



CONTRACT FOR BUILDING OF A SUPERCOMPUTER – LARGE CLUSTER II

Contracting parties:

Vysoká škola báňská – Technická univerzita Ostrava (VSB – Technical University of Ostrava)
having its registered office at: 17. listopadu 2172/15, 708 00 Ostrava – Poruba, Czech Republic

represented by: prof. RNDr. Václav Snášel, CSc., Rector

corporate ID: 61989100, tax ID: CZ61989100

Bank: ČSOB, a.s., bank account no.: 100954151/0300

(hereinafter “**VSB-TUO**”)

The European High-Performance Computing Joint Undertaking

having its registered office at: 12E, rue Guillaume Kroll, L-1882 Luxembourg

represented by [REDACTED] Executive Director

(hereinafter “**EuroHPC JU**”,

VSB-TUO and EuroHPC JU hereinafter jointly as the “**Client**”)

and

HEWLETT-PACKARD s.r.o.

having the registered office at: Za Brumlovkou 1559/5, 140 00 Praha 4 – Michle Czech Republic

represented by: [REDACTED] Statutory Representative

recorded in the Register of Companies held by Municipal Court in Prague,

File C, Insert 1974

Corporate ID: 17048851, Tax ID: CZ17048851

Bank: Československá obchodní banka, a.s., bank account no.: Bank account 722513/0300

(hereinafter the “**Supplier**”)

concluded this contract on today’s date under Section 2586 et seq. of Czech Act No. 89/2012 Coll.,
Civil Code, as amended (hereinafter the “**Civil Code**”),

(hereinafter the “**Contract**”)

The contracting parties, aware of their obligations contained herein and intending to be bound by this Contract, have agreed on the following wording of the Contract:

1. INTRODUCTORY PROVISIONS

1.1 The Client declares that:

1.1.1 Vysoká škola báňská – Technical University Ostrava (hereinafter “**VSB-TUO**”) is a corporate entity, public university founded under Act No. 111/1998 Coll., on Higher Education Institutions and on Amendments and Supplements to Some Other Acts (the Higher Education Act), as amended,

1.1.2 The European High-Performance Computing Joint Undertaking (hereinafter “**EuroHPC JU**”) is a corporate entity established by the European Union Commission pursuant to Council Regulation (EU) 2018/1488 of 28th September 2018, and

1.1.3 The Client overall meets all conditions and requirements stipulated herein and is entitled to conclude this Contract and duly meet the obligations contained herein.

1.2 The Supplier declares that:

1.2.1 it is a corporate entity duly established and existing under the law of the country of its registered office; and

1.2.2 it meets all the conditions and requirements stipulated herein and is entitled to conclude this Contract and duly meet the obligations contained herein.

1.3 On 10. 04. 2020, VSB-TUO announced its intention to award a public contract entitled “**EURO_IT4I Supercomputer**” (hereinafter the “**Public Contract**”) by sending the notification of a start of a procurement procedure in the Tenders Electronic Daily under Act No. 134/2016 Coll., on Public Procurement, as amended (hereinafter the “**Act**”). Pursuant to this procurement procedure, the tender of the Supplier was selected as the most suitable for the implementation of the Public Contract under Section 122 (1) of the Act.

1.4 In relation to VSB-TUO, EuroHPC JU is a subsidy provider of the project entitled “IT4Innovations centre for European science and industry” financed from EuroHPC JU under Hosting Agreement no. 05/2019. In addition, EuroHPC JU is a contracting party in this Contract and it shall be directly financially involved in the acquisition of a supercomputer – Large Cluster and shall obtain ownership rights to the Work in such percentage as the percentage of financing by EuroHPC JU of the total price as defined in par. 12.1 hereof.

2. PURPOSE OF THE CONTRACT

2.1 The Client, as the contracting entity of the Public Contract intends to expand its facilities and equipment in the “IT4Innovations National Supercomputer Centre – Path to exascale” project, reg. no. CZ.02.1.01/0.0/0.0/16_013/0001791 (hereinafter “**P2E**”) financed from the European Fund for regional development under the Operational Programme Research, Development and Education (hereinafter the “**OP RDE**”) and in the IT4Innovations center for European science and industry project (hereinafter “**EURO-IT4I**”) financed from EuroHPC JU in accordance with Hosting Agreement no. 05/2019 (hereinafter “**EuroHPC**”).

2.2 The subject matter of the Public Contract is the acquisition of a Large Cluster which is the second one for IT4Innovations, i.e. comprehensive technical provision of research infrastructure in “*supercomputing*” (hereinafter the “**Large Cluster**”). The solution that was selected as performance under the Public Contract is a comprehensive solution for complex

calculations (supercomputer) consisting of technical equipment, program equipment, licences, design, implementation and configuration of the system, including warranty, maintenance and training.

- 2.3 The implementation of the Public Contract shall result in a creation of a computing environment with high computing performance, high data capacity, high communication speed, unified user environment, support of parallel environments and tasks, high-performing shared user data space, single administration of sources, tasks and users and comprehensive monitoring, intended for calculations of various complex tasks. The computing environment shall be, as a whole, integrated and fully operational in the VSB-TUO's infrastructure providing for sufficient power and cooling capacities.
- 2.4 The Supplier acknowledges that the adherence to the time schedule determined hereunder has an impact on the compliance with the obligations of VSB-TUO to the provider of the subsidy from OP RDE and EuroHPC JU and compliance with other educational and research and development obligations of VSB-TUO. The Supplier additionally acknowledges the possibility of a considerable damage that might be caused if the deadlines stipulated herein are not adhered to.
- 2.5 The purpose of this Contract is the implementation of the Public Contract consisting in the building of the Large Cluster, in accordance with the requirements of the Client defined herein, including the provision of ownership entitlements of the Client to the provided performance and provision of authorisations of the Client to any use and change in the supplied performance and its parts, which are copyright works, without the need for further consent of the Supplier or third parties.
- 2.6 By this Contract, the Supplier guarantees to the Client that it shall complete the assignment by the Public Contract and meet all arising conditions and obligations assumed by the Supplier in the procurement procedure of the Public Contract in line with the tender dossier, specifically Technical Assignment by the Client, contained in Annex 1 hereto, and tender of the Supplier. This guarantee is superior to other conditions and guarantees stated herein. For the avoidance of doubt, it means that:
 - 2.6.1 in the event of any uncertainty regarding the interpretation hereof, these provisions shall be interpreted to reflect, to the maximum extent possible, the purpose of the Public Contract expressed by the tender dossier of the Public Contract;
 - 2.6.2 in the event of missing provisions hereof, sufficiently specific provisions of the tender dossier of the Public Contract shall be used.
- 2.7 The Supplier shall be bound by its tender submitted to the Client in the procurement procedure for the Public Contract which shall have subsidiary application in the definition of mutual relationships arising herefrom.

3. SUBJECT MATTER OF THE CONTRACT

- 3.1 By this Contract, the Supplier undertakes to complete the work for the Client consisting in:
 - 3.1.1 building and supply of the Large Cluster, as a system for complex calculations, as it is defined in Annex 1 hereto, containing the Basic Technical Specification (composed of the Design of the Technical Solution, Technical Parameters of the Proposal and Technical Assignment of the Client), including the supply of all components, software, user, operational and administrator documentation and project management (hereinafter the "**System**");
-

3.1.2 implementation of the System in the infrastructure of the data room of VSB-TUO, IT4Innovations, specified in Annex 1 to the tender dossier of the Public Contract (hereinafter the “**Data Room**”), including the provision of functional connection to the electrical power supply and cooling and training of employees of VSB-TUO (the implementation and training hereinafter as the “**Implementation**”);

(the System including the Implementation hereinafter as the “**Work**”);

all this in line with the specification stated in the Basic Technical Specification, and after its completion, in line with the Detailed Technical Specification, specified below; the Work involves the provision of warranty for quality of the Work in accordance with article 11 and other obligations of the Supplier in accordance with article 7 hereof.

- 3.2 As part of the provision of performance hereunder, the Supplier undertakes to prepare, pursuant to the Basic Technical Specification, a detailed technical specification of the Work (hereinafter the “**Detailed Technical Specification**”) that shall specify, in detail, individual parts of the Work as well as the procedure and conditions in its performance, including the definition of the required cooperation of the Client and other issues the specification of which is necessary for the provision of performance hereunder (e.g. par. 7.4 hereof). The Detailed Technical Specification shall contain a structured summary of all documentation of individual used system components, especially product sheets and manuals for the components.
- 3.3 After it is approved by the Client, the Detailed Technical Specification shall become an integral part hereof and a binding specification of the Work. The preparation of the Detailed Technical Specification shall be considered as part of performance hereunder in terms of the contractual obligations of the Supplier and the Supplier shall be not entitled to other remuneration than the agreed price of the Work for its preparation. The preparation and approval of the Detailed Technical Specification shall not decrease the quality of the Work, shall not change the agreed price under par. 12.1 hereof, deadlines for binding milestones of the time schedule for works listed in Annex 2 hereto, or the scope of the maximum cooperation of the Client listed in Annex 3 hereto; however, the possibility of the contracting parties to conclude an amendment hereto under par. 20.1 hereof for this purpose shall be not affected.
- 3.4 Where this Contract uses the “**Technical Specification**” term, it means both the Basic Technical Specification and the Detailed Technical Specification, if approved by the Client, and provisions of the Detailed Technical Specification, after approved by the Client, prevail over provisions of the Basic Technical Specification.
- 3.5 The Supplier additionally undertakes to prepare, as part of the Work, the specification of acceptance tests defining the specification of the manner of accepting individual parts of the Work, if, given their nature, their handover and acceptance should be preceded by acceptance tests, primarily verification of due Implementation (hereinafter the “**Specification of Acceptance Tests**”). The Specification of Acceptance Tests shall be prepared in line with the requirements of the Client for acceptance procedures and acceptance tests listed herein, primarily Annex 9 hereto. The requirements for acceptance tests shall primarily reflect individual binding milestones defined in the time schedule in accordance with Annex 2 hereto.
- 3.6 The performance provided by the Supplier includes the transfer of ownership rights to all elements of the Work to the Client as well as the transfer of intellectual property rights that are transferable (e.g. patents etc.) for the case that it is necessary to exercise these intellectual property rights to ensure proper and smooth use of the Work. In addition, the performance provided by the Supplier includes the provision of an authorisation to the Client to use and change intangible parts of the Work where the intellectual property rights are not transferable
-

(e.g. copyright), in the scope stipulated hereby. The provision or transfer of the rights in accordance with this paragraph shall not apply to the warranty conditions for the Work.

- 3.7 By this Contract, the Client undertakes to provide the Supplier with necessary cooperation in the provision of performance hereunder in the scope determined in Annex 3 hereto.
- 3.8 The Client undertakes to pay the agreed price for the provided performance to the Supplier hereunder in a due and timely manner, under the conditions stipulated herein.
- 3.9 The Supplier undertakes to provide performance hereunder itself or using its subcontractors. Any additional change in a subcontractor or scope of performance entrusted to a subcontractor shall be approved by the Client in writing in advance, unless the performance originally entrusted to a subcontractor is provided by the Supplier itself. The Client undertakes that it shall not withhold or refuse its consent without a serious reason that it shall immediately communicate to the Supplier; in such case, the parties shall act in good faith in order to approve the change without unnecessary delays.

4. TIME AND PLACE OF PERFORMANCE

- 4.1 The Supplier undertakes to provide all performance that form the Work and hand it over to the Client in line with the deadlines stated in the time schedule appended in Annex 2 hereto.
 - 4.2 The Supplier undertakes to provide the Implementation at an agreed place of performance in the period agreed in the time schedule.
 - 4.3 The detailed time schedule of provision of performance hereunder, containing, among other things, a detailed time schedule of works in the Data Room, shall be specified in the Detailed Project Plan that shall be prepared by the Supplier and it becomes binding for the contracting parties by its acceptance in line with the procedure listed in par. 6.3 hereof (if this Contract uses the “**Time Schedule**” it means the relevant part of the Detailed Project Plan which contains a detailed time schedule in Annex 2 hereto). In the event of any conflict between the detailed Project Plan and the Detailed Technical Specification, the Detailed Project Plan shall prevail, unless the Detailed Technical Specification is approved by an amendment to the Contract with the will of the parties to derogate from the Detailed Project Plan, primarily in the following cases:
 - 4.3.1 if obstacles on the part of the Client arise during the performance of the Work that, taking into account all circumstances, objectively prevent the Supplier to continue the due performance of the Work;
 - 4.3.2 if an interruption of work by the Supplier during the performance of the Work was caused by obstacles excluding the obligation to provide indemnification under Section 2913 (2) of the Civil Code.
 - 4.4 The detailed Project Plan must adhere to and be prepared in accordance with the project management rules as determined in Annex 8 hereto.
 - 4.5 The Supplier shall be entitled to provide the performance before the agreed deadline only with the prior consent of an authorised person of the Client.
 - 4.6 The place of performance shall be the Data Room at the registered office of VSB-TUO, unless the Detailed Technical Specification explicitly stipulates otherwise. The access to the Data Room during the System Implementation shall be provided to the Supplier on a one-off basis, or repeatedly, depending on circumstances, always for a pre-determined period in line with an agreement between the Supplier and VSB-TUO. The request of the Supplier for granting of access to the data Room shall be sent to VSB-TUO no later than 2 working days in advance.
-

The contracting parties shall prepare a protocol on the provided access to the Data Room and termination of access of the Supplier to the Data Room which shall contain no less than the data and time of start and termination of access, definition of things entrusted to the Supplier for use for the purpose of due performance hereunder, or installed and left by the Supplier in the Data Room, including the description of their condition and definition of other facts having an impact on the possibility of the Supplier to duly start the performance of the subject matter hereof. The contracting parties undertake to sign the written protocol prepared in accordance with this paragraph hereof and prepare at least 4 copies, one original copy always for the Supplier, one for EuroHPC JU, one for VSB-TUO and one possibly for another supplier to VSB-TUO. For the avoidance of doubts, the condition of the place of performance and any return of things entrusted to the Supplier shall be subject to acceptance in accordance with article 6 hereof or protocol certification in a manner analogic in accordance with this paragraph hereof. Over the period when the Supplier has reserved access to the data Room for the Implementation of the System, VSB-TUO and its other Suppliers shall abstain from the access to the Data Room without prior agreement of VSB-TUO with the Supplier. After the completion of the Implementation, the access of the Supplier shall be provided as needed after prior agreement with VSB-TUO; authorised persons of the contracting parties may conclude a separate agreement for this purpose regarding the mode of access to the Data Room for purposes of meeting of the Supplier's obligations arising from the Contract.

- 4.7 If the nature of a specific part of performance hereunder allows it, it shall be possible to provide it or its part by remote access. If the Contract does not explicitly stipulate whether a specific part of performance should be carried out by the Supplier at a specific place, or by remote access, and the nature of the performance allows both options, VSB-TUO shall be entitled to select between these two manners at its discretion without an impact on the agreed price hereunder. If the performance or its part is to be provided by remote access, VSB-TUO shall enable such remote access to the Supplier. The costs relating to the remote access are borne by the Supplier, unless the Contract stipulates otherwise.

5. CHANGE MANAGEMENT

- 5.1 Any of the contracting parties is authorised to propose changes in the specification of the Work in writing during the Contract term. If a change in the specification is proposed by the Client, the Supplier shall be obliged to make every effort to accept the change in the specification. The Client shall not be obliged to accept a change in the specification proposed by the Supplier.
- 5.2 At a written call of the Client, the Supplier undertakes to assess the consequences of proposed changes in the specification of the Work within 10 working days, which shall include an assessment of impacts of these changes on the price and scope of the Work, agreed deadlines for performance, scope of necessary cooperation and any other relevant aspects of the contractual relationship (hereinafter the "**Assessment of Consequences**"). If the Assessment of Consequences requires additional costs or if its preparation could have a negative impact on the meeting of the obligations of the Supplier hereunder, the Supplier shall prepare the Assessment of Consequences pursuant to a written agreement with the Client on the payment of costs for the preparation of the Assessment of Consequences and adjustment to other contractual conditions that the preparation of the Assessment of Consequences may affect.
- 5.3 Any changes in the specification of the Work shall be agreed in the form of a written amendment hereto in accordance with par. 20.1 hereof adjusting contractual conditions in line with the Assessment of Consequences, unless the Contract stipulates otherwise.
- 5.4 Detailed rules of change management are stipulated in Annex 8 hereto.
-

- 5.5 The Supplier acknowledges that any change to the Contract shall not be in conflict with the binding rules of legal regulations regarding the public procurement law or rules of financing and eligibility of costs from OP RDE and EuroHPC. Where the contracting parties agree with an expansion of the Work or another change in the specification of the Work in the interest of achieving the purpose hereof and provision of the requested features of the Work under this article 5 hereof, while maintaining or reducing the original price for the Work, it shall not constitute an unjust enrichment of the Supplier.

6. ACCEPTANCE AND HANDOVER OF PERFORMANCE

- 6.1 The handover and acceptance of the Work shall adhere to the acceptance procedure defined in this article 6 hereof. For the avoidance of doubts, the contracting parties declare that provisions of article 6 shall not apply to performance listed in Annex 6 hereto.
- 6.2 The Supplier undertakes to hand over the subject matter of acceptance to the Client for acceptance in a time period not threatening the binding milestones for the provision of the performance determined in the Time Schedule. Each individual milestone in accordance with the Time Schedule is a moment as of which the stage of the Work implementation should adhere to the provision of the Work hereunder.
- 6.3 Regardless other provisions hereof, the following rules shall apply to the preparation of any documents pursuant to individual milestones created by the Supplier hereunder:
- 6.3.1 The Supplier undertakes to consult the preparation of documents with the Client on an ongoing basis and deliver the relevant document to the Client in a period determined in the Time Schedule.
- 6.3.2 The Client undertakes to send a notification of its consent with the specific document or its comments to the document to the Supplier within 5 working days from its delivery. If the Client does not raise any objections or does not make any comments to the document in the determined period or does not communicate its consent with the specific document to the Supplier, the contracting parties treat the document in this version as duly delivered and binding for the contracting parties after the expiration of this period and signing of the handover protocol in accordance with par. 6.3.4.
- 6.3.3 If the Client raises its objections or makes comments regarding the document in accordance with par. 6.3.2, the Supplier undertakes to present the adjusted document to the Client for acceptance.
- 6.3.4 After acceptance or expiry of the time limit for raising of objections or comments to the document in accordance with article 6.3.2, the contracting parties undertake to confirm the handover and acceptance of the document by preparation of a written handover protocol. The acceptance takes effect by signing of the handover protocol by both contracting parties. The signing of the handover protocol cannot be refused by a contracting party, otherwise, it is deemed that the document is binding and approved after the expiration of the 3rd working day after the call made by the other contracting party for the signing of the handover protocol.
- 6.4 The Work acceptance procedure shall include ongoing acceptance tests that shall be based on the Specification of Acceptance Tests. The Specification of Acceptance Tests shall allow for a verification whether the Work has the required features agreed herein. The Client reserves the right to perform its own verifications tests, in cooperation with the Supplier, that do not arise from Annex 9 hereto and/or the Specification of Acceptance Tests and the result of which has
-

no impact on the final acceptance of the Work itself. For the avoidance of doubt, the contracting parties state that the acceptance procedure of the Work will be divided into two parts, which will reflect the gradual delivery and implementation of the System pursuant to Annex 2 hereto. Each part of the acceptance procedure shall be governed by the principles set forth in this Article of the Contract and shall be independent of each other, except that the acceptance procedure of the System within the meaning of milestone no. 11 of Annex 2 hereto must be completed and agreed on before the acceptance procedure within the meaning of milestone no. 16 of Annex 2 hereto. Within the scope of the Technical Specification, the Supplier is obliged to describe what impacts the implementation of milestone no. 14 will have on the already functional part of the System operated by the Client.

- 6.5 Unless the contracting parties agree otherwise, underlying documents, scenarios and examples for acceptance tests shall be prepared by the Supplier in cooperation with the Client, taking into account acceptance procedures according to par. 6.4 hereof.
- 6.6 The Supplier shall call upon the Client to take part in the acceptance procedure no later than 5 working days before its commencement, unless the contracting parties agree otherwise. The Supplier shall prepare a written record of acceptance tests and shall primarily state in it whether the tests proved errors and defects, including a description of these errors and defects. The Client shall be provided with copies of all documents prepared in relation to performed acceptance tests.
- 6.7 If the Work or its part in the meaning of milestone no. 8 meets the acceptance criteria of acceptance tests in accordance with Annex 9 hereto and the Specification of Acceptance Tests, the Supplier undertakes to provide the Client to accept this performance no later than on the day following the completion of acceptance tests and the Client undertakes to accept it. The contracting parties undertake to prepare an acceptance protocol on the acceptance which shall state, among other things, that the Work was handed over free of defects and accepted by the Client without reservations, or whether it was accepted with reservations. The handover and acceptance shall take effect when the acceptance protocol is signed by the Client.
- 6.8 The Work as a whole shall be considered as definitively handed over and accepted after its last part so far not handed over is duly handed over and accepted. The contracting parties undertake to prepare a final acceptance protocol on the acceptance of the Work as a whole, together with the written protocol on termination of access of the Supplier to the Data Room of the VSB-TUO, as stated in par. 4.6 hereto. Provisions of par. 6.7 shall apply mutatis mutandis.
- 6.9 No later than on the day of signing of the final acceptance protocol, the Supplier shall be obliged to hand over the documentation of the Work specified in Annex 1 hereto and par. 3.1.1 hereof to the Client.
- 6.10 Part of the Work involving the training of the Client's staff by the Supplier shall be considered as accepted by the Client after the handover of the attendance sheet to the Client. The Supplier shall present the attendance sheet on the completion of training for all persons who were supposed to attend it, the attendance sheet shall include the list of training attendants who shall sign it after the completion of the relevant training.

7. OTHER OBLIGATIONS OF THE SUPPLIER

- 7.1 The Supplier further undertakes:
- 7.1.1 to provide performance hereunder in a due and timely manner, it shall be predominantly obliged to deliver the Work for the acceptance procedure, taking into
-

- account deadlines stipulated herein for the handover of the Work and in such form to meet the required acceptance criteria;
- 7.1.2 to provide performance or its parts hereunder with due care corresponding to conditions agreed herein;
 - 7.1.3 to notify the Client in time of all imminent defects in its performance or potential outages and provide the Client with all information necessary for the performance hereunder;
 - 7.1.4 to immediately inform the Client in writing of obstacles preventing performance of the subject matter hereof and other activities relating to the performance of the subject matter hereof;
 - 7.1.5 to notify the Client of potential risks of damage within 3 working days after the Suppliers learns such fact and take such measures in a due and timely manner according to its capabilities that shall entirely eliminate or mitigate the risk of damage;
 - 7.1.6 in providing performance hereunder, to proceed with due care and use “*best practice*” processes;
 - 7.1.7 to provide cooperation in order to obtain all necessary authorisations of the Client or, in accordance with the instruction of the Client, make any necessary announcements to the public administration authorities to provide for the fulfilment of the Contract’s purpose;
 - 7.1.8 to provide necessary cooperation to the Client or third parties designated by the Client; the cooperation of the Supplier shall be provided primarily in preparation and implementation of the System and other systems of the Client in the Data Room, operations of the System and other systems of the Client and optimisation of their operations in the Client’s infrastructure. This paragraph hereof additionally contains the specification of technical conditions and interface, organisational and project cooperation necessary for the implementation of other public contracts of the Client and for operations of supplied systems. The Supplier acknowledges that any dispute between the Supplier and another supplier of the Client shall be decided by the Client;
 - 7.1.9 to facilitate the presence of the Client’s representative and representative of another supplier of the Client in any stage of the Work installation in the Data Room;
 - 7.1.10 to inform the Client on the meeting of its obligations hereunder and other important facts that may have an impact on the exercise of rights and meeting of the obligations of the contracting parties;
 - 7.1.11 to ensure that all persons involved in the meeting of its obligations arising herefrom who shall be present in the premises or at workplaces of VSB-TUO, adhere to applicable legal regulations regarding occupational health and safety, sanitary regulations, fire prevention regulations at workplaces and environmental protection, and all internal regulations of VSB-TUO about which the Supplier was informed by VSB-TUO; for the avoidance of doubts, the Supplier shall be responsible for the qualification and physical fitness of all its employees and employees of its subcontractors, if any. At a request of the Client and everywhere where suggested by the generally applicable regulations, the Supplier shall be additionally obliged to present a qualification card for performed works and an authorisation for production
-

- of selected technical devices to the Client at the request of the Client or supervisory bodies;
- 7.1.12 to protect intellectual property rights of VSB-TUO and third parties;
- 7.1.13 when performing this Contract, to proceed in line with the Detailed Project Plan and project management rules in accordance with Annex 8 hereto;
- 7.1.14 to notify the Client of potential unsuitability of instructions by the Client in justified cases;
- 7.1.15 to observe safety principles and ensure safety observation in terms of processes, organisation and technical issues;
- 7.1.16 to behave in such manner to prevent danger to/damage to property or danger to health of persons and prevent these risks;
- 7.1.17 to behave in such manner to prevent any damage to the already installed facilities and equipment, or impact on the operations, functionality and safety or other operated systems of VSB-TUO;
- 7.1.18 to systematically and right from the start ensure such measures to mitigate safety risks, to prevent unauthorised access to or intrusion in the System or to data contained in the System, and ensure granting of authorised access only in the necessary scope and for a necessary period of time; and
- 7.1.19 to ensure that all parts of the Work meet the parameters in the Technical Specification in respect of climatic conditions in the Data Room and energy intensity of their operations, the Supplier explicitly undertakes to refund all actual and provable costs to VSB-TUO that are incurred as a result of a breach of this provision; this obligation shall apply over the entire period of use of the relevant part of the Work by VSB-TUO.
- 7.2 The Supplier acknowledges that the remuneration of all costs incurred in connection with the completion of the Work and provision of warranty in line with the conditions agreed in article 11 hereof shall be an integral part of the performance and confirms that these were included in the price of the Work hereunder. The Client reserves the right to inspect the status of the Work implementation and the entire course of performance hereunder and monitor its quality. For this purpose, inspection days shall be organised. Minutes shall be prepared from each inspection day agreed by both VSB-TUO and Supplier. The frequency of inspection days shall be determined pursuant to an agreement of both VSB-TUO and Supplier.
- 7.3 The Supplier additionally undertakes to maintain an insurance policy (or policies, further used in singular for simplification) valid and effective over the entire term hereof, including the warranty period for the Work the subject matter of which shall be insurance of liability for damage caused by the Supplier to a third party (mainly, but not exclusively to the Client) by breach of its contractual obligations, insurance of liability for detriment caused by operations of the Supplier and a product defect. In addition, the insurance must cover any bodily harm, death and potentially other natural rights of third parties. The limit of insurance benefits arising from the insurance policy must not be lower than CZK 150,000,000 and the level of co-insurance of the Supplier higher than 10% of the total amount of insurance benefits. At request, the Supplier shall be obliged to present the insurance policy or certificate of insurance meeting the above requirements to the Client by no later than 30 days after the delivery of a written call of the Client.
-

- 7.4 The Supplier additionally undertakes to ensure compliance with occupational health and safety in relation to the subject matter hereof in cases when required in accordance with currently applicable legal regulations in line with Act No. 309/2006 Coll. stipulating further requirements for health and safety at work in labour relations and concerning occupational health and safety protection in activities or services provided outside labour relations (Act on Further Requirements on Occupational Health and Safety), as amended (hereinafter the “**Act on Further Requirements on Occupational Health and Safety**”). After identification of any risks relating to the performance of the Supplier, the Supplier additionally undertakes to inform the Client of these risks immediately, no later than on the following working day.
- 7.5 For the avoidance of doubts, the performance by the Supplier shall additionally include the performance through a subcontractor in accordance with par. 3.9 hereof.
- 7.6 If the provision of performance hereunder results in production of waste, the Supplier undertakes to remove all waste produced in the performance of the agreed activities in line with Act No. 185/2001 Coll., on Waste and Amending Some Other Laws, as amended, and other relating regulations. The Supplier undertakes to remove the waste produced as part of performance hereunder by no later than within 5 working days from the production of waste; the waste shall not be left in the Data Room for more than 24 hours. The Supplier shall be obliged to proceed similarly in the liquidation of packaging.
- 7.7 When providing performance hereunder, the Supplier undertakes to proceed in a manner preventing damage to the Work and property of the Client or third parties at the place of performance or its vicinity. For this purpose, the Supplier primarily undertakes to maintain order during the completion of the Work at the place of its implementation and its vicinity at its own expense.

8. SPECIFIC OBLIGATIONS RELATING TO THE FUNDING SCHEME

- 8.1 The Supplier undertakes to provide the managing body of OP RDE, i.e. the Czech Republic – Ministry of Education, Youth and Sports (hereinafter as “**MO OP RDE**”), inspections by EuroHPC JU, European Anti-Fraud Office (OLAF), European Public Prosecutor’s Office (EPPO), European Court of Auditors (ECA), or other supervisory bodies of the public administration in accordance with applicable legal regulations, with access to the parts of tenders, contracts and relating documents that are subject to protection under special legal regulations in cases when it is necessary in respect of the requirements of legal regulations (e.g. Section 8 to 10 of Act No. 255/2012 Coll., on Inspection (Inspection Code), as amended, European Regulations 883/2013, 2185/96, 2017/1939, Article 287 of the Treaty on the Functioning of the European Union (TFEU) and Article 257 of EU Financial Regulation 2018/1046).
- 8.2 The Supplier undertakes to provide an authorisation to the MO OP RDE, EuroHPC JU and any other bodies stated above, in a manner determined in the prior paragraph, to inspect its subcontractors, if any.
- 8.3 The Supplier undertakes to provide all entities authorised to make inspection of projects, the funds of which are used to pay for the services and supplies hereunder, with the possibility to inspect documents relating to the performance under the contract, for a period not shorter than 3 years from the closing of MO OP RDE and EuroHPC JU under Article 90 of Council Regulation (EC) No 1083/2006, unless the Czech law stipulates a longer time period.
- 8.4 The Supplier acknowledges that the Client shall be obliged to meet the requirements for publicity in programmes of structural funds stipulated in Article 90 of Council Regulation (EC) No 1828/2006 and Publicity Policy in OP RDE, or publicity of EuroHPC in accordance with
-

Hosting Agreement No. 05/2019 that the Client shall deliver to the Supplier at request, in all relevant documents relating to the procurement procedure or policy.

- 8.5 The Supplier undertakes to provide all necessary cooperation in the performance of financial control under Act No. 320/2001 Coll. on Financial Control in Public Administration and on the Amendment to some Acts, as amended, and to provide any necessary cooperation in the performance of checks, audits, and investigations by bodies as per par. 8.1 hereof. The Supplier grants the Client's staff and outside personnel authorised by the Client the appropriate right of access to sites and premises where the Contract is performed and to all the information, including information in electronic format, needed to conduct such checks and audits. The Supplier ensures that the information is readily available at the moment of the check or audit and, if so requested, that information is handed over in an appropriate format. On the basis of the findings made during the audit, a provisional report is drawn up. The Client or its authorised representative will send it to the Supplier, who has 30 days following the date of receipt to submit observations. The Supplier will receive the final report within 60 days following the expiry of that deadline to submit observations. The Supplier undertakes to provide for the meeting of this obligation in its subcontractors.

9. PREVENTION OF MORAL DEPRECIATION OF TECHNOLOGIES

- 9.1 The contracting parties acknowledge that, given the term of this Contract and development in information and communication technologies, moral depreciation or another unsuitability of technologies that are to be supplied as part of the performance hereunder may occur during the Contract term (hereinafter the "**Moral Depreciation**").
- 9.2 Morally depreciated technology hereunder is the technology which is, at the time of the verification of its up-to-date character in accordance with par. 9.3 hereof, on the lower technological development level than as of the day of inclusion in the Technical Specification, and in respect of which it can be objectively assumed that if the Supplier or another reasonably acting entity were preparing the Technical Specification at the time of the inspection, it would not include this technology and would give preference to an adequately more advanced technology. Morally depreciated technology includes the cases when:
- 9.2.1 it is possible to purchase more advanced products (it means products having substantially better technical parameters) for the price not exceeding the bid price; or
- 9.2.2 production of the original technology was discontinued, or provision of associated services was discontinued, primarily support of the product by the manufacturer or discontinued production of spare parts; or
- 9.2.3 the original technology is not compatible with other technologies used for the implementation of the Work.
- 9.3 The Client shall be entitled to call upon the Supplier in writing to verify the up-to-date character of technologies that are to be supplied in accordance with the Technical Specification, and verify whether they are morally depreciated, any time during the Contract term, however no later than achievement of milestone "*5. Discussion and Acceptance of the Detailed Technical Specification*". The Supplier shall be obliged to communicate the results of the inspection to the Client in the period of 5 working days.
- 9.4 If there are morally depreciated technologies in the Technical Specification at the moment of the verification, the Supplier undertakes to replace the obsolete technology by another technology that shall meet the criteria of not being morally depreciated.
-

- 9.5 If the technology is replaced in accordance with par. 9.4 hereof, the Supplier shall be obliged to ensure all other adjustments to the Technical Specification in order to maintain proper functioning of the Work.
- 9.6 In relation to the replacements of morally depreciated technologies, the Supplier shall not be entitled to increase the price of any part of the Work. The replacement of morally depreciated technologies shall be part of the Work hereunder, specifically the System Implementation section.
- 9.7 If there is a conflict between the contracting parties regarding the issues relating to article 9 hereof, the contracting parties undertake to remove these conflicts by an expert appraisal of a court expert in the field of cybernetics, computing, who shall be agreed by the contracting parties for this purpose. It shall be an obligation of the Supplier to order and acquire the appraisal from the expert agreed by the Client and the Supplier.
- 9.8 The costs of the acquisition of an expert appraisal shall be borne by the party which was not successful in the dispute in accordance with par. 9.7 hereof according to the expert appraisal. However, when it is not possible to clearly determine which party was successful in the dispute, the Supplier and VSB-TUO shall bear the costs in equal parts.
- 9.9 If the morally depreciated technologies are replaced before the approval of the Detailed Technical Specification, the Supplier shall not be entitled to request an adjustment to the Time Schedule hereunder. Provisions of article 5 shall apply to this adjustment mutatis mutandis.

10. OWNERSHIP RIGHTS AND COPYRIGHT

- 10.1 The Client shall acquire the ownership rights to the Work as a whole and its individual parts which are movable assets on the date of the acceptance of the Work as a whole in accordance with par. 6.8 hereof by the Client, i.e. by achievement of milestone "18. Comprehensive Supply of the System", pursuant to a written acceptance protocol signed by authorised persons of the contracting parties. With respect to the gradual handover of the System to the Client, the Client acquires ownership rights to the part of the System within the meaning of milestone no. 8 on the day of signing the acceptance protocol related to this milestone. The danger of damage to the Work as well as all things handed over shall be transferred to the Client on the day when it acquires the ownership right to the Work, respectively on the day of acceptance of the part of the System within the meaning of milestone no. 8, in accordance with the prior part of this paragraph. Until the transfer of the ownership right to the Work, the Supplier shall provide for the protection of the Work or the part of the System and all its parts against damage. The Supplier shall select the type and scope of use of protective means in a manner sufficient to the protection of the Work before damage or destruction and continuously ensure the protection until the transfer of the danger of damage to the Work. The Client shall provide necessary cooperation to the Supplier to meet the obligation in accordance with par. 10.1 hereof. The Supplier undertakes to inform the Client of the suitability of selected protective means or potential damage to them during inspection days in accordance with par. 7.2 hereof. If there are comments of the Client to the suitability or sufficiency of used protective means, the Supplier undertakes to adapt the protective means to the requirements of the Client. In addition, the Client undertakes to the Supplier to provide all available records in the event of a damage to the Work before its completion which might suggest the cause of the damage (records of access to the Data Room and adjacent premises, CCTV footage of the Data Room and adjacent premises, records from the measurement and regulation system etc.) and provide the Supplier with cooperation during the Work completion in ensuring prevention of potential damage. For the avoidance
-

of doubts, all protective means shall be removed as of the day of the acceptance of the Work as a whole in accordance with par. 6.8 hereof, unless the Client stipulates otherwise.

- 10.2 The Supplier undertakes to ensure the provision of the right to the Client, in respect of each part of the System, to use these parts of the System in line with their intended purpose and ensure that the right to further modify the System is not affected by the claims of the Supplier.
- 10.3 The Supplier declares that it is the rightful owner of the Intellectual property rights to all Products and/or their components delivered under this contract, and that it is entitled to assign or licence those rights in accordance with the terms of this Contract. If those Intellectual property rights are the property of third parties, the Supplier guarantees that it has requested and obtained those third parties' written authorisation to grant to the Client the assignment or licence of their Intellectual property rights to the extent as provided under this Contract. The Supplier shall be solely responsible for taking the Supplier are to be performed, to ensure the opposability to third parties of the assignments or licences granted to the Client by the Supplier or by such third parties.
- 10.4 The Supplier guarantees that none of the Products, Documentation or other protected material delivered, whether or not developed in execution of this Contract, infringes any third party's Intellectual property rights.
- 10.5 Each Party shall inform the other Party of the existence or threat of any third party's action or claim alleging an infringement of its Intellectual property rights by the Client's use of any Products, Documentation or other protected material delivered under this Contract, provided such use is made in conformity with the terms of this Contract.
- 10.6 In the event of such a dispute or threat thereof, the Supplier undertakes to conduct all litigation, arbitration or negotiations for settlement, in its own name as well as in the Client's name, at its own and sole expense. The Client agrees to provide the Supplier with all information and assistance that may reasonably be required, at the Supplier's own and sole expense. However, the Client reserves the right to decide to conduct its own defence or to negotiate its own settlement, at its own discretion. The Supplier will be responsible for any payment arising out of any settlement or judgement following such a dispute or threat, except for the payment of a settlement made by the Client without the Supplier's written consent. Such consent may not be withheld without reasonable grounds. If the infringement of a third party's Intellectual property right on a Product and its Documentation is declared in a judgement, arbitration sentence or party settlement, or if such is likely to happen, the Supplier agrees to (1) either procure for the Client the right to continue using the Product and its Documentation, (2) either replace them with substantially equivalent non-infringing Products, or, if none of the foregoing is available, (3) grant to the Client a credit in the amount corresponding to the purchase price of the proportion of the Product which can no longer be used. The Supplier will not be responsible under the present guarantee for any third party claiming an infringement of its Intellectual property rights based on (1) the Client's use of Products in combination with equipment not delivered by the Supplier, if such combined use is the cause of the claimed infringement, or (2) the Client's use of any Product and Documentation delivered hereunder in a form other than the one delivered by the Supplier, if such change in form is the cause of the claimed infringement.
- 10.7 The guarantee against third party claims is due by the Supplier until five (5) years following the end of the Contract, or until five (5) years following the last use by the Client of the Product and its Documentation delivered by the Supplier, whichever period ends last.
-

11. WARRANTY AND LIABILITY FOR DEFECTS

- 11.1 The Supplier shall provide a warranty that the Work as a whole has functional properties and is capable to be used for purposes determined herein or hereunder as of the handover date. With respect to the fact that the Supplier will first deliver the part of the System within the meaning of milestone no. 8, the contracting parties declare that the warranty within the meaning of the previous sentence will also apply to the delivered part of the System within the meaning of milestone no. 8.
- 11.2 The Supplier shall provide a warranty for quality and warranty that the Work shall maintain its properties in accordance with par. 11.1 hereof over the warranty period. The warranty period shall start on the day of acceptance of the part of the System within the meaning of milestone no. 11 and shall last for a period of 5 years after the achievement of milestone “11. *Successful performing of Acceptance Tests of the part of the System*”.
- 11.3 Binding conditions of warranty provision over the warranty period in accordance with the prior paragraph shall be determined in Annex 6 hereto. For the avoidance of doubts, the warranty shall be part of the subject matter of performance hereunder and shall be provided without the claim of the Supplier for financial consideration.
- 11.4 Defects of the Work shall be removed in line with the following rules:
- 11.4.1 The Supplier shall start the removal of category A defect, i.e. defect that fully or significantly prevents the use of the Work, immediately after it is reported, and it shall remove the defect within 24 hours after it is reported or provides a substitute solution acceptable for VSB-TUO.
- 11.4.2 The Supplier shall start to remove the category B defect, i.e. defect that does not prevent the use of the Work, however significantly limits it, or that creates a risk of impossibility to use the Work, no later than within 2 hours after it is reported, and it shall remove the defect within 72 hours after it is reported or provides a substitute solution acceptable for VSB-TUO, unless the Supplier and VSB-TUO agree otherwise.
- 11.4.3 The Supplier shall start to remove the category C defect, i.e. defect that is neither category A nor category B defect, by no later than 24 hours after it is reported and the date of defect removal shall not exceed 7 working days after it is reported, unless the contracting parties agree otherwise in writing.
- 11.4.4 A substitute solution for the category A defect shall be the reported category B defect and a substitute solution for category B defect shall be the reported category C defect; an admissible solution shall be only the substitute solution that actually allows for a change in the category of the defects, unless the contracting parties agree otherwise through their authorised persons in justified cases.
- 11.4.5 If the removal of a defects of any category requires a shutdown of the Work, the Supplier shall make a written request for the shutdown of the Work. The defects shall be removed only if the request is approved by VSB-TUO and the shutdown takes place on the date determined by VSB-TUO. The period between the filing of the request for the shutdown and start of the Work shutdown shall not be included in the period for defect removal.
- 11.4.6 The Supplier shall remove category A and category B defects 24 hours a day and 7 days a week. Category C defects shall be removed only on working days from 8:00 am to 4:00 pm, unless the parties agree otherwise.
-

- 11.4.7 For the avoidance of doubts, significantly limiting are those defects that cause the worsening of the Work parameters by more than 20% as compared to parameters validated during the acceptance of the Work.
 - 11.4.8 For the avoidance of doubts, defects that create a risk of impossibility to use the Work are those defects that cause a loss of redundancy or availability of the Work in parts where the tender conditions explicitly determine this redundancy or availability.
 - 11.4.9 The Client shall be entitled to change the category of the defect during the defect duration pursuant to identification of new facts.
 - 11.4.10 The Client shall be entitled to report several concurrent defects as new defects of another category than the existing defects.
 - 11.4.11 If there are any doubts regarding the categorisation of the defects, it shall be resolved by VSB-TUO.
 - 11.4.12 The contracting parties state that the Client shall not be obliged to report defects to the Supplier within certain period from their identification, however it undertakes to report the defects identified without unnecessary delays. The reporting or failure to report the defect by the Client shall not be decisive for the obligation of the Supplier to remove the defects that occurred in the Work and shall not relieve the Supplier of this obligation.
 - 11.4.13 The period from reporting of category A or category B defect until its removal shall not be included in the warranty period.
 - 11.4.14 No more than 8 cases of category A defect may happen in 365 calendar days; defects in the Work caused by the Client intentionally or due to its gross negligence, or when the Client intentionally or due to its gross negligence makes it possible for a third party to cause a defect in the Work, shall not be included in the maximum number of defects.
 - 11.4.15 In 365 days, there may be no more than 16 cases of category B defects; the defects in the Work caused by the Client intentionally or due to its gross negligence, or when the Client intentionally or due to its gross negligence makes it possible for a third party to cause a defect in the Work, shall not be included in the maximum number of defects.
 - 11.4.16 In 365 calendar days, there may be no more than 264 hours of the Work shutdowns, of which there may be 120 hours of unplanned shutdowns. The shutdown of the Work is every, including started, hour, when the Work is not usable by the user due to planned or unplanned repairs, updates or similar activities on the Work. This should ensure the overall availability of the Work. A shutdown of the Work includes a defect in accordance with prior paragraphs of this article hereof, if the nature and the scope of the defect results in a malfunction of the Work.
 - 11.4.17 The first period of 365 calendar days for the determination of the total number of defects of individual categories in accordance with par. 11.4.14 and 11.4.15 and shutdowns in accordance with par. 11.4.16 of this article hereof shall start from milestone "11. Successful performing of Acceptance Tests of the part of the System". Other periods, each of 365 calendar days, immediately succeed the end of the previous period. The numbers of defects in the Work caused by the Client intentionally or due to its gross negligence, or when the Client intentionally or due to
-

its gross negligence makes it possible for a third party to cause a defect in the Work, shall not be included in the cumulative number of defects.

- 11.5 The Supplier declares that all the Work supplied hereunder shall be free of legal defects and undertakes to the Client to provide all defence and protection against the claims of third parties as well as indemnify the Client in full if a third party makes a copyright or another claim arising from the legal defects of the supplied performance. If a claim of a third party, regardless its eligibility, resulted in a temporary or permanent restraining order or restrictions to the use of the Work or its substantial part, the Supplier undertakes to provide for a substitute solution and reduce the impacts of such situation, the claims of the Client for indemnification shall not be affected.
- 11.6 For the avoidance of doubts, the obligation of the Supplier to remove the defect in the Work shall not be extinguished when it was caused by the Client by an incorrect use or another manner intentionally or due to a gross negligence or intentionally or due to gross negligence made it possible to a third party to cause the defect in the Work; this shall not affect the liability of the one who caused the defect by the breach of its legal obligations for the damage caused in this regard. For the avoidance of doubts, the obligation of Client to compensate for the damage shall be excluded in cases when the Client did not cause the defect in the Work intentionally or due to gross negligence.
- 11.7 For the avoidance of doubts, the liability for defects shall be governed by Section 2615 and Section 2099 and following provisions of the Civil Code, unless the Contract stipulates otherwise.

