INVITATION TO TENDER

ECMWF/2014/216

for the
Acquisition of
Disk Systems and Servers

Volume II:

SPECIFICATION OF REQUIREMENTS

14 July 2014

TRADEMARKS

All names or descriptions used in this Invitation To Tender (ITT) which are trademarks, trade or brand names, or other references to proprietary products are hereby acknowledged as the property of their respective owners. No entry, term or definition in this ITT should be regarded as having any implication as to the validity or otherwise of any trademark.

The appearance of any proprietary name or reference in this document should not in itself be taken to imply a preference for one product over another unless specifically stated otherwise.

Table of Contents

1. Introduction

1.1. What is ECMWF?

The European Centre for Medium-Range Weather Forecasts (ECMWF) is an independent intergovernmental organisation supported by 34 states. ECMWF is both a research institute and a 24/7 operational service, producing and disseminating numerical weather predictions to its Member States.

ECMWF operates one of the largest supercomputing facilities of its type in Europe and holds the largest meteorological data archive in the world. General information about ECMWF can be found at

http://www.ecmwf.int/en/about.

1.2. Scope of the Invitation To Tender

To maintain and enhance parts of its infrastructure, ECMWF wishes to select a vendor from which it will acquire **Servers** and **Disk Systems** and associated infrastructure for a period of five years. It is expected that between 1.5 and 2 Million Pounds would be spent to acquire and support such equipment each year. The percentage of this money allocated to servers is expected to vary significantly over the years. As an estimate, the amount spent on Servers would be about 30% of the total, where the amount spent on Disk Systems and related infrastructure would represent around 70% of the total.

ECMWF would prefer to award a single contract covering both servers and disk storage. However ECMWF reserves the right to award two separate contracts, one covering the acquisition of **Disk Systems** (Lot 1), the other covering the acquisition of **Servers** (Lot 2). For this reason ECMWF would accept bids covering either one or both of these Lots.

While this Invitation To Tender covers the **Initial Acquisition** of equipment, as defined in sections 3 and 4 of this document, successful Tenderer(s) will be required to sign a call-off contract, based on the Draft contract found in Volume III, allowing the acquisition of further equipment over the following five years with a possibility of extension, and be prepared to stand by discount levels.

1.3. Definitions and units.

The words or phrases below have the meaning ascribed to them when using **bold** font style.

Disk A rewritable storage device where data are recorded persistently by various electronic, magnetic, optical, or mechanical changes to a surface layer of one or more rotating disks, or by the use of integrated circuit assemblies as memory to store data persistently.

Disk System A self-contained piece of equipment, connectable to computer hosts, and able to provide to these hosts access to one or more logical disk, or LUN. *Such equipment could for example include controller units, to which disk containers such as disk trays are connected. A controller unit would include one or more disk controllers, in charge of organising disk space and distributing it across connected servers, providing read and write cache and protecting the storage against disk or power failures. A disk container would include a collection of physical disk devices.*

Initial Acquisition Acquisition from successful Tenderer(s) of a first lot of **Disk Systems** and a first lot of **Servers** complying with the requirements defined in sections 3 and 4 of this document.

LUN A logical disk, that provides an area of usable storage capacity on one or more physical Disk components in a Disk System. This logical disk can be connected to one or more hosts, which can use standard SCSI protocol commands to position, read and write data on the LUN.

Server A computer in a network, capable of running a multiprocessing operating system and used to provide services to other computers in the network.

System Disk Systems and/or Servers.

Usable Capacity For Disk Systems, Usable Capacity is defined as the amount of storage space that can be presented to attached hosts, not counting any space required for redundancy, mirroring, or tiering. *For example, if a Disk System was composed of 64 1TB disks, configured as 6 8+2 RAID6 arrays and four hot spare disks, the Usable Capacity provided by that Disk System would consist of 48TB*.

Works All tangible items and software products, and software licences furnished and all services to be performed by the Tenderer.

Storage and bandwidth units used in this document.

Storage Units

MB	1000x1000 Bytes.
MiB	1024x1024 Bytes.
GB	1000x1000x1000 Bytes.
GiB	1024x1024x1024 Bytes.
TB	1000x1000x1000x1000 Bytes.
TiB	1024x1024x1024x1024 Bytes.

Bandwidth and throughput units

MiB/s	1024x1024 Bytes/sec.
GiB/s	1024*1024*1024 Bytes/sec.
Gib/s	1024*1024*1024 bits/sec.

Power units

kWh	1,000 watt-hours

1.4. Background to this acquisition

ECMWF operates multiple computing suites as shown in Figure 1.

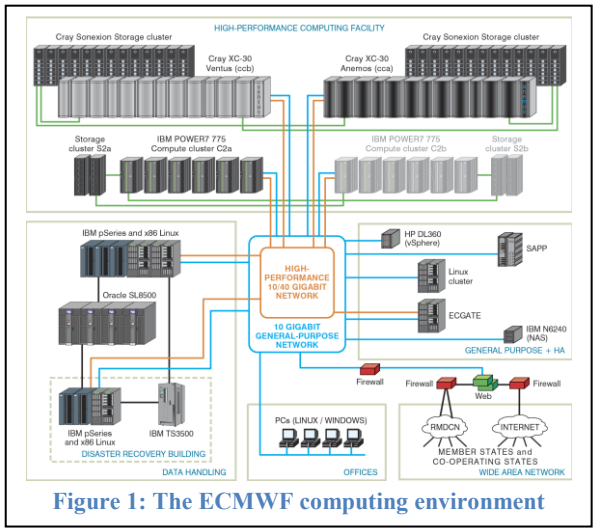

Figure 1: The ECMWF computing environment

1.4.1 ECMWF's supercomputers

ECMWF's High Performance Computing Facility is in the middle of being migrated from a service based on IBM POWER7 technology to a Cray XC30 based service. These are used to run large mathematical models allowing the Centre to predict the weather worldwide over periods of several weeks. Note that this ITT does not cover the provision of any equipment related to ECMWF's supercomputers.

