FP6 INCO contract No. 026343

CX-CMCS

Centre of Excellence for Computational Modelling of Complex Systems

Deliverable D04

Equipment Tendering and Procurement Report

Author(s): Antun Balaz, Aleksandar Belic, Aleksandar Bogojevic

Status –Version: Final – a

Date: October 3, 2006

Distribution - Type: Public

Code: CX-CMCS-Deliverable-D04

Abstract: Deliverable D04 – “Equipment Tendering and Procurement Report” is a public document. The deliverable includes detailed information regarding the equipment tendering and procurement related to the reinforcing of the research capacity of the Scientific Computing Laboratory, Institute of Physics Belgrade (SCL).

Document Revision History

Date	Issue	Editor	Summary of main changes
September 27, 2006	-	Antun Balaz, Aleksandar Belic, Aleksandar Bogojevic	draft document created
October 2, 2006		Antun Balaz, Aleksandar Belic, Aleksandar Bogojevic	added contributions by D. Arsenovic and S. Vrhovac
October 3, 2006	Final -a	Antun Balaz, Aleksandar Belic, Aleksandar Bogojevic	Final document

Preface

The basic strategic objective of the CX-CMCS proposal is to transform the Scientific Computing Laboratory (SCL) into a centre of excellence, i.e. to decisively increase the quality of research conducted at SCL, and make it a preferred WB research partner for EU institutions working in the fields of simulation of complex systems and of GRID technology.

SCL is a unit of the Institute of Physics in Belgrade. The Institute contributes more than 10% of the total scientific output of Serbia and constantly ranks among the best R&D institutions in the region. SCL has 14 staff members, and participates in several international and national projects, including FP6 project SEE-GRID and Cost action P10. SCL defines the current state of the art in high performance computing in WBC with its PARADOX cluster (64+2 processors with aggregate speed Rmax=0.21 Tflops).

The proposed CX-CMCS SSA aims to reinforce research capacity at SCL by: hiring young researchers, providing of training and mobility for the research staff, and upgrading the computing infrastructure. The success of this endeavour will be measured through a benchmarking exercise to be performed in the project's last year. Our networking partners (4 from EU and 3 from Serbia) have been carefully selected to provide the skills and expertise necessary to reinforce the research potential of SCL through training and joint research. The proposed equipment upgrade (storage element, high throughput switch, and upgrade of RAM) will make it possible to tackle even the most complex GRID applications allowing SCL to become a key regional player in deployment and use of emerging GRID technology. CX-CMCS plans to set up an International Advisory Board whose expertise will help SCL develop a long term strategy and facilitate integration into ERA.

CX-CMCS aims to be a living example that it is possible to bridge the "digital divide" between countries and regions having high tech ICT technologies and those that do not.

Strategic objectives

The basic strategic objective of the CX-CMCS proposal is to transform SCL into a centre of excellence, i.e. to decisively increase the quality of research conducted at SCL, and make it a preferred WB research partner for EU institutions working in the fields of simulation of complex systems and of GRID technology.

Centres of excellence do not exist in a vacuum, however. In order for SCL to achieve and maintain a status of excellence, the proposed SSA aims to positively effect the research environment in Serbia at several levels: SCL's immediate R&D environment (the national partners in this proposal), the high performance computing segment, and the national R&D system as a whole.

Specific objectives

The specific objectives for the current SSA proposal have been formulated through an analysis of the following key points:

- Wider developmental objectives of Serbia and Montenegro and the West Balkan region pertaining to research and development (as presented in the Action Plan adopted at the Ministerial conference in Thessaloniki in June 2003);
- Existing strengths and weaknesses at SCL an the high performance computing sector in Serbia including: professional resources, material resources, financial and organizational resources, principle impediments;
- Assessment of availability of graduate students and young researchers that could be newly employed at SCL.
- Assessment of indirect social impacts of the process of strengthening of SCL and its efficient integration into a wider European R&D effort.

