Deploying the Dell M1000e Blade Chassis with 10G Ethernet Pass-Through Modules (PTM) in a Highly Virtualized and Converged Fabric Environment:

Focusing on the Network

Victor Lama
Fabric Specialist G500

EXECUTIVE SUMMARY

Today's heavily virtualized data centers demand a lot of the network infrastructure. It must be very stable, scalable, resilient, and efficient in its forwarding decisions. The applications which provide the services leveraged by an organization to support its business drivers all rest on the ability of the network to sustain and accommodate workloads of widely varying characteristics. The network is in fact the engine that enables all the end systems – servers, storage, applications and clients. Therefore, unlike in days past, the network can no longer be treated as an afterthought when planning a data center build.

The Challenges in Today's Data Center

A typical large enterprise data center can run hundreds of applications, each with its own compute, storage and network requirements. Planning a data center build to support such applications is further complicated by the fact that many organizations do not have documented baseline resource requirements based on pre-deployment testing for each of these applications. Furthermore, the number of clients that come to depend on the application's services typically increase greatly after the application is placed into production. An additional challenge is presented by the bandwidth requirements placed on the network infrastructure when the applications are running on virtualized servers – oftentimes to the tune of anywhere between 5 and 10 virtual machines per processor core. That equates to up to 80 VMs per dual-socket, quad-core server!

With the advent of LAN and SAN convergence, the data center network's complexity has increased exponentially. Unified fabrics present a long-term investment that may eventually lead to a large reduction in operational and capital expenditures. But far from simplifying things, FCoE adds significant complexity to deploying and maintaining the data center network. Access layer switching appliances that were running only one protocol stack (TCP/IP) are now running three – TCP/IP, FCoE and FC. In addition to the complexity, FCoE imposes a demand on edge and inter-switch link bandwidth utilization that typically varies between 20 and 50 percent. This requires the deployment of a carefully planned QoS scheme across the entire data center. These complications call for a simplified, yet robust, architecture that will mitigate these challenges.

Service-oriented architectures are the basis for public and private cloud computing networks. The services offered to consumers of cloud-based services must be highly available, reliable, secure, and meet pre-defined Service Level Agreements (SLAs) and general user-experience expectations. Consistency in the deployment of network policies (QoS, security, VLANs, routing, etc) and the management and portability of those policies as virtual machines are migrated throughout the data center are key to sustaining client SLAs. The resiliency offered by automated virtual machine migration and orchestration make it imperative that the server access architecture be uniform with regard to the network switch's features and functionality and bandwidth capabilities. Such uniformity also contributes to easing the management burden by providing predictable traffic flows and deterministic failover.

Just imagine the case in which a VM running a mission critical application on a rack mount server that enjoys a dedicated network access port for its CNA is migrated to a blade server with a 4- or 8-to-1 oversubscription ratio. Now add the bandwidth utilization of the FCoE traffic. How will that impact the performance of the application? How will the client's experience and promised SLA be impacted? One can never really know unless extensive failover testing is done on each and every application, which is typically not the case in any data center that hosts hundreds of applications.

With the rising cost of maintaining large enterprise data centers, it is highly desirable to construct server farms that leverage server consolidation with smaller footprints that also require less power and cooling. This requirement can easily be met by deploying blade servers when possible.

The demands made on today's data center networks are much greater and more difficult to maintain than at any time in the past. Data center server access architectures must be uniform and provide sufficient edge and core bandwidth to support the very wide spectrum of applications and services on which the business depends for its survival and success.

The Dell M1000e as a Consolidated Computing Solution

The Dell M1000e blade chassis can support up to 16 half-height or 8 full-height server blades in a compact 10U form factor. The consolidated footprint allows for denser server farm designs and reduced space consumption, which can result in considerable savings for organizations that pay rent at co-location and disaster recovery facilities. Compared to deploying rack mount servers, the blade chassis also offers reduced power and cooling consumption in a streamlined management paradigm.

Selecting the Correct Network I/O Blade

Central to the deployment of the blade chassis are the I/O modules that will interconnect the server CNAs with the external switch fabric. While chassis-based FCoE switch solutions, such as the Dell M8428-k, can offer considerable benefits over a top-of-rack FCF (FCoE Forwarder) deployment, it is not the right fit for every organization. It depends on the extent to which the architects and design engineers would like to unify the network fabric, as well as previous technology investments that may have already been made.

For those deployments in which a top-of-rack or end-of-row FCF approach is the adopted standard, the Dell M1000e offers a KR-based, CEE-enabled 10G Ethernet Pass-Through Module (PTM). The Dell 10G PTM enables an alternative blade server access architecture that alleviates complexity and yields a more manageable data center while retaining many of the advantages of blade servers and server virtualization.

Deploying 10G PTMs result in a 1:1 correlation between a blade server's 10G CNA port and a 10G top-of-rack switch port. This dedicated-connection-per-server approach provides the necessary bandwidth to support highly virtualized, converged networks by removing the oversubscription and added latency inherent in a switch-based chassis I/O solution. Moreover, it provides parity with the access approach taken with rack mount servers, resulting in a server access architecture that is simplified, scalable, and predictable in terms of traffic patterns and network policy applicability. The KR-based passive midplane and KR-enabled PTM are designed to support 40G backplane Ethernet for future growth and an extended lifecycle.

