

Answer to the Bidders queries

No	RFP Specification	Vendor Query	Justification
5.1 1	Should be configured with Dual Controllers. The controllers should use multi-core processors, and the total number of cores per controller should be greater than or equal to 24	<p>Request for Modifications: Should be configured with Dual Controllers. The controllers should use multi-core processors, and the total number of cores across dual controllers should be greater than or equal to 24.(The storage solution should only be built upon an industry-standard open processor architecture. Eg.Intel or AMD).</p> <p>Justification: The RFP requires each controller to have 24 cores or more. Currently, you are using IBM V5030, which has only 6 Intel cores per controller. The requirement for 24 cores per controller is excessive, as only proprietary processors offer this specification. We suggest specifying Intel or AMD processors (open architecture) to facilitate a fair comparison between open architecture processors (e.g., Intel or AMD) and proprietary processors (e.g., Huawei).</p> <p>Allocating 24 cores per controller can lead to underutilization of resources, as many storage environments do not demand such a high number of cores. Efficient core management and optimized software can handle typical storage workloads effectively with fewer cores.</p> <hr/> <p>While the requirement for 24-core controllers might seem to ensure high performance, it is unnecessary and excessive for this use case. Our storage system, with 12-core controllers, provides the necessary performance and reliability without the drawbacks of over-provisioning. By leveraging optimized core utilization, advanced I/O processing, and efficient cache management, we deliver a balanced, cost-effective, and high-performing storage solution. Because of all these factors we request you to modify this with</p> <p>"Should be configured with Dual Controllers. The controllers should use multi-core processors, and the total number of cores per controller should be greater than or equal to 12"</p> <p>Below are the points that should be considered.</p> <p>Resource Over-Provisioning: Allocating 24 cores per controller can lead to underutilization of resources. In many storage environments, the actual workload does not demand such a high number of cores, resulting in wasted computational power and energy. Efficient core management and optimized software can handle typical storage workloads with fewer cores effectively</p> <p>Performance Metrics: Our workload sizing & performance testing (done with internal tools) indicates that storage systems with 12-core per controller, consistently meet performance expectations.</p>	<p>The requirement for each controller to have 24 cores or more is designed to ensure optimal performance and scalability of the storage solution. A higher number of cores allows for better task allocation, significantly enhancing overall efficiency and performance. This ensures that the system can handle multiple operations simultaneously without degradation, which is critical in high-demand environments.</p> <p>Specifying a high core count also future-proofs the storage solution, accommodating growth in data workloads and technological advancements. It provides necessary headroom for future requirements, ensuring long-term viability.</p> <p>Furthermore, the requirement does not exclude the use of processor architectures like Intel or AMD. These processors utilize the x86 architecture, which is itself a proprietary standard. Vendors can still offer solutions with these processors that meet the specified core requirements, facilitating fair comparison while setting a high-performance standard.</p>

No	RFP Specification	Vendor Query	Justification
5.13	The total cache capacity of the system should be 64 GB, and the cache capacity of any controller should be greater than 32 GB	<p>Request for Modifications: The total cache capacity of the system should be 512 GB, and the cache capacity of any controller should be 256 GB</p> <p>Justification: A higher controller cache helps in improving performance, reducing latency, and enhancing data reliability.</p> <hr/> <p>A higher controller cache helps in improving performance, reducing latency, and enhancing data reliability. Hence we request you to modify this to</p> <p>"The total cache capacity of the system should be 512 GB, and the cache capacity of any controller should be 256 GB"</p>	64GB Cache controller is sufficient to provide required IOPs with low latency. And vendors are welcome to propose higher cache with higher IOPs.
5.15	The single storage array shall be offered with 25 TB usable capacity (Physical) after RAID6. (Before compression and De-duplication)	<p>Request for Modifications: The single storage array shall be offered with the single storage array shall be offered with 25TB using NVMe SSD/ NVMe flash drives</p> <p>Justification: NVMe SSDs leverage a more efficient command set and parallel processing capabilities, offering significantly higher data transfer speeds and lower latency compared to traditional SATA or SAS-based SSDs.</p> <p>NVMe SSDs leverages a more efficient command set and parallel processing capabilities and offers significantly higher data transfer speeds and lower latency compared to traditional SATA or SAS-based SSDs. This results in faster read and write operations, reduced application loading times, and overall improved system responsiveness</p> <p>Hence please modify this to</p> <p>"The single storage array shall be offered withThe single storage array shall</p>	SAS SSD is sufficient to meet required IOPs. Vendors are welcome to propose
5.16	120,000 IOPS with roposed solution. Note: Other Specification for capacity and performance calculation. Avg Block Size (KB) - 8, Avg Latency	Note: Other Specification for capacity and performance calculation. Avg Block Size (KB) - 8, Avg Latency (ms) - 1, Read-Write Ratio - 70.00%at 0% Cache hit ratio	0% Cache hit ratio can be considered.
4 x 10 Gbit/s Ethernet interfaces, 8 x 1		<p>Request for Modifications: 4 × 10 Gbit/s Ethernet interfaces for system.</p> <p>Justification: Currently, you are using only FC connectivity, but the RFP requests 4 × 10 Gbit/s Ethernet interfaces, 8 × 1 Gbit/s Ethernet interfaces (multi-mode optical modules), and 8 × 16 Gbit/s FC interfaces for the system.</p> <p>The requirement for 4 × 10 Gbit/s and 8 × 1 Gbit/s Ethernet interfaces is specific to certain storage systems only.</p>	The requirement for 4 x 10 Gbit/s and 8 x 1 Gbit/s Ethernet interfaces ensures robust and scalable network connectivity for future expansion. While current usage may rely on FC connectivity, 10 Gbit/s Ethernet ports are essential for connecting more application servers with IP SAN, supporting high data

