

**Responded to Pre Bid Query : RFP/1265 - Servers
PROCUREMENT, INSTALLATION, COMMISSIONING AND MAINTENANCE OF SERVERS**

Sl. No	RFP Page No	RFP Clause No	Existing Clause	Query/Suggestions	Banks Response
1	60	A. Blade Server (Category 1) -> Clause # 10	Minimum 3 PCIe 5.0 Type based x16 Slots supporting Ethernet, FC adapters / CNA Card	<p>Kindly modify the clause as " Minimum 2 PCIe 4.0/5.0 slots supporting Ethernet, FC adapters / CNA Card. All the PCI slots to be fully populated so that the server can provide maximum ethernet, FC / Converged Ethernet bandwidth"</p> <p>Justification: Our Blade Server Architecture has 2 CNA cards, which works using PCI 4.0. This provides card level redundancy, at the same time providing both Ethernet & FC Adapter functionality and up to 200 Gig cumulative network bandwidth using 2 CNA Cards. This architecture ensures both redundancy and performance at the same time allowing to create up to 512 programmable virtual interfaces (vNIC or vHBAs) as required, without any limitations to the Bank.</p>	Please refer corrigendum No. 3
2	60	A. Blade Server (Category 1) > Clause # 11	Minimum One (1) USB 3.0 port for connecting US	<p>Kindly modify the clause as "Minimum One (1) USB 3.0 port for connecting USB device such as Keyboard/Mouse"</p> <p>Justification: Connecting USB 'drive' to the server can cause security issues like virus or other malware to come to the server & hence not recommended. In earlier generation servers, USB ports were used to install the Operating System or drivers etc. Now a days, the same can be achieved more securely using virtual Media mounted to the server in a secure fashion, which eliminates the need to connect a USB Drive to the server. Some times, customer wants to connect Keyboard/Mouse/Video (KVM) console directly to the Server, which is supported in our server using can server using KVM dongle cable, which has USB port connector for a keyboard and mouse.</p>	<p>No Change, as per RFP</p> <p>USB port is required for connecting Keyboard and Mouse.</p>
3	60	A. Blade Server (Category 1) -> Clause # 12	ACPI 5.1 Compliant, PCIe 5.0 Compliant; WOL Support; Microsoft/VMware/RHEL Logo Certifications; USB 3.0 support or better.	<p>Request you to modify the clause as " ACPI 5.1 or higher Compliant, PCIe 4.0 / 5.0 Compliant; WOL Support or equivalent management feature to remotely power on the Server; Microsoft/VMware/RHEL Logo Certifications; USB 3.0 support or better."</p> <p>Justification: The ACPI version 5.1 is released around 2014. Our server supports newer ACPI 6.5 & hence request you to modify the clause. Also, in our servers, while PCI 4.0 is used, it offers up to 200 Gig Converged Bandwidth to each Blade Server. Since there is no performance or redundancy challenges, kindly allow PCI 4.0/5.0 Compliant. The traditional WOL (Wake on LAN) has many disadvantages and shortcomings. Mainly Security consideration while implementing WOL solution since it uses the concept of direct broadcast. The magic packets used for WOL do not provide any confirmation after reaching target devices, and also, since the magic packets are broadcast and not routed, the WOL does not work outside the local subnet. Now a days these shortcomings are addressed using advanced features implemented in the Server Management software, which allows to remotely power ON & power off the server. Hence, kindly allow the changes as requested.</p>	Please refer corrigendum No. 3
4	62	B. Rack Server (Category 2) => Clause # 5	Memory => Should support scalability up to 4 TB without having to replace the existing DIMMs	<p>Request you to modify as "Should support scalability up to 3 TB without having to replace the existing DIMMs"</p> <p>Justification: In case of AMD based server, the dual socket architecture supports maximum of 24 Memory slots. Since the ask on day 1 is 2 TB, which is provided with 16 x 128 GB Memory modules, when fully populating all 24 Memory slots, the server supports maximum of 3 TB RAM.</p>	Please refer corrigendum No. 3
5	62	B. Rack Server Category 2 => Clause # 9	The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum four cards to provide card level redundancy also Should support RDMA over Converged Ethernet (RoCE).	<p>Request you to modify as "The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum two cards to provide card level redundancy also Should support RDMA over Converged Ethernet (RoCE)"</p> <p>Justification: Now a days server NIC card comes with 4 x 25 Gig ports. This ensures that with 2 cards, we can provide 8 x 25 Gig bandwidth, at the same time providing card level redundancy</p>	No Change, as per RFP
6	62	B. Rack Server Category 2 => Clause # 9	One (1) NIC or Equivalent port dedicated for Remote Management of the server hardware. Should support RDMA over Converged Ethernet (RoCE).	<p>Request you to modify the clause as "One (1) NIC or Equivalent port dedicated for Remote Management of the server hardware."</p> <p>Justification: Management of the server does not require RDMA over Converged Ethernet (RoCE). This is asked specifically for Rack Server Category 2 only. (Category 3, 4, 5 etc does not ask for RoCE for management,) kindly remove this point.</p>	Please refer corrigendum No. 3

7	62	B. Rack Server Category 2 => Caluse # 10	Expansion Slots -> Minimum 6 PCIe 5.0 Type based x16 Slots supporting Ethernet, FC adapters	Request you to modify as " Minimum 6 PCIe 4.0 / 5.0 Type based x8/x16 Slots supporting Ethernet, FC adapters " Justification: Our server supports mix of PCI 4.0/5.0 interfaces, at the same time offers all the functionality asked in the RFP including the quantity of adapters, speed for network & HBA adapters.	Please refer corrigendum No. 3
8	63	Rack Server Category 3 => Caluse # 3	Processor -> 96 cores each socket, dual socket-192 cores (AMD EPYC 9655 or Intel Xeon 6972P)	Request you to modify as " 96 cores each socket, dual socket-192 cores (AMD EPYC 9654 or Intel Xeon 6972P) Justification: The AMD 9655 CPU requires DDR5-6400 DIMM memory modules. As per RFP specification, the server need to be populated with 3 TB RAM on day 1 & scalable up to 4 TB RAM. This requires 256 GB Memory module to be populated (as AMD Server only supports 24 DIMM slots), however 256 GB DDR5-6400 memory module is not available in the market today. Kindly amend this clause to allow AMD 9654 CPU, as asked in Category 1,2,4,5,6 in this RFP.	No Change, as per RFP
9	63	C. Rack Server Category 3 => Caluse # 9	Ethernet Controller => The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum four cards to provide card level redundancy also Should support RDMA over Converged Ethernet (RoCE).	Request you to modify as "The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum two cards to provide card level redundancy also Should support RDMA over Converged Ethernet (RoCE)" Justification: Now a days server NIC card comes with 4 x 25 Gig ports. This ensures that with 2 cards, we can provide 8 x 25 Gig bandwidth, at the same time providing card level redundancy	No Change, as per RFP
10	63	C. Rack Server Category 3 => Caluse # 10	Expansion Slots -> Minimum 6 PCIe 5.0 Type based x16 Slots supporting Ethernet, FC adapters	Request you to modify as " Minimum 6 PCIe 4.0 / 5.0 Type based x8/x16 Slots supporting Ethernet, FC adapters " Justification: Our server supports mix of PCI 4.0/5.0 interfaces, at the same time offers all the functionality asked in the RFP including the quantity of adapters, speed for network & HBA adapters.	Please refer corrigendum No. 3
11	64	D. Rack Server Category 4 => Caluse # 10	Ethernet Controller => The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum four cards to provide card level redundancy also Should support RDMA over Converged Ethernet (RoCE).	Request you to modify as "The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum two cards to provide card level redundancy also Should support RDMA over Converged Ethernet (RoCE)" Justification: Now a days server NIC card comes with 4 x 25 Gig ports. This ensures that with 2 cards, we can provide 8 x 25 Gig bandwidth, at the same time providing card level redundancy	No Change, as per RFP
12	64	D. Rack Server Category 4 => Caluse # 11	Expansion Slots -> Minimum 6 PCIe 5.0 Type based x16 Slots supporting Ethernet, FC adapters	Request you to modify as " Minimum 6 PCIe 4.0 / 5.0 Type based x8/x16 Slots supporting Ethernet, FC adapters " Justification: Our server supports mix of PCI 4.0/5.0 interfaces, at the same time offers all the functionality asked in the RFP including the quantity of adapters, speed for network & HBA adapters.	Please refer corrigendum No. 3
13	65	E: GPU Rack Servers Type 1 (Category 5) => Caluse # 4	GPU -> Server should support three NVIDIA H100 Tensor Core GPU card or Intel® Data Center GPU Max 1100 or AMD Instinct™ MI250X. On day one server should be populated with two cards and all required licenses (including Nvidia AI enterprise / equivalent) for GPU slicing with VMware, Nvidia enterprise tech support, training, inferencing etc.	Request you to modify the clause as " Server should support two NVIDIA H100 Tensor Core GPU card or Intel® Data Center GPU Max 1100 or AMD Instinct™ MI250X. On day one server should be populated with two cards and all required licenses (including Nvidia AI enterprise / equivalent) for GPU slicing with VMware, Nvidia enterprise tech support, training, inferencing etc. Justification : Server have 6 PCI slot, after configuring 2 GPU , server can accomodate only 4 card (2 NIC card and 2 HBA card). There is no slot for third GPU	Please refer corrigendum No. 3
14	65	E: GPU Rack Servers Type 1 (Category 5) => Caluse # 14	Ethernet Controller => The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum four cards to provide card level redundancy also Should support RDMA over Converged Ethernet (RoCE).	Request you to modify as "The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum two cards to provide card level redundancy also Should support RDMA over Converged Ethernet (RoCE)" Justification: Now a days server NIC card comes with 4 x 25 Gig ports. This ensures that with 2 cards, we can provide 8 x 25 Gig bandwidth, at the same time providing card level redundancy. Also Server have 6 PCI slot, if we configure 2 GPU then server can accomodate only 4 card (2 NIC card and 2 HBA card).	Please refer corrigendum No. 3
15	65	E: GPU Rack Servers Type 1 (Category 5) => Clause # 15	Expansion Slots -> Minimum 6 PCIe 5.0 Type based x16 Slots supporting Ethernet, FC adapters	Request you to modify as " Minimum 6 PCIe 4.0 / 5.0 Type based x16 Slots supporting Ethernet, FC adapters"	Please refer corrigendum No. 3

16	66	F: GPU Rack Servers Type 2 (Category 6) => Sever Type	General Specification	<p>The 8 GPU Server based on Nvidia H200 GPU (SXM) is designed based on Nvidia HGX Reference Architecture. While the main specification of the server such as 8 x H200 GPU, Nvlink, Nvswitch etc is common across all the Server OEM's, since this is an appliance approach, for each OEM, the server comes with different CPU, Memory, Disk, NIC card etc. Hence request you to modify the specification as. While there is slight differences in these configs, since all the OEM architecture is based on Nvidia HGX reference architecture, they all can be used for AI workloads. Request you to make the changes to allow other OEM's to participate</p> <p>Kindly modify as "Each Node must be configured with Nvidia 8 x H200 141 GB GPUs (SXM) connected via Nvidia Nvlink with NV Switch. The total aggregate memory per node from the GPUs should be at least 1128 GB. Solution should come with all required licenses (including Nvidia AI enterprise / equivalent) for GPU slicing with VMware, Nvidia enterprise tech support, training, Inferencing etc. Each Node should be populated with 2 Intel / AMD CPU's with 64 Cores each Socket (128 cores in total), 3 TB RAM, Boot Drive, 8 x NVMe Drives for Data, 8 x 400 Gig NIC ports for GPU-GPU East-West communication (Intel GPU Backend Network), 200 Gig Network for North-South/Front-end network & separate Management network ports</p>	No Change, as per RFP
17	66	F: GPU Rack Servers Type 2 (Category 6) => Server Type. Clause # 3	Server Type -> 64 Cores each socket, Dual socket- 128 cores) AMD 9554 or Intel Xeon Platinum 8592V	<p>Request you to modify as "64 Cores each socket, Dual socket- 128 cores) AMD 9554 or AMD 9575F or Intel Xeon Platinum 8592V</p> <p>Justification: Different OEM's Nvidia 8 GPU HGX Server uses different CPU's. Kindly accommodate the change for us to participate. All the CPU's mentioned here are dual socket 64 Core CPU's (128 Cores in total)</p>	Please refer corrigendum No. 3
18	66	F: GPU Rack Servers Type 2 (Category 6) => Clause # 10 => Memory	Each Server should be installed with minimum 3 TB Memory DDR5 or higher RDIMM ; Should support scalability up to 4 TB without having to replace the existing DIMMs	<p>AI Servers are purpose build appliance and as per best practices, all the slot has to be fully populated for better performance of all the memory channel. AMD based Servers supports maximum of 24 Memory slots per Server. The 8 GPU AMD based server supports only max of 128 GB RAM per slot, this can go up to 3 TB RAM. Hence request you to modify the clause as the server "Each Server should be installed with minimum 3 TB Memory DDR5 or higher RDIMM," which is currently supported in AMD based 8 GPU Servers. Request you to remove the scalability up to 4 TB, and amend this clause for wider participation.</p>	Please refer corrigendum No. 3
19	66	F: GPU Rack Servers Type 2 (Category 6) => Clause # 10 => Internal Storage	Minimum 2 x 480 GB M2 SSD drives or higher having capability to be used as Mirror Disk with Above RAID controller for installing the operating system hypervisor and 8 * 7.6 TB NVMe drives per node.	<p>Request you to modify as "Minimum 2 x 480 GB M2 SSD drives or higher having capability to be used as Mirror Disk with Above RAID controller for installing the operating system, and The system should provide additional 2 x 1.92 or higher NVMe drives per node"</p> <p>Justification: The internal drives in a 8 GPU HGX Appliance from each OEM has different capacity & quantity of drives qualified. Please allow the OEM to quote the NVME disk capacity qualified for thier server. In case of an AI deployment, the data is stored on external storage and is accessed by the servers/GPU's as required. The local NVME drives is used only for a temporary caching of this data & having higher cache size does not help in terms of AI workload or performance. Please allow the OEM to quote the NVME drives as per theri qualification.</p>	Please refer corrigendum No. 3
20	66	F: GPU Rack Servers Type 2 (Category 6) => Caluse # 13 => HBA	Each server must have dedicated 4 x 32G dedicated fiber channel SAN Storage connectivity distributed across two cards	<p>Please remove this clause. The 8 GPU Server for AI use cases typically accesses the external storage via Ultra Ethernet of 100 or 200 Gig speed for North South Traffic including Storage. Since AI Applications need access to huge amount of data for training and fine tuning, 32 Gig FC HBA is a limiting factor. Hence Our 8 GPU Server is designed with Nvidia Network Adapter offering 2 x 200 Gig ethernet speed. All NVIDIA reference architecture also do not recommend FC HBA, as FC HBA becomes a limiting factor due to BW's.</p>	Please refer corrigendum No. 3

