

ESG Lab Review

Cisco MDS 9148 Multilayer Fabric Switch

Date: March 2010 **Author:** Tony Palmer, Senior Engineer/Analyst, ESG Lab

Abstract: This ESG Lab review documents hands-on testing of the Cisco MDS 9148 8-Gbps Multilayer Fabric Switch with a focus on scalability, flexibility and cost-efficiency.

The Challenges

Faced with exploding data growth and increasing demands thanks to server consolidation and virtualization as well as a need to better protect and manage shared information assets, IT managers in organizations of all sizes are increasingly leveraging a shared storage approach using SAN technology. In the enterprise, this has meant Fibre Channel switches and directors, an option often out of reach of smaller businesses. In fact, when ESG asked IT professionals in small- to medium-size businesses to name the business initiatives which would have the greatest impact on IT spending decisions over the next 12-18 months, the number one response was cost reduction, with simplification of business processes and risk management a close second and third.¹

The Solution: Cisco MDS 9148 Multilayer Fabric Switch

The MDS 9148 is a 48-port, non-blocking Fibre Channel switch that combines the NX-OS operating system used in Cisco's Nexus series of data center Ethernet switches and MDS 9000 SAN switches and directors with configuration wizards for simplified deployment and management. It is offered with a pay-as-you-grow licensing model supporting from 16 to 48 ports of 8-Gbps storage area network (SAN) connectivity.

Figure 1. The Cisco MDS 9148 Multilayer Fabric Switch



Priced as an entry level switch, the MDS 9148 provides many of the enterprise features that are included in the MDS 9500 Series Multilayer Directors and MDS 9200 Series Multilayer Fabric Switches:

- 8-Gbps, non-blocking throughput on all 48 ports.
- FlexAttach technology to enable transparent server deployment without the need to reconfigure the SAN.
- N-Port Virtualization (NPV) technology and N-port ID virtualization (NPIV) support for virtual environments.
- VSAN support for security and fault isolation.
- PortChannels for aggregating up to 16 physical ports into one logical channel

¹ Source: ESG Research Report, 2010 IT Spending Intentions Survey, January 2010.

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by [Company Name].


- In Service Software Upgrade to enable firmware upgrades without removing the switch from service.
- Dual hot-swappable power supplies and fans to support high availability.
- RADIUS and TACACS+ port security, Fibre Channel Security protocol, Secure FTP, Secure Shell Version 2, and role-based access control for comprehensive security controls.
- The 9148, like all MDS switches and directors, blocks corrupted and malformed frames from entering a fabric –protecting servers and applications.

ESG Lab validated the performance and enterprise-class capabilities of the MDS 9148 during a day of hands-on testing at Cisco's campus in San Jose, California.

Performance/Scalability

Like the previous generation MDS 9124, the Cisco MDS 9148 is a fully non-blocking fabric switch. ESG Lab tested performance to validate that the 9148 was able to pass traffic at 8 Gbps on all 48 ports simultaneously. Testing was executed using three 16-port JDS Uniphase load testers connected to all 48 ports of one MDS 9148 at 8 Gbps as seen in Figure 2.

Figure 2. ESG Lab Performance Test Bed



The load tests were executed using JDSU Xgig Maestro software, which was used to verify the rate at which all channels were transmitting and receiving data as well as latency. The load testers were configured to transmit on odd-numbered ports and receive on even-numbered ports at 99.9% of wire speed.

