Dell EMC + VMware Cloud Infrastructure Platform for NFV

VMware vCloud NFV 1.5 – Dell EMC ScaleIO Design Guide

Service Provider Solutions Group
April 2017

A Dell EMC Design Guide

Revisions

Date	Description
April 2017	Initial release

Table of contents

Executive Summary

This document provides detailed design guidance for creating a VMware vCloud NFV Platform for the purpose of Network Functions Virtualization (NFV) based on VMware best practice and real-world scenarios. This solution leverages the Dell EMC ScaleIO Software Defined Storage solution. The platform dramatically simplifies data center operations, delivering enhanced agility, rapid innovation, better economics and scale.

Audience

This document is for those who are responsible for the implementation of the vCloud NFV Platform with ScaleIO. This document is based on the reference architecture described in the vCloud NFV Reference Architecture v1.5 document.

This document assumes that the audience has some understanding of the VMware and ScaleIO components, which are used in the solution. The audience should have access to the installation and configuration guides of the respective components.

Document Structure

This document is divided into the sections listed in Table 1 below.

Section	Description
vCloud NFV Detail Design	This section contains the design details for all the components of the vCloud NFV platform with ScaleIO
Capacity Planning and Sizing	Capacity planning and sizing guidelines.
Monitoring and Logging	Metrics and dashboards for monitoring the platform are described in this section

Table 1 Document Structure

1 Pre-Requisites

1.1 Software Requirements

Table 2 depicts two types of components:

- Required – The solution relies on these components and will not function as planned without them.
- Recommended – These components provide useful additional capabilities. These capabilities are discussed in this document. Alternative or third party components can be used where appropriate.

Component	Version	Required in Solution	Functional Block
VMware vSphere			
VMware ESXi	6.0 U2	Required	NFVI
VMware vCenter Server Appliance	6.0 U2	Required	VIM
VMware vSphere Replication	6.1.1	Recommended	NFVI
VMware vSphere Data Protection	6.1.2	Recommended	NFVI Ops
EMC ScaleIO	2.0.0.3	Recommended	NFVI
VMware vRealize Operations Insight			
VMware vRealize Operations Advanced	6.2.1	Required	NFVI Ops
VMware vRealize Log Insight	3.3.1	Required	NFVI Ops
VMware vCloud Director for Service Providers	8.10	Required	VIM
VMware NSX			
VMware NSX for vSphere	6.2.2	Required	NFVI
VMware NSX Manager	6.2.2	Required	VIM
EMC ScaleIO			
MDM	2.0.0.3	Required	NFVI
SDS	2.0.0.3	Required	NFVI
SDC	2.0.0.3	Required	NFVI

Table 2 VMware Software Requirements

1.2 Hardware Requirements

1.2.1 Compute Resources

The compute resources are the physical servers on which the hypervisor is installed; these are the server nodes that contribute CPU and Memory capacity to the workload cluster for pooling the resources. Additionally, these nodes must have sufficient bandwidth and

redundancy for the network connectivity of the workloads they host. All hardware used must be on the VMware Hardware Compatibility List[1] (HCL).

1.2.2 Storage Resources

This reference architecture leverages Dell EMC ScaleIO as the shared storage solution.

ScaleIO is a software-only solution that uses existing local disks and LANs so that the host can realize a virtualized SAN with all the benefits of external storage. ScaleIO software turns existing local internal storage into internal shared block storage. ScaleIO software components are installed in the application hosts and inter-communicate using a standard LAN to handle the application I/O requests sent to ScaleIO block volumes.

The ScaleIO virtual SAN software consists of three software components:

- Meta Data Manager (MDM) - Configures and monitors the ScaleIO system. The MDM can be configured in a redundant Cluster Mode, with three members on three servers, or in Single Mode on a single server.
- ScaleIO Data Server (SDS) - Manages the capacity of a single server and acts as a back-end for data access. The SDS is installed on all servers contributing storage devices to the ScaleIO system.
- ScaleIO Data Client (SDC) - SDC is a lightweight device driver situated in each host whose applications or file system requires access to the ScaleIO virtual SAN block devices. The SDC exposes block devices representing the ScaleIO volumes that are currently mapped to that host.

ScaleIO components are designed to work with a minimum of three server nodes. The physical server node, running VMware vSphere, can host other workloads beyond the ScaleIO virtual machine. ScaleIO is implemented as a software layer that takes over the existing local storage on the servers. This software layer combines the local storage with storage from the other servers in the environment, and presents logical units (LUNs) from this aggregated storage for use by the virtual environment. These LUNs are presented using the iSCSI protocol and are then usable as data stores within the environment.

The software sits between the disks and the file system on the same layer as a logical volume manager. Virtual machines continue to process I/O to VMDKs within a datastore, but this datastore is now being provided by the ScaleIO software instead of by the local disks. In a vSphere environment, ScaleIO is implemented as a separate virtual machine. The software components are installed on the ScaleIO virtual machine.

Protection domain - A large ScaleIO storage pool can be divided into multiple protection domains, each of which contains a set of SDSs. ScaleIO volumes are assigned to specific protection domains.

---

[1] VMware Hardware Compatibility Guide (HCL) : http://www.vmware.com/resources/compatibility/search.php

Storage pool - A storage pool is a subset of physical storage devices in a protection domain. Each storage device belongs to one (and only one) storage pool. When a protection domain is generated, by default it has one storage pool.

1.2.3 Network Resources

Each ESXi host in the cluster should have a network configuration to cater to the redundancy and performance needs of the platform. At a minimum, there should be no single point of failure by providing redundant network controllers and Ethernet ports.

The Ethernet ports of the ESXi hosts need to be connected in a redundant configuration to the physical switches providing alternate paths in case of hardware failure. VLANs are configured to segregate network workloads such as VMware vSphere vMotion traffic, ScaleIO Virtual SAN traffic and host management traffic.

1.3 Supporting Components

Product	Description
Directory Server	Centralized authentication source for management components
DNS Server	Provide forward and reverse lookup service to all platform components
NTP Server	Time sync service to all components
SMTP Server	Used to send email notifications from platform as a result of events and alarms
SNMP Server	Used to send SNMP alerts to external monitoring systems
SFTP/FTP Server	Used for NSX Manager backups
NFS Server	Used for vCloud Director transfer space

Table 3 Supporting Components

2 vCloud NFV Detail Design

2.1 NFV Infrastructure Design

The vCloud NFV Infrastructure components are the ESXi hosts that provide the underlying resources for the virtualized network functions (VNFs). The Virtual SAN is also provides the storage resources for the platform while NSX caters to the network requirements. In this section, we look at the design for the NFV Infrastructure (NFVI) and its components.

2.1.1 Cluster Design

The NFV Infrastructure platform contains three clusters, the management cluster, resource cluster and edge cluster. This architectural best practice allows for efficient resource management, clear demarcation between resource providers and resource consumers, establish security boundaries and design different levels of availability based on cluster workloads.

All the hosts in a cluster should have identical configuration and specifications for efficient management of resources. For better resource management the management components are deployed in the management cluster, the VMware NSX Edge devices for the VNFs are deployed in the edge cluster and the VNFs are deployed in the resource cluster. NSX Edge devices used by the management components are deployed in the management cluster.

Two vCenter servers and two NSX manager instances are deployed in the management cluster. The first instance includes the management components and provides networking services such as load balancing to the components in the management cluster. The second instance is used to manage the VNFs deployed in the resource cluster and provide networking services such as routing by deploying NSX Edge devices in the edge cluster. Each vCenter Server points to a load balanced pair of external Platform Services Controller (PSC) instances. The PSC design is explained in more detail in section 2.2.1.

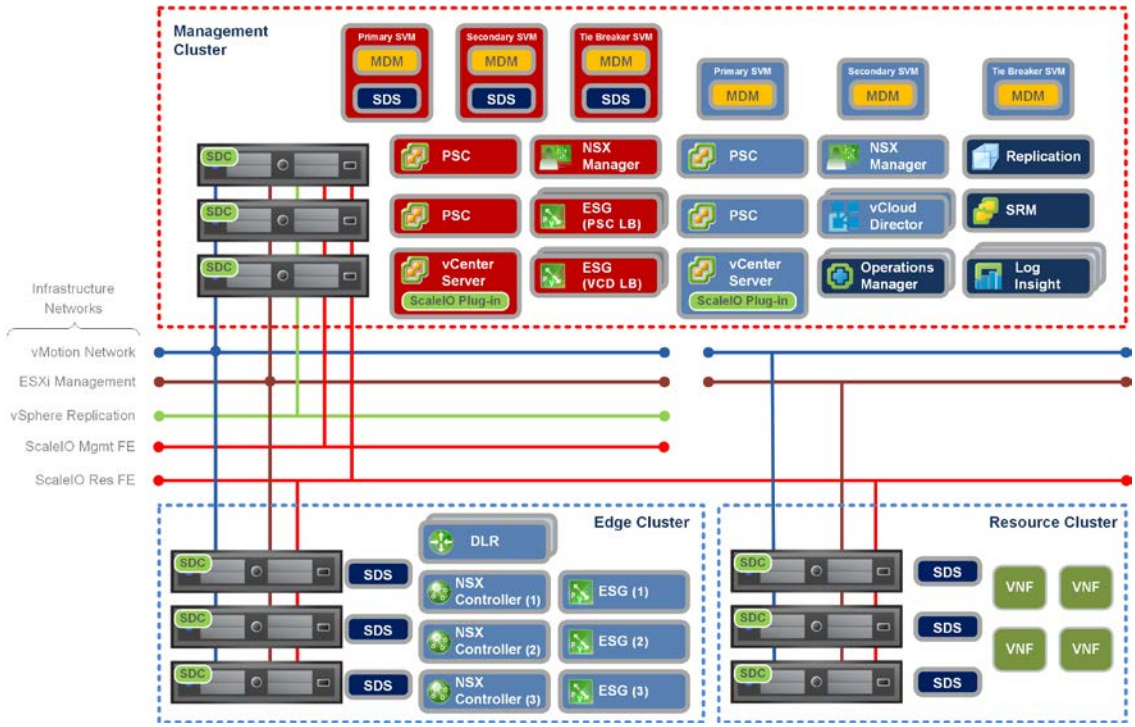

Figure 1 vCloud NFV Cluster Design

- **Management Cluster** - The management cluster leverages VMware vSphere High Availability and VMware vSphere Distributed Resource Scheduler, and requires specific configuration parameters. Table 4 below lists the parameters pertaining to the management cluster.

Parameter	Value
vSphere High Availability	
vSphere HA	Enabled
Host Monitoring	Enabled
Admission Control	Enabled
Admission Control Policy	Host failures to tolerate = 1 3 Node ScaleIO configuration supports 1 host failure
VM Monitoring	Disabled
Monitoring Sensitivity	High
Datastore Heart-beating	Automatically select datastores accessible from the host
vSphere Distributed Resource Scheduler	
vSphere DRS	Enabled
Automation Level	Fully Automated
Migration Threshold	2
Virtual Machine Automation	Disabled
Power Management	Disabled
Enhanced vMotion Compatibility	
EVC	Enabled2
Affinity and Anti-Affinity Rules	

Recommended Anti-Affinity Rules (VMs should be on separate hosts)	NSX Edge for VCD load balancer (active, standby) NSX Edge for PSC load balancer (active, standby) Management PSCs (psc1, psc2) Resource PSCs (psc1, psc2) vCloud Director cells (cell1, cell2) vRealize Operations Manager (master, replica) vRealize Log Insight (master, worker1, worker2)
Recommended Affinity Rules	SVM (Master MDM) tied to a ESXi host SVM (Slave MDM) tied to a ESXi host SVM (Tiebraker MDM) tied to a ESXi host

Table 4 Management Cluster Settings

- **Edge Cluster** - The edge cluster leverages vSphere HA and vSphere DRS and requires specific configuration parameters. Table 5 lists the parameters pertaining to the edge cluster.

Parameter	Value
vSphere High Availability	
vSphere HA	Enabled
Host Monitoring	Enabled
Admission Control	Enabled
Admission Control Policy	Host failures to tolerate same as
VM Monitoring	Disabled
Monitoring Sensitivity	High
Datastore Heart-beating	Automatically select datastores accessible from the host
vSphere Distributed Resource Scheduler	
vSphere DRS	Enabled
Automation Level	Partially Automated
Virtual Machine Automation	Disabled
Power Management	Disabled
Enhanced vMotion Compatibility	
EVC	Enabled2
Affinity and Anti-Affinity Rules	
Recommended Anti-Affinity Rules (VMs should be on separate hosts)	NSX Controllers (controller1, controller2, controller3) NSX DLR (active, standby) NSX Edge VNF routing (esg1, esg2, esg3)
Recommended Affinity Rules	None

Table 5 Edge Cluster Settings

- **Resource Cluster** - The resource cluster leverages vSphere HA and vSphere DRS and requires specific configuration parameters. Table 6 below lists the parameters pertaining to the resource cluster.

Parameter	Value
vSphere High Availability	
vSphere HA	Enabled
Host Monitoring	Enabled
Admission Control	Enabled
Admission Control Policy	Host failures to tolerate same as ScaleIO Virtual SAN
VM Monitoring	Disabled
Monitoring Sensitivity	High
Datastore Heart-beating	Automatically select datastores accessible from the host
vSphere Distributed Resource Scheduler	
vSphere DRS	Enabled
Automation Level	Partially Automated
Virtual Machine Automation	Disabled
Power Management	Disabled
Enhanced vMotion Compatibility	
EVC	Enabled
Affinity and Anti-Affinity Rules	
Recommended Anti-Affinity Rules	VNF workloads as defined by vendor
Recommended Affinity Rules	VNF workloads as defined by vendor SVM (SDS) tied to a particular host

Table 6 Resource Cluster Settings

VMware recommends evaluating any performance impact of enabling Enhanced vMotion Compatibility (EVC) on VNF workloads in the resource cluster. For more detailed information, refer to the document "Impact of EVC on Application Performance[2]".

VMware recommends enabling the EVC mode and setting this to the processor vendor of the CPUs of the ESXi hosts in the cluster. It is not recommended to have hosts from mixed CPU vendors (Intel / AMD) in the same cluster. For details on selecting the EVC mode, refer to the KB1003212[3].

2.1.2 Network Design

The vCloud NFV platform consists of infrastructure networks and tenant networks. Infrastructure networks are the host level networks that connect hypervisors to physical networks. The infrastructure network traffic consists of vSphere vMotion traffic, EMC ScaleIO virtual SAN traffic and host management traffic.

Tenant networks are the overlay networks created by NSX and physical VLAN networks, such as the Management VLAN, to connect VMs to hypervisor. The vCloud NFV platform relies on tenant networks to connect management components such as vCenter server, vCloud Director, NSX Manager, vRealize Operations Manager, vRealize Log Insight etc. Tenant networks are also used to provide connectivity to VNFs.

[2] Impact of Enhanced vMotion Compatibility on Application Performance : http://www.vmware.com/files/pdf/techpaper/VMware-vSphere-EVC-Perf.pdf
[3] Enhanced vMotion Compatibility (EVC) processor support : http://kb.vmware.com/kb/1003212

All ESXi hosts in the vCloud NFV platform are configured with two VMware vSphere Distributed Switch (VDS) devices, which provide a consistent network configuration across multiple hosts and is a requirement of NSX for vSphere. One VDS is used for tenant networks and the other VDS maintains the infrastructure networks. For the management cluster, the Tenant VDS hosts the Management VLAN network.

The hypervisor communicates through VMkerne port groups on the VDS while virtual machines connect to Virtual Machine port groups and are labeled for easy identification.

The ESXi host physical NICs are used as uplinks to connect the VDS to the physical network switches. Each ESX host should have a minimum of two 10Gbps Ethernet NICs. Each of the two VDS on the ESXi host will then have one uplink each.

For high availability, VMware recommends four NICs for each ESXi host. This allows each VDS to be assigned two uplinks for high availability. The two uplinks are spread across two physical network devices, one of these can be an expansion card and the other can be an on-board controller so there is no single point of failure. However, for consistent performance ensure that the network devices are the same chipset family and vendor.

All ESXi Physical NICs will connect to Layer 2 or Layer 3 Managed switches on the physical network. At least two 10 Gigabit Ethernet switches with enough ports for all the physical NICs of all the ESXi hosts are required.

Table 7 lists the VDS configuration parameters. Since the VXLAN traffic frames are slightly larger because of encapsulation, the MTU for each VDS must be set to 1600 bytes for use with NSX. For the best performance, the same MTU size should be set throughout the network.

Specification	Value
MTU	9200 Bytes
Teaming Mode	IEEE 802.3ad, LACP
Segment IDs	5000 - 7999

Table 7 VDS Configuration Parameters

Network I/O Control (NIOC) is used to prioritize the network traffic over the two-shared uplinks of each VDS. In case of contention, the NIOC share value determines the bandwidth allocation of the networks on the VDS. Refer to the NIOC performance evaluation guide[4] for more details on NIOC configuration. Table 10 lists the recommended NIOC shares for this reference architecture.

2.1.2.1 Physical Networks

The physical networks consists of physical network switches and computer servers, in our case, it is the ESXi hypervisor servers. Figure 2 is the physical network topology.

