



# **Deliverable 2.2:**

## **Specifications and Offer Documents for Computing Facility**

Project no. 316165

Project acronym:  
**CCQCN**

Project full title:  
**Crete Center for Quantum Complexity and Nanotechnology**

REGPOT-2012-2013-1

SEVENTH FRAMEWORK PROGRAMME

**Deliverable 2.2**  
Specifications and Offer Documents for Computing Facility

Due date of deliverable: November 2013

Dissemination level: PU

## TABLE OF CONTENTS

1. Introduction .....	4
2. Computing equipment to be procured .....	5
3. The procurement process and the procurement timeline .....	7
APPENDIX I: Specifications of HPC Computing Facility .....	9
General / Operational Specifications .....	10
Compute & Network infrastructure.....	11
Storage subsystem .....	19
Virtualization nodes .....	22

# Specifications and Offer Documents for Computing Facility Crete Center for Quantum Complexity and Nanotechnology (CCQCN)

## 1. Introduction

Deliverable 2.2 comprises one of the deliverables of Work Package 2 (WP2), the objective of which is to successfully procure, purchase and install new equipment and infrastructure that will greatly enhance the experimental and computational capabilities of the CCQCN Center and, thus, make it a center of excellence, very competitive on a European (and global) scale.

WP2 comprises two Tasks. Task 1 (T1-WP2) focuses on the installation of new research equipment, which will strengthen the infrastructure for performing advanced experiments. Task 2 (T2-WP2) focuses on the installation of a high performance computer facility, which is essential for the success of the project.

The present deliverable (D2.2) presents the specifications of the high-performance computing facility, and describes the procurement process, which is currently running.

Because of the rigorous (and, thus, lengthy) procurement process followed by the University of Crete (under Greek and EU law compliance), the offer documents (that is, offers by the perspective suppliers/vendors) will be presented after the full completion of the procurement process (in a *D2.2 Supplement* file).

It should be mentioned that the in-situ operation of a high performance computing (HPC) facility is essential for the success of the CCQCN project; research in Crete is currently inhibited by the lack of capability to perform computations on problems that demand serious computer power, such as those that demand bases with large number of states, large electromagnetic calculations, etc. The HPC facility addresses this problem and enables the researchers that work in the Center as well as various collaborators and visitors to advance their research.

## 2. Computing equipment to be procured

As described in the DoW (WP2), the following main categories of computing equipment will be procured:

### I. Hardware architecture and components

Hardware components comprise high density – blade nodes with dual CPU configurations (8 cores per CPU) and sufficient RAM (at least 64 GB per node). The storage hardware implements the Lustre or the GPFS filesystem model with redundant, mixed SATA/SAS HDD enclosures and the corresponding file servers in HA setup. One portion of the storage will be dedicated to metadata and object storage. The theoretical peak performance is of the order of 20 TFlops for the total of CPU and GPU's nodes.

### II. Networking architecture and equipment

Storage-to-nodes interconnect, are based on state-of-the-art network architecture with either 10GigE or InfiniBand switched fabric communications links, for fast data exchange between the compute and storage nodes, eliminating at the same time any chance of bottleneck between storage and compute nodes. Each node group takes advantage of the relatively low cost but quite efficient 10GigE switched network for message passing and computational tasks. Selecting this kind of networking technologies over other communication links, is a prerequisite for latency sensitive HPC applications, which this cluster intends to run, and at the same time reflects the current design trends for this kind of clusters. Moreover, message-passing and storage-passing traffic demands yield increased throughput, especially for the IOPS driven applications, this way, maximize value for money profit. In order to isolate the various management tasks (H/W management, machine events etc.) from the data network, we will use a separate network with the appropriate hardware components.

