Creating a VMware Software-Defined Data Center

REFERENCE ARCHITECTURE
VERSION 1.5

vmware®

Table of Contents

Figures

Executive Summary

This reference architecture describes an implementation of a software-defined data center (SDDC) using VMware vCloud® Suite Enterprise 5.8, VMware NSX™ for vSphere® 6.1, VMware IT Business Management Suite™ Standard Edition 1.1, and VMware vCenter™ Log Insight™ 2.0 to create an SDDC. This SDDC implementation is based on real-world scenarios, user workloads, and infrastructure system configurations. The configuration uses industry-standard servers, IP-based storage, and 10-Gigabit Ethernet (10GbE) networking to support a scalable and redundant architecture.

An overview of the solution and the logical architecture as well as results of the tested physical implementation are provided. Consult with your VMware representative as to how to modify the architecture to suit your business needs.

Audience

This document will assist enterprise architects, solution architects, sales engineers, field consultants, advanced services specialists, and customers who are responsible for infrastructure services. This guide provides an example of a successful deployment of an SDDC.

Overview

Expectations for IT to deliver applications and services at speed and scale are greater than ever. The VMware SDDC enables companies to evolve beyond outdated, hardware-centric architectures and to create an automated, easily managed platform that embraces all applications, for fast deployment across data centers and clouds. With an SDDC, users see a predefined list of services available to them and can have these services created instantly upon request, while still enabling IT to control and secure these services. This reference architecture delivers a working configuration that provides users with the on-demand access they need while ensuring that IT maintains the control and security it requires.

FEATURES	DESCRIPTION
Secure user portal	Entitled users utilizing a Web-based portal can request IT services—known as "blueprints"—from a service catalog provided by VMware vCloud Automation Center™.
Blueprints	Blueprints define the attributes associated with items in the service catalog. These items might include virtual, physical, or cloud machines as well as other IT services such as load balancers, firewall rules, runtime policies, and billing policies..
Security groups	Security groups enable administrators to organize objects dynamically, manually, or both in order to apply a common security policy to those objects.
Security policies	Security policies enable administrators to specify guidelines to control introspection services and firewall rules.
Distributed firewall	NSX for vSphere provides a distributed firewall service that operates at the VMware ESXi™ kernel level. This enables firewall rule enforcement in a highly scalable manner without creating bottlenecks common to physical and virtual firewall appliances. With this reduced overhead, the service can perform at true line rate with minimal CPU overhead.
Logical routing – Distributed routing	The distributed routing capability in the NSX for vSphere platform provides an optimized and scalable way of handling traffic between virtual machines or other resources within the data center. Traditionally, virtual machines connected to different subnets must communicate with one another through an external router. In this manner, all virtual machine–to–virtual machine communication crossing subnets must pass through a router. The distributed routing on the NSX for vSphere platform prevents this traditional, nonoptimized traffic flow by providing hypervisor-level routing functionality. Each hypervisor has a routing kernel module that performs routing between the logical interfaces (LIFs) defined on that distributed router instance.
Logical switching	The logical switching capability in the NSX for vSphere platform enables users to spin up isolated logical L2 networks with the same flexibility and agility as they have had with virtual machines.
Backup and restore	VMware vSphere Data Protection™ Advanced provides the ability to back up and restore entire virtual machines as well as application-level services such as Microsoft SQL Server.
Monitoring	vCenter Log Insight provides intelligent operations management via VMware vCenter Operations Manager™ and real-time log management via vCenter Log Insight.
Cost modeling and comparison	IT Business Management enables administrators and users to see what their resources cost the business and to compare their private cloud costs with several public or hybrid cloud providers such as VMware vCloud Air™.
Extensibility	vCloud Automation Center provides out-of-the-box integration with VMware vCenter Orchestrator™ and many third-party solutions.

Table 1. Key Features of This Solution

VMware Software Components

This configuration uses the following VMware software components:

PRODUCT	VERSION	DESCRIPTION
vCloud Suite Enterprise	5.8	Comprehensive suite of products used to deliver the SDDC. In this architecture, users leverage the following components of vCloud Suite Enterprise 5.8: VMware vSphere Enterprise Plus Edition™ 5.5 Update 2, vCloud Automation Center 6.1, vCloud Automation Center Application Services 6.1, vCenter Orchestrator 5.5.2.1, VMware vSphere Data Protection 5.8, and vCenter Operations Manager 5.8.3.
VMware vCenter Server™	5.5 U2	Central platform for managing and configuring the ESXi hypervisor. VMware vSphere Web Client is the centralized point of administration for compute clusters and all networking services provided by NSX for vSphere.
NSX for vSphere	6.1	NSX for vSphere exposes a complete suite of simplified logical networking elements and services including logical switches, routers, firewalls, load balancers, virtual private network (VPN), QoS, monitoring, and security.
IT Business Management	1.1	IT Business Management provides transparency into costs of cloud environments and IT services. Infrastructure teams can understand the costs of supplying private and public cloud environments, while CIOs and IT executives can understand the costs of supplying IT services.
vCenter Log Insight	2.0	Real-time log management and log analysis with machine learning–based intelligent grouping, high-performance search, and better troubleshooting across physical, virtual, and cloud environments.

Table 2. Components

Architectural Overview

This design uses three cluster types, each with its own distinct function. It provides a management plane that is separate from the user workload virtual machines. In addition, it leverages an edge cluster, which provides dedicated compute resources for network services such as edge routers; these provide access to the corporate network and the Internet. This design simplifies the network configuration by eliminating the need to trunk a large number of VLANs to all hosts. Virtual machine–to–virtual machine and virtual machine–to–edge traffic utilizes the NSX for vSphere distributed logical router, which is implemented as a kernel module in each ESXi host. Virtual machines are secured on the network, using the NSX for vSphere distributed firewall, which is also implemented as a kernel module. This enables firewall rules to be enforced before any traffic is put on the wire.

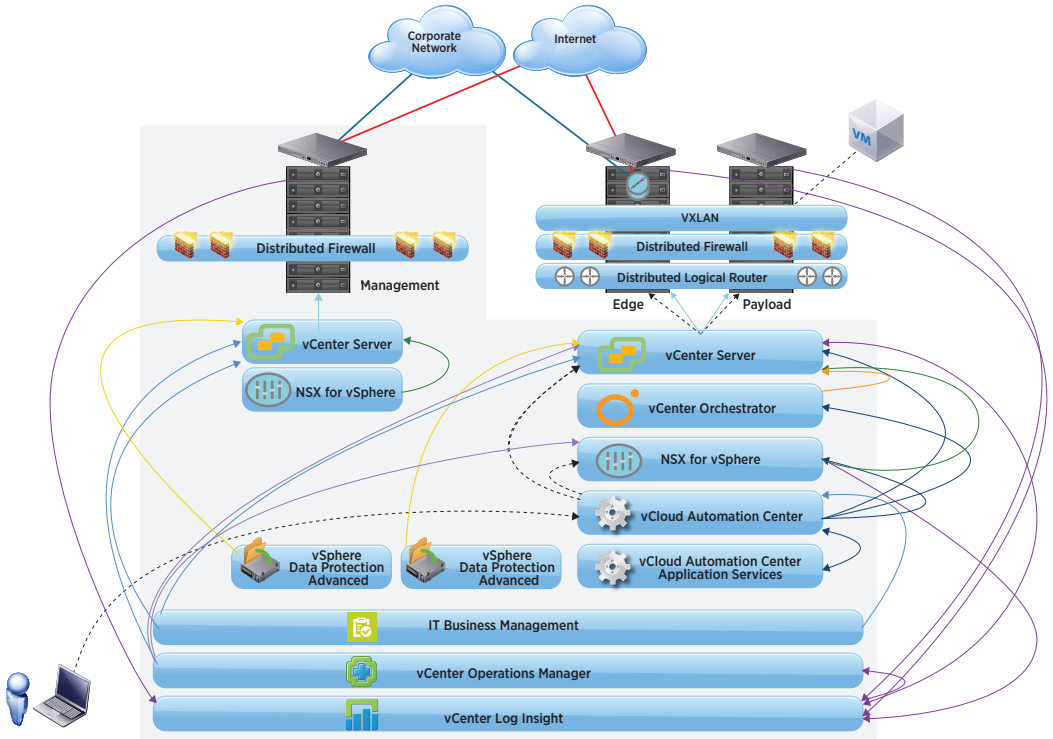

Figure 1. Wire Map

Management Cluster

The management cluster contains the management and monitoring solutions for the entire design. A single management cluster can support multiple pods of edge and payload clusters. The minimum number of hosts required is three, but it will scale out as the number of edge and payload pods increases.

