

# **Dell EMC Validated System for Virtualization — Design Guide for VMware with iSCSI SAN Infrastructure**

A Dell Architecture Guide

## Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your computer.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

© 2016 Dell Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. 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.

# Contents

<b>1 Executive Summary.....</b>	<b>6</b>
<b>2 Audience.....</b>	<b>7</b>
<b>3 Why Dell EMC Validated System for Virtualization?.....</b>	<b>8</b>
Differentiated “Dell EMC Validated System” approach addresses challenges and limitations.....	8
<b>4 Solution overview and design principles.....</b>	<b>10</b>
The Dell EMC Validated System for Virtualization.....	10
Availability.....	11
Flexibility.....	11
Best practices in a VMware virtualized infrastructure.....	12
Virtual switch configuration best practicesDell Compellent best practicesNetwork best	
PracticesBackup and recovery best practices.....	12
Comprehensive.....	13
<b>5 Customer adoption scenarios.....</b>	<b>14</b>
Infrastructure Capacity Cube.....	14
Dell Performance Analysis Collection Kit (DPACK).....	15
Dell System Builder.....	16
Sample Configurations.....	17
iSCSI Customer Defined Scale Unit - Small.....	17
iSCSI Customer Defined Scale Unit - Medium.....	18
iSCSI Customer Defined Scale Unit - Large.....	19
<b>6 Deployment options for Dell EMC Validated System for Virtualization.....</b>	<b>21</b>
Management Software.....	22
Dell Active System Manager (ASM).....	22
Dell Foglight for Virtualization Standard Edition.....	23
Dell vRanger.....	23
Dell DR4300 Disk Backup Appliance.....	23
Dell DR6300 Backup-to-Disk Appliance.....	23
VMware vRealize Operations Manager.....	23
VMware vCenter Server Appliance.....	23
VMware vSphere 6.0.....	24
Dell Active System Manager (ASM).....	24
Dell Foglight for Virtualization Standard Edition.....	24
Dell vRanger.....	24
Dell DR4300 Disk Backup Appliance.....	24
Dell DR6300 Backup Appliance.....	25
VMware vRealize Operations Manager.....	25
VMware vCenter Server Appliance.....	25

VMware vSphere 6.0.....	25
Dell PowerEdge Rack Servers.....	25
Dell PowerEdge R630.....	25
Dell PowerEdge R730 .....	25
Dell PowerEdge FX2.....	25
Dell PowerEdge FC430.....	26
Dell PowerEdge FC630.....	26
Dell SC Storage Solutions.....	26
Dell SC Storage SCv2020.....	26
Dell SC Storage SC4020.....	26
Dell SC Storage SC9000.....	26
Dell Networking.....	26
Dell Networking N4032.....	26
Dell Networking S3048-ON.....	27
Dell Networking S4048-ON.....	27
Dell Networking Z9100-ON.....	27
<b>7 iSCSI architecture design.....</b>	<b>28</b>
Compute Architecture.....	28
Network Architecture.....	29
MLAG configuration at Dell Networking N4032.....	29
Virtual Link Trunking (VLT) at Dell Networking S4048-ON.....	29
Network Connectivity FX2.....	30
Network Configuration PowerEdge Rack Servers.....	31
Network Configuration for LAN traffic with QLogic and NPAR.....	32
Network Configuration for LAN Traffic with Intel/Qlogic Network Daughter Card and Distributed	
Virtual Switch.....	33
vSwitch setup.....	34
Storage Architecture.....	34
Storage Fabric configuration.....	35
Storage connectivity.....	36
Storage connectivity PowerEdge Rack.....	37
Storage configuration.....	37
Drive and RAID configuration for SC4020 and SCv2020.....	38
Management Architecture.....	39
Data Protection Architecture.....	41
<b>8 Scaling the solution.....</b>	<b>42</b>
Scaling out with Customer Defined Scale Unit (CDSU).....	42
Scaling out compute and storage independently.....	43
Scaling network infrastructure.....	44
<b>9 Lifecycle solution management.....</b>	<b>46</b>
Deploy.....	46
Operate.....	47
End-to-end view from virtual to physical.....	47

Infrastructure remediation.....	48
End-to-end view from virtual to physical.....	49
Infrastructure remediation.....	49
Maintain.....	50
Repurpose.....	51
Backup and Recovery with Dell vRanger.....	52
Capacity planning (Foglight).....	53
<b>10 Testing and validation.....</b>	<b>55</b>
Design principle validation — high availability.....	55
Design principle validation — load balanced.....	56

# Executive Summary

This Dell Architecture guide outlines the Dell EMC Validated System for Virtualization with VMware for Fibre Channel environments. This solution allows customers to scale from 100 to 3000 virtual machines. Solution options encompass a combination of solution models for Fibre Channel environments including rack servers and Dell converged platforms, Dell mid-tier enterprise storage (SC4020) and low-end (SCv2020). Architecture design, configuration and implementation include considerations for the key components of the architecture required for a vSphere 6 implementation on a Dell infrastructure.

## Audience

The audience for this paper includes, but is not limited to the following individuals:

- Sales engineers (technologists and architects)
- Field consultants
- Partner engineering team members
- Customers
- Individuals interested in deploying an optimized and validated systems stack with Dell technology and VMware vSphere

## Why Dell EMC Validated System for Virtualization?

With the increase of business demands and the decrease of IT budgets, customers face unprecedented pressure to improve efficiency and lower costs. The current operational model of delivering IT services, which involves procuring technology from best of breed technology providers and managing them in silos, proves to not only be time consuming but problematic. In this approach, customers are typically burdened to manually make design decisions, validate various components, set up and configure components and manage the environment in an ongoing fashion by engaging multiple vendors for assistance. Across the end-to-end infrastructure lifecycle, these elements increase in complexity and in cost for customers.

A growing number of organizations recognize the benefits of a converged system in that it efficiently combines the compute, storage and networking components for effortless deployment and postdeployment system management. However, most organizations have not yet realized the goal of implementing a converged system because they demand a higher degree of flexibility and simplicity that moves beyond existing converged offerings. As it is, existing reference architecture solutions do provide some degree of flexibility, but do not offer manageability or scalability benefits. The Dell EMC Validated System for Virtualization bridges this gap by offering a tested and validated integrated system that is highly flexible, scalable, and driven using end-to-end automation throughout the infrastructure lifecycle.

The Dell EMC Validated System for Virtualization is the industry's most flexible converged system to date, with choice in building blocks of compute, storage and networking tested and validated to integrate and operate together in support of virtualized environment. The System can also be right-sized to a high degree to fit your needs, from as small as two compute nodes with associated storage to as large as multiple high-density racks, scaled in increments as little as a single node or drive.

### Differentiated “Dell EMC Validated System” approach addresses challenges and limitations

To provide IT services faster, while lowering costs and streamlining operations, Dell engineered the Dell EMC Validated System for Virtualization. This groundbreaking system enables you to achieve greater operational efficiencies and savings, and unparalleled management simplicity, by giving you more power than ever to define and design it.

The System can also be deployed, with options ranging from “do-it-yourself” with a deployment guide to a system integrated at a merge center or on site by Dell or by your own integration vendor. The Dell EMC Validated System for Virtualization system is:

- Built on our best-of-breed products that are designed for virtualization across the ecosystem
- Tested, validated, and fully integrated, yet flexible enough to be tailored for your organization, removing risk and accelerating your time to value
- Driven by automation, to simplify ongoing management and drive IT agility
- Optimized and protected with investments Dell has made in software that make our solutions easier to manage and deliver ongoing results
- Delivered with Dell's global reach, exceptional execution, and delivery, which provides consistent deployment, management, and maintenance in every region of the world

The Dell EMC Validated System for Virtualization system is designed to provide detailed guidance and measurable benefits to customers and help solve challenges with IT infrastructure. It addresses the following key requirements of customers:

- **Reduce Capital Expenditure (CAPEX):** How do we reduce CAPEX by starting small and growing large?
- **High flexibility and less complexity:** Flexibility in configuration, scaling, and implementation choices based on specific outcomes desired.
- **Lower operational costs:** How do we lower our overall Total Cost of Ownership (TCO) by optimizing the infrastructure resources and reducing management complexity?



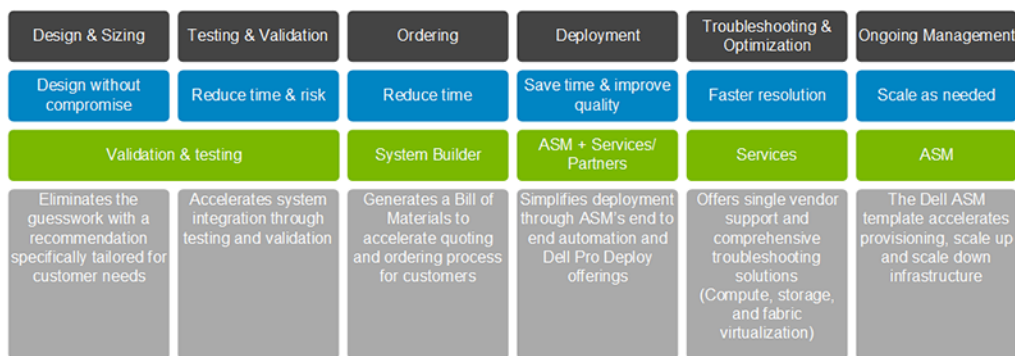
- **Optimized environment:** How do we deploy and configure an optimized environment, based on best practices, to deliver the best experience to our end users?
- **Meet business Service Level Agreement (SLA) under all conditions:** How to design, deploy, and manage a solution that handles failures without causing disruption?

By offering various design choices and guidance on choosing the right components, the Dell EMC Validated System for Virtualization takes the guesswork out of solution design and reduces the enormous time it takes to procure, validate, and integrate components. They can be designed to start small, based on the customer's initial requirements, and grow, based on customer's ongoing requirements, which reduce the initial investment required for infrastructure deployment.

The Dell EMC Validated System for Virtualization goes through rigorous validation by Dell engineers to ensure that the system works as designed. Since the systems are designed with an assumption that disruptions are bound to happen, they are tested for failures at every level to ensure that they have enough capacity to account for failures when they occur.

With the Dell Active System Manager (ASM), the Dell EMC Validated System for Virtualization further simplifies ongoing management. Whether it is configuring the system based on Dell and VMware best practices, scaling the system as business demands grow, updating the system as new hardware updates come along, ensuring compliance of the system by managing drift, or repurposing the system as business needs change, ASM, when combined with the Dell EMC Validated System for Virtualization, drives IT agility and reduces ongoing IT costs.

The Dell EMC Validated System for Virtualization also provides a single point of support for the complete system including hardware and software through Dell ProSupport Plus. ProSupport Plus resolves issues faster when they occur and reduces the risk of severe issues and outages.

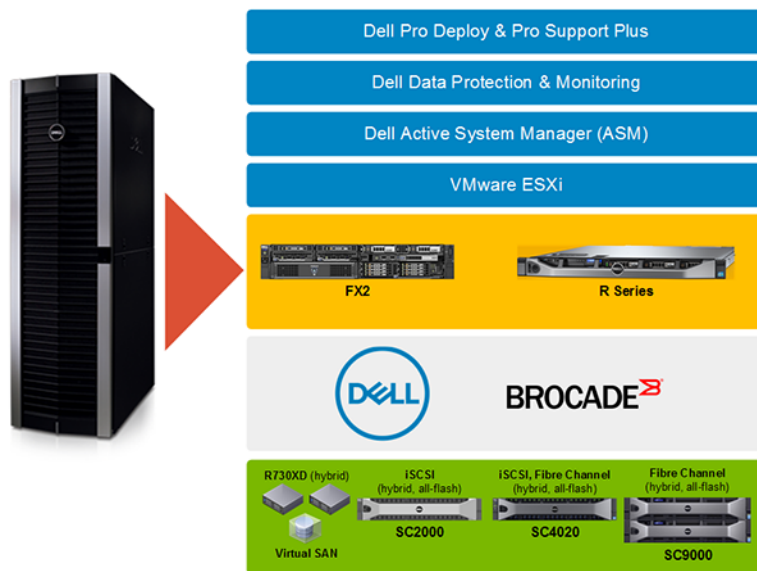


**Figure 1. Dell EMC Validated System for Virtualization helps IT teams improve efficiency**

## Solution overview and design principles

The Dell EMC Validated System for Virtualization provides a scalable, tested, highly available, automated, and fully supported integrated system with built-in monitoring and data protection for a general-purpose virtualized infrastructure. Built with design flexibility in mind that caters to various business and technical requirements, the Dell EMC Validated System for Virtualization is designed in accordance with the best practices and recommendations for virtualizing enterprise applications and to maintain a baseline level of performance.

### The Dell EMC Validated System for Virtualization



The Dell EMC Validated System for Virtualization is designed for Dell's proven server, storage, Networking, and management platforms. The design is also based on unique customer requirements such as:

- Capacity
- Number of workloads
- key workload characteristics (CPU, memory, input/output operations per second, latency, and throughput)
- Server form factor (racks, FX2, blades)
- Fabric (FC, iSCSI)
- Density

The common design principle of each deployment of the Dell EMC Validated System for Virtualization system is covered in this topic.

Topics:

- [Availability](#)
- [Flexibility](#)
- [Best practices in a VMware virtualized infrastructure](#)
- [Comprehensive](#)

# Availability

Loss of productivity due to unplanned downtime is a grave concern for most IT customers. Dell EMC Validated System for Virtualization mitigates this concern by “Designing in for Failure”. The underlying principle is to expect outages to occur during the life of the system.

In each scenario, the recommended configuration considers performance under a degraded state, a result of component failure (SAN controller or Server Network Interface Card) or system level (Server or Network) failure. In each configuration, there is enough capacity to account for failures while ensuring that workload performance is not impacted.

- If a node fails at the compute level, the VMware vSphere cluster built on PowerEdge platforms provides infrastructure availability that ensures that the VMs are migrated to other available physical hosts. This ensures that there is little or no disruption to the workloads.
- The network layer is redundant with two Dell Networking 10GbE switches to eliminate a single point of failure in the solution.
- The storage arrays provide redundancy at various levels within the array.

**NOTE:** See the [iSCSI Architecture Design](#) section for implementation details and the [Testing and Validation](#) section for an example of system behavior under simulated failures.

While high availability reduces points of failure and limits the probability of downtime, it cannot completely eliminate downtime unless sound fault tolerant implementation is in place. Designing a fully fault tolerant system depends on three key customer requirements:

- **Recovery Time Objective** How quickly must the system recover?
- **Recovery Point Objective** How much data can you afford to lose?
- **Cost** How much investment do we need to implement to meet these objectives?

There are multiple architecture patterns that can be implemented to support these requirements. These include:

- Back up and Restore
- Pilot Light
- Full working low-capacity Standby
- Multisite Active-Active Systems

Dell SC Storage (Dell SC4020 and Dell SCv2020) and VMware Virtual SAN can be architected using Dell SC Live volumes and VMware Metro Storage cluster technology to support pilot light, low capacity standby and multisite active-active architectural patterns. The actual implementation details are outside the scope of this document. See the [Storage Center Synchronous Replication and Live Volume](#) Solutions Guide for details on how Dell SC Series live volumes enables auto failure for zero RPO/RTO and stretching of compute and storage between multiple sites.

The Dell EMC Validated System for Virtualization comes with a built-in backup and recovery solution using Dell vRanger and data recovery appliances. For all deployment options, it enables simple and extremely cost effective mechanism to protect data. See the [Data Protection Architecture](#) section for details.

# Flexibility

When considering the size of the desired infrastructure, customers rarely know how much compute and storage they need. The requirements are tied to business needs, application requirements, and the pace of growth, which can change over time. Too often, customers are forced to choose small, low-cost solutions that require a fork-lift upgrade later or buy a large, high-density solution and pay for capacity that is not immediately needed.