### III. Storage

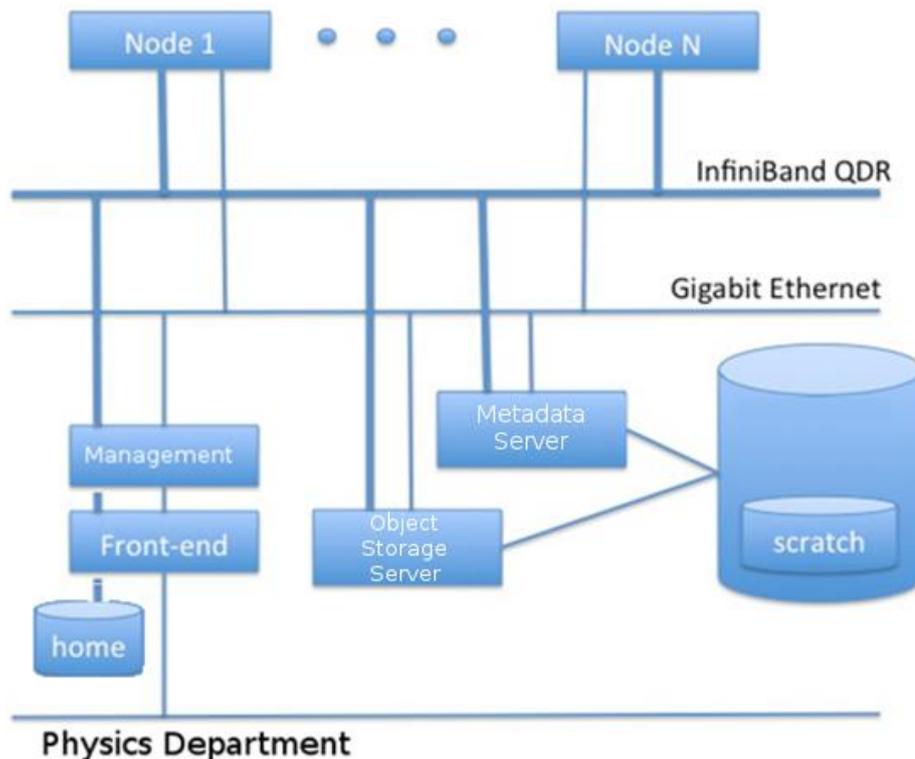
The CCQCN HPC system needs to ensure data high availability (HA), performance and scalability at the same time. In HPC environments, a parallel and distributed filesystem is usually deployed in order to satisfy these conditions. Of the most popular open source filesystems in HPC environments are GPFS and Lustre. To implement such filesystems, relatively low-cost/large capacity storage for the actual data and a fast, reduced capacity meta-data storage, are needed. To ensure redundancy, storage enclosures with multiple controllers and network interfaces for object and meta-data storage, multiple file servers in HA configuration, are needed. Moreover, to ensure data integrity and performance and avoid single point of failure at the same time, each storage enclosure needs to be equipped with redundant data interfaces, RAID-6 disk configuration and a set of hot spare disks.

#### IV. Cluster Management and queuing system

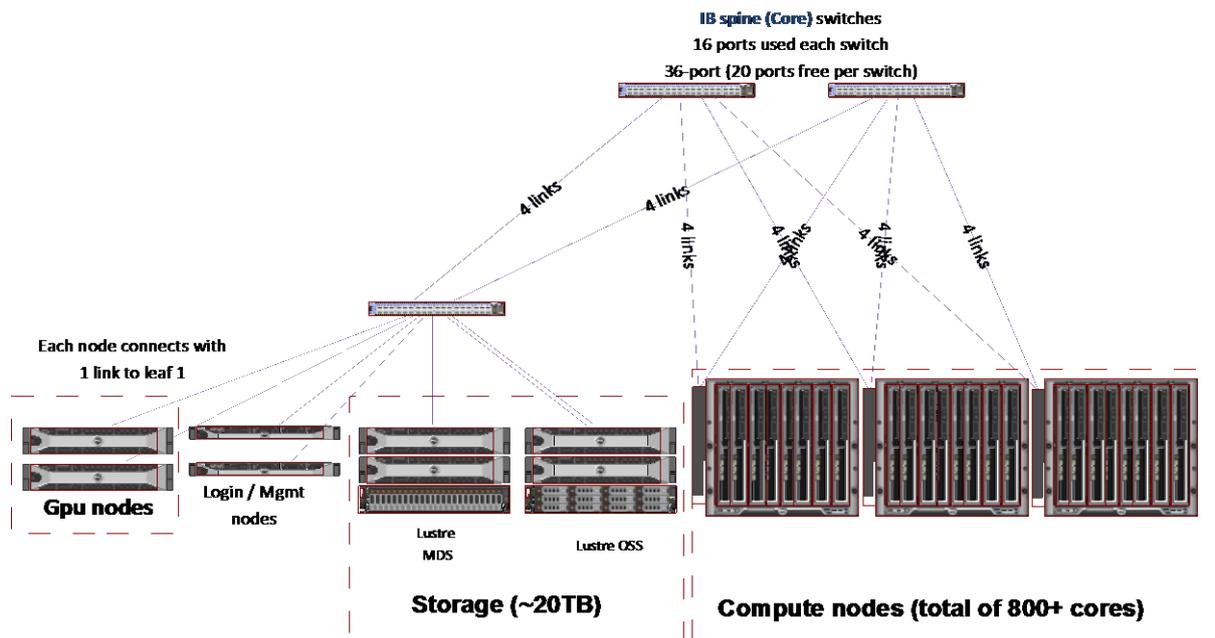
Batch queuing systems using the most widely acceptable open source tools for job submission and workload management, will be implemented. This scheme implies that there will be a redundant set of 2 job submission nodes and again, a redundant set of 2 management servers. Management servers, will also commit any administrative tasks, such as operating system deployment, image management, software deployment etc.

