Reference Architecture: VMware vCloud Suite with ThinkSystem Servers

Last update: 7 June 2019
Version 1.2
Configuration Reference Number: CLDVW04XX73

Describes reference architecture for VMware vCloud using network shared storage

Contains performance data for sizing recommendations

Includes details about hybrid cloud connectivity to Amazon Web Services

Contains detailed bill of materials for servers, networking, and storage

Mike Perks

Chandrakandh Mouleeswaran

LENOVO PRESS

Table of Contents

1 Introduction

This document describes the reference architecture for the VMware vCloud Suite and Lenovo® ThinkSystem servers, storage, and networking. This document also covers components required to be used for integrating an on-premise VMware vRealize cloud with Amazon AWS public clouds.

The intended audience of this document is IT professionals, technical architects, sales engineers, and consultants to assist in planning, designing, and implementing a private or hybrid cloud using VMware products.

This reference architecture covers the following VMware products:

- vSphere 6.7u1 (and vSphere 6.5 U2), which provides compute virtualization
- vCloud Suite 7.5, which provides a VMware vSphere-based private cloud using vRealize Suite products and additional products to support vSphere Data Protection and Availability
- vRealize Suite 7.5, which provides cloud management capabilities for private, public and hybrid clouds with support for multiple hypervisors
- AWS Server Migration Service Connector 1.0.12.50 which supports migration of virtual machines from on-premise vSphere cloud to AWS pubic cloud.

This document provides an overview of the business problem and business value provided by VMware vCloud Suite. A description of customer requirements is followed by an architectural overview of the solution and a description of the logical components. The operational model describes the architecture for deploying into small to medium Enterprises. Performance and sizing information is provided with the best practices and networking considerations for implementing VMware vCloud Suite. The last section features detailed Bill of Materials configurations for the Lenovo ThinkSystem servers, storage, and networking hardware that are used in the solution.

See also the Reference Architecture for VMware Software Defined Data Center: lenovopress.com/lp0661 which uses VMware vSAN for hyperconverged storage and NSX for software defined networking.

2 Business problem and business value

The following section provides a summary of the business problems that this reference architecture is intended to help address, and the value that this solution can provide.

2.1 Business problem

With rising costs and complexity, it is becoming increasingly harder to manage IT infrastructure in a data center. As it changes over time, the infrastructure becomes more fragile and more difficult to know the impacts of making changes.

Overlaid on the infrastructure issues are the business demands to both reduce costs and at the same time provide more flexible applications that can meet the end-user demands for stability, performance, availability, and easier upgradability.

2.2 Business value

VMware vCloud Suite provides all the software needed for building an Enterprise infrastructure that is flexible, easy to manage and easy to change for future needs. Together with the addition of policy driven configuration and on demand provisioning, vCloud Suite makes it easier to manage, extend and upgrade the underlying infrastructure.

The Lenovo solution for VMware vCloud Suite provides businesses with an affordable, interoperable, and reliable industry-leading cloud solution to manage all of their virtualized workloads. Built around the latest Lenovo ThinkSystem servers, storage, and networking hardware, this reference architecture provides VMware sizing recommendations, deployment best practices, and validated configurations for compute, management, networking, and storage.

3 Requirements

This section describes the functional and non-functional requirements for this reference architecture.

3.1 Functional requirements

The following section describes the functional requirements that are needed for typical cloud deployments. Figure 1 shows a simplified use-case model for cloud deployments.