A single vCenter Server instance manages the resources in the management cluster. Additional vCenter Server instances are used to manage edge and payload clusters.

The management cluster also contains common core infrastructure. This includes Microsoft Active Directory, a Microsoft SQL Server cluster, vCenter Operations Manager, IT Business Management, and vCenter Log Insight.

VMware NSX Manager™ instances, one for each vCenter Server, are deployed into the management cluster. NSX for vSphere components, such as VMware NSX Controller™ instances, are also deployed for and in the management cluster.

All vCloud Automation Center components are also deployed in the management cluster.

Edge Cluster

The edge cluster simplifies physical network switch configuration. It is used to deliver networking services to the payload cluster (user-workload) virtual machines. All external networking, including corporate and Internet, for user-workload virtual machines is accessed via the edge cluster. The minimum cluster size is three hosts, but it can scale depending on the volume of services required by the payload cluster virtual machines.

Payload Clusters

The payload clusters are the simplest of the three types; they run user-workload virtual machines. Payload cluster networking is completely virtualized using NSX for vSphere. A single transport zone exists between all payload clusters and the edge cluster. A single NSX for vSphere distributed logical router exists between all clusters. This gives any virtual machine on any host the ability to communicate with any other virtual machine on any host in any cluster—if NSX for vSphere distributed firewall rules permit—without incurring any layer 3 routing penalties. The ESXi host handles all layer 3 routing decisions. When traffic must leave a host, it is encapsulated in an NSX for vSphere packet and is sent to the destination host via layer 2, where the destination host delivers the packet to the destination virtual machine.

Physical Component Details

Compute

The following table lists the recommended physical server configuration:

COMPONENT	SPECIFICATION
CPU	24GHz – 2 x 2.0GHz six-core CPUs (12 total cores)
Memory	128GB ECC RAM
Internal storage	6GB SD card boot device
Network interface cards	2 x 10Gb
Power supplies	Redundant
Fans	Redundant

Table 3. Component Specifications

All physical server hardware, regardless of cluster, utilizes the same configuration for ease of management and to guarantee resource availability as the solution grows. For large deployments, each cluster type should be placed into its own rack with its own top-of-rack (ToR) switches. This is discussed in detail in the "Network" section.

Storage

The management cluster utilizes two 500GB NFS datastores and one 5GB VMware vSphere VMFS iSCSI datastore (to store raw device mapping (RDM) files) in addition to one 10GB and one 50GB iSCSI RDM. The NFS storage serves all management virtual machines; the iSCSI RDMs are used to support the shared-storage requirement of a Microsoft SQL Server failover cluster.

The edge cluster utilizes two 500GB NFS datastores, which serve the NSX Controller instances for the edge and payload clusters as well as VMware NSX Edge™ devices.

The payload clusters utilize at least two NFS datastores. The size and number depend on user application I/O needs. Two 500GB datastores are used for this reference architecture.

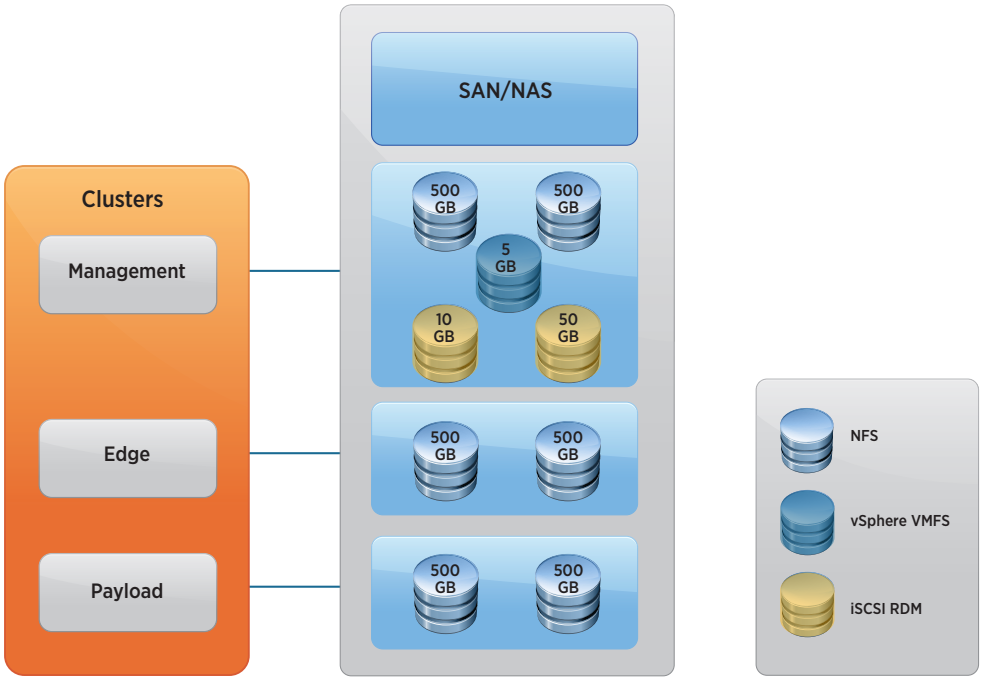

Figure 2. Storage

Network

Each rack contains a pair of 10GbE ToR switches. Each host has one 10GbE port connected to each ToR switch. The hosts utilize the VMware vSphere Distributed Switch™ and load-based teaming (LBT) on all port groups except for the ones carrying the NSX for vSphere overlay, which uses media access control (MAC) hashing because LBT is not supported for the overlay port groups.

802.1.Q trunks are used for carrying a small number of VLANs—for example, NSX for vSphere, management, storage, and VMware vSphere vMotion® traffic. The switch terminates and provides default gateway functionality for each respective VLAN; that is, it has a switch virtual interface (SVI) for each VLAN. Uplinks from the ToR switch to the aggregation layer are routed point-to-point links. VLAN trunking on the uplinks—even for a single VLAN—is not allowed. A dynamic routing protocol (OSPF, ISIS, or BGP) is configured between the ToR and aggregation layer switches. Each ToR switch advertises a small set of prefixes, typically one per VLAN or subnet that is present. In turn, it will calculate equal cost paths to the prefixes received from other ToR switches.

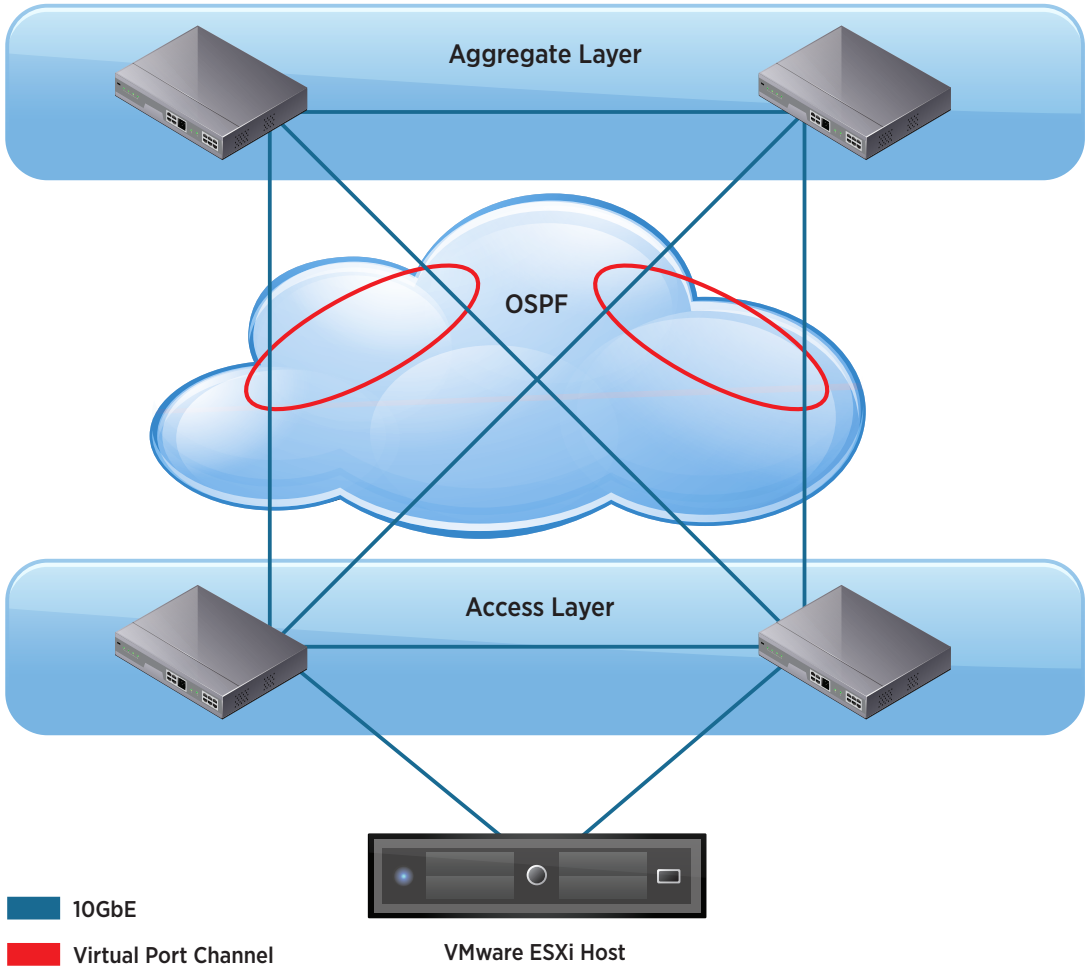

Provides default gateway for all VLANs. VLANs are not trunked to aggregate layer switches.

