



ANDHRA PRADESH TECHNOLOGY SERVICES LIMITED

(Government of AP Undertaking)

(CERT-In Empanelled, ISO 9001:2015, ISO 27001:2013 Certified)

3rd Floor, R&B Building, Opp. Indira Gandhi Municipal Stadium,

MG Road, Labbipet, Vijayawada-520010, Andhra Pradesh, India.

Ph.0866-2468102 | md_apt@ap.gov.in | <https://www.apt.gov.in>

Lr. No. ITC51-15021/27/2025-PROC-APTS

Dt. 27-01-2026

To
All prospective bidders

Sir,

Sub: APTS – Proc 2 - Procurement of IT Infrastructure for the Data Lake Project of the Real Time Governance Society, GoAP - Pre-bid Minutes issued - Reg.

Ref: 1. RFP Ref. No. ITC51-15021/27/2025-PROC-APTS, Dt. 30-12-2025

2. Pre-bid conference held on 07-01-2026

* * *

Please refer to the subject and references cited. APTS has invited bids on 30-12-2025 for “Procurement of IT Infrastructure for the Data Lake Project of the Real Time Governance Society, GoAP”. Pre bid conference was held on 07-01-2026.

The specifications, terms and conditions of the bid document were discussed during the pre-bid conference and following are the amendments/clarifications given to the bid document after discussions:

#	Pg. No	Section & clause No	As Per RFP	Suggestions / Requests by the bidders	Amendment/Clarification issued
1.	8	A.3. BDD S.No-13	Last date/time for Sale of document: 13-01-2026, 2.00 PM		Modified as: Last date/time for Sale of document: 03-02-2026, 2.00 PM
2.	8	A.3. BDD S.No-14	Bid closing date/time: 13-01-2026, 3.00 PM		Modified as: Bid closing date/time: 03-02-2026, 3.00 PM

3.	8	A.3. BDD S.No-15	Pre-qualification bid opening date & time: 13-01-2026, 3.30 PM		Modified as: Pre-qualification bid opening date & time: 03-02-2026, 3.30 PM
4.	9	A.3. BDD S.No-25	Earnest Money Deposit (EMD): Rs. 1,00,00,000/- (Rs. One Crore only)	Request to Reduce EMD Amount	As per RFP.
5.	9	A.3. Bid Data Sheet-S.no-31	Contract period: The contract period, from the date of signing of the agreement, includes 3 Years Warranty period.	Please confirm warranty type as 'comprehensive OEM warranty' for 3 years for all supplied items with back-to-back OEM support and spares availability for entire warranty. Also clarify whether warranty start is governed by installation date (as per Section D) vs contract signing date (as per Bid Data Sheet) to avoid ambiguity.	Clarified as: Warranty: The warranty shall be a comprehensive OEM-backed warranty for 3 years for all supplied hardware, software, licenses, optics, and accessories, with back-to-back OEM support and assured spares availability for the full warranty term. Warranty for all items shall commence from the date of successful Final Acceptance Test (FAT) certification issued by RTGS for the scope of work of this RFP, and not from the dates of delivery, installation, or contract signing. OEM support shall be available 24x7, and spares must be arranged on a next-business-day, 4-hour response basis (or better) to meet the uptime and SLA requirements
6.	10	B.1. Eligibility Criteria-S.no-2	AP GST Registration: Bidder should have at least one office with GST Registration in any of the 26 districts of AP.	We request RTGS to remove this clause as IBM doesn't have any registered office in AP.	Clarified as: This condition is for the bidder. Not for OEM.
7.	11	B.1. Eligibility Criteria-S.no-5	Financial Turnover Criteria: The Bidder should have a minimum annual average turnover of Rs. 100 Crore (Rs. Twenty Three Crore) in the last three financial years i.e. 2022-23, 2023-24 & 2024-25.	Current Requirement: Minimum annual average turnover of ₹100 Crore over the last three financial years (2022-23, 2023-24, 2024-25). Relaxation Requested:	As per RFP.

				Consideration of total turnover instead of average annual turnover for 3 years including current year (2023-24, 2024-25, 2025-26). Request for amendment of Annual Turnover to Minimum 50 Crore for the average of last three financial years	
8.	11	B.1. Eligibility Criteria- S.no-7	Project Experience Criteria The bidder should have executed work orders involving supply of Servers during the last 3 financial years (i.e. 2022-23, 2023-24 & 2024-25) as below: a) At least one work order with value of Rs. 50 Crore or more. b) At least two work orders with value of Rs. 27 Crore or more. c) At least three work orders with value of Rs. 20 Crore or more.	We request you to amend this clause as below: The Bidder/OEM should have executed work orders involving supply of Servers during the last 3 financial years (i.e. 2022-23, 2023-24 & 2024-25) as below: a) At least one work order with value of Rs. 50 Crore or more. b) At least two work orders with value of Rs. 27 Crore or more. c) At least three work orders with value of Rs. 20 Crore or more. By this amendment it will be allowing more bidders to participate as large project are executed by OEMs. We request you to kindly consider our request.	As per RFP.
9.	15	D.1. S.no-6	Scope of Work: Successful bidder has to supply, install & maintain all the items including re-installation of Operating system and other applications incase gets corrupted. In case, the supplied items are down and not working, same need to be repaired and restored for normal functioning as per agreed Service Level Requirements. Failing which penalty will be recovered from Performance Security.	1. Since the bare-metal server operating system is not within the scope of the RFP, please clarify whether the installation and reinstallation of the OS, in case it gets corrupted, are under the supplier's scope. 2. The term installation mentioned in the scope refers only to rack mounting of the server and installation of the latest firmware, configuration of management interface. Installation of the operating system and any other application software is not within the supplier's scope.	Modified as: The Successful Bidder shall supply, install, configure, and maintain all hardware and bundled software components in the BoM, and shall restore them to normal functioning as per the Service Level Requirements. 'Installation' includes full rack-mounting, cabling, firmware/BIOS configuration and tuning for 100G/200G RDMA throughput, and installation of the base OS / hypervisor required to bring the platform to a 'hardware-ready' state for application deployment. Re-installation of the base OS

				3. Please confirm whether the supply and installation of the operating system for virtual machines (VMs) in the virtualization infrastructure are within the scope of the RFP. Kindly confirm whether 'maintain' during warranty includes comprehensive onsite support (parts + labor) for entire BoM with 24x7 support and defined response/resolution SLAs. Also clarify whether OS/Application media/licenses are in bidder scope or will be provided by Purchaser (since re-installation responsibility is mentioned).	/ hypervisor on bare-metal servers and virtualization hosts, in case of corruption or failure, is within the bidder's scope. Application software and VM guest operating systems (and their licenses) shall be provided by RTGS or by a separate application SI, but the Infra bidder must cooperate in re-installing them on the supplied infrastructure when required. Maintenance during warranty shall be comprehensive onsite support (parts + labor) for the entire BoM with 24x7 support and response/resolution SLAs as defined in the RFP; Penalties for non-compliance shall be recovered from the Performance Security as specified.
10.	15	D.1. S.no-7	Scope of Work: (i) Warranty period as specified in the BoM, will start from the date of Installation. (ii) Incase of Site Not Ready (SNR Certificate to submit) warranty will start after 30days from the date of delivery.	Request clarification/amendment: If installation is delayed due to Purchaser/User site readiness beyond 30 days, warranty should commence from actual installation/commissioning date (or mutually agreed Go-Live/AT date), not automatically after 30 days of delivery, to ensure full usable warranty period.	As per RFP.
11.	15	D.1. S.no-9	Scope of Work: Bidder/OEM shall provide service desk portal for call logging and monitoring & resolution. The original call log for all the logged calls of complaints & calls closed status should be sent by email to APTS/Department on monthly basis for monitoring.	Please clarify whether the Service Desk portal referred in Clause 9 is only the OEM's support portal for logging and tracking hardware warranty-related incidents, with the Bidder's role limited to coordination and facilitation with OEM. Kindly confirm that no separate service desk tool or dedicated manpower (onsite/remote) is expected from the Bidder during the warranty period, and that monthly call log reports can be generated from the OEM portal and shared with APTS/Department.	Add as: Service Desk portal & reporting: The 'Service Desk portal' in this clause refers to the OEM's / Bidder's support portal used for logging, tracking, and closing incidents related to the supplied infrastructure. No separate service desk tool is mandated in addition to this portal. APSDC/RTGS will operate the central NOC / helpdesk; however, the Successful Bidder shall ensure that all hardware and warranty incidents are logged through the

					OEM/Bidder portal and shall act as the single point of contact for coordination with OEMs to meet the defined response and resolution SLAs (including 24×7 and 4-hour response for critical components). The Bidder shall generate and share a monthly SLA / call-log report from the OEM/Bidder portal, showing all calls raised, status, and closure details, with APTS/ Department for monitoring.
12.	15	D.1. S.no-10	Scope of Work: Along with the above mentioned call log, a date wise abstract of calls logged and repair status within SLA and outside SLA shall be provided on the dashboard by Bidder/OEM.	SLA reports will be extracted from the EMS provided by RTGS department as OEM /Bidder cannot be configure or feasible to provide a separate dashboard to provide the SLA reports from the OEM Portal.	Modified as: SLA reporting & reconciliation: The central SLA dashboard for runtime/uptime shall be provided through the RTGS Enterprise Monitoring System (EMS). The Successful Bidder / OEM is not required to deploy a separate web dashboard, but shall provide an automated monthly SLA reconciliation report (CSV/Excel) exported from the OEM/Bidder portal, containing ticket IDs, components affected, timestamps, and resolution details. This report shall be mappable to RTGS incident IDs and will be used by RTGS to validate SLA performance and to compute any applicable penalties without manual data re-entry.
13.	15	D.1. S.no-12	Scope of Work: At its discretion, APTS/Department may enter into Annual Maintenance Contract with the supplier after completion of warranty period. AMC Payments will be released to the supplier quarterly after successful maintenance of the equipment under AMC for that period. Service Level Agreement mentioned in this RFP during	We request RTGS to consider Annual Advance payment term.	As per RFP.

			warranty period is also applicable for AMC period.		
14.	16	D.2. Item-1-- CL 01- Servers for HDFS Data Nodes – S.no-2	CPU: Latest series Processor/s with 32 Physical Cores with Hyperthreading (64 logical cores) are to be supplied with provision to expand up to 64 Physical Cores.	Request to clarify if 32 physical cores to be provided from server as overall core or from single CPU? Need to mention Single socket Justification: Efficient Utilization of Compute resources & IO's and better cost efficiency. Can be scaled to 64 physical cores by just replacing the CPU. Also reducing CPU to CPU latency in dual socket scenario.	Modified as: CPU (CL-01 – HDFS Data Nodes): Each CL-01 server shall be populated with latest-generation processor 1 x 64 physical cores CPU with Hyper-Threading. Platforms may use one or more such CPUs per node as per OEM design, but each CPU must provide at least 64 physical cores. The overall platform shall be PCIe Gen 5.0 x16–based and sized to support AI-scale HDFS workloads and NVMe-over-Fabric, as per the technical requirements.
15.	16	D.2. Item-1- CL 01 - Servers for HDFS Data Nodes (Bare metal servers) -S.no-3	Memory: 512GB DDR5 Registered DIMM (RDIMM) scalable up to 8TB.	Change request : Please change this statement to : 512GB DDR5 Registered DIMM (RDIMM) scalable up to 6TB. Justification: Currently, only Intel (single OEM) platforms qualify for the 8 TB memory requirement. We request to kindly amend this specification to 6 TB memory, which will allow qualification of both AMD and Intel CPU–based servers, ensuring fair competition and wider OEM participation. From a practical deployment perspective, server memory utilization will not exceed 6 TB under any circumstances, even in peak workload scenarios. Therefore, specifying 6 TB fully meets the operational and scalability requirements while avoiding unnecessary vendor lock-in. Scalable upto 4TB	Modified as: Memory (CL-01 – HDFS Data Nodes): Each CL-01 server shall be populated with 512 GB DDR5 RDIMM, scalable up to 6 TB per node. Platforms from any OEM (Intel or AMD) that support at least 6 TB memory scalability per node on PCIe Gen 5.0 architecture will be considered compliant; configurations limited below 6 TB scalability shall be treated as non-responsive.

