

Fiber Channel DAS Storage					
OEM & Model/Part Number:					
Reference Number	Requirement (Note: JBOD Configuration with Basic Level Raid Functionality. These are minimum requirements that a procurement shall meet, the exact numbers & quantities will be supplied on the individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Solution Shall Meet a minimum of the following IOPS (I/Os per second) 200 per drive.				
2	Solution Shall Meet a minimum of the following throughput (Mbits per second) 400 per drive.				
3	Solution Shall Meet the following minimum usable total disk space: 1 TB.				
4	Shall Provide dual ported drives with connections to drive enclosure(s).				
5	Solution Shall Provide the capability of supporting a JBOD configuration of the following Drive Type (Protocol/Interface): Fibre Channel.				
6	Overall Solution Shall Provide a multiple physical drive count that will meet requirements for IOPS, throughput and capacity (specified at Delivery Order Level).				
7	Solution Shall Provide the ability to allow non-disruptive updates or upgrades to drive firmware when used with an External Raid controller that supports this function.				
8	Solution Shall Provide ability to configure disk drives by an External Raid controller and meet IOPs, Throughput and Capacity requirements above.				

SAS DAS Storage					
OEM Model/Part Number:					
Reference Number	Requirement (Note: JBOD Configuration with Basic Level Raid Functionality. These are minimum requirements that a procurement shall meet, the exact numbers & quantities will be supplied on the individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Solution Shall Meet a minimum of the following IOPS (I/Os per second) 180 per drive (15K) and 130 per drive (10K).				
2	Solution Shall Meet a minimum of the following throughput (Mbits per second) 300 per drive.				
3	Solution Shall Meet the following minimum usable total disk space: 2 TB.				
4	Shall Provide dual ported drives with connections to drive enclosure(s).				
5	Solution Shall Provide drive(s) that support data transfer speeds up to 6 Gbit/s.				
6	Solution Shall Provide the capability of supporting a JBOD configuration of the following Drive Type (Protocol/Interface): SAS				
7	Overall Solution Shall Provide a multiple physical drive count that will meet requirements for IOPS, throughput and capacity (specified at Delivery Order Level).				
8	Solution Shall Provide the ability to allow non-disruptive updates or upgrades to drive firmware when used with an External Raid controller that supports this function.				
9	Solution Shall Provide ability to configure disk drives by an External Raid controller and meet IOPs, Throughput and Capacity requirements above.				

SATA DAS Storage					
OEM Model/Part Number:					
Reference Number	Requirement (Note: JBOD Configuration with Basic Level Raid Functionality. These are minimum requirements that a procurement shall meet, the exact numbers & quantities will be supplied on the individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Solution Shall Meet a minimum of following IOPS (I/Os per second) 90 per drive at 7200 RPM.				
2	Solution Shall Meet a minimum of the following throughput (Mbits per second) 200 per drive.				
3	Solution Shall Meet the following minimum usable total disk space: 2 TB.				
4	Solution Shall Provide drive(s) that support data transfer speeds up to 6 Gbit/s.				
5	Solution Shall Provide the capability of supporting a JBOD configuration of the following Drive Type (Protocol/Interface): Enterprise Level SATA (Nearline SAS or NL-SAS interface).				
6	Overall Solution Shall Provide a multiple physical drive count that will meet requirements for IOPS, throughput and capacity (specified at Delivery Order Level).				
7	Solution Shall Provide the ability to allow non-disruptive updates or upgrades to drive firmware when used with an External Raid controller that supports this function.				
8	Solution Shall Provide ability to configure disk drives by an External Raid controller and meet IOPs, Throughput and Capacity requirements above.				

SSD DAS Storage					
OEM Model/Part Number:					
Reference Number	Requirement (Note: JBOD Configuration with Basic Level Raid Functionality. These are minimum requirements that a procurement shall meet, the exact numbers & quantities will be supplied on the individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Solution Shall Meet a minimum of the following IOPS (I/Os per second) 4000 per Drive.				
2	Solution Shall Meet a minimum of the following throughput (Mbits per second): write=250/read=350. Vendor shall provide Performance numbers using 4K blocks, doing both Sequential and Random Reads and Writes.				
3	Solution Shall Provide the following, an equally balanced throughput for read and writes operations of a minimum of 250 Mb/s. Vendor shall provide Performance numbers using 4K blocks, doing both, Sequential and Random Reads and Writes.				
4	Solution Shall Provide the following minimum usable total disk space: 400 GB.				
5	Solution Shall Provide support for the following type of drive in a JBOD configuration: SSD.				
6	Solution Shall Provide the capability of supporting one of the following Drive Types in a JBOD configuration with parallel NAND flash chips: SSD (Flash-Memory MLC (Multi-Level Cell) or SLC (Single Level Cell) DRAM or RAM based).				
7	Flash-Memory MLC (Multi-Level Cell) will support minimum of 8,000 write operations.				
8	Flash-Memory SLC (Multi-Level Cell) will support minimum of 80,000 write operations.				
9	Overall Solution Shall Provide a multiple physical drive count that will meet requirements for IOPS, throughput and capacity (specified at Delivery Order Level).				
10	Solution Shall Provide the ability to allow non-disruptive updates or upgrades to drive firmware when used with an External Raid controller that supports this function.				
11	Solution Shall Provide ability to configure disk drives by an External Raid controller and meet IOPs, Throughput and Capacity requirements above.				

Low Performance SAN Storage					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 20,000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum aggregate data rate of 2000 MBytes per second for all read or write or combination of read/write operations through the storage subsystem.				
3	Shall provide the minimum usable native disk capacity of 25 TB without compression or deduplication.				
4	Shall be capable of supporting a minimum of two (2) 8-Gbit/s Fibre Channel ports per controller. All Fibre Channel ports shall be capable of autosensing slower speeds of 2 and 4 Gbit/s.				
5	Shall be capable of supporting a minimum of two (2) 1GbE iSCSI ports per controller.				
6	Shall be capable of supporting a minimum of one (1) 10GbE iSCSI ports per controller.				
7	Shall provide, excluding space on Storage System Operating System (OS) Drives, the initial data storage capacity must be distributed evenly among multiple shelves and support mirroring between shelves.				
8	Shall provide the disk failure redundancy configuration to meet the capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
9	Shall provide the disk drive technologies to meet the capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
10	Shall provide capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
11	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				

12	Shall provide a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).			
13	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.			
14	Shall provide a minimum of two storage controllers in a module for controller based storage solution, which are load balanced and provide automatic failover including the ability to maintain access to all data through a controller failure.			
15	Shall provide the ability to upgrade a controller on the storage system without replacing the storage subsystem and the upgrade of the controller shall be non-destructive to the data, requiring no data migration, no reconfiguration and no LUN remapping to any connecting Host.			
16	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.			
17	Shall provide the ability to expand and scale Host side connectivity and backend capacity, independently in order to increase performance.			
18	Shall provide storage subsystem support boot from SAN.			
19	Shall support the ability to manually change LUN UDID value.			
20	Shall support the ability to create LUNs and volumes across multiple spindles.			
21	Shall support non-disruptive LUN and volume expansion.			
22	Shall support ability to present a minimum of a 15TB LUN/volume to a host via iSCSI or FC.			
23	Shall provide the ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.			
24	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".			
25	Minimum shall be capable of supporting Fibre Channel (FC), Internet Small Computer System Interface (iSCSI) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.			

Attachment G - Storage Arrays / Storage Appliances Specifications

26	Shall provide the ability to expand Fibre Channel ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.			
27	Shall provide the ability to expand iSCSI ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.			
28	Shall be capable of supporting native 10GbE Fibre Channel over Ethernet (FCoE) host connections.			
29	Shall be capable of supporting 10GbE Ethernet ports			
30	Shall be capable of supporting Jumbo frame and LACP protocol.			
31	Solution shall support a minimum of 256 SCSI initiators (host connections).			
32	Shall provide load balance multi path failover software and multi mount points when using Ethernet.			
33	Shall provide the capability of synchronous and asynchronous replication with write order fidelity.			
34	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
35	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
36	Shall provide the ability to allow source update during data migration. Online data migration from other storage array and data migration is transparent to users.			
37	Shall be capable of maintaining redundancy and the performance metrics requirements as stated above during snapshot and clone operations.			
38	Power Switches shall be covered to prevent inadvertent activation.			
39	Shall provide role-based security or audit trail logging for access to storage.			
40	Shall provide a single master management interface to manage multiple storage subsystem of the same type in data center implementations.			
41	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.			
42	Shall have a graphical user interface (GUI) or command (CLI), or a Wizard to automate the process to provision a large number of Disk Groups and LUNs.			

Attachment G - Storage Arrays / Storage Appliances Specifications

43	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
44	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
45	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Medium Performance SAN Storage					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 50,000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum aggregate data rate of 3500 MBytes per second for all read or write or combination of read/write operations through the storage subsystem.				
3	Shall provide the minimum usable native disk capacity of 50 TB without compression or deduplication.				
4	Shall be capable of supporting a minimum of four (4) 8-Gbit/s Fibre Channel ports per controller. All Fibre Channel ports shall be capable of autosensing slower speeds of 2 and 4 Gbit/s.				
5	Shall be capable of supporting a minimum of four (4) 1GbE iSCSI ports per controller.				
6	Shall be capable of supporting a minimum of two (2) 10GbE iSCSI ports per controller.				
7	Shall provide, excluding space on Storage System Operating System (OS) Drives, the initial data storage capacity must be distributed evenly among multiple shelves and support mirroring between shelves.				
8	Shall provide the disk failure redundancy configuration to meet the capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
9	Shall provide the disk drive technologies to meet the capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
10	Shall provide capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
11	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				

12	Shall provide a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).			
13	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.			
14	Shall provide a minimum of two storage controllers in a module for controller based storage solution, which are load balanced and provide automatic failover including the ability to maintain access to all data through a controller failure.			
15	Shall provide the ability to upgrade a controller on the storage system without replacing the storage subsystem and the upgrade of the controller shall be non-destructive to the data, requiring no data migration, no reconfiguration and no LUN remapping to any connecting Host.			
16	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.			
17	Shall provide the ability to expand and scale Host side connectivity and backend capacity, independently in order to increase performance.			
18	Shall provide storage subsystem support boot from SAN.			
19	Shall support the ability to manually change LUN UDID value.			
20	Shall support the ability to create LUNs and volumes across multiple spindles.			
21	Shall support non-disruptive LUN and volume expansion.			
22	Shall support ability to present a minimum of a 15TB LUN/volume to a host via iSCSI or FC.			
23	Shall provide the ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.			
24	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".			
25	Minimum shall be capable of supporting Fibre Channel (FC), Internet Small Computer System Interface (iSCSI) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.			

Attachment G - Storage Arrays / Storage Appliances Specifications

26	Shall provide the ability to expand Fibre Channel ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
27	Shall provide the ability to expand iSCSI ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
28	Shall be capable of supporting native 10GbE Fibre Channel over Ethernet (FCoE) host connections.				
29	Shall be capable of supporting 10GbE Ethernet ports.				
30	Shall be capable of supporting Jumbo frame and LACP protocol.				
31	Solution shall support a minimum of 256 SCSI initiators (host connections).				
32	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
33	Shall provide the capability of synchronous and asynchronous replication with write order fidelity.				
34	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
35	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
36	Shall provide the ability to allow source update during data migration. Online data migration from other storage array and data migration is transparent to users.				
37	Shall be capable of maintaining redundancy and the performance metrics requirements as stated above during snapshot and clone operations.				
38	Power Switches shall be covered to prevent inadvertent activation.				
39	Shall provide role-based security or audit trail logging for access to storage.				
40	Shall provide a single master management interface to manage multiple storage subsystem of the same type in data center implementations.				
41	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
42	Shall have a graphical user interface (GUI) or command (CLI), or a Wizard to automate the process to provision a large number of Disk Groups and LUNs.				

Attachment G - Storage Arrays / Storage Appliances Specifications

43	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
44	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
45	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

High Performance SAN Storage					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 100,000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum aggregate data rate of 7000 MBytes per second for all read or write or combination of read/write operations through the storage subsystem.				
3	Shall provide the minimum usable native disk capacity of 75 TB without compression or deduplication.				
4	Shall be capable of supporting a minimum of four (4) 8-Gbit/s Fibre Channel ports per controller. All Fibre Channel ports shall be capable of autosensing slower speeds of 2 and 4 Gbit/s.				
5	Shall be capable of supporting a minimum of four (4) 1GbE iSCSI ports per controller.				
6	Shall be capable of supporting a minimum of two (2) 10GbE iSCSI ports per controller.				
7	Shall provide in the case of a Frame storage solution, having a minimum of two storage backend ports/disk adapters in a module or a card that are load balanced and provide automatic failover between them.				
8	Shall provide capability of isolating subsystem resources (Ports, Host I/F and Disk) and dedicate a defined amount of resources to a particular host.				
9	Shall provide the native capability to monitor or maintain the IOPs requirements on individual applications running on the same storage subsystem. The monitoring is to ensure the performances of existing applications are not affected when new applications are added to the storage subsystem.				
10	Shall provide, excluding space on Storage System Operating System (OS) Drives, the initial data storage capacity must be distributed evenly among multiple shelves and support mirroring between shelves.				
11	Shall provide the disk failure redundancy configuration to meet the capacity, IOPs and workload performance requirements (to be specified at delivery order level).				

Attachment G - Storage Arrays / Storage Appliances Specifications

12	Shall provide the disk drive technologies to meet the capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
13	Shall provide capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
14	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				
15	Shall provide a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				
16	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
17	Shall provide a minimum of two storage controllers in a module for controller based storage solution, which are load balanced and provide automatic failover including the ability to maintain access to all data through a controller failure.				
18	Shall provide the ability to upgrade a controller on the storage system without replacing the storage subsystem and the upgrade of the controller shall be non-destructive to the data, requiring no data migration, no reconfiguration and no LUN remapping to any connecting Host.				
19	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
20	Shall provide the ability to expand and scale Host side connectivity and backend capacity, independently in order to increase performance.				
21	Shall provide storage subsystem support boot from SAN.				
22	Shall support the ability to manually change LUN UDID value.				
23	Shall support the ability to create LUNs and volumes across multiple spindles.				
24	Shall support non-disruptive LUN and volume expansion.				
25	Shall support ability to present a minimum of a 15TB LUN/volume to a host via iSCSI or FC.				

26	Shall provide the ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.			
27	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".			
28	Minimum shall be capable of supporting Fibre Channel (FC), Internet Small Computer System Interface (iSCSI) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.			
29	Shall provide the ability to expand Fibre Channel ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.			
30	Shall provide the ability to expand iSCSI ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.			
31	Shall be capable of supporting native 10GbE Fibre Channel over Ethernet (FCoE) host connections.			
32	Shall be capable of supporting 10GbE Ethernet ports.			
33	Shall be capable of supporting Jumbo frame and LACP protocol.			
34	Solution shall support a minimum of 256 SCSI initiators (host connections).			
35	Shall provide load balance multi path failover software and multi mount points when using Ethernet.			
36	Shall provide the capability of synchronous and asynchronous replication with write order fidelity.			
37	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
38	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
39	Shall provide the ability to allow source update during data migration. Online data migration from other storage array and data migration is transparent to users.			
40	Shall be capable of maintaining redundancy and the performance metrics requirements as stated above during snapshot and clone operations.			
41	Power Switches shall be covered to prevent inadvertent activation.			
42	Shall provide role-based security or audit trail logging for access to storage.			

Attachment G - Storage Arrays / Storage Appliances Specifications

43	Shall provide a single master management interface to manage multiple storage subsystem of the same type in data center implementations.				
44	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
45	Shall have a graphical user interface (GUI) or command (CLI), or a Wizard to automate the process to provision a large number of Disk Groups and LUNs.				
46	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
47	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
48	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Fibre Fabric SAN Switch - Small to Medium					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide support for a minimum of 96 FC (Fiber Channel) ports in a single chassis (switch enclosure).				
2	Shall provide support for Fiber Channel Protocol, minimum port speed will be 8 Gbit/s Fiber Channel Ports and all Fibre Channel ports shall be capable of autosensing slower speeds.				
3	Shall support a minimum of 256 Gbit/per second of internal bandwidth in a single chassis (switch enclosure).				
4	Shall support no more than 4:1 over subscription at 8 Gbit/per second.				
5	Shall support N Port Virtualization with NPV or similar technology.				
6	Shall provide Redundant Fabric Connectivity between Hosts and Storage device.				
7	Shall provide Redundant Switch Control (Supervisor) Cards.				
8	Shall provide redundancy in all switch components with no single point of failure and non-disruptive to operations for all switch component replacements or repairs or firmware and microcode upgrades or updates.				
9	Solution shall provide that all components be easily accessible without having to de-install, disconnect or remove other components so that service can be accomplished without interruption to normal switch operations.				
10	Shall provide Fabric Isolation, Local SAN Fabric shall be isolated from replication SAN Fabric.				
11	Shall provide the ability to isolate data path from Host to Storage within a fabric, and configure throughput for specific applications.				
12	Solution shall be able to provision additional Fibre ports required to fully connect solution components (storage subsystems, inter-switch links, management hosts, replication hardware, etc) in addition to host(s) without interruption to normal operations.				