[4] Performance Evaluation of NIOC : http://www.vmware.com/files/pdf/techpaper/Network-IOC-vSphere6-Performance-Evaluation.pdf

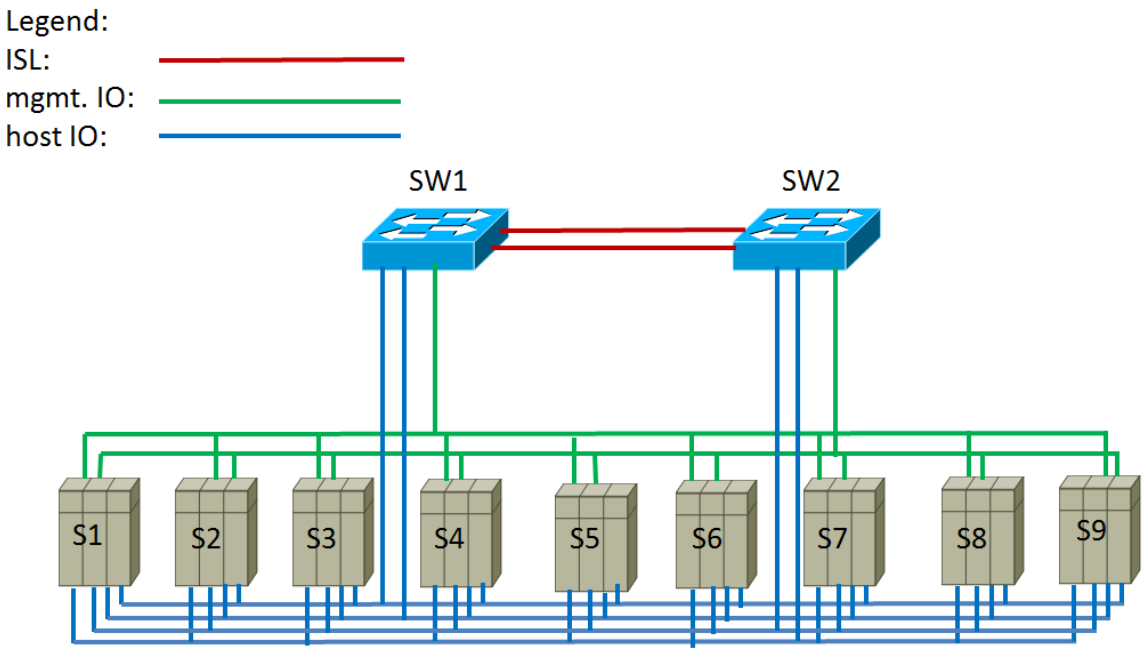

Figure 2 Physical Network Topology

Each server has six 10GbE dual-port NICs (four out of these six are used), two for management IO NIC bonding and four for host IO NIC bonding. See Table 8 for a detail port map.

	NIC Port	Switch	SW Pport	Bond	
server1	p1p1	sw1	Te0/64	po64	
iDrac: 172.16.104.10	p1p2	sw2	Te0/64	po64	
	p2p1	sw1	Te0/65		SIO-1
	p2p2	sw2	Te0/65		SIO-2
	p4p1	sw1	Te0/80	po80	
	p4p2	sw2	Te0/80	po80	
	p5p1	sw1	Te0/81	po80	
	p5p2	sw2	Te0/81	po80	
server2	p1p1	sw1	Te0/66	po66	
iDrac: 172.16.104.11	p1p2	sw2	Te0/66	po66	
	p2p1	sw1	Te0/67		SIO-1
	p2p2	sw2	Te0/67		SIO-2
	p4p1	sw1	Te0/82	po82	
	p4p2	sw2	Te0/82	po82	
	p5p1	sw1	Te0/83	po82	
	p5p2	sw2	Te0/83	po82	
server3	p1p1	sw1	Te0/68	po68	
iDrac: 172.16.104.12	p1p2	sw2	Te0/68	po68	
	p2p1	sw1	Te0/69		SIO-1
	p2p2	sw2	Te0/69		SIO-2
	p4p1	sw1	Te0/84	po84	
	p4p2	sw2	Te0/84	po84	
	p5p1	sw1	Te0/85	po84	
	p5p2	sw2	Te0/85	po84	
server4	p1p1	sw1	Te0/70	po70	
iDrac: 172.16.104.13	p1p2	sw2	Te0/70	po70	
	p2p1	sw1	Te0/71		SIO-1

	p2p2	sw2	Te0/71		SIO-2
	p4p1	sw1	Te0/86	po86	
	p4p2	sw2	Te0/86	po86	
	p5p1	sw1	Te0/87	po86	
	p5p2	sw2	Te0/87	po86	
server5	p1p1	sw1	Te0/72	po72	
iDrac: 172.16.104.14	p1p2	sw2	Te0/72	po72	
	p2p1	sw1	Te0/73		SIO-1
	p2p2	sw2	Te0/73		SIO-2
	p4p1	sw1	Te0/88	po88	
	p4p2	sw2	Te0/88	po88	
	p5p1	sw1	Te0/89	po88	
	p5p2	sw2	Te0/89	po88	
server6	p1p1	sw1	Te0/74	po74	
iDrac: 172.16.104.15	p1p2	sw2	Te0/74	po74	
	p2p1	sw1	Te0/75		SIO-1
	p2p2	sw2	Te0/75		SIO-2
	p4p1	sw1	Te0/90	po90	
	p4p2	sw2	Te0/90	po90	
	p5p1	sw1	Te0/91	po90	
	p5p2	sw2	Te0/91	po90	
server7	p1p1	sw1	Te0/76	po76	
iDrac: 172.16.104.16	p1p2	sw2	Te0/76	po76	
	p2p1	sw1	Te0/77		SIO-1
	p2p2	sw2	Te0/77		SIO-2
	p4p1	sw1	Te0/92	po92	
	p4p2	sw2	Te0/92	po92	
	p5p1	sw1	Te0/93	po92	
	p5p2	sw2	Te0/93	po92	
server8	p1p1	sw1	Te0/78	po78	
iDrac: 172.16.104.17	p1p2	sw2	Te0/78	po78	
	p2p1	sw1	Te0/79		SIO-1
	p2p2	sw2	Te0/79		SIO-2
	p4p1	sw1	Te0/94	po94	
	p4p2	sw2	Te0/94	po94	
	p5p1	sw1	Te0/95	po94	
	p5p2	sw2	Te0/95	po94	
server9	p1p1	sw1	Te0/96	po96	
iDrac: 172.16.104.18	p1p2	sw2	Te0/96	po96	
	p2p1	sw1	Te0/97		SIO-1
	p2p2	sw2	Te0/97		SIO-2
	p4p1	sw1	Te0/98	po98	
	p4p2	sw2	Te0/98	po98	
	p5p1	sw1	Te0/99	po98	
	p5p2	sw2	Te0/99	po98	

Table 8 Physical Network Connection Port Map

2.1.2.2 Virtual Network

The virtual networks, which are used to connect VMs to hypervisor, consist of two VDS switches per each cluster, management and host VDS. Each cluster has three hypervisor servers. Figure 3 is a diagram for the management VDS in one cluster. The host IO VDS is similar to the management VDS except it has four LACP uplinks.

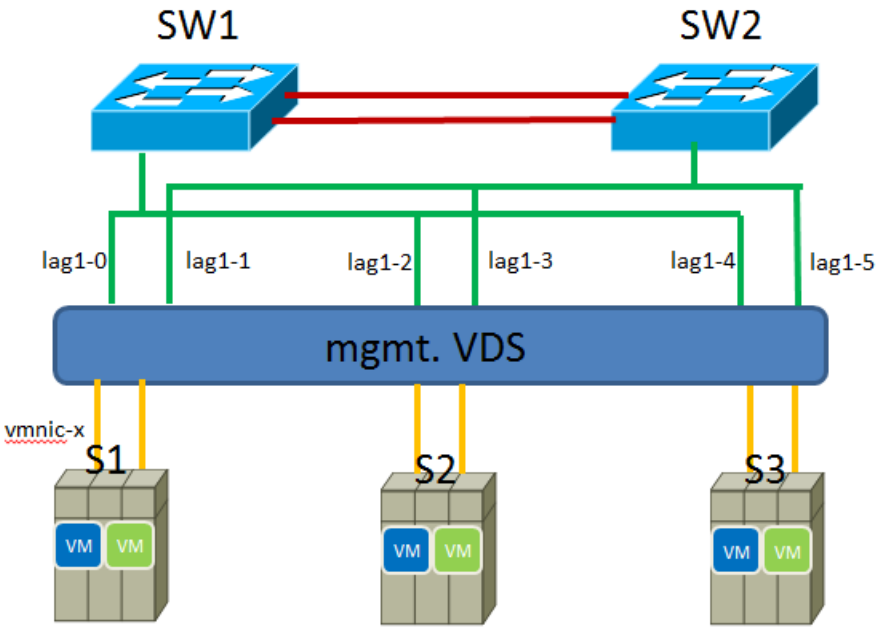

Figure 3 Virtual Network, management VDS in one cluster

On switch side, a bonded interface is configured with the following:

1. Physical Interface
   interface TenGigabitEthernet 2/0/3
   description link to R630-1 vmnic7 68:05:ca:38:ae:cb
   fabric isl enable
   fabric trunk enable
   channel-group 10 mode active type standard
   lacp timeout long
   no shutdown

2. Port channel interface
   interface Port-channel 10
   vlag ignore-split
   description NIC bonding interface for R730xd-1
   switchport
   switchport mode trunk
   switchport trunk allowed vlan add 40
   switchport trunk tag native-vlan
   spanning-tree shutdown
   no shutdown
   !

In VDS, by default LACP is not added automatically, you have to explicitly add LACP configuration to enable this feature.

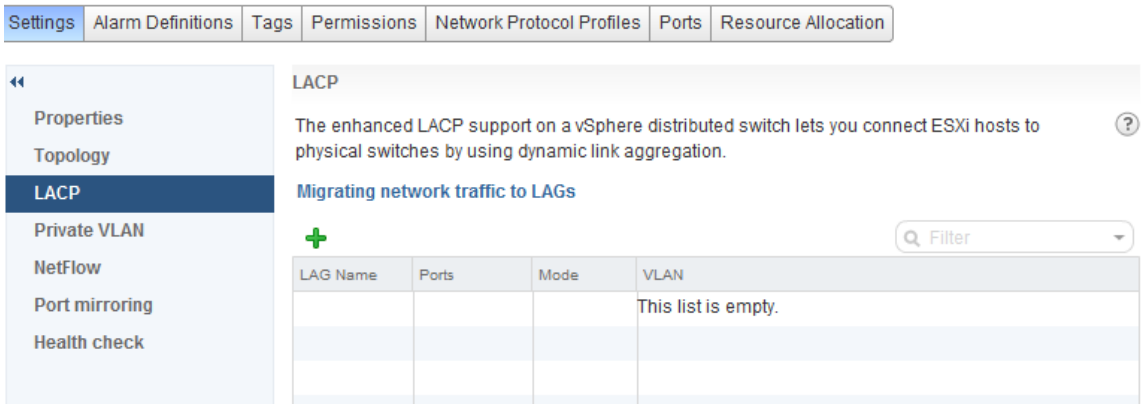

Figure 4 LACP Settings – Empty list

No LAG is been created, click the **+** sign to add LAG configuration.

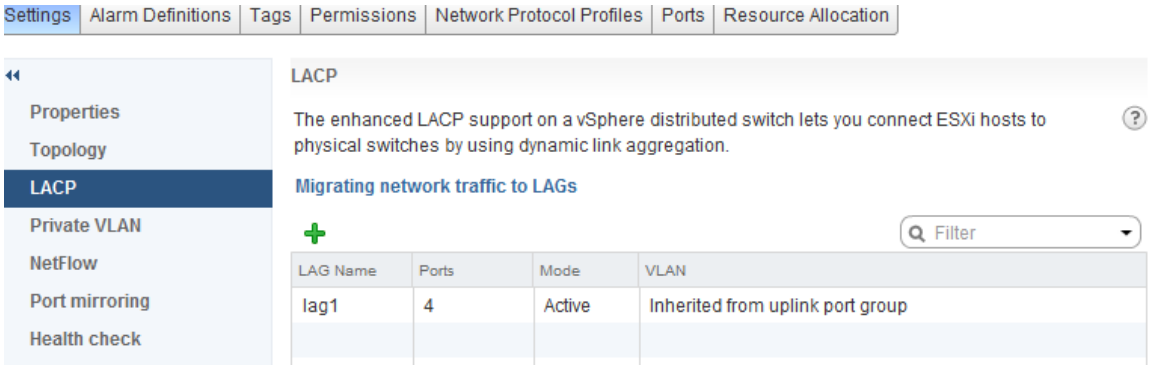

Figure 5 LACP Settings – Populated list

2.1.2.3 Infrastructure Networks

Each ESXi host has multiple VMkernel port groups configured for each infrastructure network. The infrastructure networks are

- vMotion Network - This is the network for vSphere vMotion traffic.
- ScaleIO Virtual SAN Network - This is the network for ScaleIO Virtual SAN shared storage traffic.
- ESXi Management - This is the network for ESXi host management traffic

VMware recommends that the vSphere vMotion be terminated at the ToR switch, as it is host-to-host traffic only. The infrastructure networks of the resource cluster are also segregated on to separate VLAN networks for security and performance. Table 9 lists the infrastructure VDS configuration for each of the three clusters.

Port Group	Type	Teaming policy	VLAN	Cluster
ESXi Management	VMkernel	Load Based	20	Management
vMotion Network	VMkernel	Explicit Failover	30	Management
ScaleIO Virtual SAN Network	VMkernel	Explicit Failover	40	Management
ESXi Management	VMkernel	Load Based	20	Edge
vMotion Network	VMkernel	Explicit Failover	30	Edge
ScaleIO Virtual SAN Network	VMkernel	Explicit Failover	40	Edge
ESXi Management	VMkernel	Load Based	25	Resource
vMotion Network	VMkernel	Explicit Failover	35	Resource
ScaleIO Virtual SAN Network	VMkernel	Explicit Failover	45	Resource

Table 9 Infrastructure VDS Configuration

The NIOC share values are configured at the VDS level. Table 10 lists the recommended I/O parameters for the Infrastructure VDS for each of the three clusters.

Network	Limit	Shares	NIC Shares	Share Value	Cluster
ESXi Management traffic	Unlimited	Normal	Normal	50	Management
vMotion traffic	Unlimited	Normal	Normal	100	Management
ScaleIO Virtual SAN traffic	Unlimited	Normal	Normal	100	Management
ESXi Management traffic	Unlimited	Normal	Normal	50	Edge
vMotion traffic	Unlimited	Normal	Normal	100	Edge
ScaleIO Virtual SAN traffic	Unlimited	Normal	Normal	100	Edge
ESXi Management traffic	Unlimited	Normal	Normal	50	Resource
vMotion traffic	Unlimited	Normal	Normal	100	Resource
ScaleIO Virtual SAN traffic	Unlimited	Normal	Normal	100	Resource

Table 10 Infrastructure VDS NIOC Parameters

2.1.2.4 Tenant Networks

Tenant networks are used to interconnect the VMs of the vCloud NFV platform. These are configured on a dedicated tenant VDS in each of the clusters. The tenant networks are

- VNF Network - This is the VXLAN based network for VNF to VNF communication
- Management VLAN - This is the VLAN based network for management component communication

Table 11 lists the recommended I/O parameters for the tenant VDS for each of the three clusters.

Network	Limit	Shares	NIC Shares	Share Value	Cluster
Management VLAN	Unlimited	Normal	Normal	60	Management
Management VLAN	Unlimited	Normal	Normal	50	Edge
VNF Network	Unlimited	Normal	High	100	Resource

Table 11 Infrastructure VDS NIOC Parameters

VNF Network - East-West traffic from the VNFs are handled by VXLAN Tunnel Endpoint (VTEP) Logical Switches that can span across separate VDS instances over the entire transport zone. These logical switches are consumed by the Telco cloud by mapping them to vCloud Director external networks. Depending on the VNF network topology, VNFs connect to one or more Org Networks that are in turn connected to the external networks.

North-South traffic flow is implemented by connecting the logical switch to the Distributed Logical Router (DLR) and the NSX Edge for NFV traffic to the external network. When stateful services such as firewall, load-balancing, VPN and NAT are required, an active/standby NSX Edge pair is deployed. When NSX Edge is used solely to provide routing function such as for VNFs, an OSPF ECMP configuration can be deployed to provide additional resilience and fault tolerance.

DLR is deployed as an active-standby HA configuration while three NSX Edge devices are deployed to provide routing services, and are configured with ECMP OSPF peering.

Since a DLR and NSX Edge cannot be connected directly to each other, a transit network is used for this purpose. Anti-Affinity rules are configured so that the DLR active-standby pairs are on separate hosts. Anti-Affinity rules are created to keep the NSX Edge devices on separate hosts as well.

Figure 6 shows the vCloud Director networks and logical switches for East-West traffic and the NSX Edge devices deployed in the edge cluster for dynamic routing of the VNF network for North-South traffic.

