

RETURN BIDS TO:
RETOURNER LES SOUMISSIONS À:

Bid Receiving - PWGSC / Réception des
soumissions - TPSGC

11 Laurier St. / 11, rue Laurier

Place du Portage, Phase III

Core 0A1 / Noyau 0A1

Gatineau, Québec K1A 0S5

Bid Fax: (819) 997-9776

LETTER OF INTEREST
LETTRE D'INTÉRÊT

Comments - Commentaires

Vendor/Firm Name and Address

Raison sociale et adresse du
fournisseur/de l'entrepreneur

Issuing Office - Bureau de distribution

Acquisition Branch, STAMS, ITSPD / Direction générale
des acquisitions, SGAST, DASIT
Computer Hardware Division
Div. de l'équipement informatique
Place du Portage, Phase III, 4C2
11 Laurier Street/11, rue Laurier
Gatineau
Québec
K1A 0S5

Title - Sujet RFI ANNEX A: TECH SPEC SAN/NAS	
Solicitation No. - N° de l'invitation EN740-12RFII/A	Date 2012-05-15
Client Reference No. - N° de référence du client EN740-12RFII	GETS Ref. No. - N° de réf. de SEAG PW-\$\$EJ-461-24409
File No. - N° de dossier 461ej.EN740-12RFII	CCC No./N° CCC - FMS No./N° VME
Solicitation Closes - L'invitation prend fin at - à 02:00 PM on - le 2012-06-01	
Time Zone Fuseau horaire Eastern Daylight Saving Time EDT	
F.O.B. - F.A.B. Plant-Usine: <input type="checkbox"/> Destination: <input checked="" type="checkbox"/> Other-Autre: <input type="checkbox"/>	
Address Enquiries to: - Adresser toutes questions à: O'Sullivan, Patrick	Buyer Id - Id de l'acheteur 461ej
Telephone No. - N° de téléphone (819) 956-8340 ()	FAX No. - N° de FAX (819) 956-1156
Destination - of Goods, Services, and Construction: Destination - des biens, services et construction: DEPARTMENT OF PUBLIC WORKS AND GOVERNMENT SERVICES CANADA PORTAGE III 0A1 11 LAURIER ST Gatineau Quebec K1A0S5 Canada	

Instructions: See Herein

Instructions: Voir aux présentes

Delivery Required - Livraison exigée See Herein	Delivery Offered - Livraison proposée
Vendor/Firm Name and Address Raison sociale et adresse du fournisseur/de l'entrepreneur	
Telephone No. - N° de téléphone Facsimile No. - N° de télécopieur	
Name and title of person authorized to sign on behalf of Vendor/Firm (type or print) Nom et titre de la personne autorisée à signer au nom du fournisseur/ de l'entrepreneur (taper ou écrire en caractères d'imprimerie)	
Signature	Date

REQUEST FOR INFORMATION REGARDING NMSO MASS STORAGE SYSTEMS

DRAFT ANNEX A: TECHNICAL SPECIFICATIONS - SAN / NAS

TABLE OF CONTENTS

A. 1 Background and Purpose of this Request for Information (RFI)	2
A. 2 Nature of Request for Information	2
A. 3 Nature and Format of Responses Requested	2
A. 4 Response Costs	2
A. 5 Treatment of Responses	2
A. 6 Contents of this RFI	3
A. 7 Format of Responses	4
A. 8 Enquiries	4
A. 9 Submission of Responses	4

Annex A: Draft Technical Specifications SAN / NAS

REQUEST FOR INFORMATION REGARDING

NMSO MASS STORAGE SYSTEMS

DRAFT ANNEX A: TECHNICAL SPECIFICATIONS - SAN / NAS

A. 1 Background and Purpose of this Request for Information (RFI)

Canada intends to add new Categories to the Mass Storage Systems National Master Standing Offer (NMSO E60EJ-11000S/XXX/EJ) by June 2, 2013.

The purpose of this RFI is to seek feedback and determine if the content of Draft Annex A is representative of current market technologies. It is further intended as a means for Respondents to highlight where the Draft Technical Specifications could potentially be improved to facilitate delivery of solutions to Client Departments.

A. 2 Nature of Request for Information

This is not a bid solicitation. This RFI will not result in the award of any contract. As a result, potential suppliers of any goods or services described in this RFI should not reserve stock or facilities, nor allocate resources, as a result of any information contained in this RFI. Nor will this RFI result in the creation of any source list. Therefore, whether or not any potential supplier responds to this RFI will not preclude that supplier from participating in any future procurement. Also, the procurement of any of the goods and services described in this RFI will not necessarily follow this RFI. This RFI is simply intended to solicit feedback from industry with respect to the matters described in this RFI.

A. 3 Nature and Format of Responses Requested

Respondents are requested to provide their comments, concerns and, where applicable, alternative recommendations regarding how the requirements or objectives described in this RFI could be satisfied. Respondents are also invited to provide comments regarding the content, format and/or organization of any draft documents included in this RFI. Respondents should explain any assumptions they make in their responses.

A. 4 Response Costs

Canada will not reimburse any respondent for expenses incurred in responding to this RFI.

A. 5 Treatment of Responses

- (a) **Use of Responses:** Responses will not be formally evaluated. However, the responses received may be used by Canada to develop or modify procurement strategies or any draft documents contained in this RFI. Canada will review all responses received by the RFI closing date. Canada may, in its discretion, review responses received after the RFI closing date.
- (b) **Review Team:** A review team composed of representatives of the client (where applicable) and PWGSC will review the responses. Canada reserves the right to hire any independent consultant, or use any Government resources that it considers necessary to review any response. Not all members of the review team will necessarily review all responses.

- (c) **Confidentiality:** Respondents should mark any portions of their response that they consider proprietary or confidential. Canada will handle the responses in accordance with the *Access to Information Act*.
- (d) **Follow-up Activity:** Canada may, in its discretion, contact any respondents to follow up with additional questions or for clarification of any aspect of a response. In addition, Canada will provide all Respondents and interested industry associations with a revised draft requirement, so that interested suppliers are able to determine how industry feedback has been incorporated. Respondents are asked to indicate the language of preference for a subsequent draft in their responses.

A. 6 Contents of this RFI

- (a) This RFI contains draft Annex A: Technical Specifications - SAN / NAS. This document remains a work in progress and respondents should not assume that new clauses or requirements will not be added to any bid solicitation that is ultimately published by Canada. Nor should respondents assume that none of the clauses or requirements will be deleted or revised. Comments regarding any aspect of the draft document are welcome.
- (b) This RFI also contains the following questions addressed to the industry:
- (i) PWGSC is exploring the feasibility of publishing qualified products on the Computer Acquisition Guide (CAG) (<http://computer.pwgsc.gc.ca>) using the Microcomputer configurator model, instead of in the spreadsheet format currently used for Mass Storage. An example of how the configurator works can be seen by selecting *Servers category B1.0 and B2.1* at:
<http://computer.pwgsc.gc.ca/index-eng.cfm?af=ZnVzZWFiZjdlbWljcm8uYm9vZHVjdHNfc2VydMvYmNtZHR5PTMmUmVxdWVzdFRpbWVvdXQ9OTAwJmxhbmc9ZW5n>
 An example of the current spreadsheet format can be seen at:
http://computer.pwgsc.gc.ca/index-eng.cfm?fuseaction=capacite_storage.archive.category&lang=eng&f=gr05cat02. Note that we are required to display the spreadsheets in alternate formats in order to comply with Common Look and Feel Accessibility policies.
 It is anticipated that displaying qualified storage products in the configurator format will provide the procurement group with greater control over web publishing, reducing lead times for suppliers submitting substitutions. However, the configurator would require that a lower level of detail be published than what currently exists for SAN within the spreadsheet format, please see Servers link above.
- (ii) PWGSC is considering the longer-term feasibility of incorporating a ratings analysis similar to that currently used for the server NMSO categories, where Offerors could obtain points for non-mandatory features that add value to their product offering. Features would be validated by an independent third party (at a small cost to the Offeror) and would be incorporated into the ranking of qualified products and associated Call-up limits, as a best value (rather than simply lowest price) calculation. A current ratings analysis example can be seen at:
<http://canada.nstl.com/SipssWebSites/Report2008/s2012/ServerFeatures.asp?catg=410&SysID=41030>
- (iii) PWGSC is considering structuring categories of products offered (e.g. high-end enterprise storage) into 6 sections, with top ranked systems in the category having maximum Call-up limits set as follows: Storage Platform - \$400K, Disk -

\$400K; Fabric- \$200K; Virtualization - \$200K; NAS Gateway - \$200K; Other Related Options (financially unevaluated) - \$ TBD.

Feedback is requested on these proposals. Respondents may submit additional comments on other aspects of the NMSO if desired.

A. 7 Format of Responses

- (a) **Cover Page:** If the response includes multiple volumes, respondents are requested to indicate on the front cover page of each volume the title of the response, the solicitation number, the volume number and the full legal name of the respondent.
- (b) **Title Page:** The first page of each volume of the response, after the cover page, should be the title page, which should contain:
 - (i) the title of the respondent's response and the volume number;
 - (ii) the name and address of the respondent;
 - (iii) the name, address and telephone number of the respondent's contact;
 - (iv) the date; and
 - (v) the RFI number.
- (c) **Numbering System:** Respondents are requested to prepare their response using a numbering system corresponding to the one in this RFI. All references to descriptive material, technical manuals and brochures included as part of the response should be referenced accordingly.
- (d) **Number of Copies:** Canada requests that respondents submit 2 copies of their responses.

A. 8 Enquiries

Because this is not a bid solicitation, Canada will not necessarily respond to enquiries in writing or by circulating answers to all potential suppliers. However, respondents with questions regarding this RFI may direct their enquiries to:

Contracting Authority: Patrick O'Sullivan

E-mail Address: patrick.o'sullivan@pwgsc-tpsgc.gc.ca

Telephone: (819) 956-8340

A. 9 Submission of Responses

- (a) **Time and Place for Submission of Responses:** Suppliers interested in providing a response should deliver it to the following location by the time and date indicated on page 1 of this document:

 Department of Public Works and Government Services Bid Receiving Unit
 Portage III, 0A1
 11 Laurier Street
 Gatineau, Quebec K1A 0S5

Responses should not be sent directly to the Contracting Authority.
- (b) **Responsibility for Timely Delivery:** Each respondent is solely responsible for ensuring its response is delivered on time to the correct location.
- (c) **Bid Receiving Unit Address Solely for Delivery of Responses:** The above address is only for bid submission. No other communications are to be forwarded to this address.

Solicitation No. - N° de l'invitation

EN740-12RFII/A

Amd. No. - N° de la modif.

Buyer ID - Id de l'acheteur

461ej

Client Ref. No. - N° de réf. du client

EN740-12RFII

File No. - N° du dossier

461ejEN740-12RFII

CCC No./N° CCC - FMS No./N° VME

-
- (d) **Identification of Response:** Each respondent should ensure that its name and return address, the solicitation number and the closing date appear legibly on the outside of the response.

ANNEX A: DRAFT TECHNICAL SPECIFICATIONS (SAN / NAS)

SAN 1.0 CATEGORY

The following describes the configuration and features of a **mid-range SAN** storage solution.

1.1 Storage Platform

1.1.1 Capacity and Platform

Each storage platform must meet the following capacity and platform requirements:

- (a) the hard disk drive technologies and densities must be commercially available, meaning that the Manufacturer is continuing to manufacture and ship them to customers generally;
- (b) the hard disk drive technologies and densities must be tested and fully supported within the storage platform by the storage platform Manufacturer;
- (c) it must include industry-standard hard disk drives operating at either 4Gbps for Fibre Channel (FC) drives or 6Gbps for Serial Attach SCSI (SAS) drives;
- (d) it must also include industry-standard Serial Advanced Technology Attachment (SATA) revision 3.0 or Nearline SAS (NL-SAS) hard disk drives operating at 6Gbps. This may be achieved either by:
 - i. using the same shelves as either the FC or SAS disk drives, or
 - ii. using specialized shelves for these drive types;
- (e) the available drive options must include at least seven (7) from the following list:
 - drives with 4Gbps (for FC) or 6 Gbps (for SAS) interfaces and 15000 RPM rotational speed:
 - i) 300GB
 - ii) 450GB
 - iii) 600GB
 - drives with 4Gbps (for FC) or 6 Gbps (for SAS) interfaces and 10000 RPM rotational speed:
 - iv) 300GB
 - v) 450GB
 - vi) 600GB
 - vii) 900GB
 - drives with 6Gbps for SATA or NL-SAS interfaces and 7200 RPM rotational speed:
 - viii) 1TB
 - ix) 2TB
 - x) 3TB
 - solid state drives (SSD) based on Single Level Cell (SLC) or enterprise-class Multi-Level Cell (eMLC) technology
 - xi) 100GB
 - xii) 200GB
 - xiii) 300GB
 - xiv) 400GB
 - xv) 600GB
- (f) it must accommodate a minimum of 224 hard disk drives

- (g) It must be packaged in a standard 19" rack mount form factor (NOTE: it is understood that standard rack depth will be increased when "high density" disk shelves are provided); and
- (h) It must include lights or an LCD panel for power, activity and fault indications.

1.1.2 Cooling

Each storage platform must meet the following cooling requirements:

- (a) it must provide sufficient cooling for a fully populated cabinet at the mandatory minimum storage capacity;
- (b) all cooling for the system controller(s) as well as all hard disk drives must be redundant and monitored for failure by the storage platform hardware;
- (c) it must allow hot swapping of failed cooling fans;
- (d) the cooling system within the storage platform itself must be fully redundant; and
- (e) in the event of a component failure, the cooling system must allow continued operation of the storage platform until service can be performed.

1.1.3 Drives and Shelves

Each storage platform must meet the following drives and shelves requirements:

- (a) the hard disk drives must include a minimum of either 4Gbps for Fibre Channel (FC) or 6Gbps for Serial Attach SCSI (SAS) with dual attached interfaces;
- (b) it must provide minimum of 4 active connections to the mandatory 224 hard disk drives. Bandwidth must be allocated evenly to the total number of physical drives over several channels;
- (c) it must provide fully redundant back-end paths to all hard disk drives. A channel failure must not interrupt access to attached disk drives;
- (d) it must allow hot addition of storage shelves without needing to power the storage platform down and without interrupting access to existing drives and redundant arrays of inexpensive disk (RAID) groups;
- (e) it must include as many back-end channels as necessary to support all the back-end shelves of disks so that a shelf component replacement or failure does not interrupt access to adjacent shelves in the platform.
- (f) the hard disk drives in the storage platform must be fully hot pluggable while the storage platform is operational. There must be no loss of data if a hard drive is removed, assuming the drive is part of a fault-tolerant configuration in the platform;
- (g) it must rebuild a replaced hard disk drive automatically and without user intervention when it is inserted, assuming it is replacing a hard disk drive that was part of a fault-tolerant configuration; and
- (h) it must allow hard disk drives to be allocated as global hot spares, which must automatically rebuild the contents of a failed hard disk drive in any fault-tolerant RAID set. This process must be fully automatic whenever a disk failure occurs in a fault-tolerant RAID set.

1.1.4 Power

Each storage platform must meet the following power requirements:

- (a) it must provide sufficient power to operate a fully populated system with all boards and cache installed, and the maximum number of hard disk drives installed;
- (b) the power supplies must be fully redundant, allowing uninterrupted operation of the storage platform in the event of a power supply failure, until service can be performed. Redundancy may be achieved either by using:
 - i. a second power supply, or
 - ii. an N+1 approach; and
- (c) each AC power supply must connect independently to a discrete AC power source.

1.1.5 Controllers

Each storage platform must meet the following controller requirements:

- (a) it must include dual redundant active/active storage controllers for handling both I/O to the attached host systems as well as disk I/O and RAID functionality;
- (b) it must be redundant, so that the surviving controller automatically recovers controller subsystem failures, and service to attached hosts is continued without disruption;
- (c) each storage controller within the storage platform must have access to all 224 of the mandatory hard disk drives in order to assign, configure, protect and share those drives;
- (d) the storage controllers must allow configuration of hard disk drives within the storage platform as:
 - i. RAID5 stripes with parity;
 - ii. RAID6 stripes with dual parity;
 - iii. RAID1; and
 - iv. RAID0+1 stripes with mirroring or RAID1+0 striped mirrors (aka RAID10).
- (e) it must allow the creation and addressing of up to 2048 simultaneous logical drives; and
- (f) it must simultaneously support all RAID types from 1.1.5 (d) within the storage platform.