12. PRICE AND PAYMENT CONDITIONS

- 12.1 The price for the performance shall be specified in detail in Annex 4 hereto, and the price for the Work has been agreed between the contracting parties for the amount listed in the "Total" line in the table in Annex 4 hereto.
- 12.2 Prices herein shall be determined as final and shall not be exceeded. Prices can be changed only due to changes in legal regulations stipulating the amount of the value added tax, in the scope stipulated by a change in the legal regulations.
- 12.3 For the avoidance of doubts, the price in accordance with par. 12.1 hereof shall include all costs relating to the due provision of performance hereunder and completion of all relating activities and supplies defined herein and in the Technical Specification, including the fee for the provision of all licences and authorisations hereunder, accommodation and travel costs of the employees of the Supplier involved in the provision of the performance hereunder (transportation to the place of performance, accommodation at the place of performance, meal allowances). The price as stated in par. 12.1 hereof additionally includes the fee for meeting of all obligations of the Supplier agreed herein and meeting of all obligations of the Supplier stipulated by law or rules of OP RDE /EuroHPC. The price as stated in par. 12.1 hereof does not include only the performance that the Client explicitly undertook to provide as its cooperation.
- 12.4 The Work as a whole supplied in a due and timely manner shall be handed over at the moment of acceptance, i.e. signature of the final acceptance protocol by both contracting parties.
- 12.5 The price of the Work shall be paid in accordance with the payment schedule in Annex 4 hereto, after the due completion of the Work as determined by milestone "18. Comprehensive Supply of the System", respectively after completion of milestone no. 11, in line with the conditions listed in Annex 4 hereto. The condition for the invoicing of the price for the Work
-

shall be the acceptance of the Work by the Client, or acceptance of the delivered part of the System within the meaning of milestone no. 11, respectively. The Client shall not make any advance payments.

- 12.6 The price for the Work hereunder shall be paid by a bank transfer, pursuant to invoices (tax documents) that the Supplier shall be obliged to issue under the conditions listed herein. The first invoice shall be issued to EuroHPC JU and the second invoice shall be issued to VSB-TUO. The invoice to EuroHPC JU shall be issued for an amount of EUR 5,130,000 and the invoice to VSB-TUO shall be issued in Czech currency by translation of EUR 9,725,000 net of VAT using the exchange rate of the Czech National Bank as of the provision of performance, however for no more than CZK 251,800,000 net of VAT. The issued invoice shall be sent by the Supplier to EuroHPC JU immediately, however no later than 5 working days after signing the acceptance protocol of the Work within the meaning of the milestone no. 11 in accordance with par. 6.7 hereof by authorised representatives of the contracting parties. The issued invoice shall be sent by the Supplier to VSB-TUO immediately, however no later than 5 working days after signing the final acceptance protocol of the Work within the meaning of the milestone no. 18 in accordance with par. 6.7 hereof by authorised representatives of the contracting parties.
- 12.7 The invoice, and therefore the price for the Work, shall be due in 30 calendar days from the date of the invoice issuance by the Supplier in respect of VSB-TUO, and 60 calendar days from the invoice issuance by the Supplier in respect of EuroHPC JU.
- 12.8 All invoices shall contain the particulars of a tax document stipulated in Act No. 235/2004 Coll., on Value added Tax, as amended (hereinafter the "VAT Act"), shall contain the Contract number, and an appended signed acceptance protocol confirming the procedural acceptance of the Work. The invoices issued to VSB-TUO in accordance with par. 12.6 hereof shall be issued by the Supplier in accordance with Annex 4 hereto with the calculation of the tax base and calculation of the tax amount, it shall contain the name of the project OP RDE and reg. no. of project OP RDE in accordance with par. 2.1 hereof. The invoice issued to EuroHPC JU shall not contain the calculation of the VAT as the performance is exempt from VAT under Section 68 (8) of the VAT Act and Council Regulation (EU) 2018/1488 of 28 September 2018 establishing the European High-Performance Computing Joint Undertaking, Article 18. The entitlement for exemption from value added tax results from these regulations. If the invoice or its annexes do not contain the determined particulars or do not list correctly stated required information, the Client shall be authorised to return it to the Supplier before its due date with the list of missing particulars or incorrect information (it additionally applies to incorrect information listed in the Statement of performance). In such case, the maturity period shall be interrupted, and a new maturity period shall start on the date of delivery of the corrected invoice to the Client.
- 12.9 Monetary amounts shall be paid by a bank transfer to the account of the other contracting party listed in the invoice. The monetary amount shall be considered as paid on the day when it was deducted from the sender's account in favour to the recipient's account.
- 12.10 The Supplier acknowledges that VSB-TUO shall pay the amount in line with the received invoice only to the published bank accounts in the register of payers and identified entities. If the Supplier does not have a published bank account, VSB-TUO undertakes to pay only part of the price to the Supplier in the amount of the tax base and it shall pay the part of the price in the amount of the value added tax after the publication of the relevant account in the register of payers and identified entities without being delayed in the settlement of its payable.
- 12.11 If the Supplier becomes an unreliable payer under Act No. 235/2004 Coll., on Value Added Tax, as amended, VSB-TUO shall pay only part of the price in the amount of the tax base to the
-

Supplier. VSB-TUO shall pay the part of the price in the amount of the value added tax only after a written documentation by the Supplier on its payment to the relevant tax administrator, without VSB-TUO being delayed in the settlement of its payable.

13. AUTHORISED PERSONS

- 13.1 Each of the contracting parties shall designate an authorised person or a deputy of the authorised person. The authorised persons shall represent the contracting party in contractual, business and technical matters relating to the performance hereunder.
- 13.2 The authorised persons shall be authorised to perform, on behalf of the parties, all acts or legal acts during acceptance procedures hereunder and prepare amendments hereto after written approval to persons authorised to bind the parties (statutory bodies) or their representatives acting based on power of attorney.
- 13.3 The authorised persons shall not be authorised to activities that would directly result in a change herein or the subject matter hereof.
- 13.4 Names of the authorised persons shall be listed in Annex 7 hereto and their roles shall be stipulated by this Contract.
- 13.5 The contracting parties shall be entitled to change the authorised persons; however, they are obliged to notify the other contracting party of such change. A deputy of the authorised persons shall be authorised in writing with the listed scope of authorisation.

14. COOPERATION AND MUTUAL COMMUNICATION

- 14.1 The contracting parties undertake to cooperate and mutually communicate all information necessary for the due meeting of their obligations. The contracting parties shall be obliged to inform the other contracting party on all facts that are or may be important for due performance hereunder.
 - 14.2 The contracting parties shall be obliged to meet their obligations arising herefrom to prevent any delays in the meeting of individual deadlines and delays in adhering to the maturity of individual financial payables.
 - 14.3 All communication between the contracting parties shall be the task of the authorised persons in accordance with article 13 hereof, statutory bodies of the contracting parties or employees authorised in writing.
 - 14.4 All notifications between the contracting parties relating to this Contract, or that are to be made based on this Contract, shall be made in writing and delivered to the other party either in person or by a registered letter or another form of registered mail to the address listed on the title page hereof, unless stipulated or agreed otherwise between the contracting parties. If the communication in accordance with the previous sentence has no impact on the validity and effectiveness of the Contract, delivery by fax or email to numbers and addresses listed in Annex 7 hereto shall be allowed.
 - 14.5 If the Contract requests a delivery of a certain document in writing, it shall be delivered in a printed version and in an electronic (digital) form as MS Word 2003 or higher MS Word version, MS Excel 2003 or higher MS Excel version, or editable PDF document on an agreed medium.
 - 14.6 The contracting parties undertake that they shall inform the other contract party of changes in its postal address, fax number or email address by no later than within 3 working days.
-

- 14.7 The Supplier undertakes to, in the period of 10 days from the delivery of a justified written request of the Client for an exchange in the authorised person of the Supplier with whom the Client was not satisfied for objective reasons, replace this authorised person by another suitable person with adequate qualifications.

15. PROTECTION OF INFORMATION

- 15.1 The contracting parties are aware of the fact that as part of the meeting the obligations hereunder:

15.1.1 they may mutually knowingly or by an omission provide information that is treated as confidential under Section 1730 of the Civil Code (hereinafter the “**Confidential Information**”);

15.1.2 their employees and person on similar positions may gain access, through conscious activities of the other party or its omission, to the Confidential Information of the other party.

- 15.2 The contracting parties undertake that none of them shall provide access to a third party to the Confidential Information that it received from the other contracting party while providing performance hereunder. This shall not apply when its provision to a third party is necessary to meet the obligations arising herefrom.

- 15.3 Third parties as defined in par. 15.2 shall not include:

15.3.1 the employees of the contracting parties and persons on similar positions;

15.3.2 bodies of the contracting parties and their members;

15.3.3 in relation to the Confidential Information of the Client, subcontractors of the Supplier; and

15.3.4 in relation to the Confidential Information of the Supplier MB OP RDE, management body of EuroHPC JU and external suppliers to the Client;

provided that they take part in the performance hereunder or performance relating to performance hereunder, the Confidential Information shall be disclosed to them solely for this purposes and the disclosure of the Confidential Information shall be in the scope absolutely necessary to the fulfilment of its purpose and under identical conditions in relation to the obligation of confidentiality of these person, as stipulated for the contracting parties herein, primarily in respect of par. 15.6 hereof.

- 15.4 The contracting parties undertake to fully maintain confidentiality and protect the Confidential Information arising herefrom and from relevant legal regulations, primarily obligations arising from Act No. 110/2019 Coll., on Personal Data processing. The contracting parties undertake to inform all persons in this respect who shall take part in the performance hereunder on their parts of the above obligations of confidentiality and protection of the Confidential Information and undertake to ensure the meeting of these obligations by all persons taking part in the performance hereunder in a suitable manner.

- 15.5 If the information provided by the Client or third parties and necessary for the performance hereunder contains data subject to special protection under Act No. 110/2019 Coll., on Personal Data processing, and/or Regulation (EU) 2016/679 (General Data Protection Regulation) and Regulation (EU) 2018/1725, the Supplier undertakes to provide for the meeting of all reporting obligations that the above act requires, and obtain prescribed consents of personal data subjects forwarded for processing.
-

- 15.6 All Confidential Information shall remain in the sole ownership of the transferring party and the receiving party shall make identical efforts to maintain its confidentiality and its protection as if it were its own Confidential Information. Except for the scope that is necessary for the performance hereunder, both parties undertake to abstain from duplications of the Confidential Information of the other party, abstain for transferring it to a third party or its own employees and representatives, except for those who need to know it to be able to provide performance hereunder. Both parties additionally undertake not to use of the Confidential Information of the other party in another manner than for the purpose of performance hereunder.
- 15.7 Unless the contracting parties agree otherwise in writing, all information shall be considered as implicitly confidential that is or might be part of the trade secret, i.e. for example, but not limited to, descriptions or parts of descriptions of technological processes and formulas, technical formulas and technical know-how, information on operational methods, procedures and working methods, business or marketing plans, concepts and strategies or their parts, proposals, contracts, agreements or other arrangements with third parties, information on results of operations, on relationships with business partners, labour issues and all other information the disclosure of which by the receiving party could damage the transferring party.
- 15.8 If the Confidential Information is provided in writing or in the form of text files on electronic data carriers (media), the transferring part shall be obliged to notify the receiving party about the confidentiality of such material by marking it at least on the title page or from page of the medium. Any absence of such notification however shall not result in the extinguishment of the obligation to protect the information provided in such manner.
- 15.9 Regardless the above provisions, the following information shall not be the Confidential Information:
- 15.9.1 the information that became publicly known and its publication did not breach the obligations of the receiving contracting party or legal regulations;
 - 15.9.2 the information that the receiving party had provably legally available before the conclusion hereof, if such information was not subject to another contract for protection of information previously concluded between the contracting parties;
 - 15.9.3 the information that results from a procedure during which the receiving party concludes it independently and is able to document it by its records or the Confidential Information of a third party;
 - 15.9.4 after signing this Contract, provided to the receiving party by a third party which is not restricted in handling the information; and
 - 15.9.5 the information published in the profile of the Client, as the contracting authority of the Public Contract under Section 219 of the Act, if the Supplier did not prove that the Client is not obliged to publish using the procedure in accordance with par. 15.12 hereof.
- 15.10 The breach of the obligation of confidentiality by a contracting party shall be additionally the cases when this obligation is breached by any of persons listed in par. 15.3 hereof to which a contracting party provided the Confidential Information of the other contracting party.
- 15.11 The termination hereof for any reason shall not affect provisions of this article 15 hereof and they shall be effective after the termination of the Contract for 5 years after the end of the warranty period for the Work.
-

- 15.12 The Supplier shall state to the Client in writing within 5 working days from signing hereof what parts hereof, including its annexes, constitute the business secret of the Supplier under Section 504 of the Civil Code or another information protected in line with special legal regulations with reasons for such classification. The Supplier acknowledges that this procedure shall not apply in relation to the amount of the actually paid price for the performance hereunder and the list of subcontractors of the Supplier. Parts hereof constituting the business secret or other information protected in line with special regulations shall be secured against reading (redacted etc.) before the publication of the Contract. In the event of any dispute between the Client and the Supplier regarding the scope of these undisclosed parts of the Contract, the Client shall be authorised to publish only indisputable parts hereof and the remaining parts after this issue between the parties is resolved by a procedure for resolution of disputes.

16. SANCTIONS CLAUSE

- 16.1 If the Client is delayed in the payment of the invoiced issued to VSB-TUO, the Supplier shall be entitled to request a payment of the default interest in the amount stipulated by legal regulations.
- 16.2 The contracting parties have agreed that:
- 16.2.1 if the Supplier is delayed in the achievement of any of the milestones 3, 5, 7, 11 and 18 listed in the Time Schedule, the Client shall be entitled for a contractual fine of CZK 40,000 for each, including commenced, day of the delay;
 - 16.2.2 if the Supplier is delayed in the removal of warrant category A defect, the Client shall be entitled for the contractual fine of CZK 4,000 for each, including commenced, hour of the delay;
 - 16.2.3 if the Supplier is delayed in the removal of warrant category B defect, the Client shall be entitled for the contractual fine of CZK 700 for each, including commenced, hour of the delay;
 - 16.2.4 if the Supplier is delayed in the removal of warrant category C defect, the Client shall be entitled for the contractual fine of CZK 2,000 for each, including commenced, day of the delay;
 - 16.2.5 if the cumulative number of all warranty category A defects is exceeded in accordance with par. 11.4.14 hereof, the Client shall be entitled for the contractual fine of the amount determined by the sum of CZK 20,000 per each defect in addition to the determined maximum number;
 - 16.2.6 if the cumulative number of all warranty category B defects is exceeded in accordance with par. 11.4.15 hereof, the Client shall be entitled for the contractual fine of the amount determined by the sum of CZK 4,000 per each defect in addition to the determined maximum number;
 - 16.2.7 if the cumulative number of all Work shutdowns is exceeded in accordance with par. 11.4.16 hereof, and if it is not concurrently a defect of any of the above categories in accordance with par. 11.4 hereof, the Client shall be entitled for the contractual fine in the amount determined by the sum of CZK 4,000 for each, including commenced, hour of shutdown in addition to the maximum number;
 - 16.2.8 if the Supplier does present the insurance policy or insurance certificate in accordance with par. 7.3 hereof in time, or it is proved that it was not insured any time throughout the period stated in the first sentence of this paragraph in line with
-

- this Contract, the Client shall be entitled to request a contractual fine of CZK 250,000 for each started 14 successive days of the delay in the presentation of any insurance policy, if the insurance continued in the time of the delay, or for each started 14 successive days of the period when the Supplier was not insured in conflict with this Contract;
- 16.2.9 if the Supplier breaches the obligation to liquidate or remove waste and packaging in a due manner and before deadlines stipulated by par. 7.6 hereof, the Client shall be entitled for a contractual fine of CZK 10,000 for each, including commenced, day of the delay in the compliance with determined deadlines;
- 16.2.10 if the Supplier does not duly perform activities in accordance with conditions stipulated in par. 7.4 hereto, primarily if it breaches the obligation to inform the Client after identification of any risks relating to the performance immediately, however no later than on the following working day, the Client shall be entitled to a contractual fine of CZK 10,000 for each, including commenced, day of the delay in the compliance with such notification obligation or an individual case of breach of the obligation in line with the conditions of par. 7.4 hereof; and
- 16.2.11 if the Supplier breaches any of the obligations stipulated in paragraphs 7.1.8, 7.1.9, 7.1.10, 7.1.11, 7.1.13, 7.1.15, 7.1.17, 7.1.18 or 7.1.19 hereof, the Client shall be entitled for a contractual fine of CZK 10,000 per each individual case of breach of such obligation; if the breach of the obligation is a continuing defective condition, the individual case is each started day of defective condition.
- 16.3 If the Supplier breaches the obligation arising herefrom regarding the protection of business secret and the Confidential Information in accordance with article 15 hereof, the Client shall be entitled to request a payment of a contractual fine of CZK 500,000 for every breach of such obligation from the Supplier.
- 16.4 The contracting parties explicitly limit the total aggregate amount of contractual fines, the payment of which the Client is entitled to request, to the total of 10% of the price for the Work net of VAT in accordance with 12.1 hereof.

17. INDEMNIFICATION

- 17.1 VSB-TUO and the Supplier shall be liable for the damage caused under applicable legal regulations and hereunder, regardless of the fact whether the liability should have been or was the subject matter of an insurance policy in accordance with par. 7.3 hereto. Both aforementioned parties undertake to make maximum efforts to prevent damage and reduce caused damage. For the avoidance of doubts, the Client shall be entitled to request indemnification only of the actual damage under Section 2952 of the Civil Code.
- 17.2 None of the parties shall be liable for the damage caused as a result of factually incorrect or otherwise erroneous assignment received from the other party. If the VSB-TUO provided the Supplier with an erroneous assignment and the Supplier, given its obligation to provide performance with due care, could and should identify the incorrectness of such assignment, may invoke the provisions of the preceding sentence only if it notified the Client of the erroneous assignment and the Client insisted on the original assignment.
- 17.3 None of the contracting parties shall be obliged to compensate for the damage caused by the breach of its contractual obligations arising herefrom, the fulfilment of which was prevented due to some of the obstacles under Section 2913 (2) of the Civil Code which exclude enforcement to fulfil the indemnification obligation.
-

- 17.4 The contracting parties undertake to notify the other contracting party without undue delays of circumstances excluding the liability preventing the due performance hereunder from being carried out. The contracting parties undertake to make maximum efforts to avert and overcome the circumstance excluding the liability.
- 17.5 Any indemnification shall be paid in the currency valid in the territory of the Czech Republic, and the decisive rate for the translation to this currency shall be the exchange rate of the Czech National Bank as of the day when the damage occurred.
- 17.6 In derogation from Section 2050 of the Civil Code, the parties have agreed that agreement on any contractual fine shall not affect the right for indemnification for the damage caused by the breach of obligations to which the contractual fine relates, and the claim for the indemnification may be made independently from the contractual fine up to the amount limited by this Contract.
- 17.7 The sanction and the indemnification for the caused damage shall be payable within 30 calendar days from the date of delivery of the written call to payment together with the relevant invoice to the contracting party which is obliged to pay the relevant sanction or indemnification.
- 17.8 The contractual fine hereunder shall not be included in the indemnification arising from the breach of obligations stipulated herein, and these claims may be made independently in full and the contracting parties have agreed that the amount of the indemnification and contractual fines paid by one contracting party shall not exceed 110% of the price for the Work net of VAT as stated in par. 16.4 and par. 17.9 hereof. If any contractual fine is decreased by a court, the right of the Client for indemnification in the amount by which the damage exceeds the amount determined by the court as reasonable shall be maintained in the scope defined herein.
- 17.9 The contracting parties have agreed to restrict the right for indemnification that may arise to one contracting party during the performance hereunder, to the total amount equalling 100% of the price for the Work net of VAT in accordance with par. 12.1 hereof. However, this shall not affect Section 2898 of the Civil Code.
- 17.10 For the avoidance of doubt, the EuroHPC JU will not be responsible for any loss or damage caused to the Supplier during or as a consequence of performance of the Contract, unless the loss or damage was caused by wilful misconduct or gross negligence of the EuroHPC JU.

18. VALIDITY AND EFFECTIVENESS OF THE CONTRACT

- 18.1 The Contract shall become valid on the date when it is signed by both contracting parties. It shall take effect on the day when it is published in the register of contracts under relevant provisions of Act No. 340/2015 Coll., on Special Conditions for the Effectiveness of Certain Contracts, the Disclosure of these Contracts and the Register of Contracts (Act on the Register of Contracts), as amended (hereinafter the "Act on the Register of Contracts"). For this case, the contracting parties state that this contract contains business secret under Section 504 of the Civil Code, and the business secret between the contracting parties is the provision that was redacted by the contracting parties in accordance with par. 15.12 hereof. Given this provision hereof, the contracting parties have agreed that the Client shall provide for the publication of the Contract under the Act on the Register of Contracts and shall notify the Supplier about the publication of the Contract in the register of contracts.
- 18.2 Each contracting party shall be entitled to withdraw from the Contract only for the reasons stipulated herein.
-

- 18.3 The Client shall be entitled to withdraw from the Contract in the event of:
- 18.3.1 the delay of the Supplier in the handover of the Work for the period exceeding 30 days as compared to the deadline for the performance stipulated herein, if the Supplier does not provide for rectification in the additional reasonable period that the Client provides in a written call to meet the obligation, this period shall not be shorter than 15 days from the delivery of such call;
 - 18.3.2 there is breach of the Contract for which it shall be entitled to request a contractual fine exceeding CZK 500,000 hereunder; the provision of this paragraph however shall not affect contractual fines in accordance with par. 16.2.2 to 16.2.7 hereof and shall not apply when the fines are imposed;
 - 18.3.3 cumulative limits stipulated in par. 11.4.14 and/or par. 11.4.15 and/or par. 11.4.16 hereof are exceeded three times;
 - 18.3.4 the total amount of the contractual fines to the payment of which the Client would be entitled hereunder amounts to CZK 1.500,000 hereunder; and
 - 18.3.5 the obligation to have a concluded insurance policy and present it to the Client at its call under the conditions stipulated in par. 7.3 hereof is breached, if the Supplier does not provide for rectification in the additional period provided by the Client which must not be shorter than 5 working days.
- 18.4 The Supplier shall be entitled to withdraw from the Contract if the Client is delayed in the payment of any due amount hereunder for more than 60 days, if the Client does not provide for rectification in an additional reasonable period of time that the Supplier provides to it for meeting of obligations in a written call, and this time period shall not be shorter than 15 days from the delivery of such call.
- 18.5 Each of the contracting parties is entitled to withdraw from the Contract in writing, if:
- 18.5.1 the property of the other contracting party is declared bankrupt, the contracting party itself files a debtor's petition for initiation of insolvency proceedings or the insolvency petition is rejected as the property is not sufficient to pay for the costs of the insolvency proceedings under Act No. 182/2006 Coll., on Bankruptcy and Settlement (Insolvency Law), as amended; or
 - 18.5.2 The other party is placed in liquidation.
- 18.6 The withdrawal from the Contract shall take effect on the day of delivery of a written notification on withdrawal to the other party.
- 18.7 The Contracting parties have agreed that, in the event of a withdrawal from the Contract, they shall return all provided performance, unless the Contract stipulates otherwise, or the parties agree otherwise upon the withdrawal or in a period of 3 weeks after the withdrawal from the Contract. For the avoidance of doubts, any costs relating to a disassembly of part of the Work that is to be returned to the Supplier using the procedure stipulated in this paragraph shall be borne solely by the Supplier.
- 18.8 The termination of this Contract shall not affect provisions hereof relating to licences, warranties, claims arising from liability for defects, claims arising from liability for damage and claims from contractual fines, if they arose before the termination hereof, provisions on protection of information, and other provisions and claims the nature of which indicates that they are to continue after the termination hereof.
-

19. SETTLEMENT OF DISPUTES

- 19.1 The rights and obligations of the contracting parties not explicitly stipulated herein shall be governed by the Civil Code and relating applicable legal regulations.
- 19.2 The contracting parties undertake to make maximum efforts to eliminate mutual disputes arising hereunder or in relation hereto, including the disputes regarding the interpretation hereof and seek the resolution of disputes first amicably through negotiation of authorised persons or authorised representatives.
- 19.3 If the disputed issue is not resolved in accordance with par. 19.2 hereof within 60 days from the delivery of a call to amicable settlement of a dispute sent by any contracting party to the other contracting party, the dispute shall be resolved with definitive effect by a competent general court of the Czech Republic.

20. FINAL PROVISIONS

- 20.1 This Contract is a complete agreement of the contracting parties on the subject matter hereof. This Contract may be changed only by a written agreement of the contracting parties in the form of numbered amendments hereto, signed by persons authorised to act on behalf of the contracting parties. Any changes hereto shall be made in compliance with the Act.
- 20.2 The contracting parties have explicitly agreed that they exclude the use of Section 557 of the Civil Code.
- 20.3 If any of the provisions hereof proved to be invalid or unenforceable or became invalid or unenforceable after the conclusion hereof, this fact shall not result in invalidity or unenforceability of other provisions hereof, unless the mandatory provisions of legal regulations indicate otherwise. The contracting parties undertake to replace such invalid or unenforceable provision by a valid and enforceable provisions with the content as close as possible to the purpose of the invalid or unenforceable provision.
- 20.4 The contracting parties have agreed that they do not wish that any rights and obligations are derived from the existing or future practice established between the contracting parties or practices maintained in general or in the sector relating to the subject matter hereof in addition to explicit provisions hereof, unless the Contract explicitly stipulates otherwise. In addition to the above, the contracting parties confirm that they are not aware of any relevant commercial practices or existing established practice.
- 20.5 An offset against the receivables of the Supplier shall not be allowed. The contracting parties exclude, in relation to the receivables of the Client hereunder or in relation hereto, an application of Section 1987 (2) of the Civil Code and they agree that even an uncertain and/or indefinite receivable is eligible for an offset, however only to the moment of a possible filing of an action for performance hereunder. For the purposes of an offset of a receivable, the contracting parties undertake to first settle their disputes by a settlement, before they proceed to a separate unilateral offset of a receivable.
- 20.6 The Supplier shall not be entitled to offset any receivable arising herefrom against any receivable of a Client.
- 20.7 All rights and obligations arising herefrom shall be transferred, if the nature of these rights and obligations does not exclude it, to legal successors of the contracting parties.
- 20.8 The Supplier shall not be entitled to assign monetary claims in respect of the Client to a third party without the prior written consent of the Client.
-

20.9 The following annexes form an integral part hereof:

- Annex 1: Basic Technical Specification
- Annex 2: Time Schedule
- Annex 3: Cooperation of the Client
- Annex 4: Specification of the price and payment schedule
- Annex 5: Licence conditions for the selected SW
- Annex 6: Conditions of warranty provision
- Annex 7: Authorised persons
- Annex 8: Project management and similar change management rules
- Annex 9: Requirements for acceptance tests

20.10 This contract shall be prepared in 5 counterparts of which the Supplier shall receive 2 copies and the Client shall receive 3 copies.

The contracting parties declare that they have read the Contract, agree with its content and in witness thereof they add their signatures.

Client

Supplier

In Ostrava, on _____

In Prague, on _____

Vysoká škola báňská – Technická univerzita Ostrava

prof. RNDr. Václav Snášel, CSc., Rector

HEWLETT-PACKARD s.r.o.

_____, Statutory Representative

In Luxembourg, on _____

**European High-Performance Computing Joint
Undertaking**

_____, Executive Director

Annex 1 - 1

Basic Technical Specification

Basic Technical Specification – Design of Technical Solution

- SPEC_1 The supplier offers a comprehensive computing system solution designed for HPC, i.e. a complex of computing, storage, network and other systems and software solution, including implementation and integration into the contracting authority's data center (hereinafter referred to as the *solution* or *EURO_IT4I system*).
- SPEC_2 Offered EURO_IT4I system allows effective execution of a large number of simultaneous computing jobs of all phases of the job life cycle (i.e. preparation, pre-processing, computing, post-processing) of various types (especially parallel, but also serial, batch, and interactive) of large number of users. Offered system allows secure and reliable user data storage, fast access to data, and effective administration of systems, components, resources and services.
- SPEC_3 Offered EURO_IT4I system provides powerful computing resources, which are easily accessible by users and their jobs. The solution provides features, services, and functions required for effective system operation and administration by the contracting authority. The offered solution is balanced and parameters and structure of individual subsystems considers other subsystems.
- SPEC_4 EURO_IT4I system delivery includes all systems, equipment, accessories, licenses, documentation, design, implementation, and other work, training, etc. necessary to meet the contracting authority's requirements.
- SPEC_5 The solution considers disposition and limitations resulting from the environment and conditions of the contracting authority's data center.
- SPEC_6 The solution does not contain any limitations preventing or restricting the operation of EURO_IT4I system by the contracting authority in required, expected, or rational scope. EURO_IT4I system is intended for approximately 3,000 users.
- SPEC_7 The offered solution as a whole completely fulfills the contracting authority's requirements. Required functionality and features are realistically functional and usable in operation of the solution; required parameters are realistically achievable. Fulfilment of the contracting authority's requirements is not conditioned.
The fulfillment of the contracting authority's requirements is not based solely on functionality, properties, or parameters of individual components.
- SPEC_8 The functionality, features, and parameters of the solution are stated for the proposed/delivered configuration designed for everyday use. The functionality, properties, and parameters are not conditioned.
- SPEC_9 The solution meets all contracting authority's technical requirements simultaneously. All required features, functions, and parameters will be achieved using a single production setup of all the components in the solution. Fulfillment of the requirements is not conditioned on changing the setup or interconnection of the components.
-

- SPEC_10 To the maximum extent possible, the offered solution is autonomous, independent of external systems and services, and self-sufficient without the need for additional equipment, systems, or services.
- SPEC_11 Design and implementation of the solution ensures reliable, secure, powerful, and efficient operation of EURO_IT4I system in the contracting authority's data center.


EURO_IT4I System Components

- SPEC_12 EURO_IT4I system includes a *Compute cluster*. A Compute cluster consists of Compute nodes connected by a Compute network — a high-speed, low latency network. A Compute cluster is designed for computing users' jobs.
- SPEC_13 The Compute cluster contains four partitions:
- Universal compute partition
 - Accelerated compute partition
 - Data analytics compute partition
 - Cloud infrastructure compute partition
- SPEC_14 The *Universal compute partition* consists of standard servers without accelerators such as GPUs or FPGAs and is based on x86 CPU architecture to provide quick accessibility for the users and their existing codes.
- SPEC_15 The *Accelerated compute partition* delivers most of the compute power usable for HPC but also excellent performance in HPDA and AI workloads, especially in the learning phase of Deep Neural Networks. The Accelerated compute partition consists of heavily GPU-accelerated nodes.
- SPEC_16 The *Data analytics compute partition* will be oriented on supporting huge memory jobs by implementing a NUMA SMP system with large cache coherent memory.
- SPEC_17 The Cloud infrastructure compute partition supports both the research and operation of the Infrastructure/HPC as a Service. The Cloud infrastructure compute partition is intended for provision and operation of cloud technologies like OpenStack and Kubernetes.
- SPEC_18 EURO_IT4I system includes *Login nodes* – nodes for user access, job and data preparation, code compilation and debugging, result processing, and data transfer.
- SPEC_19 EURO_IT4I system includes *Visualization nodes* – nodes for remote visualization and user data modeling using hardware-accelerated graphics applications.
- SPEC_20 EURO_IT4I system includes data storages called *Storages*. Storages are designed for storing and sharing data. Storages are a comprehensive solution of storage devices, I/O servers (e.g. file servers), network, and necessary software. The storages provides required data services.
- SPEC_21 EURO_IT4I system includes *HOME storage*. *HOME storage* is a file storage designed for storing users' operating system settings and users' applications on the supercomputer.
- SPEC_22 EURO_IT4I includes *SCRATCH storage*. *SCRATCH storage* is a powerful file storage designed for short-term users' job data. Compute nodes use SCRATCH storage intensively.
-

- SPEC_23 EURO_IT4I system includes *INFRA storage*. *INFRA Storage* is a file storage designed for storing and sharing infrastructure data of the supercomputer. The storage is used for storing system images of nodes, logs, infrastructure service data, application software, scheduler data, etc.
- SPEC_24 EURO_IT4I system includes *PROJECT Network Gateways*. *PROJECT Network Gateways* are servers intended for connecting EURO_IT4I system with PROJECT data storage to allow access to PROJECT data storage. PROJECT data storage is the contracting authority's external storage designed for multi-year user project data and sharing of the data between IT4Innovations supercomputers.
- SPEC_25 EURO_IT4I system includes *Data management nodes*. *Data management nodes* are designed for data transfer between storages and are used primarily to transfer data from/to supercomputer via parallel transfers to relieve the load on Login nodes. The nodes also serve to directly access the storages from external locations (Internet) by the users.
- SPEC_26 EURO_IT4I system includes infrastructure and management nodes (hereinafter referred to as *Infrastructure nodes*). *Infrastructure nodes* are designed to provide infrastructure services to EURO_IT4I system (e.g. DHCP, DNS, LDAP, provisioning, license servers, schedulers, monitoring, logging, etc.) and for management and administration of the supercomputer, resources, and services.
- SPEC_27 EURO_IT4I system includes a data *Backup* solution.
- SPEC_28 EURO_IT4I system includes *Network infrastructure*, i.e. a network connection of the components and systems to achieve required functionality, provide access to individual services, and provide performance, availability, and security of the services.
- SPEC_29 The network infrastructure consists of cluster *Compute network*, *LAN*, *other networks*, as designed by the supplier, and integration into contracting authority's WAN.
Compute network connects Compute nodes of the compute cluster, Login nodes, Visualization nodes, Data management nodes, and other servers.
LAN provides communication between components inside the supercomputer.
Integration into WAN network of the contracting authority provides connection to contracting authority's WAN network and Internet connectivity for the supercomputer.
- SPEC_30 EURO_IT4I system includes the solution and infrastructure for installation and operation of the supercomputer in the contracting authority's data center (hereinafter referred to as *Infrastructure for operation in data center*). This includes racks and accessories required for placement of the supercomputer, the solution for power supply, and cooling of the supercomputer, systems, interfaces, and connections to the infrastructure of contracting authority's data center.
- SPEC_31 EURO_IT4I system includes all necessary software and licenses.
- SPEC_32 EURO_IT4I system delivery includes all required equipment and systems needed for required functionality, performance, and effective operation of the computing system (including those not explicitly stated in this document).
-

- SPEC_33 For implementation of EURO_IT4I system, the supplier will not use contracting authority's equipment, systems, infrastructure, and services unless explicitly stated in this document and only for purpose and scope stated in this document.

EURO_IT4I System Diagram



The EURO_IT4I system diagram is a schematic/simplified illustration of EURO_IT4I system; it does not demonstrate either a complete system or all or exact interconnections. Grey-colored parts are not a part of the EURO_IT4I system delivery and are supplied by the contracting authority.

Compute Cluster/Partitions

- SPEC_34 The following benchmark will be used for Compute cluster performance measurement:
- High Performance LINPACK <<http://www.netlib.org/benchmark/hpl/>>
- Applied benchmark implementation will fully comply with benchmark specifications. If an optimized version of the benchmark will be used, it will be supplied to the contracting authority no later than at the time of acceptance tests.
- SPEC_35 Rmax computing performance is determined by executing a High Performance LINPACK benchmark running parallel on all Compute nodes of the given partition (a single benchmark instance on the whole partition).
- SPEC_36 Computing performance is specified for the offered configuration designed for normal operation. The determined computing performance is not conditioned in any way (e.g. by using a specific processor mode not suitable for long-term system operation; or by declaring expected efficiency not guaranteed by the supplier). Declared computing performance will be proven by corresponding acceptance tests during implementation.
- SPEC_37 The sum of LINPACK Rmax computing performance of the individual compute partitions is more than 8600 TFLOPS. Detailed information is stated in mandatory annex of this offer <Technical parameters of the tender.xlsx> (Annex 1-2 of the Contract).
-

- SPEC_38 (I) Supplier specifies the theoretical computing performance R_{peak} of individual Compute node for all compute partitions. Detailed information is stated in mandatory annex of this offer <Technical parameters of the tender.xlsx> (Annex 1-2 of the Contract).
- SPEC_39 (I) Supplier states R_{max} computing performance according to SPEC_35 for each compute partition. Detailed information is stated in mandatory annex of this offer <Technical parameters of the tender.xlsx> (Annex 1-2 of the Contract).

Compute Nodes

- SPEC_40 Each Compute node meets the following requirements:
- x86-64 architecture
 - 64-bit operating system, Linux
 - Compute network connection
 - LAN connectivity
- SPEC_41 Compute node RAM is DDR4 with ECC.
- SPEC_42 RAM is evenly distributed (in terms of throughput, capacity, and access time) on processors and CPU cores of Compute node. RAM consists of memory modules of the same type (size, rank, etc.), and is distributed evenly, with the same configuration, on memory controllers and memory channels of Compute node. All memory channels of every processor in a node are utilized.
- SPEC_43 Compute nodes are intended exclusively for computing. The supplier doesn't use any Compute node to provide other functionality.
- SPEC_44 All Compute nodes in the same partition have identical hardware and software (e.g. firmware) configuration and operate using the same configuration (frequency, timing, feature setting).
- SPEC_45 (I) Supplier states below in the offer the hardware platform of the solution and detailed configuration of Compute nodes including processor model.
- SPEC_46 (I) Supplier provides in the offer below the name and version of Compute node operating system.

Universal Partition – Universal Nodes

- SPEC_47 Universal partition provides LINPACK R_{max} computing performance more than 1935 TFLOPS using CPUs only.
- SPEC_48 Each Universal node meets the following requirements:
- Two CPU sockets/processors per node
 - No accelerators
 - For CPUs only, node theoretical computing performance R_{peak} of 4000 GFLOPS, at a minimum, in double precision
 - at least 2GiB of RAM per CPU core
 - at least 256GiB of RAM per node
 - RAM DDR4, at least 3000MT/s
 - diskless; network boot from central image repository
-

- Compute network connection 100Gb/s, at a minimum
- LAN connectivity requirements

HPE Apollo 2000 Gen10 Plus System – general description

The density-optimized HPE Apollo 2000 Gen10 Plus System with a shared infrastructure chassis and support for up to four HPE ProLiant Gen10 Plus independent, hot-plug servers that allows user to maximize rack space. With rear-aisle serviceability access, server nodes can be serviced without impacting operation of other nodes in the same chassis to provide increased server up-time. It delivers the flexibility to tailor the system to the precise needs of your demanding high-performance computing (HPC) workload with the right compute, flexible I/O, and storage options. The system can be deployed with a single server, leaving room to scale as customer's needs grow. The system has been purpose-built to address the requirements of HPC applications such as manufacturing, oil and gas, life sciences, and financial services.

Density-optimized, scale-out compute for HPC and AI workloads

System Performance and Optimization

The HPE Apollo 2000 Gen10 Plus system delivers 4x density than 1U servers—four 1U servers in a 2U chassis. This allows you to maximize your data center footprint.

Expanded power capability with up to 3000W power supplies that deliver N+N power availability. High-speed infrastructure built on the latest industry standard components provides high performance to accelerate application.

Support for the full stack of 2nd generation AMD® EPYC™ 7000 series processors in the HPE ProLiant XL225n Gen10 Plus server.

Flexible infrastructure offers multiple storage options, eight memory channels and 3200 MT/s memory, PCIe Gen4, and support for processors over 240W for improved application performance.

Comprehensive Server Security and Management

Only HPE offers the industry standard server with firmware anchored into silicon with HPE iLO5 and Silicon Root of Trust. Tied into the Silicon Root of Trust is the AMD Secure Processor, a dedicated security processor embedded in the AMD EPYC system on a chip (SoC).

Secure recovery recovers firmware to last known good state if code is compromised to limit system disruption.

Firmware Runtime Validation provides a daily firmware check and alert of compromised code so that issues are contained rather than impacting the system.

HPE Performance Cluster Manager (HPCM) is a fully integrated system management software offering administrators all the functionalities they need to manage their clusters.

Shared infrastructure enables more efficient system management.

Comprehensive manageability including an extensive set of tools for node to rack management.

Flexible Scale-out Building Block

Allows you to scale too many servers so that customers have future proof expansion capacity.

Provides storage and I/O flexibility to help optimize for performance or economy to get to the right compute for the right workload.

Multiple storage chassis options from 0 to 24 SFF SAS/SATA drives and support for NVMe drives.

HPE Apollo 2000 Chassis


HPE Apollo 2000 chassis can reduce its data center footprint due to the density optimized design of the Apollo 2000 Gen10 Plus which reduce overall fabric and space cost for an overall lower TCO. This system can support up to four 1U servers in a 2U chassis which is twice the compute density when compared to a traditional 1U server. A 42U rack fits up to 20 Apollo 2000 series chassis accommodating up to 80 servers per rack.

The 2000 chassis offers multiple back plane options including 0 and 8 SFF SAS/SATA/NVMe good for a variety of HPC applications and 24 SFF SAS/SATA drives ideal for enterprise workloads. Each chassis comes with 5 dual-rotor fans as default and 2 additional fans can be added for extra cooling if needed. The chassis also has an integrated RCM option. And whether you require dynamic power throttling or true N+N redundancy, the Apollo 2000 Gen10 Plus has the right power options with support for up to two 1600W/2200W/3000W power supplies. The chassis can also be managed by the HPE Advanced Power Manager (HPE APM) - an optional rack level manager.

HPE ProLiant XL225n Gen10 Plus Server


The HPE ProLiant XL225n Gen10 Plus server deployed in the HPE Apollo 2000 Gen10 system delivers a flexible scale-out architecture with higher density than traditional rack mount systems and firmware-level server security with a flexible scale-out architecture for your high-performance computing (HPC) workloads. Each server node can be serviced individually without affecting the operation of other nodes sharing the same chassis to provide increased server uptime.

The HPE ProLiant XL225n Gen10 Plus server is a 2P/1U half width server tray built on the AMD EPYC 7002 processors with 8-64 cores, up to 3.4 GHz CPU speed and power ratings +240W. The tray has 8 memory channels per socket and 3200 MT/s memory, and supports 16 DIMMs per processor, two PCIe Gen4 slots, an OCP3.0 slot and multiple networking options including high speed fabric. Storage options include, 0 to 6 SFF SAS/SATA or 2 SFF SAS/SATA/NVME per node.



HPE Apollo 2000 Gen10 Plus standard configuration options:

Specifications	HPE Apollo 2000 Gen10 Plus System
Processor	Intel® Xeon® Scalable Processor Family AMD® EPYC® 7000 Series Processor Family (POI)
Processors / cores / speed	Up to two Intel® Xeon® Scalable Processors per server node. Up to 200+W Up to two AMD® EPYC® Scalable Processors per server node. Up to 200+W
Memory (type, channels, slots)	Supports up to 3200MT/s DDR4 SmartMemory 8 channels, 16 slots
Persistent Memory	TBD (Currently out of plan)
Drives Bays	Dependent on chassis selection XL220n - 6 SFF HDD/SSD or up to 6 NVMe SSD option XL290n - 12 SFF HDD/SSD or up to 12 NVMe SSD options XL225n - 6 SFF HDD/SSD or up to 2 NVMe SSD options XL220n, XL290n & XL225n - 2 internal 22110 M.2 optional kit w/ HW RAID NVMe
Network controller*	Intel: OCP 3.0 Form factor AMD: Intel 1x1Gb + OCP 3.0 Form Factor
GPU Support	(2) FH/FL (DW or SW) on XL290r – No GPU Support on AMD model up to 4 x single width accelerator support
Infrastructure management	HPE iLO 5 Management (standard), Intelligent Provisioning (standard), UEFI, iLO Advanced (optional), HPE OneView Advanced (optional)
Power Supply – Hot Plug	Two 1600W/2200/3000W HPE Apollo Platform Manager option for rack level management
Storage Controller	HPE Dynamic Smart Array S100i; optional HPE Smart Array PCIe card
Warranty (parts, labor, onsite support)	3/3/3



Accelerated Partition - Accelerated Nodes

SPEC_49 Accelerated partition provides LINPACK Rmax computing performance more than 4500 TFLOPS.

SPEC_50 Each Accelerated node meets the following requirements:

- Two CPU sockets/processors per node
- PCI-Express gen4, supported by all key components including CPUs, GPUs and Compute network adapters
- For CPUs only, node theoretical computing performance Rpeak of 2400 GFLOPS, at a minimum, in double precision
- 512GiB of RAM, DDR4, at least 3000MT/s
- LAN connectivity requirements

SPEC_51 Each Accelerated node contains eight GPU accelerators of the same type.