21	66	F: GPU Rack Servers Type 2 (Category 6) => Caluse # 14	<p>The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum four cards to provide card level redundancy and should support RDMA over Converged Ethernet (RoCE)</p>	<p>Please modify as " The bandwidth required for network per server is minimum Storage Network :- 2x200G per Server Back end :- 8x400G per server (GPU to GPU communication) User Network : 2x200G per server OOB Mgmt :- 1G Port per server and should support RDMA over Converged Ethernet (RoCE). Also, the server need to have 8 x 400 Gig ethernet ports which are mapped to the 8 x H200 GPU's for non blocking GPU-to-GPU communication between the servers.</p> <p>Justification: For NVIDIA HGX systems, GPU to GPU connectivity is over 400G Fabric, Storage & user traffic is on 200G connectivity, 25G will be a limiting factor to achieve full performance of the system, Each GPU by default is shipped with 400G NIC Card, same will be provided by default. Different OEM's have different implementation of network uplink. In our server, 200 Gig port across 2 cards are provided to ensure the required bandwidth and redundancy. Each server has 8 x H200 GPU's and all of them need to communicate to GPU's in other servers in a non blocking manner. To facilitate this, each server need to have separate 8 x 400 Gig Ethernet ports for GPU-GPU communication.</p>	Please refer corrigendum No. 3
22	67	GPU Rack Servers Type 2 (Category 6) => Caluse # 15	<p>Minimum 12 PCIe 5.0 Type based x16 Slots supporting Ethernet, FC adapters for each server.</p>	<p>Please change this clause as "Minimum 10 PCIe 5.0 Type based x16 Slots supporting Ethernet for each server." Justification: The 8 GPU server is a AI specialised server which comes with required PCI slots</p>	Please refer corrigendum No. 3
23	67	General compliance	<p>For all the categories of the servers mentioned above i.e. from Category 2 to Category 6, the bidder must comply with below specifications. Bidder(s) must submit their response in yes or no only, any compliance with qualified statement shall be treated as non-compliance</p>	<p>Please modify as "For all the categories of the servers mentioned above i.e. from Category 2 to Category 5, the bidder must comply with below specifications. Bidder(s) must submit their response in yes or no only, any compliance with qualified statement shall be treated as non-compliance Justification: The 8 GPU Rack Server is used for AI Workload, and only AI Specific OS certification applies to the same - such as Ubuntu, Redhat Linux and typically runs container workload on Redhat Openshift / K8s / Bare metal systems.</p>	No Change, as per RFP
24	67	General compliance	<p>ACPI 5.1 Compliant, PCIe 5.0 Compliant; WOL Support; Microsoft/VMware/RHEL Logo Certifications.</p>	<p>Request you to modify the clause as " ACPI 5.1 or higher Compliant, PCIe 4.0 / 5.0 Compliant; WOL Support or equivalent management feature to remotely power on the Server; Microsoft/VMware/RHEL Logo Certifications; USB 3.0 support or better for Category 2 to 5". Since Category 6 is 8 GPU Server for AI appliance purpose built using Nvidia HGX reference architecture, support for general purpose OS such as Microsoft Windows & VmWare is not supported. This is not a limiting factor for AI Workload, as almost all customers deploy AI workload on Linux OS such as Ubuntu/Redhat Linux with Redhat Openshift Container platform on top.</p>	Please refer corrigendum No. 3
25	67	General compliance	<p>Should be Compatible with Latest Windows server, Red Hat Linux/Openshift and VMware ESXi Server version 7.0 U3/VCF 4.5.2 and all later/upgraded/higher versions. Also, the supplied hardware should support all version upgrades coming in next 7 years. If hardware supplied by selected OEM is not compatible with these releases during 7 years from date of commissioning, bidder need to replace hardware with compatible VMware, Windows server, Red Hat Linux/Openshift release without any additional cost to the Bank</p>	<p>Please modify this clause for Category 6 GPU Server as "Should be Compatible with Red Hat Linux/Openshift. Also, the supplied hardware should support all version upgrades coming in next 7 years. If hardware supplied by selected OEM is not compatible with these releases during 7 years from date of commissioning, bidder need to replace hardware with compatible VMware, Windows server, Red Hat Linux/Openshift release without any additional cost to the Bank" Justification: Since Category 6 is 8 GPU Server specifically built for workload based on Nvidia HGX reference architecture, support for general purpose OS such as Microsoft Windows & VmWare is not required. This is not a limiting factor for AI Workload, as almost all customers deploy AI workload on Linux OS such as Ubuntu/Redhat Linux with Redhat Openshift Container platform on top.</p>	No Change, as per RFP
26	70	G. HARDWARE - BLADE ENCLOSURE -> Clause # 16	<p>Blade Server Ethernet Interconnect Each enclosure must have total 16 x 25G ports with pair of redundant network modules with 8 x 25G uplink ports respectively in each. i.e. 8 x 25G uplink bandwidth must be maintained even after 50% of interconnect failure and without using multi-chassis aggregation.</p>	<p>Request you to change to "Blade Server Ethernet Interconnect/ Converged Network Interconnect / Unified Fabric Switch " Also request you to modify the detailed specification as " "Each enclosure must have total 16 x 25G ports or higher with pair of redundant network modules / Converged Network Interconnect / Unified Fabric Switch with 8 x 25G uplink ports respectively in each. i.e. 8 x 25G uplink bandwidth must be maintained even after 50% of interconnect failure and without using multi-chassis aggregation."</p>	Please refer corrigendum No. 3

27	70	G. HARDWARE - BLADE ENCLOSURE -> Clau	Blade Server FC Interconnect "The enclosure must have redundant Fibre Channel Interconnect modules with minimum 8 x 32Gbps Uplink Ports to the SAN Switch i.e. 4 x 32G uplink bandwidth must be maintained even after one interconnect failure. Each module should be fully licensed to use all available ports in case of both Multi chassis and Standalone chassis scenario."	Request you to change to "Blade Server FC Interconnect/Converged Network Interconnect / Unified Fabric Switch" Also "The enclosure must have redundant Fibre Channel Interconnect modules / Converged Network Interconnect / Unified Fabric Switch with minimum 8 x 32Gbps Uplink Ports to the SAN Switch i.e. 4 x 32G uplink bandwidth must be maintained even after one interconnect failure. Each module should be fully licensed to use all available ports in case of both Multi chassis and Standalone chassis scenario."	Please refer corrigendum No. 3
28	84	Appendix-F : Indicateive Price Bid : SI No VIII	Nvidia H200 -> 24 Qty	Nvidia H200 used in 8 GPU server Category 6, comes as a GPU Tray with 8 GPU's. Hence please clarify 24 Quantity indicated 3 x 8 GPU Tray as spare ? Please confirm	No Change, as per RFP
29	84	Appendix-F : Indicateive Price Bid : SI No IX	Intel® Data Center GPU Max 1100	In category 5 Server, there is option given to support either Nvidia H100 or Intel Max 1100 or AMD Instinct MI 250X GPU. We plan to quote server with Nvidia H100 GPU along with the server. With this, spare should also be H100 = 72 Quantity only. And Not Intel Max 1100 GPU. please confirm.	No Change, as per RFP
30	84	Appendix-F : Indicateive Price Bid : SI No X	AMD Instinct MI250X	In category 5 Server, there is option given to support either Nvidia H100 or Intel Max 1100 or AMD Instinct MI 250X GPU. We plan to quote server with Nvidia H100 GPU along with the server. With this, spare should also be H100 = 72 Quantity only. And Not AMD Instinct MI 250X GPU. please confirm.	No Change, as per RFP
31	19	13. PERIOD OF BID VALIDITY AND VALIDITY OF PRICE QUOTED IN REVERSE AUCTION (RA): (i)	Price quoted by the Bidder in Reverse auction shall remain valid for duration of 36 calendar months from the date of Purchase order.	Price quoted by the Bidder in Reverse auction shall remain valid for duration of 12 calendar months from the date of Purchase order. Justification: The price from component suppliers is dynamic	No Change, as per RFP
32	80	4. Term of the Project - Project Schedule; Milestones and delivery locations	Delivery of all equipment should be within 6 weeks and installation, testing, commissioning within 10 weeks from date of placing of order.	Change Request: Delivery of all equipment should be within 16 weeks and installation, testing, commissioning within 20 weeks from date of placing of order. Justification: Components, especially GPUs, have a high lead time	No Change, as per RFP
33	56	Appendix-B. Bidder's Eligibility Criteria	6. The Bidder (including its OEM, if any) should either be Class-I or Class-II local supplier as defined under this RFP.	Please remove this clause. Justification: There are no OEMs for Blade Servers in India	No Change, as per RFP
34	96	Appendix-L Other Terms and Penalties 5.	On-site comprehensive warranty and AMC: The warranty will be from 60 months from date of installation or 63 months from date of delivery, whichever is later and AMC for further 24 months	Change Request: On-site comprehensive warranty and AMC: The warranty will be from 60 months from date of installation or 63 months from date of delivery, whichever is earlier and AMC for further 24 months Justification: Installation delays due to banks activity Windows is not in our control.	Please refer corrigendum No. 3
35	96	Appendix-L Other Terms and Penalties 3.	Delivery of all equipment should be within 6 weeks from date of placing of the order. In the event of the any or all equipment(s) not being delivered, installed, tested and commissioned within a period of 10 weeks from date of Purchase Order, a penalty of 0.5 percent of the total cost of equipment for each week or part thereof the delay, subject to maximum amount of ten (10) percent of the total cost of equipment will be charged to vendor. This amount of the penalty so calculated shall be deducted at the time of making final payment after successful installation and commissioning of hardware.	Change Request: Delivery of all equipment should be within 16 weeks from date of placing of the order. In the event of the any or all equipment(s) not being delivered, installed, tested and commissioned within a period of 20 weeks from date of Purchase Order, a penalty of 0.5 percent of the total cost of equipment for each week or part thereof the delay, subject to maximum amount of ten (10) percent of the total cost of equipment will be charged to vendor. This amount of the penalty so calculated shall be deducted at the time of making final payment after successful installation and commissioning of hardware. Justification: Components, especially GPUs, have a high lead time	No Change, as per RFP
36	96	Appendix-L Other Terms and Penalties 4.	OEMs engineering team should be available within one day on call in case of any issues or upon intimation of bank.	Change Request: OEMs engineering team should be available within five working days on call in case of any issues or upon intimation of bank. Justification: Engineering teams operate in a different time zone	No Change, as per RFP
37	22	AWARD CRITERIA AND AWARD OF CONTRACT	Applicability of Preference to Make in India, Order 2017 (PPP-MII Order)	Blade servers, blade chassis and GPU servers are not manufactured in India and will not qualify under MII.	No Change, as per RFP
38	60	Blade Servers - Internal Storage	For 240 servers(120 at each site) out of 1775, minimum 2x 480 GB M2 SSD drives or higher having capability to be used as Mirror Disk with Above RAID controller for installing the operating system hypervisor and 4 x 3.xx TB high performance NVMe.	Can you please elaborate on this requirement. Do we have to configure 6 drives per server - 2 x 480GB SSD for OS and 4 x 3.xx TB additional drives	No Change, as per RFP yes, bidder need to configure 6 drives per server - 2 x 480GB SSD for OS and 4 x 3.xx TB additional drives
39	60	Host Bus Adaptor(HBA)/Converged Network Adaptor(CNA)	Each blade server must be equipped with at least two 32G dedicated Fibre Channel SAN storage connectivity ports, with redundancy, supporting both Fibre Channel (FC) and NVMe over Fabrics (NVMeoF) protocols.	Blade servers can support maximum 1 x dual port 32GB FC HBA card per server. It will not be possible to configure any additional FC HBA card in blade server.	Please refer corrigendum No. 3