No	RFP Specification	Vendor Query	Justification
5.17	Gbit/s Ethernet interfaces (multi-mode optical modules) for system.	<p>A single management port per controller provides adequate access and control for administrators to perform essential tasks efficiently. It offers a cost-effective and straightforward solution that meets the needs of the storage environments. Hence please modify this so that it will bring clarity.</p> <p>"Per controller it should support 2 × 25 Gbit/s SFP /Ethernet interfaces & 1 Gbit/s Ethernet interfaces."</p>	transfer rates. This forward-thinking approach accommodates growth and future-proofs the infrastructure. The inclusion of 1 Gbit/s interfaces provides flexibility and redundancy, enhancing reliability and performance. In summary, this clause prepares the storage system for future needs, ensuring a scalable and adaptable network infrastructure
5.19	Apart from the above IO module the storage should support adding another additional IO module per controller, for future expansions	This is OEM specific architecture and for wider specification please remove this clause.	Considering the business environment, applications requirement may increase in the future. Should have the capability to expand another IO module to connect more hosts in the future.
5.21	Should support RAID 5, RAID 6, and RAID-TP. RAID-TP is able to tolerate simultaneous failure of three disks.	<p>Request for Modifications: Should support either RAID 5/ RAID 6 (or equivalent/better)/ RAID-TP.</p> <p>Justification: The RFP requires support for RAID 5, RAID 6, and RAID-TP, which can tolerate the simultaneous failure of three disks. RAID-TP has the highest overhead and potentially the lowest performance due to triple parity calculations. We support DRAID6, which tolerates two simultaneous drive failures and is a highly suitable choice due to its balanced approach to performance, rebuild speed, and fault tolerance. Today's hardware components are significantly more reliable, reducing the likelihood of simultaneous multiple drive failures. Our storage also provides proactive failure warnings through advanced monitoring systems and predictive analytics, allowing for proactive measures to be taken.</p> <p>RAID-TP has the highest overhead and potentially the lowest performance due to triple parity calculations and hence we do support DRAID6, which tolerates two simultaneous drive failures and is a highly suitable choice due to its balanced approach to performance, rebuild speed, and fault tolerance. Also because of the Increased Reliability where today's hardware components are significantly more reliable than in the past, reducing the likelihood of simultaneous multiple drive failures. Apart from this our storage provides proactive Failure warnings through advanced monitoring systems and predictive analytics. It can detect potential drive failures before they occur, allowing for proactive measures to be taken, such as replacing a failing drive proactively. So please modify this with</p> <p>"Should support either RAID 5/ RAID 6 (or equivalent/better)/ RAID-TP"</p>	Supporting three disk failure is advantage for more reliability. At least RAID 6 should be considered with required capacity.

