

Dell Networking FC Flex IOM: Infrastructure & Network Convergence w/ FCoE

Dell Networking Solutions Engineering
January 2014

Revisions

Date	Description	Authors
12/15/13	Version 1 - Initial Document	Humair Ahmed
01/16/14	Version 1.1 - Minor updates	Humair Ahmed

Copyright © 2014-2017 Dell Inc. or its subsidiaries. All Rights Reserved.

Except as stated below, no part of this document may be reproduced, distributed or transmitted in any form or by any means, without express permission of Dell.

You may distribute this document within your company or organization only, without alteration of its contents.

THIS DOCUMENT IS PROVIDED "AS-IS", AND WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED. IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE SPECIFICALLY DISCLAIMED. PRODUCT WARRANTIES APPLICABLE TO THE DELL PRODUCTS DESCRIBED IN THIS DOCUMENT MAY BE FOUND AT: <http://www.dell.com/learn/us/en/vn/terms-of-sale-commercial-and-public-sector-warranties>

Performance of network reference architectures discussed in this document may vary with differing deployment conditions, network loads, and the like.

Third party products may be included in reference architectures for the convenience of the reader. Inclusion of such third party products does not necessarily constitute Dell's recommendation of those products. Please consult your Dell representative for additional information.

Trademarks used in this text: Dell™, the Dell logo, Dell Boomi™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, Force10™ and Vostro™ are trademarks of Dell Inc. EMC VNX®, and EMC Unisphere® are registered trademarks of Dell. Other Dell trademarks may be used in this document. Cisco Nexus®, Cisco MDS®, Cisco NX-OS®, and other Cisco Catalyst® are registered trademarks of Cisco System Inc. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™ and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. VMware®, Virtual SMP®, vMotion®, vCenter® and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation. Broadcom® and NetXtreme® are registered trademarks of QLogic is a registered trademark of QLogic Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and/or names or their products and are the property of their respective owners. Dell disclaims proprietary interest in the marks and names of others.


Table of contents

- Revisions.....2
- Overview4
- Convergence in the Data Center7
 - FCoE7
 - Convergence with FCoE using Dell FC Flex IOM8
- Why Dell MXL/IOA with FC Flex IOM11
- Conclusion12

Overview

In today's Enterprise and Data Center environments, Ethernet has become the interconnection medium of choice for LANs and networking, while Fibre Channel (FC) has dominated on the storage front in storage networking. While Ethernet has seen much success owed to the constant evolution and low cost of Ethernet technologies along with massive investment in infrastructure and management, FC has seen incredible success in the storage networking market due to its intrinsic nature of providing lossless traffic behavior and wide acceptance from storage and adapter vendors. For Ethernet, 10 GbE is quickly becoming the standard with 40 GbE following closely behind; FC has standardized on 8G with 16G slowly being considered within the industry. Today, with Ethernet and FC both having achieved much success in their respective domains, a need has arisen to simplify, consolidate, and make more efficient the Enterprise and Data Center networks of the future.

Further, the fact that both Ethernet and FC have succeeded in their own separate domains has resulted in the proliferation of different types of Peripheral Component Interconnect Express cards (PCIe) for servers to handle different types of traffic, separate cabling to handle Ethernet and FC, additional expense in infrastructure for separate Ethernet and FC hardware, and more power and cooling requirements.



Fortunately, with the evolution of technology and new protocols like Data Center Bridging (DCB), I/O consolidation is enabling the converged networks of the future where Ethernet and FC traffic are able to use the same physical infrastructure. In simple terms, I/O consolidation is the ability of a switch or a host adapter to use the same physical infrastructure to carry different types of traffic with different traffic characteristics and requirements.

10 GbE, the development and advancement of PCIe adapters, and the Data Center Bridging (DCB) protocol have made converged networking possible. This unified approach to a converged environment is having several positive effects in Enterprises and Data Centers:

1. Reduction in cabling and optics
2. Lower number of PCIe adapters needed in servers which allows for greater consolidation of servers and denser blade chassis systems
3. Consolidation of switches
4. Less maintenance and points of failure
5. Cost savings in infrastructure hardware and power/cooling

Typically on a standalone server, this consolidation of hardware can have the effect as shown below in Figure 1 where the adapter requirements are cut by 50%. An additional benefit is seen in a blade server chassis where more servers can be consolidated in a smaller space due to the space saved by fewer PCIe adapter requirements.