40	60	Ethernet Controller /Converged Network Adaptor(CNA)	The bandwidth required for network per server is a minimum of 2 x 50G per server that must be partitioned across a minimum of two cards to provide card level redundancy with a minimum of 50G bandwidth available per card. Should support RDMA over Converged Ethernet (RoCE)	Blade servers can be configured with 2 cards of dual port 10/25Gbps to provide required ethernet bandwidth.	No Change, as per RFP
41	60	Expansion Slots	Minimum 3 PCIe 5.0 Type based x16 Slots supporting Ethernet, FC adapters / CNA Card	Blade servers has mix of Gen 4 and Gen 5 slots - 2 slots of Gen 4 and 1 slot of Gen 5.	Please refer corrigendum No. 3
42	62	Rack Servers (Category 2) - Processor	64 cores each socket, dual socket-128 cores (AMD 9554 or Intel Xeon Platinum 8592V)	We don't have 8592V processor. Instead, we can quote for 8592+ 64C 1.9Ghz processor. Request you to please change this point.	Please refer corrigendum No. 3
43	62	Rack Servers (Category 2) - Memory	Should support scalability up to 4 TB without having to replace the existing DIMMs	To provide 4Tb memory scalability, we have to configure 128GB DIMMs which are commercialy expensive. Also, with AMD processors it will be challenging to provide 4TB memory scalability. Request you to change this point to 2Tb or 3Tb scalability.	Please refer corrigendum No. 3
44	62	Rack Servers (Category 2) - Expansion Slots	Minimum 6 PCIe 5.0 Type based x16 Slots supporting Ethernet, FC adapters	2U rack server supports Up to 8 x PCIe Gen4 or up to 4 x PCIe Gen5 slots. Request you to change the clause accordingly.	Please refer corrigendum No. 3
45	63	Rack Servers (Category 3) - Processor	96 cores each socket, dual socket-192 cores (AMD EPYC 9655 or Intel Xeon 6972P)	Intel 6972P Granite rapids processor and AMD 9655 processor is still not getting shipped from DELL. The tentative timeline is April 2025. Request you to please consider another processor for this requirement.	No Change, as per RFP
46	63	Rack Servers (Category 3) - Memory	Should support scalability up to 4 TB without having to replace the existing DIMMs	To provide 4Tb memory scalability, we have to configure 128GB DIMMs which are commercialy expensive. Also, with AMD processors it will be challenging to provide 4TB memory scalability. Request you to change this point to 2Tb or 3Tb scalability.	Please refer corrigendum No. 3
47	63	Rack Servers (Category 3) - Expansion Slots	Minimum 6 PCIe 5.0 Type based x16 Slots supporting Ethernet, FC adapters	2U rack server supports Up to 8 x PCIe Gen4 or up to 4 x PCIe Gen5 slots. Request you to change the clause accordingly.	Please refer corrigendum No. 3
48	64	Rack Servers (Category 4) - Internal Storage	8 * 7.6 TB NVMe drives per node (vSAN ESA certified nodes)	Can you confirm what is meant by vSAN ESA certified nodes	No Change as per RFP vSAN Express Storage Architecture(ESA) certified nodes are x86 servers that are pre-configured, tested and certified for VMware Hyper-Converged Infrastructure Software.
49	64	Rack Servers (Category 4) - Expansion Slots	Minimum 6 PCIe 5.0 Type based x16 Slots supporting Ethernet, FC adapters	2U rack server supports Up to 8 x PCIe Gen4 or up to 4 x PCIe Gen5 slots. Request you to change the clause accordingly.	Please refer corrigendum No. 3
50	65	GPU Rack Servers Type 1 (Category 5) - Processor	64 Cores each socket, Dual socket- 128 cores) AMD 9554 or Intel Xeon Platinum 8592V	We don't have 8592V processor. Instead, we can quote for 8592+ 64C 1.9Ghz processor. Request you to please change this point.	Please refer corrigendum No. 3
51	65	GPU Rack Servers Type 1 (Category 5) - Memory	Should support scalability up to 4 TB without having to replace the existing DIMMs	To provide 4Tb memory scalability, we have to configure 128GB DIMMs which are commercialy expensive. Also, with AMD processors it will be challenging to provide 4TB memory scalability. Request you to change this point to 2Tb or 3Tb scalability.	Please refer corrigendum No. 3
52	65	GPU Rack Servers Type 1 (Category 5) - Ethernet Controller	The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum four cards to provide card level redundancy and should support RDMA over Converged Ethernet (RoCE)	Server will not have enough PCIe slots after populating 2 GPU cards in the server. Request you to please change ethernet cards required to 4 x 25G ports per server.	Please refer corrigendum No. 3
53	65	GPU Rack Servers Type 1 (Category 5) - Expansion Slots	Minimum 6 PCIe 5.0 Type based x16 Slots supporting Ethernet, FC adapters	2U rack server supports Up to 8 x PCIe Gen4 or up to 4 x PCIe Gen5 slots. Request you to change the clause accordingly.	Please refer corrigendum No. 3
54	66	GPU Rack Servers Type 2 (Category 6) - Processor	64 Cores each socket, Dual socket- 128 cores) AMD 9554 or Intel Xeon Platinum 8592V	We don't have 8592V processor. Instead, we can quote for 8592+ 64C 1.9Ghz processor. Request you to please change this point.	Please refer corrigendum No. 3
55	66	GPU Rack Servers Type 2 (Category 6) - Memory	Should support scalability up to 4 TB without having to replace the existing DIMMs	To provide 4Tb memory scalability, we have to configure 128GB DIMMs which are commercialy expensive. Also, with AMD processors it will be challenging to provide 4TB memory scalability. Request you to change this point to 2Tb or 3Tb scalability.	Please refer corrigendum No. 3

56	66	GPU Rack Servers Type 2 (Category 6) - Expansion Slots	Minimum 12 PCIe 5.0 Type based x16 Slots supporting Ethernet, FC adapters for each server	Server with 8 GPU supports upto 10 x16 Gen5 (x16 PCIe) full-height, half-length. Request you to please change this point	Please refer corrigendum No. 3
57	70	Hardware - Blade Enclosure	In case of multi-chassis aggregation (not more than 4 chassis in one set), each master enclosure must have redundant master modules to eliminate single point of failure. The multi-chassis setup should have total 40 x 25G uplink ports with 20 x 25G uplink ports for each master enclosure i.e. 20 x 25G uplink bandwidth must be maintained even after 50% of master interconnect failure for multi-chassis set.	DELL support upto 10 chassis in multi chassis aggregation. Limiting 4 chassis in one chassis will give advantage to one OEM as they will have to quote lesser number of switches whereas in DELL case we will have to provide more switches to comply on this point	No Change, as per RFP
58	81	Payment schedule	50% + taxes of the Servers (hardware, software and warranty) will be released on delivery of hardware.	We request you to revise this clause as "75% + taxes of the Servers (hardware, software and warranty) will be released on delivery of hardware."	No Change, as per RFP
59	81	Payment schedule	The remaining 50% + taxes of the Servers (hardware, software and warranty) will be released on commissioning, verification of bill of material by Bank/CDAC/Third party and submission of PBG.	We request you to revise this clause as "The remaining 25% + taxes of the Servers (hardware, software and warranty) will be released on commissioning, verification of bill of material by Bank/CDAC/Third party and submission of PBG.	No Change, as per RFP
60	62	B. Rack Server Category 2 => Caluse # 9	The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum four cards to provide card level redundancy also Should support RDMA over Converged Ethernet (RoCE).	Request you to modify as "The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum two cards to provide card level redundancy also Should support RDMA over Converged Ethernet (RoCE)" Justification: Now a days server NIC card comes with 4 x 25 Gig ports. This ensures that with 2 cards, we can provide 8 x 25 Gig bandwidth, at the same time providing card level redundancy	No Change, as per RFP
61	63	Rack Server Category 3 => Caluse # 3	Processor -> 96 cores each socket, dual socket-192 cores (AMD EPYC 9655 or Intel Xeon 6972P)	Request you to modify as " 96 cores each socket, dual socket-192 cores (AMD EPYC 9654 or Intel Xeon 6972P)" Justification: The AMD 9655 CPU requires DDR5-6400 DIMM memory modules. As per RFP specification, the server need to be populated with 3 TB RAM on day 1 & scalable up to 4 TB RAM. This requires 256 GB Memory module to be populated (as AMD Server only supports 24 DIMM slots), however 256 GB DDR5-6400 memory module is not available in the market today. Kindly amend this clause to allow AMD 9654 CPU, as asked in Category 1,2,4,5,6 in this RFP.	No Change, as per RFP
62	63	C. Rack Server Category 3 => Caluse # 9	Ethernet Controller => The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum four cards to provide card level redundancy also Should support RDMA over Converged Ethernet (RoCE).	Request you to modify as "The bandwidth required for network per server is minimum 8 x 25 G per server that must be partitioned across minimum two cards to provide card level redundancy also Should support RDMA over Converged Ethernet (RoCE)" Justification: Now a days server NIC card comes with 4 x 25 Gig ports. This ensures that with 2 cards, we can provide 8 x 25 Gig bandwidth, at the same time providing card level redundancy	No Change, as per RFP
63	72	H. Rack Enclosure	Rack Sizing	Please Specify the Dimension of the Rack & qty	No Change, as per RFP The bidder needs to supply standard 42 U Racks as per the specifications of servers being quoted.
64	72	H. Rack Enclosure	Environmental Monitoring	Please specify controller based or IPDU based environmental monitoring solution	No Change, as per RFP Any one of the monitoring solution.
65	73	H. Rack Enclosure	Intelligent Power Distribution Units (PDUs)	Please let us know type of the PDU need to consider Single Phase or three phase IPDU	No Change, as per RFP Three Phases
66	73	H. Rack Enclosure	Intelligent Power Distribution Units (PDUs)	Switching with outlet level monitoring or only outlet level monitoring required	No Change, as per RFP Any one of the outlet level monitoring is ok.
67	59	Blade Server(Category 1)	Each Server should be installed with minimum 1.5 TB Memory DDR5 with 4800 MT/s memory speed	5th Gen Intel® Xeon® Scalable Processors support DDR5 4400MT/s , 4800 MT/s and upto 5600 MT/s. Request SBI to consider DDR5 4400 MT/s and higher memory speeds	No Change, as per RFP
68	62	Rack Server(Category 2)	Each Server should be installed with minimum 2 TB Memory DDR5 or higher RDIMM with 4800 MT/s memory speed	5th Gen Intel® Xeon® Scalable Processors support DDR5 4400MT/s , 4800 MT/s and upto 5600 MT/s. Request SBI to consider DDR5 4400 MT/s and higher memory speeds	No Change, as per RFP
69	63	Rack Server(Category 3)	Each Server should be installed with minimum 3 TB Memory DDR5 or higher RDIMM with 4800 MT/s memory speed	5th Gen Intel® Xeon® Scalable Processors support DDR5 4400MT/s , 4800 MT/s and upto 5600 MT/s. Request SBI to consider DDR5 4400 MT/s and higher memory speeds	No Change, as per RFP
70	55	3	The Bidder must have an average turnover of minimum Rs. 250 crore during last 03 (three) financial year(s) i.e. FY 2021-22, FY 2022-23 and FY 2023-24.	The Bidder must have an average turnover of minimum Rs. 250 Crore during last 03 (three) financial year(s) from the business of IT / ITeS i.e. FY2021-22, FY2022-23 and FY2023-24.	No Change, as per RFP

