



Sl #	Section / Page No	Clause Reference	Query from Bidder / Reasons for change	Response by IIITD
1	Page 4 - Master Node	Point 4 - Storage	Kindly change it to "2 x 1.92 TB NVMe M.2 or U.2 or U.3 SSDs in RAID 1 array", since we do not over any M.2 NVMe disk option in the server.	Bidders can offer M.2/U.2/U.3 NVMEs
2	Page 4 - Master Node	Point 12 - OS Support	Kindly remove support for Ubuntu version 20.4.	We need support for ubuntu 20.4 or higher
3	Page 5 - GPU Node	Point 7 - Network 1	Kindly amend "4x OSFP ports serving 8x single-port NVIDIA ConnectX-7 VPI" to "8 x single port Nvidia ConnectX-7 VPI" since this is proprietary to some vendor.	Bidder can offer deviated hardware to meet the benchmarks asked in annexures they can mention the deviation in remarks
4	Page 5 - GPU Node	Point 7 - Network 1	"2x dual-port QSFP112 NVIDIA ConnectX-7 VPI" ? What is the expectation ? Is it 2 x dual port NDR 200G ?	
5	Page 5 - GPU Node	Point 8 - Network 2	"100G Ethernet NIC". What is it going to be used for ? Are transceivers needed ?	
6	Page 5 - GPU Node	Point 12 - Power Supplies	Kindly change it to "N+N or N+2", since HPE provides only N+2 configuration Power Supplies. This is to ensure that even after failure of 2 x PSUs, the system does not throttle.	
7	Page 5 - GPU Node	Point 4 - OS Support	Kindly relax the support for Ubuntu 20.4.	We need support for ubuntu 20.4 or higher
8	Page 9 - Network Switch	Ports	"All ports on front" - what is the expectation ? The Nvidia switch which shall be provided will have Power to Port Air flow, hence the Power supplies will be in the front side and the ports in the back side of the rack. This helps in directly connecting the server ports which are in the back to the IB ports in the back. This is the standard practice and hope this is fine.	Accepted
9	Page 7 - PFS	Point 1 - Storage features	Kindly delete "The solution should deliver at least 4 nines (99.99) uptime reliability". 4 nines is an Enterprise class feature and NOT prevalent in the HPC space. HPC Storage is built with sufficient NSPOF (No single point of failure), hence we request 4 nines requirement to be relaxed.	No Changes
10	Page 8 - PFS	Point 1 - Storage features	Kindly change it to "The solution should be software defined storage or Hardware appliance and run on multiple industry standard platforms to provide multiple vendor choice and avoid vendor lock-in". Alternately, a OEM specific Hardware appliance also can be offered". Reason for this request is that HPE Lustre storage is built as an integrated and modular appliance.	Deviation is allowed,you can mention the deviation in remarks
11	Page 8 - PFS	Point 2 - Storage Mgmt	Kindly remove "All tuning parameters must be accessible through the GUI" since most of the tuning parameters are over CLI rather than GPU.	Accepted