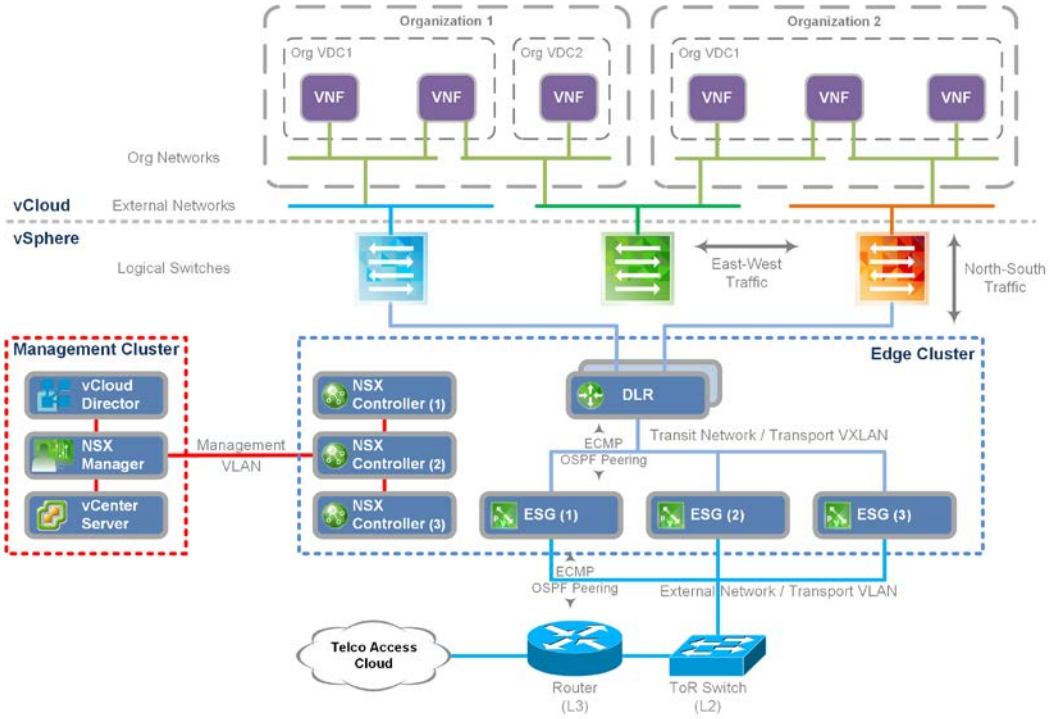

Figure 6 VNF Networking

Management VLAN - All the management nodes local to the site are interconnected using the Management VLAN network across all the three clusters. Figure 7 shows the management VLAN and the management components that utilize this network.

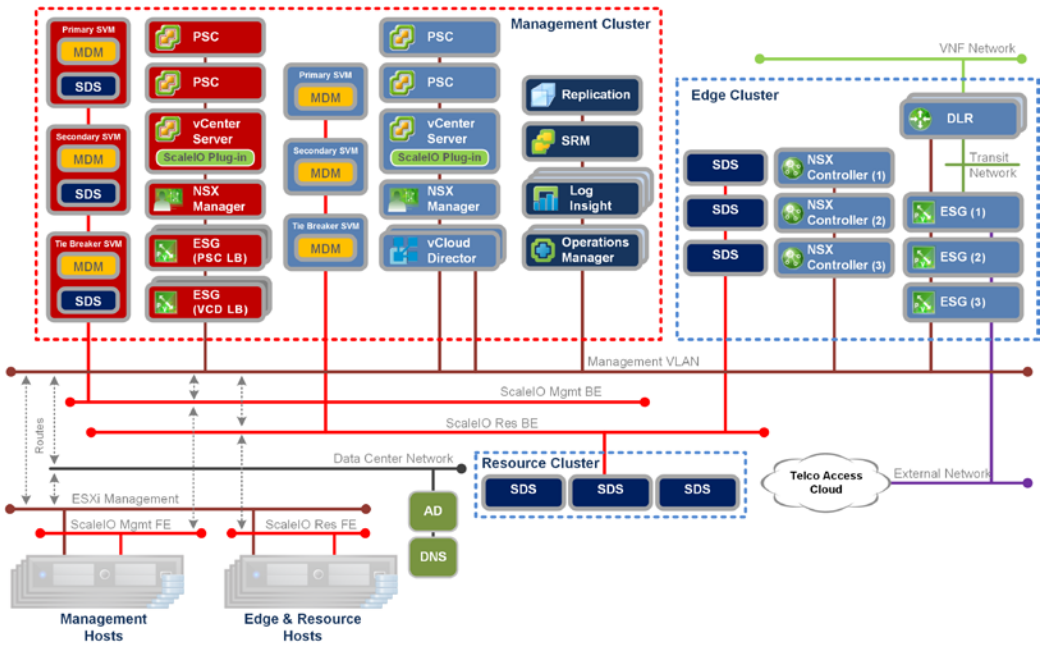

Figure 7 Management VLAN Tenant Network

2.1.2.5 Datacenter Network

The datacenter network is the physical network on which NFVI support services such as Active Directory for authentication, DNS for name resolution, NTP for time synchronization and SMTP for email notifications are connected.

These are shared components used by both the management components and the ESXi hosts. This network is to be routed at the physical network with the “Management VLAN” and “ESXI Management” networks so that these services can be consumed.

2.1.2.6 Network Summary

Table 12 lists the indicative VLAN IDs and port groups and their functions for the three clusters. The network names and VLAN ID are for the purpose of describing the network architecture in this document; replace these with the actual datacenter network configuration.

VLAN ID	Function	Port Group	Network Type
Management Cluster			
20	ESXi management	ESXi Management	Infrastructure
30	vSphere vMotion	vMotion Network	Infrastructure
40	ScaleIO Virtual SAN	Scale IO Network	Infrastructure
50	Management VLAN	Management VLAN	Tenant
60	vSphere Replication	Replication Network	Infrastructure
80	External Network	External Network	Tenant
Edge Cluster			
20	ESXi management	ESXi Management	Infrastructure
30	vSphere vMotion	vMotion Network	Infrastructure
40	ScaleIO Virtual SAN	Scale IO Network	Infrastructure
50	Management VLAN	Management VLAN	Tenant
Resource Cluster			
25	ESXi management	ESXi Management	Infrastructure
35	vSphere vMotion	vMotion Network	Infrastructure
45	ScaleIO Virtual SAN	Scale IO Network	Infrastructure
70 (VTEP)	VXLAN Transport (NSX Overlay Network)	Created when VXLAN software switch is created	Tenant

Table 12 VLAN IDs, Function, and Port Groups

2.1.3 Storage Design

This section discusses the design for shared storage solution based on EMC ScaleIO. The ESXi hosts in all the clusters are connected to a dedicated VLAN for EMC ScaleIO SAN traffic. The ScaleIO components (MDM, SDS, and SDC), and an iSCSI target, are installed on dedicated ScaleIO virtual machines (SVMs). The SDS adds the ESXi hosts to the ScaleIO to be used for storage, thus enabling the creation of volumes. Using iSCSI targets, the volumes are exposed to the ESXi via an iSCSI adapter. ScaleIO volumes must be mapped both to the SDC and to iSCSI initiators. This ensures that only authorized ESXis can see the targets. Reliability is enhanced by enabling multipathing, either automatically or manually. Before starting to deploy ScaleIO, ensure that the following prerequisites are satisfied:

- The management network and Virtual Machine Port Group on all the ESXis that are part of the ScaleIO system must be configured.
- Devices that are to be added to SDS must be free of partitions.
- One datastore is created from one of the local devices for all the ESXis. This datastore is needed when deploying SVMs

ScaleIO supports the following network configurations –

- A single data storage network
- Two or more data networks, each on separate IP subnets
- A single IP data network using several NIC-bonding configurations, or vSwitch load balancing

Each host in the ScaleIO virtual SAN cluster must be configured with a VMkernel port group and a Virtual Machine port group and enabled for ScaleIO virtual SAN on the “Infrastructure vDS” distributed switch. Since the VDS share uplinks, NIOC will be used to prioritize ScaleIO virtual SAN storage traffic.

ScaleIO components are designed to work with a minimum of three server nodes. When all SDSs in a Protection Domain have one HDD and one SSD associated with them then two storage pools should be defined – High performance storage pool consisting of SSD drives for latency sensitive workloads and Capacity storage pool consisting of HDD drives for non-sensitive workloads.

VMware recommends that when the disks are connected to a RAID controller, each disk must be configured as a standalone RAID-0.

Two formats of storage can be used with the solution – virtual machine disk (VMDK) or raw device mapping (RDM). This solution recommends using VMDK or VMFS based storage to leverage all the benefits of vsphere.

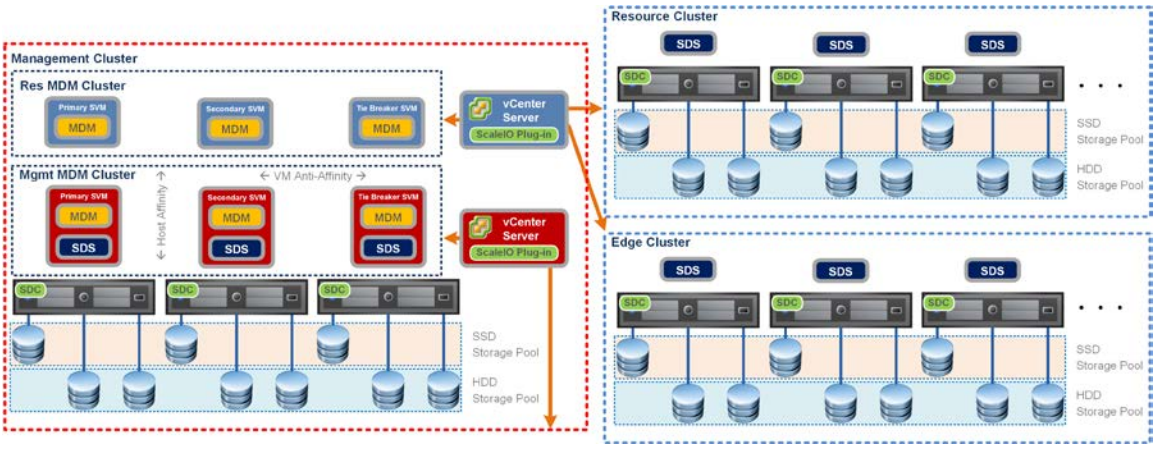

Figure 8

In order to configure Scale IO, these steps need to be followed -

- Prepare the ScaleIO environment by configuring each ESXi host in the cluster
- Register the ScaleIO plug-in to the vSphere Web Client
- Upload the OVA template to the ESXi host
- Deploy the ScaleIO system from the vSphere Web Client using the ScaleIO plug-in
- Create volumes with required capacity from the ScaleIO system and map the volumes to the ESXi hosts
- Create datastores by scanning the ScaleIO LUNs from ESXi hosts
- Install the ScaleIO GUI to manage the system

For more details, refer to EMC ScaleIO 2.0 user guide.

VMware recommends enabling sparse VM swap files feature for efficient utilization of available usable storage capacity. This feature ensures that VM swap files are created as sparse files instead of thick provisioned. This advanced setting needs to be set on each ESXi host that is in the VSAN cluster. The setting is called SwapThickPrivisionDisabled, and is disabled by default.

2.2 Virtualized Infrastructure Manager Design

2.2.1 vCenter Server

For this reference architecture, two vCenter server instances are deployed, each with a pair of external load balanced PSCs. One vCenter instance is used for managing the management cluster while the other vCenter instance is used for managing the resource and edge cluster. The two PSCs of each vCenter Server are linked and joined to a single SSO domain. The PSCs of the management vCenter Server are not linked to the PSCs of the resource vCenter Server. The vCenter Server instance hosting the resource cluster is used in the configuration of vCloud Director. The two PSC pairs are load balanced by a single NSX Edge in a highly available configuration.

The vCenter Server will use the embedded database, which reduces management overhead and licensing costs of deploying an external database and its supported operating system. vSphere HA will be used as the redundancy mechanism for vCenter server.

Figure 9 shows the vCenter Server deployment in the management cluster.

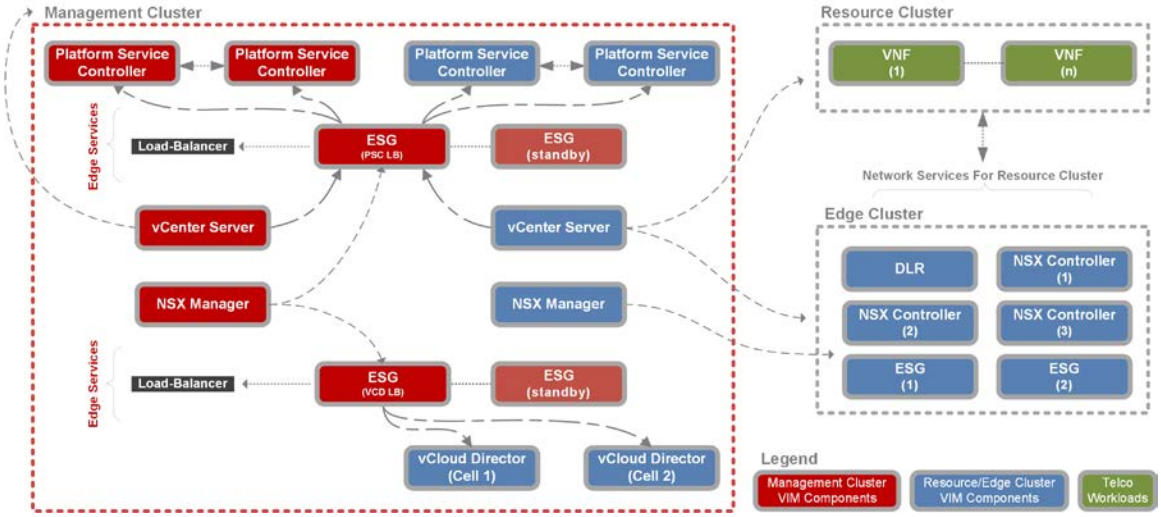

Figure 9 Management Cluster VIM Components

VMware recommends using the vCenter Server virtual appliance; this is pre-configured and enables faster deployment. The Windows based vCenter Server with external PSC can also be used; however, this document refers to the appliance-based deployment.

The resource cluster can be sized and scaled independently of the management cluster where the number and size (and thus capacity) are more or less fixed. vCenter Server appliance supports up to 1,000 hosts or 10,000 virtual machines at full vCenter Server scale.

The management cluster can be scaled up while the Edge and resource clusters can be scaled out to increase capacity as needed. Scaling decisions will need to be evaluated carefully to ensure there are no bottlenecks at other infrastructure components.

The deployment size for the management cluster vCenter Server instance is selected as small, this supports up to 150 hosts and 3,000 virtual machines. This allows for handling the current management cluster workloads as well as accommodates future scalability. The resource cluster vCenter Server deployment size is selected as large, this supports up to 1,000 hosts and 10,000 virtual machines.

2.2.2 NSX for vSphere Manager

Two NSX Manager instances are deployed and connected to respective vCenter server instances. One is for the management cluster itself while the other is for the network services of the VNFs deployed in the resource cluster. The two NSX Managers use the management network “Management VLAN” for their communication.

The NSX Manager instance deployed for the management cluster is responsible for deploying NSX Edge for load balancing the PSCs and another for load balancing the vCloud Director cells.

The second NSX Manager instance is used to deploy NSX Edge and DLRs for the VNFs. While the VNFs are deployed in the resource cluster, the network components are deployed in the edge cluster.

vSphere HA will be used as the redundancy mechanism for the NSX Manager instances. NSX Manager has a built-in service to schedule backups for configuration information to an FTP/SFTP server. Figure 10 shows the NSX Manager instances and their components.

2.2.3 vCloud Director

vCloud Director is deployed in a clustered configuration by using two vCloud Director instances for redundancy. After the first instance is deployed a response file is created, this is then used to deploy the second vCloud Director cell. Each vCloud Director instance is referred to as a cell, both vCloud cells point to the same vCenter Server used for managing the resource cluster. NSX Edge is used to provide the load balancing service for the vCloud Director cells.

To facilitate file transfers in a multi-cell environment, a shared storage volume must be configured and made accessible to all cells in a vCloud Director Server group. The necessary volume size varies on the expected number of concurrent uploads. Following an upload, the data moves to the designated organization virtual datacenter and the data no longer resides on the NFS volume. The recommended starting size for the NFS transfer volume is 250GB. Transferred images can be large, so monitor this volume and increase the size if necessary. Figure 10 shows the vCloud Director cell deployment architecture.

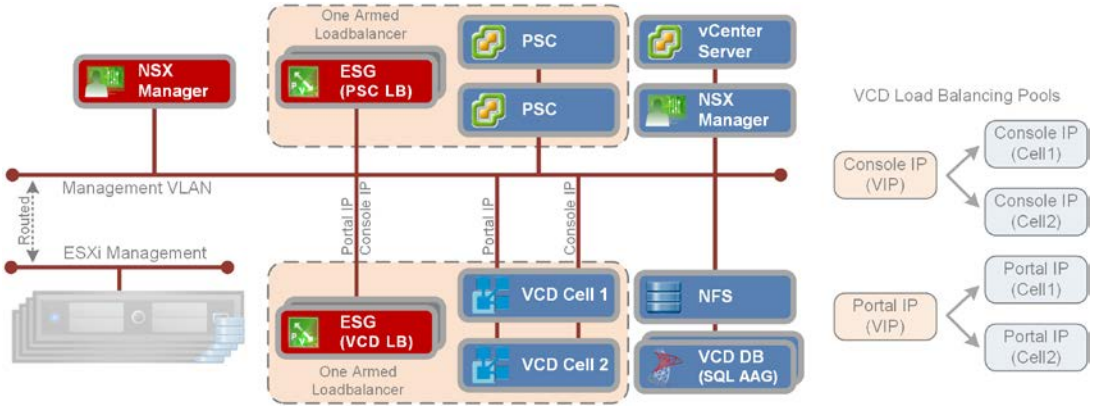

Figure 10 vCloud Director Cells Architecture

The vCloud Director cells are deployed with an external highly available clustered Microsoft SQL database. VMware recommends configuring Microsoft SQL database along with Always Availability Group for database high availability.