The outlined analysis has resulted in the following specific objectives, each of which directly leads to a set of measurable and directly verifiable sub-objectives.

Objective 1 – Enhance quality of R&D at SCL

- **Sub-objective 1.1**: Set up an International Advisory Board for the new centre of excellence;
- **Sub-objective 1.2**: Establish a framework for more efficient management of research at SCL by developing a flexible, problem oriented R&D plan that will successfully integrate that research into a wider European effort.
- **Sub-objective 1.3**: Develop a specific set of benchmarks for tracking the quality of R&D at SCL, and perform a benchmarking exercise.
- **Sub-objective 1.4**: Devise and implement a long term strategy for achieving and maintaining research excellence.
- **Sub-objective 1.5**: Insure viability of SCL as a centre of excellence beyond the project lifetime by finding other sources of funding.

Objective 2 – Expand and mobilize human resources

- **Sub-objective 2.1**: Recruit and employ young researchers; develop explicit career plans for the newly employed researchers.
- **Sub-objective 2.2**: Enhance working conditions for young researchers by setting up an R&D environment at SCL that is integrated into ERA, providing challenging research problems, state of the art equipment, and enhanced mobility.

Objective 3 – Reinforce existing S&T capacities at SCL

- **Sub-objective 3.1**: Maintain and upgrade existing S&T equipment and high-tech infrastructure.
- **Sub-objective 3.2**: Improve the availability and reliability of SCL's computing resources, determine and implement optimal strategies for their use.

Objective 4 – Enhance mobility and integration into ERA

- **Sub-objective 4.1**: Network with EU, regional and national partner institutions through exchange of personnel, research results and joint numerical experiments; participate in joint RTD activities within these networks.
- **Sub-objective 4.2**: Host scientists from EU for training and research.
- **Sub-objective 4.3**: Organize training of graduate students and young researchers through short-term missions at EU institutions.

Objective 5 – Contribute to the reinforcing of ICT capacities at the national level

- **Sub-objective 5.1**: Reinforce the quality of research in SCL's immediate R&D environment, by strengthening their human capacity through stipends, yearly visits, and by conducting joint research activities.
- **Sub-objective 5.2**: Reinforce human capacity in Serbia's high performance computing sector by training young researchers to be employed at national research institutions and hi-tech companies.
- **Sub-objective 5.3**: Contribute to the national R&D system by developing a set of recommendations for policy makers at national and local levels for fostering growth of research excellence in a rapidly changing high-tech environment.

.

The three year CX-CMCS project kicked-off on July 1, 2006. The project plans to issue the following deliverables:

Deliverable No	Deliverable title	Delivery date	Nature	Dissemination level
D01	CX-CMCS Web site	M1	R	PU
D02	Career development plan for newly employed young researchers	M2	R	CO
D03	CX-CMCS International Advisory Board	M3	O	PU
D04	Equipment tendering and procurement report	M3	R	PU
D05	Inauguration meeting report	M4	R	PU
D06	Mobility and training plan	M6	R	PU
D07	CX-CMCS Brochure	M6	R	PU
D08	12M Progress reports	M12, M24	R	PU
D09	CX-CMCS Promotional video material	M15	O	PU
D10	Benchmark procedures for quality assessment of RTD centres of excellence	M18	R	PU
D11	SCL research quality assessment	M24	R	PU
D12	Proceedings of International dissemination workshop	M30	R	PU
D13	Strategy of long term sustainable growth of research excellence in transition	M30	R	PU
D14	Scientific computing landscape of Serbia	M33	R	PU
D15	Presentation of policy papers to decision makers	M34	R	PU
D16	Final project report	M36	R	PU

Legend: R = Report, O = Other, PU = Public, CO = Confidential (only for members of the consortium incl. EC).

Table of contents

Table of figures

References

[1] Project CX-CMCS – 026343 – Annex I – Description of Work

Executive summary

What is the focus of this Deliverable?