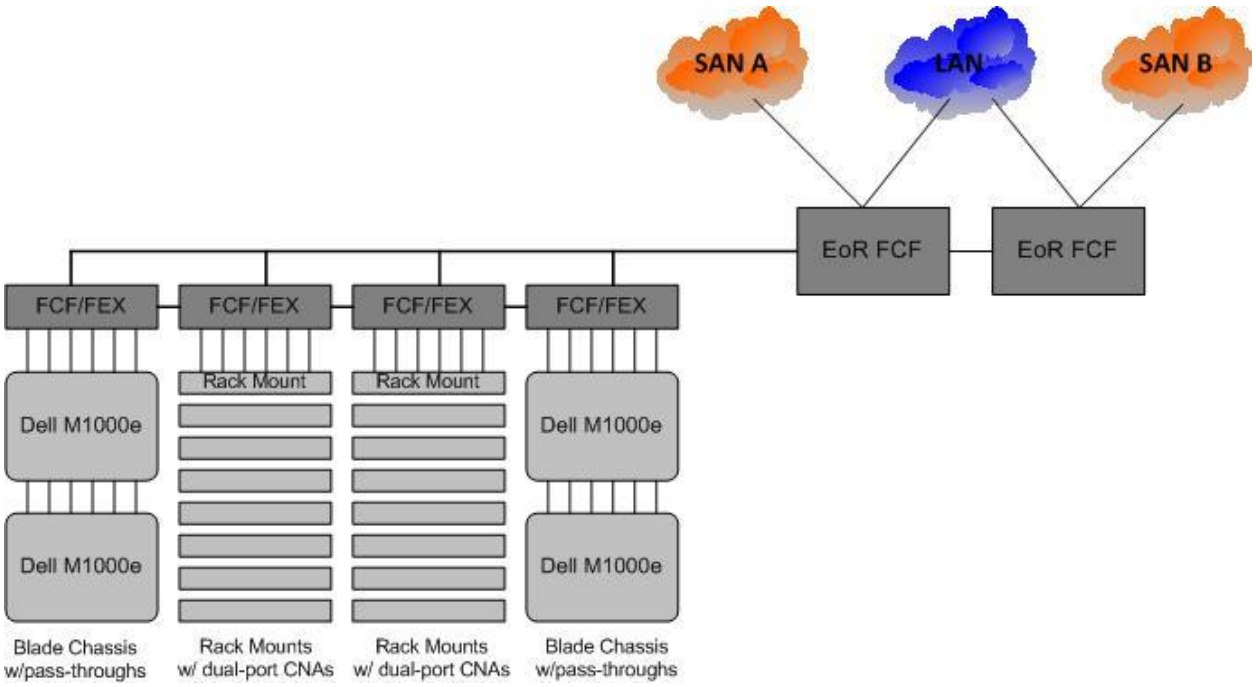

Figure 1 Uniform Server Access Architecture

The Advantages of Deploying the Dell 10G PTM Blade

Simplified Deployment and Management

Using a Dell 10G PTM instead of an FCoE pass-through switch dramatically decreases the number of management points by a factor of 2 switches for each chassis deployed. There is no software to upgrade, no complex QoS/DCB configuration to manage, no worries about whether FIP snooping is operating correctly, no CAM tables to analyze, and troubleshooting is made easier by the fact that each server CNA port is directly mapped to a ToR switch port. Just slide the 10G PTM into the chassis, connect the external cables to the top-of-rack switch and the installation is complete.

Robust and Flexible

The Dell 10G PTM is optimized for virtualized servers and unified fabrics. It can support legacy IP Ethernet, Converged Enhanced Ethernet (CEE), and FCoE. This allows for seamless migration to unified fabrics without the need to swap out the pass-through modules or upgrade the software because they organically possess the technology to pass these types of traffic. The 10G server access speed provides the necessary bandwidth to support highly virtualized servers and the transparent nature of the 10G PTM lends itself to SR-IOV (Single Root I/O Virtualization), which allows the virtual machine environment to be managed by an external top-of-rack switch that supports the VEPA (Virtual Ethernet Port Aggregator – 802.1Qbg) standard or Cisco's VN-Tag (802.1Qbh). It also preserves the ability to partition a server NIC into separate physical PCI-e functions with Dell's innovative NPAR technology.

Economical

A top-of-rack approach also allows the Dell 10G Ethernet PTM to leverage inexpensive Twinax cable with low-power (~1.5 W at each end) SFP+ transceivers. Twinax cables offer a 90% reduction in costs as compared to optical cables and optical SFP+ transceivers. The short cables will be confined to the equipment cabinet and terminated at the top-of-rack FEX/switch. Therefore, no extra inter-rack cable management is necessary. There is also the added savings of purchasing inexpensive pass-through modules as opposed to full-featured FCoE pass-through bridges.

For more detailed information on the Dell 10G Ethernet Pass-Through Module, please go to:

http://www.dell.com/us/business/p/dell-ethernet-passthru-10gb/pd

SUMMARY

Deploying the Dell M1000e Blade Chassis with the Dell 10G Ethernet Pass-Through Module provides the following benefits:

- Uniform server access architecture across the data center with consistent subscription ratios per server
- Sufficient server access bandwidth to support heavily virtualized and converged networks
- Highly simplified deployment and management of blade server clusters
- Simplified troubleshooting through the 1:1 mapping of CNAs to ToR switch ports
- Support for control and management of virtual servers through VEPA and VN-Tag
- Support for Dell NIC Partitioning (NPAR)
- Drastically reduced capital costs, as compared to deploying full-featured Ethernet bridges