Each vCloud Director cell requires two IP addresses, one for the VM Console access and the other for the vCloud Director portal access. Both of these are load balanced across the two

cells. When VNFs are deployed in vCloud Director, they are connected directly to the external vCloud Director network. The external vCloud Director network is managed by NSX and employs dynamic routing to create a highly available network for the VNFs. In case there is any failure of an NSX Edge, traffic is immediately routed through the remaining NSX Edges. The ECMP OSPF Peering ensures that bandwidth utilization across the active NSX Edges is optimal while ensuring high availability of the network routing infrastructure. As the infrastructure scales, additional NSX Edges can be deployed to cater to the increased workload.

2.3 Operations Management Design

2.3.1 vRealize Operations Manager

The vRealize Operations Manager appliance is deployed in a master-replica configuration for high availability. The appliance has all the services required by vRealize Operations Manager hence allows for an architecture that can be scaled easily by adding additional instances.

The appliance deployment size is selected as small with four vCPUs, 16 GB RAM and 84 GB HDD storage space. This size assumes a data retention period of 6 months for 50 VMs, 12 hosts and 3 datastores. VMware recommends sizing the appliance as per the exact data retention requirements using the vRealize Operations Manager sizing guide[5].

First, a single master node is deployed then a second replica node is deployed to form the cluster. The data is replicated and switch over happens automatically in case the master fails. Anti-affinity rules ensure that the nodes are always deployed on separate hosts.

VMware vRealize Operations Management Pack listed under the monitoring section in this document are installed to retrieve various performance and health parameters from the vCloud NFVI platform. If additional Management Packs are installed, the resource requirements of the vRealize Operations Manager appliance may need to be increased.

2.4 vRealize Log Insight

VMware recommends deploying one vRealize Log Insight master node and two worker nodes. This gives the best performance and high availability configuration. The integrated load balancer of the cluster is enabled and used to ensure that load is balanced fairly amongst the available nodes. All the nodes should be deployed on the same Layer2 network and clients should point to the FQDN of the load balancer.

The initial vRealize Log Insight appliance deployment size is kept at default with 132 GB of disk space provisioned. 100 GB of the disk space is used to store raw data. The vRealize Log Insight appliance should be sized based on the IOPS, syslog connections and events per

[5] vRealize Operation Manager 6.1 Sizing Guidelines: http://kb.vmware.com/kb/2130551

second. For more details on sizing the appliance, refer to the vRealize Log Insight sizing guide[6].

Additional sizing considerations, such as the number of vSphere vCenter Servers supported by a single instance of vRealize Log Insight are documented in the vRealize Log Insight Configuration Limits[7].

[6] Sizing the Log Insight virtual appliance: http://pubs.vmware.com/log-insight-30/topic/com.vmware.log-insight.getting-started.doc/GUID-284FC5F4-B832-47A7-912E-D407A760CAE4.html
[7] vRealize Log Insight Configuration Limits: http://pubs.vmware.com/log-insight-33/topic/com.vmware.log-insight.administration.doc/GUID-0601A373-4B74-4B93-8C39-DA85F1D34FD4.html

3 Backup and Restore

3.1 vSphere Data Protection

This section of the document covers the backup and recovery of the management components of the vCloud NFV platform. For the purpose of this reference architecture, this document will cover vSphere Data Protection (VDP) as the backup solution however; supported third Party backup solutions may also be used.

The VDP appliance is deployed on a separate datastore than the Scale IO virtual SAN datastore of the protected workloads in the management cluster. The appliance is connected to the management VLAN for communication with the Management vCenter Server. Connectivity through vCenter Server provides vSphere Data Protection with visibility to all VMware ESXi servers, and therefore to the virtual machines that must be backed up. The VMware vSphere web client interface is used to select, schedule, configure and manage backups and recoveries of virtual machines.

Logical Design

Figure 11 shows the logical design of the vSphere Data Protection solution.

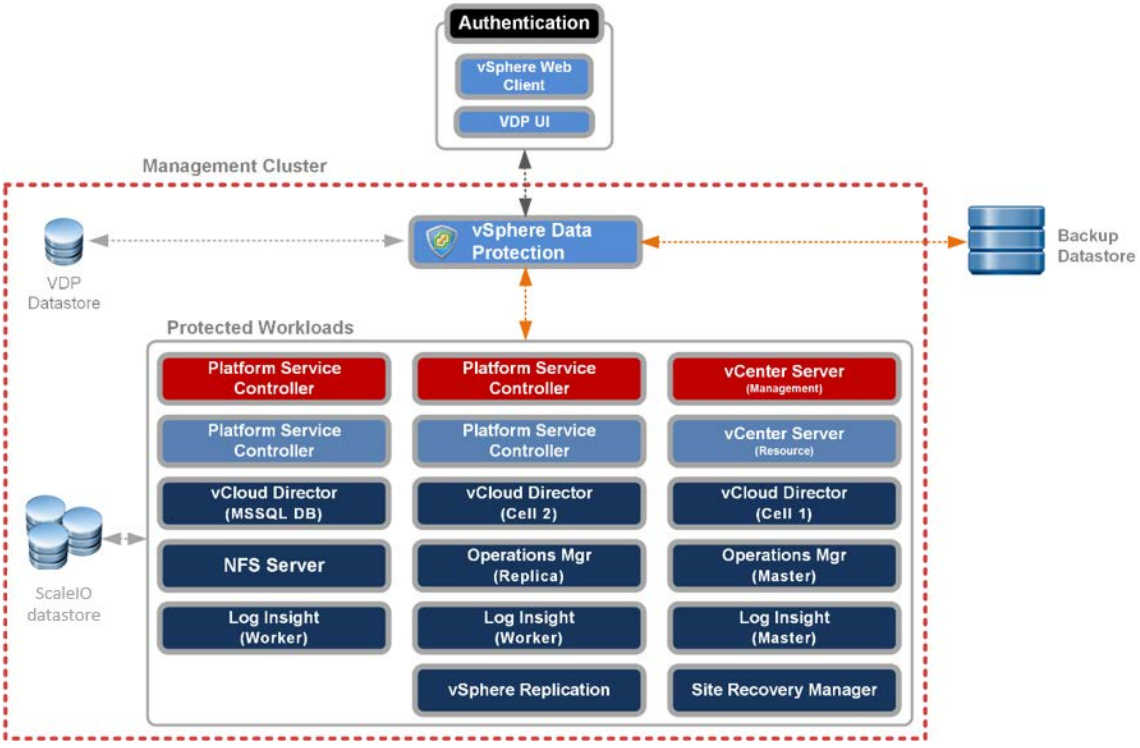

Figure 11 vSphere Data Protection Logical Design

The VDP appliance communicates with the vCenter server to make a snapshot of a virtual machine's .vmdk files. Deduplication takes place within the appliance by using a variable-length deduplication technology. To increase the efficiency of image level backups, VDP

utilizes the Changed Block Tracking[8] (CBT) feature, which greatly reduces the backup time of a given virtual machine image and provides the ability to process a large number of virtual machines within a particular backup window.

Storage Design

The backup datastore stores all the production data that is required in a disaster recovery event or data loss to recover the backup up management components based on a recovery point objective (RPO).

It is important to choose the target location and meet the minimum performance requirements to mitigate such a scenario. There are two options when choosing the target storage location,

- Option 1: Store Backup Data on the Same ScaleIO virtual SAN Datastore

  Simpler to manage with no dependency on storage administrator

  Can take full advantage of vSphere capabilities

  The risk is that if the destination datastore is unrecoverable, you will lose the ability to recover your data.

- Option 2: Store Backup Data on a Dedicated Storage

  If the ScaleIO virtual SAN storage becomes unavailable, you can recover your data, because your backup data is not located on the same-shared storage.

  Separate management and backup workloads.

  The backup schedule does not impact the management cluster storage performance, because the backup storage is completely separated.

  The drawback is that you might be required to install and configure a dedicated storage volume for backups

vSphere Data Protection generates a significant amount of I/O, especially when performing multiple, concurrent backups. The storage platform must be able to handle this I/O. If the storage does not meet the performance requirements, it is possible for backup failures to occur and for error messages to be generated. VMware recommends using a separate dedicated storage volume for best performance.

Backup Policies

VMware recommends using the HotAdd transport mechanism for faster backups and restores and less exposure to network routing, firewall and SSL certificate issues when taking image backups of entire virtual machine. Refer to the HotAdd best practices[9] for more details.

---

[8] Enable CBT on VMs : http://kb.vmware.com/kb/1020128
[9] HotAdd Best Practices : http://kb.vmware.com/kb/2048138

Even when vSphere Data Protection uses Changed Block Tracking technology to optimize the success rate to back up data, it is crucial to avoid any window where the management components storage is in high demand to avoid any business impact.

The retention policies are the properties of a backup job, therefore it is important to group virtual machines by business priorities and the retention requirements set by the business level. For this reference architecture, vSphere Data Protection will only backup the management components deployed in the management cluster.

The section below lists the vCloud NFV management components and their backup strategies.

- **ESXi Hosts** – The ESXi hosts are not backed up, instead their configuration data can be exported and imported back on a newly installed server. Alternatively, host profiles may be used to restore the configuration of the hosts to their initial configured state.

  Reference: http://kb.vmware.com/kb/2042141

- **vCenter Server with External Platform Services Controller** – The vCloud NFV platform uses a pair of load balanced platform services controller instances for each vCenter server. The PSC instances are replicated while the vCenter server has an embedded database and points to the PSC load balancer virtual IP. The vCenter server and its corresponding PSCs must be backed up at the same time. If all the components fail at the same time, the PSC must be restored first. vSphere Data Protection is used to take a full image level backup of both the PSCs and vCenter Server.

  Reference: http://kb.vmware.com/kb/2110294, http://pubs.vmware.com/vsphere-60/topic/com.vmware.vsphere.install.doc/GUID-7223A0A9-9E3C-4093-8121-9BD8B3DB793F.html

- **NSX Manager** – The NSX manager has a built in backup and restore mechanism. All the configuration data can be backed up on a schedule to an FTP server. The NSX Manager backup contains all of the NSX configuration, including controllers, logical switching and routing entities, security, firewall rules, and everything else that you configure within the NSX Manager UI or API. The vCenter database and related elements like the virtual switches need to be backed up separately.

  Reference: http://pubs.vmware.com/NSX-62/topic/com.vmware.nsx.admin.doc/GUID-72EFCAB1-0B10-4007-A44C-09D38CD960D3.html

- **vCloud Director** – vCloud Director introduces new backup and restore challenges because of the additional layer of abstraction and the relationships with objects in vSphere. Coordinated backups between vCenter Server and vCloud Director are needed to ensure complete restoration of vCloud objects.

The vCloud Director database resides on an MS SQL database in the management cluster. The vCloud Director Database server, vCloud Director cells and NFS server share will be backed using the same schedule as vCenter Server and NSX Manager.

When restoring vCloud Director, the database must be restored before the vCloud Director cells, and in sync with vCenter to avoid inconsistencies. It is important that the database be kept in a consistent state.

- **vRealize Operations Manager** – vRealize Operations Manager single or multi-node clusters can be backed up and restored by using vSphere Data Protection or other backup tools. You can perform full, differential and incremental backups and restores of virtual machines. All nodes need to be backed up and restored at the same time.

  Reference: http://pubs.vmware.com/vrealizeoperationsmanager-62/topic/com.vmware.vcom.core.doc/GUID-315B1700-0536-4037-AB14-56EC27EA8970.html

- **vRealize Log Insight** – The vCloud NFV platform utilizes a three-node vRealize Log Insight cluster. The entire cluster needs to be backed up and restored at the same time. For detailed information, see the link below.

  Reference: https://pubs.vmware.com/log-insight-30/topic/com.vmware.log-insight.administration.doc/GUID-2BD249BB-24A5-4FEF-9AE0-765B88EE2FFE.html

- **vSphere Replication** – The vSphere Replication appliance is backed up using vSphere Data Protection using an image level backup of the entire appliance. When an image is restored and the appliance powered on, the data replication resumes after a few minutes.

- **Site Recovery Manager** – The Site Recovery Manager instance is deployed on a Windows machine along with an embedded database where all the configuration information is stored. This database can be backed up and restored as detailed in the link below.

  Reference: http://pubs.vmware.com/srm-61/topic/com.vmware.srm.install_config.doc/GUID-E1FC67CD-48E3-4B25-AA1D-8A2408F5B517.html

- **vSphere Data Protection** – The vSphere Data Protection appliance has a checkpoint and rollback mechanism built in. By default, VDP keeps two system checkpoints. If you roll back to a checkpoint, all backups and configuration changes taken since the checkpoint was taken will be lost when the rollback is completed. The first checkpoint is created when VDP is installed. Subsequent checkpoints are created by the Maintenance service. This service is disabled for the first 24 to 48 hours of VDP operation.

Reference: http://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vmware-data-protection-administration-guide-61.pdf

Monitoring

CPU, memory, network and disk performance and capacity will be monitored by vRealize Operations Manager. Events and log information are sent to vRealize Log Insight. Capacity can also be viewed via the Reports tab/VDP capacity.

4 Capacity Planning & Sizing

4.1 Sizing Guidelines

4.1.1 Cluster Sizing

The following section lists the cluster configuration for each of the three clusters. In addition to this, the sizing considerations for each of the management components already explained in the design sections needs to be factored into the calculations.

- **Management Cluster** – The management cluster sizing is based on the number of management nodes in the cluster. This has to take into account factors such as multi node topologies of management components e.g. components with external databases, PSC location (internal / external) etc. The management cluster capacity needs to be proportionately increased where components have high availability requirements e.g. active-active NSX controllers, active-passive NSX Edge, HA pairs etc.

  In addition to this, vCenter Server uses admission control to ensure that sufficient resources are available in a cluster to provide failover protection and to ensure that virtual machine resource reservations are respected. Admission control imposes constraints on resource usage and violation will prevent virtual machines from being powered on.

  Other factors such as allowed host failures and number of hosts, which can be taken offline for maintenance also need to be factored when planning for capacity.

  It is recommended to have all the hosts in the cluster with the same specification and preferably same vendor / model / family. This will result in a balanced cluster, and the load will be balanced across the cluster in a predictable way.

  Last, but not the least, resource overheads of virtualization required by the hypervisor need to be taken into account. As a best practice, it is not recommended to design the architecture with overcommitted resources for the management nodes.

- **Edge Cluster** – The edge cluster is used to host the NSX network infrastructure for the North-South VNF networks. The NSX components deployed in the edge cluster are:

  NSX Controllers – From a high availability point of view, three NSX Controllers are deployed together and capacity for all three needs to be accounted for in the clusters in which they are deployed. Note that the NSX Controller nodes are deployed by the NSX Manager and their size is not customizable.

  Distributed Logical Router – Distributed logical routers are deployed as active-passive high availability pairs by the NSX Manager. The size of the DLR is not user selectable.

NSX Edge – NSX Edges are deployed in an active-active configuration with ECMP OSPF peering for high availability. Multiple NSX Edge instances can be deployed for performance. For stateful services such as load balancing, NSX Edge is deployed as an active-standby pair.

The number of NSX devices in the edge cluster will be determined by the expected North-South VNF network traffic. This document provides guidelines for one set of such devices; these can be scaled as per the requirements of the VNF workloads.

- **Resource Cluster** – The resource cluster is used to deploy the VNFs. Since it is impossible to calculate the VNF sizes before knowing which VNFs will be deployed, this document suggests an approach to arrive at the resource cluster size once the VNF size is known. VNF vendors are expected to provide the sizing requirements for their VNFs.

  This guideline aims to build a scalable architecture that can be scaled at all tiers as additional VNFs are deployed without compromising the functional aspects of the platform.

  If the VNF sizing requirements are not known, tools such as VMware Capacity Planner can be used to analyze the Telco workloads and make a suitable sizing recommendation.

  The vCloud NFV platform uses an Elastic Provider vDC configuration in vCloud Director, which means that additional ESXi host clusters could be added when there is a need to scale the cluster's capacity. This is a key business benefit as it allows one to optimize TCO/ROI and gradually scale investment over a period of time.

  The number of hosts in the cluster also needs to take into account operational aspects of the workloads for example when a host is placed in maintenance mode.

4.1.2 Storage Sizing

For sizing the ScaleIO virtual SAN datastore of each cluster, we need to identify the workloads in each cluster and add their disk requirements. Some buffer should be added to the overall capacity for operational tasks such as for VM templates, cloning, swap files, log files, snapshots etc. These will depend on the specific VNF configuration and requirements during VNF onboarding.

The minimum configuration required for Scale IO virtual SAN is three ESXi hosts. However, this smallest environment has important restrictions. In a three-node MDM cluster - Master MDM, Slave MDM and Tiebreaker - there are two copies of the repository and can withstand one MDM cluster failure. However, in order to avoid a single point of failure a five-node MDM cluster – Master MDM, two Slave MDM and two Tiebreaker – has three copies of the repository and can withstand two MDM cluster failure.

The Scale IO virtual SAN cluster can be configured to tolerate a number of host failures; this makes a mirror copy of the data in the datastore resulting in reduced usable capacity.

For the purpose of this reference architecture, VMware recommends four nodes in each cluster. This ensures there are enough nodes to meet the availability requirements and to allow for a rebuild of the components after a failure. For detailed guidelines, refer to the Scale IO user guide and installation guide.

As best practice, VMware recommends having a dedicated storage for backups. The sizing of this storage depends on the number of virtual machines to be backed up, frequency of backup, retention policy and amount of changed blocks since last backup.