The focus of this deliverable is a detailed technical report about equipment tendering and upgrading related to the reinforcement of SCL research infrastructure.

What is next in the process to deliver the CX-CMCS results?

The deliverable and workflow progress is described in the project Annex-I – Description of Work [1]. The next deliverable in the activity A2.3 (Equipment upgrade) will be deliverable D04-b to be delivered after the procurement and installation of the new equipment is finished.

What are the deliverable contents?

This deliverable contains the detailed technical presentation of original plans for equipment upgrade and of the actual upgrade specifications.

Conclusions

Though delayed, the upgrading of SCL distributed computer infrastructure planned by the CX-CMCS project is appropriately modified and in the process of implementation.

1. Plan of equipment upgrade

At the start of the CX-CMCS project, SCL had at its disposal state of the art computing and communication resources, and partial support was planned within WP2 for maintaining and upgrading of existing S&T equipment and high-tech infrastructure (see Sub-objective 3.1 of the CX-CMCS Annex-I Description of Work). The activity A2.3 (Equipment upgrade) deals with the upgrade of the existing computing infrastructure.

Figure 1 –AEGIS01-PHY-SCL Grid cluster at SCL

Reinforcement of SCL's research capacity through CX-CMCS funds includes a plan to upgrade the existing hi-tech infrastructure by getting rid of its three existing weaknesses:

- Lack of adequate storage space to tackle an important array of modern GRID applications;
- Inadequate number of stackable high throughput switches to optimally link all the existing nodes into a more versatile cluster;
- Inadequate RAM for fine-grained applications that are difficult to parallelise (e.g. weather forecasting, fault propagation in materials, searching through distributed databases such as GRID based medical databases, human genome databases, etc.).

The upgrade plan of the PARADOX cluster had envisaged the following equipment to be procured from CX-CMCS funds:

- Storage element (5 TBytes capacity with possibility of further upgrade, fibre channel interface, less than 5ms seek time, two storage processors, Raid 0-5 capability, with appropriate software for SAN management, rack chassis, 3kVA UPS)
- Network switch (stackable high throughput switch, 48 port, layer 2, 10/100/1000, LAN switching, bandwidth: 156 Gbps or better, throughput: 115 Mpps or better)
- RAM (33 GB in 512MB modules, FCC registered DDR RAM)

However, due to the long delay between the submission of the proposal (May 2005) and its actual start (July 2006), the situation on the ICT equipment market has drastically changed in terms of prices, availability of components and their performances. This allowed us to plan procurement of storage elements of larger capacities than planned, as well as more RAM modules of higher capacity for the allotted amount of funds from the CX-CMCS budget.

Figure 2 – No. 3 Rack with Worker nodes

Another change to the equipment upgrade plan due to the above described delay is that the continuing need for high throughput LAN urged SCL to procure network switches from other funding sources. Instead of one layer 2 network switch with 48 Gbps ports, SCL procured three 3Com 3870 layer 3 stackable 48 port Gbps high throughput switches. These switches are stacked into an aggregated switching facility of 144 ports, with the Gigabit uplink connection to Serbian National Research and Education Network. In addition to this, three UPS units of 10 kVA capacity (APC RT10000) were also procured from other funding sources, since the maintenance of the PARADOX cluster requires stable environment.

The procurement of network switches and UPS units from other sources allowed SCL to use all resources allotted for the equipment upgrade for the procurement of storage elements and RAM upgrade.

One of the prime objectives of CX-CMCS is to aggressively seek other sources of funding to ensure sustainability of SCL's funding beyond the project's lifetime. As was already mentioned, even before the project started additional funds were secured for purchase of network switches and UPS units. Due to intense media relations and outreach activities of the CX-CMCS project, we made contact with regional AMD representatives, and have secured their in-kind donation in CPUs. Ten AMD CPUs were therefore obtained free of charge: 6 x AMD Opteron DualCore 285 2.6Ghz CPUs which are deployed in storage servers, and 4 x AMD Opteron DualCore 8218 2.6Ghz CPUs which are deployed in the storage element.

