



# Dell Validated System for EMC ScaleIO

Version 1.0

Dell Engineering  
June 2016

## Revisions

Date	Description
June 2016	Initial release 1.0

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.  
Copyright © 2016 Dell Inc. All rights reserved. Dell and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

# Table of contents

Revisions.....	2
1 Audience and Purpose.....	5
2 Executive summary .....	6
3 Why Dell .....	7
4 Why EMC ScaleIO .....	8
5 Customer Use Case and Workload .....	10
6 Solution Overview .....	12
6.1 Dell PowerEdge R730XD .....	15
6.2 Dell Switches.....	15
6.3 Dell PowerEdge R630 .....	15
7 Overview of Software-Defined Storage with EMC ScaleIO .....	16
7.1.1 EMC ScaleIO Software Components .....	16
7.1.2 EMC ScaleIO Storage Definitions.....	17
8 Infrastructure Overview .....	20
8.1 Compute Infrastructure .....	20
8.1.1 EMC ScaleIO .....	20
8.2 Management Infrastructure .....	21
8.2.1 VMware vCenter Server .....	21
8.2.2 EMC ScaleIO GUI .....	23
8.2.3 EMC ScaleIO Gateway .....	23
8.2.4 EMC ScaleIO Plugin.....	24
8.2.5 OpenManage Essentials.....	24
8.3 Network Infrastructure.....	25
9 Configuration Guidance.....	30
9.1 Compute Node Configurations.....	30
9.1.1 Mixed Configuration.....	31
9.1.2 Standard Configuration.....	32
9.2 Management Node Configuration.....	32
9.3 Base configurations and scaling.....	33
9.4 Best Practices.....	34
9.4.1 Compute.....	34



9.4.2 Storage .....	34
9.4.3 Network .....	34
9.4.4 Storage Pool.....	35
9.4.5 Network Scaleout .....	35
10 Component Specifications .....	38
10.1 Software Components .....	39
10.2 BIOS, Firmware and Drivers .....	39
11 Conclusion .....	41
A Deployment Guidelines .....	42
B Additional resources.....	45
B.1 Storage components.....	45
B.2 Virtualization components .....	45
B.3 Dell servers and switches .....	46
B.4 Other Dell resources .....	46
C Abbreviations and Acronyms .....	47

# 1 Audience and Purpose

The audience for this paper includes, but is not limited to, sales engineers, technologists and architects, field consultants, partner engineering team members, customers, and anyone else interested in high level guidance for deploying an optimized and validated solution stack with the best of Dell and EMC ScaleIO technology.



## 2 Executive summary

The Dell Validated System for EMC ScaleIO is intended to help customers that are building out a ScaleIO environment to simplify their implementation's design, architecture, configuration, and deployment. Customers want Dell to speed their time to value, and to decrease the amount of time it takes to get up and running with an optimized ScaleIO rack level solution.

In response to customer demands, the Dell Validated System (VS) for EMC ScaleIO is designed to be modular and flexible. The system's modularity and flexibility allow customers to build with validated building blocks at rack scale:

- Compute/Storage
- Network
- Management Stack

Customers can take the building blocks that make sense for them. Customers have flexibility to scale based on their own workloads and use cases.

This document covers at a high level the architecture and deployment of running ScaleIO software on Dell PowerEdge servers at rack scale. This document is intended to explain, in general terms, how Dell VS for EMC ScaleIO can work for customers. It is not meant to be a planning guide for specific software-defined storage (SDS) use cases.

Dell recommends that customers also consider cost, reliability and serviceability, plus key performance indicators (KPIs) such as input/output operations per second (IOPS), latency and throughput.

For customers needing further guidance, please contact Dell or EMC ScaleIO sales teams.

### 3 Why Dell

Dell believes in the value of validated systems because these systems offer flexibility and choice depending on each customer's consumption model. Dell has been architecting validated systems for many years. This experience has given Dell a proven validation and certification process to provide customers peace of mind that their implementation works. Dell now has collaborated with EMC ScaleIO to validate the system's configuration, architecture, and rack level design, ensuring that customers achieve optimal performance and scale. This collaboration allows customers to run EMC ScaleIO on award-winning Dell PowerEdge Servers, with flexibility across the network & management framework.

Customers are starting down the journey towards a hyper-converged environment to help drive down their total cost of ownership (TCO). The first step for most customers is to understand SDS technology. Because SDS is a new technology and expertise is limited, customers struggle to get value out of their initial project. Customers with limited resources do not want to spend their time testing and developing server configurations and network architecture; this can be costly and complicated. Customers do want to build with validated, tested building blocks to help decrease time to value by simplifying design, architecture, configurations, and deployment. It is not in the customer's best interest to maintain SDS design and hardware validation labs where their own engineers assume storage vendor responsibility to test, qualify, integrate and provide problem resolution. This is where Dell and EMC ScaleIO can add value.

The Dell Validated System (VS) for EMC ScaleIO provides the only engineering-validated system focused on ScaleIO running on industry-standard Dell PowerEdge servers at rack scale. Customers believe SDS offers potential benefits such as storage solution flexibility, increased performance and reduced cost but they still need help getting there. Dell is here to help so customers can utilize Dell's server technology and expertise to simplify all the front end work.

Customers are realizing SDS is a new technology bringing new complexity; it is not simply a matter of loading software on a server. Through 2017, over 30% of SDS implementations will not realize the expected 40% to 50% reduction in TCO because of poor hardware selection and implementation.<sup>1</sup> Dell and EMC ScaleIO can help customers simplify the design, planning, testing, and architecture of their SDS implementations to help deliver the TCO they expect for their projects.

---

<sup>1</sup> Gartner September 2015- Why Hardware Matters to the Success of Your Software-Defined Storage Deployment



## 4 Why EMC ScaleIO

EMC ScaleIO is software that creates a server-based SAN environment built from server direct-attached storage (DAS) to deliver flexible, scalable performance and capacity on demand. ScaleIO converges server compute and storage resources into a single-layer architecture, aggregating capacity and performance, scaling to thousands of nodes. Performance matters, and every server in the ScaleIO cluster is used in processing I/O operations, providing all I/O and throughput available to any application running on the cluster. Its decentralized data flow is key to the linearly scalable performance of ScaleIO, and results in a highly parallel system, eliminating bottlenecks. The high application performance that results enables customers to meet service level agreements, with enhanced system control and manageability.

Additionally, ScaleIO provides software-defined storage designed with enterprise-grade resiliency. ScaleIO delivers scale-out, block storage using commodity hardware. ScaleIO provides the scalability, elasticity, flexibility, and performance needed to meet the demands of performance hungry applications. It is designed to massively scale from three to thousands of nodes, combining HDDs and SSDs to create virtual pools of block storage defined by performance requirements. Unlike most traditional storage systems, as the number of servers grows, so do throughput and IOPS (Input/Output per Second). The scalability of performance is linear with regard to the growth of the deployment. Whenever the need arises, additional storage and compute resources (that is, additional servers and drives) can be added modularly. Storage and compute resources grow together so the balance between them is maintained. Storage growth is therefore always automatically aligned with application needs.

ScaleIO is designed and optimized to transform your server and storage performance, and to address your critical needs. Customers receive countless advantages:

- **Massive Scalability:** The Dell VS for EMC ScaleIO can scale up to 16 nodes in a single rack. Clusters can scale to 1000+ nodes in multiple racks.
- **Supreme Elasticity:** Storage and compute resources can be increased or decreased whenever the need arises. The system automatically rebalances data “on the fly” with no downtime. Additions and removals can be done in small or large increments. No capacity planning or complex reconfiguration due to interoperability constraints is required, which reduces complexity and cost.
- **Extreme Performance:** Customers can achieve 8x better IOPS performance than traditional storage<sup>2</sup>. Performance scales linearly and cost/performance rates improve with growth. Performance optimization is automatic. Whenever rebuilds are needed, they occur in the background with minimal or no impact to running applications.
- **Compelling Economics:** Customers can expect an as much as 50% more efficient <sup>3</sup>power/cooling/space compared to their traditional storage infrastructures when replacing their storage server plus application server. Also, with the ability to add individual nodes in the Dell VS

---

<sup>2</sup> Based on EMC ScaleIO Testing and Validation versus traditional SAN environments

<sup>3</sup> Based on EMC ScaleIO Testing and Validation versus traditional SAN environments



for EMC ScaleIO cluster, you pay as you grow, and therefore you do not need to pre-pay for infrastructure you may not need.



## 5 Customer Use Case and Workload

Customers are beginning to take the journey towards a hyper-converged infrastructure, and for many the first step starts with software defined storage (SDS). SDS promises to deliver modern storage and data services as software-based capabilities that can leverage server platforms to improve storage economics, as well as to provide data mobility, and build towards hyper-converged infrastructure.

The biggest misconception about SDS is that the underlying hardware is not important because many of the features and reliability and performance enhancements are now provided by software. Some SDS software vendors treat server hardware as an afterthought, providing one-size-fits-all reference architectures that will not help Infrastructure and Operations (I&O) leaders to reduce costs or maintain required performance compared to traditional storage arrays.<sup>4</sup> To obtain cost savings, software-defined technology is heavily dependent on hardware innovation and efficiency. Poorly selected and right sized industry-standard hardware can result in SDS deployments that do not match the cost and performance requirements of the organization.<sup>5</sup> Dell and EMC ScaleIO have engineered the Dell VS for EMC ScaleIO to provide the most efficient server configurations to help customers reach the performance and scaling requirements to run their applications utilizing SDS.

When implementing an SDS solution, it is important to define the use case and the workload that it will be supporting, the customer's IT environment, and consumption model and scope.

- **Workload and Storage Requirements**
  - What workload and application are running?
  - What are the requirements of the workload?
  - Is this customer an early adopter, or risk-averse?
- **Customer IT Environment**
  - Are there processes in place that favor one solution over another?
  - What are the customer's skill set and level of expertise?
  - What is the support & development model?
- **Consumption Model & Scope**
  - What standards are in place? Consider OS, virtualization, backup, and so on.
  - What technologies are currently deployed?
  - Is this DIY? Build with validated building blocks? Appliance?

---

<sup>4</sup> Gartner Why Hardware Matters to the Success of Your Software-Defined Storage Deployment Published: 8 September 2015

<sup>5</sup> Gartner Why Hardware Matters to the Success of Your Software-Defined Storage Deployment Published: 8 September 2015

The Dell VS for EMC ScaleIO is for customers interested in purchasing SDS software as a first step toward SDS/HCI offerings that require both granularity and scale. The system is designed for operational efficiency by driving lower TCO for enterprise applications (Oracle, SAP, SQL, and Exchange) as well as Cloud or next-generation applications such as NoSQL, Splunk, and OpenStack. Additional value is added by providing rack level validation with modular (optional) building blocks, creating a new consumption model for SDS on a validated, industry-leading compute platform.