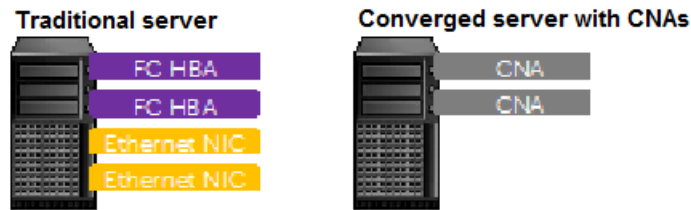


Figure 1 Four NICs/HBAs replaced with two CNAs

In Figure 2, the traditional LAN and SAN network is displayed. Note how the LAN network has its own Ethernet switches and the SAN network has its own FC switches. The server connects to both the LAN and SAN network. Also shown are two FC fabrics for high availability. With the advent of new technologies that allow both network and storage traffic to run effectively on one network, this degree of separation is no longer needed or practical.

As shown in Figure 3, I/O consolidation allows for a converged network where less hardware is needed and both LAN and SAN traffic traverse the same physical infrastructure. The different types of traffic in this case are LAN traffic which is typically lossy and storage traffic which is typically lossless. Not only is there consolidation in devices and cabling but also a consolidation in PCI-e adapters required on the server to carry different types of traffic. This is especially important in a blade server chassis since the saved space allows for greater density in terms of server blades.

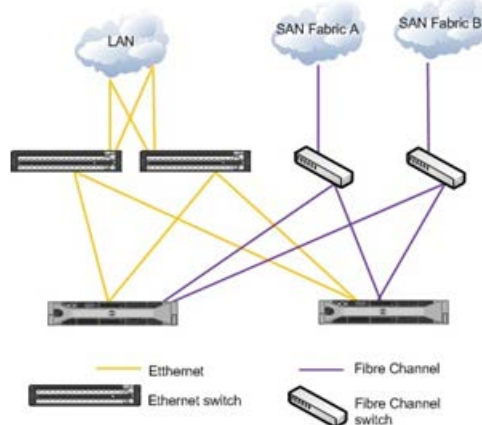


Figure 2 Traditional LAN and SAN network

In Figure 3, Dell PowerEdge rack servers are being utilized and convergence is occurring at the ToR converged switch. Although the network has been converged and hardware, cabling, maintenance, power, and cooling requirements have been minimized, additional consolidation can still be done. To provide an extra level of consolidation, Dell's PowerEdge blade servers can be utilized with Dell's MXL/IOA blade switches with FC Flex IO Modules to provide the converged switch capability as shown in Figure 4. This solution provides the additional benefit of infrastructure convergence where there is no longer a need for separate rack servers with ToR switches providing the Ethernet-FC bridging, because the servers and blades are now consolidated within the Dell M1000e chassis; convergence is now occurring within the chassis via Dell MXL/IOA w/ FC Flex IOM. This level of consolidation has further eliminated an extra layer of cabling, minimized power and cooling requirements, and simplified management.