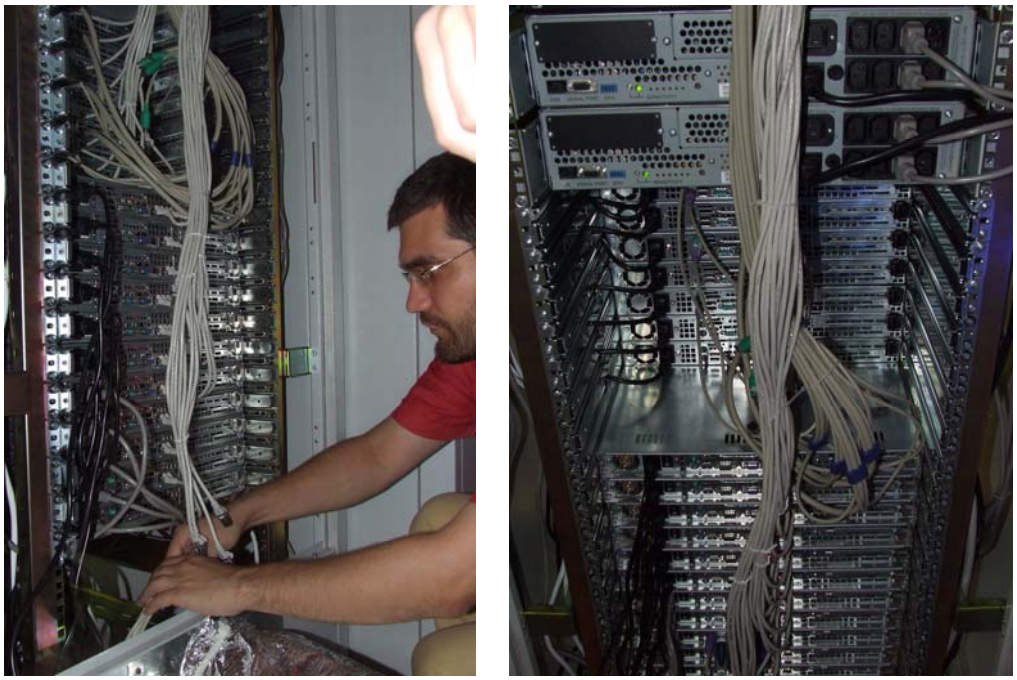

Figure 3 – Details of equipment integration after the upgrading of stackable switches

2. Technical specifications of the equipment to be procured

After careful consideration of the components available on the ICT market, we decided to procure:

- 3 storage units, each of the aggregated storage capacity of 9 TB;
- 8 storage servers providing access to storage (six dual-core CPUs free of charge as a part of AMD in-kind donation);
- 1 storage element enabling data analysis (four dual-core CPUs free of charge as a part of AMD in-kind donation);
- 60 GB DDR2 ECC RAM in 1 GB modules (266 MHz PC2100);
- 10 GB DDR2 ECC RAM in 1 GB modules (667 MHz PC2-5300).

Detailed technical specification of the first three items is given as follows:

1) Storage units

- **Chassis: 4U ATX 16+2 SATA - Redundant 650W**
  - Chassis 4U rackmount for EATX motherboard (max 12" x 13");
  - Redundant triple power supply (2+1) 650W PFC;
  - Size 7" (177 mm) H x 19" (482 mm) W x 27" (685 mm) D;
  - 16+2 bays for 3.5" SATA HDD;
  - 1 x 5.25" bay for floppy and CD/DVD slim;
  - 5 x 80mm fans for system cooling and 2 x 80mm rear fans to extract used up air;
  - Front panel button for Power on/off, reset, power LED, HDD LED, LAN LED, 2 USB and 1 serial RJ45 port
- **Motherboard: 1 x Dual Xeon 5000P Server**
  - Motherboard EATX format 13.05" x 12", (33,1cm x 30,5cm) (W x H);
  - 2 x 771-pin LGA socket supporting Intel® Xeon® 64-bit 5100/5000 (Woodcrest/Dempsey) processors;
  - 1333/1066/667 MHz bus system;
  - chipset Intel® 5000P;
  - 8 DIMM socket 240 pin, supporting up to 32GB memory DDR2 667/533 Mhz Fully Buffered DIMM ECC;
  - 2 PCI-Express 8x, 1 PCI-express 4x (on 8x slot), 1 slot SEPC (Supermicro PCI-e Power connector for 2U riser card active), 2 PCI-X slot 133MHz 64-bit, 1 PCI-X slot 100MHz 64-bit; optional support for IPMI 2.0;
  - On board devices: SATA2 ESB2 controller supporting up to 6 HD in RAID 0,1,5,10 mode;
  - LAN card: Gigabit Intel® PRO1000 (ESB2/Gilgal) 82563EB dual port;
  - VGA ATI ES1000 16MB;
  - connectors: 6 SATA ports, 1 IDE channel, floppy, 2 LAN, 2 ports + 3 header USB 2.0 (max 5 USB 2.0), 1 VGA, 2 PS2 for keyboard and mouse, 1+0 serial ports, 8Mb Flash EEPROM Phoenix® BIOS;
  - supports up to four 3pin fans and four
  - 4pin tachometric fans.
- **Processors: 2 x Xeon 5130 2.00Ghz 4MB 1333Mhz Act-1U**
  - 2 x Intel® Xeon™ Woodcrest Dual Core Processor 2.00Ghz, FC -LGA6 771-Pin socket, 65W;

  - external cache: 4MB, bus frequency 1333Mhz, with 64-bit support, XDBIT (Execute Disable Bit), Virtualization Technology and DBS (Demand-Based switching);
  - Boxed model inclusive of active 1U heatsink.
- **RAM: 4 x DDR2-667 Reg. ECC FBDIMM 1 GB Module**
- **RAID controller for OS disks: 1 x 3Ware RAID 8006-2**
  - Full size 64/32/bit PCI EIDE Controller for Workstations;
  - Supports up to 2 Serial ATA.
  - RAID Levels: 0, 1, JBOD.
  - Supports Hot-Swap, Hot spare, DMI, S.M.A.R.T.
  - Windows XP/2000/NT4/Me/9x; RedHat Linux 6.2/7.1/7.2; SuSE Linux 6.2/7.2/7.3 operating systems support.
- **RAID controller for storage disks: 1 x 16 port PCI-X Multi Lan Raid**
  - Serial ATA controller with connection for up to 16 SATA or SATAII drives- multilane model;
  - PCI-X on 64Bit/133MHz Bus. Intel IOP 80331 I/O processor;
  - SO-DIMM socket for up to 1GB memory SDRAMM DDR333 ECC;
  - Support for RAID 0,1,(10), 3, 5, 6 and JBOD;
  - Multiple RAID, array roaming online, on line RAID migration (level and stripe), simultaneous on line expansion and migration, background start up which allows immediate system usage;
  - Supports automatic insertion or removal and automatic rebuilding;
  - Ethernet support;
  - Compatible OS: Windows2000/XP/Server 2003, RedHat Linux, SuSE Linux, FreeBSD.
- **1 x Backup Battery**
  - Backup battery for SATA raid controller.
- **HDD for OS: 2 x WD 160GB SATA II 7.200 RPM 16MB RAID**
  - Western Digital caviar RE SATAII: Serial ATA interface HDD, capacity 160GB;
  - Speed: 7200RPM, average access time: 8.9 ms;
  - Internal buffer: 16MB;
  - MTBF: 1 million hours;
  - Reliability: 24x7x365 in Duty Cycle.
- **HDD for disk storage: 16 x SEAGATE 750GB SATAII 7.200 RPM 16MB ES**
  - Serial ATA interface HDD, capacity: 750GB;
  - Speed: 7200RPM, average access time: 8.2 ms;
  - Internal buffer: 16MB;
  - MTBF: 1.2 million hours;
  - Vibration tolerant;
  - Reliability: 24x7x365 in Duty cycle.
- **1 x 2 bays 1" SATA Backplane**
- **4 x Backplane 4 bays 1" SATA Multilane**
- **LAN: 1 x Intel Pro 1000GT Server Quad Port 8494GT**
- **Management: 1 x IPMI Kit SIM1U Modul**
  - Card for remote control according to IPMI 2.0 specifications, it allows to monitor and intervene on the main parameters (CPU and system's temperature, fan status, power voltage);
  - It provides a remote feed management both by interfacing with the OS in use (shut down) and independently (reset, power down, power up, power cycle);
  - It permits the control of stand-by and/or crashed system;
  - Features: SOL (serial over LAN) for console re-direction and remote access to BIOS, Virtual Media over LAN (Virtual USB Floppy/CD and drive re-direction), LAN alert SNMP Trap, event's log;
  - IPMI functions are accessed via dedicated Gbps LAN or Gbos LAN integrated on motherboard;
  - Can use client Java (Windows/Linux), Web interface or CLI;
  - IPMI functions are password protected.
- **1 x Slide Rail Kit**
  - Slide rail kit to turn chassis tower into rack.
- **1 x 3 years On-Centre Warranty for Server**