71	55	5	Bidder should have experience of five years in supply, installation, commissioning and maintenance of minimum 250 blade/rack Servers. (Combining all orders).	Bidder should have experience of five years in supply, installation, commissioning and maintenance of minimum 250 blade/rack Servers. (Combining all orders with one order should be min 100 Nos of single order). The Client Should be Govt / BSFI / PSU.	No Change, as per RFP
72	81	13	50% + taxes of the Servers (hardware, software and warranty) will be released on delivery of hardware. The remaining 50% + taxes of the Servers (hardware, software and warranty) will be released on commissioning, verification of bill of material by Bank/CDAC/Third party and submission of PBG. No payment will be made on part delivery of ordered hardware. Payment for AMC will be made quarterly in arrears.	Request to change payment terms as below:- 1) 70% + Taxes of the Servers (hardware, software and warranty) will be released on delivery of hardware. 2) 30% + Taxes of the Servers (hardware, software and warranty) will be released on commissioning, verification of bill of material by Bank/CDAC/Third party and submission of PBG.	No Change, as per RFP
73	80	4	Delivery of all equipment should be within 6 weeks and installation, testing, commissioning within 10 weeks from date of placing of order.	Delivery of all equipment should be within 10 weeks and installation, testing, commissioning within 14 weeks from date of placing of order.	No Change, as per RFP
74	59	Appendix-C/ Technical & Functional Specifications/ A. Blade Servers (Category 1)/ 5.Memory	Should support scalability up to 2 TB without having to replace the existing DIMMs.	Existing Blade Servers in SBI offers memory scalability upto 4TB. Request to modify the clause to have identical solution with similar performance for future as well. Hence, request to modify the clause as follows: Should support scalability upto 4TB without having to replace the existing DIMMs.	No Change, as per RFP
75	60	Appendix-C/ Technical & Functional Specifications/ A. Blade Servers (Category 1)/8. Host Bus Adaptor(HBA)/ Converged Network Adaptor(CNA)	Each blade server must be equipped with at least two 32G dedicated Fibre Channel SAN storage connectivity ports, with redundancy, supporting both Fibre Channel (FC) and NVMe over Fabrics (NVMeoF) protocols.	Considering the uniformity with the Bank's existing setup. Request to modify the clause as follows : "Each blade server must be equipped with at least two 32G dedicated Fibre Channel SAN storage connectivity ports, with redundancy, supporting both Fibre Channel (FC) and NVMe over Fabrics (NVMeoF) protocols. Bidder must quote dedicated FC protocol supported card and this must not be a shared adapter with ethernet/SCSI/FCoE."	No Change, as per RFP
76	60	Appendix-C/ Technical & Functional Specifications/ A. Blade Servers (Category 1)/ 9.Ethernet Controller / Converged Network Adaptor(CNA)	The bandwidth required for network per server is a minimum of 2 x 50G per server that must be partitioned across a minimum of two cards to provide card level redundancy with a minimum of 50G bandwidth available per card. Should support RDMA over Converged Ethernet (RoCE).	As Bank's existing servers are currently configured with 100G bandwidth per card i.e 4x50G per server with card level redundancy . To keep uniformity with existing SBI setup we recommend to change the existing clause as follow: 4x50G per server across two cards with card level redundancy . Minimum 100G bandwidth per card	No Change, as per RFP
77	60	Appendix-C/ Technical & Functional Specifications/ A. Blade Servers (Category 1)/ 13.Operating System	Should be Compatible with Latest Windows server, Red Hat Linux and VMware ESXi Server version 7.0 U3/VCF 4.5.2 and all later/upgraded/higher versions. Also, the supplied hardware should support all version upgrades coming in next 7 years. If hardware supplied by selected OEM is not compatible with these released during 7 years from date of commissioning, bidder need to replace hardware with compatible VMware, Windows server, Red Hat Linux, Redhat Openshift release without any additional cost to the Bank.	Hardware Offered supports latest generation of Operating systems like Windows, Linux, Vmware, etc. Upgrade / higher versions of Software are released by the respective software OEMs and they must provide support and compatibility with existing hardware. Hence request you remove this clause, as software is not part of this RFP.	No Change, as per RFP
78	62	Appendix-C/ Technical & Functional Specifications/ C. Rack Servers (Category 2)/ 8. Host Bus Adaptor(HBA)/ Converged Network Adaptor(CNA)	Each server must be equipped with four 32G dedicated Fibre Channel SAN storage connectivity ports, distributed across two cards, and supporting both Fibre Channel (FC) and NVMe over Fabrics (NVMeoF) protocols.	As the Rack server comes dedicated ethernet and FC card Hence we recommend to remove Converge Network Adapter in this clause	No Change, as per RFP
79	63	Appendix-C/ Technical & Functional Specifications/ C. Rack Servers (Category 3)/ 8. Host Bus Adaptor(HBA)/ Converged Network Adaptor(CNA)	Each server must be equipped with four 32G dedicated Fibre Channel SAN storage connectivity ports, distributed across two cards, and supporting both Fibre Channel (FC) and NVMe over Fabrics (NVMeoF) protocols.	As the Rack server comes dedicated ethernet and FC card Hence we recommend to remove Converge Network Adapter in this clause	No Change, as per RFP
80	64	Appendix-C/ Technical & Functional Specifications/ D: Rack Servers (Category 4) / 5. Memory	Should support scalability up to 2 TB without having to replace the existing DIMMs	Existing Blade Servers in SBI offers memory scalability upto 4TB. Request to modify the clause to have identical solution with similar performance for future as well. Hence, request to modify the clause as follows: Should support scalability upto 4TB without having to replace the existing DIMMs.	No Change, as per RFP
81	64	Appendix-C/ Technical & Functional Specifications/ D: Rack Servers (Category 4) / 7. Internal Storage	8 * 7.6 TB NVMe drives per node (vSAN ESA certified nodes)	vSAN Certified nodes are preconfigured from factory. While in this case these nodes will be configured onsite in SBI Datacenter, hence request to change the existing clause as : vSAN ESA compatible configuration instead of vSAN ESA certified node	No Change, as per RFP

82	64	Appendix-C/ Technical & Functional Specifications/ D: Rack Servers (Category 4) / 7. Internal Storage	8 * 7.6 TB NVMe drives per node (vSAN ESA certified nodes)	SBI is currently using vSAN on Bladed Architecture for better consolidation & higher performance. This eliminates the need for external network switches upto 60 nodes in a vSAN cluster. For similar solution on Rack Servers will require Total 480 ports on external TOR switch, that needs to be procured additionally. Blade Solution provides many savings for the Bank commercially and technically. Hence we suggest to use Blade Solution for vSAN Nodes.	No Change, as per RFP
83	64	Appendix-C/ Technical & Functional Specifications/ D: Rack Servers (Category 4) / 9. HBA	Each server must have dedicated 2 x 32G dedicated fiber channel SAN Storage connectivity distributed across two cards	vSAN nodes are configured using internal drives for Storage requirement and don't require any external storage. Hence, we recommend to remove HBA card from the vSAN nodes	No Change, as per RFP
84	65	Appendix-C/ Technical & Functional Specifications/ E: GPU Rack Servers Type 1 (Category 5)/4. GPU	Server should support three NVIDIA H100 Tensor Core GPU card or Intel® Data Center GPU Max 1100 or AMD Instinct™ MI250X. On day one server should be populated with two cards and all required licenses (including Nvidia AI enterprise / equivalent) for GPU slicing with VMware, Nvidia enterprise tech support, training, inferencing etc.	NVIDIA H100 is much higher in performance compared to Intel Max 1100 and AMD MI250X. Hence, we recommend to change : Server should support four NVIDIA H100 GPUs or Intel Gaudi 3 or AMD MI300 GPUs.	No Change, as per RFP
85	66	Appendix-C/ Technical & Functional Specifications/ F: GPU Rack Servers Type 2 (Category 6)/3. Server type	64 Cores each socket, Dual socket- 128 cores) AMD 9554 or Intel Xeon Platinum 8592V	In GPU server the computation is handled by the GPUs through its cores and memory. CPU processing is very low. Hence request to modify the Processor as follows : 48 cores each socket, Dual socket - 96 cores AMD 9455 or Intel Xeon Platinum 8558	No Change, as per RFP
86	69	Appendix-C/ Technical & Functional Specifications/ G: HARDWARE - BLADE ENCLOSURE/ 2. Interconnect Support	Should support simultaneous housing of Ethernet, FC, iSCSI or CNA offering Hot Pluggable & Redundancy as a feature. Enclosure should have a minimum of 6 Interconnect Bays populated in case of 4 chassis group, in case of standalone chassis per chassis 2 Interconnect Bays need to provide in redundancy and Server to server communication for servers in same chassis must happen over these switches. These switches must be internal or external to server and should not be shared across multiple chassis.	Request to add following for better understanding : Each Chassis must be offered with dedicated Network Modules in redundancy. Server to Server communication should be in 1:1 non blocking	No Change, as per RFP
87	81	Payment Schedule	50% + taxes of the Servers (hardware, software and warranty) will be released on delivery of hardware. The remaining 50% + taxes of the Servers (hardware, software and warranty) will be released on commissioning, verification of bill of material by Bank/CDAC/Third party and submission of PBG. No payment will be made on part delivery of ordered hardware. Payment for AMC will be made quarterly in arrears.	80% on Delivery, 10% on PBG & 10% on sign off.	No Change, as per RFP
88	21	17.ii	All the Bidders who qualify in the evaluation process shall have to participate in the online reverse auction to be conducted by Bank's authorized service provider on behalf of the Bank.	Request Bank to confirm if there is H1 Elimination before RA ?	No change, as per RFP As per GeM's rules.
89	32	32.ix	Service Provider shall be agreeable for on-call/on-site support during peak weeks (last and first week of each month) and at the time of switching over from PR to DR and vice-versa.	For OnCall Support : Will Bank give the Partner Remote Access for Trouble shooting ? What is the frequency of switching over from PR toDR and viceversa ?	No change, as per RFP. Bank will give the Remote Access for Trouble shooting only to Technical Support of OEM. The frequency of switching over from PR to DR and viceversa is periodic as well as on demand, as and when required.
90	61	16	5 years Warranty and 2 years AMC - 24x7 comprehensive onsite support from OEM with maximum 2 hours response time with 6 hours Call to Resolution including part replacement, access to OEM support portal, OEM technical support on 24X7X365 basis. Highest Level of Proactive and Reactive support covering Half	Half yearly Firwarware and Proactive Health analysis can be done remotely as well ? OR Its an Onsite Activity tobe done by the OEM ?	No change, as per RFP. The activity to be done onsite.
91	79	Description of Services -1	6. The price discovered through this RFP will be valid for State Bank Group (SBG)	Request the Point of Valid for the Entire State Bank Group top be Deleted (Because : The Price discovered in this RFP Is basis the Qty And is very specific to this RFP) We request this cannot be applied to Group companies of SBI which may have smaller qty of Servers Requirement	No Change, as per RFP
92	80	Integration / Migration Requirements with existing systems	Yes, New hardware should be integrated with old platform without any additional cost to the Bank.	Request Clarification on what is Integration with Old platform ? As this is Greenfield expansion as understood from the RFP	No Change, as per RFP New hardware should be integratable with existing vcenter.

93	80	Scalability Requirements	The rate determined through this RFP will be applicable to other departments of the bank, as well as its subsidiaries, joint ventures, offices, and foreign subsidiaries. These entities may place separate purchase orders for additional quantities beyond the 125% limit specified in this RFP. Each entity will have its own Service Level Agreement (SLA) and payment cycle based on this RFP & SLA.	Request the Point of Valid for the Entire State Bank Group top be Deleted (Because : The Price discovered in this RFP is basis the Qty And is very specific to this RFP) We request this cannot be applied to Group companies of SBI which may have smaller qty of Servers Requirement	No Change, as per RFP
94	81	12. Training	Training to minimum 5 Bank officials and two vendor partners for 5 days regarding daily operations, troubleshooting and Management of Hardware supplied.	Can this Training be Virtual Training ? Incase Onsite : Will this training be at GITC Belapur	No change, as per RFP. The training will be in Bidder/OEM's Lab/office. However bank may ask for virtual or on premise training.
95	78	Description of Services, 3	In the event of Bank engaging the services of CDAC/any other party for inspection and testing of the supplied material, the bidder should ensure the presence of OEM engineer and successfully demonstrate that all equipment, software and services under this RFP have been delivered	Please share the testing criteria	No change, as per RFP. The hardware should comply with technical specifications mentioned in Appendix-C of this RFP
96	79	Description of Services, 7	Bank will procure minimum 40% servers of the quantity mentioned in first year	please share the timelines for procurement of remaining 60% quantity	No change, as per RFP. Bank may order remaining hardware within three years from the date of Purchase order
97	98	Appendix-L , Other Terms and Penalties: 6.e,	Vendor shall ensure that the full configuration of the Equipment is available to the Bank in proper working condition viz. uptime of 99.99% of the time on a 24*7 on quarterly basis	Requesting this to change on half yearly basis.	No change as per RFP
98	110	RESPONSIBILITIES OF SERVICE PROVIDER, 6.4	Service Provider shall report the incidents, including cyber incidents and those resulting in disruption of service and data loss/ leakage immediately but not later than one hour of detection.	Please let us know the resolution time line for the same.	No change as per RFP
99	80	Integration / Migration Requirements with existing systems, pt 6	Yes, New hardware should be integrated with old platform without any additional cost to the Bank	For VMWARE new license will not be integrated with old licensed infra, for that case who will provide the required license ? Please clarify.	No change as per RFP Bank will arrange the required licenses
100	80	Performance Requirements, Pt 7	Uptime requirement is 99.99% on quarterly basis for each rack server, server chassis including servers	Requesting this to change on half yearly basis.	No change as per RFP
101				Any network device or connectivity is required to supply or configure for the connection or Bank will provide that, Please confirm	No change as per RFP Networking switches will be arranged by bank.
102	108	4.2.2	Service Provider has the requisite technical and other competence, sufficient, suitable, qualified and experienced manpower/personnel and expertise in providing the Services to the Bank.	Any Residential Enginner required as manpower in DC & DR? if Yes, please let us know the number of resource required.	No change as per RFP
103	3	Tender Fees	Rs. 25,000/-	Request bank to Remove the Tender Fees	No Change, as per RFP
104	4	EMD	Rs. 2,00 Crore	Request bank to relax the EMD amount upto 50 Lakhs	No Change, as per RFP
105	5	Bank Guarantee	Rs 5,00,00,000/- (Rs Five Crore only) Performance Guarantee	Request bank to Relax the PBG upto 1 Crore with 7 Years & 3 months Validity period	No Change, as per RFP
106	19	PERIOD OF BID VALIDITY AND VALIDITY OF PRICE QUOTED IN REVERSE AUCTION (RA):	ii. Price quoted by the Bidder in Reverse auction shall remain valid for duration of 36 calendar months from the date of Purchase order.	Request bank to Modify the clause " Price quoted by the Bidder in Reverse auction shall remain valid for duration of 3 calendar months from the date of Purchase order."	No Change, as per RFP
107	55	Eligibility Clause - Point 6	The Bidder (including its OEM, if any) should either be Class-I or Class-II local supplier as defined under this RFP.	Request you to kindly remove this clause	No Change, as per RFP
108	84	Commercial Bid	Point 2 - Comprehensive annual Maint - 8 to 12%	Request you to modify the clause " This cost should be in the range upto 12 % p.a. of the Product cost".	No Change, as per RFP
109	80	Scope of Work and Payment Schedule	Point 4- Delivery of all equipment should be within 6 weeks and installation, testing, commissioning within 10 weeks from date of placing of order.	Request bank to modify the delivery upto 10 weeks and installation, testing, commissioning within 14 weeks from date of placing of order.	No Change, as per RFP
110	97	Appendix-L (Point C)	Point 6 -C - Response time with in two Hours	Request to change "Response Time within 4 Hours"	No Change, as per RFP
111	97	Appendix-L (Point D)	Point 6 -D - Resolution time - Within four hours	Request to change " Resolution time within 6 Hours"	No Change, as per RFP
112	98	Appendix-L (Point h)	Point 6 -C - Response time with in 30 minutes	Request to change "Response Time upto 2 Hours"	No Change, as per RFP
113	59	Appendix-C/ Technical & Functional Specifications/ A. Blade Servers (Category 1)/ 5.Memory	Should support scalability up to 2 TB without having to replace the existing DIMMs.	Existing Blade Servers in SBI offers memory scalability upto 4TB. Request to modify the clause to have identical solution with similar performance for future as well. Hence, request to modify the clause as follows: Should support scalability upto 4TB without having to replace the existing DIMMs.	No Change, as per RFP

