# Dell Hybrid Cloud Platform Reference Architecture for SAP and SAP HANA workloads

A Dell Reference Architecture for SAP and SAP HANA workload built on Dell PowerEdge FX2 and PowerEdge R930 with Dell Active System Manager, VMware vSphere 6, and VMware vRealize

Dell Global SAP Center of Excellence
May 2016

A Dell Reference Architecture

# Revisions

| Date | Description |
|---|---|
| May 2016 | Initial release |

# Table of contents

# 1 Introduction

This section describes the objective of the Dell Hybrid Cloud Platform Reference Architecture for SAP and SAP HANA and the scope of this document. It includes the target audience for the document, the challenges that the architecture addresses, and some of the significant benefits that are achieved by using the reference architecture, including information on time and effort savings metrics by using a highly automated cloud platform.

## 1.1 Scope

This document outlines the reference architecture for a workload specific cloud infrastructure solution that is suitable for both a private and a hybrid cloud platform. This reference architecture is for SAP and SAP HANA virtualized workloads based on the Dell PowerEdge FX2 converged chassis and PowerEdge R930 modular servers, ultra-low latency LAN and SAN switches, and the Dell SC series high performance enterprise storage system.

The solution outlined in the reference architecture is:

- Designed for a VMware vSphere 6 infrastructure.
- Optimized with Active System Manager; Dell's unified management and open automation framework.
- The architecture is integrated with VMware vRealize Orchestrator for flexibility, vRealize Automation for self-service and governance capability, and vRealize Operations to automate IT Operations management.
- Highly reliable, with integrated monitoring across the physical and virtual infrastructure.
- Comes built-in with data protection, replication, and recovery.
- Utilizing Active System Manager's integration to both vCenter and vRealize, to create a fully featured and highly optimized private and hybrid cloud platform with end-to-end automation capabilities that span across the physical, the virtual, and the cloud.
- Scalable and a highly available infrastructure that is designed, according to the best practices and recommendations, for virtualizing enterprise applications and providing them to business users, in a cloud model.
- Providing guidance on the choice of components and configurations, with the rationale for certain design decisions, and how they deliver measurable value to customers.

**Note:** Deployment and implementation steps for the proposed solution architecture are outside the scope of this reference architecture but may be covered in associated documents.

## 1.2 Audience

The audience for this reference architecture includes, but is not limited to, sales engineers, technologists, architects, field consultants, partner engineering team members, customers, and anyone else interested in deploying an optimized and validated solution stack by using the best of Dell and VMware technology.

## 1.3 Customer challenges

With ever increasing business demands and growing volumes of data, most customers are under unprecedented pressure to improve efficiency and lower costs. However, the current operational model of delivering IT services, which involves procuring technology from technology providers can prove not only to be time consuming but also problematic for a number of reasons.

Customers are typically burdened to make design decisions, validate various components, set up and configure components manually, and manage the environment in an ongoing fashion by engaging multiple vendors for support. All of these elements across the end-to-end infrastructure and application life cycle add up to increased complexity and ongoing costs for customers.

The Dell Hybrid Cloud Platform for SAP and SAP HANA workloads with VMware vRealize is designed to provide detailed guidance and measurable benefits to customers and help solve the challenges mentioned above.

The architecture addresses the following key requirements that we have heard from customers:

- **High flexibility and less complexity:** Flexibility in configuration, scaling, and implementation choices based on specific desired outcomes.
- **Reduction in Capital Expenditure (CAPEX):** Reduce CAPEX by starting at the right scale, growing in pre-defined units, and paying as you go.
- **Lower operational costs:** Lower your overall Total Cost of Ownership by optimizing the infrastructure resources and reducing management complexity.
- **Optimized environment:** Deploy and configure an optimized environment, based on best practices, to deliver a complete and seamless experience to your end users.
- **Meet business SLA's under all conditions:** Design, deploy, and manage a solution that handles failures without causing disruption.

## 1.4 Benefits of the reference architecture solution

By offering well-planned design choices and providing the rationale for choosing the right components, the Dell reference architecture takes the "guess work" out of solution design.

Dell Solution Engineering teams have designed and thoroughly validated the reference architecture solution. The reference architecture approach provides:

- Design principles that are central to each targeted solution.
- Architectural design based on best practices.
- Automated deployment and lifecycle management of the solution including pre-built templates, automated firmware updates, and scaling of physical resources.
- Flexibility to address specific customer needs by further building on or modifying the base reference architecture, if and when required.

### 1.4.1 Tangible, measurable, and profound benefits

Using the wizard-driven automation and pre-built templates of Active System Manager and Dell's unified management and automation solution that anchor the reference architecture, you can:

- Onboard, configure, and deploy new cloud infrastructure in under 3 minutes of hands-on administrator time, and **95% faster and with 77% fewer steps** than with manual processes.
- Deploy a new IT service in **just 6 steps and under 30 seconds** of admin time using a pre-built automation template.
- Provision bare metal hardware and deploy a new ESXi cluster **80% faster and in 71% fewer steps** than one of the leading competitors[1].

These automation capabilities deliver speed, reliability, responsiveness, and consistency to provisioning, deployment, scaling, and de-provisioning of resources in the cloud environment.

### 1.4.2 State of the art financing solutions

The challenges of deploying a cloud platform are not only technical.  They also include paying for the capital and operational expenditures, often in the face of uncertain financial environments. Fortunately, Dell Blueprint Reference Architectures are ideally suited for the scale-ready payment systems that are available from Dell Financial Services that allow customers to control how and when they pay – based on their forecasted usage, their deployment schedule, or their actual usage.

The Dell scale ready payment solutions for enterprise include the following highly flexible options:

- **Pay as You Grow** enables customers to install all of their required technology now and pay for their technology, over time, based on their forecasted usage.
- **Provision and Pay** allows customers to grow their technology solutions, over time, by matching their payments to their deployment schedule.
- **Scale on Demand** enables customers to install their required technology now and pay for it over time, based on their actual measured usage.

These flexible and innovative financing solutions allow Dell to share the risk, convert CAPEX to OPEX, and allow the customer to align their technology acquisitions to their cash flows.

---

[1] Based on independent testing performed at a third-party testing laboratory.  To read more, please visit ASM on Dell Tech Center.

## 1.5 Why Dell for SAP and SAP HANA cloud reference architecture?

### Powerful technology

Single server configurations for SAP HANA, ranging in size from 128GB to 3TB RAM, are based on the Dell PowerEdge R930 server platform to provide a consistent experience and a solid base for future expansion, without forcing "rip and replace" as system needs evolve. Dell's PowerEdge R930 with Intel E7 processors is certified for SAP HANA single server and scale-out solutions, and includes everything needed to support your SAP HANA solution.

### Virtualization for development and production

Dell single server solutions are also certified by SAP to run SAP HANA virtualized. A virtualized HANA implementation allows for a more economical application development and production environment.

### Enterprise class availability

Highly redundant hardware configuration utilizing both fibre channel SAN technology and a highly available multi-node design, the Dell scale out solutions for SAP HANA are engineered to provide resiliency and easy expandability from 4TB up to 32TB and beyond.

### High performance

Automated tiering of data with Dell Compellent storage software provides the quickest access to the data sets most needed for analysis.

### Disaster recovery

Dell solutions for SAP HANA are all tested and able to provide remote disaster recovery through SAP HANA system replication or storage based replication, leveraging the unique features found in Dell Compellent SAN storage.

### Modular growth

The Dell scale out solution is designed to grow from 4 TB up to 32 TB in modular increments, without disruption to the existing system, preserving your investment as your system changes.

### SAP applications in-memory

Application deployment is now supported on Dell solutions for SAP HANA enabling production business applications to enjoy the speed and performance of in-memory computing, without the need for multiple compute environments.

# 2 Solution overview

This document describes the reference architecture for a workload specific SAP and SAP HANA cloud infrastructure solution, which is suitable for both a private or hybrid cloud platform.

The reference architecture is based on:

- Dell PowerEdge FX2 converged system
- Dell PowerEdge R930 HANA servers
- Latest Dell LAN, SAN, and storage technology
- VMware vSphere and vCenter Server 6.0 with components of the vRealize suite
- Blue Medora Management Plugins for SAP and SAP HANA
- Dell Active System Manager, as an overall unified management and automation framework

The reference architecture provides a scalable and highly-available SAP and SAP HANA infrastructure, with built-in monitoring and data protection, which is designed according to the best practices and recommendations for virtualizing enterprise applications and providing them to business users in a cloud or Infrastructure-as-a-Service (IaaS) model. A high level architecture diagram of the key components of the solution is shown below.

**Figure 1 Components of the SAP and SAP HANA Cloud Platform Reference Architecture**

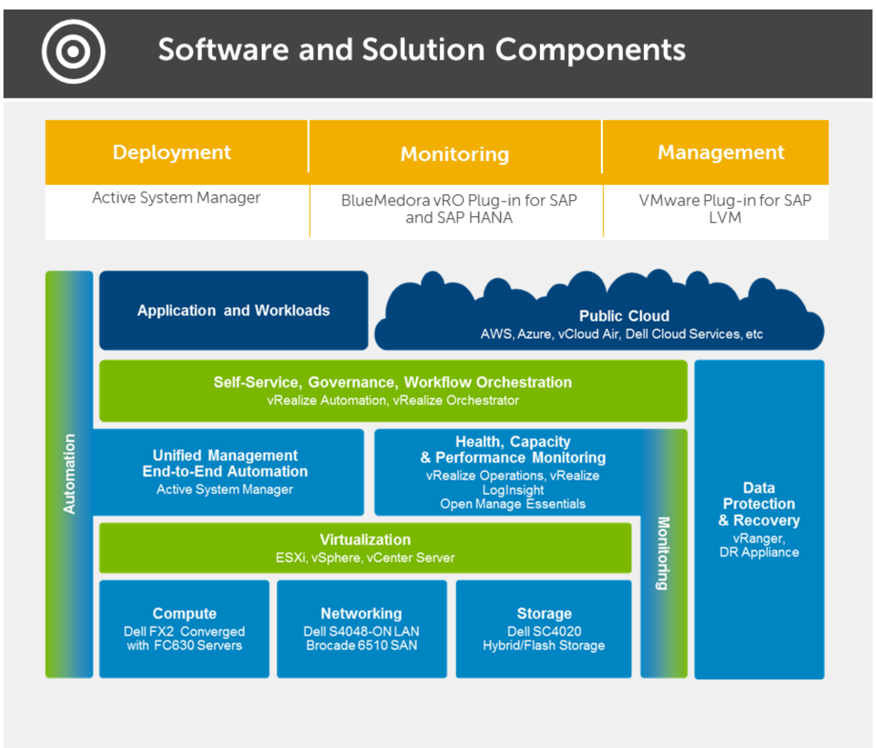


For the three main operations tasks in a SAP and SAP HANA specific cloud infrastructure, we built and tested:

- **Deployment** – Dell ASM for SAP HANA on PowerEdge R930 servers and vRealize Suite components for workload specific rollouts
- **Monitoring** – vRealize Operations with BlueMedora plugins for SAP and SAP HANA
- **Management** – SAP NetWeaver Landscape Virtualization Management with the VMWARE plug-in

This SAP and SAP HANA platform reference architecture, consistent with Dell's virtualization reference architecture for VMware is:

- Built on a hardware platform consisting of the Dell PowerEdge FX2 converged chassis and modular Dell PowerEdge R930 HANA servers; ultra-low latency LAN and SAN switches with the Dell S4048-ON Open Networking switches and Brocade 6510 SAN switches; and two Dell SC4020 high performance enterprise storage systems for general SAP workloads and SAP HANA workloads.
- Equipped with a virtualization layer that is based on VMware vSphere 6.0 and ESXi hypervisor, with VMware vCenter Server as the centralized virtualization management platform.
- An end to end solution with the health, capacity, and performance monitoring from the SAP and SAP HANA workloads to the physical and virtual infrastructure, which is provided by VMware vRealize Operations in conjunction with BlueMedora Plugin for SAP and SAP HANA.
- Integrated with Dell Open Manage Essentials (OME) that offers the hardware monitoring and call home features for hardware fault resolution.
- Built with data protection, replication, and recovery services that comes with the deep VMware integration offered by Dell vRanger, and a Dell Data Protection appliance that provides local, remote, and cloud-based backup with inline deduplication.
- Built with additional vRanger possibilities of direct backup and restore methodology to the Dell Data Protection appliance, for SAP HANA.

The VMware vRealize Suite is a purpose-built cloud management solution for heterogeneous data centers and hybrid cloud environments. The VMware vRealize suite built into this solution offers:

- vRealize Automation (vRA) is included to provide a self-service catalog, request and approvals framework, and policy and governance capabilities.
- vRealize Automation is designed for the Software Defined Data Center (SDDC) and has the control, governance, and multi-tenancy needed for hybrid and multi-cloud deployments.
- vRealize Orchestrator provides a flexible orchestration framework that can be used for automating technical tasks or integrating with business processes such as helpdesk and service ticketing. vRO is also the means by which Dell Active System Manager integrates into the vRealize Suite.

Dell Active System Manager (ASM) provides an overall unified management console and end-to-end automation framework that offers, among others, the following broad capabilities:

- Hardware discovery
- Onboarding
- Initial configuration and inventory management to physical resource pooling
- Bare metal provisioning, and firmware compliance to virtual infrastructure provisioning
- Deployment of hypervisors and configuration of hosts and clusters too virtual machine

- Application and workload provisioning

Together with its tight VMware integration, ASM provides a truly complete end-to-end management and automation experience that:

- Creates a fully featured, deeply automated, and highly optimized private and hybrid cloud platform, with end-to-end automation capabilities from the physical to the virtual to the cloud.
- Offers an end-to-end automation capabilities, that span from the bare metal physical infrastructure, inclusive of servers, networking, and storage; to the virtual infrastructure, including hypervisor deployment and cluster configuration; to operating system and VM deployment; all the way to application-level provisioning, for a complete Infrastructure as a Service (IaaS) cloud platform with unparalleled automation capabilities.

## 2.1 Public and hybrid cloud capabilities

Hybrid clouds provide an optimum balance of private and public cloud capabilities, including secure on-demand access to shared but private resources, together with the flexibility to deploy or move workloads offsite to public clouds, to meet specific needs.  This model provides both security and scalability, balances capital and operational expenses, and delivers the benefits of providing for peak demand capacity and disaster protection.

This reference architecture describes a hardware and software stack that is designed as a highly optimized private cloud platform, and which, by providing a flexible management and automation framework from Dell and VMware, allows the incorporation and consumption of public cloud resources, and provides the foundation for a hybrid cloud as well.

With vRealize Automation, users and administrators can provision virtual machines into private resource polls managed by Active System Manager, or into public clouds, including any vCloud Air cloud service provider.  VMware vCloud Air is a public cloud platform built on the foundation of vSphere that is compatible with your on-premise data center, which includes infrastructure, disaster recovery, and various application service offerings, allowing you to deploy your workloads into the public cloud.

Further, with the optional vCloud Connector, administrators can perform virtual machine migration from the on-premise private cloud to vCloud Air, and set up catalog synchronization between private and public templates.

## 2.2 Hardware stack overview

The compute platform for SAP workloads for this reference architecture is the Dell PowerEdge FX, a 2U converged chassis that combines the density and efficiencies of blades with the simplicity and cost benefits of rack-based systems. It includes in-chassis switching that reduces Top of Rack port usage and cabling complexity.

For SAP HANA workloads, the compute platform is Dell's world record performance leading PowerEdge R930, which is capable to carry HANA systems from 128 GB to 3 TB and beyond.

The figure below shows a typical layout of the hardware components of the reference architecture.