What customers need is the flexibility to operate at any scale. The Dell EMC Validated System for Virtualization solution provides the best of both requirements, meaning no trade-offs on convenience and flexibility of rack deployments or the different levels of performance that various storage solutions provide.

To deliver flexibility, the Dell EMC Validated System for Virtualization is designed around the concept of a Customer Defined Scale Unit (CDSU). CDSU describes the resources contained in a balanced architecture where the compute, Networking, and storage are built and sized appropriately for each other based on the customer’s initial requirement.

## Designing with Customer Defined Scale Units

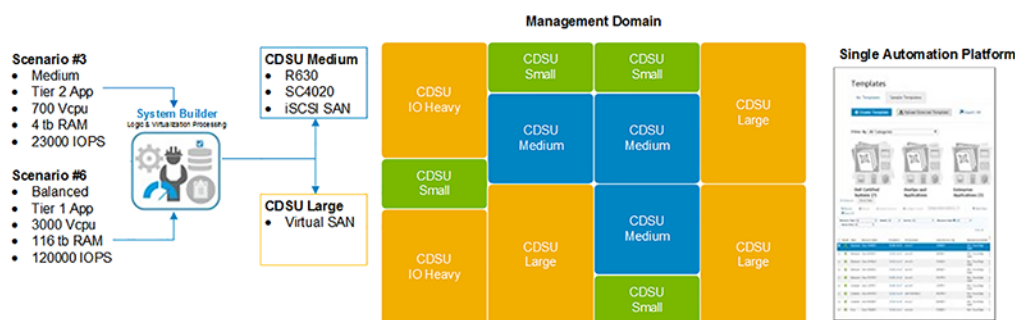


Figure 2. Customer Defined Scale Unit

Dell's proven server, storage, Networking, and management solutions each play a vital role in delivering the concept of CDSU. As an example, the modular Dell PowerEdge FX2 architecture combines the density and efficiencies of blades with the simplicity and small growth size of rack-based systems. The Dell SC Series of highly modular array components allows for easy scalability and hardware reuse for storage subsystems, while the Dell Z9100 switch can be used to create a simple yet scalable modular network. What really differentiates the Dell EMC Validated System for Virtualization solution from other integrated system offerings is the ability to configure, manage, scale, and repurpose different scale units irrespective of the underlying technology from a single management console, Dell's Active System Manager (ASM).

The flexible Customer Defined Scale Units, provide:

- **Flexibility to start small and grow in smaller chunks:** Customers can start small with four servers and a single storage array, then scale them either individually by adding more compute nodes, storage enclosures, or both compute and storage to maintain constantly changing workload demands.
- **Reduction of OPEX and operational complexity:** Adopting a single standardized automation platform transforms the model from an inefficient, overprovisioning platform to a pay-as-you-grow, flexible architecture model that allows for a significant reduction in OPEX and simplifies operational complexity.

## Best practices in a VMware virtualized infrastructure

The Dell EMC Validated System for Virtualization incorporates the best practices for a VMware virtualized infrastructure.

### Virtual switch configuration best practices

Virtual switches in the solutions are designed based on the following VMware recommendations:

- The architecture maintains standard and distributed virtual switches depending on the configuration
- Uses distributed virtual switches for Intel X520 NICs that are recommended for bandwidth allocation, which ensures appropriate bandwidth is allocated to VMware and ensures redundancy for management port
- Using Network I/O Control (NIOC) is recommended with distributed virtual switches to help shape the traffic and ensure that network traffic is prioritized as needed
- Using hardware iSCSI offload with Broadcom network adapters optimizes the host CPU utilization and improves host efficiency while network controllers handle iSCSI traffic distribution

## Dell Compellent best practices

The solutions recommend enabling virtual ports for the controllers. Virtual port mode provides port and storage controller redundancy by connecting multiple active ports to each Fibre Channel or Ethernet switch. In virtual port mode, each physical port has a virtual port associated with it (Virtual WWN for Fibre Channel). Servers target only the virtual ports. During normal conditions, all ports process I/O. If

a port or storage controller failure occurs, a virtual port moves to another physical port in the same fault domain. When the failure is resolved and ports are rebalanced, the virtual port returns to the preferred physical port.

## Network best Practices


The solutions are configured with Virtual Link Trunking (VLT). VLT allows two physical switches to represent a single logical switch. With physical links configured as a port-channel, the connection of two individual switches configured with VLT logically groups it as a single entity for the access switches that connect to the VLT domain. Both the VLT peers have their own configuration, control and data planes. VLT allows creating port-channels across two switches, eliminating the need for spanning tree blocking states across its connectivity and effectively utilizing its entire VLT links connected to the access or Top of Rack (ToR) switches

## Backup and recovery best practices

The Dell vRanger solution is installed on a physical server to take advantage of Virtual Disk Development kit (VDDK) and (LAN free) SAN backup. The solution ensures that the volumes are directly mounted in vRanger and backed up through the SAN fabric, without impacting the LAN fabric.

## Comprehensive

The Dell EMC Validated System for Virtualization solution brings together multiple Dell technologies to deliver the industry's most comprehensive integrated solution. As shown, the Dell EMC Validated System for Virtualization solution can eliminate or reduce various elements in the infrastructure lifecycle across design, sizing, planning, deployment, testing, and maintenance through comprehensive testing and validation and best practices. The Dell Active System Manager (ASM) provides end-to-end automation and a single point of assistance from the Dell ProDeploy and Dell ProSupport service offerings.



**Figure 3. The Dell EMC Validated System for Virtualization comprehensive solution**

## Customer adoption scenarios

Business requirements typically trigger the starting point of any new infrastructure project. Each company has a set of unique requirements which may include:

- **New initiative or market:** Improve customer experience, product introduction, or taking an existing product into a new market
- **Governance and Control:** An organization's set of internal IT processes and standards
- **Regulatory Requirements:** Regulatory compliance is an organization's adherence to laws, regulations, guidelines and specifications relevant to its business. Violations of regulatory compliance regulations may result in legal punishment.
- **Information Sensitivity:** The control of access to information or knowledge that might result in loss of an advantage or level of security if disclosed to others
- **Acquisition/Merger/Divestiture:** Due to the complex nature of these types of activities, it is common to set up a dedicated infrastructure

In addition to the business requirements, IT departments have their own triggers, which may include:

- **Transform IT Delivery Mode:** A want to move from resource silos to holistic services composed through template-based provisioning and orchestration
- **Infrastructure outside of support and maintenance:** For products that are out of Support, renewing the Support contract can often be expensive. When products transition through End of Support the IT organization's risk increases, for example a product may no longer receive security updates.
- **Performance and or Capacity limitations:** As existing resources are consumed, the capacity may not exist to upgrade to new features, for example upgrading to vSphere 6.0 update 2
- **Hardware upgrade costs prohibitive:** Although under Support and maintenance the underlying hardware may not have capacity for upgrades or be expensive because of limited availability (no longer manufactured)

Whether the project is to Support a new business initiative or driven by the IT organization, the requirements must be translated into infrastructure requirements. This topic presents tools used to determine the infrastructure requirements and provides sample configurations.

Topics:

- [Infrastructure Capacity Cube](#)
- [Dell Performance Analysis Collection Kit \(DPACK\)](#)
- [Dell System Builder](#)
- [Sample Configurations](#)

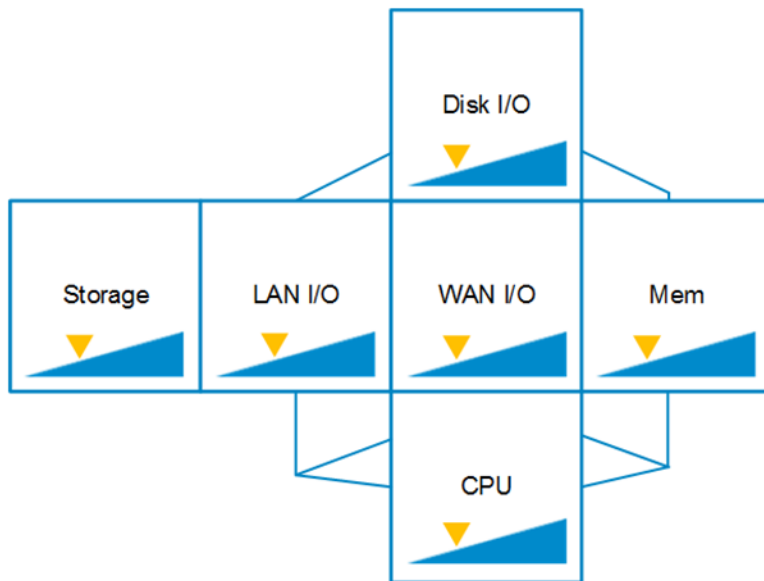
## Infrastructure Capacity Cube

To take advantage of the enhanced capabilities from the Dell portfolio of infrastructure products there are common attributes, which assist initial sizing and capacity planning. Although each customer's requirements are different, they are all anchored to a common set of attributes, which include:

- CPU
- Memory
- Storage
- Disk I/O

- LAN I/O
- WAN I/O

Depending on the workload or application requirements, it is common for Infrastructure teams to enter the design phase with one of these attributes. But to correctly size the infrastructure the other five attributes must be considered. Each attribute is important to optimize the infrastructure. Dell has created the Infrastructure Capacity Cube to help visualize the interdependencies.



**Figure 4. Dell Infrastructure Capacity Cube**

For new requirements, it is important work with the application developers and independent software vendors (ISVs) to characterize your profile using the Dell Infrastructure Capacity Cube. Alternatively, utilization metrics can be gathered from existing infrastructure to help size the replacement platform.

## Dell Performance Analysis Collection Kit (DPACK)

For the past five years, Dell has offered the Dell Performance Analysis Collection Kit (DPACK), an industry standard method of impartially documenting customer infrastructure performance from the perspective of the server. The kit chronicles performance across five sides of the Infrastructure Capacity Cube.

DPACK is a point-in-time performance-capture technology donated to the IT community by Dell to promote awareness of actual computer performance needs. The tool specifically helps users understand physical or virtual server workload performance at the individual machine, workgroup or data center level — all without requiring any dedicated experience or any expense. The complimentary technology measures the requests demanded of the operating system and applications, aggregates the metrics with the entire computing environment being measured, and provides a hardware-agnostic perspective of overall computer resource consumption.

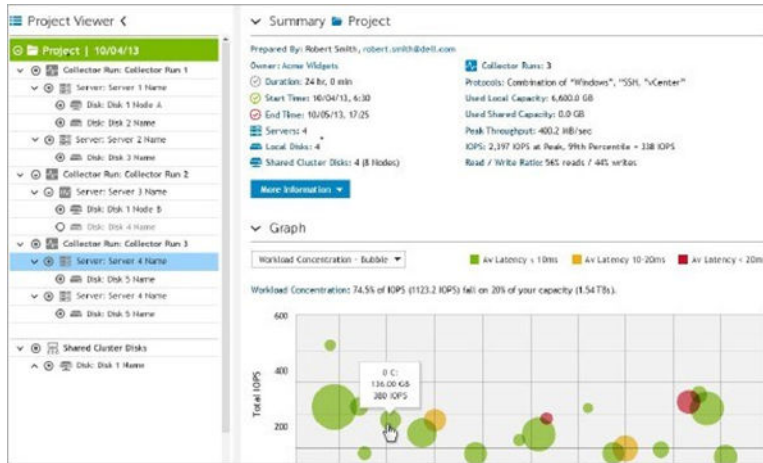
The core benefits of DPACK are:

- Agent-free and remote performance monitoring that does not require installation or change control
- Individual machine-to-datacenter level performance analysis
- Performance capture, which provides a hardware-agnostic evaluation of resources needed to accommodate core sizing components of IT environments such as:
  - Capacity
  - Disk throughput and I/O
  - Memory utilization

- Aggregate CPU cycles consumed

DPACK also includes additional diagnostic information including:

- Disk queue
- System latency
- Application I/O transfer size



**Figure 5. Dell Performance Analysis Collection Kit (DPACK)**

DPACK can be incorporated into ongoing performance and capacity planning to help customers make informed decisions for future enhancements, upgrades, or expansions.

## Dell System Builder

Dell System Builder provides a user interface that asks key questions about what the customer is trying accomplish including IT policies, configuration needs, and workload requirements. Based on these requirements and customer preference, Dell sales teams and channel partners can use the System Builder to create a tailored yet validated, integrated system that can be ordered to include hardware, software, and all services.



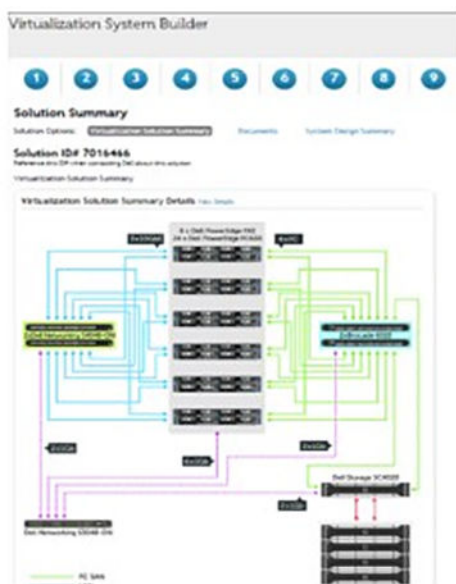


Figure 6. Dell System Builder

## Sample Configurations

This section presents sample configurations created using the Dell System Builder. Each of the sample configurations is based on nine questions presented within the System Builder interface. These examples use values from the Infrastructure Capacity Cube to determine the size of the recommended infrastructure.

