Integrating Blade Solutions with EqualLogic™ SANs

Comprehensive Connectivity Strategies

Tony Ansley

Dell Product Group

www.dell.com/equallogic

www.equallogic.com

www.dell.com/blades

Table of Contents

Introduction

In recent years, there has been a major shift in the enterprise data center towards blade server solutions as part of an overall emphasis on consolidation, virtualization, and green initiatives. Blade enclosures offer a number of unique advantages over typical rack-mounted servers that help simplify IT data centers, including:

- Increased density
- Rapid deployment
- Reduced power and cooling requirements
- Reduced cabling
- Simplified management
- Integrated advanced I/O module architecture

Along with server consolidation, administrators are also focusing efforts on storage consolidation. As the migration to blade server solutions takes place, a parallel deployment around consolidated, SAN-based storage solutions is often considered.

Dell™ EqualLogic™ PS Series iSCSI storage arrays offer compelling benefits that make them an ideal companion to blade server deployments. A storage area network (SAN) based on EqualLogic arrays offers a fundamental change in the way enterprises think about purchasing and managing storage. Built on patented peer storage architecture, these solutions offer enterprise-class performance and reliability, intelligent automation, and seamless virtualization of storage for simplified storage management. The PS Series combines an intelligent, automated management framework and a comprehensive set of enterprise data services with a fault-tolerant hardware architecture that supports all major non-mainframe operating systems. It delivers a modular and cost-effective SAN solution that leverages unique scale-out architecture so that it can be deployed in appropriate increments for small and medium businesses, while also being cost-effective for large enterprises that require capacity and high-end performance.

Combining scale-out SAN and blade technologies can be a challenge given that typical I/O module architectures for a blade enclosure have “fan in” capabilities focused on connecting the individual blades within the enclosure, but have limited “fan out” capabilities. Typical Ethernet I/O modules provide anywhere from 4 to 8 ports for connection to networks outside the blade enclosure. This requires some up-front planning for an EqualLogic SAN infrastructure, which by design, is not an integral part of the blade solution. EqualLogic SANs make heavy usage of an Ethernet infrastructure to enable key functionality, such as its storage virtualization and load balancing features. While a complete discussion of the EqualLogic architecture is beyond the scope of this paper, each EqualLogic array will require multiple Ethernet ports on an Ethernet-based SAN infrastructure. Since blade solutions provide limited numbers of external ports, connecting to a SAN infrastructure will require one or more external switches to connect the blade enclosures to the EqualLogic storage solution.

There are several options that administrators have for using the two technologies together. Each has its strengths and weaknesses that require trade-offs either in scalability, performance, or cost. This paper will explore the three major architecture options for connecting a blade enclosure with an EqualLogic SAN, and provide insight into the pros and cons of each.

Multi-Tiered: Maximum Scalability/Maximum Performance

One solution to the scalability challenges presented by any blade solution and scale-out SANs consists of an Ethernet switch infrastructure that is tiered, such that all arrays and a set of external switches are located on one tier and all blades enclosures with their integrated I/O modules are located on a separate tier. These tiers are then connected together in such a way that allows storage requests from the blade servers to the EqualLogic SAN to be processed, and data transferred between the two tiers. Figure 1 below, illustrates the overall SAN architecture.

Note: For clarity, only the connections to the active array controller are shown; ports from the passive array controller should be connected to the array tier switches in a similar fashion.

