simula

Request for Tender

Invitation to tender for
Phase 3 of the eX³ project
Additional compute nodes
for the heterogeneous high performance computing

Simula Research Laboratory AS eX³ ITT 04/2021

simula

simula

Abbreviations

AI	Artificial Intelligence
ARM	Advanced RISC Machines
B	Byte, synonymous with the ISO unit ‘octet’ and as such is 8 bits long.
B2F	Back-To-Front
BBU	Battery Backup Unit
BMC	Baseboard Management Controller
CM	Cluster Management
Core	A core is the processing unit within a CPU that receives instructions and performs calculations, or actions, based on those instructions.
CPU	Central Processing Unit, a chip this days consisting of multiple processing elements or cores.
CRU	Customer Replaceable Unit
CUDA	CUDA (Compute Unified Device Architecture), a parallel programming framework by graphics processor manufacturer Nvidia.
DIMM	Dual Inline Memory Module
DP	Dual-Processor
ECC	Error-correcting code memory (ECC-memory)
Firmware	A specific class of computer software that provides the low-level control for the device's specific hardware.
FPGA	Field Programmable Gate Array
GHS	Global Hot Spare
GPGPU	General Purpose Graphical Processing Unit
GPU	Graphical Processing Unit
HBA	Host Bus Adapter
HCA	Host Channel Adapter
HIC	Host Interface Card
HPC	High Performance Computing
IB	Infiniband Interconnect Network as described by the InfiniBand® Trade Association
IOPS	IO operations per second as measured on file-system level
IPMI	Intelligent Platform Management Interface protocol used by BMC
IPU	Intelligence Processing Unit
ISA	Instruction Set Architecture

LUN	Logical UNit, a set of disks combined to one logical unit
MD	Multiple-Devices aka Linux Software RAID
MI	Machine Intelligence
ML	Machine Learning
MPI	Message Passing Interface
NAS	Network Attached Storage
NBD	Next Business Day
NPU	Neural Processing Unit
NTB	Non-Transparent Bridging is part of PCIsig standard version 2 or later
NVDIMM	None-Volatile Dual Inline Memory Module
NVMe	NVM Express (NVMe) or Non-Volatile Memory Host Controller Interface Specification (NVMHCIS)
OCP	Open Compute Platform
OEM	Original Equipment Manufacturer
OpenCL	OpenCL (Open Computing Language) is a framework for writing programs that execute across heterogeneous platforms consisting of central processing units (CPUs), graphics processing units (GPUs), digital signal processors (DSPs), field-programmable gate arrays (FPGAs) and other processors or hardware accelerators
OpenMP	OpenMP (Open Multi-Processing) is an application programming interface (API) that supports multi-platform shared memory multiprocessing programming in C, C++, Fortran and several other programming languages.
PCIe	Peripheral Component Interconnect Express
PDU	Power Distribution Unit
PSU	Power Supply Unit
RAID	Redundant Array of {Inexpensive|Independent} Disks
RISC	Reduced Instruction Set Computer architecture
RPS	Redundant Power Supply
SSA-K	Norwegian Government Standard Contract for Acquisitions for Software and Hardware
top-bin	Refers to product binning in semiconductor device fabrication in which top-bin refers to chips that pass all tests including higher GHz speeds, lower voltage, more functional cores etc. and therefore can be sold at a higher price point.
TPU	Tensor Processing Unit
URL	Universal Resource Locator

simula

Terminology

This document aims to use simple terminology[1]. Some terms are used interchangeably.

Customer	Client, Buyer, procurer
Contractor	Vendor, Supplier
DIFI	Agency for Public Management and eGovernment (Norw.: Direktoratet for forvaltning og IKT)
DOFFIN	Doffin is a Norwegian database of public procurements.
ESPD	ESPD is a self-declaration form used in public procurement procedures by public buyers and businesses in the EU/ EEC.
FOA[2]	The Norwegian Public Procurement Regulation (Norw.: Forskrift om offentlige anskaffelser)
FVL	Norwegian Public Administration Act of 10 February 1967, No. 9
ITT	Invitation to tender
KOFA	The Norwegian Complaints Board for Public Procurement (Norw.: Klagenemnd for offentlige anskaffelser)
LOA[3]	Puplic Procurement Act (Norw.: Lov om Offentlige Anskaffelser)
Mercell	SRL uses the Mercell tender management tool to manage this procedure.
NDA	Non-Disclosure Agreement
offer	tender, proposal, bid, submission
RFP	Request for Price/Procurement
RFI	Request for Information
RFQ	Request for quotation/ quote
RFT	Request For Tender
RFx	RFI, RFQ, RFP or RFT
supplier	vendor, seller, contractor, potential supplier, respondent
tender	a generic term used to describe making an approach to market (‘going out to