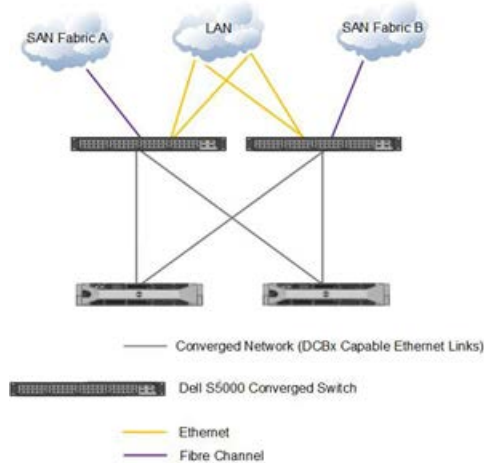


Figure 3 Converged LAN/SAN Network/Infrastructure w/ FC Flex IO

Convergence in the Data Center

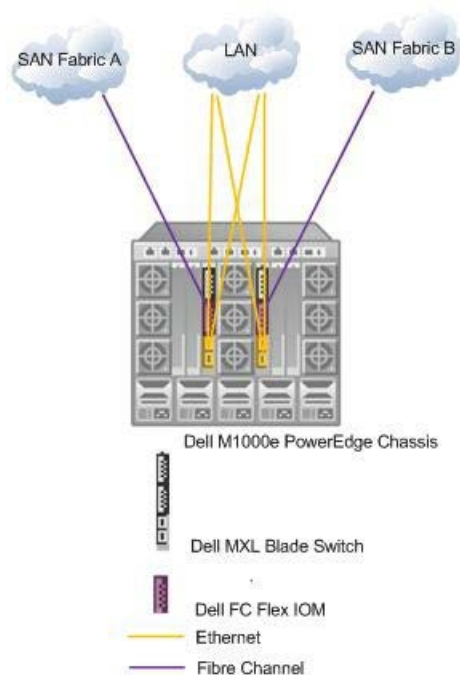


Figure 4 Converged LAN/SAN Network/Infrastructure w/ FC Flex IO

FCoE

Since FC is the dominant storage protocol in the data center, Fibre Channel over Ethernet (FCoE) has become the obvious choice for those looking to migrate to a converged network while keeping the back-end FC SAN intact.

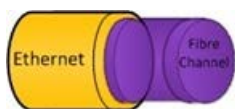


Figure 5 FC frame encapsulated within Ethernet

FCoE can be seen as just another upper-layer protocol on top of Ethernet as shown in Figure 6 below. FC frames are encapsulated in Ethernet frames as shown in Figure 7. Since the maximum FC frame size is 2148 bytes and is larger than the maximum Ethernet payload size of 1500 bytes, 'baby' jumbo Ethernet frames (2500 bytes) must be used for FCoE.

FCoE allows for I/O consolidation by transporting FC frames encapsulated within Ethernet frames. The encapsulation layer is FCoE, and SCSI is mapped over FC. This results in only Ethernet NICs that support FCoE to be needed on the server. Now Ethernet NICs and FC HBAs can be consolidated into one adapter called a Converged Network Adapter (CNA). In Figure 5, the convergence occurs at the blade level where FC is de-encapsulated from the FCoE packet and sent upstream to the FC SAN. The backend SAN remains unchanged allowing for an easy transformation to a converged network via FCoE.



Figure 6 Upper-level protocols over Ethernet



Figure 7 FC Encapsulation in Ethernet

Convergence with FCoE using Dell FC Flex IOM

The Dell FC Flex IO module provides the additional functionality to a Dell MXL/IOA blade switch to enable it to act as a NPIV Proxy Gateway (NPG) capable of bridging between Ethernet and FC. In NPG mode, the Dell MXL/IOA w/ FC Flex IOM does not consume a fabric Domain ID or become part of the switched fabric but instead simply acts as a gateway to the fabric by de-encapsulating FC packets from FCoE and forwarding the frames to the existing backend FC SAN.



Thus, the Dell FC Flex IO Module, once installed in a MXL/IOA blade switch, allows for an easy transition to a converged infrastructure and network via FCoE.

Up to 2 FC Flex IOMs can be inserted per MXL/IOA blade switch allowing for up to 32G of FC throughput per module or 64G throughput per MXL/IOA. All ports support 2/4/8G FC and SFP+ optics. Further, Dell MXL/IOA w/ FC Flex IOM provides flexibility in terms of other Flex IO Modules it can be paired with as shown in Figure 8 below.