1.1.6 Cache

Each storage platform must meet the following cache requirements:

- (a) it must include at least 8GB of dedicated I/O cache per storage controller;
- (b) the cache on the storage controller must perform both read and write I/O operations;
- (c) the write cache must be mirrored and must utilize error detection and correction logic that will detect and recover from any memory errors without data loss or interruption to service; and
- (d) the write data within the cache on the storage controllers must be protected by one of these two methods:
 - i. a battery that allows the cache contents to be held intact for a minimum of 48 hours. The caches must then complete their write operations to disk when power is restored; or
 - ii. all pending write data must be automatically written to disk before the disk system is powered off, and the platform must provide sufficient battery power to complete this function.

1.1.7 I/O Ports and Connectivity

Each storage platform must meet the following requirements for I/O ports and connectivity:

- (a) it must include a minimum of 2 storage controllers that may be replaced in the event of a controller failure;
- (b) it must provide a minimum of 4 fibre channel ports for connectivity to Intel and Open System host computers;
- (c) all 4 fibre channel ports must be independent ports operating at 8Gbps each and operate in either point-to-point or loop modes of operation;
- (d) each of the 4 fibre ports must support full fabric login and must have a unique fibre channel World Wide Name;
- (e) it must provide simultaneous connectivity to any combination of 128 or more Intel and UNIX hosts using dual fibre channel host bus adapters in each host;
- (f) it must provide host bus adapters including any required software drivers for all supported operating systems; and
- (g) it must offer "no single point of failure" connectivity options, for both failover as well as load balancing under all of the mandatory operating system environments. This may be provided using add-on failover software packages or using native Operating System facilities
- (h) it must provide two (2) native 10GbE connections for FCoE host connectivity that meet the ANSI T11 FC-BB-5 Fibre Channel Over Ethernet (FCoE) standards for the encapsulation of FC packets over Full Duplex and Lossless Ethernet networks and must be compliant with the following IEEE standards:

-
- i. 802.1Ae;
 - ii. 802.1Qbb;
 - iii. 802.1Qaz which defines:
 - a. enhanced transmission selection (ETS); and
 - b. data center bridging exchange (DCBX).

1.1.8 Hosts

Each storage platform must meet the following requirements for host connectivity:

- (a) it must connect to Intel-based host computers running the following:
 - i. Windows 2008 R2 with Hyper-V in 32-bit and 64-bit configurations;
 - ii. Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11 in 32-bit and 64-bit configurations; and
 - iii. VMWare ESX Server 3.5X and 4.X;
- (b) it must connect to the following UNIX and Open Systems hosts simultaneously, in addition to the previously listed Intel systems:
 - i. Oracle Solaris 10 and 11 systems;
 - ii. HP-UX 11i v.X systems;
 - iii. IBM AIX v6.X and v7.X systems;
- (c) Support of additional platform types and operating systems is desirable, but not mandatory.

1.1.9 Clustering

Each storage platform must meet the following requirements for clustering:

- (a) it must directly support clustering under the following host operating environments:
 - i. Windows 2008 R2 in 32-bit and 64-bit configurations;
 - ii. Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11 in 32-bit and 64-bit configurations; and
 - iii. VMWare ESX Server 3.5X and 4.X with shared access to the same logical unit numbers (LUNs) for Vmotion; and
- (b) it must directly support clustering under the following host operating environments:
 - i. MC/Serviceguard for HP-UX;
 - ii. PowerHA for AIX; and
 - iii. Oracle Solaris Cluster for Solaris

1.1.10 Software and Additional Capabilities

The storage platform must include the following software functionalities and additional capabilities. Furthermore, these must be entirely storage platform-based functionality and must not require any software or assistance from host systems on the SAN:

- (a) it must provide LUN-masking functionality. This means it must mask or limit visibility of logical drive configurations within the storage platform to only specific hosts connected to the storage platform;
- (b) it must synchronously replicate logical volumes remotely via extended fibre channel networks over dark fibre or dense wavelength division multiplexing (DWDM) / coarse wavelength division multiplexing (CWDM).;
- (c) it must asynchronously replicate logical volumes remotely via extended fibre channel networks over dark fibre, DWDM/CWDM, or TCP/IP.;
- (d) it must perform up to 4 concurrent host-less point-in-time snapshot copies of any logical volume that may be reassigned to any other host on the SAN.;
- (e) it must perform up to 2 concurrent host-less full block data copies of any logical volume that may be reassigned to any other host on the SAN.;
- (f) it must allow online firmware upgrades to be made without disrupting the operation of the platform; and

-
- (g) it must perform sub-LUN auto-tiering of data written to the storage platform.

1.1.11 Management

The storage platform must provide the following management capabilities:

- (a) it must provide a comprehensive graphical user interface (GUI) based management system that allows real-time monitoring of all components in the platform and reports degradation of components and failures;
- (b) the GUI interface must either be a Windows-based application included with the system or a WEB or Java-based embedded function accessible using a standard browser;
- (c) it must connect to an IP-based network either through a direct Ethernet connection on the platform or through an in-band connection via a fibre-attached host;
- (d) it must issue SNMP traps or SMTP mail in the event of device degradation or failure;
- (e) the GUI interface must show all installed hardware and its current operational status; and
- (f) the GUI interface must monitor the full performance of the storage array, including:
 - i. disk, LUN or RAID group I/O's per second for both read and write requests;
 - ii. cache utilization and hit rate statistics; and
 - iii. queuing information for disks, LUNs or RAID sets.

1.2 FABRIC

1.2.1 Fibre Channel Switch

The storage platform must operate with 8 Gbps 24 and 48 port fibre channel fabric switches, which must be fully supported and warranted by the storage platform Manufacturer. The fibre channel switches must meet the following requirements:

- (a) they must operate with fibre channel fabrics and must be capable of full fibre channel zoning across switched fabrics;
- (b) they must support a minimum of 512 active enabled unique zones at a time per fibre channel fabric;
- (c) they must be available in both stand-alone and rack mountable configurations. A rack mounting kit that is applied to a stand-alone switch will be accepted;
- (d) they must operate at 8GB/s and must be fully populated with small form factor pluggable optical media modules for shortwave operation;
- (e) they must provide lights or indicators for power and port status for all fibre channel ports;
- (f) they must provide a 1Gbps Ethernet interface and must be manageable using TCP/IP as the transport protocol;
- (g) they must provide redundant cooling and power;
- (h) they must fully comply with the following ANSI T-11 standards:
 - i. FC-PH,
 - ii. FC-PH-2,
 - iii. FC-PH-3,
 - iv. FC-AL,
 - v. FC-AL-2,
 - vi. FC-FLA,
 - vii. FC-FG,
 - viii. FC-GS-2,
 - ix. FC-PLDA,
 - x. FC-VI, and
 - xi. FC-SW-2;
- (i) they must support fibre channel class 2 and 3 connections;
- (j) they must provide full fabric support as per the ANSI standards specified at (1.2.1 h);
- (k) they must support cascading by connecting 2 or more switches together to form a single fabric that is compliant with the ANSI standards specified at (1.2.1 h);
- (l) they must include a comprehensive GUI-based management system that allows real-time monitoring of all components in the platform and to report failures or degraded components;

-
- (m) the GUI interface must be an embedded WEB or Java-based function accessible using a standard browser;
 - (n) they must generate SNMP traps in the event of a degraded condition in the switch;
 - (o) the GUI interface must show the current operational status for all installed hardware components;
 - (p) the GUI interface must allow configuration of all aspects of the fibre channel switches including:
 - i. the name,
 - ii. the domain ID,
 - iii. the passwords and user accounts for management,
 - iv. the IP addressing,
 - v. the modes of operation of the ports,
 - vi. all zone and path information, and
 - vii. any other parameters critical to the operation of the switch; and
 - (q) The GUI interface must provide complete performance monitoring allowing a storage administrator to view:
 - i. the number of frames per second, with a breakdown of which were good frames and which were error frames,
 - ii. the % utilization of fibre channel port,
 - iii. the operational speed of fibre channel ports,
 - iv. the mode of operation of fibre channel port (e.g. F-port, N-port, E-port), and
 - v. the throughput in frames as well as MB per second.

1.3 VIRTUALIZATION

1.3.1 Virtualization

The storage platform must include a virtualization solution that meets the following requirements:

- (a) it must be either:
 - i. be manufactured by the same Manufacturer as the base storage platform defined in 1.1; or
 - ii. be sold under the name of the same Manufacturer (sometimes referred to as rebranding) as the base storage platform, but only if that Manufacturer warrants, supports and maintains the solution.
- (b) it must be a discrete and independent device(s) that does not rely on any components, functionality or software from the base storage platform defined at 1.1;
- (c) it must be packaged in an industry-standard 19" rack mount form factor and must include all accessories, cables and hardware required to mount and power the unit in an industry-standard 19" rack;
- (d) it must have redundant and hot swappable power and cooling for all electronic components of the solution. Fully redundant pairs of equipment with fixed cooling and power that allow swapping of an entire portion of the solution without interrupting host access to virtualized storage is acceptable,
- (e) it must include 4 X 8Gbps host side fibre channel ports for connections to a client host side SAN fabric;
- (f) it must include 4 X 8Gbps host side fibre channel ports for connections to the mandatory base storage platform as well as potential existing client third party storage platforms
- (g) it must virtualize the base storage platform and must simultaneously support third party storage platforms from at least 5 of the following Manufacturers:
 - i. Dell,
 - ii. EMC,
 - iii. Hitachi Data Systems,
 - iv. HP,
 - v. IBM,
 - vi. Network Appliance, and
 - vii. Oracle.

-
- (h) it must either provide thin provisioning capabilities or allow virtual pools of storage to be created from one or more of the mandatory base platform(s) as well as any of the supported third party storage platforms simultaneously;
 - (i) it must present block storage from one or more base storage platforms as well as any of the supported additional third party storage platforms in the form of LUNs to the mandatory supported host platforms. The underlying details of the capacity from the base and third party systems must be masked from the host systems so that LUNs may be comprised of any or all of the underlying capacity;
 - (j) it must allow full block copies of LUNs to be created over local SAN connections between the supported base and third party systems and the dynamic relocation of LUNs without interrupting host access, without a loss of data, and without changing the addressing of those LUNs to the hosts; and
 - (k) it must allow synchronous and/or asynchronous copies of LUNs to be created over distance-extended SAN connections between the supported base and third party systems for the purpose of disaster recovery between any of the supported storage platforms.

1.4 NAS Gateway

1.4.1 Capacity and Platform

The storage platform must include a Network Attached Storage Gateway. The NAS Gateway must meet the following requirements:

- (a) it must either:
 - i. be manufactured by the same Manufacturer as the base storage platform defined in 1.1; or
 - ii. be sold under the name of the same Manufacturer (sometimes referred to as rebranding) as the base storage platform, but only if that Manufacturer warrants, supports and maintains the solution
- (b) it must be a discrete and independent device(s) that does not rely upon any components, functionality or software from the base storage platform defined at 1.1; however, the capacity this NAS Gateway will address and share may be provided by the base storage platform defined at 1.1;
- (c) it must address and share a minimum of 32TB of raw data storage while also adhering to all other minimums;
- (d) it must be fully compatible with and supported by the base storage platform defined at 1.1. Use of this NAS Gateway with the base storage platform must not preclude the base storage platform from also servicing other fibre channel block attached hosts at the same time;
- (e) it must include sufficient cooling for a fully populated configuration. All cooling for the NAS Gateway must be redundant and monitored for failures by the NAS Gateway;
- (f) it must allow hot swapping of failed cooling fans; and
- (g) it must be packaged in an industry-standard 19" rack mount form factor and must include all accessories, cables and hardware required to mount and power the unit in an industry-standard 19" rack.

1.4.2 Power

Each Network Attached Storage Gateway must meet the following power requirements:

- (a) it must provide sufficient power to operate a fully loaded system with all boards and components installed;
- (b) it must be fully redundant, allowing the NAS Gateway to continue operating without interruption in the event of a power supply failure, until service can be performed. Redundancy may be achieved either by using:
 - i. a second power supply, or
 - ii. an N+1 approach; and
- (c) each AC power supply must connect independently to a discrete AC power source.

1.4.3 Controllers and RAID

Each Network Attached Storage Gateway must provide capacity that is RAID protected by the base storage platform. The following RAID formats must be supported by the NAS Gateway when presented to it by the base storage platform defined at 1.1:

- i. RAID5 stripes with parity;
- ii. RAID6 stripes with dual parity;
- iii. RAID1; and
- iv. RAID0+1 stripes with mirroring or RAID1+0 striped mirrors (aka RAID10).

1.4.4 NAS Processor Unit

Each Network Attached Storage Gateway must meet the following requirements for the NAS processor unit(s):

- (a) it must include a micro-kernel operating system designed for providing file services to CIFS and NFS via the included Ethernet interfaces. The micro-kernel operating system may be either a Linux or Unix-based operating system;
- (b) it must load the micro-kernel operating system from a fault-tolerant medium that is either RAID protected, or duplicated and included, in a second NAS processor unit that may assume operation in the event of a failure to load the operating system at boot time;
- (c) it must contain 2 separate redundant clustered processor units or "heads" that operate in an active / active fashion providing network services to clients for CIFS and NFS. In the event of a failure of one of the processor units, the remaining unit must assume the IP address and identity of the failed processor unit and must continue to provide service to clients on the network automatically;
- (d) the processor units must both be attached via a total aggregate minimum of 4 X 4Gbps fibre channel interfaces to the base storage platform; and
- (e) the processor units in the NAS Gateway must contain an aggregate minimum of either 8 X 1Gbps or 2 X 10Gbps Ethernet interfaces for TCP/IP client access.

1.4.5 Software and Additional Capabilities

Each Network Attached Storage Gateway must meet the following requirements for software functionality and additional capabilities:

- (a) it must include all client access licenses for end user workstations to access and use the shared file systems via CIFS or NFS, with no requirement for additional fees or licensing;
- (b) it must fully integrate, in mixed mode or native mode, Microsoft Active directory environments and must be manageable as a Windows server in those environments using native Microsoft tools for viewing and managing sessions, shares and open files;
- (c) it must support snapshot functionality for all shared file systems allowing an administrator to create point-in-time copies of all files for the purpose of recovering deleted files; and
- (d) it must include and be licensed for NDMP to facilitate backups of the shared file systems to fibre channel attached backup targets.

1.4.6 Management

Each Network Attached Storage Gateway must meet the following requirements for management capabilities:

- (a) it must be manageable remotely via an Ethernet interface and must provide an intuitive GUI-based interface for day-to-day operations;
- (b) it must include a simple and intuitive installation system allowing non-technical operators to easily configure and provision the unit for operation on a network with only a basic knowledge of TCP/IP addresses and volume and file system management;
- (c) it must provide GUI-based functionality to:
 - i. create and manage volumes and file systems across RAID sets;

-
- ii. assign and manage user permissions for CIFS and NFS users to volumes and files;
 - iii. view attributes of volumes including space usage and file information;
 - iv. configure all user-assigned parameters required for operation of the system;
 - v. monitor utilization of network interfaces, processors and disk subsystems to gauge the load on those items;
 - vi. backup all locally hosted data to a locally-attached tape drive or provide an agent or facility for a remote console to initiate this process directly from the NAS disk to a backup target; and
 - vii. load balance file shares across either of the 2 processor units as needed and allow an administrator to manually failover file shares if required from 1 processor unit to the other.
- (d) the GUI management system must manage and operate both processor units as a single entity, allowing a single session to facilitate all management functions described here.

SAN 2.0 CATEGORY

The following describes the configuration and features of a **small Enterprise SAN** storage solution.