SPEC_52 Each offered NVIDIA A100 GPU accelerator of accelerated nodes meets the following requirements:

- CUDA technology support
- Theoretical computing performance Rpeak of 17 TFLOPS in double precision, at a minimum
- 40GiB HBM memory

SPEC_53 The GPU accelerators in accelerated node will be connected using high-speed connections/network. Connections between any two GPU accelerators in the node will provide theoretical throughput of at least 250GB/s in each direction simultaneously. The requirement will be met for simultaneous, parallel communication of all GPUs (GPU pairs) in the node. Each GPU accelerator will provide theoretical throughput of at least 250GB/s in each direction simultaneously.


NVIDIA 8-GPU A100 technical details:


8-GPU	
GPUs	8x NVIDIA A100
HPC and AI Compute FP64/TF32*/FP16*/INT8*	156TF/2.5PF*/5PF*/10POPS*
Memory	320 GB
NVIDIA NVLink	3rd generation
NVIDIA NVSwitch	2nd generation
NVIDIA NVSwitch GPU-to-GPU Bandwidth	600 GB/s
Total Aggregate Bandwidth	4.8 TB/s

SPEC_54 Each Accelerated node contains PCIe gen4 switches enabling the connection of GPU accelerators and Compute network adapters and CPUs.

SPEC_55 Each Accelerated node provides Compute network connection 800Gb/s (4x200Gb/s). Compute network adapters will be evenly distributed across the CPU sockets, connected directly to the PCIe switches where the GPU accelerators are connected.

SPEC_56 Accelerated node provides unified memory spanned over node’s GPUs memory and node’s RAM memory. The unified memory is accessible using CUDA API.





Data Analytics Partition – Data Analytics Node

- SPEC_57 The Data analytics partition provides the LINPACK Rmax computing performance of 40 TFLOPS.
- SPEC_58 EURO_IT4I system contains one Data analytics node.
- SPEC_59 The Data analytics node meets the following requirements:
- NUMA SMP architecture, cache coherent memory access
 - Thirty two CPU sockets/processors
 - 24TiB RAM, DDR4, at least 2900MT/s, coherent memory
 - Compute network connection 400Gb/s (4x100Gb/s or 2x200Gb/s)
 - LAN connectivity requirements
- SPEC_60 The Data analytics node is a single system allowing usage of all the resources by a single job, under a single instance of running OS. The Data analytics node will provide on-demand software partitioning using technologies like cgroups and cpusets to allow maximum utilization by multiple jobs/users with smaller requirements.
- SPEC_61 The Data analytics node provides fast memory access; the latency of remote NUMA node memory access is not greater than six times the latency of local NUMA node memory access, idle latencies are considered. The memory latency will be evaluated using Intel Memory Latency Checker or similar tool.
- SPEC_62 Software implementation of the memory coherency is not offered.

Data Analytics Partition will realized by **HPE SuperDome Flex** system.

HPE Superdome Flex – general description

The HPE Superdome Flex server is a compute breakthrough to power critical applications, accelerate data analytics and tackle HPC and AI workloads holistically. It delivers an unmatched combination of flexibility, performance and reliability for critical environments of any size. A unique modular architecture and unparalleled scale allow to start small and grow at customer site. Leveraging its in-memory design and groundbreaking performance, your business can process and analyze growing quantities of data at extraordinary speed. HPE Superdome Flex Server safeguards these vital workloads with superior reliability, availability, serviceability (RAS) and end-to-end security. The HPE Pointnext Services portfolio, broad partner ecosystem, and mission-critical expertise complement the capabilities and value of the platform to help ensure your move to HPE Superdome Flex is a success.

Power critical apps, accelerate data analytics, and tackle HPC and AI workloads holistically

Outpace evolving data demands with optimum flexibility

- Unique modular scale-up architecture utilizing 4-socket building block
- Designed for environments of all sizes – never outgrow, no over-provisioning
- HPE OneView management, Redfish APIs

Process and analyze growing data at extraordinary speed

- 4–32 sockets Intel® Xeon™ Scalable Processors as a single system
- 768GB–48TB shared memory for in-memory computing
- Unbounded I/O to support a wide variety of workloads

Safeguard mission-critical workloads

- Extreme Superdome RAS to achieve highest service levels
- HPE Serviceguard for Linux for high availability and disaster recovery
- HPE Pointnext services and expertise to lower risk

Designed with Memory-Driven Computing Principles

Outpace Evolving Data Demands with Optimum Flexibility

The unique modular design of HPE Superdome Flex allows to start small, at 4 sockets, and grow at your own pace—without sacrificing performance. Scaling seamlessly up to 32 sockets, in 4-socket building blocks, it can give you the compute power you need right now, no matter how much data you have or how fast it is growing. This flexible architecture helps you avoid overprovisioning and disruptive upgrades, with all the cost and complexity those carry.

The 4-socket modular building block offers a cost-efficient entry point for mission-critical workloads, and the ability to scale up at your own pace to 32 sockets with choice of either economical Gold or high-end Platinum Intel Xeon Scalable processors. HPE Superdome Flex starts at 768 GB and can expand to 48 TB of shared memory, offering plenty of room for growth. Customer can choose high performance DRAM only, or a combination of DRAM and HPE Persistent Memory, to meet the individual requirements of your workload. With support for HPE's unique x86 hard partitioning technology, HPE nPars, you can isolate workloads and/or consolidate multiple workloads onto a single managed complex.

Process and Analyze ever-growing Data at Extraordinary Speed

HPE Superdome Flex delivers groundbreaking performance at any scale, giving you the compute power you need for the most demanding workloads. With its innovative scalable architecture that connects individual 4-socket building blocks to one another in a point-to-point fashion, processors accessing data residing in memory in another block will always travel using a single “hop”—no matter if a processor at the top of the rack is accessing data from memory at the bottom. In addition, high bandwidth technology across all the sub-systems—processors, memory, interconnectivity, and I/O—means you can achieve the performance demanded to accelerate analytics and power critical apps.

Safeguard Mission-critical Workloads

HPE Superdome Flex is built for mission-critical availability and features proven Superdome reliability, availability and serviceability (RAS) capabilities not present in other standard x86 servers. These capabilities include a unique Firmware First approach to help you contain errors at the firmware level before any interruption can occur at the OS layer. In addition, the Analysis Engine offers best-in-class predictive fault handling and initiates self-repair without operator assistance. HPE Superdome Flex offers advanced and unique resiliency capabilities across every subsystem— memory, I/O, processor, and fabric—for prompt error detection and system self-healing. To help ensure business continuity, support for HPE Serviceguard for Linux® clustering technology is also provided.

Key Features

Unique modular architecture that scales seamlessly in 4-socket increments, from 4- to 32-sockets in a single-system. With up to 32 sockets/896 cores, customer has plenty of headroom to scale.

The system is built using 4-socket, 5U chassis that are cabled together. A fully configured Superdome Flex can support up to eight chassis, or 32-sockets.

Cost-efficient entry point for mission-critical workloads at 4 sockets, and the ability to scale up to 32 sockets with choice of either economical Gold or high-end Platinum Intel Xeon Scalable processors, 1st or 2nd generation.

48 DIMM slots of DDR4 memory per chassis, providing a large memory footprint for the most demanding in-memory workloads. Memory starts at 768GB and can grow up to 48TB of shared memory in a fully configured Superdome Flex.

Choice of high performance DRAM only, or a combination of DRAM and HPE Persistent Memory available in 128, 256 and 512GB kits featuring Intel® Optane™ DC Persistent Memory to meet individual workload requirements.

Designed for the future based on Memory-Driven Computing design principles to boost analytics performance.

Unbounded I/O with support for up to 128 PCIe standup cards.

16 half-height IO slots, or 8 full-height plus 4 half-height IO slots, per four-socket chassis.

Base IO includes 4-drive bays, two 1GbE and two 10GbE NIC ports, four USB ports.

Built-in DVD.

Choice of management including open Redfish software ecosystem, OpenStack and HPE OneView. Supports HPE’s unique x86 hard partitioning (HPE nPars) to isolate workloads and/or consolidate multiple workloads onto a single managed complex.

Analysis Engine for better diagnostics and mission-critical reliability.


HPE's “Firmware First” approach contains errors at the firmware level, including memory errors, before any interruption can occur at the Operating System layer.

HPE Serviceguard for Linux (SGLX) high availability and disaster recovery clustering solution protects your deployment from a multitude of infrastructure and application faults across physical or virtual environments over any distance.


HPE Superdome Flex at a glance

Processor	<ul style="list-style-type: none"> Up to 4 Intel® Xeon® Scalable Processors per chassis 1st and 2nd generation Gold or Platinum (81XX, 61XX, 82XX or 62XX) Up to 28 core per processor, max cores 896 (112 per chassis) Supports up to 8 chassis
Memory	<ul style="list-style-type: none"> 768 GB to 48 TB memory capacity per system 48 DIMM slots of DDR4 memory per 4-socket chassis Choice of high performance DRAM only, or with 2nd gen 62xx and 82xx processors, a combination of DRAM and Persistent Memory
Partitioning	<ul style="list-style-type: none"> Support for HPE nPars hard partitions in 48 different configurations
Storage	<ul style="list-style-type: none"> Four (4) internal drives bays for HDDs or SSDs to be used for boot or data. Supports both external SAS and Fibre Channel storage arrays when configured with SAS and Fibre Channel HBAs. For supported storage arrays see: www.hpe.com/storage/spock
Networking and I/O	<ul style="list-style-type: none"> Up to 16 half-height IO slots or 8 full-height + 4 half-height IO slots per four-socket chassis 1Gb, 10Gb, 10/25Gb*, 100Gb**
Operating Environment	<ul style="list-style-type: none"> Red Hat Enterprise Linux (RHEL) SUSE Linux Enterprise Server (SLES) Oracle Linux, Oracle VM VMware, Microsoft Windows Server <p style="text-align: right;">**Requires transceivers or DAC</p>
Form Factor	<ul style="list-style-type: none"> 5U chassis per 4 sockets



Superdome Flex Chassis Front View (without bezel)



Superdome Flex Chassis Rear View




HPE Superdome Flex: Chassis Internal Top View



Cloud Partition – Cloud Nodes

- SPEC_63 Cloud partition provides the LINPACK Rmax computing performance of 131 TFLOPS.
- SPEC_64 Each Cloud node meets the following requirements:
- Two CPU sockets/processors per node
 - 256GiB of RAM, DDR4, at least 3000MT/s
 - 2 local NVMe disks with capacity of at least 480GB in RAID1, NVMe PCIe 4x lanes
 - Compute network connection 100Gb/s, at a minimum
 - LAN connectivity requirements
- SPEC_65 The Cloud node has CPUs with the same architecture, core count, and frequency as the Universal nodes.
- SPEC_66 Each Cloud node has two dedicated (extra) Ethernet ports (10Gb/s each) connected to two Ethernet networks/switches (for redundancy) providing VLANs to allow better utilization of the VMs running on the platform and the possibility to create user defined networks. These networks/switches are considered as a part of the LAN infrastructure.




Compute Network

- SPEC_67 Compute network uses a RDMA technology with the interconnection throughput of min. 100Gb/s and interconnection latency of 10 microseconds, at a maximum. The equipment connection to Compute network and inside Compute network will meet the above-mentioned requirements for connections.
- SPEC_68 Compute network complies with the following characteristics:
The Compute network topology is fully non-blocking Fat Tree. Compute network consists of two layers of switches, one layer of switches (leaf switches) is intended for connecting endpoints, the second layer of switches (spine switches) interconnects leaf switches; the layers are mutually bipartite connected. Tree topology of Compute network is balanced.
- SPEC_69 Compute network uses effective routing optimized for the topology and the configuration of Compute network.
- SPEC_70 For the selected network topology/configuration (with regard to switch configuration and routing mechanism), connection of Accelerated nodes to Compute network provides optimal communication.
-

Fat Tree topology is offered, so corresponding ports of Accelerated nodes (e.g. port 1 of each Accelerated node) are connected to minimal number of leaf switches.

- SPEC_71 The solution provides an effective MPI communication of Compute nodes.
- SPEC_72 Offered compute network supports and provide an IP protocol.
- SPEC_73 A system-wide LINPACK benchmark running parallel on all Compute nodes of the Compute cluster (utilizing only CPUs) will be possible.
- SPEC_74 (I) Supplier states below the technology, topology, throughput, latency of Compute network, and the equipment configuration.
- SPEC_75 (I) Supplier states in the offer configuration and characteristics of proposed Compute network relevant for proposed network topology.



Login Nodes

- SPEC_76 EURO_IT4I system includes 4x Login nodes.
- SPEC_77 Each Login node meets the following requirements:
- Physical node
 - x86-64 architecture
 - Two CPU sockets/processors per node
 - For CPUs only, node theoretical computing performance Rpeak of 2400 GFLOPS, at a minimum, in double precision
 - 256GiB RAM DDR4, at least 3000MT/s, ECC
 - 2 local SSD disks with capacity of at least 480GB in RAID1
 - Hot-swap disks
 - Compute network connection 100Gb/s, at a minimum
 - LAN connectivity requirements
 - Redundant, hot-swap power supply units, redundant power supplies
- SPEC_78 The Login nodes will provide user access implementing an SSH2 protocol and services for SCP and SFTP file transfer.
- SPEC_79 The Login nodes uses the identical technology and instruction set of processors as Universal nodes.
- SPEC_80 The Login nodes have the same hardware configuration and operate in identical operating settings (frequency, latency, properties settings).
-


SPEC_81 The Login nodes are intended exclusively for providing user access and operations. The supplier will not use any Login node to provide other functionality.

SPEC_82 (I) Supplier states below the hardware platform of the solution and detailed configuration of Login nodes including processor model.

Login nodes will be realized by **HPE ProLiant DL385 Gen10 Plus** platform.

HPE ProLiant DL385 Gen10 Plus Server – general description

The HPE ProLiant DL385 Gen10 Plus server offers the 2nd generation AMD® EPYC™ 7000 Series processor delivering up to two time the performance of the prior generation. With up to 128 cores (per 2-socket configuration), 32 DIMMs for memory up to 3200 MT/s, the HPE ProLiant DL385 Gen10 Plus server delivers low cost virtual machines (VMs) with unprecedented security. Equipped with PCIe Gen4 capabilities, HPE ProLiant DL385 Gen10 Plus offers improved data transfer rates and higher networking speeds. Combined with a better balance of processor cores, memory and I/O makes the HPE ProLiant DL385 Gen10 Plus the ideal choice for virtualization, and memory-intensive and HPC workloads.

	
Processor	<ul style="list-style-type: none"> • Up to two (2) AMD EPYC™ 7002 series processors • Up to 64 cores with up to 225W • Up to 256 MB cache
Memory	<ul style="list-style-type: none"> • HPE DDR4 SmartMemory up to 3200MT/s • 32 DIMM slots, up to 4.0 TB (32 x 128 GB RDIMM) • Up to 8.0 TB memory (AMD EPYC processors can support up to 4TB of memory each)
HDD and Controller	<ul style="list-style-type: none"> • 28 SFF, 20 LFF HDD/SDD, 16 NVMe drives . • Up to 428.4 TB internal storage • Choice of HPE Smart Array Essential or Performance controllers for performance or additional features
Networking and I/O	<ul style="list-style-type: none"> • Choice of OCP plus standup adapters • Up to 8 PCIe Gen4 slots • Supports up to three double-wide GPUs
Management	<ul style="list-style-type: none"> • HPE iLO 5 • HPE OneView

Flexible Design

- Adaptable chassis, including the modular drive bays that can be configured with up to 28 SFF, up to 20 LFF or up to 16 NVMe drive options along with support for up to 3 double wide GPU options
- Faster application performance with increased memory speeds up to 3200 MT/s.
- Redesigned HPE Smart Array Essential and Performance RAID Controllers for performance or additional features including the ability to operate in both SAS and HBA mode.
- Choice of OCP 3.0 or PCIe standup adapters which offers a choice of networking bandwidth and fabric, making it scalable for changing business requirements.
- Supports a wide range of operating systems making it suitable for most environments.

- Higher bandwidth and improved data transfer rates with PCIe Gen4 capabilities.


Security Innovations

- Silicon Root of Trust is an immutable fingerprint in the iLO silicon. The silicon root of trust validates the lowest level firmware to BIOS and software to ensure a known good state.
- Tied into the silicon root of trust is the AMD Secure Processor, a dedicated security processor embedded in the AMD EPYC™ system on a chip (SoC). The security processor manages secure boot, memory encryption and secure virtualization.
- Run Time Firmware Validation validates the iLO and UEFI/BIOS firmware at runtime. Notification and automated recovery is executed on detection of compromised firmware.
- If system corruption has been detected Server System Restore will automatically alert iLO Amplifier Pack to initiate and manage the system recovery process, avoiding lasting damage to your business by quickly restoring firmware to the factory settings or the last known authenticated safe setting.

Automation

- HPE iLO 5 monitors servers for ongoing management, service alerting, reporting and remote management to resolve issues quickly and keep your business running from anywhere in the world.
- HPE OneView is an automation engine that transforms compute, storage, and networking into software-defined infrastructure to automate tasks and speed business process implementations.
- HPE InfoSight provides built-in AI predicts problems before they happen, proactively resolves issues, and continually learns as it analyzes data—making every system smarter and more reliable.
- HPE iLO RESTful API feature provides iLO RESTful API extensions to Redfish allowing you to take full advantage of the wide range of value-added API features, and integrate easily with leading orchestration tools.



HPE ProLiant DL385 G10 Plus front view:



HPE ProLiant DL385 G10 Plus back view:



HPE ProLiant DL385 G10 Plus top view:





SPEC_83 (I) Supplier provides the name and version of the Login node operating system in the Login node detailed configuration above.

Visualization Nodes

- SPEC_84 EURO_IT4I system includes 2x Visualization nodes.
- SPEC_85 Each Visualization node meets the following requirements:
- Physical node
 - x86-64 architecture
 - Two sockets/processors per node
 - For CPUs only, node theoretical computing performance Rpeak of 2400 GFLOPS, at a minimum, in double precision
 - 256GiB DDR4 RAM, at least 3000MT/s, ECC
 - one powerful GPU card with OpenGL support
 - 2 local SSD disks with capacity of at least 480GB in RAID1
 - Hot-swap disks
 - Compute network connection 100Gb/s, at a minimum
 - LAN connectivity requirements
 - Redundant, hot-swap power supply units, redundant power supplies
- SPEC_86 The Visualization nodes provides remote hardware-accelerated visualization with OpenGL support.
- SPEC_87 Visualization node GPU cards meets the following requirements:
- Theoretical computing performance Rpeak of 12 TFLOPS, at a minimum, in single precision
 - At a minimum 24GiB GPU memory
 - GPU memory is a continuous space with uniform addressing and uniform performance.
- SPEC_88 Visualization nodes will be fully integrated into EURO_IT4I system in the same scope as Compute nodes. Specifically, this includes user account integration, authentication, integration and accessing Visualization node services via Scheduler, and storage accessibility.
- SPEC_89 Visualization nodes uses the identical technology and instruction set of processors and as Universal nodes.
- SPEC_90 Visualization nodes have the same hardware configuration and will operate in identical operating settings (frequency, latency, properties settings).
- SPEC_91 Visualization nodes are intended exclusively for user data visualization. The supplier will not use any Visualization node to provide other functionality.
- SPEC_92 (I) Supplier states below the hardware platform of the solution and detailed configuration of Visualization nodes including processor model.

Visualization nodes will be realized by **HPE ProLiant DL385 Gen10 Plus** server.
The HPE ProLiant DL385 Gen10 Plus server is in detail described in Login nodes section.





SPEC_93 (I) Supplier provides the name and version of the Visualization node operating system in the Visualization node detailed configuration above.

SPEC_94 (I) Supplier states the name, version, and description of the Visualization node software in the Visualization node detailed configuration above.

Storages

SCRATCH Storage


SPEC_95 SCRATCH storage is a file storage intended exclusively for storing EURO_IT4I users' data.


SPEC_96 SCRATCH storage is independent of other storages.


SPEC_97 From the user's point of view, SCRATCH file storage behaves as a single, continuous space with a single namespace.


SPEC_98 SCRATCH file storage is accessible, mounted using a native protocol, on all Compute nodes, Login nodes, Visualization nodes, and Data Management nodes.


- SPEC_99 On clients, SCRATCH file storage is accessible on the */scratch* path.
- SPEC_100 SCRATCH file storage has a net available capacity on file-system level of 1PB (1×10^{15} byte).
- SPEC_101 SCRATCH file storage allows storing of 500 million files.
- SPEC_102 SCRATCH storage provides long-term sustainable sequential read performance for the 1MiB block size of more than requested 500GB/s (500×10^9 byte/s). The required performance is achievable from Compute nodes.
- SPEC_103 SCRATCH storage provides long-term sustainable sequential write performance for the 1MiB block size of 350GB/s (350×10^9 byte/s). The required performance is achievable from Compute nodes.
- SPEC_104 SCRATCH storage provides long-term sustainable random I/O performance for the block size of 4KiB and 80%/20% read/write mode of 5 million IOPs. The required performance is achievable from Compute nodes.
- SPEC_105 SCRATCH storage is flash based. All data (including metadata) are stored on NVMe disks. The disks are suitable for their designation and expected load.
- SPEC_106 (I) For SCRATCH storage, the supplier states in mandatory annex of this offer <Technical parameters of the tender.xlsx> (Annex 1-2 of the Contract) the net available capacity; long-term sustainable sequential performance for the 1MiB block size; and long-term sustainable random I/O performance with a block size of 4KiB and 80%/20% read/write mode.
- SPEC_107 (I) For SCRATCH storage, the supplier provides below the storage architecture, configuration of the equipment, type and number of disks, RAID level, number of disks in RAID, number of spare disks, and software description.




























HOME Storage

- SPEC_108 HOME storage is a file storage intended exclusively for storing EURO_IT4I users' data.
 - SPEC_109 HOME storage is independent of other storages.
 - SPEC_110 From the user's point of view, HOME file storage behaves as a single, continuous space with a single namespace.
 - SPEC_111 HOME file storage is accessible, mounted using a native protocol, on all Compute nodes, Login nodes, Visualization nodes, and Data Management nodes.
 - SPEC_112 On clients, HOME file storage is accessible on the */home* path.
 - SPEC_113 HOME file storage has a net available capacity on file-system level of 25TB (25x10¹²byte).
 - SPEC_114 HOME file storage allows storing of 500 million files.
 - SPEC_115 HOME storage provides long-term sustainable sequential performance for the 1MiB block size of 1.2GB/s (1.2x10⁹ byte/s). The required performance is achievable from Compute nodes.
 - SPEC_116 HOME storage provides long-term sustainable random I/O performance for the block size of 4KiB and 80%/20% read/write mode of 7,000 IOPs. The required performance is achievable from Compute nodes.
 - SPEC_117 HOME storage is flash based. All data (including metadata) are stored on SSD disks. The disks are suitable for their designation and expected load.
-

- SPEC_118 HOME storage does not use Lustre file system technology.
- SPEC_119 (I) For HOME storage, the supplier states in mandatory annex of this offer <Technical parameters of the tender.xlsx> (Annex 1-2 of the Contract) the net available capacity; long-term sustainable sequential performance for the 1MiB block size; and long-term sustainable random I/O performance with a block size of 4KiB and 80%/20% read/write mode.
- SPEC_120 (I) For HOME storage, the supplier provides below the storage architecture, configuration of the equipment, type and number of disks, RAID level, number of disks in RAID, number of spare disks and software description.





HPE Proliant DL325 Gen10 Plus server is described below in Data Management nodes section.

HPE MSA 2050 SAN Storage – general description


The flash-ready HPE MSA 2050 SAN Storage system is designed for affordable application acceleration that is ideal for small and remote office deployments. The HPE MSA 2050 SAN Storage system will give customer the combination of simplicity and flexibility to grow now and into the future, and advanced features you may not expect to find in an entry-priced

array. Start small and scale as needed with any combination of solid state disks (SSD), high-performance enterprise, or lower-cost midline SAS-based drives.



HPE MSA Storage has been the industry-leading entry storage Fibre Channel platform for the past eight years with nearly 500,000 storage systems sold worldwide. Now the HPE MSA 2050 SAN Storage system delivers up to two times higher performance than the previous generation at the same price, delivering in excess of 200,000 IOPS.

HPE MSA Storage Family



Features

The HPE MSA 2050 Storage system architecture maximizes performance.

- All HPE MSA 2050 Storage models includes:
 - Includes SFF or LFF array chassis, depending on the model.
 - Two MSA SAN or SAS controllers, depending on the model:
 - Four host ports per controller.
 - Each SAN controller supports 8 Gb FC, 16 Gb FC, 1 GbE iSCSI or 10 GbE iSCSI host connectivity.
 - Each SAS controller supports 12 Gb SAS host connectivity.
- 8 GB cache per controller.
- Battery-free cache backup with super-capacitors and compact flash

The HPE MSA 2050 SAN Controller allows to customer to create your own Combo Controller by mixing FC and iSCSI SFPs.

Storage Management Utility V3 (SMU)—The MSA management GUI brings a new modern look-and-feel to array management.

Thin Provisioning allows storage allocation of physical storage resources once they are consumed by an application. Thin Provisioning also allows over-provisioning of physical storage pool resources, allowing ease of growth for volumes without predicting storage capacity upfront.

All models feature a wide variety of drives: high-performance SSD drives, enterprise-class SAS, and SAS midline drives.

Will support a maximum of seven disk enclosures (either LFF and/or SFF). Add-on enclosures can either be HPE MSA 2050 LFF Disk Enclosures or HPE MSA 2050 SFF Disk Enclosures.

Can grow incrementally to a maximum of 96 LFF, 192 SFF drives, or a combination of SFF and LFF enclosures up to the maximum of eight total enclosures.

Virtual storage disks groups can be spanned across multiple enclosures.

Virtual Storage RAID levels supported: 1, 5, 6, and 10.

RAID 0 supported for read cache only. SSD read cache is a feature that extends the MSA controller cache. Read cache is most effective for workloads that are high in random reads. A maximum of two SSDs per pool can be added for read cache.

Maximum hard drive counts vary by RAID levels: maximum of two drives for RAID level 1; maximum of 16 drives for RAID levels 5, 6, and 10.

Multiple disk groups can be aggregated into a single storage pool.

Storage pools allow data on a given LUN to span across all drives in a pool. When capacity is added to a system, customer also gets a performance benefit of the additional spindles.

The maximum LUN size is 140 TB (128 TiB).

Snapshot enhancements for virtual storage, including performance improvements, hierarchical snapshots, and simplified resource management. Administrators can monitor and optionally control snapshot space usage.

Prioritize data by assigning appropriate affinity level (Performance, No Affinity, or Archive).

Customer can configure 512 TiB capacity per virtual pool by enabling large pool support.

Non-disruptive, online controller code upgrade (requires multi-pathing software).

Upgradable by design. Owners of an HPE MSA 2040, MSA 2042, or MSA 1040 array are able to do data-in-place upgrades to the new HPE MSA 2050 array. This unique ability protects earlier investments in drives and JBODs.

Product at a Glance

HPE MSA 2050 SAN Storage Description

HPE MSA 2050 SAN Storage

Array	
Access Type	Block
Form Factor	2U, SFF, or LFF
Number of Controllers per Array	2
Minimum/Maximum Host Ports	8
FC Host Connectivity	8 Gb or 16 Gb
iSCSI Host Connectivity	1 Gb or 10 Gb
SAS Host Connectivity	6 Gb or 12 Gb
Cache per Array:	
Maximum Read Cache per Array	8 TB
Data (read/write) Cache Plus System Memory per Array	16 GB
Pool Capacity (with large pool support)	562 TB (512 TiB)
RAID Levels Supported	RAID 0 (read cache only), 1, 5, 6, and 10
Enclosures	
Expansion Drive Enclosures	0 to 7 enclosures
LFF/SFF Array/Enclosure Nixing	Supported

HPE MSA 2050 SAN Storage

Maximum Number of Drives per Array Enclosure	24 SFF / 12 LFF
Maximum Number of Drives per Drive Enclosure	24 SFF / 12 LFF
Drive Enclosure Interface Type	6 Gb SAS
Drives	
Maximum Total HDDs per Array	192 SFF / 96 LFF
Maximum Total SSDs per Array	192 SFF / 96 LFF
Maximum Raw Capacity per Array Enclosure	76.8 TB SFF / 168 TB LFF
Maximum Raw Capacity per Drive Enclosure	76.8 TB SFF / 168 TB LFF
Maximum Raw Capacity per Array	614.4 TB SFF / 1344 TB LFF

Drive Capacities

SFF SSDs	800 GB, 960 GB, 1.6 TB, 1.92 TB, 3.2 TB
LFF SSDs	800 GB, 960 GB, 1.92 TB
SFF HDDs	15K: 600 GB, 900 GB 10K: 600 GB, 1.2 TB, 1.8 TB, 2.4 TB
LFF HDDs	7.2K: 4 TB, 6 TB, 8 TB, 10 TB, 12 TB, 14 TB
SEDs	SSDs: 800 GB, 1.6 TB SFF 10K HDD: 1.2 TB LFF 7.2K HDD: 4 TB

Software Features

Thin Technologies	Thin Provisioning, Space Reclamation, Thin Rebuild
Tiering	Performance Tier, Standard Tier, Archive Tier
Replication	Snapshots (512), Volume Copy, Remote Snaps
Quality of Service	Virtual Tier Affinity

Additional Features

Maximum Number of Volumes	512
Maximum Number of Snapshots	512
Maximum Number of Hosts	512
Maximum Number of Initiators	1024
Customer Self-Installable, Self-Repairable, and Self-Upgradeable	Yes
File Services	Yes, with HPE Storage File Controller

INFRA Storage

SPEC_121 INFRA storage is a file storage intended for storing supercomputer infrastructure services data and contracting authority’s data necessary for providing HPC services to users.


SPEC_122 INFRA storage is independent of other storages.


SPEC_123 The supplier assumes the following use of INFRA storage:

- Node installation and boot images;
- Scheduler data and logs;
- Application software provided by contracting authority.

Storage parameters were determined based on Technical requirements specification provided by contracting authority.

- SPEC_124 The supplier will not use INFRA storage for purposes other than stated in SPEC_123.
- SPEC_125 INFRA file storage has a net available capacity on file-system level of 25TB (25x10¹²byte).
- SPEC_126 INFRA file storage allows storing of 500 million files.
- SPEC_127 INFRA storage provides long-term sustainable sequential performance for the 1MiB block size of 1.2GB/s (1.2x10⁹ byte/s). The required performance is achievable from Compute nodes and from Infrastructure nodes independently.
- SPEC_128 INFRA storage provides long-term sustainable random I/O performance for the block size of 4KiB and 80%/20% read/write mode of 7,000 IOPs. The required performance is achievable from Compute nodes and from Infrastructure nodes.
- SPEC_129 INFRA storage is flash-based. All data (including metadata) are stored on SSD disks. The disks are suitable for their designation and expected load.
- SPEC_130 INFRA file storage allows partitioning of data capacity into logical partitions/file systems of the required size and making these partitions/file systems available only to selected nodes.
- SPEC_131 INFRA storage solution is available to all EURO_IT4I system's nodes.
- SPEC_132 INFRA storage makes available:
- Node installation images – on Infrastructure nodes providing remote installation service or remote network boot service;
 - Scheduler data and logs – on Infrastructure nodes running Scheduler;
 - Contracting authority's application software – on Compute, Login, and Visualization nodes, on Data management nodes, and on selected Infrastructure nodes.
- For each of the above-mentioned nodes, a data path to INFRA storage exists with throughput of 10Gb/s, at a minimum.
- SPEC_133 INFRA storage does not use Lustre file system technology.
- SPEC_134 (I) For INFRA storage, the supplier states in mandatory annex of this offer <Technical parameters of the tender.xlsx> (Annex 1-2 of the Contract) the net available capacity; long-term sustainable sequential performance for the 1MiB block size; and long-term sustainable random I/O performance with a block size of 4KiB and 80%/20% read/write mode.
- SPEC_135 (I) For INFRA storage, the supplier provides below the storage architecture, configuration of the equipment, type and number of disks, RAID level, number of disks in RAID, number of spare disks, and software description.
-





HPE Proliant DL325 Gen10 Plus server is described below in Data Management nodes section.
HPE MSA 2050 SAN Storage is described in HOME storage section above.

File Storage

- SPEC_136 The file storage provides shared network file system services.
- SPEC_137 File storage clients provides standard file system functionalities (in the POSIX way).
- SPEC_138 On the client side, the file storage is transparently integrated into the operating system, allows standard file operations and implement the usual semantics of native file systems, supports a native OS file system interface (API) and integrates operating system users as file system users.
- SPEC_139 The file storage meets the following requirements:
- Unicode support in file names;
 - Long file names support;
-

- Access control, standard Unix permissions (read, write, execute; user, group, others) and extended ACL;
- User quotas, limits for disk capacity and number of files configurable individually for each user;
- Group quotas, limits for disk capacity and number of files configurable individually for each group;
- Reporting of used capacity and number of files for individual users and groups;
- Support for files larger than 1TB;
- Symbolic links support;
- File locking support.

User and group quotas are not expected to be enabled on INFRA storage file systems.

SPEC_140 Single, continuous space with a single namespace means that the user uses a single namespace to access the files in the storage, and all the capacity and properties of the file storage are available within that single namespace.

SPEC_141 Storage services does not negatively affect each other. The requested sequential and random I/O performance is achievable under simultaneous parallel load on all storages.

Requirements – Availability, Redundancy

SPEC_142 The storage solution provides high availability. The storage does not contain any component whose outage could cause a failure of the storage services (no “single point of failure” exists). Storage solution components, specifically disks, power supply units, disk array controllers, switches, and servers are redundant and hot-swappable without storage services outage. The offered disk arrays only consists of dual-port disks and there are multiple independent data paths for all disks.

SPEC_143 Failure or outage of any single storage server or network device does not cause a failure of storage services.
During a failure or outage of a storage server or network device, the storage performance can be lower than required.

SPEC_144 The storage solution is designed for long-term heavy load.

SPEC_145 (I) In the offer above, the supplier provides the details about the solution for high availability and redundancy of the supplied systems and equipment.

Disk Redundancy

SPEC_146 The storage provides sufficient redundancy so that a failure of any two disks will not cause data loss.

SPEC_147 The storage provides recovery from disk failure, i.e. re-ensuring the required data redundancy (e.g. RAID group reconstruction using hot-spare disks). Recovery from disk failure occurs automatically, without operator’s intervention.

SPEC_148 Recovery from disk failure, i.e. re-ensuring the required data redundancy, will be completed within 48 hours of disk failure. During recovery, the storage performance can be temporarily lower than required.

- SPEC_149 The storage configuration ensures required data redundancy after a failure of any two disks of the storage, without operator's intervention.
- SPEC_150 Each disk array or similar equipment of the storage solution provides a reserve capacity or spare disks in the number or capacity of
- $\max\left(\frac{1}{24} * \text{number_of_disks}; 2\right)$ disks, at a minimum, for each type of disk in the disk array.
- The $\max(a; b)$ function returns the larger value from a, b. The result is rounded up to the nearest integer.
- The *number_of_disks* value is the number of disks of the same type in the disk array. The value includes all disks - the disks with data and parity data, as well as spare disks and reserve capacity.
- SPEC_151 (I) On the other respective places of this offer, the supplier provides the details about the storage redundancy, number and type of disks, RAID level, number of disks in RAID, number of spare disks, and storage software solution description.

Storage Capacity

- SPEC_152 The data storage capacity (size) is stated as the net usable capacity, i.e. the storage capacity actually usable by the user at the highest service level provided.
The file storage capacity is the capacity of the file system provided by the storage.
- SPEC_153 The storage net usable capacity is stated for offered configuration designed for normal operation.
- SPEC_154 All required/offered capacities (including the metadata capacity required for storing the requested number of files) are achievable simultaneously.
- SPEC_155 Determination of the net usable capacity does not take into account the features of the system or its component, which could provide larger storing capacity under conditions that cannot be ensured (compression, deduplication, etc.).
- SPEC_156 Determination of the net usable capacity does not take into account the features of the system or its component, which could allow allocating more space than physically possible or practically feasible without further action (oversubscription).
- SPEC_157 Storage capacity is specified using prefixes of decimal multiples.
Gigabyte (GB) 10^9 byte
Terabyte (TB) 10^{12} byte
Petabyte (PB) 10^{15} byte
- SPEC_158 The net available storage capacity is determined on a suitable storage client system by a suitable system tool demonstrating the size/storage capacity.
- SPEC_159 The tools used to determine the capacity will provide reliable information and work with a known data block size or a known and accurate unit.
In Linux OS, the file system capacity in bytes can be displayed using the `df -B1` command, value "Available".
-

Storage Performance

- SPEC_160 The storage performance (sequential performance, I/O performance, metadata performance) is stated as the long-term performance realistically sustainable by the user from storage clients at the highest service level provided.
- SPEC_161 The file storage performance is the long-term, realistically sustainable performance of operations performed on file storage systems from file storage clients.
- SPEC_162 The storage performance is stated for the offered configuration designed for normal operation.
- SPEC_163 The storage performance is determined for highly occupied storage.
- SPEC_164 The storage performance is determined for non-privileged user operations.
- SPEC_165 The performance is not determined based on the assumption of specific, favorable conditions or specific, advantageous measurement modes (e.g. cache operation), unless such conditions or modes are explicitly required.
- SPEC_166 The storage performance is specified using prefixes of decimal multiples.

Performance Measurement

- SPEC_167 The supplier will demonstrate compliance with the storage performance requirements by running performance tests (benchmarks) within the acceptance tests (performance measurements).
- SPEC_168 The performance tests will be run on test servers configured as file storage clients.
- SPEC_169 The test server disk caches will be cleared before each storage performance measurement.
- SPEC_170 The performance measurement will be carried out by procedures and under conditions that correspond to normal operation – normal provision of storage services. No action will be taken prior to and during measurement of the storage performance that would affect the measurement result.
- SPEC_171 (I) In the delivery, the supplier will state the proposal of the performance measurement implementation.

Using FIO Benchmark for Performance Measurement

- SPEC_172 FIO Benchmark will be used for Performance Measurement fully in compliance with SPEC_172 of “Technical requirements specification for EURO_IT4I system” provided by contracting authority.

Sequential Performance Measurement

- SPEC_173 Sequential Performance Measurement will be done fully in compliance with SPEC_173 of “Technical requirements specification for EURO_IT4I system” provided by contracting authority.
-

Random I/O Performance Measurement

SPEC_174 Random I/O Performance Measurement will be done fully in compliance with SPEC_174 of “Technical requirements specification for EURO_IT4I system” provided by contracting authority.

Metadata Performance Measurement

SPEC_175 In order to verify the stability of the file storage metadata operations and to determine the file storage metadata performance, the long-term sustainable performance of the file storage metadata operations will be measured as acceptance tests.

SPEC_176 Metadata Performance Measurement will be done fully in compliance with SPEC_176 of “Technical requirements specification for EURO_IT4I system” provided by contracting authority.

Data Management Nodes

SPEC_177 EURO_IT4I system includes two servers operating exclusively as Data management nodes – nodes designed for data transfer between supercomputer’s storages and external storages.

SPEC_178 Each Data management node meets the following requirements:
Physical node
x86-64 architecture
At least one processor, at least 16 CPU cores total;
At least 128GiB RAM, operated in DDR4 mode, at least 2666MT/s with ECC;
The RAM consists of memory modules of the same parameters, evenly distributed across all server memory;
Theoretical throughput of processor(s) to RAM at least 160GB/s (in the offered server configuration);
At least 2 local SSD disks with the capacity of 240GB, at a minimum, in RAID1;
Hot-swap disks;
Redundant, hot-swap power supply units, redundant power supplies.
Compute network connection 100Gb/s, at a minimum
LAN connection requirements

SPEC_179 On Data management nodes, HOME and SCRATCH file storages will be accessible, mounted via native protocol.

SPEC_180 The Data management nodes will provide protocols for SFTP and SCP file transfer and ensure transfer from/to HOME and SCRATCH. On Data management nodes, the contracting authority will install and operate GridFTP and other services.

SPEC_181 Data management nodes will be integrated into EURO_IT4I system. Specifically, this includes user account integration, authentication, and integrating and accessing Data management nodes services via Scheduler. Data management nodes will be clients of the Scheduler.

SPEC_182 (I) Supplier provides the description and detailed configuration of Data management nodes including processor model.

Data Management Nodes will be realized by **HPE ProLiant DL325 Gen10 Plus** server.

HPE ProLiant DL325 Gen10 Plus Server – general description

Building on HPE ProLiant as the intelligent foundation for hybrid cloud, the HPE ProLiant DL325 Gen10 Plus server offers the 2nd generation AMD® EPYC™ 7002 Series processor delivering up to two times the performance of the prior generation.

The HPE ProLiant DL325 Gen10 Plus delivers increased value to through intelligent automation, security, and optimization. With more cores, increased memory bandwidth, enhanced storage, and PCIe Gen4 capabilities, the HPE ProLiant DL325 Gen10 Plus offers two-socket performance in a one-socket 1U rack profile. The HPE ProLiant DL325 Gen10 Plus, with the AMD EPYC single-socket architecture, enables businesses to acquire an enterprise-class processor, memory, I/O performance, and security without having to purchase a dual processor.

	[REDACTED]
Processor	<ul style="list-style-type: none"> • One AMD EPYC™ 7002 series processor • Up to 64 cores with up to 225W • Up to 256 MB cache
Memory	<ul style="list-style-type: none"> • HPE DDR4 SmartMemory up to 3200MT/s • 16 DIMM slots, up to 2.0 TB (16 x 128 GB RDIMM) • Up to 4.0 TB memory (AMD EPYC processors can support up to 4TB of memory each)
HDD and Controller	<ul style="list-style-type: none"> • Up to 12LFF/24SFF Smart Drives HDD/SSD, 24 SFF NVMe • Up to 367.2 TB internal storage • Choice of HPE Smart Array Essential or Performance controllers for performance or additional features
Networking and I/O	<ul style="list-style-type: none"> • Choice of OCP plus standup adapters • Up to 3 PCIe 4.0 slots • Supports 2 single-width active GPUs
Management	<ul style="list-style-type: none"> • HPE iLO 5 • HPE OneView (optional)

Delivering 2P Performance at 1P Economics

Supports industry standard technology utilizing the 2nd generation AMD EPYC 7002 series processor family with up to 64 cores, PCIe Gen4 capabilities, and HPE SmartMemory up to 3200 MT/s DDR4 memory.

Dense, modular chassis that can accommodate up to 24 SFF SAS/SATA HDD/SSD, up to 12 LFF SAS/SATA HDD/SSD, or up to 24 NVMe drive options as well as up to 3 PCIe 4.0 slots.

Redesigned HPE Smart Array Essentials and Performance Controllers allow you the flexibility to choose the 12 Gbps controller to optimize your environment.

Options for OCP or PCIe standup adapters offer a choice of networking bandwidth and fabric, enabling you to scale to the changing needs of your business.

Support for PCIe Gen4, resulting in faster networking speed and higher bandwidth for network intensive applications.

Faster application performance with increased memory speeds up to 3200 MT/s.


Security Innovations

Silicon Root of Trust is an immutable fingerprint in the iLO silicon. The silicon root of trust validates the lowest level firmware to BIOS and software to ensure a known good state.

Tied into the silicon root of trust is the AMD Secure Processor, a dedicated security processor embedded in the AMD EPYC™ system on a chip (SoC). The security processor manages secure boot, memory encryption and secure virtualization.

Run Time Firmware Validation validates the iLO and UEFI/BIOS firmware at runtime. Notification and automated recovery is executed on detection of compromised firmware. If system corruption has been detected Server System Restore will automatically alert iLO Amplifier Pack to initiate and manage the system recovery process, avoiding lasting damage to your business by quickly restoring firmware to the factory settings or the last known authenticated safe setting.

HPE ProLiant DL325 G10 Plus front view:



HPE ProLiant DL325 G10 Plus back view:



HPE ProLiant DL325 G10 Plus top view:





SPEC_183 (I) Supplier provides the name and version of the Data management node operating system in detailed configuration above.

PROJECT Network Gateways

PROJECT data storage is the contracting authority storage designed for storing multi-year user project data and sharing of the data between IT4Innovations supercomputers.

PROJECT data storage is not yet implemented and the contracting authority does not know the exact equipment specifications.

PROJECT Network Gateways are servers intended for connecting EURO_IT4I system Compute Network with PROJECT data storage Access Network (100Gb/s Ethernet network).

SPEC_184 The delivery includes four PROJECT Network Gateways.