				Justification: Intel comes with 16DIMMs(256*16) configuration & AMD 24DIMMs(256*24).	
16.	16	D.2. Item-1-CL 01-Servers for HDFS Data Nodes – S.no-4	Memory Protection: Advanced ECC with multi-bit error protection.	Request to amend to · Advanced ECC with multi-bit error protection or equivalent. Justification - Request to amend in order to allow other OEMs to participate as every OEMs will have its own terminology for memory features and functionality	Modified as : Memory Protection – Advanced ECC: Server memory shall support Advanced ECC with multi-bit error protection or equivalent technology. ‘Equivalent’ means the memory subsystem must at minimum provide Single Device Data Correction (SDDC) or similar capability, allowing correction of all bit errors arising from the failure of a single DRAM device on a DIMM, so as to prevent data corruption or system crashes during long-running AI and analytics workloads.
17.	16	D.2. Item-1-CL 01-Servers for HDFS Data Nodes - S.no-4	Memory Protection: Online spare	Request to amend to · Online spare or equivalent. Justification - Request to amend in order to allow other OEMs to participate as every OEMs will have its own terminology for memory features and functionality.	Modified as: Memory Protection – Online spare: ‘Online spare’ or equivalent fault-tolerant memory feature shall be supported. ‘Equivalent’ includes OEM mechanisms such as rank sparing, fault-resilient memory modes, or similar technologies that can automatically retire faulty memory areas and continue operation without unplanned downtime, in order to meet the required cluster availability targets.
18.	16	D.2. Item-1-CL 01-Servers for HDFS	Internal Storage: 2 x 960GB M.2 SSD in RAID 1 Configuration for OS and logs. 2 x 7.68TB NVME Drives 10 x 16TB SATA/SAS/SSD Drives	Please clarify if 16TB is SATA/SAS Hard disk or SAS/SATA SSD? Change request: Please change the statement as 15TB of RAW NVME storage Please change the statement as 160TB of SATA/SAS/SSD RAW storage	Modified as: Internal Storage (CL-01 – HDFS Data Nodes): Each CL-01 server shall provide hybrid internal storage as specified: (a) 2 × 960 GB M.2 SSD in RAID-1 for OS and logs; (b) 2 × 7.68

		Data Nodes-S.no-6		Justification: By revising the specification to define a raw storage requirement, all OEMs will be provided the flexibility to propose disk options available within their respective portfolios. This ensures broader OEM participation, avoids vendor-specific constraints, and still fully meets the intended capacity and performance objectives. SSD drives instead of 7.68TB NVMe drives. Justification: Since 16TB comes in LFF/3.5 Drives (SAS/SATA/SSD), we cannot provide NVMe drives, it has to be SSD/SAS only.	TB NVMe drives for high-IOPS tier; and (c) 10 × 16 TB SAS/SATA HDD or SSD drives for capacity tier. The 16 TB devices refer to 3.5-inch (LFF) SAS/SATA drives; bidders may offer higher-performance media (e.g., enterprise SSD) in this tier provided the minimum raw capacity per node is not reduced and the solution continues to support concurrent management of both NVMe and SAS/SATA devices via a suitable controller. Requests to remove or downgrade the NVMe tier are not accepted; NVMe-based tiering is mandatory to meet the IOPS and throughput requirements of the target AI and analytics workloads.
19.	16	D.2. Item-1-CL 01-S.No. 8	Network Cards: • 1x 1GbE 4-Port Adapter • Ethernet 10/25Gb 2-port SFP28 Adapter		Modified as: Network Cards: • Dedicated 1GbE out-of-band management interface for remote server administration. • Dual-port 200Gbps RDMA-capable (RoCE v2) QSFP NIC for ultra-low-latency, high-throughput data workloads
20.	16	D.2. Item-1-CL 01-Servers for HDFS Data Nodes-S.no-9	FC HBA: 1x Dual Port 32Gbps FC HBA	Change request: Request you to kindly remove Justification: As there is no SAN requirement in the proposed solution, we kindly request that the SAN-specific specification be removed, as it is not applicable to the intended architecture. Need to be removed FC HBA	Deleted

				Justification: Since no external storage is required	
21.	16	D.2. Item-1- CL 01- Servers for HDFS Data Nodes- S.no-10	Interfaces: USB 2.0/3.0 – 4 Nos.	USB 2.0/3.0 - 3 Nos. Justification - USB ports are generally used for Keyboard or mouse connectivity whereas these servers are for HDFS data nodes which does not require any peripheral connectivity as they would be kept in a DC rack. A min of 3 USB ports are standard and will be sufficient. Not all the OEMs will have 4 USB ports and hence request to amend for other leading OEMs to qualify & participate.	Modified as: Interfaces: USB 2.0/3.0 – 3 Nos.
22.	16	D.2. Item-1- CL 01- Servers for HDFS Data Nodes - S.no-11	Bus Slots: Should have • PCI-Express 5.0 x16 slots – 8 Nos. • PCIe 5.0x8 slots – 2 Nos.	Request to amend to - PCI-Express 4.0 / 5.0 x8 or x16 slots - 6 Nos. Justification - As per RFP requirement for NIC/HBA cards only 2 PCIe 4.0 x8 slots will be required on server. Hence total 6 PCIe slots are more than sufficient for future scalability. This is restricting other OEMs to comply and hence request to amend for other leading OEMs to qualify & participate. Change request: PCI-Express 5.0 x16 slots – 2 Nos. Justification: The specified PCIe slot configuration is specific to a single OEM. To enable broader OEM participation, we kindly request that this requirement be removed, as it does not impact the functional or operational requirements of the department. PCIe 5.0 Slots minimum 3 or above Justification: Since required cards in the server are two for ethernet connections 10 slots is on higher end.	Modified as: Bus Slots (CL-01 – HDFS Data Nodes): Each CL-01 server shall be based on PCIe Gen 5.0 x16 architecture and shall provide a minimum of 3 × PCI-Express 5.0 slots (x16 electrical) for high-speed adapters (e.g., dual 100G/200G RDMA NICs and future accelerators). Additional PCIe slots may be provided as per OEM design, but platforms limited to PCIe Gen 4.0 or with fewer than 3 usable PCIe 5.0 x16-electrical slots shall be treated as non-compliant.

23.	16	D.2. Item-1- CL 01- Servers for HDFS Data Nodes - S.no-12	Power Supply: 2 x 1600 Watts Platinum Hot Swap Power Supplies	Request to amend to - 2 x 1400 Watts or higher Platinum Hot Swap Power Supplies Justification - Every OEM will have its own power design and accordingly power supply capacity. Since 1600W is OEM specific we request to amend the clause to allow other OEMs to participate. Required power supply for the entire BOM Justification: Since overall power consumption with single CPU, memroy & HDD will not exceed 1500W.	Modified as: Power Supply (CL-01 – HDFS Data Nodes): CL-01 servers shall be equipped with redundant, hot-plug, high-efficiency ($\geq 96\%$) power supplies in N+1 configuration, sized as per OEM guidelines to support the full load of all CPUs, NVMe devices, and NICs with one PSU failed. A specific wattage value (e.g., 1600 W or 1400 W) is not mandated, but bids must include OEM sizing documents showing that the proposed PSU configuration can sustain worst-case power draw under N+1 conditions without derating or performance throttling.
24.	16	D.2. Item-1- CL 01- S.no-14	Industry Standard Compliances: ACPI 6.3 Compliant	ACPI 6.1 Compliant in all mentioned servers. Justification: Since 6.3 is not available for AMD OEMs.	Modified as: Industry Standard Compliances: ACPI 6.1 Compliant (applicable for all servers)
25.	17	D.2. Item-1- CL 01- Servers for HDFS Data Nodes- S.no-16	Latest OS and Latest Virtualization Software Support: Shall support Microsoft Windows Server 2019 or above/RHEL 7 or above/Ubuntu 18 or above/VMware ESXi 7 or above	Request to amend to - Shall support Microsoft Windows Server 2019 or above/RHEL 9 or above/Ubuntu 22 or above/VMware ESXi 8 or above Justification - As per requirement latest processor to be supplied which come with latest gen servers. Also, Latest OS and Latest Virtualization Software Support is mentioned in the clause heading. Hence, we request to consider latest OS & hypervisor support or at least one gen older version.	Modified as: Latest OS and Latest Virtualization Software Support: Servers shall support current-generation operating systems and hypervisors, including at minimum: Microsoft Windows Server 2019 or above, RHEL 9 or above, Ubuntu 22.04 LTS or above, and VMware ESXi 8.x or above, along with the necessary drivers for 100G/200G RDMA NICs and PCIe Gen 5.0 storage. Backward compatibility with earlier supported versions (RHEL 8, ESXi 7, etc.) is acceptable where required by RTGS, but

					lack of support for these latest stable versions shall be treated as non-compliant.
26.	17	D.2. Item-1- CL 01- Servers for HDFS Data Nodes- S.no-17	GPU support: System should support NVIDIA's latest computational accelerators and graphics accelerators.	Need to be removed. Justification: Since GPU are not getting installed.	No Change. As per RFP as we are asking only support.
27.	18	D.2. Item 2: CL-02 - Ozone Nodes (Bare metal servers)- S.no-2	Platform: Offered Platform shall be a dense platform supporting 24 x Large Form factor drives so that it can be used for multi-purpose uses like Software defined storage, Scale out databases like Mongo DB, Cassandra, Data lake like Ezermal Data Fabric, Hadoop Data Lake etc.	Change to 12LFF Justification: On the latest GenPCI5.0 max supported LFF is 12, which is mix of NLSAS/SATA/SSD/No NVMe	No change. As per RFP.
28.	18	D.2. Item 2: CL-02 - Ozone Nodes (Bare metal servers)- S.no-3	CPU: Minimum two numbers of latest Scalable Processors supporting up to 48 cores per processor and shall be offered with minimum 16 Physical Cores per processor with hyper threading.	Request to amend to - Minimum two numbers of latest Scalable Processors supporting up to 28cores per processor and shall be offered with minimum 16 Physical Cores per processor with hyper threading. Justification - We meet the requirement of 16 physical cores per processor and request to amend the scalable cores to 28 cores. The reason being the purpose of the server is to provide high internal storage capacity and not high compute. The dense server will not support high no of cores and hence request to amend to allow other leading OEMs to comply and participate. Moreover, the server processors are not upgraded in practical scenario post deployment.	Modified as: CPU (CL-02 – Ozone Nodes): Each CL-02 Ozone node shall be populated with a minimum of two latest-generation scalable processors, each providing at least 32 physical cores with Hyper-Threading. The platform must support latest generation CPUs (≥32 cores per socket) and to ensure long-term scalability of the object-storage tier.

				Need more clarity on the CPU cores which is required by default. Justification: Since it's not very clear how many cores are required by default, not able to provide exact processor.	
29.	18 & 19	D.2. Item 2: CL-02 - Ozone Nodes (Bare metal servers)- S.no-6	HDD Bays: 1. Offered Server shall be supplied with 24 _ X 16 TB NL SAS drives and 3x3.84 TB NVMe SSD drives. 2. Offered server shall be supplied with at-least Dual M.2 NVMe boot drives of 960GB each without consuming any drive slot in Raid 1.	Change request: 384TB of RAW NL-SAS storage along with 11TB of RAW NVME storage. Justification: By revising the specification to define a raw storage requirement, all OEMs will be provided the flexibility to propose disk options available within their respective portfolios. This ensures broader OEM participation, avoids vendor-specific constraints, and still fully meets the intended capacity and performance objectives. 10*24TB NL SAS+2*7.68 SSD Justification: On the latest GenPCI5.0 max supported LFF is 12, which is mix of NLSAS/SATA/SSD/No NVMe.	No change. As per RFP.
30.	19	D.2. Item 2: CL-02 - Ozone Nodes (Bare metal servers)- S.no-8	Controller: 1. Boot drives shall be configured in Raid 2. Vendor shall offer additional Raid controller card for LFF drives with following minimum features: 3. PCIe 4.0 based and shall have minimum of 8GB Flash Back write cache (FBWC) 4. Shall have minimum of 32 ports or Lanes. 5. Shall be supported to work in both Raid and HBA mode.	Request to amend to - 4. Shall have minimum of 16 ports or Lanes. 6. Shall support RAID 0,1,5,6, 10, 50, 60. 7. Shall support at-least 64 number of Volumes / Virtual disks. 8. Shall support extending logical drive/Online Capacity Expansion, background surface scan / background consistency check, instant secure erase, Migration of raid / stripe size as well as modification of write cache policy.	No change. As per RFP.

			6. Shall support RAID 0,1,5,6, 10, 50, 60, Raid 1 with triple copy and Raid 10 with Triple copy. 7. Shall support at-least 64 number of Volumes. Shall support extending logical drive, background surface scan, instant secure erase, Migration of raid / stripe size as well as modification of write cache policy.	Justification - The specifications are specific to one OEM based on spec sheet. We request to amend to allow other leading OEMs to comply and participate.	
31.	19	D.2. Item 2: CL-02 - Ozone Nodes (Bare metal servers)- S.no-9	Raid Controller Certification & other Critical features: 1. Offered H/W raid controller shall be certified or validated with both VMWARE VSAN and Microsoft Space direct. 2. Offered H/W raid controller shall support at-least 64 Logical drives and shall support rapid re-build and parity initialization. 3. Offered H/W raid controller shall support SSD accelerations for I/O requests so that attached SSD disks can speak to controllers directly while bypassing the controller firmware layer.	Request to amend to – 1.Offered H/W raid controller shall be certified or validated with both VMWARE VSAN and Microsoft Space direct. 2. Offered H/W raid controller shall support at-least 64 Logical drives and shall support re-build and parity / full initialization. 3. Offered H/W raid controller shall support SSD accelerations for I/O requests so that attached SSD disks can speak to controllers directly while bypassing the controller firmware layer or Non-RAID (Pass Through Mode). Justification - Request to remove VMware vSAN and MS Space direct certification as this is not related to the requirement of base metal physical servers. Also request to consider other equivalent feature for other leading OEMs to comply and participate.	Modified as: 1. Offered H/W raid controller shall be certified or validated with both VMWARE VSAN and Microsoft Space direct. Certification or validation with VMware vSAN and Microsoft Storage Spaces Direct is not mandatory; however, equivalent enterprise-class OS/hypervisor support (e.g., RHEL, ESXi, KVM) and published hardware compatibility lists shall be provided. 2. Offered H/W raid controller shall support at-least 64 Logical drives and shall support re-build and parity / full initialization. 3. Offered H/W raid controller shall support SSD accelerations for I/O requests so that attached SSD disks can speak to controllers directly while bypassing the controller firmware layer or Non-RAID (Pass Through Mode).
32.	19	D.2. Item 2: CL-02 - S.No.11	Networking features: 1Gb 4-port network adaptors and Offered Server shall be supplied with at-least 2 x 25GbE Ports.		Modified as: Network Cards:

					- Dedicated 1GbE out-of-band management interface for remote server administration. - Dual-port 200Gbps RDMA-capable (RoCE v2) QSFP NIC for ultra-low-latency, high-throughput data workloads
33.	19	D.2. Item 2: CL-02 - Ozone Nodes (Bare metal servers)- S.no-11	Bus Slots: 1. Minimum 6 x PCI-Express 5.0 slots, All the slot shall be at-least x16. 2. Offered server shall also support 2 x OCP ports at PCI 5.0 Speed.	Request to amend to - 1. Minimum 4 x PCI-Express 4.0 or higher slots, All the slot shall be at-least x16. 2. Offered server shall also support 1 x OCP port at PCI 4.0 or higher Speed. Justification - The RFP requirement of 1 no of 1GbE quad port card and 1 no of 25G card can be easily met with 1 OCP slot and 1 PCIe slot itself. Hence, we request to amend the clause which does not compromise on the solution point of view and allows other leading OEMs to comply and participate. Change request: 3 no PCIe 5.0 x16 Justification: The specified PCIe slot configuration is specific to a single OEM. To enable broader OEM participation, we kindly request that this requirement be removed, as it does not impact the functional or operational requirements of the department. PCIe 5.0 Slots minimum 3 or above Justification: Since required cards in the server are two for ethernet connections 10 slots is on higher end.	Modified as: Bus Slots (CL-02 – Ozone Nodes): Each CL-02 server shall be based on PCIe Gen 5.0 x16 architecture and shall provide at least 4 × PCI-Express 5.0 slots (x16 electrical) for high-speed adapters (e.g., dual 100G/200G RDMA NICs and future accelerators). In addition, the server shall support at least 1 × OCP slot at PCIe 5.0 speed. Platforms limited to PCIe Gen 4.0 or with fewer than 4 usable PCIe 5.0 x16-electrical slots shall be treated as non-compliant, as PCIe 5.0 bandwidth (32 GT/s per lane) is mandatory to support 100G/200G NICs and NVMe tiering without saturating the bus.
34.	19	D.2. Item 2: CL-02 -	Power Supply: 1. Should support hot plug Platinum redundant low halogen power supplies.	Request to amend to - 1. Should support hot plug Platinum redundant power supplies.	Modified as:

