



Dell Blueprint for Virtualization: VMware vSphere 6.0 Update 1 on Dell PowerEdge FX2 for up to 1000 Virtual Machines

A Dell Reference Architecture

Dell Global Solutions Engineering
October 2015

A Dell Reference Architecture

Revisions

Date	Description
June 2015	Initial release
August 2015	<ul style="list-style-type: none">• Updated Network architecture guidance from programmable MUX mode to VLT mode• Added Management infrastructure section• Updated sizing information
September 2015	<ul style="list-style-type: none">• Updated LAN network switches to S4048-ON• Updated vSphere 5.5 to vSphere 6.0 Update 1.0• Changed port group load balancing algorithm• Updated out-of-band switch to S3048-ON• Updated solution with ASM 8.1.1• Updated solution Firmware versions• Dell vRanger and DR series appliance for data protection, backup and recovery

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Copyright © 2015 Dell Inc. All rights reserved. Dell and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.



Table of contents

1	Scope	5
2	Audience.....	6
3	Customer Challenges Addressed by This RA.....	7
3.1	Benefits of the Reference Architecture.....	7
4	Introduction	9
4.1	Solution Component Specification in this RA.....	9
4.2	PowerEdge FX2	10
4.3	PowerEdge F630	11
4.4	Dell PowerEdge FN410S	11
4.5	Dell Networking S4048-ON Switch	11
4.6	Dell Storage SC4020 array	11
4.7	Brocade 6510	12
5	Design Principles.....	13
5.1	Single Platform for Any Scale.....	13
5.2	Effective Utilization	14
5.3	Reduction in Management Complexity	14
5.4	Business Continuity and Quality of Service.....	15
6	Solution Architecture	17
6.1	Storage Architecture and Configuration.....	20
6.1.1	Drive and RAID Configuration	20
6.1.2	Additional Storage Enclosures for Increased Capacity	20
6.1.3	Storage Fabric Configuration.....	20
6.1.4	Enabling Redundancy.....	21
6.2	Network Architecture and Configuration	23
6.2.1	Network Design with mVLT	23
6.2.2	Server LAN Configuration.....	24
6.2.3	NPAR and Hypervisor Network Configuration for LAN Traffic	25
6.3	Scaling out Compute and Storage.....	27
6.3.1	Medium-sized Scale-unit	28
6.3.2	Large-sized Scale-unit	30
6.4	Management Components and Architecture.....	33



6.4.1 Dell Active System Manager 8.1.1.....	33
6.4.2 Dell Compellent Enterprise Manager	34
6.4.3 Dell Compellent Replay Manager	34
6.4.4 Dell OpenManage integration for VMware vCenter	34
6.4.5 Dell vRanger Pro.....	34
6.4.6 Dell DR Series Backup Appliance	36
6.4.7 Dell Foglight for Virtualization	36
6.4.8 Management Architecture.....	36
6.5 Data Protection, Replication, and Recovery Architecture	38
7 Solution Specification	40
8 Solution Validation.....	43
8.1 Validated Firmware Revisions of Components.....	43
8.2 Storage Verification	44
9 Solution Summary	47
A Additional Resources.....	48
B Solution Component Details	49
B.1 Dell PowerEdge FX2 Architecture.....	49
B.1.1 Dell PowerEdge FC430 Server	51
B.1.2 Dell PowerEdge FM120x4 Server	51
B.1.3 Dell PowerEdge FC830 Server	52
B.1.4 Dell PowerEdge FC630 Server	53
B.1.5 Dell PowerEdge FN410S.....	54
B.2 Dell Networking S4048-ON Switch	54
B.3 Dell Storage SC4020 Array.....	55
B.4 Brocade 6510.....	55
B.5 VMware vSphere 6.0 Update 1	55



1 Scope

This paper focuses on designing a general purpose virtual infrastructure solution with VMware vSphere 6.0 Update 1 on the Dell PowerEdge FX2 architecture with Dell SC4020 storage array. This paper presents a scalable and highly available infrastructure that is designed according to the best practices and recommendations for virtualizing enterprise applications. It provides detailed guidance on choice of components and configurations and why certain design decisions were made and how they deliver measurable value to customers. The verification of the infrastructure for component failover is also described. Deployment and implementation steps for the proposed solution architecture are outside the scope of this paper.



2 Audience

The audience for this paper includes, but is not limited to, sales engineers (technologists, architects), field consultants, partner engineering team members, customers and anyone else interested in deploying an optimized & validated solutions stack with Dell technology & VMware vSphere.



3 Customer Challenges Addressed by This RA

With ever increasing business demands and growing data volumes, most customers are under unprecedented pressure to improve efficiency and lower costs. However, current operational model of delivering IT services which involves procuring technology from best of breed technology providers can prove not only to be time consuming but problematic for a number of reasons. In this approach, customers are typically burdened to make design decisions, validate various components, set up and configure components manually and manage the environment in an ongoing fashion by engaging multiple vendors for support. All of these elements across the end to end infrastructure life cycle add up to increased complexity and ongoing costs for customers.

This Reference architecture based on PowerEdge FX2 and Dell Storage SC4020 is designed to provide detailed guidance and measurable benefits to customers and help solve challenges mentioned above.

It addresses the following key requirements of customers.

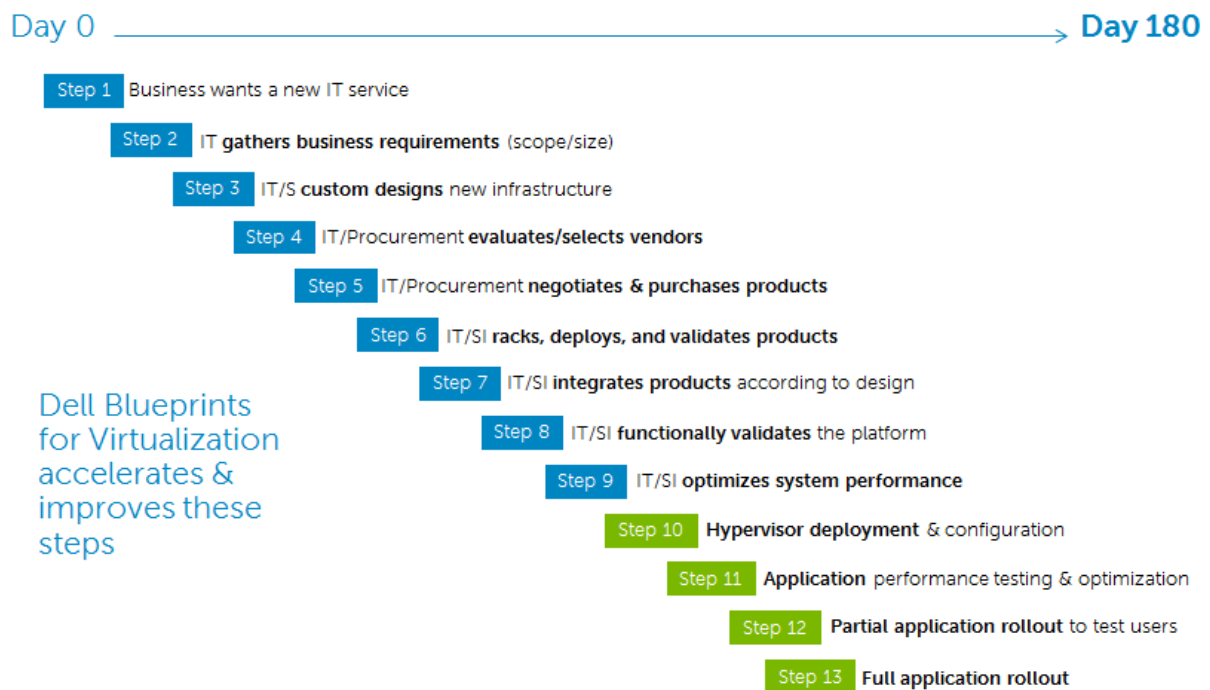
- **Reduce Capital Expenditure (CAPEX):** How do we reduce CAPEX by starting small and growing big?
- **High flexibility and less complexity:** Flexibility in configuration, scaling & in implementation choices based on specific outcomes desired
- **Lower operational costs:** How do we lower our overall TCO by optimizing the infrastructure resources and reducing management complexity?
- **Optimized environment:** How do we deploy and configure an optimized environment based on best practices to deliver the best experience to our end users?
- **Meet business SLA's under all conditions:** How to design, deploy and manage a solution that handles failures without causing disruption?

3.1 Benefits of the Reference Architecture

By offering various design choices and guidance on choosing right components, this Dell reference architecture takes the "guess work" out of solution design, which reduces the enormous time it takes to procure, validate and integrate the components.



Figure 1 How the Dell reference architectures help IT teams improve efficiency



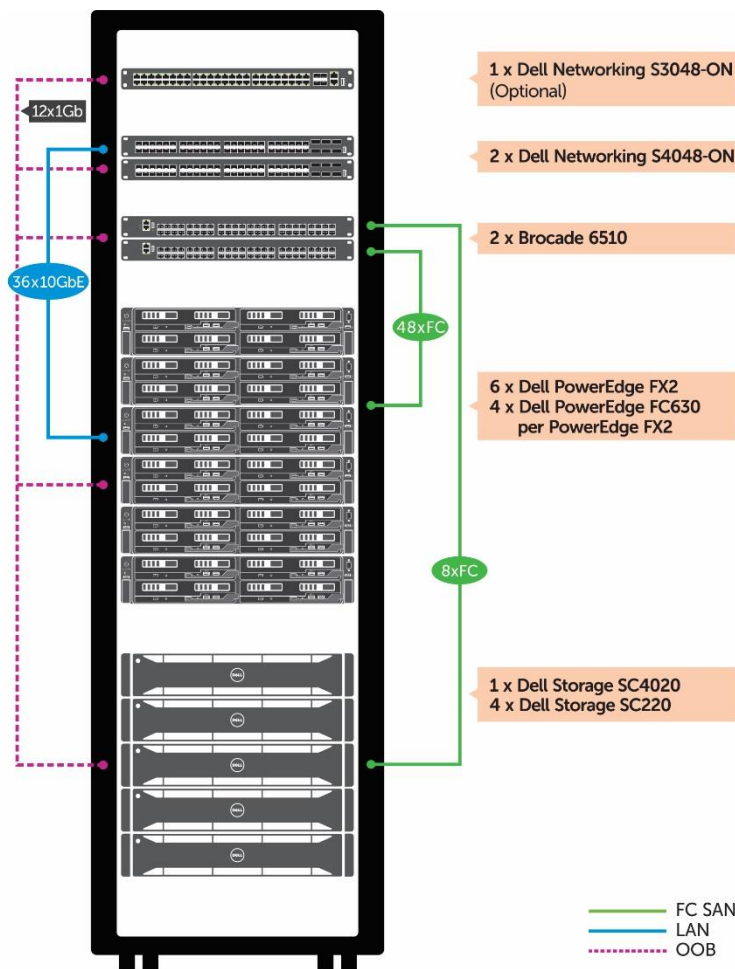
Dell Solution Engineering teams have thoroughly validated and verified the reference architecture solution. The reference architecture is intended to provide

- Design principles that are central to the reference architecture
- Architectural design based on best practices
- Automated deployment and life cycle management of the solution including firmware updates and scaling of physical resources
- Flexibility to address specific customer needs

4 Introduction

This reference architecture presents a virtualized infrastructure that is highly available and scalable. This paper describes some key considerations for designing such a solution. It also provides details of different components used in implementing a virtualized solution built with Dell servers, storage, and networking hardware. Figure 2 shows the rack elevation view of the solution components used in this architecture.

Figure 2 Solution components - Rack elevation



4.1 Solution Component Specification in this RA

The PowerEdge FX2 architecture, Dell Storage SC4020 storage array and Dell Networking S4048-ON switch offer a wide-range of features and configuration options. This enables greater flexibility in choosing the components and configuration that is relevant for the solution architecture being designed. Table 1 provides the solution specification for the components used in this architecture.

The subsections that follow Table 1 provide technical and product overview of components included in this reference architecture.

Table 1 Specification of components used in this reference architecture

Component	Details	
Virtualization Infrastructure	Dell PowerEdge FX2 PowerEdge FN410S; 10GbE SFP+ IOA Qlogic QLE 2562 DP PCIe Add-on FC HBA for storage connectivity	
Virtualization Hosts	Dell PowerEdge FC630 Servers	
	Processor	2 x Intel Xeon E5-2660v3 Family
	Memory	128 GB; 8 x 16GB DDR4 DIMMs
	OS Volume	2 x 16 GB SD cards in internal SD module
	Network	Qlogic 57840 quad-port NDC
	OS	VMware vSphere 6.0 Update 1
Storage Arrays	Dell Storage SC4020 with 6x 400 GB (WI) SSD, 6 x 1.6 TB (RI) SSD and 12 x 10K SAS drives	
Infrastructure Deployment	Dell Active System Manager (ASM) 8.1.1	

The choice of components in Table 1 is supported by the design principles described in Section 5 and then the configuration choices are supported by the solution architecture presented in Section 6.

Appendix B provides details about the components used in this solution.

[Deployment Guide: VMware vSphere 6.0 Update 1 on Dell PowerEdge FX2 With ASM](#) provides step-by-step instructions on how to implement this solution and deploy using ASM 8.1.1

4.2 PowerEdge FX2

The PowerEdge FX2 is a 2U hybrid rack-based computing platform that combines the density and efficiencies of blades with the simplicity and cost benefits of rack-based systems. With an innovative modular design that accommodates IT resource building blocks of various sizes — compute, storage, networking and management — the FX2 enables data centers to construct their infrastructures with greater flexibility. There are two versions of the FX2 chassis — the “switched” configuration, PowerEdge FX2s, supports up to eight low-profile PCI Express® (PCIe) 3.0 expansion slots, and the “unswitched” FX2 configuration, a lower-cost alternative that does not offer expansion slots. The FX2 enclosure also offers I/O modules to several I/O aggregators that can simplify cabling, improve East/West traffic within the server, and enable LAN/ SAN convergence — reducing cost and complexity

Please find more details on FX2 and its architecture in Appendix B.



4.3 PowerEdge F630

An impressively powerful two-socket FX converged architecture server, the Dell PowerEdge FC630 is designed to be a workhorse for data centers looking for new levels of efficiency and density in an incredibly small footprint. Powered by up to two 18-core Intel® Xeon® E5-2600 v3 processors, each FC630 has 24 DIMMs of memory, two 2.5-inch or eight 1.8-inch front access drives, a 10Gb SNA and access to two PCI Express® (PCIe) expansion slots in the shared chassis

Please find more details on FC630 and its architecture in Appendix B.

4.4 Dell PowerEdge FN410S

PowerEdge FX2 supports multiple networking options for the server connectivity to the top-of-rack (ToR) switches and / or to the data center core. Using the IOA configuration simplifies connectivity by as much as 8-to-1, greatly reducing cabling complexity. The FN IOA is a Plug-n-Play networking device providing a low touch experience for server administrators. Most set up and networking functions are automated making for minimal touch for basic to advanced features.

Please find more details on FN410S in Appendix B.