The following schematics delineate the HPC network architecture and the HPC infrastructure.

Schematic 1: HPC network architecture



Schematic 2: HPC infrastructure (indicative)



The detailed set of specifications of the computing equipment is presented in the Appendix.

The total cost is budgeted at 465,000 Euros. The procurement process and the timeline are described in the following section.

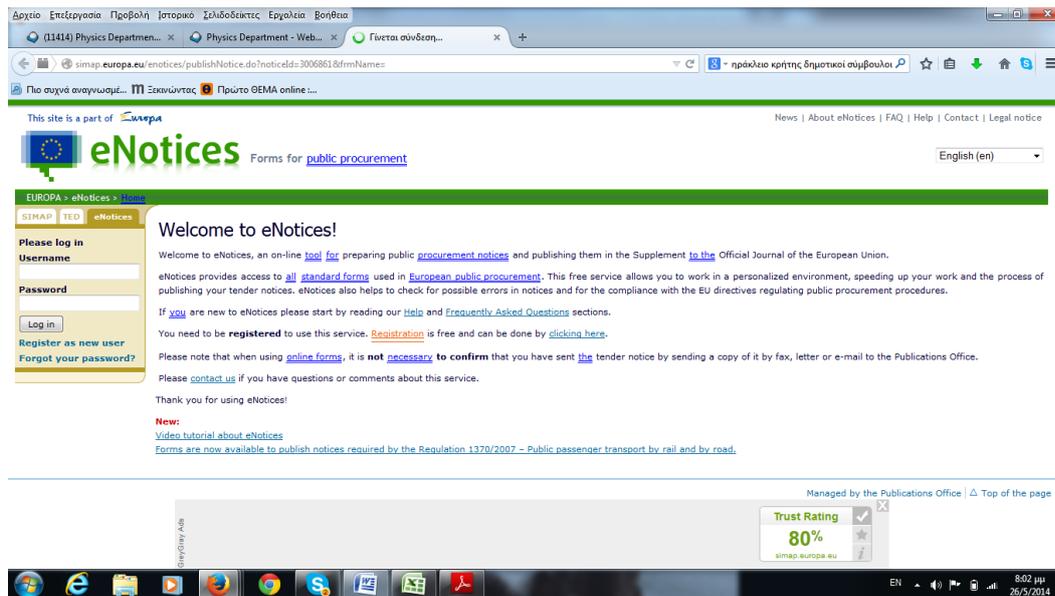
### 3. The procurement process and the procurement timeline

Procurement of the HPC computing equipment is currently in progress. Under Greek law, an international bid process is currently under way; the equipment specifications and the rules and regulations of the international bid process need to be openly announced for 52 days.

The public announcement is presented in:

- The Greek Government's Gazette (Εφημερίδα της Κυβερνήσεως, Τεύχος Διακηρύξεων Δημοσίων Συμβάσεων), in the Public Procurements Issue (it was announced on February 5, 2014).
- The Greek press (announced on February 5 and February 7, 2014)
- The University of Crete procurement announcements system at <http://www.elke.uoc.gr> and <http://www.uoc.gr> (announced on February 5, 2014)

d) The Publications Office of the European Union via the eNotices service:



In which it was announced on January 31, 2014 (No: 2014-014762).

On April 9, 2014, the bid offers will be opened and the evaluation procedure will start. The HPC facility is expected to be operational by Fall 2014. Official language of the bid process is Greek (a necessary condition under Greek Law). It should be mentioned that the procurement process has three necessary steps: a) the opening of the bid offers and the examination/evaluation of the necessary legal documents for participating in the bid process, b) the opening of the technical offer (compliance with specs), and c) the opening of the economic offer. After each step, the evaluation committee must submit each specific report to the University's Research Committee for confirmation; this Committee meets once a month. If vendors have objections to the specific evaluation/compliance reports, they can submit their objections for discussion and deliberation, after each step. The successful completion of these steps, lead to the grading of each technical and economic offer (and their aggregate grade) and subsequent ranking of each offer.