		Ozone Nodes (Bare metal servers)- S.no-12	2. Server shall be configured with at-least dual 1600W power supply.	2. Server shall be configured with at-least dual 1400W or higher power supply. Justification - Every OEM will have its own power design and accordingly power supply capacity. Since 1600W is OEM specific we request to amend the clause to allow other OEMs to participate. Change request: server should provided with dual redundant power supplies. Justification: Server power requirements vary across OEMs, and the specified 1600W power supply is specific to a single OEM. To enable broader OEM participation, we kindly request that this requirement be amended, as it does not impact the functional or performance requirements of the department.	Power Supply (CL-02 – Ozone Nodes): CL-02 servers shall support hot-plug, redundant Platinum or Titanium power supplies with N+1 redundancy. The server shall be configured with two (or more) identical PSUs sized as per OEM guidelines to safely support the full load of all CPUs, drives, NVMe tier, and NICs even under single-PSU failure, without derating or performance throttling. ‘Low-halogen’ compliance is required as per environmental standards. A fixed wattage value (e.g., 1600W or 1400W) is not mandated, but the bidder must submit OEM power-sizing documentation demonstrating that the proposed PSU configuration is sufficient for the worst-case configuration under N+1 conditions.
35.	19	D.2. Item 2: CL-02 - Ozone Nodes (Bare metal servers)- S.no-13	Fans Redundant hot-plug system fans. At-least 5 - Dual rotor fan shall be supplied.	Request to amend to - Redundant hot-plug system fans. At-least 5 - Dual rotor / Standard fans shall be supplied. Justification - Dual rotor is a specific term used on one of the OEM spec sheet. Hence, we request to amend the clause to allow other OEMs to participate.	Modified as: Fans (CL-02 – Ozone Nodes): CL-02 servers shall include redundant, hot-plug system fans, with a minimum of 5 high-performance fans per node. Fans shall be sized to provide sufficient static pressure and airflow to cool a fully populated 24-bay LFF chassis and the associated CPUs/NICs under continuous, high-density workloads, ensuring no thermal throttling under worst-case conditions. OEM may implement this using dual-rotor or equivalent high-performance fan designs.
36.	19 & 20	D.2. Item 2: CL-02 - Ozone	Industry standard Compliance: • ACPI 6.4 Compliant • PCIe 5.0 Compliant • WOL Support	Request to amend to - • ACPI 6.4 Compliant • PCIe 5.0 Compliant • WOL Support	Modified as: Industry standard Compliance (CL-02 – Ozone Nodes): CL-02 servers shall

		Nodes (Bare metal servers)- S.no-14	• Microsoft® Logo certifications • PXE Support • USB 3.2 Gen1 Compliant • SMBIOS 3.4 • Redfish API • IPMI 2.0 • Secure Digital 4.0 • TPM 2.0 support • Advanced Encryption Standard (AES) • Triple Data Encryption Standard (3DES) • SNMP v3 • TLS 1.2 • DMTF Systems Management Architecture for Server Hardware Command Line (SMASH CLP) • One Button Secure Erase	• Microsoft® Logo certifications • PXE Support • USB 3.0 Gen1 • Compliant SMBIOS 3.3 • Redfish API • IPMI 2.0 • Secure Digital 4.0 • TPM 2.0 support • Advanced Encryption Standard (AES) • Triple Data Encryption Standard (3DES) • SNMP v3 • TLS 1.2 • DMTF Systems Management Architecture for Server Hardware Command Line (SMASH CLP) • One button Secure Erase Justification - The specifications are specific to one OEM based on its spec sheet. Also, 3DES and SD 4.0 are obsolete, hence we request to amend to allow other leading OEMs to comply and participate.	comply with the following minimum standards: ACPI 6.1 or higher, PCIe 5.0, Wake-on-LAN (WOL), Microsoft® logo certifications, PXE boot, USB 3.0/3.2 Gen1, SMBIOS 3.3 or higher, Redfish API, IPMI 2.0, Secure Digital 4.0, TPM 2.0, Advanced Encryption Standard (AES-256) for management and data-at-rest features, SNMP v3, TLS 1.2 or higher (TLS 1.3 preferred), and DMTF SMASH-CLP. A secure-erase / one-button secure erase capability conforming to NIST SP 800-88 (or equivalent) shall be provided for rapid, policy-driven sanitization of local storage during decommissioning;
37.	20	D.2. Item 2: CL-02 – Ozone Nodes (Bare metal servers)- S.no-15	Security: • UEFI Secure Boot and Secure Start support • Immutable Silicon Root of Trust • FIPS 140-2 validation • Common Criteria certification • Configurable for PCI DSS compliance • Advanced Encryption Standard (AES) and Triple Data Encryption Standard (3DES) on browser • Support for Commercial National Security Algorithms (CNSA) • Granular control over remote access interfaces	• UEFI Secure Boot support • Immutable Silicon Root of Trust • FIPS 140-2 validation • Common Criteria certification • Configurable for PCI DSS compliance • Advanced Encryption Standard (AES) and Triple Data Encryption Standard (3DES) on browser • Support for Commercial National Security Algorithms (CNSA) • Granular control over remote access interfaces • Smart card (PIV/CAC) and Kerberos based 2-factor Authentication	Modified as: Security (CL-02 – Ozone Nodes): CL-02 servers shall provide a hardware-anchored, standards-based security stack including at least: UEFI Secure Boot / Secure Start; immutable silicon Root of Trust; platform or management firmware validated to FIPS 140-2 or 140-3 and Common Criteria (where applicable); configurability to support PCI-DSS-aligned hardening; AES-based cryptographic protection (AES-256 or stronger) for management and data-at-rest features;

			• Smart card (PIV/CAC) and Kerberos based 2-factor Authentication • Tamper-free updates – components digitally signed and verified • Secure Recovery – recover critical firmware to known good state on detection of compromised FW • Ability to rollback firmware • Secure erase of NAND • TPM (Trusted Platform Module) 2.0 option • Bezel Locking Kit Chassis Intrusion detection	• Tamper-free updates – components digitally signed and verified • Secure Recovery – recover critical firmware to known good state on detection of compromised FW • Ability to rollback firmware • Secure erase of NAND • TPM (Trusted Platform Module) 2.0 option • Bezel Locking Kit • Chassis Intrusion detection Justification – The specifications are specific to one OEM based on its spec sheet. Also, 3DES are obsolete, hence we request to amend to allow other leading OEMs to comply and participate.	support for Commercial National Security Algorithms (CNSA) or equivalent modern cipher suites; granular control over remote access interfaces; smart-card (PIV/CAC) and Kerberos-based 2-factor authentication; tamper-free updates with digitally signed firmware components; secure recovery to a known-good firmware state and ability to roll back firmware; secure erase of local non-volatile storage; TPM 2.0; and physical protections including bezel lock and chassis-intrusion detection. OEM-named implementations (e.g., SafeBIOS, Sure Start) are acceptable as long as they are backed by recognized certifications (FIPS / Common Criteria) and meet the above capabilities.
38.	21	D.2. Item 2: CL-02 - Ozone Nodes (Bare metal servers)- S.no- 20(8 th Point)	Embedded Remote Management and firmware security: 8. Remote console sharing up to 6 users simultaneously during pre-OS and OS runtime operation, Console replay - Console Replay captures and stores for replay the console video during a server's last major fault or boot sequence. Microsoft Terminal Services Integration, 128-bit SSL encryption and Secure Shell Version 2 support. Should provide support for AES and 3DES on browser. Should provide remote firmware update functionality. Should provide support for Java free graphical remote console.	Request to remove 3DES as its obsolete and has been replaced by AES, hence we request to amend to allow other leading OEMs to comply and participate.	Modified as: Embedded Remote Management & Firmware Security – Console & Crypto: The embedded remote management controller shall support: remote console sharing by multiple users during pre-OS and OS runtime; console replay that captures and stores console video for replay of the last major fault or boot sequence; integration with Microsoft Terminal Services; 128-bit or stronger TLS encryption and SSHv2; browser-based access using modern AES-secured cipher suites (AES-256 or stronger); remote firmware update; and a Java-free graphical remote console. Console Replay (or functionally equivalent 'last crash video' capability)

					is mandatory, as it serves as a black-box record for root-cause analysis and SLA enforcement.
39.	22	D.2. Item 2: CL-02 - Ozone Nodes (Bare metal servers)- S.no- 20(13 th Point)	Embedded Remote Management and firmware security: 13. One-button Secure Erase designed to decommission/repurpose servers.	Request to remove One-button as its specific to one OEM.	Modified as: Embedded Remote Management and firmware security: 13. One-button Secure Erase (or equivalent) designed to decommission/repurpose servers.
40.	22	D.2. Item 2: CL-02 - Ozone Nodes (Bare metal servers)- S.no- 20(15 th Point)	Embedded Remote Management and firmware security: 15. Workload Performance Advisor - For Providing server tuning recommendations to improve server performance.	Request to amend to - 15. Workload Performance Advisor / BIOS Performance and Workload Tuning Guide- For Providing server tuning recommendations to improve server performance. Justification - Workload Performance Advisor is a OEM specific term and request to amend to allow other leading OEMs to comply with equivalent feature.	Modified as: Embedded Remote Management & Firmware Security – Workload tuning: The platform shall provide an automated workload performance advisory engine integrated with the server management controller, which monitors system utilization and provides real-time tuning recommendations (e.g., BIOS, power, and I/O settings) specifically for Big Data / HDFS / Ozone and AI/ML workloads. The capability must be implemented as a live, telemetry-driven advisory feature accessible through the management interface.
41.	22	D.2. Item 2: CL-02 - Ozone Nodes (Bare	Server Management: Zero Touch Provisioning (ZTP) using SSDP with remote access	Request to amend to - Zero Touch Provisioning (ZTP) using SSDP with remote access Justification - While we meet the Zero Touch Provisioning for server deployment, request to	Modified as: Server Management: Zero Touch Provisioning (ZTP)

		metal servers)- S.no-21		remove SSDP as it is specific to OEM. Hence request for amendment for larger OEM participation.	
42.	23	D.2. Item 2: CL-02 - Ozone Nodes (Bare metal servers)- S.no-22	Cloud Enabled Monitoring and Management: - Secure connection from customer sites to HPE cloud service - Unified Identity & Access Management - Manages and controls servers regardless of physical location - Subscription-based entitlement - Efficient Device Onboarding - Firmware Update Awareness with Intelligent delta-only based updates - Set Group firmware Baseline and Compliance monitoring and notification - Group based firmware management that can be scheduled or on-demand - Remote Site management with low bandwidth/high latency network connectivity - Role-based access and views for managed customer environments - GUI and Rest APIs for core features	Request to amend to - - Secure connection from customer sites to OEM cloud service. - Firmware Update Awareness with Intelligent delta-only based updates / Efficient policy driven firmware management with automation Justification - We request to amend the clause as these are one OEM specific and restrict other leading OEMs participation. Change request: - Manages and controls servers regardless of physical location - Efficient Device Onboarding - Role-based access and views for managed customer environments - GUI and Rest APIs for core features Justification: The highlighted points are specific and proprietary in nature. We kindly request that these be removed to enable broader participation.	Modified as: Cloud Enabled Monitoring and Management: - Secure connection from customer sites to OEM cloud service - Unified Identity & Access Management - Manages and controls servers regardless of physical location - Subscription-based entitlement - Efficient Device Onboarding - Firmware Update Awareness with Intelligent delta-only based updates / Efficient policy driven firmware management with automation - Set Group firmware Baseline and Compliance monitoring and notification - Group based firmware management that can be scheduled or on-demand - Remote Site management with low bandwidth/high latency network connectivity - Role-based access and views for managed customer environments. - GUI and Rest APIs for core features
43.	23	D.2. Item 3: CL-03 – S.no -2	CPU: Latest series Processor/s with 32 Physical Cores with Hyperthreading are to be supplied with provision to expand up to 64 Physical Cores.	Request to clarify if 32 physical cores to be provided from server as overall core or from single CPU? Need to mention Single socket Justification:	Modified as: CPU (CL-03 – Management / Control Nodes): Each CL-03 server shall use latest-generation scalable processors with a minimum of 64 physical cores per CPU with Hyper-Threading, on