1.4.2 ECMWF's Data Handling System.

ECMWF maintain a large database of weather-related information. The Data Handling System (DHS) currently holds 70 petabytes of data, the bulk of it being pre-compressed and scientific in nature. This data is stored in a tiered environment controlled by the High Performance Storage System (HPSS). Most data is stored on tape media, with only 3% of the data residing on disk. HPSS is highly distributed software, making use of many x86 servers to transfer data between disks or tape drives and users archiving or accessing the DHS data, via a high performance network.

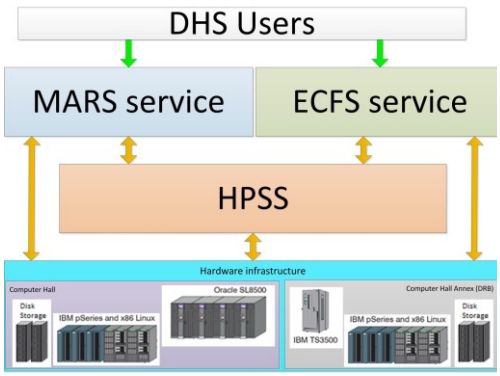

Figure 2: The Data Handling System

Two ECMWF-written applications (MARS and ECFS) provide easy-to-use interfaces to DHS users.

Tape drives and disk storage are connected to the x86 servers through several Storage Area Networks (SAN) switches. It is expected that more SAN switches will be acquired as part of the call-off contracts signed with the Tenderer selected as a result of this ITT.

1.4.3 Other Linux and Windows clusters.

Several other clusters provide facilities allowing users and operational suites to run programs not suitable for the super computers, such as data acquisition, data visualisation, data transfer between the Centre and external organisations, etc., or to provide some Windows based service. Some of the servers and disk storage that are expected to be acquired under the call-off contracts with the winning Tenderer will be used to satisfy new requirements in this area.

1.4.4 Desktops

Staff and consultants working at ECMWF access all computing facilities via x86 desktop machines running Windows or Linux, and connected to the Local Area Network.

1.4.5 Local Area Network

The Local Area Network (LAN) connects all our computing systems. The LAN is divided into a high-performance network (HPN) and a general-purpose network (GPN). The HPN transfers large amounts of operational data between the supercomputer and the data handling system. The GPN provides connectivity to the supercomputer, the DHS, other large servers, user desktops and laptops and via firewalls to the outside world. Most servers to be acquired in the context of this ITT will be connected to the HPN, through dual 10Gib/s interfaces.

1.5. Organisation of this document

The remainder of this document is organised as follows:

- Section 2 contains some general considerations including the timetable for the procurement;
- Section 3 contains technical requirements for disk storage procurement;
- Section 4 contains technical requirements for servers.

2. General Considerations

2.1. Tender requirements

Points of specification are categorised by the bold notations **M, H, D** or **R** to the left of the pertinent section:

M	Denotes a MANDATORY requirement: a requirement that must be adhered to, or a performance requirement that must be met in order that the tendered solution can be accepted by ECMWF. ECMWF will not consider a tendered solution that fails to meet a mandatory specification requirement (marked **M**) unless the Tenderer offers valid reasons why the feature in question is either unnecessary for, or irrelevant to the tendered solution.
H	Denotes a HIGHLY DESIRABLE requirement. The degree to which a tender meets the highly desirable requirements (marked **H**) will be a key factor that will be taken into account in selecting the winning tender. If such a feature cannot be supplied initially, the Tenderer should state whether, and if so when, it will be added to the installed system.
D	Denotes a DESIRABLE feature. The extent to which any tender offers features listed as desirable (marked **D**) will be one of the factors taken into account in selecting the successful tender. If such a feature cannot be supplied initially, the Tenderer should state whether, and if so when, it will be added to the installed system.
R	Denotes a REQUEST for information. A response must be given to all such requests. Requests for information (marked **R**) are intended to prompt a description of the construction, philosophy or operation of the tendered solution in areas that are regarded as being of particular importance. A clear response to such requests will be of assistance to ECMWF in the tender evaluation process.

Tenderers should note that responses to points of specification that are met must include sufficient detail to explain the way in which the requirement is met; a simple expression, such as “compliant” or “agreed”, will not normally suffice.

Any additional features not listed in the ITT as requirements, which the Tenderer feels may be relevant, should be supported by descriptive material.

M 1. When bidding for both lots, **Disk Systems** and **Servers**, Tenderers must ensure that all points of specification in this ITT are addressed. When bidding for the **Disk System** lot only, all points in sections 2 and 3 must be addressed. When bidding for the **Server** lot only, all points in sections 2 and 4 must be addressed. Please note that partial/incomplete responses will not be considered.

2.2. Schedule

The closing date and time for this ITT is **14:00 local UK time on Tuesday 26 August 2014**. Following receipt of replies to this ITT, ECMWF envisages the following schedule for the implementation of the project:

September 2014	Evaluation of tenders.
11-16 September 2014	Tenderers may be invited to present their company and Tender at ECMWF's premises.
21 September-Mid November 2014	Contract negotiations with Tenderers.
October-December 2014	Approval of the selection recommendations by ECMWF internal bodies.
Mid December 2014	Contract signature.
January 2015	Installation and acceptance of the equipment part of the **Initial Acquisition**.

2.3. Presentations and references

M 2. Tenderers must be prepared, if requested by ECMWF, to present their tender at ECMWF. The exact date and time of the presentation will be made known following receipt of the tenders, but is expected to be on 11, 12, 15 or 16 September.

R 3. Tenderers shall provide names and contact details of at least three existing customers, preferably within Europe, who are already using equipment of the kind proposed in response to this ITT and to whom reference may be made. Customers using the equipment in an environment similar to that envisaged at ECMWF are preferred.

R 4. Tenderers shall provide information about the number of company staff working worldwide, in Europe and in the UK. They will indicate how many of these work in support for **Disk Systems** and/or how many work in support for **Servers**.

2.4. Contract

M 5. Volume III contains a draft agreement giving the terms and conditions under which ECMWF wishes to acquire equipment. Tenderers must state their acceptance or non-acceptance of the individual conditions contained in the draft agreement according to the rules given in section 11.3 of Volume I of this ITT.

M 6. Tenderers must complete and return as part of their tender the "Costing and critical information spreadsheet", copy of which is found in Annex A.