2) Storage servers

- **Chassis: 1 x 1U - 4 x SATA - 500W**
  - 1U Rackmount ATX Chassis;
  - 500W PFC Cold-Swap Power Supply;
  - dimensions 1,7"(43 mm) H x 17,2"(437 mm) W x 25,6"(650 mm) D;
  - 1 x 5" external drive bay for slim CD-ROM + 1 x 3" slim floppy drive;
  - 4 x Hot-Swap SATA HDD;
  - Power On-Off button on the front panel, LED indicators for Power-On, HDD, LAN, fan, overheat;
  - Supports max. motherboard size - 12" x 10" E-ATX.
- **Motherboard: 1 x Dual Opteron Dual Core - nForce 2200 - Server - HT**
  - Motherboard Extended ATX 12” x 13.05" (30.5cm x 33.2cm) (W x H);
  - Two 940-pin socket supporting AMD Opteron™ processor series 200 (also Dual Core);
  - nVidia nForce Pro 2200 (CK804), nVidia nForce Pro 2050 (CKIO4) chipset;
  - 8 sockets DIMM 184 pin supporting up to 32GB memory DDR 400/333/266 Mhz buffered registered ECC in single/dual channel configuration;
  - 1 HTX slot, 1 PCI-Express slots x16 with nVidia SLI technology, 1 PCI-Express slot x4, 3 PCI slot 33MHz 32-bit slot, support for IPMI 2.0 controller;
  - On board devices: SATA2 nVidia nForce Pro 2200 e nVidia nForce Pro 2050 controller 3 Gbps, supporting up to 8 SATA HD in RAID 0, 1, 10, audio AC'97;
  - 2 x Broadcom gigabit (onboard) Intel Pro Network Adapter, PCI, 10BASE-T/100BASE-TX/1000BASE-T, RJ-45 connector;
  - Connectors: 8 SATA channels, 2 IDE channel, floppy, 2 LAN, up to 8 USB 2.0 (4 ports + 2 header), keyboard, mouse, 1 + 1 serial, 1 parallel, Line In/Line Out/Microfono, 4Mb Flash EEPROM AMI® BIOS;
  - Supports up to eight 4pin fans.
- **Processors: 2 x AMD Opteron DualCore 285 2,6Ghz 2MB 1000Mhz**
  - 2 x AMD Opteron Dual Core 285, with 2,6Ghz up to 2-way scalability;
  - Cache size L1: 64KB shared data, 64KB instruction (per core), 1MB cache L2 (per core);
  - Bus frequency 800Mhz or 1000Mhz;
  - Integrated DDR Memory controller;
  - Simultaneous 32 & 64 bit computing.
- **RAM: 4 x 1 GB LP PC3200 Reg. ECC DDR Module**
- **HDD: 1 x HITACHI 80GB SATA II 7.200 RPM**
  - Hitachi Deskstar 7K80: Serial ATA II interface, 80GB capacity;
  - 7200 RPM speed, average seek time 8.5 ms;
  - 8MB buffer size.
- **1 x 4 bays 1" SATA Backplane**
- **1 x 2 years On-Centre Warranty**