Figure 3. Physical Network Connections

Software-Defined Data Center Component Details

In this section, we will define the VMware software components and their configuration in enabling this solution.

All Microsoft Windows installations utilize Windows Server 2012 R2 Standard Edition. Because this architecture utilizes the Failover Clustering feature, Microsoft SQL Server Enterprise Edition 2012 SP1 is required.

COMPONENT	NUMBER DEPLOYED	DEPLOYED LOCATION
Microsoft SQL Server	2 (clustered)	Management cluster
PostgreSQL	2 (clustered)	Management cluster
vCenter Server	2	Management cluster
ESXi hosts	10	3 management cluster, 3 edge cluster, 4 payload cluster
vCloud Automation Center	1 (in a redundant distributed configuration composed of eight virtual machines)	Management cluster
vCloud Automation Center Application Services	1	Management cluster
vCenter Orchestrator	2 (clustered)	Management cluster
vSphere Data Protection	2	1 management cluster, 1 payload cluster
NSX Manager	2	Management cluster
NSX for vSphere controllers	6	3 management cluster, 3 edge cluster
vCenter Operations Manager	1	Management cluster
IT Business Management	1	Management cluster
vCenter Log Insight	1	Management cluster

Table 4. SDDC Component Details

vSphere Data Center Design

vSphere Enterprise Plus Edition is the core that enables the SDDC. All ESXi hosts are stateful installs—that is, the ESXi hypervisor is installed to local disks.

Management Cluster

ATTRIBUTE	SPECIFICATION
ESXi version	5.5 Update 2
Number of hosts	3
Number of CPUs per host	2
Number of cores per CPU	6
Core speed	2.0GHz
Memory	128GB
Number of network adapters	2 x 10Gb

Table 5. Management Cluster Details

The cluster leverages VMware vSphere High Availability (vSphere HA) and VMware vSphere Distributed Resource Scheduler™ (vSphere DRS). vSphere HA is set to monitor both hosts and virtual machines. Its admission control policy utilizes a single failover host, guaranteeing sustainability with one node failure. A failover host is used because vSphere DRS relies on vCenter Server if the vCenter Server virtual machine that resides on the host that fails vSphere DRS is not available to move virtual machines to other hosts—sometimes referred to as "defragging"—to enable vSphere HA to restart the vCenter Server virtual machine. vSphere DRS is set to fully automated mode.

VLAN ID	FUNCTION
970	ESXi management
980	vSphere vMotion
1000	IP storage (iSCSI)
1020	IP storage (NFS)
1680	Management virtual machines (vCenter Server, Microsoft SQL Server, vCloud Automation Center, etc.)
3000	Microsoft Failover Clustering heartbeat
3001	NSX for vSphere overlay (VXLAN)

Table 6. Management Cluster VLAN IDs and Functions

STORAGE	TYPE	FUNCTION
NFSMGT01	NFS	500GB management virtual machine datastore
NFSMGT02	NFS	500GB management virtual machine datastore
RDMMAPPING01	VMFS over iSCSI	5GB VMFS for RDM files
Physical RDM 1	RDM over iSCSI	10GB RDM in physical compatibility mode – Microsoft Cluster Service quorum
Physical RDM 2	RDM over iSCSI	50GB RDM in physical compatibility mode (clustered shared Microsoft SQL Server data)

Table 7. Management Cluster Storage

ATTRIBUTE	SPECIFICATION
Number of CPUs	4
Processor type	VMware virtual CPU
Memory	16GB
Number of network adapters	2
Network adapter type	VMXNET3
Number of disks	3 30GB (C:\) – VMDK 50GB (D:\) – RDM (physical mode) 10GB (Q:\) – RDM (physical mode)
Operating system	Windows Server 2012 R2

Table 8. Microsoft SQL Server Cluster Configuration

For more information on how to configure Microsoft Cluster Service (MSCS) in a vSphere environment, see VMware Knowledge Base article 1037959 and the setup for Microsoft clustering guide.

ATTRIBUTE	SPECIFICATION
vCenter Server version	vCenter Server 5.5 Update 2 installable
Quantity	2 (1 for management cluster, 1 for edge and payload clusters)
Number of CPUs	4
Processor type	VMware virtual CPU
Memory	16GB (for management cluster) 32GB (for edge and payload clusters)
Number of network adapters	1
Network adapter type	VMXNET3
Number of disks	1 100GB (C:\) – VMDK
Operating system	Windows Server 2012 R2

Table 9. VMware vCenter Server Configuration

To increase database resiliency, database clustering via Microsoft SQL Server is used. To use Microsoft SQL Server and the Linked Mode feature of vCenter Server, the Windows installable version of vCenter Server must be used.

The following components are also installed on the vCenter Server instance as part of the VMware vCenter Server installation process. All components requiring a Microsoft SQL Server database—that is, vCenter Server and VMware vSphere Update Manager™—must have their databases located on the Microsoft SQL Server cluster.

COMPONENT	DESCRIPTION
VMware vCenter Single Sign-On™	Authentication broker service
vSphere Web Client	Web version of the VMware vSphere Client™ used to manage vSphere and the only way to configure new features of vSphere 5.1 and later
vCenter Inventory Service	Used to manage the vSphere Web Client inventory objects and property queries that the client requests when users navigate the vSphere environment
vCenter Server	The main component that enables the centralized management of all ESXi hosts and the virtual machines that run on them
vSphere Update Manager	Automates patch management and eliminates manual tracking and patching of vSphere hosts

Table 10. VMware vCenter Server Components

ATTRIBUTE	SPECIFICATION
Data center object	WDC
Linked Mode	Enabled

Table 11. VMware vCenter Server Datacenter Configuration

PORT GROUP	VLAN ID	FUNCTION
MGMT	970	ESXi management
vMotion	980	vSphere vMotion
iSCSI	1000	IP storage (iSCSI)
NFS	1020	IP storage (NFS)
VMMGMT	1680	Management virtual machines
MSCSHB	3000	Microsoft Failover Clustering heartbeat
	3001	NSX for vSphere overlay (VXLAN)

Table 12. Management Cluster Virtual Switch Port Groups, VLAN IDs, and Function

Edge Cluster

ATTRIBUTE	SPECIFICATION
ESXi version	5.5 Update 2
Number of hosts	3
Number of CPUs per host	2
Number of cores per CPU	6
Core speed	2.0GHz
Memory	128GB
Number of network adapters	2 x 10Gb

Table 13. Edge Cluster Details

The cluster leverages vSphere HA and vSphere DRS. vSphere HA is set to monitor both hosts and virtual machines with its admission control policy set to a percentage of cluster resources reserved—33 percent for a three-node cluster—guaranteeing the sustainability of one node failure. vSphere DRS is set to fully automated mode.

Payload Cluster

ATTRIBUTE	SPECIFICATION
ESXi version	5.5 Update 2
Number of hosts	4
Number of CPUs per host	2
Number of cores per CPU	6
Core speed	2.0GHz
Memory	128GB
Number of network adapters	2 x 10Gb

Table 14. Payload Cluster Details

The cluster leverages vSphere HA and vSphere DRS. vSphere HA is set to monitor both hosts and virtual machines with its admission control policy set to a percentage of cluster resources reserved—25 percent for a four-node cluster—guaranteeing the sustainability of one node failure. vSphere DRS is set to fully automated mode.