**Figure 2 Rack Elevation View of Reference Architecture Hardware Components**

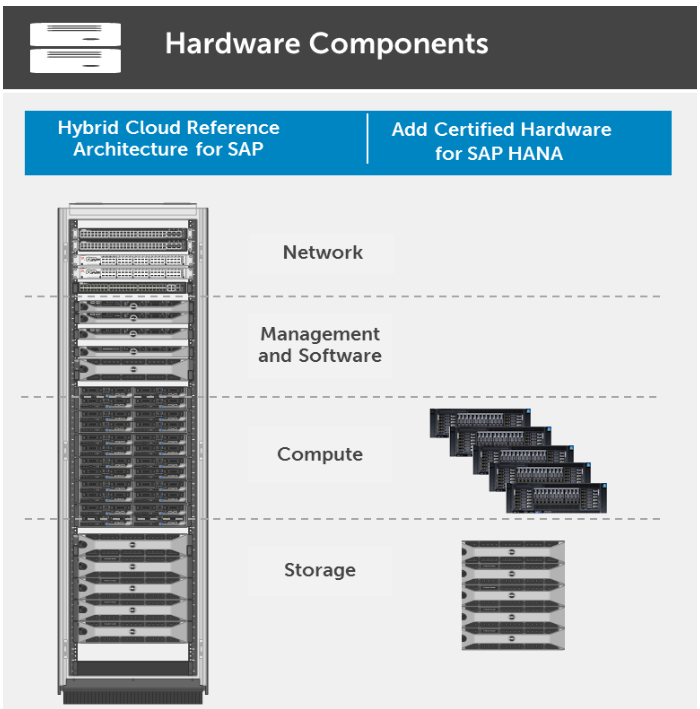


For the network fabric, the ultra-low latency S4048-ON switches from Dell, combined with the in-chassis switching of the FX2 architecture, create a very efficient IP network layer. A Brocade Fibre Channel SAN fabric delivers a high-performance dedicated storage network with quality-of-service and support for advanced SAN management features.

The Dell Storage SC4020 all-in-one array is a highly virtualized storage system that is flash-optimized for affordable all-flash or hybrid SSD/HDD deployment. With high IOPs and low latency it delivers great value in scalability, efficiency, and high performance for private cloud applications.

These hardware components are described in greater detail in the Hardware Solution Components section and in the Sizing Guidelines section.

## 2.3 Software solution components

The tables below list the primary software components for the reference architecture.

Table 1 Dell software components

| Dell Software | Version | Notes |
|---|---|---|
| Dell Active System Manager | 8.1.1 | Unified management and automation software |
| Dell ASM Plugin for VMware vRealize Orchestrator | 1.0 | ASM plugin for vRealize Orchestrator. Enables vRealize Automation integration. |
| Dell Open Manage Essentials | 2.0.1 | Software to monitors physical infrastructure |
| Dell Enterprise Manager 2015 | 1.10 | Storage management console |
| Dell vRanger Pro | 7.3 | Backup server for virtual infrastructure. |
| | | |
| **Optional Dell Software:** | | |
| Dell OpenManage integration for VMware vCenter | 3.0 | Plugin for vCenter Server |
| Dell Storage vSphere Web Client Plugin | 3.0 | Plugin for vSphere Server |
| Dell Compellent Solution Pack for VMware vCenter Operations Manager | 1.0 | Plugin for vRealize Operations |

Table 2 VMware software components

| VMware Software | Version | Notes |
|---|---|---|
| VMware vCenter Server | 6.0 | vSphere management server |
| VMware vSphere ESXi | 6.0 | Server hypervisor |
| VMware vRealize Automation | 7.x | VMware cloud management and infrastructure, and service catalog |
| VMware vRealize Orchestrator | 6.0.1 | Orchestration engine |
| | | |
| **Optional VMware Software:** | | |
| VMware vCloud Connector | 2.7.1 | Connection software for private-public cloud migration |
| VMware vRealize Operations | 6.0.1 | Operations management. Optional in lieu of Dell Foglight. |
| VMware vRealize Business | 6.1.0 | Used for chargeback and showback. Optional in lieu of Foglight. |
| VMware vRealize Log Insight | 2.5 | Log analytics and management. Optional in lieu of Foglight. |
| VMware vCenter Site Recovery Manager | 6.0 | Data recovery. Optional in lieu of vRanger. |

Table 3 Additional software components

| BlueMedora | Version | Notes |
|---|---|---|
| SAP plugin for vRealize Operation | | SAP monitoring |
| SAP HANA plugin for vRealize Operation | | SAP HANA monitoring |
| | | |
| **SAP** | | |
| LVM – NetWeaver Landscape Virtualization Manager* | 2.1 | SAP system management |

* Not part of delivery. The software must be downloaded through the SAP software portal.

### 2.3.1 Dell primary software components

The following Dell software products are standard components of the reference architecture:

- Dell Active System Manager
- Dell Active System Manager plugin for VMware vRealize Orchestrator
- Dell Open Manage Essentials
- Dell Enterprise Manager
- Dell vRanger Pro
- Dell DR Series Backup Appliance

The following sections focus primarily on automation, monitoring, and cloud platform management software and therefore do not describe the embedded hardware element managers, such as servers and switches.

#### 2.3.1.1 Dell Active System Manager

Active System Manager (ASM) is Dell's 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.

ASM delivers:

- Rapid time to value – By accelerating growth and innovation with fast and responsive IT automation.
- Superior ease of use – By optimizing IT and simplifying enterprise systems management.
- Unrivaled flexibility – By increasing options with an open and extensible architecture.

ASM enables system administrators to deploy new services from pre-built templates:

- In 6 steps
- Within 30 seconds of hands-on time
- And is 80% faster than similar products

**Figure 3 Active System Manager**

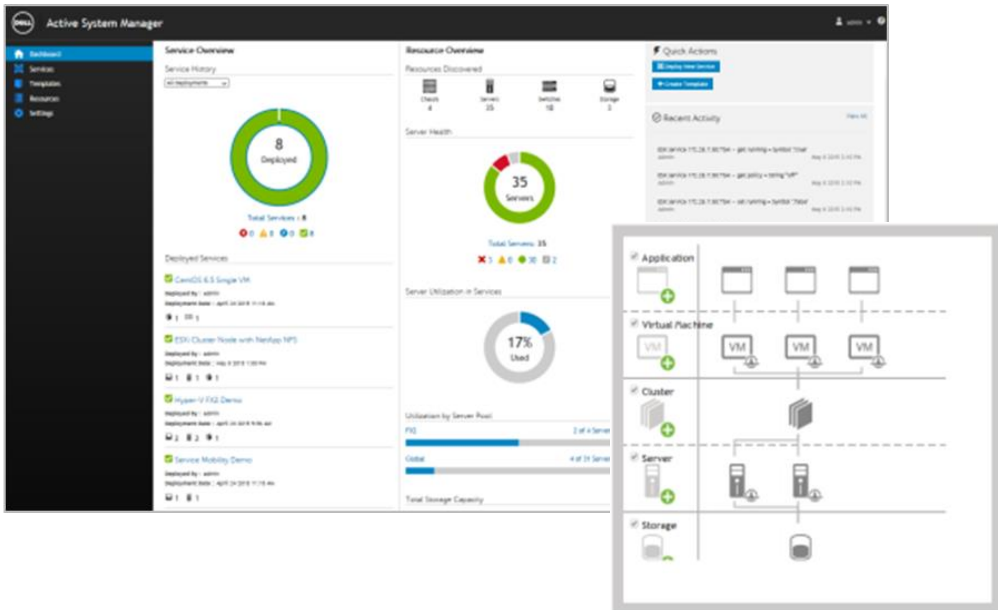


ASM takes a top-down, service-centric approach to IT automation that spans across both physical and virtual infrastructure. ASM encompasses everything from servers, switches, and storage, to hypervisors and clusters, to virtual machines, operating systems, and applications.

ASM provides a unified and comprehensive user experience for the initial deployment and the ongoing lifecycle management of shared or converged infrastructure.

ASM enables IT administrators to accelerate service delivery and improve efficiency in a private cloud platform in a number of ways, including the ability to:

- Get up and running quickly, with rapid discovery and initial configuration, and wizard-driven automation of server, chassis, and IO network component onboarding.
- Efficiently manage infrastructure lifecycle, with comprehensive firmware management, compliance monitoring, and updates.
- Define and provision unique or complex physical, virtual, and hybrid services, as reusable templates
- Easily deploy, manage, and scale IT services, with automated provisioning. Scale physical and virtual resources up and down, and decommission services to recover resources.
- Define and manage user access levels and permissions, with role based access control.
- Respond quickly with dynamic, on-demand provisioning of resources from shared pools.

In short, ASM provides a strong foundation for a private or hybrid cloud environment.

### 2.3.1.2 Dell Active System Manager plugin for VMware vRealize Orchestrator

ASM supports a plugin for vRealize Orchestrator (vRO) and, through this vRO integration, ASM can publish ASM templates to the vRealize Automation (vRA) service catalog.

The ASM plugin for vRealize Orchestrator:

- Allows ASM automation templates to be called from vRO orchestration workflows.
- Enables automation methods supported by ASM, including provisioning, configuration, capacity scaling, and deprovisioning, to be embedded into vRO workflows.
- Extends the automation and orchestration capabilities of vRO to bring ASM automation into business process workflows, to improve responsiveness and simplify complex IT processes.
- Enables integration with vRealize Automation, so that automation templates for infrastructure and workloads can be published to the vRA service catalog, for a self-service IT experience.

The figure below shows the Active System Manager plugin for vRealize Orchestrator.

**Figure 4 Dell Active System Manager plugin for VMware vRealize Orchestrator**

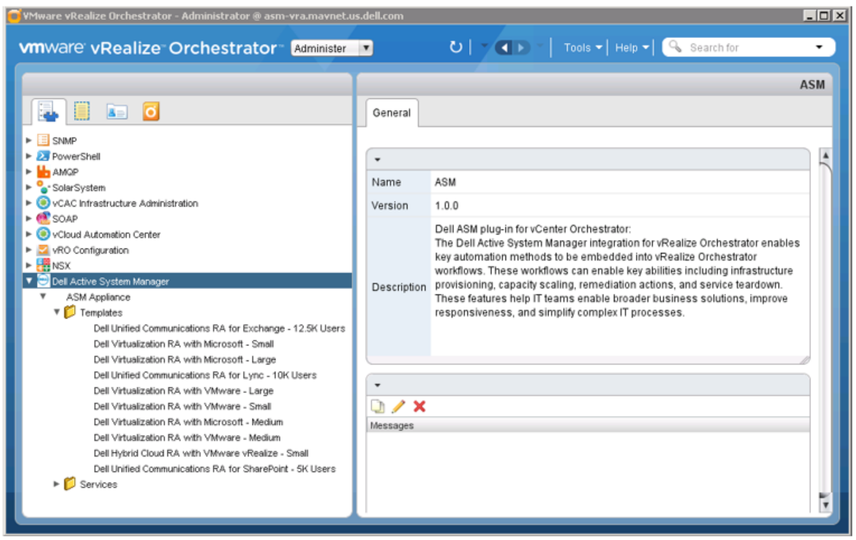


### 2.3.1.3 Dell Open Manage Essentials

Dell Open Manage Essentials (OME) is a one-to-many management console for monitoring Dell data center equipment including server, storage, and networking devices. It also provides lifecycle management of Dell PowerEdge servers. It is the easiest and simplest way for system administrators to maximize Dell systems uptime and health.

Some of the key features of OME include:

- Monitoring health status, events, and inventory for Dell PowerEdge servers, storage, and networking devices.
- Providing hardware level control and management for Dell PowerEdge server, blade chassis, and internal storage arrays.

- Enabling deeper management and control of Dell Blade chassis, Dell storage, Dell networking devices through context-sensitive link and launch of their respective element management tools.

Open Manage Essentials integrates with the following solutions:

- **Dell SupportAssist**
  Enables automatic hardware failure notifications sent securely, to Dell technical support, for intelligent analysis and diagnosis that can lead to optimized availability and reduced manual intervention.  This is available as part of Dell ProSupport, at no extra cost.
- **Dell Repository Manager**
  Allows precise control of firmware revisions.
- **Dell OpenManage integration for VMware vCenter**
  Enables agent-free discovery, monitoring, deployment, configuration and updates, within VMware vCenter. Enables comprehensive management in physical and virtualized environments, and also enriched health monitoring of Dell servers and storage.

**Figure 5 Dell Open Manage Essentials**

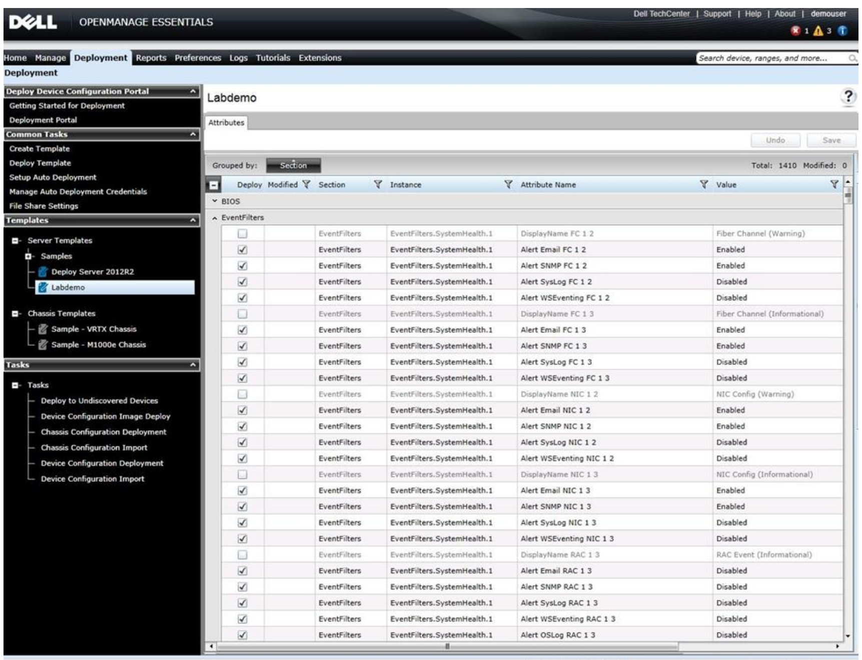


In this reference architecture, VMware vCenter serves as the primary monitoring interface for the overall virtual and physical infrastructure, with hardware health and status being linked to it through the OME integration.  In addition, OME is available for detailed hardware monitoring and fault analysis, as needed.

With the Dell Open Manage Essentials integration into VMware vCenter, you can:

- Monitor, manage, configure, and deploy Dell PowerEdge servers; all within your vCenter console.

- Experience the unified and comprehensive management of physical and virtual infrastructure with Dell OpenManage integration for VMware vCenter and the advanced management features of Dell PowerEdge servers.
- Enables easy scaling and a streamlined process for applying updates to multiple Dell hosts in a single workflow, by providing unique cluster-level hardware views directly within vCenter.

#### 2.3.1.4 Dell Enterprise Manager

Dell Enterprise Manager 2015 is a comprehensive management console for managing Dell Storage Center (SC) series storage products. Enterprise Manager extends the capability of individual Storage Center web interfaces to deliver centralized management of all local and remote Storage Center environments. It also enables tight integration with leading technology providers including VMware.

Enterprise manager provides detailed capacity and performance monitoring and storage-based chargeback features that automatically calculate storage costs, based on the actual space consumed by applications.

**Figure 6 Dell Enterprise Manager**


#### 2.3.1.5 Dell vRanger Pro

Dell vRanger software provides high-speed backup and recovery of virtual infrastructures built using VMware and physical Windows Server environments. It also provides high-speed replication for VMware. The software protects entire virtual environments within minutes, by detecting and backing up new VMs automatically, and delivering safe and scalable data protection to even the largest VMware and Hyper-V environments.

With vRanger, you can locate and restore individual files in seconds, even if they are buried in virtual and physical backups, from a single, intuitive interface. Dell vRanger, when combined with Dell Storage DR appliances, provide complete backup, replication, de-duplication, and recovery capabilities for the virtual infrastructure environment.

**Figure 7 Dell vRanger Pro**

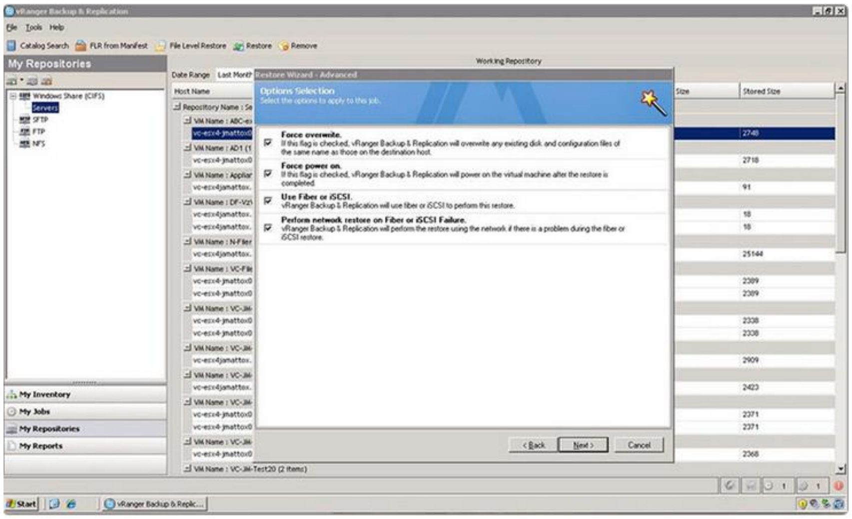


vRanger provides high-speed, resource-efficient backup, replication and recovery of virtual machine images, and supports backup and recovery of Windows physical servers, files and folders. vRanger also delivers maximum storage savings when paired with the Dell DR deduplication appliance.