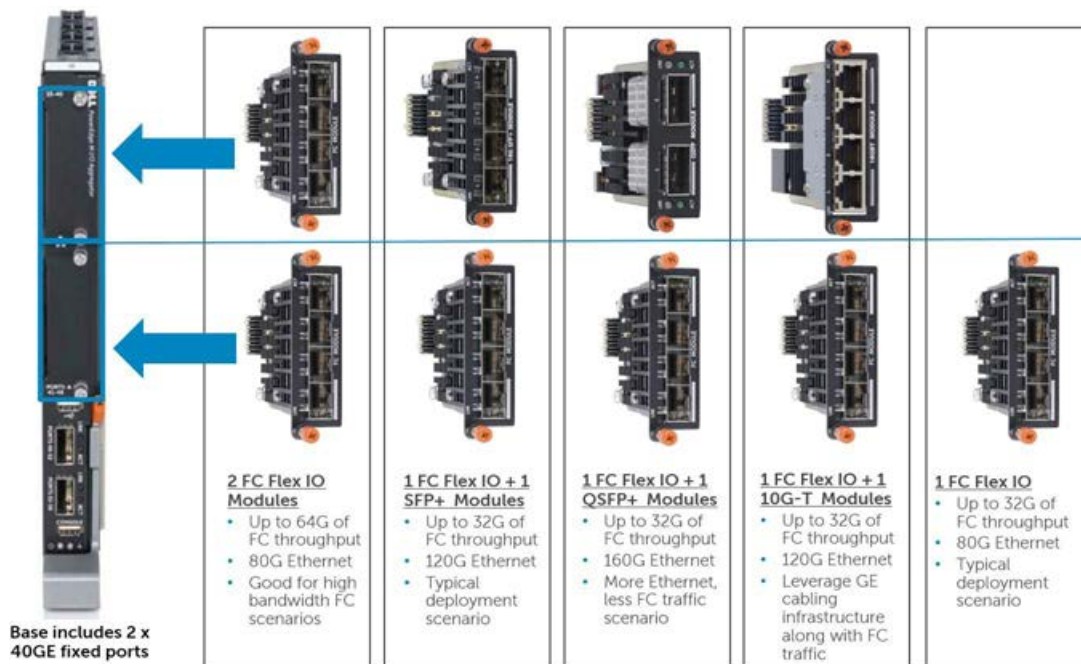


Figure 8 Multiple Flex IOM combinations with FC Flex IOM

For ease of deployment the Dell MXL/IOA will detect the FC Flex IOM and automatically set the Dell MXL/IOA switch to NPIV mode. FC Flex IOM takes the convergence from the ToR down to the blade level consolidating infrastructure while still providing the benefits of network convergence and leveraging the backend FC SAN. Figure 9 demonstrates the extra consolidation when moving from ToR convergence to convergence at the blade.

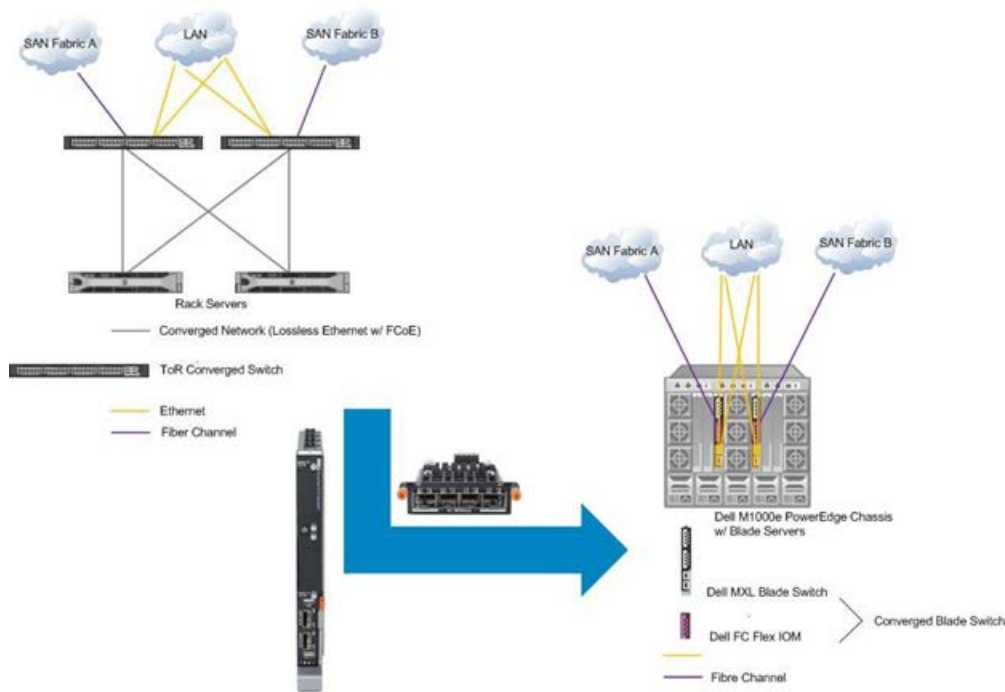


Figure 9 Infrastructure and network convergence with Dell MXL/IOA blade via FC Flex IOM

Figure 9 below displays a topology of a non-converged network with blade servers. Note how blade switches are needed both for the LAN and SAN. In the below topology, Dell MXL blade switches are employed in Fabric A of the M1000e chassis and Brocade M5424 FC switches are employed in Fabric B. Not only is there more blade switches required for the different networks, but there is also twice the number of adapters needed on the servers to support both Ethernet and FC.