13	Shall provide fully redundant configuration, shall have the ability to retain a copy of the current configuration and the previous configuration.				
14	Solution shall provide cover to all power switches to prevent inadvertent activation.				
15	Shall support ability to be managed from a single Fabric management tool which can access all components in the SAN Fabric switch.				
16	Shall provide the support of role-based security for access and management of the switch.				
17	Shall support authenticable access with logging (for audits).				
18	Solution shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage Fabric(s) in a data center implementation from a single master management console.				
19	Shall provide the ability to gather performance statistics for Inter-Switch Links (ISLs), host and storage device connections, and traffic between specific Fiber Channel sources and destinations (flows).				
20	Shall provide the ability to set performance thresholds based on manual entry or calculated based on previous measurements.				
21	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential switch system problems, failures and switch resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.				
22	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Fabric Switch Subsystem, such as switch cards, ports, ports GBICs, Fans or Power Supplies and similar components.				
23	Shall provide a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Fibre Fabric SAN Switch - Large					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide support for a minimum of 288 FC (Fiber Channel) ports in a single chassis (switch enclosure).				
2	Shall provide support for Fiber Channel Protocol, minimum port speed will be 8 Gbit/s Fiber Channel Ports and all Fibre Channel ports shall be capable of autosensing slower speeds.				
3	Shall support a minimum of 256 Gbit/per second of internal bandwidth in a single chassis (switch enclosure).				
4	Shall support no more than 4:1 over subscription at 8 Gbit/per second.				
5	Shall support N Port Virtualization with NPV or similar technology.				
6	Shall provide Redundant Fabric Connectivity between Hosts and Storage device.				
7	Shall provide Redundant Switch Control (Supervisor) Cards.				
8	Shall provide redundancy in all switch components with no single point of failure and non-disruptive to operations for all switch component replacements or repairs or firmware and microcode upgrades or updates.				
9	Solution shall provide that all components be easily accessible without having to de-install, disconnect or remove other components so that service can be accomplished without interruption to normal switch operations.				
10	Shall provide Fabric Isolation, Local SAN Fabric shall be isolated from replication SAN Fabric.				
11	Shall provide the ability to isolate data path from Host to Storage within a fabric, and configure throughput for specific applications.				
12	Solution shall be able to provision additional Fibre ports required to fully connect solution components (storage subsystems, inter-switch links, management hosts, replication hardware, etc) in addition to host(s) without interruption to normal operations.				

13	Shall provide fully redundant configuration, shall have the ability to retain a copy of the current configuration and the previous configuration.				
14	Solution shall provide cover to all power switches to prevent inadvertent activation.				
15	Shall support ability to be managed from a single Fabric management tool which can access all components in the SAN Fabric switch.				
16	Shall provide the support of role-based security for access and management of the switch.				
17	Shall support authenticable access with logging (for audits).				
18	Solution shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage Fabric(s) in a data center implementation from a single master management console.				
19	Shall provide the ability to gather performance statistics for Inter-Switch Links (ISLs), host and storage device connections, and traffic between specific Fiber Channel sources and destinations (flows).				
20	Shall provide the ability to set performance thresholds based on manual entry or calculated based on previous measurements.				
21	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential switch system problems, failures and switch resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.				
22	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Fabric Switch Subsystem, such as switch cards, ports, ports GBICs, Fans or Power Supplies and similar components.				
23	Shall provide a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Low Performance NAS Storage					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 15000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum of 1000 MBytes per second throughput.				
3	Shall provide the minimum usable disk capacity of 25 TB . This initial data storage capacity must be distributed evenly among two or more shelves in order for future local mirroring between shelves.				
4	Shall be capable of supporting a minimum of two (2) 1GbE Ethernet ports per controller.				
5	Shall be capable of supporting a minimum of one (1) 10GbE Ethernet ports per controller.				
6	Shall provide, excluding space on Storage System Operating System (OS) Drives, the initial data storage capacity must be distributed evenly among multiple shelves and support mirroring between shelves.				
7	Shall provide the disk failure redundancy configuration to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
8	Shall provide the appropriate disk drive technologies to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
9	Shall provide the capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
10	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				

11	Shall provide a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				
12	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
13	Shall provide having a minimum of two storage controllers in a module for controller based storage solution, which are load balanced and provide automatic failover including the ability to maintain access to all data through a controller failure.				
14	Shall provide the ability to upgrade a controller on the storage system without replacing the storage subsystem. The upgrade of the controller shall be non-destructive to the data, requiring no data migration, no reconfiguration and no LUN remapping to any connecting Host.				
15	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
16	Shall provide the ability to expand and scale Host side connectivity and backend capacity independently in order to increase performance.				
17	Shall support the ability to create volumes across multiple spindles.				
18	Shall support dynamic volume expansion.				
19	Shall be able to present a minimum of 15TB volume to connecting hosts.				
20	Shall provide ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.				
21	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".				
22	Shall provide the capability of supporting Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
23	Shall be capable of supporting Jumbo frame and LACP protocol.				
24	Shall provide the ability to expand Ethernet ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
25	Shall be capable of supporting 10GbE Ethernet ports.				

26	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
27	Shall provide the capability of synchronous and asynchronous replication with write order fidelity.				
28	Shall provide the maintaining redundancy and performance metrics requirements as stated above during snapshot and clone operations.				
29	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times.				
30	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times.				
31	Shall provide the ability to allow source update during data migration. Online data migration from other storage devices and data migration is transparent to users.				
32	Power Switches shall be covered to prevent inadvertent activation.				
33	Shall support role-based security or audit trail logging for access to storage.				
34	Shall provide a single master management interface to manage multiple storage subsystems of the same type in data center implementations.				
35	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
36	Shall have a graphical user interface (GUI) or command line interface (CLI), or a Wizard to automate the process to provision a large number of Disk Groups or Volumes.				
37	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
38	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
39	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Medium Performance NAS Storage					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 35000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum of 2500 MBytes per second throughput.				
3	Shall provide the minimum usable disk capacity of 50 TB . This initial data storage capacity must be distributed evenly among two or more shelves in order for future local mirroring between shelves.				
4	Shall be capable of supporting a minimum of four (4) 1GbE Ethernet ports per controller.				
5	Shall be capable of supporting a minimum of two (2) 10GbE Ethernet ports per controller.				
6	Shall provide, excluding space on Storage System Operating System (OS) Drives, the initial data storage capacity must be distributed evenly among multiple shelves and support mirroring between shelves.				
7	Shall provide the disk failure redundancy configuration to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
8	Shall provide the appropriate disk drive technologies to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
9	Shall provide the capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
10	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				

11	Shall provide a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				
12	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
13	Shall provide having a minimum of two storage controllers in a module for controller based storage solution, which are load balanced and provide automatic failover including the ability to maintain access to all data through a controller failure.				
14	Shall provide the ability to upgrade a controller on the storage system without replacing the storage subsystem. The upgrade of the controller shall be non-destructive to the data, requiring no data migration, no reconfiguration and no LUN remapping to any connecting Host.				
15	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
16	Shall provide the ability to expand and scale Host side connectivity and backend capacity independently in order to increase performance.				
17	Shall support the ability to create volumes across multiple spindles.				
18	Shall support dynamic volume expansion.				
19	Shall be able to present a minimum of 15TB volume to connecting hosts.				
20	Shall provide ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.				
21	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".				
22	Shall provide the capability of supporting Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
23	Shall be capable of supporting Jumbo frame and LACP protocol.				
24	Shall provide the ability to expand Ethernet ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
25	Shall be capable of supporting 10GbE Ethernet ports.				

26	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
27	Shall provide the capability of synchronous and asynchronous replication with write order fidelity.				
28	Shall provide the maintaining redundancy and performance metrics requirements as stated above during snapshot and clone operations.				
29	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times.				
30	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times.				
31	Shall provide the ability to allow source update during data migration. Online data migration from other storage devices and data migration is transparent to users.				
32	Power Switches shall be covered to prevent inadvertent activation.				
33	Shall support role-based security or audit trail logging for access to storage.				
34	Shall provide a single master management interface to manage multiple storage subsystems of the same type in data center implementations.				
35	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
36	Shall have a graphical user interface (GUI) or command line interface (CLI), or a Wizard to automate the process to provision a large number of Disk Groups or Volumes.				
37	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
38	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
39	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

High Performance NAS Storage					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 75000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum of 4000 MBytes per second throughput.				
3	Shall provide the minimum usable disk capacity of 75 TB . This initial data storage capacity must be distributed evenly among two or more shelves in order for future local mirroring between shelves.				
4	Shall be capable of supporting a minimum of four (4) 1GbE Ethernet ports per controller.				
5	Shall be capable of supporting a minimum of two (2) 10GbE Ethernet ports per controller.				
6	Shall provide, excluding space on Storage System Operating System (OS) Drives, the initial data storage capacity must be distributed evenly among multiple shelves and support mirroring between shelves.				
7	Shall provide the disk failure redundancy configuration to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
8	Shall provide the appropriate disk drive technologies to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
9	Shall provide the capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
10	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				

11	Shall provide a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).			
12	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.			
13	Shall provide having a minimum of two storage controllers in a module for controller based storage solution, which are load balanced and provide automatic failover including the ability to maintain access to all data through a controller failure.			
14	Shall provide the ability to upgrade a controller on the storage system without replacing the storage subsystem. The upgrade of the controller shall be non-destructive to the data, requiring no data migration, no reconfiguration and no LUN remapping to any connecting Host.			
15	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.			
16	Shall provide the ability to expand and scale Host side connectivity and backend capacity independently in order to increase performance.			
17	Shall support the ability to create volumes across multiple spindles.			
18	Shall support dynamic volume expansion.			
19	Shall be able to present a minimum of 15TB volume to connecting hosts.			
20	Shall provide ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.			
21	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".			
22	Shall provide the capability of supporting Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.			
23	Shall be capable of supporting Jumbo frame and LACP protocol.			
24	Shall provide the ability to expand Ethernet ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.			
25	Shall be capable of supporting 10GbE Ethernet ports.			

26	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
27	Shall provide the capability of synchronous and asynchronous replication with write order fidelity.				
28	Shall provide the maintaining redundancy and performance metrics requirements as stated above during snapshot and clone operations.				
29	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times.				
30	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times.				
31	Shall provide the ability to allow source update during data migration. Online data migration from other storage devices and data migration is transparent to users.				
32	Power Switches shall be covered to prevent inadvertent activation.				
33	Shall support role-based security or audit trail logging for access to storage.				
34	Shall provide a single master management interface to manage multiple storage subsystems of the same type in data center implementations.				
35	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
36	Shall have a graphical user interface (GUI) or command line interface (CLI), or a Wizard to automate the process to provision a large number of Disk Groups or Volumes.				
37	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
38	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
39	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Low Performance NAS Storage Modular					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 15000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum of 1000 MBytes per second throughput.				
3	Shall provide the minimum usable disk capacity of 25 TB . This initial data storage capacity must be distributed evenly among two or more shelves in order for future local mirroring between shelves.				
4	Shall be capable of supporting a minimum of two (2) 1GbE Ethernet ports per controller.				
5	Shall be capable of supporting a minimum of one (1) 10GbE Ethernet ports per controller.				
6	Shall provide the disk failure redundancy configuration to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
7	Shall provide the appropriate disk drive technologies to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
8	Shall provide the capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
9	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				
10	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				
11	Shall natively provide Single Global Namespace file system from the storage device.				

Attachment G - Storage Arrays / Storage Appliances Specifications

12	Shall provide the capability of maintaining Single Global Namespace within the Data Center when scale beyond the maximum capacity of the storage device.				
13	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
14	Shall provide having a minimum of two storage controllers or nodes in a subsystem that are load balanced and provide automatic failover including the ability to maintain access to all data through a controller or node failure.				
15	Shall provide adding, upgrading or replacing storage subsystem components while adhering to the required storage subsystem uptime requirement, redundancy and needing no data migration, no reconfiguration and no LUN remapping to any connecting Host.				
16	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
17	Shall provide the ability to expand and scale Host side connectivity and backend capacity independently in order to increase performance.				
18	Shall support the ability to create volumes across multiple spindles.				
19	Shall support dynamic volume expansion.				
20	Shall provide ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.				
21	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".				
22	Shall provide the capability of supporting Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
23	Shall be capable of supporting Jumbo frame and LACP protocol.				
24	Shall provide the ability to expand Ethernet ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
25	Shall be capable of supporting 10GbE Ethernet ports.				
26	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
27	Shall provide the capability of synchronous and asynchronous replication.				

28	Shall be capable of maintaining redundancy and the performance metrics requirements as stated above during snapshot and clone operations.				
29	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all time during the migration process.				
30	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
31	Shall provide the ability to allow source update during data migration. Online data migration from other storage devices and data migration is transparent to users.				
32	Power Switches shall be covered to prevent inadvertent activation.				
33	Shall support role-based security or audit trail logging for access to storage.				
34	Shall provide a single master management interface to manage multiple storage subsystems of the same type in data center implementations.				
35	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
36	Shall have a graphical user interface (GUI) or command line interface (CLI), or a Wizard to automate the process to provision a large number of Disk Groups or Volumes.				
37	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
38	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
39	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Medium Performance NAS Storage Modular					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 35000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum of 2500 MBytes per second throughput.				
3	Shall provide the minimum usable disk capacity of 50 TB . This initial data storage capacity must be distributed evenly among two or more shelves in order for future local mirroring between shelves.				
4	Shall be capable of supporting a minimum of four (4) 1GbE Ethernet ports per controller.				
5	Shall be capable of supporting a minimum of two (2) 10GbE Ethernet ports per controller.				
6	Shall provide the disk failure redundancy configuration to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
7	Shall provide the appropriate disk drive technologies to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
8	Shall provide the capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
9	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				
10	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				
11	Shall natively provide Single Global Namespace file system from the storage device.				

12	Shall provide the capability of maintaining Single Global Namespace within the Data Center when scale beyond the maximum capacity of the storage device.				
13	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
14	Shall provide having a minimum of two storage controllers or nodes in a subsystem that are load balanced and provide automatic failover including the ability to maintain access to all data through a controller or node failure.				
15	Shall provide adding, upgrading or replacing storage subsystem components while adhering to the required storage subsystem uptime requirement, redundancy and needing no data migration, no reconfiguration and no LUN remapping to any connecting Host.				
16	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
17	Shall provide the ability to expand and scale Host side connectivity and backend capacity independently in order to increase performance.				
18	Shall support the ability to create volumes across multiple spindles.				
19	Shall support dynamic volume expansion.				
20	Shall provide ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.				
21	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".				
22	Shall provide the capability of supporting Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
23	Shall be capable of supporting Jumbo frame and LACP protocol.				
24	Shall provide the ability to expand Ethernet ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
25	Shall be capable of supporting 10GbE Ethernet ports.				
26	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
27	Shall provide the capability of synchronous and asynchronous replication.				

28	Shall be capable of maintaining redundancy and the performance metrics requirements as stated above during snapshot and clone operations.				
29	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all time during the migration process.				
30	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
31	Shall provide the ability to allow source update during data migration. Online data migration from other storage devices and data migration is transparent to users.				
32	Power Switches shall be covered to prevent inadvertent activation.				
33	Shall support role-based security or audit trail logging for access to storage.				
34	Shall provide a single master management interface to manage multiple storage subsystems of the same type in data center implementations.				
35	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
36	Shall have a graphical user interface (GUI) or command line interface (CLI), or a Wizard to automate the process to provision a large number of Disk Groups or Volumes.				
37	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
38	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
39	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

High Performance NAS Storage Modular					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 75000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum of 4000 MBytes per second throughput.				
3	Shall provide the minimum usable disk capacity of 75 TB . This initial data storage capacity must be distributed evenly among two or more shelves in order for future local mirroring between shelves.				
4	Shall be capable of supporting a minimum of four (4) 1GbE Ethernet ports per controller.				
5	Shall be capable of supporting a minimum of two (2) 10GbE Ethernet ports per controller.				
6	Shall provide the disk failure redundancy configuration to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
7	Shall provide the appropriate disk drive technologies to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
8	Shall provide the capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
9	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				
10	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				
11	Shall natively provide Single Global Namespace file system from the storage device.				

12	Shall provide the capability of maintaining Single Global Namespace within the Data Center when scale beyond the maximum capacity of the storage device.				
13	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
14	Shall provide having a minimum of two storage controllers or nodes in a subsystem that are load balanced and provide automatic failover including the ability to maintain access to all data through a controller or node failure.				
15	Shall provide adding, upgrading or replacing storage subsystem components while adhering to the required storage subsystem uptime requirement, redundancy and needing no data migration, no reconfiguration and no LUN remapping to any connecting Host.				
16	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
17	Shall provide the ability to expand and scale Host side connectivity and backend capacity independently in order to increase performance.				
18	Shall support the ability to create volumes across multiple spindles.				
19	Shall support dynamic volume expansion.				
20	Shall provide ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.				
21	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".				
22	Shall provide the capability of supporting Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
23	Shall be capable of supporting Jumbo frame and LACP protocol.				
24	Shall provide the ability to expand Ethernet ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
25	Shall be capable of supporting 10GbE Ethernet ports.				
26	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
27	Shall provide the capability of synchronous and asynchronous replication.				