4.2 Sizing Design

4.2.1 Management Cluster

Table 15 lists the management components along with their clustered nodes that are deployed in the management cluster.

System	vCPU	Memory	Storage	Description
Platform Services Controller	2	2 GB	30 GB	For management cluster vCenter Server
Platform Services Controller	2	2 GB	30 GB	For management cluster vCenter Server
Platform Services Controller	2	2 GB	30 GB	For resource cluster vCenter Server
Platform Services Controller	2	2 GB	30 GB	For resource cluster vCenter Server
vCenter Server (small)	4	16 GB	82 GB	Management cluster
vCenter Server (large)	16	32 GB	82 GB	Resource and edge cluster
NSX Manager	4	16 GB	60 GB	For management cluster workloads
NSX Manager	4	16 GB	60 GB	For edge and resource cluster workloads
Edge Service Gateway (active) (Large)	2	1 GB	2.15 GB	Active Load Balancer for VCD cells
Edge Service Gateway (standby) (Large)	2	1 GB	2.15 GB	Standby Load Balancer for VCD cells
Edge Service Gateway (active) (Large)	2	1 GB	2.15 GB	Active Load Balancer for management and resource PSCs
Edge Service Gateway (standby) (Large)	2	1 GB	2.15 GB	Standby Load Balancer for management and resource PSCs
Log Insight (master)	4	8 GB	132 GB	Deployed in management cluster
Log Insight (replica)	4	8 GB	132 GB	Deployed in management cluster

Log Insight (replica)	4	8 GB	132 GB	Deployed in management cluster
Operations Manager (active)	4	16 GB	84 GB	Deployed in management cluster
Operations Manager (passive)	4	16 GB	84 GB	Deployed in management cluster
vCloud Director (cell1)	2	4 GB	22 GB	Deployed in management cluster
vCloud Director (cell2)	2	4 GB	22 GB	Deployed in management cluster
vCloud Director Database (AAG master)	4	16 GB	250 GB	Shared database for vCloud Director Cells (MSSQL)
vCloud Director Database (AAG replica)	4	16 GB	250 GB	Shared database for vCloud Director Cells (MSSQL)
NFS Server	2	4 GB	10 GB	vCloud Director shared transfer space
vSphere Data Protection	4	10 GB	6 TB	Space includes backup data storage space
Primary MDM – Mgmt	1	1 GB	8 GB	SVM deployed in management cluster
Secondary MDM – Mgmt	1	1 GB	8 GB	SVM deployed in management cluster
Tiebreaker - Mgmt	1	1 GB	8 GB	SVM deployed in management cluster
Primary MDM – Res/Edge	1	1 GB	8 GB	SVM deployed in management cluster
Secondary MDM – Res/Edge	1	1 GB	8 GB	SVM deployed in management cluster
Tiebreaker – Res/Edge	1	1 GB	8 GB	SVM deployed in management cluster
TOTAL	88	208 GB	7.55 TB	

Table 13 Management Cluster Sizing

4.2.2 Edge Cluster

Considering that VNF workloads are network intensive we deploy the X-Large configuration for the Edge Service Gateway. Three of these are deployed with ECMP configuration for high availability.

Table 14 Edge Cluster Sizing

System	vCPU	Memory	Storage	Description
NSX Controller 1	4	4 GB	22 GB	For resource cluster workloads
NSX Controller 2	4	4 GB	22 GB	For resource cluster workloads
NSX Controller 3	4	4 GB	22 GB	For resource cluster workloads
Distributed Logical Router (Active)	1	512 MB	1.15 GB	For resource cluster workloads
Distributed Logical Router (Passive)	1	512 MB	1.15 GB	For resource cluster workloads
Edge Service Gateway (active) X-Large	4	1 GB	2.15 GB	For resource cluster workloads (ECMP)
Edge Service Gateway (active) X-Large	4	1 GB	2.15 GB	For resource cluster workloads (ECMP)
Edge Service Gateway (active) X-Large	4	1 GB	2.15 GB	For resource cluster workloads (ECMP)
SVM – SDS	1	1 GB	8 GB	SVM deployed on edge cluster
SVM – SDS	1	1 GB	8 GB	SVM deployed on edge cluster
SVM – SDS	1	1 GB	8 GB	SVM deployed on edge cluster
TOTAL	29	19 GB	99 GB	

Table 14 lists the NSX Edge devices deployed in the edge cluster. These should cater to most VNFs however should be sized appropriately as per the VNF workload requirements in the resource cluster. As VNF workload requirements increase, both the edge cluster capacity and the NSX Edge devices deployed in it can be scaled.

4.2.3 Resource Cluster

The sizing requirements for the VNFs will be captured as part of the VNF On-boarding process. NEPs are expected to provide the entire infrastructure requirement details of a VNF so that resource requirements can be planned and sized as required. The vCloud NFV platform allows scaling the resource cluster to accommodate VNF capacity requirements.

5 VNF Onboarding

This section lists the key considerations when onboarding VNFs on to the vCloud NFV platform. Since each VNF has its own specific requirements, these guidelines may need to be adjusted as required.

5.1 Capacity Requirements

The first step is to identify the capacity requirements of the VNF. The target NFVI needs to have sufficient capacity for the successful deployment and operation of the VNF. Some of the points to be taken into account when planning for the NFVI capacity required for hosting a VNF are:

- Number of active nodes of the VNF
- Nodes required for high availability and redundancy
- VNFs deployed for scaling for transient and seasonal peaks
- Capacity reservation for failover nodes
- Capacity required for operational overheads
- Capacity reserved for future growth

NEPs are expected to provide the above information for their VNFs before actual onboarding so that capacity can be planned and provided for in advance.

5.2 Resource Requirements

Once the capacity in terms of number of nodes has been identified, the resource requirements of these nodes need to be assessed. Each category of node may have its own resource requirement e.g. nodes used for scaling may be sized differently than the primary nodes.

In addition to the resource requirements of the VNFs, the vCloud NFV platform needs to be configured for the efficient allocation of resources. Resource allocation is done by first creating one or more provider VDCs to aggregate hardware resources of the NFVI. OrgVDCs are then used to allocate the pooled resources to the VNFs. The choice of whether to deploy multiple VNFs in a single orgVDC or have separate orgVDCs etc. will need to be determined by the requirements provided by the NEP for the particular VNF.

NEPs should provide data about number of CPU cores, RAM and storage required by the VNF.

5.3 Operational Requirements

When onboarding VNFs, their operational aspects also need to be considered in terms of capacity and performance requirements. Some of the operational tasks that may impact the capacity availability and need to be taken into account at the time of VNF onboarding are:

- Space for temporary machine images such as snapshots
- Space for maintenance and backup operations
- Log retention policies
- Scale up and scale out for VNF nodes
- VNF templates in catalog

A key operation is the deployment of VNFs from the catalog. This may be either for deploying a new VNF instance or for scaling existing VNF. vCloud Director allows creating catalogs and creating vApp templates in the catalog. A vApp is a group of related VMs along with all configuration information required by the application deployed inside the vApp. Deploying a vApp template from catalog is the fastest way to deploy one or more VNFs. This can be automated using the vCloud Director APIs.

VNF Managers can make use of the vCloud Director APIs to dynamically scale the VNFs by deploying additional VNFs from the catalog. The catalog can contain a vApp template for an entire VNF stack or vApps with individual VNF components. This gives the flexibility to scale specific nodes of the VNF rather than deploying the entire stack.

5.4 High Availability Requirements

The high availability requirement of the deployed VNF workloads needs to be factored when calculating capacity requirements and capacity availability of the NFVI platform. The size, quantity and type (active/passive) of the redundant nodes influences storage, network and compute resource consumption.

VNF Workloads may have anti affinity requirements such as ensuring that nodes of an HA group are never placed on the same physical server. This puts additional constraints on resource management.

A new feature of vCloud Director is the ability of tenants to set affinity / anti-affinity rules for VNFs deployed in their vCloud Organization. Because of this, service providers need to closely monitor the resource utilization and ensure both high availability requirements and resource requirements of VNFs are met.

In the case of disaster recovery, the VNFs may not exist, however enough resources need to be reserved to ensure that when VNFs are failed over, sufficient resources are available to allow them to be powered on. This is done by ensuring sufficient resource reservations are in place at both the resource cluster as well as vCloud Director.

5.5 Security Requirements

When VNFs are on boarded, necessary security requirements need to be identified. Multi-tenancy is achieved by creating organizations in vCloud Director; Resource partitioning can

be addressed by separating the VNFs in vCloud Director by placing them in separate containers such as providerVDCs and orgVDCs.

For network security, the vCloud NFV platform supports various network configurations such as creation for firewall rules for both east west and north-south traffic. Firewall rules can be configured at both the perimeter as well as internally by micro-segmentation using NSX. Depending on the network requirements, isolated networks may be created for VNF interconnects and segregation of workload traffic based on role e.g. management traffic and data traffic.

5.6 Network Requirements

The most important factors when on boarding a VNF is its networking requirements. The vCloud NFV platform is designed to be flexible to meet the network needs of the VNFs being deployed. Some of the network considerations when onboarding a VNF are:

- Network routing between nodes
- External / MPLS network connectivity
- Network isolation and segregation requirements
- Firewall configuration
- WAN / VPN connectivity
- Connectivity to management systems
- Latency and bandwidth requirements

5.7 VMware Tools Requirement

VMware Tools is a suite of utilities installed in the VNF that enhances the performance of the virtual machines guest operating system and improves management of the virtual machine. Without VMware Tools installed in your guest operating system, guest performance lacks important functionality. VMware recommends installing VMware Tools to eliminate or improves these issues:

- Ability to take quiesced snapshots of the VNF
- Synchronize guest OS time with ESXi host
- Support for guest-bound calls created with the VMware VIX API
- Device drivers to optimize mouse operations and improve sound, graphics and networking performance.
- Guest OS customization support from within vCloud Director
- Scripting that helps automate guest operating system operations
- vRealize Operations Manager end point operations

Although the guest operating system can run without VMware Tools, many VMware features are not available until you install VMware Tools. For example, if you do not have VMware Tools installed in your virtual machine, you cannot use the shutdown or restart options from

the toolbar. You can use only the power options. For an overview of VMware Tools please refer to the KB340[10].

5.8 Onboarding Process

A sequence of important steps involved in VNF onboarding process is provided below:

- Setup connectivity to external networks
- Deploy / configure NSX Edge for north-south traffic
- Configure transit networks
- Deploy / configure logical routers for north-south traffic
- Create logical switch for east-west traffic
- Create vCloud Director organization
- Create catalog for the organization
- Set OVF environment parameters, if any
- Import the VNF into the catalog
- Review VNF sizing and network parameters in the catalog
- Create organizationVDC and allocate resources
- Deploy VNF as a vApp from catalog
- Review network mapping and connectivity of VNF
- Configure anti-affinity rules if required in vCloud Director
- Power ON VNFs
- Review end-to-end network connectivity

[10] Overview of VMware Tools : http://kb.vmware.com/kb/340

6 Supporting Components

The vCloud NFV platform relies on the following supporting components.

- **Directory Service** - This is used to provide single-sign-on authentication services to all the management components.
- **DNS** - This is to provide a domain naming service for all the management components. Some management components such as vCloud Director require its host name to be fully (forward and reverse lookup) resolvable.
- **NTP** - As a best practice, all the management components, including the physical ESXi servers should have their clocks synchronized to a single time source. VMware recommends a dedicated NTP server be setup for this purpose.
- **SMTP** - An SMTP server is valuable in sending email notifications from the platform to various administration teams. This is best used with vRealize Operations Manager to send email notifications based on advanced metrics, events and health parameters.

7 Monitoring & Logging

The monitoring and logging blocks of the vCloud NFV platform is comprised of vRealize Log Insight and vRealize Operations Manager. CPU, memory, network, and disk performance and capacity are monitored by vRealize Operations Manager and syslog events are sent to vRealize Log Insight. vRealize Log Insight Content Packs give it the ability to parse product specific information from the respective product logs. Similarly, vRealize Operations Manager Management Packs allow product specific metrics, alerts, events and dashboards to be used for monitoring. These are available for download from VMware Solution Exchange.

7.1 Logging

vRealize Log Insight is the central repository for log files from all the management components. vRealize Log Insight in turn uses content packs to parse the log files to extract information. This information from the log files can in turn be used as a data source for vRealize Operations Management. Table 15 lists the vCloud NFV platform components and logs that are sent to vRealize Log Insight.

Refer to the “vRealize Operations Management Pack for Log Insight” installation and configuration guide for instructions on how to install and configure the management pack for vRealize Log Insight. This management pack allows vRealize Operations Manager to query Log Insight for logs of the selected inventory object. For this, the Launch in Context for Log Insight feature needs to be enabled as described in the document.

System	Logs	Content Pack
ESXi	syslog host.d vpxa.log	vSphere Content Pack for Log Insight
vCenter Server	syslog events, tasks, alarms	vSphere Content Pack for Log Insight
EMC ScaleIO	logs, traces	
NSX Manager	NSX Manager syslogs Distributed firewall logs NSX Edge syslog	NSX for vSphere
vCloud Director	syslog Apache log4j	vCloud Director
vSphere Data Protection	syslog	none
Microsoft Windows 2012		Microsoft Windows
Microsoft SQL Server OR Oracle Database		Microsoft SQL Content Pack for Oracle Databases
vRealize Orchestrator	syslog	none
vRealize Operations Manager	syslog	vRealize Operations Manager
Site Recovery Manager	SRM Server Logs	none
vSphere Replication	syslog	none

Table 15 Component Log Collection

If logs are to be retained for a longer period of time, an external NFS share can be mounted for log archiving. The size of this share would need to be determined based on the logs to be archived and the period of retention.

ScaleIO Logging -

- Retrieve ESXi logs – Collect logs from these folders –
  - /var/log
  - /scratch/log
- Retrieve ScaleIO components logs – Log in to SVM and run the following script for each component –
  - /opt/emc/scaleio/<scaleio component>/diag/get_info.sh -f
    where <scaleio component> is mdm or sds

Refer to ScaleIO 2.0 user guide for more details about collecting logs

7.2 Monitoring

vRealize Operations Manager is used as the monitoring component for the vCloud NFV platform. The following management packs are used to gather metrics for management components and provide out of the box events, alerts and dashboards.

Detailed information about installing, configuring and using these management packs are provided in the documentation accompanying the management packs and are not repeated here.

Management Pack	Version	Description
vCenter Server	Bundled	Monitors entire data center including clusters, hosts, storage, networking, virtual machines etc.
NSX for vSphere	3.0.2	Discover, analyze and graphically represent the broad number of virtual networking services available within NSX for vSphere. Quickly identify configuration, health or capacity problems within virtual NSX networks
vCloud Director	4.0	This adapter monitors the health of supported vCloud Director entities and sends early warning smart alerts for monitored Provider vCloud Director resources
vRealize Log Insight	1.0.1	Access the unstructured log data for any component of RA environment by allowing to launch in context
Network Devices	1.0.1	Not only provides insight into virtual network layer but also provides information of the physical layer network devices like switches and routers
OpenStack	1.5	Provides out of the box dashboards, reports, inventory views, and alerts complete with remediation actions for comprehensive operational capabilities for managing an OpenStack environment

Table 16 Monitoring Management Packs

vROps will monitor SVM in general not application level. In case of any issues, storage admin can drill down further by login into the ScaleIO GUI. The ScaleIO Graphical User Interface (GUI) enables you to review the overall status of the system, drill down to the component level and monitor these components. The various screens display different views and data that are beneficial to the storage administrator. ScaleIO GUI provides an interface to monitor the underlying ScaleIO components. There are multiple monitoring areas on the GUI:

- Dashboard
- Protection Domain
- Protection Domain servers
- Storage Pool

For more details, refer to the EMC ScaleIO 2.0 user guide

7.2.1 Metrics

The following section lists the metric categories that are collected by these management packs.

System	Metric Categories / Resource Kinds
vCenter Server	VC Server Resources - CPU usage, disk, memory, network, and summary metrics VC Datacenter - CPU usage, disk, memory, network, storage, disk space, and summary metrics VC Compute Resources - CPU usage, configuration, storage, disk space, disk, memory, network, power, and summary metrics VC Resource Pools - Configuration, CPU usage, memory, and summary metrics Host Systems - Configuration, Hardware, agent, CPU usage, datastore, disk, memory, network, storage, and summary metrics Virtual Machines - Configuration, CPU use, memory, datastore, disk, virtual disk, guest file system, network, power, disk space, storage, and summary metrics Datastores - Capacity, device, and summary metrics
NSX for vSphere	NSX Manager NSX Controller Cluster NSX Controllers NSX Transport Zone NSX Logical Router NSX Edge NSX Edge - DHCP Service NSX Edge - DNS Service NSX Edge - Firewall Service NSX Edge - IPSec VPN Service NSX Edge - L2 VPN Service NSX Edge - Load Balancer Service NSX Edge - NAT Service NSX Edge - Routing Service NSX Edge - SSL VPN Service Top of Rack Switch

	Physical Fabric
vCloud Director	Cloud Cell Organization Organization VDC Provider VDC Organization VDC Network Provider VDC Storage Policy Organization VDC Storage Policy vCloud Datastore vCloud External Network vCloud Network Pool vCloud Host vCloud vCenter vApp vCloud Virtual Machine

Table 17 Monitoring Metrics

7.2.2 Dashboards

Table 18 lists the default dashboards that are provided out of the box with the respective component management packs.