### iSCSI Customer Defined Scale Unit - Small

- Total Cores: 176
- Total Memory: 2 TB
- Total Disk Capacity: 90 TB
- Total (Tier 1) IOPs: 6000

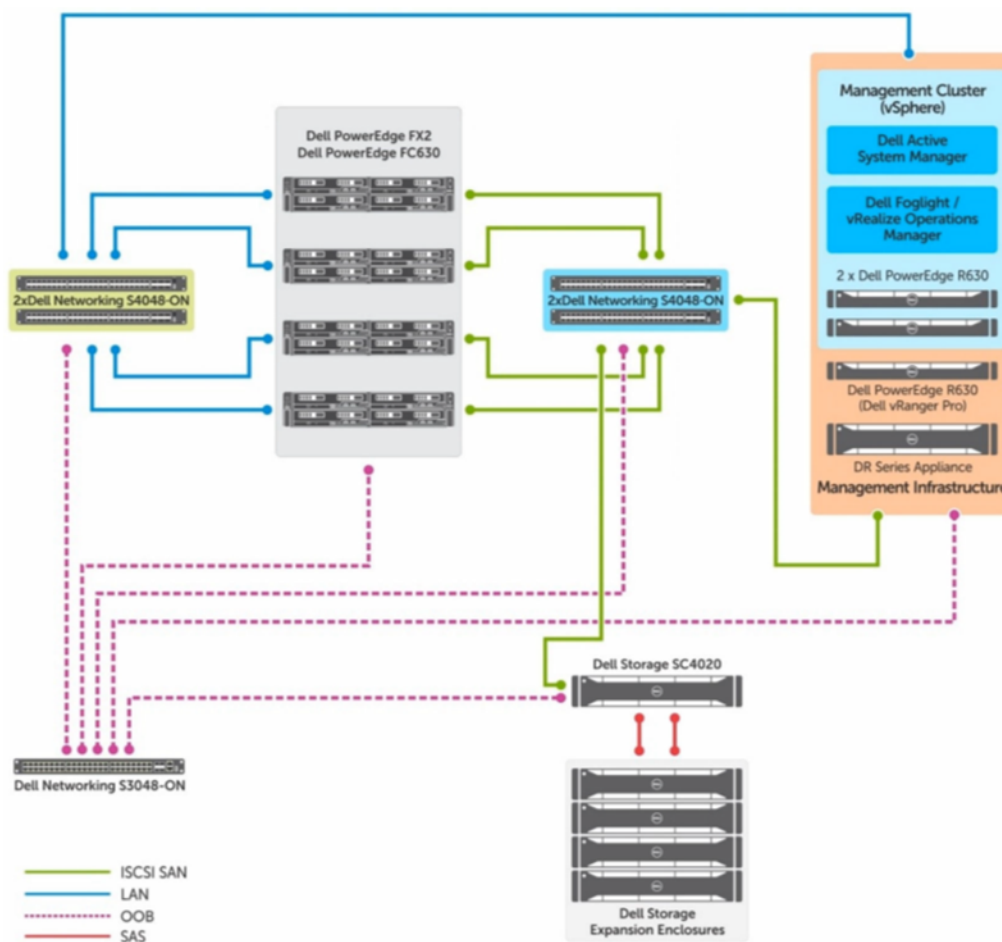


Figure 7. iSCSI Customer Defined Scale Unit - Small

## iSCSI Customer Defined Scale Unit - Medium

- Total Cores: 352
- Total Memory: 4 TB
- Total Disk Capacity: 69 TB
- Total (Tier 1) IOPs: 23515

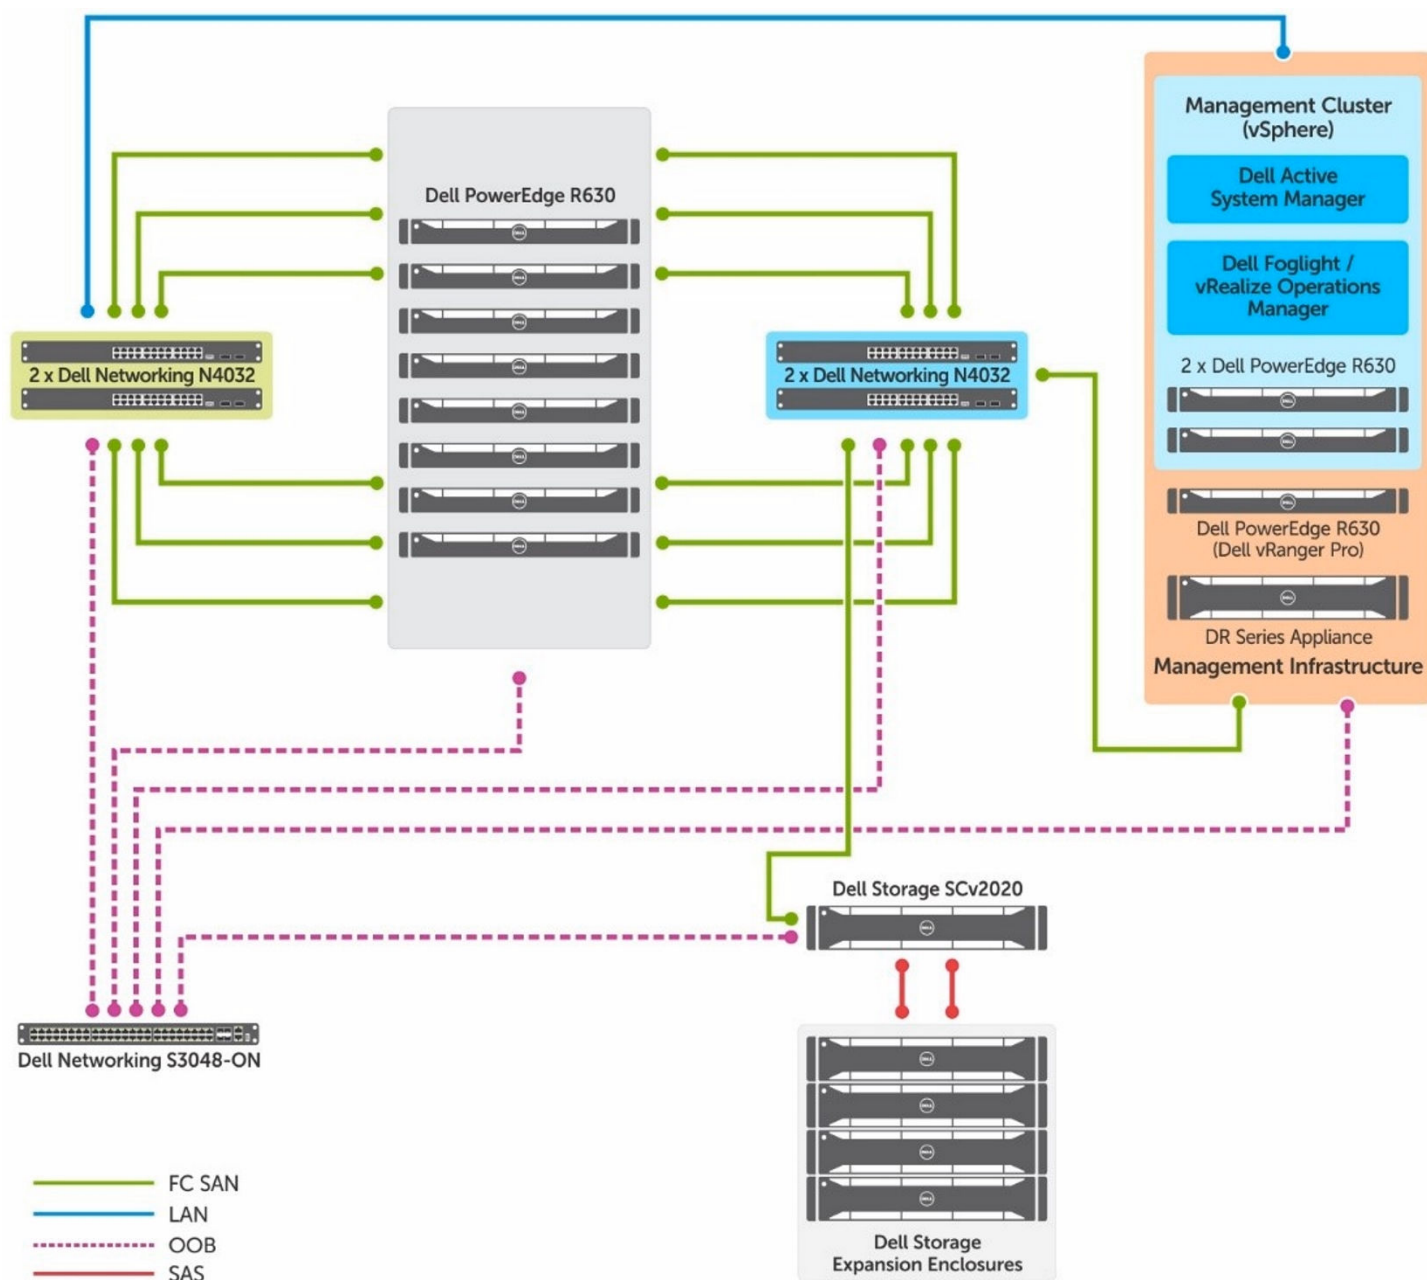


Figure 8. iSCSI Customer Defined Scale Unit - Medium

## iSCSI Customer Defined Scale Unit - Large

- Total Cores: 704
- Total Memory: 8 TB
- Total Disk Capacity: 140 TB
- Total (Tier 1) IOPs: 47030

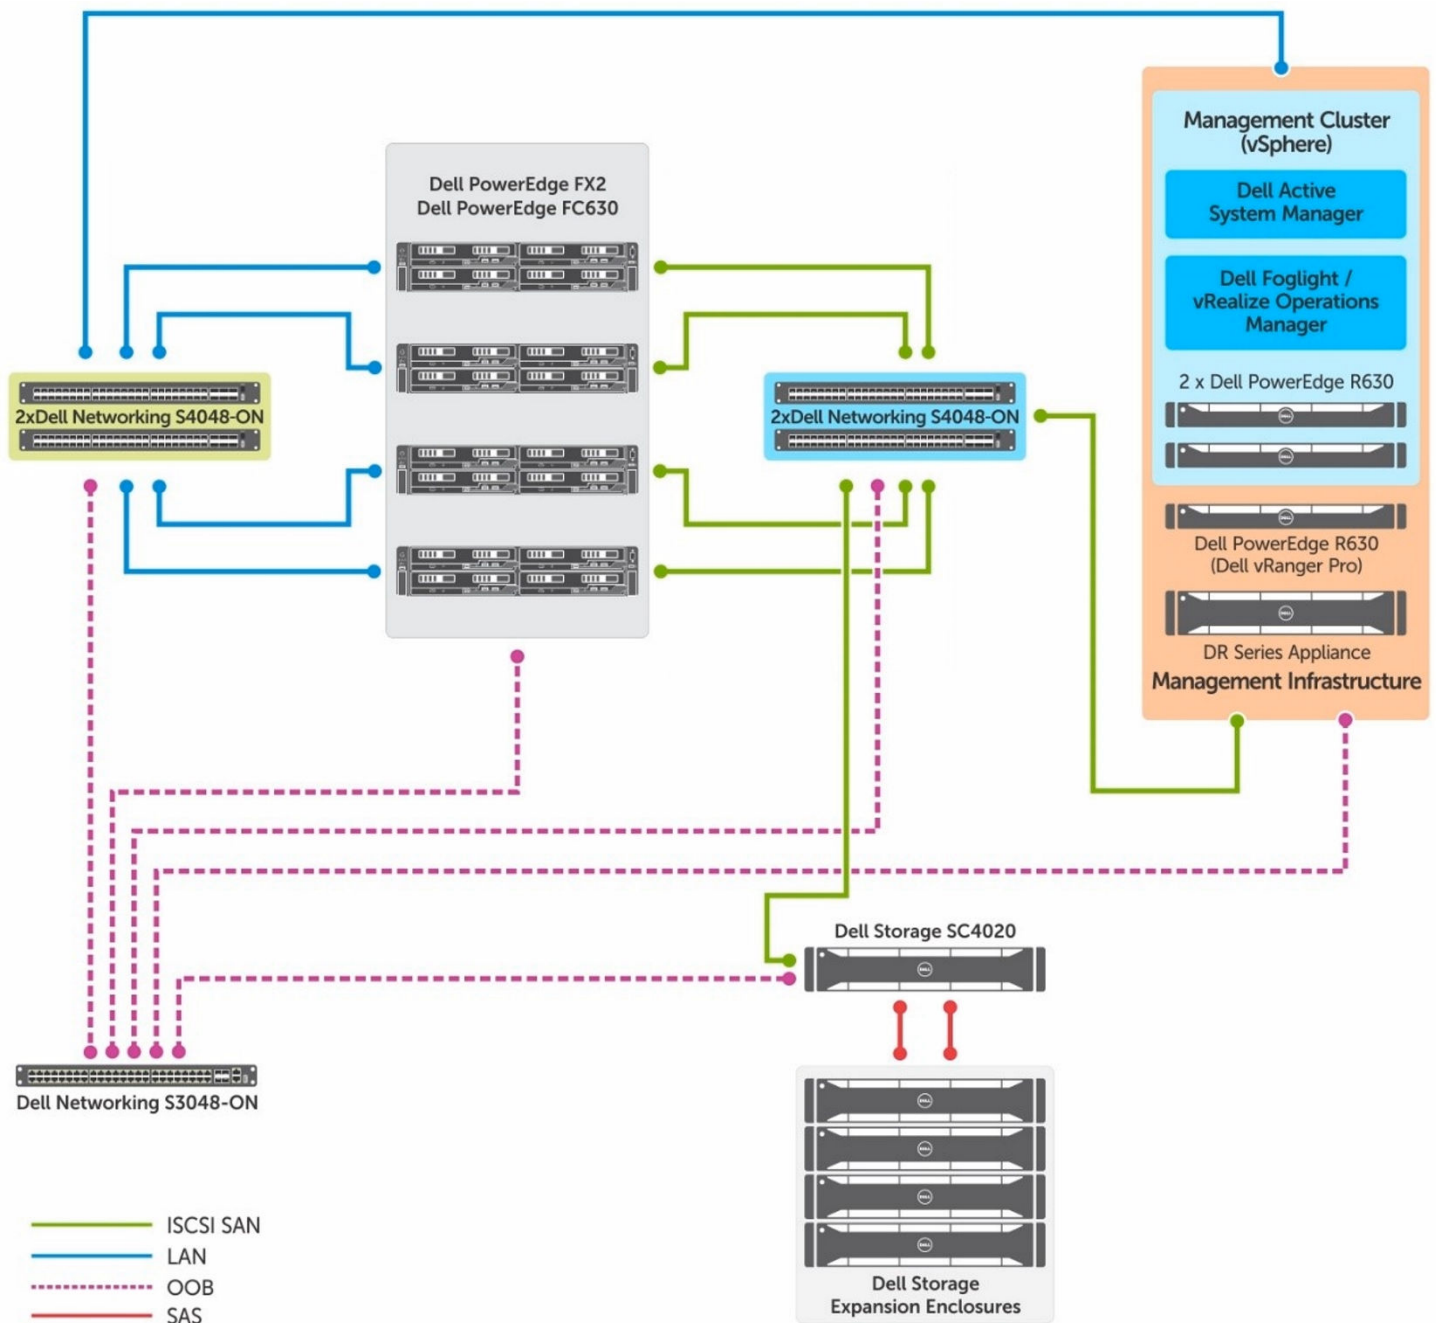
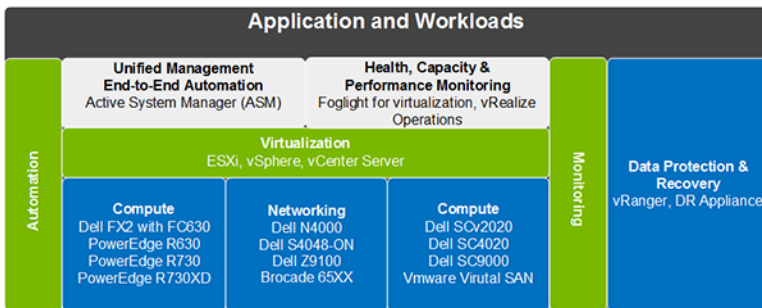


Figure 9. iSCSI Customer Defined Scale Unit - Large

# Deployment options for Dell EMC Validated System for Virtualization

The basis of the Dell EMC Validated System for Virtualization is the award winning data center portfolio comprised of servers, storage and networking products that are certified by VMware. To enhance the capabilities of the VMware framework, the Dell EMC Validated System for Virtualization includes Dell Software.



**Figure 10. Dell EMC Validated System for Virtualization**

Depending on the project requirements and the customer's technology preferences, the Dell EMC Validated System for Virtualization are available in the following deployment choices.

- Fibre Channel
- iSCSI
- Software Defined Storage using VMware Virtual SAN

The following table shows the Dell EMC Validated System for Virtualization portfolio aligned to fabric choices. The iSCSI column has been highlighted in this guide.

**Table 1. Dell EMC Validated System for Virtualization — iSCSI portfolio**

Portfolio		Fibre Channel	iSCSI	Software Defined
<b>Automation</b>	Active System Manager	X	X	X
<b>Monitoring</b>	Foglight	X	X	X
<b>Monitoring</b>	VMware vRealize Operations Manager	X	X	X
<b>Virtualization</b>	VMware vCenter	X	X	X
<b>Virtualization</b>	VMware vSphere	X	X	X
<b>Backup and recovery</b>	Dell vRanger	X	X	X
	DR4300	X	X	X
	DR6300	X	X	X

Portfolio		Fibre Channel	iSCSI	Software Defined
PowerEdge Rack	R630	X	X	
	R730	X	X	
	R730XD			X
PowerEdge FX2	FC430	X	X	
	FC630	X	X	
Storage	SCv2020		X	
	SC4020	X	X	X
	SC9000	X		X
	Virtual SAN			X
LAN Networking	N4032		SCv2020 only	
	S3048-ON for OOB	X	X	X
	S4048-ON	X	X	X
	Z-series	X	X	X
Storage Networking	Brocade 6500 series	X		
	S4048-ON	X	X	X

This section provides a brief introduction to the Dell EMC Validated System for Virtualization portfolio.