28	Shall be capable of maintaining redundancy and the performance metrics requirements as stated above during snapshot and clone operations.				
29	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all time during the migration process.				
30	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
31	Shall provide the ability to allow source update during data migration. Online data migration from other storage devices and data migration is transparent to users.				
32	Power Switches shall be covered to prevent inadvertent activation.				
33	Shall support role-based security or audit trail logging for access to storage.				
34	Shall provide a single master management interface to manage multiple storage subsystems of the same type in data center implementations.				
35	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
36	Shall have a graphical user interface (GUI) or command line interface (CLI), or a Wizard to automate the process to provision a large number of Disk Groups or Volumes.				
37	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
38	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
39	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Low Performance iSCSI Storage Modular					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 20000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum of 1000 MBytes per second throughput.				
3	Shall provide the minimum usable disk capacity of 25 TB . This initial data storage capacity must be distributed evenly among two or more shelves in order for future local mirroring between shelves.				
4	Shall be capable of supporting a minimum of two (2) 1GbE iSCSI ports per controller.				
5	Shall be capable of supporting a minimum of one (1) 10GbE iSCSI ports per controller.				
6	Shall provide the disk failure redundancy configuration to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
7	Shall provide the appropriate disk drive technologies to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
8	Shall provide capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
9	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				
10	Shall provide a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				

11	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode. upgrades or updates.			
12	Shall provide having a minimum of two storage controllers in a module for controller based storage solutions or a node, which are load balanced and provide automatic failover including the ability to maintain access to all data through a controller failure.			
13	Shall provide the ability to upgrade a controller or a node on the storage system without replacing the storage subsystem and the upgrade of the controller or the node, shall be non-destructive to the data, requiring no data migration, no reconfiguration and no LUN remapping to any connecting Host.			
14	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.			
15	Shall provide the ability to expand and scale Host side connectivity and backend capacity independently in order to increase performance.			
16	Shall provide storage subsystem support boot from SAN.			
17	Shall support the ability to manually change LUN UDID value.			
18	Shall support the ability to create volumes across multiple spindles.			
19	Shall support non-disruptive LUN and volume expansion.			
20	Shall be able to present a minimum of 15TB LUN/volume to a host via iSCSI.			
21	Shall provide the ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.			
22	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".			
23	Minimum shall be capable of supporting Internet Small Computer System Interface (iSCSI) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.			
24	Shall support the ability to expand iSCSI ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.			
25	Shall be capable of supporting 10GbE Ethernet ports.			
26	Shall be capable of supporting Jumbo frame and LACP protocol.			
27	Solution shall support a minimum of 256 SCSI initiators (host connections).			

28	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
29	Shall provide the capability of synchronous and asynchronous replication with write order fidelity.				
30	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
31	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
32	Shall be able to allow source update during data migration. Online data migration from other storage array and data migration is transparent to users.				
33	Shall be capable of maintaining redundancy and the performance metrics requirements as stated above during snapshot and clone operations.				
34	Power Switches shall be covered to prevent inadvertent activation.				
35	Shall support role-based security or audit trail logging for access to storage.				
36	Shall provide a single master management interface to manage multiple storage subsystems of the same type in data center implementations.				
37	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
38	Shall have a graphical user interface (GUI) or command (CLI), or a Wizard to automate the process to provision a large number of Disk Groups and LUNs.				
39	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
40	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
41	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar function.				

Medium Performance iSCSI Storage Modular					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 50000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum of 2500 MBytes per second throughput.				
3	Shall provide the minimum usable disk capacity of 50 TB . This initial data storage capacity must be distributed evenly among two or more shelves in order for future local mirroring between shelves.				
4	Shall be capable of supporting a minimum of four (4) 1GbE iSCSI ports per controller.				
5	Shall be capable of supporting a minimum of two (2) 10GbE iSCSI ports per controller.				
6	Shall provide the disk failure redundancy configuration to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
7	Shall provide the appropriate disk drive technologies to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
8	Shall provide capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
9	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				
10	Shall provide a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				

11	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode. upgrades or updates.				
12	Shall provide having a minimum of two storage controllers in a module for controller based storage solutions or a node, which are load balanced and provide automatic failover including the ability to maintain access to all data through a controller failure.				
13	Shall provide the ability to upgrade a controller or a node on the storage system without replacing the storage subsystem and the upgrade of the controller or the node, shall be non-destructive to the data, requiring no data migration, no reconfiguration and no LUN remapping to any connecting Host.				
14	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
15	Shall provide the ability to expand and scale Host side connectivity and backend capacity independently in order to increase performance.				
16	Shall provide storage subsystem support boot from SAN.				
17	Shall support the ability to manually change LUN UDID value.				
18	Shall support the ability to create volumes across multiple spindles.				
19	Shall support non-disruptive LUN and volume expansion.				
20	Shall be able to present a minimum of 15TB LUN/volume to a host via iSCSI.				
21	Shall provide the ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.				
22	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".				
23	Minimum shall be capable of supporting Internet Small Computer System Interface (iSCSI) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
24	Shall support the ability to expand iSCSI ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
25	Shall be capable of supporting 10GbE Ethernet ports.				
26	Shall be capable of supporting Jumbo frame and LACP protocol.				
27	Solution shall support a minimum of 256 SCSI initiators (host connections).				

28	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
29	Shall provide the capability of synchronous and asynchronous replication with write order fidelity.				
30	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
31	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
32	Shall be able to allow source update during data migration. Online data migration from other storage array and data migration is transparent to users.				
33	Shall be capable of maintaining redundancy and the performance metrics requirements as stated above during snapshot and clone operations.				
34	Power Switches shall be covered to prevent inadvertent activation.				
35	Shall support role-based security or audit trail logging for access to storage.				
36	Shall provide a single master management interface to manage multiple storage subsystems of the same type in data center implementations.				
37	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
38	Shall have a graphical user interface (GUI) or command (CLI), or a Wizard to automate the process to provision a large number of Disk Groups and LUNs.				
39	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
40	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
41	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar function.				

High Performance iSCSI Storage Modular					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 75000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum of 4000 MBytes per second throughput.				
3	Shall provide the minimum usable disk capacity of 75 TB . This initial data storage capacity must be distributed evenly among two or more shelves in order for future local mirroring between shelves.				
4	Shall be capable of supporting a minimum of four (4) 1GbE iSCSI ports per controller.				
5	Shall be capable of supporting a minimum of two (2) 10GbE iSCSI ports per controller.				
6	Shall provide the disk failure redundancy configuration to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
7	Shall provide the appropriate disk drive technologies to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
8	Shall provide capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
9	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				
10	Shall provide a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				

11	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode. upgrades or updates.			
12	Shall provide having a minimum of two storage controllers in a module for controller based storage solutions or a node, which are load balanced and provide automatic failover including the ability to maintain access to all data through a controller failure.			
13	Shall provide the ability to upgrade a controller or a node on the storage system without replacing the storage subsystem and the upgrade of the controller or the node, shall be non-destructive to the data, requiring no data migration, no reconfiguration and no LUN remapping to any connecting Host.			
14	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.			
15	Shall provide the ability to expand and scale Host side connectivity and backend capacity independently in order to increase performance.			
16	Shall provide storage subsystem support boot from SAN.			
17	Shall support the ability to manually change LUN UDID value.			
18	Shall support the ability to create volumes across multiple spindles.			
19	Shall support non-disruptive LUN and volume expansion.			
20	Shall be able to present a minimum of 15TB LUN/volume to a host via iSCSI.			
21	Shall provide the ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.			
22	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".			
23	Minimum shall be capable of supporting Internet Small Computer System Interface (iSCSI) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.			
24	Shall support the ability to expand iSCSI ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.			
25	Shall be capable of supporting 10GbE Ethernet ports.			
26	Shall be capable of supporting Jumbo frame and LACP protocol.			
27	Solution shall support a minimum of 256 SCSI initiators (host connections).			

28	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
29	Shall provide the capability of synchronous and asynchronous replication with write order fidelity.				
30	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
31	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
32	Shall be able to allow source update during data migration. Online data migration from other storage array and data migration is transparent to users.				
33	Shall be capable of maintaining redundancy and the performance metrics requirements as stated above during snapshot and clone operations.				
34	Power Switches shall be covered to prevent inadvertent activation.				
35	Shall support role-based security or audit trail logging for access to storage.				
36	Shall provide a single master management interface to manage multiple storage subsystems of the same type in data center implementations.				
37	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
38	Shall have a graphical user interface (GUI) or command (CLI), or a Wizard to automate the process to provision a large number of Disk Groups and LUNs.				
39	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
40	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
41	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar function.				

LTO Tape Library Large					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum and maximum ranges that procurement shall fit in, The exact numbers will be supplied on the individual procurement request.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support the following minimum amount of storage that will be leveraging Long Term Storage of: 75 TB (Terabytes).				
2	Shall support a minimum of 8 Tape Drives.				
3	Shall support the ability to scale to a maximum of at least 32 Tape Drives.				
4	Shall support a minimum of 120 Tape Cartridge slots.				
5	Shall support the ability to scale to a maximum of at least 640 Tape Cartridge slots.				
6	Shall provide a minimum of four (4) 8-Gbit/second Fiber Channel ports per Library I/O Control Unit and all Fiber Channel ports shall be capable of autosensing slower speeds.				
7	Shall provide a minimum of four (4) 1GbE iSCSI ports per Library I/O Control Unit.				
8	Shall provide a minimum of two (2) 10GbE iSCSI ports per Library I/O Control Unit.				
9	Shall support the capability of a minimum throughput of 6 TB (Terabytes) per hour.				
10	Shall support the capability of a minimum scalability of up to 24 TB (Terabytes) per hour.				
11	Shall support LTO Ultrium (Linear Tape-Open) tape technology format.				
12	Shall provide the appropriate Robotic access time to meet the stated IOPs and workload performance requirements (to be specified at delivery order level).				
13	Shall provide Tape Drive to read and write to either LTO4 or higher specifications and throughput of 120 MB/s (Megabytes) or higher.				
14	Shall provide the capability of Tape Drive encryption with a minimum 1 TB (Terabyte) tape capacity and throughput minimum of 120 MB/s (Megabytes) with using a minimum encryption of FIPS 140-2				
15	Shall provide the capability of Tape Library and LTO Tape Drives that can support enterprise tape encryption using a minimum of FIPS 140-2 validated encryption.				

16	Shall provide tape drives that read previously written tapes in LTO format back 2 previous LTO generations (example: LTO 2 Tapes able to be read by LTO 4 Tape Drive).			
17	Shall support iSCSI to Fibre Channel bridge device and Fiber Channel protocol connectivity.			
18	Shall provide Library iSCSI port expansion to meet storage performance requirements (to be specified at delivery order level) in the same unit(s). Capable of adding additional ports per Library I/O Control Unit.			
19	Shall provide Library Fiber Channel ports expansion to meet the storage performance requirements (to be specified at delivery order level) in the same unit(s). Capable of adding additional ports per Library I/O Control Unit.			
20	Shall be capable of supporting native 10GbE Ethernet host connections for the Library.			
21	Shall support a tape cartridge bulk load capacity of: 10			
22	Shall support the capability of expansion and scalability, including Interface cards and capacity (Tape Drives & Cartridge Slots) independently with minimum of planned downtime of 2 hours or less.			
23	Shall support an automated and manual method to check Tape Drive and Robotic integrity.			
24	Shall support the ability of the Library to recover from a Tape Drive failure and will not cause data loss.			
25	Shall support the capability of multiple SCSI initiators (host connections) to the Library.			
26	Shall Support a minimum of four nines storage subsystem uptime (or 99.99% availability, excluding planned downtime). (Note: Four Nines (99.99%) = 53 minutes of downtime on a rolling 12 month basis).			
27	Shall support Tape Library having a minimum of two I/O Control Units/Board(s)/Card(s).			
28	Shall provide redundancy in all Tape Library system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates (only exception is Library Robotic Arm).			
29	Shall support the capability of isolating subsystem resources (Ports, Interface Cards, and Tape Drives) and dedicate a defined amount of these resources to a particular host or VLAN.			
30	Power switches shall be covered to prevent inadvertent activation.			

31	Shall support role-based security for access and manage Tape Library control.				
32	Shall support authenticable access with logging (for audits).				
33	Shall support a single master management interface to manage multiple Tape Library(s) of the same type in data center implementations.				
34	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage Tape Libraries in a data center implementation from a single master management console.				
35	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential Tape Library system problems, failures and Library resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.				
36	Shall provide native capability to report and generate reports at a minimum on hardware failures or items out-of-specifications in the Tape Library & Subsystems, such as Tape Drives, Fans or Power Supplies and similar components.				
37	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

LTO Tape Library Medium to Small					
OEM Model/Part Number:					
Reference Number	Requirement	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support the following minimum amount of storage that will be leveraging Long Term Storage of: 25 TB (Terabytes).				
2	Shall support a minimum of 4 Tape Drives.				
3	Shall support the ability to scale to a maximum of at least 12 Tape Drives.				
4	Shall support a minimum of 25 Tape Cartridge slots.				
5	Shall support the ability to scale to a maximum of at least 180 Tape Cartridge slots.				
6	Shall provide a minimum of two (2) 8-Gbit/second Fiber Channel ports per Library I/O Control Unit and all Fiber Channel ports shall be capable of autosensing slower speeds.				
7	Shall provide a minimum of two (2) 1GbE iSCSI ports per Library I/O Control Unit.				
8	Shall provide a minimum of one (1) 10GbE iSCSI ports per Library I/O Control Unit.				
9	Shall support the capability of a minimum throughput of 6 TB (Terabytes) per hour.				
10	Shall support the capability of a minimum scalability of up to 24 TB (Terabytes) per hour.				
11	Shall support LTO Ultrium (Linear Tape-Open) tape technology format.				
12	Shall provide the appropriate Robotic access time to meet the stated IOPs and workload performance requirements (to be specified at delivery order level).				
13	Shall provide Tape Drive to read and write to either LTO4 or higher specifications and throughput of 120 MB/s (Megabytes) or higher.				
14	Shall provide the capability of Tape Drive encryption with a minimum 1 TB (Terabyte) tape capacity and throughput minimum of 120 MB/s (Megabytes) with using a minimum encryption of FIPS 140-2				
15	Shall provide the capability of Tape Library and LTO Tape Drives that can support enterprise tape encryption using a minimum of FIPS 140-2 validated encryption.				

16	Shall provide tape drives that read previously written tapes in LTO format back 2 previous LTO generations (example: LTO 2 Tapes able to be read by LTO 4 Tape Drive).			
17	Shall support iSCSI to Fibre Channel bridge device and Fiber Channel protocol connectivity.			
18	Shall provide Library iSCSI port expansion to meet storage performance requirements (to be specified at delivery order level) in the same unit(s). Capable of adding additional ports per Library I/O Control Unit.			
19	Shall provide Library Fiber Channel ports expansion to meet the storage performance requirements (to be specified at delivery order level) in the same unit(s). Capable of adding additional ports per Library I/O Control Unit.			
20	Shall be capable of supporting native 10GbE Ethernet host connections for the Library.			
21	Shall support a tape cartridge bulk load capacity of: 10			
22	Shall support the capability of expansion and scalability, including Interface cards and capacity (Tape Drives & Cartridge Slots) independently with minimum of planned downtime of 2 hours or less.			
23	Shall support an automated and manual method to check Tape Drive and Robotic integrity.			
24	Shall support the ability of the Library to recover from a Tape Drive failure and will not cause data loss.			
25	Shall support the capability of multiple SCSI initiators (host connections) to the Library.			
26	Shall Support a minimum of four nines storage subsystem uptime (or 99.99% availability, excluding planned downtime). (Note: Four Nines (99.99%) = 53 minutes of downtime on a rolling 12 month basis).			
27	Shall support Tape Library having a minimum of two I/O Control Units/Board(s)/Card(s).			
28	Shall provide redundancy in all Tape Library system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates (only exception is Library Robotic Arm).			
29	Shall support the capability of isolating subsystem resources (Ports, Interface Cards, and Tape Drives) and dedicate a defined amount of these resources to a particular host or VLAN.			
30	Power switches shall be covered to prevent inadvertent activation.			

31	Shall support role-based security for access and manage Tape Library control.				
32	Shall support authenticable access with logging (for audits).				
33	Shall support a single master management interface to manage multiple Tape Library(s) of the same type in data center implementations.				
34	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage Tape Libraries in a data center implementation from a single master management console.				
35	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential Tape Library system problems, failures and Library resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.				
36	Shall provide native capability to report and generate reports at a minimum on hardware failures or items out-of-specifications in the Tape Library & Subsystems, such as Tape Drives, Fans or Power Supplies and similar components.				
37	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

LTO Tape Cartridge (LTO4)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum and maximum ranges that procurement shall fit in, The exact numbers will be supplied on the individual procurement request.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Ultrium (LTO) cartridge's container dimensions to be 102.0 x 105.4 x 21.5(mm).				
2	Ultrium (LTO) tape dimensions within the cartridge container shall be an industry standard of half-inch (½ inch) wide tape.				
3	LTO tape format shall be laid out in this manner, 4 wide data bands, sandwiched between 5 narrow servo bands.				
4	LTO tape format shall be laid out on tape in the following order, 4 data bands x 11-20 wraps per head x 8 or 16 tracks per wrap.				
5	Every LTO 1 / 2 / 3 cartridge shall contain a Cartridge Memory chip capable of containing 128 blocks of memory, each block containing 32 Bytes for a total of 4 KB.				
6	Every LTO 4 / 5 cartridge shall contain a Cartridge Memory chip capable of containing 256 blocks of memory, each block containing 32 Bytes for a total of 8 KB.				
7	Every LTO Cartridge Memory chip shall be capable of being read and/or written, one block at a time by a non contacting passive RF interface.				
8	Every LTO Cartridge tape shall have a leader pin attached to the end of the tape to enable the tape drive to grasp the tape and mount it in a take-up reel inside the tape drive.				
9	Every LTO Cartridge tape shall have this leader pin held in place at the opening of the tape cartridge by a small spring when the tape cartridge is not in a tape drive.				
10	The LTO cartridge tape shall be capable of the following: 15-30 years of archival life. 5000 cartridge loads/unloads and approx. 260 full file passes.				