## 6 Solution Overview

This paper describes Dell VS for EMC ScaleIO for a general purpose private virtualization infrastructure solution based on:

- The Dell PowerEdge R730xd platform
- VMware vSphere and vCenter Server, and
- EMC ScaleIO software-defined storage.

This validated system provides a scalable and highly-available infrastructure with built-in monitoring via VMware vCenter, ScaleIO Console and Dell OpenManage Essentials. It is designed according to the best practices and recommendations for virtualizing enterprise applications. It is inclusive of the configuration servers, networking, and storage, the virtual infrastructure, including ESXi hypervisor and vCenter cluster configuration, and the ScaleIO software configuration. A high level architecture diagram of the key components of the solution is shown below.

Note: Dell has tested the full end to end solution to validate the architecture. Customers should be aware that architecture components are modular and flexible.

The Management pod is an optional component, but management via vCenter with the ScaleIO plugin is required. If a customer has an existing vCenter environment it can be utilized for management. The Dell Management pod is built upon the principle that the production and management clusters should be physically separated to minimize the risk of downtime and production impact.

For Networking, the 10GbE and 1GbE production and management network is required. Dell has tested the solution with the Dell Networking S4048-ON and S3048-ON to validate the architecture.

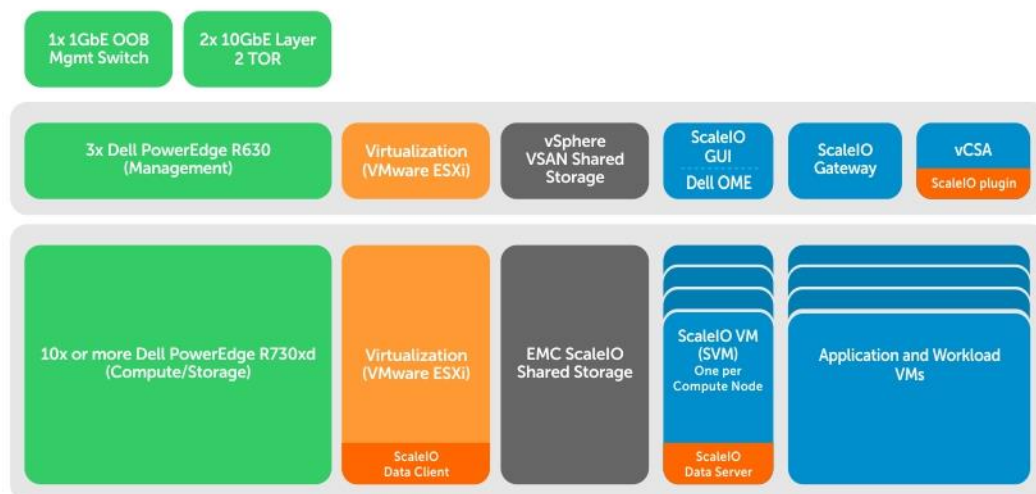


Figure 1 Solution Overview

Dell and EMC ScaleIO will support three configurations with additional configurations to be added over time. These configurations have been tested and validated within the Dell labs running at scale. Dell recommends the following three validated configurations: Mixed option 1 (6x 1.6TB SSD and 18x 1.2TB SAS HDD), Mixed option 2 (6x 800GB SSD and 18x 1.2TB SAS HDD) and Standard (24x 1.2TB SAS HDD).

Table 1 Compute/Storage Nodes

<b>ScaleIO Compute/Storage Node Configuration (10 x or more)</b>	<b>Mixed Configuration (PowerEdge R730XD)</b>	<b>Standard Configuration (PowerEdge R730XD)</b>
CPU	Dual-socket Intel 20 core CPU	Dual-socket Intel 20 core CPU
Memory	512GB	512GB
Controller	PERC H730	PERC H730
Storage	Option 1 <ul style="list-style-type: none"> <li>• 6 x 1.6TB SSD</li> <li>• 18 x 1.2TB HDD</li> </ul> Option 2 <ul style="list-style-type: none"> <li>• 6 x 800GB SSD</li> <li>• 18 x 1.2TB HDD</li> </ul>	24 x 1.2TB HDD
Network	4 x 10GbE	4 x 10GbE

Table 2 Networking (Optional)

<b>Hardware</b>	<b>Specification</b>
Network Fabric	2 x Dell S4048-ON
OOB Management	1 x Dell S3048-ON

Table 3 Management Nodes (Optional)

<b>Management Nodes (3 x or more)</b>	<b>PowerEdge R630</b>
CPU	Single-socket Intel 14 core CPU
Memory	128GB
Controller	PERC H730
Storage	Cache Tier: 1 x 400GB SSD Capacity Tier: 3 x 1.2TB HDD
Network	4 x 10GbE





## 6.1 Dell PowerEdge R730XD

The Dell PowerEdge Portfolio offers a variety of form factors and density options. For customers looking to standardize on a 2U server platform the R730xd is the perfect building block for SDS. The Dell PowerEdge R730xd server based on Intel Xeon technology provides a flexible building block. With a variety of configuration options, the R730xd allows customers to utilize a single platform for multiple ScaleIO configurations. The PowerEdge R730xd is an exceptionally flexible and scalable two-socket 2U rack server that delivers high-performance processing and a broad range of workload-optimized storage. The PowerEdge R730xd offers the latest Intel® Xeon® processor E5 2600 v4 product family, 24 DIMMs of high-performance DDR4 memory, and a broad range of local storage options including SSDs.

## 6.2 Dell Switches

The Dell Networking S4048-ON is a 1U high-density 10/40 GbE ToR switch with 48 dual-speed 1/10GbE (SFP+) ports and six 40 GbE (QSFP+) uplinks. This switch 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.

## 6.3 Dell PowerEdge R630

Dell PowerEdge R630 is a rack server with 2U performance in a compact 1U chassis. The R630 is designed for data centers that require high efficiency and density in a small footprint. The PowerEdge R630 two-socket rack server delivers uncompromising density and productivity. The processor and memory density as well as storage flexibility provide an excellent virtualization infrastructure for management of the Dell VS for the EMC ScaleIO solution.



## 7 Overview of Software-Defined Storage with EMC ScaleIO

EMC ScaleIO is a software-defined storage solution that converges storage and compute resources to form an enterprise class storage product. ScaleIO creates a storage environment that is capable of delivering exceptionally high levels of performance and scalability using commodity servers. It utilizes agnostic resources that include drive types, servers and their components, as well as various virtualization solutions.

ScaleIO software features software components that have a small footprint and are highly available. The following sections discuss ScaleIO software components and core constructs.

Note: The following sections provide a high level overview of the ScaleIO architecture. More in-depth documentation and configuration guidance are included in the ScaleIO software bundle.

### 7.1.1 EMC ScaleIO Software Components

The EMC ScaleIO virtual SAN consists of the following software components:

- **ScaleIO Virtual Machine (SVM)**  
The ScaleIO Virtual Machine is a Linux-based virtual machine hosting some of the ScaleIO software components listed below. Each vSphere host can have a dedicated SVM. The local storage on each vSphere host is directly attached to the SVM, where additional services can present the devices to the ScaleIO environment.
- **MetaData Manager (MDM)**  
The MetaData Manager configures and monitors the ScaleIO system. MDM performs ScaleIO migrations, rebuilds, and other system-related functions. The MDM is configured in a redundant Cluster Mode. To provide redundancy, an MDM cluster is comprised of Master, Slave, and Tiebreaker MDMs. The Master MDM updates the database that contains the SDS configuration and data distribution and constantly replicates to the Slave MDMs so that they can take over instantly when there is a Master MDM failure.
- **ScaleIO Data Server (SDS)**  
The ScaleIO Data Server manages the local storage devices of a single server and acts as a back-end for data access. The SDS service is installed on all servers contributing storage devices to the ScaleIO environment.
- **ScaleIO Data Client (SDC)**  
The ScaleIO Data Client is a small footprint device driver that exposes ScaleIO volumes as block devices to the application that resides on the same server where the SDC is installed. In the case of VMware deployments, the SDC driver runs within the ESXi OS and provides the hypervisor with access to the ScaleIO volumes as VMFS datastores.

Depending on the desired configuration, the software components are installed on the server node and provides a virtual SAN layer exposed to the applications residing on the servers. In the VMware vSphere environment, the MDM and SDS components are installed on the SVM, whereas the SDC is installed directly on the ESXi host as shown in Figure 2 below.



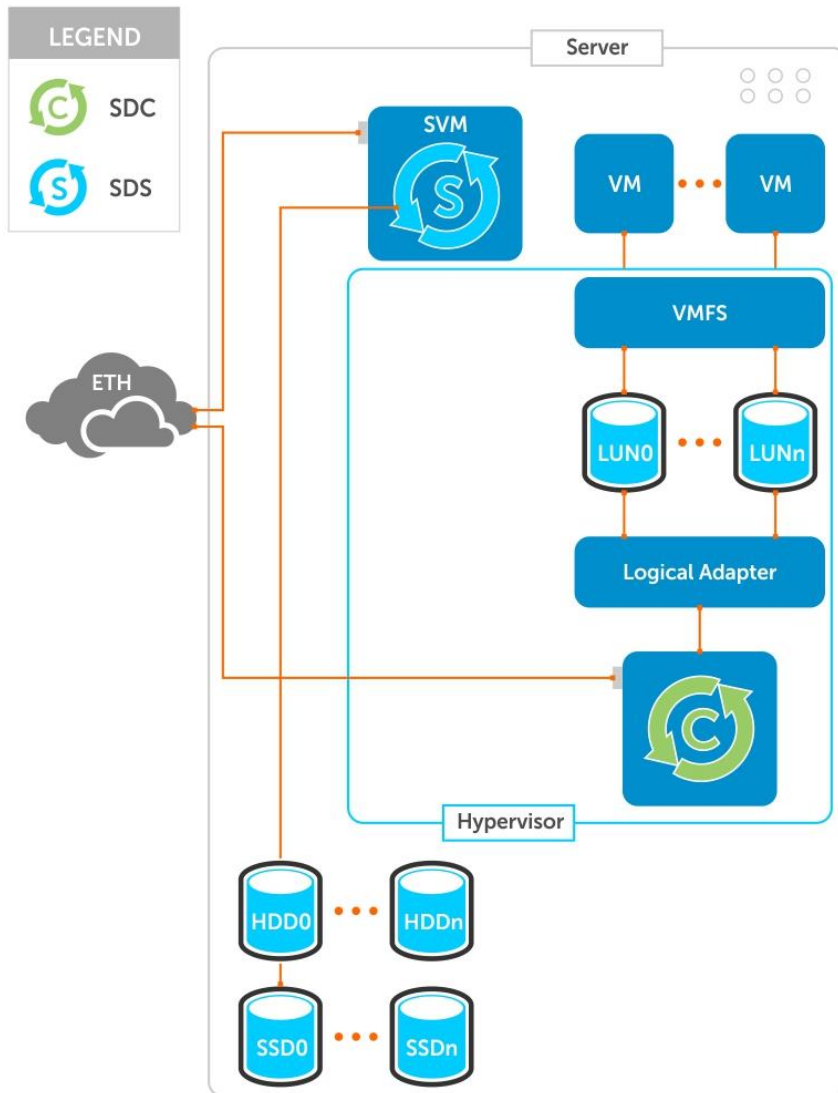


Figure 2 EMC ScaleIO ESX Implementation

## 7.1.2 EMC ScaleIO Storage Definitions

The ScaleIO system links the physical layer with the virtualization layer using:

- Protection Domains
- Storage Pools
- Fault Sets

Each of these components are described in the following sections.

