

## Reference Architecture: GPU Acceleration over Ethernet for Dell EMC PowerEdge MX7000

### Tech Note by

Ramesh Radhakrishnan  
Seamus Jones

### Summary

Many of today's demanding applications require GPU resources. This reference architecture incorporates GPUs to the PowerEdge MX infrastructure, utilizing the PowerEdge MX Scalable Fabric, Dell EMC DSS 8440 GPU Server and Liquid Command Center Software.

Request a remote demo of this reference architecture or a quote from Dell Technologies Design Solutions Experts at the [Design Solutions Portal](#)

### Background


Emerging workloads, like AI represent a powerfully uneven series of compute processes, such as data-heavy ingest and GPU-heavy data training. When coupled with the fact that these workloads can demand even more resources over time, it becomes clear this complex new paradigm demands a new type of IT infrastructure.

Dell EMC PowerEdge MX7000 modular chassis simplifies the deployment and management of today's challenging workloads by allowing IT to dynamically assign, move and scale shared pools of compute, storage and networking. It provides IT the ability to deliver fast results, not spend time managing and reconfiguring infrastructure to meet ever-changing needs. Composable GPU Infrastructure from Liquid powered by Dell Technologies expands the promise of software-defined composability for today's AI-driven compute environments and high value applications.


### GPU Acceleration for MX7000

For unique workloads like AI that require accelerated computing, the addition of GPU acceleration within the MX7000 is paramount. With Liquid, supported GPUs can be quickly added to any new or existing MX7000 compute sled, delivering the resources needed to effectively handle each step of the AI workflow including data ingest, cleaning/tagging, training, and inference. Spin-up new bare-metal servers with the exact number of GPUs required, and add or remove dynamically as needed, via Liquid software.

#### GPU-oF Delivered via Liquid Command Center



## Essential PowerEdge Components and Ethernet Cabling



**Dell EMC Networking MX9116n on Fabric A**  
 Ultra-low latency, 100GbE switching engine that interconnects pools of disaggregated GPU devices via high performance Ethernet.

**Dell EMC MX7000 Modular Chassis**  
 7U modular enclosure with eight slots. Holds dual socket MX740c compute nodes or quad socket MX840c double-width compute sleds and 12Gbs single-width storage sleds.

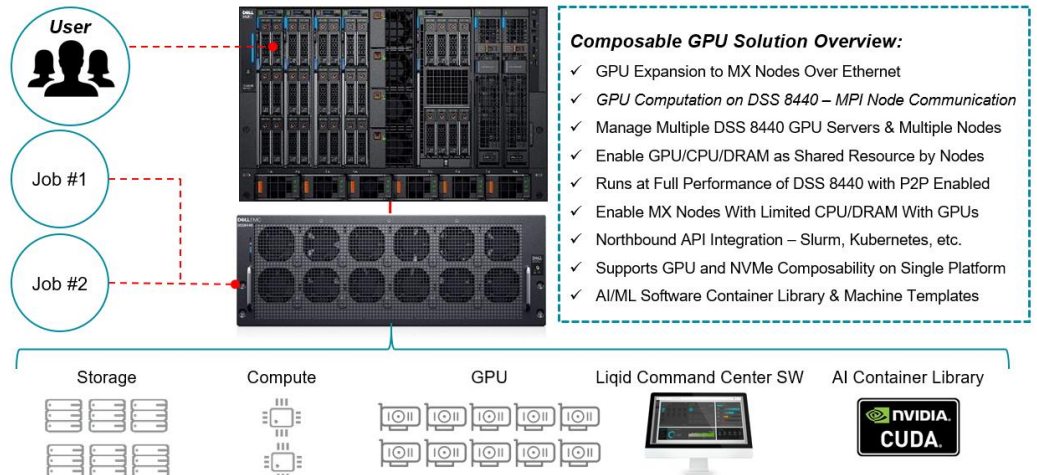
**Dell EMC DSS 8440 Server w/ Ubuntu OS**  
 Supports up to 10x GPU (FHFL) or 16x GPU (HHHL). Flexible topology design supports multiple GPU types as well as NVMe storage devices.

## Liquid Command Center Software

The first step in the GPU expansion process, is to install up to 16x HHHL or 10x FHFL GPUs into a Dell EMC DSS 8440 server. As noted in the table 1, this solution supports several GPU device options. The next step is to connect the DSS 8440 to Fabric A on the MX7000 via 100GbE.


Liquid Command Center software resides on the fabric and will discover the GPU devices in the DS8440 and enable them for utilization by the MX7000 compute nodes. The users can distribute GPU-centric jobs from any compute sled on the MX7000 to GPUs located within the DSS 8440.

GPU Expansion Over Ethernet	
<b>GPU Chassis</b>	Dell DSS 8440 with Ubuntu OS
<b>Ethernet Interconnect</b>	Fabric MX9116n (Multiple 100GbE Connections Available)
<b>GPU Expansion</b>	10x GPU (FHFL) or 16x GPU (HHHL)
<b>GPU Supported</b>	V100, A100, RTX, T4, Others
<b>Compute Node OS Supported</b>	Linux Distributions (VMware and Microsoft Windows Support Pending)
<b>Devices Supported</b>	GPU and NVMe Storage
<b>Form Factor</b>	11U MX7000 (7U) + DSS 8440 (4U)



## Accelerator Performance

To effectively demonstrate the performance of GPU accelerated MX7000 compute sleds, we tested it against DSS 8440 server with local GPUs and measured minimal to no overhead. The deep learning benchmark tests were run on the following networks: ResNet-50, ResNet-152, Inception V3, VGG-16. The DS8440 was outfitted with 8x NVIDIA Tesla RTX8000 GPUs. The results clearly demonstrate that GPU enabled MX7000 delivers unrestricted performance on various industry standard benchmarks, using accelerator optimized Dell PowerEdge infrastructure.



## Benefits of this Reference Architecture

<b>Scalable</b>	Up to 16x GPU per MX7000 Compute Blade and scalable to multi-node
<b>Sharable</b>	Enables GPU Resource Sharing Between Compute Nodes
<b>Dynamic</b>	Reconfigurable GPU-to-CPU Ratios to Match Workloads
<b>Serviceable</b>	No Downtime on MX7000 to Service GPU Pool
<b>Disaggregated</b>	Does Not Consume an IO Fabric Slot on MX7000
<b>Accelerated</b>	Supports GPU Peer-2-Peer Capability on DSS 8440
<b>Future-Proof</b>	Supports Future GPUs Including A100
<b>Availability</b>	Available Now thru <a href="#">Dell Technologies Custom Design Solutions</a> (Formerly OEM   Embedded & Edge Solutions)

## In Conclusion

GPU expansion for the MX7000 unlocks the ability to handle the most demanding compute workloads for both new and existing AI and HPC deployments. Liquid Command Center on Dell EMC PowerEdge Servers accelerates applications by dynamically composing GPU resources directly to workloads without a power cycle on the compute sled.