3) Storage element

- **Chassis: 1U**
  - Chassis support for max. motherboard size - 14.3" x 17"
  - 1000W high-efficiency power supply w/ PFC
  - Size 1.7" (43 mm) H x 17.2" (437 mm) W x 27.75" (705 mm) D;
  - 3x 3.5" hot-swap U320/160 SAS / SATA drive bays with SES2;
  - Slim DVD-ROM drive;
  - 6x heavy-duty counter-rotating fans with optimal fan speed control;
  - Front panel button for Power on/off, reset, power LED, HDD LED, 2 x LAN LED, System Overheat LED, 2 USB and 1 serial RJ45 port
- **Motherboard: 1 x Quad AMD® Opteron™ 8000 Series**
  - Quad AMD® Opteron™ 8000 Series (Socket F) Support, 1000 MHz HyperTransport Link;
  - nVidia MCP55 Pro / AMD 8132 Chipsets;
  - Dimensions:16" x 13" (40.6cm x 33.2cm)
  - Up to 64GB DDR2 667 SDRAM, Up to 32GB DDR2 533 SDRAM, Up to 64GB DDR2 400 SDRAM;
  - Intel 82546GB Dual-port Gigabit LAN/Ethernet Controller
  - 6 SATA2 3.0Gbps Ports
  - 1 (x16), 1 (x8) PCI-Express, 2 64-bit 133/100MHz PCI-X, 2 64-bit 100MHz PCI-X, 2 HyperTransport (HTX)
  - SIMSO IPMI 2.0 Support
  - 9 Fan Support with Speed Control
- **Processors: 4 x AMD Opteron DualCore 8218 2.6Ghz 2MB 1000Mhz**
  - 2 x AMD Opteron Dual Core 8218, with 2.6Ghz up to 8-way scalability;
  - Cache size L1: 64KB shared data, 64KB instruction (per core), 1MB cache L2 (per core);
  - Bus frequency 1000Mhz;
  - Integrated DDR Memory controller;
  - Simultaneous 32 & 64 bit computing.
- **RAM: 16 x DDR2-667 Reg. ECC FBDIMM 2 GB Module**
- **HDD: 2 x WD 500GB SATA II 7.200 RPM 16MB**
  - Western Digital WD500YS SATAII: Serial ATA interface HDD, capacity 500GB;
  - Speed: 7200RPM, average access time: 8.7 ms;
  - Internal buffer: 16MB;
  - MTBF: 1.2 million hours;
  - Reliability: 24x7x365 in Duty Cycle.
- **SAS HDD Backplane with SES2**
- **1 x 2 years On-Centre Warranty**

3. Further procurement and installation procedure

The aforementioned delay between proposal submission and Project start has caused additional procedural problems, connected with the changes of the rules for importing computing equipment that came in the effect in the mean time. The procedure to get exempted from the VAT has become particularly cumbersome and time consuming. Since at the time of writing of this deliverable it was not possible to estimate the time frame within which the procurement could be finished, it was decided to produce another version of this deliverable – D04-b – at the end of the procurement and installation process.