4.5 Dell Networking S4048-ON Switch

Dell Networking S4048-ON is a 1U high-density 10/40 GbE ToR switch with 48 dual-speed 1/10GbE (SFP+) ports and six 40 GbE (QSFP+) uplinks. This switch leverages a non-blocking and cut-through switching architecture to provide ultra-low-latency performance for applications. The six 40 GbE ports can be used to create either a Virtual Link Trunk (VLT) between the switches to enable traffic isolation within the solution infrastructure or to provide connectivity to the data center core network

More details on S4048-ON are included in Appendix B.

4.6 Dell Storage SC4020 array

Dell Storage SC4020 belongs to the Storage Center (SC) 4000 series arrays based on the SC8000 platform. This storage array offers multi-protocol support and virtualized multi-tier, multi-RAID-level storage policies. Each SC4020 array comes with dual redundant controllers, 24 internal drive slots, eight 8 Gb Fibre Channel (FC) or four 10 Gb iSCSI network ports and one 1 Gb port per controller for out-of-band (OOB) management traffic. The supported drive types range from Write-Intensive (WI) Solid State Drives (SSDs) to Read-Intensive (RI) SSDs to 15K, 10K, and 7.2K SAS drives in small form factor (2.5-inch). This storage array supports multi-tier data placement that improves application performance. This array supports expansion up to 192 drives by adding Dell Storage SC220 or SC200 enclosures.

The Storage Center 6.5 Operating System (OS) provides features such as block-level compression, synchronous Live Volumes to restore data in a non-disruptive manner and Active Directory (AD) Single-Sign-On (SSO). The Storage Center Manager enables easier and out-of-box web-based management of



SC4020 arrays while the Enterprise Manager and its components can be leveraged to build a secure, multi-tenant environment with the Dell Storage Center SAN.

4.7 Brocade 6510

Brocade 6510 is a 48-port 1U Gen 5 FC switch that is suitable for high-performance data requirements of server virtualization, cloud and enterprise applications. This switch can be configured in 24, 36 or 48 ports and supports 2, 4, 8, 10 or 16 Gbps speeds. This enables enterprises to start small and scale the FC infrastructure based on the growth, as needed. By using features such as Virtual Fabrics, Quality of Service (QoS) and zoning, IT departments can build multi-tenant cloud environments. Management and diagnostics features such as Monitoring and Alerting Policy Suite (MAPS), Dashboards, Flow Vision, Fabric Performance Impact (FPI) monitoring and Credit Loss Recovery help administrators and IT organizations avoid problems before they impact the SAN operations.



5 Design Principles

While designing this reference architecture based on PowerEdge FX2, Dell Storage SC4020 and Dell Networking S4048-ON, certain specific design tenets have been followed in order to address the customer requirements that are listed above.

5.1 Single Platform for Any Scale

When it comes to infrastructure sizing, rarely do customers know the end state of how much compute and storage they would need. The requirements are tied to business needs and pace of growth, which can change over time. However, customers today are often forced to make a tradeoff between low upfront cost and higher density & performance requirements. They need to make a choice upfront to either buy a small solution that needs a fork-lift upgrade later or buy a large solution and pay for capacity that they do not need right now.

What customers truly need is the flexibility to operate at any scale. This means no tradeoffs on compute between the convenience and flexibility of rack deployments and the density and performance of blade servers or the different levels of performance offered by various storage solutions.

This reference architecture for virtualization uniquely adopts a single infrastructure architecture regardless of scale -- which means no change in core technology or added complexity as the infrastructure grows.

The reference architecture offers customers:

- **Flexibility to start small and grow in smaller chunks.** The reference architecture supports deployments that can start small with a single FX2 chassis with 4 servers and a single storage array SC4020 and then scale them either individually to add more compute nodes (FX2 chassis) or storage enclosures (SC220) or both to support constantly changing workload demands.
- **Reduction of OPEX and operational complexity:** By adopting a single standardizing design across the board and transforming the resource consumption model from an inefficient over-provisioning model to a pay as you grow model, the OPEX and operational complexity can be reduced.

In order to deliver this customer value, the reference architecture has been designed around the concept of a "Scale-Unit". Each scale-unit is a self-contained, balanced "Lego block" where the computer, networking, and storage are built to be sized appropriately for each other based on the customer's initial requirement. Each building block can either be scaled individually or together, based on the customer's ongoing requirements.

Also, each component in the reference architecture has been chosen in support of this design tenet.

- Dell PowerEdge FX2s architecture's incremental, modular approach allows for greater flexibility to address workload needs and allows organizations to budget just for the required resources. It delivers both the flexibility of rack deployments and the density and performance of blade servers.
- SC4020, the storage platform used in the RA, offers a wider range of performance-enhancing features, including true multi-tier flash optimization and allows SC4020 to provide high



performance for the small, medium and large configurations referred in this RA. Customers can easily scale their storage environment by adding more enclosures to address their capacity expansion needs. The performance and scalability of SC4020 is clearly showcased in the large configuration referred to in this reference architecture, where a single SC4020 supports 24 compute nodes without any performance degradation.

- Brocade 6510 was chosen in the RA, given its optimal port density (48 ports) that can support the requirement to start with a single FX2 chassis (4 nodes – 8 ports) and grow to 6 or more chassis (24 nodes-28 ports or more) if required. However, while this reference architecture has been tested with Brocade 6510, in order to obtain more flexibility, customers can choose 6505 series in cases where they don't see the requirement to scale out beyond 4 chassis.

5.2 Effective Utilization

Designing a balanced configuration is important for customers to deliver efficiency and high quality of service without leaving any money on the table. Given that most solutions are conceived and designed in silos, they could often lead to poor utilization. For example, in order to achieve higher consolidation ratios, the solution could be designed as "top heavy", in which server resources are idle while storage resources are oversubscribed -- creating silos of inefficiency.

This Dell reference architecture has been designed as an "eco system" where all technology choices & recommendations are carefully evaluated in context to the overall design.

- In each scenario of different environment sizes of small, medium and large, the recommended configuration optimizes utilization of compute, memory and storage resources. This, coupled with the ability to start small and grow big, helps customers eliminate the need for over provisioning upfront and maximizes utilization.
- The design choice of using FNIOA in the RA over pass-through was to balance performance and utilization. While pass-through forces all traffic out to the ToR switch which is not efficient, FN IOA optimizes "East-West" traffic within the FX2 enclosure, ensuring superior performance and cost savings.

The configuration used for solution validation, including the number of cores and memory, serve only as a reference point for customers to understand the performance parameters for a given configuration. Given that different workloads have a different set of performance characteristics; customers have the complete flexibility to choose a different CPU, memory and capacity configuration as needed.

5.3 Reduction in Management Complexity

A study from IDC shows that customers spend only 30% of their overall IT budget on acquiring infrastructure, but spend 70% on managing it on an on-going basis. A key chunk of the customer's high OPEX is driven by complexity involved both in setting up and managing the virtualized environment. Many of today's tools concentrate on a single element domain, such as servers or VMs and are manual in nature -- leading to slower deployment times and inconsistent end-to-end configurations.

This reference architecture includes specific system management tools to address the above challenge.



- **Automation & life cycle management:** Dell's Active System Manager, which offers a single configuration routine in the form of a template, helps customers to achieve consistent results in an automated fashion. ASM also helps the customer automate the application of the firmware updates across the stack and manage the drift against a baseline configuration. It also allows customers to add compute and storage resources as needed to scale and optimize resource utilization, thereby providing measurable value throughout the end-to-end infrastructure life cycle.
- **Deep monitoring:** Dell's Open Manage Essentials provides a single tool to monitor server, storage and networking and help customers identify the issues and prevent any disruption.
- **Easier configuration:** By using Dell PowerEdge FN410S IOA in VLT mode simplifies the network configuration and management within the infrastructure. VLT mode in FN IOA enables low-touch and automated configuration of upstream port-channel and VLTi links.

5.4 Business Continuity and Quality of Service

Loss of productivity due to unplanned downtime is a grave concern for most IT customers. Dell RA's balanced configuration based on PowerEdge FX2, Dell Storage SC4020 and Dell Networking S4048-ON solves this concern by providing high performance under all operating conditions, including hardware failures. This means comprehensive hardware redundancy both at the topology level and at the element level. The design choices made in the RA ensure that enough resources are available at all layers of the infrastructure, and the performance of the application is not degraded by failures.

- In each scenario of different environment sizes of small, medium and large, the configuration recommended takes into account performance under degraded state, which could include a node failure, SAN controller failure and so on. In each configuration, there is enough room to account for failures, while ensuring that performance of the workload is not impacted.
- If a node fails at the compute level, infrastructure availability provided by the VMware vSphere cluster built on PowerEdge FX2 ensures that the VMs are migrated to other available physical hosts. This ensures that there is little or no disruption to the workloads.
- The reference architecture also includes the use of redundant hypervisor cards with the Dell PowerEdge FX2 platform and helps guard against a single point of failure.
- The network layer is designed to be redundant with two Dell Networking 10GBE switches to eliminate a single point of failure in the solution. FN IOAs are connected together to create a Virtual Link Trunk Interconnect (VLTi). This ensures that the east-west traffic between the compute cluster nodes within the chassis stays in the chassis, and multiple active paths to the top-of-rack (ToR) can exist. Configuring each physical host with multiple network paths ensures that there is a redundant path available even in case of a network switch failure within the solution infrastructure.
- The chosen network architecture with multiple VLT (mVLT) or multiple Link Aggregation Groups (mLAG) using the VLT mode for the FN IOAs in the RA has been chosen to provide redundancy at all layers of the stack. Virtual and physical network connections are further made resilient by using the native NIC teaming feature of VMware vSphere.
- In addition, the data belonging to the virtual infrastructure and virtualized applications should sustain a storage failure. The storage arrays used in this solution architecture provide redundancy at various levels within the array. For example, each of the storage arrays is configured with



redundant power supplies and redundant controllers connected to redundant SAN switches. This design ensures that the component level failures within the array do not impact the virtualized infrastructure and virtualized application availability.



6 Solution Architecture

This section describes the solution architecture in detail and describes how this architecture can be scaled out to increase infrastructure capacity.

The management components for this solution such as Active System Manager (ASM) for deployment automation and VMware vCenter are expected to be part of the existing data center environment and are therefore not represented in any of the solution architecture diagrams. The management infrastructure can be connected either to the customer core network or directly to the Local Area Network (LAN) fabric in this architecture.

Figure 3 illustrates the high-level scalable architecture of the basic building block.

Figure 3 Scalable architecture of the solution

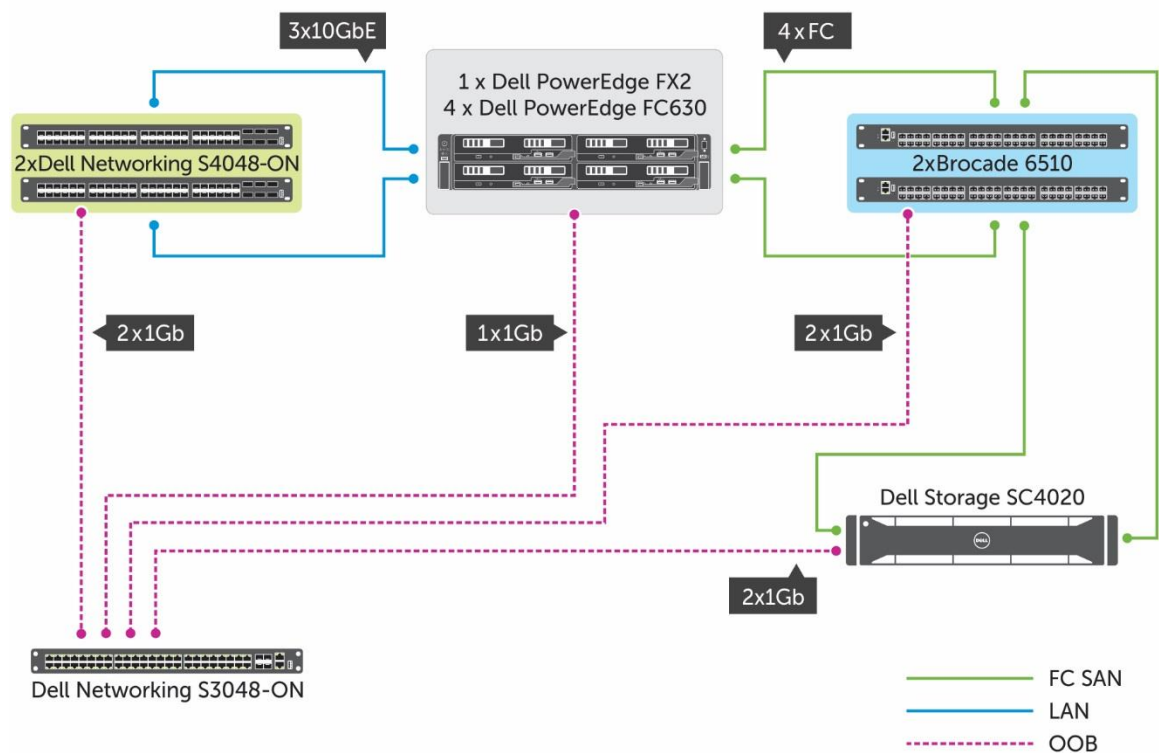


Table 2 describes the compute server configuration in this architecture.



Table 2 Compute configuration in the solution architecture

Resource	Description
Compute Nodes	4 x PowerEdge FC630 in each PowerEdge FX2
Processors	2 x Intel E5-2660v3 family processors in each FC630 server
Memory	128 GB in each FC630
Network	1 x Qlogic 57840 Quad-Port Blade Network Daughter Card (bNDC)
FC Host Bus Adapter	1 x Qlogic QLE2562 Dual-Port 8Gbps FC adapter mapped using PowerEdge FX2 PCIe expansion slots
OS Volume	2 x 16 GB SD cards in the internal SD card module

This solution architecture implements non-converged LAN and SAN fabrics.

