



# PowerEdge MX-Series Optimizations for the Software Defined Data Center

Technical Note by:

Andrew Hawthorn  
Todd Mottershead

## SUMMARY

The Software Defined Data Center is emerging as one of the leading architectural approaches for customers who wish to cut costs, increase agility and improve reliability without incurring vendor lock-in.

The new MX-Series Modular solutions from Dell EMC were designed specifically to enable SDDC by integrating key optimizations for Software Defined Storage and Software Defined Networking with the industry leading Dell EMC PowerEdge Server family.

## Background

The term Software-Defined Data Center (SDDC) defines the extension of virtualization to all data center resources. According to Wikipedia, an SDDC virtualizes “all elements of the infrastructure – networking, storage, CPU and security and delivers them as a service.” The benefits include reduced acquisition cost, reduced operating cost, increased levels of automation with each component potentially provisioned, operated and managed through an application programming interface (API).

Most customers have started the journey towards SDDC with the implementation of server based hypervisors like VMWare, Hyper-V and KVM for the compute layer and many have extended this concept to storage with Software Defined Storage (SDS) solutions. For these customers, the next stage in the journey will often be the virtualization of Network Services (SDN). The Dell EMC MX-Series has been designed specifically to assist customers with this journey.

## MX Optimizations for SDS

Software Defined Storage (SDS) solutions aggregate disk storage local to each server into a highly reliable, extremely scalable, high performance storage pool that is easy to deploy and manage. This approach costs less, performs better and has helped many customers reduce the time it takes to deploy new solutions but has historically been a poor fit for blade environments due to their low disk counts.

The complexity of managing large numbers of servers led to the development of blade systems where multiple servers could be enclosed and managed from a single chassis. The challenge for many customers as they evolve to SDS based storage is that these systems were designed for the SAN based storage technologies available and simply do not offer the local disk capacities necessary for SDS.

With drive subsystems optimized primarily for boot functions, most blade designs offer only 2 drives. In designing the MX-Series, Dell EMC Engineers took the opportunity to rethink the entire architecture and to design a solution that not only exceeds the management efficiencies of existing blade solutions but adds in key design elements that make it an ideal solution for SDS environments.

## MX SDS Capacity Enhancements

A key element of SDS optimization is capacity. Unlike other 2 socket blade solutions on the market, the Dell EMC PowerEdge MX740c offers up to 8 drives including 6 front-mount hot pluggable 2.5” drives and 2 internal SSD's installed on the optional Boot Optimized Storage System (BOSS) controller.





For customers desiring the cost efficiency of traditional disks, this allows for up to 5 x 2.4TB SAS drives for a total of 12TB of raw capacity, 1 SSD for caching and then offers the optional BOSS controller to provide space for the operating system and log files.

For customers desiring maximum capacity and performance, the system can be configured for “All Flash” operation with raw capacities up to 23TB using 6 x 3.84TB SSD’s. The optional BOSS controller can again be utilized maintain the operating system and log files allowing all hot pluggable drives to be utilized for SDS.

### Additional Expansion

For customers with even more demanding storage requirements, the optional MX5016s storage system can be added. With support for up to 16 additional hot plug SAS drives, this device can be used add drive slots to 1 server or to divide the drive slots between other servers in the enclosure to increase the capacity of an SDS solution.

### Other considerations for optimized SDS

SDS solutions are designed to accommodate network bottlenecks however, it is logical that reducing latency and increasing bandwidth can also increase the efficiency of the storage pool by reducing the replication time required to protect the pool. “All Flash” environments in particular can deliver significantly higher Input/Outputs per second (IOP’s) than traditional disks and the subsequent increase in disk activity can more quickly be processed in an environment where network latency is reduced and/or network bandwidth is increased.

Dell EMC Engineers addressed both of these elements in the MX-Series. First, all compute devices have standardized on 25Gb/s Ethernet which more than doubles the throughput available with existing 10Gb/s technologies. Second, the MX-Series offers network switching options specifically designed to accommodate the full bandwidth of the solution with no oversubscription. With industry-leading latency rates of sub-600ns for the MX5108n and Sub-500ns for the MX9116n, the MX-Series also processes network transactions faster.

### MX Optimizations for SDN

All MX switch options come enabled for SDN. Delivering on the promise of “Open Networking,” both the MX5108n and the MX9116n ship pre-configured with Dell EMC Networking OS10 but include ONIE, enabling the option for third party SDN solutions from Dell EMC Networking partners.

### Conclusion

The Dell EMC MX-Series was designed specifically to help customers facilitate their move to a Software Defined Data Center. From compute, to storage through to networking, the solution has been designed to work seamlessly and deliver uncompromised performance, reliability and efficiency. Tying all of these elements together is the new Open Manage Enterprise management framework which is built on the same API’s that Dell EMC provides to customers for custom development. This open approach allows customers to integrate the MX-Series seamlessly within their own management tools and workflows or, use the interface provided by Dell. Most importantly, this “open” approach allows customers to avoid costly vendor lock-in.



**Note:** an SDS solution can survive the failure of multiple disks but not the failure of all the disks. For this reason, Dell EMC Engineers designed the MX5016s to support a maximum of 16 drives and to allow the addition of multiple MX5016’s in a single enclosure to avoid the risk of creating a single point of failure for the entire SDS pool. This design dramatically reduces the impact of a storage sled failure. Each MX5016s has redundant, hot plug expanders and is connected to redundant SAS switches in the enclosure.