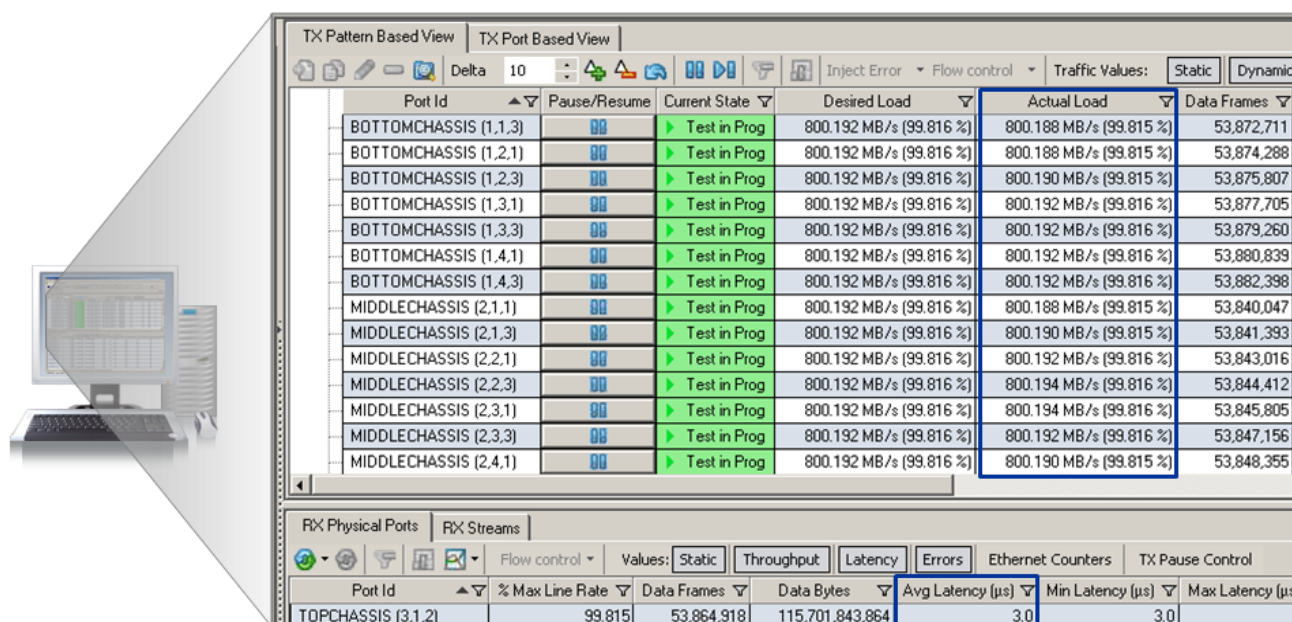
ESG Lab Testing

Three tests were run to validate performance in different configurations. The first test was run with full size frames to test throughput and was performed with a fixed payload size of 2,112 bytes. ESG Lab examined every port during the test and confirmed that all odd-numbered ports were sending and all even-numbered ports were receiving at just over 800 MB/sec. or 7.81 Gbps on each port simultaneously. It is important to note that this represents 99.8% of wire speed. The actual throughput driven by the load tester is rounded down from the configured value of 99.9% due to internal calculations.² Stated another way, ESG Lab confirmed that a single MDS 9148 delivers an aggregate throughput of more than 131 Terabytes per hour across all ports. During this first phase of performance testing, the MDS 9148 delivered perfectly consistent low latency of 3.0 microseconds per frame as shown in Figure 3.

The throughput test was repeated with a minimal, fixed payload size of 24 bytes. Small frame tests show how efficiently the switch can process the Fibre Channel protocol. Even though the payload size was almost 100 times smaller than the first test, all ports were sending or receiving at an impressive 552 MB/sec. with a consistent 1.0 microsecond average latency per frame. In this test, each port was processing 9.66 million frames per second. That's an impressive 463.68 million frames per second across the whole switch.

² "MB/sec." is calculated as 1024*1024 or 1,048,576 bytes per second. Gbps is calculated as 1024³ or 1,073,741,824 bits per second.

Figure 3. 8 Gbps Line Rate Throughput



The next two tests were designed to evaluate fair traffic allocation in the switch. A fan-out test was conducted with 15 ports transmitting to one receiving port. This test was performed to confirm that the switch would effectively divide the load equally between the fifteen simulated server ports. Each of the 15 transmitting ports was attempting to send 800 MB/sec. of data to the target port. The MDS 9148 allowed an actual load of 53.4 MB/sec. on each of the 15 transmitting ports with a full 800 MB/sec. load on the target port.

The final test was a head-of-line blocking test. Head-of-line blocking is a condition where traffic waiting to be transmitted to a congested destination port prevents or blocks traffic destined elsewhere from being transmitted. In this test, the JDSU Xgig load tester was configured with output from port 1 sent to port 2, while port 3 was configured to send output to both ports 2 and 4. The result was an even distribution across ports 2 and 4 of the traffic coming from port 3. The traffic from port 1 was throttled to 50% to accommodate the load from port 3.

Why This Matters

As a SAN fabric scales, the likelihood that users and applications will be vying for the same resources increases. Users are increasing server consolidation through server virtualization, which requires more bandwidth. ESG Lab has confirmed that the 9148's 8-Gbps, non-blocking architecture provides predictable scalability and performance across all 48 ports in the switch in both directions with minimal latency. The MDS 9148 also excelled at fair traffic allocation, ensuring that users and applications can always get to their data, with no head-of-line blocking.