				Efficient Utilization of Compute resources & IO's and better cost efficiency. Can be scaled to 64 physical cores by just replacing the CPU. Also reducing CPU to CPU latency in dual socket scenario.	a dual-socket-capable motherboard. Bidders shall offer at least 1 × 64-core CPU initially; the platform shall support expansion to higher core counts (e.g., adding a second CPU) without replacing the original processor, to protect the initial investment and provide headroom for future AI and metadata workloads.
44.	23	D.2. Item 3: CL-03 – S.no- 3	Memory: 512GB DDR5 Registered DIMM (RDIMM) scalable up to 8TB.	Scalable upto 4TB Justification: Intel comes with 16DIMMs(256*16) configuration & AMD 24DIMMs(256*24).	Modified as: Memory (CL-03 – Management / Control Nodes): Each CL-03 server shall be populated with 512 GB DDR5 RDIMM and shall be scalable up to 6 TB memory per node, in line with the overall memory scalability requirement. The motherboard shall provide sufficient DIMM slots (as per OEM design) to allow cost-effective expansion using industry-standard 128 GB / 256 GB DIMMs without requiring a platform change. Configurations that cannot scale to 6 TB memory per node shall be treated as non-compliant.
45.	23	D.2. Item 3: CL-03 – S.no- 4	Memory Protection: Advanced ECC with multi-bit error protection	Request to amend to · Advanced ECC with multi-bit error protection or equivalent. Justification - Request to amend in order to allow other OEMs to participate as every OEMs will have its own terminology for memory features and functionality.	Modified as: Memory Protection – Advanced ECC (CL-03): Control / management nodes shall support Advanced ECC with multi-bit error protection or equivalent technology. ‘Equivalent’ means the memory subsystem must, at minimum, provide Single Device Data Correction (SDDC) or similar capability so that a failure of a single DRAM device on a DIMM does not cause data

					corruption or crashes of critical metadata services.
46.	23	D.2. Item 3: CL-03 – S.no -4	Memory Protection: Online spare	Request to amend to · Online spare or equivalent. Justification - Request to amend in order to allow other OEMs to participate as every OEMs will have its own terminology for memory features and functionality.	Modified as: Memory Protection – Online spare (CL-03): ‘Online spare’ or equivalent high-availability memory feature shall be supported. ‘Equivalent’ includes OEM mechanisms such as rank sparing, fault-resilient memory modes, or similar technologies that can automatically isolate faulty memory regions while keeping the node online, in order to prevent management-node failures that could impact the entire cluster.
47.	23	D.2. Item 3: CL-03 – S.no -7	Controller: Integrated or add-on PCIe 3.0 based 8G SAS RAID Controller with RAID 0, 1, 5, 6, 10, 50, 60.	Change request: Integrated or add-on PCIe 3.0 based 4G SAS RAID Controller with RAID 0, 1, 5, 6, 10, 50, 60. Justification: The server requirement specifies NVMe drives, for which an 8 GB cache memory on the RAID controller is not required. This requirement is mentioned only for this server and not for other servers and appears to be specific to a single OEM. We kindly request that the same be amended or removed to enable broader participation.	Modified as: Controller (CL-03 – Management / Control Nodes): CL-03 servers shall be equipped with an integrated or add-on PCIe-based SAS/SATA RAID or storage controller supporting at least RAID 0/1/5/6/10/50/60 and a minimum of 8 GB flash-backed write cache (or higher). The controller may operate in RAID or pass-through/HBA mode as required, but must provide protection of write-pending data against power loss, particularly for OS, log, and metadata volumes. Controllers with smaller caches or without flash-backed protection shall not be considered equivalent.
48.	24	D.2. Item 3: CL-03 S.No. 8	Network Cards: • 1x 1GbE 4-Port Adapter • Ethernet 10/25Gb 2-port SFP28 Adapter		Modified as: Network Cards: • Dedicated 1GbE out-of-band management interface for remote server administration.

					Dual-port 200Gbps RDMA-capable (RoCE v2) QSFP NIC for ultra-low-latency, high-throughput data workloads
49.	24	D.2. Item 3: CL-03 – S.no -9	FC HBA: 1x Dual Port 32Gbps FC HBA	Change request: Kindly remove Justification: As there is no SAN requirement in the proposed solution, we kindly request that the SAN-specific specification be removed, as it is not applicable to the intended architecture	Deleted.
50.	24	D.2. Item 3: CL-03 – S.no-10	Interfaces: USB 2.0/3.0 – 4 Nos.	Request to amend to · USB 2.0/3.0 - 3 Nos. Justification - USB ports are generally used for Keyboard or mouse connectivity whereas these servers are for HDFS data nodes which does not require any peripheral connectivity as they would be kept in a DC rack. A min of 3 USB ports are standard and will be sufficient. Not all the OEMs will have 4 USB ports and hence request to amend for other leading OEMs to qualify & participate	Modified as: Interfaces: USB 2.0/3.0 – 3 Nos.
51.	24	D.2. Item 3: CL-03 – S.no-11	Bus Slots Should have · PCI-Express 5.0 x16 slots – 8 Nos. · PCIe 5.0x8 slots – 2 Nos.	Request to amend to · PCI-Express 4.0 / 5.0 8x or x16 slots - 3 Nos. Justification - The requirement is for 1U Rack Server which will supports max 3 PCIe slots. Also, the requirement of NIC / HBA can be easily met with PCIe 4.0 x8 or x16 slots, hence we request to amend the clause for allowing larger OEM participation. Change request: PCI-Express 5.0 slots – 2 Nos. Justification: The specified PCIe slot configuration is specific to a single OEM. To enable broader OEM participation, we kindly request that this requirement be removed, as it	Modified as: Bus Slots (CL-03 – Management / Control Nodes): Each CL-03 server shall be based on PCIe Gen 5.0 x16 architecture and shall provide a minimum of 3 × PCI-Express 5.0 slots (x16 electrical) for high-speed adapters (e.g., dual 100G/200G RDMA NICs and future accelerators). Additional slots may be provided as per OEM design. Platforms limited to PCIe Gen 4.0 or with fewer than 3 usable PCIe 5.0 x16-electrical slots shall be treated as non-compliant, as PCIe 5.0 bandwidth is mandatory to support

				does not impact the functional or operational requirements of the department. PCIE 5.0 Slots minimum 3 or above Justification: Since required cards in the server are two for ethernet connections 10 slots is on higher end.	100G/200G networking and NVMe-oF without I/O bottlenecks.
52.	24	D.2. Item 3: CL-03- S.no-12	Power Supply: 2 x 1600 Watts Platinum Hot Swap Power Supplies	Request to amend to - 2 x 1100 Watts or higher Platinum Hot Swap Power Supplies Justification - Every OEM will have its own power design and accordingly power supply capacity. Since 1600W is OEM specific we request to amend the clause to allow other OEMs to participate. Required power supply for the entire BOM. Justification: Since overall power consumption with single CPU, memroy & HDD will not exceed 1500W.	Modified as: Power Supply (CL-03 – Management / Control Nodes): CL-03 servers shall support hot-plug, redundant Platinum or Titanium power supplies in N+1 configuration. The server shall be configured with two (or more) identical PSUs sized, as per OEM power-sizing guidelines, to support the full load of all CPUs, drives, NVMe devices, and NICs with one PSU failed, without derating or throttling. A fixed wattage (e.g., 1600 W or 1100 W) is not prescribed, but bids must include OEM documentation demonstrating that the proposed PSU configuration safely covers worst-case power draw under N+1 conditions.
53.	24	D.2. Item 3: CL-03- S.no-14	Industry Standard Compliances: ACPI 6.3 Compliant	ACPI 6.1 Compliant in all mentioned servers. Justification: Since 6.3 is not available for AMD OEMs.	Modified as: Industry Standard Compliances: ACPI 6.1 Compliant
54.	24	D.2. Item 3: CL-03- S.no-17	GPU support: System should support NVIDIA's latest computational accelerators and graphics accelerators	Request to remove as this being an 1U server will not support latest high end computational accelerators and graphics accelerators. Need to be removed. Justification: Since GPU are not getting installed.	No Change. As per RFP

55.	26	D.2. Item 4: CL-04- Servers for KMS for Encryption (Bare metal servers)- S.no-2	CPU: Latest series Processor/s with 8 Physical Cores with Hyperthreading are to be supplied with provision to expand up to 64 Physical Cores.	Request to clarify if 8 physical cores to be provided from server as overall core or from single CPU? Need to mention Single socket Justification: Efficient Utilization of Compute resources & IO's and better cost efficiency. Can be scaled to 64 physical cores by just replacing the CPU. Also reducing CPU to CPU latency in dual socket scenario.	Modified as: CPU (CL-04 – KMS / Encryption Servers): Each CL-04 KMS server shall be equipped with latest-generation scalable processors providing at least 64 physical cores per CPU with Hyper-Threading, on a dual-socket-capable platform. At minimum, one 64-core CPU shall be supplied per node initially; the platform shall support expansion to higher core counts (e.g., adding a second CPU) without replacing the original processor, to ensure sufficient headroom for future encryption and key-management workloads.
56.	26	D.2. Item 4: CL-04- Servers for KMS for Encryption (Bare metal servers) – S.no-3	Memory: 128GB DDR5 Registered DIMM (RDIMM) scalable up to 8TB.	Change request: 512GB DDR5 Registered DIMM (RDIMM) scalable up to 6TB. Justification: Currently, only Intel (single OEM) platforms qualify for the 8 TB memory requirement. We request to kindly amend this specification to 6 TB memory, which will allow qualification of both AMD and Intel CPU-based servers, ensuring fair competition and wider OEM participation. From a practical deployment perspective, server memory utilization will not exceed 6 TB under any circumstances, even in peak workload scenarios. Therefore, specifying 6 TB fully meets the operational and scalability requirements while avoiding unnecessary vendor lock-in. Scalable upto 4TB Justification:	Modified as: Memory (CL-04 – KMS / Encryption Servers): Each CL-04 server shall be populated with 128 GB DDR5 RDIMM and shall be scalable up to 6 TB memory per node, in line with the overall memory scalability requirement. Platforms from any OEM (Intel or AMD) that support at least 6 TB memory scalability per node on PCIe Gen 5.0 architecture will be considered compliant; configurations limited below 6 TB scalability shall be treated as non-responsive. Populated 128 GB is deemed sufficient for current key-management workloads; higher initial population is optional but not mandatory.

				Intel comes with 16DIMMs(256*16) configuration & AMD 24DIMMs(256*24).	
57.	26	D.2. Item 4: CL-04- Servers for KMS for Encryption (Bare metal servers)- S.no-4	Memory Protection: Advanced ECC with multi-bit error protection	Request to amend to · Advanced ECC with multi-bit error protection or equivalent. Justification - Request to amend in order to allow other OEMs to participate as every OEMs will have its own terminology for memory features and functionality.	Modified as: Memory Protection (CL-04 – KMS / Encryption Servers): KMS servers shall support Advanced ECC with multi-bit error protection or equivalent technology. ‘Equivalent’ requires, at minimum, Single Device Data Correction (SDDC) or a similar mechanism so that failure of a single DRAM device on a DIMM does not cause key corruption or service outages in the encryption subsystem.
58.	26	D.2. Item 4: CL-04- Servers for KMS for Encryption (Bare metal servers)- S.no-4	Memory Protection: Online spare	Request to amend to · Online spare or equivalent. Justification - Request to amend in order to allow other OEMs to participate as every OEMs will have its own terminology for memory features and functionality.	Modified as: Memory Protection – ‘Online spare’ or equivalent high-availability memory feature shall be supported. ‘Equivalent’ includes OEM mechanisms such as rank sparing, fault-resilient memory modes, or similar technologies that can automatically isolate faulty memory regions while keeping the node online, in order to prevent management-node failures that could impact the entire cluster.
59.	26	D.2. Item 4: CL-04 S.No-8	Network Cards: • 1x 1GbE 4-Port Adapter • Ethernet 10Gb 2-port SFP28 Adapter		Modified as: Network Cards: • Dedicated 1GbE out-of-band management interface for remote server administration. • Dual-port 200Gbps RDMA-capable (RoCE v2) QSFP NIC for ultra-low-latency, high-throughput data workloads
60.	26	D.2.	FC HBA: 1x Dual Port 32Gbps FC HBA	Change request: Kindly remove	Deleted

		Item 4: CL-04-Servers for KMS for Encryption (Bare metal servers)- S.no-9		Justification: As there is no SAN requirement in the proposed solution, we kindly request that the SAN-specific specification be removed, as it is not applicable to the intended architecture.	
61.	26	D.2. Item 4: CL-04-Servers for KMS for Encryption (Bare metal servers)- S.no-10	Interfaces: USB 2.0/3.0 – 4 Nos.	Request to amend to · USB 2.0/3.0 - 3 Nos. Justification - USB ports are generally used for Keyboard or mouse connectivity whereas these servers are for HDFS data nodes which does not require any peripheral connectivity as they would be kept in a DC rack. A min of 3 USB ports are standard and will be sufficient. Not all the OEMs will have 4 USB ports and hence request to amend for other leading OEMs to qualify & participate.	Modified as: Interfaces: USB 2.0/3.0 – 3 Nos.
62.	26	D.2. Item 4: CL-04-Servers for KMS for Encryption (Bare metal servers)- S.no-11	Bus Slots: Should have · PCI-Express 5.0 x16 slots – 8 Nos. · PCIe 5.0x8 slots – 2 Nos.	Request to amend to PCI-Express 4.0 / 5.0 8x or x16 slots - 3 Nos. Justification - The requirement is for 1U Rack Server which will supports max 3 PCIe slots. Also, the requirement of NIC / HBA can be easily met with PCIe 4.0 x8 or x16 slots, hence we request to amend the clause for allowing larger OEM participation. Change request: PCI-Express 5.0 slots – 2 Nos. Justification: The specified PCIe slot configuration is specific to a single OEM. To enable broader OEM participation, we kindly request that this requirement be removed, as it	Modified as: Bus Slots (CL-04 – KMS / Encryption Servers): Each CL-04 server shall be based on PCIe Gen 5.0 x16 architecture and shall provide a minimum of 3 × PCI-Express 5.0 slots (x16 electrical). These slots are required for dual 100G/200G RDMA NICs and future crypto/accelerator cards. Platforms limited to PCIe Gen 4.0 or with fewer than 3 usable PCIe 5.0 x16-electrical slots shall be treated as non-compliant, as PCIe 5.0 bandwidth is mandatory to avoid bottlenecks for 100G/200G networking and hardware encryption offload.

				does not impact the functional or operational requirements of the department. PCIE 5.0 Slots minimum 3 or above. Justification: Since required cards in the server are two for ethernet connections 10 slots is on higher end.	
63.	26	D.2. Item 4: CL-04- Servers for KMS for Encryption (Bare metal servers)- S.no-12	Power Supply: 2 x 1600 Watts Platinum Hot Swap Power Supplies	Request to amend to - 2 x 1100 Watts or higher Platinum Hot Swap Power Supplies. Justification - Every OEM will have its own power design and accordingly power supply capacity. Since 1600W is OEM specific we request to amend the clause to allow other OEMs to participate. Required power supply for the entire BOM. Justification: Since overall power consumption with single CPU, memroy & HDD will not exceed 1500W.	Modified as: Power Supply (CL-04 – KMS / Encryption Servers): CL-04 servers shall support hot-plug, redundant Platinum or Titanium power supplies configured in N+1 mode. The server shall be equipped with two (or more) identical PSUs sized, as per OEM power-sizing guidelines, to support the full load of all CPUs, NICs, and any future crypto accelerators under single-PSU failure, without derating or performance throttling. A fixed wattage (e.g., 1600 W, 1400 W, or 1100 W) is not prescribed, but bidders must submit OEM documentation demonstrating that the proposed PSU configuration safely covers worst-case power draw under N+1 conditions.
64.	26	D.2. Item 4: CL-04- Servers for KMS for Encryption (Bare metal servers)- S.no-14	Industry Standard Compliances: ACPI 6.3 Compliant	ACPI 6.1 Compliant in all mentioned servers. Justification: Since 6.3 is not available for AMD OEMs.	Modified as: Industry Standard Compliances: ACPI 6.1 Compliant