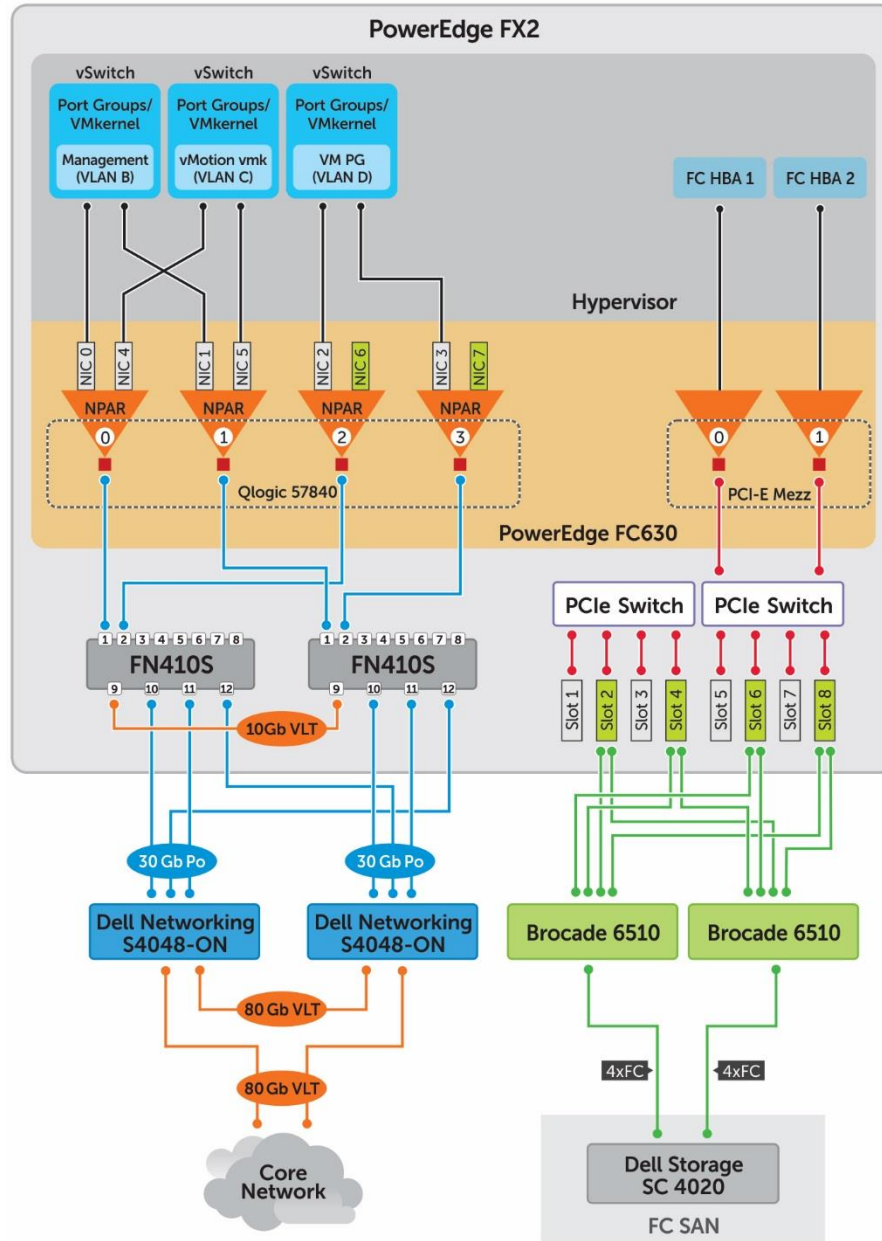
Brocade 6510 FC switches provide connectivity between the Dell Storage SC4020 FC storage arrays and the compute cluster. The storage architecture and configuration are described in detail in Section 6.1.

The end to end IO connectivity represented in Figure 4 also consists of multiple port-channel and VLT configurations for LAN connectivity. Dell Networking S4048-ON switches provide the 10 GbE network connectivity between the compute cluster and the rest of the data center. Section 6.2 describes the network architecture in detail.

Figure 4 illustrates the end to end connectivity from PowerEdge FC630 in slot-1 of the PowerEdge FX2 chassis. This includes connectivity from the LAN and SAN fabric to the VM switch in the hypervisor.



Figure 4 End to end network and Fibre channel connectivity in the solution



6.1 Storage Architecture and Configuration

The following sections provide detailed storage architecture and configuration deployed in this solution.

6.1.1 Drive and RAID Configuration

This reference architecture employs Dell Storage SC4020 array. The array is configured with a combination of Write-Intensive (WI) SSD drives, Read-Intensive (RI) SSD drives and 10K SAS drives to provide approximately 26 TB of raw disk capacity. This hybrid deployment offers both performance and capacity needed for the solution infrastructure. Table 5 provides the configuration details of the SC4020 storage array.

As mentioned in Section 5.1, the choice of hybrid drive configuration enables scaling of compute infrastructure beyond 24 nodes by providing deep IO performance. This performance capability of the hybrid configuration can be complemented by adding more capacity in terms of Dell Storage SC220. Section 6.3 provides guidance on how the compute and storage scale out can be achieved, and the considerations for the scale-out. Table 3 provides the configuration of the Dell Storage SC4020 array used in this solution.

Table 3 Dell Storage SC4020 storage array configuration

Component	Description
Disk Drives	6 x 400 GB WI SSD 6 x 1.6 TB RI SSD 12 x 1.2 TB 10K SAS drives
Storage Profile	RAID 5 with stripe width 9 RAID 6 with stripe width 10
Volumes	2 x 6 TB volumes for VM Store
Replay Profile	Standard Daily Replay Profile with one week expiration

6.1.2 Additional Storage Enclosures for Increased Capacity

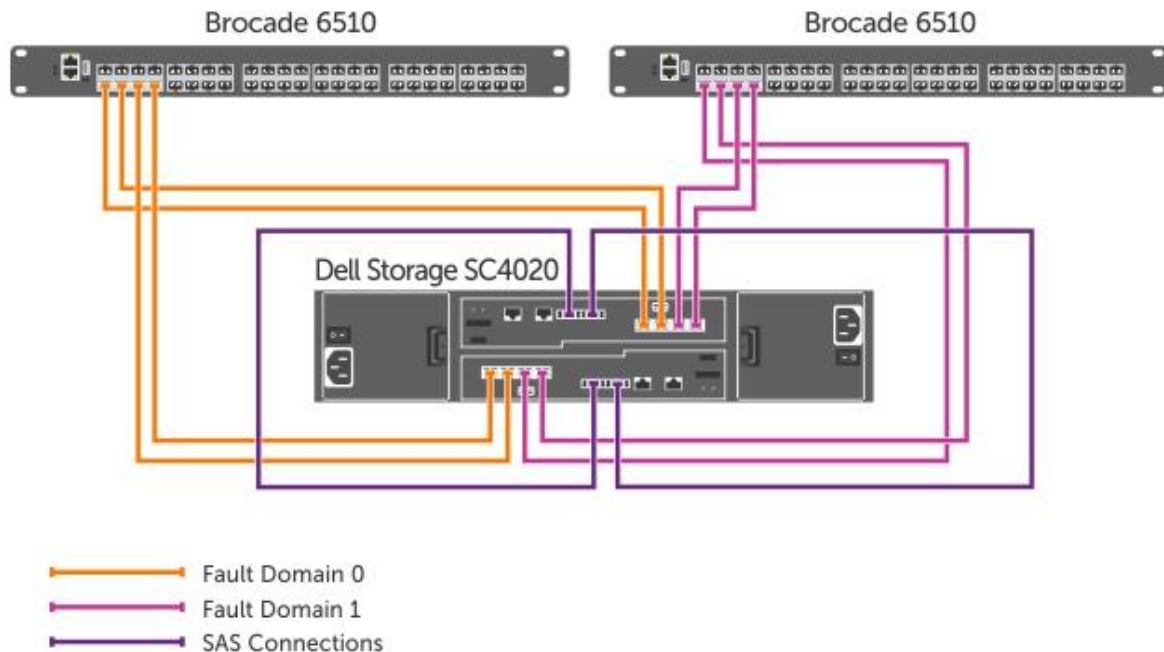
When needed, the overall disk capacity can be expanded, up to a maximum of 192 drives including the 24 drives in SC4020, by adding Dell Storage SC220 or Dell Storage SC200 enclosures. These storage arrays connect to the SAN by using 8 Gbps FC ports on the SC4020 controllers. The SAN fabric uses Brocade 6510 FC switches for connecting the compute cluster and the SC4020 storage array. Figure 5 illustrates how the SC4020 storage array is connected to the Brocade 6510 FC switches.

6.1.3 Storage Fabric Configuration

As shown in Figure 5, ports 1 and 2 from each controller connect to FC switch 1 while ports 3 and 4 connect to FC switch 2. Each FC switch represents a fault domain for the FC storage and enables port failover. Brocade 6510 switches support both 8 Gbps and 16 Gbps SFP+ transceivers. In this solution architecture, 8 Gbps transceivers are used as the SC4020 storage array supports only 8 Gbps FC.

Figure 5 Dell Storage SC4020 connectivity to FC SAN fabric



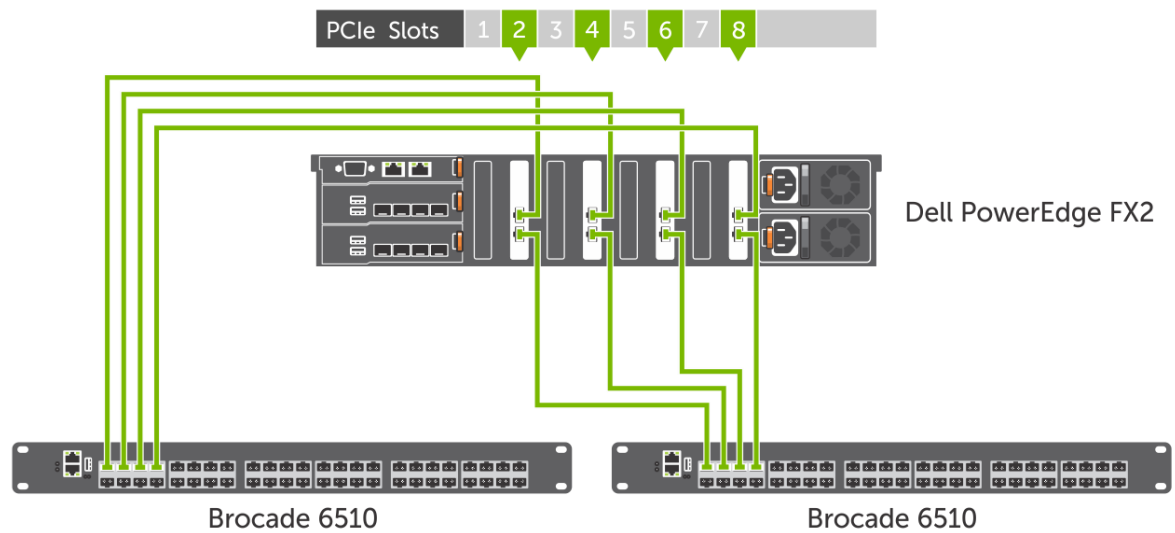


6.1.4 Enabling Redundancy

As shown in Figure 5, fault domains are implemented to maintain continuous connectivity to stored data and the server. Fault domains are established to create redundant I/O paths. These fault domains provide for continuous connectivity with no single point of failure and without loss of bandwidth and provide fault tolerance at the controller level. With Storage Center version 5 and above, the virtual ports feature is used to implement fault domain. In addition to controller level redundancy, virtual ports can also offer port level redundancy. At a minimum, Storage Center requires two ports in the same fault domain on the same controller to support port level failover. Figure 6 illustrates how the FC HBAs in PowerEdge FX2 are connected to the Brocade 6510 switches.

For the compute cluster to connect to the FC fabric, Qlogic QLE2562 FC HBAs are used. The PowerEdge FX2 architecture provides the PCIe slots that are mapped through a PCIe switch to the PowerEdge FC630 servers in the chassis. This mapping is shown in Section B.1. For the compute cluster connectivity, one QLE2562 adapter is used per PowerEdge FC630 server. Figure 6 illustrates how the compute is connected to the SAN fabric for FC connectivity.

Figure 6 Server to SAN fabric connectivity using PowerEdge FX2 PCIe slots



For HA of FC connections, ports from each FC HBA are connected to two different switches. These ports along with the FC connections from the SC4020 array are configured to be in the same zone to enable storage volume access to the compute cluster. Figure 7 shows the FC volumes mapped as datastores on the ESXi hosts within the vSphere 6.0 Update 1 cluster.

Figure 7 FC volumes mapped as data stores in vSphere cluster

ISP2532-based 8Gb Fibre Channel to PCI Express HBA

vmhba3

Fibre Channel

20:00:00:24:ff:41:15:a2 21:00:00:24:ff:41:15:a2

vmhba4

Fibre Channel

20:00:00:24:ff:41:15:a3 21:00:00:24:ff:41:15:a3

Details

vmhba3

Model: ISP2532-based 8Gb Fibre Channel to PCI Express HBA

WWN: 20:00:00:24:ff:41:15:a2 21:00:00:24:ff:41:15:a2

Targets: 4 Devices: 2 Paths: 4

View:

Devices

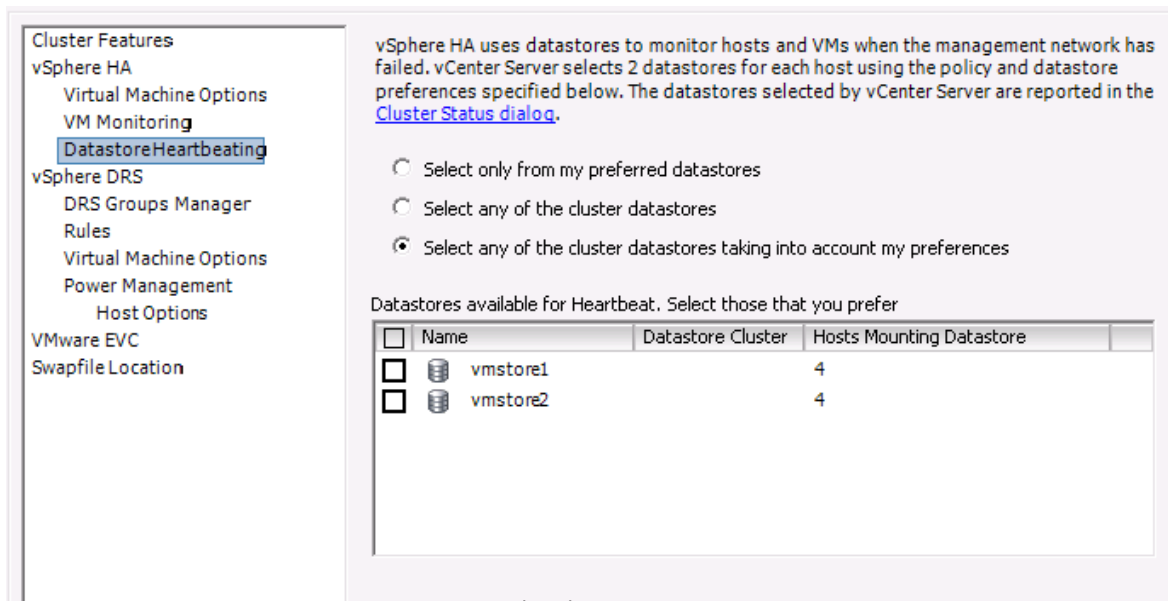
Paths

Identifier	Operational State	LUN	Type	Drive Type	Transport	Capacity	Owner
naa.6000d31000f041000000000000...	Mounted	2	disk	Non-SSD	Fibre Channel	6.00 TB	NMP
naa.6000d31000f041000000000000...	Mounted	1	disk	Non-SSD	Fibre Channel	6.00 TB	NMP

Multiple datastores within the vSphere cluster enable Datastore Heartbeating, which ensures that partitioning or isolated host networks do not trigger VM movement within the cluster. Figure 8 shows the default configuration for the Datastore Heartbeating in the vSphere cluster. By default, the vSphere cluster selects up to five datastores for the Datastore Heartbeating method.