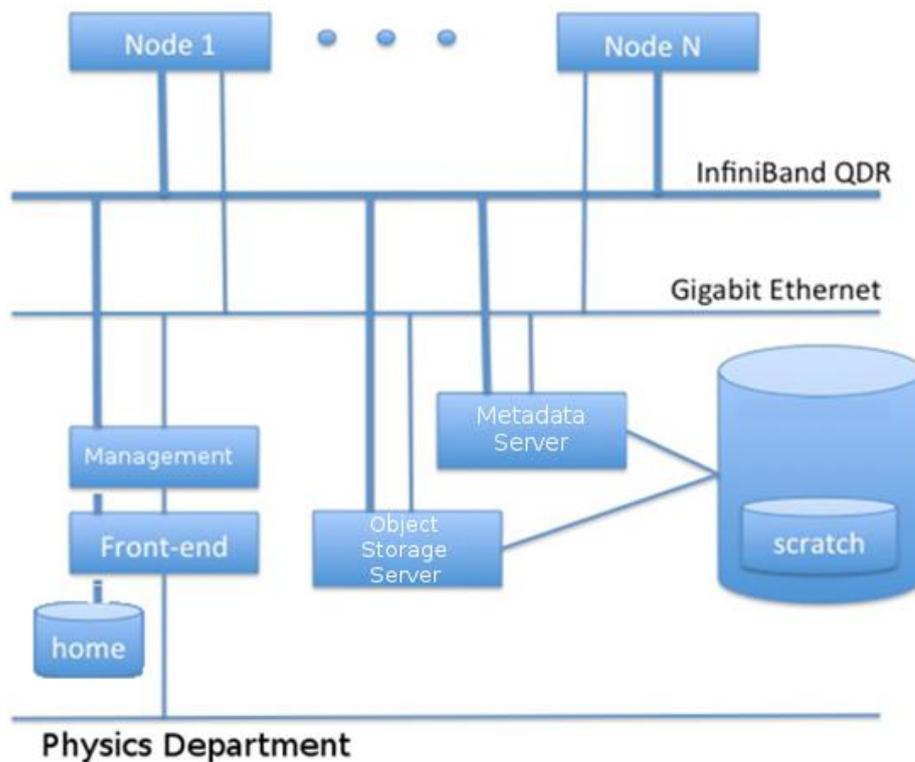
Following a procurement change (requested for Part 1 and Part 6 of the CCQCN research equipment, in order to procure a different set of equipment, because the experimental work, which was planned to be conducted on the Part 1 and Part 6 equipment can be conducted in another research facility), an additional budget for the upgrade of the HPC facility has been requested. The additional (requested) budget comprises 57,405 Euros for the enhancement of HPC GPU (graphics processing unit) computing and storage, as well as 25,600 Euros for the enhancement of HPC peripherals, support equipment, data visualization nodes, small items replacements and some consumables, and space installation modification. It should be mentioned that the total cost for both research equipment and computing equipment will remain the same as planned in the DoW.

## APPENDIX I: Specifications of HPC Computing Facility

### HPC cluster at-a-glance

- *Compute nodes:* The cluster will be composed of at least 50 dual CPU nodes and at least 2 nodes equipped with GP-GPU cards.
- *Management nodes:* Management nodes will perform maintenance and management operations on the infrastructure.
- *Access nodes:* Access nodes will provide access from the outside world to the cluster.
- *Virtualization nodes:* will be used externally. These nodes will provide the necessary infrastructure for supporting other CCQCN services.
- *Network infrastructure:* Network infrastructure will be implemented with Infiniband 4 x QDR 40 Gb/s or similar. Management network must be capable of managing the whole infrastructure. Access network will provide connectivity with the network infrastructure of the Physics Department.
- *Parallel Filesystem*

The parallel filesystem will be implemented either with Luster or GPFS



The following table presents the detailed set of specifications that a vendor/supplier submitting a bid for the procurement of the CCQCN computing equipment must comply to. The term “Mandatory” means that if “YES” the vendor must necessarily satisfy the specification (and confirm so in the column “Supplier’s Response”).

S/N	Specifications	Mandatory	Supplier’s Response	References/additional comments
<b>1</b>	<b>General / Operational Specifications</b>			
1.1.1	The proposed system must be designed for High Performance Computing	YES		
1.1.2	The proposed system must be designed either as Cluster of Workstations (COW), or Massively Parallel Processor (MPP)	YES		
1.1.3	The proposed system must provide a Parallel Filesystem available to compute nodes and clients	YES		
1.1.4	Compute nodes must be connected through high speed Infiniband connections	YES		
1.1.5	A number of login nodes must be available	YES		
1.1.6	A number of management nodes must be available	YES		
1.1.7	The infrastructure of the proposed system must contain an isolated management network capable of providing management operations, based on 1 GbE or better	YES		
1.1.8	The system must be capable of out-of-band management	YES		
1.1.9	The nodes and the storage system must provide monitoring and error detection mechanisms.	YES		
1.1.10	The proposed architecture must avoid SPOF (Single Point of Failure)	PREFERABLE		

