Express	• Default on 600+ series rack/tower, modular, and XR series • Includes all features of Basic • Expanded remote management and server life-cycle features
iDRAC9 Enterprise	• Available as an upsell on all servers • Includes all features of Basic and Express. Includes key features such as virtual console, AD/LDAP support, and more • Remote presence features with advanced, Enterprise-class, management capabilities
iDRAC9 Datacenter	• Available as an upsell on all servers • Includes all features of Basic, Express, and Enterprise. Includes key features such as telemetry streaming, Thermal Manage, automated certificate management, and more • Extended remote insight into server details, focused on high end server options, granular power, and thermal management

For a full list of iDRAC features by license tier, see Integrated Dell Remote Access Controller 9 User's Guide at Dell.com.

For more details on iDRAC9 including white papers and videos, see:

- Support for Integrated Dell Remote Access Controller 9 (iDRAC9) on the Knowledge Base page at Dell.com

Systems Management software support matrix

Table 26. Systems Management software support matrix

Categories	Features	PE mainstream
Embedded Management and In-band Services	iDRAC9 (Express, Enterprise, and Datacenter licenses)	Supported
	OpenManage Mobile	Supported
	OM Server Administrator (OMSA)	Supported
	iDRAC Service Module (iSM)	Supported
	Driver Pack	Supported
Change Management	Update Tools (Repository Manager, DSU, Catalogs)	Supported
	Server Update Utility	Supported
	Lifecycle Controller Driver Pack	Supported
	Bootable ISO	Supported
Console and Plug-ins	OpenManage Enterprise	Supported
	Power Manager Plug-in	Supported
	Update Manager Plug-in	Supported
	SupportAssist Plug-in	Supported
	CloudIQ	Supported
Integrations and connections	OM Integration with VMware Vcenter/vROps	Supported
	OM Integration with Microsoft System Center (OMIMSC)	Supported
	Integrations with Microsoft System Center and Windows Admin Center (WAC)	Supported

Table 26. Systems Management software support matrix (continued)

Categories	Features	PE mainstream
	ServiceNow	Supported
	Ansible	Supported
	Third-party Connectors (Nagios, Tivoli, Microfocus)	Supported
Security	Secure Enterprise Key Management	Supported
	Secure Component Verification	Supported
Standard operating system	Red Hat Enterprise Linux, SUSE, Windows Server 2019 or 2022, Ubuntu, CentOS	Supported (Tier-1)

13

Appendix D: Service and support

Topics:

- Why attach service contracts
- ProSupport Infrastructure Suite
- Specialty Support Services
- ProDeploy Infrastructure Suite
- Supplemental Deployment Services
- Unique Deployment Scenarios
- DAY 2 – Automation Services with Ansible
- Dell Technologies Consulting Services

Why attach service contracts

Dell PowerEdge servers include a standard hardware warranty that highlights our commitment to product quality by guaranteeing repair or replacement of defective components. While industry-leading, our warranties are limited to 1 or 3 years, depending on model, and do not cover software assistance. Call records show that failure rates for servers are roughly 1% and more commonly, customers seek Dell technical support for software-related issues like configuration guidance, troubleshooting, upgrade assistance, or performance tuning. Encourage customers to purchase ProSupport service contracts to supplement warranty coverage and ensure optimal support for both hardware and software. ProSupport provides a complete hardware guarantee beyond the original warranty period (up to 12 years: including seven years standard support and an additional five years of Post-Standard Support). Details of the ProSupport Suite and benefits are listed below.

ProSupport Infrastructure Suite

ProSupport Infrastructure Suite is a set of support services that enable customers to build the solution that is right for their organization. It is an industry-leading, enterprise-class support that aligns with the criticality of your systems, the complexity of your environment, and the allocation of your IT resources.

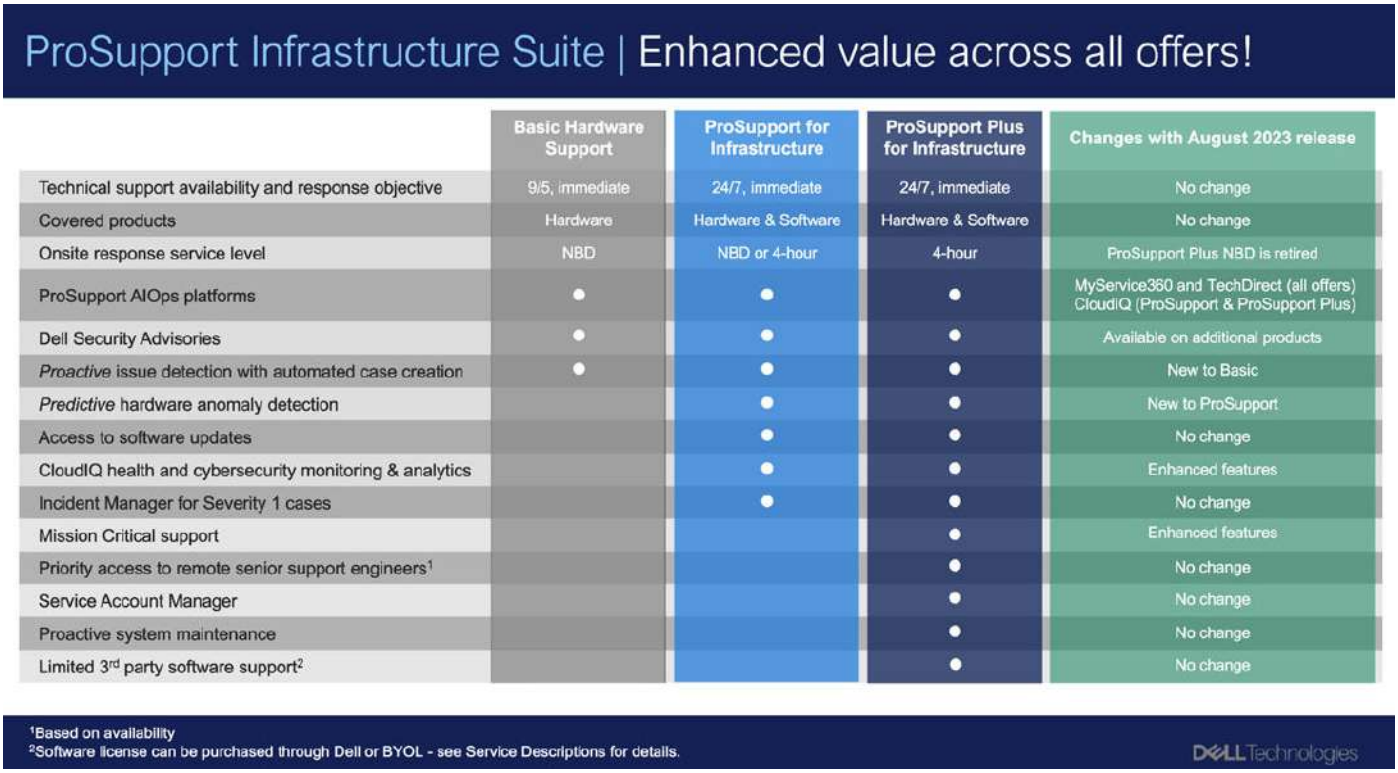


Figure 17. ProSupport Enterprise Suite

ProSupport Plus for Infrastructure

ProSupport Plus for Infrastructure is the ultimate solution for customers seeking preventative maintenance and optimal performance on their business-critical assets. The service caters to customers who require proactive, predictive, and personalized support for systems that manage critical business applications and workloads. When customers purchase PowerEdge server, we recommend ProSupport Plus, our proactive and preventative support service for business-critical systems. ProSupport Plus provides all the benefits of ProSupport, including the following "Top five reasons to buy ProSupport Plus (PSP)"

1. **Priority access to specialized support experts:** Immediate, advanced troubleshooting from an engineer that understands Dell infrastructure solutions.
2. **Mission Critical Support:** When critical (Severity 1) support issues happen, the customer is assured that we do all that we can to get them back up and running as quickly as possible.
3. **Service Account Manager:** A customer's #1 support advocate, ensuring they get the best possible proactive and predictive support experience.
4. **Systems maintenance:** On a semiannual basis, we will keep a customer's ProSupport Plus system(s) up to date by installing the latest firmware, BIOS, and driver updates to improve performance and availability.
5. **Third-party software support:** Dell is a customer's single point of accountability for any eligible third-partysoftware that is installed on their ProSupport Plus system, whether they purchased the software from us or not.

ProSupport for Infrastructure

Comprehensive 24x7 support for hardware and software – best for production, but not critical, workloads and applications. The ProSupport service offers highly trained experts around the clock and around the globe to address IT needs. We help minimize disruptions and maximize availability of PowerEdge server workloads with:

- 24x7 support through phone, chat and online
- A central point of accountability for all hardware and software issues
- Hypervisor, operating system and application support
- Dell security advisories
- Onsite response service levels 4 hour or Next Business Day options
- Proactive issue detection with automated case creation

- Predictive hardware anomaly detection
- Incident Manager assigned for Severity 1 cases
- Collaborative third-party support
- Access to AIOps Platforms - (MyService360, TechDirect, and CloudIQ)
- Consistent experience regardless of where customers are located or what language that they speak.

Basic Hardware Support

Provides reactive hardware support during normal business hours, excluding local national holidays. No software support orsoftware-related guidance. For improved levels of support, choose ProSupport or ProSupport Plus.

Specialty Support Services

Optional specialty support services complement the ProSupport Infrastructure Suite to provide additional proficiencies that are critical for modern data center operations.

Hardware coverage add-ons to ProSupport

- **Keep Your Hard Drive (KYHD), Keep Your Component (KYC), or Keep Your GPU:**

  Normally if a device fails under warranty, Dell replaces it using a one-for-one exchange process.KYHD/KYCC/KYGPU gives you the option to retain your device. It provides full control of sensitive data and minimizes security risk by letting you retain possession of failed drives, components, or GPU when receiving replacement parts without incurring additional cost.
- **Onsite Diagnosis Service:**

  Ideal for sites with non-technical staff. Dell field technician performs initial troubleshooting diagnosis onsite and transfers to Dell remote engineers to resolve the issue.
- **ProSupport Add-on for HPC:**

  Sold as an add-on to a ProSupport service contract, the ProSupport Add-on for HPC provides solution-aware support to cover the additional requirements that are required to maintain an HPC environment such as:

  - Access to senior HPC experts
  - Advanced HPC cluster assistance: performance, interoperability, and configuration
  - Enhanced HPC solution level end-to-end support
  - Remote pre-support engagement with HPC Specialists during ProDeploy implementation
- **ProSupport Add-on for Telco (Respond & Restore):**

  An add-on service designed for the top 31 TELCO customers globally, Respond & Restore provides direct access to Dell solution experts who specialize in TELCO carrier-grade support. This add-on also provides a hardware uptime guarantee, meaning if a system fails, Dell has it installed and operational within 4 hours for Severity 1 issues. Dell incurs penalties and fees if SLAs are not met.

Personalized Support and Supplemental Site-wide Expertise

- **Technical Account Manager:**

  Designated technology lead who monitors and manages the performance and configuration of specific technology sets.
- **Designated Remote Support:**

  Personalized support expert who manages all troubleshooting and resolution of IT assets.
- **Multivendor Support Service:**

  Support your third-party devices as one service plan for servers, storage, and networking (includes coverage for: Broadcom, Cisco, Fujitsu, HPE, Hitachi, Huawei, IBM, Lenovo, NetApp, Oracle, Quanta, SuperMicro and others).

Services for large enterprises

- **ProSupport One for Data Center:**

  ProSupport One for Data Center offers flexible site-wide support for large and distributed data centers with more than 1,000 assets (combined total of server, storage, networking, so forth). This offering is built on standard ProSupport features that leverage our global scale and are tailored to specific customer needs. While not for everyone, this service option offers a truly unique solution for our largest customers with the most complex environments.

  - Team of assigned Services Account Managers with remote or onsite options
  - Assigned technical and field engineers who are trained on the customer's environment and configurations.
  - On-demand reporting and recommendations that are enabled by ProSupport AIOps tools (MyService360, TechDirect, and CloudIQ)
  - Flexible onsite support and parts options that fit their operational model
  - A tailored support plan and training for their operations staff
- **ProSupport One for CSPs (Cloud Serviced Providers)**

  ProSupport One for CSPs is a unique offer that is designed for a limited set of Dell accounts purchasing Gen AI computing solutions greater than 1,000 servers and $250M in sales. PS1 for CSPs improves the entire services experience combining support, deployment (rack integration), residency services, a designated support engineer and the LOIS parts locker as one holistic bundle. Special pricing has been determined to compete effectively against competitors and provide the best customer experience. PS1 for CSPs can only be sold with XE Servers and all networking platforms (Dell and NVIDIA). All other products would be eligible for the standard PS1DC not this special bundle offer. More details on PS1 for CSPs here.
- **Logistics Online Inventory Solution (LOIS)**

  Ideal for large organizations that have their own staff to support their data center. Dell offers a service that is called Logistics Online Inventory Solution which is an onsite parts locker that provides self-maintainers with a local inventory of common replacement components. Having access to these parts lockers allows the self-maintainer to replace a failed component immediately without delay. Each replacement part would automatically initiate a replenishment of the parts inventory that is shipped the next day or delivered onsite by Dell during a regular scheduled visit (called Scheduled Onsite Service). As part of the LOIS system, customers can integrate their systems directly to Dell TechDirect using APIs to help streamline the support management process.

End-of-Life Services

- **Post Standard Support (PSS)**

  Extend service life beyond the initial seven years of ProSupport, adding up to five more additional years of hardware coverage.
- **Data Sanitization & Data Destruction**

  Renders data unrecoverable on repurposed or retired products, ensuring security of sensitive data and enabling compliance and provides NIST-compliant certification.
- **Asset Recovery Services**

  Recycle, resale, and disposal of hardware. Helps you securely and responsibly retire IT assets that are no longer needed while protecting both your business and the planet.

ProDeploy Infrastructure Suite

ProDeploy Infrastructure Suite provides various deployment offerings that satisfy a customer's unique needs. It is made up of 5 offers: ProDeploy Configuration Services, ProDeploy Rack Integration Services, Basic Deployment, ProDeploy, and ProDeploy Plus.

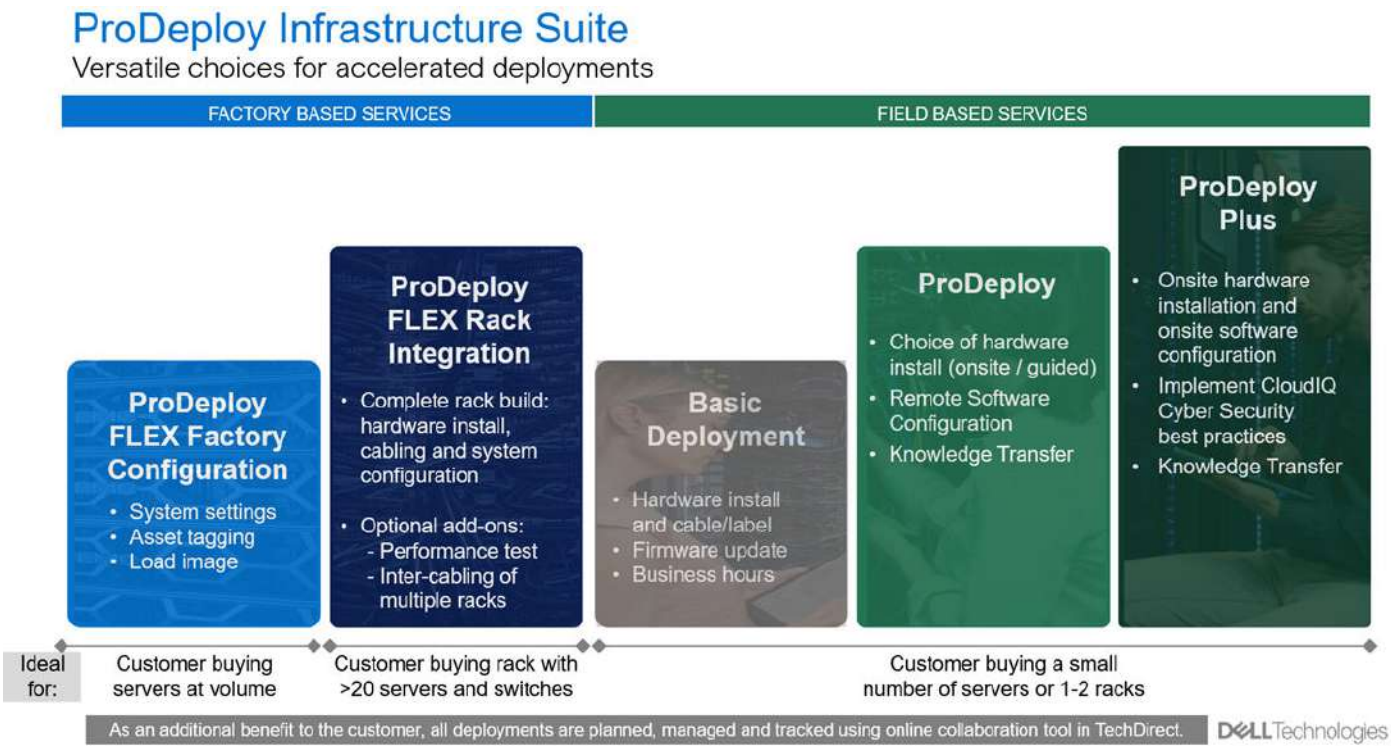


Figure 18. ProDeploy Infrastructure Suite

Factory-based Services

The new Factory Services consist of two tiers of deployment that happen prior to shipping to the Customer's site.

ProDeploy FLEX FactoryConfiguration

Ideal for Customers buying servers in volume and seeking pre-configuration prior to shipping such as: custom image, system settings, and asset tagging so it arrives ready to use out of the box. Furthermore, servers are packaged and bundled to meet specific shipping and distribution requirements for each customer location to facilitate the rollout process. Once the server is onsite , Dell can install and configure the server to the environment using any of the field-based deployment services outlined in the next section.

ProDeploy FLEX Rack Integration

Ideal for customers seeking to build out fully integrated racks prior to shipping. These rack builds include hardware install, cabling, and full system configuration. You can also add-on a factory stress test and optional on-site final rack configuration to complete the rack installation.