Edge and Payload Clusters

PORT GROUP	VLAN ID	FUNCTION
MGMT	970	ESXi management
vMotion	980	vSphere vMotion
NFS	1020	IP storage (NFS)
External	2010	External connectivity to corporate network
	3001	NSX for vSphere overlay (VXLAN)

Table 15. VMware vSphere Distributed Switch Port Groups and VLANs

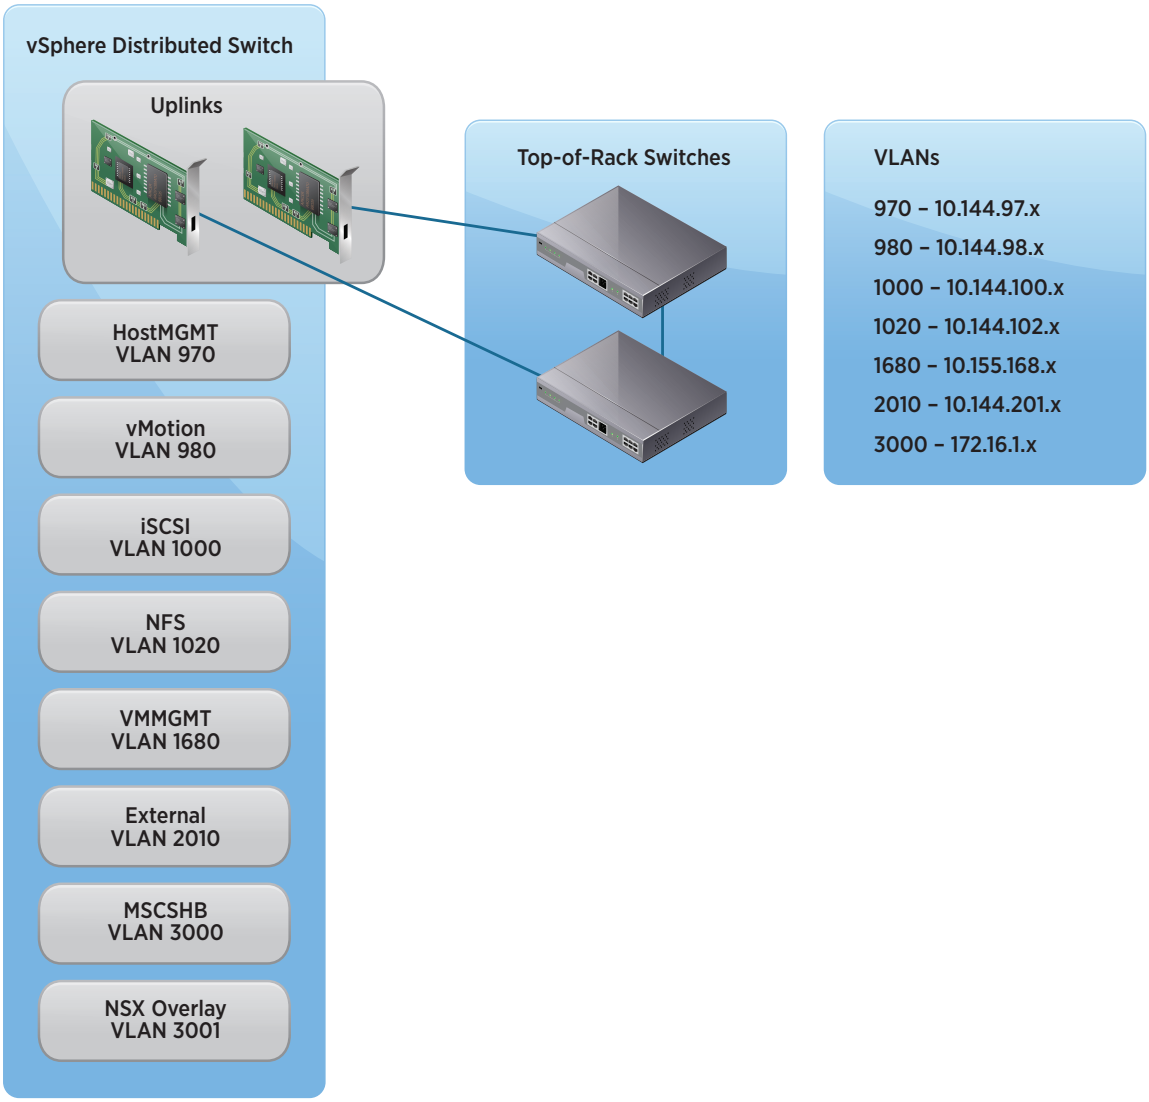

Figure 4. Logical Network

The external port group is configured with one static port, with the elastic option disabled; an NSX Edge device is all that is connected to this port group. The remaining ports are configured with static binding, with the default eight ports and the elastic option enabled.

The uplink configuration is one uplink to each physical switch. The distributed port groups use the load-based teaming algorithm for all port groups except those that carry NSX for vSphere overlay traffic; those port groups use the MAC hashing algorithm.

vCenter Orchestrator

vCenter Orchestrator is deployed using the vCenter Orchestrator appliance. For resiliency, it is set up in a cluster, with its database residing on the Microsoft SQL Server cluster. NSX for vSphere and vCloud Automation Center plug-ins are installed on both instances.

vCenter Orchestrator is a critical component in the SDDC. All vCloud Automation Center communication to NSX for vSphere is handled via vCenter Orchestrator workflows.

vSphere Data Protection Advanced

vSphere Data Protection Advanced is deployed as two separate appliances. These appliances are distributed in the OVA format. The first appliance resides in the management cluster and is responsible for backups and restores of the virtual machines residing in the management cluster. The second appliance resides in the payload cluster and is responsible for backups and restores of the virtual machines residing in the payload cluster.

The advanced edition is used to create application-aware backups of the Microsoft SQL Server cluster databases.

Two backup policies were created for the management cluster: The first backs up all virtual machines in the cluster nightly; the second is an application-specific policy to back up Microsoft SQL Server failover clusters. All databases are backed up. Retention for both policies is set to 14 days, 4 weeks, 12 months, and 7 years.

A single policy was created for the payload cluster. The backup target is the payload cluster itself. This ensures that any new virtual machines added are automatically added to the backup job. Retention for this policy is also set to 14 days, 4 weeks, 12 months, and 7 years.

NSX for vSphere

NSX for vSphere provides all the logical switches, routing, and distributed firewall services used to create this architecture. All virtual machine traffic, excluding the management cluster, is encapsulated using NSX for vSphere. All virtual machine–to–virtual machine, or east–west, traffic in the payload cluster is routed between hosts by the NSX for vSphere distributed logical router. When a request to or from the external network is serviced, it travels through an NSX Edge device, which provides all north–south routing—that is, routing to and from external networks.

NSX for vSphere has a one-to-one relationship with vCenter Server, so two NSX Manager instances are deployed, one for the management cluster vCenter Server instance and the other for the edge and payload cluster vCenter Server instance. These are both deployed in the management cluster.

NSX for vSphere utilizes controller virtual machines to implement the network control plane. The NSX Controller instances must be deployed in odd numbers to prevent a split-brain scenario. As such, three controllers per NSX for vSphere instance are deployed, with vSphere DRS rules set up to ensure that they not run on the same host. The NSX Controller instances for the management cluster are deployed into the management cluster itself. The NSX Controller instances for the edge and payload clusters are deployed into the edge cluster.

The ESXi hosts must be prepared for NSX for vSphere. The following values are used:

SPECIFICATION	VALUE
MTU	9000
Teaming mode	Source MAC
Segment IDs	5000–7999
Transport zones	1 – encompassing all clusters

Table 16. NSX for vSphere Host Preparation Values

Internal (east–west) routing is enabled using the NSX for vSphere distributed logical router. With it, all virtual machines can communicate with each other—assuming that firewall rules allow—via layer 3, with no need to access an external router. In this configuration, each ESXi host running the distributed router acts as the default gateway for virtual machines on the host. Virtual machines communicate with the host on which they are currently running, which then encapsulates the traffic in the NSX for vSphere overlay and sends it to the destination host, where the packets are decapsulated and delivered to the correct destination. This distributed routing eliminates the need to hop to an external router to communicate with virtual machines on different subnets.

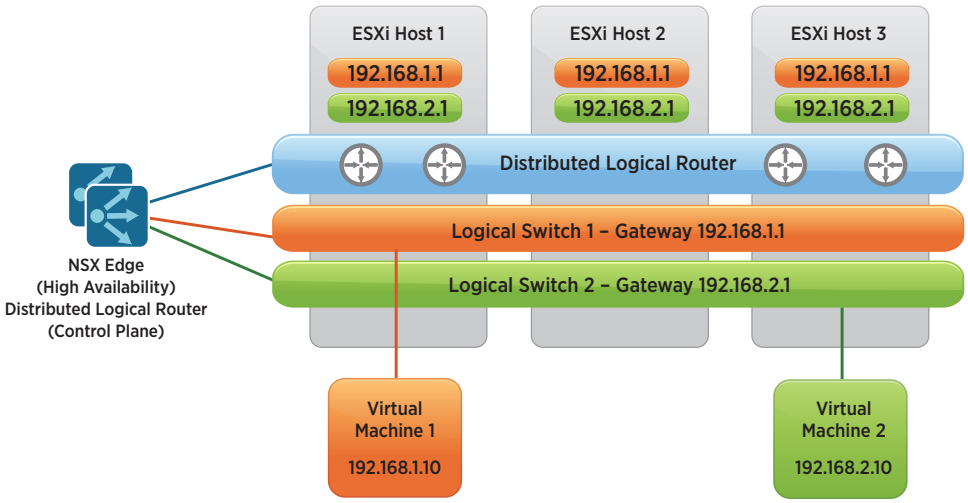

Figure 5. NSX for vSphere East–West Routing

To enable external (north–south) connectivity, an NSX Edge router is deployed in high-availability mode. One interface is connected to the external network; another is connected to a logical switch, which is also connected to the NSX for vSphere distributed logical router. Both the NSX Edge device and the distributed logical router run the OSPF dynamic routing protocol to propagate routing information upstream to the physical network and also downstream to the NSX for vSphere distributed logical router.