1 https://www.anskaffelser.no/sites/anskaffelser/files/competition_and_contracts_text.pdf
2 https://lovdata.no/dokument/SF/forskrift/2016-08-12-974
3 https://lovdata.no/dokument/NL/lov/2016-06-17-73

	tender’)
TED	TED is the European database of public procurements.
SSA	The Norwegian Government Standard Terms and Conditions for IT Procurement (Norw: Statens standardavtaler for IT-anskaffelser)

simula

Table of Contents

Appendix 1: Customer requirement specification

simula

Appendix 1 of Purchase Agreement

Customer requirement specification

1 THE CUSTOMER'S PURPOSE AND REQUIREMENTS

1.2 About Simula Research Laboratory AS

Simula Research Laboratory conducts basic and applied research and provides education in selected fields within scientific computing, software engineering, machine learning, communication systems and cybersecurity, thereby contributing to innovation in society. Simula is organized as a limited company and is owned by the Ministry of Education and Research. For more information, please refer to Simula's Annual Report 2019[4].

1.3 Background Information On The eX³ project

1.2.1 Facilitating national research on next generation HPC

Computational infrastructures are essential for an increasing number of disciplines, and it is particularly important for research that requires demanding calculations or generates large amounts of data. The project "Experimental Infrastructure for Exploration of Exascale Computing"[5], or eX³, was initiated in 2018 to develop competence and techniques for successful exploitation of the coming generation supercomputers. Such exascale computers will be able to perform at a minimum of a billion billion (10^18) floating point operations per second. This new generation of high performance computing, which will also be needed to handle the exponentially increasing, complex data that is being generated in research and in applications, is a major target for international research.

Simula's eX³ project was one of the 19 projects selected for funding out of 92 applications submitted to the most recent round of the Research Council of Norway's (RCN) program for research infrastructures. The project was to build and run a national HPC infrastructure that will help prepare Norwegian researchers, big data users, data center management, and industry for the coming exascale era of computing. The upcoming exascale computing will rely on an intricate interplay between thousands of sophisticated processing nodes, each with a large number of cores, deep memory hierarchies and equipped with accelerators, organized in complex communication topologies. The aggregated level of complexity in a system designed for billion-way concurrency represents a major challenge with many features: How to program such computers? How to port existing code, and how to reach a satisfactory

4 https://www.simula.no/sites/default/files/sim_annual_20190401_en.pdf
5 https://www.ex3.simula.no/

level of reliability and efficiency while maintaining an acceptable energy footprint? The eX³ infrastructure provides a platform for researching answers to such questions, utilizing relevant technology albeit on a smaller scale than in a full-blown exascale production system.

1.2.2 Invitation to Tender

Simula Research Laboratory hereby invites interested contractors to an open competition for tendering the third phase of purchases of equipment for the national research infrastructure eX³. These additional computational nodes will augment the existing high performance computing (HPC) facility where researchers can test their hypothesis and port their code to various platforms.

This procurement document is based on The Norwegian Government's Standard Terms and Conditions for IT Procurement SSA-K standard. Appendix 0 attached is the contracts that both parties will sign.

Herein Appendix 1 we (the customer) specify the technical requirements for this purchase.

All rows shall be ticked (Yes or No)	Yes	No
Appendix 0: A00 Draft Contract–SSA-K 2015 English	**X**	
Appendix 1: Customer requirements specification	**X**	
Appendix 2: Contractor description of the deliverables		**X**
Appendix 3: Customer technical platform		**X**
Appendix 4: Delivery date and other deadlines		**X**
Appendix 5: Approval test		**X**
Appendix 6: Administrative provisions		
Appendix 7: Total price and pricing provisions	**X**	
Appendix 8: Changes to the general contractual wording		**X**
Appendix 9: Changes subsequent to the conclusion of the Agreement		**X**
Appendix 10: Licence terms and conditions for standard software and free software		**X**
Other Appendices:		**X**

Table 1: Difi SSA-K standard appendices.