Topics:

- [Management Software](#)
- [Dell PowerEdge Rack Servers](#)
- [Dell PowerEdge FX2](#)
- [Dell SC Storage Solutions](#)
- [Dell Networking](#)

## Management Software

### Dell Active System Manager (ASM)

The Dell Active System Manager (ASM) is a unified management and automation solution that simplifies the deployment, provisioning, and management of heterogeneous resources, including both infrastructure and workloads, and delivers the benefits of automation to service-centric IT management.

For the Dell EMC Validated System for Virtualization, ASM provides broad capabilities that range from hardware discovery, onboarding, initial configuration, and inventory management to virtual infrastructure provisioning with the deployment of hypervisors and configuration of hosts and clusters. ASM can also provide virtual machine, workload, and provisioning capabilities.

## Dell Foglight for Virtualization Standard Edition

The easy-to-read dashboards and reports within the Dell Foglight for Virtualization Standard Edition provide a single-pane-of-glass view into the performance of your virtual environment. Comprehensive capacity and resource utilization tools help improve VM density and reduce operational expenditures.

Dell Foglight for Virtualization Standard Edition delivers:

- Simplified performance management
- Comprehensive capacity and resource utilization management
- Dynamic change tracking
- Insightful chargeback and showback reporting

## Dell vRanger

Data Protection | Dell vRanger software provides high-speed backup and recovery of VMware, Microsoft Hyper-V, and physical Windows Server environments and also provides high-speed replication for VMware. vRanger protects entire virtual environments in minutes, detects and backs up new VMs automatically while delivering safe, scalable data protection for even the largest VMware and Hyper-V environments. With vRanger, you can locate and restore individual files buried in virtual and physical backups, in seconds and from a single, intuitive interface.

## Dell DR4300 Disk Backup Appliance

The Dell DR4300 Disk Backup Appliance backup-to-disk appliance is optimal for most businesses, offering up to 108 tb of usable capacity and up to 22 tb/hr of throughput.

## Dell DR6300 Backup-to-Disk Appliance

The Dell DR6300 Backup-to-Disk Appliance delivers enterprise-level performance, offering up to 360 tb of usable capacity and up to 29 tb/hr of throughput.

## VMware vRealize Operations Manager

The VMware vRealize Operations Manager delivers intelligent operations management across physical, virtual, and cloud infrastructure — from VMware vSphere and Hyper-V to Amazon Web Services. The vRealize Operations Manager correlates data from applications to storage in a unified, easy-to-use management tool that provides control over performance, capacity, and configuration, with predictive analytics driving proactive action, and policy-based automation.

## VMware vCenter Server Appliance

The VMware vCenter Server Appliance provides a centralized and extensible platform for managing virtual infrastructure. The vCenter Server Appliance manages VMware vSphere environments and gives IT administrators simple and automated control over the virtual environment to deliver infrastructure with confidence. The vCenter Server Appliance also comes with a complimentary 25-OSI license of VMware vRealize Log Insight.

## VMware vSphere 6.0

VMware vSphere 6.0, the industry-leading virtualization platform, empowers users to virtualize any application with confidence, redefines availability, and simplifies the virtual data center. The result is a highly available, resilient, on-demand infrastructure that is the ideal foundation for any cloud environment.

## Dell Active System Manager (ASM)

The Dell Active System Manager (ASM) is a unified management and automation solution that simplifies the deployment, provisioning, and management of heterogeneous resources, including both infrastructure and workloads, and delivers the benefits of automation to service-centric IT management.

For the Dell EMC Validated System for Virtualization, ASM provides broad capabilities that range from hardware discovery, onboarding, initial configuration, and inventory management; to physical resource pooling, bare metal provisioning, and firmware compliance; to virtual infrastructure provisioning, with deployment of hypervisors and configuration of hosts and clusters; to virtual machine, application, and workload provisioning.

## Dell Foglight for Virtualization Standard Edition

Easy-to-read dashboards and reports provide a single-pane-of-glass view into the performance of your virtual environment. Comprehensive capacity and resource utilization tools help improve VM density and reduce operational expenditures.

The Dell Foglight for Virtualization Standard Edition delivers:

- Simplified performance management
- Comprehensive capacity and resource utilization management
- Dynamic change tracking
- Insightful chargeback and showback reporting

## Dell vRanger

Data Protection | vRanger software provides high-speed backup and recovery of VMware, Microsoft Hyper-V, 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 DR4300 Disk Backup Appliance

The Dell DR4300 Disk Backup Appliance backup-to-disk appliance is optimal for most businesses, offering up to 108 TB of usable capacity and up to 22 TB/hr of throughput.



## Dell DR6300 Backup Appliance

The Dell DR6300 Backup Appliance backup-to-disk appliance delivers enterprise-level performance, offering up to 360 TB of usable capacity and up to 29 TB/hr of throughput.

## VMware vRealize Operations Manager

VMware vRealize Operations delivers intelligent operations management across physical, virtual, and cloud infrastructure — from VMware vSphere and Hyper-V to Amazon Web Services. It correlates data from applications to storage in a unified, easy-to-use management tool that provides control over performance, capacity, and configuration, with predictive analytics driving proactive action, and policy-based automation.

## VMware vCenter Server Appliance

VMware vCenter Server provides a centralized and extensible platform for managing virtual infrastructure. vCenter Server manages VMware vSphere environments, giving IT administrators simple and automated control over the virtual environment to deliver infrastructure with confidence. vCenter Server also comes with a complimentary 25-OSI license of vRealize Log Insight.

## VMware vSphere 6.0

VMware vSphere 6.0, the industry-leading virtualization platform, empowers users to virtualize any application with confidence, redefines availability, and simplifies the virtual data center. The result is a highly available, resilient, on-demand infrastructure that is the ideal foundation for any cloud environment.

## Dell PowerEdge Rack Servers

Dell PowerEdge Rack Servers are self-contained, discrete servers with dedicated power supplies and network connections. PowerEdge Rack Servers are ideal for low-density deployments where power and cooling are constrained.

## Dell PowerEdge R630

With 2U performance packed into a compact 1U chassis, the Dell PowerEdge R630 two-socket rack server delivers uncompromising density and productivity. Part of the 13th generation of Dell PowerEdge servers, the PowerEdge R630 is ideal for compute virtualization, driving large business applications or running transactional databases.

## Dell PowerEdge R730

The incredible versatility of the Dell PowerEdge R730 server delivers outstanding functionality in just 2U of rack space. With the combination of powerful processors, large memory, fast storage options and GPU accelerator support, the PowerEdge R730 performs exceptionally well in various demanding environments.

## Dell PowerEdge FX2

Add infrastructure to the job with modular server, storage and networking blocks that combine neatly within the Dell PowerEdge FX2 converged 2U rack chassis. With an innovative modular design that accommodates IT resource building blocks of various sizes (compute, storage, networking, and management), the PowerEdge FX2 enables data center infrastructures to be built with greater flexibility.

## Dell PowerEdge FC430

A quarter-width, two-socket FX architecture server, the Dell PowerEdge FC430 offers the ultimate in shared infrastructure with the perfect balance of performance, memory and high-speed I/O ideal for midsize virtualization.

## Dell PowerEdge FC630

With powerful processors, impressively large memory and highly scalable I/O capabilities, the Dell PowerEdge FC630 excels at running applications for midsize and large businesses, such as enterprise resource planning and customer relationship management. Combining exceptional performance and functional flexibility, the PowerEdge FC630 is also ideal for large databases or substantial virtualization environments.

## Dell SC Storage Solutions

From enterprise-class Fibre Channel, iSCSI and FCoE SANs to highly-scalable NAS solutions, Dell Storage SC Series arrays provide a unified platform for the ultimate in performance, adaptability and machine-driven efficiency.

## Dell SC Storage SCv2020

The high performance and quality of the entry-level Dell SC Storage SCv2020 is perfect for small and medium businesses needing primary storage, or for project-based storage. A 2U, twenty-four 2.5" HDDs up to 168 HDDs with 1gb iSCSI, 10gb iSCSI.

## Dell SC Storage SC4020

The Dell SC Storage SC4020 belongs to the Storage Center (SC) 4000 series arrays based on the SC8000 platform. The SC Storage SC4020 array offers multiprotocol support and virtualized multi-tier, multi-RAID storage policies. Each SC Storage SC4020 array comes with dual redundant controllers, 24 internal drive slots, eight 8gb Fibre Channel (FC) or four 10gb iSCSI network ports and one 1gb port per controller for outband management.

## Dell SC Storage SC9000

The Dell SC Storage SC9000 provides an ideal solution for large-scale storage, high-end workloads and distributed enterprise environments. The latest Dell flagship array offers industry-leading capabilities of the SC Series line of systems in a fully upgraded, more scalable system using the Dell 13G server platform.

## Dell Networking

Based on open standards, Dell Networking solutions improve network performance, lower datacenter costs and remain flexible to adopt innovations.

## Dell Networking N4032

The N-Series is a family of energy-efficient and cost-effective 1 GbE and 10 GbE switches designed for modernizing and scaling network infrastructure. Twenty-four 10GbE RJ45 autosensing (10Gb/1Gb) fixed ports. Up to thirty-two 10GbE ports using breakout cables and optional QSFP+ module

## Dell Networking S3048-ON

Accelerate applications in high-performance environments with a low-latency ToR switch that features 48 x 1 GbE and 4 x 10 GbE ports, a dense 1U design and up to 260 Gbps performance. The Dell Networking S3048-ON also supports Open Network Installation Environment (ONIE) for zero-touch installation of alternate network operating systems.

## Dell Networking S4048-ON

The 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 uses a nonblocking and cut-through switching architecture to provide ultra low-latency performance for applications. The six 40 GbE ports can be used to create a Virtual Link Trunk (VLT) between the switches to enable traffic isolation within the solution infrastructure. The Dell Networking S4048-ON can also be used or as a connection to the data center core network.

## Dell Networking Z9100-ON

The Dell Networking Z9100-ON is an open networking, SDN-ready fixed switch that is purpose-built for applications in high-performance data centers and for computing environments such as big data, cloud and Web 2.0. The Dell Networking Z9100-ON offers multirate switch 10/25/40/50/100GbE options with easy on-ramp to SDN for part or all production environment needs.

# iSCSI architecture design

The architecture design must ensure that there is component level high availability. For both storage and LAN fabrics, high availability at the component level implies that there are multiple data paths from compute to storage and conversely. This ensures that the traffic does not get interrupted even if there is a component failure.

The solution architecture also ensures high availability during system maintenance such as firmware updates. The design ensures that connectivity to the network and SAN is maintained and that workloads continue to run, even during these maintenance events.

The architecture is designed to ensure that there are no bottlenecks within the solution. Processor, memory, network, disk I/O operations and disk capacity are balanced for a general purpose virtualization infrastructure. The redundancy built into the solution architecture, while enabling component high availability, should ensure that the redundant data paths, in a normal state, are load balanced. For example, in the compute infrastructure that is implemented using Dell PowerEdge R630 with quad port network adapters, the network ports connected to top of rack switches offer both high availability at the component level and load balanced paths for network traffic. The teamed host network configuration ensures that the network paths are load balanced.

Topics:

- [Compute Architecture](#)
- [Network Architecture](#)
- [Storage Architecture](#)
- [Management Architecture](#)
- [Data Protection Architecture](#)

## Compute Architecture

It is important to select the right processor and memory for a virtualization infrastructure. High processor core count and frequency are recommended as is a processor with a higher number of cores to increase the possible number of VMs per node. Higher CPU speed improves the performance of all the VMs in the server.

Processors with higher cores and speed have high-power requirements. Some high dense servers, such as the Dell PowerEdge FC430, do not support processors that have a high power consumption. The Intel Xeon Processor E5-2699 v4 is the recommended processor for all servers except the PowerEdge FC430. For the PowerEdge FC430, the Intel Xeon Processor E5-2695 v4 is recommended.

Dell PowerEdge servers support RDIMMs and LRDIMMs at both 2133MT/s and 2400 MT/s. LRDIMMs are slightly more expensive than RDIMMs and support higher density. For general-purpose virtualization solutions, RDIMMs should suffice, however 2400 MT/s is recommended. Memory can be configured in various modes in the BIOS. Optimizer mode is the default mode and is recommended for most virtualization use cases to provide the optimized memory performance. For improved reliability and resiliency, other modes such as mirror mode and Dell Fault resilient mode are available.

The solution supports both Intel and QLogic network adapters. QLogic adapters support network partitioning that aid in bandwidth allocation for various traffic types and iSCSI offload. Intel X520 adapters do not support network partitioning.

Dell servers support various system profiles in BIOS configuration. The system profile dictates how the BIOS or operating system controls the processor and memory, and balances performance and power consumption. The default profile is "Performance". Under this mode, the processor is set to maximum performance. For customers who require energy-efficiency setting, the system profile can be changed to "Performance Per Watt Optimized".

On each PowerEdge server, ESXi is installed on the Dell Internal Dual SD Module. The dual SD module is configured in mirror mode to provide redundancy for ESXi. The servers are diskless, which saves the cost of HBAs and disks, and the cost of the associated power and cooling.

## Network Architecture

The iSCSI network configuration deployed in this solution architecture uses a non-converged configuration, which means the LAN and iSCSI SAN fabrics are implemented using separate switch and adapter configurations. For both these fabrics, Dell Networking 10GbE Ethernet switches are used. The following sections provide details on the network configurations employed in this architecture.

### MLAG configuration at Dell Networking N4032

Dell Networking N4032 switches support Multi-Chassis Link Aggregation (MLAG) that supports an active-active split aggregation deployment across two switches creating a more resilient network with higher bandwidth capabilities.

The Dell Networking N4032 switches provide twenty-four 10Gb BASE-T switch ports that are used for LAN and iSCSI connectivity. The optional 40GbE ports in these switches are used to create a Link Aggregation Group (LAG) as a peer link between two switches in the MLAG domain. If the 40GbE ports are needed for core connectivity, this peer link can be configured between any non-management ports in the Dell Networking N4032 switches. Within the MLAG domain, the primary and secondary role is assigned automatically by the switch firmware.