65.	27	D.2. Item 4: CL-04- Servers for KMS for Encryption (Bare metal servers)- S.no-17	GPU support: System should support NVIDIA's latest computational accelerators and graphics accelerators.	Request to remove as this being an IU server will not support latest high end computational accelerators and graphics accelerators. Need to be removed Justification: Since GPU are not getting installed.	No Change. As per RFP.
66.	29	D.2. Item 5: CL 05 - Servers for Master Nodes (Bare metal servers)- S.no-2	CPU: Latest series Processor/s with 32 Physical Cores with Hyperthreading are to be supplied with provision to expand up to 64 Physical Cores.	Request to clarify if 32 physical cores to be provided from server as overall core or from single CPU? Need to mention Single socket Justification: Efficient Utilization of Compute resources & IO's and better cost efficiency. Can be scaled to 64 physical cores by just replacing the CPU. Also reducing CPU to CPU latency in dual socket scenario.	Modified as: CPU (CL-05 – Master Nodes): Each CL-05 Master Node shall use latest-generation scalable processors on a dual-socket-capable motherboard. At minimum, the node shall be supplied with one CPU providing at least 64 physical cores with Hyper-Threading; the platform shall support expansion to higher total core counts by adding a second CPU (e.g., 2 × 64-core), without replacing the original processor. This ensures sufficient memory bandwidth and PCIe Gen 5.0 lanes for control-plane, metadata, and AI orchestration workloads.
67.	29	D.2. Item 5: CL 05 - Servers for Master Nodes (Bare metal	Memory: 512GB DDR5 Registered DIMM (RDIMM) scalable up to 8TB.	Change request: 512GB DDR5 Registered DIMM (RDIMM) scalable up to 6TB. Justification: Currently, only Intel (single OEM) platforms qualify for the 8 TB memory requirement. We request to kindly amend this specification to 6 TB memory, which will allow qualification of both AMD and Intel	Modified as: Memory: 512GB DDR5 Registered DIMM (RDIMM) scalable up to 6 TB.

		servers) -S.no-3		CPU-based servers, ensuring fair competition and wider OEM participation. From a practical deployment perspective, server memory utilization will not exceed 6 TB under any circumstances, even in peak workload scenarios. Therefore, specifying 6 TB fully meets the operational and scalability requirements while avoiding unnecessary vendor lock-in. Scalable upto 4TB Justification: Intel comes with 16DIMMs(256*16) configuration & AMD 24DIMMs(256*24).	
68.	29	D.2. Item 5: CL 05 - Servers for Master Nodes (Bare metal servers) – S.no-4	Memory Protection: Advanced ECC with multi-bit error protection	Request to amend to · Advanced ECC with multi-bit error protection or equivalent. Justification - Request to amend in order to allow other OEMs to participate as every OEMs will have its own terminology for memory features and functionality.	Modified as: Memory Protection: Advanced ECC : Control / management nodes shall support Advanced ECC with multi-bit error protection or equivalent technology. ‘Equivalent’ means the memory subsystem must, at minimum, provide Single Device Data Correction (SDDC) or similar capability so that a failure of a single DRAM device on a DIMM does not cause data corruption or crashes of critical metadata services.
69.	29	D.2. Item 5: CL 05 - Servers for Master Nodes (Bare metal	Memory Protection: Online spare	Request to amend to · Online spare or equivalent. Justification - Request to amend in order to allow other OEMs to participate as every OEMs will have its own terminology for memory features and functionality.	Modified as: Memory Protection: Online spare : ‘Online spare’ or equivalent high-availability memory feature shall be supported. ‘Equivalent’ includes OEM mechanisms such as rank sparing, fault-resilient memory modes, or similar technologies that can automatically

		servers) -S.no-4			isolate faulty memory regions while keeping the node online, in order to prevent management-node failures that could impact the entire cluster.
70.	29	D.2. Item 5: CL 05 - S.No. 8	Network Cards: 1x 1GbE 4-Port Adapter - Ethernet 10/25Gb 2-port SFP28 Adapter		Modified as: Network Cards: - Dedicated 1GbE out-of-band management interface for remote server administration. - Dual-port 200Gbps RDMA-capable (RoCE v2) QSFP NIC for ultra-low-latency, high-throughput data workloads
71.	29	D.2. Item 5: CL 05 - -S.No-9	FC HBA: 1x Dual Port 32Gbps FC HBA	Change request: Kindly remove Justification: As there is no SAN requirement in the proposed solution, we kindly request that the SAN-specific specification be removed, as it is not applicable to the intended architecture	Deleted
72.	29	D.2. Item 5: CL 05 - Servers for Master Nodes (Bare metal servers) -S.no-10	Interfaces: USB 2.0/3.0 – 4 Nos.	Request to amend to · USB 2.0/3.0 - 3 Nos. Justification - USB ports are generally used for Keyboard or mouse connectivity whereas these servers are for HDFS data nodes which does not require any peripheral connectivity as they would be kept in a DC rack. A min of 3 USB ports are standard and will be sufficient. Not all the OEMs will have 4 USB ports and hence request to amend for other leading OEMs to qualify & participate.	Modified as: Interfaces: USB 2.0/3.0 – 3 Nos.
73.	29	D.2. Item 5: CL 05 - Servers for Master	Bus Slots: Should have · PCI-Express 5.0 x16 slots – 8 Nos. · PCIe 5.0x8 slots – 2 Nos.	Request to amend to PCI-Express 4.0 / 5.0 8x or x16 slots - 3 Nos. Justification - The requirement is for 1U Rack Server which will supports max 3 PCIe slots. Also, the requirement of NIC / HBA can be easily met with PCIe 4.0 x8 or x16 slots,	Modified as: Bus Slots (CL-05 – Master Nodes): Each CL-05 server shall be based on PCIe Gen 5.0 x16 architecture and shall provide a minimum of 3 × PCI-Express 5.0 slots

		Nodes (Bare metal servers)- S.no-11		hence we request to amend the clause for allowing larger OEM participation. Change request: PCI-Express 5.0 slots – 2 Nos. Justification: The specified PCIe slot configuration is specific to a single OEM. To enable broader OEM participation, we kindly request that this requirement be removed, as it does not impact the functional or operational requirements of the department. PCIe 5.0 Slots minimum 3 or above Justification: Since required cards in the server are two for ethernet connections 10 slots is on higher end.	(x16 electrical). These are required for dual 100G/200G RDMA NICs and future accelerators. Platforms limited to PCIe Gen 4.0 or with fewer than 3 usable PCIe 5.0 x16-electrical slots shall be treated as non-compliant, as PCIe 5.0 bandwidth is essential to support 100G/200G networking and NVMe-oF traffic without bottlenecks.
74.	29	D.2. Item 5: CL 05 - Servers for Master Nodes (Bare metal servers)- S.no-12	Power Supply: 2 x 1600 Watts Platinum Hot Swap Power Supplies.	Request to amend to - 2 x 1100 Watts or higher Platinum Hot Swap Power Supplies Justification - Every OEM will have its own power design and accordingly power supply capacity. Since 1600W is OEM specific we request to amend the clause to allow other OEMs to participate. Required power supply for the entire BOM. Justification: Since overall power consumption with single CPU, memory & HDD will not exceed 1500W.	Modified as: Power Supply (CL-05 – Master Nodes): CL-05 servers shall support hot-plug, redundant Platinum or Titanium power supplies in N+1 configuration. The node shall be equipped with two (or more) identical PSUs sized, as per OEM power-sizing guidelines, to sustain the full load of all CPUs, memory, drives, and 100G/200G NICs under single-PSU failure, without derating or throttling. A specific wattage minimum (e.g., 1600 W vs 1100 W) is not mandated, but bidders must provide OEM power-sizing evidence that the proposed PSU configuration safely covers worst-case power draw for the fully expanded configuration.
75.	29	D.2.	Industry Standard Compliances: ACPI 6.3 Compliant.	ACPI 6.1 Compliant in all mentioned servers. Justification:	Modified as:

		Item 5: CL 05 - S.no-14		Since 6.3 is not available for AMD OEMs.	Industry Standard Compliances: ACPI 6.1 Compliant.
76.	30	D.2. Item 5: CL 05 - S.no-17	GPU support: System should support NVIDIA's latest computational accelerators and graphics accelerators	Request to remove as this being an 1U server will not support latest high end computational accelerators and graphics accelerators.	No Change. As per RFP.
77.	32	D.2. Item 6: CL 06 - Servers for Worker Nodes (Bare metal servers)- S.no-2	CPU: Latest series Processor/s with 48 Physical Cores with Hyperthreading are to be supplied with provision to expand up to 64 Physical Cores.	Request to clarify if 48 physical cores to be provided from server as overall core or from single CPU? Need to mention Single socket Justification: Efficient Utilization of Compute resources & IO's and better cost efficiency. Can be scaled to 64 physical cores by just replacing the CPU. Also reducing CPU to CPU latency in dual socket scenario.	Modified as: CPU (CL-06 – Worker Nodes): Each CL-06 Worker Node shall use latest-generation scalable processors on a dual-socket-capable motherboard. At minimum, the node shall be supplied with one CPU providing at least 64 physical cores with Hyper-Threading; the platform shall support expansion to higher total core counts by adding a second CPU (e.g., 2 × 64-core) without replacing the original processor. This is required to ensure sufficient compute, memory bandwidth, and PCIe Gen 5.0 lanes for distributed processing and high-throughput NVMe-oF workloads.
78.	32	D.2. Item 6: CL 06 - Servers for Worker Nodes (Bare metal servers)- S.no-3	Memory: 512GB DDR5 Registered DIMM (RDIMM) scalable up to 8TB.	Change request: 512GB DDR5 Registered DIMM (RDIMM) scalable up to 6TB. Justification: Currently, only Intel (single OEM) platforms qualify for the 8 TB memory requirement. We request to kindly amend this specification to 6 TB memory, which will allow qualification of both AMD and Intel CPU-based servers, ensuring fair competition and wider OEM participation.	Modified as: Memory: 512GB DDR5 Registered DIMM (RDIMM) scalable up to 6 TB.

				From a practical deployment perspective, server memory utilization will not exceed 6 TB under any circumstances, even in peak workload scenarios. Therefore, specifying 6 TB fully meets the operational and scalability requirements while avoiding unnecessary vendor lock-in. Scalable upto 4TB Justification: Intel comes with 16DIMMs(256*16) configuration & AMD 24DIMMs(256*24).	
79.	32	D.2. Item 6: CL 06 - Servers for Worker Nodes (Bare metal servers)- S.no-4	Memory Protection: Advanced ECC with multi-bit error protection.	Request to amend to · Advanced ECC with multi-bit error protection or equivalent. Justification - Request to amend in order to allow other OEMs to participate as every OEMs will have its own terminology for memory features and functionality.	Modified as: Memory Protection (CL-06 – Worker Nodes): Worker-node servers shall support Advanced ECC with multi-bit error protection or equivalent technology. ‘Equivalent’ requires, at minimum, Single Device Data Correction (SDDC) or a comparable mechanism so that failure of a single DRAM device on a DIMM does not cause data corruption or crashes during long-running distributed compute and AI workloads.
80.	32	D.2. Item 6: CL 06 - Servers for Worker Nodes (Bare metal servers)- S.no-4	Memory Protection: Online spare	Request to amend to · Online spare or equivalent. Justification - Request to amend in order to allow other OEMs to participate as every OEMs will have its own terminology for memory features and functionality.	Modified as: Memory Protection – Online spare: ‘Online spare’ or equivalent high-availability memory feature shall be supported. ‘Equivalent’ includes OEM mechanisms such as rank sparing, fault-resilient memory modes, or similar technologies that can automatically isolate faulty memory regions while keeping the node online, in order to prevent

					management-node failures that could impact the entire cluster.
81.	32	D.2. Item 6: CL 06 - S.No-8	Network Cards: 1x 1GbE 4-Port Adapter - Ethernet 10/25Gb 2-port SFP28 Adapter		Modified as: Network Cards: - Dedicated 1GbE out-of-band management interface for remote server administration. - Dual-port 200Gbps RDMA-capable (RoCE v2) QSFP NIC for ultra-low-latency, high-throughput data workloads
82.	32	D.2. Item 6: S.no-9	FC HBA: 1x Dual Port 32Gbps FC HBA	Change request: Kindly remove Justification: As there is no SAN requirement in the proposed solution, we kindly request that the SAN-specific specification be removed, as it is not applicable to the intended architecture.	Deleted
83.	32	D.2. Item 6: CL 06 - Servers for Worker Nodes (Bare metal servers)- S.no-10	Interfaces: USB 2.0/3.0 – 4 Nos.	Request to amend to · USB 2.0/3.0 - 3 Nos. Justification - USB ports are generally used for Keyboard or mouse connectivity whereas these servers are for HDFS data nodes which does not require any peripheral connectivity as they would be kept in a DC rack. A min of 3 USB ports are standard and will be sufficient. Not all the OEMs will have 4 USB ports and hence request to amend for other leading OEMs to qualify & participate.	Modified as: Interfaces: USB 2.0/3.0 – 3 Nos.
84.	32	D.2. Item 6: CL 06 - Servers for Worker Nodes (Bare	Bus Slots: Should have · PCI-Express 5.0 x16 slots – 8 Nos. · PCIe 5.0x8 slots – 2 Nos.	Request to amend to PCI-Express 4.0 / 5.0 8x or x16 slots - 3 Nos. Justification - The requirement is for 1U Rack Server which will supports max 3 PCIe slots. Also the requirement of NIC / HBA can be easily met with PCIe 4.0 x8 or x16 slots, hence we request to amend the clause for allowing larger OEM participation.	Modified as: Bus Slots (CL-06 – Worker Nodes): Each CL-06 Worker Node shall be based on PCIe Gen 5.0 x16 architecture and shall provide a minimum of 3 × PCI-Express 5.0 slots (x16 electrical). These slots are required for dual