It is noted that eX^3 already is an highly experimental facility set up to make state-of-the-art technology components available for academic research purposes, and that the purpose for this invitation to tender is as before to acquire the latest top-bin technology. In phase 3 we agument our system with the latest top-bin general purpose CPUs, GPGPUs and FPGA offload engines.

In order for new contractors to get a sense of what eX^3 project is about, please reference the original tender documents from phase 1 below. These tender documents for can be found on Doffin[6] and Mercell[7] and appendixes reference can found there.

Document	Size
A00 Draft Contract - SSA-K 2015 English.pdf	195 KB
A01-A12 Customer Requirements & all - ITT-eX³-Phase-1-Heterogeneous-HPC-system.pdf	1.66 MB
Rules of the Competition.pdf	184 KB
A07A Attachment Price and Requirement Matrices.ods	981 KB
A07A Attachment Price and Requirement Matrices.xlsx	992 KB
Errata-A01-A12 Customer Requirements & all - ITT-eX³-Phase-1-Heterogeneous-HPC-system-MR1-pg70.pdf	61 KB
Errata-Rules of the Competition-MR1-pg12.pdf	108 KB

Table 2: eX³ phase 1 ITT documents

1.3 SCOPE OF THIS TENDER

For the phase 3 procurement round, the eX³ project management group wants three main technology tracks. The suggestions are based on input and wishes from several users and project partners in the consortium, as well as our own evaluation of technology available in the market.

The eX³ project team has gathered technical and pricing information such that budget and expectations are set appropriately.

Please notice that we expect that the offered systems keep to the specification requested as close as possible, but appreciate that there are some variations dependent on the vendor's.

6 https://www.doffin.no/Notice/Details/2018-398671
7 https://permalink.mercell.com/91001251.aspx

As indicated in the original tender, we sought to purchase a sufficient amount of nodes for our students and researchers to test hypothesis and port and test hybrid OpenMP/MPI/offload engine codes. We found that 4 nodes fit our requirements and budget.

1.3.1 Track 1 - FPGAs technology evaluation

4 servers[8], each with;

- Dual Processor (DP) AMD Epyc "Rome" 7302 CPUs
- 512 GB DDR4-3200 memory
- 960GB SSD Read intensive OS system disk
- 4 qty 1.92TB NVMe PCIe Gen4 U.2 Mixed for local scratch storage /work
- 3.84TB Samsung SmartSSD NVMe Gen3 w/Xilinx Kintex
- 1 qty Xilinx Alveo U250
- 1 qty Xilinx Alveo U280
- Min. 1 qty 1/10/25Gbps Ethernet
- Dual-port 200 Gbps Infiniband communication
- Mellanox Innova-2 Flex Open Programmable SmartNIC 100Gbps
- 1Gbps BMC w/ at least IPMI v2.0 and Redfish 1.0
- 4 Free PCIe Gen4 x16 or x8 slots (may vary dependent on vendor)

In summary: 4 nodes with 8 Epyc Rome CPUs, 12 Xilinx FPGAs.

The base configuration of the servers is selected to match existing I/O servers used for Infiniband network research, such that these units can facilitate two research areas (FPGAs and HPC networking).

1.3.2 Track 2 - "LUMI light" nodes

4 servers, each with;

- Dual Processor AMD Epyc "Milan" 7763 current top-bin SKU CPU

8 Keeping with Phase 1 tenders 4 nodes is a minimum requirement for porting testing hybrid OpenMP+MPI+offload engine codes.

- 2 TB DDR4-3200MHz[9]
- 1 qty 960GB OS NVMe
- 4 qty 1.92TB NVMe Gen4 U.2 scratch storage
- Min. 1 qty 1/10/25Gbps Ethernet Nvidia ConnectX-4 Lx
- Dual-port 200 Gbps Infiniband communication Nvidia Mellanox ConnectX-6 VPI
- 1Gbps BMC w/ at least IPMI v2.0 and Redfish 1.0
- Two of the nodes should be fitted with latest generation AMD Instinct and two with Nvidia A100 GPGPUs, for a total of
  - 4 qty AMD Instinct Mi100/32GB PCIe Gen4 GPUs
  - 4 qty NVIDIA Ampere A100/40GB PCIe Gen4 GPUs

In summary: 4 GPGPU enabled nodes with 8 Epyc Milan CPUs, of which two with 2 qty Nvidia A100 GPGPUs each and two nodes with 2 qty AMD Instict Mi100 GPUs each. These nodes (esp. the ones with AMD GPUs) would be a very relevant set-up to prepare for use of the LUMI supercomputer to be opened at CSC Finland[10] later this year.