### 7.1.2.1 Protection Domains

A Protection Domain contains a group of SDSs that provide backup for each other. Each SDS belongs to only one Protection Domain, and therefore by definition, each Protection Domain is a unique set of SDSs.

The recommended number of nodes in a Protection Domain is 16-32, in order to:

- Enable optimal performance
- Reduce theoretical mean time between failure issues
- Enable the system to sustain multiple failures in different Protection Domains.

Protection Domains can be defined during installation. They can also be modified post-installation.

### 7.1.2.2 Storage Pools

A Storage Pool is a set of physical storage devices in a Protection Domain. Each storage device belongs to only one Storage Pool, and therefore, Storage Pools allow the creation of storage tiers.

When a volume is created in a ScaleIO system, it is distributed to all devices within a Storage Pool. Each volume block has two copies located on two different SDSs which allows the system to maintain data availability when there is a single device failure. The data availability can be maintained even when there are multiple failures, as long as each failure took place in a different storage pool.

To provide consistent performance, Dell recommends that all devices in the Storage Pool have similar storage properties. For example, if all the SDSs in the Protection Domain have two physical drives associated with them, one HDD and one SSD, then two Storage Pools should be defined:

- A Capacity Storage Pool consisting of all HDDs in the Protection Domain, and
- A Performance Storage Pool consisting of all SSDs in the Protection Domain.

A ScaleIO Storage Pool should have no more than 300 storage devices. Storage Pools can be defined during installation. They can also be modified post-installation.

### 7.1.2.3 Fault Sets

A Fault Set contains a group of SDSs within a Protection Domain. Each group should be composed of SDSs that have a higher chance of going down together, for example SDSs that are located in the same rack. By grouping them into a Fault Set, data mirroring for all devices in the Fault Set will take place on SDSs that are outside of the Fault Set.

Fault Sets are also referred to as fault units, where a fault unit can be either a Fault Set, or an SDS not associated with a Fault Set. A fault unit can be thought of as a Fault Set of a single SDS. To enable mirroring, there must be enough capacity within at least 3 fault units.

Fault Sets can be defined during installation and SDSs can be added to them. Fault Sets must be defined and configured before adding SDSs to the system. Fault Sets can be defined using the following management tools:

- Installation manager

- CLI
- REST
- Plug-in.

Fault Sets can also be defined when adding SDS nodes after installation.



## 8 Infrastructure Overview

As shown in [Figure 1](#), Solution Architecture, the solution includes a minimum of ten (10) R730xd servers providing compute and software-defined storage via ScaleIO. Each of the R730xd servers run VMware ESXi with the ScaleIO SDC (ScaleIO Data Client) driver/VIB installed. A ScaleIO SVM is deployed on each R730xd server running SDS (ScaleIO Data Server) service. Application and workload VMs can be spawned in ScaleIO-provided, software-defined, shared storage that scales linearly as you add more R730xd compute/storage nodes.

The management infrastructure runs on three (3x) R630 servers providing compute and software-defined storage via VMware VSAN. This infrastructure runs the management software stack comprising VMware vCenter with ScaleIO Plugin, ScaleIO GUI, ScaleIO Gateway, and Dell Open Manage Essentials. This is an optional component of the solution. Customers also have an option to run these software components in their existing management infrastructure.

The following sections describe the compute & management server, as well as networking and software components that implement the full ScaleIO solution.

### 8.1 Compute Infrastructure

Dell PowerEdge R730xd is a 2U1N rack server that supports up to two Intel Xeon E5-26xx v4 processor series. The R730xd can support up to 44x physical CPU cores, up to 768GB of physical RAM, and up to 24x 2.5 drives per node, making it a perfect choice for the storage-dense hyper-converged platform.

Along with support for the new Intel Xeon E5-26xx v4 processors, Dell has introduced availability of a new R730xd configuration in which the rear backplane boot drives are connected to the onboard AHCI controller. This configuration provides a dedicated storage controller for boot and persistent storage for the EMC ScaleIO management components. Therefore the PERC H730 storage controller can be dedicated to providing storage devices to the ScaleIO environment, without interruption from the boot device. Various HDD and SSD configurations are available, providing up to 31TB of raw storage capacity to the ScaleIO environment. For more information on the different configuration options available, see the [Configuration Guidance](#) chapter.

This validated system describes a minimum of ten (10) and up to sixteen (16) R730xd servers in a single 42U rack.

#### 8.1.1 EMC ScaleIO

ScaleIO SDC is installed on each ESXi hypervisor along with the ScaleIO SVM running locally on the ESXi boot device. The SVM runs the SDS service contributing the local R730xd storage devices to the ScaleIO System. Several of the SVMs will also run the Master MDM, Slave MDM, or Tiebreaker service to handle the ScaleIO system related functions and provide failure resiliency. Protection Domains and Fault Sets are used to group SDSs into logical entities that are failure resistant. A Storage Pool will be required for each distinct storage device type, such as 10k SAS HDD devices or unique SSD models.

## 8.2 Management Infrastructure

Dell PowerEdge R630 is a rack server with 2U performance in a compact 1U chassis. The R630 is designed for data centers that require high efficiency and density in a small footprint. The PowerEdge R630 two-socket rack server delivers uncompromising density and productivity. The processor and memory density as well as storage flexibility provide an excellent virtualization infrastructure for management of the EMC ScaleIO Solution. For more information, see the [Configuration Guidance](#) chapter.

Virtual SAN requires sole access to the local disks of the host. Virtual SAN disks are not shared with another file system, such as Virtual Flash File System (VFFS), VMFS partitions, or an ESXi boot partition. R630 servers are not configured for a rear backplane boot drive via onboard AHCI controller. For this reason, ESXi is installed on the IDSDM device (Internal Dual SD Module) on R630 servers.

To comply with VMware guidelines, you must configure VSAN traces and an ESXi persistent scratch location correctly:

- If you install [ESXi on a USB or SD device](#) and allocate all local storage to Virtual SAN, you do not have any local disk or datastore available for persistent logging. Configure a [Dump Collector](#) and a [Syslog Collector](#) to direct ESXi memory dumps and system logs to a server on the network, rather than to a local disk.
- You must also configure a [persistent scratch location](#).

### 8.2.1 VMware vCenter Server

This Validated System is based on VMware vSphere 6 with ESXi and vCenter Server as the virtualization management platform. 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. This topic describes the VMware vCenter Server which is a central management point for the hypervisors and VMs.

vCenter Server is used to manage the ESXi hosts and the logical vSphere clusters. vCenter Server provides centralized control and visibility through:

- **vSphere web client:** Manages the essential functions of vSphere from any browser.
- **vCenter single sign-on:** Allows users to log in once and access all instances of vCenter Server, without the need for further authentication. VMware Single Sign-On (SSO) Service integrates multiple identity sources including Active Directory, Open LDAP, and local accounts for authentication.
- **Custom roles and permissions:** Restricts access to the entire inventory of virtual machines, resource pools, and servers by assigning users to custom roles.
- **Inventory search:** Explores the entire vCenter inventory, including virtual machines, hosts, datastores and networks, from anywhere within vCenter.



Dell VS for EMC ScaleIO leverages the VMware vCenter Server Virtual Appliance (vCSA) to manage the virtual infrastructure. vCSA provides the following functionality:

- Cloning of VMs
- Template creation
- VMware vMotion and VMware Storage vMotion
- Initial configuration of VMware Distributed Resource Scheduler (DRS) and VMware vSphere high-availability clusters.

vCSA also provides monitoring and alerting capabilities for hosts and VMs. System administrators can create and apply alarms to all managed objects in VMware vCenter Server, including:

- Data center, cluster, and host health, inventory, and performance
- Data store health and capacity
- VM usage, performance, and health
- Virtual network usage and health.