System	Dashboards
NSX for vSphere	NSX Main NSX Logical Topology NSX Object Path NSX Edge Services
vCloud Director	VCD All Metric Selector VCD Organization vDC Utilization VCD vApp Utilization VCD Mashup charts VCD Alerts VCD Troubleshooting

Table 18 Monitoring Dashboards

8 High Availability

The vCloud NFV platform architecture and design ensures that there is no single point of failure for any of the components that make up the platform. High availability of the individual components has already been covered in earlier sections. This section describes the high availability considerations in the key architectural blocks of the platform.

8.1 NFV Infrastructure

The first step in building a highly available platform is to ensure that all the hardware components are redundant. In the case of the vCloud NFV platform, this includes the physical switches and physical compute servers.

The physical switches are connected in a redundant mesh configuration to ensure there is more than one path between two points. This ensures that physical switch failures can be tolerated.

Each of the compute servers should have redundant Ethernet ports connected to the physical network in link aggregation mode.

Since this reference architecture uses Virtual SAN for storage, this simplifies the storage network topology. If third party storage devices are deployed, then the similar high availability requirements of the storage network and devices needs to be taken into account.

Each cluster has four ESXi hosts at minimum to provide the requisite capacity and redundancy for both compute and ScaleIO Virtual SAN storage resources.

ScaleIO uses a mirroring scheme to protect data against disk and node failures. When an SDS node or SDS disk fails, applications can continue to access ScaleIO volumes and the data is still available through the remaining mirrors. ScaleIO starts a rebuild process that creates another mirror for the data chunks that were lost in the failure. The surviving SDS cluster nodes carry out the rebuild process by using the aggregated disk and network bandwidth of the cluster.

8.2 Virtualized Infrastructure Manager

All the components of the virtualized infrastructure manager have high availability configuration. The NSX Manager for the management cluster is used to deploy one-armed load balancers to load balance management traffic for the vCloud Director cells and another load balancer for the PSCs. This allows immediate redirecting of traffic to the standby node in case the primary node fails.

In addition to this, vSphere HA is used to bring up the failed node on another ESXi host. Anti-affinity rules ensure that two HA nodes of the same management component are never on the same physical host. Each cluster has enough resources reserved to tolerate one host failure. The resource reservation is taken into account when planning the capacity of the clusters and ensures that there are sufficient resources to power on a failed node on another ESXi host should the need arise without impacting the performance of the nodes already

running on the ESXi host. Management components such as NSX manager and vCenter server are protected using vSphere HA.

8.3 Operations Management

The Operations Management block is comprised of vRealize Operations Manager and vRealize Log Insight. Like the virtualized infrastructure manager block, both of these components are deployed in an active-active load balanced configuration.

A Physical

Setup	Hostname	MAC	Service Tag	Device Type	IP	Vlan
Hawk	Management		HKGK0Z1	Dell Networking S4048T	172.16.101.4	Untagged
Hawk	Spine1		4NCK0Z1	Dell Networking S6010	172.16.101.5	Untagged
Hawk	Spine2		8PCK0Z1	Dell Networking S6010	172.16.101.6	Untagged
Hawk	Leaf1		5LCK0Z1	Dell Networking S6010	172.16.101.7	Untagged
Hawk	Leaf2		HNCK0Z1	Dell Networking S6010	172.16.101.8	Untagged

Table 19 Physical: Dell Networking S4048T and S6010

Setup	Hostname	MAC	Service Tag	Device Type	IP	Vlan	Hypervisor Mgmt IP	Vlan	User name	Password
Hawk	VIM1	44:A8:42:26:BD:61	6KQ4F42	PowerEdge R730	172.16.104.10	2104	172.16.105.10	U 2105	root	dellnfv
Hawk	VIM2	44:A8:42:07:FC:0B	6KP3F42	PowerEdge R730	172.16.104.11	2104	172.16.105.11	U 2105	root	dellnfv
Hawk	VIM3	44:A8:42:26:BE:6D	6KG4F42	PowerEdge R730	172.16.104.12	2104	172.16.105.12	U 2105	root	dellnfv
Hawk	Compute1	44:A8:42:26:C2:0E	6KJ6F42	PowerEdge R730	172.16.104.13	2104	172.16.105.13	U 2105	root	dellnfv
Hawk	Compute2	44:A8:42:26:BC:E7	6KL6F42	PowerEdge R730	172.16.104.14	2104	172.16.105.14	U 2105	root	dellnfv
Hawk	Compute3	44:A8:42:22:C2:D6	6KQ5F42	PowerEdge R730	172.16.104.15	2104	172.16.105.15	U 2105	root	dellnfv
Hawk	Edge1	44:A8:42:07:FC:23	6KM3F42	PowerEdge R730	172.16.104.16	2104	172.16.105.16	U 2105	root	dellnfv
Hawk	Edge2	44:A8:42:22:C5:B8	6KQ6F42	PowerEdge R730	172.16.104.17	2104	172.16.105.17	U 2105	root	dellnfv
Hawk	Edge3	44:A8:42:26:BB:07	6KN6F42	PowerEdge R730	172.16.104.18	2104	172.16.105.18	U 2105	root	dellnfv

Table 20 Physical: PowerEdge R730

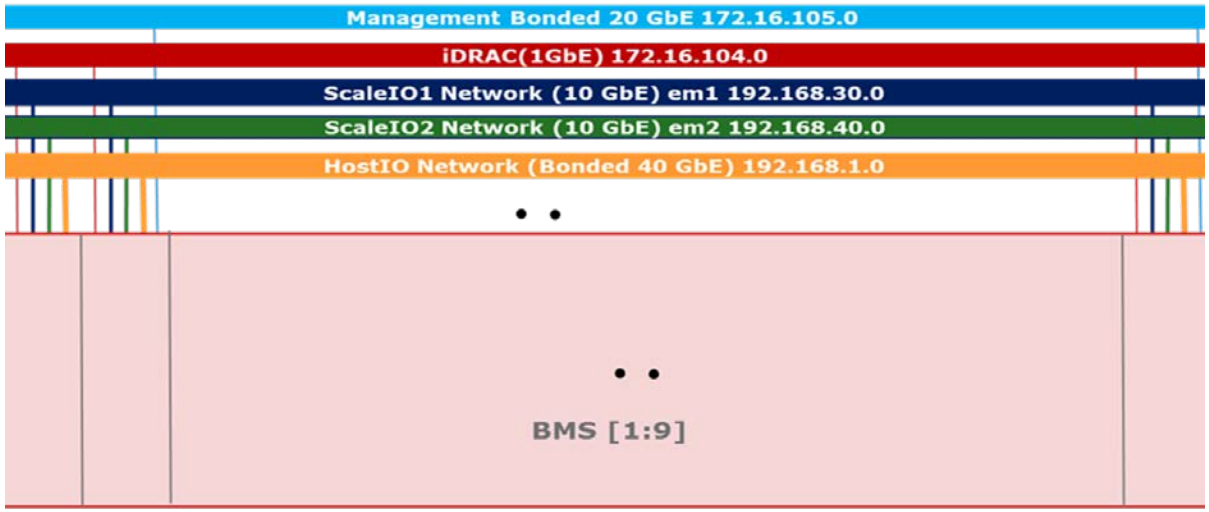
Management Bonded 20 GbE 172.16.105.0
iDRAC(1GbE) 172.16.104.0
ScaleIO1 Network (10 GbE) em1 192.168.30.0
ScaleIO2 Network (10 GbE) em2 192.168.40.0
HostIO Network (Bonded 40 GbE) 192.168.1.0
BMS [1:9]

B Rack

Model	CPU	Memory	Storage	Network	Boot from SD Card
R730	E5-2697 v3 2.60GHz 14c	128G RDIMM 2133 MHz	2x600G SAS, 1x200G SSD	6xIntel X520, Brcm 5720 embedded	2xSD
R730	E5-2697 v3 2.60GHz 14c	128G RDIMM 2133 MHz	2x600G SAS, 1x200G SSD	6xIntel X520, Brcm 5720 embedded	2xSD
R730	E5-2697 v3 2.60GHz 14c	128G RDIMM 2133 MHz	2x600G SAS, 1x200G SSD	6xIntel X520, Brcm 5720 embedded	2xSD
R730	E5-2697 v3 2.60GHz 14c	128G RDIMM 2133 MHz	2x600G SAS, 1x400G SSD	6xIntel X520, Brcm 5720 embedded	2xSD
R730	E5-2697 v3 2.60GHz 14c	128G RDIMM 2133 MHz	2x600G SAS, 1x400G SSD	6xIntel X520, Brcm 5720 embedded	2xSD
R730	E5-2697 v3 2.60GHz 14c	128G RDIMM 2133 MHz	2x600G SAS, 1x400G SSD	6xIntel X520, Brcm 5720 embedded	2xSD
R730	E5-2697 v3 2.60GHz 14c	128G RDIMM 2133 MHz	2x600G SAS, 1x200G SSD	6xIntel X520, Brcm 5720 embedded	2xSD
R730	E5-2697 v3 2.60GHz 14c	128G RDIMM 2133 MHz	2x600G SAS, 1x200G SSD	6xIntel X520, Brcm 5720 embedded	2xSD
R730	E5-2697 v3 2.60GHz 14c	128G RDIMM 2133 MHz	2x600G SAS, 1x200G SSD	6xIntel X520, Brcm 5720 embedded	2xSD
Dell Networking S4048T		Management Switch			
Dell Networking S6010		Spine			
Dell Networking S6010		Spine			
Dell Networking S6010		Leaf/ToR			
Dell Networking S6010		Leaf/ToR			
		vCloud NFV v1.5			
	vCenter - VIM and Compute	6.0.0 Update 2 Build 3634788			
	VIM ESXi	6.0.0 Update 2 Build 3620759			
	Compute ESXi	6.0.0 Update 2 Build 3620759			
	Edge ESXi	6.0.0 Update 2 Build 3620759			
	NSX Manager	6.2.2 Build 3604087			
	vRealize Operations Manager	6.2.1 Build 3774215			
	vRealize Log Insight	3.3.1 Build 3644329			
	vCloud Director	8.1 Build 3879706			
	ScaleIO	2.0.2			
	vSphere Replication	6.1.1			
	vSphere Data Protection	6.1.2			
	Site Recovery Manager	6.1.1			

Table 21 Rack

C VMs

Cluster	Name	IP Assigned	Status	Hostname Given	Username cli / webclient	Password	Remarks
	Jumphost	172.16.105.20	Completed	jumphost	Administrator	Admin123	
	Jumphost-b	172.16.105.19	Completed	jumphost-b	Administrator	Admin123	Clone of Jumphost
VIM/MGMT	External PSC	172.16.105.21	Completed	mgmt01psc01	root / Administrator@ vsphere.local	Dellnfv1!	Hosts: 172.16.105.21 mgmt01psc01.dellemc.local
	vCenter VIM	172.16.105.22	Completed	mgmt01vc01	root / Administrator@ vsphere.local	Dellnfv1!	Hosts: 172.16.105.22 mgmt01vc01.dellemc.local
	EPSC-Compute	172.16.105.23	Completed	mgmt01psc02	root / Administrator@ vsphere.local	Dellnfv1!	Hosts: 172.16.105.23 mgmt01psc02.dellemc.local
	vCenter Compute	172.16.105.24	Completed	mgmt01vc02	root / Administrator@ vsphere.local	Dellnfv1!	Hosts: 172.16.105.24 mgmt01psc02.dellemc.local
	NSX Manager-1	172.16.105.25	Completed	nsxmgr1	admin	dellnfv	NSX manager
	ESGPSCActive	172.16.105.100	Completed	esgpscact	admin	Nsxcontroller#3	Edge Services Gateway

	ESGPSCSTANDBY	172.16.105.100	Completed	esgpscsby	admin	Nsxcontroller#3	Distributed Logical Router
	ESGVCDActive	172.16.105.99	Completed	esgvcdact	admin	Nsxcontroller#3	Distributed Logical Router
	ESGVCDSTANDBY	172.16.105.99	Completed	esgvcdsby	admin	Nsxcontroller#3	Edge Services Gateway
	Windows Server 2012	172.16.105.50	Completed	win2k12addns	Administrator	Dellnfv1!	DNS
	Windows Server 2012 SQL Account	172.16.105.40	Completed	sql	vcdmgr	dellnfv	Database: vcddb Instance: MSSQLSERVER sa-Dellnfv1! Administrator-Dellnfv1!
	vClould RH6.5 Cell-1	172.16.105.51	Completed	vcdcell1	root	VMware1!	vCloud director
		172.16.105.52	Completed	vcdcell1-rc	root	VMware1!	
	vClould RH6.5 Cell-2	172.16.105.53	Completed	vcdcell2	root	VMware1!	
		172.16.105.54	Completed	vcdcell2-rc	root	VMware1!	
	vRealize-1	172.16.105.36	Completed	vropsmaster	root / admin	Dellnfv1!	
	vRealize-2	172.16.105.37	Completed	vropsreplica	root / admin		
	vRealize Loginsight3.0-1	172.16.105.33	Completed	vRealizeLoginsight-1	root / admin	Dellnfv1!	
	vRealize Loginsight3.0-2	172.16.105.34	Completed	vRealizeLoginsight-2	root / admin		
	vRealize Loginsight3.0-3	172.16.105.35	Completed	vRealizeLoginsight-3	root / admin		

	Ubuntu (Centos-6)	172.16.105.27	Completed		root	Dellnfv1!	
ScaleIO	ScaleIO gateway	172.16.105.45 **192.168.30.13** **192.168.40.13**	Completed		root / admin	Dellnfv1!	
	Master MDM	172.16.105.46 **192.168.30.14** **192.168.40.14**	Completed		root / admin	Dellnfv1!	
	Slave MDM	172.16.105.47 **192.168.30.15** **192.168.40.15**	Completed		root / admin	Dellnfv1!	
	Tie Breaker MDM	172.16.105.48 **192.168.30.16** **192.168.40.16**	Completed		root / admin	Dellnfv1!	
	SRM	172.16.105.112					
	Centos-Test vm	172.16.105.94	Completed				
	NTP	172.16.105.38	Completed				
	VRO	172.16.105.39	Completed				
	VDP	172.16.105.41	Completed				
Edge	NSX Manager-2	172.16.105.26	Completed		admin	dellnfv	NSX manager
	NSX Controller 1	172.16.105.42	Completed		admin	Nsxcontroller#3	Controllers
	NSX Controller 2	172.16.105.43	Completed		admin	Nsxcontroller#3	Controllers
	NSX Controller 3	172.16.105.44	Completed		admin	Nsxcontroller#3	Controllers

	DLR0-0	172.168.200.1 172.168.100.1 172.168.110.1	Completed		admin	Nsxcontroller#3	Distributed Logical Router
	DLR0-1	172.168.200.1 172.168.100.1 172.168.110.1	Completed		admin	Nsxcontroller#3	Distributed Logical Router
	ESG1-0	172.168.200.3 172.16.105.96 192.168.1.21	Completed		admin	Nsxcontroller#3	Edge Services Gateway
	ESG2-0	172.168.200.4 172.16.105.97 192.168.1.22	Completed		admin	Nsxcontroller#3	Edge Services Gateway
	ESG3-0	172.168.200.5 172.16.105.98 192.168.1.23	Completed		admin	Nsxcontroller#3	Edge Services Gateway
ScaleIO	SDS	172.16.105.64 **192.168.30.28** **192.168.40.28**	Completed				
	SDS	172.16.105.65 **192.168.30.29** **192.168.40.29**	Completed				
	SDS	172.16.105.66 **192.168.30.30** **192.168.40.30**	Completed				
Compute	MCM-7				admin	admin	Chassis 54 Node 7

	MCM-8				admin	admin	Chassis 54 Node 8
	Logical MCM				admin	admin	
	SSM-5						Chassis 54 Node 5
	SSM-6						Chassis 54 Node 6
	CSM-1						Chassis 54 Node 1
	CSM-2						Chassis 54 Node 2
	EMS				admin	Dellnfv1!	
ScaleIO	ScaleIO gateway	172.16.105.60 **192.168.30.21** **192.168.40.21**	Completed				
	Master MDM	172.16.105.61 **192.168.30.22** **192.168.40.22**	Completed				
	Slave MDM	172.16.105.62 **192.168.30.23** **192.168.40.23**	Completed				
	Tie Breaker MDM	172.16.105.63 **192.168.30.24** **192.168.40.24**	Completed				