11	LTO-4: Native Data Capacity 800GB; Compression Capability 2:1; Partition Capable=No; Tape Thickness=6.6 um; Tape Length=820 m; Tape Tracks=896; Wraps per Band=14; Linear Density (bits/mm)=13250				
----	--	--	--	--	--

LTO Tape Cartridge (LTO5)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum and maximum ranges that procurement shall fit in, The exact numbers will be supplied on the individual procurement request.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Ultrium (LTO) cartridge's container dimensions to be 102.0 x 105.4 x 21.5(mm).				
2	Ultrium (LTO) tape dimensions within the cartridge container shall be an industry standard of half-inch (½ inch) wide tape.				
3	LTO tape format shall be laid out in this manner, 4 wide data bands, sandwiched between 5 narrow servo bands.				
4	LTO tape format shall be laid out on tape in the following order, 4 data bands x 11-20 wraps per head x 8 or 16 tracks per wrap.				
5	Every LTO 1 / 2 / 3 cartridge shall contain a Cartridge Memory chip capable of containing 128 blocks of memory, each block containing 32 Bytes for a total of 4 KB.				
6	Every LTO 4 / 5 cartridge shall contain a Cartridge Memory chip capable of containing 256 blocks of memory, each block containing 32 Bytes for a total of 8 KB.				
7	Every LTO Cartridge Memory chip shall be capable of being read and/or written, one block at a time by a non contacting passive RF interface.				
8	Every LTO Cartridge tape shall have a leader pin attached to the end of the tape to enable the tape drive to grasp the tape and mount it in a take-up reel inside the tape drive.				
9	Every LTO Cartridge tape shall have this leader pin held in place at the opening of the tape cartridge by a small spring when the tape cartridge is not in a tape drive.				
10	The LTO cartridge tape shall be capable of the following: 15-30 years of archival life. 5000 cartridge loads/unloads and approx. 260 full file passes.				

11	<p>LTO-5: Native Data Capacity 1.5TB; Compression Capability 2:1; Partition Capable=Yes; Tape Thickness=6.4 um; Tape Length=846 m; Tape Tracks=1280; Wraps per Band=20; Linear Density (bits/mm)=15142</p>				
----	---	--	--	--	--

IP Based Deduplication Storage (Small)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide the minimum usable native disk capacity without data compression, without data Deduplication, and without single instance (no duplicate objects) of up to 10TB (Terabytes), expandable to 50TB (Terabytes).				
2	Shall provide a minimum ingest aggregate throughput of 1.8TB/hr.				
3	Solution shall support interoperability between different sized units (Extra large, Large, Medium, and Small)				
4	Shall support concurrent parallel I/O streaming.				
5	Solution shall provide the ability to support an IP based storage device that supports both Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity protocols. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
6	Solution shall provide a minimum of two 1GbE and Ethernet ports per system.				
7	Solution shall provide a minimum of two 10GbE Ethernet ports per system.				
8	Shall provide the support of Jumbo frame, LACP (802.3ad) protocol, and bandwidth aggregation.				
9	Shall support OpenStorage Technology (OST) protocol.				
10	Shall support Ethernet port expansion to meet storage performance requirements in the same unit(s) as listed above (to be specified at delivery order level) for each storage size types. Capable of adding minimum of two additional ports per system.				

11	Solution shall provide storage capacity expansion or scalability with no downtime, no forklift upgrade, no reconfiguration of the internal operating system or re-creation of the existing subsystem components, and remapping connections of the storage subsystem to connecting hosts. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
12	Shall support expansion and scalability of the Data Deduplication System without requiring a forklift upgrade.				
13	Shall support that Data residing on the Deduplication device that meets the requirement to replicate to a geo-dispersed environment (one to many).				
14	Shall provide edge to core replication. Many to one replication.				
15	Shall provide Deduplication on the hardware appliance, not as a software only solution.				
16	Shall provide a fully redundant solution and support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
17	Shall provide fast instant restore of files and fast restore of a complete file system to primary storage and to tape without the need to rehydrate and reconstitute data on external or additional disk storage.				
18	Shall integrate with _____ Enterprise backup software (to be specified at the delivery order level).				
19	Shall provide the ability to be directly attached as a NAS device for a large block sequential I/O data stream.				
20	Shall meet the requirement to have a portion of the storage for saving non-deduped data and another portion for deduped data.				
21	Shall meet the requirement to select a portion of storage for data replication.				
22	Shall provide a built-in capability of self healing that provides corruption recovery during replication.				
23	Shall integrate with backup software to backup data directly to tape from the Deduplication secondary storage device.				
24	Shall Support a minimum of four nines storage subsystem uptime (or 99.99% availability, excluding planned downtime). (Note: Four Nines (99.99%) = 53 minutes of downtime on a rolling 12 month basis).				

25	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
26	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
27	Shall provide disk based data Deduplication.			
28	Shall provide built-in capability of perform automatic data Integrity checking after data Deduplication.			
29	Shall provide secure data Deduplication algorithms.			
30	Shall provide built-in capability for self healing that provides corruption recovery when data corruption is detected after Deduplication, without Re-Deduplication of the original data again.			
31	Shall provide the ability to rebuild file system from stored metadata.			
32	Shall provide the ability to ingest data from multiple sources (from software and/or hardware) simultaneously.			
33	Shall provide cover to power switches to prevent inadvertent activation.			
34	Shall provide the support of role-based security for access and management of storage devices.			
35	Shall support authenticable access with logging (for audits).			
36	Shall support a native Management tool able to manage multiple Deduplication Storage devices of the same type in data center implementations.			
37	Shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage Deduplication Storage device(s) in a data center implementation from a single master Management console.			
38	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.			
39	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and similar components.			

40	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				
----	--	--	--	--	--

IP Based Deduplication Storage (Medium)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide the minimum usable native disk capacity without data compression, without data Deduplication, and without single instance (no duplicate objects) of 50TB (Terabytes), expandable to 100TB (Terabytes).				
2	Shall provide a minimum ingest aggregate throughput of 3TB/hr.				
3	Solution shall support interoperability between different sized units (Extra large, Large, Medium, and Small)				
4	Shall support concurrent parallel I/O streaming.				
5	Solution shall provide the ability to support an IP based storage device that supports both Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity protocols. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
6	Solution shall provide a minimum of two 1GbE and Ethernet ports per system.				
7	Solution shall provide a minimum of two 10GbE Ethernet ports per system.				
8	Shall provide the support of Jumbo frame, LACP (802.3ad) protocol, and bandwidth aggregation.				
9	Shall support OpenStorage Technology (OST) protocol.				
10	Shall support Ethernet port expansion to meet storage performance requirements in the same unit(s) as listed above (to be specified at delivery order level) for each storage size types. Capable of adding minimum of two additional ports per system.				

11	Solution shall provide storage capacity expansion or scalability with no downtime, no forklift upgrade, no reconfiguration of the internal operating system or re-creation of the existing subsystem components, and remapping connections of the storage subsystem to connecting hosts. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
12	Shall support expansion and scalability of the Data Deduplication System without requiring a forklift upgrade.				
13	Shall support that Data residing on the Deduplication device that meets the requirement to replicate to a geo-dispersed environment (one to many).				
14	Shall provide edge to core replication. Many to one replication.				
15	Shall provide Deduplication on the hardware appliance, not as a software only solution.				
16	Shall provide a fully redundant solution and support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
17	Shall provide fast instant restore of files and fast restore of a complete file system to primary storage and to tape without the need to rehydrate and reconstitute data on external or additional disk storage.				
18	Shall integrate with _____ Enterprise backup software (to be specified at the delivery order level).				
19	Shall provide the ability to be directly attached as a NAS device for a large block sequential I/O data stream.				
20	Shall meet the requirement to have a portion of the storage for saving non-deduped data and another portion for deduped data.				
21	Shall meet the requirement to select a portion of storage for data replication.				
22	Shall provide a built-in capability of self healing that provides corruption recovery during replication.				
23	Shall integrate with backup software to backup data directly to tape from the Deduplication secondary storage device.				
24	Shall Support a minimum of four nines storage subsystem uptime (or 99.99% availability, excluding planned downtime). (Note: Four Nines (99.99%) = 53 minutes of downtime on a rolling 12 month basis).				

25	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
26	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
27	Shall provide disk based data Deduplication.			
28	Shall provide built-in capability of perform automatic data Integrity checking after data Deduplication.			
29	Shall provide secure data Deduplication algorithms.			
30	Shall provide built-in capability for self healing that provides corruption recovery when data corruption is detected after Deduplication, without Re-Deduplication of the original data again.			
31	Shall provide the ability to rebuild file system from stored metadata.			
32	Shall provide the ability to ingest data from multiple sources (from software and/or hardware) simultaneously.			
33	Shall provide cover to power switches to prevent inadvertent activation.			
34	Shall provide the support of role-based security for access and management of storage devices.			
35	Shall support authenticable access with logging (for audits).			
36	Shall support a native Management tool able to manage multiple Deduplication Storage devices of the same type in data center implementations.			
37	Shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage Deduplication Storage device(s) in a data center implementation from a single master Management console.			
38	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.			
39	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and similar components.			

40	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				
----	--	--	--	--	--

IP Based Deduplication Storage (Large)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide the minimum usable native disk capacity without data compression, without data Deduplication, and without single instance (no duplicate objects) of 100TB (Terabytes), expandable to 500TB (Terabytes).				
2	Shall provide a minimum ingest aggregate throughput of 4TB/hr.				
3	Solution shall support interoperability between different sized units (Extra large, Large, Medium, and Small)				
4	Shall support concurrent parallel I/O streaming.				
5	Solution shall provide the ability to support an IP based storage device that supports both Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity protocols. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
6	Solution shall provide a minimum of two 1GbE and Ethernet ports per system.				
7	Solution shall provide a minimum of two 10GbE Ethernet ports per system.				
8	Shall provide the support of Jumbo frame, LACP (802.3ad) protocol, and bandwidth aggregation.				
9	Shall support OpenStorage Technology (OST) protocol.				
10	Shall support Ethernet port expansion to meet storage performance requirements in the same unit(s) as listed above (to be specified at delivery order level) for each storage size types. Capable of adding minimum of two additional ports per system.				

11	Solution shall provide storage capacity expansion or scalability with no downtime, no forklift upgrade, no reconfiguration of the internal operating system or re-creation of the existing subsystem components, and remapping connections of the storage subsystem to connecting hosts. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
12	Shall support expansion and scalability of the Data Deduplication System without requiring a forklift upgrade.				
13	Shall support that Data residing on the Deduplication device that meets the requirement to replicate to a geo-dispersed environment (one to many).				
14	Shall provide edge to core replication. Many to one replication.				
15	Shall provide Deduplication on the hardware appliance, not as a software only solution.				
16	Shall provide a fully redundant solution and support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
17	Shall provide fast instant restore of files and fast restore of a complete file system to primary storage and to tape without the need to rehydrate and reconstitute data on external or additional disk storage.				
18	Shall integrate with _____ Enterprise backup software (to be specified at the delivery order level).				
19	Shall provide the ability to be directly attached as a NAS device for a large block sequential I/O data stream.				
20	Shall meet the requirement to have a portion of the storage for saving non-deduped data and another portion for deduped data.				
21	Shall meet the requirement to select a portion of storage for data replication.				
22	Shall provide a built-in capability of self healing that provides corruption recovery during replication.				
23	Shall integrate with backup software to backup data directly to tape from the Deduplication secondary storage device.				
24	Shall Support a minimum of four nines storage subsystem uptime (or 99.99% availability, excluding planned downtime). (Note: Four Nines (99.99%) = 53 minutes of downtime on a rolling 12 month basis).				

25	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
26	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
27	Shall provide disk based data Deduplication.				
28	Shall provide built-in capability of perform automatic data Integrity checking after data Deduplication.				
29	Shall provide secure data Deduplication algorithms.				
30	Shall provide built-in capability for self healing that provides corruption recovery when data corruption is detected after Deduplication, without Re-Deduplication of the original data again.				
31	Shall provide the ability to rebuild file system from stored metadata.				
32	Shall provide the ability to ingest data from multiple sources (from software and/or hardware) simultaneously.				
33	Shall provide cover to power switches to prevent inadvertent activation.				
34	Shall provide the support of role-based security for access and management of storage devices.				
35	Shall support authenticable access with logging (for audits).				
36	Shall support a native Management tool able to manage multiple Deduplication Storage devices of the same type in data center implementations.				
37	Shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage Deduplication Storage device(s) in a data center implementation from a single master Management console.				
38	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.				
39	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and similar components.				

40	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				
----	--	--	--	--	--

IP Based Deduplication Storage (Extra Large)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide the minimum usable native disk capacity without data compression, without data Deduplication, and without single instance (no duplicate objects) of 500TB (Terabytes), expandable to 750TB (Terabytes).				
2	Shall provide a minimum ingest aggregate throughput of 20TB/hr.				
3	Solution shall support interoperability between different sized units (Extra large, Large, Medium, and Small)				
4	Shall support concurrent parallel I/O streaming.				
5	Solution shall provide the ability to support an IP based storage device that supports both Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity protocols. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
6	Solution shall provide a minimum of two 1GbE and Ethernet ports per system.				
7	Solution shall provide a minimum of two 10GbE Ethernet ports per system.				
8	Shall provide the support of Jumbo frame, LACP (802.3ad) protocol, and bandwidth aggregation.				
9	Shall support OpenStorage Technology (OST) protocol.				
10	Shall support Ethernet port expansion to meet storage performance requirements in the same unit(s) as listed above (to be specified at delivery order level) for each storage size types. Capable of adding minimum of two additional ports per system.				

11	Solution shall provide storage capacity expansion or scalability with no downtime, no forklift upgrade, no reconfiguration of the internal operating system or re-creation of the existing subsystem components, and remapping connections of the storage subsystem to connecting hosts. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
12	Shall support expansion and scalability of the Data Deduplication System without requiring a forklift upgrade.				
13	Shall support that Data residing on the Deduplication device that meets the requirement to replicate to a geo-dispersed environment (one to many).				
14	Shall provide edge to core replication. Many to one replication.				
15	Shall provide Deduplication on the hardware appliance, not as a software only solution.				
16	Shall provide a fully redundant solution and support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
17	Shall provide fast instant restore of files and fast restore of a complete file system to primary storage and to tape without the need to rehydrate and reconstitute data on external or additional disk storage.				
18	Shall integrate with _____ Enterprise backup software (to be specified at the delivery order level).				
19	Shall provide the ability to be directly attached as a NAS device for a large block sequential I/O data stream.				
20	Shall meet the requirement to have a portion of the storage for saving non-deduped data and another portion for deduped data.				
21	Shall meet the requirement to select a portion of storage for data replication.				
22	Shall provide a built-in capability of self healing that provides corruption recovery during replication.				
23	Shall integrate with backup software to backup data directly to tape from the Deduplication secondary storage device.				
24	Shall Support a minimum of four nines storage subsystem uptime (or 99.99% availability, excluding planned downtime). (Note: Four Nines (99.99%) = 53 minutes of downtime on a rolling 12 month basis).				

25	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
26	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
27	Shall provide disk based data Deduplication.			
28	Shall provide built-in capability of perform automatic data Integrity checking after data Deduplication.			
29	Shall provide secure data Deduplication algorithms.			
30	Shall provide built-in capability for self healing that provides corruption recovery when data corruption is detected after Deduplication, without Re-Deduplication of the original data again.			
31	Shall provide the ability to rebuild file system from stored metadata.			
32	Shall provide the ability to ingest data from multiple sources (from software and/or hardware) simultaneously.			
33	Shall provide cover to power switches to prevent inadvertent activation.			
34	Shall provide the support of role-based security for access and management of storage devices.			
35	Shall support authenticable access with logging (for audits).			
36	Shall support a native Management tool able to manage multiple Deduplication Storage devices of the same type in data center implementations.			
37	Shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage Deduplication Storage device(s) in a data center implementation from a single master Management console.			
38	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.			
39	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and similar components.			

40	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				
----	--	--	--	--	--

VTL (Virtual Tape Library) with Data Deduplication Storage (Small)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide the minimum usable disk capacity of up to 10TB (Terabytes), expandable to 50TB (Terabytes).				
2	Shall provide a minimum ingest aggregate throughput of 2TB/hr.				
3	Solution shall support a minimum of two (2) 8-Gbit/s or higher Fibre Channel ports per device. All Fibre Channel ports shall be capable of autosensing slower speeds.				
4	Solution shall meet the capability of supporting a minimum of two (2) 1GbE Ethernet ports per device.				
5	Shall provide support for a minimum of 6 emulated tape libraries.				
6	Shall provide support for a minimum of 16 tape drives.				
7	Shall provide support for a minimum of 4096 virtual tape cartridges.				
8	Solution shall support interoperability between different sized units (Extra large, Large, Medium, and Small)				
9	Shall support concurrent parallel I/O streaming.				
10	Solution shall provide support for proactively monitoring disk drives and detecting potential failures prior to the failure occurring and taking preventive actions.				
11	Solution shall provide Deduplication function and services that integrate with existing backup environment.				
12	Shall provide the disk failure redundancy to meet the above ingest throughput requirements.				
13	Solution shall support file level restore.				
14	Solution shall not require the application to be "off-line" to complete backup activities.				
15	Solution shall support compression of backup data.				
16	Shall support Open Storage Technology (OST) protocol.				