SPEC_185 Each PROJECT Network Gateway server meets the following requirements:

- Physical server;
- x86-64 architecture;
- Linux 64-bit, CENTOS 7 (or newer) or RHEL 7 (or newer) operating system;
- At least one processor, at least 3GHz base frequency, at least 16 CPU cores total;
- At least 32GiB RAM, operated in DDR4 mode, at least 2900MT/s with ECC;
- The RAM consists of memory modules of the same parameters, evenly distributed across all server memory;
- Theoretical throughput of processor(s) to RAM at least 160GB/s (in the offered server configuration);
- At least 2 local SSD disks with the capacity of 120GB, at a minimum, in RAID1;
- Hot-swap disks;
- Redundant, hot-swap power supply units, redundant power supplies.

SPEC_186 Each PROJECT Network Gateway node will provide a connection to EURO_IT4I system Compute Network with aggregated throughput of at least 100Gb/s.


SPEC_187 Each PROJECT Network Gateway node will provide a connection to PROJECT Access Network with aggregated throughput of 100Gb/s. For this purpose, the contracting authority reserves four 100Gb/s Ethernet QSFP28 type ports in edge devices of PROJECT data storage Access Network (two ports in each edge device). A part of the delivery are modules and optical cables required for connection to the edge devices of PROJECT data storage Access Network. Network cabling will be done in the Data center dropped ceiling.

SPEC_188 The supplier will cooperate with the contracting authority in:

Provisioning of Compute network adapter software stack on PROJECT Network Gateways;
Provisioning of interconnecting, addressing, and routing Compute Network and PROJECT
Access Network;
Provisioning of NFS mount-points of PROJECT file storages on the Compute nodes, Login
nodes, Visualization nodes, and Data management nodes.

SPEC_189 (I) In the offer, the supplier states count and detailed configuration of PROJECT Network Gateways.

PROJECT Network Gateways will be realized by **HPE Proliant DL325 Gen10 Plus** server.
The HPE Proliant DL325 Gen10 Plus server is in detail described in Data Management nodes section.




Infrastructure Nodes

- SPEC_190 Infrastructure nodes and their infrastructure are designed so that they ensure reliable, secure, fast, and effective operation and administration of EURO_IT4I system.
- SPEC_191 Virtualization technology is not used for implementation of infrastructure for Infrastructure nodes.
- SPEC_192 EURO_IT4I system includes 6x physical Infrastructure node. These servers provide 768GiB of RAM memory in total.
- SPEC_193 Each physical server referred in SPEC_192 meets the following requirements:
Disks in RAID with data redundancy
Hot swap disks
Redundant, hot swap power supply units, redundant power supplies
LAN connection requirement
- SPEC_194 Failure or outage of any single infrastructure node will not interrupt currently running computing jobs and operation of Compute clusters, Login nodes, Visualization nodes, storages, Data management nodes, backup, network, or supercomputer administration.
- SPEC_195 The key services will run in high-availability mode, preferably using the native mechanisms of the services.
-

SPEC_196 (I) Supplier states the number of infrastructure nodes, their purpose, and detailed configuration.

Infrastructure nodes purpose is to provide infrastructure services to EURO_IT4I system (DHCP, DNS, LDAP, license servers, schedulers, monitoring, logging, etc.) and for management and administration of the supercomputer, resources, and services.

Infrastructure Nodes will be realized by **HPE Proliant DL325 Gen10 Plus** server.
The HPE Proliant DL325 Gen10 Plus server is described in Data Management nodes section.



Backup Solution

SPEC_197 A part of the delivery will be a comprehensive solution for data backup.
EURO_IT4I system contains functionality and infrastructure, which enable creation of logically and hardware independent copies of production data on backup storage.

SPEC_198 The Backup solution particularly provides the following:
Backup of all delivered servers/nodes, with an exception of Compute nodes provisioned from images
Backup of image repository (images) for nodes provisioning (it may be a part of backup of INFRA storage)
Backup of HOME storage
Backup of INFRA storage


For HOME and INFRA storage, 80% space usage of total capacity and large amount of small files is assumed. It is assumed that 4% of storage total capacity will change daily.

Temporary server data, which is not required for server operation recovery are not backed up.


SPEC_199 Data are backed up every 24 hours between 10:00 p.m. and 6:00 a.m. (i.e. during an eight-hour period, at a maximum).


SPEC_200 The backup solution provides data recovery from the last 21 daily backups.

- SPEC_201 The backup solution provides sufficient backup capacity to meet all the requirements (with regard to the technology used and realistically achievable parameters in the given application).
- SPEC_202 Data backup and recovery priorities (from the highest to lowest priority):
1. Highest priority: Nodes critical for providing system services, INFRA storage
2. Other nodes
3. Lowest priority: HOME storage
- SPEC_203 The data backup and restore solution meets the following basic properties:
Time schedule of backups
Individual files and directories backup and restore
Files and directories' owners, permissions, extended ACLs, and attributes backup and restore
Parallel run of backups and recoveries
- SPEC_204 The backup solution uses a backup-to-disk technology.
- SPEC_205 Backup uses a dedicated, independent data storage. Backup data storage meets the requirements of Technical requirements specification provided by contracting authority.
- SPEC_206 The backup solution has a minimal negative impact on EURO_IT4I system operation and performance.
- SPEC_207 The backup system license covers all the requirements in maximum scope.
- SPEC_208 (I) Following table states the backup solution capacity and throughput sufficient to meet the backup requirements.





- SPEC_209 (I) Supplier below describes in detail the calculation of required capacity to meet the contracting authority's backup requirements.







SPEC_210 (I) Supplier below describes solution architecture, configuration of equipment, description of backup software solution, its properties, and suggested backup policy.










LAN Infrastructure


LAN - General


- SPEC_211 The LAN provides communication between the equipment in the supercomputer and connection to the contracting authority's WAN.
- SPEC_212 The LAN includes the public and private parts of the network.
- SPEC_213 The public parts of the network (hereinafter referred to as public networks) will use IPv4 and IPv6 public addresses and will provide services accessible from the Internet (Login nodes, Visualization nodes, Data management nodes).
- SPEC_214 The private parts of the network (hereinafter referred to as private networks) will use IPv4 private addresses and will be used for internal services and device management.
- SPEC_215 Devices in private networks will have access to the Internet over NAT.
NAT will be provided by a contracting authority's network device located in the contracting authority's WAN.
- SPEC_216 The network solution supports IPv6.
- SPEC_217 The LAN will be divided into several L3 networks. For each L3 network, a different L2 network will be used (VLAN or another network device). It is not desirable to use one L2 network for multiple L3 networks. LAN division will separate the following types of communication into different L3 networks:
- Services accessible from the Internet
 - Data communication between nodes (services)
 - Infrastructure nodes and services
 - Network devices management
 - Disk arrays, storage management
 - Server (BMC, IPMI, etc.) management
 - Non-IT infrastructure (power supply, cooling, etc.) management




LAN – Edge devices

- SPEC_218 Connection of the LAN with the contracting authority’s WAN will be implemented by redundant L3 switches (hereinafter referred to as Edge devices, labeled as ED1, ED2 in the figure above).
Offered edge devices meets the following requirements:
Dual power supply
Redundant server connection using multi-chassis ether channel
Edge devices interconnection throughput of 800Gb/s, at a minimum
Separate routing instances for private and public networks
Support Access Control Lists (ACL) for at least 2000 ingress and 2000 egress ACL rules configurable for each port independently.
- SPEC_219 The Edge switches have no separate control-plane.
- SPEC_220 Edge devices support user authentication using RADIUS or TACACS+ protocol, creating of different user roles for network management (operator, administrator, etc.), and issued commands logging.
- SPEC_221 Edge devices support import and export of its device configuration using TFTP, FTP, SCP, or SFTP protocol. Configuration allows saving in the form, which allows its editing using text editor.
- SPEC_222 Edge devices support IPFIX (Internet Protocol Flow Information eXport) and sFlow v5. Data flow export provides information about all IP packets or about one packet from 10 processed IP packets.







LAN – connections

- SPEC_223 LAN connection speed of devices:
Connection speed of each Login node, Visualization node and Data mgmt node is 100Gb/s.
LAN connection speed of each Infrastructure node is 10Gb/s.
The Compute nodes have direct connection to LAN, the speed of each node is 1Gb/s.
- SPEC_224 Minimum speed specified in SPEC_223 will be reachable when accessed from the contracting authority's WAN. For systems specified in SPEC_225, the minimum speed requirement will be met even in the case when any single active network device fails.
- SPEC_225 Connection of Login nodes, Visualization nodes, Data management nodes, and Infrastructure nodes to LAN is redundant.
- SPEC_226 Aggregated connection speed of Login nodes, Visualization nodes, and Data management nodes in LAN reachable from the contracting authority's WAN is 400Gb/s.
- SPEC_227 Aggregated connection speed of Compute nodes in LAN reachable from the contracting authority's WAN is more than 10Gb/s as shown on LAN schematics above.
- SPEC_228 There is direct connection of Compute nodes to LAN.

LAN – Active network devices

These requirements are common for all active network devices including Edge devices.

- SPEC_229 Active network devices support VLAN segmentation using 802.1Q standard with at least 100 VLANs and numbering from 1 to 4094.
- SPEC_230 Active network devices support IGMPv2 and IGMPv3 protocols.
- SPEC_231 Active network devices, which are L3 devices, support IPv4 multicast routing and PIM Sparse Mode and PIM Source-Specific Multicast protocols.
- SPEC_232 Active network devices are remotely controllable and centrally manageable.
- SPEC_233 Active network devices support SNMPv2 and SNMPv3 protocols enabling:
1. Reading of device state and ports utilization
2. Sending of SNMP traps for defined events
- SPEC_234 Active network devices support remote access using SSH2 protocol with at least AES256-CTR encryption protocol, HMAC-SHA256 hashing protocol and RSA key of 4096 bits.
- SPEC_235 Management interfaces of active network devices will be connected to the contracting authority's OOB network.
Serial management interfaces of edge devices will be connected to the contracting authority's OOB router (Cisco 2901/K9). The supplier will provide required cables and installation.
Ethernet management interfaces of all active network devices and of nodes described in SPEC_279 will be connected to the contracting authority's OOB network implemented by the contracting authority's OOB switch. The supplier will utilize one GBase-T Ethernet port of the contracting authority's OOB switch. The supplier will provide required cables and cabling installation as per cable routes requirements described in tender documentation.
- SPEC_236 (I) The description of the Edge devices is provided in this chapter above.
-

WAN Integration

- SPEC_237 Connection to the contracting authority's WAN/LAN will provide connectivity of 4x100Gb/s.
- SPEC_238 Each LAN edge device of EURO_IT4I system will be connected to each of the two central devices of the contracting authority's WAN/LAN using a 100Gb/s connection. For this purpose, the contracting authority reserves four ports in WAN/LAN central devices WAN1 and WAN2. The ports are equipped with QSFP-100G-SR4 modules with MPO-12 connectors. The delivery includes the modules and fiber-optic cables required for connection of edge devices to the WAN1 and WAN2 devices. The network cabling will be installed in the data center suspended ceiling.
- SPEC_239 Twelve Cat6a connections will be provided between the location of LAN edge devices and the contracting authority's WAN/LAN described in tender documentation. The supplier will provide Cat6a patch panel with RJ-45 ports at the location of LAN edge devices and connect it with existing patch panel in the contracting authority's WAN/LAN rack. Connections will be used for OOB connections mentioned in SPEC_235.
- SPEC_240 Connection of edge devices and the contracting authority's WAN devices will be realized over Ethernet and private IPv4 in separate routing instances.
- SPEC_241 Connection to the contracting authority's WAN devices will be redundant.
- SPEC_242 (I) Supplier above in the offer provides description of the WAN integration.

Response to General Technical Requirements

- SPEC_243 All servers (with the exception of Compute nodes) and disk arrays (including disk modules of disk arrays) use:
- Redundant, hot-swap power supply units
 - Redundant power supplies
 - Hot-swap disks
- SPEC_244 All offered devices and systems provide remote management.
- SPEC_245 All servers (nodes) allows replacement of CPUs, RAM memory modules, and Compute network adapters (if present) independently of other components.
- SPEC_246 All servers have remote network management, independent of OS operation, providing power supply control, reset option, graphical console and virtual media.
- SPEC_247 Disk array controllers are redundant. Failure or outage of any single disk array controller will not cause a failure of disk array services. Disk array controller cache is protected against data loss or corruption in case of power outage or any single controller failure.
- SPEC_248 All disk array controllers and server RAM modules use an error detection and correction mechanism – Error-correcting code memory (ECC memory).
- SPEC_249 The parameters and characteristics of the disks used in the solution are appropriate to their deployment and load. The warranty also applies to disk wear (including SSDs and NVMe drives) due to excessive use.
-

- SPEC_250 Normal operation and accessibility of declared capacities do not require operator's intervention.
- SPEC_251 In the case of SSH, SCP, and/or SFTP protocol support, protocol version 2 is supported.
- SPEC_252 All devices and cables will be physically labeled with unique identification (labels). The identification will be easily accessible, legible and properly recorded.

Software

- SPEC_253 A part of EURO_IT4I system is a comprehensive HPC system software solution, all software equipment required to fulfill the contracting authority's requirements (firmware, operating systems, drivers, management, and monitoring software, file systems, scheduler, etc.). All delivered software will be licensed for unlimited time, with no additional fees (at the version level applicable at the time of delivery).

Software - Environment

- SPEC_254 EURO_IT4I system will deliver a transparent, uniform, shared user environment to end users.
- SPEC_255 EURO_IT4I system will provide, to the maximum extent possible, uniform environment on all nodes accessible to end users, i.e. Login nodes, Compute nodes and Visualization nodes. User environment on nodes of the same type will be identical. User environment on nodes of different types will be identical to the maximum extent possible; environment differences arise only from differences in the nature and services of servers. Environment means especially behavior of systems, applications, names of files, directories, commands, environment variables, values of environment variables, configuration, etc. Nodes will provide en_US and C locale.
- SPEC_256 Login nodes will provide end users with a command-line and graphical interface. Login nodes will provide shell (bash and csh) and X Window System (X11).
- SPEC_257 EURO_IT4I system will provide end users with unified account and unified authentication. End user accounts will be accessible and identical on all Login, Compute, and Visualization nodes and Data management nodes.
- SPEC_258 User account implementation will use LDAP and SSSD technology. User accounts will be stored and managed in the LDAP database. User accounts and groups will be created via posixAccount and posixGroup LDAP schemas. User accounts and groups will be managed via a secure LDAP protocol. The contracting authority will manage user accounts and groups exclusively. The contracting authority provides centralized management of all IT4Innovations center users; the contracting authority will synchronize EURO_IT4I system users and groups information from their central repository to EURO_IT4I system LDAP database. User home directories will also be managed exclusively by the contracting authority. EURO_IT4I system LDAP service will be implemented using two LDAP servers with a replicated database. LDAP service will be implemented on suitable infrastructure nodes. For the LDAP service, the OpenLDAP software will be used. SSSD service will be used for name/uid, group/gid resolving on nodes providing user access. Secure LDAP protocol will be used for accessing LDAP service. Login nodes, Visualization nodes, and Data management nodes will use SSSD enumeration.
- SPEC_259 The unified user authentication will use SSH keys. The SSH keys will be managed by the contracting authority or user and will be stored in the LDAP database or user's home directory.
-

The SSH service will allow authentication of authorized user's SSH keys against user's entry in the LDAP database. It is assumed that the solution will be implemented using a configuration of SSHD and SSSD features. The supplier based on Technical requirements specification provided by contracting authority expects that there will be no end user authentication requiring password authentication.

Software - Scheduler

- SPEC_260 EURO_IT4I system will provide unified access to computing resources. For access to computing resources and for performing computational jobs on Compute nodes, Visualization nodes, and Data management nodes, EURO_IT4I system will use an advanced job scheduler and resource manager PBS Professional (hereinafter referred to as Scheduler).
- SPEC_261 Offered scheduler will effectively utilize available computing resources, taking into account specific features of Compute nodes and Compute network. Scheduler will allow advanced scheduling policies. Scheduler allows executing of batch and interactive jobs. Offered scheduler supports job priorities and job dependencies. Offered scheduler takes into account current resources availability. Scheduler will handle compute resource failures and outages. Offered scheduler allows scheduling according to resources reserved statically or dynamically.
- SPEC_262 Offered scheduler allows checking status of computational jobs. Scheduler also allows event notification (on job abort/begin/end execution) by email. Offered scheduler allows users to submit jobs and monitor their status. Scheduler also returns user jobs output back to submit location. Offered scheduler allows operators to display jobs status, scheduler's server and queues status and details of individual jobs.
- SPEC_263 Scheduler will use high availability cluster in the active-passive mode.
- SPEC_264 Scheduler will provide information about users' computational jobs resource consumption for accounting purposes. Provided information will be especially job runtime, names of allocated Compute nodes and number of allocated compute cores.
- SPEC_265 Offered scheduler allows execution of custom scripts defined by the contracting authority before and after job execution (prologue, epilogue), when submitting jobs to Scheduler and while modifying jobs (hooks).
- SPEC_266 Offered scheduler allows creating allocation of a set of nodes of specified properties available in a specified period to a specified group of users (reservation).
- SPEC_267 Offered scheduler allows limiting the number of concurrent user jobs, the number of concurrent user group jobs, the maximum duration of a job, and access to computational resources only to authorized users (ACL).
- SPEC_268 Offered scheduler takes into account the Compute network topology to allow better parallel job placement on Compute nodes with higher communication efficiency.
- SPEC_269 Offered scheduler provides command-line interface and API in the C language.
- SPEC_270 Offered scheduler is PBS compatible.
- SPEC_271 The server part of the Scheduler will be implemented on suitable Infrastructure nodes with 3GHz base frequency CPU cores.
For implementation of the solution, INFRA storage will be used.
-

- SPEC_272 Scheduler license covers all Compute nodes (execution), Login nodes (submit) Visualization nodes (execution and submit) and Data management nodes (execution).
- SPEC_273 The supplier will provide basic installation and configuration of Scheduler to demonstrate its functionality and requirements fulfillment. The contracting authority can further customize the configuration to meet his needs.

PBS Professional – detailed description:

PBS Professional offers comprehensive workload management for high performance computing and cloud environments. The workload management suite allows HPC users to simplify their environment while optimizing system utilization, improving application performance, and improving ROI on hardware and software investments. PBS Professional is the preferred solution for many of the largest, most complex clusters and supercomputers —and at the same time is the choice for smaller organizations needing HPC solutions that are easy to adopt and use.

PBS Professional includes the following features and corresponding benefits:

- The software accelerates job execution and selects optimal job placement across diverse, broadly distributed resources. With PBS Professional it is easy to create intelligent policies to manage distributed, mixed-vendor computing assets as a single, unified system.
- Fast, powerful scheduling with power and topology awareness including support for NVIDIA® and AMD accelerators.
- Scalability and throughput: Tested to 50,000+ nodes, PBS Professional scales to support millions of cores with fast job dispatch and minimal latency. The software can support 1,000,000+ jobs per day.
- **Higher utilization** - Users can run jobs - or portions of jobs - in the period immediately before a planned outage. Typically, computer systems remain unused for several hours prior to outages since insufficient time is available to complete a job before the outage starts. PBS Professional fills those holes with malleable “shrink-to-fit” jobs, allowing useful work to be accomplished during the pre-outage period when otherwise no jobs would be running on the system - and thereby **providing greater than 95 percent utilization**.
- Green Provisioning™ for automatic resource shutdown/restart to conserve energy (proven to lower one customer's energy use by up to 30 percent).
- **Flexible plugin framework** - PBS Professional offers a powerful yet easy to use plugin framework to customize implementations for meeting complex user requirements.

Software - Management

- SPEC_274 EURO_IT4I system includes tools for management of all systems and services, which are a part of the system.
- SPEC_275 EURO_IT4I system includes tools for remote management of all hardware devices (servers/nodes, disk arrays, network switches, etc.) particularly allowing for configuration and management of the devices, critical event detection and notification of such events via SMTP (email) or SNMP.
- SPEC_276 EURO_IT4I system will provide a command-line tool for remote power and reset control of all servers. Functionality will be independent of the operating system of the controlled servers (it is assumed the IPMI protocol will be used).
Functionality will be accessible on suitable Infrastructure nodes (nodes intended for management).
-

- SPEC_277 EURO_IT4I system will provide a command-line tool for remote display and manipulation of console terminals of all servers. Functionality will be independent of the operating system of the controlled servers (it is assumed the IPMI protocol will be used). Functionality will be accessible on suitable Infrastructure nodes (nodes intended for management).
- SPEC_278 EURO_IT4I system will provide a command-line tool for effective centralized remote management of node operating systems individually (single node) and in bulk (groups of nodes or all nodes). The system will provide, in particular, remote command execution; file transfer, modification and deletion; and comparison of command outputs and command return values. The system will allow parallel action execution. Remote management will also work effectively in the case some of the managed nodes are not communicating or operating correctly, and will identify and report action failures. The supplier will provide suitable server naming and server group definitions for effective use of the centralized remote management system. Functionality will be accessible on suitable Infrastructure nodes (nodes intended for management). ClusterShell software will be used for the remote management system.
- SPEC_279 The functionalities described in SPEC_276 to SPEC_278, and tools for Scheduler management will be collectively available on two Infrastructure nodes (nodes intended for management).
- SPEC_280 EURO_IT4I system will provide effective remote server installation from the central image repository. The system will allow individual (single node) or bulk (multiple nodes) installation. The system will allow parallel server installation. The server part of the remote server installation will be operated on suitable Infrastructure nodes.
- SPEC_281 The remote server installation according to SPEC_280 will be used for Compute nodes installation (with exception of Universal nodes) and will allow installation of Login nodes, Visualization nodes and Data management nodes.
- SPEC_282 EURO_IT4I system will provide network boot service from the central image repository. The system will allow individual (single node) or bulk (multiple nodes) network boot. The system will allow parallel server boot. The server part of the remote server installation will be operated on suitable Infrastructure nodes.
- SPEC_283 The network boot according to SPEC_282 will be used for Universal nodes. The network boot will be available and can be used for Accelerated nodes and Cloud nodes. Complete startup of all nodes in every compute partition using network boot will finish in less than 15 minutes.
- SPEC_284 System will provide management of images in the central image repository, in particular, creation, modification, and deletion of images. Servers of identical type will use an identical installation image; the number of images will be limited. The central image repository will have the total capacity min. 1TB. Central image repository will be implemented on INFRA storage.
- SPEC_285 EURO_IT4I system subsystems will be interconnected and integrated to provide effective operation and management of EURO_IT4I system without unnecessary obstructions and non-productive activities.
-

As a management software will be delivered **HPE Performance Cluster Manager**.

HPE Performance Cluster Manager – general description:

HPE Performance Cluster Manager is a fully integrated system management solution offering all the functionalities you need to manage your Linux®-based high performance computing (HPC) clusters. The software provides system setup, hardware monitoring and management, image management and software updates as well as power management for systems of any scale - up to 100,000 nodes. The HPE Performance Cluster Manager reduces the time and resources spent administering HPC systems - lowering total cost of ownership, increasing productivity and providing a better return on your hardware investments. HPE Performance Cluster Manager (HPCM) provides a single, consistent, and powerful management interface for all HPE HPC systems. HPE Performance Cluster Manager provides a comprehensive and customizable interface to monitor essential system metrics and initiate management actions from a single point of control. HPE Performance Cluster Manager reduces the time and resources spent administering the system by streamlining software maintenance procedures and automating repetitive tasks.

With HPE Performance Cluster Manager offers these functionalities and corresponding benefits:

- Fast System Setup - Provision thousands of cluster nodes in less than an hour. Adding new hardware does not require system shutdowns.
- Hardware management - Fine-grained centralized monitoring and management of all aspects of your cluster hardware - CPUs, memory, accelerators, cooling and networking.
- Image Management and Software Updates - Securely provision any software image on all or select cluster nodes to accommodate wide variety of requirements.
- Power Management - Collect and analyze power metrics for all hardware for better capacity planning and resource allocation.
- ISV and Open Source Software Integration - Integration with ISV and Open Source workload management and orchestration tools, system management tools for more ease of use.

HPE Performance Manager Features-at-a-Glance

System Setup:

- Provision thousands of cluster nodes in less than an hour
- Add new cluster nodes without shutdown of system
- Automatic hardware discovery
- Configuration information stored in a secure repository

Image Management and Software Updates:

- Secure software image repository supports Linux OS, middleware and applications
- Image formats: RPM, ISO, gold image
- Any software image can be provisioned on all or select cluster nodes
- Accountability built in to track changes

Power Management:

- Collect and analyze power metrics (for nodes, chassis, rack, PDU, CDU)
- Topology-aware Power On/Off
- Power capping
- Assign/account power resources for jobs

Hardware Management:



- Collect various telemetry data and store in a secure repository
 - Analyze behavior of metrics at a specific point in time or over a period of time
 - View metrics and alerts via GUI, CLI, Ganglia, Nagios, ELK
 - Setup automatic reactions to alerts
 - Broadcast commands to selected nodes
 - FRU inventory for serviceability
-

- Find differences between nodes
- Integrated firmware flashing for BIOS, BMC/iLO, CMC, network adapter/switch
- Central management is protected with secure administrative node

ISV and Open Source Software Integration:

- PBS Professional, Slurm: Launch tool, manage resources, use selected tool commands from HPCM
- Mellanox UFM: View IB switch metrics with associated cluster node metrics
- Ganglia, Nagios Core, ELK: Integrated for system monitoring, alerts and analytics
- DCGM connector: Integrated management of NVIDIA GPUs
- Ansible: HPCM stores hardware metrics in Ansible Playbook format

HPE Performance Cluster Manager (HPCM)





Software – OS and Applications

- SPEC_286 All servers/nodes use Red Hat Enterprise Linux or CENTOS operating system, major version 7 or newer.
- SPEC_287 All Universal nodes, Accelerated nodes, Cloud nodes, Login nodes, Visualization nodes, and Data management nodes will use the same distribution and the same version of the operating system – CentOS 7.
Data analytics node will use a different distribution of operating system than other Compute nodes – RedHat 7 (in accordance with SPEC_286); system libraries of these distributions have the same versions.
- SPEC_288 The operating system license will cover all nodes/servers in the delivery.
- SPEC_289 Delivered operating systems and software will provide updates, in particular security, reliability, functionality, and performance updates.
The supplier will provide and install software updates according to conditions specified in the contract.
- SPEC_290 The delivery contains the software necessary for acceptance tests execution (High Performance LINPACK, etc.) and for possible subsequent repetition of the tests during the warranty period, with the exception of the software provided by the contracting authority.
The contracting authority will provide Intel compilers and Intel MPI/MKL libraries in the latest stable version.
The contracting authority provides comprehensive management of application software, its acquisition, assembly/building, installation, and configuration.

Software – Addressing, Name Services

- SPEC_291 Address prefixes/ranges, domain name suffixes, namespaces, user accounts, and groups are managed and assigned solely by the contracting authority.
- SPEC_292 IP addressing will be in accordance with contracting authority's address policy and plan. Unless already specified in another section of the documentation, used IP address ranges will be specified after consultation with the contracting authority.
- SPEC_293 IP addressing will be static or dynamic (using DHCP).
- SPEC_294 All used IP addresses will resolve to names. IP addresses will be resolved by DNS servers provided by the supplier, used domain names suffixes will be specified after consultation with the contracting authority.
- SPEC_295 The contracting authority will have an option to change configuration of DNS services provided by the supplier and option to integrate contracting authority's own DNS services (views) to EURO_IT4I system.
-

SPEC_296 EURO_IT4I system equipment and systems will use accurate time. Time synchronization will be done using the contracting authority's NTP servers in the contracting authority's network.

Software – Monitoring

SPEC_297 EURO_IT4I system includes monitoring of accessibility and status of components and services (hereinafter referred to as availability monitoring). Availability monitoring will provide information about accessibility of all delivered equipment available in network over IP and information about accessibility/status of all relevant components/services of supplied nodes/servers, storages, etc. Relevant components/services are those that affect EURO_IT4I system services' functionality, accessibility, and/or performance. Availability monitoring will classify accessibility/state based on severity (OK, Warning, and Critical). Availability monitoring will allow configuring thresholds of monitored parameters for severity classification. Availability monitoring will provide identification of the equipment, component, or service to which availability or status information relates. Availability monitoring will provide timestamp of objects' accessibility/state change and record them in a log. Icinga2 software in clustered, highly available configuration will be used for implementation of availability monitoring. Icinga2 software will be operated on suitable Infrastructure nodes. Supplier will provide installation and configuration of the Icinga2 software. The contracting authority will complement the solution with their own set of checks and integrate the Icinga2 instance to their Icinga2 cluster. Infrastructure nodes will provide resources for operation of at least two Icinga2 instances.

SPEC_298 All important states and inaccessibility of EURO_IT4I system components and services will be reported to a single availability monitoring, according to SPEC_297.

SPEC_299 Availability monitoring as specified in SPEC_297 will particularly provide the following information:

1. Accessibility of device network interfaces in the IP networks (including accessibility of management interface, e.g. BMC interface) – for all equipment with network interface
2. Equipment power supply status – for all servers/nodes, disk arrays, and network devices
3. Equipment cooling status according to the temperature of the equipment and functionality of cooling fans/pumps – for all servers/nodes, disk arrays, and network devices
4. Status of the equipment according to its indicated status and/or the record in the equipment's log – for all servers/nodes, disk arrays, and network devices
5. Servers/node RAM status (indicates memory error occurrence)
6. Server/node CPU load status calculated for one CPU core of the server
7. Server/node memory usage status
8. Server/node status according to server/node file system capacity used
9. Server/node status according to the total number of server/node processes
10. Server/node status according to the number of server/node zombie processes
11. Server/node shared file storage (INFRA, SCRATCH, HOME, PROJECT) client mount status
12. Key server/node processes/services status (existence, basic response)
13. Availability of server/node network ports used by key server/node processes/services (accessibility)
14. Network interfaces status (Ethernet, FC, InfiniBand, etc.) of the server/node (up/down)

15. Server/node disks status according to disks' state and properties (up/down, disk failures, interface error (SAS, SATA, etc.), multipath status, etc.)
16. Status of disk array according to disk array controllers' state and the status of communication/data interfaces and disk array ports (both frontend and backend).
17. Disk array status according to the state of logical objects of disk array (RAID/volume group, logical volume)
18. Disk array status according to state of physical disks in disk array
19. Network devices status according to the state and properties of network interfaces (up/down, line speed) with respect to the expected state and configuration
20. Status of services provided in the high availability cluster and status of the high availability cluster itself

SPEC_300 EURO_IT4I system will collect performance and capacity parameters (metrics) and provide their visualization (hereinafter referred to as performance monitoring). Performance monitoring will provide data collection and charts:

For file storages:

1. current capacity usage and free capacity
2. current total number of files
3. current data throughput (read, write, total)
4. current number of I/O operations (read, write, total; metadata operations)

For SCRATCH storage - current capacity usage and throughput of physical storage (disk array, metadata/object storage target, etc.)

For Login, Visualization, Data management, and Infrastructure nodes:

1. current Ethernet interface throughput (send, receive, total)
2. current Compute network interface throughput (send, receive, total)

at a minimum, for three years, in 5 minute interval for the first month.

Performance monitoring will provide API for automated provision of data (metrics values).

Software – Logging

SPEC_301 EURO_IT4I solution systems will log and centrally store information about activities, status changes, events, etc. for a period of 3 months, at a minimum. Logs will include timestamp, system, service and user identification, and event description. Solution will provide a storage for collected data with the total capacity 20TB.

For centralized collection and processing of logs, the Elastic Stack will be used. The Elasticsearch solution will be implemented using native Elasticsearch cluster with at least three cluster members.

The Elasticsearch solution will be operated on suitable Infrastructure nodes.

For data storage of the Elasticsearch solution, Infrastructure nodes' local disks in RAID (with redundancy) will be used.

Software – Mail


SPEC_302 EURO_IT4I system will send all email communication exclusively through the contracting authority's SMTP servers.

The contracting authority will provide SMTP servers.

Software - Security

IT4Innovations National Supercomputing Center acquired the Information Security Management System Certification in accordance with ISO 27001 (ISO/IEC 27001:2013, ČSN ISO/IEC 27001:2014). EURO_IT4I system will be implemented in accordance with the internal regulations of the contracting authority.

- SPEC_303 EURO_IT4I system will provide access and services only to authorized users and systems. The system will not provide access or services to unauthorized users and systems. The system will be secured against data leaks, service misuse, and service and system breach.
- SPEC_304 Systems and services will use secure, strong passwords and secure keys, secure encryption, and secure protocols. No identical authentication data for different accounts or services will be used.
- SPEC_305 Services not required for proper operation and functionality of the solution will not be run/enabled/available on the nodes, preferably not even installed.
- SPEC_306 EURO_IT4I systems not communicate with other systems without explicit contracting authority's permission.
- SPEC_307 (I) Supplier describes the software solution and the names and number of licenses of the proposed software.




Integration into Data Center

- SPEC_308 The supplier will implement and deploy EURO_IT4I system in the contracting authority's infrastructure – IT4Innovations data center (hereinafter referred to as integration into the data center).
- SPEC_309 Integration into the data center includes all deliveries and activities the result of which will be deployment of EURO_IT4I system in IT4Innovations data center premises and infrastructure.
- SPEC_310 Integration into the data center will be carried out in accordance with applicable legislation and regulations and in accordance with the requirements and recommendations of the manufacturers of individual systems. For installed systems, revisions required by legislation and regulations will be carried out.
- SPEC_311 The supplier will respect the contracting authority's infrastructure. The infrastructure is described in Technical requirements specification provided by contracting authority.
-

- SPEC_312 Modification of the contracting authority's infrastructure necessary for implementation of EURO_IT4I system and its integration into the data center will be a part of EURO_IT4I system delivery.
- SPEC_313 The supplier will not change operational parameters of the contracting authority's data center. The offer does not use or assume data center infrastructure parameters that differ from those specified by the contracting authority's in the procurement documents.
- SPEC_314 All deliveries and activities of the integration into the data center will be discussed with and approved by the contracting authority. All use and possible modifications of the data center equipment or technologies (necessary for the integration into the data center) will be discussed with the contracting authority and the contracting entity providing data center service and must be approved by the contracting authority. Supplier agrees that the contracting authority reserves the right to change or modify technical integration proposals (routing, connection points, used materials, etc.) with respect to operating conditions and best practices of data center operation.

Placement

- SPEC_315 Offered system placement respects disposition of the contracting authority's data center premises.
- SPEC_316 All ICT equipment of EURO_IT4I system is designed for installation into racks and will be installed into racks or is equipment designed as a rack. Racks required for installation of the system are a part of the delivery.
- SPEC_317 A part of the delivery is connection of racks required for realization of EURO_IT4I system into the contracting authority's infrastructure.
- SPEC_318 EURO_IT4I system equipment will be placed within the area shown in the figure below.




- SPEC_319 The distance of the racks from obstacles (walls, columns, rows of other racks) and the way of installation of the devices in the racks will not obstruct delivery, installation, replacement, and service of all equipment.
-

- SPEC_320 The placement of all installed equipment of the delivered solution will not obstruct revision of electrical installation, cooling distribution, air-conditioning, and security systems (sensors, etc.).
- SPEC_321 Rack installation on data center floor will be stable and safe. Rack weight will be suitably distributed. Floor construction and maximum load capacity (25kN/m²) will be observed. In case the installation applies point load or excessive load on a small area, load spreader will be used. Floor tiles modifications (penetrations to the raised floor area) will be carried out in a way that does not reduce the tile's load capacity. Adjustment of the floor tiles will be carried out by the contracting authority's contractor providing the data center services.
- SPEC_322 Cabling (inside and outside the racks) will provide reliable connection, handling of equipment (e.g. when servicing the equipment), and prevent damage to the cables during handling. Penetrations to the raised floor area will be fitted with grommets.
- SPEC_323 (I) Supplier above provides a schema illustrating the proposed EURO_IT4I system placement in the contracting authority's data center.

Power Supply

- SPEC_324 EURO_IT4I system power supply and operation solution will respect the contracting authority's constraints, specifically power circuits parameters.
- SPEC_325 EURO_IT4I system will be supplied simultaneously from independent data center power supply branches (A and B).
- SPEC_326 Failure or outage of power circuit or power supply branch will not damage any equipment and will not cause danger to persons or property.
- SPEC_327 Failure or outage of any single power circuit or power supply branch will not cause:
Inaccessibility or outage of supercomputer services, with the exception of Universal and Accelerated compute nodes services
Inaccessibility or outage of Compute network services
Inaccessibility of more than 67% of Universal compute nodes
Inaccessibility of more than 67% of Accelerated compute nodes
After the restoration of power circuit or power supply branch, the power redundancy of all devices will be ensured automatically.
- SPEC_328 EURO_IT4I system is designed in such a way that under standard/even load of ICT systems the supply branches and individual phases of the power supply will be under approximately even load.
- SPEC_329 The solution takes into account the maximum input power for operation of all supplied equipment.
- SPEC_330 EURO_IT4I system allows the entire system to be shut down. Proper EURO_IT4I system shutdown will take 60 minutes, at a maximum.
- SPEC_331 (I) Supplier here below states information about the energy and weight requirements of offered solution.
-



Cooling

- SPEC_333 EURO_IT4I system cooling and operation solution respects the contracting authority's constraints, specifically data center cooling system parameters.
- SPEC_334 The offered solution is thermally neutral to the data room. The solution provides provide cooling of all heat released by EURO_IT4I system equipment.
- SPEC_335 As the cooling source, EURO_IT4I system will only use water cooling circuits (cold and hot water circuits) of the data center.
- SPEC_336 EURO_IT4I system cooling will be supplied simultaneously from independent data center water cooling circuits.
For cold water cooling, circuits SV1 and SV3 will be used simultaneously.
For hot water cooling, circuits TV1 and TV2 will be used simultaneously.
These circuits are considered the primary ones. Internal cooling loops of the supplier's equipment (e.g. CDU) are considered the secondary ones.
- SPEC_337 Failure or outage of the cooling circuits will not damage any equipment and will not cause danger to persons or property.
- SPEC_338 Failure or outage of any single data center cooling circuit will not cause inaccessibility or outage of supercomputer services.
- SPEC_339 EURO_IT4I system is designed in such a way that under standard/even load of ICT systems the cold water cooling circuits are under approximately even load.

Offered EURO_IT4I system needs more than 300 kW of cooling capacity from the hot water cooling circuits (TV1 and TV2), so it will utilize both hot water cooling circuits and cooling is designed in such a way that under standard/even load of ICT systems the hot water cooling circuits will be under approximately even load.
The even distribution of the cooling is done by logical units.
The fulfillment of the requirement regarding even load is not based on alternating/switching of cooling circuits.
- SPEC_340 Offered Universal partition uses the hot water cooling circuits (TV1 and TV2) as prevailing source of cooling and uses direct liquid cooling. Universal partition will be fully operational with primary hot cooling circuits temperature set at any value in range of 30°C to 32°C.
-


SPEC_341 Each connection of EURO_IT4I system cooling to the primary cooling circuits will allow individual remotely controlled switching between two primary cooling.


Switching between the primary cooling circuits will take 3 minutes, at the maximum. Switching will not cause inaccessibility or outage of EURO_IT4I system services and will not damage any equipment in the data center. Circuit switching will not cause undesired primary circuits connections.


Remotely controlled cooling circuits switching will be integrated into the contracting authority's Measurement and regulation system (MaR). The integration will be performed by the contracting authority in cooperation with the supplier. The offered solution is compatible with MP-Bus technology, it provides an MP-Bus interface and act as an MP-Bus Slave device controlled by MaR MP-Bus Master.


SPEC_342 In case of risk of exceeding the maximum operating temperature of the equipment or components of the equipment (e.g. in case of failure of the source cooling circuit), the affected equipment will be shut down automatically and immediately to eliminate the risk of overheating or damaging the equipment or components.


SPEC_343 (I) Supplier describes the cooling solution of EURO_IT4I system. Specifically, the supplier provides information about the cooling of Universal partition, Accelerated partition and other compute partitions and systems, about cooling load distribution, and about cooling circuits switching.












SPEC_344 (I) Supplier provides a schema and parameters of EURO_IT4I system connections to the data center cooling circuits.

HPE Adaptive Rack Cooling System Cooling Unit racks are connected to the DC cold-water cooling circuits by 2" hoses (as shown on following schematic pictures). Each unit is connected to two circuits SV1 and SV3 using a group (2 on inlet and 2 on return side) of remotely controlled automatic valves for switching between circuits. Solution includes in total 9 HPE ARCS cooling units and those will be connected by at minimum 4 groups of switching valves to ensure the cooling load split to both DC circuits.





Transport

SPEC_345 During equipment transport and installation, the load-bearing capacity of the floor on transport route will not be exceeded.

The load-bearing capacity of the floor in the data center and its access corridor (room 219 and 223) is 25kN/m².

The load-bearing capacity of the floor in the access area (room 217 and 218) is 5kN/m². For the transport of material, the supplier will temporarily install load distribution plates (e.g. plywood panels) in this area, so that the resulting load of the floor during the physical delivery does not exceed 5kN/m² and thus prevent floor damage.

Implementation and Further Actions

Implementation

- SPEC_346 A part of the delivery is a comprehensive implementation of EURO_IT4I system so that all the contracting authority's requirements are met.
- SPEC_347 The delivery includes design, delivery, installation, implementation, configuration, debugging, testing of all systems, and performance of acceptance tests.

Training

- SPEC_348 A part of the delivery will be a training in the scope and detail sufficient to acquire the knowledge necessary for the independent operation and administration of EURO_IT4I system.
- SPEC_349 The training will provide the information necessary to understand the internal functioning of the systems, hardware, and software. The training will include a thorough introduction to operating procedures and administration of EURO_IT4I system.
- SPEC_350 Duration of the training will be 25 hours, at a minimum. The number of participants will not be limited to less than 16 participants.
- SPEC_351 The training will include introduction to management framework/tools, practical demonstration and work with the real system. Fully implemented and accepted EURO_IT4I system will be used for demonstrations.
- SPEC_352 Trainers will be HPE skilled professionals. The training will be done in Czech or English.
- SPEC_353 The supplier will provide the teaching materials and presentations in English.
- SPEC_354 The schedule and detailed plan of the training will be produced in cooperation with the contracting authority's project manager.
- SPEC_355 The training will take place in the contracting authority's premises. For this purpose a conference room will be provided free of charge.
- SPEC_356 (I) Supplier includes below the frame schedule and content of the training.

Training will be structured into 5 working days with expected following frame schedule and content structure. The final agenda will be in detail finalized during prepare of Detailed technical specification phase of delivery in agreement with contracting authority taking into account knowledge base of supposed participants of the training.

Day 1: Introduction to the EURO_IT4I solution and individual subsystems
Universal Partition
Accelerated Partition
Data Analytics Node
Cloud Nodes
Compute Network and LAN
Login Nodes
Visualization Nodes
Scratch Storage

Home Storage
Infra Storage
Data Management Nodes
Project Network Gateways
Infrastructure Nodes
Backup Solution

Day 2: Management of the EURO_IT4I system
Infrastructure services configuration
Management and Administration tools
Monitoring and issue reporting
Power and Cooling Circuits
Healthcheck and Operating Parameters

Day 3: Administration of SCRATCH, HOME and INFRA Storages
Present configuration and possible modifications
Configuration of new data areas

Day 4: Job Scheduler, Backup Administration
Present configuration and possible modifications
System resources dedicated to system users
Access Environment and Authentication

Day 5: Management and Administration cont.
Practical examples
Hands-on
Q&A

Documentation

- SPEC_357 The supplier will provide a comprehensive documentation of EURO_IT4I system.
- SPEC_358 EURO_IT4I system documentation will comprehensively cover all supplied systems and will be logically structured. The documentation will include the documentation of the actual state of the system and documentation of operation procedures (operating manuals). The documentation will describe all specific (custom) modifications (settings, functionalities, etc.).
- SPEC_359 The documentation will also cover processes and procedures for system administration, regular maintenance, emergency management, and recovery.
- SPEC_360 The documentation will also include English documentation (datasheets/spec sheets, manuals, administrator, and user guides) of delivered hardware and software. Documentation will be provided in electronic form, allowing text copying.

Benchmarking support

- SPEC_361 The supplier will provide expert support for the benchmarking of the system. The optimization of benchmark performance, rule conforming execution and the submission of the results to the official lists will be performed by the supplier. We understand that the contracting authority intends to include the system in the following lists benchmarking:
1. Top500, HPCG, Graph500
-

- i. Universal Partition
 - ii. Accelerated Partition
 - iii. Data Analytics Partition
2. Green500
 - i. Universal Partition
 - ii. Accelerated Partition
3. IO500
 - i. Universal Partition
 - ii. Accelerated Partition
4. MLPerf
 - i. Accelerated Partition

EC Declaration of Conformity

SPEC_362 All delivered systems and equipment will be accompanied by EC declaration of conformity.

Waste Disposal

SPEC_363 A part of the delivery is the disposal of waste produced by implementation of the delivery. Supplier understands that the contracting authority is not obligated to store packaging and packaging material and will not do so.

Warranty and Services

The system has warranty and will be serviced for 5 years.

Annex 1 - 2 – Technical Parameters of the Offer

Supplier

HEWLETT-PACKARD s.r.o.