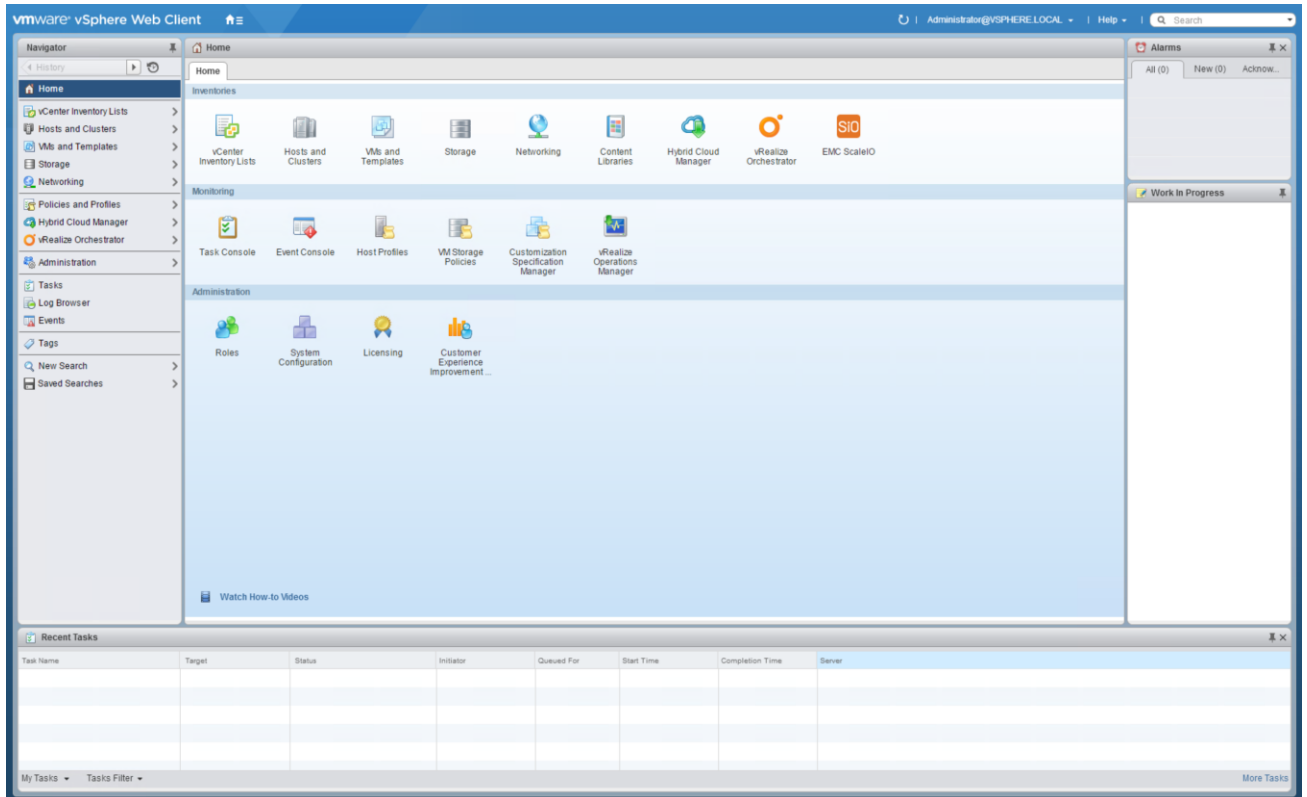


Figure 3 VMware vSphere Web Client

## 8.2.2 EMC ScaleIO GUI

The ScaleIO Graphical User Interface (GUI) makes it possible to review the overall status of the system and perform system management functions such as adding new storage devices or clearing device errors. There are various screens that display different views and data that are all designed to facilitate the ScaleIO system management. The same look-and-feel are maintained throughout the different screens and the screens are organized in a tree-like structure. The tree can be navigated to drill down and attain a more granular view of an element, when available.

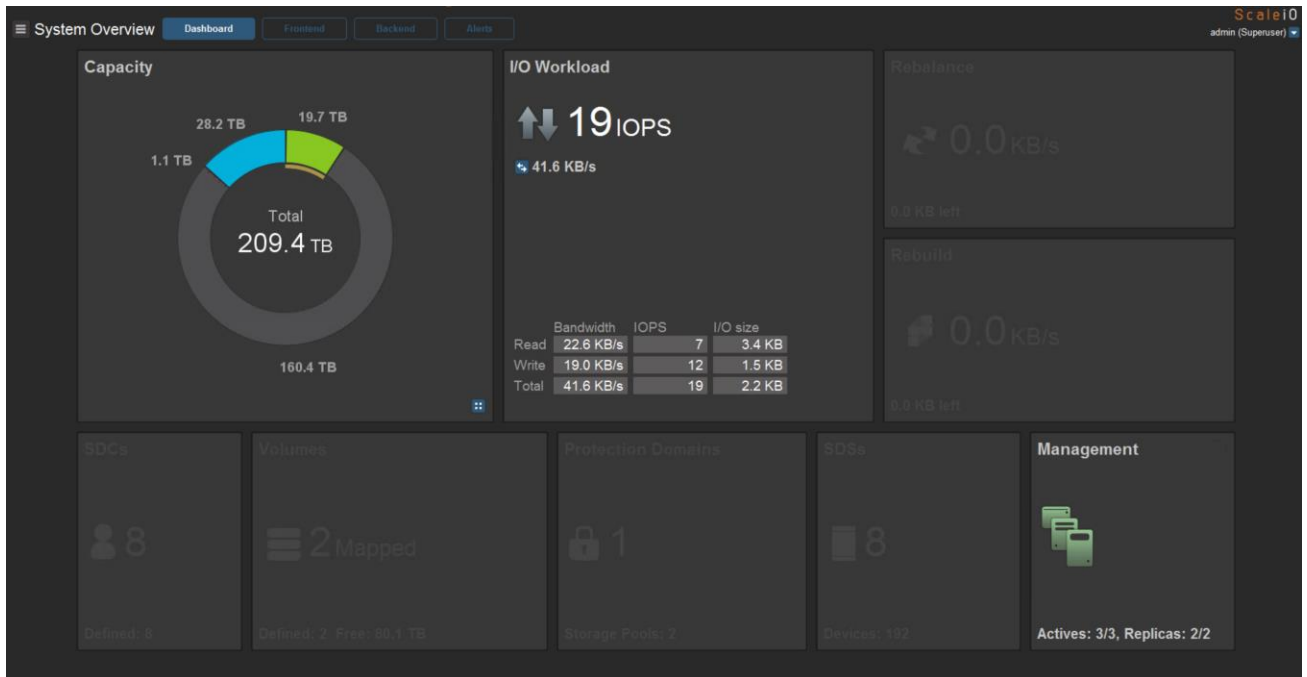


Figure 4 EMC ScaleIO GUI

## 8.2.3 EMC ScaleIO Gateway

The ScaleIO Gateway is installed as a ScaleIO VM and consists of the following:

- **Installation Manager**  
The Installation Manager is used for installing ScaleIO, upgrading and installing components, and running the ScaleIO *get-info* operation.
- **REST Gateway**  
The REST Gateway consists of a REST API that can be used for monitoring and provisioning via the REST interface.
- **SNMP trap sender**  
The SNMP trap sender is disabled by default. When enabled, it can forward alerts to a SNMP management device.

## 8.2.4 EMC ScaleIO Plugin

The ScaleIO plugin for vSphere simplifies the installation and management of the ScaleIO system in an ESXi environment. The plugin can be used for the initial ScaleIO system configuration as well as for adding additional SDS nodes.

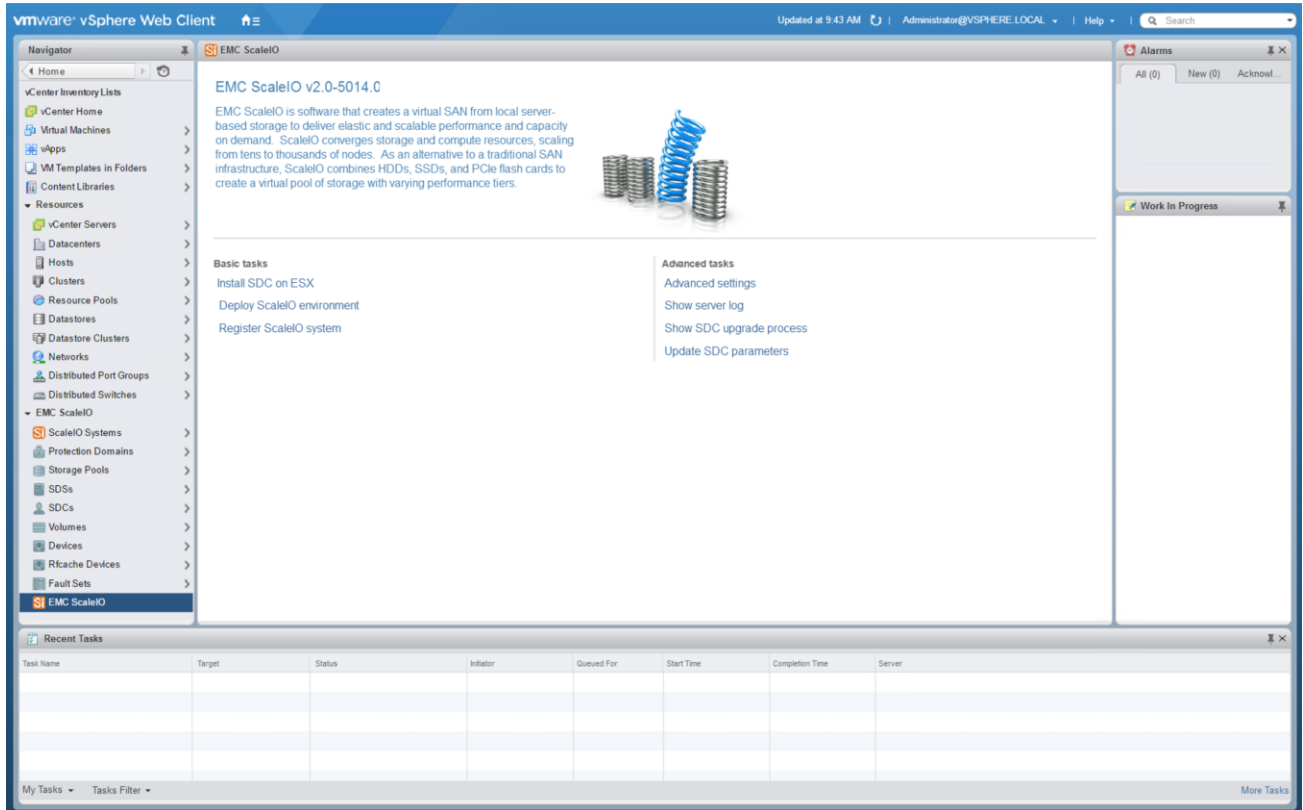


Figure 5 VMware vSphere Web Client with EMC ScaleIO Plugin Tasks

## 8.2.5 OpenManage Essentials

OpenManage Essentials (OME) is a hardware management application that provides a comprehensive view of Dell systems, devices, and components in the enterprise's network. With OpenManage Essentials, a web-based and one-to-many systems management application for Dell systems and other devices, you can:

- Discover and inventory the system
- Monitor the health of the system
- View and manage system alerts
- Perform system updates and remote tasks
- View hardware inventory and compliance reports
- Deploy or re-provision a server or chassis
- Manage the configuration baseline of a server or chassis
- Manage the virtual I/O identity of a server