2.5. Pricing

M 7. All prices must:

- Encompass the costs of the mandatory features as described in this ITT;
- Be inclusive of all equipment delivered;
- Be inclusive of shipping, delivery and site preparation charges;
- Be inclusive of installation charges and miscellaneous ancillary equipment;
- Be inclusive of all software and firmware licences relating to the **Works** over a period of 5 years, including but not limited to any cost related to licence keys required to unlock features required to satisfy the requirements set below;
- Be inclusive of all Maintenance and Support Services over a period of 5 years;
- Be exclusive of all import duties and UK taxes;
- Be quoted in Pounds Sterling;

For the **Initial Acquisition** the prices must be valid until six months after the closing date for receipt of tenders.

M 8. Items deemed desirable or highly desirable in this ITT must be priced separately where applicable. Where the pricing of such an item depends upon the implementation or installation of an item not marked **M**, this dependency must be explicitly stated.

M 9. Tenderers must provide a summary of their pricing, quoting separately any:

- Design and consultancy charges;
- Project Management charges;
- Installation charges;
- Hardware & software maintenance & support;
- Training programme.

M 10. Tenderers must guarantee to supply, in addition to the **Initial Acquisition** and at ECMWF's request, further appropriate equipment and software licences or services on the terms of the contract.

M 11. Tenderers must give guaranteed percentage levels of discount on their then current list prices for all future equipment and software licences and maintenance acquired under this agreement. However, if such a form of price guarantee is not acceptable to the Tenderer, then it must include an alternative proposal for price variation for all the components of the **Works**; such proposal must include a formula for calculating ceiling prices that offers equivalent value for money over the life of the call-off contract.

R 12. Tenderers shall provide their list prices current at the closing date of this ITT and must also explain how future list prices will be made available to ECMWF during the course of the contract.

2.6. Project Management, Installation and Acceptance Tests

Tenderers should take note of section 17 of the draft agreement (Volume III) that outlines contract management and monitoring.

Tenderers should take note of section 6 of the draft agreement (Volume III) that contains the standards of the tests which must be met by the **System** before it will be accepted by ECMWF.

ECMWF together with the Tenderer will prepare the configuration of the **System** to allow initial testing. ECMWF will be responsible for maintaining the system configuration after the installation phase.

2.7. Maintenance and support of the tendered hardware and software

M 13. Tenderers must explicitly undertake to provide, or to arrange provision of, maintenance, spare parts and support for the hardware and software acquired under this ITT in accordance with Sections 11, 12 and 13, and Exhibit 6 of the draft agreement described in Volume III for a period of not less than 5 years from the relevant acceptance date. Tenderers should indicate to what extend this will be done using the **Works** standard warranty, and where service uplifts have been used. Where this involves an arrangement with a third party, responsibility for the provision of such maintenance and support must in any case rest with the Tenderer. If a Tenderer proposes to use a third party this must be clearly stated and full details of the company providing support must be given.

M 14. For **Disk Systems** acquired in the **Initial Acquisition**, with the exclusion of the **Disks** found in the **Disk Systems**, Tenderers must commit to the provision of maintenance and support services for 24 hours per day, 7 days per week for software and hardware support with a response time of less than 4 hours. Where remedial action is required, Tenderers must commence such remedial action within these 4 hours.

M 15. For **Servers** and for the **Disks** included in the **Disk Systems** acquired in the **Initial Acquisition**, Tenderers must commit to the provision of maintenance and support services on next business day, for software and hardware support with an acknowledgment within 4 hours of ECMWF's notification. Where remedial action is required, Tenderers must commence such remedial action no later than the next business day.

R 16. Tenderers shall describe their policy regarding maintenance, spare parts and support for the hardware and software of the tendered **Systems** when reaching their end of life, after the period committed to above.

R 17. Tenderers shall describe how they will provide the maintenance specified in **M 13**, **M 14**, and **M 15**. In particular they should cover how calls are taken and tracked, any escalation procedures and the involvement and responsibilities of any third party equipment manufacturers. They shall also describe whether onsite and remote support is provided by their own engineers and employees, and at what point third party engineers and employees are expected to take over management and resolution of reported issues.

R 18. If a Tenderer depends on third party support (e.g. reseller making use of equipment manufacturer support), said Tenderer will explain how ECMWF is able to follow the status of calls entered on behalf of ECMWF with said third party. The Tenderer will also explain what methods can be used to escalate issues, and at what point direct contact between ECMWF and said third party can be established.

R 19. Tenderers shall describe what type of help and support will be made available to ECMWF once the tendered **System** has been installed and accepted. In particular, Tenderers must state which of the following will be available:

- Remote software diagnosis service;
- Corrective action to rectify software bugs and performance limiting factors;
- Response to simple telephone enquiries;
- Detailed response to more complex enquiries (telephoned/written);
- Escalation procedures for problems which cannot be fixed in a timely manner;
- Software upgrades.

R 20. Tenderers shall describe the conditions under which they expect ECMWF staff to perform remedial actions such as the installation of customer replaceable units.

2.8. Training

M 21. Tenderers must propose a training programme, preferably given at ECMWF's premises, which must provide analysts with sufficient understanding of the internal working of the hardware and software being tendered to enable them to provide effective day-to-day and emergency support, including software upgrades and changes to system configuration. Tenderers bidding for **Disk Systems** must provide training for five analysts. Tenderers bidding for **Servers** must provide training for five analysts.

2.9. Documentation

M 22. Tenderers must provide documentation and manuals for the configuration, operation and maintenance of the tendered **Systems** in a printable electronic form (preferably in PDF and HTML format). Such material must be made available prior to installation and all necessary updates must be provided free of all charges for the duration of the contract.

R 23. Tenderers shall state whether on-line access to such material is available via a web-site that would be accessible to ECMWF staff, and provide any instructions necessary for them to gain access to the web pages.

R 24. Tenderers shall provide information about their certification to any quality assurance standards (e.g. ISO 9001) and also that of any of their subcontractors if applicable.

3. Disk Systems Technical Requirements

This section specifies ECWMF's technical requirements for the **Disk Systems** in the **Initial Acquisition**. Equipment to be delivered at a later stage under the conditions of the call-off contract may or may not be subject to similar requirements, as discussed at the time of purchase.