- STANDARD SKUs for Rack Integration is available in USA only and requires:
  - 20 or more devices (R and C series servers, VxRail, and all Dell or non-Dell switches.
  - Shipping to contiguous USA.
- USE CUSTOM QUOTE for Rack Integration scenarios that require:
  - Shipment to any country or region outside USA or shipping outside contiguous USA
  - Shipping to multiple locations
  - Racks containing fewer than 20 servers
  - Any rack that includes Storage.

ProDeploy Flex | Modular deployment (built in factory, onsite or remote)

Pre -deployment	Single point of contact for project management	●
	Expanded end-to-end project management	Selectable
	Site readiness review and implementation planning	●
Deployment	Deployment service hours	24/7
	Hardware installation options [1]	Onsite, factory[2,5] or remote[3]
	System software installation and configuration options [1]	Onsite, factory[2,5] or remote[3]
	Multivendor networking deployment [4]	Onsite, factory[2,5] or remote[3]
	Onsite Deployment in remote locations	Selectable
	Onsite Deployment in challenging environments	Selectable
	Onsite Deployment with special site-based protocols or requirements	Selectable
	Install connectivity software based on Secure Connect Gateway technology	●
	Dell NativeEdge Orchestrator deployment	Selectable
	Configure 3rd party software applications and workloads [4]	Selectable
Post -deployment	Deployment verification, documentation, and knowledge transfer	●
	Configuration data transfer to Dell support	●
Online collaboration	Online collaborative environment - Planning, managing and tracking delivery process	●

[1] Hardware and Software delivery methods can be independently chosen; selecting Rack integration for software requires hardware Rack integration to also be selected.
[2] Factory Rack Integration for server and VxRail; includes associated Dell network switches; final onsite rack installation available.
[3] Remote hardware option includes project specific instructions, documentation and live expert guidance for hardware installation.
[4] Select 3rd party multivendor networking and software applications.
[5] Pair with Field Onsite Hardware service for final installation.

Figure 19. ProDeploy Flex modular services

Field-based services

- **ProDeploy Plus:**

  Elevate Infrastructure deployments with our most complete service from planning through onsite hardware installation and software configuration including the implementation of cybersecurity best practices. ProDeploy Plus provides the skill and scale that is needed to successfully perform demanding deployments in today's complex IT environments. The deployment starts with a site readiness review and implementation plan. Certified deployment experts perform the software configuration to include setup of leading operating systems and hypervisors. Dell will also configure PowerEdge software tools to include iDRAC and OpenManage system utilities and support AIOps platforms: MyService360, TechDirect, and CloudIQ. Unique to ProDeploy Plus, the cybersecurity implementation helps customers understand potential security risks and make recommendations for reducing product attack surfaces. The system is tested, validated prior to completion. The customer will also receive full project documentation and knowledge transfer to complete the process.

- **ProDeploy:**

  ProDeploy provides remote software configuration and choice of hardware installation (onsite or guided). ProDeploy is great for customers who are price sensitive or willing to participate in some portion of the deployment to include providing remote access to their network. The ProDeploy remote software implementation includes everything that is mentioned in ProDeploy Plus except it does not include the added value, cybersecurity implementation and best practices.

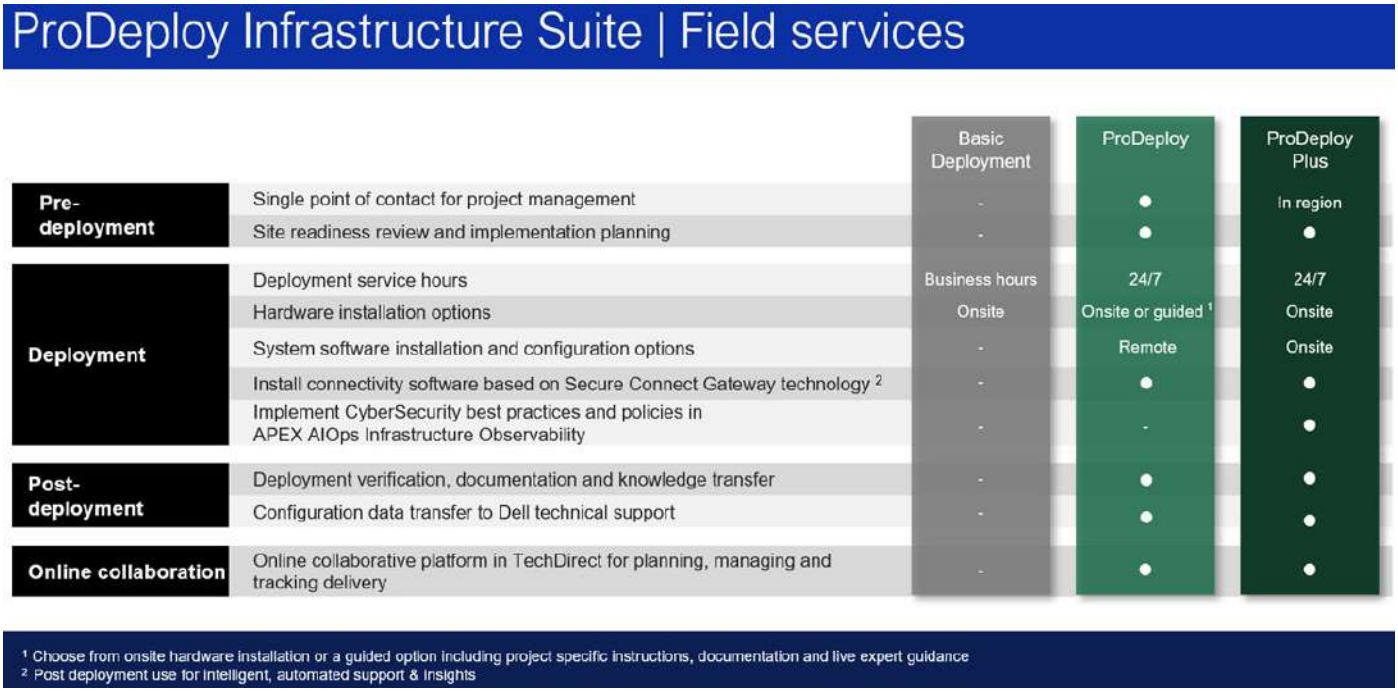


Figure 20. ProDeploy Infrastructure Suite - Field services

Supplemental Deployment Services

Additional ways to expand scope or deploy for unique scenarios.

Two Host Adder (requires PD/PDP)

Deploying new storage, compute, or networking devices may require interconnection to other servers (also called hosts). The Dell delivery team will set up four hosts per device as part of every ProDeploy service. For example, if the customer is buying two storage arrays the ProDeploy service will automatically include connectivity of four hosts each (4x2=8 total hosts per project since there are two devices). This supplemental "Two Host Adder" service provides for the configuration of additional hosts above what is already provided as part of the ProDeploy service. In many cases, customers can work with us while we set up the included hosts, so they may understand how to do the rest themselves. Always ask the customer how many hosts are being connected and sell the host adder depending on the customer's technology skillset. Note that this service applies to the connectivity of Dell devices not 3rd party devices.

Additional Deployment Services (ADT) – sold with or without PD/PDP

You can expand the scope of a ProDeploy engagement leveraging Additional Deployment Time (ADT). ADT covers additional tasks above the normal deliverables of the ProDeploy offers. ADT can also be used as a standalone service without ProDeploy. SKUs are available for both Project Management and Technical Resource Expertise. SKUs are sold as blocks of four hours remote or eight hours onsite. The delivery team can help in scoping the number of hours required for additional tasks.

Data Migration Services

Migrating data sets is no easy task. Our experts use proven tools and process to streamline data migrations and avoid compromising data. A customer project manager works with our experienced team of experts to create a migration plan. Data migration is part of every technology upgrade, platform change, and shift to the cloud. You can rely on Dell data migration services to perform a seamless transition.

Residency Services

Certified technical professionals act like an extension of your IT staff to enhance internal capabilities and resources and help you realize faster adoption and maximized ROI of new technology. Residency Services help customers transition to new capabilities quickly by leveraging specific technology skill sets. Residency experts can provide post implementation management and knowledge transfer that is related to a new technology acquisition or day-to-day operational management of the IT infrastructure.

- Global experts available to serve in-person (onsite) or virtual (remote)
- Engagements starting at 2 weeks with flexibility to adjust
- Residency is available for project management needs, and many different technology skills sets such as: Server, storage, Gen AI, networking, security, multi-cloud, data mgmt., and modern workforce application residents

Unique Deployment Scenarios

Custom Deployment Services

When a deployment is beyond the scope of the ProDeploy Infrastructure Suite, you can turn to the custom deployment services team to address complex implementation scenarios and customer unique requirements. The Dell custom deployment team is staffed with solution architects who will assist with customer scoping calls to define the project and develop the statement of work. Custom services can handle a wide range of deployments that can be performed in the factory or onsite. All custom engagement services are requested through SFDC.

ProDeploy FLEX

ProDeploy Flex is a modular service and a powerful tool for you to attach more services and improve revenue and margins. The ProDeploy Flex modular offer allows sales teams to build and better tailor services by mixing factory and field delivery options. You can also select special deployment scenarios without going to the custom order desk. FLEX is ideal for unique deployments where ProDeploy or ProDeploy Plus are not an adequate answer to the customer needs. Key features of ProDeploy FLEX :

- Build deployment quotes using modular, selectable features for both hardware and software.
- The system automatically scales pricing based on volume.
- Ideal for customers who require NativeEdge Orchestrator or edge deployments.
- Ability to add deployment services to third-party networking devices.

Deployment of HPC

High-Performance Computing (HPC) implementations require specialists that understand advanced feature sets. Dell deploys the world 's fastest systems and understands the nuances that make them perform. HPC deployments are most often scoped as custom service engagements, however we can do smaller HPC clusters under 300 nodes using a standard ProDeploy SKU. Any standard SKU for HPC deployment will be sold as one base SKU per cluster (ProDeploy for HPC Base) along with one ProDeploy for HPC Add-on for each device in the cluster (server nodes and switches).

Scope of ProDeploy for HPC:

NOTE: Available as standard SKUs in US and Canada. Custom Service would be required for all other regions.

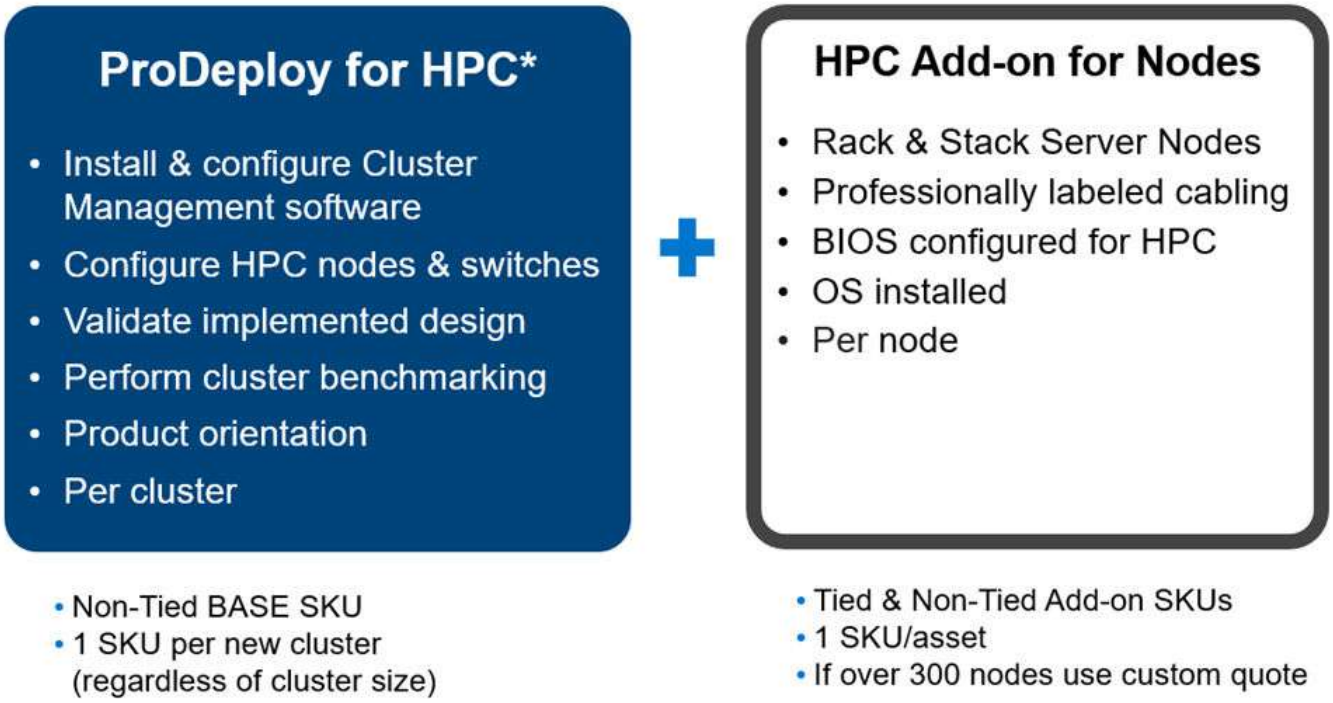


Figure 21. Standard deliverables of ProDeploy for HPC

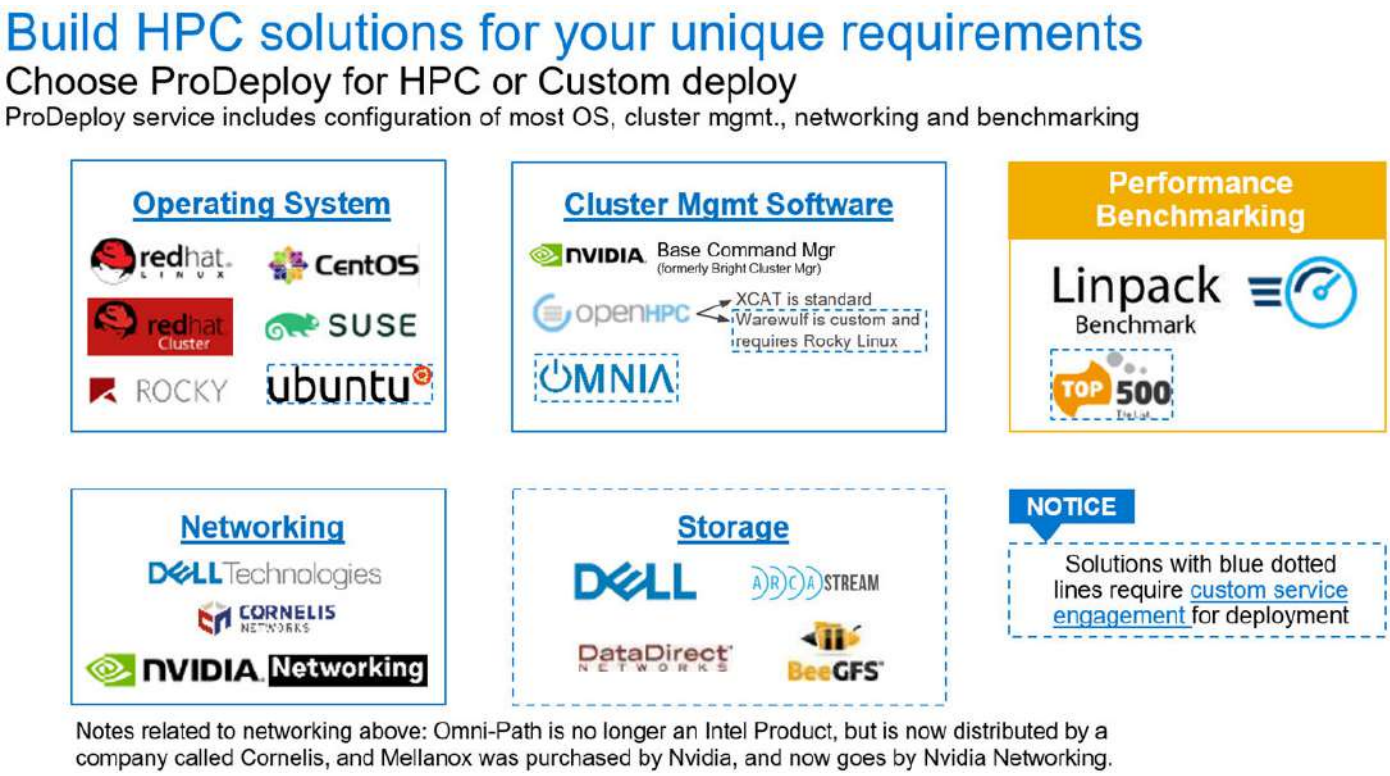


Figure 22. Visual view of HPC deployment options to include hardware and software

DAY 2 – Automation Services with Ansible

Dell solutions are built as "automation ready" with integrated APIs (Application Programming Interfaces) to allow customers to programmatically call actions on the product through code. Although Dell has published Anisble automation use cases, some customers need additional assistance with GitOps. By the end of the service, the customer will have the foundational

components required to accelerate automation and understand how the programming works together: Day 1 and Day 2 use case automation scripts (ansible modules), CI/CD tool (Jenkins), and Version control (Git).

Dell Technologies Consulting Services

Our expert consultants help customers transform faster, and quickly achieve business outcomes for the high value workloads Dell PowerEdge systems can handle. From strategy to full-scale implementation, Dell Technologies Consulting can help determine how to perform IT, workforce, or application transformation. We use prescriptive approaches and proven methodologies that are combined with portfolio and partner ecosystem of Dell Technologies to help achieve real business outcomes. From multi cloud, applications, DevOps, and infrastructure transformations, to business resiliency, data center modernization, analytics, workforce collaboration, and user experiences-we are here to help.

Dell Managed Services

Some customers prefer Dell to manage the complexity and risk of daily IT operations, Dell Managed Services utilizes proactive, AI enabled delivery operations and modern automation to help customers realize desired business outcomes from their infrastructure investments. With these technologies, our experts run, update and fine-tune customer environments aligned with service levels, while providing environment-wide and down-to-the-device visibility. There are two types of managed service offers. First the outsourcing model or CAPEX model where Dell manages the customer owned assets using our people and tools. The second is the as-a-Service model or OPEX model called APEX. In this service, Dell owns all technology and all the management of it. Many customers will have a blend of the two management types depending on the goals of the organization.