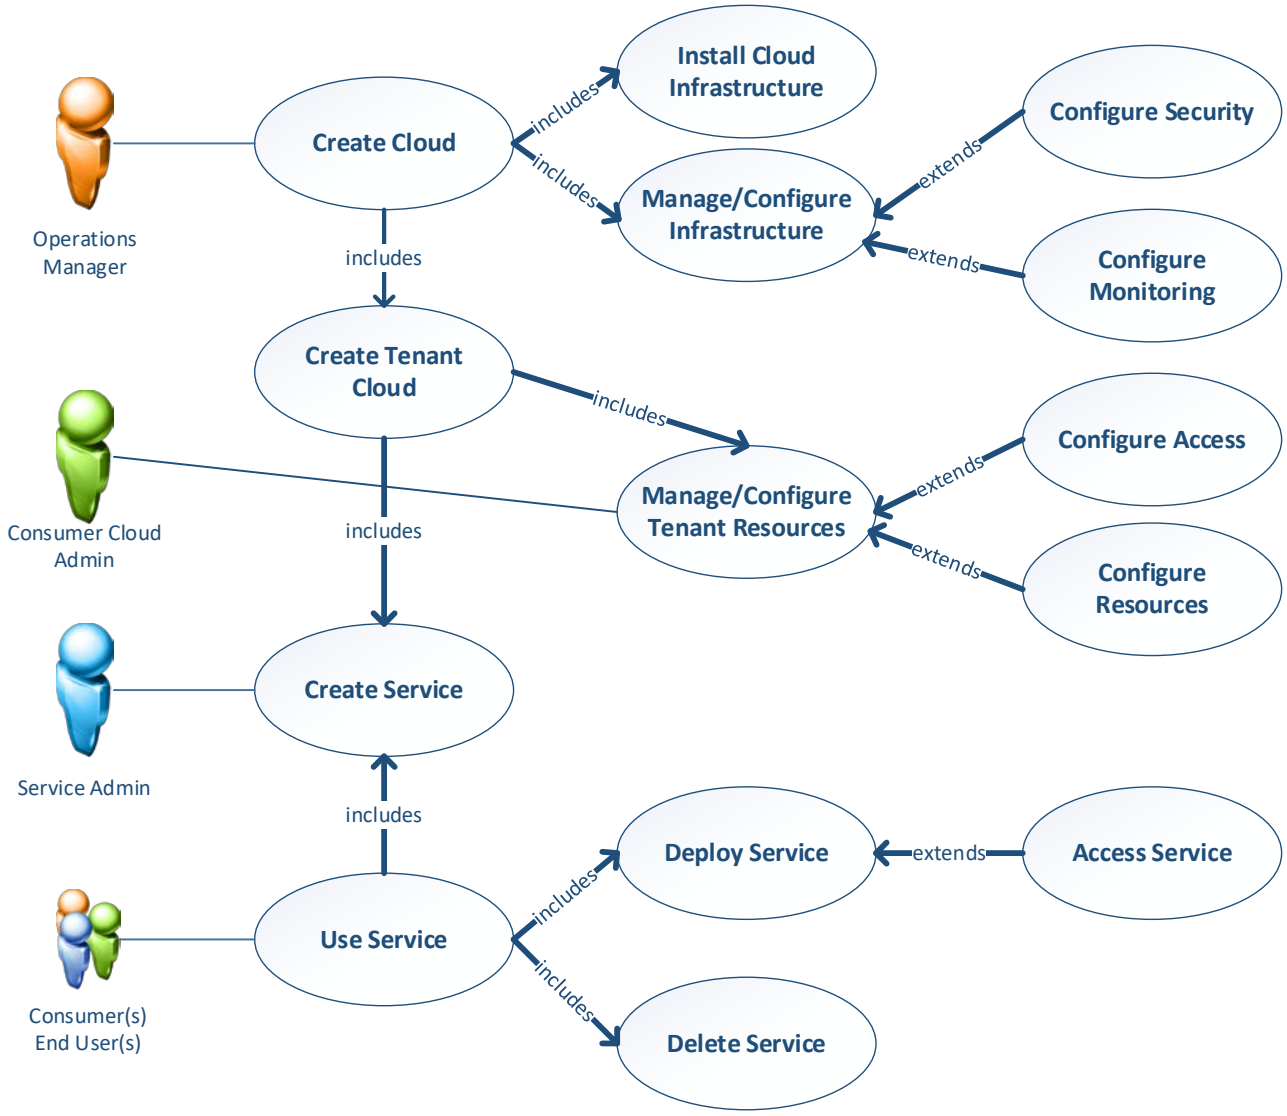

Figure 1: Use case model

Table 1 lists the functional requirements.