		metal servers)- S.no-11		Change request: PCI-Express 5.0 slots – 2 Nos. Justification: The specified PCIe slot configuration is specific to a single OEM. To enable broader OEM participation, we kindly request that this requirement be removed, as it does not impact the functional or operational requirements of the department. PCIe 5.0 Slots minimum 3 or above Justification: Since required cards in the server are two for ethernet connections 10 slots is on higher end.	100G/200G RDMA NICs and any NVMe / accelerator cards used for scratch and temporary storage. Platforms limited to PCIe Gen 4.0 or with fewer than 3 usable PCIe 5.0 x16-electrical slots shall be treated as non-compliant, as PCIe 5.0 bandwidth is necessary to support 100G/200G traffic and NVMe-oF workloads without I/O bottlenecks.
85.	32	D.2. Item 6: CL 06 - Servers for Worker Nodes (Bare metal servers)- S.no-12	Power Supply: 2 x 1600 Watts Platinum Hot Swap Power Supplies	Request to amend to - 2 x 1100 Watts or higher Platinum Hot Swap Power Supplies Justification - Every OEM will have its own power design and accordingly power supply capacity. Since 1600W is OEM specific we request to amend the clause to allow other OEMs to participate. Required power supply for the entire BOM. Justification: Since overall power consumption with single CPU, memroy & HDD will not exceed 1500W.	Modified as: Power Supply (CL-06 – Worker Nodes): CL-06 Worker Nodes shall support hot-plug, redundant Platinum or Titanium power supplies configured in N+1 mode. Each node shall be equipped with two (or more) identical PSUs sized, as per OEM power-sizing guidelines, to sustain the full load of all CPUs, memory, NVMe devices, and 100G/200G NICs under single-PSU failure, without derating or performance throttling. No fixed wattage value is prescribed; however, bidders must provide OEM power-sizing documentation showing that the proposed PSU configuration safely covers the worst-case, fully expanded configuration under N+1 conditions.
86.	32	D.2. Item 6: CL 06 - S.no-14	Industry Standard Compliances: ACPI 6.3 Compliant	ACPI 6.1 Compliant in all mentioned servers. Justification: Since 6.3 is not available for AMD OEMs.	Modified as: Industry Standard Compliances: ACPI 6.1 Compliant

87.	33	D.2. Item 6: CL 06 - S.no-17	GPU support: System should support NVIDIA's latest computational accelerators and graphics accelerators.	Request to remove as this being an 1U server will not support latest high end computational accelerators and graphics accelerators.	No Change. As per RFP.
88.	34 & 35	D.2. Item 7: CL 07 – Servers for GPU Nodes – S.no-4	Memory: 32 DIMM slots. Servers should be supplied with 512 GB memory and scalable upto 4 TB using DDR5 Registered DIMM (RDIMM) operating at 6400 MT/s Should be capable of identifying and reporting whether genuine OEM memory is installed for system reliability.	Scalable upto 3TB Justification: Since AMD can provide 24DIMMS upto 3TB, as this is a GPU Node server most of the performance will be done at GPU level.	Modified as: Memory: 24 DIMM slots. Servers should be supplied with 512 GB memory and scalable upto 3 TB using DDR5 Registered DIMM (RDIMM) operating at 6400 MT/s Should be capable of identifying and reporting whether genuine OEM memory is installed for system reliability.
89.	35	D.2. Item 7: CL 07 – Servers for GPU Nodes – S.no-5	Bus Slots: Server should support upto six PCI-Express 5.0 x16 slots.	Change request: Server should support up to five PCI-Express 5.0 x16 slots. Justification: The specified PCIe slot configuration is specific to a single OEM. To enable broader OEM participation, we kindly request that this requirement be removed, as it does not impact the functional or operational requirements of the department.	Modified as: Bus Slots: Server should support upto Five PCI-Express 5.0 x16 slots.
90.	35	D.2. Item 7: CL 07 – Servers for GPU Nodes – S.no-8	Controller: Embedded / PCIe based x16 RAID controller with 8GB Flash backed write cache, supporting RAID 0, 1, 5, 6, 10, 50, 60. Must support mix-and-match SAS, SATA, and NVMe drives to the same controller. Controller must support 6G SATA, 12G SAS, 16G NVMe. Above mentioned controller must support following: 1. Expand Logical Drive 2. Instant Secure Erase	Request to remove this clause as this is a specialized GenAI Optimized Server which requires extremely high throughput and low latency between application and storage disks. RAID controllers impacts the throughput and adds latency impacting the performance and also not supported by leading OEMs. Change request: Kindly remove: Embedded / PCIe based x16 RAID controller with 8GB Flash backed write cache, supporting RAID 0, 1, 5, 6, 10, 50, 60. 2. Instant Secure Erase	Deleted.

			3.Migrate RAID/Stripe Size 4. Modifying Cache Write Policy 5.Move Logical Drive 6. Re-enable Failed Logical Drive 7. Load balancing 8. Configurable stripe size up to 1 MB 9. Check consistency for background data integrity 10. Patrol read for media scanning and repairing 11. Rebuild	Justification: High-density GPU servers are designed with NVMe-grade storage, and therefore conventional RAID functionality is not relevant. The ‘Instant Secure Erase’ requirement is proprietary and specific to a single OEM. We kindly request that this requirement be removed to enable broader participation.	
91.	35	D.2. Item 7: CL 07 – Servers for GPU Nodes – S.no-9	Networking features: Server should be supplied with below networking cards: 1. 1Gb 4-port network adaptors 2. 10Gb 2-port Ethernet adaptor	Change request: Server should be supplied with below networking cards: 1. 10Gb 2-port Ethernet adaptor 2. 2x 200G port Justification: High-density GPU servers are designed with high-density ports; therefore, we kindly request that the requirement be amended to enable broader participation.	Modified as: Networking (CL-07 – GPU Nodes): Each CL-07 GPU node shall be supplied with: (a) at least one 10G (or 1/10G) Ethernet interface for management/auxiliary traffic, and (b) a minimum of 2 × 100G or 2 × 200G RDMA-capable Ethernet ports (RoCE v2) or equivalent, using NVIDIA ConnectX-6/ConnectX-7 (or equivalent) adapters, in line with the high-speed networking requirements. Low-speed links such as 1G/10G shall not be used as the primary data path for AI workloads; the 100G/200G ports shall be used for data-plane connectivity to the NVMe-oF fabric and the 15-PB data lake.
92.	35	D.2. Item 7: CL 07 – Servers for GPU Nodes – S.no-11	Power Supply: • Should support hot plug redundant (N+1) low halogen power supplies with minimum 96% efficiency. • Server must have dedicated power supply for System Domain and GPU Domain with redundancy.	Request to Amend as - Should support hot plug redundant (N+N) power supplies with minimum 96% efficiency. 2.Server must have dedicated power supply for System Domain and GPU Domain with redundancy. Justification - Power Supply will provide the power to complete server including GPUs.	Modified as: Power Supply (CL-07 – GPU Nodes): CL-07 GPU servers shall support hot-plug, high-efficiency (≥96%) Platinum or Titanium power supplies in a redundant configuration sized for full CPU + GPU load. An N+1 or N+N redundancy scheme is acceptable, provided that the

				Hence request to amend the clause for leading OEMs to comply and participate. Change request: Should support hot plug redundant (N+1) low halogen power supplies with minimum 96% efficiency Justification: The requirement for a dedicated power supply for the GPU appears to be specific to a single OEM. We kindly request that this requirement be removed to enable broader participation.	platform can sustain worst-case power draw of all populated GPUs and system components under a single-PSU or equivalent failure without derating or performance throttling. OEM-specific 'domain' terminology is not mandated, but the PSU design must be explicitly certified by the OEM as suitable for the proposed GPU configuration (including future higher-TDP GPUs) under N+1/N+N conditions.
93.	36	D.2. Item 7: CL 07 – Servers for GPU Nodes – S.no-13	Industry Standard Compliance: • ACPI 6.5 Compliant • Advanced Encryption Standard (AES) • Triple Data Encryption Standard (3DES) • Active Directory v1.0 • ASHRAE A3/A4 • DMTF Systems Management Architecture for Server Hardware • Command Line Protocol (SMASH CLP) • Energy Star • IPMI 2.0 • Microsoft® Logo certifications • PCIe 5.0 Compliant • PXE Support • Redfish API • Secure Digital 4.0 • SMBIOS 3.7 • SNMP v3 • TLS 1.3 • TPM 2.0 Support • UEFI Class 3 USB 3.0 Compliant	Request to amend as - • ACPI 6.4 Compliant • Advanced Encryption Standard (AES) • Triple Data Encryption Standard (3DES) • Active Directory v1.0 • ASHRAE A3/A4 • DMTF Systems Management Architecture for Server Hardware • Command Line Protocol (SMASH CLP) • Energy Star • IPMI 2.0 • Microsoft® Logo certifications • PCIe 5.0 Compliant • PXE Support • Redfish API • Secure Digital 4.0 • SMBIOS 3.3 • SNMP v3 • TLS 1.3 • TPM 2.0 Support • UEFI Class 2.7 • USB 3.0 Compliant Justification - The specifications are specific to one OEM based on its spec sheet. Also, 3DES are obsolete, hence we request to amend	Modified as: Industry Standard Compliance: • ACPI 6.1 Compliant • Advanced Encryption Standard (AES 256) • Active Directory v1.0 • ASHRAE A3/A4 • DMTF Systems Management Architecture for Server Hardware • Command Line Protocol (SMASH CLP) • Energy Star • IPMI 2.0 • Microsoft® Logo certifications • PCIe 5.0 Compliant • PXE Support • Redfish API • SMBIOS 3.3 • SNMP v3 • TLS 1.3 • TPM 2.0 Support • UEFI Class 2.7 • USB 3.0 Compliant

				to allow other leading OEMs to comply and participate.	
94.	36 &3 7	D.2. Item 7: CL 07 – Servers for GPU Nodes – S.no-15	System Security: Enhanced Server Data Security: Encryption and key management, UEFI-managed encryption, and self-encrypting drives (SED) for enhanced data-at-rest protection. Sanitize Data with One-Button Erase: Must be NIST SP 800-88 compliant. Must be FIPS 140-3, NIST SP 800-53, NIST SP 800-171, and NIST SP 800-88 Compliant. Must provides an intuitive cloud operating experience, ensuring streamlined and highly secure operations from the edge to the cloud. Physical Security Options: System maintenance switch, USB security, bezel lock, and chassis intrusion detection switch. Servers must be built using trusted supply chain to undergo rigorous inspections and checkpoints to detect and mitigate malicious microcode and counterfeit parts throughout the server build and lifecycle	Request to amend as - Enhanced Server Data Security: Encryption and key management, UEFI-managed encryption /cryptographically signed firmware, and self-encrypting drives (SED) for enhanced data-at-rest protection. Sanitize Data with Secure Erase: Must be NIST compliant. Removed One-Button Erase Must be FIPS 140-3 compliant NIST SP 800-53, NIST SP 800-171, and NIST SP 800-88 Compliant. Must provides an intuitive cloud operating experience, ensuring streamlined and highly secure operations from the edge to the cloud. Physical Security Options: System maintenance switch, USB security, bezel lock, and chassis intrusion detection switch. Servers must be built using trusted supply chain to undergo rigorous inspections and checkpoints to detect and mitigate malicious microcode and counterfeit parts throughout the server build and lifecycle Justification - The specifications are specific to one OEM based on its spec sheet. Hence, we request to amend to allow other leading OEMs to comply and participate.	Modified as: System Security (CL-07 – GPU Nodes): Enhanced server data security shall include encryption and key management, UEFI-managed encryption and/or cryptographically signed firmware, and support for self-encrypting drives (SED) for data-at-rest protection. The platform shall provide a single-action secure erase capability for local storage that is compliant with NIST SP 800-88 Rev.1 (Clear/Purge or equivalent crypto-erase). All cryptographic functions used for protection of data-at-rest or management-plane security shall use FIPS 140-3 validated cryptographic modules; security controls and logging shall support implementation of NIST SP 800-53, NIST SP 800-171, and NIST SP 800-88 requirements or their equivalent national/industry standards, as applicable. The solution shall provide an intuitive, cloud-like operating experience with streamlined, secure operations from edge to cloud, and shall support physical security options including system maintenance switch, USB port security, bezel lock, and chassis intrusion detection. Servers must be built using a trusted supply chain, undergoing documented inspections and checkpoints to detect and mitigate malicious microcode and counterfeit parts throughout the server build and lifecycle; this trusted

					supply-chain requirement is mandatory for RTGS and is not waivable.
95.	37	D.2. Item 7: CL 07 – Servers for GPU Nodes – S.no-15	System Security: Server must be offered with Trusted Platform Module (TPM) 2.0. UEFI Secure Boot and Secure Start support	Request to amend as - UEFI Secure Boot and Secure Start support Justification - The specifications are specific to one OEM based on its spec sheet. Hence, we request to amend to allow other leading OEMs to comply and participate.	Modified as: System Security: Server must be offered with Trusted Platform Module (TPM) 2.0. UEFI Secure Boot or equivalent.
96.	37	D.2. Item 7: CL 07 – Servers for GPU Nodes – S.no- 17(2 nd & 3 rd point)	Firmware security: 1. Physically Tamper-Resistant independent security system outside remote management ASIC, designed to meet FIPS 140-3 Level 3 hardware protection requirements. 2. Firmware verification and recovery— Run scheduled firmware verification scans and configure recovery actions to implement when an issue is detected.	Request to amend as - 2. Physically Tamper-Resistant independent security system outside remote management ASIC designed to meet FIPS 140-3 hardware protection requirements. 3. Firmware verification and recovery - Run scheduled /Auto firmware verification scans and configure recovery actions to implement when an issue is detected. Justification - The specifications are specific to one OEM based on its spec sheet. Hence, we request to amend to allow other leading OEMs to comply and participate.	Modified as: Firmware Security (CL-07 – GPU Nodes): Each GPU node shall implement a physically tamper-resistant, independent hardware security component, separate from the remote management ASIC, designed to meet FIPS 140-3 Level 3 physical security requirements for tamper resistance and protection of cryptographic keys. The platform shall support firmware verification and recovery in line with NIST platform-firmware resiliency guidance (protect, detect, recover), with capabilities to run scheduled or automatic firmware integrity scans and to configure recovery actions (such as rollback to a known-good image) when an issue is detected. Equivalent implementations achieving the same FIPS 140-3 Level 3 protections and automated verification/recovery behavior are acceptable, but the FIPS Level 3 physical-security objective for the GPU nodes is mandatory for this project.
97.	39	D.2. Item 7: CL 07 –	Cloud Enabled Monitoring and Management: Server Lifecycle Management:	Request to amend as - Server Lifecycle Management	Modified as: Enabled Monitoring and Management:

		Servers for GPU Nodes – S.no-20	1. Seamless firmware upgrades with Intelligent delta-only based updates. 3. Serial and parallel firmware updates with 'Stop on failure' option	1. Seamless firmware upgrades with Intelligent delta-only based updates / Efficient policy driven firmware management with automation 3. Serial and parallel firmware updates with 'Stop on failure' option / Multi System update Justification - The specifications are specific to one OEM based on its spec sheet. Hence, we request to amend to allow other leading OEMs to comply and participate.	Server Lifecycle Management: 1. Seamless firmware upgrades with Intelligent delta-only based updates/ Efficient policy driven firmware management with automation 3. Serial and parallel firmware updates with 'Stop on failure' option / Multi System update or equivalent
98.	39 & 40	D.2. Item 7: CL 07 – Servers for GPU Nodes – S.no-20	Cloud Enabled Monitoring and Management: Supportability: 1. Automatic creation of support cases for critical serviceable events. 2. Notification of critical hardware alerts via in-app and email notifications 3. On-demand firmware compliance reporting 4. Remote console using a web browser with HTML5 support	Change request: Kindly remove. Justification: All the specified features are proprietary and specific to a single OEM. We kindly request that these be removed to enable broader participation.	As per RFP.
99.	40	D.2. Item 7: CL 07 – Servers for GPU Nodes - S.no-20	Cloud Enabled Monitoring and Management: Identity & Access Management 1. Unified Identity & Access Management 2. Role-based Access Control (RBAC) 3. Single Sign-on (SSO) 4. Resource-level restriction policies (RRP)	Request to amend as - Identity & Access Management 1. Unified Identity & Access Management 2. Role-based Access Control (RBAC) 3. Single Sign-on (SSO) 4. Resource-level restriction policies (RRP) Justification - The specifications are specific to one OEM based on its spec sheet. Hence, we request to amend to allow other leading OEMs to comply and participate. Change request: Kindly remove.	Modified as: Cloud Enabled Monitoring and Management: Identity & Access Management 1. Unified Identity & Access Management 2. Role-based Access Control (RBAC) 3. Single Sign-on (SSO)

				Justification: All the specified features are proprietary and specific to a single OEM. We kindly request that these be removed to enable broader participation.	
100	40	D.2. Item 7: CL 07 – Servers for GPU Nodes - S.no-20	Cloud Enabled Monitoring and Management: Sustainability - Carbon footprint reporting Predictive AI insights to forecast future energy costs and carbon emissions. Threshold-based alerting.	Change request: Kindly remove. Justification: All the specified features are proprietary and specific to a single OEM. We kindly request that these be removed to enable broader participation.	Modified as: Cloud Monitoring – Sustainability: The cloud management stack for CL-07 GPU nodes should integrate with the overall DCIM / monitoring solution to expose power and utilization metrics that can be used for energy and carbon accounting. Any OEM-provided sustainability functions (energy/cost/emissions views, threshold-based alerts) will be treated as value-adds; equivalent integration with third-party tools (e.g., via standard APIs or export formats) is acceptable.
101	40	D.2. Item 7: CL 07 – Servers for GPU Nodes - S.no-20	Cloud Enabled Monitoring and Management: AI-driven insights with workflow policy approvals	Request to amend as - AI-driven insights. with workflow policy approvals (Removed) Justification - The specifications are specific to one OEM based on its spec sheet. Hence, we request to amend to allow other leading OEMs to comply and participate. Change request: Kindly remove. Justification: All the specified features are proprietary and specific to a single OEM. We kindly request that these be removed to enable broader participation.	Modified as: Cloud Enabled Monitoring and Management: AI-driven insights. Implementations using standard infrastructure AIOps / predictive-monitoring or third-party AI-monitoring suites are acceptable, provided they support AI-assisted insights and configurable policy-based actions.
102	40	D.2. Item 7: CL 07 – Servers for GPU	Cloud Enabled Monitoring and Management: Software must support multi- tenancy	Change request: Kindly remove. Justification: All the specified features are proprietary and specific to a single OEM. We	As per RFP.

		Nodes - S.no-20		kindly request that these be removed to enable broader participation.	
103	40	D.2. Item 7: CL 07 – Servers for GPU Nodes - S.no-20	Cloud Enabled Monitoring and Management: Software must offer map-based visibility with an intuitive view of server health via global map.	Request to amend as - Software must offer map-based visibility with an intuitive view of server health via global map or Multi-system health views for locations as infrastructure topology map. Justification - The specifications are specific to one OEM based on its spec sheet. Hence, we request to amend to allow other leading OEMs to comply and participate. Change request: Kindly remove. Justification: All the specified features are proprietary and specific to a single OEM. We kindly request that these be removed to enable broader participation.	Modified as: Cloud Enabled Monitoring and Management: Software must offer map-based visibility with an intuitive view of server health via global map or Multi-system health views for locations as infrastructure topology map.
104	40	D.2. Item 7: CL 07 – Servers for GPU Nodes - S.no-20	Cloud Enabled Monitoring and Management: Multi-vendor server monitoring	Request to remove this as this is OEM specific and restricts other leading OEMs to comply and participate.	Deleted
105	40	D.2. Item 8- TE 01 - Talend Virtualization Infra (Cluster)-S.no-1(3 rd Point)	Platform: 3.Offered Platform shall have complete independence for the scaling of both Compute and Storage platform. Scaling of both components shall be completely independent.	The Leading and modern HCI solutions offers scale-out architecture for better scalability and performance. And True HCI means the native integration of compute, storage & virtualization. So, please change this point as below. 3. Offered platform shall have flexibility to scale compute and storage as required through compute only and storage only nodes.	Modified as: Platform – Independent scaling (TE-01 Cluster): The offered platform shall support independent scalability of compute and storage resources. It must be possible to increase storage capacity without mandatorily adding usable compute cores (for example, by adding storage-only nodes, storage-only sleds, or additional disks in existing storage components), and to increase compute capacity without mandatorily adding

					storage (for example, via compute-only nodes). Architectures that force 1:1 scaling of CPU and TBs in fixed ‘monolithic’ HCI nodes, causing purchase of unwanted CPU/RAM merely to obtain extra storage, shall not be accepted.
106	41	D.2. Item 8- TE 01 - Talend Virtuali zation Infra (Cluster)–S.no- 4(7 th Point)	Capacity & Scalability: 7.Vendor shall provision additional 32 dedicated CPU Cores for Storage layer functionality as well in the overall offered solution.	As mentioned in 4.1 the requirement is only 404 cores but as per 4.4 bidders to quote 32 core processors per node, which will total up to 512 cores. So, the overheads can be taken care of using the additional cores. So, please remove this point.	Modified as: Capacity & Scalability – Storage CPU overhead: The minimum 404 vCPU cores specified for Talend workloads shall be exclusive of storage-layer overheads and Node level redundancy overheads. The bidder shall size and provide additional CPU cores and memory required for the storage/HCI/SDS layer (metadata, replication, erasure coding, deduplication, compression, etc.) over and above the Talend application requirement, and must present sizing calculations for these storage-layer CPU and memory needs in the technical proposal and compliance statement. As a guideline, this overhead shall be no less than the equivalent of 32 physical CPU cores for the proposed configuration; higher allocations based on the OEM’s best-practice sizing are acceptable and encouraged.
107	41	D.2. Item 8- TE 01 - Talend Virtuali zation Infra (Cluster)–S.no-4	Capacity & Scalability 4. Each compute node / engine shall be supplied with at-least 4 x 10Gbps IP ports, 1.5 TB memory along with dual latest generation Intel processor with minimum 2.5 Ghz clock speed and each processor with 32 Cores. Each compute node shall have at-least two number of minimum 480	Removal of Intel Justification: Since Intel & AMD can both support.	Modified as: Capacity & Scalability 4. Each compute node / engine shall be supplied with at-least 4 x 10Gbps IP ports, 1.5 TB memory along with dual latest generation Processors with minimum 2.5 Ghz clock speed and each processor with

			GB NVMe drives for boot. Offered compute node shall also support 25Gbps & 100Gbps IP ports.		32 Cores. Each compute node shall have at-least two number of minimum 480 GB NVMe drives for boot. Offered compute node shall also support 25Gbps & 100Gbps IP ports.
108	41	D.2. Item 8- TE 01 - Talend Virtuali zation Infra (Cluster) -S.no-4	Capacity & Scalability 1. Offered Platform shall be supplied with at-least /Minimum 8 number of Compute nodes(minimum total cores 404) in single cluster, scalable to at-least 16, along with 60 TB capacity using 7.68TB NVMe drives on storage layer. 4. . Each compute node / engine shall be supplied with at-least 4 x 10Gbps IP ports, 1.5 TB memory along with dual latest generation Intel processor with minimum 2.5 Ghz clock speed and each processor with 32 Cores. Each compute node shall have at-least two number of minimum 480 GB NVMe drives for boot. Offered compute node shall also support 25Gbps & 100Gbps IP ports	Removal of 8 Nodes Justification: Since more cores and storage capacity can be given in lesser nodes, request you to allow to size the solution based on total number of cores & required storage capacity.	Modified as: Capacity & Scalability 1. Offered Platform shall be supplied with at-least /Minimum 16 number of Compute nodes (minimum total cores 404) in single cluster, along with 60 TB capacity using 7.68TB NVMe drives on storage layer. 4. Each compute node / engine shall be supplied with at-least 2 x 200 Gbps front-end ports, 1.5 TB memory. Each compute node shall have at-least two number of minimum 480 GB NVMe drives for boot. Required Network interconnect switches for Cluster communication and cables should be included. Others remain same.
109	42	D.2. Item 8- TE 01 - Talend Virtuali zation Infra (Cluster) -S.no-6	Data Protection: 1. Offered platform shall use all NVMe flash drives only for storage layer. All offered NVMe Flash drives on storage layer shall be configured with Hardware raid with double parity protection or in Replication factor 3 (RF-3) for asked usable capacity, if hardware raid is not supported by the bidder. 2. Failure of any compute engine shall not reduce the overall number of	1. RF3 is superior that hardware raid with double parity. Both are not equivalent. So please mention RF3 only or change as raid with double parity protection or in Replication factor 2 (RF-2). 2. Hyperconverged infrastructure is meant to natively integrate compute & storage for better performance. If the requirement is to have 60TB usable even after a node failure, please change this point as below. 2. Failure of any compute engine shall not reduce the total usable storage.	Modified as: Data Protection – Replication / RAID: All storage for the TE-01 cluster shall be based on NVMe flash. The storage layer shall be configured with Replication Factor 3 (RF-3) (or equivalent redundancy factor) for Talend data, ensuring that the cluster can tolerate the loss of at least two nodes or two drives in different failure domains without data loss. Use of RF-2 or equivalent ‘single-failure’ protection is not

			offered drives. In case vendor is not supporting this capability, then additional one node of the same configuration in terms of compute and storage shall be provided.		permitted for production Talend data at this scale. Where hardware RAID is used for any boot or metadata volumes, it shall provide at least double-parity protection (for example RAID-6), but the primary data layer for Talend workloads shall use RF-3 or stronger SDS-level redundancy.
110	42	D.2. Item 8- TE 01 - Talend Virtuali zation Infra (Cluster)-S.no- 10	Investment Protection: Offered platform shall non-disruptively scalable to higher generation platform within the given family without any forklift upgrade. There shall be no downtime, no data re-balancing while upgrading to next generation model.	Advanced/modern HCI solutions automatically balances the data for better availability and for improved performance without disruption. No data re-balancing means lot of manual intervention required. So, please modify this point as below. Offered platform shall non-disruptively scalable to higher generation platform within the given family without any forklift upgrade. There shall be no downtime while upgrading to next generation model.	Clarified as: Offered platform shall non-disruptively scalable to higher generation platform within the given family without any forklift upgrade. There shall be no downtime, while upgrading to next generation model.
111	43	D.2. Item 8- TE 01 - Talend Virtuali zation Infra (Cluster)-S.no- 17	Replication for Storage Layer: 1. Offered self-service private cloud platform shall have flexibility to de-link the storage layer for hardware-based data replication so that, if required customer can decide the customized configuration at DR location. 2. Offered and supported hardware-based replication shall support incremental replication after resumption from Link Failure situation or during failback operations. 3. Offered and supported hardware-based replication shall have capability to create the application consistency group for replication operations. Shall have	Hardware-based replication is legacy method and meant for SAN environments. This provides limited failure-domain coverage with complex DR orchestration. So, please allow bidders to quote advanced software defined storage-based replication by changing the points as below. 1. Offered self-service private cloud platform shall have flexibility to customize configuration at DR location. 2. Offered replication shall support incremental replication after resumption from Link Failure situation or during failback operations. 3. Offered replication shall have capability to create the application consistency group for replication operations. Shall have flexibility to have 200 volumes per protection group.	Modified as: Replication for Storage Layer – SDS based: Replication for the TE-01 storage layer shall be implemented at the software-defined storage / hypervisor level (VM/vDisk or volume abstraction), not restricted to legacy array-hardware replication. The self-service private cloud platform shall allow flexible configuration at the DR site (including differing node counts or disk mixes). The replication mechanism shall support incremental replication after link failures and during failback, and shall provide application or crash-consistent protection/consistency groups across multiple VMs and volumes, with support for at least 200 volumes per group. Hardware-array replication may be offered as an additional option, but is not