The acquisition of new **Disk Systems** is mainly driven by the needs of ECMWF's DHS. As the amount of data stored in ECMWF archives increases, the amount of disk space used as a high level hierarchy tier will grow over the years. Historically, the amount of disk space available to the DHS has increased by between 40% and 50% per year.

Table 1 below provides an estimate of the disk requirements for the DHS and other servers over the next 5 years, assuming that the growth in disk space was continuing at a rate comparable to the one seen historically. All values are for Terabytes of **Usable Capacity**.

Year	Requirements for additional usable space	Replacement of obsolete Disk Systems	Total requirements for additional space
2013 (1)	*1,870 TB*		
2014 (2)	1,000 TB	150 TB	1,150 TB
2015	1,400 TB	0 TB	1,400 TB
2016	2,100 TB	0 TB	2,100 TB
2017	3,150 TB	200 TB	3,350 TB
2018	4,700 TB	450 TB	5,150 TB
2019	7,100 TB	600 TB	7,700 TB
(1)	*Total amount of space at year end.*		
(2)	*Not covered by this ITT*		

Table 1: Estimate of the disks storage requirements for the DHS and other servers.

It is expected that several independent **Disk Systems** would be acquired each year to support the new disk requirements. Each of these **Disk Systems** would be a self-contained piece of equipment, connectable to hosts, and able to provide LUN access to these hosts.

The **Disk Systems** initially acquired in the context of this tender will be used to provide block devices to the connected hosts, through Fibre Channel (FC) SAN connectivity.

3.1. Capacity and Performance

While the Tenderer selected as a result of this ITT is expected to deliver several lots of equipment to ECMWF, this ITT requires for the Tenderer to deliver equipment that will cover part of ECMWF's storage requirements for 2015. It is envisaged that two types of configurations would be installed:

- *Performance*-oriented **Disk Systems**;
- *Capacity*-oriented **Disk Systems**.

Vendors shall provide pricing information for the acquisition and support of several identical performance-oriented **Disk Systems**, as described in section 3.1.1, and one or more capacity-oriented **Disk Systems,** as described in section 3.1.2.

Tenderers are strongly encouraged to refer to the definition of "Usable Capacity" found in section 0 of this document.

3.1.1 Performance-oriented Disk Systems.

Typical usage for disks of this type is as a cache between a supercomputer environment and a tape library. The disk content retention time will be fairly short, with most of the data being overwritten every three to five days. The I/O patterns are a mixture of long sequential I/Os, where gigabytes or tens of gigabytes are read and written at a time, and short random I/Os. Experience has shown that data stored in these **Disk Systems** is very unlikely to benefit from storage tiering, with no particular datasets being accessed more frequently than others.

M 25. Tenderers shall provide a quote for the acquisition and support of two or more identical performance-oriented **Disk Systems**, providing a combined **Usable Capacity** of at least 500TB. Pricing, in accordance with **M 7**, must be inclusive of all costs related to requirements found in sections 2, 3.1.1 and 3.2.

R 26. Tenderers shall state what the **Usable Capacity** of each performance-oriented **Disk System** is.

H 27. It is highly desirable that the number of **Disk Systems** required to provide these 500TB of usable disk capacity is kept to a minimum.

M 28. Depending on the amount of usable disk capacity that it provides, each tendered performance-oriented **Disk System** must be able to sustain the IOPS and provide the aggregate bandwidth defined in Table 2, with a read:write ratio being approximately 2:1, for periods of at least five minutes. *For example, if two **Disk Systems** are tendered, each having a **Usable Capacity** of 250TB, each of these must be able to deliver 4GiB/s of bandwidth and 20,000 IOPS*. If multiple tiers of storage are proposed by the Tenderers, it must be assumed that data read from the **Disk System** resides at the lowest, less performing tier of storage.

Usable Capacity/ Disk System	Minimum Bandwidth/ Disk System	Minimum IOPS/ Disk System
100 TB or less	2.1 GiB/s	10,500
100 TB to 150 TB	2.7 GiB/s	13,500
150 TB to 200 TB	3.3 GiB/s	16,750
200 TB to 250 TB	4.0 GiB/s	20,000
More than 250TB	4.7 GiB/s	23,250

Table 2: Required bandwidth and IOPS for Performance-oriented Disk Systems

R 29. Tenderer shall state how many IOPS and how much bandwidth each performance-oriented **Disk System** can deliver over a five minutes period, with a read:write ratio being approximately 2:1, and assuming that all data read is at the lowest, less performing tier of storage.

M 30. Using no more than three streams of data transfer, it must be possible to achieve sustained transfer rates to/from a single LUN at speeds of 500MB/s during long sequential I/Os, over a 60 second period.

3.1.2 Capacity-oriented Disk Systems.

Typical usage for **Disk Systems** of this type is to be a long term cache for data accessed regularly in the archive, hence the **Disk System** can be characterised as read-mostly. Typical I/O pattern sees a slow filling of the disk space, followed by retrievals of small amount of data.

M 31. Tenderers shall provide a quote for the acquisition and support of one or more identical Capacity-oriented **Disk Systems**, providing a combined **Usable Capacity** of at least 400TB. Pricing, in accordance with **M 7**, must be inclusive of all costs related to requirements found in sections 2, 3.1.2 and 3.2.

R 32. Tenderers shall state what the **Usable Capacity** of each Capacity-oriented **Disk System** is.

H 33. It is highly desirable that the number of **Disk Systems** required to provide these 400TB of usable disk capacity is kept to a minimum.

M 34. Depending on the amount of usable disk capacity that it provides, each tendered capacity-oriented **Disk System** must be able to sustain the IOPS and provide the aggregate bandwidth defined in Table 3, with a read:write ratio being approximately 10:1, for periods of at least five minutes. *For example, if two capacity-oriented **Disk System**s are tendered, each having a **Usable Capacity** of 200TB, each*

of these must be able to deliver 1GiB/s of bandwidth and 5,000 IOPS. If multiple tiers of storage are proposed by the Tenderers, it must be assumed that data read from the **Disk System** resides at the lowest, less performing tier of storage.