We ask you to fill in the parameter values on all subsequent worksheets of the workbook.

The last worksheet of the workbook is titled "End".

Fill in/modify only the cells marked in yellow.

The supplier is not authorized to interfere with other parts of the workbook or modify the preset functions in any way.

Parameter values must be provided in accordance with the instructions, meaning, and method of measurement specified in the text of Annex 1 to the tender documents.

Values on the printed original document are considered binding.

The values or information presented here must correspond to the information stated in the technical solution proposal, which is part of the tender.

Universal partition

Parameter	Value	Minimum requested value
Number of nodes	720	
Server model	HPE Apollo2000 G10+	
Number of CPUs per node	2	2
CPU model	[REDACTED]	
CPU core count	64	
CPU frequency [GHz]	[REDACTED]	
CPU TDP [W]	[REDACTED]	
CPU Rpeak in double precision [GFLOPS]	[REDACTED]	
Node RAM size [GiB]	256	256
Node CPUs core count	128	
Node CPUs Rpeak in double precision [GFLOPS]	4000	[REDACTED]
Partition CPUs core count	92160	
Partition Rpeak in double precision, CPUs only [TFLOPS]	[REDACTED]	
Partition RAM size [TiB]	180	
Universal partition LINPACK Rmax computing performance, using CPUs only [TFLOPS]	1935	2 346,00

Accelerated partition	Parameter	Minimum requested value	
		Value	
	Number of nodes		70
	Server model	HPE Apollo6500 G10+	
	Number of CPUs per node	2	2
	CPU model	AMD 7452	
	CPU core count		32
	CPU frequency [GHz]		2,35
	CPU TDP [W]		155,00
	CPU Rpeak in double precision [GFLOPS]		1 203,20
	Node RAM size [GiB]	512	512
	Node CPUs core count		64
	Node CPUs Rpeak in double precision [GFLOPS]	2400	2406,4
	Partition CPUs core count		4480
	Partition Rpeak in double precision, CPUs only [TFLOPS]		168,448
	Partition RAM size [TiB]		35
	Number of GPUs per node	8	8
	GPU model	Nvidia A100	
	GPU Rpeak in double precision [TFLOPS]	17	19,50
	GPU HBM size [GiB]	40	40,00
	Accelerated partition LINPACK Rmax computing performance [TFLOPS]	4500	6 616,00

Data analytics partition

Parameter	Value	Minimum requested value
Number of nodes	1	1
Server model	HPE Superdome Flex	
Number of CPUs per node	32	32
CPU model	[REDACTED]	
CPU core count		24
CPU frequency [GHz]		[REDACTED]
CPU TDP [W]		[REDACTED]
CPU Rpeak in double precision [GFLOPS]		[REDACTED]
Node RAM size [TiB]	24	24
Node CPUs core count		768
Node CPUs Rpeak in double precision [GFLOPS]		[REDACTED]
Partition CPUs core count		768
Partition Rpeak in double precision, CPUs only [TFLOPS]		[REDACTED]
Partition RAM size [TiB]		24
Data analytics partition LINPACK Rmax computing performance, using CPUs only [TFLOPS]	40	40,00

Cloud partition

Parameter	Minimum requested value	Value
Number of nodes		36
Server model		HPE Apollo2000 G10+
Number of CPUs per node	2	2
CPU model		[REDACTED]
CPU core count		64
CPU frequency [GHz]		[REDACTED]
CPU TDP [W]		[REDACTED]
CPU Rpeak in double precision [GFLOPS]		[REDACTED]
Node RAM size [GiB]	256	256
Node CPUs core count		128
Node CPUs Rpeak in double precision [GFLOPS]	4000	[REDACTED]
Partition CPUs core count		4608
Partition Rpeak in double precision, CPUs only [TFLOPS]		[REDACTED] 131,00
Partition RAM size [TiB]		9
Cloud partition LINPACK Rmax computing performance, using CPUs only [TFLOPS]	131	131,00

Partitions' summary

Parameter	Minimum requested value	Value
Universal partition LINPACK Rmax computing performance, using CPUs only [TFLOPS]	1935	2346
Accelerated partition LINPACK Rmax computing performance [TFLOPS]	4500	6616
Data analytics partition LINPACK Rmax computing performance, using CPUs only [TFLOPS]	40	40
Cloud partition LINPACK Rmax computing performance, using CPUs only [TFLOPS]	131	131
The sum of LINPACK Rmax computing performance of the individual compute partitions [TFLOPS]	8600	9133

Compute network

Parameter

Compute network technology (IB HDR, etc.)

Compute network topology

Minimum requested value

Value

IB HDR

non-blocking Fat Tree

Parameter	Minimum requested value	
	Value	
SCRATCH storage		
net available capacity on file-system level [TB]	1000	1 000,00
long-term sustainable sequential read performance for the 1MiB block [GB/s]	500	1 000,00
long-term sustainable sequential write performance for the 1MiB block[GB/s]	350	350,00
long-term sustainable random I/O performance for the block size of 4KiB and 80%/20% read/write mode[Mega IOPs]	5	5,00
storage software technology	Lustre	

HOME storage	Parameter	Minimum requested value	
		Value	
	net available capacity on file-system level [TB]	25	25,00
	long-term sustainable sequential read performance for the 1MiB block [GB/s]	1,2	1,20
	long-term sustainable sequential write performance for the 1MiB block[GB/s]	1,2	1,20
	long-term sustainable random I/O performance for the block size of 4KiB and 80%/20% read/write mode[IOPs]	7000	7 000,00
	storage software technology	NFS	

INFRA storage	Parameter	Minimum requested value	
		Value	
	net available capacity on file-system level [TB]	25	25,00
	long-term sustainable sequential read performance for the 1MiB block [GB/s]	1,2	1,20
	long-term sustainable sequential write performance for the 1MiB block[GB/s]	1,2	1,20
	long-term sustainable random I/O performance for the block size of 4KiB and 80%/20% read/write mode[IOPs]	7000	7 000,00
	storage software technology	NFS	

Minimum requested value

Servers	Parameter	Value
	number of Login nodes	4
	number of Visualization nodes	2
	number of Data management nodes	2
	number of PROJECT Network Gateways	4
	number of Infrastructure nodes	6
	number of physical servers - Infrastructure nodes	6
	number of physical servers - Storages (excl. PROJECT storage)	4
	number of physical servers - Backup	2

Minimum requested value

Value

Backup

Parameter

backup disk storage net available capacity [TB]
expected backup capacity taking into account specific features of the
solution (compression, deduplication) [TB]





Software

Parameter

Value

- OS - Compute nodes
- OS - Login nodes
- OS - Visualization nodes
- OS - Data management nodes
- OS - Infrastructure nodes
- Remote visualization software
- Backup Software
- Scheduler



	Parameter	Requested value	Value
Placement	total weight of installed devices [t]		27,39
Cooling	<p>Universal partition uses the hot water cooling circuits (TV1 and TV2) as prevailing source of cooling and uses direct liquid cooling. Universal partition is fully operational with primary hot cooling circuits temperature set at any value in range of 30°C to 32°C.</p> <p>Accelerated partition uses the hot water cooling circuits (TV1 and TV2) as prevailing source of cooling. Accelerated partition is fully operational with primary hot cooling circuits temperature set at any value in range of 30°C to 32°C.</p>	yes	yes
Power	<p>maximum electrical input [kVA]</p> <p>expected average electrical input for maximal compute load [kVA]</p>		

Parameter

Value

Compliance with the requirements of the tender dossier

Please, state that the offered solution is in accordance with the requirements of the tender dossier.

Technical requirements specification for EURO_IT4I system

SPEC_1	yes	yes
SPEC_2	yes	yes
SPEC_3	yes	yes
SPEC_4	yes	yes
SPEC_5	yes	yes
SPEC_6	yes	yes
SPEC_7	yes	yes
SPEC_8	yes	yes
SPEC_9	yes	yes
SPEC_10	yes	yes
SPEC_11	yes	yes
SPEC_12	yes	yes
SPEC_13	yes	yes
SPEC_14	yes	yes
SPEC_15	yes	yes
SPEC_16	yes	yes
SPEC_17	yes	yes
SPEC_18	yes	yes
SPEC_19	yes	yes
SPEC_20	yes	yes
SPEC_21	yes	yes
SPEC_22	yes	yes
SPEC_23	yes	yes
SPEC_24	yes	yes
SPEC_25	yes	yes
SPEC_26	yes	yes
SPEC_27	yes	yes
SPEC_28	yes	yes
SPEC_29	yes	yes
SPEC_30	yes	yes
SPEC_31	yes	yes
SPEC_32	yes	yes
SPEC_33	yes	yes
SPEC_34	yes	yes
SPEC_35	yes	yes
SPEC_36	yes	yes
SPEC_37	yes	yes
SPEC_38	yes	yes
SPEC_39	yes	yes
SPEC_40	yes	yes
SPEC_41	yes	yes
SPEC_42	yes	yes
SPEC_43	yes	yes
SPEC_44	yes	yes

SPEC_45	yes	yes
SPEC_46	yes	yes
SPEC_47	yes	yes
SPEC_48	yes	yes
SPEC_49	yes	yes
SPEC_50	yes	yes
SPEC_51	yes	yes
SPEC_52	yes	yes
SPEC_53	yes	yes
SPEC_54	yes	yes
SPEC_55	yes	yes
SPEC_56	yes	yes
SPEC_57	yes	yes
SPEC_58	yes	yes
SPEC_59	yes	yes
SPEC_60	yes	yes
SPEC_61	yes	yes
SPEC_62	yes	yes
SPEC_63	yes	yes
SPEC_64	yes	yes
SPEC_65	yes	yes
SPEC_66	yes	yes
SPEC_67	yes	yes
SPEC_68	yes	yes
SPEC_69	yes	yes
SPEC_70	yes	yes
SPEC_71	yes	yes
SPEC_72	yes	yes
SPEC_73	yes	yes
SPEC_74	yes	yes
SPEC_75	yes	yes
SPEC_76	yes	yes
SPEC_77	yes	yes
SPEC_78	yes	yes
SPEC_79	yes	yes
SPEC_80	yes	yes
SPEC_81	yes	yes
SPEC_82	yes	yes
SPEC_83	yes	yes
SPEC_84	yes	yes
SPEC_85	yes	yes
SPEC_86	yes	yes
SPEC_87	yes	yes
SPEC_88	yes	yes
SPEC_89	yes	yes
SPEC_90	yes	yes
SPEC_91	yes	yes
SPEC_92	yes	yes
SPEC_93	yes	yes
SPEC_94	yes	yes

SPEC_95	yes	yes
SPEC_96	yes	yes
SPEC_97	yes	yes
SPEC_98	yes	yes
SPEC_99	yes	yes
SPEC_100	yes	yes
SPEC_101	yes	yes
SPEC_102	yes	yes
SPEC_103	yes	yes
SPEC_104	yes	yes
SPEC_105	yes	yes
SPEC_106	yes	yes
SPEC_107	yes	yes
SPEC_108	yes	yes
SPEC_109	yes	yes
SPEC_110	yes	yes
SPEC_111	yes	yes
SPEC_112	yes	yes
SPEC_113	yes	yes
SPEC_114	yes	yes
SPEC_115	yes	yes
SPEC_116	yes	yes
SPEC_117	yes	yes
SPEC_118	yes	yes
SPEC_119	yes	yes
SPEC_120	yes	yes
SPEC_121	yes	yes
SPEC_122	yes	yes
SPEC_123	yes	yes
SPEC_124	yes	yes
SPEC_125	yes	yes
SPEC_126	yes	yes
SPEC_127	yes	yes
SPEC_128	yes	yes
SPEC_129	yes	yes
SPEC_130	yes	yes
SPEC_131	yes	yes
SPEC_132	yes	yes
SPEC_133	yes	yes
SPEC_134	yes	yes
SPEC_135	yes	yes
SPEC_136	yes	yes
SPEC_137	yes	yes
SPEC_138	yes	yes
SPEC_139	yes	yes
SPEC_140	yes	yes
SPEC_141	yes	yes
SPEC_142	yes	yes
SPEC_143	yes	yes
SPEC_144	yes	yes

SPEC_145	yes	yes
SPEC_146	yes	yes
SPEC_147	yes	yes
SPEC_148	yes	yes
SPEC_149	yes	yes
SPEC_150	yes	yes
SPEC_151	yes	yes
SPEC_152	yes	yes
SPEC_153	yes	yes
SPEC_154	yes	yes
SPEC_155	yes	yes
SPEC_156	yes	yes
SPEC_157	yes	yes
SPEC_158	yes	yes
SPEC_159	yes	yes
SPEC_160	yes	yes
SPEC_161	yes	yes
SPEC_162	yes	yes
SPEC_163	yes	yes
SPEC_164	yes	yes
SPEC_165	yes	yes
SPEC_166	yes	yes
SPEC_167	yes	yes
SPEC_168	yes	yes
SPEC_169	yes	yes
SPEC_170	yes	yes
SPEC_171	yes	yes
SPEC_172	yes	yes
SPEC_173	yes	yes
SPEC_174	yes	yes
SPEC_175	yes	yes
SPEC_176	yes	yes
SPEC_177	yes	yes
SPEC_178	yes	yes
SPEC_179	yes	yes
SPEC_180	yes	yes
SPEC_181	yes	yes
SPEC_182	yes	yes
SPEC_183	yes	yes
SPEC_184	yes	yes
SPEC_185	yes	yes
SPEC_186	yes	yes
SPEC_187	yes	yes
SPEC_188	yes	yes
SPEC_189	yes	yes
SPEC_190	yes	yes
SPEC_191	yes	yes
SPEC_192	yes	yes
SPEC_193	yes	yes
SPEC_194	yes	yes

SPEC_195	yes	yes
SPEC_196	yes	yes
SPEC_197	yes	yes
SPEC_198	yes	yes
SPEC_199	yes	yes
SPEC_200	yes	yes
SPEC_201	yes	yes
SPEC_202	yes	yes
SPEC_203	yes	yes
SPEC_204	yes	yes
SPEC_205	yes	yes
SPEC_206	yes	yes
SPEC_207	yes	yes
SPEC_208	yes	yes
SPEC_209	yes	yes
SPEC_210	yes	yes
SPEC_211	yes	yes
SPEC_212	yes	yes
SPEC_213	yes	yes
SPEC_214	yes	yes
SPEC_215	yes	yes
SPEC_216	yes	yes
SPEC_217	yes	yes
SPEC_218	yes	yes
SPEC_219	yes	yes
SPEC_220	yes	yes
SPEC_221	yes	yes
SPEC_222	yes	yes
SPEC_223	yes	yes
SPEC_224	yes	yes
SPEC_225	yes	yes
SPEC_226	yes	yes
SPEC_227	yes	yes
SPEC_228	yes	yes
SPEC_229	yes	yes
SPEC_230	yes	yes
SPEC_231	yes	yes
SPEC_232	yes	yes
SPEC_233	yes	yes
SPEC_234	yes	yes
SPEC_235	yes	yes
SPEC_236	yes	yes
SPEC_237	yes	yes
SPEC_238	yes	yes
SPEC_239	yes	yes
SPEC_240	yes	yes
SPEC_241	yes	yes
SPEC_242	yes	yes
SPEC_243	yes	yes
SPEC_244	yes	yes

SPEC_245	yes	yes
SPEC_246	yes	yes
SPEC_247	yes	yes
SPEC_248	yes	yes
SPEC_249	yes	yes
SPEC_250	yes	yes
SPEC_251	yes	yes
SPEC_252	yes	yes
SPEC_253	yes	yes
SPEC_254	yes	yes
SPEC_255	yes	yes
SPEC_256	yes	yes
SPEC_257	yes	yes
SPEC_258	yes	yes
SPEC_259	yes	yes
SPEC_260	yes	yes
SPEC_261	yes	yes
SPEC_262	yes	yes
SPEC_263	yes	yes
SPEC_264	yes	yes
SPEC_265	yes	yes
SPEC_266	yes	yes
SPEC_267	yes	yes
SPEC_268	yes	yes
SPEC_269	yes	yes
SPEC_270	yes	yes
SPEC_271	yes	yes
SPEC_272	yes	yes
SPEC_273	yes	yes
SPEC_274	yes	yes
SPEC_275	yes	yes
SPEC_276	yes	yes
SPEC_277	yes	yes
SPEC_278	yes	yes
SPEC_279	yes	yes
SPEC_280	yes	yes
SPEC_281	yes	yes
SPEC_282	yes	yes
SPEC_283	yes	yes
SPEC_284	yes	yes
SPEC_285	yes	yes
SPEC_286	yes	yes
SPEC_287	yes	yes
SPEC_288	yes	yes
SPEC_289	yes	yes
SPEC_290	yes	yes
SPEC_291	yes	yes
SPEC_292	yes	yes
SPEC_293	yes	yes
SPEC_294	yes	yes

SPEC_295	yes	yes
SPEC_296	yes	yes
SPEC_297	yes	yes
SPEC_298	yes	yes
SPEC_299	yes	yes
SPEC_300	yes	yes
SPEC_301	yes	yes
SPEC_302	yes	yes
SPEC_303	yes	yes
SPEC_304	yes	yes
SPEC_305	yes	yes
SPEC_306	yes	yes
SPEC_307	yes	yes
SPEC_308	yes	yes
SPEC_309	yes	yes
SPEC_310	yes	yes
SPEC_311	yes	yes
SPEC_312	yes	yes
SPEC_313	yes	yes
SPEC_314	yes	yes
SPEC_315	yes	yes
SPEC_316	yes	yes
SPEC_317	yes	yes
SPEC_318	yes	yes
SPEC_319	yes	yes
SPEC_320	yes	yes
SPEC_321	yes	yes
SPEC_322	yes	yes
SPEC_323	yes	yes
SPEC_324	yes	yes
SPEC_325	yes	yes
SPEC_326	yes	yes
SPEC_327	yes	yes
SPEC_328	yes	yes
SPEC_329	yes	yes
SPEC_330	yes	yes
SPEC_331	yes	yes
SPEC_332	yes	yes
SPEC_333	yes	yes
SPEC_334	yes	yes
SPEC_335	yes	yes
SPEC_336	yes	yes
SPEC_337	yes	yes
SPEC_338	yes	yes
SPEC_339	yes	yes
SPEC_340	yes	yes
SPEC_341	yes	yes
SPEC_342	yes	yes
SPEC_343	yes	yes
SPEC_344	yes	yes

SPEC_345	yes	yes
SPEC_346	yes	yes
SPEC_347	yes	yes
SPEC_348	yes	yes
SPEC_349	yes	yes
SPEC_350	yes	yes
SPEC_351	yes	yes
SPEC_352	yes	yes
SPEC_353	yes	yes
SPEC_354	yes	yes
SPEC_355	yes	yes
SPEC_356	yes	yes
SPEC_357	yes	yes
SPEC_358	yes	yes
SPEC_359	yes	yes
SPEC_360	yes	yes
SPEC_361	yes	yes
SPEC_362	yes	yes
SPEC_363	yes	yes

Annex 1 - 3 - Technical requirements specification for EURO_IT4I system

1 Contracting Authority's Intention

The contracting authority intends to modernize and increase the performance of the IT4Innovations National Supercomputing Center (hereinafter referred to as IT4Innovations) computing infrastructure operated by the contracting authority.

Purpose of the "EURO_IT4I system" public procurement is to implement a new supercomputer – EURO_IT4I system – that will replace the current Salomon supercomputer.

EURO_IT4I system will be operated by the IT4Innovations National Supercomputing Center and will be used to provide services in the field of High Performance Computing.

EURO_IT4I system will be installed and operated in the contracting authority's data center located in the IT4Innovations building, at the Technical University of Ostrava.

2 Subject of the Procurement

The subject of the "EURO_IT4I Supercomputer" public procurement is the delivery of a comprehensive High Performance Computing system, i.e. a complex of computing, storage, network, and other systems and software solution, including implementation and integration into the contracting authority's data center, training, warranty, support, and other services.

3 Document Structure

The following labels are shown in the following text:

SPEC_Nr. indicates individual requirements of the procurement for easier identification.

SPEC_Nr. (I) indicates that additional information is required. The supplier must state this information in their offer.

4 EURO_IT4I System

- SPEC_1 The contracting authority requires a comprehensive computing system solution designed for HPC, i.e. a complex of computing, storage, network, and other systems and software solution, including implementation and integration into the contracting authority's data center (hereinafter referred to as the *solution* or *EURO_IT4I system*).
- SPEC_2 EURO_IT4I system must allow effective execution of a large number of simultaneous computing jobs of all phases of the job life cycle (i.e. preparation, pre-processing, computing, post-processing) of various types (especially parallel, but also serial, batch, and interactive) of large number of users. EURO_IT4I system must allow secure and reliable user data storage, fast access to data, and effective administration of systems, components, resources and services.
- SPEC_3 EURO_IT4I system must offer powerful computing resources, which are easily accessible by users and their jobs. The solution must provide features, services, and functions required for effective system operation and administration by the contracting authority. The solution must be balanced and parameters and structure of individual subsystems must consider other subsystems.



- SPEC_4 EURO_IT4I system delivery must include all systems, equipment, accessories, licenses, documentation, design, implementation, and other work, training, etc. necessary to meet the contracting authority's requirements.
- SPEC_5 The solution must consider disposition and limitations resulting from the environment and conditions of the contracting authority's data center.
- SPEC_6 The solution must not contain any limitations preventing or restricting the operation of EURO_IT4I system by the contracting authority in required, expected, or rational scope. EURO_IT4I system is intended for approximately 3,000 users.
- SPEC_7 The solution as a whole must completely fulfill the contracting authority's requirements. Required functionality and features must be realistically functional and usable in operation of the solution; required parameters must be realistically achievable. Fulfillment of the contracting authority's requirements must not be conditioned.
The fulfillment of the contracting authority's requirements must not be based solely on functionality, properties, or parameters of individual components.
- SPEC_8 The functionality, features, and parameters of the solution must be stated for the proposed/delivered configuration designed for everyday use. The functionality, properties, and parameters must not be conditioned
- SPEC_9 The solution must meet all contracting authority's technical requirements simultaneously. All required features, functions, and parameters must be achieved using a single production setup of all the components in the solution. Fulfillment of the requirements must not be conditioned on changing the setup or interconnection of the components.
- SPEC_10 To the maximum extent possible, the solution must be autonomous, independent of external systems and services, and self-sufficient without the need for additional equipment, systems, or services.
- SPEC_11 Design and implementation of the solution must ensure reliable, secure, powerful, and efficient operation of EURO_IT4I system in the contracting authority's data center.


4.1 EURO_IT4I System Components

- SPEC_12 EURO_IT4I system must include a *Compute cluster*. A Compute cluster consists of Compute nodes connected by a Compute network — a high-speed, low latency network. A Compute cluster is designed for computing users' jobs.
- SPEC_13 The Compute cluster must contain four partitions:
- Universal compute partition
 - Accelerated compute partition
 - Data analytics compute partition
 - Cloud infrastructure compute partition
- SPEC_14 The *Universal compute partition* must consist of standard servers without accelerators such as GPUs or FPGAs and should be based on x86 CPU architecture to provide quick accessibility for the users and their existing codes.

- SPEC_15 The *Accelerated compute partition* should deliver most of the compute power usable for HPC but also excellent performance in HPDA and AI workloads, especially in the learning phase of Deep Neural Networks. The Accelerated compute partition should consist of heavily GPU-accelerated nodes.
- SPEC_16 The *Data analytics compute partition* should be oriented on supporting huge memory jobs by implementing a NUMA SMP system with large cache coherent memory.
- SPEC_17 The Cloud infrastructure compute partition should support both the research and operation of the Infrastructure/HPC as a Service. The Cloud infrastructure compute partition is intended for provision and operation of cloud technologies like OpenStack and Kubernetes.
- SPEC_18 EURO_IT4I system must include *Login nodes* – nodes for user access, job and data preparation, code compilation and debugging, result processing, and data transfer.
- SPEC_19 EURO_IT4I system must include *Visualization nodes* – nodes for remote visualization and user data modeling using hardware-accelerated graphics applications.
- SPEC_20 EURO_IT4I system must include data storages called *Storages*. Storages are designed for storing and sharing data. Storages are a comprehensive solution of storage devices, I/O servers (e.g. file servers), network, and necessary software. The storages must provide required data services.
- SPEC_21 EURO_IT4I system must include *HOME storage*. *HOME storage* is a file storage designed for storing users' operating system settings and users' applications on the supercomputer.
- SPEC_22 EURO_IT4I system must include *SCRATCH storage*. *SCRATCH storage* is a powerful file storage designed for short-term users' job data. Compute nodes use SCRATCH storage intensively.
- SPEC_23 EURO_IT4I system must include *INFRA storage*. *INFRA Storage* is a file storage designed for storing and sharing infrastructure data of the supercomputer. The storage is used for storing system images of nodes, logs, infrastructure service data, application software, scheduler data, etc.
- SPEC_24 EURO_IT4I system must include *PROJECT Network Gateways*. *PROJECT Network Gateways* are servers intended for connecting EURO_IT4I system with PROJECT data storage to allow access to PROJECT data storage. PROJECT data storage is the contracting authority's external storage designed for multi-year user project data and sharing of the data between IT4Innovations supercomputers.
- SPEC_25 EURO_IT4I system must include *Data management nodes*. *Data management nodes* are designed for data transfer between storages and are used primarily to transfer data from/to supercomputer via parallel transfers to relieve the load on Login nodes. The nodes also serve to directly access the storages from external locations (Internet) by the users.
- SPEC_26 EURO_IT4I system must include infrastructure and management nodes (hereinafter referred to as *Infrastructure nodes*). *Infrastructure nodes* are designed to provide infrastructure services to EURO_IT4I system (e.g. DHCP, DNS, LDAP, provisioning, license servers, schedulers, monitoring, logging, etc.) and for management and administration of the supercomputer, resources, and services.
- SPEC_27 EURO_IT4I system must include a data *Backup* solution.

- SPEC_28 EURO_IT4I system must include *Network infrastructure*, i.e. a network connection of the components and systems to achieve required functionality, provide access to individual services, and provide performance, availability, and security of the services.
- SPEC_29 The network infrastructure consists of cluster *Compute network, LAN, other networks*, as designed by the supplier, and integration into contracting authority's WAN.
Compute network connects Compute nodes of the compute cluster, Login nodes, Visualization nodes, Data management nodes, and other servers.
LAN provides communication between components inside the supercomputer.
Integration into WAN network of the contracting authority provides connection to contracting authority's WAN network and Internet connectivity for the supercomputer.
- SPEC_30 EURO_IT4I system must include the solution and infrastructure for installation and operation of the supercomputer in the contracting authority's data center (hereinafter referred to as *Infrastructure for operation in data center*). This includes racks and accessories required for placement of the supercomputer, the solution for power supply, and cooling of the supercomputer, systems, interfaces, and connections to the infrastructure of contracting authority's data center.
- SPEC_31 EURO_IT4I system must include all necessary software and licenses.
- SPEC_32 EURO_IT4I system delivery must include all required equipment and systems needed for required functionality, performance, and effective operation of the computing system (including those not explicitly stated in this document).
- SPEC_33 For implementation of EURO_IT4I system, the supplier must not use contracting authority's equipment, systems, infrastructure, and services unless explicitly stated in this document and only for purpose and scope stated in this document.

4.2 EURO_IT4I System Diagram



The EURO_IT4I system diagram is a schematic/simplified illustration of EURO_IT4I system; it does not demonstrate either a complete system or all or exact interconnections. Grey-colored parts are not a part of the EURO_IT4I system delivery and are supplied by the contracting authority.

4.3 Requirements – Compute Cluster/Partitions

- SPEC_34 The following benchmark must be used for Compute cluster performance measurement:
- High Performance LINPACK <http://www.netlib.org/benchmark/hpl/>
- Referential benchmark implementation is not required for the measurement; optimized benchmark implementation can be used. Applied benchmark implementation must fully comply with benchmark specifications. If an optimized version of the benchmark is used, it has to be supplied to the contracting authority no later than at the time of acceptance tests.
- SPEC_35 Rmax computing performance is determined by executing a High Performance LINPACK benchmark running parallel on all Compute nodes of the given partition (a single benchmark instance on the whole partition).
- SPEC_36 Computing performance must be specified for the offered/delivered configuration designed for normal operation. The determined computing performance must not be conditioned in any way (e.g. by using a specific processor mode not suitable for long-term system operation; or by declaring expected efficiency not guaranteed by the supplier). Declared computing performance must be proven by corresponding acceptance tests during implementation.
- SPEC_37 The sum of LINPACK Rmax computing performance of the individual compute partitions must be 8600 TFLOPS, at a minimum.
- SPEC_38 (I) In the offer, the supplier must specify the theoretical computing performance R_{peak} of individual Compute node for all compute partitions.

SPEC_39 (I) In the offer, the supplier must state Rmax computing performance according to SPEC_35 for each compute partition.

4.4 Requirements – Compute Nodes

SPEC_40 Each Compute node must meet the following requirements:

- x86-64 architecture
- 64-bit operating system, Linux, see section 4.20.4 Software – OS and Applications for details
- Compute network connection
- LAN connectivity

SPEC_41 Compute node RAM must be DDR4 with ECC.

SPEC_42 RAM must be evenly distributed (in terms of throughput, capacity, and access time) on processors and CPU cores of Compute node. RAM must consist of memory modules of the same type (size, rank, etc.), and must be distributed evenly, with the same configuration, on memory controllers and memory channels of Compute node. All memory channels of every processor in a node must be utilized.

SPEC_43 Compute nodes are intended exclusively for computing. The supplier must not use any Compute node to provide other functionality.

SPEC_44 All Compute nodes in the same partition must have identical hardware and software (e.g. firmware) configuration and must operate using the same configuration (frequency, timing, feature setting).

SPEC_45 (I) In the offer, the supplier must state the hardware platform of the solution and detailed configuration of Compute nodes including processor model.

SPEC_46 (I) In the offer, the supplier must provide the name and version of Compute node operating system.

4.5 Requirements – Universal Partition – Universal Nodes

SPEC_47 Universal partition must provide LINPACK Rmax computing performance of 1935 TFLOPS, at a minimum, using CPUs only.

SPEC_48 Each Universal node must meet the following requirements:

- Two CPU sockets/processors per node
- No accelerators
- For CPUs only, node theoretical computing performance Rpeak of 4000 GFLOPS, at a minimum, in double precision
- at least 2GiB of RAM per CPU core
- at least 256GiB of RAM per node
- RAM DDR4, at least 3000MT/s
- diskless; network boot from central image repository is required
- Compute network connection 100Gb/s, at a minimum
- LAN connectivity, see section 4.17 Requirements – LAN Infrastructure for details

4.6 Requirements – Accelerated Partition - Accelerated Nodes

- SPEC_49 Accelerated partition must provide LINPACK Rmax computing performance of 4500 TFLOPS, at a minimum.
- SPEC_50 Each Accelerated node must meet the following requirements:
- Two CPU sockets/processors per node
 - PCI-Express gen 4, supported by all key components including CPUs, GPUs and Compute network adapters
 - For CPUs only, node theoretical computing performance Rpeak of 2400 GFLOPS, at a minimum, in double precision
 - 512GiB of RAM, DDR4, at least 3000MT/s
 - LAN connectivity, see section 4.17 Requirements – LAN Infrastructure for details
- SPEC_51 Each Accelerated node must contain at least eight GPU accelerators of the same type.
- SPEC_52 Each GPU accelerator of accelerated nodes must meet the following requirements:
- CUDA technology support
 - Theoretical computing performance Rpeak of 17 TFLOPS in double precision, at a minimum
 - 40GiB HBM memory
- SPEC_53 The GPU accelerators in accelerated node must be connected using high-speed connections/network. Connections between any two GPU accelerators in the node must provide theoretical throughput of at least 250GB/s in each direction simultaneously. The requirement must be met for simultaneous, parallel communication of all GPUs (GPU pairs) in the node. Each GPU accelerator must provide theoretical throughput of at least 250GB/s in each direction simultaneously.
- SPEC_54 Each Accelerated node must contain PCIe gen 4 switches enabling the connection of GPU accelerators and Compute network adapters and CPUs.
- SPEC_55 Each Accelerated node must provide Compute network connection 400Gb/s (2x200Gb/s or 4x100Gb/s). Compute network adapters need to be evenly distributed across the CPU sockets, connected directly to the PCIe switches where the GPU accelerators are connected.
- SPEC_56 Accelerated node must provide unified memory spanned over node's GPUs memory and node's RAM memory. The unified memory must be accessible using CUDA API.

4.7 Requirements – Data Analytics Partition – Data Analytics Node

- SPEC_57 The Data analytics partition must provide the LINPACK Rmax computing performance of 40 TFLOPS, at a minimum.
- SPEC_58 EURO_IT4I system must contain one Data analytics node.
- SPEC_59 The Data analytics node must meet the following requirements:
- NUMA SMP architecture, cache coherent memory access
 - Thirty two CPU sockets/processors
 - 24TiB RAM, DDR4, at least 2900MT/s, coherent memory
 - Compute network connection 400Gb/s (4x100Gb/s or 2x200Gb/s)
 - LAN connectivity, see section 4.17 Requirements – LAN Infrastructure for details

- SPEC_60 The Data analytics node must be a single system allowing usage of all the resources by a single job, under a single instance of running OS. The Data analytics node must provide on-demand software partitioning using technologies like cgroups and cpusets to allow maximum utilization by multiple jobs/users with smaller requirements.
- SPEC_61 The Data analytics node must provide fast memory access; the latency of remote NUMA node memory access must not be greater than six times the latency of local NUMA node memory access, idle latencies are considered. The memory latency will be evaluated using Intel Memory Latency Checker or similar tool.
- SPEC_62 Software implementation of the memory coherency is not permitted.

4.8 Requirements – Cloud Partition – Cloud Nodes

- SPEC_63 Cloud partition must provide the LINPACK Rmax computing performance of 131 TFLOPS.
- SPEC_64 Each Cloud node must meet the following requirements:
- Two CPU sockets/processors per node
 - 256GiB of RAM, DDR4, at least 3000MT/s
 - 2 local NVMe disks with capacity of at least 480GB in RAID1, NVMe PCIe 4x lanes
 - Compute network connection 100Gb/s, at a minimum
 - LAN connectivity, see section 4.17 Requirements – LAN Infrastructure for details
- SPEC_65 The Cloud node must have CPUs with the same architecture, core count, and frequency as the Universal nodes.
- SPEC_66 Each Cloud node must have two dedicated (extra) Ethernet ports (at least 10Gb/s each) connected to two Ethernet networks/switches (for redundancy) providing VLANs to allow better utilization of the VMs running on the platform and the possibility to create user defined networks. These networks/switches are considered as a part of the LAN infrastructure, so requirements stated in section 4.17 Requirements – LAN Infrastructure apply.

4.9 Requirements – Compute Network

- SPEC_67 Compute network must use a RDMA technology with the interconnection throughput of 100Gb/s, at a minimum, and interconnection latency of 10 microseconds, at a maximum. The equipment connection to Compute network and inside Compute network must meet the above-mentioned requirements for connections.
- SPEC_68 Compute network must comply with one of the following characteristics:
- A. Compute network is a fully non-blocking network. Compute network consists of one switching device only – the director switch. The director switch has provably fully non-blocking internal interconnection and (theoretical) full bisection bandwidth.
 - B. The Compute network topology is fully non-blocking Fat Tree. Compute network consists of two layers of switches, one layer of switches (leaf switches) is intended for connecting endpoints, the second layer of switches (spine switches) interconnects leaf switches; the layers are mutually bipartite connected. Tree topology of Compute network is balanced.
 - C. The Compute network topology is Dragonfly+. Compute network is organized into groups/islands. Each island alone (intra-group connections) provides fully non-blocking communication of its endpoints. Each island consists of two layers of switches, one layer of switches (leaf switches) is intended for connecting endpoints, the second layer of switches (spine switches) is intended for connecting to other islands; the layers are mutually bipartite connected. Each island provides connectivity (aggregated throughput)

to other islands equal or greater than 4/9 of connectivity of its endpoints. Connectivity to other islands is spread evenly. The maximum number of islands is four. Storages' requested performance (see section 4.12 Requirements – Storages), notably SCRATCH performance, can be reached from clients in any single island.

- D. The Compute network topology is Enhanced Hypercube. Each switch in the network provides the same connectivity (the number of links and their distribution and throughput) to the network. Switch connectivity is spread to all dimensions of the network. The number of links to each dimension is the same number or differs by one (i.e. for any two dimensions of the network the maximum difference of link count is one); links of the same speed are used; in the case of link count difference, lower dimensions have greater number of links than higher dimensions. The number of hypercube dimensions is less than or equal to six. For each switch, the ratio of connectivity-to-the-network to connectivity-to-endpoints is greater than or equal to 2.2. For the ratio calculation, if provided connectivity of endpoint is higher than connectivity requested in technical specification, the lower value can be used (e.g. if 100Gb/s connection to Compute network is requested and 200Gb/s speed is provided, then 100Gb/s can be used for ratio calculation).

- SPEC_69 Compute network must use effective routing optimized for the topology and the configuration of Compute network.
- SPEC_70 For the selected network topology/configuration (with regard to switch configuration and routing mechanism), connection of Accelerated nodes to Compute network must provide optimal communication.
For Fat Tree topology, corresponding ports of Accelerated nodes (e.g. port 1 of each Accelerated node) must be connected to minimal number of leaf switches.
For Dragonfly+ topology, all Accelerated compute nodes must be located in one island, corresponding ports of Accelerated nodes must be connected to minimal number of leaf switches.
For Enhanced Hypercube topology, corresponding ports of Accelerated nodes must be connected to minimal number of switches in the smallest sub-hypercube of Compute network.
- SPEC_71 The solution must provide an effective MPI communication of Compute nodes.
- SPEC_72 Compute network must support and provide an IP protocol.
- SPEC_73 A system-wide LINPACK benchmark running parallel on all Compute nodes of the Compute cluster (utilizing only CPUs) must be possible.
- SPEC_74 (I) In the offer, the supplier must state the technology, topology, throughput, latency of Compute network, and the equipment configuration.
- SPEC_75 (I) In the offer, the supplier must state configuration and characteristics of proposed Compute network relevant for proposed network topology.

4.10 Requirements – Login Nodes

- SPEC_76 EURO_IT4I system must include, at a minimum, four Login nodes.
- SPEC_77 Each Login node must meet the following requirements:
- Physical node
 - x86-64 architecture
 - Two CPU sockets/processors per node

- For CPUs only, node theoretical computing performance Rpeak of 2400 GFLOPS, at a minimum, in double precision
- 256GiB RAM DDR4, at least 3000MT/s, ECC
- 2 local SSD disks with capacity of at least 480GB in RAID1
- Hot-swap disks
- Compute network connection 100Gb/s, at a minimum
- LAN connectivity, redundant, see section 4.17 Requirements – LAN Infrastructure for details
- Redundant, hot-swap power supply units, redundant power supplies

- SPEC_78 The Login nodes must provide user access implementing an SSH2 protocol and services for SCP and SFTP file transfer.
- SPEC_79 The Login nodes must use the identical technology and instruction set of processors as Universal nodes.
- SPEC_80 The Login nodes must have the same hardware configuration and must operate in identical operating settings (frequency, latency, properties settings).
- SPEC_81 The Login nodes are intended exclusively for providing user access and operations. The supplier must not use any Login node to provide other functionality.
- SPEC_82 (I) In the offer, the supplier must state the hardware platform of the solution and detailed configuration of Login nodes including processor model.
- SPEC_83 (I) In the offer, the supplier must provide the name and version of the Login node operating system.

4.11 Requirements – Visualization Nodes

- SPEC_84 EURO_IT4I system must include two Visualization nodes, at a minimum.
- SPEC_85 Each Visualization node must meet the following requirements:
- Physical node
 - x86-64 architecture
 - Two sockets/processors per node
 - For CPUs only, node theoretical computing performance Rpeak of 2400 GFLOPS, at a minimum, in double precision
 - 256GiB DDR4 RAM, at least 3000MT/s, ECC
 - one powerful GPU card with OpenGL support
 - 2 local SSD disks with capacity of at least 480GB in RAID1
 - Hot-swap disks
 - Compute network connection 100Gb/s, at a minimum
 - LAN connectivity, redundant, see section 4.17 Requirements – LAN Infrastructure for details
 - Redundant, hot-swap power supply units, redundant power supplies
- SPEC_86 The Visualization nodes must provide remote hardware-accelerated visualization with OpenGL support.
- SPEC_87 Visualization node GPU cards must meet the following requirements:
- Theoretical computing performance Rpeak of 12 TFLOPS, at a minimum, in single precision
 - At a minimum 24GiB GPU memory

- GPU memory must be a continuous space with uniform addressing and uniform performance.

SPEC_88	Visualization nodes must be fully integrated into EURO_IT4I system in the same scope as Compute nodes. Specifically, this includes user account integration, authentication, integration and accessing Visualization node services via Scheduler, and storage accessibility.
SPEC_89	Visualization nodes must use the identical technology and instruction set of processors and as Universal nodes.
SPEC_90	Visualization nodes must have the same hardware configuration and must operate in identical operating settings (frequency, latency, properties settings).
SPEC_91	Visualization nodes are intended exclusively for user data visualization. The supplier must not use any Visualization node to provide other functionality.
SPEC_92	(I) In the offer, the supplier must state the hardware platform of the solution and detailed configuration of Visualization nodes including processor model.
SPEC_93	(I) In the offer, the supplier must provide the name and version of the Visualization node operating system.
SPEC_94	(I) In the offer, the supplier must state the name, version, and description of the Visualization node software.

4.12 Requirements – Storages

4.12.1 Requirements - SCRATCH Storage

SPEC_95	SCRATCH storage is a file storage intended exclusively for storing EURO_IT4I users' data.
SPEC_96	SCRATCH storage must be independent of other storages.
SPEC_97	From the user's point of view, SCRATCH file storage must behave as a single, continuous space with a single namespace.
SPEC_98	SCRATCH file storage must be accessible, mounted using a native protocol, on all Compute nodes, Login nodes, Visualization nodes, and Data Management nodes.
SPEC_99	On clients, SCRATCH file storage must be accessible on the <i>/scratch</i> path.
SPEC_100	SCRATCH file storage must have a net available capacity on file-system level of 1PB (1x10 ¹⁵ byte).
SPEC_101	SCRATCH file storage must allow storing of 500 million files.
SPEC_102	SCRATCH storage must provide long-term sustainable sequential read performance for the 1MiB block size of 500GB/s (500x10 ⁹ byte/s). The required performance must be achievable from Compute nodes.
SPEC_103	SCRATCH storage must provide long-term sustainable sequential write performance for the 1MiB block size of 350GB/s (350x10 ⁹ byte/s). The required performance must be achievable from Compute nodes.

- SPEC_104 SCRATCH storage must provide long-term sustainable random I/O performance for the block size of 4KiB and 80%/20% read/write mode of 5 million IOPs. The required performance must be achievable from Compute nodes.
- SPEC_105 SCRATCH storage must be flash based. All data (including metadata) must be stored on SSD or NVMe disks. The disks must be suitable for their designation and expected load.
- SPEC_106 (I) For SCRATCH storage, the supplier must state in the offer the net available capacity; long-term sustainable sequential performance for the 1MiB block size; and long-term sustainable random I/O performance with a block size of 4KiB and 80%/20% read/write mode.
- SPEC_107 (I) For SCRATCH storage, the supplier must provide in the offer the storage architecture, configuration of the equipment, type and number of disks, RAID level, number of disks in RAID, number of spare disks, and software description.

4.12.2 Requirements – HOME Storage

- SPEC_108 HOME storage is a file storage intended exclusively for storing EURO_IT4I users' data.
- SPEC_109 HOME storage must be independent of other storages with an exception of INFRA storage. HOME and INFRA storage can share a storage backend (disk array or similar) but with clear association of individual disks to the specified storage (so one disk cannot be shared by HOME and INFRA storage).
- SPEC_110 From the user's point of view, HOME file storage must behave as a single, continuous space with a single namespace.
- SPEC_111 HOME file storage must be accessible, mounted using a native protocol, on all Compute nodes, Login nodes, Visualization nodes, and Data Management nodes.
- SPEC_112 On clients, HOME file storage must be accessible on the */home* path.
- SPEC_113 HOME file storage must have a net available capacity on file-system level of 25TB (25x10¹²byte).
- SPEC_114 HOME file storage must allow storing of 500 million files.
- SPEC_115 HOME storage must provide long-term sustainable sequential performance for the 1MiB block size of 1.2GB/s (1.2x10⁹ byte/s). The required performance must be achievable from Compute nodes.
- SPEC_116 HOME storage must provide long-term sustainable random I/O performance for the block size of 4KiB and 80%/20% read/write mode of 7,000 IOPs. The required performance must be achievable from Compute nodes.
- SPEC_117 HOME storage must be flash based. All data (including metadata) must be stored on SSD or NVMe disks. The disks must be suitable for their designation and expected load.
- SPEC_118 HOME storage must not use Lustre file system technology, because it is not suitable for given purpose.