Flexibility and Ease of Management

Cisco has long offered the ability to configure multiple levels of administrative users with role-based access. FlexAttach is a feature that takes advantage of this capability to enable server administrators to move, add, or change server hardware without reconfiguring the SAN. FlexAttach uses zones to associate Cisco Virtual Port World Wide Names (VPWWN) with storage volumes. When a server is attached, its Host Bus Adapters (HBAs) are associated with the appropriate Cisco VP WWN. This functionality allows SAN administrators to protect the configuration of the switch fabric, while still allowing server administrators to replace hardware without assistance from the SAN administrators.

ESG Lab Testing

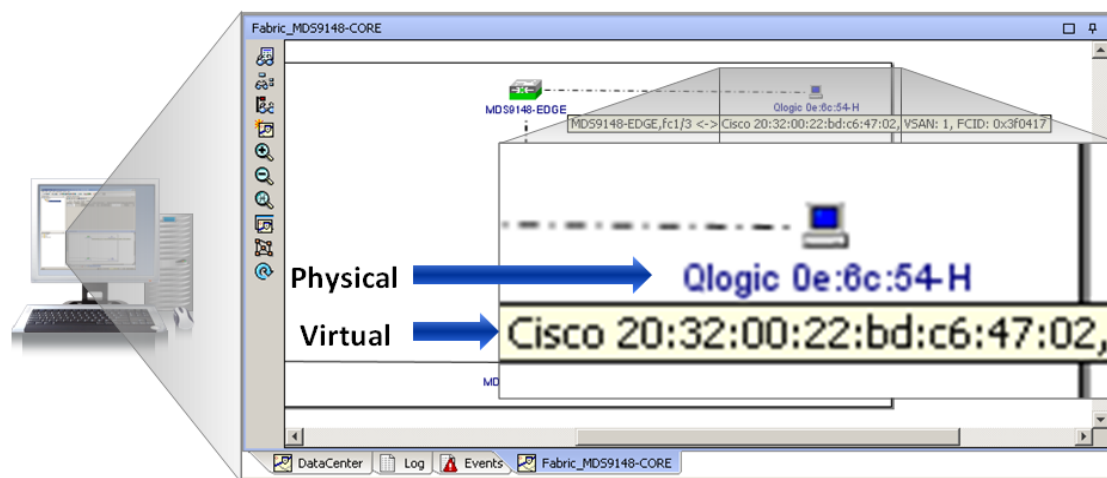
ESG Lab used the test bed illustrated in Figure 4 for the remainder of testing detailed in this report. One Windows 2003 server with a Qlogic QLA 2342 HBA installed was attached to a single MDS 9148 switch which was in turn attached to a second MDS 9148 via two inter-switch links (ISLs). The second MDS 9148 was attached to an Adaptec SANbloc-S50 disk enclosure.

Figure 4. ESG Lab Functionality Test Bed.



ESG Lab utilized a two-step configuration wizard launched from Cisco's Fabric Manager GUI to configure the 'Edge' MDS 9148 to enable FlexAttach on each port. After the configuration was completed, the new Cisco VPWWN assignment was seen under the server's HBA. Figure 5 illustrates the new configuration, showing the server's physical Qlogic HBA and the virtual Cisco HBA configured on the SAN.