Figure 23. Dell Managed Services

Managed Detection and Response (MDR)

Dell Technologies Managed Detection and Response (MDR) is powered by Secureworks Taegis XDR software platform. MDR is a managed service that secures the customer's IT environment against malicious actors and provides remediation if and when a threat is identified. When a customer purchases MDR, they will receive the following features from our team:

- Dell badge resources
- Agent rollout assistance to help deploy the Secureworks Endpoint Agent
- 24x7 threat detection & investigation
- Up to 40hrs per quarter of response and active remediation activities
- If the customer experiences a breach, we will provide up to 40hrs per year of Cyber incident response initiation
- Quarterly reviews with the customer to review the data

Dell Technologies Education Services

Build the IT skills required to influence the transformational outcomes of the business. Enable talent and empower teams with the right skills to lead and perform transformational strategy that drives competitive advantage. Leverage the training and certification required for real transformation.

Dell Technologies Education Services offers PowerEdge server training and certifications that are designed to help customers achieve more from their hardware investment. The curriculum delivers the information and the practical, firsthand skills that their team must confidently install, configure, manage, and troubleshoot Dell servers.

To learn more or register for a class today, see Education.Dell.com.

14

Appendix A: Additional specifications

Topics:

- Chassis dimensions
- Chassis weight
- NIC port specifications
- Video specifications
- USB Ports
- PSU rating
- Environmental specifications

Chassis dimensions

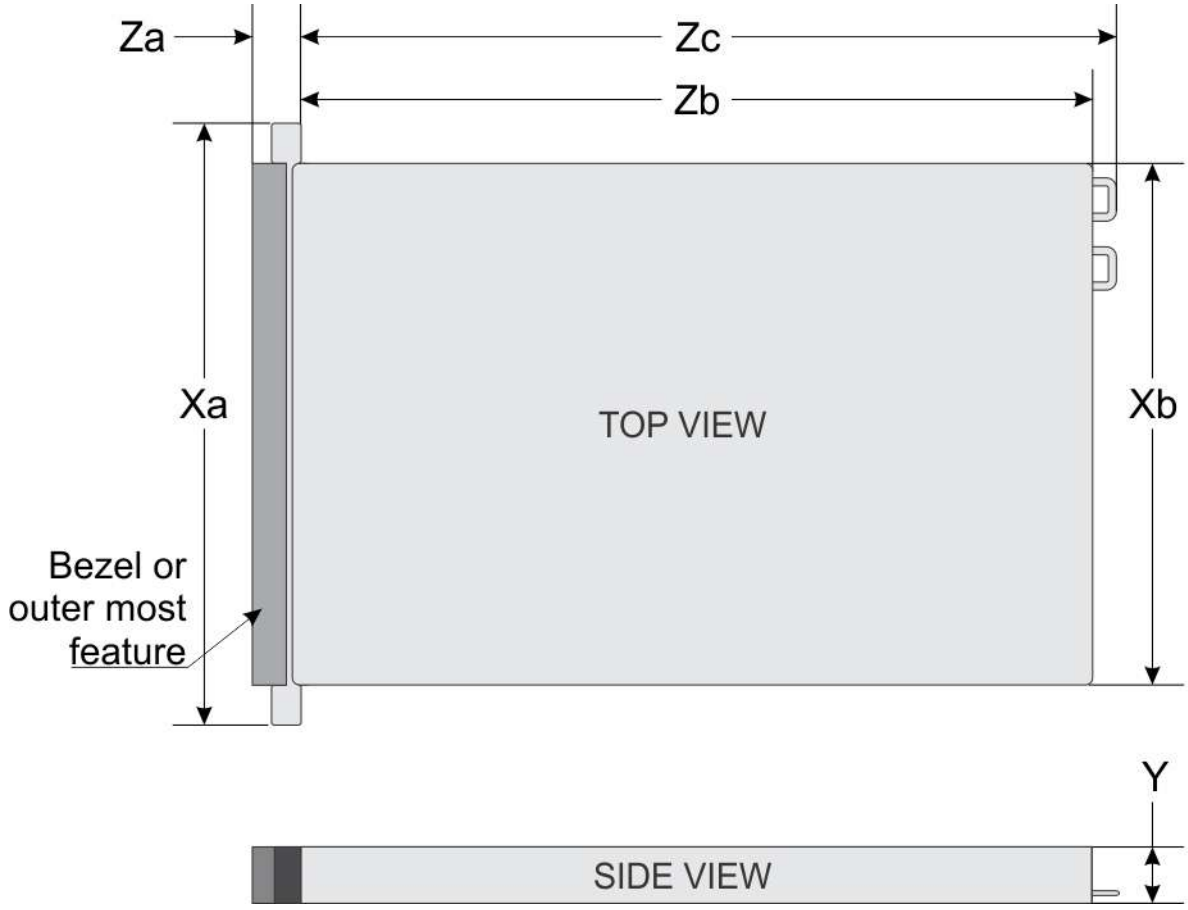


Figure 24. Chassis dimensions

Table 27. PowerEdge R360 chassis dimensions

Drives	Xa	Xb	Y	Za	Zb	Zc
12 drives	482.0 mm (18.976 inches)	434.0 mm (17.086 inches) CC:	42.8 mm (1.685 inches)	35 mm (1.378 inches)	483.82 mm (19.048 inches)	512.53 mm (20.178

Table 27. PowerEdge R360 chassis dimensions

Drives	Xa	Xb	Y	Za	Zb	Zc
		434.0 mm (17.08 inches)			Ear to PSU Surface	inches) Ear to PSU handle

NOTE: Zb is the nominal rear wall external surface where the system board I/O connectors reside.

NOTE: Xb is unique to this system and has two values for the Storage module (SM) and the Compute Module (CC) .

Chassis weight

Table 28. PowerEdge R360 system weight

System configuration	Maximum weight (with all drives/SSDs)	
A server with fully populated drives	2.5 inch drives	11.64 kg (25.68 pounds)
	3.5 inch drives	13.23 kg (29.17 pounds)
A server without drives and PSU installed	2.5 inch drives	8.36 kg (18.45 pounds)
	3.5 inch drives	9.01 kg (19.88 pounds)

NIC port specifications

The system supports up to two 10/100/1000 Mbps Network Interface Controller (NIC) ports embedded on the LAN on Motherboard (LOM) and up to four ports integrated on the optional network cards.

Table 29. NIC port specification for the system

Feature	Specifications
LOM	1 Gb x 2
Network Card (optional)	1 Gb x 4, 10 Gb x 2, 10 Gb x 4

Video specifications

The platform supports the following video resolution and refresh rates:

Table 30. Video specifications for R360

Resolution	Refresh Rate	Sestotal Freq.	Pixel Clock	DVO DisplayPort
1024 x 768	60 Hz	48.4 kHz	65.0 MHz	Yes*
1280 x 800	60 Hz	49.7 kHz	83.5 MHz	Yes*
1280 x 1024	60 Hz	64.0 kHz	108.0 MHz	Yes*
1360 x 768	60 Hz	47.71 kHz	85.5 MHz	Yes*
1440 x 900	60 Hz	55.9 kHz	106.5 MHz	Yes*
1600 x 900	60 Hz	55.54 kHz	97.75 MHz	Yes*
1600 x 1200	60 Hz	75.0 kHz	162.0 MHz	Yes*
1680 x 1050	60 Hz	64.7 kHz	119.0 MHz	Yes*
1920 x 1080	60 Hz (RB)	67.158 kHz	173.0 MHz	No

Table 30. Video specifications for R360 (continued)

Resolution	Refresh Rate	Sestotal Freq.	Pixel Clock	DVO DisplayPort
1920 x 1200	60 Hz (RB)	74.556 kHz	193.25 MHz	No

*DVO - DP is for investigation only, dependent on Nuvoton DVO capabilities to support up to 165 MHz.

*(RB) - Reduced Blanking for Digital Displays requiring less blank time. It was introduced for Signal Integrity improvements by reducing Pixel Clock rates for VGA- Analog input devices.

USB Ports



Figure 25. Front USB port

1. USB 2.0 Port



Figure 26. Rear USB ports

1. USB 2.0 Port
2. USB 3.2 Gen1

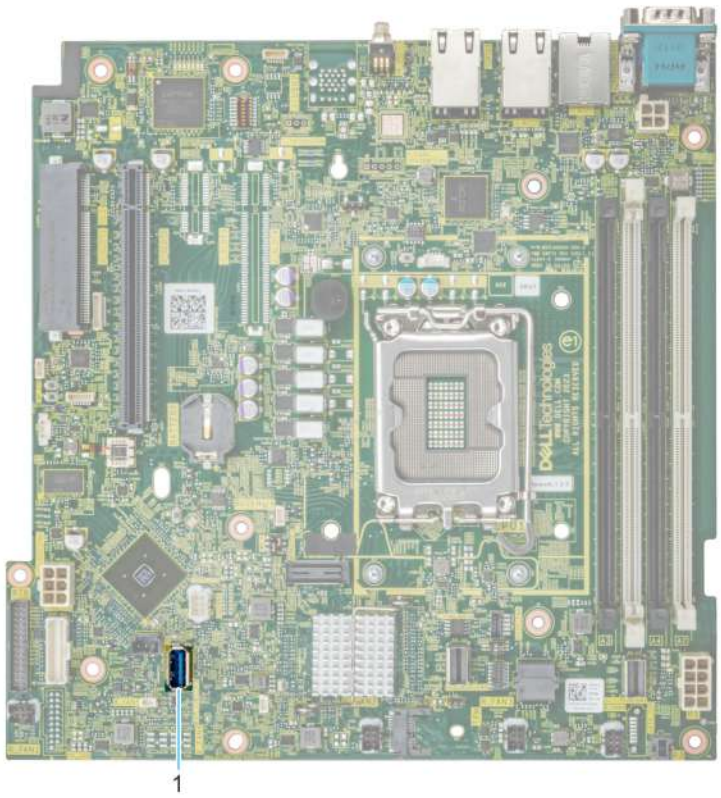


Figure 27. Internal USB Port

1. USB 3.2 Gen1

Table 31. Systems USB Specifications

Front		Rear		Internal	
USB port type	**No. of ports**	**USB port type**	**No. of ports**	**USB port type**	**No. of ports**
USB 2.0-compliant port	One	USB 2.0-compliant port	One	Internal USB 3.2 Gen1-compliant port	One
iDRAC Direct port (Micro-AB USB 2.0-compliant port)	One	USB 3.2 Gen1-compliant ports	One		

PSU rating

Below table lists the power capacity the PSUs in high/low line operation mode.

Table 32. PSUs highline and lowline ratings

—	**600 W Platinum**	**700 W Titanium**
Peak Power (Highline/-72 VDC)	600 W	700 W
Lowline/-40 VDC	600 W	NA
Highline 240 VDC	600 W	700 W
Highline 200 ~ 380VDC	NA	NA
DC -(48 to 60V)	NA	NA

The PowerEdge R360 supports up to two AC power supplies with 1+1 redundancy, autosensing, and auto switching capability.

If two PSUs are present during POST, a comparison is made between the wattage capacities of the PSUs. In case the PSU wattages do not match, the larger of the two PSUs is enabled. Also, there is a PSU mismatch warning that is displayed in BIOS, iDRAC, or on the system LCD.

If a second PSU is added at run-time, in order for that particular PSU to be enabled, the wattage capacity of the first PSU must equal the second PSU. Otherwise, the PSU is flagged as unmatched in iDRAC and the second PSU is not enabled.

Dell PSUs have achieved Titanium or Platinum efficiency levels as shown in the table below.

Table 33. PSU efficiency level

Efficiency Targets by Load						
Form factor	**Output**	**Class**	**10%**	**20%**	**50%**	**100%**
Redundant 60 mm	600W	Platinum	-	90%	94%	91%
	700W	Titanium	90%	94%	96%	91%

Environmental specifications

NOTE: For additional information about environmental certifications, refer to the *Product Environmental Datasheet* located with the *Documentation* on Dell Support.

Table 34. Continuous Operation Specifications for ASHRAE A2

-	**Allowable continuous operations**
Temperature range for altitudes <= 900 m (<= 2953 ft)	10–35°C (50–95°F) with no direct sunlight on the equipment
Humidity percent range (non-condensing at all times)	8% RH with -12°C minimum dew point to 80% RH with 21°C (69.8°F) maximum dew point
Operational altitude de-rating	Maximum temperature is reduced by 1°C/300 m (33.8°F/984 Ft) above 900 m (2953 Ft)

Table 35. Continuous Operation Specifications for ASHRAE A3

-	**Allowable continuous operations**
Temperature range for altitudes <= 900 m (<= 2953 ft)	5–40°C (41–104°F) with no direct sunlight on the equipment
Humidity percent range (non-condensing at all times)	8% RH with -12°C minimum dew point to 85% RH with 24°C (75.2°F) maximum dew point
Operational altitude de-rating	Maximum temperature is reduced by 1°C/175 m (33.8°F/574 Ft) above 900 m (2953 Ft)

Table 36. Continuous Operation Specifications for ASHRAE A4

-	**Allowable continuous operations**
Temperature range for altitudes <= 900 m (<= 2953 ft)	5–45°C (41–113°F) with no direct sunlight on the equipment
Humidity percent range (non-condensing at all times)	8% RH with -12°C minimum dew point to 90% RH with 24°C (75.2°F) maximum dew point
Operational altitude de-rating	Maximum temperature is reduced by 1°C/125 m (33.8°F/410 Ft) above 900 m (2953 Ft)

Table 37. Common Environmental Specifications for ASHRAE A2, A3, and A4

-	**Allowable continuous operations**
Maximum temperature gradient (applies to both operation and non-operation)	20°C in an hour* (36°F in an hour) and 5°C in 15 minutes (41°F in 15 minutes), 5°C in an hour* (41°F in an hour) for tape

Table 37. Common Environmental Specifications for ASHRAE A2, A3, and A4 (continued)

-	Allowable continuous operations
	NOTE: * - Per ASHRAE thermal guidelines for tape hardware, these are not instantaneous rates of temperature change.
Non-operational temperature limits	-40 to 65°C (-104 to 149°F)
Non-operational humidity limits	5% to 95% RH with 27°C (80.6°F) maximum dew point
Maximum non-operational altitude	12,000 meters (39,370 feet)
Maximum operational altitude	3,048 meters (10,000 feet)

Table 38. Maximum vibration specifications

Maximum vibration	Specifications
Operating	0.21 G_{rms} at 5 Hz to 500 Hz (all operation orientations)
Storage	1.88 G_{rms} at 10 Hz to 500 Hz for 15 minutes (all six sides tested)

Table 39. Maximum shock pulse specifications

Maximum shock pulse	Specifications
Operating	Six consecutively executed shock pulses in the positive and negative x, y, and z axis of 6 G for up to 11 ms.
Storage	Six consecutively executed shock pulses in the positive and negative x, y, and z axis (one pulse on each side of the system) of 71 G for up to 2 ms.

Particulate and gaseous contamination specifications

The following table defines the limitations that help avoid any equipment damage or failure from particulates and gaseous contamination. If the levels of particulates or gaseous pollution exceed the specified limitations and result in equipment damage or failure, you must rectify the environmental conditions. Remediation of environmental conditions is the responsibility of the customer.

Table 40. Particulate contamination specifications

Particulate contamination	Specifications
Air filtration: Conventional Data Center only	Data center air filtration as defined by ISO Class 8 per ISO 14644-1 with a 95% upper confidence limit **NOTE:** Filtering room air with a MERV8 filter, as specified in ANSI/ASHRAE Standard 127, is a recommended method for achieving the necessary environmental conditions. **NOTE:** Air entering the data center must have MERV11 or MERV13 filtration. **NOTE:** This condition applies to data center environments only. Air filtration requirements do not apply to IT equipment designed to be used outside a data center, in environments such as an office or factory floor.
Walk-Up Edge Data Center or Cabinet (sealed, closed loop environment)	Filtration is not required for cabinets that are anticipated to be opened six times or less per year. Class 8 per ISO 1466-1 filtration as defined above is required otherwise. **NOTE:** In environments commonly above ISA-71 Class G1 or that may have known challenges, special filters may be required.
Conductive dust: data center and non-data center environments	Air must be free of conductive dust, zinc whiskers, or other conductive particles.

Table 40. Particulate contamination specifications (continued)

Particulate contamination	Specifications
	NOTE: Conductive dust, which can interfere with equipment operation, can originate from various sources, including manufacturing processes and zinc whiskers that may develop on the plating of raised floor tiles. **NOTE:** This condition applies to data center and non-data center environments.
Corrosive dust: data center and non-data center environments	• Air must be free of corrosive dust. • Residual dust present in the air must have a deliquescent point less than 60% relative humidity. **NOTE:** This condition applies to data center and non-data center environments.

Table 41. Gaseous contamination specifications

Gaseous contamination	Specifications	Notes
Copper coupon corrosion rate	ISA-71 Class G1: <300 Å/month	Per ANSI/ISA71.04
Silver coupon corrosion rate	ISA-71 Class G1: <200 Å/month	Per ANSI/ISA71.04

Thermal restriction matrix

Table 42. Processor and heat sink matrix

Heat sink	Configuration No
1U HSK	All configurations

Table 43. Label reference

Label	Description
STD	Standard
LP	Low profile
FH	Full height

Table 44. Thermal restriction matrix

Configuration		Configuration 1 2.5 inch and 3.5 inch Drives With Riser	Configuration 2 2.5 inch and 3.5 inch Drives With N1- BOSS	Configuration 3 2.5 inch and 3.5 inch Drives Without Riser and N1 - BOSS	Ambient temperature
CPU TDP/ cTDP	≤ 80 W	4 x STD Fans 1U STD HSK	1 x Silver Fan, 3 x STD Fans 1U STD HSK	3 x STD Fan 1U STD HSK	35°C (95°F)
	95 W	4 x STD Fans 1U Performance HSK	1 x Silver Fan, 3 x STD 1U Performance HSK	3 x STD Fan 1U Performance HSK	35°C (95°F)

Table 45. GPU thermal restriction matrix

Configuration		Configuration 1 2.5 inch and 3.5 inch Drives With Riser	Configuration 2 2.5 inch and 3.5 inch Drives With N1- BOSS	Configuration 3 2.5 inch and 3.5 inch Drives Without Riser and N1 - BOSS	Ambient temperature
CPU TDP/ cTDP	≤ 80 W	4 x STD Fans 1U STD HSK	1 x Silver Fan, 3 x STD Fans 1U STD HSK	NA	35°C (95°F)
	95 W	4 x STD Fans 1U Performance HSK	1 x Silver Fan, 3x STD 1U Performance HSK	NA	35°C (95°F)

NOTE: When A2 GPU card is populated, PCIe slot 2 has to be populated or should have a blank installed.