2.1 Storage Platform

2.1.1 Capacity and Platform

Each storage platform must meet the following capacity and platform requirements:

- (a) the hard disk drive technologies and densities must be commercially available, meaning that the Manufacturer is continuing to manufacture and ship them to customers generally;
- (b) the hard disk drive technologies and densities must be tested and fully supported within the storage platform by the storage platform Manufacturer;
- (c) it must include industry-standard hard disk drives operating at either 4Gbps for Fibre Channel (FC) drives or 6Gbps for Serial Attach SCSI (SAS) drives;
- (d) it must also include industry-standard Serial Advanced Technology Attachment (SATA) revision 3.0 or Nearline SAS (NL-SAS) hard disk drives operating at 6Gbps. This may be achieved either by:
 - i. using the same shelves as either the FC or SAS disk drives, or
 - ii. using specialized shelves for these drive types;
- (e) the available drive options must include at least seven (7) from the following list:

- drives with 4Gbps (for FC) or 6 Gbps (for SAS) interfaces and 15000 RPM rotational speed:

- i) 300GB
- ii) 450GB
- iii) 600GB

- drives with 4Gbps (for FC) or 6 Gbps (for SAS) interfaces and 10000 RPM rotational speed:

- iv) 300GB
- v) 450GB
- vi) 600GB
- vii) 900GB

- drives with 6Gbps for SATA or NL-SAS interfaces and 7200 RPM rotational speed:

- viii) 1TB
- ix) 2TB
- x) 3TB

- solid state drives (SSD) based on Single Level Cell (SLC) or enterprise-class Multi-Level Cell (eMLC) technology

xi)	100GB
xii)	200GB
xiii)	300GB
xiv)	400GB
xv)	600GB

- (f) it must accommodate a minimum of 448 hard disk drives;
- (g) It must be packaged in a standard 19" rack mount form factor (NOTE: it is understood that standard rack depth will be increased when "high density" disk shelves are provided); and
- (h) It must include lights or an LCD panel for power, activity and fault indications.

2.1.2 Cooling

Each storage platform must meet the following cooling requirements:

- (a) it must provide sufficient cooling for a fully populated cabinet at the mandatory minimum storage capacity;
- (b) all cooling for the system controller(s) as well as all hard disk drives must be redundant and monitored for failure by the storage platform hardware;
- (c) it must allow hot swapping of failed cooling fans;
- (d) the cooling system within the storage platform itself must be fully redundant; and
- (e) in the event of a component failure, the cooling system must allow continued operation of the storage platform until service can be performed.

2.1.3 Drives and Shelves

Each storage platform must meet the following drives and shelves requirements:

- (a) the hard disk drives must include a minimum of either 4Gbps for Fibre Channel (FC) or 6Gbps for Serial Attach SCSI (SAS) with dual attached interfaces;
- (b) it must provide minimum of 8 active connections to the mandatory 448 hard disk drives. Bandwidth must be allocated evenly to the total number of physical drives over several channels;
- (c) it must provide fully redundant back-end paths to all hard disk drives. A channel failure must not interrupt access to attached disk drives;
- (d) it must allow hot addition of storage shelves without needing to power the storage platform down and without interrupting access to existing drives and redundant arrays of inexpensive disk (RAID) groups;
- (e) it must include as many back-end channels as necessary to support all the back-end shelves of disks so that a shelf component replacement or failure does not interrupt access to adjacent shelves in the platform.
- (f) the hard disk drives in the storage platform must be fully hot pluggable while the storage platform is operational. There must be no loss of data if a hard drive is removed, assuming the drive is part of a fault-tolerant configuration in the platform;
- (g) it must rebuild a replaced hard disk drive automatically and without user intervention when it is inserted, assuming it is replacing a hard disk drive that was part of a fault-tolerant configuration; and
- (h) it must allow hard disk drives to be allocated as global hot spares, which must automatically rebuild the contents of a failed hard disk drive in any fault-tolerant RAID set. This process must be fully automatic whenever a disk failure occurs in a fault-tolerant RAID set.

2.1.4 Power

Each storage platform must meet the following power requirements:

- (a) it must provide sufficient power to operate a fully populated system with all boards and cache installed, and the maximum number of hard disk drives installed;
- (b) the power supplies must be fully redundant, allowing uninterrupted operation of the storage platform in the event of a power supply failure, until service can be performed. Redundancy may be achieved either by using:
 - i. a second power supply, or
 - ii. an N+1 approach; and
- (c) each AC power supply must connect independently to a discrete AC power source.

2.1.5 Controllers

Each storage platform must meet the following controller requirements:

- (a) it must include dual redundant active/active storage controllers for handling both I/O to the attached host systems as well as disk I/O and RAID functionality;
- (b) it must be redundant, so that the surviving controller automatically recovers controller subsystem failures, and service to attached hosts is continued without disruption;
- (c) each storage controller within the storage platform must have access to all 448 of the mandatory hard disk drives in order to assign, configure, protect and share those drives;
- (d) the storage controllers must allow configuration of hard disk drives within the storage platform as:
 - i. RAID5 stripes with parity;
 - ii. RAID6 stripes with dual parity;
 - iii. RAID1; and
 - iv. RAID0+1 stripes with mirroring or RAID1+0 striped mirrors (aka RAID10).
- (e) it must allow the creation and addressing of up to 4096 simultaneous logical drives; and
- (f) it must simultaneously support all RAID types from 2.1.5 (d) within the storage platform.

2.1.6 Cache

Each storage platform must meet the following cache requirements:

- (a) it must include at least 16GB of dedicated I/O cache per storage controller;
- (b) the cache on the storage controller must perform both read and write I/O operations;
- (c) the write cache must be mirrored and must utilize error detection and correction logic that will detect and recover from any memory errors without data loss or interruption to service; and
- (d) the write data within the cache on the storage controllers must be protected by one of these two methods:
 - i. a battery that allows the cache contents to be held intact for a minimum of 48 hours. The caches must then complete their write operations to disk when power is restored; or
 - ii. all pending write data must be automatically written to disk before the disk system is powered off, and the platform must provide sufficient battery power to complete this function.

2.1.7 I/O Ports and Connectivity

Each storage platform must meet the following requirements for I/O ports and connectivity:

- (a) it must include a minimum of 2 storage controllers that may be replaced in the event of a controller failure;
- (b) it must provide a minimum of 4 fibre channel ports for connectivity to Intel and Open System host computers;
- (c) all 4 fibre channel ports must be independent ports operating at 8Gbps each and operate in either point-to-point or loop modes of operation;

-
- (d) each of the 4 fibre ports must support full fabric login and must have a unique fibre channel World Wide Name;
 - (e) it must provide simultaneous connectivity to any combination of 256 or more Intel and UNIX hosts using dual fibre channel host bus adapters in each host;
 - (f) it must provide host bus adapters including any required software drivers for all supported operating systems; and
 - (g) it must offer "no single point of failure" connectivity options, for both failover as well as load balancing under all of the mandatory operating system environments. This may be provided using add-on failover software packages or using native Operating System facilities
 - (h) it must provide four (4) native 10GbE connections for FCoE host connectivity that meet the ANSI T11 FC-BB-5 Fibre Channel Over Ethernet (FCoE) standards for the encapsulation of FC packets over Full Duplex and Lossless Ethernet networks and must be compliant with the following IEEE standards.
 - i. 802.1Ae;
 - ii. 802.1Qbb;
 - iii. 802.1Qaz which defines:
 - a. enhanced transmission selection (ETS); and
 - b. data center bridging exchange (DCBX).

2.1.8 Hosts

Each storage platform must meet the following requirements for host connectivity:

- (a) it must connect to Intel-based host computers running the following:
 - i. Windows 2008 R2 with Hyper-V in 32-bit and 64-bit configurations;
 - ii. Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11 in 32-bit and 64-bit configurations; and
 - iii. VMWare ESX Server 3.5X and 4.X;
- (b) it must connect to the following UNIX and Open Systems hosts simultaneously, in addition to the previously listed Intel systems:
 - i. Oracle Solaris 10 and 11 systems;
 - ii. HP-UX 11i v.X systems;
 - iii. IBM AIX v.6X and v.7X systems;
- (c) Support of additional platform types and operating systems is desirable, but not mandatory.

2.1.9 Clustering

Each storage platform must meet the following requirements for clustering:

- (a) it must directly support clustering under the following host operating environments:
 - i. Windows 2008 R2 in 32-bit and 64-bit configurations;
 - ii. Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11 in 32-bit and 64-bit configurations; and
 - iii. VMWare ESX Server 3.5X and 4.X with shared access to the same logical unit numbers (LUNs) for Vmotion; and
- (b) it must directly support clustering under the following host operating environments:
 - i. MC/Serviceguard for HP-UX;
 - ii. PowerHA for AIX; and
 - Oracle Solaris Cluster for Solaris

2.1.10 Software and Additional Capabilities

The storage platform must include the following software functionalities and additional capabilities. Furthermore, these must be entirely storage platform-based functionality and must not require any software or assistance from host systems on the SAN:

-
- (a) it must provide LUN-masking functionality. This means it must mask or limit visibility of logical drive configurations within the storage platform to only specific hosts connected to the storage platform;
 - (b) it must synchronously replicate logical volumes remotely via extended fibre channel networks over dark fibre or dense wavelength division multiplexing (DWDM) / coarse wavelength division multiplexing (CWDM).;
 - (c) it must asynchronously replicate logical volumes remotely via extended fibre channel networks over dark fibre, DWDM/CWDM, or TCP/IP.;
 - (d) it must perform up to 4 concurrent host-less point-in-time snapshot copies of any logical volume that may be reassigned to any other host on the SAN.;
 - (e) it must perform up to 2 concurrent host-less full block data copies of any logical volume that may be reassigned to any other host on the SAN.;
 - (f) it must allow online firmware upgrades to be made without disrupting the operation of the platform; and
 - (g) it must perform sub-LUN auto-tiering of data written to the storage platform.

2.1.11 Management

The storage platform must provide the following management capabilities:

- (a) it must provide a comprehensive graphical user interface (GUI) based management system that allows real-time monitoring of all components in the platform and reports degradation of components and failures;
- (b) the GUI interface must either be a Windows-based application included with the system or a WEB or Java-based embedded function accessible using a standard browser;
- (c) it must connect to an IP-based network either through a direct Ethernet connection on the platform or through an in-band connection via a fibre-attached host;
- (d) it must issue SNMP traps or SMTP mail in the event of device degradation or failure;
- (e) the GUI interface must show all installed hardware and its current operational status; and
- (f) the GUI interface must monitor the full performance of the storage array, including:
 - i. disk, LUN or RAID group I/O's per second for both read and write requests;
 - ii. cache utilization and hit rate statistics; and
 - iii. queuing information for disks, LUNs or RAID sets.

2.2 Fabric

2.2.1 Fibre Channel Switch

The storage platform must operate with 8 Gbps 24 and 48 port fibre channel fabric switches, which must be fully supported and warranted by the storage platform Manufacturer. The fibre channel switches must meet the following requirements:

- (a) they must operate with fibre channel fabrics and must be capable of full fibre channel zoning across switched fabrics;
- (b) they must support a minimum of 512 active enabled unique zones at a time per fibre channel fabric;
- (c) they must be available in both stand-alone and rack mountable configurations. A rack mounting kit that is applied to a stand-alone switch will be accepted;
- (d) they must operate at 8GB/s and must be fully populated with small form factor pluggable optical media modules for shortwave operation;
- (e) they must provide lights or indicators for power and port status for all fibre channel ports;
- (f) they must provide a 1Gbps Ethernet interface and must be manageable using TCP/IP as the transport protocol;
- (g) they must provide redundant cooling and power;
- (h) they must fully comply with the following ANSI T-11 standards:
 - i. FC-PH,
 - ii. FC-PH-2,

-
- iii. FC-PH-3,
 - iv. FC-AL,
 - v. FC-AL-2,
 - vi. FC-FLA,
 - vii. FC-FG,
 - viii. FC-GS-2,
 - ix. FC-PLDA,
 - x. FC-VI, and
 - xi. FC-SW-2;
- (i) they must support fibre channel class 2 and 3 connections;
 - (j) they must provide full fabric support as per the ANSI standards specified at (2.2.1 h);
 - (k) they must support cascading by connecting 2 or more switches together to form a single fabric that is compliant with the ANSI standards specified at (2.2.1 h);
 - (l) they must include a comprehensive GUI-based management system that allows real-time monitoring of all components in the platform and to report failures or degraded components;
 - (m) the GUI interface must be an embedded WEB or Java-based function accessible using a standard browser;
 - (n) they must generate SNMP traps in the event of a degraded condition in the switch;
 - (o) the GUI interface must show the current operational status for all installed hardware components;
 - (p) the GUI interface must allow configuration of all aspects of the fibre channel switches including:
 - i. the name,
 - ii. the domain ID,
 - iii. the passwords and user accounts for management,
 - iv. the IP addressing,
 - v. the modes of operation of the ports,
 - vi. all zone and path information, and
 - vii. any other parameters critical to the operation of the switch; and
 - (q) The GUI interface must provide complete performance monitoring allowing a storage administrator to view:
 - i. the number of frames per second, with a breakdown of which were good frames and which were error frames,
 - ii. the % utilization of fibre channel port,
 - iii. the operational speed of fibre channel ports,
 - iv. the mode of operation of fibre channel port (e.g. F-port, N-port, E-port), and
 - v. the throughput in frames as well as MB per second.

2.3 Virtualization

2.3.1 Virtualization

The storage platform must include a virtualization solution that meets the following requirements:

- (a) it must be either:
 - i. be manufactured by the same Manufacturer as the base storage platform defined in 2.1; or
 - ii. be sold under the name of the same Manufacturer (sometimes referred to as rebranding) as the base storage platform, but only if that Manufacturer warrants, supports and maintains the solution.
- (b) it must be a discrete and independent device(s) that does not rely on any components, functionality or software from the base storage platform defined at 2.1;
- (c) it must be packaged in an industry-standard 19" rack mount form factor and must include all accessories, cables and hardware required to mount and power the unit in an industry-standard 19" rack;
- (d) it must have redundant and hot swappable power and cooling for all electronic components of the solution. Fully redundant pairs of equipment with fixed cooling and power that allow swapping of an entire portion of the solution without interrupting host access to virtualized storage is acceptable,
- (e) it must include 4 X 8Gbps host side fibre channel ports for connections to a client host side SAN fabric;

-
- (f) it must include 4 X 8Gbps host side fibre channel ports for connections to the mandatory base storage platform as well as potential existing client third party storage platforms
 - (g) it must virtualize the base storage platform and must simultaneously support third party storage platforms from at least 5 of the following Manufacturers:
 - i. Dell,
 - ii. EMC,
 - iii. Hitachi Data Systems,
 - iv. HP,
 - v. IBM,
 - vi. Network Appliance, and
 - vii. Oracle.
 - (h) it must either provide thin provisioning capabilities or allow virtual pools of storage to be created from one or more of the mandatory base platform(s) as well as any of the supported third party storage platforms simultaneously;
 - (i) it must present block storage from one or more base storage platforms as well as any of the supported additional third party storage platforms in the form of LUNs to the mandatory supported host platforms. The underlying details of the capacity from the base and third party systems must be masked from the host systems so that LUNs may be comprised of any or all of the underlying capacity;
 - (j) it must allow full block copies of LUNs to be created over local SAN connections between the supported base and third party systems and the dynamic relocation of LUNs without interrupting host access, without a loss of data, and without changing the addressing of those LUNs to the hosts; and
 - (k) it must allow synchronous and/or asynchronous copies of LUNs to be created over distance-extended SAN connections between the supported base and third party systems for the purpose of disaster recovery between any of the supported storage platforms.