SPEC_119 (I) For HOME storage, the supplier must state in the offer the net available capacity; long-term sustainable sequential performance for the 1MiB block size; and long-term sustainable random I/O performance with a block size of 4KiB and 80%/20% read/write mode.

SPEC_120 (I) For HOME storage, the supplier must provide in the offer the storage architecture, configuration of the equipment, type and number of disks, RAID level, number of disks in RAID, number of spare disks and software description.

4.12.3 Requirements – INFRA Storage

SPEC_121 INFRA storage is a file storage intended for storing supercomputer infrastructure services data and contracting authority's data necessary for providing HPC services to users.

SPEC_122 INFRA storage must be independent of other storages with an exception of HOME storage. See SPEC_109 for details.

SPEC_123 The contracting authority assumes the following use of INFRA storage:

- Node installation and boot images;
 - Scheduler data and logs;
 - Application software provided by contracting authority.
- Based on the assumed use the required storage parameters were determined.

SPEC_124 The supplier can use INFRA storage for purposes other than stated in SPEC_123; however, the supplier must adequately increase the storage parameters and ensure compliance with the parameters requested by the contracting authority.

SPEC_125 INFRA file storage must have a net available capacity on file-system level of 25TB (25x10¹²byte).

SPEC_126 INFRA file storage must allow storing of 500 million files.

SPEC_127 INFRA storage must provide long-term sustainable sequential performance for the 1MiB block size of 1.2GB/s (1.2x10⁹ byte/s). The required performance must be achievable from Compute nodes and from Infrastructure nodes independently.

SPEC_128 INFRA storage must provide long-term sustainable random I/O performance for the block size of 4KiB and 80%/20% read/write mode of 7,000 IOPs. The required performance must be achievable from Compute nodes and from Infrastructure nodes.

SPEC_129 INFRA storage must be flash-based. All data (including metadata) must be stored on SSD or NVMe disks. The disks must be suitable for their designation and expected load.

SPEC_130 INFRA file storage must allow partitioning of data capacity into logical partitions/file systems of the required size and making these partitions/file systems available only to selected nodes.

SPEC_131 INFRA storage solution must be available to all EURO_IT4I system's nodes.

SPEC_132 INFRA storage must make available:

- Node installation images – on Infrastructure nodes providing remote installation service or remote network boot service;
- Scheduler data and logs – on Infrastructure nodes running Scheduler;
- Elasticsearch data (in the case INFRA storage is used for this purpose) – on Infrastructure nodes running Elasticsearch;

- Contracting authority's application software – on Compute, Login, and Visualization nodes, on Data management nodes, and on selected Infrastructure nodes.
For each of the above-mentioned nodes, a data path to INFRA storage must exist with throughput of 10Gb/s, at a minimum.

- SPEC_133 INFRA storage must not use Lustre file system technology, because it is not suitable for given purpose.
- SPEC_134 (I) For INFRA storage, the supplier must state in the offer the net available capacity; long-term sustainable sequential performance for the 1MiB block size; and long-term sustainable random I/O performance with a block size of 4KiB and 80%/20% read/write mode.
- SPEC_135 (I) For INFRA storage, the supplier must provide in the offer the storage architecture, configuration of the equipment, type and number of disks, RAID level, number of disks in RAID, number of spare disks, and software description.

Requirements specified in the following sections are common for all file storages – HOME, SCRATCH and INFRA storages.

4.12.4 Requirements – File Storage

- SPEC_136 The file storage must provide shared network file system services.
- SPEC_137 File storage clients must provide standard file system functionalities (in the POSIX way).
- SPEC_138 On the client side, the file storage must be transparently integrated into the operating system, must allow for standard file operations and implement the usual semantics of native file systems, must support a native OS file system interface (API) and integrate operating system users as file system users.
- SPEC_139 The file storage must meet the following requirements:
- Unicode support in file names;
 - Long file names support;
 - Access control, standard Unix permissions (read, write, execute; user, group, others) and extended ACL;
 - User quotas, limits for disk capacity and number of files configurable individually for each user;
 - Group quotas, limits for disk capacity and number of files configurable individually for each group;
 - Reporting of used capacity and number of files for individual users and groups;
 - Support for files larger than 1TB;
 - Symbolic links support;
 - File locking support.
- User and group quotas are not expected to be enabled on INFRA storage file systems.
- SPEC_140 Single, continuous space with a single namespace means that the user uses a single namespace to access the files in the storage, and all the capacity and properties of the file storage are available within that single namespace.
- SPEC_141 Storage services must not negatively affect each other. The requested sequential and random I/O performance must be achievable under simultaneous parallel load on all storages.

4.12.5 Requirements – Availability, Redundancy

SPEC_142 The storage solution must provide high availability. The storage must not contain any component whose outage could cause a failure of the storage services (no “single point of failure” must exist).

Storage solution components, specifically disks, power supply units, disk array controllers, switches, and servers must be redundant and hot-swappable without storage services outage.

A redundant disk array is not considered a “single point of failure” if the disk array only consists of dual-port disks and there are multiple independent data paths for all disks.

SPEC_143 Failure or outage of any single storage server or network device must not cause a failure of storage services.

During a failure or outage of a storage server or network device, the storage performance can be lower than required.

SPEC_144 The storage solution must be designed for long-term heavy load.

SPEC_145 (I) In the offer, the supplier must provide the details about the solution for high availability and redundancy of the supplied systems and equipment.

4.12.6 Requirements – Disk Redundancy

SPEC_146 The storage must provide sufficient redundancy so that a failure of any two disks will not cause data loss.

SPEC_147 The storage must provide recovery from disk failure, i.e. re-ensuring the required data redundancy (e.g. RAID group reconstruction using hot-spare disks). Recovery from disk failure must occur automatically, without operator’s intervention.

SPEC_148 Recovery from disk failure, i.e. re-ensuring the required data redundancy, must be completed within 48 hours of disk failure. During recovery, the storage performance can be temporarily lower than required.

SPEC_149 The storage configuration must ensure required data redundancy after a failure of any two disks of the storage, without operator’s intervention.

SPEC_150 Each disk array or similar equipment of the storage solution must provide a reserve capacity or spare disks in the number or capacity of

$\max\left(\frac{1}{24} * \text{number_of_disks}; 2\right)$ disks, at a minimum, for each type of disk in the disk array.

The $\max(a; b)$ function returns the larger value from a, b . The result is rounded up to the nearest integer.

The *number_of_disks* value is the number of disks of the same type in the disk array. The value includes all disks - the disks with data and parity data, as well as spare disks and reserve capacity.

SPEC_151 (l) In the delivery, the supplier must provide the details about the storage redundancy, number and type of disks, RAID level, number of disks in RAID, number of spare disks, and storage software solution description.

4.12.7 Definition – Storage Capacity

SPEC_152 The data storage capacity (size) is required and must be stated as the net usable capacity, i.e. the storage capacity actually usable by the user at the highest service level provided. The file storage capacity is the capacity of the file system provided by the storage.

SPEC_153 The storage net usable capacity must be stated for offered/delivered configuration designed for normal operation.

SPEC_154 All required/offered capacities (including the metadata capacity required for storing the requested number of files) must be achievable simultaneously.

SPEC_155 Determination of the net usable capacity must not take into account the features of the system or its component, which could provide larger storing capacity under conditions that cannot be ensured (compression, deduplication, etc.).

SPEC_156 Determination of the net usable capacity must not take into account the features of the system or its component, which could allow allocating more space than physically possible or practically feasible without further action (oversubscription).

SPEC_157 Storage capacity is specified using prefixes of decimal multiples.
Gigabyte (GB) 10^9 byte
Terabyte (TB) 10^{12} byte
Petabyte (PB) 10^{15} byte

SPEC_158 The net available storage capacity is determined on a suitable storage client system by writing to the storage until it is full, or reading the entire storage, or by a suitable system tool demonstrating the size/storage capacity.

SPEC_159 The tools used to determine the capacity must provide reliable information and work with a known data block size or a known and accurate unit.
In Linux OS, the file system capacity in bytes can be displayed using the `df -B1` command, value "Available".

4.12.8 Definition – Storage Performance

SPEC_160 The storage performance (sequential performance, I/O performance, metadata performance) is required and must be stated as the long-term performance realistically sustainable by the user from storage clients at the highest service level provided.

SPEC_161 The file storage performance is the long-term, realistically sustainable performance of operations performed on file storage systems from file storage clients.

SPEC_162 The storage performance must be stated for the offered/delivered configuration designed for normal operation.

SPEC_163 The storage performance must be determined for highly occupied storage (see the benchmark specifications for details).

- SPEC_164 The storage performance must be determined for non-privileged user operations.
- SPEC_165 The performance must not be determined based on the assumption of specific, favorable conditions or specific, advantageous measurement modes (e.g. cache operation), unless such conditions or modes are explicitly required.
- SPEC_166 The storage performance is specified using prefixes of decimal multiples.

4.12.9 Performance Measurement

- SPEC_167 The supplier must demonstrate compliance with the storage performance requirements by running performance tests (benchmarks) within the acceptance tests (performance measurements).
- SPEC_168 The performance tests must be run on test servers configured as file storage clients.
- SPEC_169 The test server disk caches must be cleared before each storage performance measurement.
- SPEC_170 The performance measurement must be carried out by procedures and under conditions that correspond to normal operation – normal provision of storage services. No action must be taken prior to and during measurement of the storage performance that would affect the measurement result.
- SPEC_171 (I) In the delivery, the supplier must state the proposal of the performance measurement implementation.

4.12.10 Using FIO Benchmark for Performance Measurement

- SPEC_172 The file storage sequential performance and random I/O performance must be measured with fio version 3.13 in client-server mode. Other/newer version of the application can be used only with the approval of the contracting authority.
Fio is an open source tool (GPL license version 2) for I/O benchmarking and testing, available at <https://github.com/axboe/fio>.

Accessibility of the measured file storage is verified on each server dedicated to file storage measurement (test servers according to SPEC_168) and the fio tool in server mode is run:
`fio --server`

The measurement is initiated from the selected server – the fio tool is run in client mode:
`time fio \
 --client=machinefile jobfile \
 --output-format=normal,json |& \
tee fio.out`

In the *machinefile* file, the names of all test servers for storage measurement are specified, one name per line.

Jobfile – the fio tool configuration file – contains the test description according to the specified measurement. The *jobfile* file is stored on the server initiating the measurement.

Fio program must be executed by an unprivileged user.

The file storage directory used to store test measurement files (the `directory` setting in the *jobfile*) must be empty before the measurement.

The test files must be prepared in the directory before the sequential read performance measurement and before random I/O performance measurement. To create the test files, run the `fiio` tool in the client mode and set the `create_only=1` parameter in the respective *jobfile*. By modifying the `bs` parameter, you can change the block size, as well.

Test directories and files must be owned by an unprivileged user.

The `fiio` output on the server that initiated the measurement (i.e. the *fiio.out* file) will be used to determine the result. The measurement summary information provided under “All clients” will be used – in JSON output format.

In JSON output, the last item of the `client_stats` attribute is used, the value of the `jobname` attribute of this item is `All clients`, this item is hereinafter referred to as the summary.

The value of the `bw_bytes` attribute of the `read` or `write` (depending on the operation type) attribute of the summary represents the respective sequential performance (throughput) in bytes per second.

To measure the random I/O performance, the sum of the values of the `iops` attribute of the `read` and `write` attributes of the summary represents the total number of I/O operations per second (IOPS).

For SCRATCH storage performance measurement, multiple simultaneously running instances of `fiio` benchmark can be used, in this case the same timing and the same runtime of instances and accurate calculation must be guaranteed.

4.12.11 Sequential Performance Measurement

SPEC_173 The long-term sustainable sequential performance must be measured with the `fiio` tool according to SPEC_172.

The storage sequential performance is the lower of the two values – file storage sequential read performance and file storage sequential write performance.

For file storage sequential write performance, use the following configuration file – *fiio jobfile*:

```
[global]
rw=write
bs=1M
create_on_open=1
time_based
runtime=1h
numjobs=NUMJOBS

[name_of_file_storage]
directory=/path_to_filesystem/test/fio
filesize=FILESIZE
```

For file storage sequential read performance, use the following configuration file – *fiio jobfile*:

```
[global]
rw=read
bs=1M
time_based
runtime=1h
```

```
numjobs=NUMJOBS
;create_only=1

[name_of_file_storage]
directory=/path_to_filesystem/test/fio
filesize=FILESIZE
```

Replace the NUMJOBS and FILESIZE parameters with suitable values.

The NUMJOBS parameter indicates the number of parallel tasks performed on the test server; the value is an integer.

The FILESIZE parameter indicates the size of individual files that the job is working with; the value is the size according to the fio tool syntax (e.g. 30T).

Total size of the files with which the measurement works from one test server, i.e. the FILESIZE * NUMJOBS value, must be greater than 20 times the memory size of each test server.

Total size of the files with which the measurement works from all used test servers, i.e. the FILESIZE * NUMJOBS * number_of_test_servers value must be greater than

- for SCRATCH storage: 900TB;
- for INFRA storage: 20TB;
- for HOME storage: 20TB.

4.12.12 Random I/O Performance Measurement

SPEC_174 The long-term sustainable random I/O performance must be measured with the fio tool according to SPEC_172.

Use the following configuration file – fio *jobfile*.

```
[global]
rw=randrw
rwmixread=80
bs=4k
time_based
runtime=1h
numjobs=NUMJOBS
;create_only=1

[name_of_file_storage]
directory=/path_to_filesystem/test/fio
filesize=FILESIZE
```

Replace the NUMJOBS and FILESIZE parameters with suitable values.

The NUMJOBS parameter indicates the number of parallel tasks performed on the test server; the value is an integer.

The FILESIZE parameter indicates the size of individual files that the job is working with; the value is the size according to the fio tool syntax (e.g. 30T).

Total size of the files with which the measurement works from one test server, i.e. the FILESIZE * NUMJOBS value, must be greater than 20 times the memory size of each test server.

Total size of the files with which the measurement works from all used test servers, i.e. the FILESIZE * NUMJOBS * number_of_test_servers value must be greater than

- for SCRATCH storage: 900TB;
- for INFRA storage: 20TB;
- for HOME storage: 20TB.

4.12.13 Metadata Performance Measurement

SPEC_175 In order to verify the stability of the file storage metadata operations and to determine the file storage metadata performance, the long-term sustainable performance of the file storage metadata operations must be measured as acceptance tests.

SPEC_176 The long-term sustainable metadata performance must be measured using the mdtest tool version 3.2.1 <https://github.com/hpc/ior>. Other/newer version of the application can be used only with the approval of the contracting authority.

For measurement, use the following command:

```
time \  
mpirun -n $NPROC -machinefile \  
mdtest \  
-C -T -r \  
-F \  
-d /path_to_filesystem/test/mdtest \  
-I $FILES_PER_DIR \  
-i $ITERATIONS \  
-u \  
-z $TREE_DEPTH -b $BRANCHING_FACTOR \  
-L
```

where *machinefile* is the file containing the names of the test servers (as specified in SPEC_167), one name per line,

the NPROC, FILES_PER_DIR ITERATIONS, TREE_DEPTH, BRANCHING_FACTOR variables are positive integers determined so that:

- The FILES_PER_DIR value is greater than or equal to 100;
- The number of files created in each iteration is 10^8 , at a minimum, i.e. $NPROC * FILES_PER_DIR * BRANCHING_FACTOR^{TREE_DEPTH} \geq 10^8$;
- The test run time is 1 hour, at a minimum.

The result is the values at the “File creation” and “File stat” lines, in the “Mean” column, of the “SUMMARY rate” table.

4.13 Requirements – Data Management Nodes

SPEC_177 EURO_IT4I system must include, at a minimum, two servers operating exclusively as Data management nodes – nodes designed for data transfer between supercomputer’s storages and external storages.

SPEC_178 Each Data management node must meet the following requirements:

- Physical node
- x86-64 architecture
- At least one processor, at least 16 CPU cores total;
- At least 128GiB RAM, operated in DDR4 mode, at least 2666MT/s with ECC;
- The RAM must consist of memory modules of the same parameters, evenly distributed across all server memory;
- Theoretical throughput of processor(s) to RAM at least 160GB/s (in the offered server configuration);
- At least 2 local SSD disks with the capacity of 240GB, at a minimum, in RAID1;
- Hot-swap disks;
- Redundant, hot-swap power supply units, redundant power supplies.
- Compute network connection 100Gb/s, at a minimum

- LAN connection, redundant, see section 4.17 Requirements – LAN Infrastructure for details

- SPEC_179 On Data management nodes, HOME and SCRATCH file storages must be accessible, mounted via native protocol.
- SPEC_180 The Data management nodes must provide protocols for SFTP and SCP file transfer and ensure transfer from/to HOME and SCRATCH. On Data management nodes, the contracting authority will install and operate GridFTP and other services.
- SPEC_181 Data management nodes must be integrated into EURO_IT4I system. Specifically, this includes user account integration, authentication, and integrating and accessing Data management nodes services via Scheduler. Data management nodes must be clients of the Scheduler.
- SPEC_182 (I) In the offer, the supplier must provide the description and detailed configuration of Data management nodes including processor model.
- SPEC_183 (I) In the offer, the supplier must provide the name and version of the Data management node operating system.

4.14 Requirements – PROJECT Network Gateways

PROJECT data storage is the contracting authority storage designed for storing multi-year user project data and sharing of the data between IT4Innovations supercomputers.

PROJECT data storage is not yet implemented and the contracting authority does not know the exact equipment specifications.

PROJECT Network Gateways are servers intended for connecting EURO_IT4I system Compute Network with PROJECT data storage Access Network (100Gb/s Ethernet network).

- SPEC_184 The delivery must include, at a minimum, four PROJECT Network Gateways.
- SPEC_185 Each PROJECT Network Gateway server must meet the following requirements:
- Physical server;
 - x86-64 architecture;
 - Linux 64-bit, CENTOS 7 (or newer) or RHEL 7 (or newer) operating system;
 - At least one processor, at least 3GHz base frequency, at least 16 CPU cores total;
 - At least 32GiB RAM, operated in DDR4 mode, at least 2900MT/s with ECC;
 - The RAM must consist of memory modules of the same parameters, evenly distributed across all server memory;
 - Theoretical throughput of processor(s) to RAM at least 160GB/s (in the offered server configuration);
 - At least 2 local SSD disks with the capacity of 120GB, at a minimum, in RAID1;
 - Hot-swap disks;
 - Redundant, hot-swap power supply units, redundant power supplies.
- SPEC_186 Each PROJECT Network Gateway node must provide a connection to EURO_IT4I system Compute Network with aggregated throughput of at least 100Gb/s.
- SPEC_187 Each PROJECT Network Gateway node must provide a connection to PROJECT Access Network with aggregated throughput of at least 100Gb/s. For this purpose, the contracting authority reserves four 100Gb/s Ethernet QSFP28 type ports in edge devices of PROJECT data storage Access Network (two ports in each edge device). A part of the delivery are modules and optical cables required for

connection to the edge devices of PROJECT data storage Access Network. Network cabling in the Data center dropped ceiling is required.

Location of PROJECT data storage and cable routes are described in section 8.6 WAN/LAN Network.

SPEC_188 The supplier must cooperate with the contracting authority in:

- Provisioning of Compute network adapter software stack on PROJECT Network Gateways;
- Provisioning of interconnecting, addressing, and routing Compute Network and PROJECT Access Network;
- Provisioning of NFS mounpoints of PROJECT file storages on the Compute nodes, Login nodes, Visualization nodes, and Data management nodes.

SPEC_189 (I) In the offer, the supplier must state count and detailed configuration of PROJECT Network Gateways.

4.15 Requirements – Infrastructure Nodes

SPEC_190 Infrastructure nodes and their infrastructure must be designed so that they ensure reliable, secure, fast, and effective operation and administration of EURO_IT4I system.

SPEC_191 Virtualization technology can be used for implementation of infrastructure for Infrastructure nodes; therefore, Infrastructure nodes can be realized by virtual servers. Nevertheless, the virtualization must not negatively affect the solution.

SPEC_192 EURO_IT4I system must include at least six physical Infrastructure nodes or at least three physical host servers of virtualization infrastructure for Infrastructure nodes. These servers must provide at least 384GiB of RAM memory in total.

SPEC_193 Each physical server referred in SPEC_192 must meet the following requirements:

- Disks in RAID with data redundancy
- Hot swap disks
- Redundant, hot swap power supply units, redundant power supplies
- LAN connection, see section 4.17 Requirements – LAN Infrastructure for details

SPEC_194 Failure or outage of any single infrastructure node must not interrupt currently running computing jobs and operation of Compute clusters, Login nodes, Visualization nodes, storages, Data management nodes, backup, network, or supercomputer administration.

SPEC_195 The key services must be run in high-availability mode, preferably using the native mechanisms of the services.

SPEC_196 (I) In the offer, the supplier must state the number of infrastructure nodes, their purpose, and detailed configuration.

4.16 Requirements – Backup Solution

SPEC_197 A part of the delivery must be a comprehensive solution for data backup.

Redundancy is not considered a backup. EURO_IT4I system must contain functionality and infrastructure, which enable creation of logically and hardware independent copies of production data on backup storage.

- SPEC_198 The Backup solution must particularly provide the following:
- Backup of all delivered servers/nodes, with an exception of Compute nodes provisioned from images
 - Backup of image repository (images) for nodes provisioning (it may be a part of backup of INFRA storage)
 - Backup of HOME storage; SCRATCH storage is not subject to backup.
 - Backup of INFRA storage

For HOME storage, 80% space usage of total capacity and large amount of small files is assumed. It is assumed that 4% of HOME storage total capacity will change daily.

For INFRA storage, 80% space usage of total capacity and large amount of small files is assumed. It is assumed that 4% of INFRA storage total capacity will change daily.

It is not necessary to back up temporary server data, which is not required for server operation recovery.

- SPEC_199 Data must be backed up every 24 hours between 10:00 p.m. and 6:00 a.m. (i.e. during an eight-hour period, at a maximum).

- SPEC_200 The backup solution must provide data recovery from the last 21 daily backups.

- SPEC_201 The backup solution must provide sufficient backup capacity to meet all the requirements (with regard to the technology used and realistically achievable parameters in the given application).

- SPEC_202 Data backup and recovery priorities (from the highest to lowest priority):
1. Highest priority: Nodes critical for providing system services, INFRA storage
 2. Other nodes
 3. Lowest priority: HOME storage

- SPEC_203 The data backup and restore solution must meet the following basic properties:
- Time schedule of backups
 - Individual files and directories backup and restore
 - Files and directories' owners, permissions, extended ACLs, and attributes backup and restore
 - Parallel run of backups and recoveries

- SPEC_204 The backup solution must use a backup-to-disk technology. It is assumed that a technology for effective data storage will be used (e.g. compression or de-duplication).

- SPEC_205 Backup must use a dedicated, independent data storage. Backup data storage must meet the requirements in sections 4.12.5 Requirements – Availability, Redundancy and 4.12.6 Requirements – Disk Redundancy.

- SPEC_206 The backup solution must have a minimal negative impact on EURO_IT4I system operation and performance.

- SPEC_207 The backup system license must cover all the requirements in maximum scope.
- SPEC_208 (I) In the offer, the supplier must state the backup solution capacity and throughput which must be sufficient to meet the backup requirements.
- SPEC_209 (I) In the offer, the supplier must describe in detail the calculation of required capacity to meet the contracting authority's backup requirements.
- SPEC_210 (I) In the offer, the supplier must state solution architecture, configuration of equipment, description of backup software solution, its properties, and suggested backup policy.

4.17 Requirements – LAN Infrastructure

4.17.1 LAN - General

- SPEC_211 The LAN must provide communication between the equipment in the supercomputer and connection to the contracting authority's WAN.
- SPEC_212 The LAN must include the public and private parts of the network.
- SPEC_213 The public parts of the network (hereinafter referred to as public networks) must use IPv4 and IPv6 public addresses and must provide services accessible from the Internet (Login nodes, Visualization nodes, Data management nodes).
- SPEC_214 The private parts of the network (hereinafter referred to as private networks) must use IPv4 private addresses and be used for internal services and device management.
- SPEC_215 Devices in private networks must have access to the Internet over NAT.
NAT will be provided by a contracting authority's network device located in the contracting authority's WAN.
- SPEC_216 The network solution must support IPv6.
- SPEC_217 The LAN must be divided into several L3 networks. For each L3 network, a different L2 network must be used (VLAN or another network device). It is not desirable to use one L2 network for multiple L3 networks. LAN division must separate the following types of communication into different L3 networks:
- Services accessible from the Internet
 - Data communication between nodes (services)
 - Infrastructure nodes and services
 - Network devices management
 - Disk arrays, storage management
 - Server (BMC, IPMI, etc.) management
 - Non-IT infrastructure (power supply, cooling, etc.) management

4.17.2 LAN – Edge devices

- SPEC_218 Connection of the LAN with the contracting authority's WAN must be implemented by redundant L3 switches (hereinafter referred to as Edge devices, labeled as ED1, ED2 in the figure 2). Edge devices must meet the following requirements:
- Dual power supply
 - Redundant server connection using multi-chassis ether channel
 - Edge devices interconnection throughput of 800Gb/s, at a minimum
 - Separate routing instances for private and public networks
 - Support Access Control Lists (ACL) for at least 2000 ingress and 2000 egress ACL rules configurable for each port independently.
- SPEC_219 In the case of separate control-plane of Edge devices, the edge devices must allow usage of HSRP, VRRP or similar router redundancy technology, where keepalive communication is done over IP addresses of subnet (private) different from floating IP address from public IP range, therefore not wasting public IP address for keepalive communication.
- SPEC_220 Edge devices must support user authentication using RADIUS or TACACS+ protocol, creating of different user roles for network management (operator, administrator, etc.), and issued commands logging.
- SPEC_221 Edge devices must support import and export of its device configuration using TFTP, FTP, SCP, or SFTP protocol. Configuration must allow saving in the form, which allows its editing using text editor.

SPEC_222 Edge devices must support IPFIX (Internet Protocol Flow Information eXport), or NetFlow v9, or sFlow v5. Data flow export must provide information about all IP packets (NetFlow standard, full NetFlow), or at least about one packet from 10 processed IP packets (sFlow, sampled NetFlow standards, sampling rate 1:10).

4.17.3 LAN – connections

SPEC_223 Minimum LAN connection speed of devices:
Minimum connection speed of each Login node, Visualization node and Data management node is 100Gb/s.
Minimum LAN connection speed of each Infrastructure node is 10Gb/s.
In the case of direct connection of Compute nodes to LAN, minimum speed of each node must be 1Gb/s.

SPEC_224 Minimum speed specified in SPEC_223 must be reachable when accessed from the contracting authority's WAN. For systems specified in SPEC_225, the minimum speed requirement must be met even in the case when any single active network device fails.

SPEC_225 Connection of Login nodes, Visualization nodes, Data management nodes, and Infrastructure nodes to LAN must be redundant.

SPEC_226 Aggregated connection speed of Login nodes, Visualization nodes, and Data managements nodes in LAN reachable from the contracting authority's WAN must be 400Gb/s.

SPEC_227 Aggregated connection speed of Compute nodes in LAN reachable from the contracting authority's WAN must be at least 10Gb/s.

SPEC_228 In case there is no direct connection of Compute nodes to LAN, the solution must include a redundant LAN gateway solution to ensure IP connectivity for the devices in Compute network. LAN gateway solution must provide connection 2x100Gb/s to Compute network.

4.17.4 LAN – Active network devices

These requirements are common for all active network devices including Edge devices.

SPEC_229 Active network devices must support VLAN segmentation using 802.1Q standard with at least 100 VLANs and numbering from 1 to 4094.

SPEC_230 Active network devices must support IGMPv2 a IGMPv3 protocols.

SPEC_231 Active network devices, which are L3 devices, must support IPv4 multicast routing and PIM Sparse Mode and PIM Source-Specific Multicast protocols.

SPEC_232 Active network devices must be remotely controllable and centrally manageable.

SPEC_233 Active network devices must support SNMPv2 and SNMPv3 protocols enabling:
1. Reading of device state and ports utilization
2. Sending of SNMP traps for defined events

SPEC_234 Active network devices must support remote access using SSH2 protocol with at least AES256-CTR encryption protocol, HMAC-SHA256 hashing protocol and RSA key of 4096 bits.

- SPEC_235 Management interfaces of active network devices must be connected to the contracting authority's OOB network.
Serial management interfaces of edge devices must be connected to the contracting authority's OOB router (Cisco 2901/K9). The supplier must provide required cables and installation.
Ethernet management interfaces of all active network devices and of nodes described in SPEC_279 must be connected to the contracting authority's OOB network implemented by the contracting authority's OOB switch. The supplier may utilize, at a maximum, one GBase-T Ethernet port of the contracting authority's OOB switch. The supplier must provide required cables and cabling installation. Cable routes are described in section 8.6 WAN/LAN Network.
- SPEC_236 (I) In the offer, the supplier must provide description of the Edge devices.

4.18 Requirements – WAN Integration

- SPEC_237 Connection to the contracting authority's WAN/LAN must provide connectivity of 4x100Gb/s.
- SPEC_238 Each LAN edge device of EURO_IT4I system must be connected to each of the two central devices of the contracting authority's WAN/LAN using a 100Gb/s connection. For this purpose, the contracting authority reserves four ports in WAN/LAN central devices WAN1 and WAN2. The ports are equipped with QSFP-100G-SR4 modules with MPO-12 connectors. The delivery must include the modules and fiber-optic cables required for connection of edge devices to the WAN1 and WAN2 devices. It is required to install network cabling in the data center suspended ceiling.

The WAN/LAN infrastructure of the contracting authority and cable routes are described in Chapter 8.6 WAN/LAN Network.
- SPEC_239 Twelve Cat6a connections must be provided between the location of LAN edge devices and the contracting authority's WAN/LAN described in Chapter 8.6 WAN/LAN Network. The supplier must provide Cat6a patch panel with RJ-45 ports at the location of LAN edge devices and connect it with existing patch panel in the contracting authority's WAN/LAN rack. Connections may be used for OOB connections mentioned in SPEC_235.
- SPEC_240 Connection of edge devices and the contracting authority's WAN devices must be realized over Ethernet and private IPv4 in separate routing instances.
- SPEC_241 Connection to the contracting authority's WAN devices must be redundant.
- SPEC_242 (I) In the offer, the supplier must provide description of the WAN integration.

4.19 General Technical Requirements

- SPEC_243 All servers (with the exception of Compute nodes) and disk arrays (including disk modules of disk arrays) must use:
- Redundant, hot-swap power supply units
 - Redundant power supplies
 - Hot-swap disks
- SPEC_244 All devices and systems must provide remote management.

- SPEC_245 All servers (nodes) must allow replacement of CPUs, RAM memory modules, and Compute network adapters (if present) independently of other components.
- SPEC_246 All servers must have remote network management, independent of OS operation, providing power supply control, reset option, graphical console and virtual media.
- SPEC_247 Disk array controllers must be redundant. Failure or outage of any single disk array controller must not cause a failure of disk array services. Disk array controller cache must be protected against data loss or corruption in case of power outage or any single controller failure.
- SPEC_248 All disk array controllers and server RAM modules must use an error detection and correction mechanism – Error-correcting code memory (ECC memory).
- SPEC_249 The parameters and characteristics of the disks used in the solution must be appropriate to their deployment and load. The warranty also applies to disk wear (including SSDs and NVMe drives) due to excessive use.
- SPEC_250 Normal operation and accessibility of declared capacities must not require operator's intervention.
- SPEC_251 In the case of SSH, SCP, and/or SFTP protocol support, protocol version 2 must be supported.
- SPEC_252 All devices and cables must be physically labeled with unique identification (labels). The identification must be easily accessible, legible and properly recorded.

4.20 Requirements – Software

- SPEC_253 A part of EURO_IT4I system must be a comprehensive HPC system software solution, i.e. all software equipment required to fulfill the contracting authority's requirements (firmware, operating systems, drivers, management, and monitoring software, file systems, scheduler, etc.).

All delivered software must be licensed for unlimited time, with no additional fees (at the version level applicable at the time of delivery).

4.20.1 Software - Environment

- SPEC_254 EURO_IT4I system must deliver a transparent, uniform, shared user environment to end users.
- SPEC_255 EURO_IT4I system must provide, to the maximum extent possible, uniform environment on all nodes accessible to end users, i.e. Login nodes, Compute nodes and Visualization nodes. User environment on nodes of the same type must be identical. User environment on nodes of different types must be identical to the maximum extent possible; environment differences arise only from differences in the nature and services of servers. Environment means especially behavior of systems, applications, names of files, directories, commands, environment variables, values of environment variables, configuration, etc. Nodes must provide en_US and C locale.
- SPEC_256 Login nodes must provide end users with a command-line and graphical interface. Login nodes must provide shell (bash and csh) and X Window System (X11).

- SPEC_257 EURO_IT4I system must provide end users with unified account and unified authentication. End user accounts must be accessible and identical on all Login, Compute, and Visualization nodes and Data management nodes.
- SPEC_258 User account implementation must use LDAP and SSSD technology. User accounts must be stored and managed in the LDAP database. User accounts and groups must be created via posixAccount and posixGroup LDAP schemas. User accounts and groups must be managed via a secure LDAP protocol. The contracting authority will manage user accounts and groups exclusively. The contracting authority provides centralized management of all IT4Innovations center users; the contracting authority will synchronize EURO_IT4I system users and groups information from their central repository to EURO_IT4I system LDAP database. User home directories will also be managed exclusively by the contracting authority.
EURO_IT4I system LDAP service must be implemented using two LDAP servers with a replicated database. LDAP service must be implemented on suitable infrastructure nodes. Reservation of these nodes for running the LDAP service is not required. For the LDAP service, the OpenLDAP software is preferred.
SSSD service must be used for name/uid, group/gid resolving on nodes providing user access. Secure LDAP protocol must be used for accessing LDAP service. Login nodes, Visualization nodes, and Data management nodes must use SSSD enumeration.
- SPEC_259 The unified user authentication must use SSH keys. The SSH keys will be managed by the contracting authority or user and must be stored in the LDAP database or user's home directory. The SSH service must allow authentication of authorized user's SSH keys against user's entry in the LDAP database. It is assumed that the solution will be implemented using a configuration of SSHD and SSSD features. The contracting authority expects that there will be no end user authentication requiring password authentication.

4.20.2 Software - Scheduler

- SPEC_260 EURO_IT4I system must provide unified access to computing resources. For access to computing resources and for performing computational jobs on Compute nodes, Visualization nodes, and Data management nodes, EURO_IT4I system must use an advanced job scheduler and resource manager (hereinafter referred to as Scheduler).
- SPEC_261 Scheduler must effectively utilize available computing resources, taking into account specific features of Compute nodes and Compute network. Scheduler must allow advanced scheduling policies. Scheduler must allow executing of batch and interactive jobs. Scheduler must support job priorities and job dependencies. Scheduler must take into account current resources availability. Scheduler must handle compute resource failures and outages. Scheduler must allow scheduling according to resources resolved statically or dynamically.
- SPEC_262 Scheduler must allow checking status of computational jobs. Scheduler must allow event notification (on job abort/begin/end execution) by email. Scheduler must allow users to submit jobs and monitor their status. Scheduler must return user jobs output back to submit location. Scheduler must allow operators to display jobs status, scheduler's server and queues status and details of individual jobs.
- SPEC_263 Scheduler must use high availability cluster, at least in the active-passive mode.
- SPEC_264 Scheduler must provide information about users' computational jobs resource consumption for accounting purposes. Required information is especially job runtime, names of allocated Compute nodes and number of allocated compute cores.

- SPEC_265 Scheduler must allow execution of custom scripts defined by the contracting authority before and after job execution (prologue, epilogue), when submitting jobs to Scheduler and while modifying jobs (hooks).
- SPEC_266 Scheduler must allow creating allocation of a set of nodes of specified properties available in a specified period to a specified group of users (reservation).
- SPEC_267 Scheduler must allow limiting the number of concurrent user jobs, the number of concurrent user group jobs, the maximum duration of a job, and access to computational resources only to authorized users (ACL).
- SPEC_268 Scheduler must take into account the Compute network topology to allow better parallel job placement on Compute nodes with higher communication efficiency.
- SPEC_269 Scheduler must provide command-line interface and API in the C language.
- SPEC_270 Scheduler must be PBS compatible (reasons being: integration into the contracting authority's current systems, their knowledge and operation of their own integration, visualization, and management tools, and minimization of the number of operated software platforms).
- SPEC_271 The server part of the Scheduler must be implemented on suitable Infrastructure nodes; at least 3GHz base frequency CPU cores are required. Reservation of these nodes for running the service is not required.
For implementation of the solution, INFRA storage can be used.
- SPEC_272 Scheduler license must cover all Compute nodes (execution), Login nodes (submit) Visualization nodes (execution and submit) and Data management nodes (execution).
- SPEC_273 The supplier must provide basic installation and configuration of Scheduler to demonstrate its functionality and requirements fulfillment. The contracting authority will further customize the configuration to meet his needs.

4.20.3 Software - Management

- SPEC_274 EURO_IT4I system must include tools for management of all systems and services, which are a part of the system.
- SPEC_275 EURO_IT4I system must include tools for remote management of all hardware devices (servers/nodes, disk arrays, network switches, etc.) particularly allowing for configuration and management of the devices, critical event detection and notification of such events via SMTP (email) or SNMP.
- SPEC_276 EURO_IT4I system must provide a command-line tool for remote power and reset control of all servers. Functionality must be independent of the operating system of the controlled servers (it is assumed the IPMI protocol will be used).
Functionality must be accessible on suitable Infrastructure nodes (nodes intended for management).
- SPEC_277 EURO_IT4I system must provide a command-line tool for remote display and manipulation of console terminals of all servers. Functionality must be independent of the operating system of the controlled servers (it is assumed the IPMI protocol will be used).

Functionality must be accessible on suitable Infrastructure nodes (nodes intended for management).

- SPEC_278 EURO_IT4I system must provide a command-line tool for effective centralized remote management of node operating systems individually (single node) and in bulk (groups of nodes or all nodes). The system must provide, in particular, remote command execution; file transfer, modification and deletion; and comparison of command outputs and command return values. The system must allow parallel action execution. Remote management must also work effectively in the case some of the managed nodes are not communicating or operating correctly, and must identify and report action failures. The supplier must provide suitable server naming and server group definitions for effective use of the centralized remote management system. Functionality must be accessible on suitable Infrastructure nodes (nodes intended for management). ClusterShell software must be used for the remote management system.
- SPEC_279 The functionalities described in SPEC_276 to SPEC_278, and tools for Scheduler management must be collectively available on at least two Infrastructure nodes (nodes intended for management).
- SPEC_280 EURO_IT4I system must provide effective remote server installation from the central image repository. The system must allow individual (single node) or bulk (multiple nodes) installation. The system must allow parallel server installation. The server part of the remote server installation must be operated on suitable Infrastructure nodes. Reservation of these nodes for running the service is not required.
- SPEC_281 The remote server installation according to SPEC_280 must be used for Compute nodes installation (with exception of Universal nodes, see SPEC_282) and must allow (but is not required for) installation of Login nodes, Visualization nodes and Data management nodes. In case the remote server installation is not used for installation of Login nodes, Visualization nodes, and/or Data management nodes, another suitable solution must be used which provides identical installation and configuration of the nodes (even for all subsequent installation).
- SPEC_282 EURO_IT4I system must provide network boot service from the central image repository. The system must allow individual (single node) or bulk (multiple nodes) network boot. The system must allow parallel server boot. The server part of the remote server installation must be operated on suitable Infrastructure nodes. Reservation of these nodes for running the service is not required.
- SPEC_283 The network boot according to SPEC_282 must be used for Universal nodes. The network boot must be available and can be used for Accelerated nodes and Cloud nodes. Complete startup of all nodes in every compute partition using network boot must finish in 15 minutes.
- SPEC_284 System must provide management of images in the central image repository, in particular, creation, modification, and deletion of images. Servers of identical type must use an identical installation image; the number of images will be limited. The central image repository must have the total capacity of 1TB, at a minimum. Central image repository can be implemented on INFRA storage.

SPEC_285 EURO_IT4I system subsystems must be interconnected and integrated to provide effective operation and management of EURO_IT4I system without unnecessary obstructions and non-productive activities.

4.20.4 Software – OS and Applications

SPEC_286 All servers/nodes must use Red Hat Enterprise Linux or CENTOS operating system, major version 7 or newer. Servers used solely for implementation of Backup solution can use Microsoft Windows Server operating system, version 2016 or newer.

SPEC_287 All Universal nodes, Accelerated nodes, Cloud nodes, Login nodes, Visualization nodes, and Data management nodes must use the same distribution and the same version of the operating system.
Data analytics node can use different distribution of operating system than other Compute nodes (in accordance with SPEC_286); in such a case, system libraries of these distributions must have the same versions.

SPEC_288 The operating system license must cover all nodes/servers in the delivery.

SPEC_289 Delivered operating systems and software must provide updates, in particular security, reliability, functionality, and performance updates.
The supplier must provide and install software updates according to conditions specified in the contract.

SPEC_290 The delivery must contain the software necessary for acceptance tests execution (High Performance LINPACK, etc.) and for possible subsequent repetition of the tests during the warranty period, with the exception of the software provided by the contracting authority. The contracting authority will provide Intel compilers and Intel MPI/MKL libraries in the latest stable version.

The contracting authority provides comprehensive management of application software, its acquisition, assembly/building, installation, and configuration.

The contracting authority further prefers the following software:

- For management, configuration and automation - configuration management software Puppet or Ansible
- For file versioning – GIT

4.20.5 Software – Addressing, Name Services

SPEC_291 Address prefixes/ranges, domain name suffixes, namespaces, user accounts, and groups are managed and assigned solely by the contracting authority.

SPEC_292 IP addressing must be in accordance with contracting authority's address policy and plan. Unless already specified in another section of the documentation, used IP address ranges must be specified after consultation with the contracting authority.

SPEC_293 IP addressing can be static or dynamic (using DHCP).

SPEC_294 All used IP addresses must resolve to names. IP addresses must be resolved by DNS servers provided by the supplier, used domain names suffixes must be specified after consultation with the contracting authority.

- SPEC_295 The contracting authority requires an option to change configuration of DNS services provided by the supplier and option to integrate contracting authority's own DNS services (views) to EURO_IT4I system.
- SPEC_296 EURO_IT4I system equipment and systems must use accurate time. Time synchronization must be done using the contracting authority's NTP servers in the contracting authority's network.

4.20.6 Software – Monitoring

- SPEC_297 EURO_IT4I system must include monitoring of accessibility and status of components and services (hereinafter referred to as availability monitoring). Availability monitoring must provide information about accessibility of all delivered equipment available in network over IP and information about accessibility/status of all relevant components/services of supplied nodes/servers, storages, etc. Relevant components/services are those that affect EURO_IT4I system services' functionality, accessibility, and/or performance. Availability monitoring must classify accessibility/state based on severity (OK, Warning, and Critical). Availability monitoring must allow configuring thresholds of monitored parameters for severity classification. Availability monitoring must provide identification of the equipment, component, or service to which availability or status information relates. Availability monitoring must provide timestamp of objects' accessibility/state change and record them in a log. Icinga2 software in clustered, highly available configuration must be used for implementation of availability monitoring. Icinga2 software must be operated on suitable Infrastructure nodes. Reservation of this node for running the service is not required. Supplier must provide installation and configuration of the Icinga2 software. The contracting authority will complement the solution with their own set of checks and integrate the Icinga2 instance to their Icinga2 cluster. Infrastructure nodes must provide resources for operation of at least two Icinga2 instances.
- SPEC_298 All important states and inaccessibility of EURO_IT4I system components and services must be reported to a single availability monitoring, according to SPEC_297.
- SPEC_299 Availability monitoring as specified in SPEC_297 must particularly provide the following information:
1. Accessibility of device network interfaces in the IP networks (including accessibility of management interface, e.g. BMC interface) – for all equipment with network interface
 2. Equipment power supply status – for all servers/nodes, disk arrays, and network devices
 3. Equipment cooling status according to the temperature of the equipment and functionality of cooling fans/pumps – for all servers/nodes, disk arrays, and network devices
 4. Status of the equipment according to its indicated status and/or the record in the equipment's log – for all servers/nodes, disk arrays, and network devices
 5. Servers/node RAM status (indicates memory error occurrence)
 6. Server/node CPU load status calculated for one CPU core of the server
 7. Server/node memory usage status
 8. Server/node status according to server/node file system capacity used
 9. Server/node status according to the total number of server/node processes
 10. Server/node status according to the number of server/node zombie processes

11. Server/node shared file storage (INFRA, SCRATCH, HOME, PROJECT) client mount status
12. Key server/node processes/services status (existence, basic response)
13. Availability of server/node network ports used by key server/node processes/services (accessibility)
14. Network interfaces status (Ethernet, FC, InfiniBand, etc.) of the server/node (up/down)
15. Server/node disks status according to disks' state and properties (up/down, disk failures, interface error (SAS, SATA, etc.), multipath status, etc.)
16. Status of disk array according to disk array controllers' state and the status of communication/data interfaces and disk array ports (both frontend and backend).
17. Disk array status according to the state of logical objects of disk array (RAID/volume group, logical volume)
18. Disk array status according to state of physical disks in disk array
19. Network devices status according to the state and properties of network interfaces (up/down, line speed) with respect to the expected state and configuration
20. Status of services provided in the high availability cluster and status of the high availability cluster itself

SPEC_300 EURO_IT4I system must collect performance and capacity parameters (metrics) and provide their visualization (hereinafter referred to as performance monitoring). Performance monitoring must provide data collection and charts:

For file storages:

1. current capacity usage and free capacity
2. current total number of files
3. current data throughput (read, write, total)
4. current number of I/O operations (read, write, total; metadata operations)

For SCRATCH storage - current capacity usage and throughput of physical storage (disk array, metadata/object storage target, etc.)