## Compute & Network infrastructure

S/N	Specifications	Mandatory	Supplier's Response	References
<b>2.1</b>	<b>Compute nodes</b>			
2.1.1	Compute nodes should be of the same architecture	YES		
2.1.2	Compute nodes should be capable of MPI calculations	YES		
2.1.3	Compute nodes should be capable of OpenMP calculations	YES		
2.1.4	Minimum number of compute nodes	≥ 50		
2.1.5	Compute nodes should be arranged in "Blade" chassis	YES		
2.1.6	Minimum number of CPU per node	≥ 2		
2.1.7	Minimum number of cores per processor	≥ 10		
2.1.8	Processors must have been announced 6 months ago at most, before the date of the current tender	YES		
2.1.9	Processors should support 64-bit instruction set and AVX (Advanced Vector Extensions) extensions	YES		
2.1.10	Processor specifications based on SPEC bench Suite: <ul style="list-style-type: none"> <li>• CINT2006Rate ≥ 810</li> <li>• CFP2006Rate ≥ 610</li> </ul>	YES		
2.1.11	Minimum score of the cluster in TFlop/s	≥ 20		
2.1.12	Minimum RAM per node (GB)	≥ 96		
2.1.13	Each node should be equipped with DDR3 1600 ECC SDRAM or better	YES		
2.1.14	Low consumption RAM modules	PREFERABLE		
2.1.15	Each node should contain scratch disk(s) to elevate the	YES		

	overall performance of the cluster			
2.1.16	Available network interface for each node: GbE or 10GBE and at least on HCA (Host Channel Adapter) for Infiniband interconnect	YES		
2.1.17	HCA should be compatible with OpenFrabrics RDMA	YES		
2.1.18	Compute nodes must be compatible with Open source Operating systems like CentOS 6.x or similar	YES		
<b>2.2</b>	<b>Compute nodes with GP-GPU card</b>			
2.2.1	Model name and manufacturer	YES		
2.2.2	Nodes should be of the same manufacturer as in §2.1	PREFERABLE		
2.2.3	Number of nodes	≥ 2		
2.2.4	Minimum number of CPU per node	≥ 2		
2.2.5	Minimum number of cores per processor	≥ 8		
2.2.6	Processors must have been announced at most 6 months ago, before the date of the current tender	YES		
2.2.7	Processors should support 64-bit instruction set and AVX (Advanced Vector Extensions) extensions	YES		
2.2.8	Processor specifications based on SPEC bench Suite: <ul style="list-style-type: none"> <li>• CINT2006Rate ≥ 540</li> <li>• CFP2006Rate ≥ 450</li> </ul>	YES		
2.2.10	Minimum RAM per node (GB)	≥ 64		
2.2.11	Each node should be equipped with DDR3 1600 ECC SDRAM or better	YES		
2.2.12	Low consumption RAM modules	PREFERABLE		

2.2.13	Minimum number of GP-GPU cards per node	≥ 1		
2.2.14	Minimum number of cores per processor per GP-GPU card	≥ 50		
2.2.15	GP-CPU card score (Double precision floating point ):	≥ 1 (TFlop)		
2.2.16	GP-GPU card RAM type: GDDR5	YES		
2.2.17	Memory size (GB)	≥ 6		
2.2.18	Memory bandwidth(GB/s)	≥ 200		
2.2.19	GP-GPU card cooling type	Active/Passive		
2.2.20	Number of 1GbE (1000Base-T) interfaces per node	≥ 2		
2.2.21	NIC must support Pre-execution Environment (PXE)	YES		
2.2.22	H/W RAID controller with RAID levels RAID 1, RAID 0, RAID 10	YES		
2.2.23	Number of HDD per node	≥ 2		
2.2.24	Total raw HDD capacity per node (GB)	≥ 500		
2.2.25	HDD interface type: SATA	YES		
2.2.26	HDD rpms	≥ 7200		
2.2.27	Hot-swap HDD	YES		
2.2.28	Each node must be equipped with h/w management controller network card, with support for SSH, TELNET, HTTP, HTTPS and SNMP	YES		
2.2.29	Redundant PSU N+1 or better	YES		
2.2.30	Hot-swap PSU	YES		
2.1.31	Nodes must be compatible with Open source Operating systems like CentOS 6.x or similar	YES		
<b>2.3</b>	<b>Login node</b>			
2.3.1	Model name and manufacturer	YES		
2.3.2	Number of nodes	≥ 1		
2.3.3	All nodes must be of the same architecture	YES		