2.4 NAS Gateway

2.4.1 Capacity and Platform

The storage platform must include a Network Attached Storage Gateway. The NAS Gateway must meet the following requirements:

- (a) it must either:
 - i. be manufactured by the same Manufacturer as the base storage platform defined in 2.1; or
 - ii. be sold under the name of the same Manufacturer (sometimes referred to as rebranding) as the base storage platform, but only if that Manufacturer warrants, supports and maintains the solution
- (b) it must be a discrete and independent device(s) that does not rely upon any components, functionality or software from the base storage platform defined at 2.1; however, the capacity this NAS Gateway will address and share may be provided by the base storage platform defined at 2.1;
- (c) it must address and share a minimum of 64TB of raw data storage while also adhering to all other minimums;
- (d) it must be fully compatible with and supported by the base storage platform defined at 2.1. Use of this NAS Gateway with the base storage platform must not preclude the base storage platform from also servicing other fibre channel block attached hosts at the same time;
- (e) it must include sufficient cooling for a fully populated configuration. All cooling for the NAS Gateway must be redundant and monitored for failures by the NAS Gateway;
- (f) it must allow hot swapping of failed cooling fans; and
- (g) it must be packaged in an industry-standard 19" rack mount form factor and must include all accessories, cables and hardware required to mount and power the unit in an industry-standard 19" rack.

2.4.2 Power

Each Network Attached Storage Gateway must meet the following power requirements:

- (a) it must provide sufficient power to operate a fully loaded system with all boards and components installed;
- (b) it must be fully redundant, allowing the NAS Gateway to continue operating without interruption in the event of a power supply failure, until service can be performed. Redundancy may be achieved either by using:
 - i) a second power supply, or
 - ii) an N+1 approach; and
- (c) each AC power supply must connect independently to a discrete AC power source.

2.4.3 Controllers and RAID

Each Network Attached Storage Gateway must provide capacity that is RAID protected by the base storage platform. The following RAID formats must be supported by the NAS Gateway when presented to it by the base storage platform defined at 2.1:

- i. RAID5 stripes with parity;
- ii. RAID6 stripes with dual parity;
- iii. RAID1; and
- iv. RAID0+1 stripes with mirroring, or RAID1+0 striped mirrors (aka RAID10).

2.4.4 NAS Processor Unit

Each Network Attached Storage Gateway must meet the following requirements for the NAS processor unit(s):

- (a) it must include a micro-kernel operating system designed for providing file services to CIFS and NFS via the included Ethernet interfaces. The micro-kernel operating system may be either a Linux or Unix-based operating system;
- (b) it must load the micro-kernel operating system from a fault-tolerant medium that is either RAID protected, or duplicated and included, in a second NAS processor unit that may assume operation in the event of a failure to load the operating system at boot time;
- (c) it must contain 2 separate redundant clustered processor units or "heads" that operate in an active / active fashion providing network services to clients for CIFS and NFS. In the event of a failure of one of the processor units, the remaining unit must assume the IP address and identity of the failed processor unit and must continue to provide service to clients on the network automatically;
- (d) the processor units must both be attached via a total aggregate minimum of 4 X 4Gbps fibre channel interfaces to the base storage platform; and
- (e) the processor units in the NAS Gateway must contain an aggregate minimum of either 12 X 1Gbps or 4 X 10Gbps Ethernet interfaces for TCP/IP client access.

2.4.5 Software and Additional Capabilities

Each Network Attached Storage Gateway must meet the following requirements for software functionality and additional capabilities:

- (a) it must include all client access licenses for end user workstations to access and use the shared file systems via CIFS or NFS, with no requirement for additional fees or licensing;
- (b) it must fully integrate, in mixed mode or native mode, Microsoft Active directory environments and must be manageable as a Windows server in those environments using native Microsoft tools for viewing and managing sessions, shares and open files;
- (c) it must support snapshot functionality for all shared file systems allowing an administrator to create point-in-time copies of all files for the purpose of recovering deleted files; and
- (d) it must include and be licensed for NDMP to facilitate backups of the shared file systems to fibre channel attached backup targets.

2.4.6 Management

Each Network Attached Storage Gateway must meet the following requirements for management capabilities:

- (a) it must be manageable remotely via an Ethernet interface and must provide an intuitive GUI-based interface for day-to-day operations;
- (b) it must include a simple and intuitive installation system allowing non-technical operators to easily configure and provision the unit for operation on a network with only a basic knowledge of TCP/IP addresses and volume and file system management;
- (c) it must provide GUI-based functionality to:
 - i. create and manage volumes and file systems across RAID sets;
 - ii. assign and manage user permissions for CIFS and NFS users to volumes and files;
 - iii. view attributes of volumes including space usage and file information;
 - iv. configure all user-assigned parameters required for operation of the system;
 - v. monitor utilization of network interfaces, processors and disk subsystems to gauge the load on those items;
 - vi. backup all locally hosted data to a locally-attached tape drive or provide an agent or facility for a remote console to initiate this process directly from the NAS disk to a backup target; and
 - vii. load balance file shares across either of the 2 processor units as needed and allow an administrator to manually failover file shares if required from 1 processor unit to the other.
- (d) the GUI management system must manage and operate both processor units as a single entity, allowing a single session to facilitate all management functions described here.

SAN 3.0 CATEGORY

The following describes the configuration and features of a **scalable Enterprise SAN** storage solution.

3.1 Storage Platform

3.1.1 Capacity and Platform

Each storage platform must meet the following capacity and platform requirements:

- (a) the hard disk drive technologies and densities must be commercially available, meaning that the Manufacturer is continuing to manufacture and ship them to customers generally;
- (b) the hard disk drive technologies and densities must be tested and fully supported within the storage platform by the storage platform Manufacturer;
- (c) it must include industry-standard hard disk drives operating at either 4Gbps for Fibre Channel (FC) drives or 6Gbps for Serial Attach SCSI (SAS) drives;
- (d) it must also include industry-standard Serial Advanced Technology Attachment (SATA) revision 3.0 or Nearline SAS (NL-SAS) hard disk drives operating at 6Gbps. This may be achieved either by:
 - i. using the same shelves as either the FC or SAS disk drives, or
 - ii. using specialized shelves for these drive types;
- (e) the available drive options must include at least seven (7) from the following list:

- drives with 4Gbps (for FC) or 6 Gbps (for SAS) interfaces and 15000 RPM rotational speed:

- i) 300GB
- ii) 450GB
- iii) 600GB

- drives with 4Gbps (for FC) or 6 Gbps (for SAS) interfaces and 10000 RPM rotational speed:

- iv) 300GB
- v) 450GB
- vi) 600GB
- vii) 900GB

- drives with 6Gbps for SATA or NL-SAS interfaces and 7200 RPM rotational speed:

- viii) 1TB
- ix) 2TB
- x) 3TB

- solid state drives (SSD) based on Single Level Cell (SLC) or enterprise-class Multi-Level Cell (eMLC) technology

- xi) 100GB
- xii) 200GB
- xiii) 300GB
- xiv) 400GB
- xv) 600GB

- (f) it must accommodate a minimum of 960 hard disk drives;
- (g) It must be packaged in a standard 19" rack mount form factor (NOTE: it is understood that standard rack depth will be increased when "high density" disk shelves are provided); and
- (h) It must include lights or an LCD panel for power, activity and fault indications.

3.1.2 Cooling

Each storage platform must meet the following cooling requirements:

- (a) it must provide sufficient cooling for a fully populated cabinet at the mandatory minimum storage capacity;
- (b) all cooling for the system controller(s) as well as all hard disk drives must be redundant and monitored for failure by the storage platform hardware;
- (c) it must allow hot swapping of failed cooling fans;
- (d) the cooling system within the storage platform itself must be fully redundant; and
- (e) in the event of a component failure, the cooling system must allow continued operation of the storage platform until service can be performed.

3.1.3 Drives and Shelves

Each storage platform must meet the following drives and shelves requirements:

- (a) the hard disk drives must include a minimum of either 4Gbps for Fibre Channel (FC) or 6Gbps for Serial Attach SCSI (SAS) with dual attached interfaces;
- (b) it must provide minimum of 16 active connections to the mandatory 960 hard disk drives. Bandwidth must be allocated evenly to the total number of physical drives over several channels;
- (c) it must provide fully redundant back-end paths to all hard disk drives. A channel failure must not interrupt access to attached disk drives;
- (d) it must allow hot addition of storage shelves without needing to power the storage platform down and without interrupting access to existing drives and redundant arrays of inexpensive disk (RAID) groups;
- (e) it must include as many back-end channels as necessary to support all the back-end shelves of disks so that a shelf component replacement or failure does not interrupt access to adjacent shelves in the platform.

- (f) the hard disk drives in the storage platform must be fully hot pluggable while the storage platform is operational. There must be no loss of data if a hard drive is removed, assuming the drive is part of a fault-tolerant configuration in the platform;
- (g) it must rebuild a replaced hard disk drive automatically and without user intervention when it is inserted, assuming it is replacing a hard disk drive that was part of a fault-tolerant configuration; and
- (h) it must allow hard disk drives to be allocated as global hot spares, which must automatically rebuild the contents of a failed hard disk drive in any fault-tolerant RAID set. This process must be fully automatic whenever a disk failure occurs in a fault-tolerant RAID set.

3.1.4 Power

Each storage platform must meet the following power requirements:

- (a) it must provide sufficient power to operate a fully populated system with all boards and cache installed, and the maximum number of hard disk drives installed;
- (b) the power supplies must be fully redundant, allowing uninterrupted operation of the storage platform in the event of a power supply failure, until service can be performed. Redundancy may be achieved either by using:
 - i. a second power supply, or
 - ii. an N+1 approach; and
- (c) each AC power supply must connect independently to a discrete AC power source.

3.1.5 Controllers

Each storage platform must meet the following controller requirements:

- (a) it must include dual redundant active/active storage controllers for handling both I/O to the attached host systems as well as disk I/O and RAID functionality;
- (b) it must be redundant, so that the surviving controller automatically recovers controller subsystem failures, and service to attached hosts is continued without disruption;
- (c) each storage controller within the storage platform must have access to all 960 of the mandatory hard disk drives in order to assign, configure, protect and share those drives;
- (d) the storage controllers must allow configuration of hard disk drives within the storage platform as:
 - i. RAID5 stripes with parity;
 - ii. RAID6 stripes with dual parity;
 - iii. RAID 1; and
 - iv. RAID0+1 stripes with mirroring or RAID1+0 striped mirrors (aka RAID 10).
- (e) it must allow the creation and addressing of up to 4096 simultaneous logical drives; and
- (f) it must simultaneously support all RAID types from 3.1.5 (d) within the storage platform.

3.1.6 Cache

Each storage platform must meet the following cache requirements:

- (a) it must include at least 32GB of dedicated I/O cache per storage controller;
- (b) the cache on the storage controller must perform both read and write I/O operations;
- (c) the write cache must be mirrored and must utilize error detection and correction logic that will detect and recover from any memory errors without data loss or interruption to service; and
- (d) the write data within the cache on the storage controllers must be protected by one of these two methods:
 - i. a battery that allows the cache contents to be held intact for a minimum of 48 hours. The caches must then complete their write operations to disk when power is restored; or
 - ii. all pending write data must be automatically written to disk before the disk system is powered off, and the platform must provide sufficient battery power to complete this function.

3.1.7 I/O Ports and Connectivity

Each storage platform must meet the following requirements for I/O ports and connectivity:

- (a) it must include a minimum of 2 storage controllers that may be replaced in the event of a controller failure;
- (b) it must provide a minimum of 4 fibre channel ports for connectivity to Intel and Open System host computers;
- (c) all 4 fibre channel ports must be independent ports operating at 8Gbps each and operate in either point-to-point or loop modes of operation;
- (d) each of the 4 fibre ports must support full fabric login and must have a unique fibre channel World Wide Name;
- (e) it must provide simultaneous connectivity to any combination of 512 or more Intel and UNIX hosts using dual fibre channel host bus adapters in each host;
- (f) it must provide host bus adapters including any required software drivers for all supported operating systems; and
- (g) it must offer "no single point of failure" connectivity options, for both failover as well as load balancing under all of the mandatory operating system environments. This may be provided using add-on failover software packages or using native Operating System facilities
- (h) it must provide four (4) native 10GbE connections for FCoE host connectivity that meet the ANSI T11 FC-BB-5 Fibre Channel Over Ethernet (FCoE) standards for the encapsulation of FC packets over Full Duplex and Lossless Ethernet networks and be compliant with the following IEEE standards:
 - i. 802.1Ae;
 - ii. 802.1Qbb;
 - iii. 802.1Qaz which defines:
 - a. enhanced transmission selection (ETS); and
 - b. data center bridging exchange (DCBX).

3.1.8 Hosts

Each storage platform must meet the following requirements for host connectivity:

- (a) it must connect to Intel-based host computers running the following:
 - i. Windows 2008 R2 with Hyper-V in 32-bit and 64-bit configurations;
 - ii. Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11 in 32-bit and 64-bit configurations; and
 - iii. VMWare ESX Server 3.5X and 4.X;
- (b) it must connect to the following UNIX and Open Systems hosts simultaneously, in addition to the previously listed Intel systems:
 - i. Oracle Solaris 10 and 11 systems;
 - ii. HP-UX 11i v.X systems;
 - iii. IBM AIX v.6X and v.7X systems;
- (c) Support of additional platform types and operating systems is desirable, but not mandatory.

3.1.9 Clustering

Each storage platform must meet the following requirements for clustering:

- (a) it must directly support clustering under the following host operating environments:
 - i. Windows 2008 R2 in 32-bit and 64-bit configurations;
 - ii. Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11 in 32-bit and 64-bit configurations; and
 - iii. VMWare ESX Server 3.5X and 4.X with shared access to the same logical unit numbers (LUNs) for Vmotion; and
- (b) it must directly support clustering under the following host operating environments:
 - i. MC/Serviceguard for HP-UX ;

- ii. PowerHA for AIX; and
- iii. Oracle Solaris Cluster for Solaris

3.1.10 Software and Additional Capabilities

The storage platform must include the following software functionalities and additional capabilities. Furthermore, these must be entirely storage platform-based functionality and must not require any software or assistance from host systems on the SAN:

- (a) it must provide LUN-masking functionality. This means it must mask or limit visibility of logical drive configurations within the storage platform to only specific hosts connected to the storage platform;
- (b) it must synchronously replicate logical volumes remotely via extended fibre channel networks over dark fibre or dense wavelength division multiplexing (DWDM) / coarse wavelength division multiplexing (CWDM).;
- (c) it must asynchronously replicate logical volumes remotely via extended fibre channel networks over dark fibre, DWDM/CWDM, or TCP/IP.;
- (d) it must perform up to 4 concurrent host-less point-in-time snapshot copies of any logical volume that may be reassigned to any other host on the SAN.;
- (e) it must perform up to 2 concurrent host-less full block data copies of any logical volume that may be reassigned to any other host on the SAN.;
- (f) it must allow online firmware upgrades to be made without disrupting the operation of the platform; and
- (g) it must perform sub-LUN auto-tiering of data written to the storage platform.

3.1.11 Management

The storage platform must provide the following management capabilities:

- (a) it must provide a comprehensive graphical user interface (GUI) based management system that allows real-time monitoring of all components in the platform and reports degradation of components and failures;
- (b) the GUI interface must either be a Windows-based application included with the system or a WEB or Java-based embedded function accessible using a standard browser;
- (c) it must connect to an IP-based network either through a direct Ethernet connection on the platform or through an in-band connection via a fibre-attached host;
- (d) it must issue SNMP traps or SMTP mail in the event of device degradation or failure;
- (e) the GUI interface must show all installed hardware and its current operational status; and
- (f) the GUI interface must monitor the full performance of the storage array, including:
 - i. disk, LUN or RAID group I/O's per second for both read and write requests;
 - ii. cache utilization and hit rate statistics; and
 - iii. queuing information for disks, LUNs or RAID sets.