No	RFP Specification	Vendor Query	Justification
5.23	<p>Non-disruptive upgrade should be supported. Modular software design is used, and more than 90% components are in user mode. Controllers do not need to be restarted during the upgrade.</p>	<p>Every Storage vendor provide detailed instructions and preferred methods on firmware upgrades, including the need for controller reboot. It is important that adherence to the vendor's policies is followed in order to ensure a smooth and reliable upgrade of the firmware process. Apart from this,</p> <p>1) Certain critical firmware updates may require the controllers to be rebooted to ensure the proper functioning of the upgraded firmware and to implement the new features or fixes.</p> <p>2) Restarting the controllers after a firmware upgrade can help in ensuring the integrity and stability of the updated firmware. It allows the system to initialize with the new firmware in a controlled manner, minimizing the risk of potential issues or inconsistencies in the firmware implementation.</p> <p>3) Restarting the controllers post-firmware upgrade enables the system to recalibrate and optimize performance parameters based on the updated firmware, ensuring that the storage system operates at its full potential with the new firmware features.</p>	<p>For critical applications cannot have controller restart for upgrades, which will avoid service interruptions.</p>
5.27	<p>Should support QoS to control traffic by LUN, LUN group, or host. QoS policies should include upper limit control and minimum performance assurance, which can be configured by IOPS or bandwidth. In</p>	<p>Please remove "In addition, the upper limit control policy includes the burst configuration, and the minimum performance assurance policy includes the latency configuration" as latency is dependent on many external factors like HBA speed, SAN configuration and IO Queues etc and hence setting latency is not logical. Also our storage architecture, allows to set IOPS or bandwidth which can indirectly influence latency.</p>	<p>QoS preferentially guarantees the performance of critical services. Upper limit control is for non-critical services, and minimum performance assurance is for core services. Burst is similar to traffic accumulation of carriers. When the performance is under the upper limit, the excess performance can be saved and used to exceed the upper limit when the traffic increases sharply.</p>
5.28	<p>The snapshot function should meet the following requirements:</p> <ol style="list-style-type: none"> 1. ROW lossless snapshot mode is used. 2. Supports at least 20,000 snapshots for a single LUN and at least 250K snapshots for a system. 3. The system creates a snapshot every 3 seconds. Visualized management interfaces are provided. Data can be restored using a snapshot at any point in time, without losing the snapshot data at other points in time. 4. Snapshots do not compromise performance. System performance will not decrease as the number of 	<p>Request for Modifications: Please remove the statement 'Supports at least 20,000 snapshots for a single LUN and at least 250K snapshots for a system.'</p> <p>Justification: Please remove the statement 'Supports at least 20,000 snapshots for a single LUN and at least 250K snapshots for a system.' These numbers are excessively high and impractical for most use cases.</p> <p>Additionally, managing 250K snapshots would be extremely cumbersome.</p>	<p>Snapshots are mainly used to protect data against logical faults. Should be capable to support large number of snapshots for creating high-density snapshots at a low interval and minimize data loss. Anyway vendors can mention their supported values, and above numbers are mentioned to avoid ambiguity.</p>

No	RFP Specification	Vendor Query	Justification
	<p>not decrease as the number of snapshots increases. System latency stays shorter than 1 ms with snapshot enabled.</p> <p>5. Cascading snapshots are supported. Snapshots can be created for other snapshots. Secure snapshots are supported, that is, snapshots cannot be deleted.</p>	<p>Please remove the statement 'Supports at least 20,000 snapshots for a single LUN and at least 250K snapshots for a system.' These numbers are excessively high and impractical for most use cases. Additionally, managing 250K snapshots would be extremely cumbersome</p>	
5.29	<p>The clone function should meet the following requirements:</p> <ol style="list-style-type: none"> 1. The system supports the clone function, which provides an entity copy for a snapshot and a source LUN. 2. Immediately available clones can be created. Clone consistency groups, cascading clones, and forward and reverse synchronization are supported. 3. Entity copies are supported after splitting. 	<p>Please remove this "entity copy for a snapshot" as instead of a clone a flashcopy can be also taken for a snapshot and hence clone is not mandatorily needed.</p> <p>Every vendor has their own way to provide specific functionality and hence Please remove this clause as this is vendor specific. Our storage architecture supports several features with snapshots and flashcopy which can help to implement required functionality.</p>	<p>Snapshot is for data protection. Clone is required to map LUN to the testing environment and do LUN copy for testing purposes.</p>
5.32	<p>The proposed storage should provide dedicated multipathing software (not the multi-pathing software of the operating system) to support failover and load balancing. The multi-pathing software should run on Windows.</p>	<p>We do support integration with OS native multipathing for better administration, performance and management, so please change this to "The proposed storage should provide multi-pathing software or integration with native OS multipath drivers to support failover and load balancing. The Storage should support this integration with Windows, Linux, AIX, Solaris, and other mainstream OSs"</p>	<p>Storage dedicated multipathing software will increase the performance (IOPs) of the storage compared to using third party multipathing software. This feature is required.</p>
5.34	<p>Should monitor the service life of SSDs and displays the wear degree and estimated remaining service life of each SSD.</p>	<p>Please remove this as this is OEM specific and does not bring value addition.</p>	<p>Should be able to estimate and predict the end of service life of SSDs to avoid sudden disk failures. This feature is required.</p>