114	60	Appendix-C/ Technical & Functional Specifications/ A. Blade Servers (Category 1)/8. Host Bus Adaptor(HBA)/ Converged Network Adaptor(CNA)	Each blade server must be equipped with at least two 32G dedicated Fibre Channel SAN storage connectivity ports, with redundancy, supporting both Fibre Channel (FC) and NVMe over Fabrics (NVMeoF) protocols.	Considering the uniformity with the Bank's existing setup. Request to modify the clause as follows : "Each blade server must be equipped with at least two 32G dedicated Fibre Channel SAN storage connectivity ports, with redundancy, supporting both Fibre Channel (FC) and NVMe over Fabrics (NVMeoF) protocols. Bidder must quote dedicated FC protocol supported card and this must not be a shared adapter with ethernet/SCSI/FCoE."	No Change, as per RFP
115	60	Appendix-C/ Technical & Functional Specifications/ A. Blade Servers (Category 1)/ 9. Ethernet Controller / Converged Network Adaptor(CNA)	The bandwidth required for network per server is a minimum of 2 x 50G per server that must be partitioned across a minimum of two cards to provide card level redundancy with a minimum of 50G bandwidth available per card. Should support RDMA over Converged Ethernet (RoCE).	As Bank's existing servers are currently configured with 100G bandwidth per card i.e 4x50G per server with card level redundancy . To keep uniformity with existing SBI setup we recommend to change the existing clause as follow: 4x50G per server across two cards with card level redundancy . Minimum 100G bandwidth per card	No Change, as per RFP
116	4	12	EMD- Rs. 2,00,00,000/-	As per GEM guidelines company is having more than 500 cr is exempted for EMD so please confirm.	No Change as per RFP EMD needs to be submitted by bidder
117	5	13	Rs 5,00,00,000/- (Rs Five Crore only) - Performance Security in form of BG should be valid for Seven year(s) and three months from the effective date of the Contract.	5 Cr PBG for 5 years to cover the warranty period . For AMC period take 1 cr PBG prior to completion of to 5 cr pbq .On receiving the 1cr PBG release 5 cr PBG.	Please refer Corrigendum No 3
118	81	13	Payment - 50% against Delivery .The remaining 50% + taxes of the Servers (hardware, software and warranty) will be released on commissioning, verification of bill of material by Bank/CDAC/Third party and submission of PBG. The remaining 50% + taxes of the Servers (hardware, software and warranty) will be released on commissioning, verification of bill of material by Bank/CDAC/Third party and submission of PBG.	Please consider 80% against Delivery and 20% against installation , on verification by bank and submission of PBG. Request you to drop verification by CDAC or Third party .	No Change, as per RFP
119	81	13	AMC Payment will be Quarterly	Please consider AMC Payment as monthly instead quarterly or quarterly in advance.	No Change, as per RFP
120	56	6	The Bidder (including its OEM, if any) should either be Class-I or Class-II local supplier as defined under this RFP.	As per market knowledge none of the leading OEM has Class I or ClassII product in the blade server category. You have also asked the same in the primary category. We request you to drop this clause so that we can participate.	No Change, as per RFP
121	79	2	The prices discovered through this RFP will be valid for a minimum period of three years from the date of initial purchase order.	As the period of purchase is too long i.e. 3 year, request you to incorporate the dollar escalation clause as follows. If the exchange rate with respect to the dollar (USD) increases by more than 3 % from the date of purchase order then SBI will pay extra to the tune of the percentage increase beyond 3%. The dollar rate will be fetched from the RBI website.	No Change, as per RFP
122	46	49	49. TERMINATION FOR CONVENIENCE: i. The Bank, by written notice of not less than 90 (ninety) days, may terminate the Contract, in whole or in part, for its convenience, provided same shall not be invoked by the Bank before completion of half of the total Contract period (including the notice period). ii. In the event of termination of the Agreement for the Bank's convenience, Service Provider shall be entitled to receive payment for the Services rendered (delivered) up to the effective date of termination.	This is not a service contract rather it is a capex contract wherein all the equipment are delivered immediately on order and warranty/ AMC support is provided. Due to this this clause is not applicable and so request you to drop this clause .	No Change, as per RFP
123	28	iv	. Service Provider shall provide and implement patches/ upgrades/ updates for Products (software/ firmware/ OS) as and when released by Service Provider/ OEM free of cost. Service Provider should bring to notice of the Bank all releases/ version changes	OS supply and installation is not part of the scope so request you to please confirm.	No Change, as per RFP. OS like RHEL, VMware ESXI, Windows OS are out of scope however bidder has to provide firmware, softwares for hardware management etc.

124	75		All devices should have dual power sources and also connected through different color Power Cables to easily identify that the devices are connected to both sources. The bidder is responsible for providing a complete infrastructure solution. The bank will solely provide the space, power, and cooling for the equipment. All other items, including servers, PDUs (power distribution units), SFPs (transceivers), racks, cables (end to end network cabling, SAN cabling with material and effort) etc. must be supplied, installed, and made ready for use by the bidder. Power requirements need to be submitted by the bidder in the format below for all required racks.	Request SBI Bank to provide clarity on following (a) Top of the rack switch fully populated with required SFP is to be provided by Bank (b) All required SAN switch fully populated with required SFP is to be provided by Bank (c) All required management switch with required port connected to SBI network will be provided by the Bank (d) All required power supply up to RACK /PDU will be provided by the Bank (e) In above para asked for SFPs- understand these are for server side – yes Bidder (TCPL) will factor (f) Also confirm that all required Operating System, middleware will be provided by the BANK	No Change, as per RFP. Understanding is correct.
125	78	2	2. To provide all necessary hardware and software required to make the solution work strictly as per technical specifications. The specifications given are minimum. Bidders can quote equivalent or higher technical specifications to meet the Bank's requirements. However, no weightage would be given for higher configurations	Any other software or installation of that is out of scope please confirm.	No Change, as per RFP
126	80	6	Integration / Migration Requirements with existing systems Yes, New hardware should be integrated with old platform without any additional cost to the Bank.	Any migration of data will be taken care by bank and is not part of the current RFP scope Bank has to ensure upgradation of any firmware or related hardware from which the migration has to happen if so required.	No Change, as per RFP
127	37		Service Provider shall report the incidents, including cyber incidents and those resulting in disruption of service and data loss/ leakage immediately but not later than one hour of detection.	This is not related to the server we are supplying. Please clarify the same.	No Change, as per RFP
128	100/101		Future additions of Hardware / Software: (a) The Bank would have the right to: i. Shift supplied systems to an alternative site of its choice. ii. Disconnect / connect / substitute peripherals such as printers, etc. or devices or any equipment / software acquired from another vendor. iii. Expand the capacity / enhance the features / upgrade the hardware / software supplied, either from Vendor, or third party, or developed in-house.	a) Shifting will be taken care of by bidder at mutually agreed rate. b) This is not relevant, request you to please drop this clause . c) the patches, upgrades etc as per standard warranty of OEM will be supplied as per the scope .	No Change, as per RFP
129	97		Beyond two hours up to six hours- Rs. 10,000 per instance	Request you to drop response , resolution of various incidence instead keep only SLA measurement and associated penalty and in that reduce the quantum as follows .	No Change, as per RFP
130	97		Beyond six hours up to twelve hours- Rs. 25,000 per instance	Request you to drop response , resolution of various incidence instead keep only SLA measurement and associated penalty and in that reduce the quantum as follows .	No Change, as per RFP
131	97		Beyond twelve hours up to twenty four hours - Rs. 1,00,000 per instance	Request you to drop response , resolution of various incidence instead keep only SLA measurement and associated penalty and in that reduce the quantum as follows .	No Change, as per RFP
132			General	We understand from the RFP all the Rack and within rack connectivity will be provided by the bidder. Please confirm all the required space, fiber runner path way, civil work (if required) will be provided by the bank.	No Change, as per RFP. The Bidder need to provide Rack, PDU, power cables, Network cables to connect to TOR switches and within chassis connectivity etc. Datacentre facilities like fiber cable tray, power, cooling will be provided by the bank.
133		1	5% of Total Project Cost	Request you to consider 2 % instead of 5% of Total Project Cost	No Change, as per RFP
134		2	7% of Total Project Cost	Request you to consider 3 % instead of 7% of Total Project Cost	No Change, as per RFP
135		3	10% of Total Project Cost. Also Bank reserves the right to terminate the contract and invoke the PBG.*	Request you to consider 5% instead of 10% of Total Project Cost	No Change, as per RFP

136	59	A. Blade Servers (Category 1) Sr. No 3 Processor	AMD 9274F or Intel Xeon Gold 6542Y (24 cores each socket, dual socket-48 cores)	<p>The processors selected by the bank from Intel and AMD for this RFP appear to be not directly comparable, which could inadvertently create an imbalance in the evaluation process. The bank has specified the Intel Xeon Gold 6542Y with 24 cores at 2.9 GHz, while the selected AMD processor is the AMD EPYC 9274F, which offers 24 cores at 4.05 GHz. This significant difference in base clock speed provides an inherent advantage to one processor OEM, potentially affecting the fairness of the competition.</p> <p>To ensure a level playing field and a fair comparison of equivalent processor capabilities, we propose the bank consider the AMD EPYC 9254. The AMD EPYC 9254 matches the Intel Xeon Gold 6542Y in terms of 24 cores and a 2.9 GHz base clock, providing a more balanced and equitable evaluation framework.</p> <p>Below is a detailed comparison of the processors for reference:</p>	No change as per RFP
137				<p>The proposed AMD EPYC 9254 aligns closely with the Intel Xeon Gold 6542Y in terms of # of Cores, Base Clock Speed, performance metrics, ensuring a fair and unbiased evaluation of both processor OEMs.</p> <p>We respectfully request the bank to reconsider the selection of the AMD processor and allow the inclusion of the AMD EPYC 9254 for evaluation. This adjustment will promote transparency, fair competition, and a more accurate assessment of both solutions' capabilities.</p>	No change as per RFP
138	59	A. Blade Servers (Category 1) Sr. No 4 Chipset	Latest Enterprise Server Class Chipset	<p>Latest Enterprise Server Class Chipset or as required by AMD EPYC Processors.</p> <p>AMD EPYC™ CPUs feature an advanced System on Chip (SoC) architecture, which integrates all key functionalities directly within the processor, eliminating the need for a separate chipset. This innovative design delivers significant advantages for enterprise and data center environments.</p>	No change as per RFP
139	59	A. Blade Servers (Category 1) Sr. No 5 Memory	Each Server should be installed with minimum 1.5 TB Memory DDR5 with 4800 MT/s memory speed.	<p>Each Server should be installed with minimum 1.5 TB Memory DDR5 capable of sustaining an Operational/Runtime Speed of 4800 MT/s</p> <p>It is evident from the RFP, that bank have a huge requirement of Memory, in this case, Memory throughput becomes of utmost importance for the Server Performance.</p> <p>The critical importance of CPU memory bandwidth is evident from the fact that leading OEMs optimize server configurations with 1 DIMM per Channel (1DPC) for SPEC benchmarks. This approach ensures maximum utilization of memory bandwidth, as 1DPC configurations reduce latency and minimize signal interference, allowing the CPU to operate at peak performance.</p>	No change as per RFP
140				<p>When 2 Memory DIMMs per channel are populated, the memory runtime speed reduces to 4400 MT/s. Whereas, when only 1 DIMM per Memory channel is populated, there is no impact on memory speed and the memory would run at 4800 MT/s.</p> <p>Please Refer page no. 18 and 19 https://www.hpe.com/psnow/doc/a50010242enw – Server Memory population rules for Intel Xeon Servers.</p> <p>As the bank is asking for 4800 MT/s memory speed along with high memory capacity and memory scalability, any memory speed reduction during system runtime should be avoided. Memory bandwidth/throughput is directly proportionate to runtime memory speed and no. of memory channels populated. This is very important consideration from application performance perspective. Any reduction in memory throughput would have major impact on overall server performance and due to the same reason, all the SPEC benchmarks by the server OEMs are done only with 1 memory DIMM populated per memory channel.</p>	No change as per RFP
141	59	A. Blade Servers (Category 1) Sr. No 5 Memory	Should provide Advanced Memory Protection features like multi-bit error correction, memory mirroring, and memory spare for higher reliability.	<p>Should provide Advanced Memory Protection features like multi-bit error correction, memory mirroring, and memory spare</p> <p>or</p> <p>Advanced Memory Protection features like Advanced Memory Device Correction (AMDC) and post-package repair (PPR) capability for higher reliability.</p> <p>Processor OEMs employ unique methodologies and terminologies to implement memory RAS (Reliability, Availability, and Serviceability) features. AMD EPYC™ processors, in particular, offer cutting-edge technologies such as Advanced Memory Device Correction (AMDC) and Post-Package Repair (PPR) to ensure superior memory reliability and maximize system uptime.</p> <p>https://www.amd.com/content/dam/amd/en/documents/epyc-business-docs/white-papers/modernize-data-center-virtualization-with-amd-epyc-processors.pdf</p> <p>Hence request bank to incorporate the above in the said RFP</p>	Please refer Corrigendum No 3