Figure 8 Datastore Heartbeating configuration



Redundant FC switches and datastore configuration in the vSphere cluster ensure that the storage in the infrastructure is highly available. The redundant controllers and power supplies in the storage array along with the fault domains within the Storage Center configuration ensure that the storage can sustain component level failures and provide HA of the data stored.

6.2 Network Architecture and Configuration

The following sections provide detailed network architecture and configuration deployed in this solution.

6.2.1 Network Design with mVLT

The network architecture in this solution employs Virtual Link Trunking (VLT) between the two TOR switches and between the IOAs in the PowerEdge FX2 chassis. In a non-VLT environment, redundancy requires idle equipment which increases switch costs. Besides, the idle equipment adds value only in the event of an equipment failure. On the other hand, in a VLT environment, all paths are active -- utilizing bandwidth and switches to their fullest potential. This doubles the throughput, thus increasing performance and adding immediate value.

VLT technology allows a server or bridge to uplink a single trunk into more than one Dell Networking S4048-ON switch, and to remain unaware of the fact that the single trunk is connected to two different switches. The switches, a VLT-pair, appear as a single switch for a connecting bridge or server. Both links from the bridge network can actively forward and receive traffic. VLT provides a replacement for Spanning Tree Protocol (STP) based networks by providing both redundancy and full bandwidth utilization using multiple active paths.

Major benefits of VLT technology are:

- Dual control plane on the access side that lends resiliency
- Full utilization of the active LAG interfaces
- Rack-level maintenance is hitless and one switch can be kept active at all times.


6.2.2 Server LAN Configuration

Furthermore, the mVLT architecture employed in this solution provides multiple active paths between the compute chassis and the TOR switches. mVLT essentially is a port-channel VLT between two VLT domains. As shown in Figure 9, two VLT domains – one between the IOAs in the PowerEdge FX2 architecture and another between the Dell Networking S4048-ON switches - exist in the network architecture. Implementing mVLT within the network architecture enables a completely loop-free layer 2 (L2) network topology, while ensuring that the East-West traffic within the chassis stays in the chassis. By restricting the East-West traffic to the chassis, the available bandwidth to the TOR switches can be efficiently used. In this architecture, the total available uplink bandwidth from a PowerEdge FX2 chassis to the TOR switches is 60Gbps. For a general-purpose virtual infrastructure, the 60Gbps uplink bandwidth, along with the 10GbE non-blocking switching architecture at the TOR switches provides sufficient bandwidth for the virtualized applications running in this solution infrastructure.

The two Dell PowerEdge FN410S IOA in the PowerEdge FX2 architecture provide the top-of-rack (ToR) connectivity for the PowerEdge FC630 servers. Each IOA provides eight internal server facing ports and four external ports. FN IOA provides the low-touch configuration mode called VLT mode in which Port 9 from each IOA in the chassis form a Virtual Link Trunk Interconnect (VLTi). Ports 10, 11, and 12 form port-channel 128 which in turn connects to the TOR switches. In this architecture, Ports 10 and 11 from FN IOA 1 connect to TOR 1 while Port 12 is connected to TOR 2. On FN IOA 2, Ports 10 and 11 connect to TOR 2 while Port 12 is connected to TOR 1.

Figure 9 illustrates this connectivity.

Figure 9 Server LAN configuration



Since Ports 10, 11 and 12 from each IOA are connected to two different Dell Networking S4810 switches, a VLTi between the S4048-ON switches is mandatory. This provides complete redundancy for the network

traffic from the chassis to the ToR switches. These ports are combined using port channel configuration and enable up to 60 Gbps of bandwidth between PowerEdge FX2 architecture and the ToR switches.

As shown in Figure 9, two 40 GbE ports from each Dell Networking S4810 switch are connected together to create a VLT. This provides up to 80 GbE bandwidth between the ToR switches. The remaining 40 GbE ports on each ToR switch can be used to connect to the data center core network. The ToR configuration in this architecture uses the management ports on both the switches for the VLT backup link.

6.2.3 NPAR and Hypervisor Network Configuration for LAN Traffic

Each Dell PowerEdge FC630 server in the PowerEdge FX2 architecture provides four 10 GbE ports through a Qlogic 57840 bNDC. To ensure the bandwidth prioritization for different traffic classes such as host management, vMotion and VM network, switch independent network partitioning (NPAR) is configured. By using NPAR, each port on the bNDC is partitioned into two logical partitions. The Qlogic blade NDC adapters allow setting a maximum bandwidth limitation to each partition. Setting maximum bandwidth at 100 will prevent the artificial capping of any individual traffic type during periods of non-contention. For customers with specific requirements, NPAR maximum bandwidth settings may be modified to limit the maximum bandwidth available to a specific traffic type, regardless of contention. These adapters also allow relative bandwidth assignments for each partition. Table 4 shows the NPAR schema and the relative bandwidth assignments on each partition.

Table 4 NPAR schema and relative bandwidth weights

bNDC Port	Network Partition Label	Traffic Class	Relative Bandwidth	Maximum Bandwidth
Port 0	NIC 0	Host Management	30	100
	NIC 4	vMotion	70	
Port 1	NIC 1	Host Management	30	
	NIC 5	vMotion	70	
Port 2	NIC 2	VM Network	40	
	NIC 6	NA	60	
Port 3	NIC 3	VM Network	40	
	NIC 7	NA	60	

Figure 10 illustrates how NPAR is deployed on each physical host. As shown in Figure 10, network partitions NIC6 and NIC7 are not deployed to carry any host or VM network traffic. The relative bandwidth constraints are used only when there is a need for a traffic class to be given priority. For isolating different traffic classes from a host, Virtual LAN (VLAN) is deployed. The virtual LAN configuration is needed on the FN IOA and TOR switches to ensure end to end connectivity.



Figure 10 Host LAN configuration using NPAR

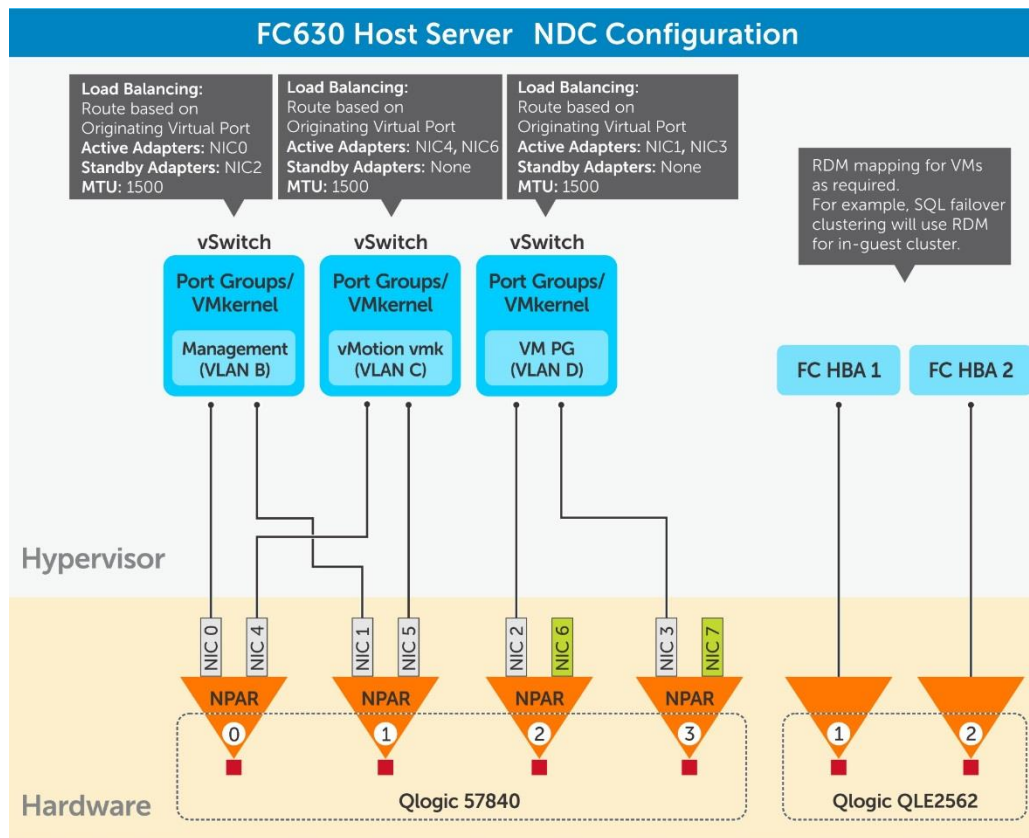
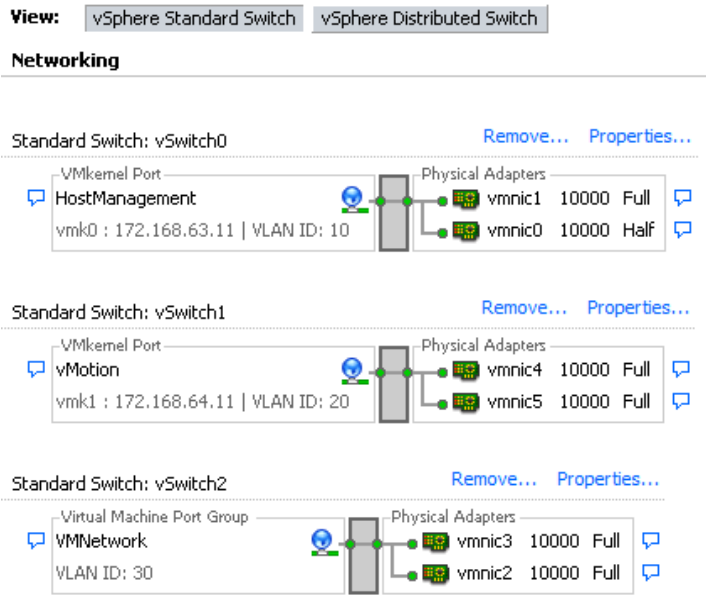


Figure 11 shows the host-level implementation of virtual switches as illustrated in Figure 10.

Figure 11 vSwitch implementation on every physical host



For each vSwitch that is created on the ESXi host, the default load balancing algorithm “Route based on originating virtual port” is used. In this setting, any single virtual network adapter will use only one physical adapter port at a time. This option is selected as it is easy to configure and provides good load balancing across VMs, especially in the case of a large number of VMs.

The Dell Networking S3048-ON 1 Gb switch in this architecture is an optional component and can be replaced by an existing out-of-band (OOB) management switch in the customer’s existing data center environment.

The network architecture design described in this section implements the key design principles and considerations described in Section 4 to provide highly available LAN fabric that is optimized for virtual and enterprise application traffic.

6.3 Scaling out Compute and Storage

The architecture shown in Figure 3 represents the basic building block for a scalable and highly available virtual infrastructure that is built by using VMware vSphere. This architecture can be scaled out to expand the overall capacity of the solution. This can be done either by adding compute or storage or both. For expanding the storage, the Dell Storage SC220 or SC200 array can be attached to the Dell Storage SC4020 arrays. The basic building block shown in Figure 5 can be considered a small scale-unit of VMware vSphere 6.0 Update 1 on PowerEdge FX2.

The following sections provide details on the small, medium and large deployments of the solution architecture and how the small building block can be scaled out to medium and large-sized deployments by adding compute and/or storage components. The decision to add compute and/or storage depends on the need that arises by virtualizing applications on this infrastructure. When there is a need for additional storage capacity for the virtualized applications in this smaller deployment, the architecture has

the flexibility to add just the storage. This can be achieved by adding a Dell Storage SC220 enclosure to the existing Dell Storage SC4020. When needed, the compute and storage can be added without the need to add more LAN or SAN fabric switches.

This reference architecture for VMware vSphere 6.0 Update 1 on PowerEdge FX2 classifies the size of the scale-unit based on how much compute and storage capacity is available. Table 5 provides this information.

Table 5 Different scale-unit sizes for the virtual infrastructure

Scale-unit size	Number of compute	Storage
Small	4 x PowerEdge FC630	1 x Dell Storage SC4020
Medium	8 x PowerEdge FC630	1 x Dell Storage SC4020 with 1 x Dell Storage SC220
Large	24 x PowerEdge FC630	1 x Dell Storage SC4020 with 4 x Dell Storage SC220

Each deployment size described in Table 7 assumes that there is additional disk capacity needed with scaling out of compute. Therefore, each scale-unit includes additional Dell Storage SC220. The following sections provide examples of how the small scale-unit shown in Figure 3 can be scaled to medium and large VMware vSphere deployments.

6.3.1 Medium-sized Scale-unit

A medium-sized scale-unit of VMware vSphere 6.0 Update 1 on PowerEdge FX2 using the components described in Section 4 can be built by adding a second PowerEdge FX2 architecture with four PowerEdge FC630s and one Dell Storage SC220 enclosure for additional storage capacity to the small building block shown in Figure 3.

Figure 12 Scale out of the solution (medium-sized) to add compute and storage capacity



