

Scope of work:

- i. The supplier is responsible for carrying out and completing the work following the tender documents and under the direction of the Scientist-In-Charge. Any work not explicitly covered in the specifications will be executed as decided by the Scientist-In-Charge
- ii. The Scientist-In-Charge has the authority to interpret the meaning of the specifications and drawings and to accept or reject any work or material that does not meet the requirements of the specifications or applicable codes and standards.
- iii. All warranty and support must be serviced directly by the OEM or should be from an Authorized System Integrator Partner who is Authorized to Support the product quoted.
- iv. CDFD requires a Single Point of Contact (SPOC) from OEM/Vendor who is responsible for all issues between CDFD and the OEM.
- v. All the server components should be verified and recommended by the motherboard manufacturer.
- vi. All the software deployment will be in the vendor's scope of work, and vendors are requested to include any software deployment charges in their proposal.
- vii. The storage nodes should be of dense rack form factor designed for cluster solution and should be in the lowest footprint and the lowest power consumption.
- viii. The complete installation will undergo the continuous running test for 2 days after initial testing and commissioning. The supplier will rectify defects, and the test will be repeated. After successful completion, the equipment will be taken over.
- ix. During the guarantee and defect liability period, the supplier will provide the necessary material and personnel for maintenance and repairs at no additional cost. They will attend to problems within 48 hours of receiving the complaint and take corrective action immediately.
- x. The supplier will train CDFD personnel at no extra cost during the installation and testing period before virtual completion and taking over.
- xi. The installation, testing, commissioning, and handing over of the High-Speed Storage system should be completed within the stipulated or extended time allowed by the Institute.
- xii. The supplier/manufacturer is responsible for timely service and calibration of the machine to ensure successful installation and satisfactory operation.
- xiii. The supplier will provide and pay for all necessary testing equipment. The supplier will rectify defects found during initial testing at their own cost. Initial testing will be done in the presence of Scientist-In-Charge, and commissioning will be done after their approval.
- xiv. All the work shall be carried out strictly and per the specifications given in the tender to the total satisfaction of the Scientist In-charge.
- xv. The work under the contract will be open to inspection and supervision by CDFD. The Tenderer must be present or have a representative present during working hours and with reasonable notice to receive instructions.

Technical Specification for 2PB high speed secondary data storage

Sr. No	Description	Compliance (Yes/No)
1.	Overall architecture	
1.1	The solution must provide software-defined storage (SDS) that allows delivery of enterprise-class storage (available on the self) services using the latest x86/IBM Power nodes (Hardware Agnostics) connected using Ethernet-based interconnect and without any dependency on specialized storage interconnect (FC, IB, etc. or hardware components).	
1.2	The storage cluster builds using multiple storage nodes together, presenting a single storage pool to provide high-bandwidth data access for multiple clients in parallel. There should be no single point of failure.	
1.3	The solution should be able to support the different generations of x86/IBM processors and appliance models in the same cluster for investment protection over the life of the proposed solution. It also supports different types (HDD, SSD, NVMe) and size of disks. The ability to mix and match various models should be supported in the same storage cluster and unlike RAID there should not be any requirement for identical storage disks.	
1.4	The solution shall provide the ability to expand the storage capacity by hot-add storage node(s) (scale-out) or hot-plug disks to existing nodes (scale-up) at a later point in time. The storage provisioning should be via GUI and with the no disruption to the workloads already running on the platform.	
1.5	In the case of cluster expansion, the existing data should automatically balance across available controllers/nodes without any administrative intervention.	
1.6	The storage solution must offer mature interfaces for enterprise file and object storage. It must provide a single platform for a complete range of access methods including Object storage (S3) and traditional access methods such as File Services (NFS, CIFS/SMB). These services should be natively available or using OEM provided gateways. This function must be enabled in the proposed solution from day one.	
1.7	Possibility to associate different types and sizes of storage media to various performance domains (SSD, spinning drives) and failure domain (drives, nodes, racks, raw, DC) and then multiple storage containers (volumes/LUN/bucket) can be configured with available data protection.	
1.8	The proposed solution must support data security, data protection, migration, self-healing, and rebalancing from day one and without additional license cost.	
1.9	The proposed software solution should be open-sourced or licensed software which is PERPETUAL in nature and should have NO dependency on the underlying hardware.	
1.10	The solution can be deployed on standard Linux distros (at least Red Hat, SUSE, Ubuntu, Rocky or CentOS).	
1.11	The solution should provide a single unified management console for managing the entire SDS environment and its associated software components along with scalability in near future.	
1.12	The solution should be easily integrated with the existing CDFD computational infrastructure.	

2.	Data Protection	Compliance (Yes/No)
2.1	The solution should support data protection methodology using data replication (two or more copies of data within the cluster) and Erasure Coding or equivalent technology.	
2.2	The data must be striped across multiple storage devices and nodes. On-disk or node failure, storage shall redistribute data across as many drives and nodes as possible to take advantage of multiple spindles/SSD/NVMe.	
2.3	Must tolerate a failure of at least two disks across any node or even simultaneous 2 nodes failure in the cluster and continue to serve data to end-users and applications without any data loss.	
2.4	Storage solution should sustain full two-node failure, one or more components (disk, NIC), etc. On failure of one node, there should not be a reduction of storage capacity by more than 18%.	
2.5	Provides data protection in a storage efficient fashion using flexible/configurable Erasure coding (i.e. 4+2P, 8+2P, or equivalent), preferably without any dependencies on hardware RAID and regardless of object size.	