2.3.4	Number of processor per node	≥ 2		
2.3.5	Number of active cores per processor	≥ 8		
2.3.6	Processors must have been announced at most 6 months ago, before the date of the current tender	YES		
2.3.7	Nodes must support hardware virtualization (Intel-VT Virtualization, AMD-V Virtualization or similar)	YES		
2.3.8	Processor specifications based on SPEC bench Suite: <ul style="list-style-type: none"> <li>• CINT2006Rate ≥ 540</li> <li>• CFP2006Rate ≥ 450</li> </ul>	YES		
2.3.10	Minimum RAM per node (GB)	≥ 64		
2.3.11	Each node should be equipped with DDR3 1600 ECC SDRAM or better	YES		
2.3.12	Low consumption RAM modules	PREFERABLE		
2.3.13	NIC must support PXE	YES		
2.3.14	H/W RAID controller with RAID levels RAID 1, RAID 0, , RAID 6, RAID 10 with cache protection	YES		
2.3.15	Number of HDD per node	≥ 6		
2.3.16	Total raw HDD capacity per node (TB)	≥ 3		
2.3.17	HDD type: NL-SAS	YES		
2.3.18	RPMS	≥ 7200		
2.3.19	Hot-swap HDD	YES		
2.3.20	Each node must be equipped with h/w management controller network card, with support for SSH, TELNET, HTTP, HTTPS and SNMP	YES		
2.3.21	Redundant PSU N+1 or better	YES		
2.3.22	Hot-swap PSU	YES		
2.3.23	Chassis 19'', 2U	YES		

2.3.24	Nodes must be compatible with Open source Operating systems like CentOS 6.x or similar	YES		
2.1.25	Nodes must be compatible with Open source Operating systems like CentOS 6.x or similar	YES		
<b>2.4</b>	<b>Management nodes</b>			
2.4.1	Model name and manufacturer	YES		
2.4.2	Number of nodes	≥ 1		
2.4.4	Number of processor per node	≥ 2		
2.4.5	Number of active cores per processor	≥ 8		
2.4.6	Processors must have been announced at most 6 months ago, before the date of the current tender	YES		
2.4.7	Nodes must support hardware virtualization (Intel-VT Virtualization, AMD-V Virtualization or similar)	YES		
2.4.8	Processor specifications based on SPEC bench Suite: <ul style="list-style-type: none"> <li>• CINT2006Rate ≥ 540</li> <li>• CFP2006Rate ≥ 450</li> </ul>	YES		
2.4.10	Minimum RAM per node (GB)	≥ 64		
2.4.11	Each node should be equipped with DDR3 1600 ECC SDRAM or better	YES		
2.4.12	Low consumption RAM modules	PREFERABLE		
2.4.13	NIC must support PXE	YES		
2.4.14	H/W RAID controller with RAID levels RAID 1, RAID 0, RAID 10	YES		
2.4.15	Number of HDD	≥ 2		
2.4.16	Total raw HDD capacity per node (GB)	≥ 500		
2.4.17	HDD type: SATA	YES		
2.4.18	RPMS	≥ 7200		

2.4.19	Hot-swap HDD	YES		
2.4.20	Each node must be equipped with h/w management controller network card, with support for SSH, TELNET, HTTP, HTTPS and SNMP	YES		
2.1.21	Each node must have at least 6 GbE network interface or better	YES		
2.4.22	Redundant PSU N+1 or better	YES		
2.4.23	Hot-swap PSU	YES		
2.4.24	Nodes must be compatible with Open source Operating systems like CentOS 6.x or similar	YES		
2.4.25	Rack mounted chassis 19", 2U	YES		
2.4.26	Management nodes must be equipped with the official management software for the cluster and network infrastructure	YES		
<b>2.5</b>	<b>Infiniband interconnect</b>			
2.5.1	Infiniband interconnect 4xQDR-IB (40Gbit/s) or better in terms of latency and bandwidth	YES		
2.5.2	BER (Maximum Bit Error Rate)	$\leq 10^{-12}$		
2.5.3	Cabling type (Copper, Fiber etc.)	YES		
2.5.4	Managed core/spine switches	YES		
2.5.5	Core/spine switches must support SNMP v2/v3	YES		
2.5.6	Adaptive Routing and Congestion Management support	YES		
2.5.7	Infiniband network topology Full Fat Tree 1:1 non-blocking	YES		
<b>2.6</b>	<b>Management Network</b>			
2.6.1	Management network must be instantiated as an	YES		

	independent and isolated from the Infiniband network			
2.6.2	Management network must be based on GbE or better	YES		
2.6.3	Management network should be designed in a way that fits the requirements of the infrastructure	YES		
2.6.4	Describe the type and number of the switches	YES		
2.6.5	Describe the type and number of interfaces	YES		
2.6.6	Switches must contain the necessary number of ports required for supporting the infrastructure	YES		
2.6.7	Switching capacity should satisfy the theoretical maximum bandwidth of the infrastructure	YES		
Switch specifications				
2.6.8	Describe the non-blocking maximum performance	YES		
2.6.9	Fully managed	YES		
2.6.10	Layer 2 switching	YES		
2.6.11	Support for IEEE 802.1p (VLAN prioritization).	YES		
2.6.12	Maximum number of VLAN IDs	≥ 1.000		
2.6.13	Maximum number of MAC addresses	≥ 8.000		
2.6.14	Support for IEEE 802.1ad	YES		
2.6.15	IEEE 802.1X	YES		
2.6.16	IEEE 802.1AB Link Layer Discovery Protocol (LLDP)	YES		
2.6.17	LLDP ḡ CDP compatibility	YES		
2.6.18	Spanning Tree/MSTP, RSTP	NAI		
2.6.19	Jumbo Frames	NAI		
2.6.20	Radius server support	YES		
2.6.21	SNMPv1, v2c, and v3	YES		
2.6.22	Remote monitoring (RMON)	YES		
2.6.23	Management VLAN	YES		
2.6.24	Support for 10GbE ports	PREFERABLE		
2.6.25	Each switch must contain at least 2 10GbE ports SFP+ for	PREFERABLE		