Figure 1: Scalable SAN Architecture for Blades

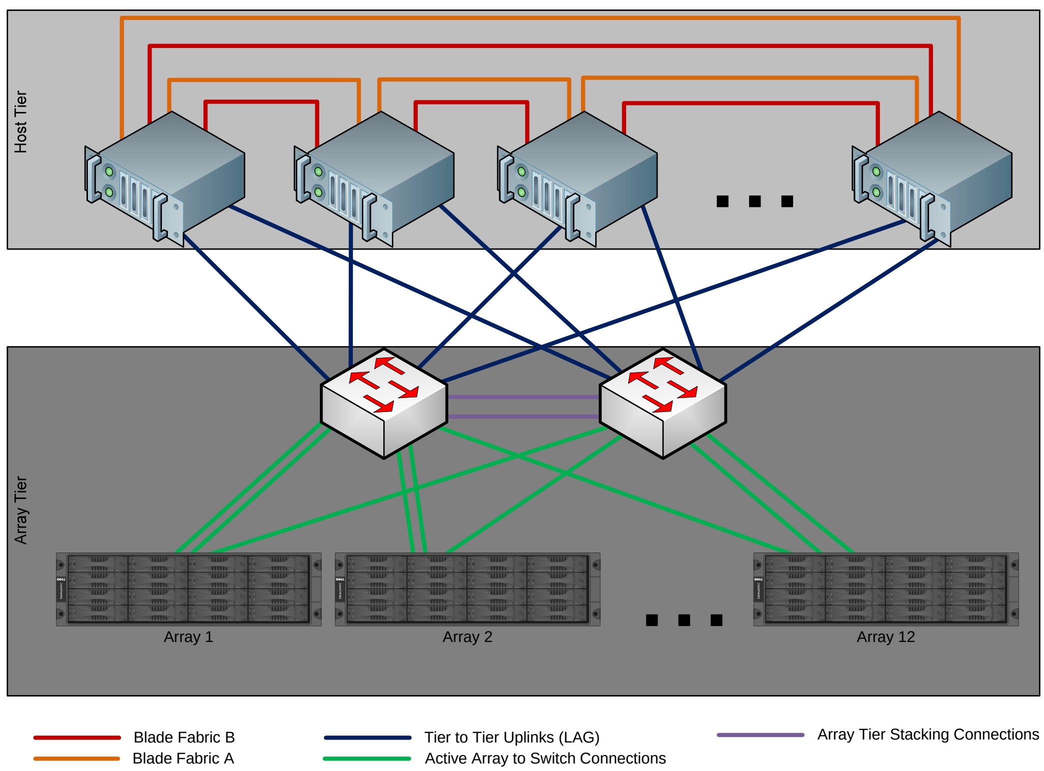

The basic concept of this design will make use of both available stacking ports, and up-linking functionality of the blade enclosure I/O modules and the external switches. Stacking ports will be used to link switches within the same tier, and each tier will then be connected together by connecting the blade enclosure's Ethernet I/O modules in the host tier to the external switches in the array tier with up-linking ports configured into Link Aggregation Groups (LAG) to provide redundancy and bandwidth scalability.

This design strategy has several advantages including:

- Allowing each tier to grow independently
- Allowing an existing Ethernet infrastructure to be used for the array tier
- Each tier can consist of switches from a different vendor than those in the other tier

From a scalability perspective, this architecture allows the administrator to easily add additional blade enclosures to the host tier by inserting the enclosure switches into the stacking loop without bringing the SAN infrastructure down for maintenance. The number of blade enclosures that can be attached to the SAN is only limited by the stacking abilities of the enclosure's Ethernet I/O module design. Likewise, by using stackable switches within the array tier, additional arrays can be added to the existing switches if ports are available, or by adding additional switches to the stack loop in the same way as the host tier.

To support a full 12-array SAN, 36 ports per switch are required to attach the arrays and the ports required for the uplinks between the array tier and the host tier. Depending on the switch designs and the number of ports available for the uplink LAG, it is possible to support a full EqualLogic SAN with just two 48-port stacking switches. It is recommended that all external ports, up to the switches' maximum supported LAG size, be used to link the two tiers to provide the best performance within the SAN solution.

From a performance perspective, this architecture is designed to isolate the SAN management and data balancing traffic inherent in an EqualLogic SAN within the array tier, and leverage the full stacking bus bandwidth of the array tier switches for this additional management traffic. This allows the host tier switches to manage only the traffic generated by specific application requests, potentially removing some of the higher-end feature requirements from the host tier switches such as larger port buffer allotment. The number of arrays that could be supported by the array tier switches is only limited by the stacking infrastructure's bandwidth. With a recommendation of 1 Gbps of stacking bandwidth for each active port on each array, a full 12-array SAN, with 3 active ports per array, would require approximately 36Gbps of total stacking bandwidth between the switches in the array tier. In actuality, depending on the applications performance requirements, some oversubscription may be acceptable. Testing the total solution can determine where stacking link oversubscription begins to negatively impact the application.

Single Tiered with Pass-Through

Another flexible option for integrating blades enclosures with an EqualLogic SAN is to use blade "pass-through" I/O modules. Many blade solution vendors provide a simple pass-through I/O module to allow the network controllers on the individual blade servers to connect directly to external switches. Typically, these pass-through I/O modules are significantly less expensive than a switch I/O module, and depending on the external switches used, could be a lower cost alternative to the dual-tier strategy mentioned earlier.