1.3.3 Track 3 - "HGX-1" DL node

We are requesting the latest implementation of Project Olympus HGX-1[11] GPU server with DP Milan processors and 8 qty Nvidia A100/80GB[12] GPUs.

It should be one (1) NVIDIA DGX A100/640GB[13].

This system should have 2 TB memory, OS: 2* 1.92TB M.2 NVME RI drives Internal Storage: 30 TB (8* 3.84 TB) U.2 NVMe Mixed Usage drives, 2* 10/25/100/200Gbps Ethernet, and 8* 200 Gbps Infiniband communication.

In summary: This system will be a natural and very powerful next step to follow up on our very popular DGX-2 system (V100 GPUs) currently in operation.

9 Installed in a Performance Optimized manner (1DPC)
10 https://www.csc.fi/
11 http://files.opencompute.org/oc/public.php?service=files&t=b18fa4e83bb3ecc4fffa179edab0ce39&download
12 https://images.nvidia.com/data-center/a100/a100-datasheet.pdf
13 https://images.nvidia.com/aem-dam/Solutions/Data-Center/nvidia-dgx-a100-datasheet.pdf

Note! - Alternative OEM implementations to Nvidia DGX A100/640G can be offered, but must include at least the same specifications as the original 2nd generation Nvidia DGX A100/80GB.

Dependent on the specification any OEM system offered may get a reduced score based on functionality. Contractors may include in A07 spreadsheet both Nvidia and OEM offerings.

Some of the systems that may be offered as an alternative are

- Supermicro SuperServer SYS-420GO-NART with DP AMD Epyc Milan and 8 qty A100/80GB
- Supermicro SuperServer SYS-420GP-TNAR w/ DP Intel Xeon IceLakeSP and 8 qty A100/80GB (with reduced score)
- Inspur NF5488A5 w/ DP AMD Epyc Milan and 8 qty A100/80GB
- Inspur NF5488M5-D w/ DP Intel Xeon IceLakeSP og stk A100/80GB (with reduced score)
- Gigabyte G492-ZD0 w/ DP AMD Epyc Milan og 8 stk A100/40GB
- 2 qty Dell XE8545 (with reduced score, but scale-out additional CPUs and memory may weight up for scale-up GPGPU's)

  Or any other vendors OCP HGX-1 implementations based on DP AMD Epyc Milan 7763 and 8 qty Nvidia A100/80GB.

2 TENDER REQUIREMENTS AND DEFINITIONS

2.1 Scope of the contract

Scope of contract is outlined in chapter 3 below. The hardware should be delivered with 3 year warranties, with options for 4 and 5 years.

For further information, please refer to phase 1 tender documents.

2.2 Budget for the procurement

The phase 3 extension is financed by the Research Council of Norway (RCN), together with in-kind contributions from partners in the eX³ consortium.

For obvious reason as specified in Rules of the Competition we will not reveal the budget.

2.3 Splitting up deliverables

It will not be possible to apply for delivery of only parts of the contract, although the use of subcontractors is allowed.

2.4 Use of Subcontractors

We do not expect the need for subcontractors for this tender.

But, when using subcontractors, or any resources belonging to other companies, should be in accordance with Public Procurement Regulations § 16-10.

For additional information, please refer to original phase 1 “Rules of the Competition” tender documents section 5.5.

2.5 Procurement announcement

This procurement is announced in Mercell (www.mercell.com) and consecutively in DOFFIN (www.doffin.no) and in TED (www.ted.europa.eu).

The permalink to the documents is https://permalink.mercell.com/155095285.aspx

3 FUNCTIONAL HARDWARE

The below sections include a summary of the expected configuration. Offer should include Appendix A07 spreadsheet filled out with the offers. Please refer to the attached A07 example spreadsheet.

3.1 Offload engine enabled computational nodes

All 8 systems in section 3.1.1-3.1.3 below shall be the same with the exception of the processors, HCAs, memory capacity and the offload engine offered. As specified in phase 1 ITT, tenderer should strive to use common hardware. With this we mean that chassis, motherboards, IO cards, disks, memory, etc. should be the same unless otherwise is stated. This means that servers with AMD Epyc “Rome” processors should be revision 2, i.e. support both “Milan” and “Rome”.