Table 46. Fan Location

Chassis	Configuration	Fan Quantity	Fan Location
2.5 inch and 3.5 inch Drives	With Riser Card	4	Fan 1 (STD), Fan 2, Fan 3, Fan 4
2.5 inch and 3.5 inch Drives	With N1-BOSS Module	4	Fan 1 (Silver), Fan 2, Fan 3, Fan 4
2.5 inch and 3.5 inch Drives	Without Riser and N1- BOSS Module	3	Fan 2, Fan 3, Fan 4

Thermal air restrictions

ASHRAE A3/A4 environment

- The operating temperature is for a maximum altitude of 950 m for ASHRAE A3/A4 Cooling.
- If over 950 m, we need to do the ambient temperature de-rating.
- CPU TDP greater than 80 W is not supported.
- BOSS-N1 (M.2) module is not supported.
- A2 GPU card is not supported.
- Non-Dell qualified peripheral cards and/or peripheral cards greater than 25 W are not supported.
- Two PSUs are required in redundant mode.

15

Appendix A. Standards compliance

The system conforms to the following industry standards.

Table 47. Industry standard documents

Standard	URL for information and specifications
ACPIAdvance Configuration and Power Interface Specification, v6.4	ACPI
Ethernet IEEE Std 802.3-2022	IEEE Standards
MSFT WHQL Microsoft Windows Hardware Quality Labs	Windows Hardware Compatibility Program
IPMI Intelligent Platform Management Interface, v2.0	IPMI
DDR5 Memory DDR5 SDRAM Specification	JEDEC Standards
PCI Express PCI Express Base Specification, v5.0	PCIe Specifications
PMBus Power System Management Protocol Specification, v1.2	Power System Management Protocol Specification
SAS Serial Attached SCSI, 3 (SAS-3) (T10/INCITS 519)	SCSI Storage Interfaces
SATA Serial ATA Rev. 3.3	SATA IO
SMBIOS System Management BIOS Reference Specification, v3.3.0	DMTF SMBIOS
TPM Trusted Platform Module Specification, v1.2 and v2.0	TPM Specifications
UEFI Unified Extensible Firmware Interface Specification, v2.7	UEFI Specifications
PI Platform Initialization Specification, v1.7	
USB Universal Serial Bus v2.0 and SuperSpeed v3.0 (USB 3.1 Gen1)	USB Implementers Forum, Inc. USB
NVMe Express Base Specification. Revision 2.0c	NVMe
NVMe Command Set Specifications 1. NVM Express NVM Command Set Specification. Revision 1.1c 2. NVM Express Zoned Namespaces Command Set. Revision 1.0c 3. NVM Express® Key Value Command Set. Revision 1.0c	
NVMe Transport Specifications 1. NVM Express over PCIe Transport. Revision 1.0c 2. NVM Express RDMA Transport Revision. 1.0b 3. NVM Express TCP Transport. Revision 1.0c	
NVMe NVM Express Management Interface. Revision 1.2c	
NVMe NVMe Boot Specification. Revision 1.0	

16

Appendix C Additional resources

Table 48. Additional resources

Resource	Description of contents	Location
Installation and Service Manual	This manual, available in PDF format, provides the following information: • Chassis features • System Setup program • System indicator codes • System BIOS • Remove and replace procedures • Diagnostics • Jumpers and connectors	Dell.com/Support/Manuals
Getting Started Guide	This guide ships with the system, and is also available in PDF format. This guide provides the following information: • Initial setup steps	Dell.com/Support/Manuals
Rack Installation Guide	This document ships with the rack kits, and provides instructions for installing a server in a rack.	Dell.com/Support/Manuals
System Information Label	The system information label documents the system board layout and system jumper settings. Text is minimized due to space limitations and translation considerations. The label size is standardized across platforms.	Inside the system chassis cover
QR code for system resources	This code on the chassis can be scanned by a phone application to access additional information and resources for the server, including videos, reference materials, service tag information, and Dell contact information.	Inside the system chassis cover
Enterprise Infrastructure Planning Tool (EIPT)	The Dell online EIPT enables easier and more meaningful estimates to help you determine the most efficient configuration possible. Use EIPT to calculate the power consumption of your hardware, power infrastructure, and storage.	Dell.com/calc

Dell EMC PowerEdge R450

Technical Guide

Part Number: E74S
Regulatory Type: E74S001
December 2021
Rev. A02

DELL Technologies

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

© 2021 Dell Inc. or its subsidiaries. All rights reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Contents

Chapter 1: System overview 5
Key workloads 5
New technologies 5

Chapter 2: System features and generational comparison 7

Chapter 3: Chassis views and features 9
Chassis views 9
Front view of the system 9
Rear view of the system 9
Inside the system 10
Quick Resource Locator 11

Chapter 4: Processor 12
Processor features 12
Supported processors 12

Chapter 5: Memory subsystem 13
Supported memory 13
Memory speed 14

Chapter 6: Storage 15
Storage controllers 15
Supported drives 15
External storage 16

Chapter 7: Networking 17
Overview 17
Supported OCP cards 17
OCP form factors 18
OCP NIC 3.0 vs. rack Network Daughter Card comparisons 18

Chapter 8: Expansion cards and expansion card risers 19
Expansion card installation guidelines 19

Chapter 9: Power, thermal, and acoustics 23
Power 23
Thermal design 24
PowerEdge R450 acoustics 24
Acoustical performance 25

Chapter 10: Rack, rails, and cable management 27
Rails information 27

Sliding Rails ... 27
Static Rails ... 28
Cable Management Arm ... 28
Strain Relief Bar ... 28

Chapter 11: Supported Operating Systems ... 29

Chapter 12: Dell EMC OpenManage systems management ... 30
Server and Chassis Managers ... 31
Dell EMC consoles ... 31
Automation Enablers ... 31
Integration with third-party consoles ... 31
Connections for third-party consoles ... 31
Dell EMC Update Utilities ... 31
Dell resources ... 31

Chapter 13: Dell Technologies Services ... 33
Dell EMC ProDeploy Enterprise Suite ... 33
Dell EMC ProDeploy Plus ... 34
Dell EMC ProDeploy ... 34
Basic Deployment ... 34
Dell EMC ProDeploy for HPC ... 34
Dell EMC Server Configuration Services ... 35
Dell EMC Residency Services ... 35
Dell EMC Data Migration Service ... 35
Dell EMC ProSupport Enterprise Suite ... 35
Dell EMC ProSupport Plus for Enterprise ... 36
Dell EMC ProSupport for Enterprise ... 36
Dell EMC ProSupport One for Data Center ... 37
ProSupport for HPC ... 37
Support Technologies ... 37
Services for Data Security ... 38
Dell Technologies Education Services ... 39
Dell Technologies Consulting Services ... 39
Dell EMC Remote Consulting Services ... 39
Dell EMC Managed Services ... 39

Chapter 14: Appendix A. Additional specifications ... 40
Chassis dimensions ... 40
Chassis weight ... 41
Video specifications ... 41
USB Ports ... 42
PSU rating ... 42
Environmental specifications ... 43
Thermal restrictions ... 44

Chapter 15: Appendix B. Standards compliance ... 47

Chapter 16: Appendix C Additional resources ... 48

1

System overview

The Dell™ PowerEdge™ R450 is Dell's latest 2-socket, 1 U rack servers designed to run complex workloads using highly scalable memory, I/O, and network options. The systems feature 3rd Generation Intel® Xeon® Processor Scalable Family , up to 16 DIMMs, PCI Express® (PCIe) 4.0 enabled expansion slots, and a choice of network interface technologies to cover NIC. The R450 is a general-purpose platform capable of handling demanding workloads and applications, such as data warehouses, ecommerce, databases, and high-performance computing (HPC).

Topics:

- Key workloads
- New technologies

Key workloads

The target workloads for the PowerEdge R450 include light duty virtualization, small IT infrastructure, and small business specific workloads.

New technologies

Table 1. New technologies

Technology	Detailed Description
3rd Generation Intel® Xeon® Processor	Core count: Up to 24 per processor supported on R450 UPI speed: Up to 3x UPIs/Socket at 10.4 GT/s or 11.2 GT/s Maximum number of PCIe lanes: Integrated 64 PCIe 4.0 lanes @ 16 GT/s PCIe Gen4 Maximum TDP: 185 W
3200 MT/s DDR4 Memory **NOTE:** Maximum memory speed depends on the processor specification. See Supported Processors table for the processor type and the supported memory speed.	Maximum 8 DIMMs per processor and 16 DIMMs in total Supports DDR4 ECC RDIMM, with ECC up to 2933 MT/s on R450 as the memory speed depends on the processor supported by the system.
Flex IO	LOM, 2x 1 Gb with BCM5720 LAN controller Rear IO with 1 Gb Dedicated Management Network Port, USB 3.0 x1, USB 2.0 x1 and VGA port Serial port option OCP Mezz 3.0 (supported by x16 PCIe lanes)
Dedicated PERC	Front storage module PERC with front PERC 10.5 and PERC 11
Software RAID	OS RAID/S150
Power Supplies	60 mm dimension is the new PSU form factor on 15G design • Platinum 600 W AC/HVDC

Table 1. New technologies (continued)

Technology	Detailed Description
	• Platinum 800 W AC/HVDC • 1100 W HVDC

2

System features and generational comparison

The following table shows the comparison between the PowerEdge R450 with the PowerEdge R440.

Table 2. Features compared to previous version

Feature	PowerEdge R450	PowerEdge R440
Processor	Up to two 3rd Generation Intel Xeon Scalable processors with up to 24 cores per processor	2nd Generation Intel® Xeon® Scalable Processor Family
Processor Interconnect	Intel Ultra Path Interconnect (Intel UPIx3)	Intel Ultra Path Interconnect (Intel UPI)
Memory	16x RDIMM/No NVDIMM	16x DDR4 RDIMM, LRDIMM
Storage Drives	Up to 4 x 3.5-inch SAS/SATA (HDD/SSD) max 64 TB Up to 8 x 2.5-inch SAS/SATA (HDD/SSD) max 61.44 TB	Up to 10 x 2.5" SAS/SATA (HDD/SSD) with up to 4 NVMe SSD max 76.8 TB or Up to 4 x 3.5 SAS/SATA HDD max 64 TB
Storage Controllers	Internal controllers: PERC H355, PERC H345, PERC H745, PERC H755, HBA355i, S150 Internal Boot: Internal Dual SD Module or Boot Optimized Storage Subsystem (BOSS-S1): HWRAID 2 x M.2 SSDs or USB External PERC (RAID): PERC H840, HBA355e	Internal Controllers: PERC H330, H730P, H740P, HBA330 External Controllers: H840, 12 Gbps SAS HBA Software RAID:S140 Boot Optimized Storage Subsystem (BOSS):HWRAID 2 x M.2 SSDs 240GB, 480GB Internal Dual SD Module
PCIe Slots	2 x PCIe Gen4 slots	2 x PCIe Gen3 (x16/x16)
Embedded NIC (LOM)	2 x 1GbE	2 x 1GbE + OCP riser: LRC 2 x 1GbE or 2 x 10GbE
Networking options (OCP 3.0)	OCP3 x16 Gen4	OCP2 x16 Gen3
I/O Ports	Front ports: • 1 x iDRAC Direct (Micro-AB USB) port • 1 x USB 2.0 • 1x VGA Rear ports: • 1 x USB 2.0 • 1 x Serial (optional) • 1 x USB 3.0 • 2 x Ethernet • 1 x VGA • 1 x Dedicated iDRAC network port Internal Port: 1 x USB 3.0 (optional)	Front ports: • 1 x iDRAC Direct (Micro-AB USB) port • 1 x USB 2.0 • 1x VGA Rear ports: • 1 x USB 2.0 • 1 x Serial (optional) • 2 x USB 3.0 • 1 x Dedicated iDRAC network port • Up to 2 x PCIe Gen 3 slots all x16

Table 2. Features compared to previous version (continued)

Feature	PowerEdge R450	PowerEdge R440
Rack Height	1U	1U
Power Supplies	1100 W DC / -48–(-60) V 800 W Platinum AC/240 V HVDC 600 W Platinum AC/240 V HVDC	Bronze 450W (Cabled PSU) Platinum 550W (Hot plug PSU with full redundancy option)
System Management	iDRAC9 iDRAC Direct iDRAC Service Module Quick Sync 2 wireless module	iDRAC9 iDRAC9 Direct iDRAC REST API with Redfish Quick Sync 2 BLE/wireless module
Internal GPU	No GPU support	No GPU support
Availability	Hot-plug Drives Hot-plug Redundant Power Supplies IDSDM	Hot-plug Drives Hot-plug Redundant Power Supplies IDSDM

3

Chassis views and features

Topics:

- Chassis views

Chassis views

Front view of the system



Figure 1. Front view of 4 x 3.5-inch drive system



Figure 2. Front view of 8 x 2.5-inch drive system

Rear view of the system

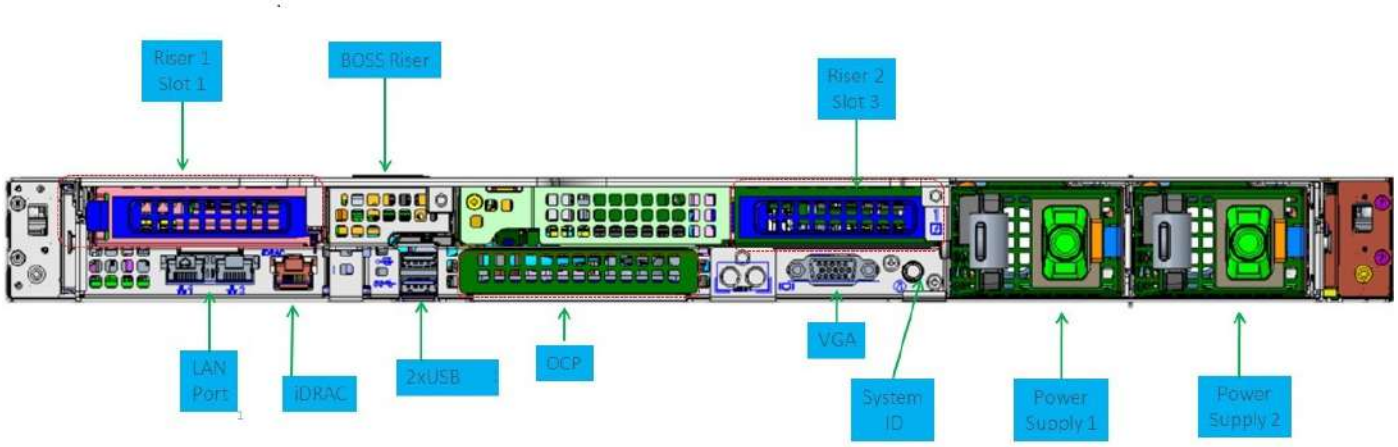


Figure 3. Rear view of the system

Inside the system

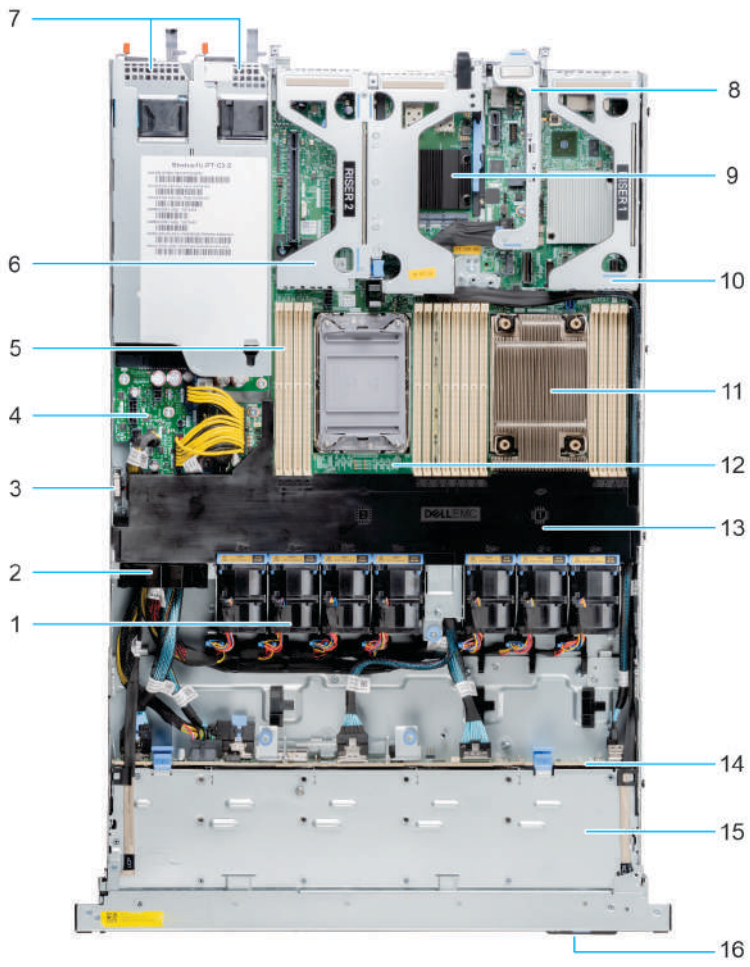


Figure 4. Inside the system

1. Fan
2. Cable retention latch
3. Intrusion switch
4. Power interposer board
5. Memory module slots
6. Riser 2c
7. PSU 1 and PSU 2
8. BOSS riser
9. OCP
10. Riser 1
11. Heat sink
12. System board
13. Air shroud
14. Drive backplane
15. Backplane cover
16. Information tag

Quick Resource Locator



Figure 5. Quick Resource Locator for R450

4

Processor



Topics:

- Processor features
- Supported processors

Processor features

The 3rd Generation Xeon® Scalable Processors stack is the next generation data center processor offering with the latest features, increased performance, and incremental memory options. This latest generation Xeon Scalable processor supports usages from entry designs that are based on Intel Xeon Silver processors to advanced capabilities offered in new Intel Xeon Platinum processor.

The following lists the features and functions that are in the upcoming 3rd Generation Intel® Xeon Scalable Processor offering:

- Faster UPI with 3 Intel Ultra Path Interconnect (Intel UPI) at 11.2 GT/s (supported in gold and platinum options)
- More, faster I/O with PCI Express 4 and up to 64 lanes (per socket) at 16 GT/s
- Enhanced Memory Performance with support for up to 2933 MT/s DIMMs

Supported processors

Following tables list the various Processor SKUs supported.