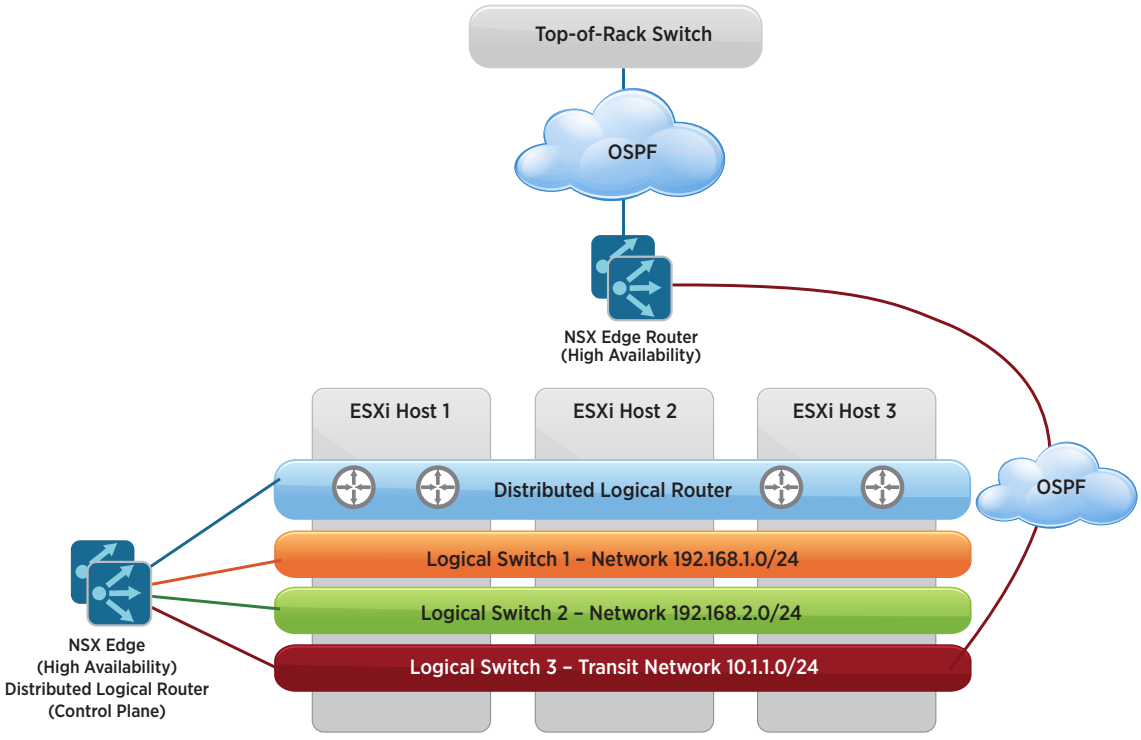

Figure 6. NSX for vSphere North–South Routing

By default, the distributed firewall—which like the distributed logical router is a kernel module on the ESXi hosts—allows all traffic. This enables the rules to be evaluated and decisions to be made in real time before a packet is placed on the network. There are several ways to create firewall rules. In this architecture, security groups were created based on virtual machine name. The appropriate firewall rules were created in a security policy in NSX Service Composer and are assigned to a security group. This automates the firewall rules applied to virtual machines as they are added or removed from security groups. To ensure that human error cannot make the vCenter Server and Active Directory servers unavailable, they were added to the exclusion list. The following security groups, policies, and rules were created for the environment.

SECURITY GROUP	SERVER ROLE	SECURITY POLICY	SOURCE GROUP	DESTINATION GROUP	SERVICE
N/A	N/A	Common	Any	NTP Servers	NTP
			Any	Management	Syslog
			Any	Domain Controllers	DHCP
			Any	Domain Controllers	Active Directory
			Any	Domain Controllers	DNS
vDPA	vSphere Data Protection Appliances	vDPA	Any	vDPA	TCP:8453
			Any	vDPA	TCP:28001,29000
			vCenter Servers	vDPA	Any
PostgreSQL	PostgreSQL Virtual Appliance	PostgreSQL	PostgreSQL	PostgreSQL	Any
			vCAC Appliances	PostgreSQL	PostgreSQL
			Any	PostgreSQL	SSH
			Any	PostgreSQL	TCP:5480
vCAC Appliances	vCAC CAFÉ Appliances	vCAC Appliances	Any	vCAC Appliances	HTTP/HTTPS
			Any	vCAC Appliances	SSH
			Any	vCAC Appliances	TCP:5480
			vCAC Appliances	PostgreSQL	PostgreSQL
			vCAC Appliances	vCAC Managers	HTTPS
			vCAC Appliances	vCenter Servers	Any
			vCAC Appliances	vCO Servers	VMware-VCO
vCAC IaaS	vCAC IaaS	vCAC Windows Servers	Any	vCAC IaaS	HTTP/HTTPS
			vCAC IaaS	SQL Servers	MSSQL
			vCAC IaaS	vCAC Managers	HTTPS
			Any	vCAC IaaS	RDP
			vCAC IaaS	Domain Controllers	Any
vCAC Managers	vCAC Manager, DEM, and Agent Servers	vCAC Windows Servers	Any	vCAC Managers	HTTPS
			vCAC Managers	SQL Servers	MSSQL
			vCAC Managers	Domain Controllers	Any

SECURITY GROUP	SERVER ROLE	SECURITY POLICY	SOURCE GROUP	DESTINATION GROUP	SERVICE
vCO Servers	vCenter Orchestrator	vCO Servers	vCO Servers	SQL Servers	MSSQL
			vCO Servers	vCenter Servers	Any
			vCAC Appliances	vCO Servers	VMware-VCO
			Any	vCO Servers	VMware-VCO
Management	vCenter Operations Manager, vCenter Log Insight, VMware IT Business Management	Management	Any	Management	HTTP/HTTPS
			Any	Management	TCP:5480
			Any	Management	Syslog
			Management	vCenter Servers	Any
			Management	Management	Any
			Management	Any	FTP
			Management	Domain Controllers	Any
SQL Servers	Microsoft SQL Servers	SQL Servers	vCenter Servers	SQL Servers	MSSQL
			vCO Servers	SQL Servers	MSSQL
			vCAC IaaS	SQL Servers	MSSQL
			vCAC Managers	SQL Servers	MSSQL
			vCAC IaaS	SQL Servers	RPC
			vCAC Managers	SQL Servers	RPC
			vCAC IaaS	SQL Servers	DTC
			vCAC Managers	SQL Servers	DTC
			SQL Servers	SQL Servers	Any
			SQL Servers	vDPA	TCP:29000
			SQL Servers	Domain Controllers	Any
N/A	N/A	Default	Any	Any	Block All

Table 17. Security Groups, Policies, and Firewall Rules

vCloud Automation Center

vCloud Automation Center empowers IT to accelerate the delivery and ongoing management of personalized, business-relevant infrastructure, application, and custom services while improving overall IT efficiency. Policy-based governance and logical application modeling ensures that multivendor, multicloud services are delivered at the right size and service level for the task to be performed. Full life-cycle management ensures that resources are maintained at peak operating efficiency, and release automation enables multiple application deployments to be kept in sync through the development and deployment process.

vCloud Automation Center provides the portal to the business users who request services. This architecture utilizes the distributed architecture of vCloud Automation Center and uses NSX for vSphere to create a highly available environment by load-balancing multiple instances of the vCloud Automation Center components.