Systems shall have the fastest memory available installed in *a performance optimized fashion* appropriate for memory bandwidth sensitive HPC workloads, which is one DIMM Per Channel (1DPC).

Please refer to “Memory Population Guidelines for AMD EPYC™ 7002[14] and 7003 Processors”[15]. We will run McCalpin Stream[16], BabelStream[17] and in-house benchmarks to verify proper memory configuration and GPGPU performance.

Systems which we have found to meet our requirements are for example Dell PowerEdge R7525, Lenovo ThinkSystem SR665, Supermicro A+ Server 2024US-TRT, Gigabyte R282-Z94, ao.

Any other system similar to the above should be acceptable.

3.1.1 Four Xilinx FPGA enabled nodes

14 https://developer.amd.com/wp-content/resources/56301_1.0.pdf
15 https://www.abacus.cz/prilohy/_5100/5100789/AMD%20Memory%20Pupulation%20Guidelines%20for%20EPYC%207003%20Series%20Processors%2056873_0.70.pdf
16 https://www.cs.virginia.edu/stream/
17 http://uob-hpc.github.io/BabelStream/

Qty	Description
4	2u GPU chassis*
8	AMD Epyc 7302 “Rome” DP/UP - 16C/32T 2.45G 256M 280W SP3
4	Xilinx Alveo U250 PCIe Gen3 x16 (passive cooling)
4	Xilinx Alveo U280 PCIe Gen4 x8 (passive cooling)
64	32GB DDR4-3200MHz RDIMM (512GB RAM per node 1DPC)
4	Samsung 960GB NVMe SSDs or equivalent (OS system disk)
4	Samsung SmartSSD Computational Storage Drive – 3.84GB 2.5" NVMe Gen 3 Xilinx KU15P FPGA
16	Samsung PM1733 V5 TLC EAGLE 1.920GB NVMe PCIe Gen4 x4 or equivalent (for /work local scratch space)
4	Dual Port Mellanox ConnectX-6 VPI P/N: MCX653106A-HDAT
4	Single or Dual port Melanox ConnectX-4 Lx 10/25Gbps SFP28 (OCP slot)
4	Mellanox Innova-2 Flex Open for Application Acceleration, dual-port SFP28, 25GbE, KU15P, 8GB, No Crypto, PCI4.0 x8, active heat sink, tall bracket P/N: MNV303611A-EDLT
4	1Gbps BMC IPMI 2.0, Redfish 1.1 capability
	At least 3 additional PCIe Gen4 x8 or x16 slots available[18]
8	Redundant PSUs with enough power for max configurations
4	5m TwinAX 25Gbps cables compatible with Mellanox CX-4

Table 3: Specification of FPGA nodes

18 CPU dependent: Some vendors include 2 slots with PCI Gen3 x16 and x8, which is acceptable.

3.1.2 "Top-bin" Nvidia A100/40GB GPGPU enabled nodes

Qty	Description
2	2u GPU chassis*
4	AMD Epyc "Milan" 7763 DP/UP - 64C/128T 2.45G 256M 280W SP3
4	Nvidia A100/40GB HBM2 PCIe Gen4 x16 (80GB HBM2e if available)
64	128GB DDR4-3200MHz RDIMM (2TB RAM per node 1DPC)
2	Samsung 960GB NVMe SSDs or equivalent (OS system disk)
8	Samsung PM1733 V5 TLC EAGLE 1.920GB PCIe Gen4 x4 or equivalent (for /work local scratch space)
2	Dual Port Mellanox ConnectX-6 VPI P/N: MCX653106A-HDAT
2	Single or Dual port Mellanox ConnectX-4 Lx 10/25Gbps SFP28 (OCP slot)
2	1Gbps BMC IPMI 2.0, Redfish 1.1 capability
	At least 4 additional PCIe Gen4 x8 or x16 slots available[19]
4	Redundant PSUs with enough power for max configurations
2	3m TwinAX 25Gbps cables compatible with Mellanox CX-4