Key features and benefits of vRanger Pro include:

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

#### 2.3.1.6 Dell DR Series Backup Appliance

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

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

**Figure 8 Dell DR Series Backup Appliance**

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

### 2.3.2 VMware Primary software components

The following VMware software products are standard components of the reference architecture:

- VMware vCenter Server
- VMware vRealize Automation
- VMware vRealize Orchestrator
- VMware vRealize Operations
- Blue Medora vRealize Operations Plugins
- Blue Medora vRealize Operations Plugin for SAP
- Blue Medora vRealize Operations Plugin for SAP HANA

#### 2.3.2.1 VMware vCenter Server

VMware vCenter Server provides a centralized platform for managing all your VMware vSphere environments so that you can automate and deliver a virtual infrastructure with confidence. vCenter Server is used to manage the ESXi hosts and logical vSphere clusters.

vCenter Server provides centralized control and visibility through:

- **vSphere web client**
  Manages the essential functions of vSphere from any browser anywhere in the world.
- **vCenter single sign-on**
  Allows users to log in once and access all instances of vCenter Server, without the need for further authentication.
- **Custom roles and permissions**
  Restricts access to the entire inventory of virtual machines, resource pools, and servers by assigning users to custom roles. Users with appropriate privileges can create these custom roles, such as night-shift operator or backup administrator.

- **Inventory search**
  Explores the entire vCenter inventor, including virtual machines, hosts, data stores and networks, from anywhere within vCenter.

**Figure 9 vRealize Automation with ASM Template Integration**

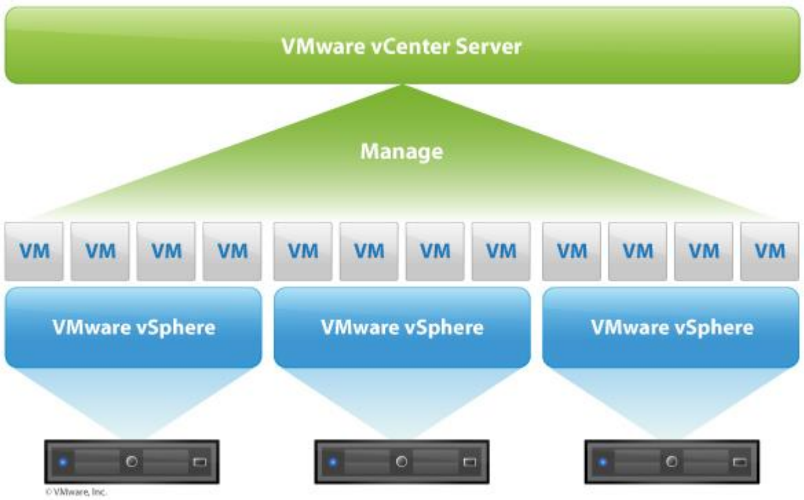


vCenter Server integrates with vRealize Orchestrator to automate more than 800 tasks using out-of-the-box workflows or by assembling workflows with the easy drag-and-drop interface. Additionally, Dell Active System Manager integrates with vSphere virtualization platforms and ESXi hosts, through a centralized vCenter Server instance, to integrate management and automation of the physical and virtual environments.

### 2.3.2.2 VMware vRealize Automation

VMware vRealize Automation (vRA) helps IT departments with delivery and ongoing management of infrastructure, applications, and custom services. vRealize Automation can be used to build both enterprise and service provider cloud solutions.

Supporting implementations for private clouds, public clouds, and hybrid clouds, vRA can be used in conjunction with existing or new business processes and tools, to leverage a very flexible, highly automated, and intuitive self-service experience for the end user.

Some of the key features of vRA are listed below:

- Delivers infrastructure, application, and custom services through a unified IT service catalog.
- Meets specific business needs at the right service level with personalized, policy-based governance.
- Delivers governance and control to enable hybrid cloud deployments.

**Figure 10 vRealize Automation with ASM Template Integration**

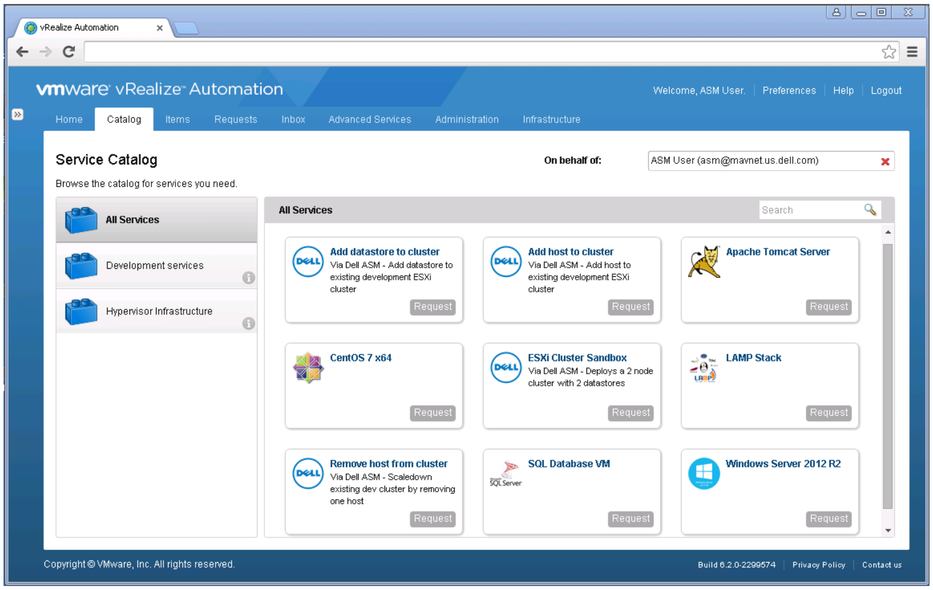


VMware vRA makes it easy to maintain control over provisioned physical and virtual workloads using the centralized provisioning structure, governance, and infrastructure management capabilities. For example, the lifecycle management of a virtual machine workload is controlled using vRA, through built-in processes for request, approval provisioning, management, reclamation, and decommissioning.

### 2.3.2.3 VMware vRealize Orchestrator

VMware vRealize Orchestrator (vRO) simplifies the automation of complex IT tasks and integrates with other VMware vRealize Suite components to adapt and extend service delivery and operational management, effectively working with existing infrastructure, tools and processes.

Some of the key features of vRO are:

- Allows administrators to develop complex automation tasks in a graphical workflow manner.
- Quickly access and launch workflows from the VMware vSphere client using various components of the VMware vRealize Suite or other triggering mechanisms.
- Create custom workflows, for many cloud deployments while interacting with governance, provisioning, documentation, monitoring, service ticketing, and IP address management (IPAM) systems.

vRO can leverage the features provided by the Active System Manager API to create end-user or administrator-specific services. The Active System Manager plugin to vRO allows a business analyst or IT architect to embed ASM templates and automation methods into orchestration workflows, and therefore integrate seamlessly into an existing business processes, such as requests, approvals, help desk, and self-service. This integration saves time and manual effort, and improves responsiveness and consistency.

In addition, the ASM integration with vRO allows an IT administrator to build templates for common or standardized IT services and infrastructure requests in ASM and publish them to the vRA service catalog, for end user self-service fulfillment.

**Figure 11 vRealize Orchestrator with ASM Plugin**

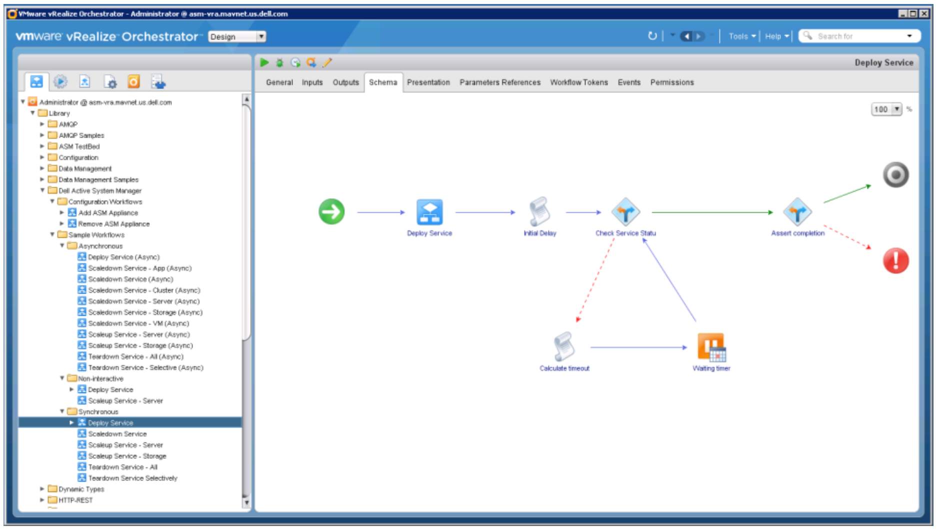

## 2.3.2.4 VMware vRealize Operations

VMware vRealize Operations delivers intelligent operations management from applications to infrastructure across physical, virtual, and cloud environments.

Some of the key features of VMware vRealize Operations are:

- **Intelligent IT Operations**
  Maximize capacity utilization and enable optimum performance and availability of applications and infrastructures, across vSphere, Hyper-V, Amazon, and physical hardware. Proactively identify and solve emerging issues, with predictive analytics and smart alerts. Manage capacity and resolve resource contention with intelligent workload placement.
- **Unified IT Management**
  Gain comprehensive visibility in one place and across applications, storage, and network devices with an open and extensible platform, which is supported by third-party management packs for solutions from SAP, and other hardware and software providers.
- **Policy Based Automation**
  Increase efficiency by streamlining key IT processes with out-of-the-box and customizable policies, guided remediation and automated enforcement of standards. Optimize performance, capacity, and compliance while retaining full control.

**Figure 12 vRealize Operations**

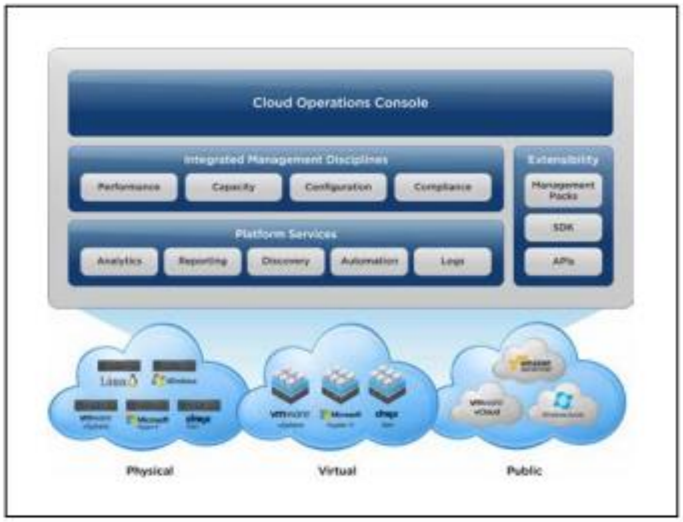


### 2.3.2.5 Blue Medora vRealize Operations Plugins

Blue Medora's products extend VMware vRealize Operations (vROps) by providing a better understanding of application, database, compute, storage, and network performance, demand and capacity.

### 2.3.2.6 Blue Medora vRealize Operations Plugin for SAP

Blue Medora vRealize Operations plugin for SAP extends the value of vRealize Operations with SAP data collected via SAP. VMware customers, who have implemented SAP platform for monitoring and managing their application infrastructures are capturing and storing a wealth of data that can be unlocked by the powerful predictive analytic capabilities that VMware Realize Operations (vROps) brings to the enterprise. With the vROps Management Pack for SAP, customers can gain comprehensive visibility and insight into the performance, capacity and health of their SAP workloads that run on VMware.

**Figure 13 Blue Medora vRealize Operations Plugin for SAP**

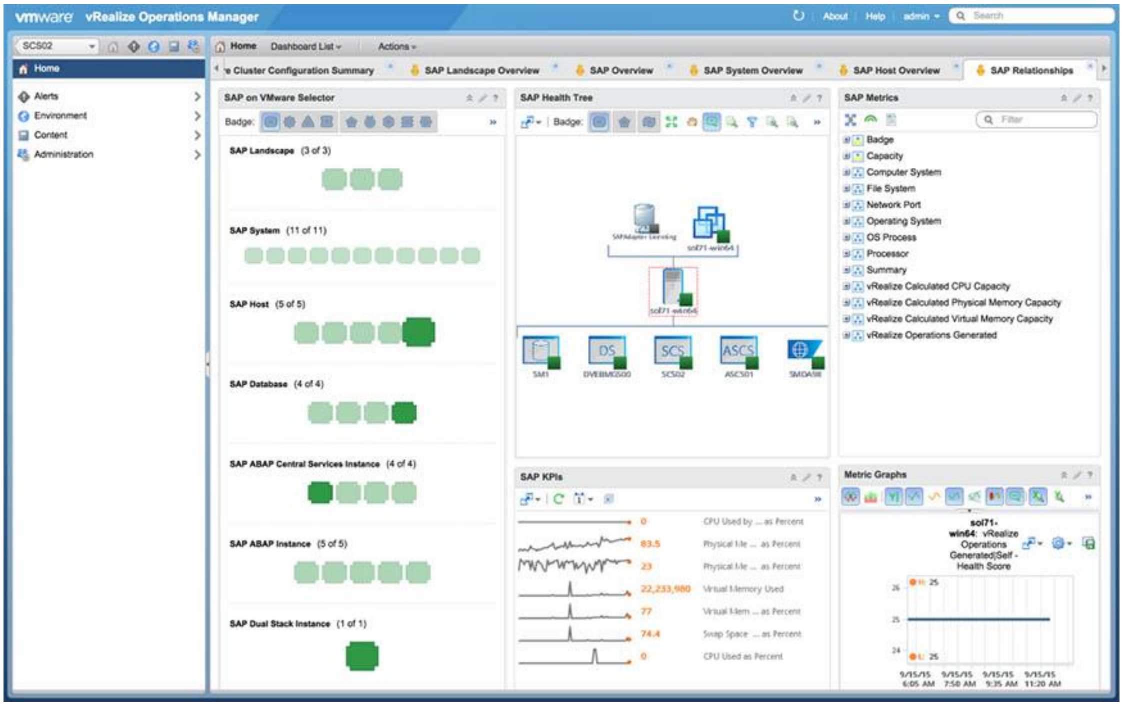


## 2.3.2.7 Blue Medora vRealize Operations Plugin for SAP HANA

SAP HANA is deployable as an on-premise appliance, in the cloud, or can be virtualized on VMware vSphere. It is a revolutionary in-memory database platform that is best suited for performing real-time analytics, and developing and deploying real-time applications.

The Blue Medora vROps Management Pack for SAP HANA connects to any SAP HANA instance running on a dedicated appliance, such as Dell PowerEdge R930, in the cloud, or can be virtualized with VMware. It has the ability to consume a large number of SAP HANA-focused health, performance, and availability metrics.