	uplinks or high-bandwidth connections to servers			
<b>2.7</b>	<b>Access network</b>			
2.7.1	Access network must be instantiated as an independent and isolated from the Infiniband network	YES		
2.7.2	Management network must be based on GbE or better	YES		
2.7.3	Management network should be designed in a way that fits the requirements of the infrastructure and be capable of interconnect with the network infrastructure of the Physics Department	YES		
2.7.4	Number of switches	≥ 1		
2.7.5	Number of active ports per switch	≥ 48		
2.7.6	Ports 1000Base-T	YES		
2.7.7	Describe the maximum non-blocking μέγιστη απόδοση	YES		
2.7.8	Fully managed	YES		
2.7.9	Layer 2 & Layer 3 switching	YES		
2.7.10	Support for IEEE 802.1p (VLAN prioritization).	YES		
2.7.11	Maximum number of VLAN	≥1000		
2.7.12	Maximum number of MAC addresses	≥ 16.000		
2.7.13	Support for IEEE 802.1ad	YES		
2.7.14	IEEE 802.1X	YES		
2.7.15	IEEE 802.1AB Link Layer Discovery Protocol (LLDP)	YES		
2.7.16	LLDP ή CDP compatibility	YES		
2.7.17	Spanning Tree/MSTP, RSTP	NAI		
2.7.18	Jumbo Frames	NAI		
2.7.19	Radius server support	YES		
2.7.20	SNMPv1, v2c, and v3	YES		
2.7.21	Remote monitoring (RMON)	YES		
2.7.22	Management VLAN	YES		
2.7.23	Support for 10GbE ports	YES		
2.7.24	Each switch must contain at least 2 10GbE ports SFP+ for uplinks or high-bandwidth connections to servers	NAI		

## Storage subsystem

S/N	Specifications	Mandatory	Supplier's Response	References
<b>3.1</b>	<b>Storage subsystem</b>			
3.1.1	Proposed storage system must be in production at the time of the current tender	YES		
3.1.2	Storage system must contain redundant HDD controllers with multiple H/W RAID support (0, 0+1, 1, 10, 6, 60)	YES		
3.1.3	Each storage unit must contain multiple controllers with failover operations	YES		
3.1.4	Controllers, must support cache memory (Read/Write) with Write Back / Write Through operations and data protections in case of unexpected power failure	YES		
3.1.5	Each storage unit must have certain number of hot-spare disks	YES		
3.1.6	HDDs in storage units must be hot-swapped	YES		
3.1.7	Each storage unit must be equipped with redundant PSUs N+1 or better	YES		
3.1.8	Describe in detail the architecture of the storage subsystem	YES		
3.1.9	HDD failures should not lead to data loss	YES, provide documentation		
3.1.10	Storage subsystem is divided into distinct Object and Metadata storage units.	YES		
Storage system must include disk subsystem attached at least to two servers. Storage subsystem is defined as follows:				
3.1.11	Storage units should support FC 8Gbps, FCoE, 10GE iSICSI, SAS	PREFERABLE		
3.1.12	Number of controllers	≥ 2		
3.1.13	Cache per controller	≥ 4 GB		
3.1.14	Each storage unit must support SSD, SAS, Near Line SAS or a combination of those	YES		
3.1.15	Maximum number of HDD support	≥ 150		