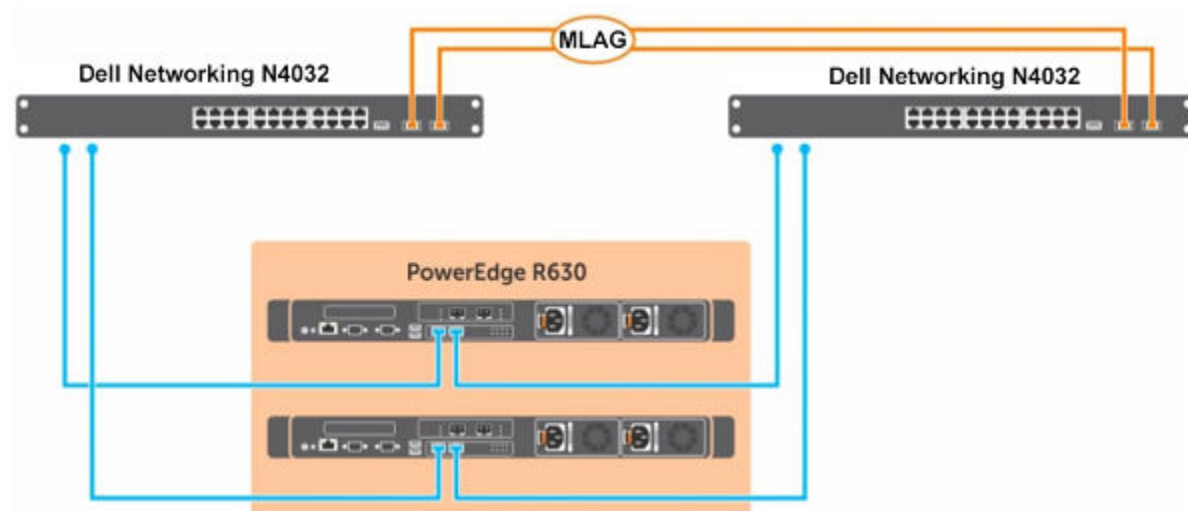


Figure 11. N4032 MLAG Configuration

### Virtual Link Trunking (VLT) at Dell Networking S4048-ON

The network architecture employs Virtual Link Trunking (VLT) between the two TOR switches. In a non-VLT environment, redundancy requires idle equipment which increases switch costs. This 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. A VLT pair of switches appears as a single switch to 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



**Figure 12. S4048-ON configured with VLT**

The Dell Networking S4048-ON switches provide six 40GbE uplink ports. The VLTi configuration in this architecture uses 2 x 40GbE ports from each ToR switch and thus provides 80GbE data path between the ToR switches. The remaining 4 x 40GbE ports provide either data center core connectivity and/or for connecting multiple virtualized infrastructure PODs — referred to as Customer Defined Scale Units. See [Scaling the Solution](#) for additional information.

**NOTE:** For the iSCSI SAN fabric, the MALG and VLT network configuration is not implemented.

## Network Connectivity FX2

- **Connectivity between hypervisor hosts and network switches:** The compute cluster hypervisor hosts PowerEdge FX server nodes and connect to the S4048-ON switches through the Dell Networking FN410 modules in the PowerEdge FX2 blade chassis. The PowerEdge compute rack servers (R730 and R630) and R630 management servers connect directly to the S4048-ON and N4032 switches.
- **Connectivity between the Dell PowerEdge FX server nodes and Dell Networking FN410:** The internal architecture of PowerEdge FX2 chassis provides connectivity between an Intel or Qlogic 10G Network Daughter Card (NDC) in each PowerEdge FX2 server node and the internal ports of the FN410 module. The FN410 has eight 10Gb internal ports. Four FC630 server nodes connect to four of the internal ports in each FN410. Eight FC430 server nodes connect to all eight of the internal ports in each FN410.
- **Connectivity between the Dell PowerEdge FN401 and S4048-ON switches:** The two Dell PowerEdge FN410S IOAs in the PowerEdge FX2 architecture provide the top-of-rack (ToR) connectivity for the PowerEdge FX2 server nodes. Each IOA provides four external ports. FN IOA provides a 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, as shown in the following figure.




Figure 13. Dell PowerEdge FN401 and Dell Networking S4048-ON switches

## Network Configuration PowerEdge Rack Servers

- Connectivity between the compute and management PowerEdge rack servers and Dell Networking S4080-ON switches: The compute and management rack servers have two 10Gb connections to Dell Networking S4048 switches through one Dual Port 10Gb Network Daughter card (NDC).





Figure 14. S4080-ON Connectivity between Compute and Management PowerEdge Servers

- Connectivity between the compute and management PowerEdge rack servers and Dell Networking N4032 switches: The compute and management rack servers have two 10Gb connections to the Dell Networking N4032 switches through one Dual Port 10Gb BASE-T Network Daughter card.



**Figure 15. N4032 Connectivity between Compute and Management PowerEdge Servers**

## Network Configuration for LAN traffic with QLogic and NPAR

Each Dell PowerEdge FX2 server node or PowerEdge rack server is configured with a dual port NDC. If QLogic NDCs are used, customers can utilize switch independent network partitioning (NPAR) to ensure the bandwidth prioritization for different traffic classes such as host management, vMotion and VM network. By using NPAR, each port on the bNDC is partitioned into four logical partitions. The QLogic blade NDC adapters allow setting a maximum bandwidth limitation to each partition. Setting the maximum bandwidth at 100 prevents the artificial capping of any individual traffic type during periods of noncontention. For customers with specific requirements, the 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. The following table shows the NPAR schema and the relative bandwidth assignments on each partition.

**Table 2. NPAR Schema and Relative Bandwidth Weights**

Server Physical Port	Network Partition Label	Traffic Class	Relative Bandwidth	Maximum Bandwidth
0	NIC0	Host Management	10	100
	NIC2	vMotion	40	100
	NIC4	VM Network	50	100
	NIC6	NA		
2	NIC1	Host Management	10	100
	NIC3	vMotion	40	100
	NIC5	VM Network	450	100
	NIC7	NA		

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



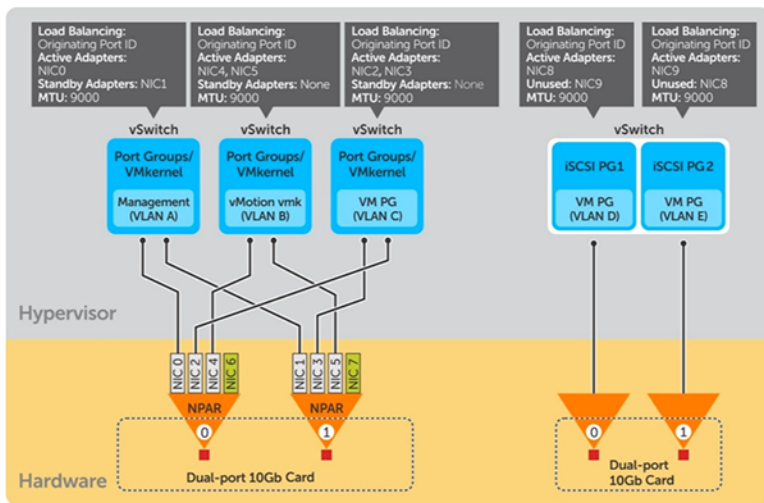


Figure 16. Host LAN Configuration Using NPAR

① **NOTE:** For configurations implemented using Dell Networking N4032 switches, Qlogic 57810 10Gb dual port (BASE-T) are used.

## Network Configuration for LAN Traffic with Intel/Qlogic Network Daughter Card and Distributed Virtual Switch

If Intel NDC are used, customers can achieve bandwidth prioritization for different traffic classes such as host management, vMotion, and VM network using VMware Distributed Virtual Switches (DVS). This DVS configuration applies to Qlogic network adapters as well.

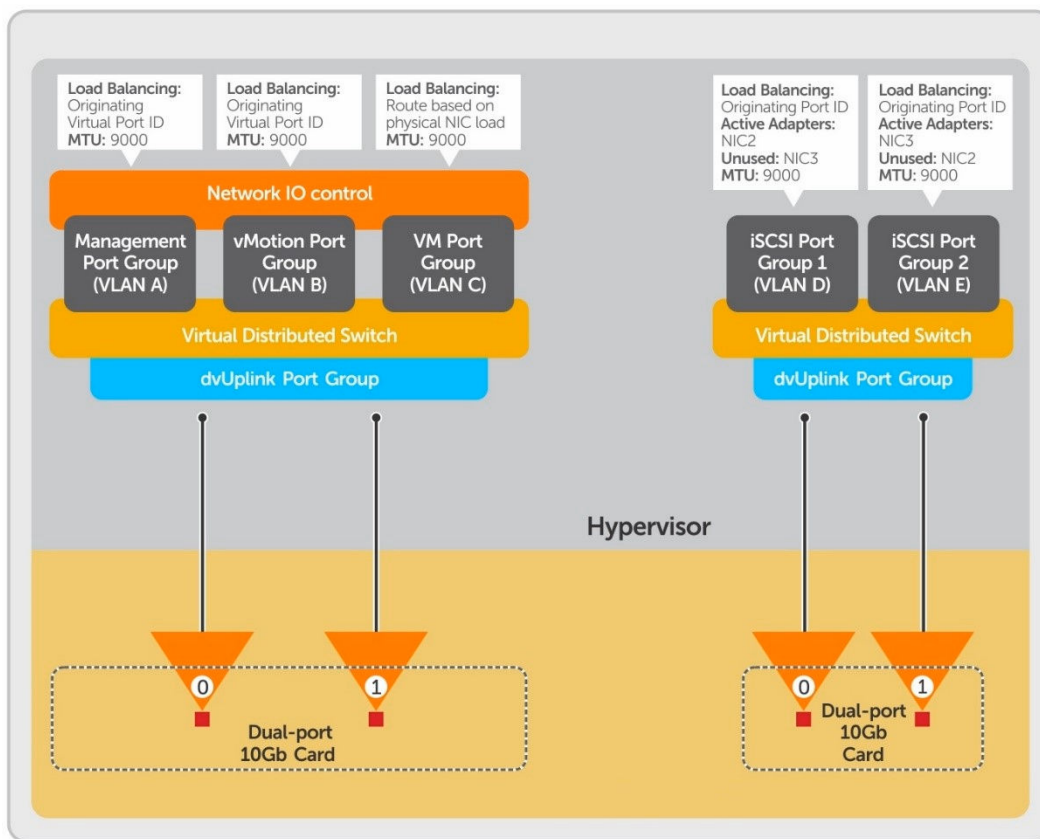


Figure 17. VMware vSphere Distributed Switch

## vSwitch setup

You can use either standard virtual switches or distributed virtual switches. Standard virtual switches are easy to deploy and configure. Additional licensing is not required. For each vSwitch that is created on the ESXi host, the load balancing algorithm is configured to be "Originating Port ID". This enables multiple active ports and utilization of the full network bandwidth available to the host. This setting must be configured on all port groups that connect using the same set of uplinks. Active and standby adapters are configured as shown in the following figure.

VMware vSphere Distributed Switch (VDS) can be configured, managed and monitored from a central interface. The VDS provides:

- Simplified virtual machine network configuration
- Enhanced network monitoring and troubleshooting capabilities
- Support for network bandwidth partitioning when NPAR is not available

**NOTE:** Currently, Intel network adapters do not support NPAR-based vSwitch configuration. Therefore, when using Intel adapters, it is assumed that the solution implements a Distributed Virtual Switch configuration.


As shown in the preceding diagrams, iSCSI connectivity is implemented using either standard or distributed virtual switch configurations.

## Storage Architecture

Dell Storage Center provides a highly available platform for the virtual data stores. The Dell Storage SCv2020 and Dell Storage SC4020 storage arrays enable high availability and redundant data paths by implementing fault domains.

## Storage Fabric configuration

This solution is configured with two fault domains for high availability. For the Dell Storage SC4020 arrays, port 1 from each controller connects to FC switch 1, while port 2 connects to FC switch 2, as shown in the following figure. Each iSCSI switch represents a fault domain for the iSCSI storage and enables port failover.



**Figure 18. Dell Storage SC4020 Arrays**

Similar to the SC4020, the SCv2020 has two controllers, each with two 10GbE (base-T) iSCSI ports. The SCv2020 is connected to the iSCSI SAN switches as shown in the following figure with two fault domains.

For the iSCSI solution infrastructure built using SCv2020 arrays, Dell Networking N4032 switches are used for the iSCSI SAN fabric.

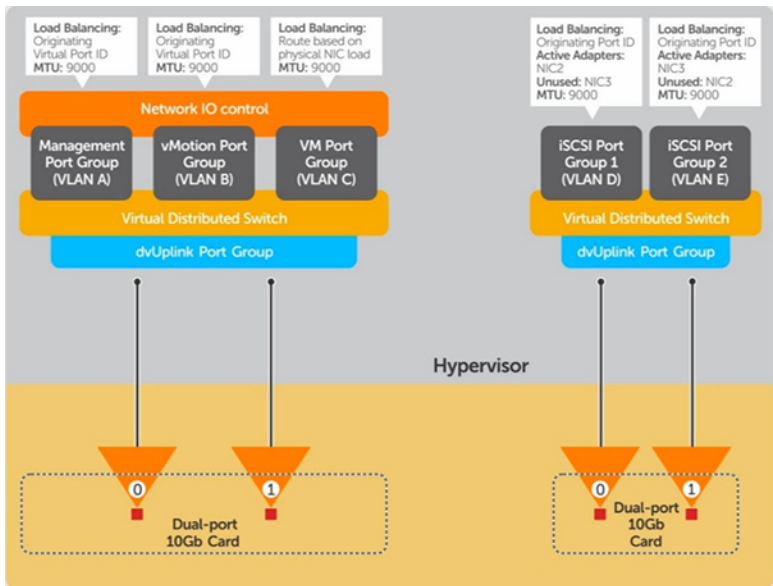


Figure 19. Dell Storage SCv2020 Arrays

## Storage connectivity

To connect the compute cluster to the iSCSI fabric, Qlogic57810/Intel X520 dual port NICs are used. The PowerEdge FX2 architecture provides the PCIe slots that are mapped through a PCIe switch to the PowerEdge FX server nodes in the chassis. The following figure illustrates how the FC630 servers are connected to the SAN fabric. If the FX2 chassis is populated with FC430s, all the PCIe slots are populated with Qlogic HBAs.




Figure 20. Server to SAN Fabric connectivity using PowerEdge FX2 PCIe slots

## Storage connectivity PowerEdge Rack

For management servers and rack compute servers, each server is configured with a QLogic 57810 dual port NICs, with each port connecting to the iSCSI SAN switch as shown. For a configuration implementing SCv2020 iSCSI storage arrays, Dell Networking N4032 switches are used.

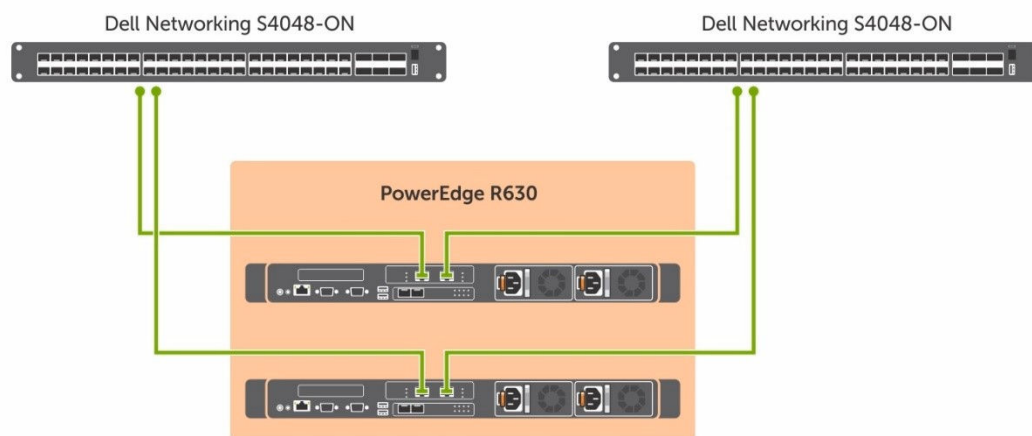


Figure 21. Server to iSCSI switch

## Storage configuration

For high availability of iSCSI connections, ports from each iSCSI NIC are connected to two different switches. These ports, along with the FC connections from the SC array, are configured to be in the same zone to enable storage volume access to the compute cluster. The following figure shows the FC volumes mapped as data stores on the ESXi hosts within the vSphere 6 cluster.