12	Page 8 - PFS	Point 3 - Storage Performance	Kindly change it to "All data distributed evenly across the storage cluster delivering great performance, without any storage hot spots or I/O bottlenecks". Reason for this change is that different Storage architectures have different mechanisms to handle metadata.	Deviation is allowed,you can mention the deviation in remarks
13	Page 8 - PFS	Point 3 - Storage Performance	We request IIITD to review the performance numbers for the 500TB storage. These numbers are very aggressive and is not demonstratable with 1 x GPU nodes. We will need a huge number of GPU nodes to demonstrate this performance. Also the overall Storage solution cost will increase significantly if the sizing is done purely based on performance sizing.	Demonstrate to >3 node.
			The solution shall provide a minimum bandwidth of at least 400 GB/second of reads (100%) and 400 GB/sec writes (100%). Kindly confirm if it is 400GB/s or 400Gbps (Read / Write).	It is 400GB/s(Read/Write)
			The solution shall provide a performance of 4K IOPS of at least 10 Million IOPs (reads). The figures are very aggressive, hence we recommend 1M Read IOPS (4k). Typically, about 8-10 nodes will be needed to demonstrate this performance. However, since we will have limited node count for demonstration, the performance demonstration shall be on a best effort basis. Kindly confirm.	Deviation is allowed,you can mention the deviation in remarks
14	Page 8 - PFS	Point 4 - Data Protection	Kindly delete the following, since we do not support snapshots. They degrade the performance of the system.	Deviation is allowed,you can mention the deviation in remarks
			The solution shall support snapshots and clones of an entire filesystem without performance degradation	
			Snapshots and clones occur instantaneously and are differential after the first instance without any performance degradation.	
			Support to take snapshots and clones on a per file system basis.	
			Ability to recover the entire filesystem or individual file from snapshots.	
15	Page 20 - Benchmarks	MLPerf 3.1 Training	Since the RFP requirement is for 8 x GPUs, kindly restrict the GPU benchmarks to 8 x H100 GPUs only.	Deviation is allowed,you can mention the deviation in remarks. Submit the benchmarks for one node
16	Page 20 - Benchmarks	MLPerf 3.1 Training	Kindly remove Stable Diffusion v2 benchmark. We have not run this.	
17	Page 21 - Benchmarks	MLPerf 3.1 Inference	Kindly relax these benchmarks. We have not run these benchmarks on a 8 x H100 System.	
18	Page 22 - Benchmarks	MLPerf Storage	Kindly relax these benchmarks and restrict the storage benchmarks to throughput and IOPS.	
19	Storage Features: Clause: The solution should be software/hardware defined storage and run on multiple industry standard platforms to provide multiple vendor choice and avoid vendor lock-in.		Request panel to change the clause from software/hardware to Software defined Storage	No changes as long as benchmarks are meeting

20	Clause: Optionally The solution should consist of one native filesystem for simplicity, performance, and reliability - and not be built on top of another filesystem like ext4, XFS, or ZFS which adds complexity, slows performance and degrades reliability		Request panel to remove word Optionally added to the above clause.	No changes as long as benchmarks are meeting
21	Clause: The solution should be scale-out storage that is fully integrated and may require separate metadata servers, administrative servers, or protocol gateways.		Request to change this clause to deny requiring the separate metadata servers. Having separate metadata servers in the cluster creates metadata hotspots as all the requests from the application will first go to a metadata server thus creating a bottleneck and performance hotspots. This defies the overall reason of having a distributed architecture	Deviation is allowed. In remarks deviation should be noted
22	Storage Performance: Clause: All data and metadata are distributed evenly across the storage cluster delivering performance, without any storage hot spots or I/O bottlenecks as asked in benchmarks in Annexures		Request the panel to amend the clause allowing distributed architecture and simultaneously considering the Annexures:	Deviation is allowed. In remarks deviation should be noted
23	This is important because without distributed metadata and data there would be hotspots of metadata causing performance impacts. Clause: The solution shall provide a performance of 4K IOPS of at least one Million IOPS (reads)		Request the panel to amend the clause to “The solution shall provide a performance of 4K IOPS of at least 10 Million IOPS (reads)”	Bidders can offer higher IOPS. We will note it for QCBS evaluation
24	Data Protection Features Clause: Optionally the system shall support encryption at rest and in flight		Request panel to amend the clause “System shall support encryption at rest and in flight”. And remove word “Optionally”. It is important that the solution must ensure security of data throughout its lifecycle	Deviation is allowed. In remarks deviation should be noted
25	Storage Performance The solution shall provide a minimum bandwidth of at least 400 Gb/second of reads (100%) and 400 Gb/sec writes (100%).		Request panel to amend the clause “The solution shall provide a minimum bandwidth of at least 400 GB/second of reads (100%) and 400 GB/sec writes (100%)	It is 400GB/s(Read/Write)