3.2 Fabric

3.2.1 Fibre Channel Switch

The storage platform must operate with 8 Gbps 24 and 48 port fibre channel fabric switches, which must be fully supported and warranted by the storage platform Manufacturer. The fibre channel switches must meet the following requirements:

- (a) they must operate with fibre channel fabrics and must be capable of full fibre channel zoning across switched fabrics;
- (b) they must support a minimum of 512 active enabled unique zones at a time per fibre channel fabric;
- (c) they must be available in both stand-alone and rack mountable configurations. A rack mounting kit that is applied to a stand-alone switch will be accepted;

-
- (d) they must operate at 8GB/s and must be fully populated with small form factor pluggable optical media modules for shortwave operation;
 - (e) they must provide lights or indicators for power and port status for all fibre channel ports;
 - (f) they must provide a 1Gbps Ethernet interface and must be manageable using TCP/IP as the transport protocol;
 - (g) they must provide redundant cooling and power;
 - (h) they must fully comply with the following ANSI T-11 standards:
 - i. FC-PH,
 - ii. FC-PH-2,
 - iii. FC-PH-3,
 - iv. FC-AL,
 - v. FC-AL-2,
 - vi. FC-FLA,
 - vii. FC-FG,
 - viii. FC-GS-2,
 - ix. FC-PLDA,
 - x. FC-VI, and
 - xi. FC-SW-2;
 - (i) they must support fibre channel class 2 and 3 connections;
 - (j) they must provide full fabric support as per the ANSI standards specified at (3.2.1 h);
 - (k) they must support cascading by connecting 2 or more switches together to form a single fabric that is compliant with the ANSI standards specified at (3.2.1 h);
 - (l) they must include a comprehensive GUI-based management system that allows real-time monitoring of all components in the platform and to report failures or degraded components;
 - (m) the GUI interface must be an embedded WEB or Java-based function accessible using a standard browser;
 - (n) they must generate SNMP traps in the event of a degraded condition in the switch;
 - (o) the GUI interface must show the current operational status for all installed hardware components;
 - (p) the GUI interface must allow configuration of all aspects of the fibre channel switches including:
 - i. the name,
 - ii. the domain ID,
 - iii. the passwords and user accounts for management,
 - iv. the IP addressing,
 - v. the modes of operation of the ports,
 - vi. all zone and path information, and
 - vii. any other parameters critical to the operation of the switch; and
 - (q) The GUI interface must provide complete performance monitoring allowing a storage administrator to view:
 - i. the number of frames per second, with a breakdown of which were good frames and which were error frames,
 - ii. the % utilization of fibre channel port,
 - iii. the operational speed of fibre channel ports,
 - iv. the mode of operation of fibre channel port (e.g. F-port, N-port, E-port), and
 - v. the throughput in frames as well as MB per second.

3.3 Virtualization

3.3.1 Virtualization

The storage platform must include a virtualization solution that meets the following requirements:

- (a) it must be either:
 - i. be manufactured by the same Manufacturer as the base storage platform defined in 3.1; or

-
- ii. be sold under the name of the same Manufacturer (sometimes referred to as rebranding) as the base storage platform, but only if that Manufacturer warrants, supports and maintains the solution.
 - (b) it must be a discrete and independent device(s) that does not rely on any components, functionality or software from the base storage platform defined at 3.1;
 - (c) it must be packaged in an industry-standard 19" rack mount form factor and must include all accessories, cables and hardware required to mount and power the unit in an industry-standard 19" rack;
 - (d) it must have redundant and hot swappable power and cooling for all electronic components of the solution. Fully redundant pairs of equipment with fixed cooling and power that allow swapping of an entire portion of the solution without interrupting host access to virtualized storage is acceptable;
 - (e) it must include 4 X 8Gbps host side fibre channel ports for connections to a client host side SAN fabric;
 - (f) it must include 4 X 8Gbps host side fibre channel ports for connections to the mandatory base storage platform as well as potential existing client third party storage platforms
 - (g) it must virtualize the base storage platform and must simultaneously support third party storage platforms from at least 5 of the following Manufacturers:
 - i. Dell,
 - ii. EMC,
 - iii. Hitachi Data Systems,
 - iv. HP,
 - v. IBM,
 - vi. Network Appliance, and
 - vii. Oracle.
 - (h) it must either provide thin provisioning capabilities or allow virtual pools of storage to be created from one or more of the mandatory base platform(s) as well as any of the supported third party storage platforms simultaneously;
 - (i) it must present block storage from one or more base storage platforms as well as any of the supported additional third party storage platforms in the form of LUNs to the mandatory supported host platforms. The underlying details of the capacity from the base and third party systems must be masked from the host systems so that LUNs may be comprised of any or all of the underlying capacity;
 - (j) it must allow full block copies of LUNs to be created over local SAN connections between the supported base and third party systems and the dynamic relocation of LUNs without interrupting host access, without a loss of data, and without changing the addressing of those LUNs to the hosts; and
 - (k) it must allow synchronous and/or asynchronous copies of LUNs to be created over distance-extended SAN connections between the supported base and third party systems for the purpose of disaster recovery between any of the supported storage platforms.

3.4 NAS Gateway

3.4.1 Capacity and Platform

The storage platform must include a Network Attached Storage Gateway. The NAS Gateway must meet the following requirements:

- (a) it must either:
 - i. be manufactured by the same Manufacturer as the base storage platform defined in 3.1; or
 - ii. be sold under the name of the same Manufacturer (sometimes referred to as rebranding) as the base storage platform, but only if that Manufacturer warrants, supports and maintains the solution
- (b) it must be a discrete and independent device(s) that does not rely upon any components, functionality or software from the base storage platform defined at 3.1; however, the capacity this NAS Gateway will address and share may be provided by the base storage platform defined at 3.1;

- (c) it must address and share a minimum of 64TB of raw data storage while also adhering to all other minimums;
- (d) it must be fully compatible with and supported by the base storage platform defined at 3.1. Use of this NAS Gateway with the base storage platform must not preclude the base storage platform from also servicing other fibre channel block attached hosts at the same time;
- (e) it must include sufficient cooling for a fully populated configuration. All cooling for the NAS Gateway must be redundant and monitored for failures by the NAS Gateway;
- (f) it must allow hot swapping of failed cooling fans; and
- (g) it must be packaged in an industry-standard 19" rack mount form factor and must include all accessories, cables and hardware required to mount and power the unit in an industry-standard 19" rack.

3.4.2 Power

Each Network Attached Storage Gateway must meet the following power requirements:

- (a) it must provide sufficient power to operate a fully loaded system with all boards and components installed;
- (b) it must be fully redundant, allowing the NAS Gateway to continue operating without interruption in the event of a power supply failure, until service can be performed. Redundancy may be achieved either by using:
 - i. a second power supply, or
 - ii. an N+1 approach; and
- (c) each AC power supply must connect independently to a discrete AC power source.

3.4.3 Controllers and RAID

Each Network Attached Storage Gateway must provide capacity that is RAID protected by the base storage platform. The following RAID formats must be supported by the NAS Gateway when presented to it by the base storage platform defined at 1.1:

- (a) RAID5 stripes with parity;
- (b) RAID6 stripes with dual parity;
- (c) RAID1; and
- (d) RAID0+1 stripes with mirroring or RAID1+0 striped mirrors (aka RAID10).

3.4.4 NAS Processor Unit

Each Network Attached Storage Gateway must meet the following requirements for the NAS processor unit(s):

- (a) it must include a micro-kernel operating system designed for providing file services to CIFS and NFS via the included Ethernet interfaces. The micro-kernel operating system may be either a Linux or Unix-based operating system;
- (b) it must load the micro-kernel operating system from a fault-tolerant medium that is either RAID protected, or duplicated and included, in a second NAS processor unit that may assume operation in the event of a failure to load the operating system at boot time;
- (c) it must contain 2 separate redundant clustered processor units or "heads" that operate in an active / active fashion providing network services to clients for CIFS and NFS. In the event of a failure of one of the processor units, the remaining unit must assume the IP address and identity of the failed processor unit and must continue to provide service to clients on the network automatically;
- (d) the processor units must both be attached via a total aggregate minimum of 6 X 8Gbps fibre channel interfaces or 4 X 10Gbps Ethernet to the base storage platform; and
- (e) the processor units in the NAS Gateway must contain an aggregate minimum of 4 X 10Gbps Ethernet interfaces for TCP/IP client access;

3.4.5 Software and Additional Capabilities

Each Network Attached Storage Gateway must meet the following requirements for software functionality and additional capabilities:

- (a) it must include all client access licenses for end user workstations to access and use the shared file systems via CIFS or NFS, with no requirement for additional fees or licensing;
- (b) it must fully integrate, in mixed mode or native mode, Microsoft Active directory environments and must be manageable as a Windows server in those environments using native Microsoft tools for viewing and managing sessions, shares and open files;
- (c) it must support snapshot functionality for all shared file systems allowing an administrator to create point-in-time copies of all files for the purpose of recovering deleted files; and
- (d) it must include and be licensed for NDMP to facilitate backups of the shared file systems to fibre channel attached backup targets.

3.4.6 Management

Each Network Attached Storage Gateway must meet the following requirements for management capabilities:

- (a) it must be manageable remotely via an Ethernet interface and must provide an intuitive GUI-based interface for day-to-day operations;
- (b) it must include a simple and intuitive installation system allowing non-technical operators to easily configure and provision the unit for operation on a network with only a basic knowledge of TCP/IP addresses and volume and file system management;
- (c) it must provide GUI-based functionality to:
 - i. create and manage volumes and file systems across RAID sets;
 - ii. assign and manage user permissions for CIFS and NFS users to volumes and files;
 - iii. view attributes of volumes including space usage and file information;
 - iv. configure all user-assigned parameters required for operation of the system;
 - v. monitor utilization of network interfaces, processors and disk subsystems to gauge the load on those items;
 - vi. backup all locally hosted data to a locally-attached tape drive or provide an agent or facility for a remote console to initiate this process directly from the NAS disk to a backup target; and
 - vii. load balance file shares across either of the 2 processor units as needed and allow an administrator to manually failover file shares if required from 1 processor unit to the other.
- (d) the GUI management system must manage and operate both processor units as a single entity, allowing a single session to facilitate all management functions described here.

SAN 4.0 CATEGORY

The following describes the configuration and features of an **high-end Enterprise SAN** storage solution.

4.1 Storage Platform

4.1.1 Capacity and Platform

Each storage platform must meet the following capacity and platform requirements:

- (a) the hard disk drive technologies and densities must be commercially available, meaning that the Manufacturer is continuing to manufacture and ship them to customers generally;
- (b) the hard disk drive technologies and densities must be tested and fully supported within the storage platform by the storage platform Manufacturer;

-
- (c) it must include industry-standard hard disk drives operating at either 4Gbps for Fibre Channel (FC) drives or 6Gbps for Serial Attach SCSI (SAS) drives;
 - (d) it must also include industry-standard Serial Advanced Technology Attachment (SATA) revision 3.0 or Nearline SAS (NL-SAS) hard disk drives operating at 6Gbps. This may be achieved either by:
 - i. using the same shelves as either the FC or SAS disk drives, or
 - ii. using specialized shelves for these drive types;
 - (e) the available drive options must include at least seven (7) from the following list:

- drives with 4Gbps (for FC) or 6 Gbps (for SAS) interfaces and 15000 RPM rotational speed:

- i) 300GB
- ii) 450GB
- iii) 600GB

- drives with 4Gbps (for FC) or 6 Gbps (for SAS) interfaces and 10000 RPM rotational speed:

- iv) 300GB
- v) 450GB
- vi) 600GB
- vii) 900GB

- drives with 6Gbps for SATA or NL-SAS interfaces and 7200 RPM rotational speed:

- viii) 1TB
- ix) 2TB
- x) 3TB

- solid state drives (SSD) based on Single Level Cell (SLC) or enterprise-class Multi-Level Cell (eMLC) technology

- xi) 100GB
- xii) 200GB
- xiii) 300GB
- xiv) 400GB
- xv) 600GB

- (f) it must accommodate a minimum of 1024 hard disk drives;
- (g) It must be packaged in a standard 19" rack mount form factor (NOTE: it is understood that standard rack depth will be increased when "high density" disk shelves are provided); and
- (h) It must include lights or an LCD panel for power, activity and fault indications.

4.1.2 Cooling

Each storage platform must meet the following cooling requirements:

- (a) it must provide sufficient cooling for a fully populated cabinet at the mandatory minimum storage capacity;
- (b) all cooling for the system controller(s) as well as all hard disk drives must be redundant and monitored for failure by the storage platform hardware;
- (c) it must allow hot swapping of failed cooling fans;
- (d) the cooling system within the storage platform itself must be fully redundant; and
- (e) in the event of a component failure, the cooling system must allow continued operation of the storage platform until service can be performed.

4.1.3 Drives and Shelves

Each storage platform must meet the following drives and shelves requirements:

- (a) the hard disk drives must include a minimum of either 4Gbps for Fibre Channel (FC) or 6Gbps for Serial Attach SCSI (SAS) with dual attached interfaces;
- (b) it must provide minimum of 32 active connections to the mandatory 1024 hard disk drives. This bandwidth must be allocated evenly to the total number of physical drives over several channels;
- (c) it must provide fully redundant back-end paths to all hard disk drives. A channel failure must not interrupt access to attached disk drives;
- (d) it must allow hot addition of storage shelves without needing to power the storage platform down and without interrupting access to existing drives and RAID groups;
- (e) it must include as many back-end channels as necessary to support all the back-end shelves of disks so that a shelf component replacement or failure does not interrupt access to adjacent shelves in the platform;
- (f) the hard disk drives in the storage platform must be fully hot pluggable while the storage platform is operational. There must be no loss of data if a hard drive is removed, assuming the drive is part of a fault-tolerant configuration in the platform;
- (g) it must rebuild a replaced hard disk drive automatically and without user intervention when it is inserted, assuming it is replacing a hard disk drive that was part of a fault-tolerant configuration; and
- (h) it must allow the allocation of hard disk drives as global hot spares, which must automatically rebuild the contents of a failed hard disk drive in any fault-tolerant RAID set. This process must be fully automatic whenever a disk failure occurs in a fault-tolerant RAID set.

4.1.4 Power

Each storage platform must meet the following power requirements:

- (a) it must provide sufficient power to operate a fully populated system with all boards and cache installed, and the maximum number of hard disk drives installed;
- (b) the power supplies must be fully redundant, allowing uninterrupted operation of the storage platform in the event of a power supply failure, until service can be performed. Redundancy may be achieved either through:
 - i. use of a second power supply, or
 - ii. through an N+1 approach; and
- (c) each AC power supply must connect independently to a discrete AC power source.

4.1.5 Controllers

Each storage platform must meet the following controller requirements:

- (a) it must provide a multiprocessor architecture that provides multiple processors for handling both I/O to the attached host systems as well as disk I/O and RAID functionality. A minimum of 32 processors must be available within the storage platform to handle both front and back end processing;
- (b) it must use separate controllers for front end SAN connectivity versus back end connectivity and must support full redundancy and hot swapping of all controllers;
- (c) it must provide redundancy, so that processor subsystem failures can be recovered by surviving processors which can continue to service the attached hosts without disruption;
- (d) it must have access to all 1024 of the mandatory hard disk drives in order to assign, configure, protect and share those drives. Implementations involving small discrete stacked storage units with individual controllers for smaller groups of drives will not be considered compliant with this requirement;
- (e) the storage controllers must allow configuration of hard disk drives within the storage platform as either:
 - i. RAID5 stripes with parity;
 - ii. RAID6 stripes with dual parity;

- iii. RAID1; and
- iv. RAID0+1 stripes with mirroring or RAID1+0 striped mirrors (aka RAID10).
- (f) it must simultaneously support all RAID types from 4.1.5(e) within the storage platform; and
- (g) it must allow the creation and allocation of a minimum of 16000 logical units to connected hosts.

4.1.6 Cache

Each storage platform must meet the following cache requirements:

- (a) it must include at least 384GB of dedicated I/O cache that may be shared between all storage processors. It is understood and accepted that a small portion of this memory is used for storing platform specific software as required;
- (b) it must perform both read and write I/O operations;
- (c) the write cache must be mirrored and must utilize error detection and correction logic that will detect and recover from any memory errors without data loss or interruption to service;
- (d) it must be serviceable without disruption to the operation of the storage platform so that failed portions of the cache may be replaced hot;
- (e) the write data within the cache on the storage controllers must be protected by either one of these two methods:
 - i. a battery that allows the cache contents to be held intact for a minimum of 48 hours. The caches must then complete their write operations to disk when power is restored; or
 - ii. all pending write data must be automatically written to disk before the disk system is powered off, and the platform must provide sufficient battery power to complete this function.