			flexibility to have more than 256 volumes per consistency group.		mandatory if equivalent or superior SDS-level replication is provided.
112	43	D.2. Item 8- TE 01 - Talend Virtuali zation Infra (Cluster)-S.no- 18(1 st Point)	Ransomware Detection: 1.The offered platform shall have in-built inline data-adaptive ransomware detection engine for the volumes being used for Virtual machines.	In-built ransomware detection engine will limit the bidders and this is favoring to one OEM. So, please remove this point or change as below. 1. The offered platform shall have in-built or external ransomware detection engine for the Virtual machines.	Modified as: Ransomware Detection: 1.The offered platform shall have in-built or external inline data-adaptive ransomware detection engine for the volumes being used for Virtual machines. If the offered solution doesn't support in-built ransomware detection, bidder must include external OEM supported ransomware detection and protection software with required hardware and software licenses for 3 Years 24/7 support from OEM with 2 hours response time, for a minimum 200TB of capacity and 1000 VMs as part of their solution. Bidder should integrate the solution with the offered HCI solution.
113	43	D.2. Item 8- TE 01 - Talend Virtuali zation Infra (Cluster)-S.no- 18(3 rd Point)	Ransomware Detection: 3. It shall be possible to select a specific set of volumes or group of volumes to be enabled for Ransomware detection for block data being used for virtual machines.	In-built ransomware detection engine will limit the bidders and this is favoring to one OEM. So, please remove this point.	Modified as: Ransomware Detection – Volume level: Ransomware detection and protection shall be mandatory for the TE-01 environment. It shall be possible to enable ransomware/anomaly detection for a selected set of volumes or groups of volumes used by Talend VMs, so that RTGS can prioritise protection of high-value datasets without incurring cluster-wide overhead. This capability may be provided either as an in-built feature of the storage/HCI platform or via an external OEM-supported ransomware

					detection and protection solution, supplied with all required hardware, software licences, and 3-year 24×7 OEM support. External solutions must integrate with the storage/virtualization layer at volume or VM/vDisk level (for example via APIs, sensors, or agents) to provide effective detection and alerting.
114	45	D.2. Item 8-TE 01 - Talend Virtualization Infra (Cluster)-S.no-23(1 st point)	Data Backup Solution: 1. The offered Hyper-visor Shall have an inbuilt native data backup solution which shall be able to protect VM and Hosts to Target Storage provider and shall provide the following functionalities: a. Shall support CIFS, NFS, S3 from Storage providers as a backup target.	In-built data backup solution will limit the bidders to quote leading HCI solutions. So, please allow bidders to quote the best third-party backup solution or provide the details of existing backup solution in SDC. Bidders can provide required licenses to use the same.	No Change, as per RFP. Please read 23.2 specifications for solutions which doesn't support inbuilt native backup solution.
115	46	D.2. Item 8-TE 01 - Talend Virtualization Infra (Cluster)-S.no-31(2 nd Point)	Pass Thru Support: 2. The offered Hypervisor shall support the passthrough support for GPU and USB devices.	Please change this point as below for wider participation. 2. The offered solution shall support the GPU passthrough and USB devices.	No Change as per RFP.
116	59 & 60	E.27. - S.no-4(g)	Preliminary examination of Bids: Initial Bid scrutiny will be held and incomplete details as given below will be treated as non- responsive. If Proposals, g. With lesser validity period	We request RTGS to clarify, does RTGS referring it to 90 days from date of submission or beyond 90 days of validity.	Clarified as: If any bidder submits bid with bid validity less than the bid validity period mentioned in the tender document, then it will be treated as non responsive bid.

117	60	E.29. Bid validity	Bid validity: The offer submitted by the Bidders should be valid for minimum period of days as mentioned in Bid data Sheet from the date of submission of bid. If required, purchaser may ask the bidder to extend the bid validity along with the extension of the EMD validity. Failure to extend the bid validity and EMD validity will be treated as non-responsive and will not be considered for further evaluation.	Price for hardware components are totally dependant on dollars rate and OEM pricing, therefore extension beyond 90 days from date of submission may not be feasible unless demanded in corrigendum before submission.	Clarified as: In rare cases, where tender process is delayed due to unforeseen reasons, extension of bid validity may be requested. Generally it is on mutual consent.
118	62	E.36 S.no-2	Purchaser's right to vary quantities at time of award: Purchaser reserves the right to place the repeat orders at the quoted price, in addition to the Quantity for which bid has been called for. However, this condition will not create any right to the bidder to demand such repeat order. During the validity of the contract period thereof, the bidder should be ready to supply any no. of devices as requested.	We request RTGS to confirm the expected delta change in quantities, or quantify it in percentage, extreme reduction in quantities may negatively impact the bidders commercially.	Clarified as: +/- 25%.
119	67	F.4- S.no-1	User license and patent rights: 1. The Vendor shall provide licenses for all software products, whether developed by it or acquired from others. In the event of any claim asserted by a third party for software piracy, the vendor shall act expeditiously to extinguish such claim. If the vendor fails to comply and the Purchaser is required to pay compensation to a third party resulting from such software piracy, the vendor shall be responsible for compensation including all expenses, court costs and lawyer fees. The Purchaser will give	(i) Provided that APTS uses such goods, software package or part thereof for the purpose and manner specified in the RFP to be read with licensing conditions. (ii) APTS shall bring any such instance of claim to the notice of the Vendor without admitting the same and without any delay and render reasonable cooperation to the Vendor to defend APTS and Vendor against such claim.	As per RFP.

			notice to the vendor of such claim, if it is made, without delay. 2.The Vendor shall indemnify the purchases against all third party claims of infringement of patent, trademark or industrial design rights arising from use of the goods, software package or any part thereof.		
120	71 &7 2	F.16 S.no-4	Spare parts: The Vendor shall ensure availability of spares in stock at his nearest service centre for immediate delivery such spare parts as: (a) are necessary for a minimum of 5 years of operation after installation at the Purchaser's sites (b) are necessary to comply with specifications.	We request RTGS to confirm contract period & list of expected spare parts in order to work on realistic budget covering spares cost for additional 2 years.	Modified as: Spare parts – Availability & EoL: The OEM shall guarantee availability of spare parts for all offered models for a minimum of 5 years from the date of installation, in line with typical enterprise hardware lifecycle practices. End-of-Life / End-of-Support products shall not be offered. The selected vendor must submit an EoL/EoS Commitment Letter from the OEM confirming this 5-year spares and support availability. To meet the uptime and SLA requirements, critical spares must be stocked to support Next Business Day (NBD) replacement as a minimum service level.
121	72	F.20 S.no-4	Maintenance service: The vendor will accomplish preventive and breakdown maintenance activities to ensure that all hardware, and firmware execute without defect or interruption for at least required up time.	We request RTGS to confirm how many PM to perform in contract period of 3 years.	Modified as: Maintenance service – Preventive maintenance: The vendor shall carry out both preventive and breakdown maintenance to ensure all hardware and firmware meet the required uptime. Preventive Maintenance (PM) for IT hardware under this contract shall be performed once every quarter (4× per year), i.e., approximately every 90 days, for the entire contract period (warranty

					plus AMC as applicable). Each PM visit shall include, at minimum, physical inspection and dust cleaning, verification of environmental conditions, firmware and hardware health checks, and validation of alarms and logs, with detailed PM reports submitted to RTGS.
122	73	F.19 S.no-6	Payment: The annual maintenance and repair cost as per separate agreement if any, shall be paid in equal quarterly installments at the end of each quarter as per the rates quoted and agreed.	We request RTGS to confirm the AMC period post warranty of 3 years, as currently quoted rates may not be valid for separate agreement unless confirmed through this requirement.	Clarified as: AMC rates are not asked in the tender. Hence this condition is not applicable.
123	74	F.23.	Assignment: The Vendor shall not assign, in whole or in part, its obligations to perform under this Contract, except with the prior written consent from Purchaser.	However, Vendor shall be allowed to assign the receivables under the Contract with the Purchaser, to any Financial Institution or Bank, with prior intimation to the Purchaser.	As per RFP.
124	76	F.36 S.no-1	Termination for default: The Purchaser, without prejudice to any other remedy for breach of Contract, by written notice of default sent to the Vendor, may terminate the Contract in whole or in part: - if the Vendor fails to deliver any or all of the Goods/services within the time period(s) specified in the contract, or within any extension thereof granted by the Purchaser pursuant to Clause F.25 or - if the Vendor fails to perform any other obligation(s) under the Contract or	Right to Terminate the contract to be exercised after serving at least 30 days notice to cure the issue i.e. if despite receipt of notice of not less than 30 days Vendor fails to rectify the default only then notice for termination to be sent.	As per RFP.

			c. if the Vendor, in the judgment of the Purchaser has engaged in corrupt or fraudulent practices in competing for or in executing the Contract.		
125	76	F.39 S.no-1	Termination for convenience: Purchaser, may at any time by giving 30 days written notice to the Vendor, terminate the Contract, in whole or in part, for its convenience. The notice of termination shall specify that termination is for the Purchaser/Purchaser's convenience, the extent to which performance of the Vendor under the Contract is terminated, and the date upon which such termination becomes effective.	Notice period for termination for convenience should be at least 6 months and that should be without prejudice to the rights of the Vendor to claim the costs and damages it incurs due to exercising of such right by the APTS.	As per RFP.
126	78	G.1	Delivery and Installation period: Successful Bidder shall deliver the goods/services, install and commission the same within Six (6) Weeks from the date of issue of Notification of Award at the AP State Data Centre or any other place as indicated by the Department.	Request to amend this clause as- we want 8 weeks of time from the date of issue of notification of award at the ap state data centre or any other place as indicated by the department. Six weeks may not be sufficient we request to amend it to 12 weeks - 6 weeks for HW delivery and 6 weeks for installation. Successful Bidder shall deliver the goods/services, install and commission the same within 12 weeks from the date of issue of Notification of Award at the AP State Data Centre or any other place as indicated by the Department. Successful Bidder shall deliver the goods/services, install and commission the same within Fifteen [15] weeks from the date of issue of Notification of Award at the AP	Modified as : Delivery and Installation period: Successful Bidder shall deliver the goods/services and complete installation and commissioning within twelve (12) weeks from the date of issue of Notification of Award at the AP State Data Centre or any other place as indicated by the Department. As a guideline, hardware delivery should normally be completed within eight (8) weeks, with the remaining period used for installation, configuration, and acceptance testing. Item-wise staggered deliveries may be allowed where operationally justified, but overall completion for all items shall not exceed 12 weeks, and any Liquidated Damages (LD) under G.2/G.3 shall be computed with reference to these item-wise contractual delivery dates.

				State Data Centre or any other place as indicated by the Department. Justification: Minimum delivery timeline of material by OEM is 8 Weeks and installation for min 6 weeks. Request amendment: Please revise delivery & installation timeline from 6 weeks to 10 weeks (or allow item-wise staggered delivery) considering OEM manufacturing/import lead times and logistics. Also clarify whether LD (G.2/G.3) will be applied on total schedule or item-wise/staggered PO basis.	
127	80	G.4.3.	Service Level Requirements		Add: The Bidder/OEM must have local technical capability, including availability of qualified L2/L3 engineers within AP/Telangana States, capable of supporting high-speed 100G/200G Ethernet fabric, NVMe-over-Fabric, and AI infrastructure, as part of their eligibility and resource deployment plan.
128	79- to 81	G.4 & G.5	G.4-Runtime & Uptime Requirements G5- Penalty for failure to maintain during warranty period	M/s. Orange Business: Since no dedicated O&M manpower is sought in this procurement, kindly clarify monitoring responsibility during warranty: (a) Will RTGS/APSDC provide EMS/helpdesk/NOC and only call logging & OEM coordination expected from bidder? OR (b) Bidder must provide/operate EMS monitoring and 24x7 monitoring manpower. If bidder scope includes EMS, please specify minimum features, integration points and acceptance criteria.	Modified as: Runtime/Uptime – Monitoring responsibility: Bidder/OEM shall be fully responsible for meeting the runtime and uptime requirements and associated penalties specified in G.4 and G.5. The central EMS/helpdesk/NOC will be provided and operated by APSDC/RTGS; however, the successful bidder must deploy and integrate the OEM's monitoring / call-home / predictive-support software (for example, vendor-supplied server/storage monitoring tools) with the RTGS/APSDC NOC. Such

					tools shall generate automated alerts, health telemetry, and, where supported, auto case-logging to the OEM, so that incidents are promptly detected and addressed. No dedicated 24x7 on-site O&M manpower is included in this RFP, but the bidder shall ensure remote support processes, escalation and response mechanisms are in place to honour the uptime SLAs.
129	81	G.6	Payment Terms and Conditions: Upon submission of Delivery Challan & Installation cum Acceptance Test Report - 90% of items cost as per Contract / PO for which DCs & IR cum AT reports submitted.	90% on equipment delivery, 5% on installation & 5% on Satisfactory Performance Certificate. Upon submission of Delivery Challan - 80% of items cost as per Contract / PO for which DCs & IR. Installation cum Acceptance Test Report - 10% of items cost as per Contract / PO Justification: OEMs takes advance payment from Vendors. Hence it will be burden on vendors. Kindly consider this change. Please clarify in detail what is required in Acceptance Test Report , as per bidder understanding Supply and Install with power on Test is sufficient for this payment terms.	As per RFP
130			General Point: Kindly insert Limitation of Liability Clause	Most importantly there is no clause for exclusion of indirect losses etc. and no cap on limitation of liability. We propose clause as mentioned below to be included: (i) Neither Party; not its subsidiaries or its affiliates will be liable to the other Party, whether in contract, tort (including negligence), strict liability or otherwise, for loss of business, revenue, profits, loss of	As per RFP.

			goodwill or reputation; or indirect, consequential, or special loss, arising in connection with any order, product, service, related documentation, information and/or intended use thereof, even if a Party has been advised, knew or should have known of the possibility of such damages. ii. Subject to the above and notwithstanding anything to the contrary elsewhere contained herein, the maximum aggregate liability of the Supplier for all claims under or in relation to this agreement shall be regardless of the form of claims shall be limited to the annual amount to be paid to the Vendor by Purchaser as of the date when the liability arose, under the applicable statement of work/Purchase Order that gives rise to such liability.	
--	--	--	--	--

All other terms and conditions remain unaltered.

Yours sincerely

Manager (Proc 2)