Table 4: Specification of 2 "top-bin" DP AMD Epyc nodes with dual Nvidia A100

19 CPU dependent: Some vendors include 2 slots with PCI Gen3 x16 and x8, which is acceptable.

3.1.3 “Top-bin” AMD Instinct Mi100 GPGPU enabled nodes

Qty	Description
2	2u GPU chassis*
8	AMD Epyc 7763 “Milan” DP/UP - 64C/128T 2.45G 256M 280W SP3
4	AMD Instinct Mi100/32GB HBM2 PCIe Gen4 x16
64	128GB DDR4-3200MHz RDIMM (2TB RAM per node 1DPC[20])
4	Samsung 960GB NVMe SSDs(OS system disk) or equivalent
16	SAMSUNG PM1733 V5 TLC EAGLE 1.920GB PCIe Gen4 x4 (for /work local scratch space) or equivalent
4	Dual Port Mellanox ConnectX-6 VPI P/N: MCX653106A-HDAT
4	Single or Dual port Melanox ConnectX-4 Lx 10/25Gbps SFP28 (OCP slot) ConnectX-5 Lx
4	1Gbps BMC IPMI 2.0, Redfish 1.1 capability
	At least 4 additional PCIe Gen4 x8 or x16 slots available
8	Redundant PSUs with enough power for max configurations
2	3m TwinAX 25Gbps cables compatible with Mellanox CX-4

Table 5: Specification of 2 “top-bin” DP AMD Epyc nodes with dual AMD Mi100

* Example systems: Dell PowerEdge R7525, Lenovo ThinkSystem SR665, Supermicro A+ Server 2024US-TRT, Gigabyte R282-Z94, or equivalent.

20 Performance optimized for 3200MT/s. Alternative bank should be free for NVDIMMs.

3.2 OCP “Olympus” HGX-1 AI/DL node

The offer should include one (1) Nvidia A100/640GB GPGPU server with below specification. Alternative OEM systems with minimum the same configurations can be offered and will be considered.

3.2.1 Nvidia DGX A100/640GB DL node specification

Qty	Description
1	Nvidia DGX A100/640GB System, 8x A100/80GB, 2TB DDR4-3200 Registered RAM, 30TB NVMe, EDU REG P/N # DGXA-2530F+P2EDI00
1	Nvidia DGX A100/640GB 3 Year Support, EDU REG
1	Nvidia DGX A100, Mandatory On-site installation and configuration
2	10M optical 100Gbps QSFP+ cables Mellanox

Table 6: Specification of OCP HGX-1 configuration

3.3 Tools and Libraries

Any vendor specific machine learning/AI SDKs/toolkits should be bundled with the respective systems offered.

3.3.1 Xilinx SDKs and AI/ML Toolkits

Xillinx SDKs/toolkits should be bundled with the systems offered. Minimum 2 concurrent seat licenses for tools and SDK should be included. Preferably floating FlexLM licenses.[21]

3.3.2 Nvidia tools for Deep Learning and AI/ML

Any vendor specific SDKs and Nvidia frameworks should be downloadable with the Nvidia DGX A100 system, or equivalent OEM systems.

3.3.3 System firmware upgrade functionality

Systems should include easy firmware upgrade tools for Linux, and be compatible Ubuntu 18.04.5LTS or higher.

21 https://www.xilinx.com/support/licensing_solution_center.html

Examples of such tools are Dell DSU, Supermicro SUM, Lenovo System Updates, etc.

Vendors which support Linux Vendor Firmware Service partially or fully will get plus scores during evaluation.

4 PROGRESS AND TIME SCHEDULE

Deadline	Dato
Tender procurement announcement	April 26th 2021
Tender deadline (fixed)	May 7th 2021, 13:00 CEST
Tenders opening (fixed)	May 7th 2021, 13:00 CEST
Award of contract	May 7th 2021, 16:00 CEST
Contract signing	May 7th 2021, 16:00 CEST
Customer takes delivery	4-8 weeks after signature

Table 7: Tentative time schedule.

4.1 Installation

Vendor should provide one highly skilled engineer to work alongside Simula's engineer for 2 days.

5 ACCEPTANCE AND APPROVAL TESTS.

The Nvidia DGX A100 or equivalent should pass Nvidia's DGX System Field Diagnostics test.[22]

The other 8 nodes will be set in production immediately and be verified with various benchmarks and internal.

Any faults discovered will be raised through the respective vendors customer support channels.

6 Qualification criteria

Please refer to sections 5 and 6 of phase 1 "Rules of the Competition Phase 3".

The contract will be awarded to the tenderer which delivered the best price and quality, according to the specified award criteria.