4.1.7 I/O Ports and Connectivity

Each storage platform must meet the following requirements for I/O ports and connectivity:

- (a) it must utilize a system of installable cards and slots for populating the system with the customer desired combination of fibre channel, FCoE and FICON ports;
- (b) it must provide a minimum of 32 fibre channel ports for connectivity to Intel and Open System host computers;
- (c) it must provide a minimum of 8 x 10Gb FCoE ports for connectivity to Intel and Open System host computers;
- (d) it must support WAN connectivity for the purposes of data mirroring to another like storage platform at a separate physical location through all of the following:
 - i. Extended SAN fibre channel,
 - ii. Fibre channel over IP (FCIP), and
 - iii. FCoE;
- (e) it must be able to provide a minimum of 32 FICON ports for mainframe connectivity;
- (f) each of the 32 fibre channel ports must meet the following requirements:
 - i. each must operate at a minimum of 8Gbps;
 - ii. each must be compliant with ANSI T11 standards for fibre channel;
 - iii. each must support full fabric login;
 - iv. each must have its own unique fibre channel world wide name (WWN);
 - v. each must act as independent ports providing aggregate separate bandwidth to host computers;
 - vi. each must have the ability to be configured for active/active failover with appropriate host based load balancing and failover software;
- (g) it must provide simultaneous connectivity to 512 or more Intel hosts using dual fibre channel host bus adapters in each host;
- (h) it must provide host bus adapters including any required software drivers for all supported operating systems;

-
- (i) it must offer “no single point of failure” connectivity options, for both failover as well as load balancing under all of the mandatory operating system environments. This may be provided using add-on failover software packages or using native Operating System facilities, such as HP-UX PVLlinks.
 - (j) it must provide no fewer than eight (8) 10GbE ports for FCoE host connectivity that meet the ANSI T11 FC-BB-5 Fibre Channel Over Ethernet (FCoE) standards for the encapsulation of FC packets over Full Duplex and Lossless Ethernet networks and be compliant with the following IEEE standards:
 - i. 802.1Ae;
 - ii. 802.1Qbb;
 - iii. 802.1Qaz which defines:
 - a. enhanced transmission selection (ETS); and
 - b. data center bridging exchange (DCBX).

4.1.8 Hosts

Each storage platform must meet the following requirements for host connectivity:

- (a) it must connect to Intel-based host computers running the following:
 - i. Windows 2008 R2 with Hyper-V in 32-bit and 64-bit configurations;
 - ii. Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11 in 32-bit and 64-bit configurations; and
 - iii. VMWare ESX Server 3.5X and 4.X;
- (b) it must connect to the following UNIX and Open Systems hosts simultaneously, in addition to the previously listed Intel systems:
 - i. Oracle Solaris 10 and 11 systems;
 - ii. HP-UX 11i v.X systems;
 - iii. IBM AIX v.6X and v.7X systems;
- (c) it must connect to IBM AS400 systems;
- (d) it must support FICON mainframe connectivity and must emulate 3390-3, 3390-9, 3390-27 and 3390-54 modes;
- (e) support of additional platform types and operating systems is desirable but not mandatory.

4.1.9 Clustering

Each storage platform must meet the following requirements for clustering:

- (a) it must directly support clustering under the following host operating environments:
 - i. Windows 2008 R2 in 32-bit and 64-bit configurations;
 - ii. Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11 in 32-bit and 64-bit configurations; and
 - iii. VMWare ESX Server 3.5X and 4.X with shared access to the same logical unit numbers (LUNs) for Vmotion; and
- (b) it must directly support clustering under the following host operating environments:
 - i. MC/Serviceguard for HP-UX;
 - ii. PowerHA for AIX; and
 - iii. Oracle Solaris Cluster for Solaris

4.1.10 Software and Additional Capabilities

The storage platform must provide the following software functionalities and additional capabilities. Furthermore, these must be entirely storage platform-based functionality and must not require any software or assistance from host systems on the SAN:

- (a) it must provide LUN-masking functionality, meaning it masks or limits visibility of logical drive configurations within the storage platform to only specific hosts connected to the storage platform;

- (b) it must synchronously replicate logical volumes remotely via extended fibre channel networks over dark fibre or DWDM/CWDM.;
- (c) it must asynchronously replicate logical volumes remotely via extended fibre channel networks over dark fibre, DWDM/CWDM, or TCP/IP.;
- (d) it must perform up to 8 concurrent host-less point-in-time snapshot copies of any logical volume that may be reassigned to any other host on the SAN.;
- (e) it must perform up to 8 concurrent host-less full block data copies of any logical volume that may be reassigned to any other host on the SAN.;
- (f) it must allow online firmware upgrades to be made without disrupting the operation of the platform.

4.1.11 Management

The storage platform must provide the following management capabilities:

- (a) it must provide a comprehensive GUI-based management system that allows real-time monitoring of all components in the platform and reports degradation of components and failures;
- (b) several management packages may be required in order to provide the required functionality. This is acceptable provided that these packages may all be executed from one dedicated management console system and that the packages coexist and function properly together;
- (c) the GUI interface must be a Windows-based application included with the system or may be a WEB or Java-based embedded function accessible using a standard browser;
- (d) it must connect to an IP-based network either through a direct Ethernet connection on the platform or through an in-band connection via a fibre-attached host;
- (e) it must provide the ability to issue SNMP traps or SMTP mail in the event of a storage platform device degradation or failure;
- (f) the GUI interface must offer complete visibility to all installed hardware and its current operational status;
- (g) the GUI interface must perform full performance monitoring of the storage array including:
 - i. disk, LUN or RAID group I/O's per second for both read and write requests;
 - ii. cache utilization and hit rate statistics;
 - iii. queuing information for disks, LUNs or RAID sets; and
 - iv. I/O throughput statistics on a per interface basis for fibre channel, FICON, ESCON and remote link connections to remote replication storage platforms.
- (h) the GUI interface must also provide the following:
 - i. it must explicitly identify bottlenecks allowing a storage administrator to take corrective action;
 - ii. it must allow configuration of all aspects of the storage platform including the controllers, cache, interfaces, drives and RAID configurations as well as logical drives and associated permissions;
 - iii. it must be able to update all firmware and software resident on the storage platform as an integrated function and must be able to activate this new software non-disruptively for all dual attached host platforms;
 - iv. it must have integrated control for the mandatory snapshot and remote replication features of the storage platform allowing the creation, assignment, configuration and destruction of snapshots and remote LUN replicas;
 - v. it must have a facility for graphically displaying the SAN and FICON connections from attached hosts to their target volumes to clearly illustrate to a storage administrator the relationship between hosts and LUNs;
 - vi. it must have provisions for the creation, configuration, allocation and management of 3390 mainframe type volumes for mainframes;
 - vii. it must manage both LUN allocation to hosts as well as switch the zone and path from a single interface for the storage platform and all fibre channel switches. This may be supplied as a separate product from the management tool both must be compatible and have launch support for the storage platform management tool; and

- viii. it must feature full SAN topology mapping showing all attached systems, switches and fibre connected hosts as well as physical and logical path information including fibre channel zoning;
- (i) it must have an array-based LUN-masking feature that allows assigning explicit permissions between specific logical drives and specified SAN attached hosts. This must be configured and enforced at the storage array; and
- (j) it must include the array-based ability to expand RAID sets or logical volumes presented to hosts. This must be available through the GUI interface as part of the RAID and logical volume management capabilities. Please note that this is referring to the ability to expand the size of RAID sets and LUNs at the hardware level and does not require a utility on the host to manipulate this space at the operating system level.

4.2 Fabric

The following describes the configuration and features of fibre channel switches.

4.2.1 Fibre Channel Switch

The storage platform must operate with 8Gbps 32 and 64 port fibre channel fabric switches, which must be fully supported and warranted by the storage platform Manufacturer. The fibre channel switches must provide the following capabilities:

- (a) they must operate with fibre channel fabrics and must be capable of full fibre channel zoning across switched fabrics;
- (b) they must support a minimum of 512 active enabled unique zones at a time per fibre channel fabric;
- (c) they must be available in both stand-alone and rack mountable configurations. A rack mounting kit that is applied to a stand-alone switch will be accepted;
- (d) they must operate at 8GB/s and must be fully populated with small form factor pluggable optical media modules for shortwave operation;
- (e) they must provide lights or indicators for power and port status for all fibre channel ports;
- (f) they must provide a 10/100/1000 Ethernet interface and must be manageable using TCP/IP as the transport protocol;
- (g) they must provide redundant cooling and power;
- (h) they must fully comply with the following ANSI T-11 standards:
 - i. FC-PH,
 - ii. FC-PH-2,
 - iii. FC-PH-3,
 - iv. FC-AL,
 - v. FC-AL-2,
 - vi. FC-FLA,
 - vii. FC-FG,
 - viii. FC-GS-2,
 - ix. FC-PLDA,
 - x. FC-VI, and
 - xi. FC-SW-2;
- (i) they must support fibre channel class 2 and 3 connections;
- (j) they must provide full fabric support as per the ANSI standards specified at (4.2.1 h);
- (k) they must support cascading by connecting 2 or more switches together to form a single fabric that is compliant with the ANSI standards specified at (4.2.1 h);
- (l) they must include a comprehensive GUI-based management system that allows real-time monitoring of all components in the platform and to report failures or degraded components;
- (m) the GUI interface must be an embedded WEB or Java-based function accessible using a standard browser;
- (n) they must generate SNMP traps in the event of a degraded condition in the switch;
- (o) the GUI interface must show the current operational status for all installed hardware components;
- (p) the GUI interface must allow configuration of all aspects of the fibre channel switches including:

-
- i. the name,
 - ii. the domain ID,
 - iii. the passwords and user accounts for management,
 - iv. the IP addressing,
 - v. the modes of operation of the ports,
 - vi. all zone and path information, and
 - vii. any other parameters critical to the operation of the switch; and
 - (q) the GUI interface must provide complete performance monitoring allowing a storage administrator to view:
 - i. the number of frames per second, with a breakdown of which were good frames and which were error frames,
 - ii. the % utilization of fibre channel port,
 - iii. the operational speed of fibre channel ports,
 - iv. the mode of operation of fibre channel port (e.g. F-port, N-port, E-port), and
 - v. the throughput in frames as well as MB per second.

4.2.2 Director Class Fibre Channel Switches

The storage platform must operate with a 256 port director class fibre channel fabric switch, which must be fully supported and warranted by the storage platform Manufacturer. The director class fibre channel switches must provide the following capabilities:

- (a) the full 256 ports must be connected through a non-blocking backplane architecture;
- (b) they must allow all ports to be simultaneously active and send data without traversing any hops or Inter Switch Links, either obvious or embedded;
- (c) they must operate with fibre channel fabrics and must be capable of full fibre channel zoning across switched fabrics;
- (d) they must support a minimum of 1024 active enabled unique zones at a time per fibre channel fabric;
- (e) they must be provided in a rack mountable configuration;
- (f) they must operate at 8GB/s and must be fully populated with small form factor pluggable optical media modules for shortwave operation;
- (g) they must support optional long wave small form factor pluggable fibre channel optical media modules or blades, with these modules preinstalled for creating long distance connections to a minimum of 30KM without repeaters or extenders;
- (h) they must provide a minimum of four 10GbE interfaces which meet the ANSI T11 FC-BB-5 Fibre Channel Over Ethernet (FCoE) standards for the encapsulation of FC packets over Full Duplex and Lossless Ethernet networks and must comply with the following IEEE standards:
 - i. 802.1Ae;
 - ii. 802.1Qbb;
 - iii. 802.1Qaz which defines:
 - a. enhanced transmission selection (ETS); and
 - b. data center bridging exchange (DCBX).
- (i) they must provide lights or indicators for power and port status for all fibre channel ports;
- (j) they must provide a 10/100/1000 Ethernet interface and must be manageable using TCP/IP as the transport protocol
- (k) they must provide the following redundant components:
 - i. cooling and power,
 - ii. memory and processors,
 - iii. fibre ports and associated port circuitry connected into the backplane;
- (l) they must fully comply with the following ANSI T-11 standards:
 - i. FC-BB-5
 - ii. FC-PH,
 - iii. FC-PH-2,
 - iv. FC-PH-3,
 - v. FC-AL,

- vi. FC-AL-2,
- vii. FC-FLA,
- viii. FC-FG,
- ix. FC-GS-2,
- x. FC-PLDA,
- xi. FC-VI, and
- xii. FC-SW-2;

- (m) they must support fibre channel class 2 and 3 connections;
- (n) they must provide full fabric support as per the ANSI standards specified at (4.2.1 h);
- (o) they must support cascading by connecting 16 or more switches together to form a single fabric that is compliant with the ANSI standards specified at (4.2.1 h);
- (p) they must include a comprehensive GUI-based management system that allows real-time monitoring of all components in the platform and to report failures or degraded components;
- (q) the GUI interface must be an embedded WEB or Java-based function accessible using a standard browser;
- (r) they must provide full failure monitoring for all components and must be thermally monitored;
- (s) they must provide alerting via SNMP and the GUI console to advise a storage administrator of a failure or degraded condition;
- (t) the GUI interface must show the current operational status for all installed hardware components;
- (u) the GUI interface must allow configuration of all aspects of the fibre channel switches including:
 - i. the name,
 - ii. the domain ID,
 - iii. the passwords and user accounts for management,
 - iv. the IP addressing,
 - v. the modes of operation of the ports,
 - vi. all zone and path information, and
 - vii. any other parameters critical to the operation of the switch.
- (v) the GUI interface must provide complete performance monitoring allowing a storage administrator to view:
 - i. the number of frames per second, with a breakdown of which were good frames and which were error frames,
 - ii. the % utilization of fibre channel port,
 - iii. the operational speed of fibre channel ports,
 - iv. the mode of operation of fibre channel port (e.g. F-port, N-port, E-port), and
 - v. the throughput in frames as well as MB per second.
- (w) they must accept a new firmware or microcode upgrade non-disruptively.

4.3 Virtualization

4.3.1 Virtualization

The storage platform must include a virtualization solution that is manufactured or OEM rebranded, warranted, supported and maintained by the Manufacturer of the base storage platform defined at 4.1. The virtualization solution must provide the following capabilities:

- (a) it must be a discrete and independent device(s) that does not rely upon any components, functionality or software from the base storage platform defined at 4.1;
- (b) it must be packaged in an industry-standard 19" rack mount form factor and must include all accessories, cables and hardware required to mount and power the unit in an industry-standard 19" rack;
- (c) it must have redundant and hot swappable power and cooling in all electronic components of the solution. Fully redundant pairs of equipment with fixed cooling and power that allow swapping of an entire portion of the solution without interrupting host access to virtualized storage is acceptable;
- (d) it must include 32 X 8Gbps host side fibre channel ports for connections to a client host side SAN fabric;

-
- (e) it must include 32 X 8Gbps host side fibre channel ports for connections to the mandatory base storage platform as well as potential existing client third party storage platforms;
 - (f) it must support the base storage platform from the offeror and must simultaneously support third party mainstream storage platforms from at least 5 of the following Manufacturers:
 - i. Dell,
 - ii. EMC,
 - iii. Hitachi Data Systems,
 - iv. HP,
 - v. IBM,
 - vi. Network Appliance, and
 - vii. Oracle.
 - (g) it must either provide thin provisioning capabilities or allow virtual pools of storage to be created from one or more of the mandatory base platform(s) as well as any of the supported third party storage platforms simultaneously;
 - (h) it must present block storage from one or more base storage platforms as well as any of the supported additional third party storage platforms in the form of LUNs to the mandatory supported host platforms. The underlying details of the capacity from the base and third party systems must be masked from the host systems so that LUNs may be comprised of any or all of the underlying capacity;
 - (i) it must allow full block copies of LUNs to be created over local SAN connections between the supported base and third party systems and the dynamic relocation of LUNs without interrupting host access, without a loss of data and without changing the addressing of those LUNs to the hosts; and
 - (j) it must support the creation of synchronous and/or asynchronous copies of LUNs over distance-extended SAN connections between the supported base and third party systems for the purpose of disaster recovery between any of the supported storage platforms.

