



CoreModule 2.0 with L4

Benefits and Features

Performance and flexibility

- Modular chassis with eight compute sleds in 7U
- Two NVIDIA L4 GPUs per sled with 24GB ECC memory each
- MX760C sled with two Intel Xeon 4th or 5th Gen CPUs
- PCIe passthrough and SR-IOV for virtual GPU functionality
- Native chassis integration and management with Dell OME
- Redundant power, cooling and network connectivity for minimal downtime

Learn More

EMEA Sales

+44 (0)20 7960 2400
emeasales@amulethotkey.com

N America Sales

+1 (212) 269 9300
ussales@amulethotkey.com

APJ Sales

+61 409 930 884
apsales@amulethotkey.com

Flexible, Compact & Secure

- AI model development, training and validation
- Simulation and Modelling
- Data analytics and Scientific research
- Virtual Desktop Infrastructure

Maximum performance, simple scalability

Dual NVIDIA L4 tensor core GPUs and dual Intel 4th or 5th Gen Xeon Scalable CPUs in eight independent compute sleds bring huge performance potential and space to grow.

All inside the chassis

No cables, no new boxes, no extra management tools - CoreModule fits completely with the MX7000 rear expansion bays using the internal PCIe connection and works with the standard chassis management controller.



Flexibility for evolving needs

Designed for a range of compute requirements, the direct PCIe connection between GPU and the host CPU makes it easy to run a range of workloads with maximum performance and no data contention.

Easily upgrade existing MX7000 deployments

CoreModule L4 uses the "Fabric B" rear expansion bays in MX7000 and can be fitted to any existing MX7000 chassis or upgrade a system currently using CoreModule T4.

www.amulethotkey.com/support

MX7000 CoreModule L4 Brief - Feb 2024