17	Solution shall provide as backup targets for backup sources such as storage area networks, direct attached storage, network attached storage, and massive array of independent disks.				
18	Shall provide fast instant restore of files and fast restore of a complete file system to primary storage and to tape without the need to rehydrate and reconstitute data on external or additional disk storage.				
19	Solution shall support ingesting multiple backup job streams to multiple disk targets.				
20	Solution shall provide the ability to rebuild file system from stored Metadata.				
21	Shall provide a fully redundant solution and support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
22	Shall Support a minimum of four nines storage subsystem uptime (or 99.99% availability, excluding planned downtime). (Note: Four Nines (99.99%) = 53 minutes of downtime on a rolling 12 month basis).				
23	Solution shall provide the ability to support minimum of Fibre Channel (FC), Internet Small Computer System Interface (iSCSI) connectivity simultaneously on a single device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
24	Shall provide the ability to support Fibre Channel port expansion to meet the performance requirements as listed above (to be specified at delivery order level) in the same unit(s). Capable of adding additional ports per device.				
25	Solution shall meet support of Ethernet port expansion to meet the performance requirements as listed above (to be specified at delivery order level) in the same unit(s). Capable of adding additional ports per device.				
26	Solution shall support virtual tape cartridges asynchronous replication.				
27	Solution shall support virtual tape libraries asynchronous replication.				
28	Shall provide edge to core replication. Many to one replication.				
29	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				

30	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
31	Solution shall provide cover to power switches to prevent inadvertent activation.				
32	Solution shall provide the support of role-based security for access and management of the VTL storage.				
33	Shall support authenticable access with logging (for audits).				
34	Solution shall support an IP based single master VTL Storage Management Tool able to manage multiple VTL Storage devices of the same type in data center implementations.				
35	Solution shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage VTL storage device(s) in a data center implementation from a single master management console.				
36	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.				
37	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and similar components.				
38	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

VTL (Virtual Tape Library) with Data Deduplication Storage (Medium)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide the minimum usable disk capacity of 50TB (Terabytes), expandable to 100TB (Terabytes).				
2	Shall provide a minimum ingest aggregate throughput of 4TB/hr.				
3	Solution shall support a minimum of two (2) 8-Gbit/s or higher Fibre Channel ports per device. All Fibre Channel ports shall be capable of autosensing slower speeds.				
4	Solution shall meet the capability of supporting a minimum of two (2) 1GbE Ethernet ports per device.				
5	Shall provide support for a minimum of 6 emulated tape libraries.				
6	Shall provide support for a minimum of 16 tape drives.				
7	Shall provide support for a minimum of 4096 virtual tape cartridges.				
8	Solution shall support interoperability between different sized units (Extra large, Large, Medium, and Small)				
9	Shall support concurrent parallel I/O streaming.				
10	Solution shall provide support for proactively monitoring disk drives and detecting potential failures prior to the failure occurring and taking preventive actions.				
11	Solution shall provide Deduplication function and services that integrate with existing backup environment.				
12	Shall provide the disk failure redundancy to meet the above ingest throughput requirements.				
13	Solution shall support file level restore.				
14	Solution shall not require the application to be "off-line" to complete backup activities.				
15	Solution shall support compression of backup data.				
16	Shall support Open Storage Technology (OST) protocol.				

17	Solution shall provide as backup targets for backup sources such as storage area networks, direct attached storage, network attached storage, and massive array of independent disks.				
18	Shall provide fast instant restore of files and fast restore of a complete file system to primary storage and to tape without the need to rehydrate and reconstitute data on external or additional disk storage.				
19	Solution shall support ingesting multiple backup job streams to multiple disk targets.				
20	Solution shall provide the ability to rebuild file system from stored Metadata.				
21	Shall provide a fully redundant solution and support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
22	Shall Support a minimum of four nines storage subsystem uptime (or 99.99% availability, excluding planned downtime). (Note: Four Nines (99.99%) = 53 minutes of downtime on a rolling 12 month basis).				
23	Solution shall provide the ability to support minimum of Fibre Channel (FC), Internet Small Computer System Interface (iSCSI) connectivity simultaneously on a single device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
24	Shall provide the ability to support Fibre Channel port expansion to meet the performance requirements as listed above (to be specified at delivery order level) in the same unit(s). Capable of adding additional ports per device.				
25	Solution shall meet support of Ethernet port expansion to meet the performance requirements as listed above (to be specified at delivery order level) in the same unit(s). Capable of adding additional ports per device.				
26	Solution shall support virtual tape cartridges asynchronous replication.				
27	Solution shall support virtual tape libraries asynchronous replication.				
28	Shall provide edge to core replication. Many to one replication.				
29	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				

30	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
31	Solution shall provide cover to power switches to prevent inadvertent activation.				
32	Solution shall provide the support of role-based security for access and management of the VTL storage.				
33	Shall support authenticable access with logging (for audits).				
34	Solution shall support an IP based single master VTL Storage Management Tool able to manage multiple VTL Storage devices of the same type in data center implementations.				
35	Solution shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage VTL storage device(s) in a data center implementation from a single master management console.				
36	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.				
37	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and similar components.				
38	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

VTL (Virtual Tape Library) with Data Deduplication Storage (Large)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide the minimum usable disk capacity of 100TB (Terabytes), expandable to 500TB (Terabytes).				
2	Shall provide a minimum ingest aggregate throughput of 6TB/hr.				
3	Solution shall support a minimum of four (4) 8-Gbit/s or higher Fibre Channel ports per device. All Fibre Channel ports shall be capable of autosensing slower speeds.				
4	Solution shall meet the capability of supporting a minimum of four (4) 1GbE Ethernet ports per device.				
5	Shall provide support for a minimum of 12 emulated tape libraries.				
6	Shall provide support for a minimum of 32 tape drives.				
7	Shall provide support for a minimum of 8192 virtual tape cartridges.				
8	Solution shall support interoperability between different sized units (Extra large, Large, Medium, and Small)				
9	Shall support concurrent parallel I/O streaming.				
10	Solution shall provide support for proactively monitoring disk drives and detecting potential failures prior to the failure occurring and taking preventive actions.				
11	Solution shall provide Deduplication function and services that integrate with existing backup environment.				
12	Shall provide the disk failure redundancy to meet the above ingest throughput requirements.				
13	Solution shall support file level restore.				
14	Solution shall not require the application to be "off-line" to complete backup activities.				
15	Solution shall support compression of backup data.				
16	Shall support Open Storage Technology (OST) protocol.				

17	Solution shall provide as backup targets for backup sources such as storage area networks, direct attached storage, network attached storage, and massive array of independent disks.				
18	Shall provide fast instant restore of files and fast restore of a complete file system to primary storage and to tape without the need to rehydrate and reconstitute data on external or additional disk storage.				
19	Solution shall support ingesting multiple backup job streams to multiple disk targets.				
20	Solution shall provide the ability to rebuild file system from stored Metadata.				
21	Shall provide a fully redundant solution and support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
22	Shall Support a minimum of four nines storage subsystem uptime (or 99.99% availability, excluding planned downtime). (Note: Four Nines (99.99%) = 53 minutes of downtime on a rolling 12 month basis).				
23	Solution shall provide the ability to support minimum of Fibre Channel (FC), Internet Small Computer System Interface (iSCSI) connectivity simultaneously on a single device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
24	Shall provide the ability to support Fibre Channel port expansion to meet the performance requirements as listed above (to be specified at delivery order level) in the same unit(s). Capable of adding additional ports per device.				
25	Solution shall meet support of Ethernet port expansion to meet the performance requirements as listed above (to be specified at delivery order level) in the same unit(s). Capable of adding additional ports per device.				
26	Solution shall support virtual tape cartridges asynchronous replication.				
27	Solution shall support virtual tape libraries asynchronous replication.				
28	Shall provide edge to core replication. Many to one replication.				
29	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				

30	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
31	Solution shall provide cover to power switches to prevent inadvertent activation.				
32	Solution shall provide the support of role-based security for access and management of the VTL storage.				
33	Shall support authenticable access with logging (for audits).				
34	Solution shall support an IP based single master VTL Storage Management Tool able to manage multiple VTL Storage devices of the same type in data center implementations.				
35	Solution shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage VTL storage device(s) in a data center implementation from a single master management console.				
36	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.				
37	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and similar components.				
38	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

VTL (Virtual Tape Library) with Data Deduplication Storage (Extra Large)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide the minimum usable disk capacity of 500TB (Terabytes), expandable to 750TB (Terabytes).				
2	Shall provide a minimum ingest aggregate throughput of 8TB/hr.				
3	Solution shall support a minimum of four (4) 8-Gbit/s or higher Fibre Channel ports per device. All Fibre Channel ports shall be capable of autosensing slower speeds.				
4	Solution shall meet the capability of supporting a minimum of four (4) 1GbE Ethernet ports per device.				
5	Shall provide support for a minimum of 12 emulated tape libraries.				
6	Shall provide support for a minimum of 32 tape drives.				
7	Shall provide support for a minimum of 8192 virtual tape cartridges.				
8	Solution shall support interoperability between different sized units (Extra large, Large, Medium, and Small)				
9	Shall support concurrent parallel I/O streaming.				
10	Solution shall provide support for proactively monitoring disk drives and detecting potential failures prior to the failure occurring and taking preventive actions.				
11	Solution shall provide Deduplication function and services that integrate with existing backup environment.				
12	Shall provide the disk failure redundancy to meet the above ingest throughput requirements.				
13	Solution shall support file level restore.				
14	Solution shall not require the application to be "off-line" to complete backup activities.				
15	Solution shall support compression of backup data.				
16	Shall support Open Storage Technology (OST) protocol.				

17	Solution shall provide as backup targets for backup sources such as storage area networks, direct attached storage, network attached storage, and massive array of independent disks.				
18	Shall provide fast instant restore of files and fast restore of a complete file system to primary storage and to tape without the need to rehydrate and reconstitute data on external or additional disk storage.				
19	Solution shall support ingesting multiple backup job streams to multiple disk targets.				
20	Solution shall provide the ability to rebuild file system from stored Metadata.				
21	Shall provide a fully redundant solution and support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
22	Shall Support a minimum of four nines storage subsystem uptime (or 99.99% availability, excluding planned downtime). (Note: Four Nines (99.99%) = 53 minutes of downtime on a rolling 12 month basis).				
23	Solution shall provide the ability to support minimum of Fibre Channel (FC), Internet Small Computer System Interface (iSCSI) connectivity simultaneously on a single device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
24	Shall provide the ability to support Fibre Channel port expansion to meet the performance requirements as listed above (to be specified at delivery order level) in the same unit(s). Capable of adding additional ports per device.				
25	Solution shall meet support of Ethernet port expansion to meet the performance requirements as listed above (to be specified at delivery order level) in the same unit(s). Capable of adding additional ports per device.				
26	Solution shall support virtual tape cartridges asynchronous replication.				
27	Solution shall support virtual tape libraries asynchronous replication.				
28	Shall provide edge to core replication. Many to one replication.				
29	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				

30	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
31	Solution shall provide cover to power switches to prevent inadvertent activation.				
32	Solution shall provide the support of role-based security for access and management of the VTL storage.				
33	Shall support authenticable access with logging (for audits).				
34	Solution shall support an IP based single master VTL Storage Management Tool able to manage multiple VTL Storage devices of the same type in data center implementations.				
35	Solution shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage VTL storage device(s) in a data center implementation from a single master management console.				
36	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.				
37	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and similar components.				
38	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Low Performance Unified Storage					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 20,000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum of 1000 MBytes per second throughput.				
3	Shall provide the minimum usable disk capacity of 25 TB . This initial data storage capacity must be distributed evenly among two or more shelves in order for future local mirroring between shelves.				
4	Shall be capable of supporting a minimum of two (2) 8-Gbit/s Fibre Channel ports per controller. All Fibre Channel ports shall be capable of autosensing slower speeds of 2 & 4 Gbit/s.				
5	Shall be capable of supporting a minimum of two (2) 1GbE Ethernet or iSCSI ports per controller.				
6	Shall be capable of supporting a minimum of one (1) 10GbE iSCSI or Ethernet ports per controller.				
7	Shall provide, excluding space on Storage System Operating System (OS) Drives, the initial data storage capacity must be distributed evenly among multiple shelves and support mirroring between shelves.				
8	Shall provide the disk failure redundancy configuration to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
9	Shall provide the appropriate disk drive technologies to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
10	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				

11	Shall provide capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
12	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				
13	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
14	Shall provide a minimum of two storage controllers in a module for controller based storage solution, which are load balanced and provide automatic failover including the ability to maintain access to all data through a controller failure.				
15	Shall provide the ability to upgrade a controller on the storage system without replacing the storage subsystem. The upgrade of the controller shall be non-destructive to the data, requiring no data migration, no reconfiguration and no LUN remapping to any connecting Host.				
16	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
17	Shall support adding, upgrading or replacing storage subsystem components while adhering to the required storage device uptime requirement, redundancy and needing no data migration.				
18	Shall provide the ability to expand and scale Host side connectivity and backend capacity independently in order to increase performance.				
19	Shall provide storage subsystem support boot from SAN.				
20	Shall support a changeable LUN UDID value.				
21	Shall support the ability to create LUNs and Volumes across multiple spindles.				
22	Shall support dynamic LUN and Volume expansion.				
23	Shall be able to present a minimum of 15TB volume to connecting hosts.				

24	Shall provide the ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.				
25	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".				
26	Shall provide the capability of supporting Fibre Channel (FC), Internet Small Computer System Interface (iSCSI), Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
27	Shall support the ability to expand Fibre Channel ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
28	Shall be capable of supporting Jumbo frame and LACP protocol.				
29	Shall support the ability to expand iSCSI ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
30	Shall be capable of supporting native 10GbE Fibre Channel over Ethernet (FCoE) host connections.				
31	Shall be capable of supporting 10GbE Ethernet ports.				
32	Solution shall support a minimum of 256 SCSI initiators (host connections).				
33	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
34	Shall provide the capability of synchronous and asynchronous replication with write order fidelity.				
35	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
36	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
37	Shall be able to allow source update during data migration. Online data migration from other storage devices and data migration is transparent to users.				
38	Shall be capable of maintaining redundancy and the performance metrics requirements as stated above during snapshot and clone operations.				

39	Power Switches shall be covered to prevent inadvertent activation.				
40	Shall support role-based security or audit trail logging for access to storage.				
41	Shall provide a single master management interface to manage multiple storage subsystems of the same type in data center implementations.				
42	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
43	Shall have a graphical user interface (GUI) or command (CLI), or a Wizard to automate the process to provision a large number of Disk Groups, Volumes and LUNs.				
44	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
45	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
46	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Medium Performance Unified Storage					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 50,000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum of 2500 MBytes per second throughput.				
3	Shall provide the minimum usable disk capacity of 50 TB . This initial data storage capacity must be distributed evenly among two or more shelves in order for future local mirroring between shelves.				
4	Shall be capable of supporting a minimum of four (4) 8-Gbit/s Fibre Channel ports per controller. All Fibre Channel ports shall be capable of autosensing slower speeds of 2 & 4 Gbit/s.				
5	Shall be capable of supporting a minimum of four (4) 1GbE Ethernet or iSCSI ports per controller.				
6	Shall be capable of supporting a minimum of two (2) 10GbE iSCSI or Ethernet ports per controller.				
7	Shall provide, excluding space on Storage System Operating System (OS) Drives, the initial data storage capacity must be distributed evenly among multiple shelves and support mirroring between shelves.				
8	Shall provide the disk failure redundancy configuration to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
9	Shall provide the appropriate disk drive technologies to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
10	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				

11	Shall provide capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
12	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				
13	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
14	Shall provide a minimum of two storage controllers in a module for controller based storage solution, which are load balanced and provide automatic failover including the ability to maintain access to all data through a controller failure.				
15	Shall provide the ability to upgrade a controller on the storage system without replacing the storage subsystem. The upgrade of the controller shall be non-destructive to the data, requiring no data migration, no reconfiguration and no LUN remapping to any connecting Host.				
16	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
17	Shall support adding, upgrading or replacing storage subsystem components while adhering to the required storage device uptime requirement, redundancy and needing no data migration.				
18	Shall provide the ability to expand and scale Host side connectivity and backend capacity independently in order to increase performance.				
19	Shall provide storage subsystem support boot from SAN.				
20	Shall support a changeable LUN UDID value.				
21	Shall support the ability to create LUNs and Volumes across multiple spindles.				
22	Shall support dynamic LUN and Volume expansion.				
23	Shall be able to present a minimum of 15TB volume to connecting hosts.				

24	Shall provide the ability to use “Thin Provisioning” or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.				
25	Shall provide that storage subsystem performance will not be affected when using “Thin Provisioning”.				
26	Shall provide the capability of supporting Fibre Channel (FC), Internet Small Computer System Interface (iSCSI), Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
27	Shall support the ability to expand Fibre Channel ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
28	Shall be capable of supporting Jumbo frame and LACP protocol.				
29	Shall support the ability to expand iSCSI ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
30	Shall be capable of supporting native 10GbE Fibre Channel over Ethernet (FCoE) host connections.				
31	Shall be capable of supporting 10GbE Ethernet ports.				
32	Solution shall support a minimum of 256 SCSI initiators (host connections).				
33	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
34	Shall provide the capability of synchronous and asynchronous replication with write order fidelity.				
35	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
36	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
37	Shall be able to allow source update during data migration. Online data migration from other storage devices and data migration is transparent to users.				
38	Shall be capable of maintaining redundancy and the performance metrics requirements as stated above during snapshot and clone operations.				