Table 3. Supported processors for R450

Processor	Clock Speed (GHz)	Cache (M)	UPI (GT/s)	Cores	Threads	Turbo	Memory Speed (MT/s)	Memory Capacity	TDP	R450
5318Y	2.1	36	11.2	24	48	Turbo	2933	6TB	165W	Supported
5317	3	18	11.2	12	24	Turbo	2933	6TB	150W	Supported
5315Y	3.2	12	11.2	8	16	Turbo	2933	6TB	140W	Supported
4316	2.3	30	10.4	20	40	Turbo	2666	6TB	150W	Supported
4314	2.4	24	10.4	16	32	Turbo	2666	6TB	135W	Supported
4310	2.1	18	10.4	12	24	Turbo	2666	6TB	120W	Supported
4309Y	2.8	12	10.4	8	16	Turbo	2666	6TB	105W	Supported

5

Memory subsystem

The R450 supports up to 16 DIMMs, with up to 1024GB of memory and speeds of up to 3200MT/s.

The R450 supports registered (RDIMMs) which use a buffer to reduce memory loading and provide greater density, allowing for the maximum platform memory capacity. Unbuffered DIMMs (UDIMMs) are not supported.

Topics:

- Supported memory
- Memory speed

Supported memory

The table below lists the memory technologies supported by the platform.

Table 4. Supported memory technologies

Feature	PowerEdge R450 (DDR4)
DIMM type	RDIMM
Transfer speed	2933 MT/s
Voltage	1.2 V (DDR4)

The following table lists the supported DIMMs for the R450 at launch. For information on memory configuration, see the *Dell EMC PowerEdge R450 Installation and Service Manual* at www.dell.com/poweredgemanuals.

Table 5. Supported DIMMs

DIMM Speed (MT/s)	DIMM Type	DIMM Capacity (GB)	Ranks per DIMM	Data Width	DIMM Volts
3200	RDIMM	8 GB	1R	8	1.2 V
3200	RDIMM	16 GB	2R	8	1.2 V
3200	RDIMM	32 GB	2R	8	1.2 V
3200	RDIMM	64 GB	2R	4	1.2 V

NOTE: Maximum memory speed depends on the processor specification. The DIMM can support upto 3200 MT/s, but the memory may not support this memory speed. See Supported Processors table for the processor type and the supported memory speed.

Memory speed

Table 6. DIMM performance details

DIMM type	Ranks per DIMM and Data Width	DIMM Capacity (GB)		Speed (MT/s) ;Voltage (V);
				DIMM Per Channel (DPC)
				1 DPC
				1.2 V
RDIMM	SRx8	8 GB	16 GB	D: 2933
	SRx4	16 GB	32 GB	
	DRx8	16 GB	32 GB	D: 2933
	DRx4	32 GB	64 GB	

NOTE: Maximum memory speed depends on the processor specification. The DIMM can support upto 3200 MT/s, but the memory may not support this memory speed. See Supported Processors table for the processor type and the supported memory speed.

6

Storage

Topics:

- Storage controllers
- Supported drives
- External storage

Storage controllers

Dell's RAID controller options offer performance improvements, including the fPERC solution. fPERC provides a base RAID HW controller without consuming a PCIe slot by using a small form factor and high-density connector to the base planar.

15G PERC Controller offerings will be a heavy leverage of 14G PERC family. The Value and Value Performance levels will carry over to 15G from 14G. New to 15G, is the Harpoon-based Premium Performance tier offering. This high-end offering will drive IOPs performance and enhanced SSD performance.

15G PERC Controller offerings will be a heavy leverage of 14G PERC family. The Value and Value Performance levels will carry over to 15G from 14G. New to 15G, is the Harpoon-based Premium Performance tier offering. This high-end offering will drive IOPs performance and enhanced SSD performance.

Table 7. PERC Series Controller Offerings

Performance Level	Controller and Description
Entry	S150 (SATA) SW RAID SATA
Value	H355, H345, HBA355 (Internal/External)
Value Performance	H745, H755

Supported drives

The table shown below lists the internal drives supported by the R450.

Table 8. Supported drives

Form Factor	Type	Speed	Rotational Speed	Capacities
2.5-inch	SATA SSD	6 Gb	N/A	480 GB, 960 GB, 1.92 TB, 3.84 TB
2.5-inch	SAS	12 Gb	10 K	600 GB, 1.2 TB, 2.4 TB
2.5-inch	SAS	12 Gb	15 K	900 GB
2.5-inch	SAS SSD	12 Gb	N/A	480 GB, 800 GB, 960 GB, 1.6 TB, 1.92 TB, 3.84 TB, 6.4 TB, 7.68 TB
3.5-inch	SATA	6 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16 TB
3.5-inch	SAS	12 Gb	7.2 K	2 TB, 4 TB, 8 TB, 12 TB, 16TB
M.2	SATA SSD	6 Gb	N/A	240 GB, 480 GB

Table 8. Supported drives (continued)

Form Factor	Type	Speed	Rotational Speed	Capacities
uSD	N/A	N/A	uSD	16 GB, 32 GB, 64 GB

External storage

The R450 supports the external storage device types listed in the table below.

Table 9. Supported external storage devices

Device Type	Description
External Tape	Supports connection to external tape products
NAS/IDM appliance software	Supports NAS software stack
JBOD	Supports connection to 12GB MD/ME - series JBODs

7

Networking

Topics:

- Overview
- Supported OCP cards

Overview

PowerEdge offers a wide variety of options to get information moving to and from our servers. Industry best technologies are chosen, and systems management features are added by our partners to firmware to tie in with iDRAC. These adapters are rigorously validated for worry-free, fully supported use in Dell servers.

The PowerEdge Server Adapter Matrix posted to knowledge portal is the central repository for PowerEdge NIC, HBA and HCA information. The matrix covers:

- Part Numbers, Tied SKUs and Customer Kits
- Server Compatibility and Support
- Optics and Cable Support
- Systems Management
- Adapter Features
- Spec Sheet Links

This document is updated as changes happen, so be sure to bookmark it rather than downloading an offline copy to stay with the latest information.

NOTE: This is a direct download link to an .XLSX and may not open in a tab as expected depending on your browser.

Supported OCP cards

Table 10. OCP support list

Form Factor	Vendor	Port type	Port speed	Port count
OCP 3.0	Intel	SFP+	10GbE	2
OCP 3.0	Broadcom	BT	1GbE	4
OCP 3.0	Broadcom	BT	10GbE	2
OCP 3.0	Broadcom	SFP28	25GbE	2
OCP 3.0	Broadcom	SFP28	25GbE	4
OCP 3.0	Broadcom	SFP+	10GbE	2
OCP 3.0	QLogic	BT	10GbE	2
OCP 3.0	QLogic	SFP+	10GbE	2
OCP 3.0	QLogic	SFP28	25GbE	2
OCP 3.0	Intel	BT	1GbE	4
OCP 3.0	Intel	BT	10GbE	2
OCP 3.0	Intel	SFP+	10GbE	4

Table 10. OCP support list (continued)

Form Factor	Vendor	Port type	Port speed	Port count
OCP 3.0	Intel	SFP28	25GbE	2
OCP 3.0	Mellanox	SFP28	25GbE	2

OCP form factors

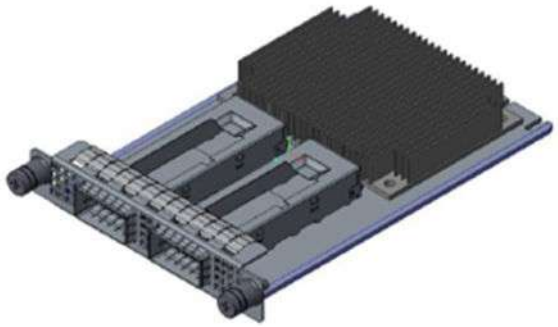


Figure 6. OCP 3.0 Small Card Form Factor (LS)

Table 11. OCP 3.0 Feature List

Feature	OCP 3.0
Form factor	SFF and LFF
PCIe Gen	Gen4
Max PCIe width	x16
Max of ports	4
Port type	BT/SFP/SFP+/SFP28/SFP56
Max port speed	100Gbe
NC-SI	Yes
WoL	Yes
Power consumption	15 W~150 W

OCP NIC 3.0 vs. rack Network Daughter Card comparisons

Table 12. OCP 3.0, 2.0, and rNDC NIC comparison

Form Factor	Dell rNDC	OCP 2.0 (LOM Mezz)	OCP 3.0	Notes
PCIe Gen	Gen 3	Gen 3	Gen 4	Supported OCP3 are SFF (small form factor)
Max PCIe Lanes	x8	Up to x16	Up to x16	See server slot priority matrix
Shared LOM	Yes	Yes	Yes	This is iDRAC port redirect
Aux Power	Yes	Yes	Yes	Used for Shared LOM

8

Expansion cards and expansion card risers

NOTE: When an expansion card is not supported or missing, riser the iDRAC and Lifecycle Controller logs an event. This does not prevent your system from booting. However, if a F1/F2 pause occurs with an error message, see *Troubleshooting expansion cards* section in the *Dell EMC PowerEdge Servers Troubleshooting Guide* at www.dell.com/poweredgemanuals.

Topics:

- Expansion card installation guidelines

Expansion card installation guidelines

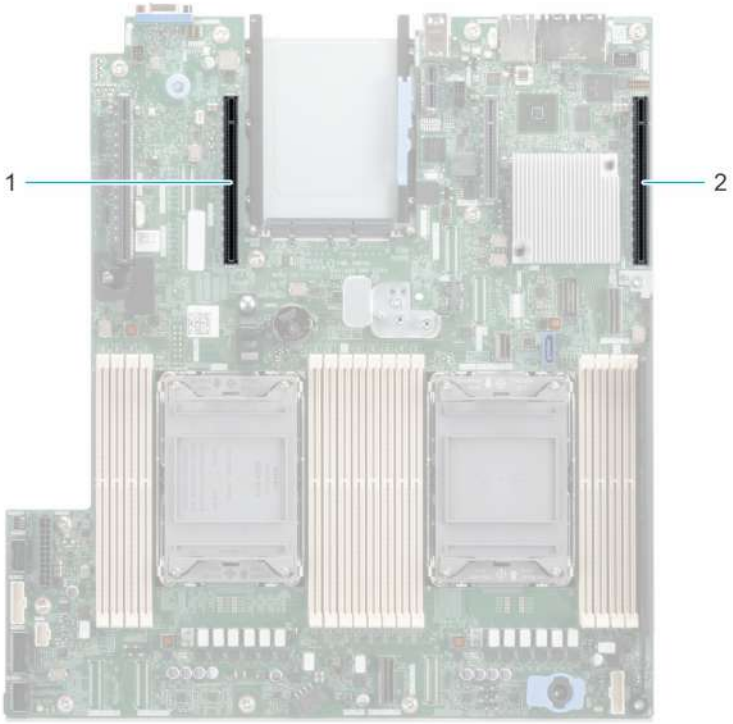


Figure 7. Expansion card slot connectors

1. Riser 2C (Slot 3)
2. Riser 1 (Slot 1)

The following table describes the expansion card riser configurations:

Table 13. Expansion card riser configurations

Configurations	Expansion card risers	PCIe Slots	Controlling processor	Height	Length	Slot width
Config0. with 1x LP	R1	1	Processor 1	Low Profile	Half length	x16
Config1. with 2x LP	R1	1	Processor 1	Low Profile	Half length	x16
	R2c	3	Processor 2	Low Profile	Half length	x16

Table 13. Expansion card riser configurations (continued)

Configurations	Expansion card risers	PCIe Slots	Controlling processor	Height	Length	Slot width
Config2. with 0x LP	NA	NA	NA	NA	NA	NA



Figure 8. Riser 1

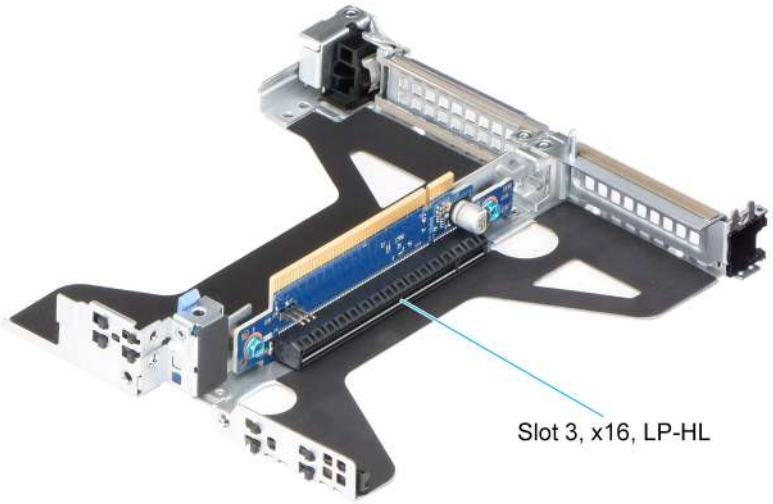


Figure 9. Riser 2c

NOTE: The expansion-card slots are not hot-swappable.

The following table provides guidelines for installing expansion cards to ensure proper cooling and mechanical fit. The expansion cards with the highest priority should be installed first using the slot priority indicated. All the other expansion cards should be installed in the card priority and slot priority order.

Table 14. Configuration 0: R1

Card type	Slot priority	Maximum number of cards
Dell Serial port module (LP)	1	1

Table 14. Configuration 0: R1 (continued)

Card type	Slot priority	Maximum number of cards
Mellanox (NIC: 25 Gb, LP)	1	1
Broadcom (NIC: 25 Gb)	1	1
Intel (NIC: 25 Gb)	1	1
Mellanox (NIC: 25 Gb)	1	1
Qlogic (NIC: 25 Gb)	1	1
SolarFlare (NIC: 25 Gb)	1	1
Broadcom (NIC: 10 Gb)	1	1
Intel (NIC: 10 Gb)	1	1
Qlogic (NIC: 10 Gb)	1	1
Broadcom (NIC: 1 Gb)	1	1
Intel (NIC: 1 Gb)	1	1
Broadcom (OCP: 25 Gb)	Integrated slot	1
Intel (OCP: 25 Gb)	Integrated slot	1
Marvell (OCP: 25 Gb)	Integrated slot	1
Mellanox (OCP: 25 Gb)	Integrated slot	1
Mellanox (OCP 3: 25 Gb)	Integrated slot	1
SolarFlare (OCP: 25 Gb)	Integrated slot	1
Broadcom (OCP: 10 Gb)	Integrated slot	1
Marvell (OCP: 10 Gb)	Integrated slot	1
Intel (OCP: 10 Gb)	Integrated slot	1
Broadcom (OCP: 1 Gb)	Integrated slot	1
Intel (OCP: 1 Gb)	Integrated slot	1
Dell External PERC Adapter	1	1
PERC H355, Front	Integrated slot	1
Dell BOSS S1 Module	Integrated slot	1
Intel (PCIe SSD AIC)	1	1
Samsung (PCIe SSD AIC)	1	1

Table 15. Configuration 1: R1+R2c

Card type	Slot priority	Maximum number of cards
Broadcom (NIC: 25 Gb)	3, 1	2
Intel (NIC: 25 Gb)	3, 1	2
Mellanox (NIC: 25 Gb)	3, 1	2
Mellanox (NIC: 25 Gb, LP)	1	1
Qlogic (NIC: 25 Gb)	3, 1	2
SolarFlare (NIC: 25 Gb)	3, 1	2
Broadcom (NIC: 10 Gb)	3, 1	2
Intel (NIC: 10 Gb)	3, 1	2

Table 15. Configuration 1: R1+R2c (continued)

Card type	Slot priority	Maximum number of cards
Qlogic (NIC: 10 Gb)	3, 1	2
Broadcom (NIC: 1 Gb)	3, 1	2
Intel (NIC: 1 Gb)	3, 1	2
Broadcom (OCP: 25 Gb)	Integrated slot	1
Intel (OCP: 25 Gb)	Integrated slot	1
Marvell (OCP: 25 Gb)	Integrated slot	1
Mellanox (OCP: 25 Gb)	Integrated slot	1
Mellanox (OCP 3: 25 Gb)	Integrated slot	1
SolarFlare (OCP: 25 Gb)	Integrated slot	1
Broadcom (OCP: 10 Gb)	Integrated slot	1
Marvell (OCP: 10 Gb)	Integrated slot	1
Intel (OCP: 10 Gb)	Integrated slot	1
Broadcom (OCP: 1 Gb)	Integrated slot	1
Intel (OCP: 1 Gb)	Integrated slot	1
Dell External PERC Adapter	3, 1	2
PERC H355, Front	Integrated slot	1
Dell BOSS S1 Module	Integrated slot	1
Intel (PCIe SSD AIC)	3, 1	2
Samsung (PCIe SSD AIC)	3, 1	2

Table 16. Configuration 3: No Riser

Card type	Slot priority	Maximum number of cards
Broadcom (OCP: 25 Gb)	Integrated slot	1
Intel (OCP: 25 Gb)	Integrated slot	1
Marvell (OCP: 25 Gb)	Integrated slot	1
Mellanox (OCP: 25 Gb)	Integrated slot	1
Mellanox (OCP 3: 25 Gb)	Integrated slot	1
SolarFlare (OCP: 25 Gb)	Integrated slot	1
Broadcom (OCP: 10 Gb)	Integrated slot	1
Intel (OCP: 10 Gb)	Integrated slot	1
Marvell (OCP: 10 Gb)	Integrated slot	1
Broadcom (OCP: 1 Gb)	Integrated slot	1
Intel (OCP: 1 Gb)	Integrated slot	1
Dell BOSS S1 Module	Integrated slot	1
PERC H355, Front	Integrated slot	1

9

Power, thermal, and acoustics

Topics:

- Power
- Thermal design
- PowerEdge R450 acoustics

Power

Table 17. Power tools and technologies

Feature	Description
Power Supply Units(PSU) portfolio	Dell's PSU portfolio includes intelligent features such as dynamically optimizing efficiency while maintaining availability and redundancy. Find additional information in the Power supply units section.
Industry Compliance	Dell's servers are compliant with all relevant industry certifications and guide lines, including 80 PLUS, Climate Savers and ENERGY STAR.
Power monitoring accuracy	PSU power monitoring improvements include: ● Dell's power monitoring accuracy is currently 1%, whereas the industry standard is 5%. ● More accurate reporting of power. ● Better performance under a power cap.
Power capping	Use Dell's systems management to set the power cap limit for your systems to limit the output of a PSU and reduce system power consumption. Dell is the first hardware vendor to leverage Intel Node Manager for circuit-breaker fast capping.
Systems Management	iDRAC Enterprise and Datacenter provides server- level management that monitors, reports and controls power consumption at the processor, memory and system level. Dell OpenManage Power Center delivers group power management at the rack, row and data center level for servers, power distribution units and uninterruptible power supplies.
Rack infrastructure	Dell offers some of the industry's highest-efficiency power infrastructure solutions, including ● Power distribution units (PDUs). ● Uninterruptible power supplies (UPSs). ● Energy Smart containment rack enclosures. ● Find additional information at: http://content.dell.com/us/en/enterprise/power- and- cooling-technologies- components- rack- infrastructure.aspx.