Usable Capacity/ Disk System	**Minimum Bandwidth/ Disk System**	**Minimum IOPS/ Disk System**
100 TB or less	0.4 GiB/s	2,160
100 TB to 200 TB	0.6 GiB/s	3,000
200 TB to 400 TB	1.0 GiB/s	5,000
More than 400TB	1.2 GiB/s	6,000

Table 3: Required bandwidth and IOPS for Capacity-oriented Disk Systems

R 35. Tenderer shall state how many IOPS and how much bandwidth each capacity-oriented **Disk System** can deliver over a 5 minutes period, with a read:write ratio being approximately 10:1, and assuming that all data read is at the lowest, less performing tier of storage.

3.2. All Disk Systems.

H 36. In order to simplify the support and configuration of the **Disk Systems**, ECMWF considers as highly desirable that all **Disk Systems** tendered have a commonality of management interfaces.

Tenderers must provide answers to requirements of the whole section 3.2 and its subsections for each type of Disk System tendered.

R 37. Tenderers shall describe when the type of tendered **Disk Systems** first became generally available, and whether an end of life date for said equipment has already been defined.

3.2.1 Costing and configurations.

R 38. Vendors shall provide information allowing ECMWF to calculate the cost of other **Disk Systems** of identical type.

R 39. As already mentioned, it is not clear that usage of a **Disk System** internal tiering would be beneficial to ECMWF. However, Tenderers should describe any available facilities within the **Disk System** to support tiering of data.

R 40. Tenderers shall list the type of **Disk** that can be installed in the tendered **Disk System**. They will explain to what extent devices of multiple types can be mixed in a **Disk System** (*e.g.: a configuration can be composed of five trays of 10K RPM SAS disks, one tray of 15KRPM SAS disks,*

one tray being a mix of SSD and 15KRPM SAS disks). Tenderers shall also describe how an administrator can ensure that specific data sets are exclusively stored on disks of a specific type.

R 41. If a tendered **Disk System** is not fully populated (*e.g. more disk containers such as trays can be added, additional disks can be installed in existing containers, additional memory can be added on the controllers, etc.*), the Tenderers will provide pricing information allowing ECMWF to evaluate the cost uplifts linked to upgrading the configuration.

R 42. Tenderers shall describe from where the **disks** used in their **Disk System** are sourced (*e.g. Western Digital, Seagate, ...*) and what grade of devices are used (*e.g. enterprise, commodity, …*).

R 43. Tenderers shall provide references to the relevant information, if published, for benchmark results published on the Storage Performance Council's website for the tendered equipment.

3.2.2 Storage allocation

M 44. It must be possible to divide the disk storage found in each **Disk System** into multiple logical disk volumes (or LUN devices, LUNs), the size of which are defined by the administrator.

R 45. Tenderers will describe any limits imposed by the **Disk System** on the size and numbers of the LUNs, in particular minimal and maximal sizes of each LUN, and maximum number of LUNs that can be configured.

R 46. Tenderers will describe recommendations regarding the LUN sizing and numbers, which would affect the performance of the **Disk System**.

3.2.3 Host definitions and storage mapping

M 47. It must be possible to define a minimum of 30 hosts (servers) on each **Disk System**, able to access LUNs concurrently. The Tenderer shall describe how these hosts are defined on the **Disk System** (*e.g. a host could be defined as a collection of FC or iSCSI ports, whose addresses are selected or defined by the storage administrator*).

M 48. It must be possible to connect and use servers running RedHat Enterprise Linux (RHEL) 6.4 to each **Disk System.**

R 49. Tenderers shall describe the type of hosts and operating systems versions that can be successfully connected to each **Disk System**, via Fibre Channel, iSCSI or SAS interfaces.

R 50. Tenderers will describe any limits imposed by the **Disk System** on the number and type of hosts that can be connected to it.

M 51. It must be possible to connect (map) multiple LUNs provided by a **Disk System** to a single host.

M 52. It must be possible to transfer mapping of LUNs between hosts while preserving the LUN contents.

H 53. It is highly desirable for LUNs connected to RHEL hosts to be manageable through the use of the multipathd driver. Should the tendered **Disk Systems** require other drivers, these shall be described

by the vendor.

D 54. It is desirable to be able to map individual LUNs to multiple hosts concurrently (shared LUNs).

3.2.4 High Availability

M 55. Each tendered **Disk System** must be "Highly Available". ECMWF defines Highly Available to mean that it has no single point of failure, such that if any component of a **Disk System** fails, either through hardware or software malfunction, or a part of the **Disk System** fails, the services being provided by the **Disk System** must continue automatically, without manual intervention.

M 56. Each tendered **Disk System** must provide service availability, including scheduled maintenance, of at least 99.99% over the period of a year.

R 57. Tenderers shall describe how the tendered **Disk Systems** will be Highly Available. This should include a description of the hardware and software specifically included to ensure that the **Disk Systems** will be Highly Available as specified in **M 55** above and also:

- how a single disk failure is handled;
- how a second disk failure is handled, and under which circumstances this may lead to loss of data;
- how further disk failures are handled, and under which circumstances this may lead to loss of data;
- how controller failure is handled;
- how a disk container (*e.g. tray*) failure is handled;
- how internal connectivity failure is handled;
- how connectivity failures (SAN or Ethernet) are handled;
- how new versions of software and firmware are installed on the **Disk System** components, without impacting access from the attached Hosts to the tendered **Disk System**.

Where hardware and software features described are optional and would be available to ECMWF at extra cost, this needs to be clearly indicated.

R 58. Tenderers shall describe how much spare capacity is included in the **Disk System** being proposed, solely for the purpose of protecting it against disk failure.

H 59. It is highly desirable that all components of each **Disk System** are redundant and hot-pluggable. Tenderers shall describe components that do not satisfy this requirement.

H 60. It is highly desirable that support engineers are able to carry out diagnostic and repair work on any **Disk System** without interruption of service. Tenderers shall describe situations where it is not possible to perform this work without interruption of service.

H 61. It is highly desirable that the **Disk Systems** have features such that data availability and integrity is not compromised by multiple concurrent **Disk** failures.

R 62. Tenderers shall describe whether protection is offered by RAID levels or other forms of protection,

and which forms of array protection are offered by the **Disk System**. Tenderers will describe the conditions where a disk failure results in loss of data.