Table 1: Functional requirements

Requirement name	Description
Virtualization	Solution supports compute, storage, and network virtualization
Monitoring, event and capacity management	Monitors the health of the cloud infrastructure, collection and management of exception events, and capacity planning
Self-service automation	Solution provides on boarding, provisioning, and management of services and VMs from a service catalog
Approval and workflow	Provides the capability to approve, modify, deny, and delegate service requests
Cloud administration	Provides capabilities to administer a cloud environment, such as adding storage or computational resources in the cloud pool or defining new segregated networks
Image management	Provides capabilities to create VMs, establish version control, search for and compare images, and delete images from the virtual images templates repositories
Service management	Provides capabilities to create services, establish version control, search for services, and delete services from the service templates catalog repositories
Access and authorization Controls	Provides the capabilities to create users and groups and to establish authorization to certain features in the cloud, such as tenant cloud administration, service developer, and user service requester
Virtual Machine Migration	Migrate applications, virtual machine and templates between private and public clouds.
Migrate Security Policies	Migrate network and security policies such as firewall rules to public cloud and vice versa,
Network Extension	Retain virtual machines network properties (L2 and L3) across clouds.
Catalog Management	Maintain common catalog for templates across clouds.

3.2 Non-functional requirements

Table 2 lists the non-functional requirements that are needed for typical cloud deployments.

Table 2: Non-functional requirements

Requirement name	Description
Backup/Recovery	Solution support for integrated backup
Ease of installation	Reduced complexity for solution deployment
Ease of management/operations	Simple management of infrastructure and cloud software
Supportability	Available vendor support
Scalability	Solution components scale with increase in number of concurrent users, VMs/services provisioned per minute or per hour
Flexibility	Solution supports variable deployment methodologies
Security	Solution provides ways to secure customer data
Reliability, availability, and serviceability (RAS)	High availability and resiliency of cloud management and managed infrastructure

4 Architectural overview

This section gives an architectural overview of vCloud Suite products. Figure 2 gives an overview of how those products are deployed into shared edge and compute, management, and additional compute clusters.

This separation of function into these clusters allows for scaling in larger environments.