Once the FC Flex IOM is added to the Dell MXL blade switch to allow it to behave as a NPIV Proxy Gateway as shown in Figure 10, the topology is greatly simplified with immediate benefit realized visually in terms of hardware requirements.

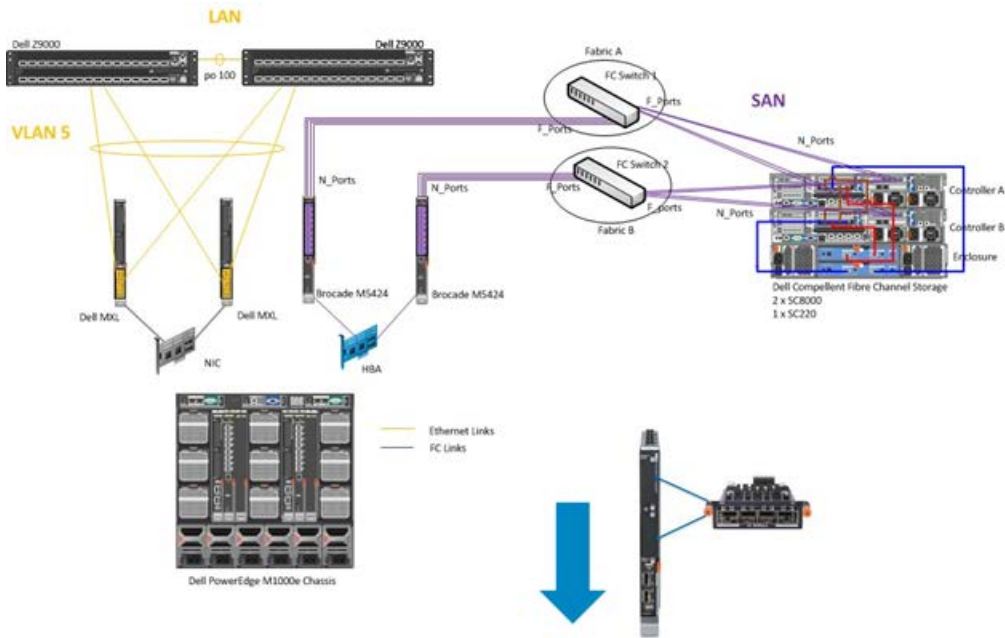


Figure 10 Non-converged network with Dell MXL and Brocade MS425

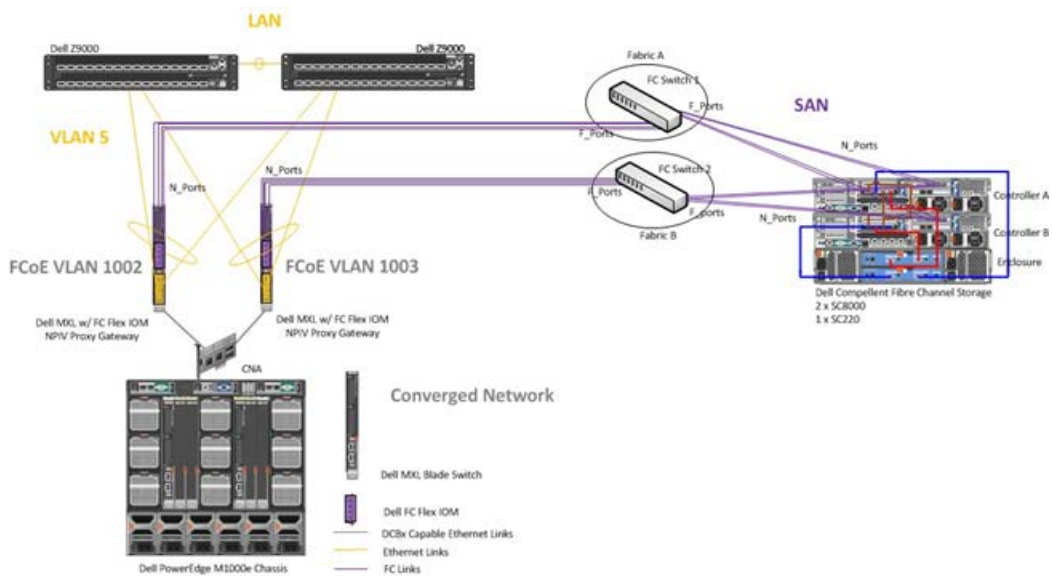


Figure 11 Converged network with Dell MXL w/FC Flex IOM

Why Dell MXL/IOA with FC Flex IOM

Beyond allowing for network and infrastructure convergence, the Dell FC Flex IOM provides unique and clear benefits when compared to options available from any other vendor. Below are some key take-aways of why Dell MXL/IOA w/ FC Flex IOM is ahead of the competition in features and flexibility.

Provides network & infrastructure convergence at the blade level (bridges Ethernet and FC)

- As shown above in comparing Figure 10 and Figure 11, unlike the Brocade M5424, Dell MXL/IOA w/ FC Flex IOM provides for convergence at the blade level splitting Ethernet and FC traffic at the chassis edge.

Provides highest # of server facing ports

	MXL/IOA w/ FC Flex IOM	Brocade M5424 (non-converged)	Brocade M8428-K	HP VC Flex Fabric 10/24	IBM Virtual Fabric 10G+ QL VF Extension Module
# of server facing ports	32	16	16	16	16

Provides highest # of uplinks and the most flexibility

	MXL/IOA w/ FC Flex IOM	Brocade M5424 (non-converged)	Brocade M8428-K	HP VC Flex Fabric 10/24	IBM Virtual Fabric 10G+ QL VF Extension Module
# of uplinks and possible combinations	2x40GE + 8x8GFC <i>OR</i> 4x40GE + 4x8GFC <i>OR</i> 2x40GE + 4x10GE + 4x8GFC <i>OR</i> 2x40GE + 4x10Gbase-T + 4x8GFC	8x10GE + 4x8GFC	8x8GFC	8x10GE <i>OR</i> 4x10GE + 4x8GFC <i>OR</i> 6x10GE + 2x8GFC	Up to 10x10GE Ext Module: 6x8GFC

- Provides flexibility in features/functionality required

- If only layer 2 functionality is required, FC Flex IOM can be used with the Dell IOA blade switch to provide Ethernet-FC bridging capability. In addition, the Dell IOA blade switch requires Zero-touch configuration. Simply insert the blade switch into the Dell M1000e chassis and it functions as a NPIV Proxy Gateway switch with the default configuration provided.