R 63. Tenderers shall describe how the performance of the **Disk System** is affected by a disk failure, and also any performance impact of a rebuild operation.

R 64. Tenderers shall describe whether the tendered **Disk System** implements support for the T10-Data Integrity Extension standard.

3.2.5 Interfaces

M 65. Each **Disk System** must include at least 4 FC interfaces, each of a nominal bandwidth of 8Gib/s or above.

H 66. It is highly desirable that the FC adapters on each **Disk System** are able to work at speeds of 16 Gib/s.

R 67. Tenderers will describe data connectivity options available on the tendered **Disk Systems**.

M 68. Each **Disk System** must be able to handle multipath I/Os, where connections between the **Disk System** and a host are made through two or more totally independent physical routes or paths.

D 69. It is desirable that all paths can be operated in active/active mode.

R 70. Tenderers shall describe how the I/O load is balanced across the various paths available between a host and each **Disk System**, and whether these paths are operating in active/passive or active/active mode.

- When the host is connected to the **Disk System** through two FC ports;
- When the host is connected to the **Disk System** through four FC ports.

R 71. Tenderers shall describe:

- How the performance of the **Disk System** is affected by the loss of connectivity on one of the FC paths used to connect the subsystems to hosts.
- The impact on hosts of the loss of connectivity on one of the FC paths used to connect the subsystems to hosts, besides performance. Of particular interest is whether this loss would be transparent to an application or not.

D 72. Each **Disk System** should include one or more Ethernet interfaces to allow connection with the ECMWF's management network for control and out of band management purposes.

3.2.6 System installation, operational characteristics and environmentals

R 73. Tenderers shall provide details of site preparation, dimensions, weight, and access requirements for the tendered **Disk Systems**.

ECMWF's Computing Environment includes a large scale Uninterruptible Power Supply (UPS) which protects the power supply to the equipment in ECMWF's computer halls. The provision of separate external UPS protection for the tendered **Disk Systems** is not required.

R 74. Tenderers shall provide details of power and environmental requirements for the tendered **Disk Systems**. The equipment might be supplied from unconditioned mains power. If equipment is supplied from unconditioned mains, Tenderers shall state whether any consequential impact on equipment availability or reliability other than unavailability during power outage is to be expected. Tenderers shall provide details of any special actions that have to be taken following a power outage to the **Disk Systems** in order to restore the service. The time required to restore the service should be stated and whether this can be achieved automatically or would require manual intervention.

R 75. Tenderers shall indicate how many kWh their **Disk System** consumes when in operation.

M 76. In no circumstances must a partial or complete loss of power of up to 12 hours result in the loss of data that has been acknowledged as successfully written to the **Disk System.** Tenderers shall describe how their equipment is protected against such failure. Of particular interest are the methods used to protect data just sent to the **Disk System** and which may only be residing in volatile storage.

D 77. It is desirable that the tendered **Disk Systems** are rack mountable in industry standard 19" racks.

R 78. Tenderers shall describe whether the tendered **Disk Systems** include racking, or whether it is expected that ECMWF will provide the racking required to install the **Disk Systems**. In the latter case, Tenderers shall describe the requirements for these racks, such as number of Rack Units (RU) that will be occupied by the tendered systems and number and type of power sockets required.

3.2.7 Administrative interfaces.

3.2.7.1 General

H 79. It is highly desirable that all administrative functions (with possibly the exception of performance monitoring) are made available through a command line interface, accessible from Linux servers connected through Ethernet.

R 80. The Tenderer shall describe the interfaces available to manage and monitor each tendered **Disk System**.

3.2.7.2 System monitoring and control

M 81. The tendered **Disk Systems** should be designed such that, in normal circumstances, no manual intervention by operations staff is required to keep it operational at a high level of efficiency.

H 82. It is highly desirable that the tendered **Disk Systems** should maintain a comprehensive set of log files. Items to be logged should include the following type of events. Tenderers are asked to describe whether these and/or other events are logged:

- Hardware malfunctions;

- System exception conditions;
- All events that could have a bearing on system security;
- Use of secure accounts (root, system administrator, etc.);
- Configuration changes;
- Any events related to failovers and other actions performed by the **Disk System**.

D 83. The tendered **Disk Systems** should allow the definitions of multiple user accounts, and allow these accounts to have different levels of administrative privilege. *For example, some accounts should be able to use the full range of commands, including the creation and allocation of disk resources and hosts while some other accounts should only be able to monitor the system and its performance.*

M 84. The tendered **Disk Systems** must notify the system operator when any component has failed.

D 85. It is desirable that tools are available, that allow the identification of potential hotspots. *For example, such tools should identify physical components that are particularly busy, LUNs that are being heavily accessed, hosts creating significant loads.* Such information should be available both in real time *(what is happening now?)* and historically *(this LUN has started being heavily used since ...).* Tenderers shall describe what tools are available and the scope of such tools and how this information can be exported for analysis.

R 86. The Tenderer shall describe all tools available that enable the **Disk Systems** to be monitored and controlled (including remote control of system power modules), indicating which of these tools will be provided within the **Disk Systems** as tendered and which are optionally available and the additional costs, if any, of such tools.

D 87. ECMWF has a monitoring system which is used to monitor the health of all major systems in operation using SNMP. It is desirable that any tendered solution can be integrated with this monitoring system.

D 88. The tendered **Disk System** should support SNMP traps and should be configurable so that the failure of any System component causes a trap.

R 89. Tenderers shall describe any extensions to SNMP to allow monitoring and/or management of non-standard parameters.

R 90. Tenderers shall describe the methods for up/downloading **Disk System** configuration and software.

H 91. It is highly desirable that the tendered Disk Systems provide tools allowing backing up the Disk System configuration (e.g. LUN layouts, Host mappings, …). Tenderers will describe how this operation is made, and how this backup can be restored onto the **Disk System**.

R 92. Tenderers shall describe the management activities, which cannot be achieved remotely, such as those that must be done via a terminal connected to the console serial port or by physical manipulation.

D 93. The tendered **Disk System** should log events using the Syslog protocol according to RFC 3164.

Tenderers should provide an overview over the events that will be logged.