This solution, shown in Figure 2, has some advantages and disadvantages. As mentioned, depending on the switch vendor used, this could be a lower cost option. Most blade solution vendors provide switch I/O modules from more than one vendor, but the vendor list is almost always limited and typically more expensive than the external equivalent from the same vendor. By using pass-through modules, the administrator has a much broader selection of switch vendors available at a variety of price points. Also, since this solution relies on external switches for SAN connectivity, the administrator can leverage an existing network infrastructure to host the SAN traffic. This allows for easier management of the SAN by integrating the SAN into the same management tools as those used for the general purpose network traffic.

Note: For clarity, the actual connections between the enclosure and the switch have been abstracted. Actual connections from enclosure pass-through ports should be evenly distributed between all switches within the SAN to provide full redundancy. Additionally, only the connections to the active array controller are shown; ports from the passive array controller should be connected to the array tier switches in a similar fashion.

Figure 2: Pass-Through I/O Modules with Single Stack

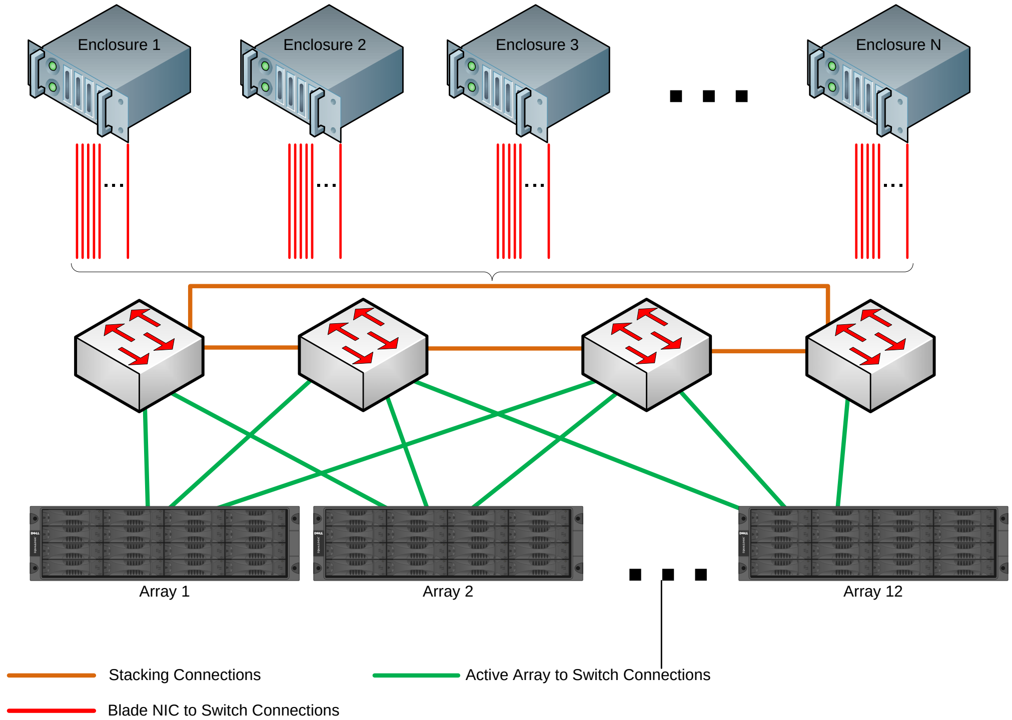

One major disadvantage to this implementation strategy is the number of cables that must be used to connect the individual blades to the external switch infrastructure. This is especially true if the solution needs to be fully redundant with two NICs on each blade. As Figure 2 shows, this solution can scale to the full potential in numbers of blade enclosures and EqualLogic arrays; however, it will require a large number of switch ports and Ethernet cables between each enclosure and the switch infrastructure. For example, a single 16-blade enclosure would require 32 cables coming out of the enclosure and would thus require 32 ports on the SAN switch infrastructure in addition to the ports required for the arrays.

Looking at the scalability of the solution, as with the multi-tiered solution mentioned earlier, the external stackable switches can be expanded as additional blade enclosures or arrays are added to the solution. Due to the large port requirements for each blade enclosure, the SAN infrastructure will typically require more than two switches to support more than a few blade enclosures and a full 12 array SAN.