4.4 NAS Gateway

4.4.1 Capacity and Platform

The storage platform must include a Network Attached Storage Gateway that is manufactured or OEM rebranded, warranted, supported and maintained by the Manufacturer of the base storage platform defined at 4.1. The NAS Gateway must provide the following capabilities:

- (a) it must be a discrete and independent device(s) that does not rely upon any components, functionality or software from the base storage platform defined at 4.1; however, the capacity this NAS Gateway will address and share may be provided by the base storage platform defined at 4.1;
- (b) it must address and share a minimum of 128TB of raw data storage while also adhering to all other minimums;
- (c) it must be fully compatible with and supported by the base storage platform defined at 4.1. Use of this NAS Gateway with the base storage platform must not preclude the base storage platform from also servicing other fibre channel block attached hosts at the same time;
- (d) it must include sufficient cooling for a fully populated configuration. All cooling for the NAS Gateway must be redundant and monitored for failures by the NAS Gateway;
- (e) it must allow hot swapping of failed cooling fans; and
- (f) it must be packaged in an industry-standard 19" rack mount form factor and must include all accessories, cables and hardware required to mount and power the unit in an industry-standard 19" rack.

4.4.2 Power

Each Network Attached Storage Gateway must meet the following power requirements:

- (a) it must provide sufficient power to operate a fully loaded system with all boards and components installed;
- (b) it must be fully redundant allowing the NAS Gateway to continue operation without interruption in the event of a power supply failure, until service can be performed. Redundancy may be achieved either through:
 - i. use of a second power supply, or
 - ii. through an N+1 approach; and
- (c) each AC power supply must connect independently to a discrete AC power source.

4.4.3 Controllers and RAID

Each Network Attached Storage Gateway must meet the following requirements for controllers and RAID:

- (a) it must provide capacity that is RAID protected by the base storage platform. The following RAID formats must be supported by the NAS Gateway when presented to it by the base storage platform defined at 4.1:
 - i. RAID5 stripes with parity;
 - ii. RAID6 stripes with dual parity;
 - iii. RAID1; and
 - iv. RAID0+1 stripes with mirroring, or RAID1+0 striped mirrors (aka RAID10).

4.4.4 NAS Processor Unit

Each Network Attached Storage Gateway must meet the following requirements for the NAS processor unit(s):

- (a) it must include a micro-kernel operating system designed for providing file services to CIFS and NFS clients via the included Ethernet interfaces. The micro-kernel operating system may be either a Linux or Unix based operating system;
- (b) it must load the micro-kernel operating system from a fault-tolerant medium that is either RAID protected, or duplicated and included, in a second NAS processor unit that may assume operation in the event of a failure to load the operating system at boot time;
- (c) it must contain 2 separate redundant clustered processor units or "heads" that operate in an active / active fashion providing network services to clients for CIFS and NFS. In the event of a failure of one of the processor units, the remaining unit must assume the IP address and identity of the failed processor unit and must continue to provide service to clients on the network automatically;
- (d) the processor units must both be attached via a total aggregate minimum of 8 X 8Gbps fibre channel interfaces to the base storage platform; and
- (e) the processor units in the NAS Gateway must contain an aggregate minimum of 8 X 10Gbps Ethernet interfaces for TCP/IP client access.

4.4.5 Software and Additional Capabilities

Each Network Attached Storage Gateway must meet the following requirements for software functionality and additional capabilities:

- (a) it must include all client access licenses for end user workstations to access and use the shared file systems via CIFS or NFS with no requirement for additional fees or licensing;
- (b) it must fully integrate in mixed mode or native mode Microsoft Active directory environments and must be manageable as a Windows server in those environments using native Microsoft tools for viewing and managing sessions, shares and open files;
- (c) it must support snapshot functionality for all shared file systems allowing an administrator to create point-in-time copies of all files for the purpose of recovering deleted files; and
- (d) it must include and be licensed for NDMP to facilitate backups of the shared file systems to fibre channel attached backup targets; and

-
- (e) it must support asynchronous file level replication over TCP/IP to another NAS Gateway of the same type for data recovery or data distribution purposes.

4.4.6 Management

Each Network Attached Storage Gateway must meet the following requirements for management capabilities:

- (a) it must be manageable remotely via an Ethernet interface and must provide an intuitive GUI-based interface for day-to-day operations;
- (b) it must include a simple and intuitive installation system allowing non-technical operators to easily configure and provision the unit for operation on a network with only a basic knowledge of TCP/IP addresses and volume and file system management;
- (c) it must provide GUI-based functionality to perform the following:
 - i. create and manage volumes and file systems across RAID sets;
 - ii. assign and manage user permissions for CIFS and NFS users to volumes and files;
 - iii. view attributes of volumes including space usage and file information;
 - iv. configure all user-assigned parameters required for operation of the system;
 - v. monitor utilization of network interfaces, processors and disk subsystems to gauge the load on those items
 - vi. backup all locally hosted data to a locally-attached tape drive or provide an agent or facility for a remote console to initiate this process directly from the NAS disk to a backup target; and
 - vii. load balance file shares across either of the 2 processor units as needed and allow an administrator to manually failover file shares if required from 1 processor unit to the other; and
- (d) the GUI management system must manage and operate both processor units as a single entity allowing a single session to facilitate all management functions described here.

SAN 5.0 CATEGORY

The following describes the configuration and features of a “**Cloud Ready**” storage solution.

5.1 Storage Platform

5.1.1 Capacity and Platform

Each storage platform must meet the following capacity and platform requirements:

- (a) the proposed “cloud ready” storage platform must either:
 - i. be validated for the Brocade CloudPlex architecture;
 - ii. be incorporated into a Cisco Unified Computing System (UCS) design;
 - iii. be incorporated into a Cisco Validated Design (CVD) architecture;
 - iv. be incorporated into an HDS Unified Compute Platform (UCP) design;
 - v. be incorporated into an HP Reference Architecture for Cloud design ; or
 - vi. incorporate the IBM Active Cloud Engine into the design.
- (b) the proposed “cloud ready” storage platform must be currently deployed within a large enterprise as part of an established unified cloud computing blueprint;
- (c) the hard disk drive technologies and densities must be commercially available, meaning that the Manufacturer is continuing to manufacture and ship them to customers generally;
- (d) the hard disk drive technologies and densities must be tested and fully supported within the storage platform by the storage platform Manufacturer;
- (e) it must include industry-standard hard disk drives operating at either 4Gbps for Fibre Channel (FC) drives or 6Gbps for Serial Attach SCSI (SAS) drives;

- (f) it must also include industry-standard Serial Advanced Technology Attachment (SATA) revision 3.0 or Nearline SAS (NL-SAS) hard disk drives operating at 6Gbps. This may be achieved either by:
- i. using the same shelves as either the FC or SAS disk drives, or
 - ii. using specialized shelves for these drive types;
- (g) the available drive options must include at least seven (7) from the following list:

- drives with 4Gbps (for FC) or 6 Gbps (for SAS) interfaces and 15000 RPM rotational speed:

- i) 300GB
- ii) 450GB
- iii) 600GB

- drives with 4Gbps (for FC) or 6 Gbps (for SAS) interfaces and 10000 RPM rotational speed:

- iv) 300GB
- v) 450GB
- vi) 600GB
- vii) 900GB

- drives with 6Gbps for SATA or NL-SAS interfaces and 7200 RPM rotational speed:

- viii) 1TB
- ix) 2TB
- x) 3TB

- solid state drives (SSD) based on Single Level Cell (SLC) or enterprise-class Multi-Level Cell (eMLC) technology

- xi) 100GB
- xii) 200GB
- xiii) 300GB
- xiv) 400GB
- xv) 600GB

- (h) it must accommodate a minimum of 240 hard disk drives;
- (i) It must be packaged in a standard 19" rack mount form factor (NOTE: it is understood that standard rack depth will be increased when "high density" disk shelves are provided); and
- (j) It must include lights or an LCD panel for power, activity and fault indications.

5.1.2 Cooling

Each storage platform must meet the following cooling requirements:

- (a) it must provide sufficient cooling for a fully populated node or cabinet;
- (b) all cooling for the system controller(s) as well as all hard disk drives must be redundant and monitored for failure by the storage platform hardware;
- (c) it must allow hot swapping of failed cooling fans;
- (d) the cooling system within the storage platform itself must be fully redundant; and
- (e) in the event of a component failure, the cooling system must allow continued operation of the storage platform until service can be performed.

5.1.3 Drives and Shelves

Each storage platform must meet the following drives and shelves requirements:

-
- (a) the hard disk drives must include a minimum of either 4Gbps for Fibre Channel (FC) or 6Gbps for Serial Attach SCSI (SAS) with dual attached interfaces;
 - (b) when fully configured, it must provide an aggregate minimum of 4 active connections to the mandatory 240 hard disk drives. This bandwidth must be allocated evenly to the total number of physical drives over several channels;
 - (c) it must provide fully redundant back-end paths to all hard disk drives. A channel failure must not interrupt access to attached disk drives;
 - (d) it must allow hot addition of nodes or storage shelves without needing to power the storage platform down and without interrupting access to existing drives and RAID groups;
 - (e) it must utilize redundant hot-pluggable components so that a node or shelf component replacement or failure does not interrupt access to adjacent nodes or shelves in the platform;
 - (f) the hard disk drives in the storage platform must be fully hot pluggable while the storage platform is operational. There must be no loss of data if a hard drive is removed, assuming the drive is part of a fault-tolerant configuration in the platform;
 - (g) it must rebuild a replaced hard disk drive automatically and without user intervention when it is inserted, assuming it is replacing a hard disk drive that was part of a fault-tolerant configuration; and
 - (h) it must allow the allocation of hard disk drives as hot spares, which must automatically rebuild the contents of a failed hard disk drive in any fault-tolerant RAID set. This process must be fully automatic whenever a disk failure occurs in a fault-tolerant RAID set.

5.1.4 Power

Each storage platform must meet the following power requirements:

- (a) it must provide sufficient power to operate a fully populated system with all boards and cache installed, and the maximum number of hard disk drives installed;
- (b) the power supplies must be fully redundant, allowing uninterrupted operation of the storage platform in the event of a power supply failure, until service can be performed. Redundancy may be achieved either through:
 - i. use of a second power supply, or
 - ii. through an N+1 approach; and
- (c) each AC power supply must connect independently to a discrete AC power source.

5.1.5 Controllers

Each storage platform must meet the following controller requirements:

- (a) it must include dual redundant active/active storage controllers for handling both I/O to the attached host systems as well as disk I/O and RAID functionality;
- (b) it must be redundant, so that the surviving controller automatically recovers controller subsystem failures, and service to attached hosts is continued without disruption;
- (c) each storage controller within the storage platform must have access to all 240 of the mandatory hard disk drives in order to assign, configure, protect and share those drives;
- (d) the storage controllers must allow configuration of hard disk drives within the storage platform as:
 - i. RAID5 stripes with parity;
 - ii. RAID6 stripes with dual parity;
 - iii. RAID1; and
 - iv. RAID0+1 stripes with mirroring or RAID1+0 striped mirrors (aka RAID10).
- (e) it must allow the creation and addressing of up to 4096 simultaneous logical drives; and
- (f) it must simultaneously support all RAID types from 5.1.5(d) within the storage platform; and
- (g) it must enable sub-LUN auto-tiering when appropriate drive types are selected, for a minimum of two tiers.

5.1.6 Cache

Each storage platform must meet the following cache requirements:

-
- (a) it must include at least 32GB of I/O cache or flash cache that may be shared between all nodes or storage processors. It is understood and accepted that a small portion of this memory is used for storing platform specific software as required;
 - (b) it must perform both read and write I/O operations;
 - (c) the write cache must be mirrored and must utilize error detection and correction logic that will detect and recover from any memory errors without data loss or interruption to service;
 - (d) it must be serviceable without disruption to the operation of the storage platform so that failed portions of the cache may be replaced without interruption of service;
 - (e) the write data within the cache on the storage controllers must be protected by either one of these two methods:
 - i. a battery that allows the cache contents to be held intact for a minimum of 48 hours. The caches must then complete their write operations to disk when power is restored; or
 - ii. all pending write data must be automatically written to disk before the disk system is powered off, and the platform must provide sufficient battery power to complete this function.

5.1.7 I/O Ports and Connectivity

Each storage platform must meet the following requirements for I/O ports and connectivity:

- (a) it must include a minimum of 2 storage controllers that may be replaced in the event of a controller failure;
- (b) it must provide a minimum of 32 fibre channel ports for connectivity to Intel and Open System host computers;
- (c) it must contain a total of 8 X 10Gbps Ethernet interfaces which provide FCoE connectivity to Intel and Open System host computers;
- (d) it must meet the ANSI T11 FC-BB-5 Fibre Channel Over Ethernet (FCoE) standards for the encapsulation of FC packets over Full Duplex and Lossless Ethernet networks and be compliant with the following IEEE standards:
 - i. 802.1Ae;
 - ii. 802.1Qbb;
 - iii. 802.1Qaz which defines:
 - a. enhanced transmission selection (ETS); and
 - b. data center bridging exchange (DCBX).

5.1.8 Hosts

Each storage platform must meet the following requirements for host connectivity:

- (a) it must connect to Intel-based host computers running the following:
 - i. Windows 2008 R2 with Hyper-V in 32-bit and 64-bit configurations;
 - ii. Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11 in 32-bit and 64-bit configurations; and
 - iii. VMWare ESX 3.5X and 4.X;
- (b) it must connect to the following UNIX and Open Systems hosts simultaneously, in addition to the previously listed Intel systems:
 - i. Oracle Solaris 10 and 11 systems;
 - ii. HP-UX 11i v.X systems;
 - iii. IBM AIX v.6X and v.7X systems;
- (c) support of additional platform types and operating systems is desirable but not mandatory.

5.1.9 Software and Additional Capabilities

The storage platform must provide the following software functionalities and additional capabilities. Furthermore, these must be entirely storage platform-based functionality and must not require any software or assistance from host systems on the SAN:

- (a) it must perform up to 8 concurrent host-less point-in-time snapshot copies of any logical volume that may be reassigned to any other host on the SAN.;
- (b) it must perform up to 8 concurrent host-less full block data copies of any logical volume that may be reassigned to any other host on the SAN.; and
- (c) it must allow minor firmware version upgrades to be made online without disrupting the operation of the platform; and
- (d) it must provide secure multi-tenancy with provisions for chargeback.

5.1.11 Management

The storage platform must provide the following management capabilities:

- (a) it must provide a comprehensive graphical user interface (GUI) based management system that allows real-time monitoring of all components in the platform and reports degradation of components and failures;
- (b) the GUI interface must either be a Windows-based application included with the system or a WEB or Java-based embedded function accessible using a standard browser;
- (c) it must connect to an IP-based network either through a direct Ethernet connection on the platform or through an in-band connection via a fibre-attached host;
- (d) it must issue SNMP traps or SMTP mail in the event of device degradation or failure;
- (e) the GUI interface must show all installed hardware and its current operational status; and
- (f) the GUI interface must monitor the full performance of the storage array, including:
 - i. disk, LUN or RAID group I/O's per second for both read and write requests;
 - ii. cache utilization and hit rate statistics; and
 - iii. queuing information for disks, LUNs or RAID sets.

5.2 Fabric

The following describes the configuration and features of fibre channel switches which includes FCoE support.