**Figure 14 Blue Medora vRealize Operations Plugin for SAP HANA**

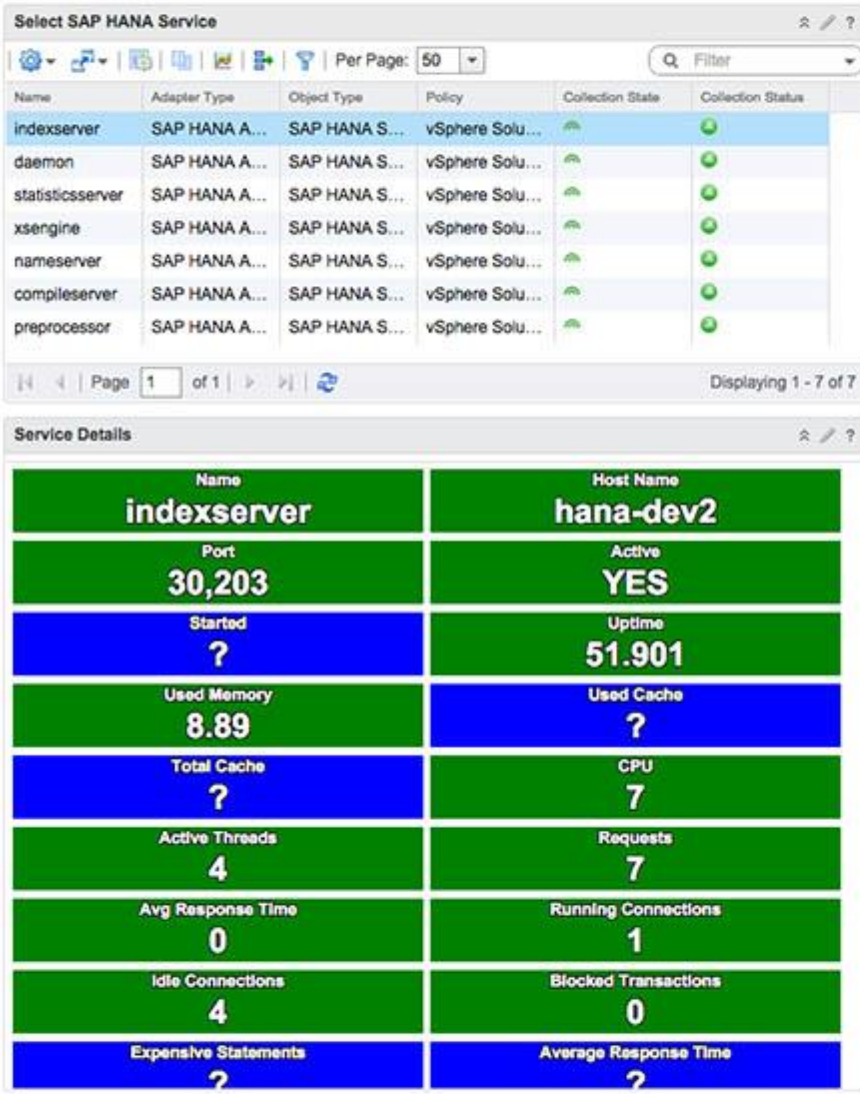

## 2.4 Hardware solution components

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

**Figure 15 Hardware solution components - rack elevation**

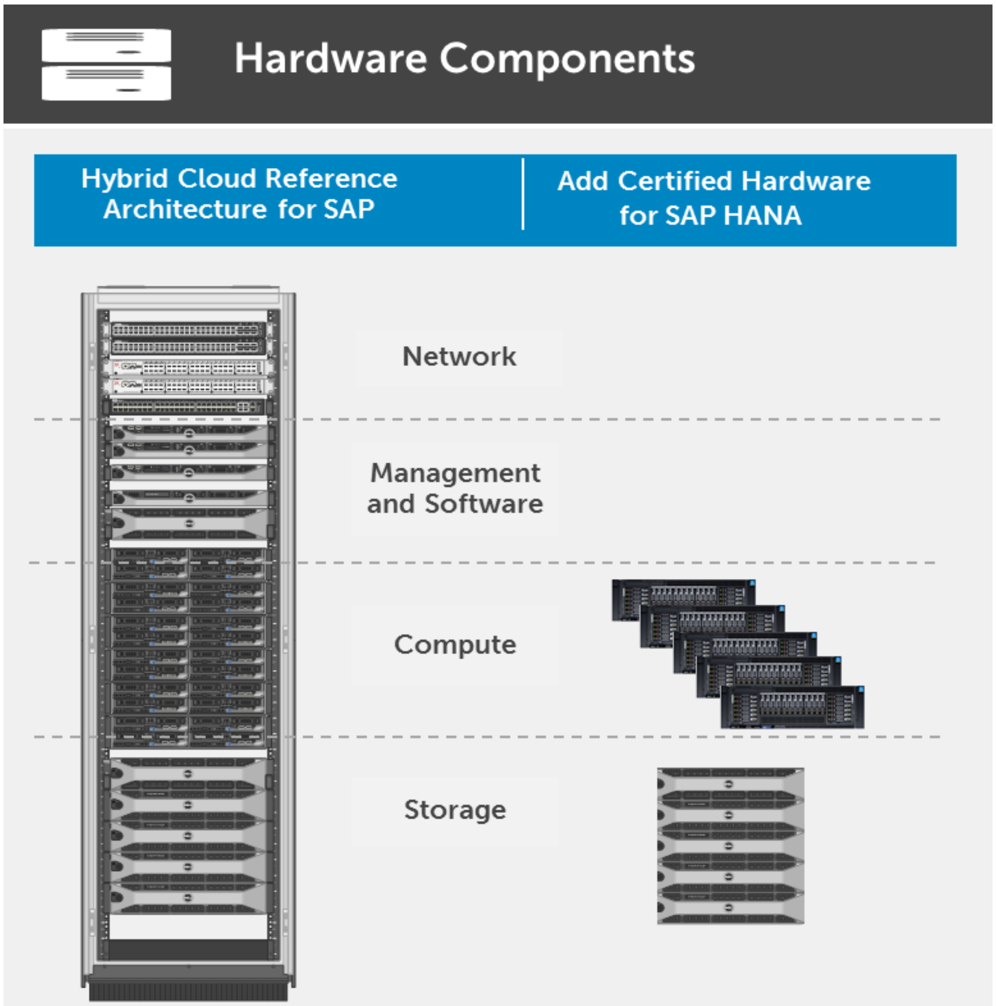


The Dell PowerEdge FX2 architecture, Dell PowerEdge R930 HANA servers, Dell Storage SC4020 and Dell Networking S4048 offer a wide-range of features and configuration options. This enables greater flexibility in choosing the components and configuration that is relevant for the solution architecture being designed. Table 1 provides the solution specification for the components used in this architecture.

The following sections provide a technical and product overview of the components included in this reference architecture:

- PowerEdge FX2
- PowerEdge F630
- Dell PowerEdge FN410S
- Dell PowerEdge R930 HANA Server
- Dell Networking S4048-ON Switch
- Dell Storage SC4020 Array
- Brocade 6510 SAN Switch

The table below outlines the specifications of components used in the reference architecture.

**Table 4 Specification of components used in this reference architecture**

| Component | Details | |
|---|---|---|
| Virtualization Infrastructure | Dell PowerEdge FX2<br>PowerEdge FN410S; 10GbE SFP+ IOA<br>Qlogic QLE 2652 DP PCIe Add-on FC HBA for storage connectivity | |
| Virtualization Hosts | **Dell PowerEdge FC630 Servers - SAP** | |
| | Processor | 2 x Intel Xeon E5-2699v3 Family or 2 x Intel Xeon E5-2643 v3 |
| | Memory | 256 GB; 16 x 16GB DDR4 DIMMs |
| | OS Volume | 2 x 16 GB SD cards in internal SD module |
| | Network | Qlogic 57840 quad-port NDC |
| | OS | VMware ESXi 6 |
| | **Dell PowerEdge R930 Servers – SAP HANA** | |
| | Processor | 4 x Intel Xeon E7-8890v3Family |
| | Memory | 2 TB,3 TB ; 64/96 x 32 GB DDR4 DIMMs |
| | OS Volume | 2 x 16 GB SD cards in internal SD module |
| | Network | 2 x Intel 520 DA Dual Port |
| | OS | VMware ESXi 6 |
| Storage Arrays | **SAP**<br>Dell Storage SC4020 with 6x 400 GB (WI) SSD, 6 x 1.6 TB (RI) SSD and 12 x 10K SAS drives<br>**SAP HANA** | |

| Component | Details |
|---|---|
| | Dell Storage SC4020 with 12x 400 GB (WI) SSD and 36 x 10K 1.8 TB SAS drives (#needs to be calculated) |
| Infrastructure Deployment | Dell Active System Manager (ASM) 8 |

The choice of components in Table 1 is supported by the design principles described earlier and then the configuration choices are supported by the solution architecture section that follows.

For more information about the components used in this solution, see Appendix B.

### 2.4.1 PowerEdge FX2

The PowerEdge FX2 is a 2U hybrid rack-based computing platform that combines the density and efficiencies of blades with the simplicity and cost benefits of rack-based systems. With an innovative modular design that accommodates IT resource building blocks of various sizes — compute, storage, networking and management — the FX2 enables data centers to construct their infrastructures with greater flexibility.

There are two versions of the FX2 chassis:

- **The switched configuration**: PowerEdge FX2s, in this configuration support up to eight low-profile PCI Express® (PCIe) 3.0 expansion slots
- **The unswitched configuration**: In this configuration, FX2 provides a lower-cost alternative that does not offer expansion slots.

The FX2 enclosure also offers I/O modules to several I/O aggregators that can simplify cabling, improve East/West traffic within the server, and enable LAN/ SAN convergence — reducing cost and complexity.

For more information about FX2 and its architecture, see Appendix B.

### 2.4.2 PowerEdge F630

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

For more information about FC630 and its architecture, see Appendix B.

### 2.4.3 Dell PowerEdge FN410S

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

For more information about FN410S, see Appendix B.

### 2.4.4 Dell PowerEdge R930 HANA server

The Dell PowerEdge R930 rack server is specifically designed for the most demanding enterprise applications and features industry-leading internal storage and memory scalability to optimize application performance. Leveraging the latest Intel® Xeon® processor E7 v3 product family with (up to 24 cores per processor), the four-socket R930 flexibly scales to optimize transaction, operations, and significantly reduce latency.

For more information about R930 and its architecture, see Appendix B.

### 2.4.5 Dell Networking S4048-ON switch

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

For more information about S4048-ON, see Appendix B.

### 2.4.6 Dell Storage SC4020 array

Dell Storage SC4020 belongs to the Storage Center (SC) 4000 series arrays based on the SC8000 platform.

The SC4020 storage array offers:

- Multi-protocol support and virtualized multi-tier, multi-RAID-level storage policies.
- Dual redundant controllers
- 24 internal drive slots
- Eight 8 Gb Fibre Channel (FC) or four 10 Gb iSCSI network ports
- One 1 Gb port per controller for out-of-band (OOB) management traffic
- Drive types ranging from Write-Intensive (WI) Solid State Drives (SSDs) to Read-Intensive (RI) SSDs to 15K, 10K, and 7.2K SAS drives in small form factor (2.5-inch).
- Multi-tier data placement that improves application performance.
- Expansion up to 192 drives by adding Dell Storage SC220 or SC200 enclosures.

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

### 2.4.7 Brocade 6510 SAN switch

Brocade 6510 is a 48-port 1U Gen 5 FC switch that is suitable for high-performance data requirements of server virtualization, cloud and enterprise applications. This switch can be configured in 24, 36 or 48 ports and supports 2, 4, 8, 10 or 16 Gbps speeds. This enables enterprises to start small and scale the FC infrastructure based on the growth.

By using features such as Virtual Fabrics, Quality of Service (QoS) and zoning, IT departments can build multi-tenant cloud environments. Management and diagnostics features such as Monitoring and Alerting Policy Suite (MAPS), Dashboards, Flow Vision, Fabric Performance Impact (FPI) monitoring and Credit Loss Recovery help administrators and IT organizations avoid problems before they impact the SAN operations.

# 3 Design principles

While designing this reference architecture, certain specific design tenets have been followed in order to address the customer requirements.

## 3.1 Single platform for any scale

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

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

This reference architecture for SAP and SAP HANA virtualized workloads uniquely adopts single infrastructure architecture regardless of scale, which means no change in core technology or added complexity as the infrastructure grows.

The reference architecture offers customers:

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

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

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

- The Dell PowerEdge FX2s architecture's incremental, modular approach allows for greater flexibility to address workload needs and allows organizations to budget just for the required resources. It delivers both the flexibility of rack deployments, and the density and performance of blade servers.

- For SAP HANA, Dell PowerEdge R930 can scale-up from just 128 GB RAM up to 4 TB, in one system, for running business applications and can scale-up from 128GB RAM to 2TB RAM, for business warehouses.
- SC4020, the storage platform used in the reference architecture, offers a wider range of performance-enhancing features, including true multi-tier flash optimization and allows SC4020 to provide high performance for the small, medium and large configurations referred in this reference architecture. Customers can easily scale their storage environment by adding more enclosures to address their capacity expansion needs. The performance and scalability of SC4020 is clearly showcased in the large configuration referred to in this reference architecture, where a single SC4020 supports 24 compute nodes without any performance degradation.
- Brocade 6510 is chosen in the reference architecture, given its optimal port density (48 ports) that can support the requirement to start with a single FX2 chassis ( 4 nodes – 8 ports) and grow to 6 or more chassis ( 24 nodes-28 ports or more ), if required. However, while this reference architecture has been tested with Brocade 6510, in order to obtain more flexibility, customers can choose 6505 series, in cases where they do not see the requirement to scale out beyond 4 chassis.

## 3.2 Effective utilization

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

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

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

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

## 3.3 Reduction in management complexity

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

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

- **Automation and life cycle management:**
  Dell Active System Manager, which offers a single provisioning interface in the form of template-based automation, helps customers to achieve consistent results in an automated fashion. ASM also helps the customer automate the application of the firmware updates across the stack and manage the drift against a baseline configuration. It also allows customers to add compute and storage resources as needed, to scale and optimize resource utilization, thereby providing measurable value throughout the end–to-end infrastructure life cycle.
- **Deep monitoring:**
  Dell's Open Manage Essentials provides a single tool to monitor server, storage, and networking and help customers identify the issues and prevent any disruption.

## 3.4 Business continuity and quality of service

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

- In each scenario of different environment sizes, the recommended configuration takes into account the performance under degraded state, which could include a node failure, SAN controller failure and so on. In each configuration, there is enough room to account for failures, while ensuring that performance of the workload is not impacted.
- If a node fails at the compute level, infrastructure availability provided by the VMware vSphere cluster built on PowerEdge FX2 ensures that the VMs are migrated to other available physical hosts. This ensures that there is little or no disruption to the workloads.
- The reference architecture also includes the use of redundant hypervisor cards with the Dell PowerEdge FX2 platform and helps guard against a single point of failure.
- The network layer is designed to be redundant with two Dell Networking 10GBE switches to eliminate a single point of failure in the solution. FN IOAs are connected together to create a Virtual Link Trunk (VLT) interconnect. This ensures that the east-west traffic between the compute cluster nodes within the chassis stays in the chassis, and multiple active paths to the top-of-rack

(ToR) can exist. Configuring each physical host with multiple network paths ensures that there is a redundant path available even in case of a network switch failure within the solution infrastructure.

- The chosen network architecture with multiple VLT (mVLT) or multiple Link Aggregation Groups (mLAG) using the VLT mode for the FN IOAs in the RA provides redundancy at all layers of the stack. Virtual and physical network connections are further made resilient by using the native NIC teaming feature of VMware vSphere.
- In addition, data belonging to the virtual infrastructure and virtualized applications should sustain a storage failure. The storage arrays used in this solution architecture provide redundancy at various levels within the array. For example, each of the storage arrays is configured with redundant power supplies and redundant controllers connected to redundant SAN switches. This design ensures that the component level failures within the array do not impact the virtualized infrastructure and virtualized application availability.

# 4 Sizing guidelines

In the following sections the sizing guidelines based on SAPS and SAP HANA needed capacity is documented.

## 4.1 SAP sizing guidelines

For SAP workloads, the most used performance indicator is SAPS (http://global.sap.com/campaigns/benchmark/measuring.epx)

SAP Application Performance Standard (SAPS) is a hardware-independent unit of measurement that describes the performance of a system configuration in the SAP environment. It is derived from the Sales and Distribution (SD) benchmark, where 100 SAPS is defined as 2,000 fully business processed order line items per hour.

In technical terms, this throughput is achieved by processing 6,000 dialog steps (screen changes), 2,000 postings per hour in the SD Benchmark, or 2,400 SAP transactions.

In the SD benchmark, fully business processed means the full business process of an order line item:

- Creating the order
- Creating a delivery note for the order
- Displaying the order, changing the delivery
- Posting a goods issue
- Listing orders
- Creating an invoice.