Table 22 VMs

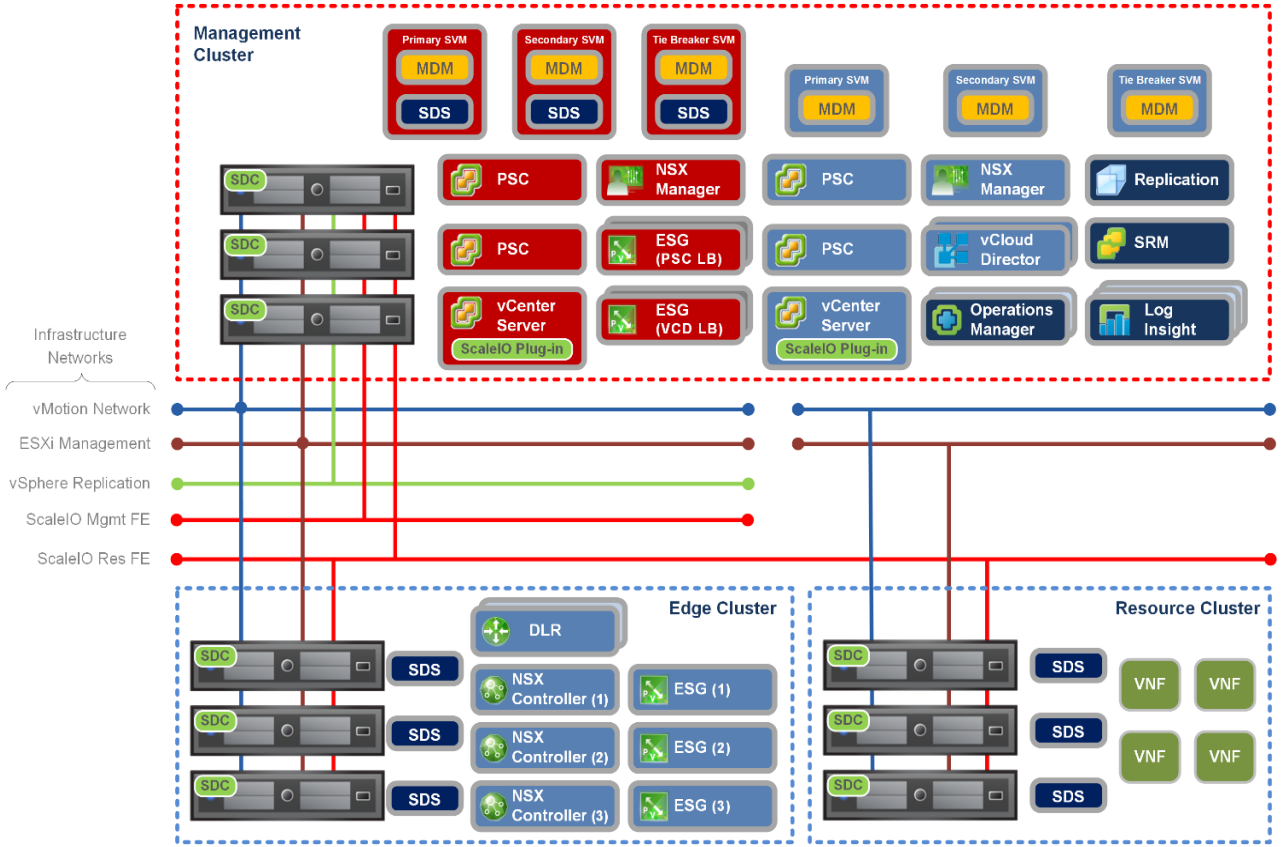

Figure 12 Management cluster

D Scale I/O systems

ScaleIO System -1		scaleio-mgmt				
vCentre	**Cluster**	**SDCs**	**ScaleIO Components (Total 4 SVM's)**	**Management IP**	**Data IP**	**UN/PW**
VC01(172.16.105.22)	**Mgmt Cluster**	ESX- 172.16.105.10	ScaleIO-GW	172.16.105.45	**192.168.30.13** **192.168.40.13**	root/Dellnfv1!
			MDM1 & SDS	172.16.105.46	**192.168.30.14** **192.168.40.14**	root/Dellnfv1!
		ESX- 172.16.105.11	**MDM2** & SDS	172.16.105.47	**192.168.30.15** **192.168.40.15**	root/Dellnfv1!
		ESX- 172.16.105.12	**TB1**& SDS	172.16.105.48	**192.168.30.16** **192.168.40.16**	root/Dellnfv1!
ScaleIO System -2		**scaleio-res**				
vCentre	**Cluster**	**SDCs**	**ScaleIO Components (Total 7 SVM's)**	**Management IP**	**Data IP**	**UN/PW**
VC02 (172.16.105.24)	**Resource Cluster**	ESX- 172.16.105.16	ScaleIO-GW	172.16.105.60	**192.168.30.21** **192.168.40.21**	root/Dellnfv1!
			MDM1 & SDS	172.16.105.61	**192.168.30.22** **192.168.40.22**	root/Dellnfv1!

		ESX- 172.16.105.17	**MDM2** & SDS	172.16.105.62	**192.168.30.23** **192.168.40.23**	root/Dellnfv1!
		ESX- 172.16.105.18	**TB1**& SDS	172.16.105.63	**192.168.30.24** **192.168.40.24**	root/Dellnfv1!
	Edge Cluster	ESX- 172.16.105.13	SDS	172.16.105.64	**192.168.30.28** **192.168.40.28**	root/Dellnfv1!
		ESX- 172.16.105.14	SDS	172.16.105.65	**192.168.30.29** **192.168.40.29**	root/Dellnfv1!
		ESX- 172.16.105.15	SDS	172.16.105.66	**192.168.30.30** **192.168.40.30**	root/Dellnfv1!

Table 23 Scale I/O systems

E Bonding

	vmnic#	nic port	switch	sw port	bond	
server1	4	p1p1	sw1	Te0/65		SIO-1
iDrac: 172.16.104.10	5	p1p2	sw2	Te0/65		SIO-2
	6	p2p1	sw1	Te0/64	po64	MGMT
	7	p2p2	sw2	Te0/64	po64	
	10	p4p1	sw1	Te0/81	po80	HostIO
	11	p4p2	sw2	Te0/81	po80	
	12	p5p1	sw1	Te0/80	po80	
	13	p5p2	sw2	Te0/80	po80	
server2	4	p1p1	sw1	Te0/67		SIO-1
iDrac: 172.16.104.11	5	p1p2	sw2	Te0/67		SIO-2
	6	p2p1	sw1	Te0/66	po66	MGMT
	7	p2p2	sw2	Te0/66	po66	
	10	p4p1	sw1	Te0/83	po82	HostIO
	11	p4p2	sw2	Te0/83	po82	
	12	p5p1	sw1	Te0/82	po82	

	13	p5p2	sw2	Te0/82	po82	
server3	4	p1p1	sw1	Te0/69		SIO-1
iDrac: 172.16.104.12	5	p1p2	sw2	Te0/69		SIO-2
	6	p2p1	sw1	Te0/68	po68	MGMT
	7	p2p2	sw2	Te0/68	po68	
	10	p4p1	sw1	Te0/85	po84	HostIO
	11	p4p2	sw2	Te0/85	po84	
	12	p5p1	sw1	Te0/84	po84	
	13	p5p2	sw2	Te0/84	po84	
server4	4	p1p1	sw1	Te0/71		SIO-1
iDrac: 172.16.104.13	5	p1p2	sw2	Te0/71		SIO-2
	6	p2p1	sw1	Te0/70	po70	MGMT
	7	p2p2	sw2	Te0/70	po70	
	10	p4p1	sw1	Te0/87	po86	HostIO
	11	p4p2	sw2	Te0/87	po86	
	12	p5p1	sw1	Te0/86	po86	

	13	p5p2	sw2	Te0/86	po86	
server5	4	p1p1	sw1	Te0/73		SIO-1
iDrac: 172.16.104.14	5	p1p2	sw2	Te0/73		SIO-2
	6	p2p1	sw1	Te0/72	po72	MGMT
	7	p2p2	sw2	Te0/72	po72	
	10	p4p1	sw1	Te0/89	po88	HostIO
	11	p4p2	sw2	Te0/89	po88	
	12	p5p1	sw1	Te0/88	po88	
	13	p5p2	sw2	Te0/88	po88	
server6	4	p1p1	sw1	Te0/75		SIO-1
iDrac: 172.16.104.15	5	p1p2	sw2	Te0/75		SIO-2
	6	p2p1	sw1	Te0/74	po74	MGMT
	7	p2p2	sw2	Te0/74	po74	
	10	p4p1	sw1	Te0/91	po90	HostIO
	11	p4p2	sw2	Te0/91	po90	
	12	p5p1	sw1	Te0/90	po90	

	13	p5p2	sw2	Te0/90	po90	
server7	4	p1p1	sw1	Te0/77		SIO-1
iDrac: 172.16.104.16	5	p1p2	sw2	Te0/77		SIO-2
	6	p2p1	sw1	Te0/76	po76	MGMT
	7	p2p2	sw2	Te0/76	po76	
	10	p4p1	sw1	Te0/93	po92	HostIO
	11	p4p2	sw2	Te0/93	po92	
	12	p5p1	sw1	Te0/92	po92	
	13	p5p2	sw2	Te0/92	po92	
server8	4	p1p1	sw1	Te0/79		SIO-1
iDrac: 172.16.104.17	5	p1p2	sw2	Te0/79		SIO-2
	6	p2p1	sw1	Te0/78	po78	MGMT
	7	p2p2	sw2	Te0/78	po78	
	10	p4p1	sw1	Te0/95	po94	HostIO
	11	p4p2	sw2	Te0/95	po94	
	12	p5p1	sw1	Te0/94	po94	

	13	p5p2	sw2	Te0/94	po94	
server9	4	p1p1	sw1	Te0/97		SIO-1
iDrac: 172.16.104.18	5	p1p2	sw2	Te0/97		SIO-2
	6	p2p1	sw1	Te0/96	po96	MGMT
	7	p2p2	sw2	Te0/96	po96	
	10	p4p1	sw1	Te0/99	po98	HostIO
	11	p4p2	sw2	Te0/99	po98	
	12	p5p1	sw1	Te0/98	po98	
	13	p5p2	sw2	Te0/98	po98	

Table 24 Bonding

F Miscellaneous Used IPs

IP Address/Range	Use
172.16.105.69-172.16.105.93	Used for VxLAN
172.16.105.101-172.16.105.110	Used for VxLAN
172.16.105.111	IP for Loginsight HA
172.16.105.30	vmotion mgmt 10
172.16.105.31	vmotion mgmt 11
172.16.105.32	vmotion mgmt 12

Table 25 Miscellaneous used IPs

G Wiring

Legend

*= 40/10 GbE Breakout (Range: x-x+4)
17-25 GbE connection (for setup ONLY)
Cat 6 Ethernet Cable
Optical Cables 10BgE to Switch1
Optical Cables 10GbE to Switch2
Uplink Twinax 10Gbe
40 Gbe Link
10 Gbe

42	Management Switch: http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/Dell-Force10-S4048T-SpecSheet.pdf																								42
	1	1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47	
	M	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48	
41	Switch1: http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/Dell_Networking_S6010_Spec_Sheet.pdf																								41
	M	1*	9*	17*	25*	33*	41*	49*	57*	65*	73*	81*	89*	97*	105	113*	121*	3								
	M	5*	13*	21*	29*	37*	45*	53*	61*	69*	77*	85*	93*	101*	109*	117*	125*	3								
40	Switch2: http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/Dell_Networking_S6010_Spec_Sheet.pdf																								40
	M	1*	9*	17*	25*	33*	41*	49*	57*	65*	73*	81*	89*	97*	105	113*	121*	4								
	M	5*	13*	21*	29*	37*	45*	53*	61*	69*	77*	85*	93*	101*	109*	117*	125*	4								
39																									39

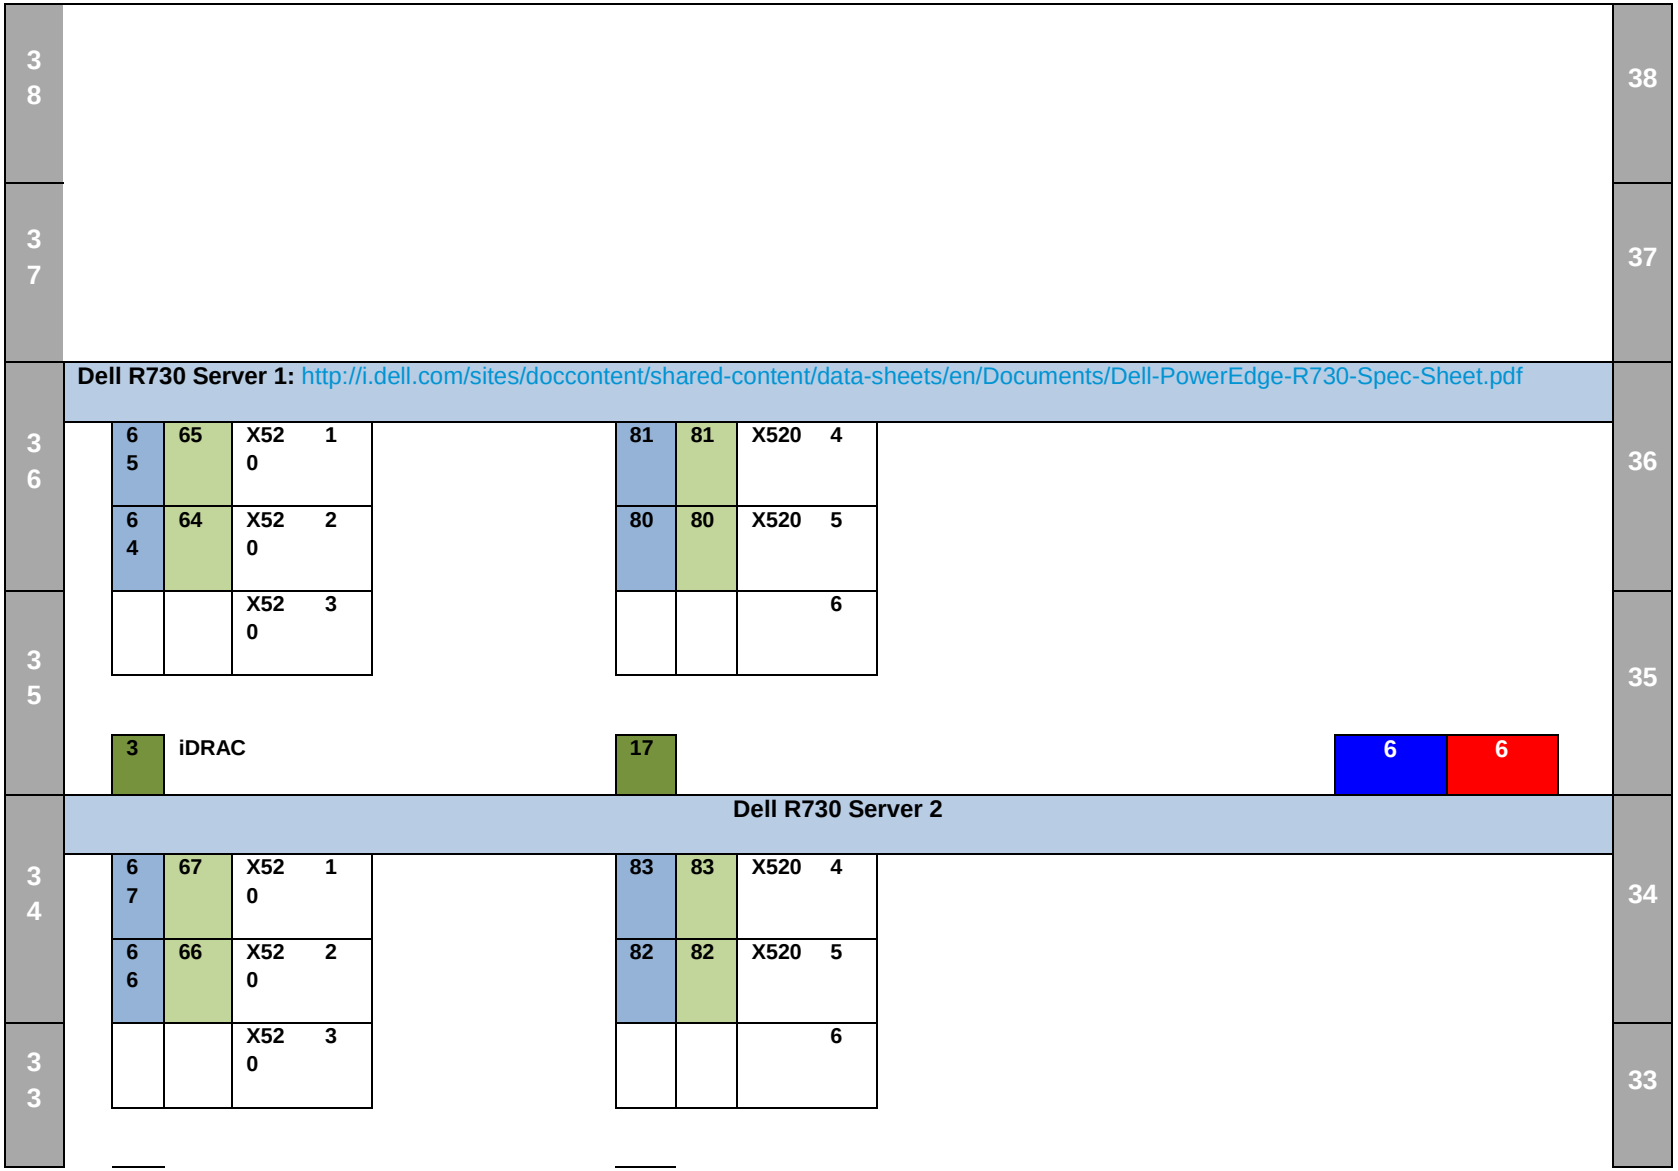
38
37
36
35
34
33
Dell R730 Server 1: http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/Dell-PowerEdge-R730-Spec-Sheet.pdf
65 65 X520 1
64 64 X520 2
X520 3
81 81 X520 4
80 80 X520 5
6
3 iDRAC
17
6 6
Dell R730 Server 2
67 67 X520 1
66 66 X520 2
X520 3
83 83 X520 4
82 82 X520 5
6