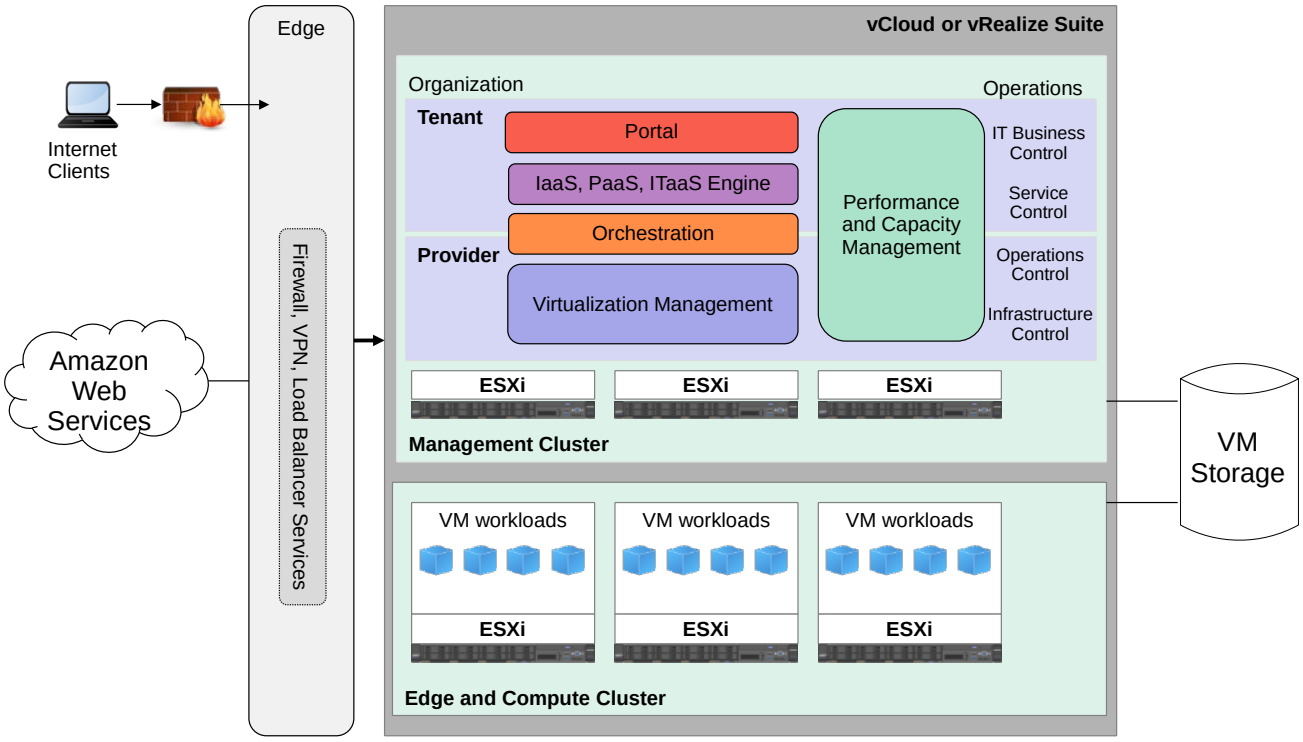

Figure 2: Conceptual design of vCloud Suite

The management cluster runs the components required to support vCloud Suite and is used for management, monitoring, and infrastructure services. A management cluster provides resource isolation which helps these services to operate at their best possible performance level. A separate cluster can satisfy an organization's policy to have physical isolation between management and production hardware and a single management cluster is required for each physical location.

The shared edge and compute cluster supports virtualized infrastructure services as well as network devices that provide interconnectivity between environments. It provides protected capacity by which internal data center networks connect via gateways to external networks. Networking edge services and network traffic management occur in this cluster and all external facing network connectivity ends in this cluster. The shared edge and compute cluster also supports the delivery of all other (non-edge) customer workloads and there can be one or more compute clusters, depending on the customer environment. Multiple compute clusters can be for different organizations or tenants, different workload types, or to spread the load in a large enterprise.

5 Component model

This section describes the component model for VMware vCloud Suite and optionally extending it into public clouds with hybrid cloud connections. Lastly the HyTrust suite of software is described which provides additional security protection features.

5.1 vCloud Components

Figure 3 shows an overview of the major components of the VMware vCloud Suite.