2.6	Should have the capability to change the protection level (Replication Factor-2 to RF-3, RF2 to Erasure Coding) per storage volume without any data loss or unavailability of data.	
2.7	Provide data integrity by Checksums or equivalent technique. The system should have file system integrity and data integrity checks built in to prevent data loss due to bit rot and other soft errors.	

3.	Security	Compliance (Yes/No)
3.1	The system must be able to support Write Once Read Many (WORM) capabilities.	
3.2	Should provide data encryption at rest and on the wire (in transit).	
4.	Performance	
4.1	Read and Write performance to scale with expansion of clusters (adding nodes, drives) linearly	
4.2	SDS solution to provide minimum of 4.5GBps read bandwidth and 2.5GBps of write bandwidth	
4.3	Solution to provide at least 90% of Read IOPS within 20ms but not limited to.	
5.	Broad Specifications of File storage	
5.1	Provides native POSIX file system compatibility as well as support for common protocols including CIFS/SMB, NFS V3, and NFS V4. Optionally LINUX FUSE support.	
5.2	Support for multiple File System Abstraction Layer (FSAL) through NFS-GANESHA or equivalent.	
6.	Management Software	
6.1	Proposed storage must be managed and monitored via integrated web-based UI, CLI & RESTful APIs.	
6.2	Should be capable to deliver deployment automation, provisioning, capacity management, storage tier optimization, reporting, data protection/replication configuration and management, storage resource management through easy-to-use web-based administrative GUI.	
6.3	The portal must have comprehensive reporting capabilities including <ul style="list-style-type: none"> • Performance monitoring • Capacity utilization per storage pool, node, disk, and per site. • Performance monitoring on latency, throughput, and replication progress. • Diagnostic information, such as node and disk recovery status. 	
7	File Gateway	
7.1	Web based UI should have following functions Authentication Listing of shares and files Download and Upload Files Create / Delete Directory	
7.2	The system must support authentication, by Integrating with Windows Active Directory (AD) / LDAP	
7.3	The system should be able to block undesired files by their extension/mime type. For example .exe etc....	
7.4	Secure Remote File Access- Access data through an end-to-end secure channel over web	
7.5	Two-Factor Authentication - OTP based verification along with the standard user authentication	
7.6	Audit Report and Access Logs - Complete traceability of all operations including platform and file resources	

8	Server/Storage Nodes (minimum of 6 numbers) for 2PB usable capacity	Compliance (Yes/No)
8.1	Solution to be based preferable from reputed hardware OEMs like IBM, DELL, LENOVO, HPE or equivalent.	
8.2	The proposed solution (hardware and software) must be based on reference architecture (RA) validated. Documentary proof should be attached describing the RA.	

8.3	Dual Intel Xeon Ice lake processor with 16 MB or higher L3 cache or equivalent IBM power series. Each storage heavy node must have at least one logical CPU (vCPU) core (or 2 GHz or higher) per connected HDD (per HDD).	
8.4	Each storage nodes must have at least 2GB of RAM/HDD and additionally 32 GB RAM per node (i.e. 32 GB + 2GB/HDD). The RAM must be 128GB ECC DDR4-3200 MHz or higher.	
8.5	Each node must have at least 2Nos x 25 Gbps NIC (for data), 2 Nos x 25Gbps NIC (client/public access), and 1 No x1Gbps (Management) Ethernet ports NIC. (i.e. per node 4 x 25 Gbps and 1x 1Gbps Ethernet port)	
8.6	Each node/JBOD must have a Platinum power supply. Redundant and hot-swappable power and fan in N+N/N+1 mode.	
8.7	All the disks drive including SSDs must be hot-swappable, replaceable.	
8.8	The storage disks are connected through JBOD. Each connectivity must be through PCI Express 3.0 (or equivalent dedicated slot) based 12 Gbps SAS-3 (12Gb/s SAS, 6Gb/s SATA) compliant HBA or equivalent RAID controller with the required number of ports to avoid the performance bottleneck. The accessories required to connect and mount the JBOD and any other accessories required are to be provided by the bidder.	
8.9	Each node must have redundant M.2 disks (configured using RAID 1) of 240 GB or higher as a boot drive. Pre-installed OS with required licenses if any.	
8.10	Additional resources (CPU, RAM, storage) are required to run any additional services (protocol/gateway, management services, etc...) other than storage services to be calculated and provide extra on each node.	
8.11	Required number of 14TB NL-SAS or higher capacity drivers (7.2K RPM) for data storage and SSDs/NVME for caching/metadata/journaling storage.	
8.12	25G switch should be provided with minimum 4 ports of 40/100Gbps.	
8.13	Complete Solution should fit in 20 - 42 U form factor standard rack.	
8.14	Any component required to fulfill the solution should be included.	
8.15	Complete solution should be integrated with OEM rack.	