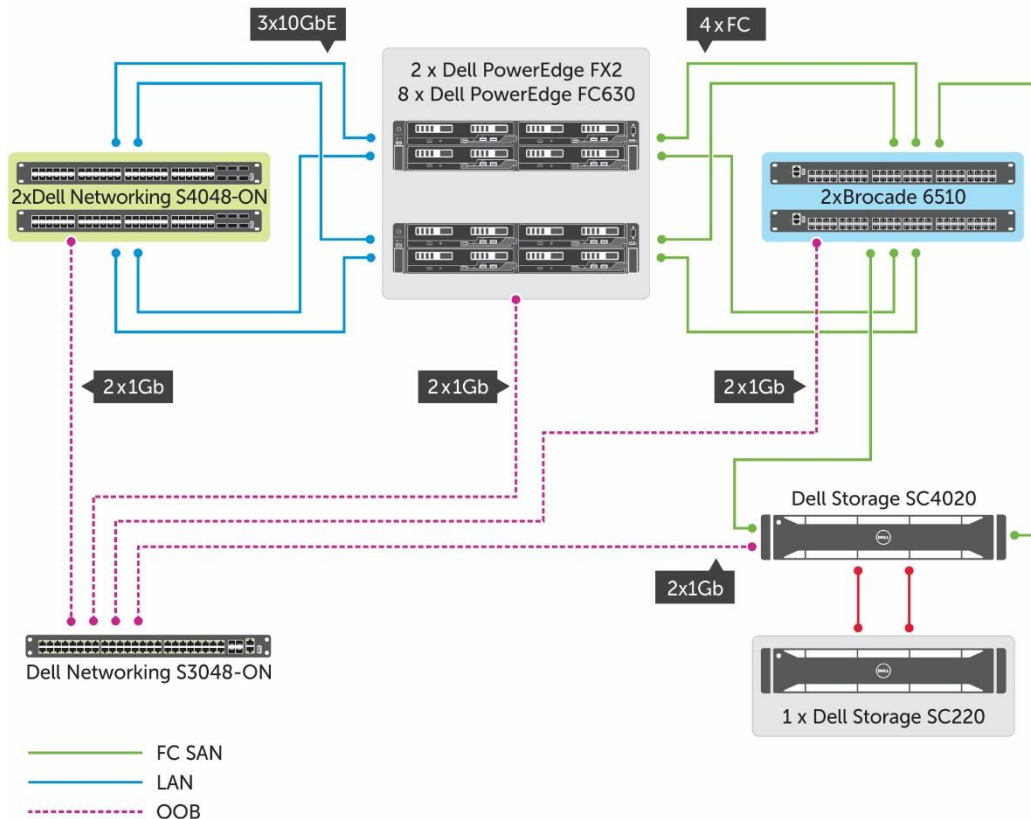
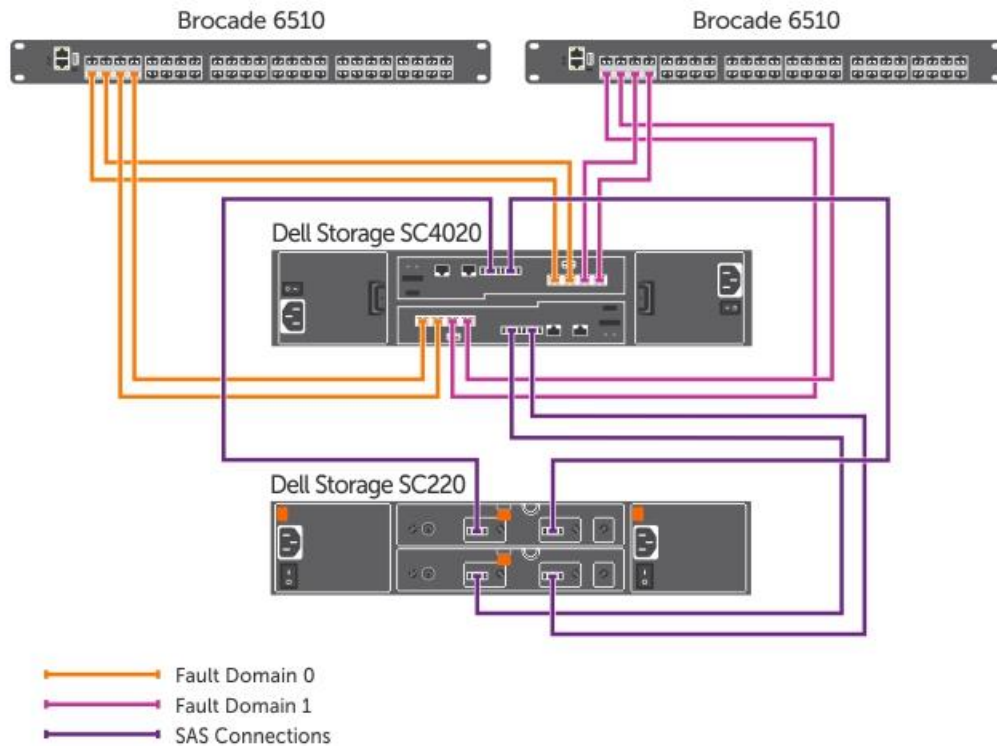


Figure 12 shows the high-level architecture of the solution with two PowerEdge FX2 architecture and a Dell Storage SC220 enclosure added to the SC4020 array. The drives in the Dell Storage SC220 enclosure are added to the same disk folder as the drives in the SC4020 array. This enables storage capacity expansion for the virtual infrastructure.

Figure 13 illustrates the storage connectivity in a medium-sized deployment with two PowerEdge FX2 architecture. A Dell Storage SC220 enclosure gets connected to the SC4020 storage array using the SAS cabling to the controllers.

Figure 13 Storage expansion for increasing the overall infrastructure storage capacity (medium-sized)



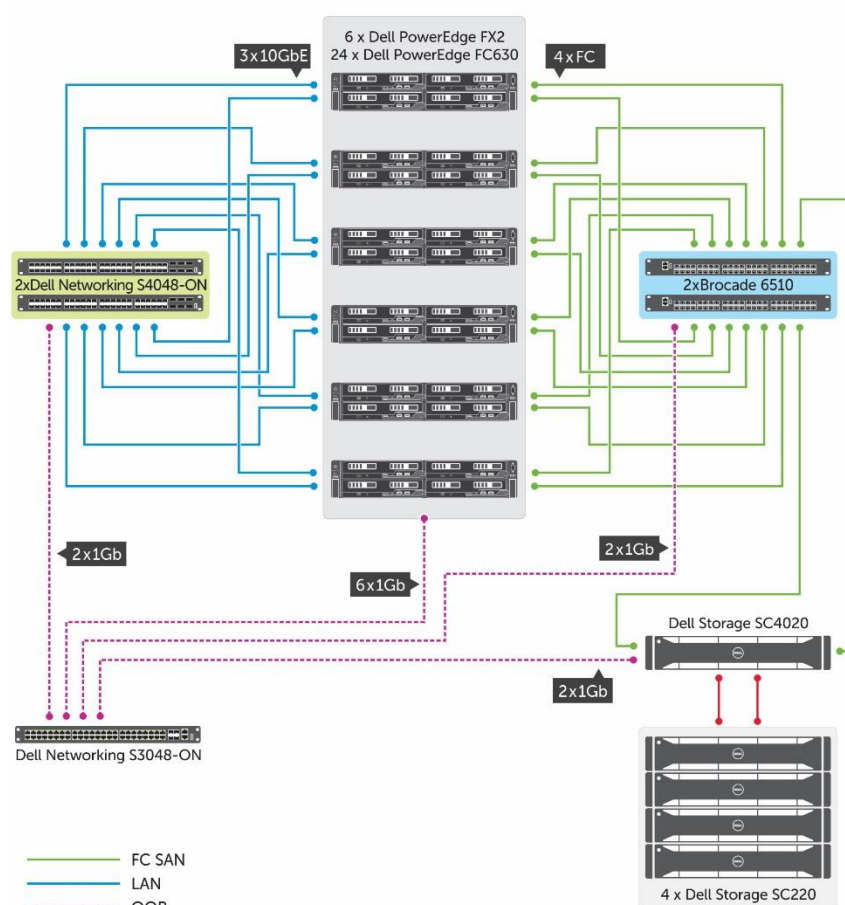
This scale out does not need any additional switches to be added to the LAN or SAN fabrics. Dell Networking S4048-ON switches provide enough ports for PowerEdge FX2 architecture uplinks. By default, Brocade 6510 switches are configured with 24 port licenses and therefore need no further investment.

6.3.2 Large-sized Scale-unit

Similar to a medium-sized infrastructure for VMware vSphere 6.0 Update 1 on PowerEdge FX2, a large-sized deployment can be implemented by adding six PowerEdge FX2 architecture with four PowerEdge FC630s and four Dell Storage SC220 enclosures for additional storage capacity.

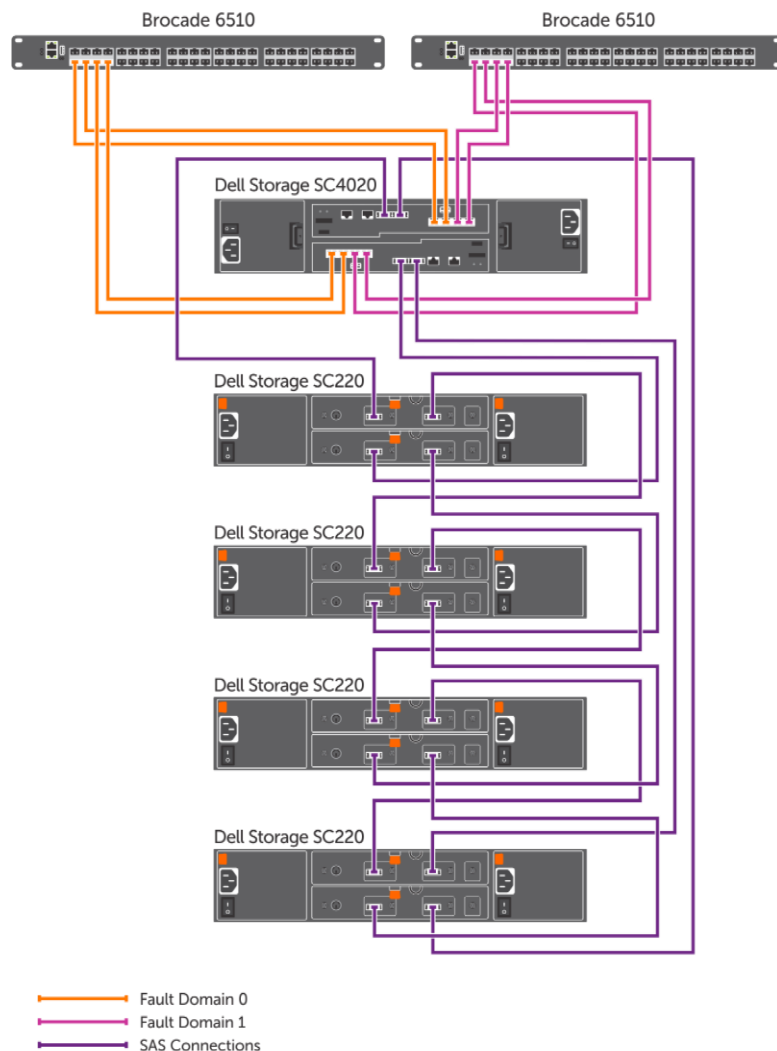
Figure 14 shows the high-level architecture of the solution for the large-sized VMware vSphere 6.0 Update 1 solution on PowerEdge FX2.

Figure 14 Scale out of the solution (medium-sized) to add compute and storage capacity



The drives in the Dell Storage SC220 enclosure are added to the same disk folder as the drives in the SC4020 array. This enables storage capacity expansion for the virtual infrastructure. Figure 16 illustrates the storage connectivity in a large-sized deployment with six PowerEdge FX2 architecture. The Dell Storage SC220 enclosures get connected to the SC4020 storage array using the SAS cabling to the controllers. Figure 15 shows the SAS cabling changes needed for the expansion.

Figure 15 Storage expansion for increasing the overall infrastructure storage capacity (large-sized)



The scale out does not need any additional switches to be added to the LAN or SAN fabrics. Dell Networking S4048-ON switches provide enough ports for PowerEdge FX2 architecture uplinks. By default, each Brocade 6510 switch is configured with 24 port licenses and the large-sized solution expansion requires four additional ports per switch. This port expansion can be performed by adding a port on demand (POD) license on each switch.

While this reference architecture presents a six PowerEdge FX2 based scale-unit as a large deployment, the overall solution has enough room to grow beyond this capacity. The Dell Networking S4048-ON and Brocade 6510 switches provide the port density to add more compute chassis. However, a single VMware vSphere 6.0 Update 1 cluster cannot include more than 32 nodes. Therefore, after scaling-out the architecture shown in Figure 14, it is ideal to add a small scale-unit for any further expansion.

6.4 Management Components and Architecture

In an infrastructure solution that is built with multiple components such as the architecture shown in this paper, it is essential to ensure that there are management tools available for infrastructure administrators to monitor and manage the life cycle of the components.

Solution life cycle consists of different phases such as deployment, monitoring and operations and backup and data protection. The management infrastructure components chosen should cover each of the phases involved in the solution life cycle management. Dell has a wide-range of management tools that help administrator manage the solution life cycle and for each component within the solution infrastructure.

This section describes some of these management components and provides details on how these management components can be integrated into the solution architecture shown in this paper.

The components mentioned in the following subsections is not a complete list of management capabilities offered by Dell. For a list of management capabilities that are appropriate for a specific environment, contact a Dell Services representative or Dell channel partner.

6.4.1 Dell Active System Manager 8.1.1

Active System Manager (ASM) is Dell's unified management and automation solution that simplifies the deployment, provisioning, and management of heterogeneous resources and workloads - for private cloud or any environment of shared or converged infrastructure.

ASM 8.1.1 takes a top-down service-centric approach to IT automation that spans across both physical and virtual infrastructure and encompasses everything from servers, switches, and storage, to hypervisors and clusters, to virtual machines, operating systems, and applications, ASM provides a unified and comprehensive user experience for the initial deployment and the ongoing lifecycle management of shared or converged infrastructure.

ASM 8.1.1 release of Dell's unified system management and automation platform, which enables IT administrators to accelerate service delivery and improve efficiency in a number of ways, including the ability to:

- Get up and running quickly, with enhanced discovery and initial configuration, and wizard-driven automation of server, chassis, and IO network component onboarding.
- Efficiently manage infrastructure lifecycle, with comprehensive firmware management, compliance monitoring, and updates.
- Easily deploy, manage, and scale IT services, with automated provisioning, scaling of resources up and down, and decommissioning of services to recover resources.
- Define and manage user access levels and permissions, with role based access control for users and administrators.
- Support the new 13th generation PowerEdge servers, as well as other compatibility list additions.



6.4.2 Dell Compellent Enterprise Manager

The Storage Center web interface for managing Dell Storage arrays provides the features for basic management of Storage Center arrays. Enterprise Manager extends this capability to deliver centralized management of all local and remote Storage Center environments and tight integration with leading technology providers such as VMware. By deploying Dell Compellent's Site Recovery Adapter (SRM) for VMware Site Recovery Manager, IT administrators can take advantage of the benefits provided by array-based replication over the classic SRM replication. Enterprise manager provides the storage-based chargeback feature that automatically calculates storage costs based on the actual space consumed by applications.

6.4.3 Dell Compellent Replay Manager

Replay Manager provides a VSS-integrated solution for creating application-aware Replays (snapshots) of Microsoft Exchange and SQL Server. Replay Manager leverages VSS to help ensure the integrity of application data spanning multiple volumes and streamlines Replay management with PowerShell integration and provides the ability to manage multiple server instances through a single interface. Replay Manager integrates with Data Instant Replay to ensure data remains consistent even if Microsoft Exchange or SQL Server is running as a snapshot or Replay is captured. With Replay Manager, administrators have the ability to take Replays of Windows Server volumes without having to take Windows Server applications offline.

