INDIAN INSTITUTE OF TECHNOLOGY BOMBAY
MATERIALS MANAGEMENT DIVISION
Powai, Mumbai 400076.

Ref. PR No. 1000053432 Rfx. No. 6100002732

Item Description: 1. HFT Low-Latency Compute Nodes Qty: 2
2. Banking / Markets Analytics Nodes Qty: 2
3. Shared Infrastructure / Virtualization N Qty: 1
4. High-Speed Network Fabric Qty: 1

Sr. No	Item Description	Detailed Technical Specification	Technical Compliance (Yes / No)	Additional Information (if any)
1.	HFT Low-Latency Compute Nodes	Quantity: 2 Servers Type: 2U Dual Socket Performance Servers Per Server Configuration: • CPU: • 2 × AMD EPYC 9354 / 9454 (32-48 cores each, base clock ≥ 3.2 GHz) or substantially equivalent. • Memory: • 384 GB DDR5 RAM (12 × 32 GB, full memory channel population) • Storage: • 2 × 1.92 TB NVMe SSD (Enterprise grade, PCIe Gen4, TLC NAND, Mixed Use) • Network: • 2 × 25 Gbps SFP28 NIC (Mellanox ConnectX-6 or equivalent) • DPDK / RDMA / Kernel Bypass capable • GPU: • 1 × NVIDIA L40S (48 GB) or substantially equivalent •FPGA:		

		• 1x Xilinx Alveo U50 FPGA or substantially equivalent • PCIe: • Minimum 4 × PCIe Gen4 x16 slots • Power: • Dual redundant hot-swappable PSUs Use Case: Ultra-low latency market data pipelines, kernel-bypass networking, tick-to-trade simulation, GPU inference.		
2.	Banking / Markets Analytics Nodes	Quantity: 2 Server Type: 2U Dual Socket High Core Servers Per Server Configuration: • CPU: • 2 × AMD EPYC 9554 / 9654 (64–96 cores each, , base clock ≥ 3.1 GHz) or substantially equivalent. • Memory: • 512 GB DDR5 RAM (16 × 32 GB) • Storage: • 8 × 7.68 TB Enterprise NVMe SSD (TLC NAND, PCIe Gen4, Read Intensive or Mixed Use) • GPU: • 1 × NVIDIA L40S (48 GB) per node or substantially equivalent • Network: • 2 × 25 Gbps SFP28 NIC • PCIe: • Minimum 6 × PCIe Gen4 x16 slots • Power: • Dual redundant hot-swappable PSUs Use Case: Backtesting, econometrics, ML/AI in finance, Spark/Dask/Ray workloads.		
3.	Shared Infrastructure / Virtualization Node	Quantity: 1 Server Type: 2U General Purpose Server Configuration: • CPU:		

		• 2 × AMD EPYC 9354 (32 cores each) or substantially equivalent. • Memory: • 256 GB DDR5 RAM • Storage: • 4 × 3.84 TB NVMe SSD • Network: • 2 × 10/25 Gbps NIC • Power: • Dual redundant PSUs Use Case: Proxmox / VMware / KVM, virtual desktops, jump hosts, notebooks, shared services.		
4.	High-Speed Network Fabric	Quantity: 1 Leaf Switch (Low Latency) • 1 × 32-port 25G SFP28 switch (Upgradeable to 100G uplinks) • Latency: • ≤ 800 nanoseconds port-to-port • Features: • Layer 2/3 • RDMA support • Cut-through forwarding • Includes: • All required optics (SFP28) • Network OS licenses		
5.	Warranty	3 years comprehensive on-site warranty, next business day on-site response/service.		