Thermal design

Thermal management of PowerEdge R450 delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from 10°C to 35°C (50°F to 86°F) and to extended ambient temperature ranges (see Environmental Specifications). The benefits to you are lower fan power consumption (lower server system power and data center power consumption) and greater acoustical versatility.

Table 18. Thermal design features

Features	Description
Reliability	• Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life.
Performance	• Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations
Efficiency	• 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets.
Management	• System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads.
Forward Compatibility	• Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced.

The thermal design of the PowerEdge R450 reflects the following:

- Optimized thermal design: The system layout is architected for optimum thermal design.
- System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power.
- Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system- component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, OCP.
- Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds.
- User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the PowerEdge R450 Owner's Manual at www.dell.com/poweredgemanuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com.
- Cooling redundancy: The R450 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system.
- Environmental Specifications: The optimized thermal management makes the R450 reliable under a wide range of operating environments.
- Optimized thermal design: The system layout is architected for optimum thermal design.

PowerEdge R450 acoustics

Dell EMC PowerEdge R450 is a rack-mount server appropriate for attended data center environment. However, lower acoustical output is attainable with proper hardware or software configurations. For example, the minimum configuration of R450 is quiet enough for typical office environment.

Acoustical performance

Dell EMC PowerEdge R450 is a rack-mount server appropriate for attended data center environment. However, lower acoustical output is attainable with proper hardware or software configurations.

Table 19. Acoustical configurations of R450

Configuration	Single Socket	Typical 3.5-inch	Typical 2.5-inch	Margin Rich
Processor Type	3rd Generation Intel® Xeon® Processor	3rd Generation Intel® Xeon® Processor	3rd Generation Intel® Xeon® Processor	3rd Generation Intel® Xeon® Processor
Processor TDP	105 W / 10 C	105 W / 10 C	105 W / 10 C	150 W / 24 C
Processor Quantity	1	2	2	2
RDIMM Memory	8 GB DDR4	16 GB DDR4	16 GB DDR4	32 GB DDR4
Memory Quantity	2	4	4	8
Backplane Type	4x3.5-inch	4x3.5-inch	8x2.5-inch	10x2.5-inch
HDD Type	3.5-inch 7.2 K RPM SATA	3.5-inch 7.2 K RPM SAS	2.5-inch 10 K RPM SAS	2.5-inch 10 K RPM SAS
HDD Quantity	2	4	6	10
PSU Type	800 W	800 W	800 W	1400 W
PSU Quantity	1	2	2	2
M.2	X	X	X	X
OCP	Dual Port 1 GbE	Dual Port 1 GbE	Dual Port 1 GbE	Dual Port 10 GbE
PCI 1	X	X	X	X
PCI 2	X	X	X	X
Front PERC	PERC H345, H355	PERC H345, H355	PERC H345, H355	PERC H345, H355
LOM Card	X	X	X	X
PERC	X	X	X	X

Table 20. Acoustical performance of R450 acoustical configurations

Configuration		Single Socket	Typical 3.5-inch	Typical 2.5-inch	Margin Rich
Acoustical Performance: Idle/ Operating @ 25 °C Ambient				N/A	N/A
$L_{wA,m}$ (B)	Idle	4.6	4.7	4.7	4.8
	Operating	5.2	5.2	5.2	5.3
K_v (B)	Idle	0.4	0.4	0.4	0.4
	Operating	0.4	0.4	0.4	0.4
$L_{pA,m}$ (dB)	Idle	32	33	33	35
	Operating	37	37	39	43
Prominent tones		No prominent tones in Idle and Operating		N/A	N/A
Acoustical Performance: Idle @ 28°C Ambient				N/A	N/A
$L_{wA,m}$ (B)		5.0	5.0	5.0	5.1
K_v (B)		0.4	0.4	0.4	0.4

Table 20. Acoustical performance of R450 acoustical configurations (continued)

Configuration	Single Socket	Typical 3.5-inch	Typical 2.5-inch	Margin Rich
$L_{pA,m}$ (dB)	36	36	36	37
Acoustical Performance: Max. Loading @ 35°C Ambient			N/A	N/A
$L_{wA,m}$ (B)	6.9	7.0	7.0	6.9
K_v (B)	0.4	0.4	0.4	0.4
$L_{pA,m}$ (dB)	54	55	55	53

NOTE:

- $L_{wA,m}$ The declared mean A-weighted sound power level (LwA) is calculated per section 5.2 of ISO 9296 (2017) with data collected using the methods described in ISO 7779 (2010). Data presented here may not be fully compliant with ISO 7779
- $L_{pA,m}$ The declared mean A-weighted emission sound pressure level is at the bystander position per section 5.3 of ISO 9296 (2017) and measured using methods described in ISO 7779 (2010). The system is placed in a 24U rack enclosure, 25cm above a reflective floor. Data presented here may not be fully compliant with ISO 7779.
- **Prominent tones:** Criteria of D.6 and D.11 of ECMA-74 are followed to determine if discrete tones are prominent and to report them, if so.
- **Idle mode:** The steady-state condition in which the server is energized but not operating any intended function.
- **Operating mode:** The maximum of the steady state acoustical output at 50% of CPU TDP or active HDDs per C.9.3.2 in ECMA-74 .

10

Rack, rails, and cable management

Key factors in selecting the proper rails include, Identifying:

- Type of rack in which the rails will be installed
- Spacing between the front and rear mounting flanges of the rack
- Type and location of any equipment mounted on the rear of the rack such as power distribution units (PDUs), and the overall depth of the rack

Refer the DellEMC Enterprise Systems Rail Sizing and Rack Compatibility Matrix for the following information:

- Specific details about rail types and their functionalities
- Rail adjustability ranges for various rack mounting flange types
- Rail depth with and without cable management accessories
- Rack types supported for various rack mounting flange types

Topics:

- Rails information
- Cable Management Arm
- Strain Relief Bar

Rails information

The R450 supports both sliding rails and static rails. Both rails have a slim rail design that supports the wide system chassis.

Sliding Rails

Stab-in/Drop-in sliding rails: The Stab-in/Drop-in sliding rail also provides tool-less support for 4-post racks with square or unthreaded round mounting holes including all generations of Dell racks. Additionally, these rails offer tool-less support for 4-post threaded racks with no conversion required.

The optional cable management arm (CMA) can be mounted on either the left or right side of the sliding rails without the use of tools for fast and easy deployment.



Figure 10. Drop in/Stab-in Sliding Rail mounting interface

Static Rails

ReadyRails Static Rails for 4-post and 2-post Racks:

- The ReadyRails static rails also supports tool-less installation to 4-post racks with square or unthreaded round mounting holes including all generations of Dell racks.
- The static rails support tooled mounting in 2-post (Telco) racks as well for added versatility.

NOTE: The R450 is compatible with the R440, R6515, and R6415 rails and CMA.

Cable Management Arm

The optional Cable Management Arm (CMA) for the system organizes and secures the cords and cables exiting the back of the server and unfolds to allow the server to extend out of the rack without having to detach the cables.

Some key features of the CMA include:

- Large U-shaped baskets to support dense cable loads
- Open vent pattern for optimal airflow
- Support for mounting on either side by swinging the spring-loaded brackets from one side to the other
- Utilizes hook-and-loop straps rather than plastic tie wraps to eliminate the risk of cable damage during cycling
- Includes a low-profile fixed tray to both support and retain the CMA in its fully closed position
- Both the CMA and the tray mount without the use of tools through simple and intuitive snap-in designs

The CMA can be mounted to either side of the sliding rails without the use of tools or the need for conversion. For systems with one power supply unit (PSU), it is recommended to mount on the side opposite to that of the power supply to allow easier access to it and the rear drives (if applicable) for service or replacement.



Strain Relief Bar

The optional Strain Relief Bar (SRB) for the system organizes and secures cables exiting the back of the server.

Sliding rails with optional SRB:

- Support tool-less attachment to rails
- Support two depth positions to accommodate various cable loads and rack depths
- Support cable loads and controls stress on server connectors
- Support cables can be segregated into discrete, purpose-specific bundles

11

Supported Operating Systems

The PowerEdge R450 system supports the following operating systems:

- Canonical® Ubuntu® Server LTS
- Citrix® Hypervisor®
- Microsoft® Windows Server® with Hyper-V
- Red Hat® Enterprise Linux
- SUSE® Linux Enterprise server
- VMware® ESXi®

Links to specific OS versions and editions, certification matrices, Hardware Compatibility Lists (HCL) portal, and Hypervisor support are available at Dell EMC Enterprise Operating Systems.

12

Dell EMC OpenManage systems management

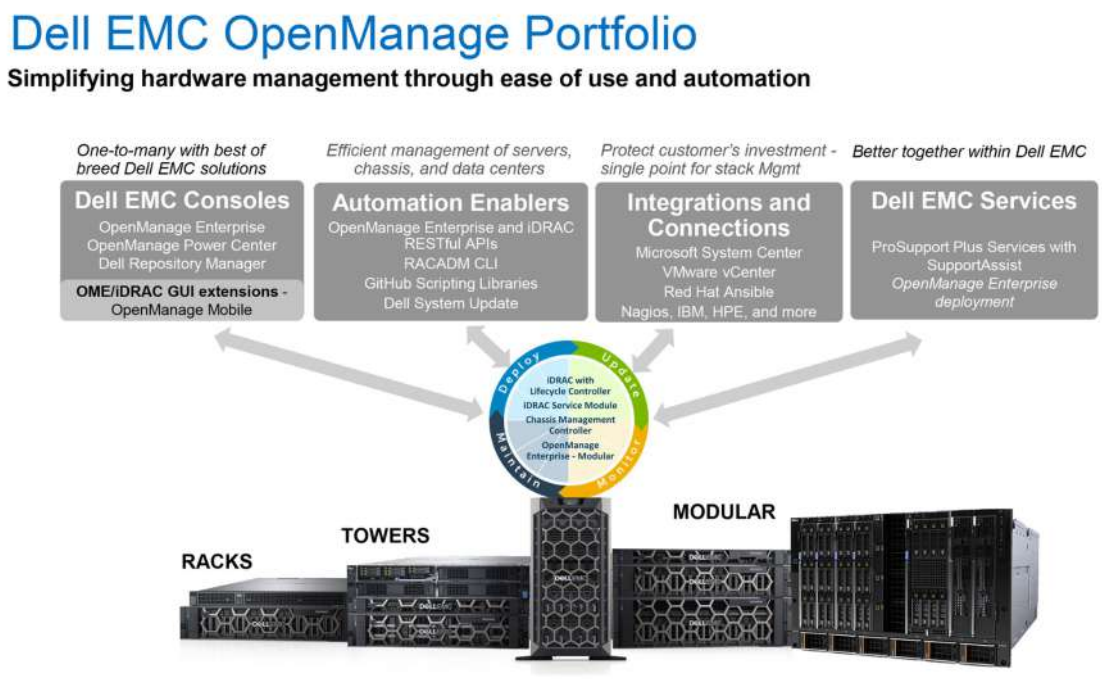


Figure 11. Dell EMC OpenManage Portfolio

Dell EMC delivers management solutions that help IT Administrators effectively deploy, update, monitor, and manage IT assets. OpenManage solutions and tools enable you to quickly respond to problems by helping them to manage Dell EMC servers effectively and efficiently; in physical, virtual, local, and remote environments, operating in-band, and out-of-band (agent-free). The OpenManage portfolio includes innovative embedded management tools such as the integrated Dell Remote Access Controller (iDRAC), Chassis Management Controller and Consoles like OpenManage Enterprise, OpenManage Power Manager plug in, and tools like Repository Manager.

Dell EMC has developed comprehensive systems management solutions based on open standards and has integrated with management consoles that can perform advanced management of Dell hardware. Dell EMC has connected or integrated the advanced management capabilities of Dell hardware into offerings from the industry's top systems management vendors and frameworks such as Ansible, thus making Dell EMC platforms easy to deploy, update, monitor, and manage.

The key tools for managing Dell EMC PowerEdge servers are iDRAC and the one-to-many OpenManage Enterprise console. OpenManage Enterprise helps the system administrators in complete lifecycle management of multiple generations of PowerEdge servers. Other tools such as Repository Manager, which enables simple yet comprehensive change management.

OpenManage tools integrate with systems management framework from other vendors such as VMware, Microsoft, Ansible, and ServiceNow. This enables you to use the skills of the IT staff to efficiently manage Dell EMC PowerEdge servers.

Topics:

- Server and Chassis Managers
- Dell EMC consoles
- Automation Enablers
- Integration with third-party consoles
- Connections for third-party consoles
- Dell EMC Update Utilities
- Dell resources

Server and Chassis Managers

- Integrated Dell Remote Access Controller (iDRAC)
- iDRAC Service Module (iSM)

Dell EMC consoles

- Dell EMC OpenManage Enterprise
- Dell EMC Repository Manager (DRM)
- Dell EMC OpenManage Enterprise Power Manager plugin to OpenManage Enterprise
- Dell EMC OpenManage Mobile (OMM)

Automation Enablers

- OpenManage Ansible Modules
- iDRAC RESTful APIs (Redfish)
- Standards-based APIs (Python, PowerShell)
- RACADM Command Line Interface (CLI)
- GitHub Scripting Libraries

Integration with third-party consoles

- Dell EMC OpenManage Integrations with Microsoft System Center
- Dell EMC OpenManage Integration for VMware vCenter (OMIVV)
- Dell EMC OpenManage Ansible Modules
- Dell EMC OpenManage Integration with ServiceNow

Connections for third-party consoles

- Micro Focus and other HPE tools
- OpenManage Connection for IBM Tivoli
- OpenManage Plug-in for Nagios Core and XI

Dell EMC Update Utilities

- Dell System Update (DSU)
- Dell EMC Repository Manager (DRM)
- Dell EMC Update Packages (DUP)
- Dell EMC Server Update Utility (SUU)
- Dell EMC Platform Specific Bootable ISO (PSBI)

Dell resources

For additional information about white papers, videos, blogs, forums, technical material, tools, usage examples, and other information, go to the OpenManage page at https://www.dell.com/openmanagemanuals or the following product pages:

Table 21. Dell resources

Resource	Location
Integrated Dell Remote Access Controller (iDRAC)	https://www.dell.com/idracmanuals
iDRAC Service Module (iSM)	https://www.dell.com/support/kbdoc/000178050/
OpenManage Ansible Modules	https://www.dell.com/support/kbdoc/000177308/
OpenManage Essentials (OME)	https://www.dell.com/support/kbdoc/000175879/
OpenManage Mobile (OMM)	https://www.dell.com/support/kbdoc/000176046
OpenManage Integration for VMware vCenter (OMIVV)	https://www.dell.com/support/kbdoc/000176981/
OpenManage Integration for Microsoft System Center (OMIMSSC)	https://www.dell.com/support/kbdoc/000147399
Dell EMC Repository Manager (DRM)	https://www.dell.com/support/kbdoc/000177083
Dell EMC System Update (DSU)	https://www.dell.com/support/kbdoc/000130590
Dell EMC Platform Specific Bootable ISO (PSBI)	Dell.com/support/article/sln296511
Dell EMC Chassis Management Controller (CMC)	www.dell.com/support/article/sln311283
OpenManage Connections for Partner Consoles	https://www.dell.com/support/kbdoc/000146912
OpenManage Enterprise Power Manager	https://www.dell.com/support/kbdoc/000176254
OpenManage Integration with ServiceNow (OMISNOW)	Dell.com/support/article/sln317784

NOTE: Features may vary by server. Please refer to the product page on https://www.dell.com/manuals for details.

13

Dell Technologies Services

Dell Technologies Services include a wide, customizable range of service choices to simplify the assessment, design, implementation, management and maintenance of IT environments and to help you transition from platform to platform. Depending on your current business requirements and the level of service right for you, we provide factory, on-site, remote, modular, and specialized services that fit your needs and budget. We'll help with a little or a lot—your choice—and provide access to our global resources.

For more information, see DellEMC.com/Services.

Topics:

- Dell EMC ProDeploy Enterprise Suite
- Dell EMC Data Migration Service
- Dell EMC ProSupport Enterprise Suite
- Dell EMC ProSupport Plus for Enterprise
- Dell EMC ProSupport for Enterprise
- Dell EMC ProSupport One for Data Center
- ProSupport for HPC
- Support Technologies
- Services for Data Security
- Dell Technologies Education Services
- Dell Technologies Consulting Services
- Dell EMC Managed Services

Dell EMC ProDeploy Enterprise Suite

ProDeploy Enterprise Suite gets your server out of the box and into optimized production—fast. Our elite deployment engineers with broad and deep experience utilizing best-in-class processes along with our established global scale can help you around the clock and around the globe. From simple to the most complex server installations and software integration, we take the guess work and risk out of deploying your new server technology.

		Basic Deployment	ProDeploy	ProDeploy Plus
Pre-deployment	Single point of contact for project management	-	●	In-region
	Site readiness review	-	●	●
	Implementation planning	-	●	●
	SAM engagement for ProSupport Plus entitled devices	-	-	●
Deployment	Deployment service hours	Business hours	24x7	24x7
	Remote guidance for hardware installation or Onsite hardware installation and packaging material removal	Onsite	Remote or Onsite	Onsite
	Install and configure system software	-	Remote	Onsite
	Install support software and connect with Dell Technologies	-	●	●
	Project documentation with knowledge transfer	-	●	●
Post-deployment	Deployment verification	-	●	●
	Configuration data transfer to Dell EMC technical support	-	●	●
	30-days of post-deployment configuration assistance	-	-	●
	Training credits for Dell EMC Education Services	-	-	●

Figure 12. ProDeploy Enterprise Suite capabilities

NOTE: Hardware installation not applicable on selected software products.

Dell EMC ProDeploy Plus

From beginning to end, ProDeploy Plus provides the skill and scale needed to successfully execute demanding deployments in today's complex IT environments. Certified Dell EMC experts start with extensive environmental assessments and detailed migration planning and recommendations. Software installation includes set up of most versions of Dell EMC SupportAssist and OpenManage system management utilities. Post-deployment configuration assistance, testing, and product orientation services are also available.