5.2.1 Fibre Channel Switch

The storage platform must operate with 8Gbps 32 and 64 port fibre channel fabric switches, which must be fully supported and warranted by the storage platform Manufacturer. The fibre channel switches must provide the following capabilities:

- (a) they must operate with fibre channel fabrics and must be capable of full fibre channel zoning across switched fabrics;
- (b) they must support a minimum of 512 active enabled unique zones at a time per fibre channel fabric;
- (c) they must be available in both stand-alone and rack mountable configurations. A rack mounting kit that is applied to a stand-alone switch will be accepted;
- (d) they must operate at 8GB/s and must be fully populated with small form factor pluggable optical media modules for shortwave operation;
- (e) they must provide a minimum of two 10GbE interfaces which meet the ANSI T11 FC-BB-5 Fibre Channel Over Ethernet (FCoE) standards for the encapsulation of FC packets over Full Duplex and Lossless Ethernet networks and must comply with the following IEEE standards:
 - i. 802.1Ae;
 - ii. 802.1Qbb;
 - iii. 802.1Qaz which defines:
 - a. enhanced transmission selection (ETS); and
 - b. data center bridging exchange (DCBX).
- (f) they must provide lights or indicators for power and port status for all fibre channel ports;

-
- (g) they must provide a 10/100/1000 Ethernet interface and must be manageable using TCP/IP as the transport protocol;
 - (h) they must provide redundant cooling and power;
 - (i) they must fully comply with the following ANSI T-11 standards:
 - i. FC-BB-5
 - ii. FC-PH,
 - iii. FC-PH-2,
 - iv. FC-PH-3,
 - v. FC-AL,
 - vi. FC-AL-2,
 - vii. FC-FLA,
 - viii. FC-FG,
 - ix. FC-GS-2,
 - x. FC-PLDA,
 - xi. FC-VI, and
 - xii. FC-SW-2;
 - (j) they must support fibre channel class 2 and 3 connections;
 - (k) they must provide full fabric support as per the ANSI standards specified at (5.2.1 h);
 - (l) they must support cascading by connecting 2 or more switches together to form a single fabric that is compliant with the ANSI standards specified at (5.2.1 h);
 - (m) they must include a comprehensive GUI-based management system that allows real-time monitoring of all components in the platform and to report failures or degraded components;
 - (n) the GUI interface must be an embedded WEB or Java-based function accessible using a standard browser;
 - (o) they must generate SNMP traps in the event of a degraded condition in the switch;
 - (p) the GUI interface must show the current operational status for all installed hardware components;
 - (q) the GUI interface must allow configuration of all aspects of the fibre channel switches including:
 - i. the name,
 - ii. the domain ID,
 - iii. the passwords and user accounts for management,
 - iv. the IP addressing,
 - v. the modes of operation of the ports,
 - vi. all zone and path information, and
 - vii. any other parameters critical to the operation of the switch; and
 - (r) the GUI interface must provide complete performance monitoring allowing a storage administrator to view:
 - i. the number of frames per second, with a breakdown of which were good frames and which were error frames,
 - ii. the % utilization of fibre channel port,
 - iii. the operational speed of fibre channel ports,
 - iv. the mode of operation of fibre channel port (e.g. F-port, N-port, E-port), and
 - v. the throughput in frames as well as MB per second.

5.2.2 Director Class Fibre Channel Switches

The storage platform must operate with a 256 port director class fibre channel fabric switch, which must be fully supported and warranted by the storage platform Manufacturer. The director class fibre channel switches must provide the following capabilities:

- (a) the full 256 ports must be connected through a non-blocking backplane architecture;
- (b) they must allow all ports to be simultaneously active and send data without traversing any hops or Inter Switch Links, either obvious or embedded;
- (c) they must operate with fibre channel fabrics and must be capable of full fibre channel zoning across switched fabrics;
- (d) they must support a minimum of 1024 active enabled unique zones at a time per fibre channel fabric;
- (e) they must be provided in a rack mountable configuration;

-
- (f) they must operate at 8GB/s and must be fully populated with small form factor pluggable optical media modules for shortwave operation;
 - (g) they must support optional long wave small form factor pluggable fibre channel optical media modules or blades, with these modules preinstalled for creating long distance connections to a minimum of 30KM without repeaters or extenders;
 - (h) they must provide a minimum of four 10GbE interfaces which meet the ANSI T11 FC-BB-5 Fibre Channel Over Ethernet (FCoE) standards for the encapsulation of FC packets over Full Duplex and Lossless Ethernet networks and must comply with the following IEEE standards:
 - i. 802.1Ae;
 - ii. 802.1Qbb;
 - iii. 802.1Qaz which defines:
 - a. enhanced transmission selection (ETS); and
 - b. data center bridging exchange (DCBX).
 - iv. 802.1AX and Link Aggregation Control Protocol (LACP)
 - (i) they must provide lights or indicators for power and port status for all fibre channel ports;
 - (j) they must provide a 10/100/1000 Ethernet interface and must be manageable using TCP/IP as the transport protocol
 - (k) they must provide the following redundant components:
 - i. cooling and power,
 - ii. memory and processors,
 - iii. fibre ports and associated port circuitry connected into the backplane;
 - (l) they must fully comply with the following ANSI T-11 standards:
 - i. FC-BB-5
 - ii. FC-PH,
 - iii. FC-PH-2,
 - iv. FC-PH-3,
 - v. FC-AL,
 - vi. FC-AL-2,
 - vii. FC-FLA,
 - viii. FC-FG,
 - ix. FC-GS-2,
 - x. FC-PLDA,
 - xi. FC-VI, and
 - xii. FC-SW-2;
 - (m) they must support fibre channel class 2 and 3 connections;
 - (n) they must provide full fabric support as per the ANSI standards specified at (4.2.1 h);
 - (o) they must support cascading by connecting 16 or more switches together to form a single fabric that is compliant with the ANSI standards specified at (4.2.1 h);
 - (p) they must include a comprehensive GUI-based management system that allows real-time monitoring of all components in the platform and to report failures or degraded components;
 - (q) the GUI interface must be an embedded WEB or Java-based function accessible using a standard browser;
 - (r) they must provide full failure monitoring for all components and must be thermally monitored;
 - (s) they must provide alerting via SNMP and the GUI console to advise a storage administrator of a failure or degraded condition;
 - (t) the GUI interface must show the current operational status for all installed hardware components;
 - (u) the GUI interface must allow configuration of all aspects of the fibre channel switches including:
 - i. the name,
 - ii. the domain ID,
 - iii. the passwords and user accounts for management,
 - iv. the IP addressing,
 - v. the modes of operation of the ports,
 - vi. all zone and path information, and
 - vii. any other parameters critical to the operation of the switch.
 - (v) the GUI interface must provide complete performance monitoring allowing a storage administrator to view:

- i. the number of frames per second, with a breakdown of which were good frames and which were error frames,
 - ii. the % utilization of fibre channel port,
 - iii. the operational speed of fibre channel ports,
 - iv. the mode of operation of fibre channel port (e.g. F-port, N-port, E-port), and
 - v. the throughput in frames as well as MB per second.
- (w) they must accept a new firmware or microcode upgrade non-disruptively.

SAN 6.0 CATEGORY

The following describes the configuration and features of a “**Scale Out NAS**” storage solution.

6.1 Storage Platform

6.1.1 Capacity and Platform

Each Scale Out NAS platform must meet the following capacity and platform requirements:

- (a) the hard disk drive technologies and densities must be commercially available, meaning that the Manufacturer is continuing to manufacture and ship them to customers generally;
- (b) the hard disk drive technologies and densities must be tested and fully supported within the storage platform by the storage platform Manufacturer;
- (c) it must include industry-standard hard disk drives operating at either 4Gbps for Fibre Channel (FC) drives or 6Gbps for Serial Attach SCSI (SAS) drives;
- (d) it must also include industry-standard Serial Advanced Technology Attachment (SATA) revision 3.0 or Nearline SAS (NL-SAS) hard disk drives operating at 6Gbps. This may be achieved either by:
 - i) using the same shelves as either the FC or SAS disk drives, or
 - ii) using specialized shelves for these drive types;
- (e) the available drive options must include at least seven (7) from the following list:

- drives with 4Gbps (for FC) or 6 Gbps (for SAS) interfaces and 15000 RPM rotational speed:

- i) 300GB
- ii) 450GB
- iii) 600GB

- drives with 4Gbps (for FC) or 6 Gbps (for SAS) interfaces and 10000 RPM rotational speed:

- iv) 300GB
- v) 450GB
- vi) 600GB
- vii) 900GB

- drives with 6Gbps for SATA or NL-SAS interfaces and 7200 RPM rotational speed:

- viii) 1TB
- ix) 2TB
- x) 3TB

- solid state drives (SSD) based on Single Level Cell (SLC) or enterprise-class Multi-Level Cell (eMLC) technology

- xi) 100GB
- xii) 200GB
- xiii) 300GB

- xiv) 400GB
- xv) 600GB

- (f) it must accommodate a minimum of 1024 hard disk drives;
- (g) When fully configured, it must provide an aggregate minimum of 32 active connections to the mandatory 1024 hard disk drives. This bandwidth must be allocated evenly to the total number of physical drives over several channels;
- (h) it must provide fully redundant back-end paths to all hard disk drives. A channel failure must not interrupt access to attached disk drives;
- (i) it must allow hot addition of nodes or storage shelves without needing to power the storage platform down and without interrupting access to existing drives and RAID groups;
- (j) it must utilize redundant hot-pluggable components so that a node or shelf component replacement or failure does not interrupt access to adjacent nodes or shelves in the platform;
- (k) the hard disk drives in the storage platform must be fully hot pluggable while the storage platform is operational. There must be no loss of data if a hard drive is removed, assuming the drive is part of a fault-tolerant configuration in the platform;
- (l) it must rebuild a replaced hard disk drive automatically and without user intervention when it is inserted, assuming it is replacing a hard disk drive that was part of a fault-tolerant configuration; and
- (m) it must allow the allocation of hard disk drives as hot spares and or virtual spares, which must automatically rebuild the contents of a failed hard disk drive in any fault-tolerant RAID set. This process must be fully automatic whenever a disk failure occurs in a fault-tolerant RAID set.
- (n) it must include a minimum of 3 storage controllers / nodes that may be replaced in the event of a controller failure; and / or a minimum of 3 nodes
- (o) It must be packaged in a standard 19" rack mount form factor (NOTE: it is understood that standard rack depth will be increased when "high density" disk shelves are provided); and
- (p) It must include lights or an LCD panel for power, activity and fault indications.
- (q) It must scale to a minimum 1 petabyte (PB) single file system.

6.1.2 NAS Processor Unit

Each Scale Out NAS platform must meet the following requirements for the NAS processor unit(s):

- (a) it must include a micro-kernel operating system designed for providing file services to CIFS and NFS via the included Ethernet interfaces. The micro-kernel operating system may be either a Linux or Unix-based operating system;
- (b) it must load the micro-kernel operating system from a fault-tolerant medium that is either RAID protected, or duplicated and included, in a second NAS processor unit that may assume operation in the event of a failure to load the operating system at boot time;
- (c) it must contain a minimum 3 separate redundant clustered processor units or "heads" or "nodes" that operate in an active / active fashion providing network services to clients for CIFS and NFS. In the event of a failure of one of the processor units, the remaining unit must assume the IP address and identity of the failed processor unit and must continue to provide service to clients on the network automatically;
- (d) the NAS units must contain an aggregate minimum of either 12 X 1Gbps or 6 X 10Gbps Ethernet interfaces for TCP/IP client access.

6.1.3 Software and Additional Capabilities

Each Scale Out NAS platform must meet the following requirements for software functionality and additional capabilities:

- (a) it must include all client access licenses for end user workstations to access and use the shared file systems via CIFS or NFS, with no requirement for additional fees or licensing;

- (b) it must fully integrate, in mixed mode or native mode, Microsoft Active directory environments and must be manageable as a Windows server in those environments using native Microsoft tools for viewing and managing sessions, shares and open files;
- (c) it must support snapshot functionality for all shared file systems allowing an administrator to create point-in-time copies of all files for the purpose of recovering deleted files; and
- (d) it must include and be licensed for NDMP to facilitate backups of the shared file systems to attached backup targets.

6.1.4 Cooling

Each Scale Out NAS platform must meet the following cooling requirements:

- (a) it must provide sufficient cooling for a fully populated node or cabinet;
- (b) all cooling for the system controller(s) as well as all hard disk drives must be redundant and monitored for failure by the storage platform hardware;
- (c) it must allow hot swapping of failed cooling fans;
- (d) the cooling system within the storage platform itself must be fully redundant; and
- (e) in the event of a component failure, the cooling system must allow continued operation of the storage platform until service can be performed.

6.1.5 Power

Each Scale Out NAS platform must meet the following power requirements:

- (a) it must provide sufficient power to operate a fully populated system with all boards and cache installed, and the maximum number of hard disk drives installed;
- (b) the power supplies must be fully redundant, allowing uninterrupted operation of the storage platform in the event of a power supply failure, until service can be performed. Redundancy may be achieved either through:
 - i. use of a second power supply, or
 - ii. through an N+1 approach; and
- (c) each AC power supply must connect independently to a discrete AC power source.

6.1.6 Controllers

Each Scale Out NAS platform must meet the following controller requirements:

- (a) it must include redundant storage controllers / nodes for handling both I/O to the attached host systems as well as disk I/O and RAID functionality;
- (b) it must be redundant, so that the surviving controller / nodes automatically recovers controller subsystem failures, and service to attached hosts is continued without disruption;
- (c) the storage controllers / nodes must allow configuration of hard disk drives within the storage platform as:
 - i. RAID5 stripes with parity;
 - ii. RAID6 stripes with dual parity;
 - iii. RAID 1; and
 - iv. RAID0+1 stripes with mirroring, RAID1+0 striped mirrors (aka RAID10) Or equivalent at the Clustered Nodes
- (d) it must simultaneously support all RAID types from 6.1.6(c) within the storage platform or equivalent at the Clustered Node; and
- (e) it must operate with SSD drives and enable auto-tiering when appropriate drive types are selected, for a minimum of two tiers.

6.1.7 Cache

Each Scale Out NAS platform must meet the following cache requirements:

- (a) it must include at least 64GB of I/O cache or flash cache that may be shared between all nodes or storage processors. It is understood and accepted that a small portion of this memory is used for storing platform specific software as required;
- (b) it must perform both read and write I/O operations;
- (c) the write cache must be mirrored and must utilize error detection and correction logic that will detect and recover from any memory errors without data loss or interruption to service;
- (d) it must be serviceable without disruption to the operation of the storage platform so that failed portions of the cache may be replaced without interruption of service;
- (e) the write data within the cache on the storage controllers must be protected by either one of these two methods:
 - i. a battery that allows the cache contents to be held intact for a minimum of 48 hours. The caches must then complete their write operations to disk when power is restored; or
 - ii. all pending write data must be automatically written to disk before the disk system is powered off, and the platform must provide sufficient battery power to complete this function.

6.1.8 Software and Additional Capabilities

The Scale Out NAS platform must provide the following software functionality and additional capabilities:

- (a) it must perform up to 8 concurrent host-less point-in-time snapshot copies of any logical volume that may be reassigned to any other host on the SAN. This must be entirely storage platform-based functionality and must not require any software or assistance from host systems;
- (b) it must allow minor firmware version upgrades to be made online without disrupting the operation of the platform; and
- (c) it must include all client access licenses for end user workstations to access and use the shared file systems via CIFS or NFS, with no requirement for additional fees or licensing;
- (d) it must fully integrate, in mixed mode or native mode, Microsoft Active directory environments and must be manageable as a Windows server in those environments using native Microsoft tools for viewing and managing sessions, shares and open files;
- (e) it must support snapshot functionality for all shared file systems allowing an administrator to create point-in-time copies of all files for the purpose of recovering deleted files; and
- (f) it must include and be licensed for NDMP to facilitate backups of the shared file systems to backup targets.

6.1.9 Management

Each Scale Out NAS platform must meet the following requirements for management capabilities:

- (a) it must be manageable remotely via an Ethernet interface and must provide an intuitive GUI-based interface for day-to-day operations;
- (b) it must include a simple and intuitive installation system allowing non-technical operators to easily configure and provision the unit for operation on a network with only a basic knowledge of TCP/IP addresses and volume and file system management;
- (c) it must provide GUI-based functionality to:
 - i. create and manage volumes and file systems across RAID sets;
 - ii. assign and manage user permissions for CIFS and NFS users to volumes and files;
 - iii. view attributes of volumes including space usage and file information;
 - iv. configure all user-assigned parameters required for operation of the system;
 - v. monitor utilization of network interfaces, processors and disk subsystems to gauge the load on those items;
 - vi. backup all locally hosted data to a locally-attached tape drive or provide an agent or facility for a remote console to initiate this process directly from the NAS disk to a backup target; and
 - vii. load balance file shares across as needed and allow an administrator to manually failover file shares if required from 1 processor / node unit to the other.

Standing Offer No. - N° de l'offre

Amd. No. - N° de la modif.

Buyer ID - Id de l'acheteur

E60EJ-11000S/B

REVISION 1.1

461ej