For Login, Visualization, Data management, and Infrastructure nodes:

1. current Ethernet interface throughput (send, receive, total)
2. current Compute network interface throughput (send, receive, total)

at a minimum, for three years, in 5 minute interval for the first month.

Performance monitoring must provide API for automated provision of data (metrics values).

The contracting authority prefers the following software:

- For metrics collection: time-series database Graphite or InfluxDB
- For chart visualization: Grafana

4.20.7 Software – Logging

SPEC_301 EURO_IT4I solution systems must log and centrally store information about activities, status changes, events, etc. for a period of 3 months, at a minimum. Logs must include timestamp, system, service and user identification, and event description. Solution must provide a storage for collected data with the total capacity of at least 20TB.

For centralized collection and processing of logs, the Elastic Stack must be used. The Elasticsearch solution must be implemented using native Elasticsearch cluster with at least three cluster members.

The Elasticsearch solution must be operated on suitable Infrastructure nodes. Reservation of these nodes for running the service is not required.

For data storage of the Elasticsearch solution, Infrastructure nodes' local disks in RAID (with redundancy) must be used.

4.20.8 Software – Mail

- SPEC_302 EURO_IT4I system must send all email communication exclusively through the contracting authority's SMTP servers.
The contracting authority will provide SMTP servers.

4.20.9 Software - Security

IT4Innovations National Supercomputing Center acquired the Information Security Management System Certification in accordance with ISO 27001 (ISO/IEC 27001:2013, ČSN ISO/IEC 27001:2014). EURO_IT4I system must be implemented in accordance with the internal regulations of the contracting authority.

- SPEC_303 EURO_IT4I system must provide access and services only to authorized users and systems. The system must not provide access or services to unauthorized users and systems. The system must be secured against data leaks, service misuse, and service and system breach.
For performance reasons, it is acceptable that patches mitigating hardware related cache side-channel attacks (Spectre, Meltdown) are not applied on servers, which provide services not affected by these vulnerabilities. Specifically, it is not necessary to use aforementioned patches on file servers and Compute servers.
- SPEC_304 Systems and services must use secure, strong passwords and secure keys, secure encryption, and secure protocols. Using default or weak passwords and/or keys is not permitted. Using identical authentication data for different accounts or services is not permitted.
- SPEC_305 Services not required for proper operation and functionality of the solution must not be run/enabled/available on the nodes, preferably not even installed.
- SPEC_306 EURO_IT4I systems must not communicate with other systems without explicit contracting authority's permission.
- SPEC_307 (I) In the offer, the supplier must describe the software solution and the names and number of licenses of the proposed software.


4.21 Requirements – Integration into Data Center

- SPEC_308 The supplier must implement and deploy EURO_IT4I system in the contracting authority's infrastructure – IT4Innovations data center (hereinafter referred to as integration into the data center).
- SPEC_309 Integration into the data center includes all deliveries and activities the result of which must be deployment of EURO_IT4I system in IT4Innovations data center premises and infrastructure.
- SPEC_310 Integration into the data center must be carried out in accordance with applicable legislation and regulations and in accordance with the requirements and recommendations of the manufacturers of individual systems. For installed systems, revisions required by legislation and regulations must be carried out.
- SPEC_311 The supplier must respect the contracting authority's infrastructure. The infrastructure is described in chapter 8 Contracting Authority Infrastructure.

- SPEC_312 Modification of the contracting authority's infrastructure necessary for implementation of EURO_IT4I system and its integration into the data center must be a part of EURO_IT4I system delivery.
- SPEC_313 The supplier must not change operational parameters of the contracting authority's data center. The offer must not use or assume data center infrastructure parameters that differ from those specified by the contracting authority's in the procurement documents.
- SPEC_314 All deliveries and activities of the integration into the data center must be discussed with and approved by the contracting authority. All use and possible modifications of the data center equipment or technologies (necessary for the integration into the data center) must be discussed with the contracting authority and the contracting entity providing data center service and must be approved by the contracting authority. The contracting authority reserves the right to change or modify technical integration proposals (routing, connection points, used materials, etc.) with respect to operating conditions and best practices of data center operation.

4.21.1 Requirements – Placement

- SPEC_315 System placement must respect disposition of the contracting authority's data center premises. The contracting authority explicitly points out the presence of columns in the data center.
- SPEC_316 All ICT equipment of EURO_IT4I system must be designed for installation into racks and must be installed into racks or must be equipment designed as a rack. Racks required for installation of the system must be a part of the delivery.
- SPEC_317 A part of the delivery is connection of racks required for realization of EURO_IT4I system into the contracting authority's infrastructure.
- SPEC_318 EURO_IT4I system equipment must be placed within the area shown in the figure below.



- SPEC_319 The distance of the racks from obstacles (walls, columns, rows of other racks) and the way of installation of the devices in the racks must not obstruct delivery, installation, replacement, and service of all equipment.
- SPEC_320 The placement of all installed equipment of the delivered solution must not obstruct revision of electrical installation, cooling distribution, air-conditioning, and security systems (sensors, etc.).
- SPEC_321 Rack installation on data center floor must be stable and safe. Rack weight must be suitably distributed. Floor construction and maximum load capacity (25kN/m²) must be observed. In case the installation applies point load or excessive load on a small area, load spreader must be used. Floor tiles modifications (penetrations to the raised floor area) must be carried out in a way that does not reduce the tile's load capacity. Adjustment of the floor tiles must be carried out by the contracting authority's contractor providing the data center services.
- SPEC_322 Cabling (inside and outside the racks) must provide reliable connection, handling of equipment (e.g. when servicing the equipment), and prevent damage to the cables during handling. Penetrations to the raised floor area must be fitted with grommets.
- SPEC_323 (I) In the offer, the supplier must provide a schema illustrating the proposed EURO_IT4I system placement in the contracting authority's data center.

4.21.2 Requirements – Power Supply

- SPEC_324 EURO_IT4I system power supply and operation solution must respect the contracting authority's constraints, specifically power circuits parameters.
- SPEC_325 EURO_IT4I system will be supplied simultaneously from independent data center power supply branches (A and B).
- SPEC_326 Failure or outage of power circuit or power supply branch must not damage any equipment and must not cause danger to persons or property.
- SPEC_327 Failure or outage of any single power circuit or power supply branch must not cause:
- Inaccessibility or outage of supercomputer services, with the exception of Universal and Accelerated compute nodes services
 - Inaccessibility or outage of Compute network services
 - Inaccessibility of more than 67% of Universal compute nodes
 - Inaccessibility of more than 67% of Accelerated compute nodes
- After the restoration of power circuit or power supply branch, the power redundancy of all devices must be ensured automatically.
- The ATS switch solution can be used only for LAN and Compute network switches power supply. In case the solution implements the ATS switch, the switch must provide remote management.
- SPEC_328 EURO_IT4I system must be designed in such a way that under standard/even load of ICT systems the supply branches and individual phases of the power supply are under approximately even load.
- SPEC_329 The solution must take into account the maximum input power for operation of all supplied equipment.
- SPEC_330 EURO_IT4I system must allow the entire system to be shut down. Proper EURO_IT4I system shutdown must take 60 minutes, at a maximum.
- SPEC_331 (I) In the offer, the supplier must include the energy calculation which must specify:
- Maximum electrical power consumption of the whole solution
 - Electrical power consumption of each rack or stand-alone equipment
- SPEC_332 (I) In the offer, the supplier must provide a schema and parameters of EURO_IT4I system connection to data center power circuit.

4.21.3 Requirements – Cooling

- SPEC_333 EURO_IT4I system cooling and operation solution must respect the contracting authority's constraints, specifically data center cooling system parameters.
- SPEC_334 The solution must be thermally neutral to the data room. The solution must provide cooling of all heat released by EURO_IT4I system equipment.
- SPEC_335 As the cooling source, EURO_IT4I system must only use water cooling circuits (cold and hot water circuits) of the data center.
- SPEC_336 EURO_IT4I system cooling will be supplied simultaneously from independent data center water cooling circuits.
- For cold water cooling, circuits SV1 and SV3 must be used simultaneously.

For hot water cooling, circuits TV1 and TV2 must be used simultaneously. These circuits are considered the primary ones. Internal cooling loops of the supplier's equipment (e.g. CDU) are considered the secondary ones.

SPEC_337 Failure or outage of the cooling circuits must not damage any equipment and must not cause danger to persons or property.

SPEC_338 Failure or outage of any single data center cooling circuit must not cause inaccessibility or outage of supercomputer services.

SPEC_339 EURO_IT4I system must be designed in such a way that under standard/even load of ICT systems the cold water cooling circuits are under approximately even load.

If EURO_IT4I system needs more than 300 kW of cooling capacity from the hot water cooling circuits (TV1 and TV2) then it must utilize both hot water cooling circuits and cooling must be designed in such a way that under standard/even load of ICT systems the hot water cooling circuits are under approximately even load.

The even distribution of the cooling is preferred by logical units, e.g. compute partitions where one circuit would serve as main for Accelerated compute partition and the other as main for Universal compute partition.

If EURO_IT4I system needs less than 300 kW of cooling capacity from the hot water cooling circuits (TV1 and TV2) then the distribution of the load between the two circuits is not required.

The fulfillment of the requirement regarding even load must not be based on alternating/switching of cooling circuits.

SPEC_340 Universal partition must use the hot water cooling circuits (TV1 and TV2) as prevailing source of cooling and must use direct liquid cooling. Universal partition must be fully operational with primary hot cooling circuits temperature set at any value in range of 30°C to 32°C.

SPEC_341 Each connection of EURO_IT4I system cooling to the primary cooling circuits must allow individual remotely controlled switching between two primary cooling circuits (see SPEC_336 for details). (For the avoidance of doubt - such a connection can be used for cooling of several devices/racks.)

Switching between the primary cooling circuits must take 3 minutes, at the maximum. Switching must not cause inaccessibility or outage of EURO_IT4I system services and must not damage any equipment in the data center. Circuit switching must not cause undesired primary circuits connections.

Remotely controlled cooling circuits switching will be integrated into the contracting authority's Measurement and regulation system (MaR). The integration will be performed by the contracting authority in cooperation with the supplier. The solution must be compatible with MP-Bus technology, it must provide an MP-Bus interface and act as an MP-Bus Slave device controlled by MaR MP-Bus Master.

SPEC_342 In case of risk of exceeding the maximum operating temperature of the equipment or components of the equipment (e.g. in case of failure of the source cooling circuit), the affected equipment must be shut down automatically and immediately to eliminate the risk of overheating or damaging the equipment or components.

- SPEC_343 (I) In the offer, the supplier must describe the cooling solution of EURO_IT4I system. Specifically, the supplier must provide information about the cooling of Universal partition, Accelerated partition and other compute partitions and systems, about cooling load distribution, and about cooling circuits switching.
- SPEC_344 (I) In the offer, the supplier must provide a schema and parameters of EURO_IT4I system connections to the data center cooling circuits.

4.21.4 Requirements – Transport

- SPEC_345 During equipment transport and installation, the load-bearing capacity of the floor on transport route must not be exceeded.
- The load-bearing capacity of the floor in the data center and its access corridor (room 219 and 223) is 25kN/m².
- The load-bearing capacity of the floor in the access area (room 217 and 218) is 5kN/m². For the transport of material, the supplier must temporarily install load distribution plates (e.g. plywood panels) in this area, so that the resulting load of the floor during the physical delivery does not exceed 5kN/m² and thus prevent floor damage.

5 Implementation and Further Actions

5.1 Implementation

- SPEC_346 A part of the delivery must be a comprehensive implementation of EURO_IT4I system so that all the contracting authority's requirements are met.
- SPEC_347 The delivery must include design, delivery, installation, implementation, configuration, debugging, testing of all systems, and performance of acceptance tests.

5.2 Training

- SPEC_348 A part of the delivery must be training in the scope and detail sufficient to acquire the knowledge necessary for the independent operation and administration of EURO_IT4I system.
- SPEC_349 The training must provide the information necessary to understand the internal functioning of the systems, hardware, and software. The training must include a thorough introduction to operating procedures and administration of EURO_IT4I system.
- SPEC_350 Duration of the training must be 25 hours, at a minimum. The number of participants must not be limited to less than 16 participants.
- SPEC_351 The training must include introduction to management framework/tools, practical demonstration and work with the real system. Fully implemented and accepted EURO_IT4I system must be used for demonstrations.
- SPEC_352 Trainers must be skilled professionals. The training must be in Czech or English.
- SPEC_353 The supplier must provide the teaching materials and presentations in English.
- SPEC_354 The schedule and detailed plan of the training must be produced in cooperation with the contracting authority's project manager.

- SPEC_355 The training must take place in the contracting authority's premises. For this purpose a conference room will be provided free of charge.
- SPEC_356 (l) In the offer, the supplier must include the frame schedule and content of the training.

5.3 Documentation

- SPEC_357 The supplier must provide a comprehensive documentation of EURO_IT4I system.
- SPEC_358 EURO_IT4I system documentation must comprehensively cover all supplied systems and must be logically structured. The documentation must include the documentation of the actual state of the system and documentation of operation procedures (operating manuals). The documentation must describe all specific (custom) modifications (settings, functionalities, etc.).
- SPEC_359 The documentation must also cover processes and procedures for system administration, regular maintenance, emergency management, and recovery.
- SPEC_360 The documentation must also include English documentation (datasheets/spec sheets, manuals, administrator, and user guides) of delivered hardware and software. Documentation must be provided in electronic form, allowing text copying.

5.4 Benchmarking support

- SPEC_361 The supplier must provide expert support for the benchmarking of the system. The optimization of benchmark performance, rule conforming execution and the submission of the results to the official lists will be performed by the supplier. The contracting authority intends to include the system in the following lists benchmarking:
1. Top500, HPCG, Graph500
 - i. Universal Partition
 - ii. Accelerated Partition
 - iii. Data Analytics Partition
 2. Green500
 - i. Universal Partition
 - ii. Accelerated Partition
 3. IO500
 - i. Universal Partition
 - ii. Accelerated Partition
 4. MLPerf
 - i. Accelerated Partition

5.5 EC Declaration of Conformity

- SPEC_362 All delivered systems and equipment must be accompanied by EC declaration of conformity.

5.6 Waste Disposal

- SPEC_363 A part of the delivery is the disposal of waste produced by implementation of the delivery. The contracting authority is not obligated to store packaging and packaging material and will not do so.

6 Warranty and Services

The system must have warranty and be serviced for 5 years.

There are mandatory requirements for the ticketing system used for warranty and support as well for the implementation phase of the system (installation): Full history of tickets processed, REST API functionality to create, update, and close the tickets, download full history of tickets including all details, metadata, and attachments. If the requirement for the specified support ticketing system prevents you to offer a complete EURO_IT4I system please state it explicitly.

7 Content Requirements of the Technical Solution Proposal

EURO_IT4I system technical solution proposal must include:

- the detailed description of EURO_IT4I system architecture
- proposed technologies
- functionality and properties of the solution
- the number of equipment and their configuration
- the number of licenses and names of software
- the means of meeting contracting authority's requirements

EURO_IT4I system technical solution proposal must be provided in an electronic form, allowing text copying.

The license terms will be added in the relevant annex to the binding model contract.

8 Contracting Authority Infrastructure

EURO_IT4I system will be installed and operated in the data room of the contracting authority's data center in the IT4Innovations building. The IT4Innovations building is located on the campus of the Technical University of Ostrava, Studentská 6231/1B, 708 00 Ostrava-Poruba.

Existing systems of the contracting authority are located in the data room.

Values mentioned in this section are indicative, real values may slightly vary based on various conditions (e.g. measurement accuracy, operating conditions, environmental conditions, etc.).

8.1 Data Center Dispositions

The data center is located on the second floor of the IT4Innovations building (room no. 223). The data center dimensions are 24.97 x 20.47 meters (511.1 square meters) with the clear height of 4.5 meters. The hall is designed as a separate, structurally undivided space.

Data room floor is systemic, antistatic raised floor with a grid of 600x600mm tiles, with height of 980mm. Data room ceiling with a grid of 600x600mm is suspended 700mm under the standard ceiling. The clear height between the systemic floor and the suspended ceiling is 2.9m.

8.1.1 Floor

Raised floor height: 980 mm – systemic solution. Load capacity limit 2,500 kg/m² (25.0 kN/m²).

Frame construction made of steel "C" profiles in 600x600 mm grid. The support posts consist of height-adjusted posts with adjustable head. The columns and horizontal frames are connected by a fixed screw connection.

Floor tiles are made of mineral-based core; fire reaction class according to EN 13501 A2. On the bottom of the tile is an AL foil. The sides of the tiles have plastic edges. The upper side of the tiles have a PVC layer. The surface is anti-static.

8.1.2 Suspended Ceiling

Below the building ceiling, there is a suspended metal ceiling made of galvanized steel 600x600x33mm, beveled edge 3mm; fitted into the hidden structure; smooth surface without perforation; UV-stable, electrostatically applied polyester varnish. Inspection cartridges supplemented with clips for tilting the cartridges down according to the suspended ceiling layout. A demounting tool is required to access the ceiling area.

Concealed hanging metal construction with U-profile and perpendicular DP-profile suspended by threaded rods.

8.2 Power Supply

Main Power Branch:

Topology: two independent power branches, 1+1 redundancy

Distribution system NN: 400/230V; 3+N+PE; 50Hz; TN-S

Distribution busbars:

Number of distribution busbars: 5pcs for each power branch

Distribution system rated current: 800A

Distribution busbars power supply system: TN-S 3x400/230V 50Hz

Distribution busbar rated power: 552kVA

Distribution system location: under raised floor

Distribution system connection method: PDB switchboard

PDB switchboard outlets specification 6x3f/32A; different configuration possible; protection: fuse with gG characteristics

IT technology connection method: IEC60309 industrial sockets

The connection points are equipped with the 3f/32A configuration and are fitted with PKY32G435 sockets. In case a different configuration of PDB switchboard outlets is required, such PDB switchboards are a part of the delivery.

8.3 Cooling

Data center cooling is provided by five separate coolant circuits – two hot-water circuits and three cold-water circuits.

Under the data center raised floor, there are three branches drawn from each of the five cooling circuits. Branches are fitted with DN65 dry disconnect couplings. The branches are located under the floor of the data center and are accessible after removing some of the floor tiles. Circuits use a mixture of 35% propylene glycol and 65% water.

Basic parameters of cooling circuits, including connection data; given values are set-points (i.e. desired or target values of MaR system), real values vary based on operating conditions:

TV1 – hot-water circuit – red:

- End fittings: Victaulic UK 76,1
- Working pressure: 4.5bar
- data center differential pressure: 50kPa
- Flow rate: 40m³/hour
- Coolant inlet temperature: 30-32°C

TV2 – hot-water circuit – yellow:

- End fittings: Victaulic UK 76,1
- Working pressure: 4.5bar
- data center differential pressure: 50kPa
- Flow rate: 40m³/hour
- Coolant inlet temperature: 30-32°C

SV1 – cold-water circuit – green:

- End fittings: Victaulic UK 76,1
- Working pressure: 4.5bar
- data center differential pressure: 100kPa
- Flow rate: 20m³/hour
- Coolant inlet temperature: 11.5°C

SV2 – cold-water circuit – blue:

- End fittings: Victaulic UK 76,1
- Working pressure: 4.5bar
- data center differential pressure: 100kPa
- Flow rate: 20m³/hour
- Coolant inlet temperature: 11.5°C

SV3 – cold-water circuit – turquoise:

- End fittings: Victaulic UK 76,1
- Working pressure: 4.5bar
- data center differential pressure: 100kPa
- Flow rate: 20m³/hour
- Coolant inlet temperature: 11.5°C

8.4 Environment

Environment in the data room has the following conditions and parameters:

- Ambient air temperature: 22°C to 25°C
- Non-condensing moisture
- Relative humidity: 20–60%
- Hypoxic atmosphere, oxygen levels are decreased to 15 % vol. to prevent fire

8.5 Data Center Access Path

Equipment required for implementing the solution can be transported in the data center using the loading dock located in the northeast part of the building. The loading dock is 2850 mm wide and 1030 mm high.

For transportation, we recommend using trucks with lift gate. Room no. 218 is accessible from the loading dock area. Equipment with dimensions up to 2410x1540 mm (height x width) can be transported through the entry door leading to room 218.

Room 218 with dimensions of 5.3m x 5.6m can be used for the removal of transport packages or as a temporary storage during transportation.

Room 218 is connected to the data center by corridor (room 219). Rooms 218 and 219 are divided by door through which objects with dimensions up to 2340x1600 mm (height x width) can be transported.

A part of this corridor's floor, specifically a length of 8.5m, has a slope of 6.5° inclination

Entrance to the data center (room 223) is in the upper section of the access corridor. Equipment with dimensions up to 2360x1520 mm (height x width) can be transported through the data center entry door.

To fit through all three door, the equipment can have dimensions of 2340x1520mm (height x width), at a maximum.


The equipment can be transported to the data center via several routes, the only limitation are the dimensions of the door stated above.

The maximum floor load capacity in the access corridor (room 219) and the data center (room 223) is 2500 kg/m² (25 kN/m²).

8.6 WAN/LAN Network

At the time of EURO_IT4I system implementation, the contracting authority's current WAN/LAN network solution will be replaced with a new infrastructure for 100Gb/s Ethernet technology. Central devices of the new WAN/LAN will be L3 switches. The central devices will be structured as multichassis with common data plane and multichassis EtherChannel support or functionally identical technology. The central devices will have HSRP or VRRP technology. The central switches of the new WAN/LAN network will provide 4 ports equipped with QSFP-100G-SR4 modules with MPO-12 connectors for connecting EURO_IT4I system edge devices with the speed of 4x100Gb/s.

The location of the WAN rack and cable routes is illustrated below.



The specific type of central switches of the new WAN/LAN network is not yet known, because the procurement of these devices has not yet been fully realized.

8.7 OOB Network

The contracting authority's current OOB network consists of:

- OOB router Cisco 2901/K9, console modules HWIC-16A (cables CAB-HD8-ASYNC)
- OOB switch Cisco WS-C3650-48TS, ports 10/100/1000BaseTX
- Cat6a patch panel with 24 RJ-45 ports, of which 12 ports are free

OOB devices are located in the WAN rack.

Annex 2

Time Schedule

No.	Milestone	Deadline
1.	The Contract takes effect	D
2.	Delivery of the Draft Detailed Project Plan	D+2 weeks
3.	Discussion and acceptance of the Detailed Project Plan	D+1 month
4.	Delivery of the Draft Detailed Technical Specification	D+1 and half months
5.	Discussion and acceptance of the Detailed Technical Specification	D+2 and half months
6.	Delivery of the Draft Specification of Acceptance Tests	D+3 months
7.	Discussion and acceptance of the Specification of Acceptance Tests	D+4 months
8.	Physical supply of the part of the System	D+4 and half months
9.	Completion of the part of the System Implementation	D+4 and half months
10.	Delivery of the part of the System's documentation	D+4 and half months
11.	Successful performing of Acceptance Tests of the part of the System	D+5 months
12.	Completion of training to the part of the System	D+5 months
13.	Physical supply of the remaining part of the System	D+7 and half months
14.	Completion of the Work Implementation	D+7 and half months
15.	Delivery of the remaining part of the System's documentation	D+7 and half months
16.	Successful performing of Acceptance Tests of the Work	D+8 months
17.	Completion of training to the Work	D+8 months
18.	Comprehensive Supply of the System	D+8 months

The discussion and acceptance of the Detailed Project Plan (milestone no. 3) involves the presentation of the Detailed Project Plan document by the Supplier to the Client, assessment of this document by the Client, notification of the Client of its consent with the document or raised objections and made comments by the Client, or a potential presentation of an adjusted document to the Client for acceptance and signing of the handover protocol, all in compliance with the rules resulting from par. 6.3 hereof. The Detailed Project Plan document appropriately describes and determines a procedure, chronological continuity, necessary sources (both human and material) and mutual dependencies of all activities and deliverables planned for a timely achievement of milestone no. 18 Comprehensive Supply of the System. The required level of detail in the discussion process is determined by the Client.

The discussion and acceptance of the Detailed Technical Specification (milestone no. 5) involves a presentation of the Detailed Technical Specification document by the Supplier to the Client, assessment of this document by the Client, notification of the Client of its consent with the document or raised objections and made comments by the Client, or a potential presentation of an adjusted document to the Client for acceptance and signing of the handover protocol, all in compliance with the rules resulting from par. 6.3 hereof. The Detailed Technical Specification document describes the technical implementation of the System and all its parts in detail. The required level of detail in the discussion process is determined by the Client.

The discussion and acceptance of the Specification of Acceptance Tests (milestone no. 7) involves a presentation of the Specification of Acceptance Tests document by the Supplier to the Client, assessment of this document by the Client, notification of the Client of its consent with the document or raised objections and made comments by the Client, or a potential presentation of an adjusted document to the Client for acceptance and signing of the handover protocol, all in compliance with the rules resulting from par. 6.3 hereof. The Specification of Acceptance Tests document determines the expected outcomes of individual acceptance tests and describes the manner of performing acceptance tests of the System in detail. The required level of detail in the discussion process is determined by the Client. For the sake of certainty, the contracting parties state that this milestone means acceptance of the Specification of Acceptance Tests of the Work as such, i.e. both gradually delivered parts of the System within the meaning of milestones no. 8 and 13.

Physical supply of the part of the System (milestone no. 8) involves the supply of the part of the System containing Universal compute partitions, Data Analytics compute partitions, Cloud infrastructure compute partitions, Visualization nodes, User's data storage SCRATCH, User's data storage HOME, PROJECT network gateways, High speed compute network, Login nodes (including Data management nodes), Management nodes (including Infrastructure nodes), System's data storage INFRA, Network infrastructure LAN (including integration of WAN), Backup system, Infrastructure integration, Software, and presentation and verification of the factual condition of the supply above. The Physical supply of the part of the System milestone may be achieved only after the Discussion and acceptance of the Detailed Technical Specification milestone.

Completion of the part of the System Implementation (milestone no. 9) is a situation when the Supplier considers the part of the System in the meaning of milestone no. 8 as complete, with all required functionality and declared parameters and shall notify the Client of this fact. After the Completion of the part of the System Implementation milestone, the Supplier must not engage in any implementation works or adjustments to the delivered part of the System, except for works and adjustments approved by the Client, and works on subsequent milestones, if necessary, but always on the basis of the notification to the Client.

It is possible to commence the acceptance tests of the part of the System in the meaning of milestone no. 11 only after the achievement of milestone no. 9 Completion of the part of the System Implementation. The failed acceptance tests mean that the System implementation has not been completed and the System implementation continues. From the moment of fulfilment of milestone no. 11 Successful performing of Acceptance Tests of the part of the System starts the warranty period within the meaning of par. 11.2 hereof.

Physical supply of the remaining part of the System (milestone no. 13) involves the supply and presentation and verification of the factual condition of all the part of the System.

Completion of the Work Implementation (milestone no. 14) is a situation when the Supplier considers the Large Cluster as a whole as complete, with all required functionality and declared parameters and shall notify the Client of this fact. After the Completion of the Work Implementation milestone, the Supplier must not engage in any implementation works or adjustments to the delivered part of the System, except for works and adjustments approved by the Client.

It is possible to commence the acceptance tests of the Work in the meaning of milestone no. 16 only after the achievement of milestone no. 14 Completion of the Work Implementation. The failed acceptance tests mean that the System implementation has not been completed and the System implementation continues.

Comprehensive Supply of the System (milestone no. 18) involves the implementation of the Work as defined in par. 3.1 hereof. To achieve milestone no. 18, it is necessary to achieve all other milestones as listed in this Annex 2 hereto which shall be checked and confirmed by the Client.

Annex 3

Cooperation of the Client

The cooperation of the Client in the provision of performance hereunder shall consist primarily in the following activities:

- Provide for cooperation and participation of responsible employees of the Client and persons knowing the environment of the Supplier in the scope necessary for the completion of the Work. The cooperation shall primarily consist in provision of consultations aiming to add information to the presented documentation in the preparation of the Detailed Project Plan, Detailed Technical Documentation and Specification of Acceptance Tests;
 - Provide access to buildings of the Client for members of the Supplier's implementation team;
 - When needed, provide for the participation of responsible employees of the Client in the performance of acceptance tests;
 - Provide information on planned changes in the Client's environment in the acceptance tests;
 - Provide information on planned change in the Client's environment, if they are going to have any impact on the performance by the Supplier and provide complete, truthful and timely information;
 - Training – for the purposes of training, it shall be possible to use the premises of VSB-TUO; VSB-TUO shall provide for the room and presentation technology (projector, whiteboard), VSB-TUO shall provide for the participation of attendants of the training on date in accordance with the Time Schedule approved by both parties.
 - Reactive service – VSB-TUO shall provide for the conditions of Reactive service at the place of installation. VSB-TUO shall provide an escalation mechanism, after agreement with the Supplier, for the classification of defects for situation when the employees of the Client and the Supplier dealing with the defects do not agree on the defect classification. If there are disputes regarding the determination of the escalation procedure, they shall be resolved by the Client.
 - Acceptance tests – cooperation in the preparation and performance of the Acceptance tests.
 - Payment, billing – the Client shall provide necessary cooperation for securing of formal correctness and acceptability of documents.
 - Insurance – cooperation in the provision of required information of the Client for conclusion of the insurance of the Supplier.
 - Premises – VSB-TUO shall provide premises for storage of materials in the necessary scope during the period of the Work implementation.
 - Data room – VSB-TUO shall provide for the cooperation of another supplier of VSB-TUO (primarily so called Building II) in the scope necessary for the integration of the Work in the Data Room.
-

Annex 4

Specification of the Price and Payment Schedule

Specification of the price

Price of the Work as in par. 12.1 of the Contract		Total price without VAT [EUR]
1.	Universal compute partitions	4 940 000,00
2.	Accelerated compute partitions	5 750 000,00
3.	Data Analytics compute partitions	370 000,00
4.	Cloud infrastructure compute partitions	250 000,00
5.	Visualization nodes	22 000,00
6.	User's data storage SCRATCH	1 310 000,00
7.	User's data storage HOME	51 000,00
8.	PROJECT network gateways	42 000,00
9.	High speed compute network	1 400 000,00
10.	Login nodes (including Data management nodes)	45 000,00
11.	Management nodes (including Infrastructure nodes)	52 000,00
12.	System's data storage INFRA	53 000,00
13.	Network infrastructure LAN (including integration of WAN)	150 000,00
14.	Backup system	50 000,00
15.	Infrastructure integration	120 000,00
16.	Software	250 000,00
Total without VAT [EUR] (price for the Work as in par. 12.1 of the Contract)		14 855 000,00

* Payment conditions are set in article 12 of the Contract.

Payment schedule

The Client shall pay the price in accordance with par. 12.1 hereof in two parts. Upon completion of milestone no. 11 pursuant to Annex no. 2 to the Contract, the EuroHPC JU shall pay the Supplier part of the price of the Work in the amount of EUR 5,130,000 based on the invoice issued by the Supplier. Upon completion of milestone no. 18 pursuant to Annex no. 2 to the Contract, the VSB-TUO shall pay the Supplier the remaining part of the price of the Work, however for no more than CZK 251,800,000 net of VAT based on the invoice issued by the Supplier.

Annex 5

Licence Conditions for the Selected Standard SW

1. Definition of the use of software products as part of the supply

- 1.1 All software shall be used on hardware of IT4Innovations National Supercomputer Centre (hereinafter "IT4Innovations") which is part of Vysoká škola báňská – Technical University Ostrava, hardware is or shall be in the sole ownership of the Client, software therefore shall not be provided for operations on external hardware. All software shall be used for academic purposes, science and research, and partially for commercial research on a contractual basis, however to the maximum extent of 20% allowed by the Communication from the Commission (EU) "Framework for State aid for research and development and Innovation" (2014/C 198/01). These are not versions intended solely for education of students.
 - 1.2 Licencing of all software supplied under the „EURO_IT4I Supercomputer“ public contract should allow for the use of these technical means and services provided by it in addition to the primary acquirer of the licence which is Vysoká škola báňská – Technical University Ostrava, which assumes the use by the employees and bodies of Vysoká škola báňská – Technical University Ostrava; and the following groups of third parties:
 - 1) **All registered users of IT4Innovations** – Registered users are the users who have an open account in IT4Innovations. The following users (especially, but not only) acquire the account
 1. Those who succeeded in the grant competition for the allocation of computing sources; or
 2. Their application was approved; or
 3. They attend training or other educational events of IT4Innovations;
 4. And met the conditions of IT4Innovations for access and use of hardware means of IT4Innovations,
 - 2) **Users provided with access to computing sources of the Large Cluster by EuroHPC JU,**
 - 3) **Contractual and other partners of the Client.**
 - 1.3 Licencing of the software supplied as part of the public contract must allow for the provision of services of the System to 3 thousand end users and administration of the System by 50 administration users (operators, administrators).
 - 1.4 The Supplier is aware of the information by the Client that the acquirer of the licence is the Client and the licence granted to the acquirer shall include an authorisation to adjust and develop the Large Cluster for the acquirer of the licence (the Client) or a third party authorised by the Client.
 - 1.5 Other conditions for the provision of licence to the software supplied by the Supplier are defined in article 10 hereto. If there are any disputes between Annex 5 hereto and article 10 hereof, the contracting parties have agreed that the licence conditions of the software supplied by the Supplier shall prevail over article 10 hereto, however, they shall always adhere to conditions in point 1 of Annex 5 hereto.
-

2. Licence conditions to the selected standard and open source SW

The main software components of the solution are:

CentOS

RedHat

HPE Performance Cluster Manager

PBS Professional

VirtualGL and TurboVNC

Lustre Software - ClusterStor

Bareos Software - Backup

... and other open source SW tools mentioned in the Design of Technical Solution in Annex 1 of this Contract.



1. **Applicability.** This end user license agreement (the "Agreement") governs the use of accompanying software, unless it is subject to a separate agreement between you and Hewlett Packard Enterprise Company and its subsidiaries ("HPE"). By downloading, copying, or using the software you agree to this Agreement. HPE provides translations of this Agreement in certain languages other than English, which may be found at: <http://www.hpe.com/software/SW Licensing>.
2. **Terms.** This Agreement includes supporting material accompanying the software or referenced by HPE, which may be software license information, additional license authorizations, software specifications, published warranties, supplier terms, open source software licenses and similar content ("Supporting Material"). Additional license authorizations are at: <http://www.hpe.com/software/SW Licensing>.
3. **Authorization.** If you agree to this Agreement on behalf of another person or entity, you warrant you have authority to do so.
4. **Consumer Rights.** If you obtained software as a consumer, nothing in this Agreement affects your statutory rights.
5. **Electronic Delivery.** HPE may elect to deliver software and related software product or license information by electronic transmission or download.
6. **License Grant.** If you abide by this Agreement, HPE grants you a non-exclusive non-transferable license to use one copy of the version or release of the accompanying software for your internal purposes only, and is subject to any specific software licensing information that is in the software product or its Supporting Material.

Your use is subject to the following restrictions, unless specifically allowed in Supporting Material:

- You may not use software to provide services to third parties.
 - You may not make copies and distribute, resell or sublicense software to third parties.
 - You may not download and use patches, enhancements, bug fixes, or similar updates unless you have a license to the underlying software. However, such license doesn't automatically give you a right to receive such updates and HPE reserves the right to make such updates only available to customers with support contracts.
 - You may not copy software or make it available on a public or external distributed network.
 - You may not allow access on an intranet unless it is restricted to authorized users.
 - You may make one copy of the software for archival purposes or when it is an essential step in authorized use.
 - You may not modify, reverse engineer, disassemble, decrypt, decompile or make derivative works of software. If you have a mandatory right to do so under statute, you must inform HPE in writing about such modifications.
7. **Remote Monitoring.** Some software may require keys or other technical protection measures and HPE may monitor your compliance with the Agreement, remotely or otherwise. If HPE makes a license management program for recording and reporting license usage information, you will use such program no later than 180 days from the date it's made available.
 8. **Ownership.** No transfer of ownership of any intellectual property will occur under this Agreement.
 9. **Copyright Notices.** You must reproduce copyright notices on software and documentation for authorized copies.
 10. **Operating Systems.** Operating system software may only be used on approved hardware and configurations.





11. 90-day Limited Warranty for HPE Software.

- HPE-branded software materially conforms to its specifications, if any, and is free of malware at the time of delivery; if you notify HPE within 90 days of delivery of non-conformance to this warranty, HPE will replace your copy. This Agreement states all remedies for warranty claims.
- HPE does not warrant that the operation of software will be uninterrupted or error free, or that software will operate in hardware and software combinations other than as authorized by HPE in Supporting Material. To the extent permitted by law, HPE disclaims all other warranties.

12. Intellectual Property Rights Infringement. HPE will defend and/or settle any claims against you that allege that HPE-branded software as supplied under this Agreement infringes the intellectual property rights of a third party. HPE will rely on your prompt notification of the claim and cooperation with our defense. HPE may modify the software so as to be non-infringing and materially equivalent, or we may procure a license. If these options are not available, we will refund to you the amount paid for the affected product in the first year or the depreciated value thereafter. HPE is not responsible for claims resulting from any unauthorized use of the software.

13. Limitation of Liability. HPE's liability to you under this Agreement is limited to the amount actually paid by you to HPE for the relevant software, except for amounts in Section 12 ("Intellectual Property Rights Infringement"). Neither you nor HPE will be liable for lost revenues or profits, downtime costs, loss or damage to data or indirect, special or consequential costs or damages. This provision does not limit either party's liability for: unauthorized use of intellectual property, death or bodily injury caused by their negligence; acts of fraud; willful repudiation of the Agreement; or any liability that may not be excluded or limited by applicable law.

14. Termination. This Agreement is effective until terminated or in the case of a limited-term license, upon expiration; however, your rights under this Agreement terminate if you fail to comply with it. Immediately upon termination or expiration, you will destroy the software and documentation and any copies, or return them to HPE. You may keep one copy of software and documentation for archival purposes. We may ask you to certify in writing you have complied with this section. Warranty disclaimers, the limitation of liability, this section on termination, and Section 15 ("General") will survive termination.

15. General.

- a. **Assignment.** You may not assign this Agreement without prior written consent of HPE, payment of transfer fees and compliance with HPE's software license transfer policies. Authorized assignments will terminate your license to the software and you must deliver software and documentation and copies thereof to the assignee. The assignee will agree in writing to this Agreement. You may only transfer firmware if you transfer associated hardware.
- b. **U.S. Government.** If the software is licensed to you for use in the performance of a U.S. Government prime contract or subcontract, you agree that, consistent with FAR 12.211 and 12.212, commercial computer software, computer software documentation and technical data for commercial items are licensed under HPE's standard commercial license.
- c. **Global Trade Compliance.** You agree to comply with the trade-related laws and regulations of the U.S. and other national governments. If you export, import or otherwise transfer products provided under this Agreement, you will be responsible for obtaining any required export or import authorizations. You confirm that you are not located in a country that is subject to trade control sanctions (currently Cuba, Iran, N. Korea, N. Sudan, and Syria) and further agree that you will not retransfer the products to any such country. HPE may suspend its performance under this Agreement to the extent required by laws applicable to either party.
- d. **Audit.** HPE may audit you for compliance with the software license terms. Upon reasonable notice, HPE may conduct an audit during normal business hours (with the auditor's costs being at HPE's expense). If an audit reveals underpayments then you will pay to HPE such underpayments.





If underpayments discovered exceed five (5) percent, you will reimburse HPE for the auditor costs.

- e. Open Source Components. To the extent the Supporting Material includes open source licenses, such licenses shall control over this Agreement with respect to the particular open source component. To the extent Supporting Material includes the GNU General Public License or the GNU Lesser General Public License: (a) the software includes a copy of the source code; or (b) if you downloaded the software from a website, a copy of the source code is available on the same website; or (c) if you send HPE written notice, HPE will send you a copy of the source code for a reasonable fee.
 - f. Notices. Written notices under this Agreement may be provided to HPE via the method provided in the Supporting Material.
 - g. Governing Law. This Agreement will be governed by the laws of the state of California, U.S.A., excluding rules as to choice and conflict of law. You and HPE agree that the United Nations Convention on Contracts for the International Sale of Goods will not apply.
 - h. Force Majeure. Neither party will be liable for performance delays nor for non-performance due to causes beyond its reasonable control, except for payment obligations.
 - i. Entire Agreement. This Agreement represents our entire understanding with respect to its subject matter and supersedes any previous communication or agreements that may exist. Modifications to the Agreement will be made only through a written amendment signed by both parties. If HPE doesn't exercise its rights under this Agreement, such delay is not a waiver of its rights.
16. **Australian Consumers**. If you acquired the software as a consumer within the meaning of the 'Australian Consumer Law' under the Australian *Competition and Consumer Act 2010* (Cth) then despite any other provision of this Agreement, the terms at this URL apply: <http://www.hpe.com/software/SWLicensing>.
17. **Russian Consumers**. If you are based in the Russian Federation and the rights to use the software are provided to you under a separate license and/or sublicense agreement concluded between you and a duly authorized HPE partner, then this Agreement shall not be applicable.



Cray Inc.

Software End User License Agreement

This Software End User License Agreement (the “EULA”) is made between Customer, as indicated on the Order, and Cray Inc. with its principal office located at 901 Fifth Ave., Suite 1000, Seattle, WA 98164 (“Cray”)

1) **Definitions.** Capitalized terms not defined in context shall have the meanings assigned to them below:

a) **“Cray Software”** means the object code version of the software that is provided or made available to Customer pursuant to an Order as well as any corrections, enhancements, and upgrades to such software that are made available to Customer from time to time, and all copies of the foregoing, with the exclusion of Open Source Software and Third-Party Software.

b) **“Documentation”** means the user manuals and documentation that Cray makes available for the Cray Software, Open Source Software, Third-Party Software, and all copies of the foregoing.

c) **“License Type”** means the metric by which the Software is licensed (e.g., by system, by user, etc.) as indicated in the applicable Order.

d) **“Open Source Software”** means software components provided by Cray that are subject to various open source software license terms as may be described in the software’s Documentation, within the provided source code of the software, or in an attribution or licensing file provided to Customer with the software.

e) **“Order”** means the document(s) by which Customer orders, purchases, procures, or otherwise obtains a right to use the Software as well as any accompany System. Any conflicting or additional terms in or accompanying an Order will not be binding on Cray unless Cray specifically accepts and acknowledges such terms in writing.

f) **“Software”** means the combination of Cray Software, Open Source Software, and Third-Party Software delivered by Cray under an applicable Order.

g) **“System”** means a particular system of Cray-branded computer hardware that may be

described in an Order. The System does not include hardware that is not branded as CRAY hardware even if provided in connection with an Order.

h) **“Third-Party Software”** means proprietary software delivered by Cray as part of the Software that is licensed by and subject to separate license terms provided by a third-party.