39	Power Switches shall be covered to prevent inadvertent activation.				
40	Shall support role-based security or audit trail logging for access to storage.				
41	Shall provide a single master management interface to manage multiple storage subsystems of the same type in data center implementations.				
42	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
43	Shall have a graphical user interface (GUI) or command (CLI), or a Wizard to automate the process to provision a large number of Disk Groups, Volumes and LUNs.				
44	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
45	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
46	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

High Performance Unified Storage					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall support a minimum of 100,000 I/Os per second during all system functions including snapshot and clone operations. Solution shall utilize appropriate disk type, spindle count, and raid type to meet I/O requirements.				
2	Shall support minimum of 5000 MBytes per second throughput.				
3	Shall provide the minimum usable disk capacity of 75 TB . This initial data storage capacity must be distributed evenly among two or more shelves in order for future local mirroring between shelves.				
4	Shall be capable of supporting a minimum of four (4) 8-Gbit/s Fibre Channel ports per controller. All Fibre Channel ports shall be capable of autosensing slower speeds of 2 & 4 Gbit/s.				
5	Shall be capable of supporting a minimum of four (4) 1GbE Ethernet or iSCSI ports per controller.				
6	Shall be capable of supporting a minimum of two (2) 10GbE iSCSI or Ethernet ports per controller.				
7	Shall provide, excluding space on Storage System Operating System (OS) Drives, the initial data storage capacity must be distributed evenly among multiple shelves and support mirroring between shelves.				
8	Shall provide the disk failure redundancy configuration to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
9	Shall provide the appropriate disk drive technologies to meet the above capacity, IOPs and workload performance requirements (to be specified at delivery order level).				
10	Shall provide capability of creating a minimum of 200 space efficient pointer based point in time copies per Storage System.				

11	Shall provide capability of creating point in time delta clone (Delta clone is a copy based on pointers to the original data, which is updated when changes occur to the original data), Clones shall be space efficient on an as needed basis with no space reservation.				
12	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				
13	Shall provide redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
14	Shall provide a minimum of two storage controllers in a module for controller based storage solution, which are load balanced and provide automatic failover including the ability to maintain access to all data through a controller failure.				
15	Shall provide the ability to upgrade a controller on the storage system without replacing the storage subsystem. The upgrade of the controller shall be non-destructive to the data, requiring no data migration, no reconfiguration and no LUN remapping to any connecting Host.				
16	Shall provide storage capacity expansion with no reconfiguration or LUN remapping of the storage device. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
17	Shall support adding, upgrading or replacing storage subsystem components while adhering to the required storage device uptime requirement, redundancy and needing no data migration.				
18	Shall provide the ability to expand and scale Host side connectivity and backend capacity independently in order to increase performance.				
19	Shall provide storage subsystem support boot from SAN.				
20	Shall support a changeable LUN UDID value.				
21	Shall support the ability to create LUNs and Volumes across multiple spindles.				
22	Shall support dynamic LUN and Volume expansion.				
23	Shall be able to present a minimum of 15TB volume to connecting hosts.				

24	Shall provide the ability to use "Thin Provisioning" or to over provision/over allocate storage capacity to hosts, allowing hosts to view more logical storage capacity than has been physically reserved on the storage array.				
25	Shall provide that storage subsystem performance will not be affected when using "Thin Provisioning".				
26	Shall provide the capability of supporting Fibre Channel (FC), Internet Small Computer System Interface (iSCSI), Network File System V4 (NFS), and Common Internet File System (CIFS) connectivity simultaneously on a single storage device. All connecting protocols shall support up to the maximum usable storage capacity of the device.				
27	Shall support the ability to expand Fibre Channel ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
28	Shall be capable of supporting Jumbo frame and LACP protocol.				
29	Shall support the ability to expand iSCSI ports non-disruptively, when adding additional ports per controller and still maintain storage performance requirements in the same system.				
30	Shall be capable of supporting native 10GbE Fibre Channel over Ethernet (FCoE) host connections.				
31	Shall be capable of supporting 10GbE Ethernet ports.				
32	Solution shall support a minimum of 256 SCSI initiators (host connections).				
33	Shall provide load balance multi path failover software and multi mount points when using Ethernet.				
34	Shall provide the capability of synchronous and asynchronous replication with write order fidelity.				
35	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
36	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
37	Shall be able to allow source update during data migration. Online data migration from other storage devices and data migration is transparent to users.				
38	Shall be capable of maintaining redundancy and the performance metrics requirements as stated above during snapshot and clone operations.				

39	Power Switches shall be covered to prevent inadvertent activation.				
40	Shall support role-based security or audit trail logging for access to storage.				
41	Shall provide a single master management interface to manage multiple storage subsystems of the same type in data center implementations.				
42	Shall provide a single sign-on integrated with Microsoft Active Directory (AD)/or LDAP to manage storage device(s) in a data center implementation from a single master management console.				
43	Shall have a graphical user interface (GUI) or command (CLI), or a Wizard to automate the process to provision a large number of Disk Groups, Volumes and LUNs.				
44	Shall have the native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information shall be capable of being transmitted to the Vendor or OEM as defined by the contract and designated VA Point of Contact.				
45	Shall have the native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans, Power Supplies and similar components.				
46	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Archive Storage (Small)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide an initial minimum usable disk capacity of 50TB (Terabytes), expandable to 100TB (Terabytes).				
2	Shall provide intelligent power management system with the ability to enter a disk power reduction mode for inactive data.				
3	Shall integrate and present as a target to leading major Enterprise backup software.				
4	Shall provide active background data integrity verification.				
5	Shall be able to proactively monitor and provide background disk scrubbing to detect potential drive failures and automatically move data to spare drive from a suspected faulty disk drive.				
6	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				
7	Shall provide a minimum aggregate ingestion throughput rate of 5.2TB/hr per system.				
8	Shall provide concurrent parallel I/O streaming.				
9	Shall support a dense storage capacity and small storage foot print by providing a minimum of 1PB (Petabyte) of storage in one cabinet.				
10	Shall support native or host/appliance based write-once-read-many (WORM) operation.				
11	Shall support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				

12	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
13	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
14	Shall provide host/appliance based or native Asynchronous replication.			
15	Shall support CIFS, NFS V4, and TCP/IP connectivity.			
16	Shall provide a minimum of (2) 8Gbit or higher speed fibre ports per system and be capable of autosensing slower speeds.			
17	Shall provide a Minimum of (2) 1GbE or higher Ethernet ports per system.			
18	Shall provide the ability to support port expansion non-disruptively to meet the storage performance requirements as listed above (to be specified at delivery order level) in the same unit(s).			
19	Power Switches shall be covered to prevent inadvertent activation.			
20	Shall provide the support of role-based security for access and management of storage devices.			
21	Shall support authenticable access with logging (for audits).			
22	Shall provide a single master management interface to manage multiple storage devices of the same type in data center implementations.			
23	Shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage storage device(s) in a data center implementation from a single master management console.			
24	Shall have a graphical user interface (GUI) or command line interface (CLI) or wizard to automate the process to provision a large number of Disk Groups, Volumes or LUNs.			
25	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.			

26	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and similar components.				
27	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Archive Storage (Medium)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide minimum usable disk capacity of 500TB (Terabytes), expandable to 1PB (Petabyte).				
2	Shall provide intelligent power management system with the ability to enter a disk power reduction mode for inactive data.				
3	Shall integrate and present as a target to leading major Enterprise backup software.				
4	Shall provide active background data integrity verification.				
5	Shall be able to proactively monitor and provide background disk scrubbing to detect potential drive failures and automatically move data to spare drive from a suspected faulty disk drive.				
6	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				
7	Shall provide a minimum aggregate ingestion throughput rate of 5.2TB/hr per system.				
8	Shall provide concurrent parallel I/O streaming.				
9	Shall support a dense storage capacity and small storage foot print by providing a minimum of 1PB (Petabyte) of storage in one cabinet.				
10	Shall support native or host/appliance based write-once-read-many (WORM) operation.				
11	Shall support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				

12	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
13	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
14	Shall provide host/appliance based or native Asynchronous replication.			
15	Shall support CIFS, NFS V4, and TCP/IP connectivity.			
16	Shall provide a minimum of (2) 8Gbit or higher speed fibre ports per system and be capable of autosensing slower speeds.			
17	Shall provide a Minimum of (2) 1GbE or higher Ethernet ports per system.			
18	Shall provide the ability to support port expansion non-disruptively to meet the storage performance requirements as listed above (to be specified at delivery order level) in the same unit(s).			
19	Power Switches shall be covered to prevent inadvertent activation.			
20	Shall provide the support of role-based security for access and management of storage devices.			
21	Shall support authenticable access with logging (for audits).			
22	Shall provide a single master management interface to manage multiple storage devices of the same type in data center implementations.			
23	Shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage storage device(s) in a data center implementation from a single master management console.			
24	Shall have a graphical user interface (GUI) or command line interface (CLI) or wizard to automate the process to provision a large number of Disk Groups, Volumes or LUNs.			
25	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.			

26	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and similar components.				
27	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

Archive Storage (Large)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Shall provide minimum usable disk capacity of 1PB (Petabyte), or more.				
2	Shall provide intelligent power management system with the ability to enter a disk power reduction mode for inactive data.				
3	Shall integrate and present as a target to leading major Enterprise backup software.				
4	Shall provide active background data integrity verification.				
5	Shall be able to proactively monitor and provide background disk scrubbing to detect potential drive failures and automatically move data to spare drive from a suspected faulty disk drive.				
6	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis).				
7	Shall provide a minimum aggregate ingestion throughput rate of 5.2TB/hr per system.				
8	Shall provide concurrent parallel I/O streaming.				
9	Shall support a dense storage capacity and small storage foot print by providing a minimum of 1PB (Petabyte) of storage in one cabinet.				
10	Shall support native or host/appliance based write-once-read-many (WORM) operation.				
11	Shall support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				

12	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
13	Solution shall provide the ability to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
14	Shall provide host/appliance based or native Asynchronous replication.			
15	Shall support CIFS, NFS V4, and TCP/IP connectivity.			
16	Shall provide a minimum of (2) 8Gbit or higher speed fibre ports per system and be capable of autosensing slower speeds.			
17	Shall provide a Minimum of (2) 1GbE or higher Ethernet ports per system.			
18	Shall provide the ability to support port expansion non-disruptively to meet the storage performance requirements as listed above (to be specified at delivery order level) in the same unit(s).			
19	Power Switches shall be covered to prevent inadvertent activation.			
20	Shall provide the support of role-based security for access and management of storage devices.			
21	Shall support authenticable access with logging (for audits).			
22	Shall provide a single master management interface to manage multiple storage devices of the same type in data center implementations.			
23	Shall provide a single sign-on integrated with Microsoft Active Directory (AD) or LDAP to manage storage device(s) in a data center implementation from a single master management console.			
24	Shall have a graphical user interface (GUI) or command line interface (CLI) or wizard to automate the process to provision a large number of Disk Groups, Volumes or LUNs.			
25	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact.			

26	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and similar components.				
27	Shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process.				

GRID Based Object Storage (Small)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Solution shall provide minimum usable native disk capacity without data compression, without data deduplication, and without single instance (no duplicate objects) of 20TB (Terabytes), expandable to 50TB (Terabytes).				
2	Solution shall provide a minimum of Enterprise grade SATA spinning disks or better as the backend storage devices to store the object data.				
3	The solution shall not include removable media, such as optical, tapes and removable disk cartridge.				
4	Solution shall support minimum of Representational State Transfer (REST), Common Internet File System (CIFS), and Network File System (NFS) protocol.				

5	<p>Solution shall provide an intelligent policy-based object storage system. Solution shall be fully automated, policy driven without manual intervention, script, or wizard based automation to manage data retention, replication, and data distribution policies.</p> <p>For example, if the policy is defined to have 1 copy of the data locally at the primary site and 1 copy of the data at the remote site. In the event when primary site is not able to communicate with the remote site due to a temporary outage at the remote site, the solution shall automatically create a temporary remote copy at a remote tertiary site intelligently or a user defined temporary remote tertiary location during the creation of the policy process to meet the policy requirement (1 local and 1 remote). When the communication to the original remote site is restored, the solution will update and re-synchronize the missing data automatically to the original remote site and automatically delete the temporary remote copy at the temporary remote tertiary location. The solution shall automatically perform self healing of the system during a non-disaster type of outage. An alert and logging mechanism shall be included in the solution to capture the sequence of activities and provide the final status via email to the designated recipient(s).</p>				
6	<p>Solution shall provide customizable policy to define the number of synchronous local data copies, the number of asynchronous remote data copies, the number of sites to store the data, the primary site location, the disaster recover (DR) site location and the temporary remote tertiary site location.</p>				
7	<p>The solution shall have the ability to provide automatic or user definable thresholds in the policy to determine temporary interruptions (such as system resource and/or communication problems) when accessing the DR site. This capability is used to determine when the remote copy of the data shall be copied to the temporary tertiary site.</p>				
8	<p>The Contractor shall provide a single Graphic User Interface (GUI) that allows more than one creations, editions, and deletions of data retention, replication, and data distribution policies in a one single process. For example, data retention, replication, and data distribution policies can be configured in one GUI interface and not to have separate individual interfaces for data retention, replication, and data distribution respectively.</p>				

9	Solution shall replicate files to the remote site based on available WAN bandwidth and remote system resources. The process shall be fully automated, policy driven without manual intervention, script or wizard based automation.			
10	Solution shall provide the necessary equipments at the secondary remote site so that the remote site can be promoted to primary site role in the event of total disaster at the primary site.			
11	In the event of a site recovery, the solution shall be capable of creating a duplicate system by copying all data from the Disaster Recovery site which will be shipped to the disaster site to be recovered.			
12	Solution shall provide storage capacity expansion or scalability with no downtime, no forklift upgrade, no reconfiguration of the internal operating system, and LUN remapping of the storage subsystem to connecting hosts. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.			
13	Solution shall support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.			
14	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
15	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
16	Solution shall provide File Transfer ingest rate to local storage solution at a minimum of 1.5 Gigabytes (GB) of data per minute (1.5GB/min).			
17	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis)			

18	Solution shall always maintain original file attributes (Create, Access dates, file permission, and security or other attributes) for all data replication process, data copy process including data restoration processes. The file attributes are to stay with the file.				
19	Solution shall provide the ability to verify data after it has been transferred from the local site to the remote location to ensure it matches with no data corruption. If verification of data fails solution shall resend the data. The copy of file at the remote site shall be identical to the copy of the image file at the local site after replication. This check shall be performed automatically.				
20	Solution shall automatically (based on policy, not manual, script or wizard based) redirect the read processes to remote storage devices/components if the local storage devices/components are not available during image file read process.				
21	Solution shall provide a solution that can automatically (based on policy, not manual, script or wizard based) redirect the write processes to remote storage node and/or backend storage if the local storage node and/or backend storage is not available during image file write process.				
22	The solution shall present a CIFS interface that is addressable by a Microsoft Active Directory fully qualified path. The Solution shall present a redundant CIFS interface. It must be presented to any windows client on the network. Any windows client on the network must be able to authenticate to the server and access to the share.				
23	The Contractor shall provide a redundant Global Namespace object based Archive system. All objects shall be retrievable via web browser (http and https) based on its object metadata definitions.				
24	Solution shall maintain mappings of the Common Internet File System (CIFS) share in the event that the primary local gateway is unavailable or offline. The solution shall continue to function automatically with no manual intervention, script or wizard based automation on local secondary gateway with the same CIFS mapping. The path used to access the CIFS shares will be the same regardless of whether the local or remote copy of the image is accessed.				

25	Solution shall provide a process automatically (based on policy, not manual, script or wizard based) update and resynchronize the recovered local site archive storage from the remote archive storage when the local archive storage is back online after a period of downtime. The process shall be fully automated, policy driven without manual intervention, script or wizard based automation			
26	Solution shall provide a process to log all file transfer success or failure.			
27	Solution shall utilize Microsoft Active Directory for all security permissions for the whole system management and data access.			
28	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system components problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact			
29	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and other similar components.			
30	Solution shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process and SMTP mail.			
31	Solution shall escalate alerts when storage space utilization falls below a user defined threshold.			
32	Solution shall provide a means for the solution to log the file transfer time from local archive to remote site as a user defined threshold. (e.g. If logging transfer time is set to 60 seconds and the actual transfer time exceeds 60 seconds it shall be logged).			
33	Solution shall support the function to access remote data from the local archive location for display purposes.			
34	Solution shall provide a means for the Solution to retrieve a copy of the data from the remote archive locations for storage at the local archive location if the local archive was not available.			
35	Solution shall log file access, create and modify dates for retention statistics and monitoring purposes.			
36	Solution shall under normal operating conditions, replicate data immediately. In no case shall the replication backlog be greater than four (4) hours of data.			