D 94. It should be possible for the tendered **Disk System** to synchronise its internal clocks from a set of designated NTP servers.

R 95. Tenderers shall describe whether the tendered equipment can be easily integrated into an Opsview/Nagios monitoring suite.

3.2.8 System software

R 96. Tenderers shall describe how and how often new releases of system level software, including firmware, would be made available to ECMWF.

M 97. Whenever new versions of system level software or firmware are released, Tenderers must guarantee to support the software or firmware release being superseded for at least a further 12 months.

H 98. It is highly desirable that the tendered **Disk System**'s operating software, including firmware, can be upgraded without interruption to the service. Tenderers must describe how this can be achieved in an operational environment.

R 99. Tenderers shall describe whether their proposed **Disk Systems** are compatible with virtualisation solutions, and if so in what configuration or usage scenarios.

R 100. Tenderers should describe any certification or support for particular filesystems.

3.3. Disk Systems Acceptance

Equipment installed by the selected **Disk System** tenderer will be subject to the acceptance tests defined in section 6 of Volume III. These tests will consist of a 24 hours period during which it will be verified that the functionalities - including those related to High Availability of the **Disk Systems** - and performance claimed in the Tender can be observed. This will be followed by a 30 days reliability test, to ensure that the equipment performs satisfactorily during that period.

3.4. Future acquisitions

While the **Disk Systems** equipment for the initial acquisition is limited to that described above, ECMWF is likely to acquire other equipment in the future, as part of the call-off contract described in Volume III.
In this context, ECMWF requests the Tenderer to describe the range of products that it would be able to provide.

R 101. Tenderers shall provide a description of the range of **Disk Systems** and SAN switching equipment and FC cabling that they would be able to provide to ECMWF in 2015. This should include information related to FC, iSCSI and SAS attached products, if available. Note that future **Disk Systems** acquisitions may be subject to sets of requirements different from those defined in sections 3.1 and 3.2, possibly less stringent.

R 102. Tenderers shall provide a description of the range of **Disk Systems** and SAN switching equipment and FC cabling that they would be able to provide to ECMWF over the life time of the call-off contract. This should include information related to FC, iSCSI and SAS attached products, if available.

R 103. For the range of products described in response to **R 101** and **R 102**, Tenderers shall describe what maintenance options will be available during the call-off contract.

R 104. Tenderers will provide any information regarding the sustainability of their product range (e.g. energy efficiency, recyclability, environmental impact on disposal, etc.).

M 105. Tenderers must agree to provide to ECMWF on an annual basis or more frequently, under non-disclosure if required, product roadmap updates on the **Disk Systems** and SAN switching equipment that they are able to provide.

4. Servers Technical Requirements

4.1. Information Provision to ECMWF and Product Availability

M 106. Tenderers must agree to provide to ECMWF on an annual basis or more frequently, under non-disclosure if required, product roadmap updates. Tenderers must also agree to inform ECMWF when a model of **Server** purchased by ECMWF is about to reach its end of life in sufficient time to allow additional systems to be purchased if ECMWF so wishes.

H 107. In order to minimise the costs of supporting **Servers**, ECMWF wishes to minimise the number of differing types of systems that it purchases, but still spread purchases over the period of the call-off contract. It is highly desirable that Tenderers can normally offer, for a period of at least 12 months from the date of General Availability, models of **Servers** that can use the same system image.

M 108. Tenderers must agree to deliver, if requested by ECMWF, a demonstration model of the offered **Servers** on free loan, technically as close as possible in all aspects to the equipment tendered for the **Initial Acquisition**. The demonstration equipment will be kept by ECMWF for a maximum of 3 months from the delivery date, during which time ECMWF will be responsible for its safe keeping.

4.2. Specific Requirements for Servers

M 109. Tenderers must provide a quote for an **Initial Acquisition** of 20 **Servers**, all of the same model, which comply with the requirements below. The quote must include the technical specifications of the model of server proposed.

Feature	**Requirement**
Processor architecture:	Processor capable of running Microsoft Windows Server 2012 64-Bit Edition and Red Hat Enterprise Linux 6.4 64-bit.
Minimum system performance:	The system must have: • SPECint_rate_base2006: no less than 208; • SPECfp_rate_base2006: no less than 184.
Memory:	64 GiB, upgradeable to at least 512 GiB.
Disk Drives:	At least 2 hot swappable SAS disks, each with a capacity of at least 500 GB.
Disk Bays:	At least 2 unoccupied drive bays with connectors to allow additional disks to be installed in these bays.

RAID for internal disks:	Support for RAID 1.
Network Interfaces:	Minimum of 2 on board 1 Gib/s BaseT Ethernet interfaces capable of Wake-on-LAN and PXE booting.
Power Supply and Management:	Fitted with hot swappable Dual Power Supplies with ACPI compliant Power Management.
USB Ports:	2 USB2 ports.
Form factor and size:	Rack mountable in industry standard 19" rack, 1 RU or 2 RU.
PCI Express slots:	At least 2 PCI Express v3 slots.
10 gigabit Ethernet:	Compatible with Intel Ethernet Server adapter PCI cards (82599ES chipset).
Fibre Channel:	Fitted with a dual port Emulex Fibre Channel 16Gib/s adapter (HBA).
Video:	VGA port.
Out-of-band Management:	Capable of providing remote access to console and remote power control capabilities (i.e. power on/off/reset.)

Table 4: Requirements for Servers

R 110. Tenderers are asked to provide any details of Linux certification and support for the tendered **Servers**, in particular which distributions of Linux have been certified.

M 111. Each **Server** needs to provide enough internal bandwidth that it is possible to transfer data between storage devices connected to the FC card and a dual port 10Gib Intel Ethernet Server adapter PCI, at rates of 12Gib/s (devices to network), or 8Gib/s (network to devices), or 14Gib/s (bidirectional). It can be assumed that the devices and network are not bottlenecks.

R 112. Tenderers are asked to provide any details regarding monitoring capabilities of the tendered **Servers**, in particular support for SNMP and Self-Monitoring, Analysis and Reporting Technology (SMART) for internal disks.

R 113. Tenderers shall provide details of power and environmental requirements for the tendered **Servers**. Tenderers will indicate how many kWh each **Server** consumes under normal operation.