6.4.4 Dell OpenManage integration for VMware vCenter

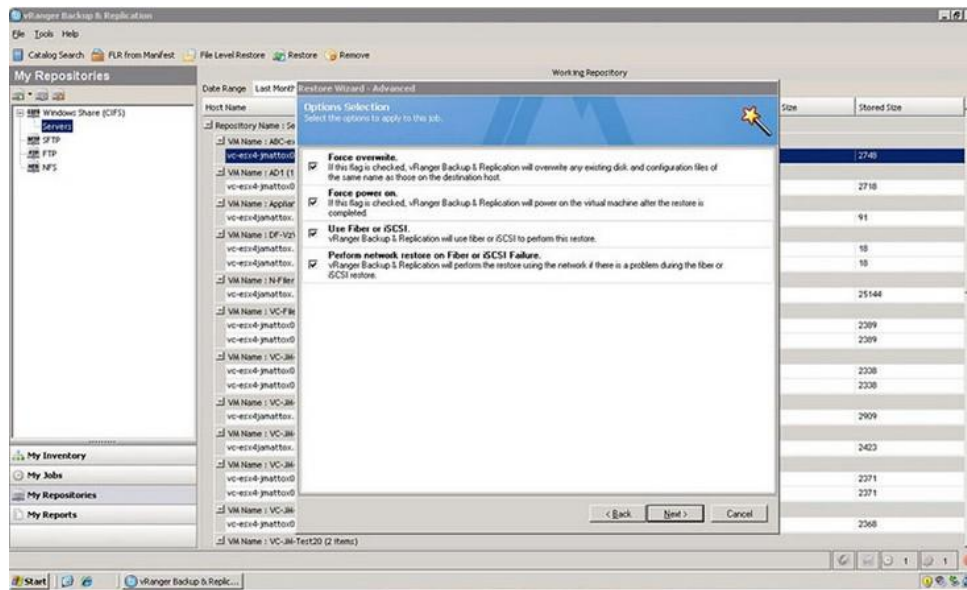
The Dell OpenManage Integration for VMware vCenter (OMIVV) provides the ability to monitor, manage, configure and deploy Dell PowerEdge Servers in the VMware vSphere infrastructure. By providing unique cluster-level hardware views directly within vCenter, the OpenManage Integration for VMware vCenter plug-in enables easy scaling and a streamlined process for applying updates to multiple Dell hosts in a single workflow. In addition, by using the six-step wizard tool, administrators can automatically and reliably update all servers in a cluster without losing any workload productivity. This integration enables a holistic view of the status of all PowerEdge FX2 architecture and PowerEdge FX630 servers in the virtual infrastructure. This includes the health and alert information to warranty information for the chassis and servers in the infrastructure.

6.4.5 Dell vRanger Pro

Dell vRanger software provides high-speed backup and recovery of virtual infrastructures built using VMware and physical Windows Server environments. It also provides high-speed replication for VMware. It protects entire virtual environments in minutes, detecting and backing up new VMs automatically, and delivers safe, scalable data protection to even the largest VMware and Hyper-V environments. With vRanger, you can locate and restore individual files in seconds — even if they are buried in virtual and physical backups — from a single, intuitive interface. Dell vRanger when combined with Dell Storage DR appliances provides complete de-duplication capabilities for the virtual infrastructure protection.



Figure 16 vRanger Pro



vRanger provides high-speed, resource-efficient backup, replication and recovery of virtual machine images, and supports backup and recovery of Windows physical servers, files and folders. In addition, vRanger delivers maximum storage savings when paired with the Dell DR deduplication appliance.

Key features and benefits of vRanger Pro include:

- Change block tracking (CBT) — eliminates the time required to scan for changed blocks in guest images on vSphere hypervisor systems which speeds backup and replication jobs.
- Instant file-level recovery (FLR) for Windows and Linux — Enables you to quickly restore a single file from a backup image in the repository through a one step process.
- Native, full catalog capability — Provides native, full catalog of every image in the backup repository, enabling immediate identification of available recovery positions. The wildcard scanning feature quickly locates backup repository files to be restored.
- One-step catalog recovery — Provides advanced search (including wildcards) and right-click recovery selection directly from the management console to speed up restore of VMs, savepoints, and hosts with native catalog.
- Advanced encryption standard (AES)-256 — Secures protected images block-by-block on the VMware host as they are read so they are also secure over the network and in the backup repository.
- Full, incremental, and differential backup — Enables a complete backup cycle for protected images that is optimized for speed and resource efficiency.
- Full, incremental, and hybrid replication — Provides the full range of options required to efficiently replicate VMware VMs over LANs and WANs.
- vSphere vMotion support — Automatically protects VMs as they move from one host to another, even when backup jobs are running.



6.4.6 Dell DR Series Backup Appliance

The Dell DR Series Disk Backup and Deduplication Appliance is a physical appliance, but it works in close conjunction with vRanger Pro to deliver maximum performance and functionality.

The DR Series systems are extremely efficient, high performance, disk-based backup and recovery appliances. The DR Series systems are simple to deploy and manage and offer unsurpassed total cost of ownership (TCO) benefits. Figure 17 shows a Dell DR Series Backup Appliance.

Figure 17 DR Series Backup Appliance



Through the use of innovative Dell deduplication and compression technology, the DR Series systems can help achieve data reduction levels up to 15:1. This reduction in data means that more backup data can be retained longer and within the same footprint.

6.4.7 Dell Foglight for Virtualization

Dell Foglight for virtualization provides insights into the health, risk and efficiency of virtual infrastructure and includes change tracking and performance impact analysis. Using Foglight, administrators can quickly identify bottlenecks and analyze network traffic at the packet level. Foglight provides a unified dashboard for overall health of the virtual infrastructure, current resource utilization and any issues that need an administrator's attention. Using Foglight, administrators and IT managers get the visibility into cost across virtualization operations. Administrators can easily chargeback cost centers based on the resource utilization at a department or business level.

6.4.8 Management Architecture

Implementing the infrastructure that manages the overall solution life cycle must be carefully planned and deployed. The components described in the above subsections represent only a subset of Dell's solution life cycle management capabilities but provide the necessary core building blocks for managing the components of the solution architecture described in this paper.

All management components described in this section can be deployed as virtual machines and these components can share the storage capacity available on the Dell Storage SC4020 array used for the virtual infrastructure. Similarly, it is recommended to deploy the same host networking configuration as described in section 6.2 on the hypervisor hosts used for deploying the management components. By doing so, administrators can leverage the existing network configuration and infrastructure. The Dell Network S4048-ON switches used in the solution architecture provide the necessary port density to connect the management infrastructure.



When virtualizing these management components, administrators must ensure that infrastructure is highly available. For storage and network, the underlying virtual infrastructure provides HA. By deploying at least two physical hosts outside the compute used for the virtual infrastructure is recommended for enabling complete HA of management components. The choice of physical servers and their configuration for the management infrastructure depends on the number of management components being deployed. PowerEdge R630 servers provide the right balance of cost and performance for virtualizing management components.

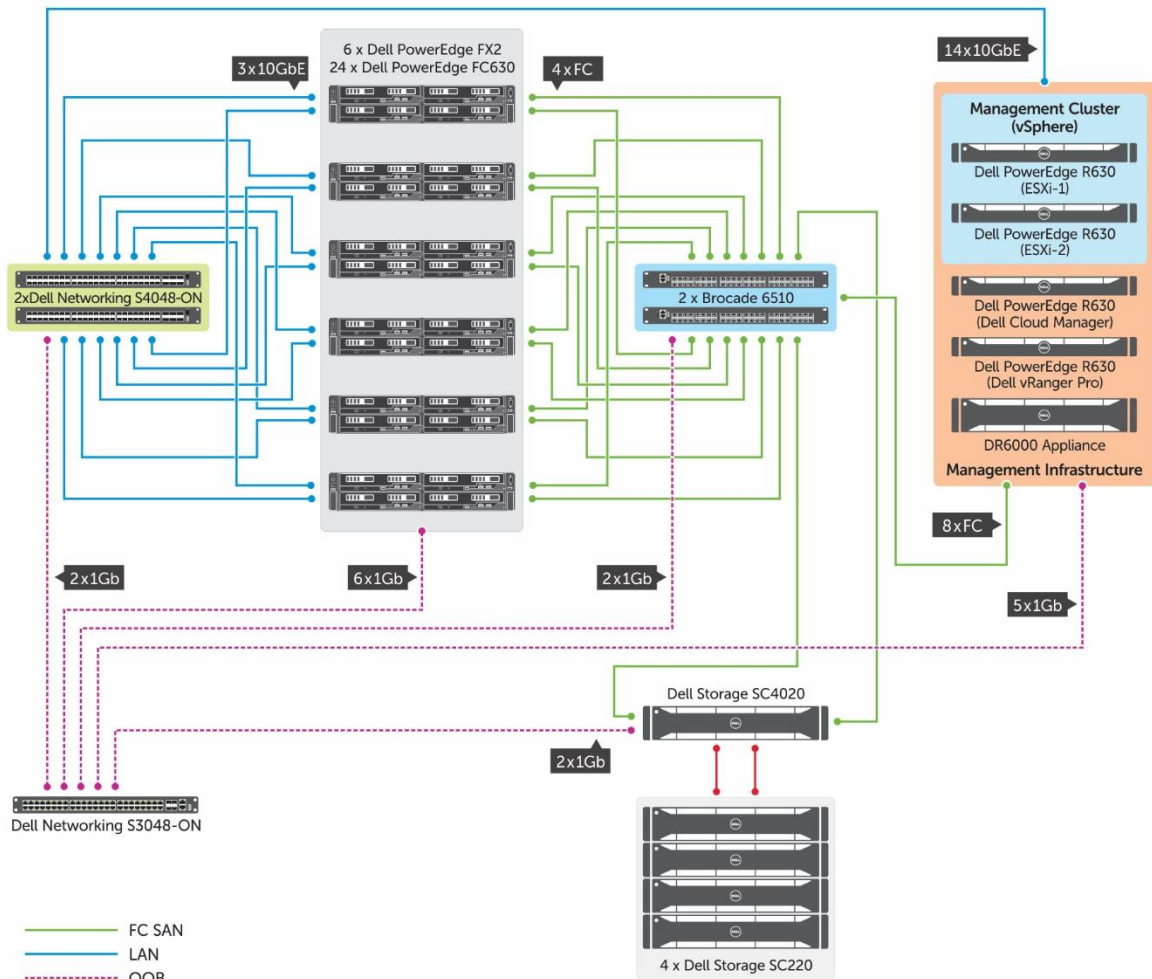
The out of band (OOB) capabilities of Dell PowerEdge Servers, Dell Storage arrays and Dell Networking enable separation of infrastructure solution component management traffic from the application or host management traffic. All management virtual machines such as ASM and FogLight communicate with the infrastructure components either via the OOB network or the host management network.

The management architecture for the solution is shown in Figure 16.

By integrating Dell vRanger Pro with Dell DR6000 backup and recovery appliance, the overall data protection performance is optimized. By using the SAN transport provided by VDDK, the data path for backup and recovery traffic does not depend on LAN anymore. It uses the Fibre Channel connectivity between the storage arrays and the Dell vRanger management server.

The management architecture for the solution is shown in Figure 18

Figure 18 FX2 Management Architecture



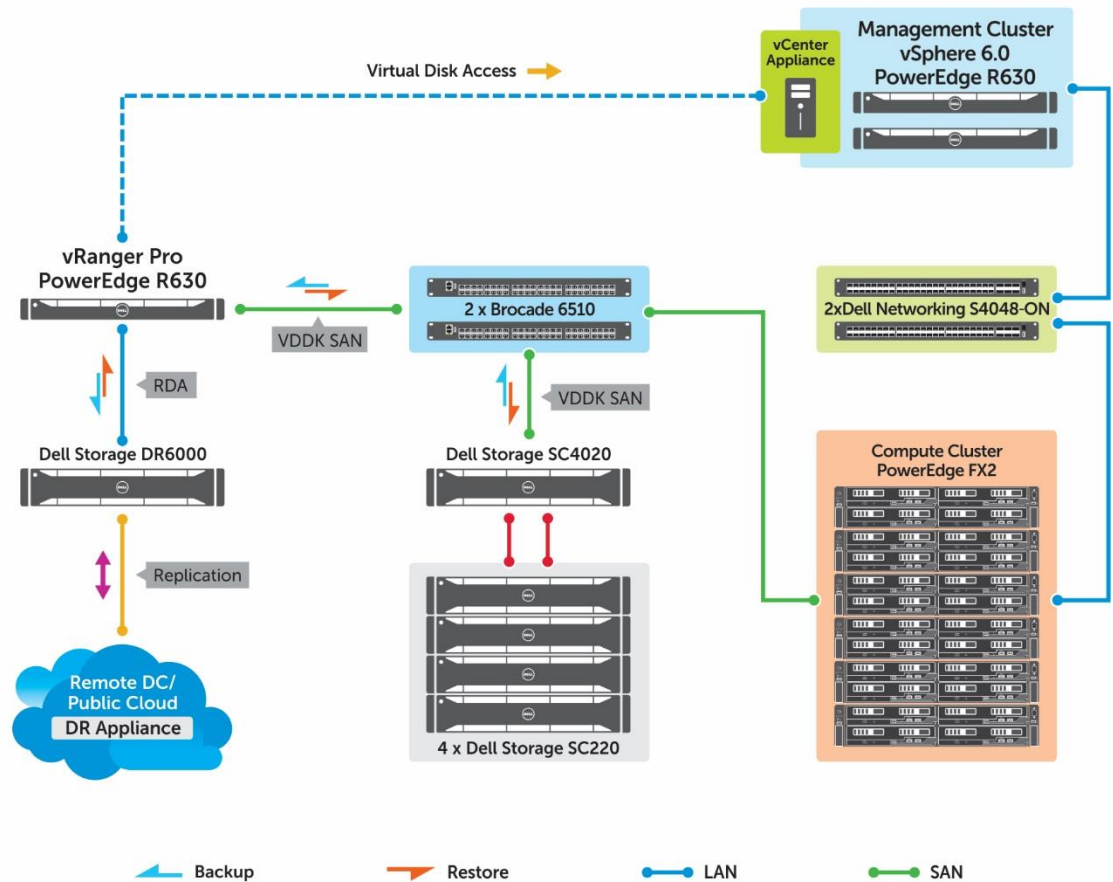
6.5 Data Protection, Replication, and Recovery Architecture

The Dell Data Protection solution for the Dell PowerEdge FX2 converged system uses the DR6000 appliance with Rapid Data Access (RDA) inline deduplication technology and the vSphere VDDK SAN Transport. The combination of RDA inline deduplication and VDDK SAN transport provides a premium data movement service that is capable of satisfying demanding backup requirements of workload at scale. Additionally, leveraging SAN for backup transport avoids taxing virtualization resources for backup purposes thereby conserving hypervisor capacity for production workloads.

For Backup, the solution leverages the deep VMware integration of Dell vRanger to acquire virtual disk read access and SAN transport by using vSphere. When virtual disk access is obtained, vRanger reads the virtual disk blocks and negotiates with RDA inline deduplication to determine if the block is unique. Only unique blocks are copied to backup storage. Additional space savings techniques such as vRanger Active Block Mapping and VMware Changed Block Tracking can be enabled to reduce the data transport volume.

Recovery is performed using RDA to read the RecoveryPoint stored in the DR6000 appliance and vSphere to establish a virtual disk write access SAN transport. For Disaster Recovery, the DR 6000 Appliance provides the capability to replicate across the WAN to a remote physical or virtual DR appliance residing in either an enterprise network or in a public cloud.

Figure 19 Data Protection Architecture



7 Solution Specification

This section provides a complete list of components used as a part of the solution architecture. This reference architecture comprises VMware vSphere 6.0 Update 1 on PowerEdge FX2 architecture, Dell Networking S4048-ON switches and Dell Storage SC4020 Storage.

Table 6 provides a list of components needed to build a virtualized infrastructure with one PowerEdge FX2 architecture.