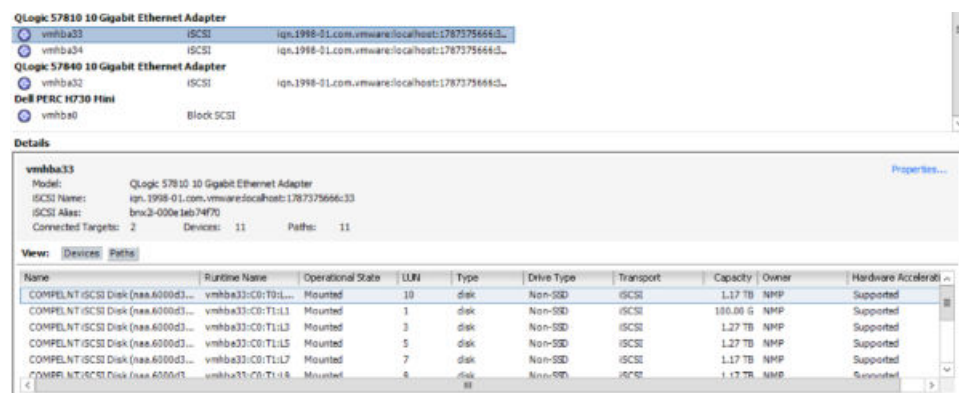
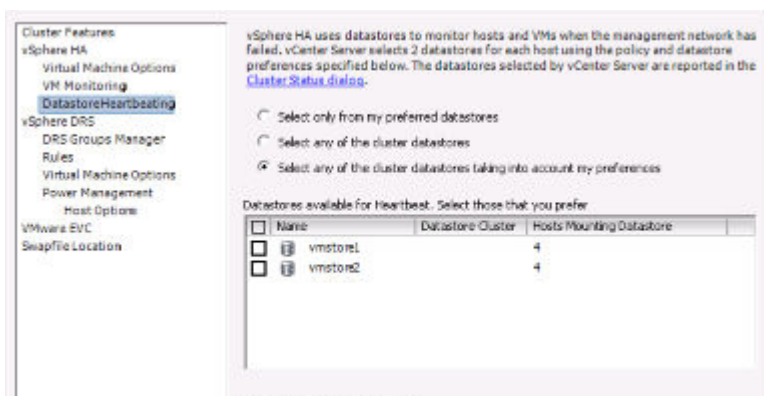


Figure 22. FC Volumes Mapped as Data Stores in vSphere Cluster

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. The following figure 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.



**Figure 23. Datastore Heartbeat 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 high availability of the data stored.

## Drive and RAID configuration for SC4020 and SCv2020

The Dell Storage Center arrays can be configured with a combination of Write-Intensive (WI) SSD drives, Read-Intensive (RI) SSD drives and 10K SAS drives. This hybrid deployment offers both performance and the capacity needed for the solution infrastructure.

**Table 3. Sample configuration of SC4020 Storage Array**

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

**Table 4. Sample configuration of SCv2020 Storage Array**

Component	Description
Disk Drives	24 x 10K SAS 1.8 TB
Storage Profile	RAID 5 with stripe width 9
	RAID 6 with stripe width 10
Volumes	2 x 6 TB volumes for VM Store

The choice of hybrid drive configuration in the SC4020 enables scaling of the compute infrastructure. This performance capability of the hybrid configuration can be complemented by adding more capacity in terms of Dell Storage SC220. See the [Scaling the Solution](#) section

for guidance on how the compute and storage scale-out can be achieved, and the considerations for the scale-out. The preceding tables provide the configuration of the Dell Storage SC4000 and SCv2000 series arrays used in this solution.

The available drive configurations of the SCv2000 series offer a flexible choice between a capacity versus performance configuration. The SCv2020 storage controller supports up to 24 2.5-inch drives and additional capacity can be added by connecting SC120 or SC100 storage enclosures. See the [Scaling the Solution](#) section for guidance on how the compute and storage scale-out can be achieved, and the considerations for the scale-out.

## Management Architecture

The [Management Software](#) section describes the recommended software components for the management infrastructure. Most of the management software except Dell vRanger Pro, the Dell DR4300, and Dell DR6300 appliances can be implemented as virtual machines and can share the storage capacity available on the Dell SC series Storage array, which is used for the virtual infrastructure. The storage and network architecture of the management servers are described in previous sections.

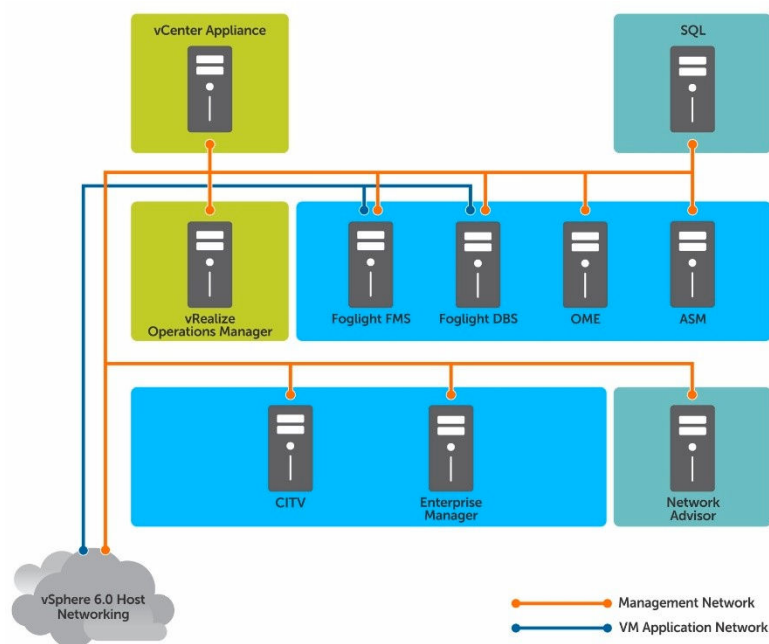
The following table provides an overview of the virtual infrastructure requirements for deploying the management components described in the [Management Software](#) section.

**Table 5. Management sizing**

Component	VMs	CPU Cores	RAM (GB)	OS (GB)	Data (GB)	NIC (Gbps)
VMware vCenter Server	1	4	12	60	50	1
Dell ActiveSystem Manager	1	4	16	80	0	1
Dell Foglight for Virtualization	2	4	4	40		1
Dell Enterprise Manager	1	2	4	40		1
Dell Storage vSphere Web Client Plugin	1	4	4	40		1
Dell Compellent Integration tools for VMware	1	2	8	40		1
Dell OpenManage Essentials	1	2	8	40	50	1
VMware vRealize Operations						

When virtualizing these management components, administrators must ensure that the infrastructure is highly available. For storage and network, the underlying virtual infrastructure provides high availability. Two 2 tb volumes are created to host the virtual machines on the Compellent array. 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 following figure illustrates the management architecture connecting to the existing LAN and SAN infrastructure used for the compute.



**Figure 24. Management Architecture**

The following table provides the specifications for the servers in the management infrastructure.

**Table 6. Server specifications for Management infrastructure**

Component	Details	
Virtualization Infrastructure for management software	2 x Dell PowerEdge R630	
	Processor	2 x Intel Xeon E5-2660v3 Family
	Memory	128 GB, 8 x 16 GB DDR4 DIMMs
	OS Volume	2 x 16 GB SD cards in internal SD module
	Network	QLogic 57810S dual-port NDC
	SAN	QLogic 57810S dual-port PCIe
	OS	VMware ESXi 6 U2
Dell vRanger Pro	1 x Dell PowerEdge R630	
	Processor	2 x Intel Xeon E5-2609v3 Family
	Memory	64 GB, 8 x 8 GB DDR4 DIMMs
	OS Volume	2 x 600 GB 15K SAS drives in RAID 1
	Network	QLogic 57810S dual-port NDC
	SAN	QLogic 57810S dual-port PCIe
	OS	Windows Server 2012 R2

For deploying the management components described in this table, two PowerEdge R630 servers are used in a VMware vSphere cluster. This virtualized implementation considers 2:1 vCPU oversubscription and no memory oversubscription. By using a vSphere cluster along with high availability and Distributed Resource Scheduler (DRS) features, the virtualized management infrastructure can be made highly available and resilient even if there is a host failure. Dell vRanger Pro is deployed on a PowerEdge R630 server and provides the backup



and recovery functionality for the virtual infrastructure. The DR4300 or DR6300 appliance is a physical host that gets connected to the existing LAN infrastructure that is used for the compute.

## Data Protection Architecture

The Dell Data Protection solution for the Dell PowerEdge FX2 converged system uses the DR series 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 satisfies demanding backup workloads at scale. Also, applying SAN for backup transport avoids taxing virtualization resources for backup purposes and conserves hypervisor capacity for production workloads.

For Backup, the solution uses the VMware integration of Dell vRanger to acquire virtual disk read access and SAN transport via vSphere. Upon accomplishing virtual disk access, 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 back up storage. More 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 DR series appliance and vSphere to establish a virtual disk write access SAN transport. For disaster recovery, the Disaster Recovery Appliance replicates from the DR series appliance across the WAN to a remote physical or virtual DR appliance that resides in an enterprise network or in a public cloud.

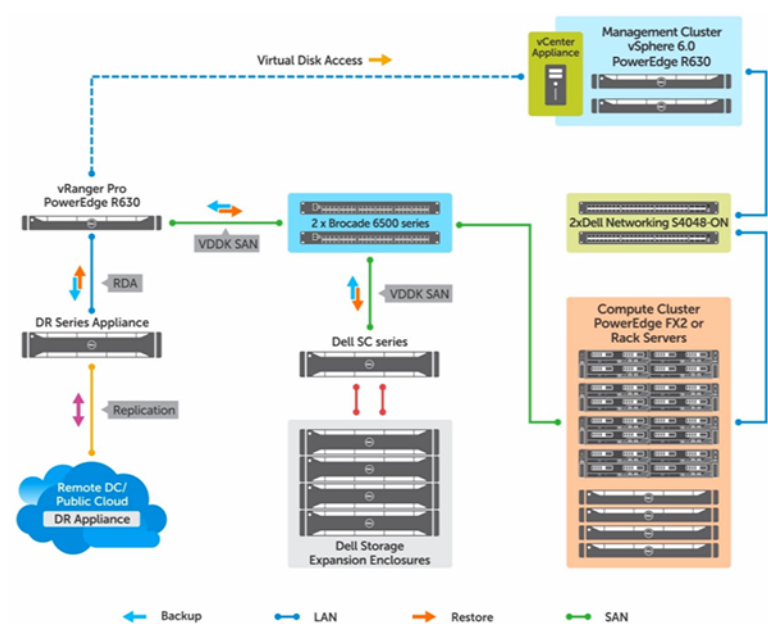


Figure 25. Data Protection Architecture

## Scaling the solution

Customers are often forced to either buy a small solution that needs a fork-lift upgrade later or buy a large solution with over provisioned capacity that may never be required.

The Dell EMC Validated System for Virtualization is built around the concept of the Customer Defined Scale Unit (CDSU). Each CDSU is a self-contained, balanced set of resources where the computer, networking, and storage are built to be sized appropriately for each other based on the customer's initial requirement. Through Dell System builder, a customer can, in their own terms, define small, medium, large Infrastructure-Units for their deployments.

Once deployed Infrastructure-Units can be expanded by adding resources, such as compute, storage and networking or more CDSUs.

Topics:

- [Scaling out with Customer Defined Scale Unit \(CDSU\)](#)
- [Scaling out compute and storage independently](#)
- [Scaling network infrastructure](#)

## Scaling out with Customer Defined Scale Unit (CDSU)

The Dell EMC Validated System for Virtualization solution allows the customer to create their own catalog of Customer Defined Scale Unit (CDSU). The Dell CDSU can be deployed individually, such as, a single site or combined within the same management domain. A benefit of the Dell CDSU is unified operating and management irrespective of the deployed architecture (PowerEdge rack servers, PowerEdge FX2 servers, iSCSI Compellent storage, Fibre Channel Compellent storage or VMware Virtual SAN).

This approach allows Infrastructure teams to provide tiered service offerings based on the attributes of each CDSU. Customers can start with any of the three deployment choices (Fibre Channel, iSCSI, or Software Defined Storage), and add more CDSUs under the same management framework. Extra capacity in terms of a CDSU can be added to existing deployments by requesting the management infrastructure not be included with the quote.

**NOTE:** To minimize disruption with upgrades, Dell recommends that customers consider regular driver and firmware updates, such as quarterly update packages the Dell Active System Manager (ASM) provides.

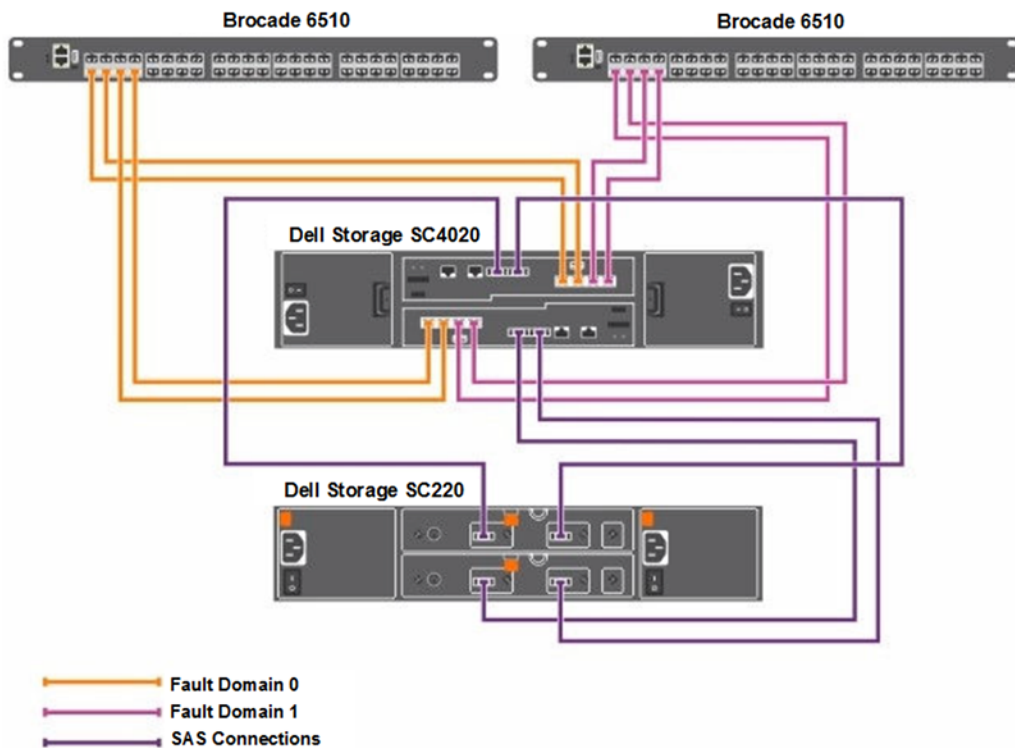


**Figure 26. Scaling out**

## Scaling out compute and storage independently

The architecture shown in the Solutions architecture sections for both racks and FX2 platforms represents the basic building block for a scalable and highly available virtual infrastructure that is built 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, storage or both. For expanding the storage, the Dell Storage SC220 or SC200 array can be attached to the Dell Storage SC4020 arrays. For Dell Storage SCv2020 arrays, SC120s can be attached to the arrays.