Dell EMC ProDeploy

ProDeploy provides full service installation and configuration of both server hardware and system software by certified deployment engineers including set up of leading operating systems and hypervisors as well as most versions of Dell EMC SupportAssist and OpenManage system management utilities. To prepare for the deployment, we conduct a site readiness review and implementation planning exercise. System testing, validation, and full project documentation with knowledge transfer complete the process.

Basic Deployment

Basic Deployment delivers worry-free professional installation by experienced technicians who know Dell EMC servers inside and out.

Dell EMC ProDeploy for HPC

HPC deployments require specialist that understand that cutting edge is yesterday's news. Dell EMC deploys the world's fastest systems and understands the nuances that make them perform. ProDeploy for HPC provides:

- Global team of dedicated HPC specialists
- Proven track record, thousands of successful HPC deployments
- Design validation, bench marking and production orientation

Learn more at http:/DellEMC.com/HPC-Services

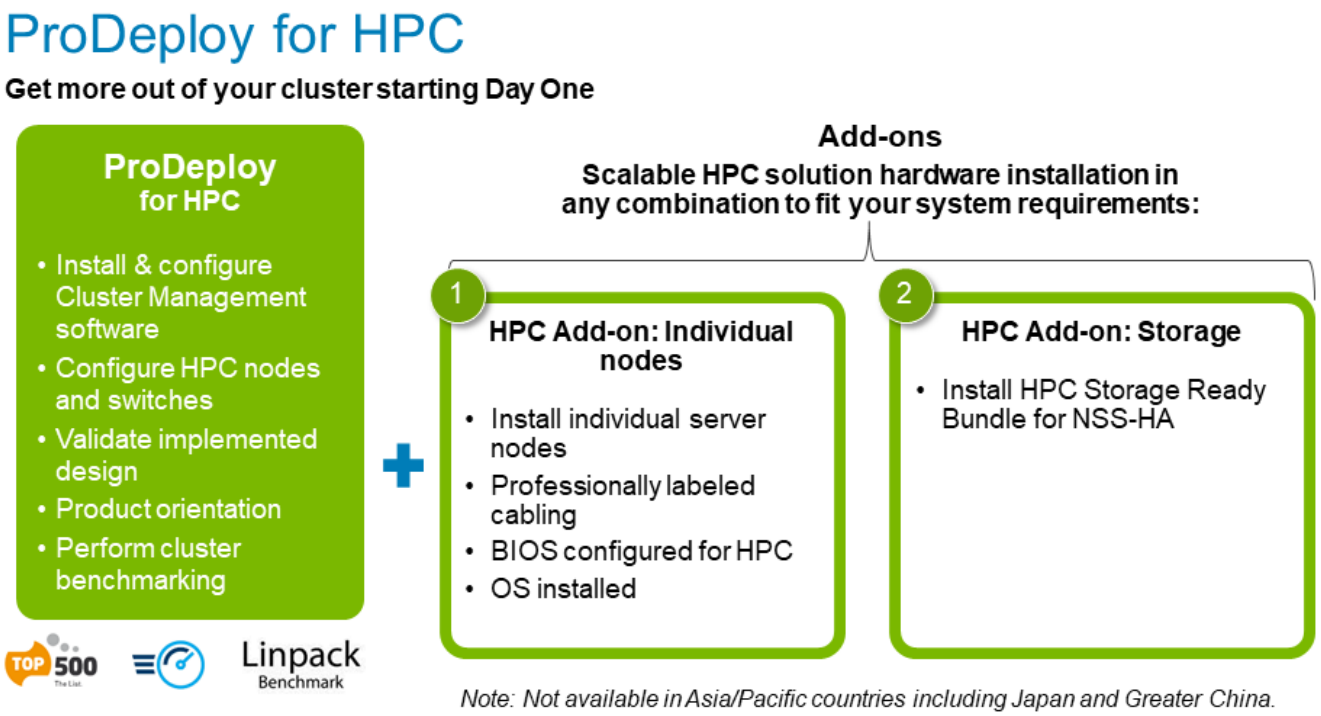


Figure 13. ProDeploy for HPC

Dell EMC Server Configuration Services

With Dell EMC Rack Integration and other Dell EMC PowerEdge Server Configuration Services, you save time by receiving your systems racked, cabled, tested, and ready to integrate into the data center. Dell EMC staff pre-configure RAID, BIOS and iDRAC settings, install system images, and even install third-party hardware and software.

For more information, see Server Configuration Services.

Dell EMC Residency Services

Residency Services helps customers transition to new capabilities quickly with the assistance of on-site or remote Dell EMC experts whose priorities and time you control. Residency experts can provide post implementation management and knowledge transfer related to a new technology acquisition or day-to-day operational management of the IT infrastructure.

Dell EMC Data Migration Service

Protect your business and data with our single point of contact to manage your data migration project. Your project manager will work with our experienced team of experts to create a plan using industry-leading tools and proven processes based on global best practices to migrate your existing files and data so your business system get up and running quickly and smoothly.

Dell EMC ProSupport Enterprise Suite

With the ProSupport Enterprise Suite, we help keep your IT systems running smoothly, so you can focus on running your business. We will help maintain peak performance and availability of your most essential workloads. ProSupport Enterprise Suite is a suite of support services that enable you to build the solution that is right for your organization.

Choose support models based on how you use technology and where you want to allocate resources. From the desktop to the data center, address everyday IT challenges, such as unplanned downtime, mission-critical needs, data and asset protection, support planning, resource allocation, software application management and more. Optimize IT resources by choosing the right support model.



Figure 14. Dell EMC ProSupport Enterprise Suite

Dell EMC ProSupport Plus for Enterprise

When you purchase your PowerEdge server, we recommend ProSupport Plus, our proactive and preventative support service for your business-critical systems. ProSupport Plus provides you with all the benefits of ProSupport, plus the following:

- An assigned Services Account Manager who knows your business and your environment
- Immediate advanced troubleshooting from an engineer who understands your PowerEdge server
- Personalized, preventive recommendations based on analysis of support trends and best practices from across the Dell Technologies infrastructure solutions customer base to reduce support issues and improve performance
- Predictive analysis for issue prevention and optimization enabled by SupportAssist
- Proactive monitoring, issue detection, notification, and automated case creation for accelerated issue resolution enabled by SupportAssist
- On-demand reporting and analytics-based recommendations enabled by SupportAssist and TechDirect

Dell EMC ProSupport for Enterprise

Our ProSupport service offers highly trained experts around the clock and around the globe to address your IT needs. We help minimize disruptions and maximize availability of PowerEdge server workloads with:

- 24x7 support through phone, chat and online
- Predictive, automated tools and innovative technology
- A central point of accountability for all hardware and software issues
- Collaborative 3rd party support
- Hypervisor, operating system and application support
- Consistent experience regardless of where you are located or what language you speak
- Optional onsite parts and labor response options including next business day or four-hour mission critical

NOTE: Subject to service offer country availability.

Enterprise Support Services
Feature Comparison

	Basic	ProSupport	ProSupport Plus
Remote technical support	9x5	24x7	24x7
Covered products	Hardware	Hardware Software	Hardware Software
Onsite hardware support	Next business day	Next business day or 4hr mission critical	Next business day or 4 hr mission critical
3rd party collaborative assistance		●	●
Automated issue detection & proactive case creation		●	●
Self-service case initiation and management		●	●
Access to software updates		●	●
Priority access to specialized support experts			●
3rd party software support			●
Assigned Services Account Manager			●
Personalized assessments and recommendations			●
Semiannual systems maintenance			●

Availability and terms of Dell Technologies services vary by region and by product. For more information, please view our Service Descriptions available on Dell.com

Figure 15. Dell EMC Enterprise Support model

Dell EMC ProSupport One for Data Center

ProSupport One for Data Center offers flexible site-wide support for large and distributed data centers with more than 1,000 assets. This offering is built on standard ProSupport components that leverage our global scale but are tailored to your company's needs. While not for everyone, this service option offers a truly unique solution for Dell Technologies largest customers with the most complex environments.

- Team of assigned Services Account Managers with remote, on-site options
- Assigned ProSupport One technical and field engineers who are trained on your environment and configurations
- On-demand reporting and analytics-based recommendations enabled by SupportAssist and TechDirect
- Flexible on-site support and parts options that fit your operational model
- A tailored support plan and training for your operations staff

ProSupport for HPC

The ProSupport for HPC provides solution-aware support including:

- Access to senior HPC experts
- Advanced HPC cluster assistance: performance, interoperability & configuration
- Enhanced HPC solution level end-to-end support
- Remote pre-support engagement with HPC Specialists during ProDeploy implementation

Learn more at DellEMC.com/HPC-Services.

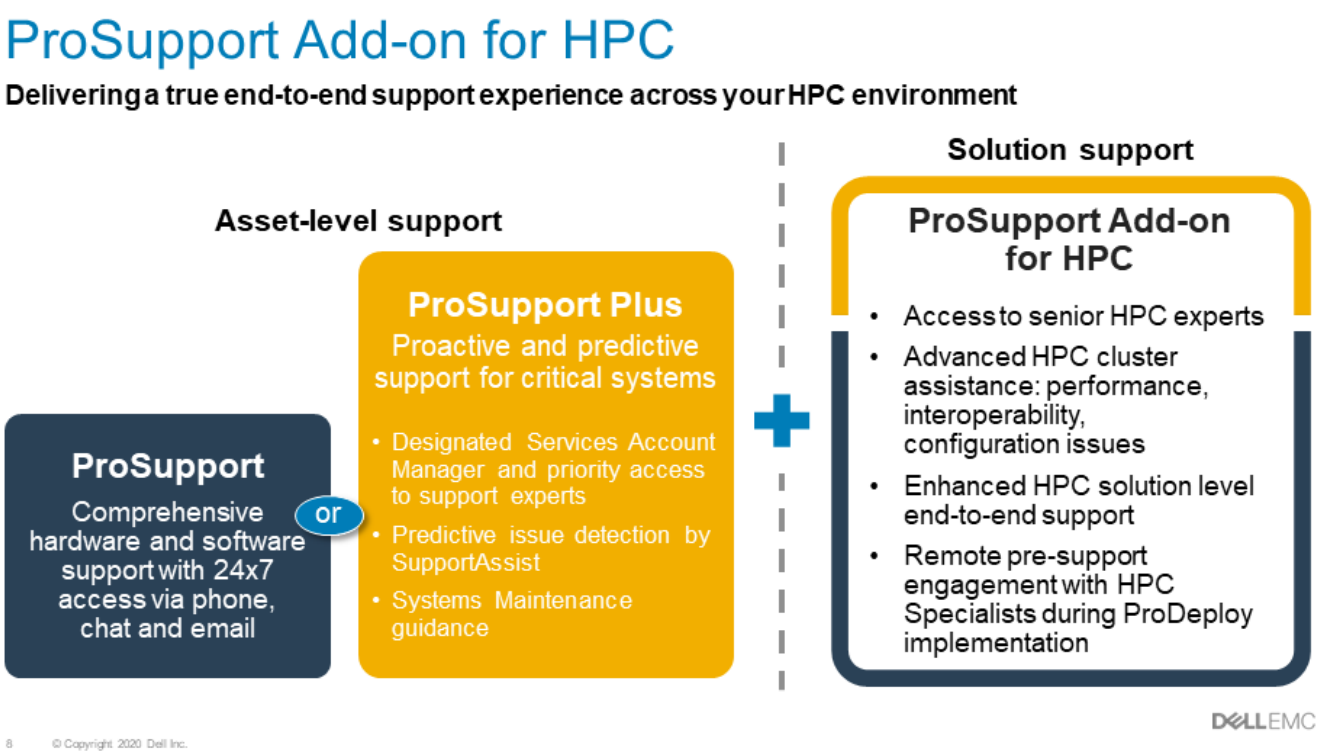


Figure 16. ProSupport for HPC

Support Technologies

Powering your support experience with predictive, data-driven technologies.

Dell EMC SupportAssist

The best time to solve a problem is before it happens. The automated proactive and predictive technology SupportAssist helps reduce steps and time to resolution, often detecting issues before they become a crisis. Benefits include:

- Value—SupportAssist is available to all customers at no additional charge
- Improve productivity—replace manual, high-effort routines with automated support
- Accelerate time to resolution—receive issue alerts, automatic case creation, and proactive contact from Dell EMC experts
- Gain insight and control—optimize enterprise devices with on-demand ProSupport Plus reporting in TechDirect, and get predictive issue detection before the problem starts

NOTE: SupportAssist is included with all support plans, but features vary based on service level agreement.

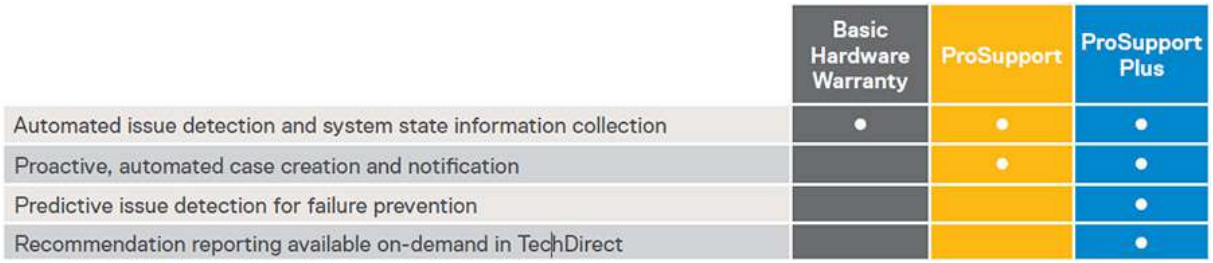


Figure 17. SupportAssist model

Get started at Dell.com/SupportAssist

Dell EMC TechDirect

Boost IT team productivity when supporting Dell EMC systems. With over 1.4 million self-dispatches processed each year, TechDirect has proven its effectiveness as a support tool. You can:

- Self-dispatch replacement parts
- Request technical support
- Integrate APIs into your help desk

Or, access all your Dell EMC certification and authorization requirements. Train your staff on Dell EMC products, as TechDirect allows you to:

- Download study guides
- Schedule certification and authorization exams
- View transcripts of completed courses and exams

Register at techdirect.dell.

Services for Data Security

As data security concerns intensify, businesses require focused security strategies to help mitigate risk. Get end to end protection throughout the life of your technology. Keep sensitive data on failed parts completely in your control with Dell EMC Keep Your Hard Drive and Keep Your Component for Enterprise or render data unrecoverable on re-purposed or retired products with Dell EMC Data Sanitization and Data Destruction for Enterprise. Promote social responsibility while maintaining data security with Dell EMC Data Sanitization for Enterprise Offsite with Asset Resale & Recycle we can help customers secure data on specific Dell EMC Server and Storage products as well as similar third-party systems. As part of this service, we remove old systems from your environment, securely sanitize data and responsibly reuse or recycle those systems to contribute to a more sustainable future. No matter your need, risk of unauthorized access to sensitive information is eliminated.

Dell Technologies Education Services

Build the IT skills required to influence the transformational outcomes of the business. Enable talent and empower teams with the right skills to lead and execute transformational strategy that drives competitive advantage. Leverage the training and certification required for real transformation.

Dell Technologies Education Services offers PowerEdge server training and certifications designed to help you achieve more from your hardware investment. The curriculum delivers the information and the practical, hands-on skills that you and your team need to confidently install, configure, manage, and troubleshoot your Dell EMC servers. To learn more or register for a class today, see LearnDell.com/Server.

Dell Technologies Consulting Services

Our expert consultants help you transform faster, and quickly achieve business outcomes for the high value workloads Dell EMC PowerEdge systems can handle.

From strategy to full-scale implementation, Dell Technologies Consulting can help you determine how to execute your IT, workforce, or application transformation.

We use prescriptive approaches and proven methodologies combined with Dell Technologies' portfolio and partner ecosystem to help you achieve real business outcomes. From multi-cloud, applications, DevOps, and infrastructure transformations, to business resiliency, data center modernization, analytics, workforce collaboration, and user experiences—we're here to help.

Dell EMC Remote Consulting Services

When you are in the final stages of your PowerEdge server implementation, you can rely on Dell EMC Remote Consulting Services and our certified technical experts to help you optimize your configuration with best practices for your software, virtualization, server, storage, networking, and systems management.

Dell EMC Managed Services

Reduce the cost, complexity, and risk of managing IT. Focus your resources on digital innovation and transformation while our experts help optimize your IT operations and investment with managed services backed by guaranteed service levels.