### 4.1.1 Using SAPS for sizing

If, as a result of using the Quick Sizer sizing tool offered by SAP, a sizing table for a particular solution suggests a configuration of 10,000 SAPS, you can check the SD two-tier benchmark table for a sample configuration. By setting the sort order in the SAPS column, you can view the benchmark results that gives guidelines for configurations that are likely to fulfill your requirements.

### 4.1.2 General purpose SAP workload sizing with SAPA

For general purpose, we recommend the E5-2699 v3 processor as you get the most SAPS for the money. In general purpose environments, SAP NetWeaver is enabled to leverage as many cores and threads as needed.

Table 5 SAPS and Memory for 2 x E5-2699v4, 256 GB Memory, 72 vCPUs, 117130 SAPS with HyperThreading enabled

| # vCPUs | # SAPS | Recommended Memory in GB |
|---|---|---|
| 1 | ~1.150 | 2-4 |
| 2 | ~2.300 | 4-8 |
| 4 | ~4.600 | 8-16 |
| 8 | ~9.200 | 16-32 |
| 16 | ~18.400 | 32-64 |
| 32 | ~36.800 | 64-128 |

### 4.1.3 SAPS single threaded applications (JAVA)

For single threaded use, we recommend the E5-2643v3 CPUs as you get the most SAPS per core.

Table 6 SAPS and Memory for 2 x E5-2643v4, 96 GB Memory, 24 vCPUs, 46500 SAPS with HyperThreading enabled*

| # vCPUs | # SAPS | Recommended Memory in GB |
|---|---|---|
| 1 | ~1.750 | 2-4 |
| 2 | ~3.500 | 4-8 |
| 4 | ~7.000 | 8-16 |
| 8 | ~14.000 | 16-32 |
| 16 | ~28.000 | 32-64 |

*in some specific use cases we recommend to disable Hyper Threading. You can contact your Dell SAP representative for further information.

## 4.2 SAP HANA sizing guideline

(Referenced from the VMware whitepaper: https://www.vmware.com/files/pdf/SAP_HANA_on_vmware_vSphere_best_practices_guide.pdf)

Quick Sizer translates business requirements into technical requirements. It deploys a proven sizing methodology developed in close cooperation with SAP technology partners. By determining your sizing requirements, you can ensure that your hardware purchase will meet your business and performance requirements. A financial benefit will lower cost and reduce total cost of ownership.

SAP's comprehensive sizing methodology helps ensure that all the required hardware components build a balanced system infrastructure. Database and application servers are adequately equipped to fulfill the business requirements. The right combination of processing performance (in SAPS), amount of disk space (in MB), disk I/O, and main memory (in MB) is proposed.

### 4.2.1 Virtual SAP HANA sizing

Sizing SAP HANA for a virtual environment is the same as sizing SAP HANA for a native environment, except for a few differences. These differences are due to the size limitations of virtual machines, virtual CPU to physical CPU mappings, and the slight overhead cost of virtualization. The basic process is to size the SAP HANA database normally (the same sizing that is done for native SAP HANA), and then make any required adjustments, to account for virtualization.

Each of these factors that are used for sizing virtual SAP HANA correctly are described in detail with examples in the sections below. These factors include:

- Maximum virtual machine size
- Virtual CPU to physical CPU mapping
- Memory overhead
- SAP HANA sizing approaches
- SAP-based SAP HANA sizing (Maximum virtual machine performance)
- SAP based SAP HANA sizing (Maximum total system performance)

#### 4.2.1.1 Maximum Virtual Machine size

The maximum virtual machine size in vSphere 5.5 is 64 vCPUs and 1 TB of RAM. This is smaller than the 4-socket PowerEdge R930 servers that are used for SAP HANA hardware appliances. It is not possible to directly map these larger sized SAP HANA hardware appliances to a virtual machine.

#### 4.2.1.2 Virtual CPU to Physical CPU mapping

The current generation of Intel processors that are used for SAP HANA have a feature called hyperthreading. Hyperthreading allows each physical core to have two execution threads. These threads share the resources of the same physical core. As a result, they do not double performance, but instead provide a 10 to 20 percent increase in performance, in most cases.

A vCPU that is created for a virtual machine is scheduled to run on a logical thread. A vCPU maps directly to a logical thread on a physical core. The default behavior for vSphere is to schedule vCPUs on separate physical CPUs, using a single execution thread on each physical core first. Once all of the physical cores have a vCPU scheduled and running, then vSphere can begin to schedule vCPUs to use the second logical execution thread on each physical core.

From a sizing perspective, vCPUs are equivalent to physical cores without hyperthreading enabled, in terms of performance. In order to obtain up to an additional 10 to 20 percent of performance with hyperthreading enabled, the number of vCPUs needs to be doubled. In other words, a physical SAP HANA server system with 64 physical cores, and with hyperthreading enabled, would provide 10 to 20 percent higher performance than a virtual SAP HANA server system (virtual machine) with 64 vCPUs and 64 execution threads, all mapped to physical cores of the example server.

#### 4.2.1.3 Memory overhead

Virtual machines have full access and use of the amount of RAM that is assigned to the virtual machine, when it is created. However, some memory is needed to run the ESXi hypervisor and manage the virtual memory assigned to the virtual machines. Many vSphere customers never notice this memory overhead due to the many memory optimization techniques that have been implemented.

Production SAP HANA virtual machines have full memory reservation. This feature does not allow for memory optimization techniques to be enabled. For this reason, it is necessary to leave a small amount of unassigned RAM on the host, to account for the memory required by the hypervisor and virtual machine memory overhead.

A very conservative estimate for the amount of RAM that needs to be unassigned for the SAP HANA virtual machines overhead is 3 to 4 percent. For example, on a system having 1 TB of RAM, approximately 30 to 40 GB would need to be left unassigned to the virtual machines. A single large virtual machine with 980 GB, or two virtual machines with 490 GB each, could be created and must have enough unassigned memory for the hypervisor and virtual machine memory overhead.

**Note:** The memory overhead is visible to the memory reservations configured for SAP HANA virtual machines. In an environment where resource over-commitment is allowed, this memory overhead does not need to be calculated and configured.

#### 4.2.1.4 SAP HANA sizing approaches

The sizing of a virtualized SAP HANA system follows the same method as the sizing of a native SAP HANA system. In this case, an initial system sizing is completed using the SAP QuickSizer or a sizing is based on an existing system. The results of these efforts is a number that represents the needed performance for the system that is to run the application that was sized. This number, called SAPS, represents a unit of performance. SAP partners use provided benchmarks to test their systems to measure their capability in terms of SAPS. Beside SAPS sizing, a memory sizing also needs to be performed. Once the needed compressed memory size is determined, the memory of a SAP HANA virtual machine can be defined.

Since SAP HANA is an in-memory platform, the memory needs of a SAP HANA system are calculated based on capacity sizing. Typically, this sizing is determined by the customer working in conjunction with SAP or partners. The tools and resources used for this effort are the same for both virtual and native environments. Once these initial sizing numbers are determined, they can be used to define the size of the virtualized SAP HANA virtual machine and the size of the vSphere host.

For more information about sizing, see the "SAP Service Marketplace" at:

https://service.sap.com/sizing

**Note:** Accessing this site requires your standard SAP support portal credentials.

The performance of CPUs continues to improve with each new release. The SAPS rating for a virtual machine depends on the performance of the physical cores on which it is running. For SAP HANA, the Intel® Xeon® Processor Intel® Xeon® Processor E7 v3 Series (Haswell-EX) are currently supported. This means that the SAPS rating for a given virtual machine might be different, depending on the processor on which it runs.

#### 4.2.1.5 Sizing for maximum virtual machine performance

In order to achieve the highest possible performance for a single virtual machine, each vCPU requires its own physical core. This means, a 1:1 ratio of vCPUs to pCPUs. For example, suppose there is one large virtual machine with 64 vCPUs, with each vCPU running on its own physical core. In this case, 64 vCPUs is equal to 64 physical CPUs without hyperthreading.

#### 4.2.1.6 Sizing for maximum total system performance

Newer physical servers have more CPU cores/threads available than the configuration maximum of vCPUs and so it may not be possible to assign all threads of a server to a single virtual machine. In order to use all logical threads in SAP HANA virtual machines, it is necessary to create multiple virtual machines. For an Intel Haswell-EX CPU with four sockets and 18 cores per socket, the system would have a total of 144 logical threads. For example, a virtual machine configuration that uses all of the threads would require four virtual machines, with each having 36 logical threads, and each virtual machine kept assigned, such that it uses its own 18 physical cores only.

### 4.2.2 Virtual SAP HANA business warehouse sizing

This table describe the number of NUMA nodes (Non-uniform memory access) in conjunction with vCPUs and memory for SAP HANA analytic (BW) systems.

Table 7 Dell PowerEdge R930 – E7v3 8890 or 8880 – per CPU 18 cores – 72 vCPUs max– no#

| #Numa Nodes | # vCPUs | # SAPS | Memory in GB |
|---|---|---|---|
| 1 | 5 | ~ | 64 |
| 1 | 10 | ~ | 128 |
| 1 | 18 | ~ | 256 |
| 2 | 24 | ~ | 384 |
| 2 | 32 | ~ | 512 |
| 3 | 48 | ~ | 768 |
| 4 | 64 | ~ | 1024 |

### 4.2.3 Virtual SAP HANA Business Suite sizing

This table describe the number of NUMA nodes in conjunction with vCPUs and memory for SAP HANA business suite systems.

Table 8 Dell PowerEdge R930 – E7v3 8890 or 8880 – per CPU 18 cores – 72 vCPUs max–no#

| #Numa Nodes | # vCPUs | # SAPS | Memory in GB |
|---|---|---|---|
| 1 | 5 | ~ | 64 |
| 1 | 8 | ~ | 128 |
| 1 | 12 | ~ | 256 |
| 1 | 15 | ~ | 384 |
| 1 | 18 | ~ | 512 |
| 2 | 36 | ~ | 1024 |

# 5 Solution architecture

This section describes the overall solution architecture, including details on the compute, storage, and networking subsystem architectures. The figure below illustrates the high-level scalable architecture of the basic building block.

**Figure 16 Scalable Architecture of the Solution**

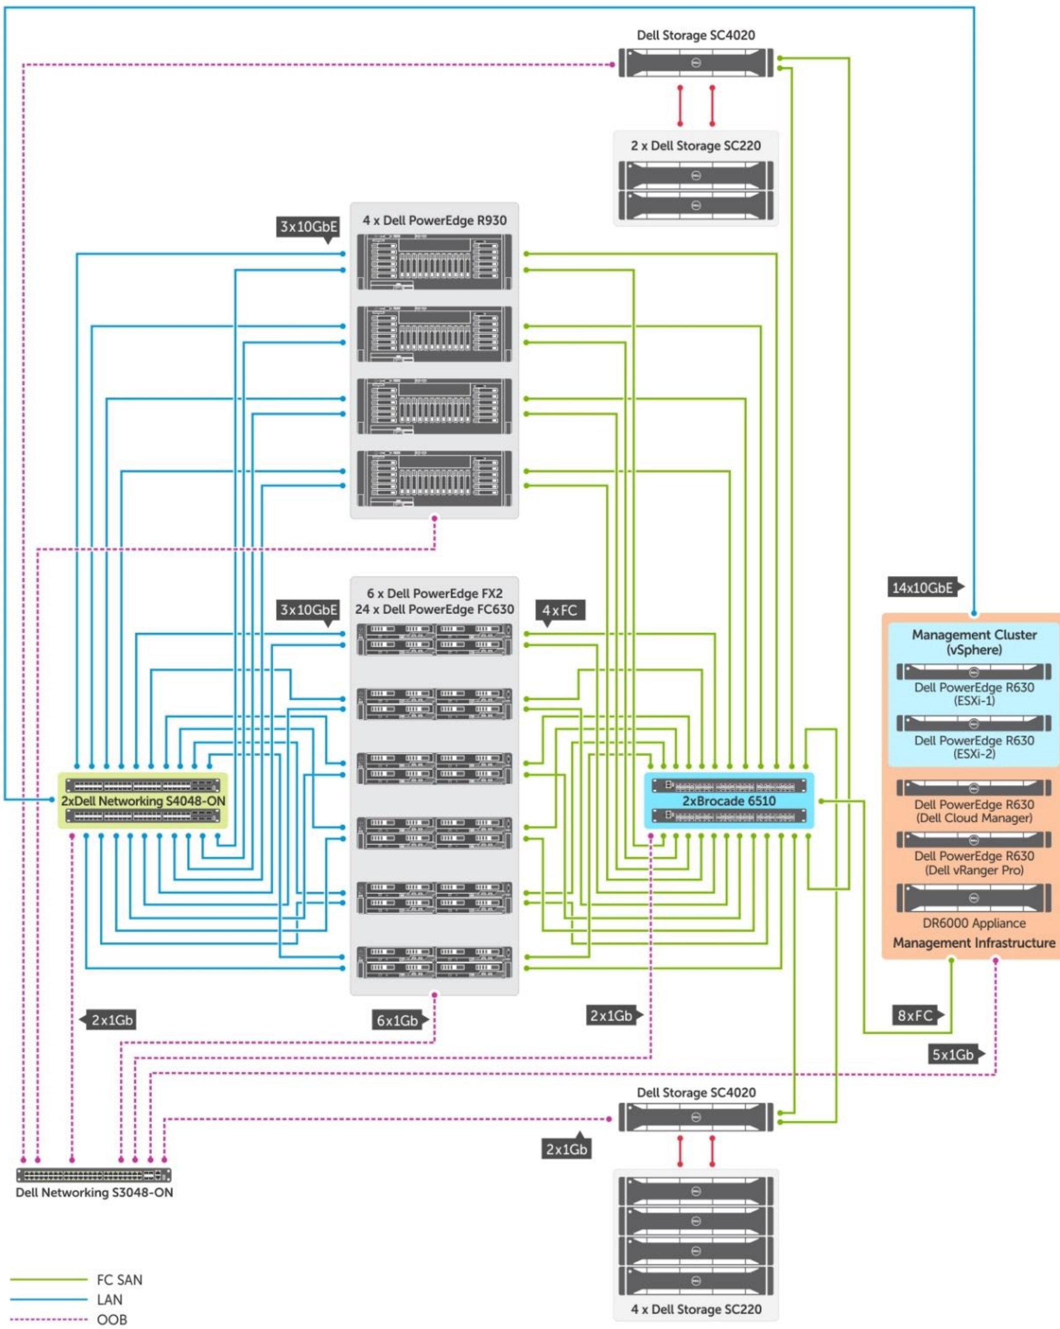

While this reference architecture presents a six PowerEdge FX2 based scale-unit as the initial deployment, the overall solution has enough room to grow beyond this capacity. The Dell Networking S4048-ON and Brocade 6510 switches provide the port density to add more compute chassis.

The following table describes the compute server configuration in this architecture:

**Table 9 SAP Compute configuration in the solution architecture**

| Resource | Description |
|---|---|
| Compute Nodes for SAP | 4 x PowerEdge FC630 in each PowerEdge FX2 |
| Compute Node for SAP HANA | 4 x PowerEdge R930 |
| Processors | 2 x Intel E5-2699v4 family processors in each FC630 server<br>4 x Intel E7-8890v3 family processor in each R930 |
| Memory | 256 GB in each FC630<br>2 TB in each R930 |
| Network | 1 x Qlogic 57840 Quad-Port Blade Network Daughter Card (bNDC) in each FC630<br>3 x Intel DualPort 520 DA 10 Gbit |
| FC Host Bus Adapter | 1 x Qlogic QLE2562 Dual-Port 8Gbps FC adapter mapped using PowerEdge FX2 PCIe expansion slots<br>2 x Qlogic QLE2660 Single Port 16 Gbit FC adapter |
| OS Volume | 2 x 16 GB SD cards in the internal SD card module |

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

The four R930 based HANA virtual Server can be used to deploy up to 16 productive HANA instances or 32 non-productive HANA instance with up to 512 GB of memory. The corresponding SC4020 Storage Array can host up to 8 productive HANA instances, with full HWCCT KPI allowance in parallel. For non-productive environments, we see between 30% to 50 % storage loads compared to productive environment and therefore the SC4020 is capable to carry up to 24 non-productive HANA instances.

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

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

The end to end IO connectivity represented in the following figure also consists of multiple port-channel and VLT configurations for LAN connectivity. Dell Networking S4048-ON switches provide the 10 GbE network connectivity between the SAP compute and SAP HANA compute cluster and the rest of the data center. Section 4.2 describes the network architecture in detail.