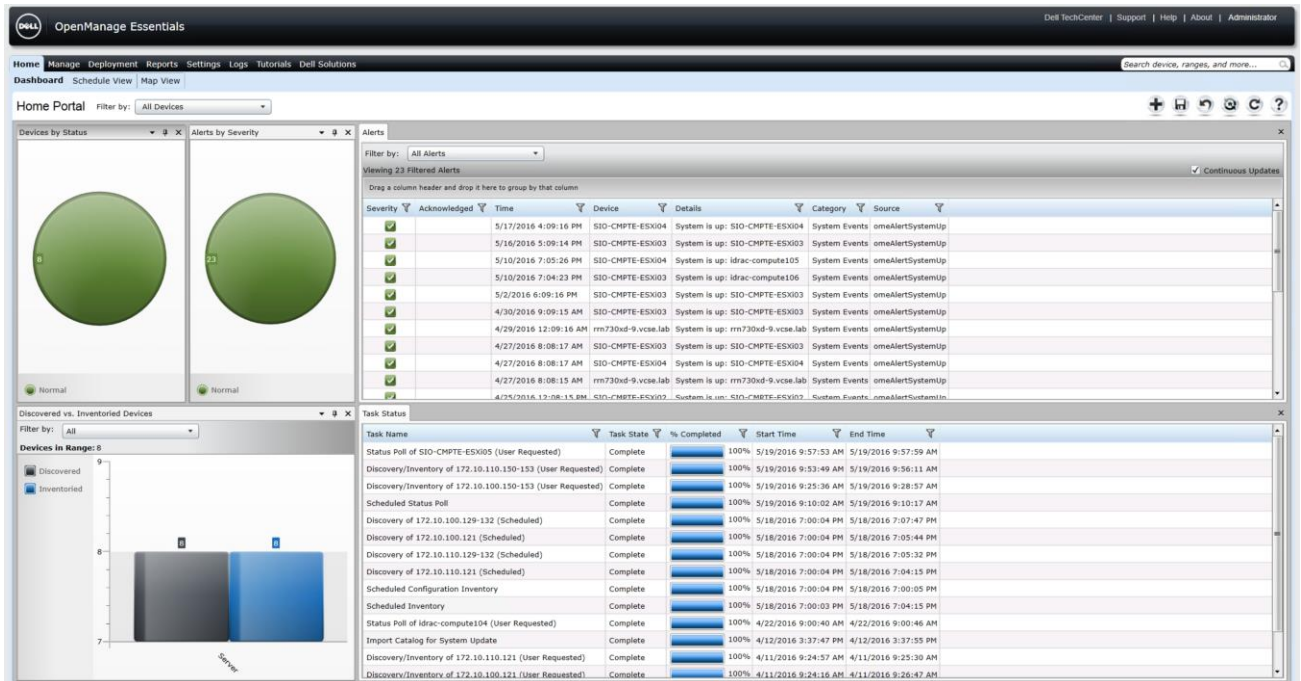


Figure 6 Dell OpenManage Essentials Dashboard

## 8.3 Network Infrastructure

The following illustration provides a high-level overview of the physical network components in the architecture.

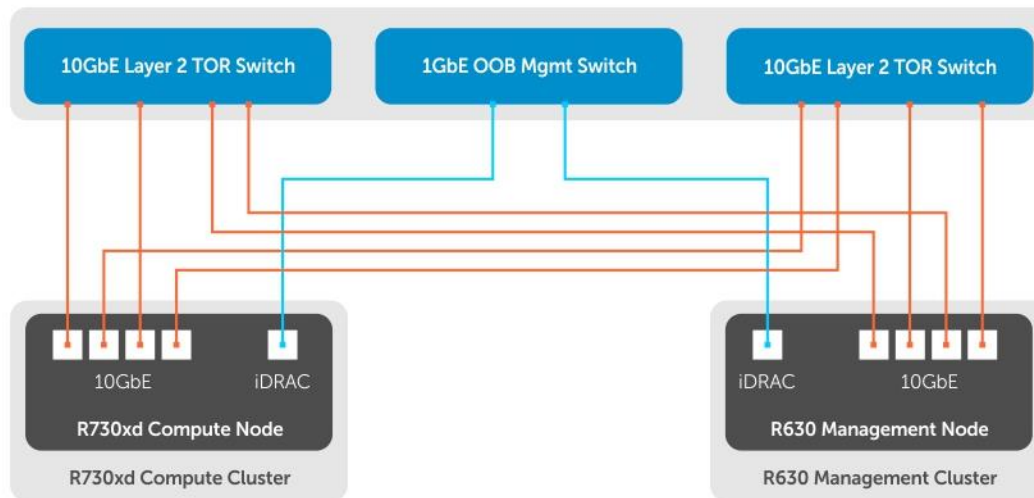


Figure 7 Physical Network Diagram

Each compute and management node is configured with five networks:

- One 1GbE network for OOB management traffic connected to OOB management switch, and
- Four 10GbE networks dedicated for data traffic connected to two layer 2 TOR switches in redundant fashion.

NOTE: Each node comes with one Intel x710 NDC. Connect each port from these NICs to separate TOR switch as shown in [Figure 7](#).

You can use any of the Layer 2 switches that meet following requirements:

- Top of Rack (TOR) Network requirements
  - Layer 2
  - 10GbE
  - IGMP Snooping and Querier enabled for the VSAN management cluster

NOTE: See [EMC ScaleIO Networking Best Practices Guide](#).

- OOB Management Network requirements
  - Layer 2
  - 1GbE

The following illustration provides an overview of the Management Node Virtual Switch configuration.

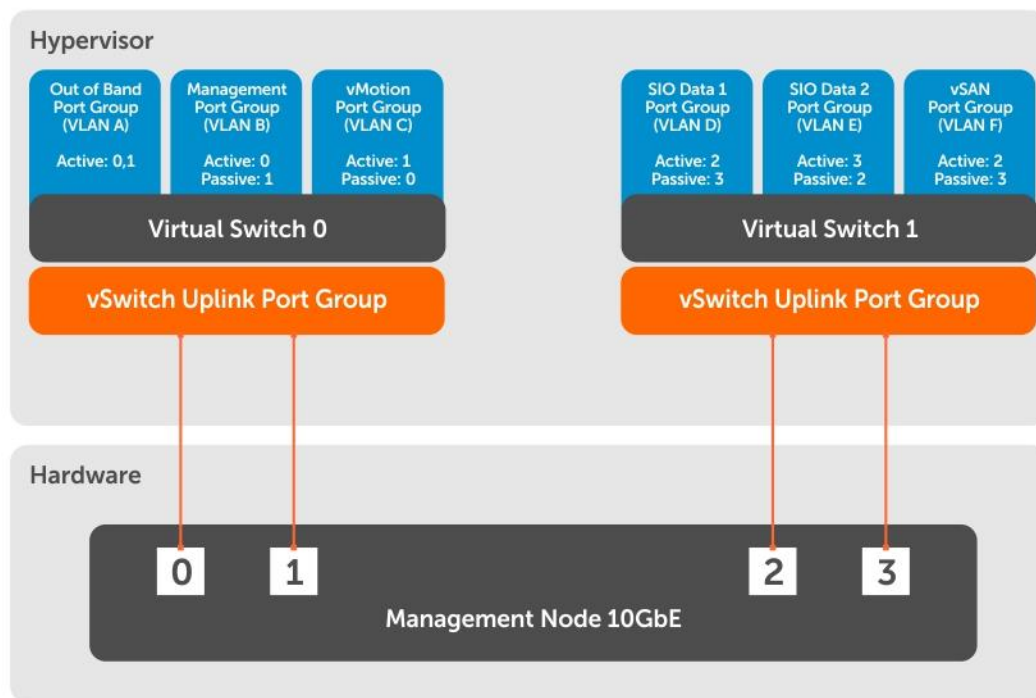


Figure 8 Management Node: Virtual Switch configuration

The R630 management nodes will be configured to run on the vSphere VSAN cluster. Additionally it will have access to the ScaleIO compute cluster providing access to ScaleIO GUI and Gateway. The Out of Band (OOB) portgroup is created to give Dell OME access to out of band management network.

Two VSSs (vSphere Standard Switch) will be configured to create the following portgroups:

Table 4 Virtual Standard Switch 0

Portgroup	VLAN ID	Active Port	Passive Port
Out of Band	A	0 and 1	None
Management	B	0	1
vMotion	C	1	0

Table 5 Virtual Standard Switch 1

Portgroup	VLAN ID	Active Port	Passive Port
ScaleIO Data 1	D	2	3
ScaleIO Data 2	E	3	2
VSAN	F	2	3

The following illustration provides an overview of the compute node Distributed Virtual Switch configuration.

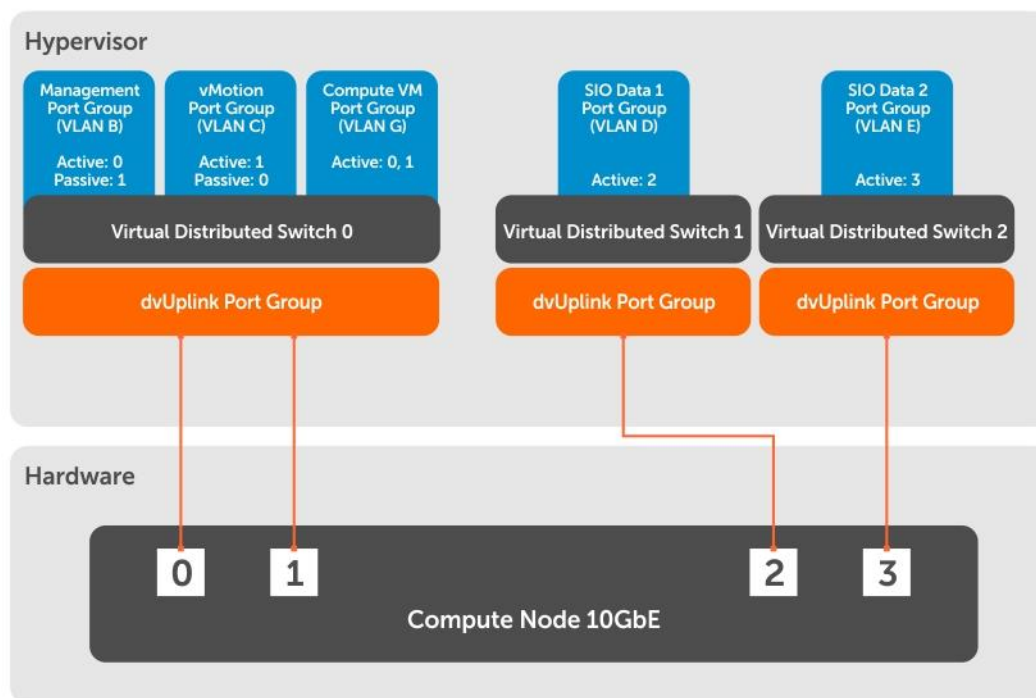


Figure 9 Compute Node: Distributed Virtual Switch configuration

The R730xd compute/storage nodes will be configured to run the ScaleIO cluster. Two dedicated portgroups are created to run ScaleIO traffic.

Three VDSs (vSphere Distributed Switch) will be configured to create following portgroups:

Table 6 Virtual Distributed Switch 0

Portgroup	VLAN ID	Active Port	Passive Port
Management	B	0	1
vMotion	C	1	0
VM	G	0 and 1	None

Table 7 Virtual Distributed Switch 1 and 2

Portgroup	VLAN ID	Active Port	Passive Port
ScaleIO Data 1	D	2	None
ScaleIO Data 2	E	3	None



## 9 Configuration Guidance

The Dell VS for EMC ScaleIO contains the following hardware and software components:

Table 8 Hardware and Software Components

Resource	Components
Compute	<ul style="list-style-type: none"><li>• Mixed Configuration</li><li>• Standard Configuration</li></ul>
Network	<ul style="list-style-type: none"><li>• 1 x 1GbE Management Switch</li><li>• 2 x 10GbE Layer2 TOR Switches</li></ul>
Storage	EMC ScaleIO Virtual SAN
Management	<ul style="list-style-type: none"><li>• VMware vSphere ESXi</li><li>• VMware VSAN</li><li>• VMware vCenter Server with EMC ScaleIO vSphere Plugin</li><li>• Dell OpenManage Essentials</li><li>• EMC ScaleIO GUI</li><li>• EMC ScaleIO Gateway</li></ul>

The Dell Support Matrix provides a list of the validated versions of Software, Firmware, and BIOS in [component specifications](#) chapter.