37	Solution shall be able to scale components such as network interface cards, adapter ports, processors, memory, and storage capacity independently.				
38	Solution shall monitor and access the remote data with the least amount of file transfer latency.				
39	Solution shall manage, control, and define policies (policy that defines the number of data replication copies, the location, and the temporary remote tertiary site location) in the Object storage device.				
40	Solution shall be capable to capture and/or log performance data such as data ingest rate, data ingest size, data replication bandwidth utilization, data replication latency, data retrieval response time and others.				
41	Solution shall support a minimum of two 1Gb Ethernet connectivity to the local production LAN.				
42	Solution shall provide redundant Power Distribution Unit (PDUs) with sufficient outlets to power all units in the rack(s).				
43	Solution equipment shall be Pre-installed and Pre-cabled in a 42U Rack prior to delivery. All cables and components installed in the rack shall be labeled.				
44	Solution shall include Pre-installed and Pre-cabled redundant top of the rack network switches in every rack provided to minimize the number of network connections to VA Core network switches. The top of the rack switches provided in the solution shall be interoperable with existing VA Core network switches.				
45	Solution shall support Internet Protocol version 6 (IPv6).				

46	<p>Solution shall support a minimum of the following file formats:</p> <p>AVI Motion Video AVI BMP Bitmap file BMP DCM Dicom file DCM HTML Web Document HTML JPG Full Color JPG MHTML MIME HTML Document MP3 Motion Video MPEG-3 MP4 Motion Video MPEG-4 MPG Motion Video MPG PDF Adobe PDF RTF Rich Text format RTF TGA Targa Image TGA TIF Scanned Document TIF TXT Text file import TXT WAV Audio file WAV</p>				
47	Solution shall make the file transfer process transparent to end users with no user intervention needed. The physical location of the data repository shall be transparent to end user.				
48	Solution shall perform automatic restoration of a corrupted file (Object/File Data and Meta Data) from a Disaster Recovery (DR) copy.				
49	Solution shall automatically resume or recover from file transfer failures.				
50	Solution shall provide capability to log events and problems and provide the ability to manually set the level of logging and reporting.				
51	Solution shall provide a role based configuration with a minimum of a system user role which allows for a system user to view current hardware status and a Super User system administration role which allows for an administrator to configure system policies, and parameters.				
52	Solution shall provide a web based single master management administrative console without the need to install client software for control of system software implementation, policies and user management protected with permission roles.				
53	Solution shall utilize Microsoft Active Directory to provide a single sign-on to the management web console and the log on security shall provide the permission to access and manage all the CIFS shares and data.				
54	Solution shall log the software version/installed hardware/firmware versions at each site.				

55	Solution shall provide interface callable from an application using HTTP/S Web Services using SOAP or REST and Get/Post methods for moving binary data. The ICD/WSDL will be included with the solution.				
56	Solution shall provide a web interface to access the metadata of the files/objects that are stored in the storage device.				
57	Solution shall provide web interface to access the system logs generated by all components. If separate logs are generated the interface will include the ability to determine what logs are available and parameters to select specific logs.				
58	Solution shall provide a web interface to access files by file path/name.				
59	Solution shall provide a web interface to store files by file path/name.				
60	Solution shall provide Web interface that provides usage statistics (metrics on all nodes –size, amount full, number of files, as examples).				
61	Solution shall provide web interface that provides access to the list of files by attributes (MIME type, date created/range, size, file name or partial file name).				
62	Solution shall provide a web interface that will export/copy files by attributes (e.g. File path or partial file path) to a specified target CFIS Share.				
63	Solution shall provide a web interface that provides file statistics (e.g., size, creation date, modify date, access date) given File name or file path, or File/Object ID specified by the solution.				
64	Solution shall provide web interface that provides image retrieval time statistics given an Object/File ID or file path.				
65	Solution shall provide a web interface that provides the “Health status” of the storage (i.e. what is the current status of each of the nodes on the network or of a particular node).				
66	Solution shall provide a web interface that provides status of devices (e.g. any disk errors or outages.)				
67	Solution shall provide web interface that provides the physical location of Storage devices given a device ID supplied by the vendor, or the list of storage devices and characteristics for all devices at a location.				

GRID Based Object Storage (Medium)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Solution shall provide minimum usable native disk capacity without data compression, without data deduplication, and without single instance (no duplicate objects) of 50TB (Terabytes), expandable to 100TB (Terabytes).				
2	Solution shall provide a minimum of Enterprise grade SATA spinning disks or better as the backend storage devices to store the object data.				
3	The solution shall not include removable media, such as optical, tapes and removable disk cartridge.				
4	Solution shall support minimum of Representational State Transfer (REST), Common Internet File System (CIFS), and Network File System (NFS) protocol.				

5	<p>Solution shall provide an intelligent policy-based object storage system. Solution shall be fully automated, policy driven without manual intervention, script, or wizard based automation to manage data retention, replication, and data distribution policies.</p> <p>For example, if the policy is defined to have 1 copy of the data locally at the primary site and 1 copy of the data at the remote site. In the event when primary site is not able to communicate with the remote site due to a temporary outage at the remote site, the solution shall automatically create a temporary remote copy at a remote tertiary site intelligently or a user defined temporary remote tertiary location during the creation of the policy process to meet the policy requirement (1 local and 1 remote). When the communication to the original remote site is restored, the solution will update and re-synchronize the missing data automatically to the original remote site and automatically delete the temporary remote copy at the temporary remote tertiary location. The solution shall automatically perform self healing of the system during a non-disaster type of outage. An alert and logging mechanism shall be included in the solution to capture the sequence of activities and provide the final status via email to the designated recipient(s).</p>				
6	<p>Solution shall provide customizable policy to define the number of synchronous local data copies, the number of asynchronous remote data copies, the number of sites to store the data, the primary site location, the disaster recover (DR) site location and the temporary remote tertiary site location.</p>				
7	<p>The solution shall have the ability to provide automatic or user definable thresholds in the policy to determine temporary interruptions (such as system resource and/or communication problems) when accessing the DR site. This capability is used to determine when the remote copy of the data shall be copied to the temporary tertiary site.</p>				
8	<p>The Contractor shall provide a single Graphic User Interface (GUI) that allows more than one creations, editions, and deletions of data retention, replication, and data distribution policies in a one single process. For example, data retention, replication, and data distribution policies can be configured in one GUI interface and not to have separate individual interfaces for data retention, replication, and data distribution respectively.</p>				

9	Solution shall replicate files to the remote site based on available WAN bandwidth and remote system resources. The process shall be fully automated, policy driven without manual intervention, script or wizard based automation.				
10	Solution shall provide the necessary equipments at the secondary remote site so that the remote site can be promoted to primary site role in the event of total disaster at the primary site.				
11	In the event of a site recovery, the solution shall be capable of creating a duplicate system by copying all data from the Disaster Recovery site which will be shipped to the disaster site to be recovered.				
12	Solution shall provide storage capacity expansion or scalability with no downtime, no forklift upgrade, no reconfiguration of the internal operating system, and LUN remapping of the storage subsystem to connecting hosts. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
13	Solution shall support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
14	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
15	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
16	Solution shall provide File Transfer ingest rate to local storage solution at a minimum of 1.5 Gigabytes (GB) of data per minute (1.5GB/min).				
17	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis)				

18	Solution shall always maintain original file attributes (Create, Access dates, file permission, and security or other attributes) for all data replication process, data copy process including data restoration processes. The file attributes are to stay with the file.				
19	Solution shall provide the ability to verify data after it has been transferred from the local site to the remote location to ensure it matches with no data corruption. If verification of data fails solution shall resend the data. The copy of file at the remote site shall be identical to the copy of the image file at the local site after replication. This check shall be performed automatically.				
20	Solution shall automatically (based on policy, not manual, script or wizard based) redirect the read processes to remote storage devices/components if the local storage devices/components are not available during image file read process.				
21	Solution shall provide a solution that can automatically (based on policy, not manual, script or wizard based) redirect the write processes to remote storage node and/or backend storage if the local storage node and/or backend storage is not available during image file write process.				
22	The solution shall present a CIFS interface that is addressable by a Microsoft Active Directory fully qualified path. The Solution shall present a redundant CIFS interface. It must be presented to any windows client on the network. Any windows client on the network must be able to authenticate to the server and access to the share.				
23	The Contractor shall provide a redundant Global Namespace object based Archive system. All objects shall be retrievable via web browser (http and https) based on its object metadata definitions.				
24	Solution shall maintain mappings of the Common Internet File System (CIFS) share in the event that the primary local gateway is unavailable or offline. The solution shall continue to function automatically with no manual intervention, script or wizard based automation on local secondary gateway with the same CIFS mapping. The path used to access the CIFS shares will be the same regardless of whether the local or remote copy of the image is accessed.				

25	Solution shall provide a process automatically (based on policy, not manual, script or wizard based) update and resynchronize the recovered local site archive storage from the remote archive storage when the local archive storage is back online after a period of downtime. The process shall be fully automated, policy driven without manual intervention, script or wizard based automation			
26	Solution shall provide a process to log all file transfer success or failure.			
27	Solution shall utilize Microsoft Active Directory for all security permissions for the whole system management and data access.			
28	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system components problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact			
29	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and other similar components.			
30	Solution shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process and SMTP mail.			
31	Solution shall escalate alerts when storage space utilization falls below a user defined threshold.			
32	Solution shall provide a means for the solution to log the file transfer time from local archive to remote site as a user defined threshold. (e.g. If logging transfer time is set to 60 seconds and the actual transfer time exceeds 60 seconds it shall be logged).			
33	Solution shall support the function to access remote data from the local archive location for display purposes.			
34	Solution shall provide a means for the Solution to retrieve a copy of the data from the remote archive locations for storage at the local archive location if the local archive was not available.			
35	Solution shall log file access, create and modify dates for retention statistics and monitoring purposes.			
36	Solution shall under normal operating conditions, replicate data immediately. In no case shall the replication backlog be greater than four (4) hours of data.			

37	Solution shall be able to scale components such as network interface cards, adapter ports, processors, memory, and storage capacity independently.				
38	Solution shall monitor and access the remote data with the least amount of file transfer latency.				
39	Solution shall manage, control, and define policies (policy that defines the number of data replication copies, the location, and the temporary remote tertiary site location) in the Object storage device.				
40	Solution shall be capable to capture and/or log performance data such as data ingest rate, data ingest size, data replication bandwidth utilization, data replication latency, data retrieval response time and others.				
41	Solution shall support a minimum of two 1Gb Ethernet connectivity to the local production LAN.				
42	Solution shall provide redundant Power Distribution Unit (PDUs) with sufficient outlets to power all units in the rack(s).				
43	Solution equipment shall be Pre-installed and Pre-cabled in a 42U Rack prior to delivery. All cables and components installed in the rack shall be labeled.				
44	Solution shall include Pre-installed and Pre-cabled redundant top of the rack network switches in every rack provided to minimize the number of network connections to VA Core network switches. The top of the rack switches provided in the solution shall be interoperable with existing VA Core network switches.				
45	Solution shall support Internet Protocol version 6 (IPv6).				

46	<p>Solution shall support a minimum of the following file formats:</p> <p>AVI Motion Video AVI BMP Bitmap file BMP DCM Dicom file DCM HTML Web Document HTML JPG Full Color JPG MHTML MIME HTML Document MP3 Motion Video MPEG-3 MP4 Motion Video MPEG-4 MPG Motion Video MPG PDF Adobe PDF RTF Rich Text format RTF TGA Targa Image TGA TIF Scanned Document TIF TXT Text file import TXT WAV Audio file WAV</p>				
47	<p>Solution shall make the file transfer process transparent to end users with no user intervention needed. The physical location of the data repository shall be transparent to end user.</p>				
48	<p>Solution shall perform automatic restoration of a corrupted file (Object/File Data and Meta Data) from a Disaster Recovery (DR) copy.</p>				
49	<p>Solution shall automatically resume or recover from file transfer failures.</p>				
50	<p>Solution shall provide capability to log events and problems and provide the ability to manually set the level of logging and reporting.</p>				
51	<p>Solution shall provide a role based configuration with a minimum of a system user role which allows for a system user to view current hardware status and a Super User system administration role which allows for an administrator to configure system policies, and parameters.</p>				
52	<p>Solution shall provide a web based single master management administrative console without the need to install client software for control of system software implementation, policies and user management protected with permission roles.</p>				
53	<p>Solution shall utilize Microsoft Active Directory to provide a single sign-on to the management web console and the log on security shall provide the permission to access and manage all the CIFS shares and data.</p>				
54	<p>Solution shall log the software version/installed hardware/firmware versions at each site.</p>				

55	Solution shall provide interface callable from an application using HTTP/S Web Services using SOAP or REST and Get/Post methods for moving binary data. The ICD/WSDL will be included with the solution.				
56	Solution shall provide a web interface to access the metadata of the files/objects that are stored in the storage device.				
57	Solution shall provide web interface to access the system logs generated by all components. If separate logs are generated the interface will include the ability to determine what logs are available and parameters to select specific logs.				
58	Solution shall provide a web interface to access files by file path/name.				
59	Solution shall provide a web interface to store files by file path/name.				
60	Solution shall provide Web interface that provides usage statistics (metrics on all nodes –size, amount full, number of files, as examples).				
61	Solution shall provide web interface that provides access to the list of files by attributes (MIME type, date created/range, size, file name or partial file name).				
62	Solution shall provide a web interface that will export/copy files by attributes (e.g. File path or partial file path) to a specified target CFIS Share.				
63	Solution shall provide a web interface that provides file statistics (e.g., size, creation date, modify date, access date) given File name or file path, or File/Object ID specified by the solution.				
64	Solution shall provide web interface that provides image retrieval time statistics given an Object/File ID or file path.				
65	Solution shall provide a web interface that provides the “Health status” of the storage (i.e. what is the current status of each of the nodes on the network or of a particular node).				
66	Solution shall provide a web interface that provides status of devices (e.g. any disk errors or outages.)				
67	Solution shall provide web interface that provides the physical location of Storage devices given a device ID supplied by the vendor, or the list of storage devices and characteristics for all devices at a location.				

GRID Based Object Storage (Large)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Solution shall provide minimum usable native disk capacity without data compression, without data deduplication, and without single instance (no duplicate objects) of 100TB (Terabytes), expandable to 200TB (Terabytes).				
2	Solution shall provide a minimum of Enterprise grade SATA spinning disks or better as the backend storage devices to store the object data.				
3	The solution shall not include removable media, such as optical, tapes and removable disk cartridge.				
4	Solution shall support minimum of Representational State Transfer (REST), Common Internet File System (CIFS), and Network File System (NFS) protocol.				

5	<p>Solution shall provide an intelligent policy-based object storage system. Solution shall be fully automated, policy driven without manual intervention, script, or wizard based automation to manage data retention, replication, and data distribution policies.</p> <p>For example, if the policy is defined to have 1 copy of the data locally at the primary site and 1 copy of the data at the remote site. In the event when primary site is not able to communicate with the remote site due to a temporary outage at the remote site, the solution shall automatically create a temporary remote copy at a remote tertiary site intelligently or a user defined temporary remote tertiary location during the creation of the policy process to meet the policy requirement (1 local and 1 remote). When the communication to the original remote site is restored, the solution will update and re-synchronize the missing data automatically to the original remote site and automatically delete the temporary remote copy at the temporary remote tertiary location. The solution shall automatically perform self healing of the system during a non-disaster type of outage. An alert and logging mechanism shall be included in the solution to capture the sequence of activities and provide the final status via email to the designated recipient(s).</p>				
6	<p>Solution shall provide customizable policy to define the number of synchronous local data copies, the number of asynchronous remote data copies, the number of sites to store the data, the primary site location, the disaster recover (DR) site location and the temporary remote tertiary site location.</p>				
7	<p>The solution shall have the ability to provide automatic or user definable thresholds in the policy to determine temporary interruptions (such as system resource and/or communication problems) when accessing the DR site. This capability is used to determine when the remote copy of the data shall be copied to the temporary tertiary site.</p>				
8	<p>The Contractor shall provide a single Graphic User Interface (GUI) that allows more than one creations, editions, and deletions of data retention, replication, and data distribution policies in a one single process. For example, data retention, replication, and data distribution policies can be configured in one GUI interface and not to have separate individual interfaces for data retention, replication, and data distribution respectively.</p>				

9	Solution shall replicate files to the remote site based on available WAN bandwidth and remote system resources. The process shall be fully automated, policy driven without manual intervention, script or wizard based automation.			
10	Solution shall provide the necessary equipments at the secondary remote site so that the remote site can be promoted to primary site role in the event of total disaster at the primary site.			
11	In the event of a site recovery, the solution shall be capable of creating a duplicate system by copying all data from the Disaster Recovery site which will be shipped to the disaster site to be recovered.			
12	Solution shall provide storage capacity expansion or scalability with no downtime, no forklift upgrade, no reconfiguration of the internal operating system, and LUN remapping of the storage subsystem to connecting hosts. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.			
13	Solution shall support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.			
14	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
15	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.			
16	Solution shall provide File Transfer ingest rate to local storage solution at a minimum of 1.5 Gigabytes (GB) of data per minute (1.5GB/min).			
17	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis)			

18	Solution shall always maintain original file attributes (Create, Access dates, file permission, and security or other attributes) for all data replication process, data copy process including data restoration processes. The file attributes are to stay with the file.				
19	Solution shall provide the ability to verify data after it has been transferred from the local site to the remote location to ensure it matches with no data corruption. If verification of data fails solution shall resend the data. The copy of file at the remote site shall be identical to the copy of the image file at the local site after replication. This check shall be performed automatically.				
20	Solution shall automatically (based on policy, not manual, script or wizard based) redirect the read processes to remote storage devices/components if the local storage devices/components are not available during image file read process.				
21	Solution shall provide a solution that can automatically (based on policy, not manual, script or wizard based) redirect the write processes to remote storage node and/or backend storage if the local storage node and/or backend storage is not available during image file write process.				
22	The solution shall present a CIFS interface that is addressable by a Microsoft Active Directory fully qualified path. The Solution shall present a redundant CIFS interface. It must be presented to any windows client on the network. Any windows client on the network must be able to authenticate to the server and access to the share.				
23	The Contractor shall provide a redundant Global Namespace object based Archive system. All objects shall be retrievable via web browser (http and https) based on its object metadata definitions.				
24	Solution shall maintain mappings of the Common Internet File System (CIFS) share in the event that the primary local gateway is unavailable or offline. The solution shall continue to function automatically with no manual intervention, script or wizard based automation on local secondary gateway with the same CIFS mapping. The path used to access the CIFS shares will be the same regardless of whether the local or remote copy of the image is accessed.				