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

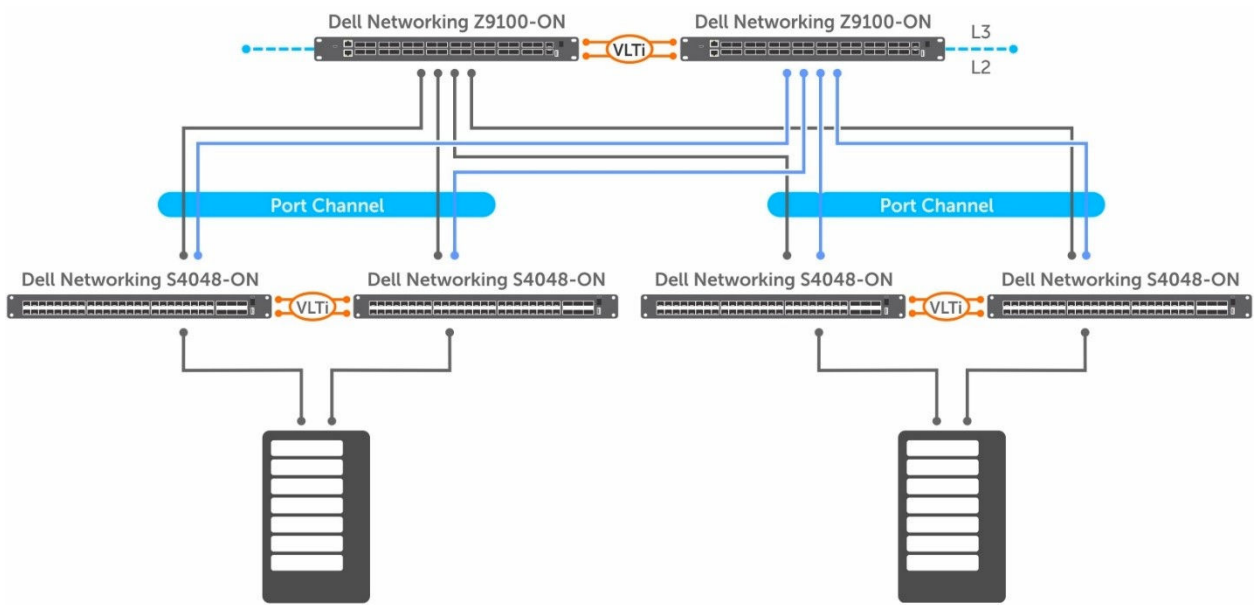


**Figure 27. Storage expansion for increasing the overall infrastructure storage capacity (medium-sized)**

This scale out does not need any additional switches 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 no further investment is needed.

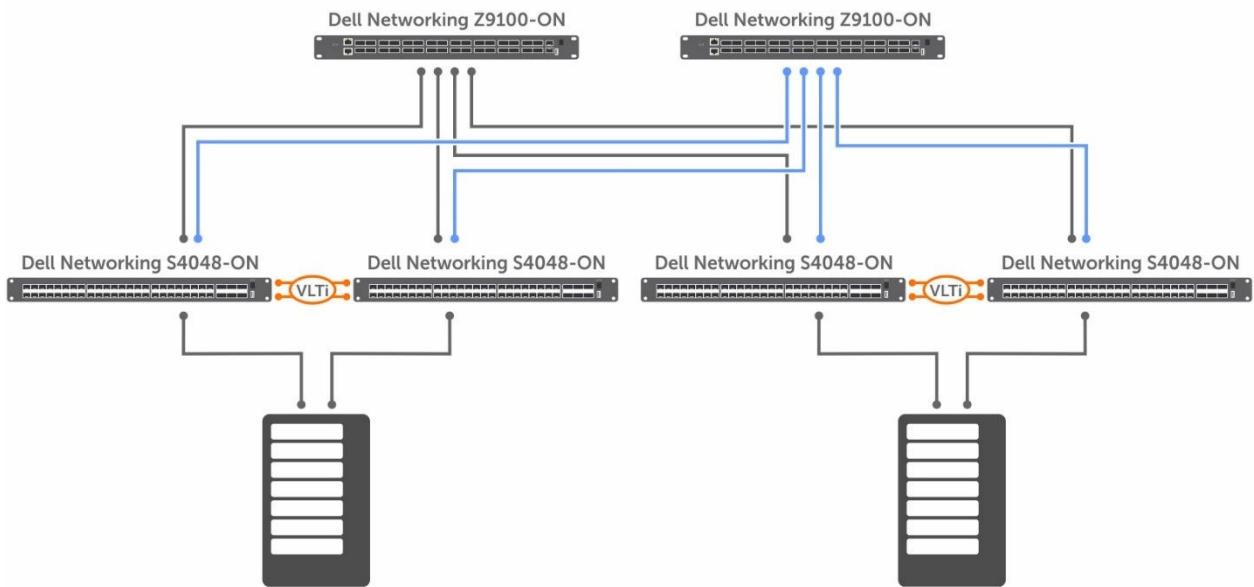
## Scaling network infrastructure

The solution can be scaled by adding multiple Customer Defined Scale Unit (CDSU) in the customer data center. The Dell Networking Z9100 switch can be used to create a simple yet scalable network. The Z9100 switches serve as the spine switches in the leaf-spine architecture. The Z9100 is a multiline rate switch supporting 10 gb, 25 gb, 40 gb, 50gb, and 100 GbE connectivity and can aggregate multiple racks with little or no oversubscription. When connecting multiple racks, using the 40 GbE uplinks from the rack, you can build a large fabric that supports multi-terabit clusters. The density of the Z9100 allows flattening the network tiers and creating an equal-cost fabric from any point to any other point in the network.



**Figure 28. Multiple infrastructure pods scaled out using leaf spine architecture**

For large domain layer-2 requirements the Extended Virtual Link Trunking (eVLT) can be used on the Z9100, as shown in the following figure. The VLT pair formed can scale in terms of hundreds of servers inside multiple racks. Each rack has four 40 GbE links to the core network providing enough bandwidth for all the traffic between each rack.



**Figure 29. Scaling with Extended Virtual Link Trunking (eVLT)**

For more information, see [Dell Networking Multi-tenant Data Center Reference Architecture](#).

# Lifecycle solution management

Topics:

- Deploy
- Operate
- Maintain
- Repurpose
- Backup and Recovery with Dell vRanger
- Capacity planning (Foglight)

## Deploy

Manual deployment involves submitting a request for provisioning server, storage, networking, and virtualization layers with various teams and then manually configuring each layer separately. Manual deployment slows down the process of end-to-end infrastructure provisioning and can impact the experience of the user and their organization. With the Dell Active System Manager (ASM), the Dell EMC Validated System for Virtualization provides a unified infrastructure management layer.

Consider a scenario where the administrators have to commission a new ESXi cluster. The administrator logs in to the Dell Active System Manager (ASM) portal and within the **Templates** section, selects the service blueprints to provision a new ESXi cluster, as shown:



**Figure 30. Dell Active System Manager (ASM)**

ASM automates the end-to-end provisioning across storage, networking, compute, and virtualization and creates an ESXi cluster in a matter of minutes. In validated testing, ASM has on-boarded, configured, and deployed new cloud infrastructure in less than three minutes of hands-on administrator time. The difference was 95 percent faster with 77 percent fewer steps than with manual processes.

ASM is distributed as a virtual appliance running a collection of interacting web applications that communicate through the REST interfaces. An extra web application, ASM-UI, also residing on the appliance, provides full browser-based GUI access to ASM. This ASM-UI

interacts with ASM through the same REST interfaces described in this guide. It is possible to effectively integrate ASM with any custom UI or access it through programs for better automation experiences. For more information, please see the [Active System Manager Release 8.2 API Reference Guide](#).

The ASM end-to-end automation does more than save time for infrastructure provisioning. By providing a single configuration routine (template) that is based on best practices, it reduces possible errors and improves overall service quality. ASM also provides a single abstraction layer for compute and storage, reducing complexity and the learning curve for administrators to enable infrastructure automation.

## Operate

### End-to-end view from virtual to physical

Monitoring is a critical component when multiple organizations are using resources from a common pool and running a diverse set of applications. Most monitoring platforms stop at the virtualization layer with little or no visibility into the underlying physical infrastructure. Dell Hybrid Cloud Platform solves this problem by providing a single end-to-end view for monitoring physical, virtual, and cloud infrastructure.

The following image shows how administrators can monitor current utilizations and the health state of related objects within Foglight. When using VMware Explorer in Foglight and an individual virtual machine is selected, Foglight captures the standard performance metrics of CPU and memory utilization, network, storage throughput, IOPS, response time, and capacity utilization.

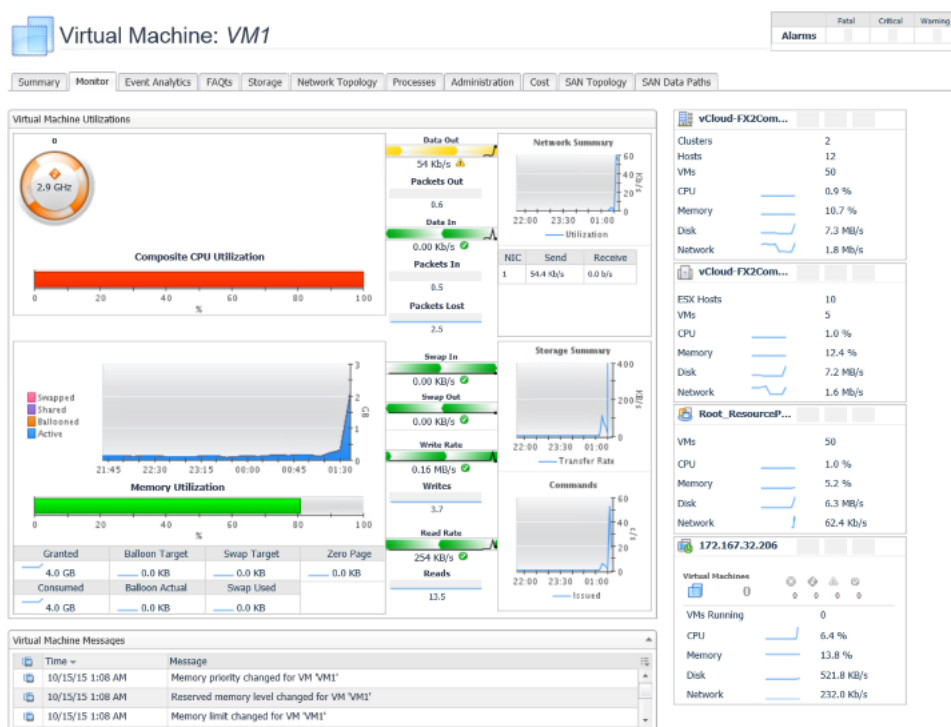
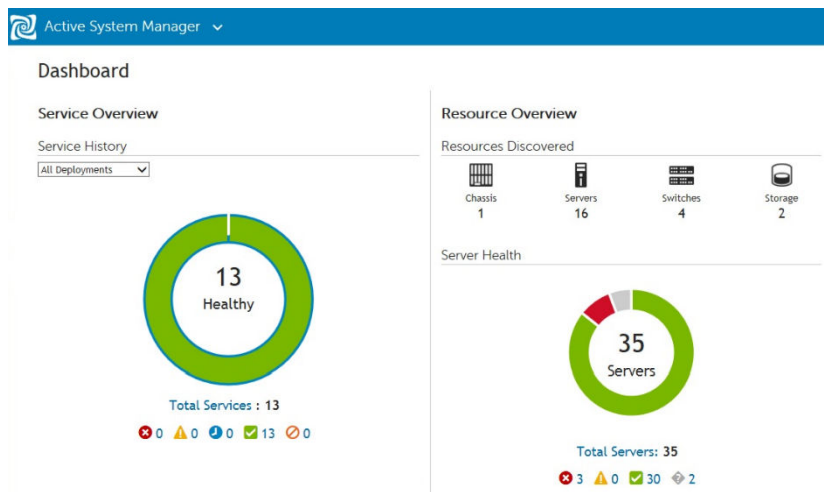


Figure 31. VMware Explorer in Foglight

# Infrastructure remediation

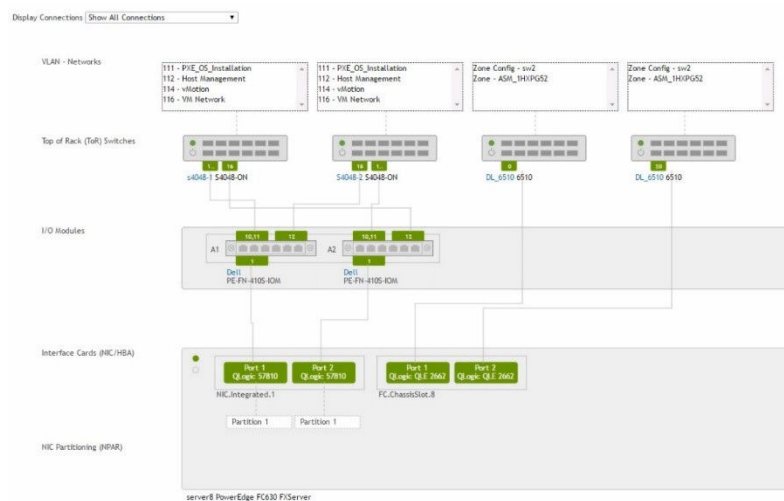
The Dell Active System Manager (ASM) provides a dashboard view of all the services running in the environment. The **Service Overview** section displays a graphical representation of the services based on the state, total number of services deployed, and state icons that represent the service state. The number next to each icon indicates how many services are in a particular state. **Resource Overview** indicates the numbers of chassis, servers, switches, and storage that have been discovered. The default view provides health status, but can be switched to utilization (servers) or capacity (storage).



**Figure 32. Dell Active System Manager (ASM)**

To further assist remediation, ASM provides a Port View to visualize the network and Fibre Channel (FC) connectivity for a particular server. The following information is displayed on the **Port View** page:

- Topology information for all networks and VLANs deployed in a service.
- Network connections between the devices.
- Health of the resources.
- Connection Details, which provide detailed information on the network devices.
- Zone and zone configuration in the VLAN-Networks list for the FC connectivity.



**Figure 33. Dell Active System Manager (ASM) Port View**



# End-to-end view from virtual to physical

Monitoring is a critical component when multiple organizations are using resources from a common pool and running a diverse set of applications. Most monitoring platforms stop at the virtualization layer with little or no visibility into the underlying physical infrastructure. Dell Hybrid Cloud Platform solves this problem by providing a single end-to-end view for monitoring physical, virtual, and cloud infrastructure.

As shown in the following screenshot, administrators can monitor current utilizations and the health state of related objects within Foglight. When a user clicks VMware Explorer in Foglight and selects an individual virtual machine, Foglight captures not only the standard performance metrics of CPU and memory utilization, but also network, storage throughput, IOPS, response time, and capacity utilization.