- If additional layer 3 features/capabilities are desired at the blade level, the FC Flex IOM can also be used with the Dell MXL blade switch to provide Ethernet-FC bridging capability.

- Flexible pricing options based on requirements

	MXL/IOA w/ FC Flex IOM	Brocade M5424 (non-converged)	Brocade M8428-K	HP VC Flex Fabric 10/24	IBM Virtual Fabric 10G+ QL VF Extension Module
List Price (includes services)	MXL+ 1FC Flex: \$15,000 IOA + 1FC Flex: \$12,000	\$9,999	\$17,999	\$18,499	\$9,999 + \$6,499

Conclusion

This whitepaper demonstrated and explained the movement from a traditional non-converged LAN/SAN network to a converged LAN/SAN network with a consolidated infrastructure using the Dell FC Flex IOM combined with the Dell MXL/IOA blade switch. This solution provides not only savings with network convergence but also the additional savings of infrastructure convergence where there is no longer a need for separate rack servers with ToR switches providing the Ethernet-FC bridging, because the servers and blades are now consolidated within the Dell M1000e chassis.

In addition, the FC Flex IOM offers flexibility as it can be used with the Dell IOA blade switch for layer 2, zero-touch configuration functionality or with the Dell MXL blade switch for additional layer 3 features. The FC Flex IOM can also be combined with other Flex IO Modules to provide increased flexibility for required uplinks.

Most importantly, this whitepaper concluded with comparisons with other vendor options available today and clearly outlined why the Dell FC Flex IOM with Dell MXL/IOA provides a clear and unique advantage in both features and flexibility. Converge and simplify your network and infrastructure today with the Dell MXL/IOA blade switch w/ FC Flex IOM. See the Dell Networking FC Flex IOM deployment guides for both Brocade and Cisco FC environments for detailed topology and configuration examples.