### 9.1 Compute Node Configurations

Currently, there are two supported configurations, Standard and Mixed. Each node should be configured as a ScaleIO Data Server (SDS) and a ScaleIO Data Client (SDC). This allows each node to contribute local storage to the ScaleIO virtual SAN and have access to the volumes carved out from the virtual SAN. Single 42 RU rack can accommodate up to 16 compute nodes. A ScaleIO Storage Pool can have a maximum of 300 storage devices and is composed of SSD or HDD devices of the same type. The image below is a logical representation of the compute cluster.

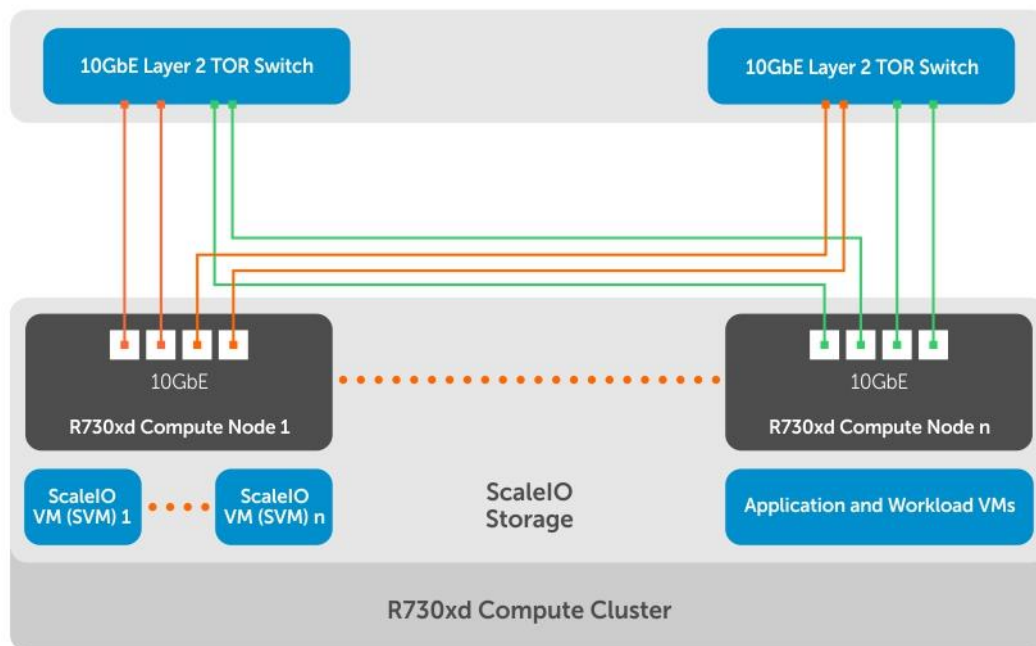


Figure 10 EMC ScaleIO Compute Cluster

### 9.1.1 Mixed Configuration

The Mixed configuration has two storage pools. The SSD storage pool is composed of six 800GB or 1.6TB SSDs that can provide ultra-high performance for IOPS intensive applications. The HDD storage pool with eighteen 1.2TB 10k SAS HDDs can accommodate applications that require large capacity storage space.




Figure 11 Mixed configuration

### 9.1.2 Standard Configuration

The Standard configuration is an all-capacity drive configuration that is suitable for applications that require large amounts of storage without the need for ultra-high performance. There are two storage pools with 12 1.2TB 10k SAS HDDs each. Multiple storage pools are needed in some cases to prevent the ScaleIO environment from exceeding the storage pool device limit.

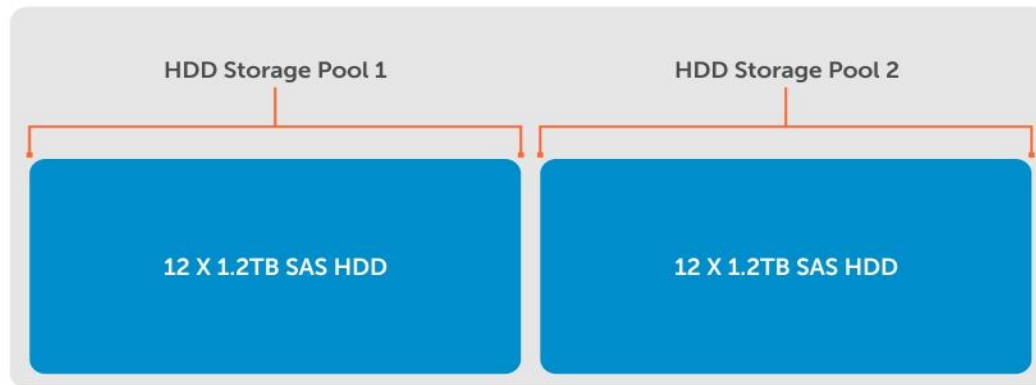


Figure 12 Standard configuration

## 9.2 Management Node Configuration

The recommended configuration for management nodes running VMware VSAN is shown below. Customers are free to use existing management infrastructure to manage this configuration. The management components run on three Dell R630 servers configured in the vSphere VSAN cluster. They connect to layer 2 TOR in redundant fashion as shown in the figure above. This highly available vSphere management cluster provides a virtualization space to run the vCenter Server Appliance with ScaleIO plugin, ScaleIO GUI and Gateway, and Dell OpenManage Essentials for monitoring the management as well as the ScaleIO compute cluster.



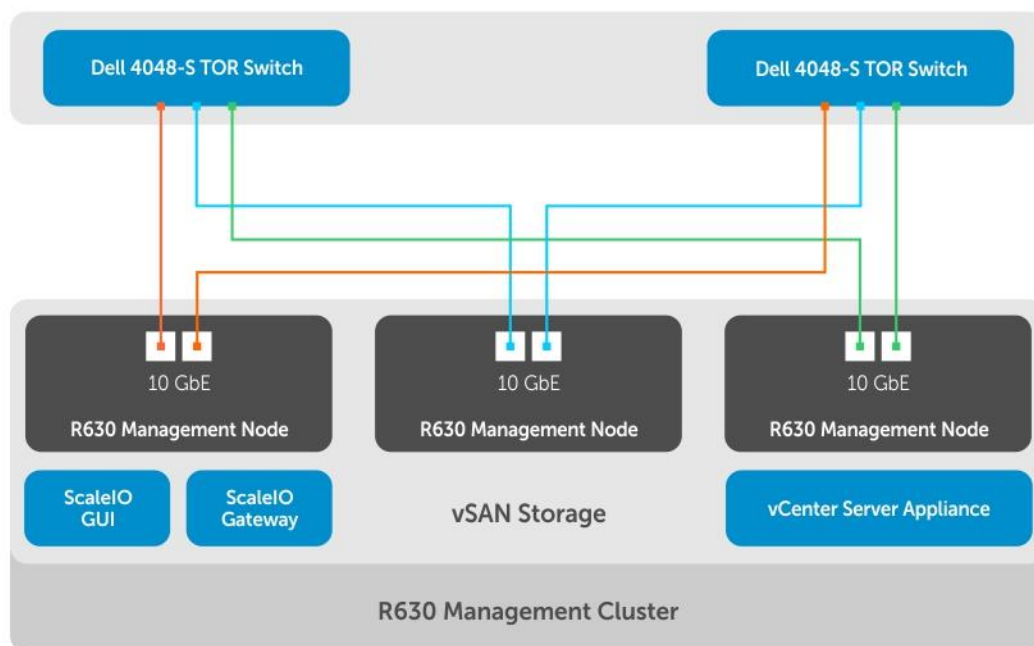


Figure 13 Management node configuration

## 9.3 Base configurations and scaling

Dell VS for EMC ScaleIO has a base configuration that is a minimum set of compute and management components, as well as fixed network resources. These components are integrated within one or more standard 42U cabinets.

Table 9 Base configuration

Hardware	Minimum set
Compute	Ten (10) 2U servers - Dell R730xd servers
Network	<ul style="list-style-type: none"> <li>• Dell Switches S4048-ON/S3048-ON (optional)</li> <li>• Any 10GbE layer 2 TOR switches (two for redundancy)</li> <li>• Any 1GbE layer 2 OOB switch</li> </ul>
Management (Optional)	<p>Three (3) 1U servers with HA support and VSAN for software scale-out running on Dell R630.</p> <p>NOTE: Customers can choose to manage the ScaleIO cluster from an existing vCenter environment.</p>

This validated system, based on EMC ScaleIO software-defined storage, is a scale-out solution that defines a maximum of 16 Dell R730xd compute nodes with various CPU, memory, and drive options.

## 9.4 Best Practices

This Validated System has been tested and validated by Dell following the best practices below.

### 9.4.1 Compute

Recommendations for selecting a CPU for Compute Nodes are as follows:

- The Intel E5-2600 v4 family of CPUs is designed for the scalable performance demands of complex, data-demanding workloads that are typical in a software defined storage virtualization environment. It supports up to 2.4x more VMs/servers and reduces operational expenses and cost per VM by up to 58%, maximizing profitability versus two-generations-old servers (E5 v2). The high storage density of the R730xd necessitates a large CPU compute capacity to fully utilize the storage space provided by the ScaleIO environment.
- The Dual Intel E5-2698 v4 CPU configuration was chosen for its high virtualization performance and core count. The 40 physical cores per server node, 2.2GHz speed, and full array of virtualization features provide the ability to scale the solution to a large number of virtual machines, on a relatively small number of R730xd compute hosts.

### 9.4.2 Storage

Recommendations for selecting storage dense Compute Nodes are as follows:

- The R730xd server provides a storage platform ideal for software defined storage clusters with the 24 x hot-swap hard drive slots in a 2 rack-unit form factor. The ScaleIO storage engine provides a resilient design that allows fault tolerance but still optimizes performance by distributing the data blocks (1MB) throughout the cluster randomly and evenly. This performance benefit is enhanced by the large storage device capacity of the R730xd platform, since the storage load can be evenly distributed over a large number of devices and cluster nodes. Each node that is added to the cluster will significantly increase the total number of storage devices, and the ScaleIO engine takes advantage by linearly increasing the storage IOPS.
- Virtualization workloads will generally require a small number of extremely high performance SSD devices and a large number of slower HDD devices to meet each unique storage performance profile requirement. ScaleIO storage pools allow the generation of different storage tiers that can be easily leveraged by the vSphere environment by placing virtual machine data on the appropriate storage performance profile. vSphere Storage vMotion features can be utilized to easily migrate data to different storage tiers as the needs arise. All of these features combine to provide a flexible scale-out approach to storage, with high performance, while eliminating performance bottlenecks and scalability limitations.