**Figure 17 End-to-end network and Fibre channel connectivity in the solution**

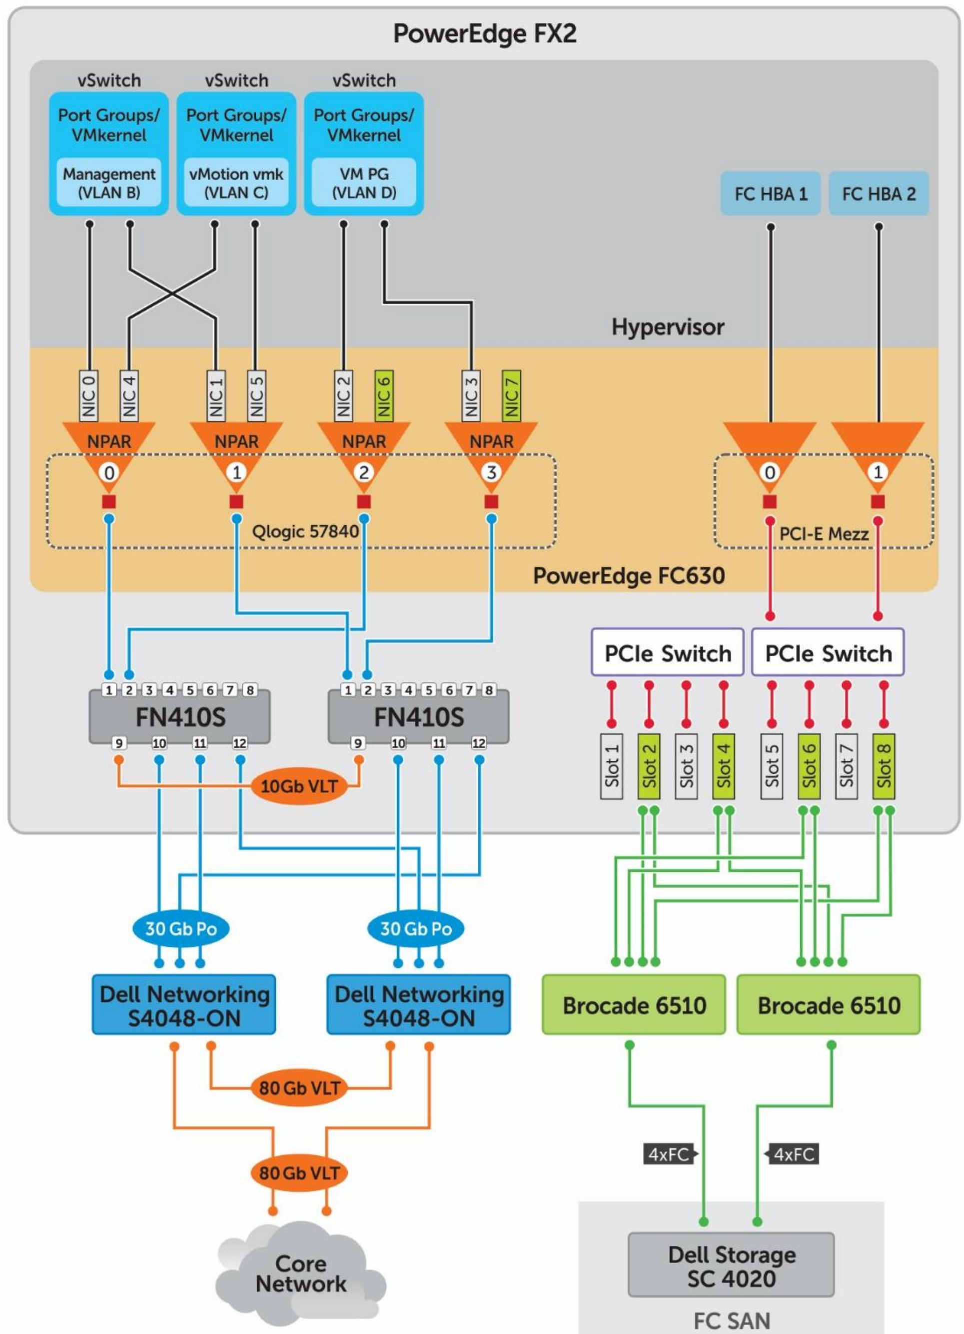

The figure above illustrates the end-to-end connectivity from PowerEdge FC630 in slot-1 of the PowerEdge FX2 chassis. This includes connectivity from the LAN and SAN fabric to the VM switch in the hypervisor.

The PowerEdge R930 Systems directly connects to the S4048-ON series switches. Dell recommends three vSwitches per R930 each with 10 Gb  ports connected.

- Client Networking (ACSC, Application Server, HANA Studio)
- HANA internode communication (only needed in virtual Scale-Out Scenarios)
- HANA Backup and recovery, and optional HANA system replication

## 5.1 Storage architecture and configuration

The following sections provide details of the storage architecture and configuration deployed in this solution.

### 5.1.1 Drive and RAID configuration for the SAP SC 4020

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

The choice of hybrid drive configuration enables scaling of compute infrastructure beyond 24 nodes by providing deep IO performance. This performance capability of the hybrid configuration can be complemented by adding more capacity in terms of Dell Storage SC220. Appendix B3  provides guidance on how the compute and storage scale out can be achieved, and the considerations for the scale-out. Table 3 provides the configuration of the Dell Storage SC4020 array used in this solution.

Table 10 Dell Storage SC4020 storage array configuration

| Component | Description |
|---|---|
| Disk Drives | 6 x 400 GB WI SSD<br>6 x 1.6 TB RI SSD<br>12 x 1.2 TB 10K SAS drives |
| Storage Profile | RAID 5 with stripe width 9<br>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 |

### 5.1.2 Drive and RAID configuration for the SAP HANA SC 4020

This reference architecture employs Dell Storage SC4020 array for HANA workload as described. The array is configured with a combination of Write-Intensive (WI) SSD drives and 10K SAS drives to provide the need KPIs for HANA and enough capacity.

Table 11 Dell Storage SC4020 storage array configuration

| Component | Description |
|---|---|
| Disk Drives | 6 x 400 GB WI SSD for up to 4 productive instances<br>12 x 400 GB WI SSD for up to 8 productive instances<br>Minimum 12 x 1.8 TB 10K SAS drives |
| Storage Profile | RAID 5 with stripe width 9 |
| Volumes | 3 volumes (HANA Data, HANA Log, HANA shared) per HANA instance |
| Replay Profile | HANA Replay profile |

### 5.1.3 Additional storage enclosures for increased capacity

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

### 5.1.4 Storage fabric configuration

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

**Figure 18 Dell Storage SC4020 connectivity to FC SAN fabric**

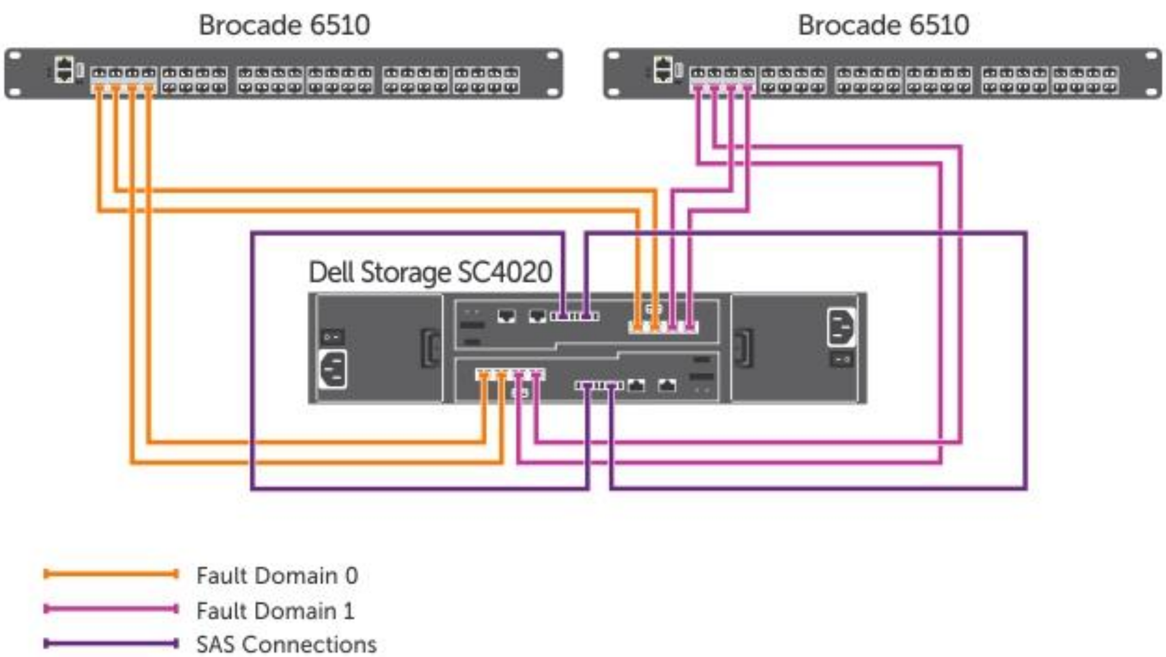


## 5.1.5 Enabling redundancy

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

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

For PowerEdge R930 we recommend using 2 x Qlogic QLE2660 FC HBAs. The configuration is analog to the FC630 configuration.

The figure below illustrates how the compute is connected to the SAN fabric for FC connectivity.

**Figure 19 Server to SAN fabric connectivity using PowerEdge FX2 PCIe slots**

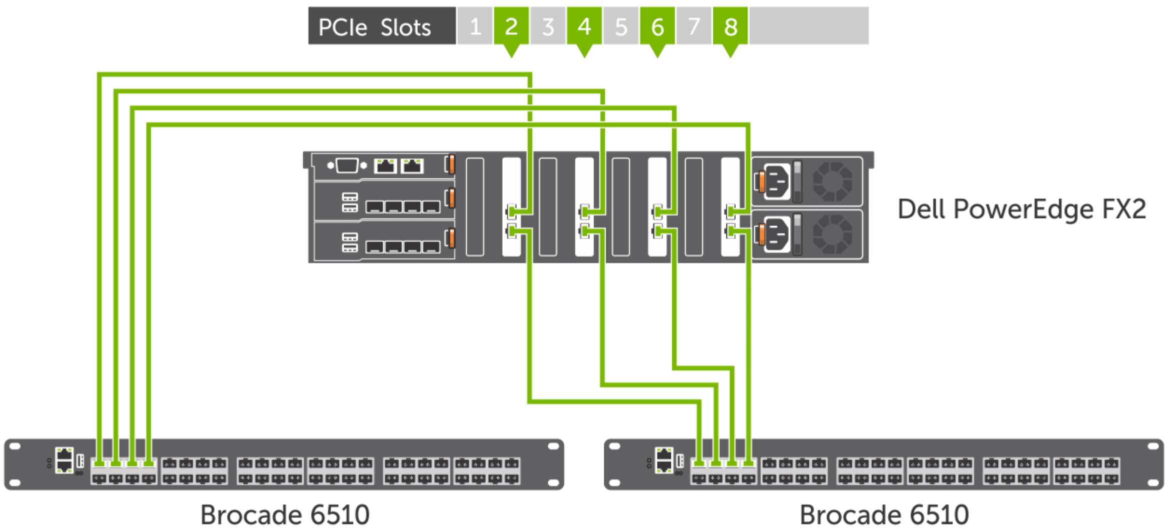


For HA of FC connections, ports from each FC HBA are connected to two different switches. These ports along with the FC connections from the SC4020 array are configured to be in the same zone to enable storage volume access to the compute cluster. The figure below shows the FC volumes mapped as datastores on the ESXi hosts within the vSphere 6 cluster.

**Figure 20 FC volumes mapped as data stores in vSphere cluster**

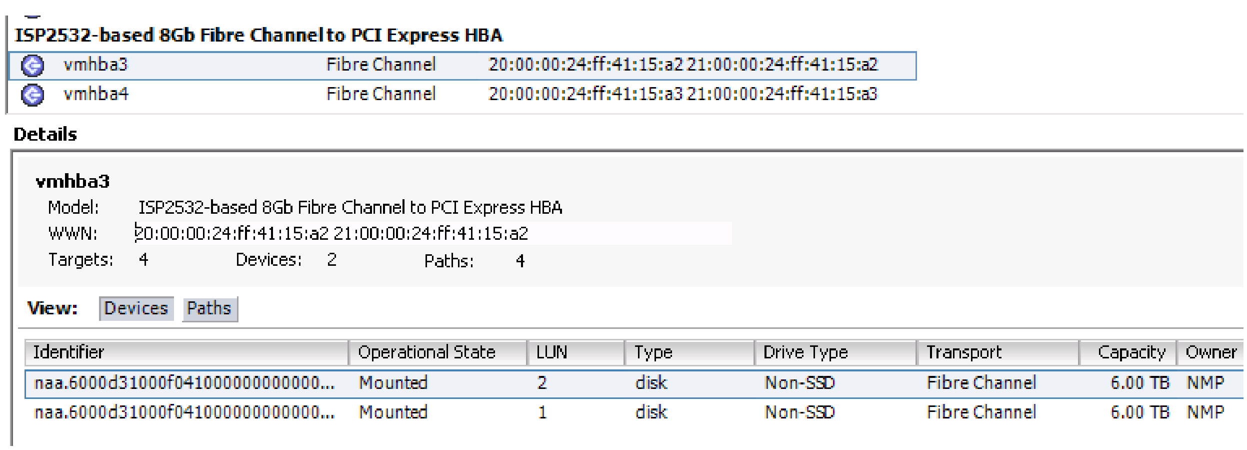


Multiple datastores within the vSphere cluster enable Datastore Heartbeating, which ensures that partitioning or isolated host networks do not trigger VM movement within the cluster. The following figure shows the default configuration for the Datastore Heartbeating in the vSphere cluster. By default,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

**Figure 21 Datastore Heartbeat Configuration**

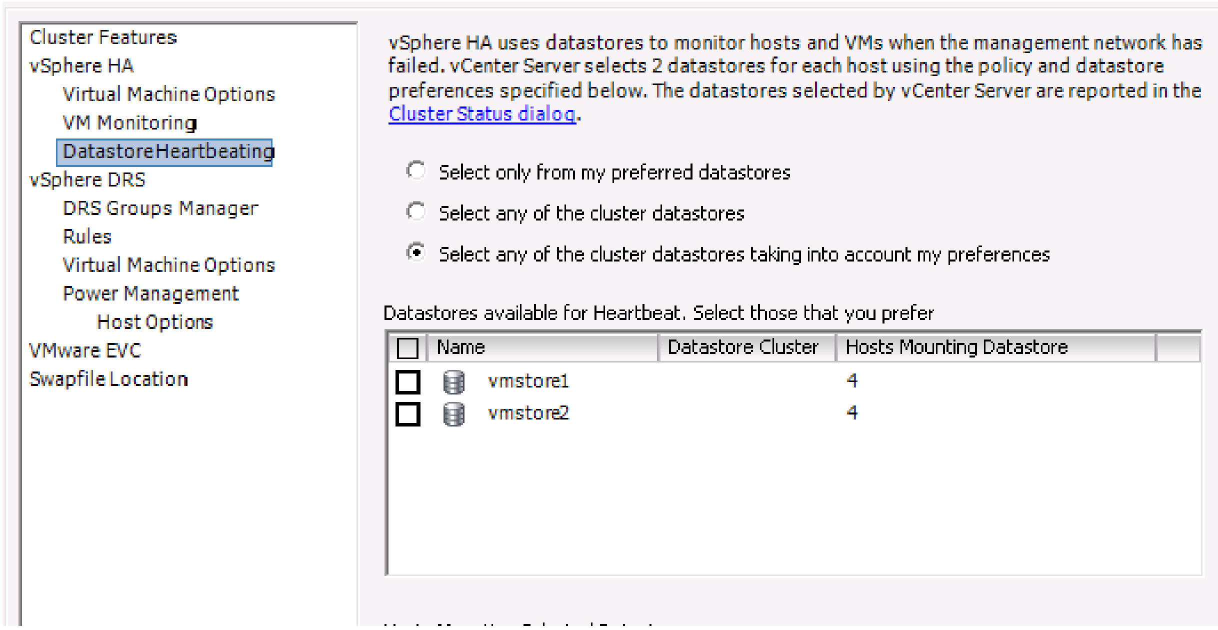


Redundant FC switches and datastore configuration in the vSphere cluster ensure that the storage in the infrastructure is highly available. The redundant controllers and power supplies in the storage array along with the fault domains within the Storage Center configuration ensure that the storage can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