142	59	A. Blade Servers (Category 1) Sr. No 5 Memory	Should support scalability up to 2 TB without having to replace the existing DIMMs.	<p>Should support scalability up to 2 TB without having to replace the existing DIMMs, also sustaining an Operational/Runtime Speed of 4800 MT/s Detailed Explanation as mentioned in SI. No. 3 above</p> <p>The critical importance of CPU memory bandwidth is evident from the fact that leading OEMs optimize server configurations with 1 DIMM per Channel (1DPC) for SPEC benchmarks. This approach ensures maximum utilization of memory bandwidth, as 1DPC configurations reduce latency and minimize signal interference, allowing the CPU to operate at peak performance. Benchmarks like SPEC are designed to evaluate a system's performance under optimal conditions, and memory bandwidth is a key factor in achieving high performance across demanding workloads such as Virtualization, AI/ML Processing, Database Management, and high-performance computing (HPC). Configuring 1DPC enables memory channels to deliver their full throughput, eliminating potential bottlenecks associated with 2DPC setups, which can introduce increased loading and reduce effective bandwidth.</p> <p>Alternatively, to get the best server performance, Bank should remove the said scalability clause, as populating all the Memory DIMMs available in 1DPC Configurations, yields peak performance of the server.</p>	No change as per RFP
143	62	B. Rack Servers (Category 2) Sr. No 3 Processor	64 cores each socket, dual socket-128 cores (AMD 9554 or Intel Xeon Platinum 8592V)	<p>The processors selected by the bank from Intel and AMD for this RFP appear to be not directly comparable, which could inadvertently create an imbalance in the evaluation process. The bank has specified the Intel Xeon Platinum 8592V with 64 cores at 2.0 GHz, while the selected AMD processor is the AMD EPYC 9554, which offers 64 cores at 3.1 GHz. This significant difference in base clock speed provides an inherent advantage to one processor OEM, potentially affecting the fairness of the competition.</p> <p>To ensure a level playing field and a fair comparison of equivalent processor capabilities, we propose the bank consider the AMD EPYC 9534. The AMD EPYC 9534 matches the Intel Xeon Platinum 8592V in terms of 64 cores and a better base Clock Speed of 2.45 GHz, providing a more balanced and equitable evaluation framework.</p> <p>Below is a detailed comparison of the processors for reference:</p>	No change as per RFP
144				<p>The proposed AMD EPYC 9534 aligns closely with the Intel Xeon Platinum 8592V in terms of # of Cores, Base Clock Speed, performance metrics, ensuring a fair and unbiased evaluation of both processor OEMs.</p> <p>We respectfully request the bank to reconsider the selection of the AMD processor and allow the inclusion of the AMD EPYC 9534 for evaluation. This adjustment will promote transparency, fair competition, and a more accurate assessment of both solutions' capabilities.</p>	No change as per RFP
145	62	B. Rack Servers (Category 2) Sr. No 4 Chipset	Latest Enterprise Server Class Chipset	<p>Latest Enterprise Server Class Chipset or as required by AMD EPYC Processors.</p> <p>AMD EPYC™ CPUs feature an advanced System on Chip (SoC) architecture, which integrates all key functionalities directly within the processor, eliminating the need for a separate chipset. This innovative design delivers significant advantages for enterprise and data center environments.</p>	No change as per RFP
146	62	B. Rack Servers (Category 2) Sr. No 5 Memory	Each Server should be installed with minimum 2 TB Memory DDR5 or higher RDIMM with 4800MT/s memory speed.	<p>Each Server should be installed with minimum 2 TB Memory DDR5 capable of sustaining an Operational/Runtime Speed of 4800 MT/s</p> <p>It is evident from the RFP, that bank have a huge requirement of Memory, in this case, Memory throughput becomes of utmost importance for the Server Performance.</p> <p>The critical importance of CPU memory bandwidth is evident from the fact that leading OEMs optimize server configurations with 1 DIMM per Channel (1DPC) for SPEC benchmarks. This approach ensures maximum utilization of memory bandwidth, as 1DPC configurations reduce latency and minimize signal interference, allowing the CPU to operate at peak performance.</p>	No change as per RFP

147				<p>When 2 Memory DIMMs per channel are populated, the memory runtime speed reduces to 4400 MT/s. Whereas, when only 1 DIMM per Memory channel is populated, there is no impact on memory speed and the memory would run at 4800 MT/s. Please Refer page no. 18 and 19 https://www.hpe.com/psnow/doc/a50010242enw – Server Memory population rules for Intel Xeon Servers.</p> <p>As the bank is asking for 4800 MT/s memory speed along with high memory capacity and memory scalability, any memory speed reduction during system runtime should be avoided. Memory bandwidth/throughput is directly proportionate to runtime memory speed and no. of memory channels populated. This is very important consideration from application performance perspective. Any reduction in memory throughput would have major impact on overall server performance and due to the same reason, all the SPEC benchmarks by the server OEMs are done only with 1 memory DIMM populated per memory channel.</p> <p>The industry is also moving from 8 memory channels to 12 memory channels to increase the memory bandwidth for in line with today and tomorrow's higher application performance needs. E.g. AMD, from it's 4th generation processor onwards and Intel with their 6th generation processors offer 12 Memory channels per processor. Most server OEMs offer only 1 DPC (24 DIMM slots) for dual processors systems to avoid any memory bandwidth reduction and lower application performance penalty related to 2 DPC.</p> <p>Hence request bank to incorporate the above in the said RFP</p> <p>Hence the bank should ask for Operational/Runtime memory speed of 4800 MT/s for best server performance and for getting highest possible memory bandwidth. As well as even after the memory upgrade the entire server memory should run at minimum 4800 MT/s.</p>	No change as per RFP
148	62	B. Rack Servers (Category 2) Sr. No 5 Memory	Should provide Advanced Memory Protection features like multi-bit error correction, memory mirroring, and memory spare for higher reliability.	<p>Should provide Advanced Memory Protection features like multi-bit error correction, memory mirroring, and memory spare or Advanced Memory Protection features like Advanced Memory Device Correction (AMDC) and post-package repair (PPR) capability for higher reliability.</p> <p>Processor OEMs employ unique methodologies and terminologies to implement memory RAS (Reliability, Availability, and Serviceability) features. AMD EPYC™ processors, in particular, offer cutting-edge technologies such as Advanced Memory Device Correction (AMDC) and Post-Package Repair (PPR) to ensure superior memory reliability and maximize system uptime.</p> <p>Hence request bank to incorporate the above in the said RFP</p>	Please refer Corrigendum No 3
149	63	C. Rack Servers (Category 3) Sr. No 5 Memory	Should support scalability up to 4 TB without having to replace the existing DIMMs	<p>Should support scalability up to 4 TB without having to replace the existing DIMMs, also sustaining an Operational/Runtime Speed of 4800 MT/s</p> <p>Detailed Explanation as mentioned in Sl. No. 13 above</p> <p>The critical importance of CPU memory bandwidth is evident from the fact that leading OEMs optimize server configurations with 1 DIMM per Channel (1DPC) for SPEC benchmarks. This approach ensures maximum utilization of memory bandwidth, as 1DPC configurations reduce latency and minimize signal interference, allowing the CPU to operate at peak performance. Benchmarks like SPEC are designed to evaluate a system's performance under optimal conditions, and memory bandwidth is a key factor in achieving high performance across demanding workloads such as Virtualization, AI/ML Processing, Database Management, and high-performance computing (HPC). Configuring 1DPC enables memory channels to deliver their full throughput, eliminating potential bottlenecks associated with 2DPC setups, which can introduce increased loading and reduce effective bandwidth.</p> <p>Alternatively, to get the best server performance, Bank should remove the said scalability clause, as populating all the Memory DIMMs available in 1DPC Configurations, yields peak performance of the server.</p>	No change as per RFP
150	64	D. Rack Servers (Category 4) Sr. No 3 Processor	AMD 9274F or Intel Xeon Gold 6542Y (24 cores each socket, dual socket-48 cores)	<p>The processors selected by the bank from Intel and AMD for this RFP appear to be not directly comparable, which could inadvertently create an imbalance in the evaluation process. The bank has specified the Intel Xeon Gold 6542Y with 24 cores at 2.9 GHz, while the selected AMD processor is the AMD EPYC 9274F, which offers 24 cores at 4.05 GHz. This significant difference in base clock speed provides an inherent advantage to one processor OEM, potentially affecting the fairness of the competition.</p> <p>To ensure a level playing field and a fair comparison of equivalent processor capabilities, we propose the bank consider the AMD EPYC 9254. The AMD EPYC 9254 matches the Intel Xeon Gold 6542Y in terms of 24 cores and a 2.9 GHz base clock, providing a more balanced and equitable evaluation framework.</p>	No change as per RFP

151				<p>The proposed AMD EPYC 9254 aligns closely with the Intel Xeon Gold 6542Y in terms of # of Cores, Base Clock Speed, performance metrics, ensuring a fair and unbiased evaluation of both processor OEMs.</p> <p>We respectfully request the bank to reconsider the selection of the AMD processor and allow the inclusion of the AMD EPYC 9254 for evaluation. This adjustment will promote transparency, fair competition, and a more accurate assessment of both solutions' capabilities.</p>	No change as per RFP
152	64	D. Rack Servers (Category 4) Sr. No 4 Chipset	Latest Enterprise Server Class Chipset	<p>Latest Enterprise Server Class Chipset or as required by AMD EPYC Processors.</p> <p>AMD EPYC™ CPUs feature an advanced System on Chip (SoC) architecture, which integrates all key functionalities directly within the processor, eliminating the need for a separate chipset. This innovative design delivers significant advantages for enterprise and data center environments.</p>	No change as per RFP
153	64	D. Rack Servers (Category 4) Sr. No 5 Memory	Each Server should be installed with minimum 1.5 TB Memory DDR5 or higher RDIMM	<p>Each Server should be installed with minimum 1.5 TB Memory DDR5 capable of sustaining an Operational/Runtime Speed of 4800 MT/s</p> <p>It is evident from the RFP, that bank have a huge requirement of Memory, in this case, Memory throughput becomes of utmost importance for the Server Performance.</p> <p>The critical importance of CPU memory bandwidth is evident from the fact that leading OEMs optimize server configurations with 1 DIMM per Channel (1DPC) for SPEC benchmarks. This approach ensures maximum utilization of memory bandwidth, as 1DPC configurations reduce latency and minimize signal interference, allowing the CPU to operate at peak performance.</p>	No change as per RFP
154				<p>When 2 Memory DIMMs per channel are populated, the memory runtime speed reduces to 4400 MT/s. Whereas, when only 1 DIMM per Memory channel is populated, there is no impact on memory speed and the memory would run at 4800 MT/s.</p> <p>Please Refer page no. 18 and 19 https://www.hpe.com/psnow/doc/a50010242enw – Server Memory population rules for Intel Xeon Servers.</p> <p>As the bank is asking for 4800 MT/s memory speed along with high memory capacity and memory scalability, any memory speed reduction during system runtime should be avoided. Memory bandwidth/throughput is directly proportionate to runtime memory speed and no. of memory channels populated. This is very important consideration from application performance perspective. Any reduction in memory throughput would have major impact on overall server performance and due to the same reason, all the SPEC benchmarks by the server OEMs are done only with 1 memory DIMM populated per memory channel.</p> <p>The industry is also moving from 8 memory channels to 12 memory channels to increase the memory bandwidth for in line with today and tomorrow's higher application performance needs. E.g. AMD, from it's 4th generation processor onwards and Intel with their 6th generation processors offer 12 Memory channels per processor. Most server OEMs offer only 1 DPC (24 DIMM slots) for dual processors systems to avoid any memory bandwidth reduction and lower application performance penalty related to 2 DPC.</p> <p>Hence request bank to incorporate the above in the said RFP</p> <p>Hence the bank should ask for Operational/Runtime memory speed of 4800 MT/s for best server performance and for getting highest possible memory bandwidth. As well as even after the memory upgrade the entire server memory should run at minimum 4800 MT/s.</p>	No change as per RFP
155	64	D. Rack Servers (Category 4) Sr. No 5 Memory	Should provide Advanced Memory Protection features like multi-bit error correction, memory mirroring, and memory spare for higher reliability.	<p>Should provide Advanced Memory Protection features like multi-bit error correction, memory mirroring, and memory spare or Advanced Memory Protection features like Advanced Memory Device Correction (AMDC) and post-package repair (PPR) capability for higher reliability.</p> <p>Processor OEMs employ unique methodologies and terminologies to implement memory RAS (Reliability, Availability, and Serviceability) features. AMD EPYC™ processors, in particular, offer cutting-edge technologies such as Advanced Memory Device Correction (AMDC) and Post-Package Repair (PPR) to ensure superior memory reliability and maximize system uptime.</p> <p>Hence request bank to incorporate the above in the said RFP</p>	Please refer Corrigendum No 3