Load-Balanced vCloud Automation Center Configuration

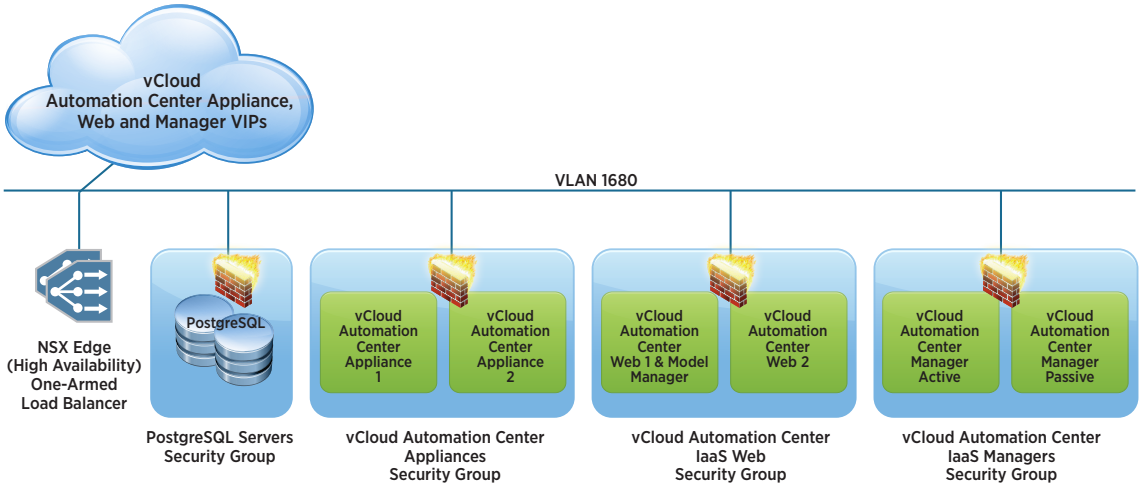

Figure 7. NSX for vSphere Load Balancer Configuration for vCloud Automation Center

To achieve the architecture shown in Figure 7, a single NSX Edge device in high-availability mode is deployed. It is configured to load-balance vCloud Automation Center appliance, Web, and vCloud Automation Center manager traffic. To achieve the preferred distributed firewall configuration, the NSX for vSphere security policies and groups discussed in the previous section were also created and dynamically applied to each virtual machine as they were built.

Because we are load-balancing Web servers that can, and will, make requests that are routed back to themselves, a registry setting that disables Windows loopback checking was created.
See VMware Knowledge Base article 2053365 for more information.

vCloud Automation Center Appliances

The vCloud Automation Center appliance is distributed as a prepackaged appliance in OVA format. For increased redundancy, two of these appliances are deployed and configured for clustering, along with a second set of vCloud Automation Center appliances, with all vCloud Automation Center services disabled to run the external PostgreSQL master–slave cluster. The NSX Edge device shown in Figure 7 is configured to load-balance the traffic to the vCloud Automation Center appliances. The two servers along with the load balancer virtual server created for the appliances share a common SSL certificate.

vSphere DRS rules were created to ensure that the vCloud Automation Center appliances run on different hosts; similarly, vSphere DRS rules were also created for the PostgreSQL cluster nodes, to ensure that they also run on different hosts.

vCloud Automation Center IaaS Web Servers

The vCloud Automation Center IaaS Web servers utilize Internet Information Services (IIS) for Windows Server. For redundancy, two IaaS Web servers are deployed. Both are active and are load balanced by the NSX Edge device shown in Figure 7. The model manager data is deployed to the first IaaS Web server only. The two servers along with the load balancer virtual server created for the IaaS Web servers share a common SSL certificate.

vCloud Automation Center IaaS Managers

The vCloud Automation Center IaaS managers utilize Microsoft Internet Information Services (IIS) for Windows Server. For redundancy, two IaaS managers are deployed. Only one is set to active; the other is passive. The NSX Edge device shown in Figure 7 is configured to load-balance the traffic, but only the currently active manager is active on the load balancer. During an outage, manual steps are taken to make the passive server active and to enable the load balancer configuration to use the now-active server. The two servers along with the load balancer virtual server created for the IaaS Manager servers share a common SSL certificate.

Distributed Execution Managers and the vSphere Agent

The distributed execution manager (DEM) and the vSphere agent run on Windows. The DEM and the vSphere agent are installed on the manager servers; these services do not support load balancing but are highly available when deployed in this configuration.

vCloud Automation Center Application Services

vCloud Automation Center Application Services is distributed as a prepackaged appliance in OVA format. It is deployed into the management cluster and integrated with vCloud Automation Center. This enables vCloud Automation Center Application Services to utilize operating system (OS) blueprints already created in vCloud Automation Center for use in creating application-specific blueprints by using its drag-and-drop interface.

Monitoring

Monitoring the performance, capacity, health, and logs in any environment is critical. But in a solution where IT gives some control to the business users, monitoring becomes mission critical to the success of user adoption.

vCenter Operations Manager

vCenter Operations Manager provides operations dashboards, performance analytics, and capacity optimization capabilities needed to gain comprehensive visibility, proactively ensure service levels, and manage capacity in dynamic virtual and cloud environments.

vCenter Operations Manager is deployed as a pair of virtual appliances in a VMware vSphere vApp™ distributed in the OVA format.

To ensure a complete picture of how the environment is running, vCenter Operations Manager is configured to monitor the management, edge, and payload vCenter Server instances. Additionally, the NSX for vSphere content pack is installed to give insight into the virtualized networking environment.

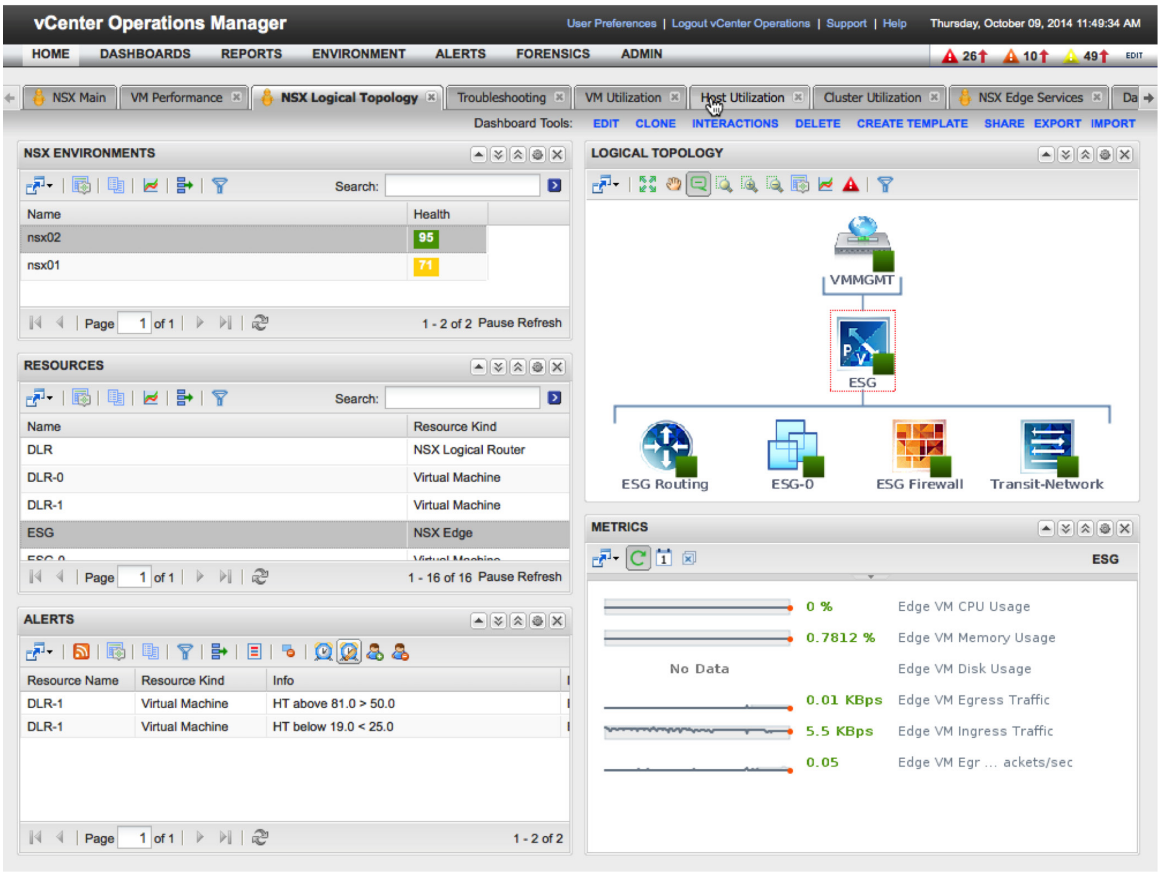

Figure 8. vCenter Operations Manager Custom UI

vCenter Operations Manager requires updates to the default monitoring settings for most organizations. For more information on how to customize vCenter Operations Manager for your specific environment, see the vCenter Operations Manager documentation.

vCenter Log Insight

vCenter Log Insight provides in-depth log analysis in an easy-to-query Web interface. It collects syslog data from ESXi hosts or any other server or device that supports syslog. There is also an installable agent for Windows that enables the collection of event logs and custom logs such as the vCenter Server and vCloud Automation Center server log files.

vCenter Log Insight is deployed as a virtual appliance.