25	Solution shall provide a process automatically (based on policy, not manual, script or wizard based) update and resynchronize the recovered local site archive storage from the remote archive storage when the local archive storage is back online after a period of downtime. The process shall be fully automated, policy driven without manual intervention, script or wizard based automation			
26	Solution shall provide a process to log all file transfer success or failure.			
27	Solution shall utilize Microsoft Active Directory for all security permissions for the whole system management and data access.			
28	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system components problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact			
29	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and other similar components.			
30	Solution shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process and SMTP mail.			
31	Solution shall escalate alerts when storage space utilization falls below a user defined threshold.			
32	Solution shall provide a means for the solution to log the file transfer time from local archive to remote site as a user defined threshold. (e.g. If logging transfer time is set to 60 seconds and the actual transfer time exceeds 60 seconds it shall be logged).			
33	Solution shall support the function to access remote data from the local archive location for display purposes.			
34	Solution shall provide a means for the Solution to retrieve a copy of the data from the remote archive locations for storage at the local archive location if the local archive was not available.			
35	Solution shall log file access, create and modify dates for retention statistics and monitoring purposes.			
36	Solution shall under normal operating conditions, replicate data immediately. In no case shall the replication backlog be greater than four (4) hours of data.			

37	Solution shall be able to scale components such as network interface cards, adapter ports, processors, memory, and storage capacity independently.			
38	Solution shall monitor and access the remote data with the least amount of file transfer latency.			
39	Solution shall manage, control, and define policies (policy that defines the number of data replication copies, the location, and the temporary remote tertiary site location) in the Object storage device.			
40	Solution shall be capable to capture and/or log performance data such as data ingest rate, data ingest size, data replication bandwidth utilization, data replication latency, data retrieval response time and others.			
41	Solution shall support a minimum of two 1Gb Ethernet connectivity to the local production LAN.			
42	Solution shall provide redundant Power Distribution Unit (PDUs) with sufficient outlets to power all units in the rack(s).			
43	Solution equipment shall be Pre-installed and Pre-cabled in a 42U Rack prior to delivery. All cables and components installed in the rack shall be labeled.			
44	Solution shall include Pre-installed and Pre-cabled redundant top of the rack network switches in every rack provided to minimize the number of network connections to VA Core network switches. The top of the rack switches provided in the solution shall be interoperable with existing VA Core network switches.			
45	Solution shall support Internet Protocol version 6 (IPv6).			

46	<p>Solution shall support a minimum of the following file formats:</p> <p>AVI Motion Video AVI BMP Bitmap file BMP DCM Dicom file DCM HTML Web Document HTML JPG Full Color JPG MHTML MIME HTML Document MP3 Motion Video MPEG-3 MP4 Motion Video MPEG-4 MPG Motion Video MPG PDF Adobe PDF RTF Rich Text format RTF TGA Targa Image TGA TIF Scanned Document TIF TXT Text file import TXT WAV Audio file WAV</p>				
47	Solution shall make the file transfer process transparent to end users with no user intervention needed. The physical location of the data repository shall be transparent to end user.				
48	Solution shall perform automatic restoration of a corrupted file (Object/File Data and Meta Data) from a Disaster Recovery (DR) copy.				
49	Solution shall automatically resume or recover from file transfer failures.				
50	Solution shall provide capability to log events and problems and provide the ability to manually set the level of logging and reporting.				
51	Solution shall provide a role based configuration with a minimum of a system user role which allows for a system user to view current hardware status and a Super User system administration role which allows for an administrator to configure system policies, and parameters.				
52	Solution shall provide a web based single master management administrative console without the need to install client software for control of system software implementation, policies and user management protected with permission roles.				
53	Solution shall utilize Microsoft Active Directory to provide a single sign-on to the management web console and the log on security shall provide the permission to access and manage all the CIFS shares and data.				
54	Solution shall log the software version/installed hardware/firmware versions at each site.				

55	Solution shall provide interface callable from an application using HTTP/S Web Services using SOAP or REST and Get/Post methods for moving binary data. The ICD/WSDL will be included with the solution.				
56	Solution shall provide a web interface to access the metadata of the files/objects that are stored in the storage device.				
57	Solution shall provide web interface to access the system logs generated by all components. If separate logs are generated the interface will include the ability to determine what logs are available and parameters to select specific logs.				
58	Solution shall provide a web interface to access files by file path/name.				
59	Solution shall provide a web interface to store files by file path/name.				
60	Solution shall provide Web interface that provides usage statistics (metrics on all nodes –size, amount full, number of files, as examples).				
61	Solution shall provide web interface that provides access to the list of files by attributes (MIME type, date created/range, size, file name or partial file name).				
62	Solution shall provide a web interface that will export/copy files by attributes (e.g. File path or partial file path) to a specified target CFIS Share.				
63	Solution shall provide a web interface that provides file statistics (e.g., size, creation date, modify date, access date) given File name or file path, or File/Object ID specified by the solution.				
64	Solution shall provide web interface that provides image retrieval time statistics given an Object/File ID or file path.				
65	Solution shall provide a web interface that provides the “Health status” of the storage (i.e. what is the current status of each of the nodes on the network or of a particular node).				
66	Solution shall provide a web interface that provides status of devices (e.g. any disk errors or outages.)				
67	Solution shall provide web interface that provides the physical location of Storage devices given a device ID supplied by the vendor, or the list of storage devices and characteristics for all devices at a location.				

GRID Based Object Storage (Extra Large)					
OEM Model/Part Number:					
Reference Number	Requirement (Note: These are minimum requirements that a procurement of this storage type shall meet, The exact numbers will be supplied on the Individual Delivery Order.)	Included in standard offering	Included as an optional feature	Customization Required	Comments
1	Solution shall provide minimum usable native disk capacity without data compression, without data deduplication, and without single instance (no duplicate objects) of 200TB (Terabytes), expandable to 500TB (Terabytes).				
2	Solution shall provide a minimum of Enterprise grade SATA spinning disks or better as the backend storage devices to store the object data.				
3	The solution shall not include removable media, such as optical, tapes and removable disk cartridge.				
4	Solution shall support minimum of Representational State Transfer (REST), Common Internet File System (CIFS), and Network File System (NFS) protocol.				

5	<p>Solution shall provide an intelligent policy-based object storage system. Solution shall be fully automated, policy driven without manual intervention, script, or wizard based automation to manage data retention, replication, and data distribution policies.</p> <p>For example, if the policy is defined to have 1 copy of the data locally at the primary site and 1 copy of the data at the remote site. In the event when primary site is not able to communicate with the remote site due to a temporary outage at the remote site, the solution shall automatically create a temporary remote copy at a remote tertiary site intelligently or a user defined temporary remote tertiary location during the creation of the policy process to meet the policy requirement (1 local and 1 remote). When the communication to the original remote site is restored, the solution will update and re-synchronize the missing data automatically to the original remote site and automatically delete the temporary remote copy at the temporary remote tertiary location. The solution shall automatically perform self healing of the system during a non-disaster type of outage. An alert and logging mechanism shall be included in the solution to capture the sequence of activities and provide the final status via email to the designated recipient(s).</p>				
6	<p>Solution shall provide customizable policy to define the number of synchronous local data copies, the number of asynchronous remote data copies, the number of sites to store the data, the primary site location, the disaster recover (DR) site location and the temporary remote tertiary site location.</p>				
7	<p>The solution shall have the ability to provide automatic or user definable thresholds in the policy to determine temporary interruptions (such as system resource and/or communication problems) when accessing the DR site. This capability is used to determine when the remote copy of the data shall be copied to the temporary tertiary site.</p>				
8	<p>The Contractor shall provide a single Graphic User Interface (GUI) that allows more than one creations, editions, and deletions of data retention, replication, and data distribution policies in a one single process. For example, data retention, replication, and data distribution policies can be configured in one GUI interface and not to have separate individual interfaces for data retention, replication, and data distribution respectively.</p>				

9	Solution shall replicate files to the remote site based on available WAN bandwidth and remote system resources. The process shall be fully automated, policy driven without manual intervention, script or wizard based automation.				
10	Solution shall provide the necessary equipments at the secondary remote site so that the remote site can be promoted to primary site role in the event of total disaster at the primary site.				
11	In the event of a site recovery, the solution shall be capable of creating a duplicate system by copying all data from the Disaster Recovery site which will be shipped to the disaster site to be recovered.				
12	Solution shall provide storage capacity expansion or scalability with no downtime, no forklift upgrade, no reconfiguration of the internal operating system, and LUN remapping of the storage subsystem to connecting hosts. Any expansion shall maintain the same redundancy, performance and efficiency of the system as the initial delivered system exhibits on all supported protocols.				
13	Solution shall support redundancy in all storage system components with no single point of failure and non-disruptive to operations for all system component replacements or repairs or firmware and microcode upgrades or updates.				
14	Solution shall provide an automatic mechanism to migrate data from existing storage to any future new storage device when the storage hardware is End-of-life (EOL) with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
15	Solution shall be able to migrate data to newer types of disk drive when replacing end-of-support disks with no impact to production operations, no downtime, and shall provide access to old and new storage data at all times during the migration process.				
16	Solution shall provide File Transfer ingest rate to local storage solution at a minimum of 1.5 Gigabytes (GB) of data per minute (1.5GB/min).				
17	Shall Support a minimum of five nines storage subsystem uptime (or 99.999% availability, excluding planned downtime). (Note: Five Nines (99.999%) = 5.3 minutes of downtime on a rolling 12 month basis)				

18	Solution shall always maintain original file attributes (Create, Access dates, file permission, and security or other attributes) for all data replication process, data copy process including data restoration processes. The file attributes are to stay with the file.				
19	Solution shall provide the ability to verify data after it has been transferred from the local site to the remote location to ensure it matches with no data corruption. If verification of data fails solution shall resend the data. The copy of file at the remote site shall be identical to the copy of the image file at the local site after replication. This check shall be performed automatically.				
20	Solution shall automatically (based on policy, not manual, script or wizard based) redirect the read processes to remote storage devices/components if the local storage devices/components are not available during image file read process.				
21	Solution shall provide a solution that can automatically (based on policy, not manual, script or wizard based) redirect the write processes to remote storage node and/or backend storage if the local storage node and/or backend storage is not available during image file write process.				
22	The solution shall present a CIFS interface that is addressable by a Microsoft Active Directory fully qualified path. The Solution shall present a redundant CIFS interface. It must be presented to any windows client on the network. Any windows client on the network must be able to authenticate to the server and access to the share.				
23	The Contractor shall provide a redundant Global Namespace object based Archive system. All objects shall be retrievable via web browser (http and https) based on its object metadata definitions.				
24	Solution shall maintain mappings of the Common Internet File System (CIFS) share in the event that the primary local gateway is unavailable or offline. The solution shall continue to function automatically with no manual intervention, script or wizard based automation on local secondary gateway with the same CIFS mapping. The path used to access the CIFS shares will be the same regardless of whether the local or remote copy of the image is accessed.				

25	Solution shall provide a process automatically (based on policy, not manual, script or wizard based) update and resynchronize the recovered local site archive storage from the remote archive storage when the local archive storage is back online after a period of downtime. The process shall be fully automated, policy driven without manual intervention, script or wizard based automation				
26	Solution shall provide a process to log all file transfer success or failure.				
27	Solution shall utilize Microsoft Active Directory for all security permissions for the whole system management and data access.				
28	Shall provide native capability to proactively monitor/manage hardware and software processing, and generate alerts on potential storage system components problems, failures and application storage resources health. This information should be capable of being transmitted to the OEM or Vendor as defined by the contract and designated VA Point of Contact				
29	Shall provide native capability to report or generate reports at a minimum on hardware failures or items out-of-specifications in the Storage Subsystem, such as Disk Drives, Fans or Power Supplies and other similar components.				
30	Solution shall have a native alert function to send failures/warnings/alerts via a "Call Home" function or similar process and SMTP mail.				
31	Solution shall escalate alerts when storage space utilization falls below a user defined threshold.				
32	Solution shall provide a means for the solution to log the file transfer time from local archive to remote site as a user defined threshold. (e.g. If logging transfer time is set to 60 seconds and the actual transfer time exceeds 60 seconds it shall be logged).				
33	Solution shall support the function to access remote data from the local archive location for display purposes.				
34	Solution shall provide a means for the Solution to retrieve a copy of the data from the remote archive locations for storage at the local archive location if the local archive was not available.				
35	Solution shall log file access, create and modify dates for retention statistics and monitoring purposes.				
36	Solution shall under normal operating conditions, replicate data immediately. In no case shall the replication backlog be greater than four (4) hours of data.				

37	Solution shall be able to scale components such as network interface cards, adapter ports, processors, memory, and storage capacity independently.			
38	Solution shall monitor and access the remote data with the least amount of file transfer latency.			
39	Solution shall manage, control, and define policies (policy that defines the number of data replication copies, the location, and the temporary remote tertiary site location) in the Object storage device.			
40	Solution shall be capable to capture and/or log performance data such as data ingest rate, data ingest size, data replication bandwidth utilization, data replication latency, data retrieval response time and others.			
41	Solution shall support a minimum of two 1Gb Ethernet connectivity to the local production LAN.			
42	Solution shall provide redundant Power Distribution Unit (PDUs) with sufficient outlets to power all units in the rack(s).			
43	Solution equipment shall be Pre-installed and Pre-cabled in a 42U Rack prior to delivery. All cables and components installed in the rack shall be labeled.			
44	Solution shall include Pre-installed and Pre-cabled redundant top of the rack network switches in every rack provided to minimize the number of network connections to VA Core network switches. The top of the rack switches provided in the solution shall be interoperable with existing VA Core network switches.			
45	Solution shall support Internet Protocol version 6 (IPv6).			
46	<p>Solution shall support a minimum of the following file formats:</p> <p>AVI Motion Video AVI BMP Bitmap file BMP DCM Dicom file DCM HTML Web Document HTML JPG Full Color JPG MHTML MIME HTML Document MP3 Motion Video MPEG-3 MP4 Motion Video MPEG-4 MPG Motion Video MPG PDF Adobe PDF RTF Rich Text format RTF TGA Targa Image TGA TIF Scanned Document TIF TXT Text file import TXT WAV Audio file WAV</p>			
47	Solution shall make the file transfer process transparent to			

48	Solution shall perform automatic restoration of a corrupted file (Object/File Data and Meta Data) from a Disaster Recovery (DR) copy.				
49	Solution shall automatically resume or recover from file transfer failures.				
50	Solution shall provide capability to log events and problems and provide the ability to manually set the level of logging and reporting.				
51	Solution shall provide a role based configuration with a minimum of a system user role which allows for a system user to view current hardware status and a Super User system administration role which allows for an administrator to configure system policies, and parameters.				
52	Solution shall provide a web based single master management administrative console without the need to install client software for control of system software implementation, policies and user management protected with permission roles.				
53	Solution shall utilize Microsoft Active Directory to provide a single sign-on to the management web console and the log on security shall provide the permission to access and manage all the CIFS shares and data.				
54	Solution shall log the software version/installed hardware/firmware versions at each site.				
55	Solution shall provide interface callable from an application using HTTP/S Web Services using SOAP or REST and Get/Post methods for moving binary data. The ICD/WSDL will be included with the solution.				
56	Solution shall provide a web interface to access the metadata of the files/objects that are stored in the storage device.				
57	Solution shall provide web interface to access the system logs generated by all components. If separate logs are generated the interface will include the ability to determine what logs are available and parameters to select specific logs.				
58	Solution shall provide a web interface to access files by file path/name.				
59	Solution shall provide a web interface to store files by file path/name.				
60	Solution shall provide Web interface that provides usage statistics (metrics on all nodes –size, amount full, number of files, as examples).				
61	Solution shall provide web interface that provides access to the list of files by attributes (MIME type, date created/range, size, file name or partial file name).				

62	Solution shall provide a web interface that will export/copy files by attributes (e.g. File path or partial file path) to a specified target CFIS Share.				
63	Solution shall provide a web interface that provides file statistics (e.g., size, creation date, modify date, access date) given File name or file path, or File/Object ID specified by the solution.				
64	Solution shall provide web interface that provides image retrieval time statistics given an Object/File ID or file path.				
65	Solution shall provide a web interface that provides the "Health status" of the storage (i.e. what is the current status of each of the nodes on the network or of a particular node).				
66	Solution shall provide a web interface that provides status of devices (e.g. any disk errors or outages.)				
67	Solution shall provide web interface that provides the physical location of Storage devices given a device ID supplied by the vendor, or the list of storage devices and characteristics for all devices at a location.				