156	64	D. Rack Servers (Category 4) Sr. No 5 Memory	Should support scalability up to 2 TB without having to replace the existing DIMMs.	<p>Should support scalability up to 2 TB without having to replace the existing DIMMs, also sustaining an Operational/Runtime Speed of 4800 MT/s Detailed Explanation as mentioned in Sl. No. 18 above</p> <p>The critical importance of CPU memory bandwidth is evident from the fact that leading OEMs optimize server configurations with 1 DIMM per Channel (1DPC) for SPEC benchmarks. This approach ensures maximum utilization of memory bandwidth, as 1DPC configurations reduce latency and minimize signal interference, allowing the CPU to operate at peak performance. Benchmarks like SPEC are designed to evaluate a system's performance under optimal conditions, and memory bandwidth is a key factor in achieving high performance across demanding workloads such as Virtualization, AI/ML Processing, Database Management, and high-performance computing (HPC). Configuring 1DPC enables memory channels to deliver their full throughput, eliminating potential bottlenecks associated with 2DPC setups, which can introduce increased loading and reduce effective bandwidth.</p> <p>Alternatively, to get the best server performance, Bank should remove the said scalability clause, as populating all the Memory DIMMs available in 1DPC Configurations, yields peak performance of the server.</p>	No change as per RFP
157	65	E. GPU Rack Servers Type 1 (Category 5) Sr. No 3 Server Type	64 Cores each socket, Dual socket- 128 cores) AMD 9554 or Intel Xeon Platinum 8592V	<p>The processors selected by the bank from Intel and AMD for this RFP appear to be not directly comparable, which could inadvertently create an imbalance in the evaluation process. The bank has specified the Intel Xeon Platinum 8592V with 64 cores at 2.0 GHz, while the selected AMD processor is the AMD EPYC 9554, which offers 64 cores at 3.1 GHz. This significant difference in base clock speed provides an inherent advantage to one processor OEM, potentially affecting the fairness of the competition.</p> <p>To ensure a level playing field and a fair comparison of equivalent processor capabilities, we propose the bank consider the AMD EPYC 9534. The AMD EPYC 9534 matches the Intel Xeon Platinum 8592V in terms of 64 cores and a better base Clock Speed of 2.45 GHz, providing a more balanced and equitable evaluation framework.</p> <p>Below is a detailed comparison of the processors for reference:</p>	No change as per RFP
158				<p>The proposed AMD EPYC 9534 aligns closely with the Intel Xeon Platinum 8592V in terms of # of Cores, Base Clock Speed, performance metrics, ensuring a fair and unbiased evaluation of both processor OEMs.</p> <p>We respectfully request the bank to reconsider the selection of the AMD processor and allow the inclusion of the AMD EPYC 9534 for evaluation. This adjustment will promote transparency, fair competition, and a more accurate assessment of both solutions' capabilities.</p>	No change as per RFP
159	65	E. GPU Rack Servers Type 1 (Category 5) Sr. No 4 GPU	<p>Server should support three NVIDIA H100 Tensor Core GPU card or Intel® Data Center GPU Max 1100 or AMD Instinct™ MI250X. On day one server should be populated with two cards and all required licenses (including Nvidia AI enterprise / equivalent) for GPU slicing with VMware, Nvidia enterprise tech support, training, inferencing etc.</p>	<p>We request bank to consider OAM based GPUs as it offers several advantages to the Bank in Data Center and AI/ML environments. These advantages stem from the design and integration benefits of the OAM form factor, which is developed to address scalability, efficiency, and density challenges specifically deployed in enterprise space like SBI.</p> <p>1. Higher Compute Density Compact Form Factor: OAM GPUs are designed for higher compute density, allowing more GPUs to be installed in a single server or chassis compared to PCIe GPUs. Optimized for Modular Systems: OAM GPUs fit into modular trays or sleds, enabling tighter integration of multiple GPUs in a smaller footprint.</p> <p>2. Superior Cooling and Thermal Management Efficient Cooling Design: OAM GPUs are optimized for better thermal dissipation through direct-to-chip cooling solutions, such as liquid cooling or vapor chamber technology, which are more effective than traditional fan-based cooling for PCIe GPUs. Reduced Thermal Bottlenecks: The design ensures uniform cooling across multiple GPUs, leading to sustained performance even under heavy workloads.</p> <p>3. High-Speed Interconnects for GPU-to-GPU Communication AMD Infinity Fabric or NVLink and Similar Technologies: OAM GPUs typically support advanced interconnects (e.g., NVIDIA NVLink or AMD Infinity Fabric) that enable ultra-high-bandwidth, low-latency communication between GPUs. Reduced PCIe Bottlenecks: Unlike PCIe GPUs that rely on shared PCIe lanes for communication, OAM GPUs allow direct peer-to-peer communication, significantly improving performance for distributed workloads.</p>	No change as per RFP

160				<p>4. Enhanced Scalability Multi-GPU Configurations: OAM modules are designed for scalability, enabling seamless integration of multiple GPUs in configurations like 8-GPU, or more, within a single chassis. Cluster Efficiency: These GPUs are ideal for large-scale AI/ML clusters or HPC systems where inter-GPU communication and scaling are critical.</p> <p>5. Unified Design for Standardization Open Compute Project (OCP): OAM is part of the Open Compute Project, promoting standardization across hardware vendors, leading to broader compatibility and ecosystem support. Vendor Agnosticism: By adhering to a common design standard, OAM GPUs enable data center operators to mix and match hardware from different manufacturers without proprietary constraints.</p> <p>6. Reduced Latency and Improved Bandwidth Direct Integration: OAM GPUs often use advanced high-bandwidth memory (HBM) and direct memory access (DMA) technologies to improve data transfer rates and reduce latency compared to PCIe-based systems. Optimized for AI/ML and HPC: This advantage is particularly significant for workloads requiring large-scale parallelism, such as neural network training or scientific simulations.</p> <p>7. Better Integration with Advanced Cooling Systems Liquid Cooling Compatibility: OAM GPUs are designed with features like direct cooling plates or heat spreaders that facilitate advanced cooling solutions, enabling higher sustained performance. Efficient Rack Cooling: In data centers, OAM systems simplify the deployment of shared cooling infrastructure, reducing operating costs and improving sustainability.</p>	No change as per RFP
161				<p>8. Tailored for AI and ML Workloads Purpose-Built for Accelerators: OAM GPUs are specifically designed for workloads that demand high parallelism, such as deep learning, AI inference, and simulations, providing higher efficiency and performance than PCIe GPUs. Customized Configurations: The flexibility of OAM systems allows for bespoke designs tailored to specific workloads, optimizing resource utilization.</p> <p>9. Enhanced Reliability and Durability Enterprise-Grade Design: OAM GPUs are built for data center environments, offering enhanced durability, longer lifecycle reliability, and better resilience to high workloads. Redundancy and Failover: The design supports features like GPU redundancy for failover, ensuring high availability in mission-critical systems. Key Use Cases for OAM GPUs AI/ML Training: Large-scale neural network training where inter-GPU communication is crucial. Data Centers: High-density GPU deployments for cloud service providers and enterprise data centers.</p> <p>Each Node must be configured with Nvidia 8 x H200 141 GB GPUs (SXM) connected via Nvidia Nvlink with NV Switch or AMD Instinct 8 x MI300X 192 GB GPUs (OAM) connected via AMD Infinity Fabric. The total aggregate memory per node from the GPUs should be at least 1128 GB. Solution should come with all required licenses (including Nvidia AI enterprise / equivalent) for GPU slicing with VMware, Nvidia enterprise tech support, training, Inferencing etc.</p> <p>Bank may reduce the # of Servers required and consolidate the GPU Rack Server requirements, thus reducing the Rack Space, Power Requirement, etc.</p>	No change as per RFP
162	65	E. GPU Rack Servers Type 1 (Category 5) Sr. No 9 Chipset	Latest Enterprise Server Class Chipset	<p>Latest Enterprise Server Class Chipset or as required by AMD EPYC Processors. AMD EPYC™ CPUs feature an advanced System on Chip (SoC) architecture, which integrates all key functionalities directly within the processor, eliminating the need for a separate chipset. This innovative design delivers significant advantages for enterprise and data center environments.</p>	No change as per RFP
163	65	E. GPU Rack Servers Type 1 (Category 5) Sr. No 10 Memory	Each Server should be installed with minimum 2 TB Memory DDR5 or higher RDIMM	<p>Each Server should be installed with minimum 2 TB Memory DDR5 capable of sustaining an Operational/Runtime Speed of 4800 MT/s It is evident from the RFP, that bank have a huge requirement of Memory, in this case, Memory throughput becomes of utmost importance for the Server Performance. The critical importance of CPU memory bandwidth is evident from the fact that leading OEMs optimize server configurations with 1 DIMM per Channel (1DPC) for SPEC benchmarks. This approach ensures maximum utilization of memory bandwidth, as 1DPC configurations reduce latency and minimize signal interference, allowing the CPU to operate at peak performance.</p>	No change as per RFP

164				<p>When 2 Memory DIMMs per channel are populated, the memory runtime speed reduces to 4400 MT/s. Whereas, when only 1 DIMM per Memory channel is populated, there is no impact on memory speed and the memory would run at 4800 MT/s. Please Refer page no. 18 and 19 https://www.hpe.com/psnow/doc/a50010242enw – Server Memory population rules for Intel Xeon Servers.</p> <p>As the bank is asking for 4800 MT/s memory speed along with high memory capacity and memory scalability, any memory speed reduction during system runtime should be avoided. Memory bandwidth/throughput is directly proportionate to runtime memory speed and no. of memory channels populated. This is very important consideration from application performance perspective. Any reduction in memory throughput would have major impact on overall server performance and due to the same reason, all the SPEC benchmarks by the server OEMs are done only with 1 memory DIMM populated per memory channel.</p> <p>The industry is also moving from 8 memory channels to 12 memory channels to increase the memory bandwidth for in line with today and tomorrow's higher application performance needs. E.g. AMD, from its 4th generation processor onwards and Intel with their 6th generation processors offer 12 Memory channels per processor. Most server OEMs offer only 1 DPC (24 DIMM slots) for dual processors systems to avoid any memory bandwidth reduction and lower application performance penalty related to 2 DPC.</p> <p>Hence request bank to incorporate the above in the said RFP</p> <p>Hence the bank should ask for Operational/Runtime memory speed of 4800 MT/s for best server performance and for getting highest possible memory bandwidth. As well as even after the memory upgrade the entire server memory should run at minimum 4800 MT/s.</p>	No change as per RFP
165	65	E. GPU Rack Servers Type 1 (Category 5) Sr. No 10 Memory	Should provide Advanced Memory Protection features like multi-bit error correction, memory mirroring, and memory spare for higher reliability.	<p>Should provide Advanced Memory Protection features like multi-bit error correction, memory mirroring, and memory spare or Advanced Memory Protection features like Advanced Memory Device Correction (AMDC) and post-package repair (PPR) capability for higher reliability.</p> <p>Processor OEMs employ unique methodologies and terminologies to implement memory RAS (Reliability, Availability, and Serviceability) features. AMD EPYC™ processors, in particular, offer cutting-edge technologies such as Advanced Memory Device Correction (AMDC) and Post-Package Repair (PPR) to ensure superior memory reliability and maximize system uptime.</p> <p>Hence request bank to incorporate the above in the said RFP</p>	Please refer Corrigendum No 3
166	65	E. GPU Rack Servers Type 1 (Category 5) Sr. No 10 Memory	Should support scalability up to 4 TB without having to replace the existing DIMMs	<p>Should support scalability up to 4 TB without having to replace the existing DIMMs, also sustaining an Operational/Runtime Speed of 4800 MT/s or</p> <p>Should support scalability up to 3 TB without having to replace the existing DIMMs, also sustaining an Operational/Runtime Speed of 4800 MT/s Detailed Explanation as mentioned in Sl. No. 8 above</p> <p>The critical importance of CPU memory bandwidth is evident from the fact that leading OEMs optimize server configurations with 1 DIMM per Channel (1DPC) for SPEC benchmarks. This approach ensures maximum utilization of memory bandwidth, as 1DPC configurations reduce latency and minimize signal interference, allowing the CPU to operate at peak performance. Benchmarks like SPEC are designed to evaluate a system's performance under optimal conditions, and memory bandwidth is a key factor in achieving high performance across demanding workloads such as Virtualization, AI/ML Processing, Database Management, and high-performance computing (HPC). Configuring 1DPC enables memory channels to deliver their full throughput, eliminating potential bottlenecks associated with 2DPC setups, which can introduce increased loading and reduce effective bandwidth.</p> <p>Alternatively, to get the best server performance, Bank should remove the said scalability clause, as populating all the Memory DIMMs available in 1DPC Configurations, yields peak performance of the server.</p>	No change as per RFP
167	66	F. GPU Rack Servers Type 2 (Category 6) Sr. No 3 Server Type	64 Cores each socket, Dual socket- 128 cores) AMD 9554 or Intel Xeon Platinum 8592V	<p>The processors selected by the bank from Intel and AMD for this RFP appear to be not directly comparable, which could inadvertently create an imbalance in the evaluation process. The bank has specified the Intel Xeon Platinum 8592V with 64 cores at 2.0 GHz, while the selected AMD processor is the AMD EPYC 9554, which offers 64 cores at 3.1 GHz. This significant difference in base clock speed provides an inherent advantage to one processor OEM, potentially affecting the fairness of the competition.</p> <p>To ensure a level playing field and a fair comparison of equivalent processor capabilities, we propose the bank consider the AMD EPYC 9534. The AMD EPYC 9534 matches the Intel Xeon Platinum 8592V in terms of 64 cores and a better base Clock Speed of 2.45 GHz, providing a more balanced and equitable evaluation framework.</p> <p>Below is a detailed comparison of the processors for reference:</p>	No change as per RFP