14

Appendix A. Additional specifications

Topics:

- Chassis dimensions
- Chassis weight
- Video specifications
- USB Ports
- PSU rating
- Environmental specifications

Chassis dimensions

The R450 has the following dimensions:

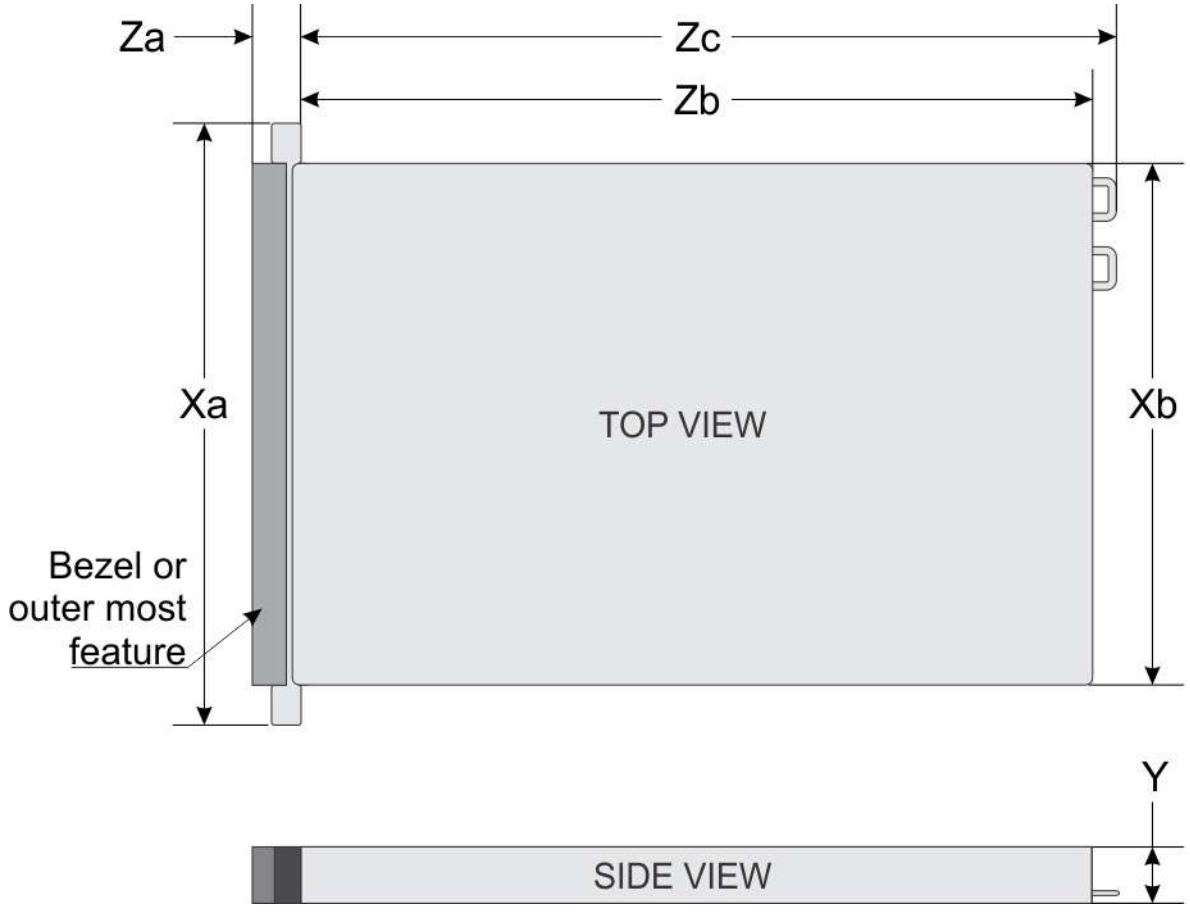


Figure 18. PowerEdge R450 Chassis dimensions

Table 22. Chassis dimensions

Drives	Xa	Xb	Y	Za	Zb	Zc
4 x 3.5-inch	482 mm	434 mm (17.08 inches)	42.8 mm(1.68	22 mm (0.866 inches) without bezel	677.8 mm (26.685 inches) (Ear to PSU surface)	712.95 mm (28.069 inches)

Table 22. Chassis dimensions (continued)

Drives	Xa	Xb	Y	Za	Zb	Zc
	(18.976 inches)		5 inches)	35.84 mm (1.41 inches) with bezel	691.07mm (27.207 inches) (Ear to butterfly L bracket housing)	(Ear to PSU handle without velcro strap)
8 x 2.5-inch	482 mm (18.976 inches)	434 mm (17.08 inches)	42.8 mm(1.685 inches)	22 mm (0.866 inches) without bezel 35.84 mm (1.41 inches) with bezel	627.03 mm (24.686 inches) (Ear to PSU surface) 640.3 mm (25.209 inches) (Ear to butterfly L bracket housing)	662.19 mm (26.070 inches) (Ear to PSU handle without velcro strap)

NOTE: Zb is the nominal rear wall external surface where the system board I/O connectors reside.

Chassis weight

Table 23. Chassis maximum weight

System configuration	Maximum weight (with all drives/SSDs/rails/bezel)
4 x 3.5-inch system	18.62 kg (41.05 pound)
8 x 2.5-inch system	16.58 kg (36.55 pound)

Video specifications

The platform supports Integrated Matrox G200 graphics controller with 16 MB of video frame buffer:

Table 24. Video Resolution and Refresh Rate

Resolution	Refresh rate (Hz)	Color depth (bits)
1024 x 768	60	8, 16, 32
1280 x 800	60	8, 16, 32
1280 x 1024	60	8, 16, 32
1360 x 768	60	8, 16, 32
1440 x 900	60	8, 16, 32
1600 x 900	60	8, 16, 32
1600 x 1200	60	8, 16, 32
1680 x 1050	60	8, 16, 32
1920 x 1080	60	8, 16, 32
1920 x 1200	60	8, 16, 32

*DVO - DP is for investigation only, dependent on Nuvoton DVO capabilities to support up to 165 MHz.Rear Panel Performance is TBD subject to final board design and losses to rear VGA connector

*(RB) - Reduced Blanking for Digital Displays requiring less blank time. This was introduced for Signal Integrity improvements by reducing Pixel Clock rates for VGA- Analog input devices.

USB Ports

All USB ports follow USB specifications.

USB 2.0 and USB 3.0 ports support maximum output current of 0.5 A and 0.9 A, respectively. The ports cannot support high power consumption devices such as CD-ROM on the rear USB port of the rear I/O board and on the right control panel USB 2.0 port.

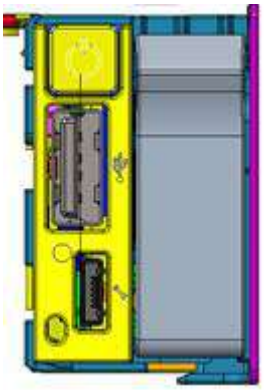


Figure 19. Front USB 2.0 Port



Figure 20. Rear USB

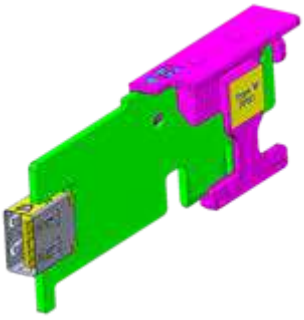


Figure 21. Internal USB

The size of the internal USB card dongle is 40 x 16 x 8 mm (L x W x H).

PSU rating

Below table lists the power capacity the PSUs in High/Low line operation mode.

Table 25. PSUs Highline and Lowline ratings

Feature	600 W Platinum	800 W Platinum	DC 1100 W
Highline	600 W	800 W	1100 W

Table 25. PSUs Highline and Lowline ratings (continued)

Feature	600 W Platinum	800 W Platinum	DC 1100 W
Lowline	600 W	800 W	1100 W
240 VDC	600 W	800 W	N/A
Highline 200-380 VDC	N/A	N/A	N/A
DC -48-60 V	N/A	N/A	1100 W

PowerEdge systems support number of AC or DC power supplies.

Dell PSUs have achieved Platinum efficiency levels as shown in the table below.

Table 26. PSU efficiency levels

Efficiency Targets by Load						
Form factor	**Output**	**Class**	**10%**	**20%**	**50%**	**100%**
Redundant 60 mm	600 W AC	Platinum	89.00%	93.00%	94.00%	91.50%
	800 W AC	Platinum	89.00%	93.00%	94.00%	91.50%
	1100W -48V	N/A	85.00%	90.00%	92.00%	90.00%

Environmental specifications

The table below details the environmental specifications for the platform. For additional information about environmental measurements for specific system configurations, see Product Safety, EMC and Environmental datasheets .

An important feature of having a broad menu of different categories is to allow the same platform model to have different operational ranges depending on the MRD defined.

A list of range categories for different configurations shall be identified by thermal team as early in the project as possible. Post release, see the Dell EMC PowerEdge R450 Technical Specifications at www.dell.com/poweredgemanuals.

Table 27. Operational Climatic Range Categories

Category A2	Allowable Operation
Temperature Ranges (For Altitude <900 meters or 2953 feet)	10 to 35°C (50 to 95°F) with no direct sunlight on the platform
Humidity Percent Ranges (Non-Condensing at all times)	8%RH with -12°C minimum dew point to 80%RH with 21°C (69.8°F) maximum dew point
Operational Altitude De-Rating	Maximum temperature is reduced by 1°C/300 meters (1.8°F/984 feet) above 900 meters (2,953 feet)

Table 28. Operational Climatic Range Categories

Category A3	Allowable Operation
Temperature Ranges (For Altitude <900 meters or 2953 feet)	5 to 40°C (41 to 104°F) with no direct sunlight on the platform
Humidity Percent Ranges (Non-Condensing at all times)	8%RH with -12°C minimum dew point to 85%RH with 24°C (75.2°F) maximum dew point
Operational Altitude De-Rating	Maximum temperature is reduced by 1°C/175meters (1.8°F/574 feet) above 900 meters (2,953 feet)

Table 29. Operational Climatic Range Categories

Category A3	Allowable Operation
Temperature Ranges (For Altitude <900 meters or 2953 feet)	5 to 45°C (41 to 113°F) with no direct sunlight on the platform
Humidity Percent Ranges (Non-Condensing at all times)	8%RH with -12°C minimum dew point to 90%RH with 24°C (75.2°F) maximum dew point
Operational Altitude De-Rating	Maximum temperature is reduced by 1°C/125 meters (1.8°F/410 feet) above 900 meters (2,953 feet)

The table below shows the requirements shared across all environmental categories:

Table 30. Shared Requirements

Allowable Operation	
Maximum Temperature Gradient (applies to both operation and non-operation)	20°C in an hour* (36°F in an hour) and 5°C in 15 minutes (9°F in 15 minutes), 5°C in an hour* (9°F in an hour) for tape hardware
Non-Operational Temperature Limits	-40 to 65°C (-40 to 149°F)
Non-Operational Humidity Limits (Non-Condensing at all times)	5% to 95%RH with 27°C (80.6°F) maximum dew point
Maximum Non-Operational Altitude	12,000 meters (39,370 feet)
Maximum Operational Altitude	3,048 meters (10,000 feet)

Table 31. Maximum vibration specifications

Maximum vibration	Specifications
Operating	0.26Grms at 5Hz to 350Hz (x, y, and z axes)
Storage	1.88Grms at 10Hz to 500Hz for 15min (all six sides tested)

Table 32. Maximum shock specifications

Maximum vibration	Specifications
Operating	Six consecutively executed shock pulses in the positive and negative x, y, and z axes of 6G for up to 11ms
Storage	Six consecutively executed shock pulses in the positive and negative x, y, and z axes (one pulse on each side of the system) of 71 G for up to 2 ms

Thermal restrictions

Table 33. Thermal Solution Configuration

Configurations		CPU	Fan Type/ Q'ty	HSK Type	Air Shroud	DIMM blank	CPU BlanK	Fan Counts	Fan Blank
SM Config	Rear Wall Config	TDP							
3.5" HDDs x4	W/O Rear HDDs	TDP <= 165W	STD Fan HPR SLVR (HPR) Fan for 165W CPU only	STD HS HPR HS for 165W CPU only	Yes	No	Only required on CPU2 for 1 Processor Config	7x Fans for 2 Processors Config 5x Fans for 1 Processor Config	Only required on Fan Slot 1 & Slot 2 for 1 Processor Config

Table 33. Thermal Solution Configuration (continued)

Configurations		CPU	Fan Type/ Q'ty	HSK Type	Air Shroud	DIMM blank	CPU BlanK	Fan Counts	Fan Blank
SM Config	Rear Wall Config	TDP							
		165W <TDP<= 220W	HPR Silver Fan (HPR)	HPR HS					
2.5" SAS/ SATA x8	W/O Rear HDDs	TDP<= 165W	STD Fan HPR SLVR (HPR) Fan for 165W CPU only	STD HS HPR HS for 165W CPU only	Yes	No	Only required on CPU2 for 1 Processor Config	7x Fans for 2 Processors Config 5x Fans for 1 Processor Config	Only required on Fan Slot 1 & Slot 2 for 1 Processor Config
		165W <TDP<= 220W	HPR Silver Fan (HPR)	STD HS					

NOTE: HDD Blanks are required for empty HDD slots.

Table 34. Thermal Restriction of 8x2.5 SAS/SATA & 4x3.5" Storage Configuration

Dell EMC PowerEdge Server Standard Operating Support (ASHRAE A2 compliant) All options supported unless otherwise noted.	Dell EMC PowerEdge Server Extended Inletient 40 ° C Operating Support (ASHRAE A3 compliant)	Dell EMC PowerEdge Server Extended Inletient 45 ° C Operating Support (ASHRAE A4 compliant)
• HPR Silver (HPR) Fan is required for CPU >165W • The Following OCP3.0 NIC only support optic cable with thermal Spec 85C and power <=1.2W ○ Intel Columbiaville DP 25GbE SFP28 8x2.5 SAS/SATA configuration only ○ Broadcom Thor QP 25G SFP28 both Configurations ○ Mellanox CX5 DP 25GbE SFP28 both configurations ○ Solarflare Medford2 DP 25GbE SFP28 both configurations • The following PCIe NIC only support optic cable with thermal spec 85°C and power <= 1.2 W: ○ Solarflare Medford2 DP 25 GbE SFP28 in 8x2.5-inch SAS/SATA configuration. ○ Broadcom 100 G 2P QSF in both configurations. ○ Mellanox CX6 DP 25 G SFP28 in both configurations. • The following PCIe NIC only support optic cable with thermal spec 85°C and power <= 2.5 W ○ Mellanox CX6 DP 100 GbE in both configurations ○ Intel Columbiaville 100 G 2P Q28 in both configurations • PCIe SSD: Intel P4800X 750G and 375G could only support in PCIe slot2 and PCIe slot3 in 4x3.5 configs.	• Not support CPU TDP > 150W • Not support BOSS M.2 Module • Not support Non-Dell qualified peripheral cards. • Not support PCIe NIC consuming power >= 25W. Ex: CX6 card • Not Support Config with RM • Not support OCP transfer rate >25G or cooling tier > 10 • Optic Cable with spec 85C is required. • Two PSUs are required. System performance may be reduced in the event of a PSU failure	A4 environment not Support

Table 34. Thermal Restriction of 8x2.5 SAS/SATA & 4x3.5" Storage Configuration

Dell EMC PowerEdge Server Standard Operating Support (ASHRAE A2 compliant) All options supported unless otherwise noted.	Dell EMC PowerEdge Server Extended Inletient 40 ° C Operating Support (ASHRAE A3 compliant)	Dell EMC PowerEdge Server Extended Inletient 45 ° C Operating Support (ASHRAE A4 compliant)
No restriction in 8x2.5 SAS/SATA config.		

15

Appendix B. Standards compliance

The system conforms to the following industry standards.

Table 35. Industry standard documents

Standard	URL for information and specifications
ACPIAdvance Configuration and Power Interface Specification, v2.0c	https://uefi.org/specsandtesttools
Ethernet IEEE 802.3-2005	https://standards.ieee.org/
HDGHardware Design Guide Version 3.0 for Microsoft Windows Server	microsoft.com/whdc/system/platform/pcdesign/desguide/serverdg.mspx
IPMI Intelligent Platform Management Interface, v2.0	intel.com/design/servers/ipmi
DDR4 Memory DDR4 SDRAM Specification	jedec.org/standards-documents/docs/jesd79-4.pdf
PCI Express PCI Express Base Specification Rev. 2.0 and 3.0	pcisig.com/specifications/pciexpress
PMBus Power System Management Protocol Specification, v1.2	http://pmbus.org/Assets/PDFS/Public/PMBus_Specification_Part_I_Rev_1-1_20070205.pdf
SAS Serial Attached SCSI, v1.1	http://www.t10.org/
SATA Serial ATA Rev. 2.6; SATA II, SATA 1.0a Extensions, Rev. 1.2	sata-io.org
SMBIOS System Management BIOS Reference Specification, v2.7	dmtf.org/standards/smbios
TPM Trusted Platform Module Specification, v1.2 and v2.0	trustedcomputinggroup.org
UEFI Unified Extensible Firmware Interface Specification, v2.1	uefi.org/specifications
USB Universal Serial Bus Specification, Rev. 2.0	usb.org/developers/docs

16

Appendix C Additional resources

Table 36. Additional resources

Resource	Description of contents	Location
Installation and Service Manual	This manual, available in PDF format, provides the following information: • Chassis features • System Setup program • System indicator codes • System BIOS • Remove and replace procedures • Diagnostics • Jumpers and connectors	Dell.com/Support/Manuals
Getting Started Guide	This guide ships with the system, and is also available in PDF format. This guide provides the following information: • Initial setup steps	Dell.com/Support/Manuals
Rack Installation Guide	This document ships with the rack kits, and provides instructions for installing a server in a rack.	Dell.com/Support/Manuals
System Information Label	The system information label documents the system board layout and system jumper settings. Text is minimized due to space limitations and translation considerations. The label size is standardized across platforms.	Inside the system chassis cover
Quick Resource Locator (QRL)	This code on the chassis can be scanned by a phone application to access additional information and resources for the server, including videos, reference materials, service tag information, and Dell EMC contact information.	Inside the system chassis cover
Energy Smart Solution Advisor (ESSA)	The Dell EMC online ESSA enables easier and more meaningful estimates to help you determine the most efficient configuration possible. Use ESSA to calculate the power consumption of your hardware, power infrastructure, and storage.	Dell.com/calc

Příloha č. 3 smlouvy – Realizační tým

Seznam techniků (členů realizačního týmu)	
Pozice	**Jméno a příjmení**
Projektový manažer	XXX
Technický specialista – kybernetická bezpečnost	XXX
Technický specialista – správa logů	XXX
Technický specialista – virtualizace	XXX
Technický specialista – ověřovací systémy	XXX
Technický specialista – ochrana digitálních dat	XXX
Technický specialista – systémy pro správu privilegovaných účtů a přístupů	XXX

Žlutě podbarvené části doplní uchazeč

Označení	Cena bez DPH v Kč	Cena bez DPH v Kč	Cena bez DPH v Kč	Cena bez DPH v Kč	Cena celkem bez DPH v Kč	DPH v zákonné výši v Kč	Cena celkem včetně DPH v Kč
	1-12 měsíc	13-24 měsíc	25-36 měsíc	37-48 měsíc	za 48 měsíců	za 48 měsíců	za 48 měsíců
Základní servisní podpora - služby	92 800 Kč	92 800 Kč	95 500 Kč	98 300 Kč	379 400 Kč	79 674 Kč	459 074 Kč
Rozšířená servisní podpora - služby	24 600 Kč	24 600 Kč	24 600 Kč	24 600 Kč	98 400 Kč	20 664 Kč	119 064 Kč
Celková cena za zabezpečení provozu					477 800 Kč	100 338 Kč	578 138 Kč

Název služby	Hodinová sazba bez DPH v Kč	Hodinová sazba včetně DPH v Kč
Hodinová sazba za služby rozšířené servisní podpory nad rámec základního rozsahu (1 hodina měsíčně, použitá sazba nesmí být vyšší než sazba použitá ve výpočtu ceny pro rozšířenou servisní podporu)	2 050 Kč	2 481 Kč

Pozn. Hodinová sazba je již zahrnuta v hodnocení nabídkové ceny, položka "Rozšířená servisní podpora", položka se tedy samostatně nehodnotí a slouží jen k jednoznačnému potvrzení hodinové sazby použité pro stanovení nabídkové ceny položky "Rozšířená servisní podpora" dle podmínek ZD.

Žlutě podbarvené části doplní dodavatel