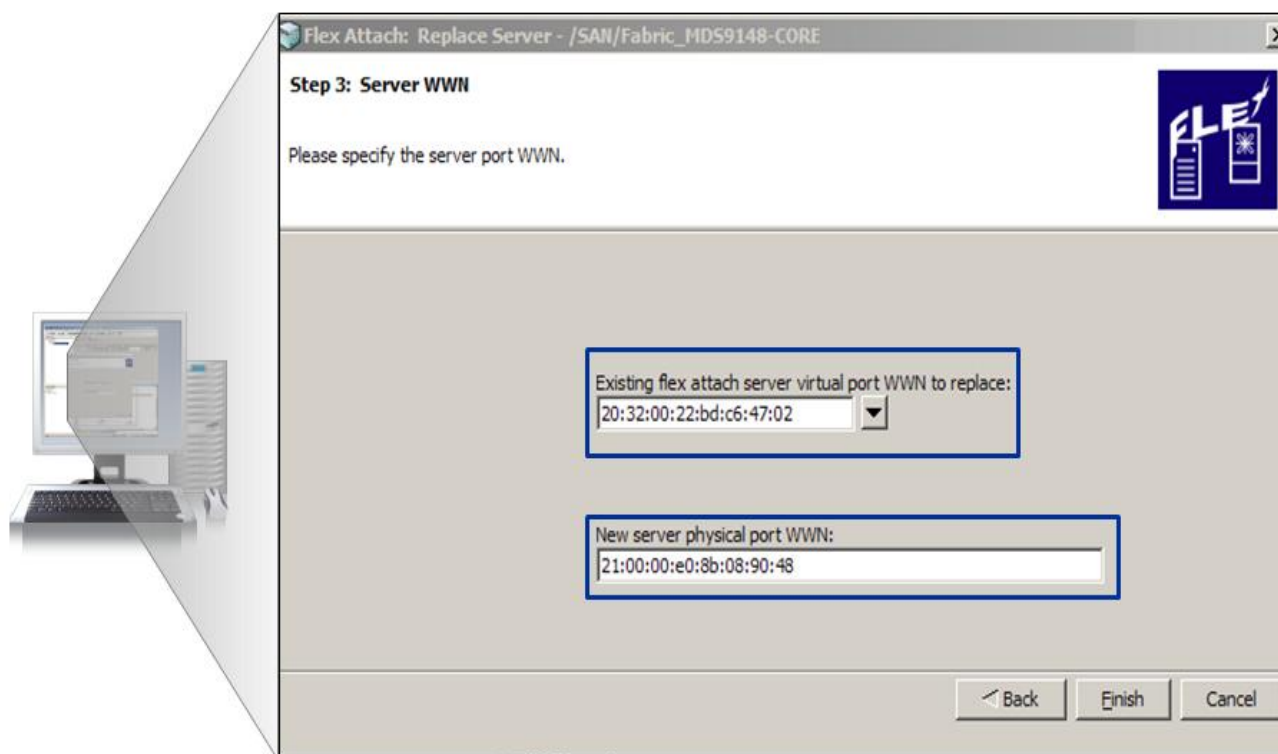
Figure 5. Cisco VPWWN Connected to Host



A hardware failure was simulated by unplugging the attached server from the SAN fabric. A new server was then introduced to the MDS 9148 and the Fabric Manager program was launched using a login that had been previously defined as a server administrator. This role had limited rights to the SAN fabric when compared to the SAN administrator

and was only allowed to add, move, or change specific servers attached to specific switches in the SAN. Using the FlexAttach functionality as shown in Figure 6, ESG Lab was able to associate the new server's HBA WWN with the VPWWN owned by the replaced server. Using Windows Disk Administrator, the new server was verified as able to mount the correct disks without changing zoning or LUN masking.

Figure 6. Replacing a Server with FlexAttach



Why This Matters

Provisioning SAN-attached servers to pools of shared storage is a complex and difficult task using traditional methods. Time and money is wasted and errors are likely as SAN and server administrators must coordinate efforts and often duplicate work as they configure each path for each host and volume. In the best case, errors are discovered and corrected when an application can't find the storage it needs. In the worst case, mis-configured high availability leaves an organization vulnerable to a hardware failure and customers face risks ranging from lost revenue to fines or even lawsuits.

Cisco FlexAttach saves time and money as it increases efficiency and reduces risk, allowing server administrators to manage server attachment to the SAN without needing to engage the SAN administrators or make any changes to SAN configuration. ESG Lab found that FlexAttach, in combination with the MDS 9148, allowed a server administrator to safely and easily replace a server using a simple, repeatable procedure.

Manageability, Serviceability, and Cost Efficiency

A number of powerful, enterprise-class capabilities enhance the manageability, serviceability, and cost efficiency of the Cisco MDS 9148 Multilayer Fabric Switch:

- **Manageability:** N-Port Virtualization, Configuration Wizards
- **Serviceability:** Online firmware upgrade, PortChannels, redundant hot swappable fans and power supplies
- **Cost-Efficiency:** Port upgrades from 16 to 48 ports in eight port increments, no additional licensing fees for enterprise-class features

This section documents ESG Lab testing of three of the enterprise-class capabilities of the MDS 9148: N-Port Virtualization, PortChannels, and In Service Software upgrades.