Table 6 Solution specification for 1 x PowerEdge FX2 architecture deployment

Component	Details	
Virtualization Infrastructure	1 x Dell PowerEdge FX2 2 x FN410S; 10GbE SFP+ IOA 4 x Qlogic QLE 2562 DP PCIe Add-on FC HBA	
Virtualization Hosts	4 x Dell PowerEdge FC630 Servers (4 per PowerEdge FX2)	
	Processor	2 x Intel Xeon E5-2660v3 Family
	Memory	128 GB; 8 x 16GB DDR4 DIMMs
	OS Volume	2 x 16 GB SD cards in internal SD module
	Network	Qlogic 57840 quad-port NDC
	OS	VMware ESXi 6.0 Update 1
Storage Arrays	Dell Storage SC4020 with 6x 400 GB (WI) SSD, 6 x 1.6 TB (RI) SSD and 12 x 10K SAS drives	
Networking	2 x Dell Networking S4048-ON for Local Area Networking (LAN) 2 x Brocade 6510 FC switches 1 x Dell Networking S3048-ON for OOB network (optional)	
Cable Requirements	7 x 10GbE SFP+ DAC cables 8 x LC-LC FC cables 2 x 40GbE QSFP+ cables	

Table 7 provides a list of components needed to build a virtualized infrastructure with two PowerEdge FX2 architecture.

Table 7 Solution specification for 2 x PowerEdge FX2 architecture deployment

Component	Details
Virtualization Infrastructure	2 x Dell PowerEdge FX2 4 x FN410S; 10GbE SFP+ IOA (2 per PowerEdge FX2) 8 x Qlogic QLE 2562 DP PCIe Add-on FC HBA (4 per PowerEdge FX2)



Component	Details	
Virtualization Hosts	8 x Dell PowerEdge FC630 Servers (4 per PowerEdge FX)	
	Processor	2 x Intel Xeon E5-2660v3 Family
	Memory	128 GB; 8 x 16GB DDR4 DIMMs
	OS Volume	2 x 16 GB SD cards in internal SD module
	Network	Qlogic 57840 quad-port NDC
	OS	VMware ESXi 6.0 Update 1
Storage Arrays	Dell Storage SC4020 with 6x 400 GB (WI) SSD, 6 x 1.6 TB (RI) SSD and 12 x 10K SAS drives Dell Storage SC220 expansion attached with 24 x 1.2 TB 10K SAS drives	
Networking	2 x Dell Networking S4048-ON for Local Area Networking (LAN) 2 x Brocade 6510 FC switches 1 x Dell Networking S3048-ON for OOB network (optional)	
Cable Requirements	14 x 10GbE SFP+ DAC cables 16 x LC-LC FC cables 2 x 40GbE QSFP+ cables	

Table 8 provides a list of components needed to build a virtualized infrastructure with two PowerEdge FX2 architecture.

Table 8 Solution specification for 6 x PowerEdge FX2 architecture deployment

Component	Details	
Virtualization Infrastructure	6 x Dell PowerEdge FX2 12 x FN410S; 10GbE SFP+ IOA (2 per PowerEdge FX2) 24 x Qlogic QLE 2562 DP PCIe Add-on FC HBA (4 per PowerEdge FX2)	
Virtualization Hosts	24 x Dell PowerEdge FC630 Servers (4 per PowerEdge FX)	
	Processor	2 x Intel Xeon E5-2660v3 Family
	Memory	128 GB; 8 x 16GB DDR4 DIMMs
	OS Volume	2 x 16 GB SD cards in internal SD module
	Network	Qlogic 57840 quad-port NDC
	OS	VMware ESXi 6.0 Update 1



Component	Details
Storage Arrays	Dell Storage SC4020 with 6x 400 GB (WL) SSD, 6 x 1.6 TB (RL) SSD and 12 x 10K SAS drives 4 x Dell Storage SC220 expansion attached with 24 x 1.2 TB 10K SAS drives
Networking	2 x Dell Networking S4048-ON for Local Area Networking (LAN) 2 x Brocade 6510 FC switches 1 x Dell Networking S3048-ON for OOB network (optional)
Cable Requirements	42 x 10GbE SFP+ DAC cables 48 x LC-LC FC cables 2 x 40GbE QSFP+ cables



8 Solution Validation

Dell Solution Engineering has tested and validated each of the configurations referred in this reference architecture. The validation entails:

- Interoperability between each of the components in the architecture
- Guidance on the firmware version used for the testing
- Example test scenario that provides guidance on sizing

The solution design in this paper employs best practices and recommendations at all infrastructure components. Any solution infrastructure must be verified for component level functionality before it can be used in production. This should involve verification of infrastructure component failover to ensure that the services offered by the solution infrastructure are highly available and resilient against component failures.

The solution infrastructure described in the paper was verified to ensure that the component level high-availability and resiliency are indeed available. There was no user-perceived down time of any services running inside the virtual machines in the infrastructure during the component failover.

8.1 Validated Firmware Revisions of Components

Each configuration specified in the reference architecture has been tested by deploying and configuring it end to end including vSphere.

The solution architecture presented in this paper was validated against the firmware version listed in Table 9 and Table 10.

Table 9 PowerEdge FC630 Servers

Device	Version
BIOS	1.2.5
CPLD	1.0.5
iDRAC8 Enterprise	2.20.20.20
Life Cycle Controller (LC) 2	2.20.20.20
Network Controller QLogic BCM57840S	7.10.18
QLogic BCM57840S Driver	7.12.2.0



Device	Version
Qlogic QLE2562 8Gb FC	03.22.00

Table 10 PowerEdge FX2s Chassis, Compellent SC4020 Storage and Switches

Device	Version
CMC	1.20.A00.201502200213
Dell FN410S IOA	9.6 (0.0)
Force10 S4048-ON	9.8
Brocade 6510	v7.3.0a
PERC H730 Mini	25.2.2-0004
Dell Storage SC4020	6.6.5

8.2 Storage Verification

For a general-purpose virtualization infrastructure, it is important to ensure that each component within the infrastructure is appropriately sized to run the virtualized applications. One of the unique aspects of the reference architecture is that it adopts a single storage architecture regardless of scale. Hence, performance of SC4020 plays a major role in the overall solution performance and influences the number of VMs that can be deployed on this infrastructure. While it is not practical to size every type of application that may run on a general-purpose infrastructure as described in this paper, it is still important to determine the overall performance that the storage subsystem can offer.

To this extent, the storage verification of the solution infrastructure was carried out to understand how Dell Storage SC4020 performs. An Online Transaction Processing (OLTP) IO load profile was simulated using IOMeter. Table 11 provides the IO profile details.



Table 11 IO profile for storage verification

IO Profile Type	Read	Write	Access Size	IO Pattern
Database – OLTP	70%	30%	8K	100% Random

For IOMeter verification, two VMs per cluster node were created and the storage verification tests were performed within the VM. Table 12 provides the VM configuration used for the IOMeter verification.

Table 12 VM configuration for IOMeter verification

VMs per Host	vCPU / VM	Memory / VM	OS VMDK	Data VMDK
2	20	60 GB	40 GB	40 GB

As shown in Table 12, a 40 GB thick provisioned VMDK was used for the IOMeter verification. These VMDK files were stored across eight 750 GB volumes created on Dell Storage SC4020. IOMeter verification was performed with varying queue depth values. This verification was performed on six different PowerEdge FX2 architecture containing a total of 24 PowerEdge FC630 servers.

Table 13 provides the data generated from the IOMeter runs on the six PowerEdge FX2 architecture.

Table 13 IOMeter performance data from six PowerEdge FX2 architecture (48 workers)

Queue Depth	1	4	16
IOPS	23714	43790	47992
MBps	185	342	374
Read MBps	129	239	262
Write MBps	55	102	112
Transactions per Second	23714	43790	47992
Average Response Time	1.93	4.2	15.33
Average Read Response Time	2.42	5.56	11.59
Average Write Response Time	0.8	1.01	24.06
Average Transaction Time	1.93	4.2	15.33



The storage performance data shown in Table 13 clearly shows that the storage subsystem in this architecture has enough IOPS and disk capacity to scale beyond an eight-node deployment.

Let us look at an example:

The solution design assumes one VM per logical processor core equivalent on each workload server. Further, other Dell lab tests show that the SC4020 can deliver 129,661 Oracle 100% read OLTP IOPS with less than 1ms latency¹, further showing that this test didn't fully saturate the SC4020 array. However, it is always recommended to size the storage infrastructure by using a real-world workload scenario that is planned for deployment.

The six PowerEdge FX2 based scale-unit can be used to deploy approximately 1,000 virtual machines each with 2 vCPUs, 4 GB memory and 40 GB VMDK. Considering an IO profile in which each VM requires 30 IOPS, the storage backend must be capable of sustaining 30,000 IOPS during peak load. The drive configuration of Dell Storage SC4020 used in this reference architecture can support more than 50,000 sustained IOPS. Therefore, with a 4:1 vCPU oversubscription and 1.5:1 memory oversubscription for the virtual machines, the six PowerEdge FX2 based scale-unit will be capable of running approximately 1,000 virtual machines

¹ Based on Dell internal testing at Dell labs using Oracle ORION toolkit in March 2014. Performance shown indicative of 12 WL drives. Oracle workload performance will vary depending upon workload I/O profile and SC4020 drive configuration, usage and manufacturing variability.



9 Solution Summary

For IT administrators, the PowerEdge FX2 architecture resolves the complexities involved in managing infrastructure by combining servers, storage and networking into an easy-to-deploy chassis. Potential errors are reduced by using unified and simplified systems management.

The design goals, as described in Section 5, will add to the solution design decisions and influence component choice. The recommended solution architecture described in Section 6 takes into account the best practices specified by Dell and VMware. The design points emphasize the use of the latest hardware and software features presented in the PowerEdge FX2 architecture to enable scalability, infrastructure HA and manageability. The infrastructure HA features ensure that the services hosted on the virtualized infrastructure are continuously available.

The ToR network architecture using a multiple Virtual Link Trunk (mVLT) ensures that there are multiple active paths from the compute to the TOR and therefore provides better resiliency. Using the VLT mode of FN IOA enables low-touch configuration for chassis networking there by reducing the complexity of network configuration. The virtual switch design and implementation with Network Partitioning (NPAR) ensures high quality-of-service (QoS) for different classes of traffic and ensures that the demands for network bandwidth are met. Hence, the bandwidth requirements for vMotion traffic are given priority, when needed, to enable seamless movement of VMs during optimization and/or planned downtime of the hypervisor hosts.

The storage architecture implemented using Dell Storage SC4020 is simple to configure and manage. In addition, the hybrid drives in the storage provide optimal performance needed for the services or virtual machines hosted on the virtual infrastructure. The tiering feature of Dell Storage SC4020 brings all the hot data to the SSD tier therefore improving the overall virtual infrastructure performance.

Dell Active System Manager 8.1.1 provides an efficient method to deploy and configure the infrastructure components using service templates. This enables faster time-to-value for the virtualized infrastructure.

To summarize, the PowerEdge FX2 architecture combined with Dell Storage SC4020 delivers the performance for a general purpose virtual infrastructure built using VMware vSphere 6.0 Update 1. The principles used in the infrastructure design make it highly available and provide the necessary ingredients for virtualizing enterprise applications. The building block approach to implementing virtual infrastructure enables reduced upfront capital expenditure while providing enough room for future growth as needed.



A Additional Resources

1. **Support.dell.com** is focused on meeting customer requirements with proven services.
2. [DellTechCenter.com](https://delltechcenter.com) is an IT Community where you can connect with Dell customers and Dell employees for the purpose of sharing knowledge, best practices and information about Dell products and installations.
3. [New features in VMware vSphere 6.0 Update 1](#)
4. Referenced or recommended Dell publications:
 - a. [Dell PowerEdge FX2](#)
 - b. [Dell Storage SC4020](#)
 - c. [Dell Storage SC220](#)
 - d. [Dell Active System Manager 8.1.1](#)
5. [Deployment Guide: VMware vSphere 6.0 Update 1 on Dell PowerEdge FX2](#)
6. [Deployment Guide: VMware vSphere 6.0 Update 1 on Dell PowerEdge FX2 with ASM](#)



B Solution Component Details

The following sections provide the technical and product overview of compute configuration options available with PowerEdge FX2 architecture.

B.1 Dell PowerEdge FX2 Architecture

PowerEdge FX2 is a 2U rack mount chassis and accommodates PowerEdge FC630, PowerEdge FC430, PowerEdge FC830, PowerEdge FM120x4 server sleds and PowerEdge FD332 storage sled. A sled is a pluggable component in a modular architecture similar to that of PowerEdge FX2. The PowerEdge FX2 architecture offers flexibility in terms of infrastructure compute, storage and networking.

Table 14 presents an overview of all component configurations supported in PowerEdge FX2 architecture.

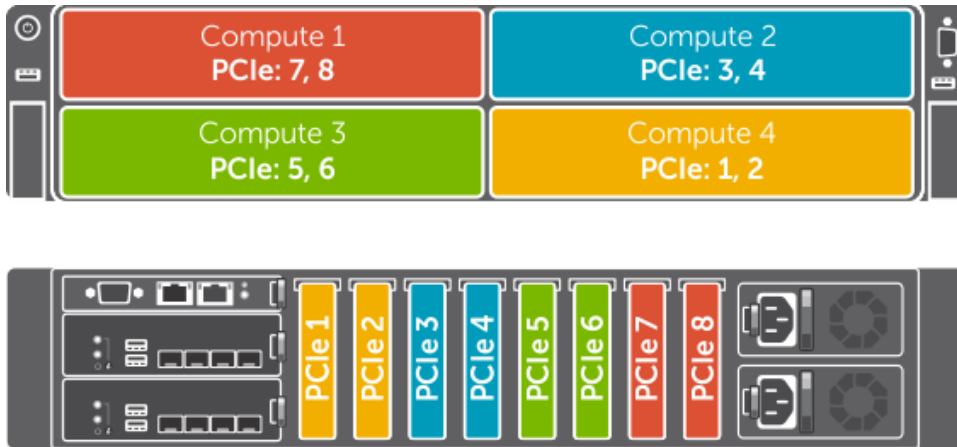
Table 14 Dell PowerEdge FX2 supported components overview

Feature	Description
Server Compatibility	PowerEdge FC630/FC430/FC830/FM120x4 servers PowerEdge FD332 storage
Form Factor	2U rack enclosure
Number of Server Sleds	Up to four FC630 Up to eight FC430 Up to two FC830 Up to four FM120x4
Number of Storage Sleds	Up to four FD332
I/O	8 PCIe slots (supporting Ethernet and FC)
Power Supplies	Up to 2 x 1600w PSU (FC630 and FC430) Up to 2x 1100w PSU (FM120x4)
Management	1 Chassis Management Controller
Network	2 x pass-through I/O Modules (IOM; 1 GbE or 10 GbE) 2 x 10 GbE SFP+ IO Aggregator (IOA; FN410s) 2 x 10 GbE 10 Base T IO Aggregator (IOA; FN410T) 2 x FC and 10 GbE combo IO Aggregator (IOA; FN2210S)

The server connectivity in a PowerEdge FX2 architecture can be extended by adding supported PCIe peripheral cards. There are eight PCIe slots available for this purpose, and they are internally mapped to servers. Figure 20 shows the internal mapping of the PCIe slots to servers in the PowerEdge FX2 architecture.