4.3. Servers Acceptance

Acceptance conditions are described in section 6 of Volume III.

4.4. Future acquisitions

While the **Server** equipment to be tendered in response to this ITT is limited to the one described in Section 4.2 above, ECMWF is likely to acquire other equipment in the future, as part of the call-off contract described in Volume III. Note that future **Server** acquisitions may be subject to different sets of requirements, possibly less stringent.

H 114. In this context, ECMWF expects to purchase servers to meet very diverse needs over the life of the contract. Consequently, it needs the Tenderer to be able to supply a range of **Server** products. For manageability and to enable hosts to function in a cluster environment all servers should be from the same original equipment manufacturer.

M 115. For future purchases of **Servers** under the call-off contract, Tenderers must agree to provide on request specimen systems on loan so that ECMWF can establish whether the specifications of the models being considered for purchase are suitable for use in the ECMWF environment. ECMWF would keep these systems for a maximum of 6 weeks per system.

R 116. Tenderers will describe the range of **Servers** that it would be able to provide at the closing date of the ITT in terms of:

Feature	**Request for Information**
Processor architecture:	Processor range available.
Processor sockets:	Minimum and Maximum number of processor sockets.
Memory:	Maximum installable memory.
Disk Drives:	Maximum internal disk capacity installable and supported disk types.
Disk Bays:	Maximum number of disk bays.
RAID for internal disks:	Supported RAID Levels.
Network Interfaces:	Maximum number of on board Ethernet interfaces and capability for Wake-on-LAN and PXE booting.
Form factor and size:	Rack unit and width of equipment.

Internal PCIe slots:	Number and type of PCIe slots.

Table 5: Request for Information for Servers

R 117. Tenderers will provide any information regarding support for virtualisation within their product range.

R 118. Tenderers will provide any information regarding the sustainability of their product range (e.g. energy efficiency, recyclability, environmental impact on disposal, etc.).

Annex A: Costing and critical information spreadsheet.

NB: Tenderers can download an excel spreadsheet including the following table from: http://www.ecmwf.int/en/about/suppliers/itt-ecmwf/2014/216-acquisition-disk-systems-and-servers.

ITT 216 costing and critical information spreadsheet			
Overall Bid.	**related reqs.**	**Tenderer Answer**	**Unit to be used.**
Tenderer is bidding for **Disk Systems**	M 1		Yes/No
Tenderer is bidding for **Servers**	M 1		Yes/No
Total Price	M 9		£
Design and consultancy charges	M 9		£
Project Management charges	M 9		£
Installation charges	M 9		£
Hardware & software maintenance & support	M 9		£
Training programme	M 9		£
Disk Storage			
Skip this section if not tendering for Disk Systems			
Performance-oriented Disk Systems			
Number of **Disk Systems** bid to reach a minimum of 5OOTB Usable capacity.	M 25		
Usable capacity per **Disk System**	R 26		TB
Sustained IOPS achievable per **Disk System**	R 29		IOPS
Sustained IO bandwidth per **Disk System**	R 29		GiB/s
Quoted price for **one** of the Bid Performance-oriented **Disk System** satisfying all mandatory requirements, including import duties, shipping, delivery, site preparation, installation, support and maintenance, software and firmware licences, licence keys, covering all costs expected over a period of five years.	M 7		£

Quoted price for optional costs required to satisfy all desirable or highly desirable requirements for **one** of the tendered Performance-oriented **Disk Systems**, including import duties, shipping, delivery, site preparation, installation, support and maintenance, software and firmware licences, licence keys, covering all costs expected over a period of five years.	M 8		£
Price for one additional Performance-oriented **Disk System** satisfying all mandatory requirements, including import duties, shipping, delivery, site preparation, installation, support and maintenance, software and firmware licences, licence keys, covering all costs expected over a period of five years.	R 38		£
Price for optional costs required to satisfy all desirable or highly desirable requirements for one additional of the bid Performance-oriented **Disk Systems**, including import duties, shipping, delivery, site preparation, installation, support and maintenance, software and firmware licences, licence keys, covering all costs expected over a period of five years.	R 38		£
Power usage when in operation	R 75		kWh
Capacity-oriented Disk Systems			
Number of **Disk Systems** bid to reach a minimum of 400TB	M 31		
Usable capacity per **Disk System**	R 32		TB
Sustained IOPS achievable per **Disk System**	R 35		IOPS
Sustained IO bandwidth per **Disk System**	R 35		GiB/s
Quoted price for **one** Capacity-oriented **Disk System** satisfying all mandatory requirements, including import duties, shipping, delivery, site preparation, installation, support and maintenance, software and firmware licences, licence keys, covering all costs expected over a period of five years.	M 7		£
Quoted price for optional costs required to satisfy all desirable or highly desirable requirements for **one** Capacity-oriented **Disk System**, including import duties, shipping, delivery, site preparation, installation, support and maintenance, software and firmware licences, licence keys, covering all costs expected over a period of five years.	M 8		£
Price for one additional Capacity-oriented **Disk System** satisfying all mandatory requirements, including import duties, shipping, delivery, site preparation, installation, support and maintenance, software and firmware licences, licence keys, covering all costs expected over a period of five years.	R 38		£

Price for optional costs required to satisfy all desirable or highly desirable requirements for one additional of the bid Capacity-oriented **Disk Systems**, including import duties, shipping, delivery, site preparation, installation, support and maintenance, software and firmware licences, licence keys, covering all costs expected over a period of five years.	R 38		£
Power usage when in operation	R 75		kWh
Future acquisitions			
Guaranteed percentage level of discount on the list price for future acquisition.	M 11		%
Guaranteed percentage level of discount on the list price for maintenance of equipment acquired in the future	M 11		%
Servers			
Skip this section if not tendering for Servers			
Quoted Price for 20 servers satisfying all mandatory requirements, including import duties, shipping, delivery, site preparation, installation, support and maintenance, software and firmware licences, licence keys, covering all costs expected over a period of five years.	M 7		£
Power usage when in operation	R 113		kWh
Future acquisitions			
Guaranteed percentage level of discount on the list price for future acquisition.	M 11		%
Guaranteed percentage level of discount on the list price for maintenance of equipment acquired in the future	M 11		%