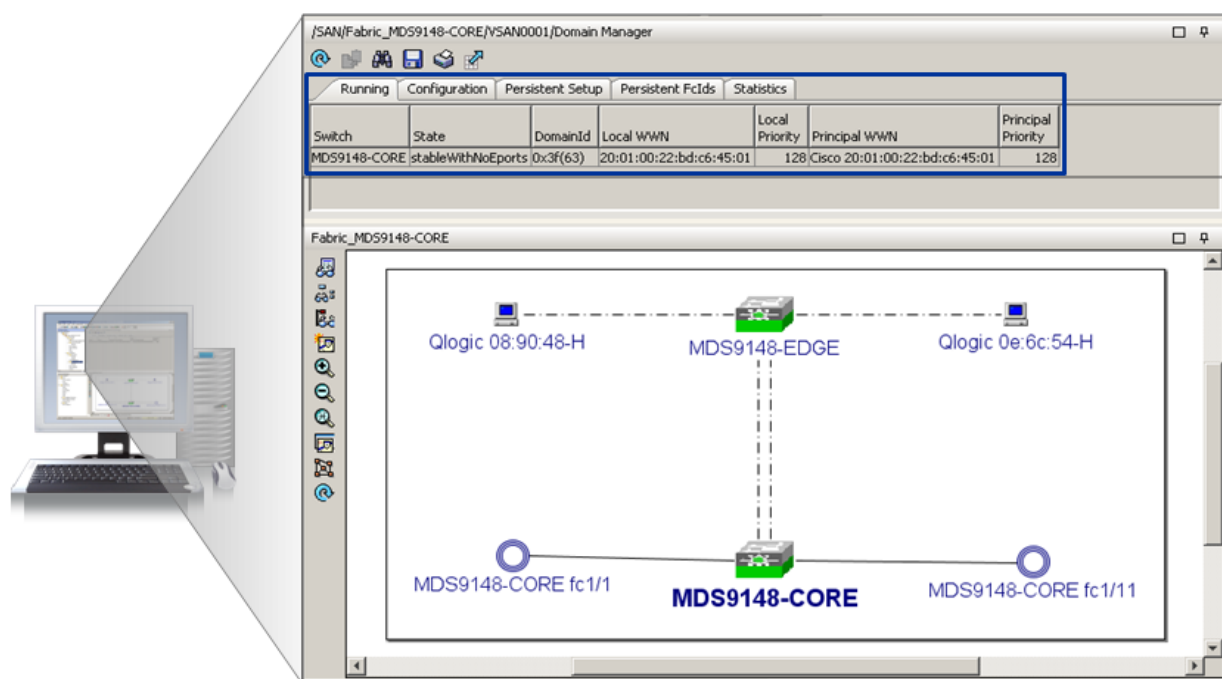
N-Port Virtualization

By default, Cisco Nexus 5000 and MDS Series switches operate in fabric mode. In fabric mode, each switch that joins a SAN is assigned a unique domain ID. Each SAN (or VSAN) supports a maximum of 239 domain IDs, so each SAN has a limit of 239 switches. In large scale, edge to core topologies, SANs may need to grow beyond this limit. N-port Virtualization (NPV) significantly reduces the number of domain IDs needed in a fabric by eliminating the need for a domain ID on edge and top of rack switches.

In NPV mode, the edge switch relays all traffic from server-side ports to the core switch. The core switch provides F-port functionality (such as login and port security) and all Fibre Channel switching capabilities. The edge switch appears as a Fibre Channel host to the core switches and as a regular Fibre Channel switch to its connected devices. As a result, the devices attached to the NPV-mode switches get their domain IDs from the core switch.

ESG Lab configured NPV using a wizard launched from the Cisco Fabric Manager application. This process automatically completed the conversion of two MDS 9148 switches to NPV devices. After the wizard was completed, both switches shared the same domain ID, as shown in Figure 7.

Figure 7. Two Switch Fabrics with One Domain ID Using NPV




PortChannels

PortChannels are used to aggregate up to 16 physical inter-switch links into a single logical path. This provides optimized bandwidth utilization across all links in the PortChannel. When the other end of the PortChannel is a modular chassis, such as an MDS 9500 director, the bundle can include any port from any ASIC in the chassis, ensuring that the channel remains active even in the event of a module failure.


ESG Lab created a PortChannel between edge and core MDS 9148 switches using two inter-switch links on ports 13 and 25 on both switches.

The industry standard IOMETER workload generator was used to create read and write traffic transmitted from two servers attached to the edge MDS 9148 switch across the SAN fabric to two disk volumes on an array attached to the MDS 9148 core switch. Port 13 on the core 9148 switch was then disabled to simulate a port failure. The load was taken on by the remaining active link (port 25) and no disruption in service was experienced on the server.