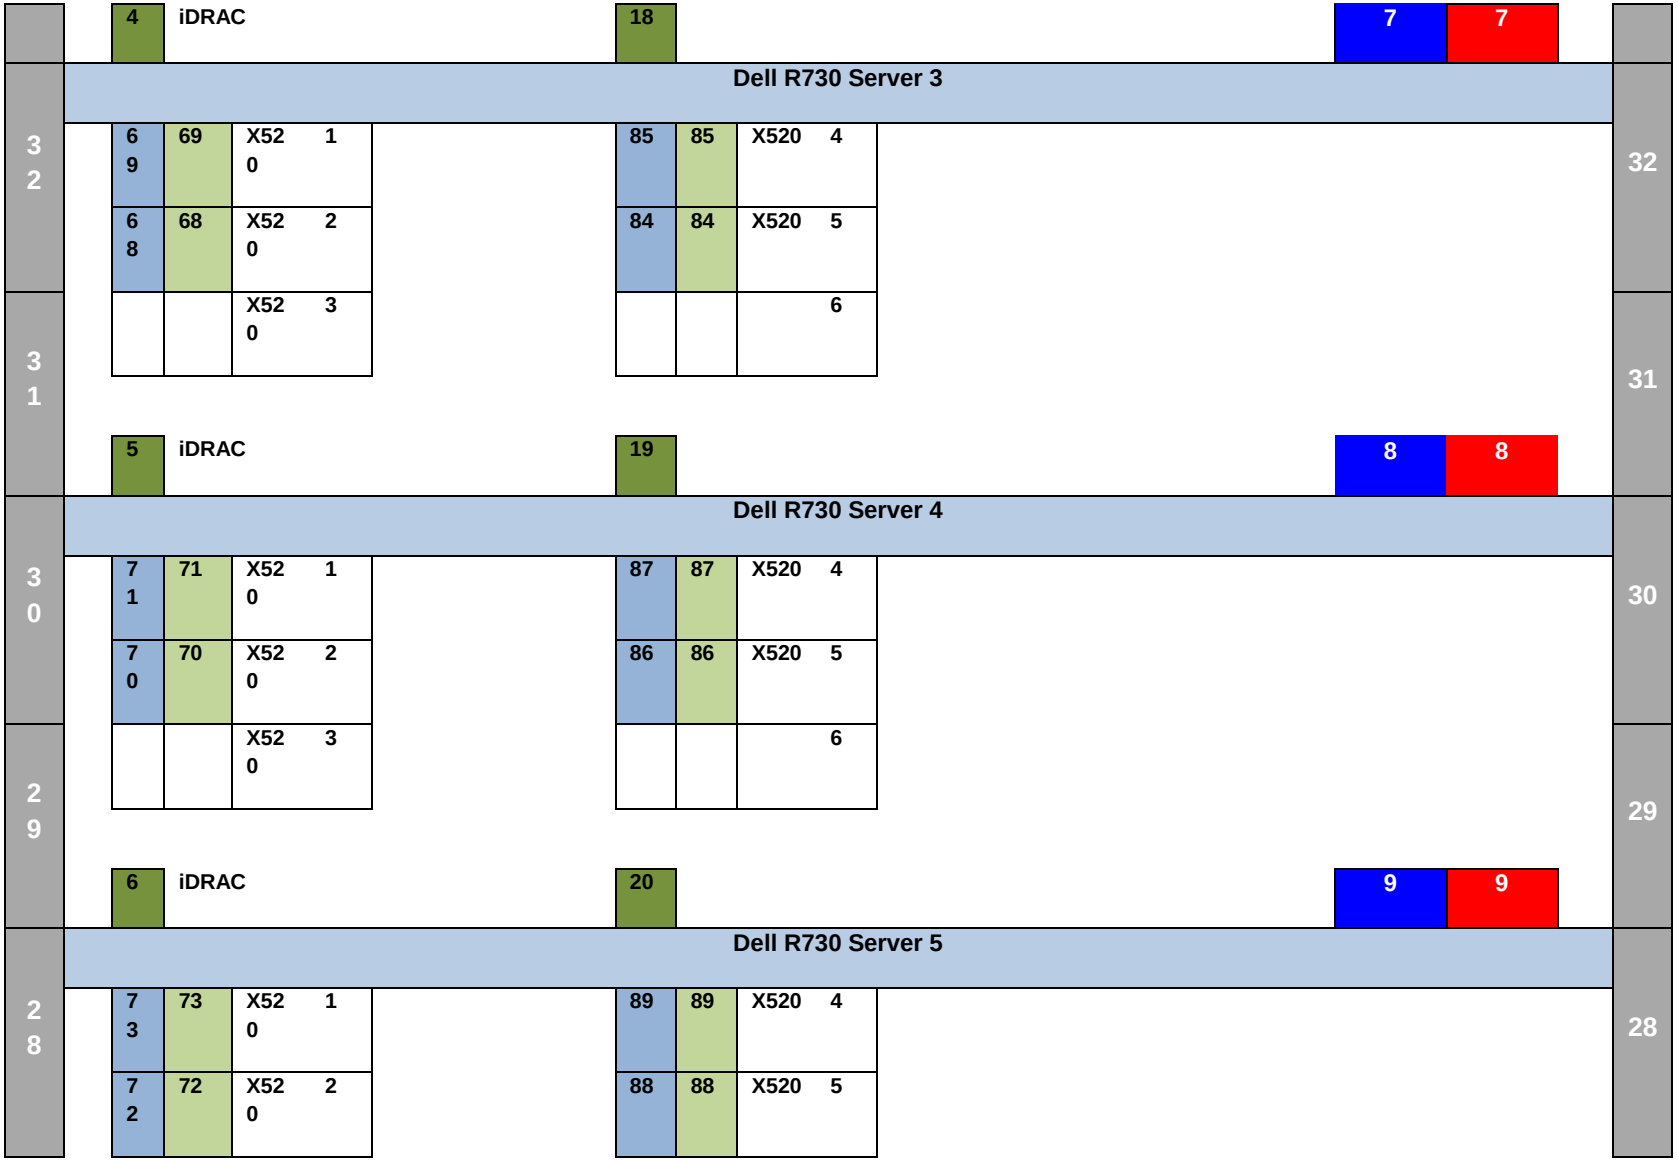
4 iDRAC
18
7 7
Dell R730 Server 3
32
31
69 69 X520 1
68 68 X520 2
X520 3
85 85 X520 4
84 84 X520 5
6
5 iDRAC
19
8 8
Dell R730 Server 4
30
29
71 71 X520 1
70 70 X520 2
X520 3
87 87 X520 4
86 86 X520 5
6
6 iDRAC
20
9 9
Dell R730 Server 5
28
73 73 X520 1
72 72 X520 2
89 89 X520 4
88 88 X520 5

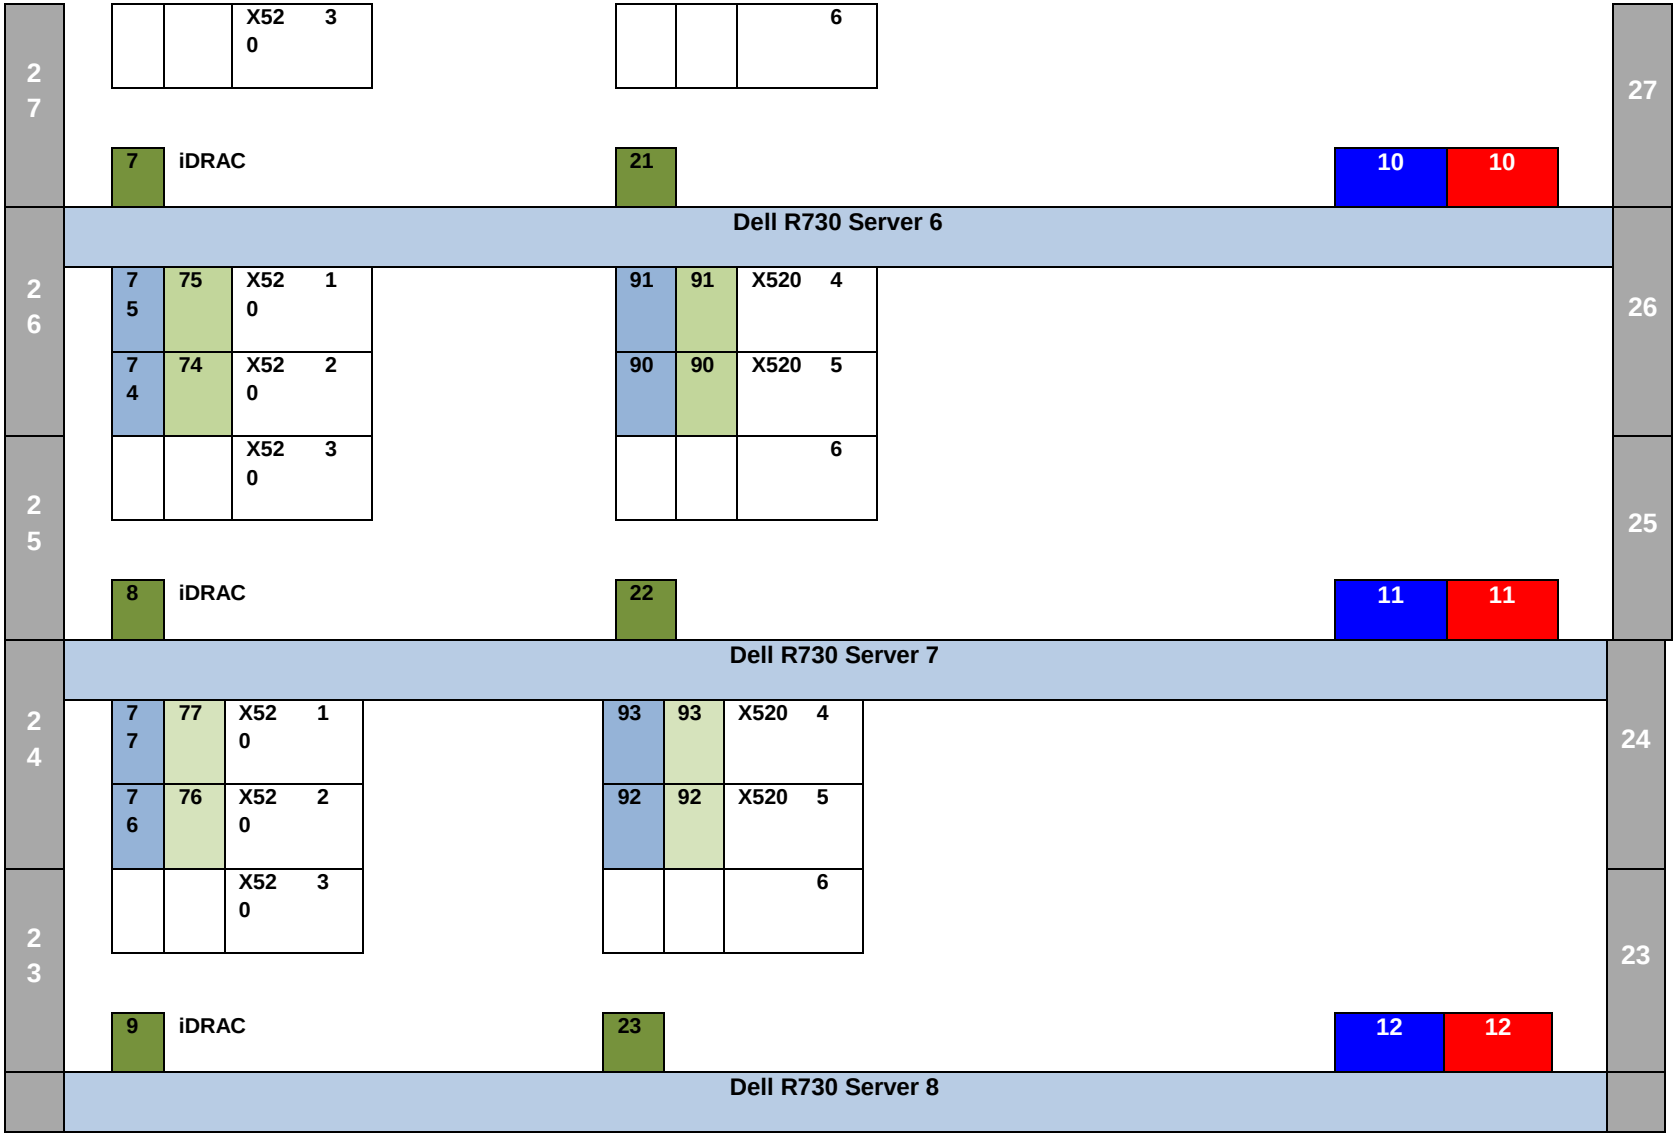
27
X520 3
6
7 iDRAC
21
10 10
Dell R730 Server 6
26
75 75 X520 10 1
91 91 X520 4
74 74 X520 20 2
90 90 X520 5
X520 3
6
25
8 iDRAC
22
11 11
Dell R730 Server 7
24
77 77 X520 10 1
93 93 X520 4
76 76 X520 20 2
92 92 X520 5
X520 3
6
23
9 iDRAC
23
12 12
Dell R730 Server 8

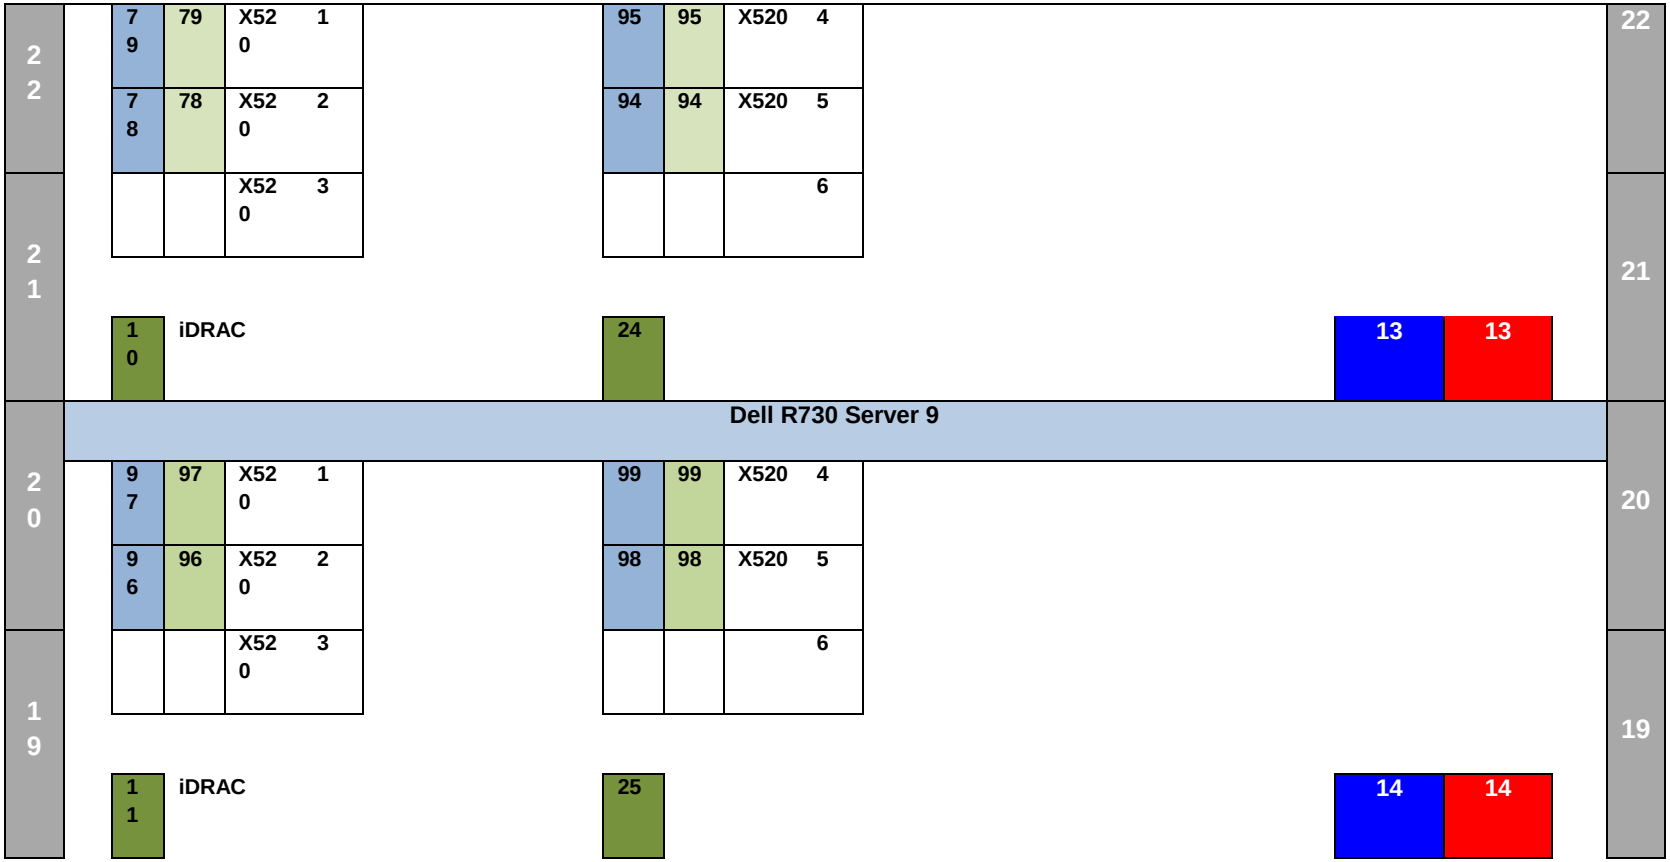

Table 26 Wiring

H Reference

Additional information can be obtained at http://www.dell.com/nfv or by e-mailing nfv@dell.com.

If you need additional services or implementation help, please contact your Dell EMC sales representative.