The Windows agent is installed on the vCenter Server instance and on vCloud Automation Center servers to collect their log files.

The syslogs of all appliances and ESXi hosts are configured to send to the vCenter Log Insight server.

Additionally, the vCloud Automation Center content pack is installed to create a dashboard view to easily monitor the vCloud Automation Center environment.

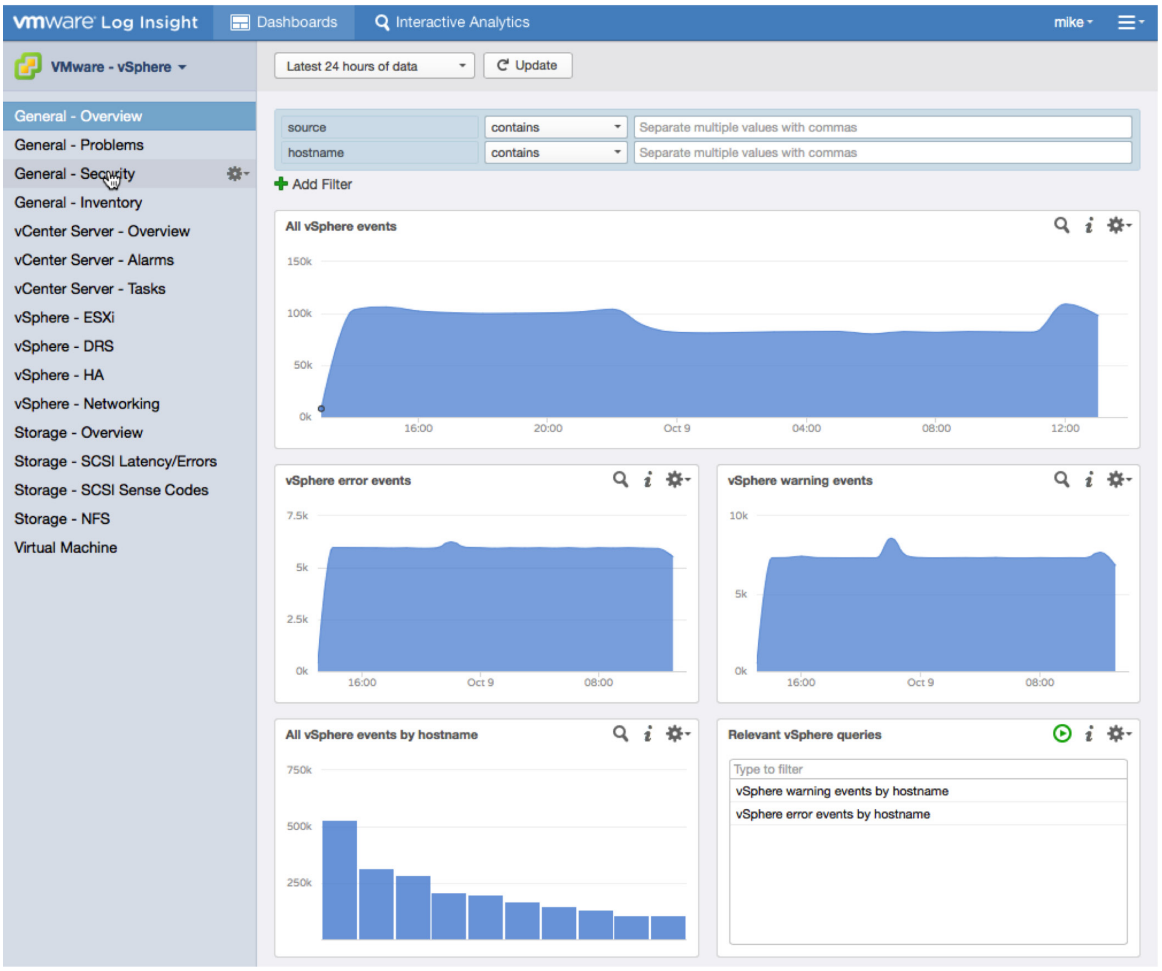

Figure 9. vCenter Log Insight

VMware IT Business Management Suite Standard Edition

IT Business Management provides transparency and control over the cost and quality of IT services. By providing a business context to the services that IT offers, IT Business Management helps IT organizations shift from a technology orientation to a service broker orientation, delivering a portfolio of IT services that align with the needs of line-of-business stakeholders.

IT Business Management is used for its deep integration with vCloud Automation Center. It is deployed as a virtual application that is packaged in the OVA file format. After it has been installed and configured to integrate with vCloud Automation Center, users who are granted access can see costing information on the services they are deploying. They can also compare those costs to hybrid or public cloud offerings from vCloud Air, Amazon AWS, and Microsoft Azure.

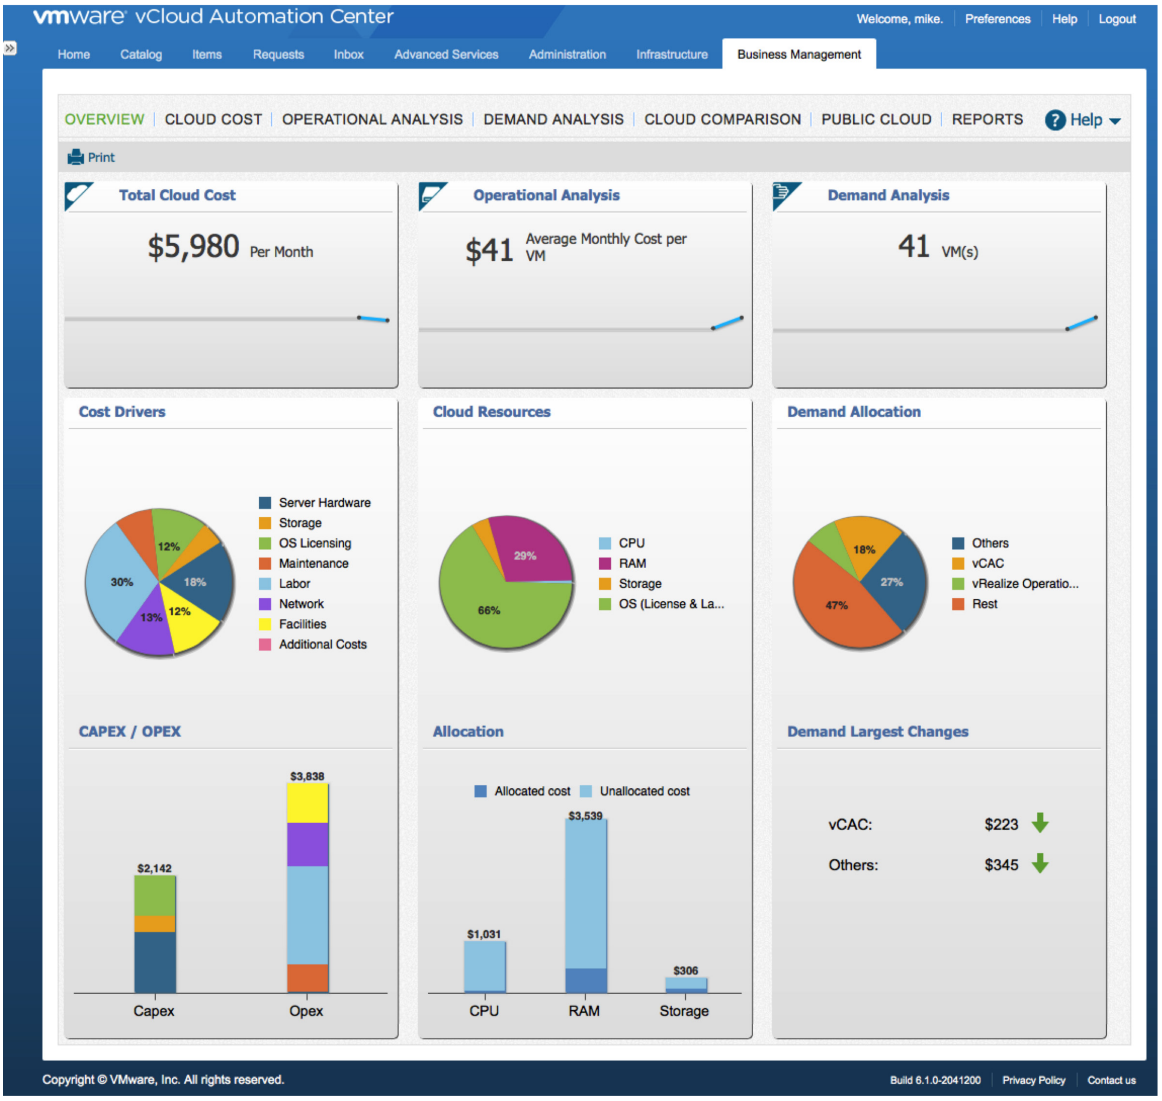

Figure 10. IT Business Management

SDDC Operational Configuration

When all of the components have been installed, they must be brought together to enable the creation of blueprints and the provision of services by authorized users. To unify the components to operate as a solution, the following configuration steps are required.