Figure 8. F-PortChannel Resilience in Action



Before



After

As illustrated in Figure 8, the red and yellow bars show server traffic while the blue and gray bars represent traffic on ports in the PortChannel. Note that when Port 13 (represented by the blue bar) was disconnected, traffic was rerouted through Port 25 (represented by the gray bar) and traffic to and from the servers (the red and yellow bars) was unaffected.


In Service Software Upgrades

During the service life of any switch, it is not unreasonable to assume that firmware will, at some point, need to be updated to either fix a bug or enable an enhanced feature. Cisco's In Service Software Upgrade (ISSU) provides the ability to perform a full NX-OS image upgrade without having to take the switch out of service.

To test both the functionality and the non-disruptive nature of ISSU, the industry standard Iometer workload generator was used to perform read and write IO from a SAN-attached server while the Software Install Wizard was launched from Cisco Device Manager. The entire software installation process took less than four minutes to complete. Iometer continued running throughout the entire process with no errors and no loss of connectivity to storage.

Figure 9 shows the Cisco MDS traffic monitor while the software install wizard was running. The red bar shows read and write traffic being generated by the server on port 3, while the yellow and blue bars represent traffic to and from the storage devices on ports 13 and 25. ESG monitored this display, as well as the IOmeter console, while the software upgrade was running to verify that there was no disruption to service.

Figure 9. ISSU in Process



Why This Matters

Paying a premium for enterprise-class switches or directors is simply not an option for most small- to medium-size businesses. The MDS 9148 is not a stripped down entry level product; it delivers the same full featured enterprise-class functionality as Cisco MDS directors. Features such as NPV and PortChannels provide entry level users with the versatility and capability to grow with their business.

Switch software updates are often required to address issues with third party drivers, to take advantage of changing standards, or to enable enhanced functionality—any business with high availability requirements for its data cannot afford downtime for these updates. ESG Lab testing has confirmed the MDS 9148's ability to perform software updates non-disruptively, with no impact to production applications or users.

The Bigger Truth

This ESG Lab review has confirmed that the NX-OS based Cisco MDS 9148 fabric switch is a very effective solution for small- to medium-size businesses looking to migrate from DAS to high availability SANs without sacrificing the enterprise-class features needed in large scale edge-to-core deployments. Using a switch-on-a-chip architecture, which ESG Lab has proven can scale up to 48 ports of non-blocking, 8-Gbps bandwidth, the MDS 9148 provides a high performance, 16-port entry for a SAN that can be easily upgraded to support future growth. First time SAN users will appreciate the user-friendly configuration wizard which guides an inexperienced user through initial set up in minimal time.

Cisco is making good on its strong commitment to support the small- to medium-size business (SMB) market with full featured products like the MDS 9148 just as it did with the MDS 9124. Enterprise-class features—including VSANs, security, non-disruptive firmware updates, and NPV—provide customers in this market segment with the versatility and capability they need to grow with their businesses. More importantly, all these features will continue to be delivered as part of the standard package, not as incremental costs.

Enterprise customers should also consider the MDS 9148 as a good choice for top-of-rack architectures enabling a core-to-edge solution, especially with existing MDS Directors as they will share the same NX-OS. Its wizard-based configuration options will make adding switches and servers to remote sites or the edge a simple process that will not require experienced SAN professionals' time. The price points could make it more attractive to include lower tiered stranded servers into the SAN environment.

ESG would recommend that Cisco continues to work with its partners to develop simple, preconfigured "SAN in a box" solutions to further aid adoption. These end-to-end solutions help reduce the complexity and ease installation. Ideally, leveraging wizard-driven configuration for the entire solution would simplify the solution implementation process even further.

Pricing is expected to be competitive for entry level configurations and ESG expects Cisco to aggressively position this product to drive market share in this space: Cisco only competes in a space in which it can be a market leader. Including enterprise-class NX-OS features such as Fabric Manager, NPV, VSANs, security, PortChannels, and comprehensive diagnostics in the base price of the switch will certainly be a differentiator. Priced correctly, ESG expects the MDS 9148 will gain significant market traction given the appropriate sales channels.

ESG Lab found that the MDS 9148 is a rock-solid, feature-rich, 8-Gbps multilayer fabric switch that cost effectively delivers the benefits of SAN-attached connectivity to small- and medium-size businesses and enterprises looking for affordably powerful edge connectivity.