## 5.2 Network architecture and configuration

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

### 5.2.1 Network design with mVLT

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

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

Major benefits of VLT technology are:

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

### 5.2.2 Server LAN configuration

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

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

The following figure illustrates this connectivity:

**Figure 22 Server LAN configuration**

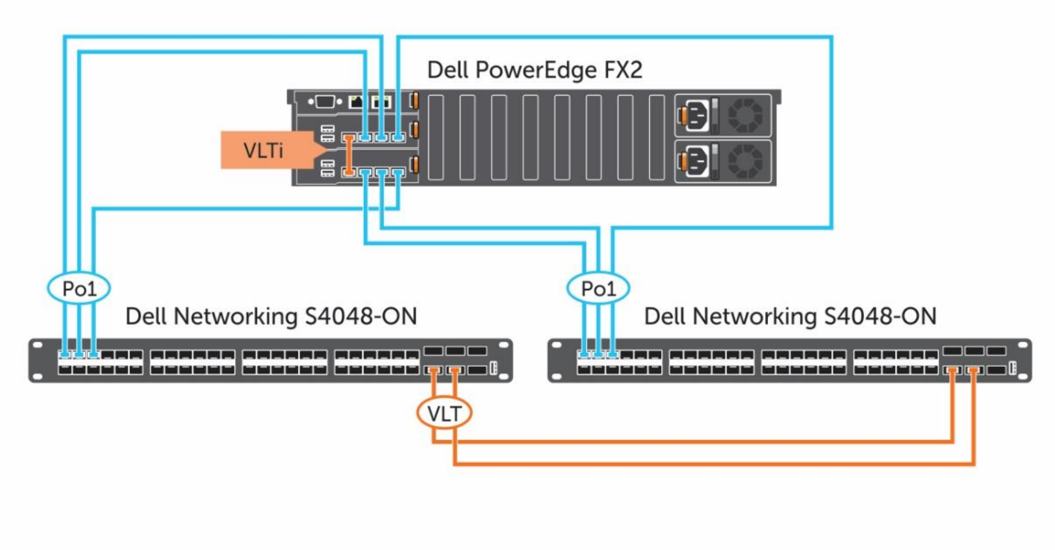


Since Ports 10, 11 and 12 from each IOA are connected to two different Dell Networking S4048-ON switches, a VLTi between the S4048-ON switches is mandatory. This provides complete redundancy for the network traffic from the chassis to the ToR switches. These ports are combined using port channel configuration and enable up to 60 Gbps of bandwidth between PowerEdge FX2 architecture and the ToR switches.

As shown in the figure above, two 40 GbE ports from each Dell Networking S4048-ON switch are connected together to create a VLT. This provides up to 80 GbE bandwidth between the ToR switches. The remaining four 40 GbE ports on each ToR switch can be used to connect to the data center core network. The ToR configuration in this architecture uses the management ports on both the switches for the VLT backup link.

### 5.2.3 NPAR and hypervisor network configuration for LAN traffic

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

NPAR maximum bandwidth settings may be modified to limit the maximum bandwidth available to a specific traffic type, regardless of contention. These adapters also allow relative bandwidth assignments for each partition. The table below shows the NPAR schema and the relative bandwidth assignments on each partition.

Table 12 NPAR schema and relative bandwidth weights

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

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

**Figure 23 Host LAN configuration using NPAR**

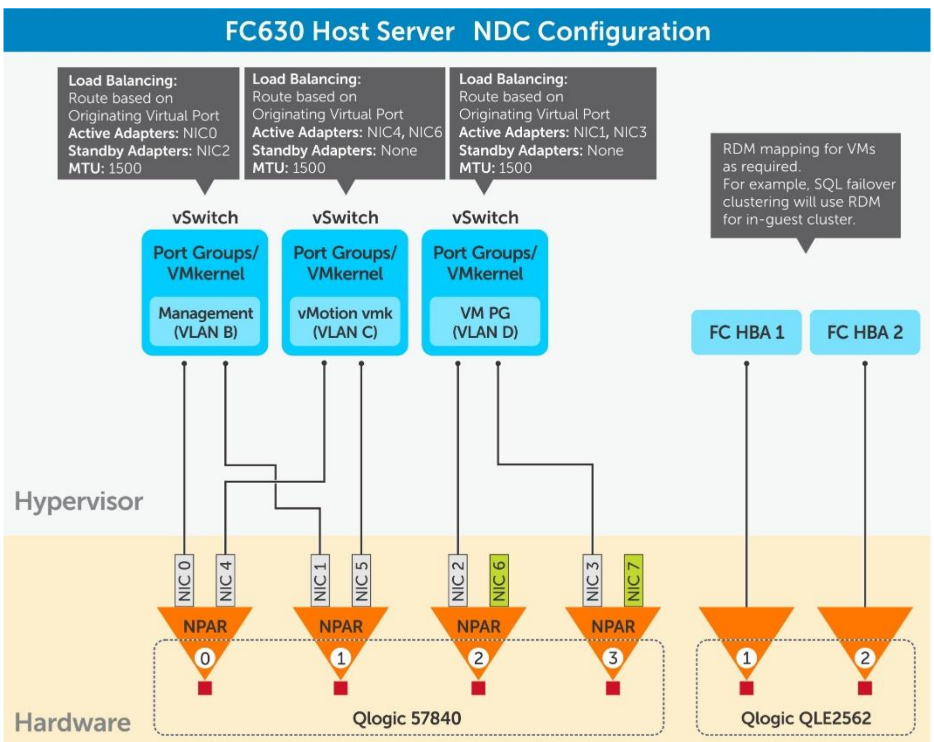


The figure below shows the host-level implementation of virtual switches as illustrated in the Host LAN diagram.

Figure 24 vSwitch implementation on every physical host

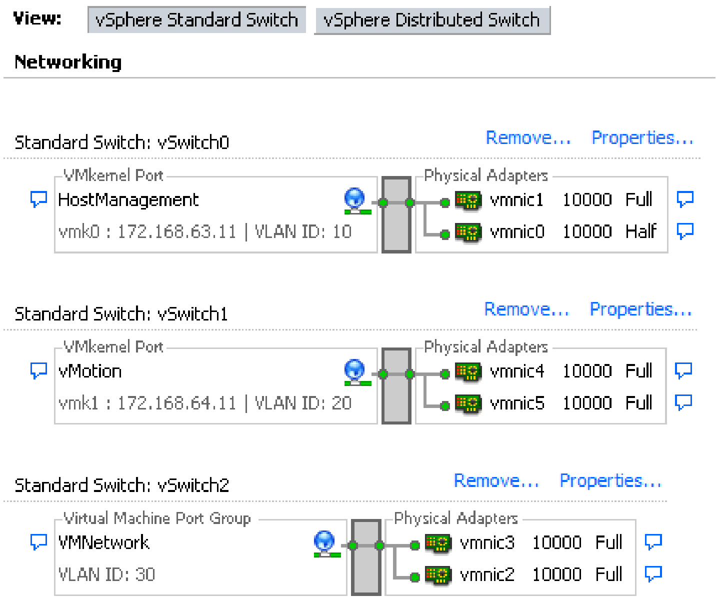


The load balancing algorithm for the virtual port groups is left to the **Route based on originating virtual port** default setting. Customers may consider changing this configuration based on the workload that is deployed on the virtualized infrastructure.

The Dell Networking S55 1 Gb switch in this architecture is an optional component and can be replaced by an existing out-of-band (OOB) management switch in the customer's existing data center environment.

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

## 5.3 Data protection, replication, and recovery architecture

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

For Backup, the solution leverages the deep VMware integration of Dell vRanger to acquire virtual disk read access and SAN transport 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 backup storage. Additional space savings techniques such as vRanger Active Block Mapping and VMware Changed Block Tracking can be enabled to reduce the data transport volume.

For all SAP and SAP HANA workloads we also recommend the full usage of Dell vRanger. For bigger HANA databases (>256 GB) we recommend to directly backup the HANA databases to the DR6000 system. The DR6000 will be connected via NFS into the HANA VM.

With each virtual HANA system deployed from our OVF template we also deliver a backup script for easy use.

**Figure 25 Backup script**

```
Welcome to SLES for SAP Applications GMC 11.1 (x86_64) - Kernel 2.6.32.24-0.2-de
fault (tty1).


delhanar03 login: root
Password:
Last login: Wed Feb 22 10:33:11 CET 2012 from console on :0
You have new mail.
delhanar03:~ # backup_hana_database.sh
```

The backup script tells you that it is prepared to mount an external storage on the system. Mount the DR6000 via NFS to e.a. /hana/autobackup. Ensure that there is enough space to perform the backup. As a rule, you need roughly four times of the size of the main memory for the backup.

**WARNING**: If your database grows over a certain amount it could become unusable as there is no space left on the data volume.

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

To restore a backup, the target HANA version must be the same or higher than the version of the sources with which the backup has been performed.

**Figure 26 Data Protection Architecture**

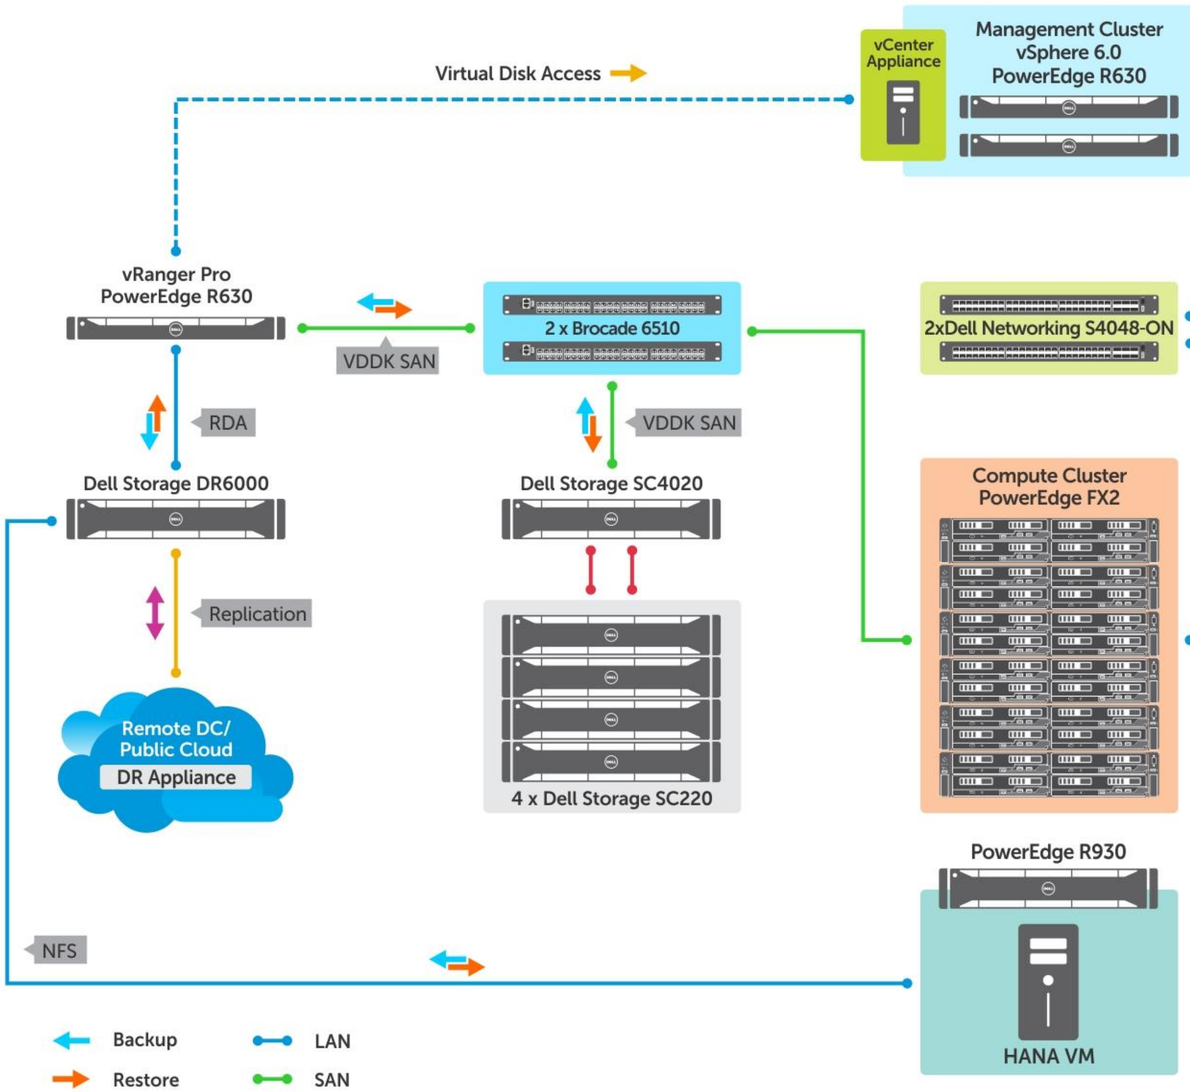

## 5.4 Management components and architecture

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

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

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

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

### 5.4.1 Management architecture

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

The Solution overview section describes the recommended software components for the management infrastructure. Most of the management software except Dell vRanger Pro and Dell DR4100 appliance can be implemented as virtual machines and can share the storage capacity available on the Dell Storage SC4020 array used for the virtual infrastructure.

Similarly, it is recommended to deploy the same host networking configuration as described in the Sizing Guidelines section on the hypervisor hosts used for deploying the management components. By doing so, administrators can leverage the existing network configuration and infrastructure. The Dell Network S4048-ON switches used in the solution architecture provide the necessary port density to connect the management infrastructure.

The following table provides an overview of the virtual infrastructure requirements for deploying the management components described in the Solution overview section.

Table 13 Management Sizing

| Component | VMs | CPU Cores | RAM (GB) | OS (GB) | Data (GB) | NIC (Gbps) |
|---|---|---|---|---|---|---|
| SQL Server 2012 | 1 | 8 | 16 | 60 | 212 | 1 |
| VMware vCenter Server | 2 | 4 | 12 | 60 | 50 | 1 |
| Dell Active System Manager | 1 | 4 | 16 | 80 | 0 | 1 |
| VMware vROperations | 1 | 4 | 4 | 40 | | 1 |
| Dell Enterprise Manager | 1 | 2 | 4 | 40 | | 1 |
| Dell Replay 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 vRA Appliance | 1 | 4 | 16 | 30 | 0 | 1 |
| VMware vRA IaaS Core | 1 | 4 | 8 | 60 | 40 | 1 |
| VMware vROrchestration | 1 | 2 | 3 | 7 | 0 | 1 |
| VMware vShield App | 4 | 2 | 1 | 0 | 5 | 1 |
| SAP Landscape Virtualization Management | 1 | 16 | 32 | 60 | 250 | 1 |

54 vCPUs and 119 GB of vMemory are required to implement the virtualized management infrastructure. Dell vRanger Pro and Dell DR6000 appliance are implemented as physical hosts. As described in Section 4.3 in this document, Dell vRanger Pro uses SAN transport to improve the backup performance for the virtual infrastructure. This functionality can only be implemented on a physical host.

The figure below illustrates the management architecture connecting to the existing LAN and SAN infrastructure used for the compute.

When virtualizing these management components, administrators must ensure that the infrastructure is highly available. For storage and network, the underlying virtual infrastructure provides HA. 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.