### 9.4.3 Network

Recommendations for selecting network interface cards for Compute Nodes are as follows:

- Hyper-converged virtualization solutions push the scale of virtualization to a new limit. Higher virtualization results in a higher utilization of available network bandwidth. A hyper-converged data

center demands that the network must be architected to offer a resilient framework that handles link, component, or device failures without impacting the application.

- Each R730xd has been configured with 4 x 10GbE NIC ports. This design is intended to separate LAN and VSAN traffic, while providing ample throughput for various workloads and storage traffic patterns. The 20GbE of LAN network throughput should prevent interaction of VM, management and vMotion traffic from causing any network disruption. However, the use of additional vSphere Network IO Control (NIOC) on the virtual distributed switch could be implemented as needed. The 20GbE of dedicated storage network throughput is easily leveraged by the ScaleIO engine. Multiple, high performance SSD drives in a ScaleIO environment have the capability to drive the throughput up to 2.5GBs per compute node. The dual 10GbE adapters are configured on redundant fabrics, providing faster rebuild time, rebalance time, and drive performance, as well as the lowest latency possible.
- Dell Networking provides a resilient, high-performance architecture that improves availability and helps meet Service Level Agreements (SLAs) more effectively. The Dell Networking S4048-ON TOR switch was selected for its ability to provide a low latency, non-blocking layer-2 network architecture. The 48 x 10GbE and 6 x 40GbE ports in a single rack unit provides flexibility in the datacenter for 1/10/40GbE uplink compatibility.

#### 9.4.4 Storage Pool

Recommendations for EMC ScaleIO storage pool configuration are as follows:

- The massive number of storage devices in a ScaleIO cluster requires careful consideration. ScaleIO 2.0 recommends no more than 300 devices in a single storage pool. For the Dell Validated System for ScaleIO, we have considered this limit by separating each node configuration into two storage pools, while still providing enough quantity of devices per storage pool to provide very high performance. Additional storage pools could be created across the cluster to isolate workloads that specific storage performance requires.

#### 9.4.5 Network Scaleout

Recommendations for network configuration for solution scaling beyond a single rack are as follows:

- A pair of S4048-ON switches support the creation of a VLT domain, which is critical to the creation of a hyper-converged network. VLT is a Dual Active Control Plane implementation, of a Multi System LAG, similar to the Cisco vPC technology, which allows the possibility of maintenance without taking the entire fabric down.
- Hyper-converged virtualization solutions push the scale of virtualization to a new limit. Higher virtualization results in a higher utilization of available network bandwidth. A hyper-converged data center demands that the network must be architected to offer a resilient framework that handles link, component, or device failures without impacting the application.
- Dell Networking provides a resilient, high-performance architecture that improves availability and helps meet Service Level Agreements (SLAs) more effectively. The Dell Networking S4048-ON TOR switch was selected for its ability to provide a low latency, non-blocking layer-2 network



architecture. The 48 x 10GbE and 6 x 40GbE ports in a single rack unit provide flexibility in the datacenter for 1/10/40GbE uplink compatibility.

- The spine and leaf architecture form a baseline for all modern data center networks primarily due to its high-availability and horizontal scaling features. It is characterized by fixed-form switching that packs high-speed interconnects with low latency and high throughput. The east-west (server-to-server/node-to-node) traffic gets full bi-sectional bandwidth over equal-cost links.
- The network is normally a mix of layer 2 switching extending to the Top-of-Rack (TOR) and layer 3 routing extending beyond that. There are however, technologies that provide layer 2 extensions between servers in the same cluster and adjacent clusters. In the same cluster, Virtual Link Trunking (VLT) is used to provide multi-switch link trunking, which can be extended all the way to the spine layer for legacy applications that need a larger layer 2 domain.




Figure 14 EMC ScaleIO scale-out networking 16 rack architecture

- With the EMC ScaleIO networking topology, the External Border Gateway Protocol (eBGP) is used between the various leaf and spine switches with Equal-Cost Multi-Path (ECMP) also implemented to provide for cross-pod connectivity. The sub-second convergence of ECMP on the Dell Networking Switches, coupled with active-active VLT links, ensures that the failure of any link will be resolved by balancing the traffic over the alternative paths.
- A pair of S4048-ON switches support the creation of a VLT domain at the top of the rack, which is critical to the creation of a hyper-converged network. VLT is a Dual Active Control Plane implementation, of a Multi System LAG, similar to the Cisco vPC technology, which allows the possibility of maintenance without taking the entire fabric down. See [Knowledge Base: How to set up Virtual Link Trunking \(VLT\) on Dell Networking Switches](#).

- The S4048-ON switches connect to Z9100-ON switches as the spine of the networking topology. The Z9100-ON switch adds multiline rate capability supporting 10GbE, 25GbE, 40GbE, 50GbE, and 100GbE. This switch combination provides for substantial growth with this initial configuration using 40GbE uplinks from the rack but is able to move to 25GbE downlinks and 100GbE uplinks in the future.
- The port count for a 16 rack datacenter pod falls within the boundaries of a single two-tier spine and leaf configuration shown in Figure 14, above. The pod consists of 1 rack of WAN edge connectivity, 1 rack of management connectivity/nodes, and 14 racks containing a total of 280 compute/storage nodes. With this allocation of resources the oversubscription would be 2.5:1.
- For further scale the following proof of concept scaling would be used. This would allow for 4 of the 16 rack pods connected together via an additional spine layer to grow the supported nodes to 1120 with the same oversubscription.

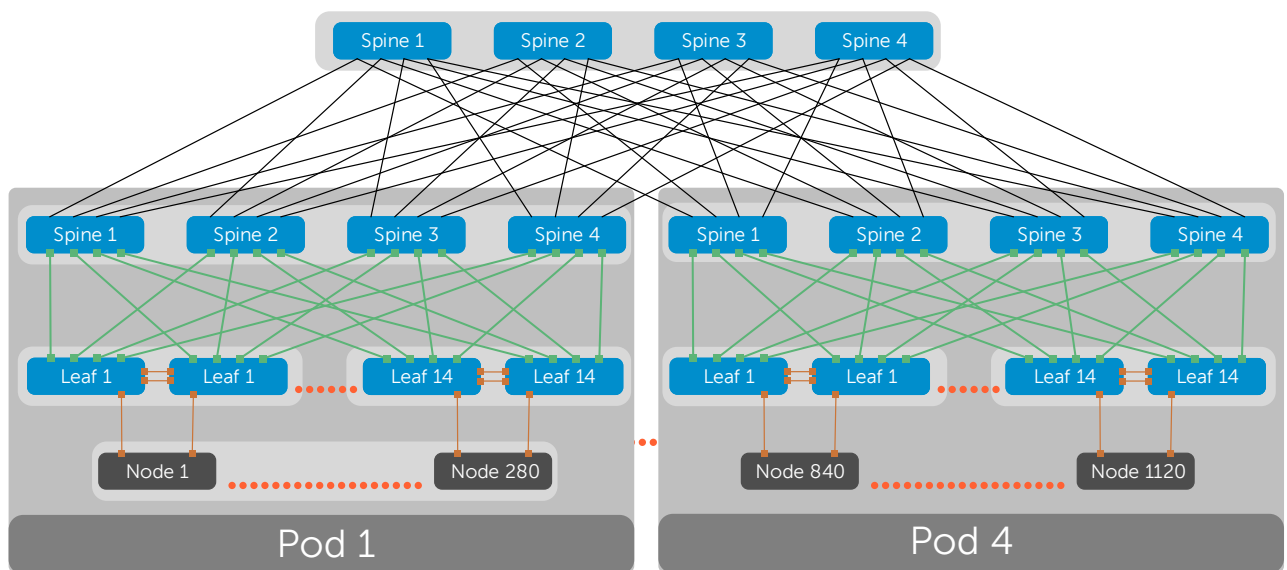


Figure 15 Hyperscale networking topology

## 10 Component Specifications

The following sections describe system infrastructure and component requirements.




Figure 16 Dell Validated System for EMC ScaleIO Rack Elevation

The above figure shows a full rack consisting of:

- Sixteen Dell R730xd servers to run ScaleIO cluster
- Three Dell R630 servers running management cluster based off VMware vSAN
- Two layer-2 10GbE TOR switches, and
- One layer-2 1GbE OOB Management switch.

## 10.1 Software Components

The tables below list the primary software components for the reference architecture of the Dell VS for EMC ScaleIO.

Table 10 Software Components

Software	Version	Notes
VMware vSphere ESXi	6.0 U1b	Lightweight hypervisor OS installed to the Dell server platforms that abstracts server hardware resources and makes them shareable by multiple virtual machines.
VMware vCenter Server Appliance (vCSA)	6.0 U1b	Provides centralized administration for the vSphere platforms and virtual machines.
EMC ScaleIO	2.0.0	Provides software-defined, distributed, shared storage consisting of protection domains, storage pools, and fault sets.
Dell OpenManage Essentials (OME)	2.1	Systems management console solution for hardware management and monitoring in multi-vendor hardware environments.

## 10.2 BIOS, Firmware and Drivers

The table below lists the primary BIOS, firmware and driver components for the Dell VS for EMC ScaleIO.

Table 11 BIOS, Firmware and Driver Components

Component	Version	Notes
PowerEdge Server BIOS	2.0.2	BIOS facilitates the hardware initialization process and transition control to the operating system.
Integrated Dell Remote Access Controller (iDRAC)	2.30.30.30	iDRAC is a systems management hardware and software solution that provides remote management capabilities, crashed system recovery, and power control functions for Dell PowerEdge systems.
PERC H730 RAID Controller	Firmware: 25.4.0.0017 ESXi lsi_mr3 Driver: 6.903.82.00 SLES 11 SP3 Driver: 06.903.07.00	12Gbps RAID controller supporting SAS or SATA hard-disk or solid-state drives, provides unsurpassed performance and enterprise-class reliability.



Component	Version	Notes
Intel x710 10Gb Ethernet Network Adapter	Firmware: 17.0.12 ESXi net-i40e Driver: 1.3.38	Intel's X710 Converged Network Ethernet Adapters are flexible and scalable for today's demanding data center and cloud environments by providing unmatched features for virtualization, SAN networking, and proven reliable performance.