168				<p>The proposed AMD EPYC 9534 aligns closely with the Intel Xeon Platinum 8592V in terms of # of Cores, Base Clock Speed, performance metrics, ensuring a fair and unbiased evaluation of both processor OEMs.</p> <p>We respectfully request the bank to reconsider the selection of the AMD processor and allow the inclusion of the AMD EPYC 9534 for evaluation. This adjustment will promote transparency, fair competition, and a more accurate assessment of both solutions' capabilities.</p>	No change as per RFP
169	66	F. GPU Rack Servers Type 2 (Category 6) Sr. No 4 GPU	<p>Each Node must be configured with Nvidia 8 x H200 141 GB GPUs (SXM) connected via Nvidia Nvlink with NV Switch. The total aggregate memory per node from the GPUs should be at least 1128 GB. Solution should come with all required licenses (including Nvidia AI enterprise / equivalent) for GPU slicing with VMware, Nvidia enterprise tech support, training, Inference etc.</p>	<p>Each Node must be configured with Nvidia 8 x H200 141 GB GPUs (SXM) connected via Nvidia Nvlink with NV Switch or AMD Instinct 8 x MI300X 192 GB GPUs (OAM) connected via AMD Infinity Fabric.</p> <p>The total aggregate memory per node from the GPUs should be at least 1128 GB. Solution should come with all required licenses (including Nvidia AI enterprise / equivalent) for GPU slicing with VMware, Nvidia enterprise tech support, training, Inference etc.</p> <p>1. Diverse Vendor Ecosystem Avoid Vendor Lock-In: Including AMD Instinct GPUs ensures you maintain flexibility and avoid over-reliance on a single vendor like NVIDIA, fostering a competitive and balanced procurement process. Broad Market Options: Encouraging multiple vendors can lead to better pricing, innovation, and support options.</p> <p>2. Price-to-Performance Advantage Cost Efficiency: AMD Instinct GPUs often deliver comparable performance at a lower price point, especially in compute-heavy workloads like AI/ML training and scientific simulations. Total Cost of Ownership (TCO): Lower initial hardware costs, coupled with efficient power usage, can reduce long-term operational costs.</p> <p>3. Open Ecosystem with ROCm ROCm (Radeon Open Compute): AMD's ROCm platform provides an open-source software stack for GPU computing, reducing licensing costs and enhancing customization opportunities. Support for Common Frameworks: ROCm supports popular frameworks like TensorFlow, PyTorch, and ONNX, ensuring compatibility with existing AI/ML workflows.</p>	No change as per RFP
170				<p>4. Industry Leadership in Memory Bandwidth High Bandwidth Memory (HBM2e): AMD Instinct GPUs typically offer higher memory bandwidth compared to equivalent NVIDIA GPUs, which can be critical for data-intensive workloads. Memory Capacity: Larger memory pools in AMD GPUs benefit applications with high dataset demands.</p> <p>5. Superior Interconnect Technologies Infinity Fabric: AMD's Infinity Fabric interconnects allow for seamless scaling of GPUs, enabling high-speed communication between nodes in multi-GPU environments. Efficient Scaling: This is particularly advantageous in data centers where scaling performance is essential for workload efficiency.</p> <p>6. Enhanced Support for Virtualized and Cloud Workloads Optimized for Virtualization: AMD Instinct GPUs are designed to excel in virtualized environments, which is increasingly critical for cloud computing and edge deployment scenarios. Energy Efficiency: Better energy efficiency contributes to greener, more sustainable data center operations.</p> <p>7. Industry Adoption and Proven Use Cases Adoption in HPC Centers: Leading HPC installations, such as the Frontier supercomputer (world's first exascale system), leverage AMD Instinct GPUs, proving their capability in the most demanding environments. AI and ML: Many enterprises and research institutions are adopting AMD GPUs for AI/ML applications, providing a solid track record.</p> <p>8. Stimulating Competition for Better Innovation Including AMD Instinct GPUs in the RFP encourages innovation across vendors, ensuring bank receive the most advanced, future-ready solutions at competitive pricing.</p> <p>MI300X Datasheet - https://www.amd.com/content/dam/amd/en/documents/instinct-tech-docs/data-sheets/amd-instinct-mi300x-data-sheet.pdf MI300X Platform Datasheet - https://www.amd.com/content/dam/amd/en/documents/instinct-tech-docs/data-sheets/amd-instinct-mi300x-platform-data-sheet.pdf</p>	No change as per RFP
171	66	F. GPU Rack Servers Type 2 (Category 6) Sr. No 9 Chipset	Latest Enterprise Server Class Chipset	<p>Latest Enterprise Server Class Chipset or as required by AMD EPYC Processors.</p> <p>AMD EPYC™ CPUs feature an advanced System on Chip (SoC) architecture, which integrates all key functionalities directly within the processor, eliminating the need for a separate chipset.</p> <p>This innovative design delivers significant advantages for enterprise and data center environments.</p>	No change as per RFP

172	66	F. GPU Rack Servers Type 2 (Category 6) Sr. No 10 Memory	Each Server should be installed with minimum 3 TB Memory DDR5 or higher RDIMM	<p>Each Server should be installed with minimum 3 TB Memory DDR5 capable of sustaining an Operational/Runtime Speed of 4800 MT/s</p> <p>It is evident from the RFP, that bank have a huge requirement of Memory, in this case, Memory throughput becomes of utmost importance for the Server Performance.</p> <p>The critical importance of CPU memory bandwidth is evident from the fact that leading OEMs optimize server configurations with 1 DIMM per Channel (1DPC) for SPEC benchmarks. This approach ensures maximum utilization of memory bandwidth, as 1DPC configurations reduce latency and minimize signal interference, allowing the CPU to operate at peak performance.</p>	No change as per RFP
173				<p>When 2 Memory DIMMs per channel are populated, the memory runtime speed reduces to 4400 MT/s. Whereas, when only 1 DIMM per Memory channel is populated, there is no impact on memory speed and the memory would run at 4800 MT/s.</p> <p>Please Refer page no. 18 and 19 https://www.hpe.com/psnow/doc/a50010242enw – Server Memory population rules for Intel Xeon Servers.</p> <p>As the bank is asking for 4800 MT/s memory speed along with high memory capacity and memory scalability, any memory speed reduction during system runtime should be avoided. Memory bandwidth/throughput is directly proportionate to runtime memory speed and no. of memory channels populated. This is very important consideration from application performance perspective. Any reduction in memory throughput would have major impact on overall server performance and due to the same reason, all the SPEC benchmarks by the server OEMs are done only with 1 memory DIMM populated per memory channel.</p> <p>The industry is also moving from 8 memory channels to 12 memory channels to increase the memory bandwidth for in line with today and tomorrow's higher application performance needs. E.g. AMD, from it's 4th generation processor onwards and Intel with their 6th generation processors offer 12 Memory channels per processor. Most server OEMs offer only 1 DPC (24 DIMM slots) for dual processors systems to avoid any memory bandwidth reduction and lower application performance penalty related to 2 DPC.</p> <p>Hence request bank to incorporate the above in the said RFP</p> <p>Hence the bank should ask for Operational/Runtime memory speed of 4800 MT/s for best server performance and for getting highest possible memory bandwidth. As well as even after the memory upgrade the entire server memory should run at minimum 4800 MT/s.</p>	No change as per RFP

174	66	F. GPU Rack Servers Type 2 (Category 6) Sr. No 10 Memory	Should provide Advanced Memory Protection features like multi-bit error correction, memory mirroring, and memory spare for higher reliability.	Should provide Advanced Memory Protection features like multi-bit error correction, memory mirroring, and memory spare or Advanced Memory Protection features like Advanced Memory Device Correction (AMDC) and post-package repair (PPR) capability for higher reliability. Processor OEMs employ unique methodologies and terminologies to implement memory RAS (Reliability, Availability, and Serviceability) features. AMD EPYC™ processors, in particular, offer cutting-edge technologies such as Advanced Memory Device Correction (AMDC) and Post-Package Repair (PPR) to ensure superior memory reliability and maximize system uptime. Hence request bank to incorporate the above in the said RFP	Please refer Corrigendum No 3
175	66	F. GPU Rack Servers Type 2 (Category 6) Sr. No 10 Memory	Should support scalability up to 4 TB without having to replace the existing DIMMs	Should support scalability up to 4 TB without having to replace the existing DIMMs, also sustaining an Operational/Runtime Speed of 4800 MT/s or Should support scalability up to 3 TB without having to replace the existing DIMMs, also sustaining an Operational/Runtime Speed of 4800 MT/s Detailed Explanation as mentioned in Sl. No. 8 above The critical importance of CPU memory bandwidth is evident from the fact that leading OEMs optimize server configurations with 1 DIMM per Channel (1DPC) for SPEC benchmarks. This approach ensures maximum utilization of memory bandwidth, as 1DPC configurations reduce latency and minimize signal interference, allowing the CPU to operate at peak performance. Benchmarks like SPEC are designed to evaluate a system's performance under optimal conditions, and memory bandwidth is a key factor in achieving high performance across demanding workloads such as Virtualization, AI/ML Processing, Database Management, and high-performance computing (HPC). Configuring 1DPC enables memory channels to deliver their full throughput, eliminating potential bottlenecks associated with 2DPC setups, which can introduce increased loading and reduce effective bandwidth. Alternatively, to get the best server performance, Bank should remove the said scalability clause, as populating all the Memory DIMMs available in 1DPC Configurations, yields peak performance of the server.	No change as per RFP
176	79	Third-Party Components -Clause 3	Should be integrated with Netapp Storage and any other brand storage and VMware softwares	NetAPP Storage Controller and 3rd party Storage interface connectivity is FC or ISCSI	No change as per RFP Both
177	80	Inegration / Migration Requirements with existing systems Clause -6	Yes, New hardware should be integrated with old platform without any additional cost to the Bank	Integration of New hardware with existing which old platform is in use(VMware Cluster or vSAN environment)	No change as per RFP Both
178	29	28. SERVICES Clause vi	Service Provider shall provide maintenance support for the Product including embedded software/ OS/ middleware etc over the entire period of Contract.	Need more clarity for OS maintenance support level	No change as per RFP Maintenance support required for Internal embedded software of physical servers
179	31	28. SERVICES Clause vii	In the event of system break down or failures at any stage, protection available, which would include the following, shall be specified. B Protection of data/ Configuration	Need Clarity protection of APPs data or system configuration data	No change as per RFP Protection required for system configuration data.
180	49	EXEMPTION OF EMD		Bid Submison for the subjective RFP is on GEM Portal. & As per GEM GTC - Bidder/vendor who have Turnover more than 500 crores they have examined to submite EMD on GEM portal. Request you to kindly confirm the same.	No change, as per RFP
181	110	RESPONSIBILITIES OF SERVICE PROVIDER, 6.4	Service Provider shall report the incidents, including cyber incidents and those resulting in disruption of service and data loss/ leakage immediately but not later than one hour of detection.	Please let us know the resolution time line for the same.	No change, as per RFP For resolution time refer RFP
182	59	A. Blade Servers (Category 1)- 2- Quantity	1775 (901 at Mumbai and 874 at Gachibowli)	Though RFP has called out number of blade servers required per location, can you also confirm number of blade chassis that will be required per location. As all OEMs has differenent number of blade servers supported per chassis, it will be important to mention the number of chassis required.	No change, as per RFP Each OEM may have different blade count per chassis.
183		Site readiness	Sonata Dependencises	1) Assurance of Site readiness at DC/DR Or Near DR. 2) Network Switthed should be in active state. 3) Confirmation of Cabeling readiness	No change, as per RFP
184	96	Appendix-L - Point No. 3	Delivery of all equipment should be within 6 weeks from date of placing of the order. In the event of the any or all equipment(s) not being delivered, installed, tested and commissioned within a period of 10 weeks from date of Purchase Order, a penalty of 0.5 percent of the total cost of equipment for each week or part thereof the delay, subject to maximum amount of ten (10) percent of the total cost of equipment will be charged to vendor.	Delivery of all equipment should be within 14 weeks from date of placing of the order. In the event of the any or all equipment(s) not being delivered, installed, tested and commissioned within a period of 18 weeks from date of Purchase Order, a penalty of 0.5 percent of the total cost of equipment for each week or part thereof the delay, subject to maximum amount of ten (10) percent of the total cost of equipment will be charged to vendor.	No change, as per RFP

185	86	Appendix-E- Scope of Work and Payment Schedule-4	Delivery of total hardware and software should be completed within 6 weeks.	Delivery of total hardware and software should be completed within 14 weeks.	No change, as per RFP
186	59	Appendix-C Technical & Functional Specifications Blade Servers(Category 1) Point 8 - Host Bus Adaptor(HBA)/Converged Ethernet Adaptor(CNA)	Each blade server must be equipped with at least two 32G dedicated Fibre Channel SAN storage connectivity ports, with redundancy, supporting both Fibre Channel (FC) and NVMe over Fabrics (NVMeoF) protocols		
187	70	Appendix-C Technical & Functional Specifications G. Hardware-Blade Enclosure 4- Blade Server Interconnect	The enclosure must have redundant Fibre Channel Interconnect modules with minimum 8 x 32Gbps Uplink Ports to the SAN Switch i.e. 4 x 32G uplink bandwidth must be maintained even after one interconnect failure. Each module should be fully licensed to use all available ports in case of both Multi chassis and Standalone chassis scenario. All required licenses to ensure interoperability with Bank's existing fabric must be provided	We seek clarity from bank for the following:- Our blade technology supports PCIe FC HBA Installation within the offered blade server. However our blade enclosure technology is more tuned for high speed and capacity network framework. So we would be able offer the SAN connectivity modules outside of the chassis meeting the downlink and uplink requirements stated in the RFP. We would also like to seek clarity on which FC switches bank is using. We need this clarity for checking the interoperability.	No change, as per RFP As on date, Cisco Switches are in use.
188	59	Appendix-C Technical & Functional Specifications Blade Servers(Category 1) Point 3 - Processor	AMD 9274F or Intel Xeon Gold 6542Y (24 cores each socket, dual socket-48 cores)	We seek clarity on Intel Gold 6542Y - The requested each CPU is 250W TDP. Considering the blades which is more dense compute architecture will have its limits on Heat dissipation. So we request to allow Intel Gold 5520+ which has higher CPU cores and lower TDP of 205 W	No change, as per RFP