

Annexure-II

Tender Specifications for 1 PB usable storage

1. The storage should be fully symmetric and distributed clustered architecture with scale out model.
2. Data should not be lost / require backup and restore during upgrade/expansion.
3. The storage should have minimum two controllers/storage nodes of the same type.
4. The storage should be upgradable to minimum 8 PB usable capacity without any disruptions/downtime to production workflow for performance.
5. The storage should support linear scalability of performance and capacity.
6. All storage nodes must be active for all storage shares, contributing in performance.
7. The storage must be configured with minimum 128 GB globally coherent Cache memory across nodes. The memory should be upgradable to minimum 512 GB.
8. All the storage nodes should be manageable through a single console.
9. The node should have hot swappable redundant fans and power supplies.
10. The complete 1PB storage solution should fit in 14U rack space. The solution must be implemented within the already existing racks, PDUs available at the IUCAA data centre. The bidder may inspect IUCAA facilities before quoting.
11. The total capacity of storage should be 1 PB usable using 12 TB or higher NL-SAS drives in RAID-6 (8+2) or equivalent with minimum of 5% capacity as global hot spare.
12. The disks in the storage should be of enterprise class disk drives with hot swap capability
13. Storage should support SSD, SAS, SATA, NLSAS disk types to create multiple tiers of storage, if required as a part of single filesystem
14. The storage should have a single file system, i.e., total space should be configured with single file system, for whatever be capacity in future (for 8PB).
15. The storage should be capable of giving minimum 3 Giga Byte Per Second throughput for read on NFS. The storage should be capable of giving minimum 1 Giga Byte Per Second throughput for write on NFS. The performance should not degrade by more than 25% on failure of a single controller/node. The performance should be demonstrated by running IOR with 1MB record size or transfer size, for a total file size should be at least thrice the cache.
16. Bidder should run benchmark using IOR process less than or equal to total the number of storage nodes,
17. Benchmark report should be submitted along with the bid.
18. The array should have a minimum of 8 X 10G Ethernet Ports with SFP+ modules across the controllers.
19. The backend internal connectivity between storage nodes should be using minimum 10/40 GigE (or better) network with no single point of failure.
20. The backend internal connectivity configured between storage nodes and disk controllers, if configured, should be redundant and there should be "No Single Point of Failure". The Complete multi-node Storage System Solution should be fully redundant, configured in High Availability mode and should NOT have any Single Point of Failure
21. All data should be striped across all storage controllers in the proposed storage system, so that performance of all controllers can be utilized for all read and write operations.
22. The array should support NFS, HTTP, FTP protocols.

23. The storage system should support non-disruptive replacement capabilities for components like Disk Drives, fans, power supplies, controllers etc.
24. The storage array should provide multiple levels of access control including role-based security and auditing capability. There should be provision to apply role-based disk quota.
25. The storage array should support connectivity to CentOS 6.0 and above, Ubuntu 12.04 and above, RHEL 6.0 and above.
26. The storage should provide a web-based tool/interface to monitor, configure and manage the storage. It should be supported on all browsers like Firefox, Chrome, Safari. The licenses should be perpetual in nature. The storage management software should provide proactive notifications for actual or impending component failure. The bidder should provide required hardware/software component required to host the storage management software, if any. The storage management software should be from OEM itself.
27. The storage should have dataset management software for high-speed scanning and indexing on offered storage. The software should scan the metadata only; it should not scan any data stored on the storage. The software should allow bi-directional movement of data between multiple file storages and should support storages from multiple OEMs. The software should have the functionality to find the duplicate data in all managed storages. The software license should be for entire capacity, and it should be perpetual in nature.
28. Necessary cables connector, cables, switches and other accessories required to mount storage to server should be provided by the bidder.
29. Storage should not require any down time and interruptions during any kind of updates and upgrades
30. Storage should support policy based archival feature.
31. Storage should support built in asynchronous data replication between two storages
32. The storage should support policy based retention and protection against accidental deletion.
33. The storage should not have a single point of failure. The storage should have self-healing design to protect against disk or node failure.
34. The storage system should be able to protect the data against simultaneous failure of two disks or one node without any data loss and data unavailability
35. The storage should support data deduplication.
36. In the event of addition of storage node to storage solution, existing data should be rebalanced across all nodes of storage automatically. This auto balancing should be done with low priority avoiding any impact to client performance.
37. Storage should have 3 years of next business day (NBD) on-site support, and warranty.