This tender has a very specific scope, and the contractor that meets the specification most closely at a reasonable price will be successful.

22 https://docs.nvidia.com/deploy/hw-field-diag/index.html Current version: dgx-diagos_1.7-6_amd64.deb

6.2.1 Mandatory requirements for participation

Requirements are per section 5 in the phase 1 procurment “Rules of the Competition” document. Briefly, the offer should include the below documentation.

- The Contractor shall also submit Tax certification no older than 6 months. Please refer to section 5.2 in above mentioned document for details
- The contractor shall submit ESPD format with part II, III and IV filled out.
- The Contractor shall have sufficient financial strength to perform the contract.
- The Contractor shall have substantial experience with similar contract performance and document these.
- The Contractor shall have sufficient capacity to deliver on the contract.

Award of contract, justification and opportunity to appeal is outlined in phase 1 procurement documents.

7 COMMERCIAL TERMS

Please refer to Phase 1 tender documents Appendix 6 and 7.

7.1 Price

Please refer to Appendix 7.

Offer should include a spreadsheet with a summery sheet with the pricing for the various components from section 3 above are listed in separate sheets as per “A07A Attachment Price and Requirement Matrices-MR.xlsx” example.

The vendor is free to use their own quotation systems format on these sheets, but the summary sheet should clearly state the total price per node type and for the complete system.

For any questions regarding format of spreadsheet please contact torel@simula.no.

7.2 Payment

For both the initial and complete configuration, the payment is net 30 days after acceptance of the system as per Norwegian Government policy or as mutually agreed upon in the final contract.

7.3 Terms of Delivery

Incoterms: Delivery is DDP (Incoterms 2010), Simula Research Laboratory AS, Attn: Tore H. Larsen, HPC department, Martin Linges Vei 25, 1325 Fornebu, Norway.

Airport destination: Gardemoen Oslo Airport (OSL) is closest. For further shipping instructions please contact torel@simula.no.

The risk for accidental causes shall pass from the Vendor to the Customer on the day of acceptance. The Vendor shall provide insurance coverage up to this day. The cost incurred may be specified in the Tender.

7.4 Invoicing

Invoicing is done when system has been accepted.

Invoice address:

Simula Research Laboratory AS
Postboks 134
1325 Lysaker
Norway

Invoice can preferably be sent electronically: invoice@simula.no

Invoice shall be marked with the following information
Project number: eX³ ITT 04/2021

Name: CEO Prof. Aslak Tveito

7.5 Administrative provisions

All communication regarding this procurement should go through the Mercell system as explained in accompanying document “Rules of the Competition” from phase 1.

For a introduction into Mercell communication module, please refer to the guide.

If bidder find the tender unclear or lacking information, and the bidder is unable to use Mercell InfoCenter system, one may send mail or e-mail to:

Simula Research Laboratory AS
c/o Tore H. Larsen
Boks 134, 1325 Lysaker

torel@simula.no

Written questions sent by mail should be marked “Questions regarding Tender eX³ ITT 4/2021”. Questions must be sent at least 5 days prior to the tender deadline date. Questions received after this date may not be answered.

Answers to questions received will be sent to all bidders by using Mercell InfoCenter messaging system referred to herein. This will ensure that all bidders automatically receive the same answer simultaneously anonymously by e-mail. No questions will be answered on an individual basis.

7.6 Erratas, additions and changes to tender

Before the end of the deadline, the Contractors have the right to make corrections, additions and changes to the tender that are not essential. Such changes will be published to all bidders simultaneously through Mercell InfoCenter messaging system, as well as connected systems Doffin and TED. Simula takes no responsibility of failure of the systems. Bidders are encouraged to actively monitor the various systems messaging system for this tender.

If erratas, additions and changes to tender is coming to late for bidders to be able to answer, Simula will consider extending the tender deadline. In such a case all bidders are informed through the messaging system and will have equal opportunity to postpone responds.

In case the bidders find mistakes or information lacking pertaining to this tender, bidders are encouraged to inform Simula in writing as soon as possible.

7.6.1 Customer Representatives

Role	Name	E-mail	Phone
Administrative responsible / PL	Are Magnus Bruaset	arem@simula.no	
CEO	Dr. Aslak Tveito	aslak@simula.no	
Tender contact	Tore H. Larsen	torel@simula.no	+47 918 33 670

Table 8: Customer contacts