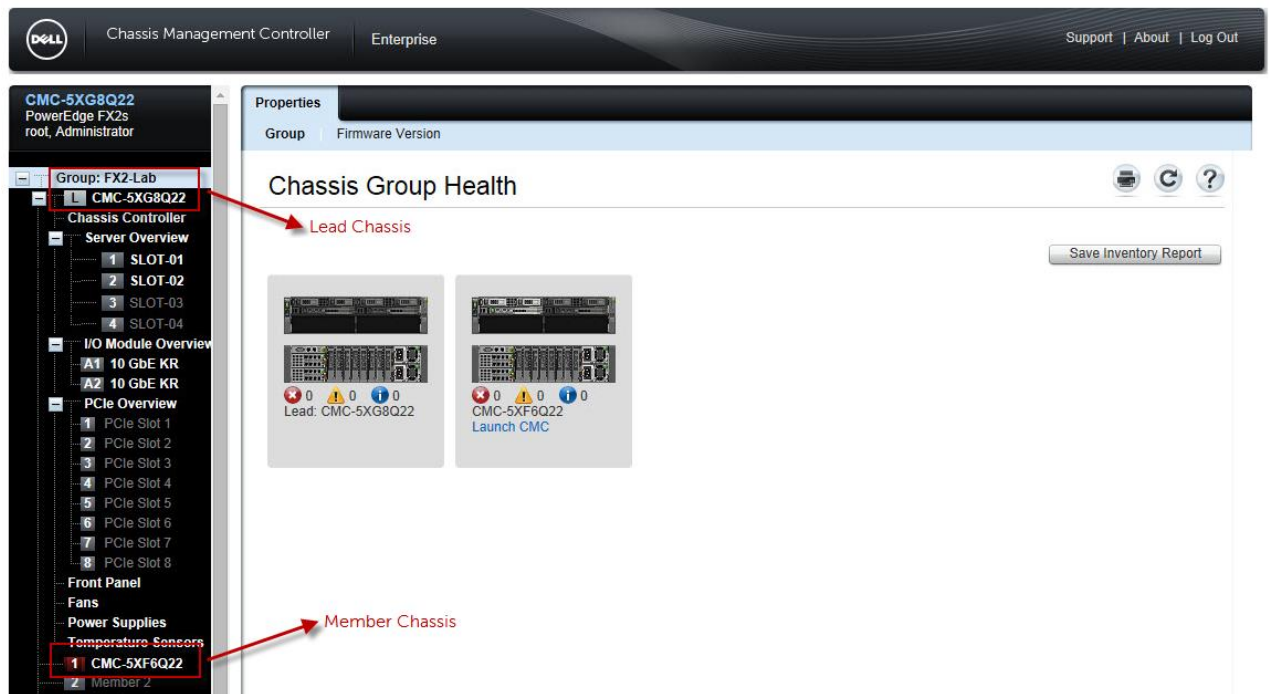


Figure 20 PCIe slot mapping in PowerEdge FX2 (with half-width servers)



The PowerEdge FX2 architecture is managed through the Chassis Management Controller (CMC). The CMC Web console provides features to manage servers and IOMs and monitor PCIe slots in a PowerEdge FX2 architecture. Figure 21 shows the group management of multiple PowerEdge FX2 architecture in the CMC web console.

Figure 21 PowerEdge FX2 CMC web console - Group Management



Using the group management feature of CMC, it is possible to manage up to 20 PowerEdge FX2 architecture from the same CMC console. This can be beneficial when there are multiple chassis



participating in an infrastructure deployment such as the virtualization infrastructure implemented in this reference architecture.

B.1.1 Dell PowerEdge FC430 Server

PowerEdge FC430 is a quarter-width server sled that supports Intel Haswell processors up to 14 processing cores and up to 256 GB physical memory. This enables PowerEdge FC430 to deliver the performance required for compute-intensive tasks. A Qlogic 57810 or Qlogic 57840 or Emulex 10 GbE blade Network Daughter Card (bNDC) used within the server provides the network connectivity through the 10 GbE pass-through IOM or 10 GbE IOA in the PowerEdge FX2 architecture. PowerEdge FC430 servers use blade mezzanine cards to connect to the PCIe add-on peripherals in the PowerEdge FX2 architecture. This connectivity can be leveraged to build optional redundancy for network connections. For operating system (OS) drives, PowerEdge FC430 supports PowerEdge RAID controller (PERC) S130.

Table 15 describes the components supported in a PowerEdge FC430 server.

Table 15 Supported components in PowerEdge FC430

	Supported in PowerEdge FC430
CPU	Up to 2 Intel Xeon E5-2600 family processors
Memory	8 DIMMs; up to 256 GB
Networking	Qlogic 57810 Dual Port 10 GbE bNDC Qlogic 57840 Quad Port 10 GbE bNDC Emulex 10GbE adapters
Storage	Up to 2 x 1.8" uSATA SSDs
Form Factor	¼ width Sled
PCIe slots	1 x PCIe slots (x8) provided by PowerEdge FX2
Systems Management	iDRAC8 Express or iDRAC8 Enterprise

B.1.2 Dell PowerEdge FM120x4 Server

PowerEdge FM120x4 is a half-width server sled that includes four Intel Atom processor based system-on-a-chip (SOC) systems. Each SOC contains a single-socket Intel Atom processor C2000 product family. This platform is ideal for large-scale static web hosting and lighter weight processing, such as batch data analytics. A fully loaded FX2 chassis can hold 16 SOC's, offering impressive density. By using eight-core processors, 128 cores and 32 DIMMs of memory can be utilized in a single 2U FX2 chassis. The low-power characteristics of the C2000 enable the FM120x4 to comply with Dell Fresh Air 2.0 requirements, enabling you to save on energy and cooling costs for your data center.

Table 16 describes the components supported in a PowerEdge FM120x4 server.



Table 16 Supported components in PowerEdge FC630

	Supported in PowerEdge FM120x4
CPU	4 x Intel Atom processor C2000 family
Memory	2 x DIMMs DDR3 memory per SOC, UDIMM only, (8 per sled)
Networking	2 x 1Gb controllers embedded in the SOC
Storage	1 x 2.5-inch front access hard drives or 2 x 1.8" SSD drives
Form Factor	1/2 width Sled
PCIe slots	No PCIe slots
Systems Management	iDRAC8 Express or iDRAC8 Enterprise

B.1.3 Dell PowerEdge FC830 Server

The PowerEdge FC830 server is a full-width server sled that supports Intel Haswell processors up to 18 processing cores and up to 1.5 TB physical memory. This enables PowerEdge FC830 to deliver the performance required for compute-intensive tasks. A Qlogic 57810 or Qlogic 57840 10 GbE or Emulex 10 GbE blade Network Daughter Card (bNDC) used within the server provides the network connectivity through the 10 GbE pass-through IOM or 10 GbE IOA in the PowerEdge FX2 architecture. PowerEdge FC830 servers use blade mezzanine cards to connect to the PCIe add-on peripherals in the PowerEdge FX2 architecture. This connectivity can be leveraged to build optional redundancy for network connections. For operating system (OS) drives, PowerEdge FC830 supports a wide variety of PowerEdge RAID controller (PERC) devices. An internal SD card module can also be used for deploying host OS on a PowerEdge FC830 server.

Table 17 describes the components supported in a PowerEdge FC830 server.

Table 17 Supported components in PowerEdge FC830

	Supported in PowerEdge FC830
CPU	Up to 4 Intel Xeon E5-4600 family processors
Memory	48 DIMMs; up to 1.5 TB ²
Networking	Qlogic 57810 Dual Port 10 GbE bNDC

² Some processor heat sink configurations may not support 768 GB. Refer to the product manual for more information.



	Qlogic 57840 Quad Port 10 GbE bNDC Emulex 10GbE adapters
Storage	PowerEdge Express Flash NVMe PCIe SSD, SATA HDD/SSD or SAS HDD/SSD Up to 8 x 1.8" SSD or 2 x 2.5"
Form Factor	1/2 width Sled
PCIe slots	4 x PCIe slots (x8) provided by PowerEdge FX2
Systems Management	iDRAC8 Express or iDRAC8 Enterprise

B.1.4 Dell PowerEdge FC630 Server

PowerEdge FC630 servers support Intel Haswell processors up to 18 processing cores and up to 768 GB physical memory. This enables PowerEdge FC630 to deliver the performance required for compute-intensive tasks. A Qlogic 57810 or Qlogic 57840 10 GbE blade Network Daughter Card (bNDC) used within the server provides the network connectivity through the 10 GbE pass-through IOM or 10 GbE IOA in the PowerEdge FX2 architecture. PowerEdge FC630 servers use blade mezzanine cards to connect to the PCIe add-on peripherals in the PowerEdge FX2 architecture. This connectivity can be leveraged to build optional redundancy for network connections. For operating system (OS) drives, PowerEdge FC630 supports a wide variety of PowerEdge RAID controller (PERC) devices.

PowerEdge FC630 servers in a PowerEdge FX2 architecture provide the necessary compute required for the virtual infrastructure. This server strikes a balance between the density provided by PowerEdge FC430 and the capacity and performance provided by PowerEdge FC830. PowerEdge FC430 servers scale up the memory up to 768 GB (unlike FC430 where the maximum supported memory is 256 GB), when needed. By using four FC630 servers in PowerEdge FX2 architecture, the solution architecture enables the same computing capacity offered by two FC830 servers in PowerEdge FX2 while ensuring that there is greater physical isolation for virtual applications. Therefore, PowerEdge FC630 is an ideal choice for the virtualized architecture presented in this paper.

For an overview of PowerEdge FC430, PowerEdge FC830 and PowerEdge FM120x4, refer to Appendix C.

Table 18 describes the components supported in a PowerEdge FC630 server.

Table 18 Supported components in PowerEdge FC630

	Supported in PowerEdge FC630
CPU	Up to 2 Intel Xeon E5-2600 family processors



Memory	24 DIMMs; up to 768 GB ³
Networking	Qlogic 57810 Dual Port 10 GbE bNDC Qlogic 57840 Quad Port 10 GbE bNDC Emulex 10GbE adapters
Storage	PowerEdge Express Flash NVMe PCIe SSD, SATA HDD/SSD or SAS HDD/SSD Up to 8 x 1.8" SSD or 2 x 2.5"
Form Factor	1/2 width Sled
PCIe slots	2x PCIe slots (x8) provided by PowerEdge FX2
Systems Management	iDRAC8 Express or iDRAC8 Enterprise

B.1.5 Dell PowerEdge FN410S

As mentioned in Table 11, PowerEdge FX2 supports multiple networking options for server connectivity to the top-of-rack (ToR) switches and / or to the data center core. These networking options provide flexible choices when designing solution infrastructure. This reference architecture employs PowerEdge FN410S for the TOR connectivity from the PowerEdge FX2 chassis.

Using the IOA configuration simplifies connectivity by as much as 8-to-1, greatly reducing cabling complexity. The FN IOA is a Plug-n-Play networking device providing a low-touch experience for server administrators. Most IOA initial setup and networking functions are automated making for minimal touch for basic to advanced features.

B.2 Dell Networking S4048-ON Switch

Dell Networking S4048-ON is a 1U high-density 10/40 GbE ToR switch with 48 dual-speed 1/10GbE (SFP+) ports and six 40 GbE (QSFP+) uplinks. This switch leverages a non-blocking and cut-through switching architecture to provide ultra-low-latency performance for applications. The six 40 GbE ports can be used to create either a Virtual Link Trunk (VLT) between the switches to enable traffic isolation within the solution infrastructure or as a connectivity to the data center core network.

The Dell S4048-ON supports the open source Open Network Install Environment (ONIE) for zero-touch installation of alternate network operating system including feature-rich Dell Networking OS and also supports Dell Networking's Embedded Open Automation Framework, which provides advanced network automation and virtualization capabilities for virtual data center environments.

³ Some processor heat sink configurations may not support 768 GB. Refer to the product manual for more information.



B.3 Dell Storage SC4020 Array

Dell Storage SC4020 belongs to the Storage Center (SC) 4000 series arrays based on the SC8000 platform. This storage array offers multi-protocol support and virtualized multi-tier, multi-RAID-level storage policies. Each SC4020 array comes with dual redundant controllers, 24 internal drive slots, eight 8 Gb Fibre Channel (FC) or four 10 Gb iSCSI network ports and one 1 Gb port per controller for out-of-band (OOB) management traffic. The supported drive types range from Write-Intensive (WI) Solid State Drives (SSDs) to Read-Intensive (RI) SSDs to 15K, 10K, and 7.2K SAS drives in small form factor (2.5-inch). This storage array supports multi-tier data placement that improves application performance. This array supports expansion up to 192 drives by adding Dell Storage SC220 or SC200 enclosures.

Using the largest capacity disk drives supported, a Dell Storage SC4020 array can support up to 1 Peta-Byte (PB) of raw disk capacity. The 32-GB controller cache (16 GB per controller) provides the necessary storage performance for most enterprise application needs.

The Storage Center 6.5 Operating System (OS) provides features such as block-level compression, synchronous Live Volumes to restore data in a non-disruptive manner and Active Directory (AD) Single-Sign-On (SSO). The Storage Center Manager enables easier and out-of-box web-based management of SC4020 arrays while the Enterprise Manager and its components can be leveraged to build a secure, multi-tenant environment with the Dell Storage Center SAN.

B.4 Brocade 6510

Brocade 6510 is a 48-port 1U Gen 5 FC switch that is suitable for high-performance data requirements of server virtualization, cloud and enterprise applications. This switch can be configured in 24, 36 or 48 ports and supports 2, 4, 8, 10 or 16 Gbps speeds. This enables enterprises to start small and scale the FC infrastructure based on the growth, as needed. By using features such as Virtual Fabrics, Quality of Service (QoS) and zoning, IT departments can build multi-tenant cloud environments. Management and diagnostics features such as Monitoring and Alerting Policy Suite (MAPS), Dashboards, Flow Vision, Fabric Performance Impact (FPI) monitoring and Credit Loss Recovery help administrators and IT organizations avoid problems before they impact the SAN operations.

B.5 VMware vSphere 6.0 Update 1

vSphere 6.0 Update 1 is an industry leading server virtualization platform that enables creation of a highly available, resilient and on-demand infrastructure for virtualizing enterprise applications or as a general purpose virtual infrastructure. vSphere 6.0 Update 1 provides improved features that are available to enterprises looking at data center consolidation and enhancing resource utilization. Apart from the scalability improvements, vSphere 6.0 Update 1 enhanced vSphere Fault Tolerance support for multi-vCPU VMs, compression support for vSphere Replication and added support for multisite content library.

The table below provides an overview of the capabilities in vSphere 6.0 Update 1.



Table 19 VMware vSphere 6.0 Update 1 Cluster capabilities

	VMware vSphere 6.0 Update 1
VM Hardware Version	Virtual Hardware 11
Number of vCPU	128
VM Memory	4 TB
Cluster Nodes	64
Maximum CPU per Host	480
Maximum Memory per Host	12 TB
Maximum VMs per Cluster	8000

The virtualization capabilities offered by vSphere 6.0 Update 1 along with the management capabilities offered by vCenter enable building an efficient platform for virtualizing enterprise applications.