## 11 Conclusion

The Dell Validated System for EMC ScaleIO is intended to help customers simplify the design, architecture, configurations, and deployment used for building out a ScaleIO environment. Customers have asked Dell to speed their time to value, by decreasing the amount of time it takes to get up and running with an optimized ScaleIO rack level solution. Dell has worked with EMC ScaleIO to validate the configuration, architecture, and rack level design to ensure that customers achieve optimal performance and scale. This allows customers to run EMC ScaleIO on award-winning Dell PowerEdge Servers with flexibility across the network and management framework. Dell and EMC ScaleIO have engineered the Dell VS for EMC ScaleIO to provide the most efficient server configurations to help customers reach the performance and scaling requirements to run their applications utilizing SDS.



# A Deployment Guidelines

When deploying the Dell Validated System for EMC ScaleIO solution, Dell recommends that you follow this sequence of steps to ensure maximum efficiency:

1. Top-of-Rack (ToR) switch Configuration
  - a. The ToR switch pair should be configured to provide high-availability in the case of a switch failure.
  - b. VLANs need to be configured appropriately on ToR switches for Management, vMotion, vSAN, ScaleIO Data1, ScaleIO Data2, out-of-band (OOB), and any other networks.
  - c. Server switch ports need to be configured appropriately for vSphere, vSAN, and ScaleIO traffic types.
  - d. Switch uplinks need to be made available for connectivity to the core network and the out-of-band management switch.
  - e. IGMP Snooping and Querier is recommended for the vSphere Virtual SAN (vSAN) network.
2. The Out-of-Band Switch (OOB) Switch should be configured with the native out-of-band VLAN to all server iDRACs. The OOB switch should have an uplink connection so that it is accessible to the ToR switch pair.
3. There are 4x 10GbE ports/Node and 1x 1GbE iDRAC for each of the Management and Compute servers for LAN, vSAN and/or ScaleIO traffic. These should be cabled appropriately to the ToR and OOB switches.
4. Configure the Dell iDRAC IP addresses for static or DHCP as necessary on each of the Management and Compute servers.
5. Dell BIOS Configuration
  - a. The Intel CPU Virtualization feature should be enabled.
  - b. Onboard SATA controller should be configured for AHCI mode.
  - c. Ensure that the Dell BIOS boot order configuration is set appropriately for IDSDM boot for the R630 Management servers and the rear SATA drive A-1 (non-RAID) for the R730xd Compute servers.
6. Configure the PERC H730 RAID controller into JBOD personality mode for all servers.
7. Install VMware ESXi using the Dell ESXi Customized ISO to each of the R630 Management and R730xd Compute servers. Ensure the correct storage and NIC driver versions in vSphere.
8. Configure each Management and Compute server ESXi hypervisor IP address, management VLAN, default gateway, and DNS address. The ESXi SSH console should be enabled on the R730xd Compute servers for ScaleIO configuration access.

9. Configure the first R630 Management ESXi hypervisor for a single node vSAN cluster datastore using the local SSD and HDD storage devices.
10. Windows Jump VM
  - a. Deploy and configure a Windows virtual machine (VM) to run the ScaleIO GUI and OpenManage Essentials (OME) to the VSAN Datastore.
  - b. OME will need access to the out-of-band and management VLANs in order to access the iDRACs and ESXi. This can be achieved through multi-homing the VM or Layer-3 routing.
  - c. This VM can be used as a Jump server to gain access to the vSphere and ScaleIO environment for many of the following steps.
  - d. Install additional vSphere and ScaleIO tools as necessary.
11. Deploy and configure the vCenter Server Appliance (vCSA) to the VSAN Datastore.
12. Configure the vCenter Management and Compute vSphere clusters and enable vSphere Virtual SAN. Configure the vSphere NTP service as needed on each host.
13. Configure vSphere Standard Virtual Switches for each R630 Management server to support vSphere, vSAN and ScaleIO traffic.
14. Configure vSphere Virtual SAN disk groups on the additional Management servers.
15. Configure each Management vSphere hosts to redirect the ESXi scratch folder to persistent storage to enable persistent vSAN trace logs.
16. Configure vSphere Distributed Virtual Switches for the Compute cluster to support vSphere and ScaleIO traffic.
17. Configure OME, inventory all server iDRACs and ESXi, and apply firmware updates to all hosts as necessary.
18. Register the ScaleIO Plug-in for vCenter and deploy the ScaleIO virtual machine (SVM) template to the VSAN Datastore. Multiple SVM templates can be created to speed the ScaleIO deployment.
19. Deploy the ScaleIO Data Client (SDC) to each Compute host.
20. Deploy the ScaleIO Environment
  - a. Install the ScaleIO Data Server (SDS) and Metadata Manager (MDM) services (as needed) to each R730xd Compute node.
  - b. IPs will need to be assigned on the Management and ScaleIO data networks (separate subnets).
  - c. SSD and HDD storage of different types will each need to be configured into individual Storage Pools.
21. Create new volume(s) from each ScaleIO Storage Pool and map them to the R730xd Compute servers as vSphere clustered storage. Add the new VMFS datastores to vSphere.



22. Perform ScaleIO performance optimizations for vSphere environments based on ScaleIO v2 Performance Fine-Tuning document recommendations.
23. The ScaleIO gateway VM can be migrated to the Management vSphere cluster.
24. Apply licensing for vSphere, vCenter, vSAN and ScaleIO.

## B Additional resources

This section provides guidance to find documentation and other support resources for components of the Dell Validated System for EMC ScaleIO.

### B.1 Storage components

Table 12 Storage

Product	Description	Link to documentation
EMC ScaleIO	Converges storage and compute resources into a single-layer architecture, aggregating capacity and performance, simplifying management, and scaling to thousands of nodes.	<a href="http://bit.ly/1uSKzW2">http://bit.ly/1uSKzW2</a>

### B.2 Virtualization components

Table 13 Virtualization

Product	Description	Link to documentation
VMware vCenter Server	Provides a scalable and extensible platform that forms the foundation for virtualization management.	<a href="http://vmw.re/1kltqhc">http://vmw.re/1kltqhc</a>
VMware vSphere ESXi	Virtualized infrastructure for Vblock Systems. Virtualizes all application servers and provides VMware High Availability (HA) and Dynamic Resource Scheduling (DRS).	<a href="http://vmw.re/1wutyZU">http://vmw.re/1wutyZU</a>
VMware Single Sign-On (SSO) Service	Provides VMware-specific authentication services.	<a href="http://vmw.re/1XaubUq">http://vmw.re/1XaubUq</a>
VMware VSAN	VMware Virtual SAN is a radically simple, enterprise-class shared storage solution for hyper-converged infrastructure optimized for today's all-flash performance.	<a href="http://vmw.re/1joENei">http://vmw.re/1joENei</a>



## B.3 Dell servers and switches

Table 14 Servers and Switches

Product	Description	Link to documentation
PowerEdge R730xd	Dell rack server that provides core compute infrastructure for the Dell VS for EMC ScaleIO.	<a href="http://dell.to/1UoCBkR">http://dell.to/1UoCBkR</a>
PowerEdge R630	Dell rack server that provides management infrastructure for the Dell VS for EMC ScaleIO.	<a href="http://dell.to/1WEYbYo">http://dell.to/1WEYbYo</a>
Dell Networking S4048-ON	The Dell Networking S-Series S4048-ON is an ultralow-latency 10/40GbE top-of-rack (ToR) switch built for applications in high-performance data center and computing environments.	<a href="http://dell.to/1Zy9iR9">http://dell.to/1Zy9iR9</a>
Dell Networking S3048-ON	High-density, 1RU 48-port 1000BASE-T+ with four 10GbE uplinks, non-blocking line-rate performance, featuring the Open Networking Install Environment (ONIE).	<a href="http://dell.to/28kzEvH">http://dell.to/28kzEvH</a>

## B.4 Other Dell resources

- [Support.dell.com](http://support.dell.com) is focused on meeting customer requirements with proven services.
- [Dell Tech Center](#) 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 product and installations.

## C Abbreviations and Acronyms

The following table lists and defines abbreviations and acronyms used in this document.

Table 15 List of abbreviations and acronyms

Abbreviation/Acronym	Definition
AHCI	Advanced Host Controller Interface
CI	Converged Infrastructure
DAS	Direct-Attached Storage
DDR	Double Data Rate
DHCP	Dynamic Host Configuration Protocol
DIMM	Dual In-Line Memory Module
DRS	(VMware) Distributed Resource Scheduler
DRS	Dynamic Resource Scheduling
eBGP	External Border Gateway Protocol
ECMP	Equal-Cost-Multi-Path
ESX/ESXi	Elastic Sky X
GbE	Gigabit Ethernet
GUI	Graphical User Interface
HA	High Availability
HCI	Hyper-Converged Infrastructure
HDD	Hard Disk Drive
I&O	Infrastructure and Operations



Abbreviation/Acronym	Definition
iDRAC	Integrated Dell Remote Access Controller
IDSDM	Internal Dual Secure Digital Module
IGMP	Internet Group Management Protocol
IOPS	Input/Output Operations Per Second
KPI	Key Performance Indicator
LAG	Link Aggregation Group
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
MDM	MetaData Manager
NDC	Network Daughter Card
NIC	Network Interface Card
NIOC	Network IO Control
NTP	Network Time Protocol
OME	OpenManage Essentials
ONIE	Open Networking Install Environment
OOB	Out-Of-Band
PERC	PowerEdge RAID Controller
QSFP	Quad Small Form-Factor Pluggable
RA	Reference Architecture
REST	Representational State Transfer



Abbreviation/Acronym	Definition
RU	Rack Unit (unit of measure)
SAN	Storage Area Network
SAS	Serial Attached SCSI
SATA	Serial AT Attachment
SDC	ScaleIO Data Client
SDS	Software-Defined Storage
SDS	ScaleIO Data Server
SFP	Small Form-Factor Pluggable
SLA	Service Level Agreement
SLES	SUSE Linux Enterprise Server
SNMP	Simple Network Management Protocol
SSD	Solid-State Drive
SSH	Secure Socket Shell/Secure Shell
SSO	Single Sign-On
SVM	ScaleIO Virtual Machine
TCO	Total Cost Of Ownership
ToR/TOR	Top Of Rack
vCSA	vCenter Server Virtual Appliance
VDS	vSphere Distributed Switch
VFFS	Virtual Flash File System



Abbreviation/Acronym	Definition
VIB	vSphere Installation Bundle
VLT	Virtual Link Trunking
VM	Virtual Machine
VMFS	Virtual Machine File System
vPC	Virtual Port Channel
VS	Validated System
VSAN	Virtual Storage Area Network
VSS	vSphere Standard Switch