NSX for vSphere Configuration

First, provision the common network resources for use within vCloud Automation Center. An NSX Edge device in the edge cluster was created for north–south routing; the OSPF dynamic routing protocol was configured between the NSX Edge device and the external physical switches. The distributed logical routing functionality was also enabled in both the edge and payload clusters for east–west routing, with dynamic routing via OSPF between the NSX Edge device and the distributed logical router. Logical switches were precreated for use in vCloud Automation Center for single-machine blueprints; they were connected to the distributed logical router. vCloud Automation Center can dynamically create logical switches in multimachine blueprints.

Tenants

A tenant is an organizational unit in a vCloud Automation Center deployment. A tenant can represent the entire organization or specific business units. A default vsphere.local tenant is created during the installation. This tenant is the only tenant that can leverage native Active Directory integration; all other tenants must bind to an Active Directory domain controller as an LDAP server. Because of this limitation and the ability to have resource reservations at a business-group level—this will be discussed later—this architecture utilizes only the default tenant.

Endpoints

Endpoints are the infrastructure sources that vCloud Automation Center consumes. In VMware vCloud Suite and in this architecture, the endpoint is vCenter Server—more specifically, the vCenter Server instance that manages the edge and payload clusters.

Fabric Groups

Fabric groups are groups of compute resources that the endpoints discover; they define the organization of virtualized compute resources. In most single-site environments, a single fabric group is created that contains all nonmanagement clusters.

Business Groups

Business groups define the users and machine prefixes and are used later to grant access to a percentage of resources. Users assigned the group manager role can create blueprints and see all machines created in the group. Support users can work for another user, and users can be entitled to request blueprints in the catalog. In most environments, business groups are created for department or business units in an organization.

Network Profiles

Network profiles define the type of connection—external, private, NAT, or routed—that a resource has. NAT and routed profiles require an external profile. External profiles connect resources to an existing network.

Reservation Policies

Reservation policies enable a user to associate one or more reservations into a policy that can be applied to a blueprint. Multiple reservations can be added to a reservation policy, but a reservation can belong to only one policy. A single reservation policy can be assigned to more than one blueprint. A blueprint can have only one reservation policy.

Reservations

A virtual reservation is a share of the memory, CPU, networking, and storage resources of one compute resource allocated to a particular business group.

To provision virtual machines, a business group must have at least one reservation on a virtual compute resource. Each reservation is for one business group only, but a business group can have multiple reservations on a single compute resource or on compute resources of different types.

A machine blueprint is the complete specification for a virtual, cloud, physical machine, or service. Blueprints determine a machine's attributes, the manner in which it is provisioned, and its policy and management settings.

In this architecture, blueprints can be either vSphere based—that is, single machine—or multimachine, which requires one or more vSphere blueprints and provisions and manages them together as a single entity. Multimachine blueprints also enable the dynamic provisioning of networks using network profiles.

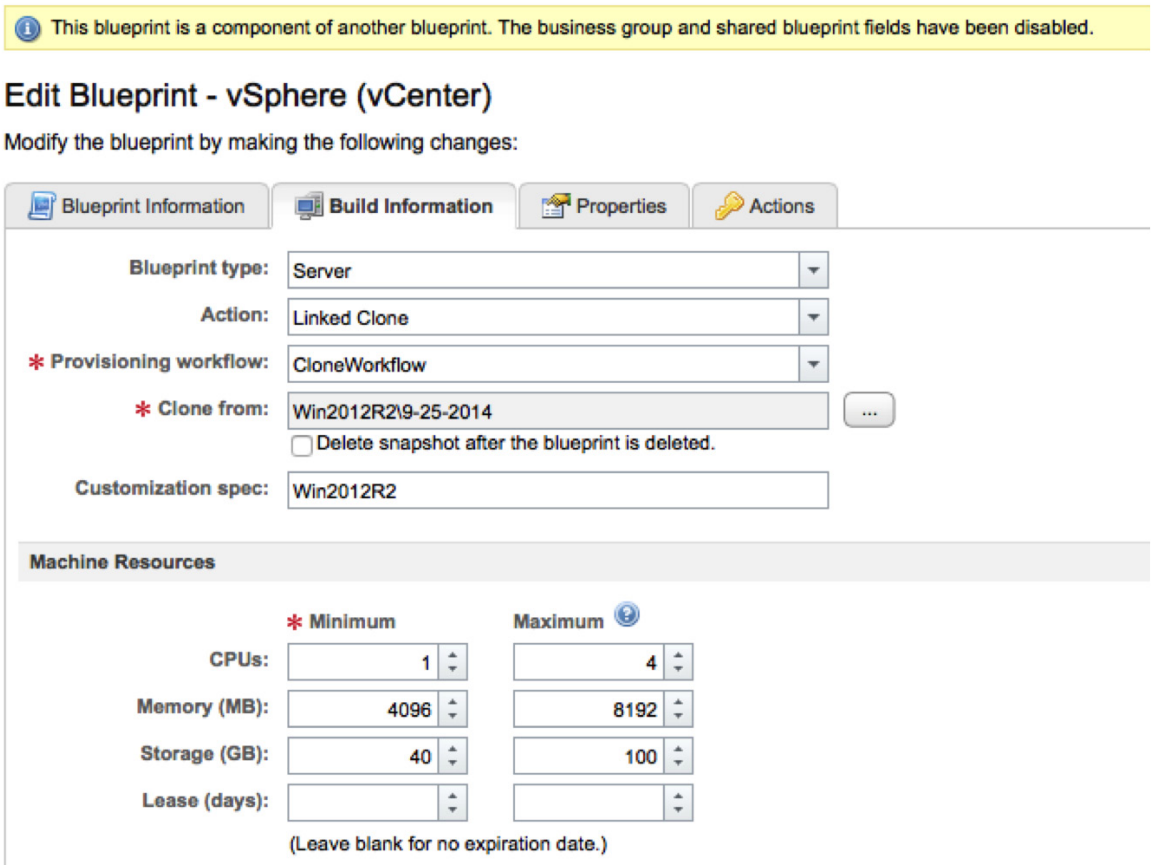

Figure 11. Blueprint Build Information

By default, the blueprint is provisioned onto a network selected in the reservation for a given business group. Custom properties can be specified to choose the network, the network profile, and other blueprint settings.

After blueprints have been created and published to the catalog, authorized users can log in to the vCloud Automation Center portal and request these resources.

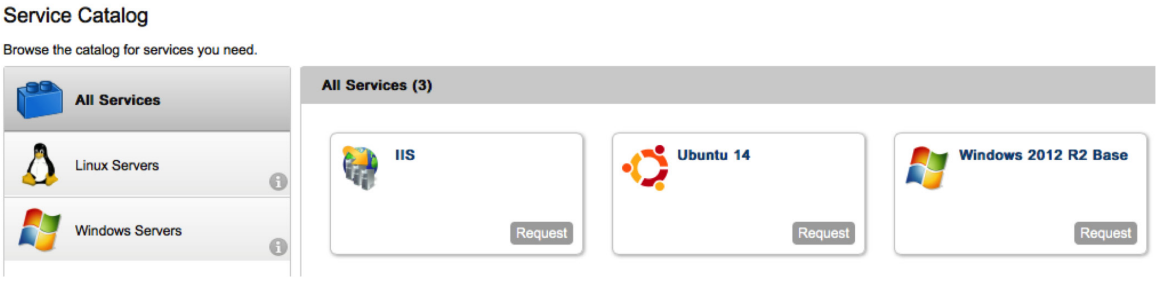

Figure 12. vCloud Automation Center Service Catalog

About the Author

Mike Brown is a senior technical marketing manager in the Cloud Infrastructure Technical Marketing group. Mike's focus is on reference architectures for VMware vCloud Suite and the software-defined data center as well as on resource management. He has multiple industry certifications, including VMware Certified Design Expert (VCDX). Follow Mike on the vSphere Blog and on Twitter @vMikeBrown.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com