2) **Software Licenses.**

a) **General.** Subject to the terms of this EULA, Cray grants to Customer, and Customer accepts from Cray, for so long as this EULA remains in effect, and unless terminated in accordance with Section 4), a non-exclusive, non-transferable (except as otherwise set forth herein) and non-sublicensable license to access and use the quantities of each item of Cray Software identified in the applicable Order within the parameters of the Documentation associated with the applicable Cray Software and License Type (the **“Cray Proprietary License”**) solely for use within the System or such other purpose stated in the Order. Unless described otherwise in the Order, Customer shall only use the Software to support its internal business operations and may not operate the Software as a hosted service, shared implementation, or as part of a software as a service solution. Unless otherwise specified in an Order, the Cray Proprietary License shall also include the rights to: i) make reasonable additional copies of the Cray Software to be used solely for non-production archival or passive disaster recovery purposes,, and ii) make and use copies of the Documentation as reasonably necessary to support Customer’s authorized users in their use of the Software.

b) **Open Source Software.** To the extent the Software is comprised of Open Source Software, Customer’s license to use, copy, distribute, or otherwise exploit each such item of the Open Source Software will be governed exclusively by the applicable Open Source Software license associated with such Open Source Software (each an **“OSS**

License”). Notwithstanding anything otherwise set forth in this EULA, Customer understands and agrees that the only warranties and representations respecting the Open Source Software are those provided in the applicable OSS Licenses, and the terms of **Section 9** (Infringement Indemnity) do not apply to any claims to the extent they arise out of the Open Source Software. Customer’s sole remedies respecting any allegations or losses arising out of the Open Source Software are those set forth in the applicable OSS Licenses.

c) Restrictions. Except and only to the extent that the exclusions and limits of this Section are prohibited by applicable law or, for Open Source Software when expressly permitted under terms of its OSS License, Customer may not reverse engineer, decompile, disassemble, or attempt to discover or modify in any way the underlying source code of the Software, or any part thereof. In addition, other than as permitted for any Open Source Software under terms of its OSS License, Customer may not: i) modify, translate, localize, adapt, rent, lease, loan, create or prepare derivative works of, or create a patent based on the Software, Documentation or any part thereof; ii) resell, sublicense or distribute the Software or Documentation; iii) provide, make available to, or permit use of the Software or Documentation, in whole or in part, by any third party (except as expressly set forth herein) without the Software’s licensor’s prior written consent; iv) use the Software or Documentation to create or enhance a competitive offering or for any other purpose which is competitive to Cray; v) remove Software that was delivered on a System from the System on which it was delivered and load such Software onto a different system or hardware component without Cray’s prior written consent; or vi) perform or fail to perform any act which would result in a misappropriation or infringement of Cray’s or its suppliers’ intellectual property rights in the Software or Documentation. Each permitted copy of the Software and Documentation made by Customer hereunder must contain all titles, trademarks, copyrights and restricted rights notices as in the original. Customer understands and agrees that the Software may work in conjunction with third party products and Customer agrees to be solely responsible for ensuring that it is properly licensed by such third parties to use the Software with such products.

3) Proprietary Rights.

a) Cray reserves any and all rights, implied or otherwise, which are not expressly granted to Customer in this EULA.

b) To the extent the Software is comprised of Third Party Software, Customer’s license to use, copy, distribute, or otherwise exploit each such item of Third Party Software will be governed exclusively by the applicable software license associated with the Third Party Software (each a **“Third Party License”**). Cray shall provide copies of each such Third Party License to Customer with the delivery of the Software. Customer is advised that copies of all Third Party Licenses are also available at <http://www.cray.com/company/policies-and-practices/software-license-agreements>.

Notwithstanding anything otherwise set forth in this EULA, Customer understands and agrees that the only warranties and representations respecting the Third Party Software are those provided in the applicable Third Party Licenses, and the terms of Section 9) (Infringement Indemnity) of this EULA do not apply to any claims to the extent they arise out of the Third Party Software. Customer’s sole remedies respecting any allegations or losses arising out of the Third Party Software are those set forth in the applicable Third Party Licenses.

c) Cray reserves the right to request, on a reasonable basis, and Customer shall provide, evidence demonstrating compliance by Customer with the terms of this EULA.

4) Termination.

a) This EULA, Cray Proprietary License, OSS Licenses, or Third-Party License terms granted hereunder may be terminated: i) by mutual written agreement of Cray and Customer; or ii) by either party for reasons arising out of a material uncured breach of this EULA by the other party, that the breaching party fails to cure to the non-breaching party’s reasonable satisfaction within 30 days following its receipt of notice of the breach.

b) Upon termination of this EULA or expiration or termination of a license for any reason, all rights granted to Customer for the applicable Software shall immediately cease and Customer shall immediately: i) cease using the applicable Software and Documentation; ii) remove all copies, installations, and instances of the applicable Software from all Customer computers and any other devices on which the Software was installed; iii) return the applicable Software to Cray together

with all Documentation and other materials associated with the Software and all copies of any of the foregoing, or destroy such items; and iv) upon Cray's request give Cray written certification made by an officer or other authorized official of Customer, within ten days, that Customer has complied with all of the foregoing obligations.

5) Survival. Any provision of this EULA that requires or contemplates execution after termination of this EULA or after termination or expiration of any of the licenses granted hereunder, is enforceable against the other party and their respective successors and assignees, notwithstanding such termination. Termination of this EULA or a license shall be without prejudice to any other remedies that the terminating party may have under law, subject to the limitations and exclusions set forth in this EULA.

6) Export. Customer acknowledges that the Software and Documentation are subject to the export control laws, rules, regulations, and restrictions including, without limitation, the United States Export Administration Regulation, and laws of other applicable foreign agencies (the "**Export Controls**"). Customer shall abide by the Export Controls. Customer hereby agrees to use the Software and Documentation in accordance with the Export Controls, and shall not export, re-export, sell, lease or otherwise transfer the Software or Documentation or any copy, portion or direct product of the foregoing in violation of the Export Controls. Customer is solely responsible for obtaining all necessary licenses or authorizations relating to the export, re-export, sale, lease or transfer of the Software or Documentation and for ensuring compliance with the requirements of such licenses or authorizations.

7) Warranties and Remedies.

a) Software Warranties- Cray Software. Cray warrants that, for a period of thirty (30) days (the "**Warranty Period**") following delivery of an item of Cray Software,

- i) the operation of the Cray Software will substantially conform to the description of operation stated in its Documentation (the "**Use Warranty**"); and
- ii) the Cray Software will not contain any viruses, worms, trojan horses, or other malicious or destructive code designed by Cray to allow unauthorized intrusion upon,

disabling of, or erasure of the Cray Software (the "**Virus Warranty**"). Nothing in this paragraph; however, will prevent or prohibit the use of a software key or similar device to manage or authorize appropriate use of the software.

b) System Warranties Excluded. Systems are warranted in accordance with the warranty terms contained in the separate agreement or Order for such System and are excluded from the warranty provisions of this EULA.

c) Remedies. Any breach of the foregoing warranties must be reported by Customer to Cray during the applicable Warranty Period. Customer's sole and exclusive remedy and Cray's sole obligation for any such breach shall be as follows:

- i) For a breach of the *Use Warranty* that impacts the use of the impacted Software, Cray will use commercially reasonable efforts to provide such patches, modifications, or other remedies as are reasonably necessary to resolve the nonconformity.
- ii) For a breach of the *Virus Warranty*, Cray shall replace the Software with a copy that is in conformance with the *Virus Warranty*.

d) Warranty Exclusions. The warranties set forth in this Section shall not apply to any nonconformity: i) that Cray cannot recreate after exercising commercially reasonable efforts to attempt to do so; ii) caused by misuse of the applicable Software or by using the Software in a manner that is inconsistent with this EULA or the Documentation; or iii) arising from the modification of the Software by anyone other than Cray.

e) Warranty Disclaimer. THE EXPRESS WARRANTIES AND REMEDIES SET FORTH IN THIS SECTION ARE THE ONLY WARRANTIES AND REMEDIES PROVIDED BY CRAY HEREUNDER. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, ALL OTHER WARRANTIES OR REMEDIES ARE EXCLUDED, WHETHER EXPRESS OR IMPLIED, ORAL OR WRITTEN, INCLUDING ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR ANY PARTICULAR PURPOSE, NON-INFRINGEMENT, SATISFACTORY QUALITY, AND ANY WARRANTIES ARISING FROM USAGE OF TRADE OR COURSE OF DEALING OR PERFORMANCE. CRAY DOES NOT WARRANT UNINTERRUPTED OR ERROR-FREE OPERATION OF THE SOFTWARE. THE SOFTWARE IS

LICENSED TO CUSTOMER ON AN "AS-IS" BASIS AND CRAY HAS NO RESPONSIBILITY OR LIABILITY FOR ANY LOSS OR DAMAGE ARISING OUT OF OR RELATED TO THE USE, MAINTENANCE, OR PERFORMANCE OF THE SOFTWARE (INCLUDING LIABILITY FOR NEGLIGENCE) EXCEPT LOSS OR DAMAGE ATTRIBUTABLE SOLELY TO THE GROSS NEGLIGENCE OR WILLFUL MISCONDUCT OF CRAY. EXCEPT AS MAY BE PROVIDED IN THIS AGREEMENT, IN NO EVENT WILL CRAY'S LIABILITY INCLUDE ANY SPECIAL, INDIRECT, INCIDENTAL, CONSEQUENTIAL, OR PUNITIVE LOSSES OR DAMAGES, INCLUDING BUT NOT LIMITED TO LOSS OF REVENUE, LOSS OF ACTUAL OR ANTICIPATED PROFITS, LOSS OF BUSINESS, LOSS OF CONTRACTS, LOSS OF GOODWILL OR REPUTATION, LOSS OF, DAMAGE TO OR CORRUPTION OF DATA EVEN IF CRAY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH POTENTIAL LOSS OR DAMAGE AND WHETHER ARISING IN OR FOR BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), BREACH OF STATUTORY DUTY, OR OTHERWISE.

8) High-Risk Disclaimer. Customer understands and acknowledges that the Software is not fault-tolerant and is not designed or intended for use in any high-risk or hazardous environment, including without limitation, the operation of nuclear power facilities, aircraft navigation, air traffic control, life support machines, weapons systems, or any other application where the failure or malfunction of any product can reasonably be expected to result in death, personal injury, severe property damage or severe environmental harm. Accordingly: a) Customer should not use the Software in any such environment; b) any use of the Software by Customer in any such environment is at Customer's own risk; c) Cray and its suppliers shall not be liable to Customer in any way for use of the Software in any such environment; and d) Cray makes no warranties or assurances, express or implied, regarding use of the Software in any such environment.

9) Infringement Indemnity.

a) Subject to the provisions of Section 7), Cray shall, at its own expense, defend any suit instituted against Customer to the extent that such claim is derived from an allegation that any Cray Software licensed to Customer hereunder constitutes an infringement of any patent issued by the United States, Canada, any EU state, or Japan, or copyright or violation or misappropriation of any third party trade secret and will indemnify Customer against

any award of damage and costs made against Customer by a final judgment of a court of last resort if it is determined therein that any such Cray Software constitutes an infringement of any patent or copyright or violation or misappropriation of any third party trade secret, provided that Customer gives Cray prompt notice in writing of any claims of infringement, violation or misappropriation, and permits Cray through Cray's counsel to defend the same and gives Cray all available information, assistance and authority as reasonably requested by Cray, and with costs of such assistance reimbursed by Cray, to enable Cray to assume such defense. To the extent not prohibited by applicable and controlling law, Cray will have control of the defense of any such suit, including appeals from any judgment therein and any negotiations for the settlement or compromise thereof, with full authority to enter into a binding settlement or compromise so long as such settlement does not include an admission of wrongdoing on the part of Customer or the waiver or release of any right of Customer except the settlement of the claim in issue. Cray will not be responsible for any compromise or settlement made or expense incurred without its consent. Should any Cray Software become, or in Cray's reasonable opinion be likely to become, the subject of any claim of infringement of any United States copyright or violation or misappropriation of any third party trade secret, Cray will have the right, at Cray's option and expense, to i) procure for Customer the right to continue using the Cray Software; ii) replace or modify the infringing Cray Software so that it becomes non-infringing; or iii) terminate Customer's right to use the infringing Cray Software upon which termination Customer agrees to promptly return the infringing Cray Software to Cray. Upon receipt of the returned Cray Software Cray will refund Customer's fees applicable to such Cray Software, reduced for depreciation determined on a straight-line five year basis.

b) Cray will have no obligation or liability to Customer with respect to any claim of infringement which results from: i) the combination or use of Cray Software with any other product, program or device not provided by Cray if such infringement would not have arisen but for such use or combination; ii) any modification by Customer or any third party of any Cray Software; iii) any specifications, designs or instructions provided to Cray by or on behalf of Customer; iv) the failure to promptly install an

update or revision provided by Cray, if such infringement could have been avoided by the use of such update or revision; or v) use of the Software in a manner not in conformance with Cray published specifications.

10) Governing Law and Venue. This EULA shall be governed by and construed in accordance with the laws of the State of Washington, without regard to the conflict of laws provisions thereof. Any action seeking enforcement of this EULA or any provision hereof shall be brought exclusively in the state or federal courts located in King County, Washington. Each party hereby agrees to submit to the jurisdiction of such courts and, to the extent permitted under applicable law, waives any objection to such jurisdiction including, without limitation, claims relating to lack of personal or subject matter jurisdiction and *forum non conveniens*. Further, each party agrees to waive any requests for a trial by jury.

11) Assignment. Except as otherwise set forth herein, Customer shall not, in whole or part, assign or transfer any part of this EULA, the licenses granted under this EULA or any other rights, interest or obligations hereunder, whether voluntarily, by contract, by operation of law or otherwise without the prior written consent of Cray. Any attempted transfer or assignment by Customer that is not permitted by this EULA shall be null and void.

12) Severability. If any provision of this EULA is ever held by a court of competent jurisdiction to be contrary to law, such provision will be enforced to the maximum extent permissible by law to effect the intent of the parties and the remaining provisions of this EULA will remain in full force and effect. Notwithstanding the foregoing, the terms of this EULA that limit, disclaim, or exclude warranties, remedies or damages are intended by the parties to be independent and remain in effect despite the failure or unenforceability of an agreed remedy. The parties have relied on the limitations and exclusions set forth in this EULA in determining whether to enter into it.

13) United States Government Rights. The Software is a "commercial item" as that term is defined at FAR 2.101. If Customer or user is the US Federal Government or any party acting on behalf of the US Federal Government (collectively "US Government"), Cray provides the Software, technology, technical data, and/or any related services in accordance with the following: If acquired

by or on behalf of any Executive Agency (other than an agency within the Department of Defense (DoD), the Government acquires, in accordance with FAR 12.211 (Technical Data) and FAR 12.212 (Computer Software), only those rights in technical data and software specifically defined in this EULA. If acquired by or on behalf of any Executive Agency within the DoD, the Government acquires, in accordance with DFARS 227.7202-3 (*Rights in commercial computer software or commercial computer software documentation*), only those rights in technical data and software customarily provided in this EULA. In addition, DFARS 252.227-7015 (*Technical Data – Commercial Items*) applies to technical data acquired by DoD agencies. Any Federal Legislative or Judicial Agency shall obtain only those rights in technical data and software customarily provided to the public as defined in this EULA. W

14) Notices. All notices provided hereunder shall be in writing and addressed to the legal department of the respective party or to such other address as may be specified in an Order or in writing by either of the parties to the other in accordance with this Section. Any notice to Cray that asserts any default by Cray under this EULA must be delivered to the attention of the Office of the General Counsel, Cray Inc., at 901 Fifth Avenue, Suite 1000, Seattle, WA 98164, USA. Except as may be expressly permitted herein, notices may be delivered personally, sent via a nationally recognized courier or overnight delivery service, or mailed by first class mail, postage prepaid. All notices, requests, demands or communications shall be deemed effective upon personal delivery or, if sent by mail, four days following deposit in the mail in accordance with this paragraph.

15) Waiver. Performance of any obligation required by a party hereunder may be waived only by a written waiver signed by an authorized representative of the other party, which waiver shall be effective only with respect to the specific obligation described therein. Any waiver or failure to enforce any provision of this EULA on one occasion will not be deemed a waiver of any other provision or of such provision on any other occasion.

16) Force Majeure. Each party will be excused from performance for any period during which, and to the extent that, it is prevented from performing any obligation or service as a result of causes beyond its reasonable control, and without its fault or negligence, including without limitation, acts of God,

strikes, lockouts, riots, acts of war, epidemics, communication line failures, power failures, and delays in Cray obtaining components from its suppliers. For added certainty, this Section shall not operate to change, delete, or modify any of the parties' obligations under this EULA, rather, it only provides for excuse of a delay in performance of such obligations.

17) Headings. The headings used in this EULA are solely for the convenience of the Parties, and shall not be used in the interpretation of the text of this EULA. Each Party has read and agreed to the specific language of this EULA; therefore no conflict, ambiguity, or doubtful interpretation shall be construed against the drafter.

18) Entire Agreement. The terms and conditions stated in this EULA constitute the complete and exclusive statement of the agreement between Customer and Cray, and supersede all prior oral and written statements of any kind whatsoever made by either party or their representatives. Any order form used by Customer in connection with this EULA will be considered to have its pre-printed clauses and statements deleted and such terms are expressly rejected by Cray. Any waivers or amendments, to be effective, must be in writing, signed by both parties.

END USER LICENSE AGREEMENT RED HAT® ENTERPRISE LINUX® AND RED HAT APPLICATIONS



PLEASE READ THIS END USER LICENSE AGREEMENT CAREFULLY BEFORE USING SOFTWARE FROM RED HAT. BY USING RED HAT SOFTWARE, YOU SIGNIFY YOUR ASSENT TO AND ACCEPTANCE OF THIS END USER LICENSE AGREEMENT AND ACKNOWLEDGE YOU HAVE READ AND UNDERSTAND THE TERMS. AN INDIVIDUAL ACTING ON BEHALF OF AN ENTITY REPRESENTS THAT HE OR SHE HAS THE AUTHORITY TO ENTER INTO THIS END USER LICENSE AGREEMENT ON BEHALF OF THAT ENTITY. IF YOU DO NOT ACCEPT THE TERMS OF THIS AGREEMENT, THEN YOU MUST NOT USE THE RED HAT SOFTWARE. THIS END USER LICENSE AGREEMENT DOES NOT PROVIDE ANY RIGHTS TO RED HAT SERVICES SUCH AS SOFTWARE MAINTENANCE, UPGRADES OR SUPPORT. PLEASE REVIEW YOUR SERVICE OR SUBSCRIPTION AGREEMENT(S) THAT YOU MAY HAVE WITH RED HAT OR OTHER AUTHORIZED RED HAT SERVICE PROVIDERS REGARDING SERVICES AND ASSOCIATED PAYMENTS.

This end user license agreement (“EULA”) governs the use of any of the versions of Red Hat Enterprise Linux, certain other Red Hat software applications that include or refer to this license, and any related updates, source code, appearance, structure and organization (the “Programs”), regardless of the delivery mechanism.

- 1. License Grant.** Subject to the following terms, Red Hat, Inc. (“Red Hat”) grants to you a perpetual, worldwide license to the Programs (most of which include multiple software components) pursuant to the GNU General Public License v.2. The license agreement for each software component is located in the software component’s source code and permits you to run, copy, modify, and redistribute the software component (subject to certain obligations in some cases), both in source code and binary code forms, with the exception of (a) certain binary only firmware components and (b) the images identified in Section 2 below. The license rights for the binary only firmware components are located with the components themselves. This EULA pertains solely to the Programs and does not limit your rights under, or grant you rights that supersede, the license terms of any particular component.
- 2. Intellectual Property Rights.** The Programs and each of their components are owned by Red Hat and other licensors and are protected under copyright law and under other laws as applicable. Title to the Programs and any component, or to any copy, modification, or merged portion shall remain with Red Hat and other licensors, subject to the applicable license. The “Red Hat” trademark and the “Shadowman” logo are registered trademarks of Red Hat in the U.S. and other countries. This EULA does not permit you to distribute the Programs or their components using Red Hat’s trademarks, regardless of whether the copy has been modified. You may make a commercial redistribution of the Programs only if (a) permitted under a separate written agreement with Red Hat authorizing such commercial redistribution, or (b) you remove and replace all occurrences of Red Hat trademarks. Modifications to the software may corrupt the Programs. You should read the information found at <http://www.redhat.com/about/corporate/trademark/> before distributing a copy of the Programs.
- 3. Limited Warranty.** Except as specifically stated in this Section 3, a separate agreement with Red Hat, or a license for a particular component, **to the maximum extent permitted under applicable law, the Programs and the components are provided and licensed “as is” without warranty of any kind, expressed or implied, including the implied warranties of merchantability, non-infringement or fitness for a particular purpose.** Red Hat warrants that the media on which the Programs and the components are provided will be free from defects in materials and manufacture under normal use for a period of 30 days from the date of delivery to you. Neither Red Hat nor its affiliates warrants that the functions contained in the Programs will meet your requirements or that the operation of the Programs will be entirely error free, appear or perform precisely as described in the accompanying documentation, or comply with regulatory requirements. **This warranty extends only to the party that purchases subscription services for the Programs from Red Hat and/or its affiliates or a Red Hat authorized distributor.**
- 4. Limitation of Remedies and Liability.** To the maximum extent permitted by applicable law, your exclusive remedy under this EULA is to return any defective media within 30 days of delivery along with a copy of your payment receipt and Red Hat, at its option, will replace it or refund the money you paid for the media. **To the maximum extent permitted under applicable law, under no circumstances will Red Hat, its affiliates, any Red Hat authorized distributor, or the licensor of any component provided to you under this EULA be liable to you for any incidental or consequential damages, including lost profits or lost savings arising out of the use or inability to use the Programs or any component, even if Red Hat, its affiliates, an authorized distributor and/or licensor has been advised of the possibility of such damages. In no event shall Red Hat’s or its affiliates’ liability, an authorized distributor’s liability or the liability of the licensor of a component provided to you under this EULA exceed the amount that you paid to Red Hat for the media under this EULA.**
- 5. Export Control.** As required by the laws of the United States and other countries, you represent and warrant that you: (a) understand that the Programs and their components may be subject to export controls under the U.S. Commerce Department’s Export Administration Regulations (“EAR”); (b) are not located in a prohibited destination country under the EAR or U.S. sanctions regulations (currently Cuba, Iran, Iraq, North Korea, Sudan and Syria, subject to change as posted by the United States government); (c) will not export, re-export, or transfer the Programs to any prohibited destination or persons or entities on the U.S. Bureau of Industry and Security Denied Parties List or Entity List, or the U.S. Office of Foreign Assets Control list of Specially Designated Nationals and Blocked Persons, or any similar lists maintained by other countries, without the necessary export license(s) or authorization(s); (d) will not use or transfer the Programs for use in connection with any nuclear, chemical or biological weapons, missile technology, or military end-uses where prohibited by an applicable arms embargo, unless authorized by the relevant government agency by regulation or specific license; (e) understand and agree that if you are in the United States and export or transfers the Programs to eligible end users, you will, to the extent required by EAR Section 740.17(e), submit semi-annual reports to the Commerce Department’s Bureau of Industry and Security, which include the name and address (including country) of each transferee; and (f) understand that countries including the United States may restrict the import, use, or export of encryption products (which may include the Programs and the components) and agree that you shall be solely responsible for compliance with any such import, use, or export restrictions.

6. **Third Party Programs.** Red Hat may distribute third party software programs with the Programs that are not part of the Programs. These third party programs are not required to run the Programs, are provided as a convenience to you, and are subject to their own license terms. The license terms either accompany the third party software programs or can be viewed at <http://www.redhat.com/licenses/thirdparty/eula.html>. If you do not agree to abide by the applicable license terms for the third party software programs, then you may not install them. If you wish to install the third party software programs on more than one system or transfer the third party software programs to another party, then you must contact the licensor of the applicable third party software programs.
7. **General.** If any provision of this EULA is held to be unenforceable, the enforceability of the remaining provisions shall not be affected. Any claim, controversy or dispute arising under or relating to this EULA shall be governed by the laws of the State of New York and of the United States, without regard to any conflict of laws provisions. The rights and obligations of the parties to this EULA shall not be governed by the United Nations Convention on the International Sale of Goods.

Copyright © 2010 Red Hat, Inc. All rights reserved. "Red Hat" and the Red Hat "Shadowman" logo are registered trademarks of Red Hat, Inc. "Linux" is a registered trademark of Linus Torvalds. All other trademarks are the property of their respective owners.

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
- b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

one line to give the program's name and an idea of what it does.
Copyright (C) yyyy name of author

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of

MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) year name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details
type `show w'. This is free software, and you are welcome
to redistribute it under certain conditions; type `show c'
for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items-- whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright
interest in the program `Gnomovision'
(which makes passes at compilers) written
by James Hacker.
```

```
signature of Ty Coon, 1 April 1989
Ty Coon, President of Vice
```

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the [GNU Lesser General Public License](#) instead of this License.

GNU GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc. <<https://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program--to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS

0. Definitions.

“This License” refers to version 3 of the GNU General Public License.

“Copyright” also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

“The Program” refers to any copyrightable work licensed under this License. Each licensee is addressed as “you”. “Licensees” and “recipients” may be individuals or organizations.

To “modify” a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a “modified version” of the earlier work or a work “based on” the earlier work.

A “covered work” means either the unmodified Program or a work based on the Program.

To “propagate” a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To “convey” a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays “Appropriate Legal Notices” to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

1. Source Code.

The “source code” for a work means the preferred form of the work for making modifications to it. “Object code” means any non-source form of a work.

A “Standard Interface” means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The “System Libraries” of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A “Major Component”, in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The “Corresponding Source” for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work's System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

3. Protecting Users' Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a) The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b) The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to “keep intact all notices”.
- c) You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.
- d) If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an “aggregate” if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

- a) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.
- b) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.
- c) Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.
- d) Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.
- e) Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A “User Product” is either (1) a “consumer product”, which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, “normally used” refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

“Installation Information” for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source. The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

7. Additional Terms.

“Additional permissions” are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a) Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b) Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- c) Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or
- d) Limiting the use for publicity purposes of names of licensors or authors of the material; or
- e) Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- f) Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered “further restrictions” within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to

notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance. However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An “entity transaction” is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party's predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

11. Patents.

A “contributor” is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor's “contributor version”.

A contributor's “essential patent claims” are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, “control” includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor's essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a “patent license” is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To “grant” such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. “Knowingly relying” means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient's use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is “discriminatory” if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

12. No Surrender of Others' Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

13. Use with the GNU Affero General Public License.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License “or any later version” applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

15. Disclaimer of Warranty.

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. Limitation of Liability.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it does.>
Copyright (C) <year> <name of author>
```

```
This program is free software: you can redistribute it and/or modify
it under the terms of the GNU General Public License as published by
the Free Software Foundation, either version 3 of the License, or
(at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License
along with this program. If not, see <https://www.gnu.org/licenses/>.
```

Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short notice like this when it starts in an interactive mode:

```
<program> Copyright (C) <year> <name of author>
This program comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
This is free software, and you are welcome to redistribute it
under certain conditions; type `show c' for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, your program's commands might be different; for a GUI interface, you would use an “about box”.

You should also get your employer (if you work as a programmer) or school, if any, to sign a “copyright disclaimer” for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see <<https://www.gnu.org/licenses/>>.

The GNU General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read <<https://www.gnu.org/licenses/why-not-lgpl.html>>.

GNU LIBRARY GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1991 Free Software Foundation, Inc.
51 Franklin St, Fifth Floor, Boston, MA 02110-1301, USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

[This is the first released version of the library GPL. It is
numbered 2 because it goes with version 2 of the ordinary GPL.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Library General Public License, applies to some specially designated Free Software Foundation software, and to any other libraries whose authors decide to use it. You can use it for your libraries, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library, or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link a program with the library, you must provide complete object files to the recipients so that they can relink them with the library, after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

Our method of protecting your rights has two steps: (1) copyright the library, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the library.

Also, for each distributor's protection, we want to make certain that everyone understands that there is no warranty for this free library. If the library is modified by someone else and passed on, we want its recipients to know that what they have is not the original version, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that companies distributing free software will individually obtain patent licenses, thus in effect transforming the program into proprietary software. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License, which was designed for utility programs. This license, the GNU Library General Public License, applies to certain designated libraries. This license is quite different from the ordinary one; be sure to read it in full, and don't assume that anything in it is the same as in the ordinary license.

The reason we have a separate public license for some libraries is that they blur the distinction we usually make between modifying or adding to a program and simply using it. Linking a program with a library, without changing the library, is in some sense simply using the library, and is analogous to running a utility program or application program. However, in a textual and legal sense, the linked executable is a combined work, a derivative of the original library, and the ordinary General Public License treats it as such.

Because of this blurred distinction, using the ordinary General Public License for libraries did not effectively promote software sharing, because most developers did not use the libraries. We concluded that weaker conditions might promote sharing better.

However, unrestricted linking of non-free programs would deprive the users of those programs of all benefit from the free status of the libraries themselves. This Library General Public License is intended to permit developers of non-free programs to use free libraries, while preserving your freedom as a user of such programs to change the free libraries that are incorporated in them. (We have not seen how to achieve this as regards changes in header files, but we have achieved it as regards changes in the actual functions of the Library.) The hope is that this will lead to faster development of free libraries.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, while the latter only works together with the library.

Note that it is possible for a library to be covered by the ordinary General Public License rather than by this special one.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Library General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- **a)** The modified work must itself be a software library.
- **b)** You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- **c)** You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- **d)** If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also compile or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- **a)** Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)
- **b)** Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- **c)** If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- **d)** Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

- **a)** Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
- **b)** Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Library General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A

FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

one line to give the library's name and an idea of what it does.
Copyright (C) year name of author

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Library General Public License as published by the Free Software Foundation; either version 2 of the license, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Library General Public License for more details.

You should have received a copy of the GNU Library General Public License along with this library; if not, write to the Free Software Foundation, Inc., 51 Franklin St, Fifth Floor, Boston, MA 02110-1301, USA.

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the library `Frob' (a library for tweaking knobs) written by James Random Hacker.

signature of Ty Coon, 1 April 1990
Ty Coon, President of Vice

That's all there is to it!

GNU AFFERO GENERAL PUBLIC LICENSE

Version 3, 19 November 2007

Copyright © 2007 Free Software Foundation, Inc. <<https://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The GNU Affero General Public License is a free, copyleft license for software and other kinds of works, specifically designed to ensure cooperation with the community in the case of network server software.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, our General Public Licenses are intended to guarantee your freedom to share and change all versions of a program--to make sure it remains free software for all its users.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

Developers that use our General Public Licenses protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License which gives you legal permission to copy, distribute and/or modify the software.

A secondary benefit of defending all users' freedom is that improvements made in alternate versions of the program, if they receive widespread use, become available for other developers to incorporate. Many developers of free software are heartened and encouraged by the resulting cooperation. However, in the case of software used on network servers, this result may fail to come about. The GNU General Public License permits making a modified version and letting the public access it on a server without ever releasing its source code to the public.

The GNU Affero General Public License is designed specifically to ensure that, in such cases, the modified source code becomes available to the community. It requires the operator of a network server to provide the source code of the modified version running there to the users of that server. Therefore, public use of a modified version, on a publicly accessible server, gives the public access to the source code of the modified version.

An older license, called the Affero General Public License and published by Affero, was designed to accomplish similar goals. This is a different license, not a version of the Affero GPL, but Affero has released a new version of the Affero GPL which permits relicensing under this license.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS

0. Definitions.

"This License" refers to version 3 of the GNU Affero General Public License.

"Copyright" also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

"The Program" refers to any copyrightable work licensed under this License. Each licensee is addressed as "you". "Licensees" and "recipients" may be individuals or organizations.

To "modify" a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a "modified version" of the earlier work or a work "based on" the earlier work.

A "covered work" means either the unmodified Program or a work based on the Program.

To "propagate" a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To "convey" a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays "Appropriate Legal Notices" to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

1. Source Code.

The "source code" for a work means the preferred form of the work for making modifications to it. "Object code" means any non-source form of a work.

A "Standard Interface" means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The "System Libraries" of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A "Major Component", in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The "Corresponding Source" for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work's System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

3. Protecting Users' Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice, keep intact all notices stating that this License and any non-

permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a) The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b) The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to "keep intact all notices".
- c) You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.
- d) If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

- a) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.
- b) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.
- c) Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.
- d) Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.
- e) Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A "User Product" is either (1) a "consumer product", which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, "normally used" refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

"Installation Information" for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source. The

information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

7. Additional Terms.

"Additional permissions" are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a) Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b) Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- c) Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or
- d) Limiting the use for publicity purposes of names of licensors or authors of the material; or
- e) Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- f) Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered "further restrictions" within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance. However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An "entity transaction" is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party's predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

11. Patents.

A "contributor" is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor's "contributor version".

A contributor's "essential patent claims" are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, "control" includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor's essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a "patent license" is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To "grant" such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. "Knowingly relying" means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient's use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is "discriminatory" if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third

party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

12. No Surrender of Others' Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

13. Remote Network Interaction; Use with the GNU General Public License.

Notwithstanding any other provision of this License, if you modify the Program, your modified version must prominently offer all users interacting with it remotely through a computer network (if your version supports such interaction) an opportunity to receive the Corresponding Source of your version by providing access to the Corresponding Source from a network server at no charge, through some standard or customary means of facilitating copying of software. This Corresponding Source shall include the Corresponding Source for any work covered by version 3 of the GNU General Public License that is incorporated pursuant to the following paragraph.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the work with which it is combined will remain governed by version 3 of the GNU General Public License.

14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU Affero General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU Affero General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU Affero General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU Affero General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

15. Disclaimer of Warranty.

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. Limitation of Liability.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it does.>  
Copyright (C) <year> <name of author>
```

```
This program is free software: you can redistribute it and/or modify  
it under the terms of the GNU Affero General Public License as  
published by the Free Software Foundation, either version 3 of the  
License, or (at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,  
but WITHOUT ANY WARRANTY; without even the implied warranty of  
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the  
GNU Affero General Public License for more details.
```

```
You should have received a copy of the GNU Affero General Public License  
along with this program. If not, see <https://www.gnu.org/licenses/>.
```

Also add information on how to contact you by electronic and paper mail.

If your software can interact with users remotely through a computer network, you should also make sure that it provides a way for users to get its source. For example, if your program is a web application, its interface could display a "Source" link that leads users to an archive of the code. There are many ways you could offer source, and different solutions will be better for different programs; see section 13 for the specific requirements.

You should also get your employer (if you work as a programmer) or school, if any, to sign a "copyright disclaimer" for the program, if necessary. For more information on this, and how to apply and follow the GNU AGPL, see <<https://www.gnu.org/licenses/>>.

The wxWindows Library Licence (WXwindows)

wxWindows Library License (WXwindows)

(Disclaimer)

Copyright (c) 1998 Julian Smart, Robert Roebling [, ...]

Everyone is permitted to copy and distribute verbatim copies of this licence document, but changing it is not allowed.

WXWINDOWS LIBRARY LICENCE
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Library General Public Licence as published by the Free Software Foundation; either version 2 of the Licence, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Library General Public Licence for more details.

You should have received a copy of the GNU Library General Public Licence along with this software, usually in a file named COPYING.LIB. If not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA.

EXCEPTION NOTICE

1. As a special exception, the copyright holders of this library give permission for additional uses of the text contained in this release of the library as licenced under the wxWindows Library Licence, applying either version 3 of the Licence, or (at your option) any later version of the Licence as published by the copyright holders of version 3 of the Licence document.
2. The exception is that you may use, copy, link, modify and distribute under the user's own terms, binary object code versions of works based on the Library.
3. If you copy code from files distributed under the terms of the GNU General Public Licence or the GNU Library General Public Licence into a copy of this library, as this licence permits, the exception does not apply to the code that you add in this way. To avoid misleading anyone as to the status of such modified files, you must delete this exception notice from such code and/or adjust the licensing conditions notice accordingly.
4. If you write modifications of your own for this library, it is your choice whether to permit this exception to apply to your modifications. If you do not wish that, you must delete the exception notice from such code and/or adjust the licensing conditions notice accordingly.

Annex 6

Conditions of Warranty Provision

3. Warranty

- 3.1 Warranty shall be provided to the Work in the scope stipulated in par. 11.2 hereto. The below provisions stipulate the conditions of warranty provision.
- 3.2 The warranty shall be provided by the Supplier at the place of the System installation, as regular and reactive warranty servicing. The warranty may be additionally provided for activities that do not require the presence at the place of the System installation, remotely using the remote access means. The activities that require the presence at the place of the System installation, shall be provided by the Supplier. The warranty cannot be provided e.g. by mere sending of a spare part and instructions for its replacement.
- 3.3 The Supplier shall provide for the sole contact point for reporting of defects and requirements for servicing of all supplied systems. The possibility to report defects shall be provided for 24 hours a day (around-the-clock) and possibility to enter a request for regular maintenance services on working days from 8:00 am to 4:00 pm.
- 3.4 Any warranty repairs shall be made in a manner reducing the impact on the operations and availability of the System.
- 3.5 If the due provision of the warranty requires shutdowns of the System or its significant part, these shall be made no more than twice in one quarter of the calendar year, unless the binding legislation stipulates otherwise. Such shutdowns, made following a time schedule, in a time window and scope approved by VSB-TUO, shall not be treated as a defect or continuance of defect.

4. After-sales services – regular

- 4.1 Regular after-sales services of the Large Cluster facility primarily involves
 - warranty maintenance services and services stipulated by the binding legislation;
 - warranty maintenance services and services determined by the manufacturer of the facility; and
 - provision and performing of updates to software(hereinafter the “**Regular Maintenance Services**”).
 - 4.2 The Regular Maintenance Services and maintenance required by warranty conditions of facilities and systems. The Regular Maintenance Services shall be provided by Regular servicing. The Regular Maintenance Services must be provided at the place of installation.
 - 4.3 Regular Maintenance Services shall be provided in accordance with the time schedule, in a time window and the scope approved by VSB-TUO, on working days between 8:00 am and 4:00 pm. VSB-TUO shall be informed on the result of the Regular servicing in writing.
 - 4.4 The time schedule of Regular servicing shall be an integral part of the user operational documentation.
 - 4.5 The Regular Maintenance Services shall additionally involve provision of new versions of the supplied software, availability of software updates and patches. Over the warranty period, the Client shall be provided with free of charge access to new versions, security patches and updates to programmes of systems by the Supplier.
-

4.6 The Regular Maintenance Services shall additionally include regular updates to programs of the system and its setting by the Supplier. Programs of the system include all supplied software (firmware, hardware drivers, operating systems, software stacks, application software, technology management software, etc.). The Supplier shall be bound to primarily perform updates that are intended to increase security, reliability, elimination of functional or performance shortcomings. Updates of the programs shall be performed with the periodicity of no more than 9 months, unless VSB-TUO stipulates a longer period or explicitly rejects performance of certain or all updates. Updates intended for the elimination of serious problems shall be performed depending on their seriousness without unnecessary delays. Updates to programs shall be coordinated with VSB-TUO and shall be subject to its approval.

4.7 **All costs of the Regular Maintenance Services (material, work, transport, waste disposal, etc.) shall be included in the price for the Work in accordance with par. 12.1 hereto.**

5. After-sales services – reactive

5.1 Reactive after-sales shall provide for the removal of defects of category A, B and C defined in par. 11.4 hereto (hereinafter the “**Reactive Services**”).

5.2 The Reactive Services shall be provided pursuant to:

- A defect identified by remote monitoring, if it is provided – VSB -TUO shall be immediately informed of defects identified in this manner by a pre-determined method. The servicing team of the Supplier shall inform VSB-TUO on further planned course of servicing.
- Defects identified directly by the Client – defects identified in this manner shall be reported by the Client to the Supplier any time using a pre-determined method. The report shall contain basic technical and contact information, in general: type and identification device, description of the defects, telephone and e-mail of the contact person of the Client.

5.3 The determination of the method of transfer of information in the Reactive Services shall be an integral part of the user operational documentation.

5.4 **For the avoidance of doubts, the price for the Reactive Services shall be included in the price for the Work in accordance with par. 12.1 hereto.**

Annex 7

Authorised Persons

On behalf of VSB-TUO:

In contractual and commercial matters:

Name and surname	[REDACTED]
Address	17. listopadu 2172/15, 708 00 Ostrava-Poruba
Email	[REDACTED]
Telephone	[REDACTED]

In technical and implementation matters:

Name and surname	[REDACTED]
Address	17. listopadu 2172/15, 708 00 Ostrava-Poruba
Email	[REDACTED]
Telephone	[REDACTED]

On behalf of EuroHPC JU:

In contractual and commercial matters:

Name and surname	[REDACTED]
Address	12E, rue Guillaume Kroll, L-1882 Luxembourg
Email	[REDACTED]
Telephone	[REDACTED]

In technical and implementation matters:

Name and surname	[REDACTED]
Address	12E, rue Guillaume Kroll, L-1882 Luxembourg
Email	[REDACTED]
Telephone	[REDACTED]

On behalf of the Supplier:

In contractual matters:

Name and surname	[REDACTED]
Address	Za Brumlovkou 1559/5, Praha 4 – Michle, 140 00
Email	[REDACTED]
Telephone	[REDACTED]

In commercial, technical and implementation matters:

Name and surname	[REDACTED]
Address	Za Brumlovkou 1559/5, Praha 4 – Michle, 140 00
Email	[REDACTED]
Telephone	[REDACTED]

Annex 8

Project Management and Detailed Rules of Change Management

Project Management

1. The Supplier is charged to prepare, manage and supervise the supply of the Work.
2. A division of the projects to stages and individual phases shall be consulted and approved by the Client and shall adhere to the Time Schedule.
3. The Supplier shall designate a person responsible for the overall coordination of the project and each phase and technological section of the project.
4. The Client undertakes to cooperate with the Supplier, provide it with necessary information and support for successful completion of all project phases in the scope hereunder. The Client shall designate responsible persons for that. These persons shall be available to the Supplier only in the scope of their labour relationship (job content, availability, working hours).
5. The Supplier shall keep documentation of all activities and facts and the documentation shall be available throughout the project and shall be handed over to the Client.

Detailed rules of change management

1. Any request for change shall be presented to the other contracting party in writing, in a provable manner.
 2. The request for change shall become the change at the moment when approved.
 3. Every change shall be approved in writing by both contracting parties before it is made, by the "*Record of approval of a request for change*" document.
 4. The change shall not disrupt objectives of the project.
 5. Both contracting parties shall be obliged to make every effort in negotiation on the request for change in order to find a mutual agreement.
 6. A negotiation on a request for change may also result in its rejection, in such case, it shall be stated in the "*Record of approval of a request for change*" document in a clear and comprehensible manner why the contracting party asked for approval rejects the request for change.
 7. The "*Record of approval of a request for change*" document shall primarily include:
 - a) Identification of the contracting party asking for the change.
 - b) A description of the request for change, as precise as possible, including drawings, schemes and other particulars, if any, that describe the requested change in detail.
 - c) Statement of the impact of the change on the monitored project parameter.
 - d) The correction mechanisms allowing to the contracting party which incorporates the request for change, to withdraw from the request for change if there are circumstance
-

preventing the integration of the request for change that were not known at the time of approval of the request for change.

8. The monitored project parameter affected by the request for change shall change at the moment of the approval of the request for change in a manner described in the "*Record of approval of a request for change*" document.
 9. The resulting change shall be reflected in all documentation, including the already handed over documentation in the form of an update to the entire document (e.g. the Detailed Technical Specification, Operational Documentation, Documentation of Training etc.).
 10. If the situation requires, an amendment hereto shall be concluded pursuant to the approved request for change. The Client shall decide on the necessity to conclude an amendment hereto with definitive effect.
 11. The resulting change shall respect the principles of sound financial management, transparency, proportionality, equal treatment and non-discrimination.
-

Annex 9

Acceptance tests requirements

All acceptance tests shall be conducted with the participation of both parties.

Proposals for acceptance tests shall respect the requirements and information of the tender dossier.

Acceptance tests shall include:

- 1) Demonstration and confirmation of actual condition of the delivered devices, licenses, etc.
 - 2) Verification of integration into Data centre
 - a) Verification of the solution's placement in respect to serviceability of the installed devices and data centre equipment
 - b) Verification of the maximum power consumption
 - c) Verification of even load on the power supply phases
 - d) Verification of the cooling solution
 - e) Verification of the environment parameters (temperature, humidity)
 - f) Verification of the system availability by power supply circuit shutdown
 - g) Verification of the system availability by cooling circuit shutdown
 - 3) Verification of required and/or declared technical parameters of the system, in particular:
 - a) Computing performance of the compute partitions
 - b) Capacity and performance of SCRATCH storage
 - c) Capacity and performance of HOME storage
 - d) Capacity and performance of INFRA storage
 - e) Capacity and performance of PROJECT storage expansion
 - f) Backup capacity
 - 4) Verification of the Compute network throughput and efficiency
 - 5) Verification of the LAN network throughput
 - 6) Verification of the system/services availability
 - a) Verification of the data storages availability
 - b) Verification of the infrastructure services availability
 - c) Verification of the LAN network availability
 - 7) Verification of the remote visualization functionality
 - 8) Verification of the data storages features and accessibility
 - 9) Verification of the backup solution
 - a) Verification of the backup policy implementation
 - b) Data backup
 - c) Data recovery
 - i) Physical server recovery
 - ii) HOME storage file recovery
 - iii) INFRA storage file recovery
 - 10) Verification of the software installation and integration, in particular
 - a) Verification of the remote centralized server/node management
 - b) Verification of the remote device management, in particular disc arrays and network devices
-

- c) Verification of the remote server installation
 - d) Verification of the network boot
 - e) Verification of the image management
 - f) Verification of the scheduler
 - g) Verification of the name/uid, group/gid resolving and the authentication
 - h) Verification of the availability monitoring
 - i) Verification of the performance monitoring
 - j) Verification of the logging
 - k) Verification of the security
- 11) Verification of the PROJECT storage integration
 - 12) Verification of reliable, stable operation of the system (stability test) by full load using scheduler-executed performance tests for the period of three days

Availability verification of the IT systems solutions shall be conducted by forced shutdown of a random server/node of the solution, power supply circuit shutdown, power cable disconnection, LAN cable disconnection, removal of disc from the data storage, etc.

Each test shall be described in detail in terms of methodology, inputs, test procedure, impact on test area and desired results. Each test must be repeatable/reproducible.