3.1.16	Support for RAID levels 0, 1, 1+0, 5, 6	YES		
3.1.17	Support for RAID migration	YES		
3.1.18	Snapshot support	YES		
3.1.19	Multi-pathing support	YES		
3.1.20	Management software must support alert notification and remote service access	YES		
3.1.21	Storage system must provide encryption support	PREFERABLE		
3.1.22	Dynamic cache data protections indefinitely	YES		
3.1.23	Nodes must be compatible with Open source Operating systems like CentOS 6.x or similar	YES		
<b>Metadata Data Storage</b>				
3.1.24	SAS HDD	YES		
3.1.25	HDD raw capacity (GB)	≥ 300		
3.1.26	Number of HDDs	≥ 12		
<b>Object Data Storage</b>				
3.1.27	NL-SAS HDDs	YES		
3.1.28	HDD raw capacity (TB)	≥ 1		
3.1.29	Number of HDDs	≥ 30		
<b>Storage subsystem servers</b>				
3.1.32	Number of units	≥ 2		
3.1.33	All units must be of the same manufacturer and be of the same specifications	YES		
3.1.34	Minimum number of processors per unit	≥ 2		
3.1.35	Minimum number of cores per processor	≥ 8		
3.1.36	Processors must have been announced at most 6 months ago, before the date of the current tender	YES		
3.1.37	Nodes must support hardware virtualization (Intel-VT Virtualization, AMD-V Virtualization or similar)	YES		
3.1.38	Processor specifications based on SPEC bench Suite: <ul style="list-style-type: none"> <li>CINT2006Rate ≥ 420</li> <li>CFP2006Rate ≥ 380</li> </ul>	YES		
3.1.40	Minimum RAM per node (GB)	≥ 64		

3.1.41	Each node should be equipped with DDR3 1600 ECC SDRAM or better	YES		
3.1.42	Available network interface for each node: GbE or 10GBE and at least on HCA (Host Channel Adapter) for Infiniband interconnect	YES		
3.1.43	HCA should be compatible with OpenFrabrics RDMA	YES		
3.1.44	Each node must have at least 4 FC8 interfaces to the storage system	YES		
3.1.45	Nodes must be compatible with Open source Operating systems like CentOS 6.x or similar	YES		
<b>3.2</b>	<b>Parallel Filesystem</b>			
3.2.1	Lustre or GPFS	YES		
3.2.2	The parallel filesystem should be built upon the architecture described at §3.1	YES		
3.2.3	Minimum capacity	≥ 35 TB (raw)		
3.2.4	Supported bandwidth	≥ 6GB/s		
3.2.5	Describe in detail the implementation of the parallel filesystem	YES		
3.2.6	Describe in detail the extension/upgrade options, the proposed system provides in conjunction with §3.1	YES		
3.2.7	The parallel filesystem must be available to compute §2.1 & §2.2 and access nodes §3.4	YES		
3.2.8	The parallel filesystem must have characteristics of Block management either as Shared-block-map, or Object based	YES		
3.2.9	The parallel filesystem must have characteristics of Metadata management, managed either by the clients (nodes) or the dedicated storage servers	YES		
3.2.10	The parallel filesystem must have characteristics of Cache Coherence and provide either distributed locking, or callbacks	YES		
3.2.11	The parallel filesytem must be deployed with Open source	YES		

	software, or all the available drivers must be provided for CentOS 6.x			
--	------------------------------------------------------------------------	--	--	--

## Virtualization nodes

S/N	Specifications	Mandatory	Supplier's Response	References
<b>4.1</b>	<b>Virtualization nodes</b>			
4.1.2	Number of nodes	≥ 3		
4.1.3	All units must be of the same manufacturer and be of the same specifications	YES		
4.1.4	Number of processors per node	≥ 2		
4.1.5	Number of cores per processor	≥ 8		
4.1.6	Processors must have been announced at most 6 months ago, before the date of the current tender	YES		
4.1.7	Nodes must support hardware virtualization (Intel-VT Virtualization, AMD-V Virtualization or similar)	YES		
4.1.8	Processor specifications based on SPEC bench Suite: <ul style="list-style-type: none"> <li>CINT2006Rate ≥ 540</li> <li>CFP2006Rate ≥ 450</li> </ul>	YES		
4.1.10	Minimum RM per node (GB)	≥ 128		
4.1.11	Η κεντρική μνήμη ανά κόμβο να είναι τεχνολογίας DDR3 1600 ECC SDRAM ή καλύτερη	YES		
4.1.12	Each node should be equipped with DDR3 1600 ECC SDRAM or better	PREFERABLE		
4.1.13	NICs must support Pre-Execution Environment (PXE)	YES		
4.1.14	HDD controller must support H/W RAID 1, RAID 0, RAID 10	YES		
4.1.15	Number of HDDs per node	≥ 2		
4.1.16	Total raw capacity per node (GB)	≥ 500		

4.1.17	HDD type: SATA	YES		
4.1.18	RPMS	≥ 7200		
4.1.19	Hot-swap HDDs	YES		
4.1.20	Each node must have a minimum of 6 GbE ports	YES		
4.1.21	Each node must be equipped with h/w management controller network card, with support for SSH, TELNET, HTTP, HTTPS and SNMP	YES		
4.1.22	Redundant PSU, N+1 or better	YES		
4.1.23	Hot-swap PSU	YES		
4.1.24	Nodes must be compatible with Open source Operating systems like CentOS 6.x or similar	YES		