Figure 27 Management Architecture

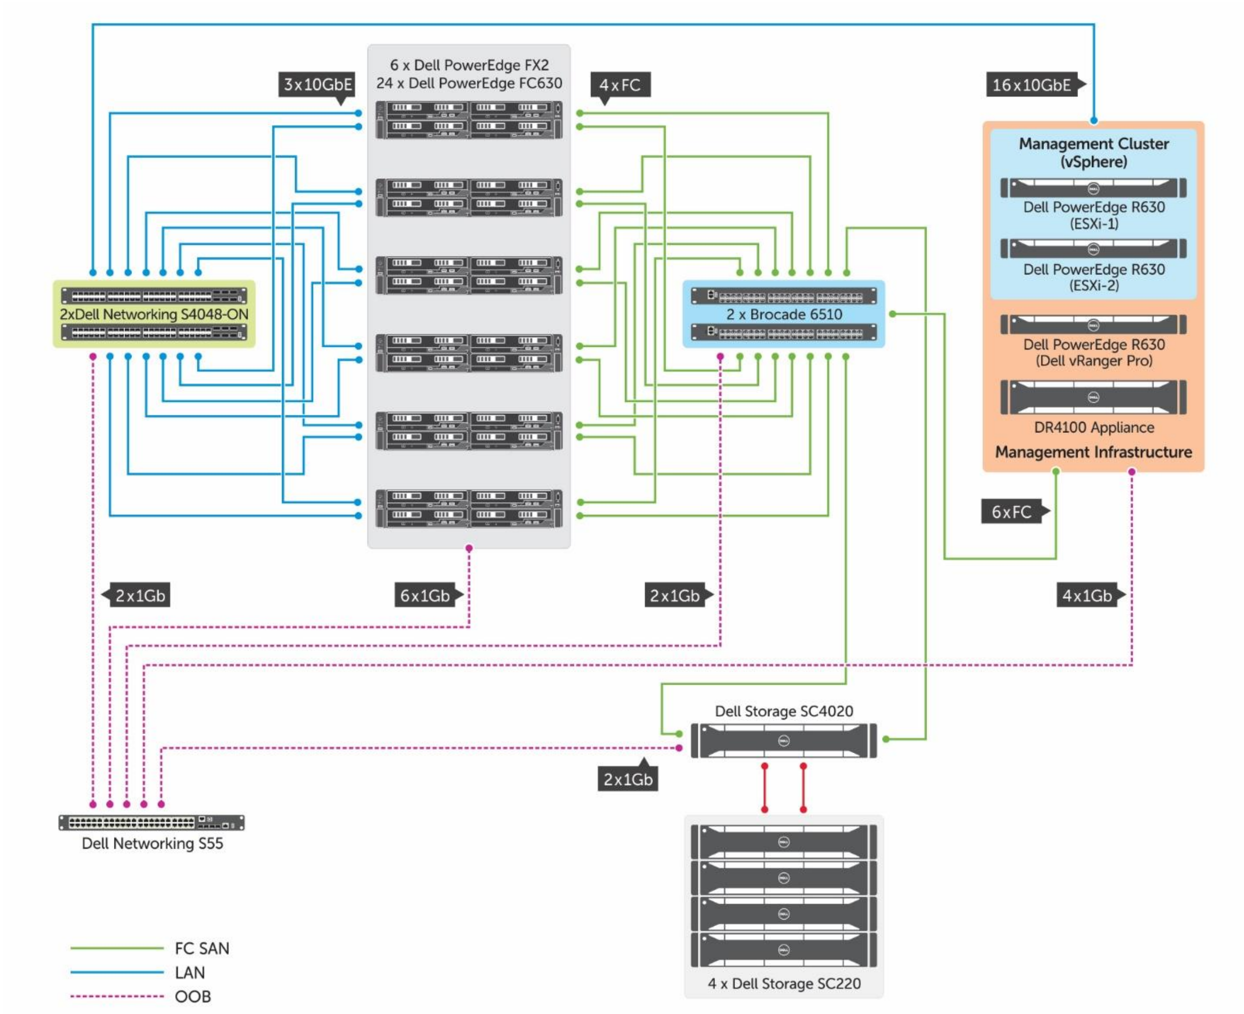


The table below provides the specifications for the servers in the management infrastructure.

Table 14 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 16GB DDR4 DIMMs |
| | OS Volume | 2 x 16 GB SD cards in internal SD module |
| | Network | Qlogic 57840S quad-port NDC |

| Component | Details | |
|---|---|---|
| | SAN | Qlogic 2562 |
| | OS | VMware ESXi 6 |
| Dell vRanger Pro | 1 x Dell PowerEdge R630 | |
| | Processor | 2 x Intel Xeon E5-2609v3 Family |
| | Memory | 64 GB; 8 x 8GB DDR4 DIMMs |
| | OS Volume | 2 x 600 GB 15K SAS drives in RAID 1 |
| | Network | Qlogic 57840S quad-port NDC |
| | SAN | Qlogic 2562 |
| | 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 HA and DRS features, the virtualized management infrastructure can be made highly available and resilient even in case of 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 DR6000 appliance is a physical host that gets connected to the existing LAN infrastructure used for the compute.

# 6 Executive summary

Customers spend enormous time to procure, validate, and integrate the components of the end-to-end infrastructure, on their own. The Dell Blueprint solutions and reference architecture provides a predesigned and validated reference architecture and not only reduces risk, but also provides an optimized solution that can be deployed and operated quickly and efficiently.

## 6.1 Why use the Dell Blueprint solutions and reference architectures?

Dell's strength is its broad product and services portfolio that can be tailored to meet customer's requirements regardless of their size, scale, or business model. But a broad portfolio does not mean that it is more complex or difficult to choose.  Dell engineering teams have performed sizing, testing, and validation to help remove the complexity and match our blueprint solutions to suit the requirements of the customer.

To address the challenges discussed in the Customer challenges section, Dell has listened to customers and designed a portfolio of reference architectures that are optimized for specific applications and workloads and yet are flexible enough to scale and adapt as needed.

These reference architectures are:

- 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.


- Optimized and protected with investments Dell has made in software that makes our solutions easier to manage and deliver ongoing results.


- Delivered with Dell's global reach and exceptional execution and delivery to provide consistent deployment, management, and maintenance in every region of the world.

# A Additional resources

- Support.dell.com is focused on meeting customer requirements with proven services.
- Dell TechCenter is an IT Community where you can connect with Dell customers and Dell employees for the purpose of sharing knowledge, best practices and information about Dell products and installations.
- Dell Active System Manager on Dell TechCenter provides information on ASM, including links to resources, technical documents, and videos
- Referenced or recommended Dell publications:
  - Dell PowerEdge FX2
  - Dell Storage SC4020
  - Dell Storage SC220
- Referenced or recommended VMware publications or websites:
  a. New features in VMware vSphere 6
  b. VMware vRealize Automation
  c. VMware vRealize Orchestrator

# B Solution component details

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

## B.1 Dell PowerEdge FX2 architecture

This table presents an overview of all component configurations supported in PowerEdge FX2 architecture.

Table 15 Dell PowerEdge FX2 supported components overview

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

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

The following figure shows the internal mapping of the PCIe slots to servers in the PowerEdge FX2 architecture:

**Figure 28 PCIe slot mapping in PowerEdge FX2 (with half-width servers)**

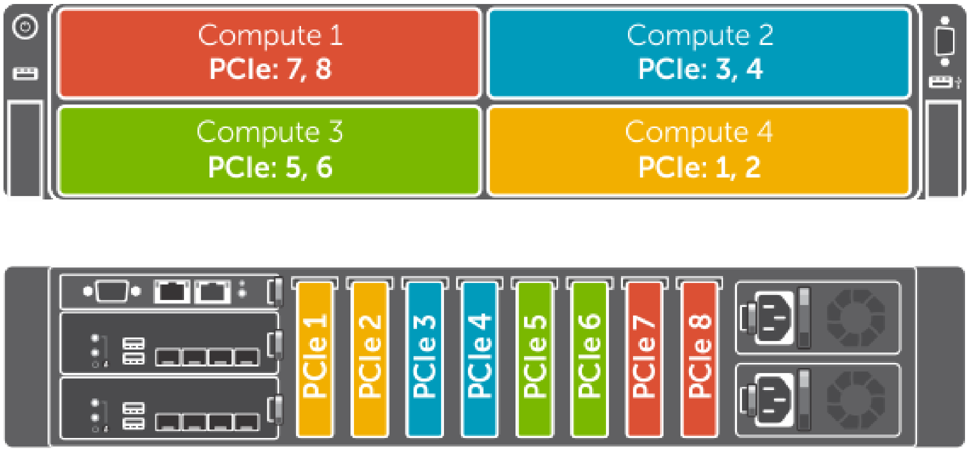


The PowerEdge FX2 architecture is managed through the Chassis Management Controller (CMC). The CMC Web console provides features to manage servers and IOMs and monitor PCIe slots in the PowerEdge FX2 architecture. The following figure shows the group management of multiple PowerEdge FX2 architecture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

**Figure 29 PowerEdge FX2 CMC web console - Group Management**

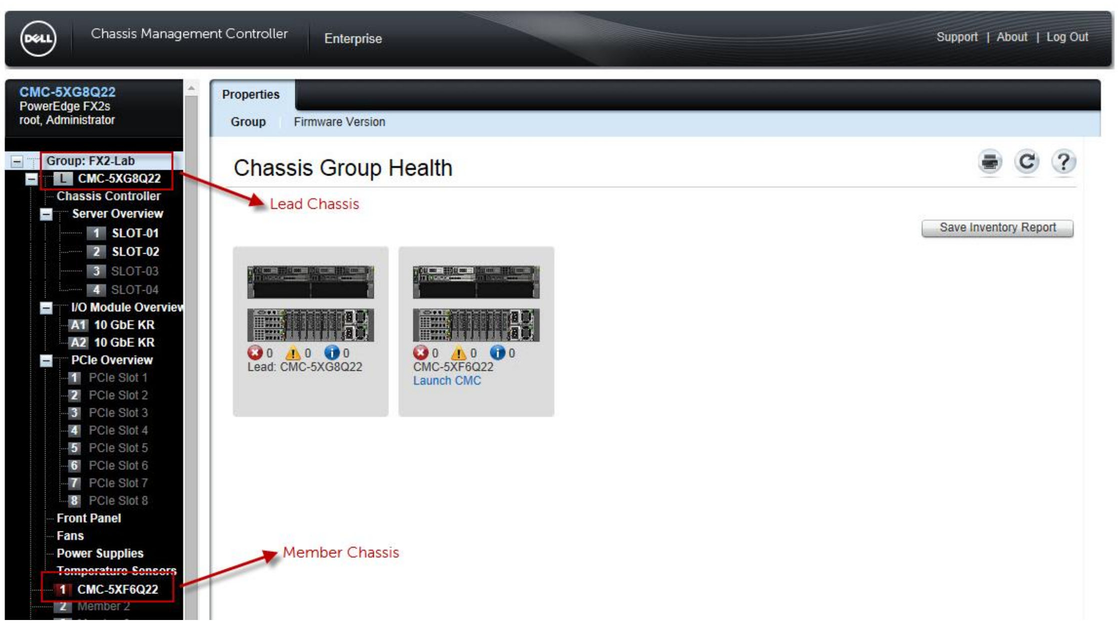


Using the group management feature of CMC, it is possible to manage up to 20 PowerEdge FX2 architectures from the same CMC console. This can be beneficial when there are multiple chassis participating in an infrastructure deployment, such as the virtualization infrastructure implemented in this reference architecture.

### B.1.1 Dell PowerEdge FC430 server

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

The following table describes the components supported in a PowerEdge FC430 server:

Table 16 Supported components in PowerEdge FC430

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

### B.1.2 Dell PowerEdge FM120x4 server

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

The following table describes the components supported in a PowerEdge FM120x4 server:

Table 17 Supported components in PowerEdge FC630

| | Supported in PowerEdge FM120x4 |
|---|---|
| CPU | 4 x Intel Atom processor C2000 family |
| Memory | 2 x DIMMs DDR3 memory per SOC, UDIMM only, (8 per sled) |
| Networking | 2 x 1Gb controllers embedded in the SOC |
| Storage | 1 x 2.5-inch front access hard drives or 2 x 1.8" SSD drives |
| Form Factor | ½ width Sled |
| PCIe slots | No PCIe slots |
| Systems Management | iDRAC8 Express or iDRAC8 Enterprise |

### B.1.3 Dell PowerEdge FC830 server

The PowerEdge FC830 server is a full-width server sled that supports Intel Haswell processors up to 18 processing cores and up to 1.5 TB physical memory. This enables PowerEdge FC830 to deliver the performance required for compute-intensive tasks. A Qlogic 57810 or Qlogic 57840 10 GbE or Emulex 10 GbE blade Network Daughter Card (bNDC) used within the server provides the network connectivity through the 10 GbE pass-through IOM or 10 GbE IOA in the PowerEdge FX2 architecture.

PowerEdge FC830 servers use blade mezzanine cards to connect to the PCIe add-on peripherals in the PowerEdge FX2 architecture. This connectivity can be leveraged to build optional redundancy for network connections. For operating system (OS) drives, PowerEdge FC830 supports a wide variety of PowerEdge RAID controller (PERC) devices. An internal SD card module can also be used for deploying host OS on a PowerEdge FC830 server.

The following table describes the components supported in a PowerEdge FC830 server:

Table 18 Supported components in PowerEdge FC830

| | Supported in PowerEdge FC830 |
|---|---|
| CPU | Up to 4 Intel Xeon E5-4600 family processors |
| Memory | 48 DIMMs; up to 1.5 TB[2] |
| Networking | Qlogic 57810 Dual Port 10 GbE bNDC<br>Qlogic 57840 Quad Port 10 GbE bNDC |

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

| | |
|---|---|
| | Emulex 10GbE adapters |
| Storage | PowerEdge Express Flash NVMe PCIe SSD, SATA HDD/SSD or SAS HDD/SSD Up to 8 x 1.8" SSD or 2 x 2.5" |
| Form Factor | ½ width Sled |
| PCIe slots | 4 x PCIe slots (x8) provided by PowerEdge FX2 |
| Systems Management | iDRAC8 Express or iDRAC8 Enterprise |

### B.1.4 Dell PowerEdge FC630 server

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

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

The following table describes the components supported in a PowerEdge FC630 server:

Table 19 Supported components in PowerEdge FC630

| | Supported in PowerEdge FC630 |
|---|---|
| CPU | Up to 2 Intel Xeon E5-2600 family processors |
| Memory | 24 DIMMs; up to 768 GB[3] |
| Networking | Qlogic 57810 Dual Port 10 GbE bNDC<br>Qlogic 57840 Quad Port 10 GbE bNDC<br>Emulex 10GbE adapters |
| Storage | PowerEdge Express Flash NVMe PCIe SSD, SATA HDD/SSD or SAS HDD/SSD Up to 8 x 1.8" SSD or 2 x 2.5" |
| Form Factor | ½ width Sled |
| PCIe slots | 2x PCIe slots (x8) provided by PowerEdge FX2 |
| Systems Management | iDRAC8 Express or iDRAC8 Enterprise |

### B.1.5 Dell PowerEdge FN410S

As mentioned in Appendix B1, PowerEdge FX2 supports multiple networking options for server connectivity to the top-of-rack (ToR) switches and/or to the data center core. These networking options provide flexible choices when designing a solution infrastructure. This reference architecture employs PowerEdge FN410S for the TOR connectivity from the PowerEdge FX2 chassis.

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

## B.2 Dell Networking S4048-ON switch

Dell Networking S4048-ON is a 1U high-density 10/40 GbE ToR switch with 48 dual-speed 1/10GbE (SFP+) ports and six 40 GbE (QSFP+) uplinks. This switch leverages a non-blocking and cut-through switching architecture to provide ultra-low-latency performance for applications. The six 40 GbE ports

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

can be used to create either a Virtual Link Trunk (VLT) between the switches to enable traffic isolation within the solution infrastructure or as a connectivity to the data center core network.

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

## B.3 Dell Storage SC4020 array

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

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

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

## B.4 Brocade 6510

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

## B.5 VMware vSphere 6

vSphere 6 is an industry leading server virtualization platform that enables creation of a highly available, resilient and on-demand infrastructure for virtualizing enterprise applications or as a general purpose virtual infrastructure. v6 provides improved features that are available to enterprises looking at data center consolidation and enhancing resource utilization. Apart from the scalability improvements, vSphere 6 enhanced vSphere Fault Tolerance support for multi-vCPU VMs, compression support for vSphere Replication and added support for multisite content library.

Table 20 The table below provides an overview of the capabilities in.VMware vSphere 6 Cluster

| | **VMware vSphere 6** |
|---|---|
| VM Hardware Version | Virtual Hardware 11 |
| Maximum Number of vCPU per VM | 128 |
| Maximum Memory per VM | 4 TB |
| Cluster Nodes | 64 |
| Maximum CPU per Host | 480 |
| Maximum Memory per Host | 12 TB |
| Maximum VMs per Cluster | 8000 |

The virtualization capabilities offered by vSphere 6 along with the management capabilities offered by vCenter enable building an efficient platform for virtualizing enterprise applications.