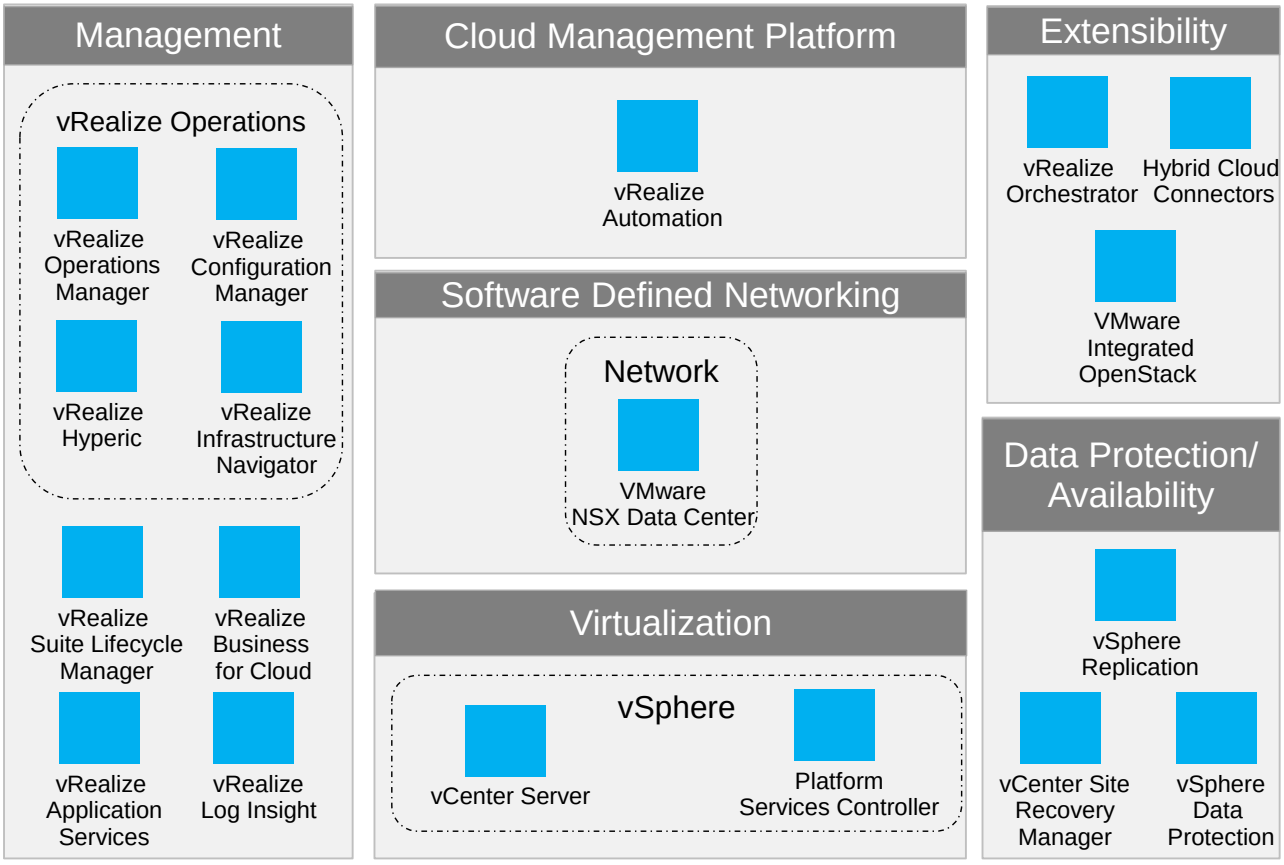

Figure 3: vCloud Suite components

The VMware vCloud Suite features the following components:

ESXi hypervisor	Provides bare-metal virtualization of servers so you can consolidate your applications on less hardware.
vCenter Server	Provides a centralized platform for managing vSphere environments and includes vSphere replication and vSphere data protection.
Platform Services Controller (PSC)	Provides a set of common infrastructure services that encompasses single sign-on (SSO), licensing, and a certificate authority (CA).

vRealize Suite Lifecycle Manager	Provides deployment options such as install, configure, import, and upgrade vRealize Suite environments and perform drift analysis and view the health of those environments
vRealize Automation	Provides a self-service, policy-enabled IT and application services catalog for deploying and provisioning of business-relevant cloud services across private and public clouds, physical infrastructure, hypervisors, and public cloud providers.
vRealize Operations	Provides a set of components for automation of operations including infrastructure health, configurations and compliance, application discovery, and monitoring of hardware and software.
• vRealize Operations Manager	Provides comprehensive visibility and insights into the performance, capacity and health of your infrastructure.
• vRealize Configuration Manager	Provides automation of configuration and compliance management across your virtual, physical, and cloud environments, which assesses them for operational and security compliance.
• vRealize Infrastructure Navigator	Provides automated discovery of application services, visualizes relationships, and maps dependencies of applications on virtualized compute, storage, and network resources.
• vRealize Hyperic	Provides monitoring of operating systems, middleware, and applications that are running in physical, virtual, and cloud environments.
vRealize Business for Cloud	Provides transparency and control over the costs and quality of IT services that are critical for private (vCloud Suite) or hybrid cloud (vRealize Suite) success.
vRealize Log Insight	Provides analytics capabilities to unstructured data and log management, which gives operational intelligence and deep, enterprise-wide visibility across all tiers of the IT infrastructure and applications. Standard for vRealize Suite.
vCenter Site Recovery Manager (SRM)	Provides disaster recovery capability with which you can perform automated orchestration and non-disruptive testing for virtualized applications by using ESXi hypervisor only. SRM is standard for vCloud Suite and optional for vRealize Suite.
vRealize Orchestrator	Provides the capability to create workflows that automate activities, such as provisioning VM, performing scheduled maintenance, and starting backups.
NSX	NSX provides virtualization of networking and is optional for vCloud Suite deployments.