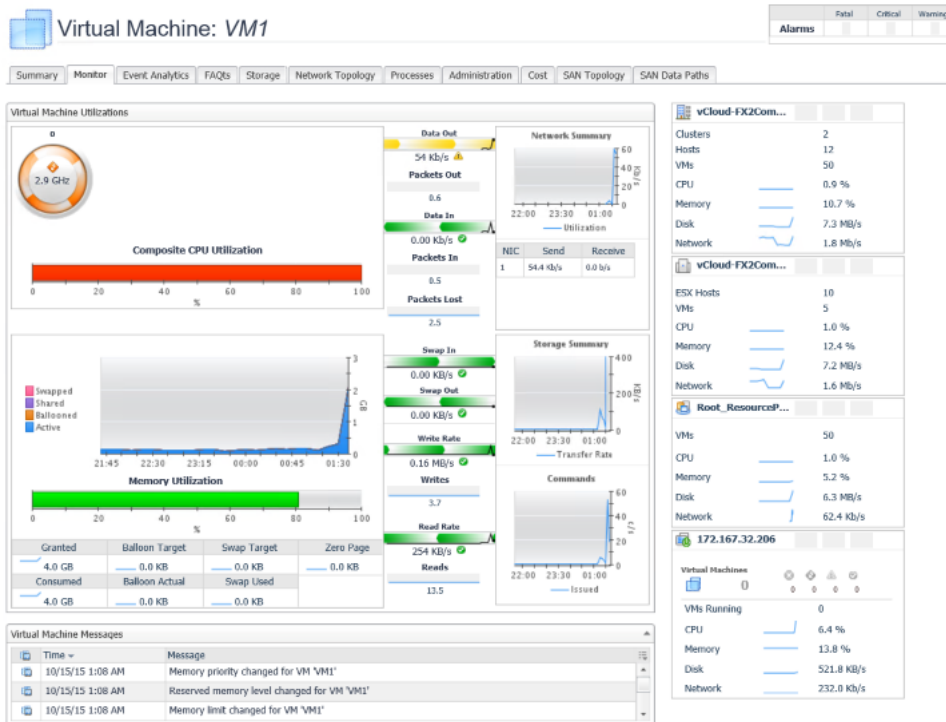
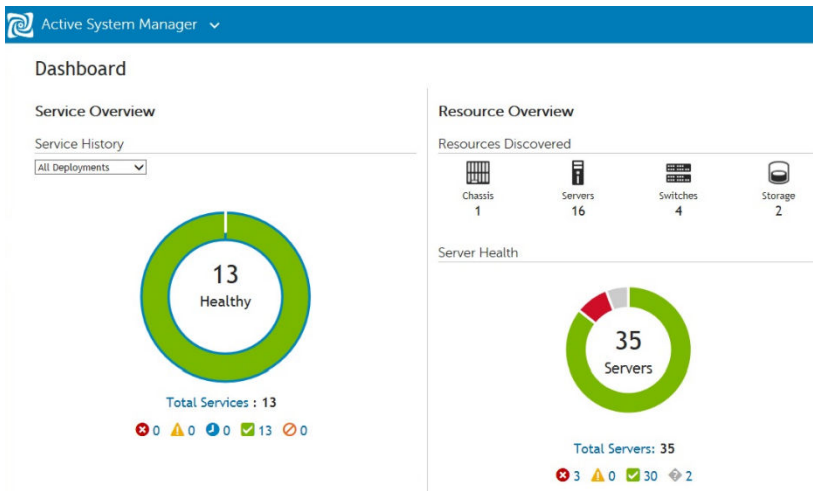


Figure 34. VMware Explorer in Foglight

## Infrastructure remediation

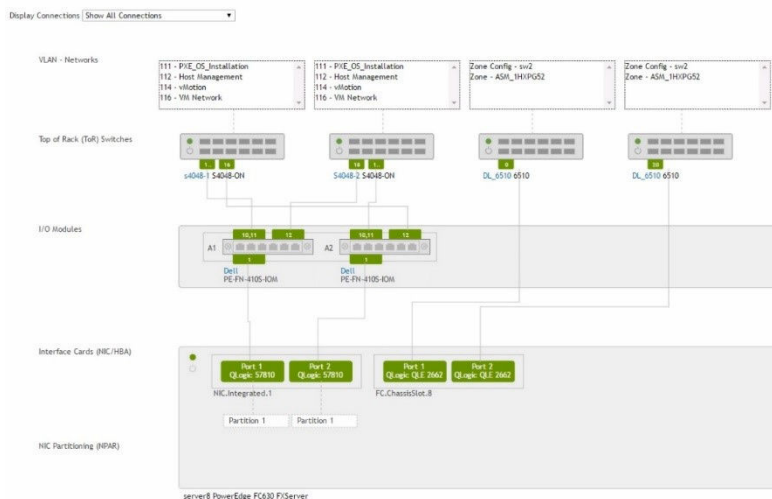
The Dell Active System Manager (ASM) provides a Dashboard view of all the services running in the environment. The Service Overview section displays a graphical representation of the services based on the state, total number of services deployed, and state icons that represent the service state. The number next to each icon indicates how many services are in a particular state. Resource Overview indicates the numbers of chassis, servers, switches, and storage that have been discovered. The default view provides health status, but can be switched to utilization (servers) or capacity (storage).



**Figure 35. Dell Active System Manager (ASM)**

To further assist remediation, ASM provides a Port View to visualize the network and Fibre Channel connectivity for a particular server. The following information is displayed on the Port View page:

- Topology information for all networks and VLANs deployed in a service.
- Network connections between the devices.
- Health of the resources.
- Connection Details, which provide detailed information on the network devices.
- Zone and zone configuration in the VLAN-Networks list for the FC connectivity.



**Figure 36. Dell Active System Manager (ASM) Port View**

## Maintain

Hidden infrastructure configuration and workload provisioning errors can impact service delivery. Even if errors are minimized during initial deployment, ongoing updates, and modifications can create unexpected hidden errors. To maintain system performance, optimization, and compliance, it is important to keep firmware updated.

With the Dell Active System Manager (ASM), administrators can verify that the firmware for the underlying hardware platform, including servers, network switches, and storage controllers, is compliant and can be easily updated when needed. ASM ensures that the

infrastructure is in a balanced state and can help prevent any disruption to service delivery. ASM also provides a streamlined experience in managing resources and associated firmware with centralized management and control.

The following figures show how cloud administrators can log in to ASM and select or schedule the firmware update, which causes minimal disruption to the VMs running in the environment.

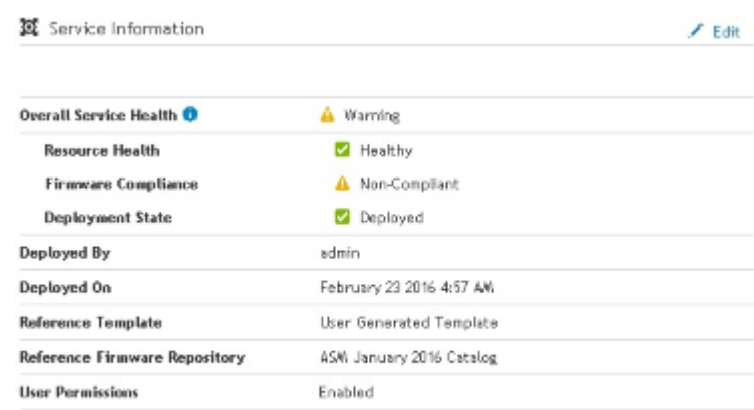


Figure 37. Dell Active System Manager (ASM) Service Information



Figure 38. Dell Active System Manager (ASM) Server Firmware Compliance Report

# Repurpose

In the Dell Active System Manager (ASM), a Server Pool is a set of servers grouped for specific use-cases such as business units or workload purposes. An administrator can also specify a set of users who can access these server pools. Server Pools, which are defined with the ASM templates, are used in the deployment of and expansion of services. Once a service has been terminated, the server resources are returned to their designated server pools.

**Resources**

Server Pool Name	Created By	Created Date	# Servers	# Users
PROD	admin	April 21 2019 12:02 PM	4	2
DEV			2	0
TEST			4	2
Development Pool			5	1
PROD			2	2

7 resources found, displaying 1 - 5. Selected items: 0

**M620's**

Hosts	State	IP Address	Asset Service Tag	Manufacturer	Model	CPUs	Memory	MBs	Location
10.204.144.22	Deployed	10.204.144.22	46182011	Dell	PowerEdge M620	2	32	0	CAC-H2000M_Server1
10.204.144.23	Deployed	10.204.144.23	46182011	Dell	PowerEdge M620	2	32	0	CAC-H2000M_Server2
10.204.144.24	Deployed	10.204.144.24	46182011	Dell	PowerEdge M620	2	32	0	CAC-H2000M_Server3
10.204.144.25	Deployed	10.204.144.25	46182011	Dell	PowerEdge M620	2	32	0	CAC-H2000M_Server4
10.204.144.40	Available	10.204.144.40	46223011	Dell	PowerEdge M620	2	128	0	CAC-H2000M_Server5
10.204.144.41	Error	10.204.144.41	70223011	Dell	PowerEdge M620	2	128	0	CAC-H2000M_Server6

4 resources found, displaying 1 - 4

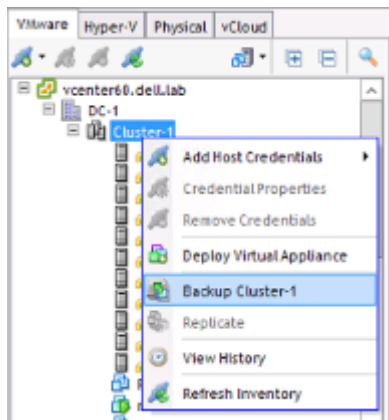
**Figure 39. Dell Active System Manager (ASM) Resources**

Server Pools can be used to segment an environment to ensure that key services have the required resources, while lower priority services are over committed.

## Backup and Recovery with Dell vRanger

With Dell vRanger, you can protect your entire core infrastructure pod compute density (from 100 to 1000 VMs) from a single pane of glass. To make sure all your important workloads are protected, vRanger and the DR appliance can support multiple simultaneous backup write streams to meet demanding backup workloads that require high data ingest performance backup write streams.

The following examples show how vRanger allows you to back up from any level of virtual inventory with complete control over what is backed up and how often.



**Figure 40. Dell vRanger backing up a complete cluster**

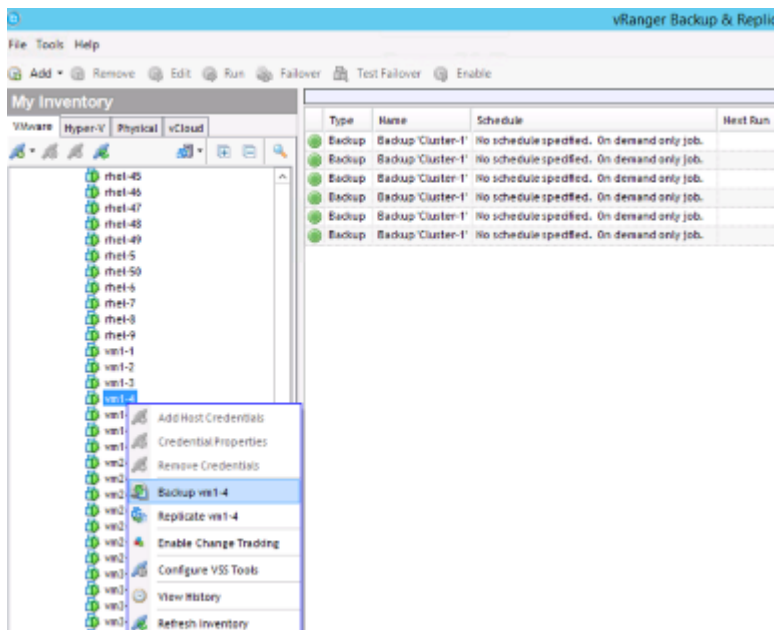


Figure 41. Dell vRanger backing up a specific VM

vRanger also integrates with vSphere Change Block Tracking (CBT), which can be combined with the proprietary Active Block Mapping (ABM) feature to minimize the amount of data transferred during backup and reduce overall backup time. Only active blocks (not idle free space or white-space blocks) that have changed since the last backup is copied to the backup repository when CBT and ABM are enabled. Further backup performance improvement can be achieved by the reduction of data movement through vRanger integration with the DR Appliance Rapid Data Access (RDA) library which performs source side deduplication. ABM, CBT and RDA collectively ensure that only active blocks which have changed and are unique (the block is not stored in the DR appliance yet) get copied during backup. The following screen shot shows how ABM can be enabled for a backup job.

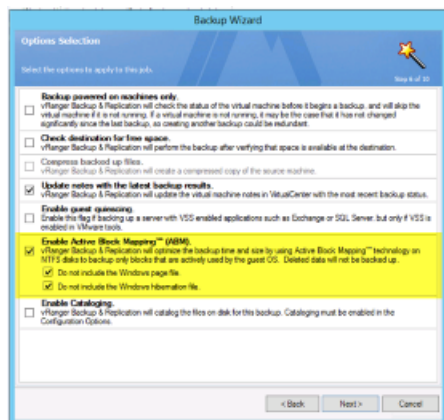


Figure 42. Dell vRanger Backup Wizard

## Capacity planning (Foglight)

In a self-service environment, with many users across different LOBs, monitoring usage trends and growth rates is important. Lack of visibility for proper planning may lead to cases where the platform does not have enough resources to manage business requirements. Dell Hybrid Cloud Platform's single monitoring pane, Foglight, solves this problem. It ensures that usage trends and growth rates are tracked to alert the administrator of current and future capacity problems.

Cloud administrators can track growth rates at the cluster level. The figure below, the cluster is underallocated for memory and has 15 months until the storage is full, whereas the CPUs are sufficient for two or more years at the current growth rate. These insights, based on past and current infrastructure usage, ensure better capacity planning and help optimize infrastructure costs, alerting you where to spend now and where you can wait.

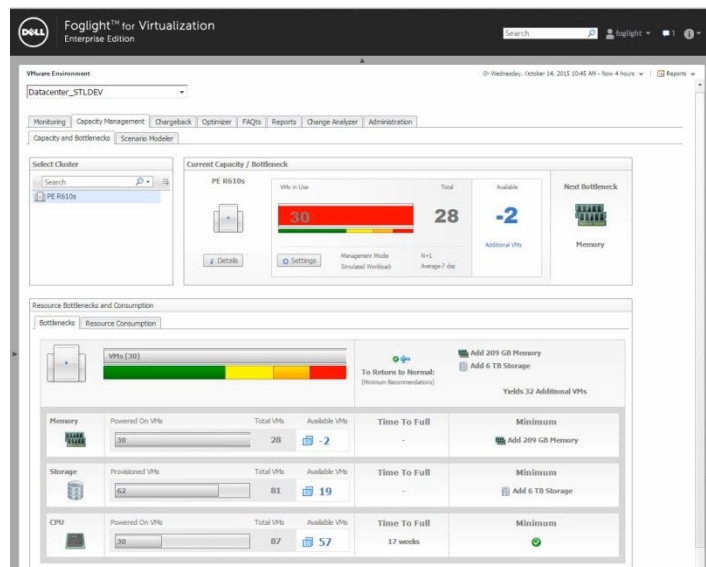


Figure 43. Foglight for Virtualization

# Testing and validation

Dell systems go through rigorous validation by Dell engineers to ensure that the system works as designed. Customers can be assured that the system is validated for various scenarios and if deployed and configured as outlined in this document, they have a high-quality virtualization infrastructure. Dell validation of the system covers three primary goals:

- 1 Follow the design principles described in the [Design Principles](#) section.
- 2 Ensure interoperability between various hardware and software components.
- 3 Ensure that system withstands sustained stress, synthetic and workload specific, within various subsystems.

For a solution infrastructure, verify that the validation of performance using synthetic and load testing tools exercises load on each subsystem in the infrastructure and not individual component silos. In this section, examples are provided to elaborate on the testing and validation methodology.

Topics:

- [Design principle validation — high availability](#)
- [Design principle validation — load balanced](#)

## Design principle validation — high availability

The system is put through various scenarios to ensure that the design principles of the architecture are achieved. One example of the design principle is high availability. Tests are performed to ensure that every aspect of the solution has redundancy including power supplies, networking ports, redundant network switches, and storage fabric. The following figure illustrates network switch redundancy. The figure shows a vCenter performance chart showing two network adapters on a server. When a switch fails, the network adapter connected to the failed switch (vmnic2) ceases to transmit. However, network traffic still continues to be transmitted in the redundant NIC (vmnic3). Upon failback of path resumes the traffic flow over all available paths.




Figure 44. LAN switch failover and failback

The following figure illustrates a similar failure on the redundant SAN fabric. Adapter vmhba32 ceases to transmit when a SAN switch fails. Adapter vmhba33 provides the necessary redundancy until the connection is re-established.




Figure 45. SAN Switch Failover and Failback

## Design principle validation — load balanced

The systems are designed so that all components are optimally utilized and that network load balancing on multiple paths is included. The solution capacity verification tests show that active paths exist within the infrastructure and are load balanced. The following figures show that the paths are actively used during solution capacity verification. The amount of load across each path may not necessarily be equal and depends on the load balancing algorithm chosen to implement within the host network configuration for the network port groups.



Figure 46. LAN path validation during Solution Capacity tests






Figure 47. SAN path Validation during Solution Capacity tests