Single Tier with Stacking

Finally, a third strategy for integrating blade enclosures into an EqualLogic SAN is to rely solely on the blade enclosure I/O modules for the SAN. This architecture, illustrated in Figure 3, allows for a single management interface for network management, similar to the single tier with pass-through strategy, while providing better scalability.

Depending on the vendor's specific I/O module features, the maximum scalability of the SAN will vary. Assuming that the I/O modules are deployed in redundant pairs and that they are stackable and have at least 4 external ports available for array connection, each enclosure could support up to 2 EqualLogic arrays. By utilizing the switch's stacking ports, additional blade enclosures can be integrated into the SAN, and each new blade enclosure would allow for the addition of two more EqualLogic arrays.

Note: For clarity, only the connections to the active array controller are shown; ports from the passive array controller should be connected to the Array Tier switches in a similar fashion.

Figure 3: Using Enclosure Switches for Single Stack

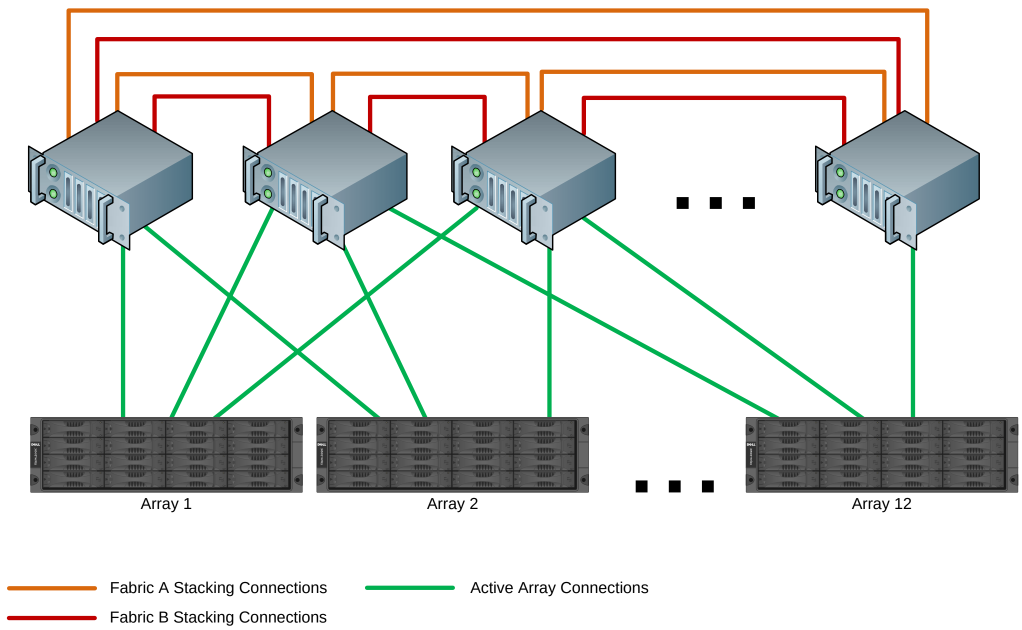

As shown in Figure 3, when additional blade enclosures are added to the solution, the administrator should ensure that the connections from each EqualLogic array are redistributed such that each of the array's active Ethernet ports are connected to up to 3 different blade switches. Also, at least one active port on each array should be attached to a switch in Fabric B if available. Note that Figure 4 displays only the active port connections of each array. An identical connection plan should be used for connecting each array's passive ports to the blade switches. This will ensure that a switch or blade enclosure failure will not cause a drop in connection to any of the arrays in the SAN.

One restriction of this design is the limitation on the number of EqualLogic arrays that can be added to the solution. Since there are a limited number of external ports available in most blade I/O module designs, to scale the storage beyond two or three arrays, the administrator would need to add additional blade enclosures. If the solution does not require a lot of servers (blades) but does require a large storage environment, the administrator should consider using external switches for the SAN infrastructure.

Conclusion

There are many ways that blade enclosures can be attached to a Dell™ EqualLogic™ SAN. This paper has reviewed three ways, but depending on the blade enclosure design features there may be other options. Each of the three architecture strategies presented here provides different trade-offs in cost, complexity, and scalability. The bottom line in designing a blade integration plan for an EqualLogic SAN revolves around a few basic items that the solution should provide:

- Full redundancy
- Enough inter-switch bandwidth to support not only host to SAN traffic, but also inter-array management and load balancing
- Enough I/O bandwidth
- Minimal latency between the host and storage arrays to meet the requirements of the host applications that are attached