VMware Integrated OpenStack (VIO)	Provides a VMware-supported OpenStack distribution (distro) that makes it easier for IT to run a production-grade, OpenStack-based deployment on top of their VMware infrastructure. For more information, see this website: vmware.com/products/openstack.
Hybrid Cloud Connectors	Allows an administrator to provide hybridization using public cloud providers such as Amazon AWS. See the next section for more information.

The vCloud Suite products also have dependencies on the following external components:

Identity source	Identity sources (Active Directory, OpenLDAP, or Local OS) or similar is required to implement and operate the vCloud Suite or vRealize Suite infrastructure.
DNS	DNS must be configured for connectivity between vCenter Server, Active Directory, ESXi hosts, and the VMs
DHCP/TFTP	PXE boot is required for vSphere Auto Deploy functionality.
Time synchronization	Accurate time keeping and time synchronization is critical for a healthy infrastructure. All components (including ESXi hosts, vCenter Server, the SAN, physical network infrastructure, and VM guest operating systems) must have accurate time keeping.
Microsoft SQL Server or Oracle database	Many of the vCloud Suite components come with embedded databases or they can use external databases such as Microsoft SQL Server or Oracle, depending on the component and the intended environment.

Other software components such as Lenovo XClarity Administrator are not shown. As well as providing management of Lenovo hardware, XClarity Administrator also has plugins for VMware vCenter, VMware vRealize Orchestrator, and VMware vRealize Log Insight which are further described in “Systems management” on page 34.

5.2 Hybrid Clouds

On-premise VMware vCloud suite connects to public clouds such as Amazon Web Services (AWS), IBM Public Cloud and OVH Public Cloud. This document only discusses AWS support.

The Amazon Elastic Compute Cloud (EC2) provides scalable computing capacity in the Amazon Web Services (AWS) public cloud by offering compute, storage, networking, software, and development tools. AWS provides Virtual Private Cloud and Dedicated Hosts for compute and different services. It supports a hybrid architecture by integrating networking, security and access control, automated workload migrations and controlling AWS from an on-premise infrastructure management tool.

AWS Server Migration Service (AWS SMS) is an agentless service to migrate on-premise workloads from vCenter to AWS. It supports incremental replications and migration can be performed faster while minimizing network bandwidth consumption and reducing server downtime. Each server volume replicated is saved as a

new Amazon Machine Image (AMI) which can be launched as an EC2 instance (VM) in the AWS cloud. The AWS Server Migration Service replicates server volumes from on-premises environment to S3 temporarily and purges them from S3 immediately after creating the Elastic Block Store (EBS) snapshots.

Table 3 describes the features supported by the AWS SMS connector. Some best practices for deployment are described in “Hybrid networking to public clouds” on page 44.

Table 3: Hybrid Cloud Features

Feature	AWS SMS Connector
Bi-Directional Migration	No
Integration with vCenter	Yes
vCenter→Connector Association	Many-Many
Connector→Public Cloud Association	1-1
vCenter linked Mode support	No
Integration with vSphere client	Use AWS Management Console
Integration with vSphere web client	Use AWS Management Console
Integration with vSphere Replication	No
Integration with vRealize Automation	No
Multi-tenant support	No
VM Management public cloud	Use AWS Management Console
VM Management on-premise	No
Migration to all public cloud regions	Limited Currently
Copy/Migrate Templates to public cloud	Yes (AMI)
Deploy VM from Template to public cloud	Yes (AMI)
Live Migration to public cloud	Yes
Cold Migration to public cloud	Yes
Bulk Migration to public cloud	Yes
Layer 2 Extension	No
Migrate VM to on-premise	No (OVA download)
Offline Data Transfer	No
Common Content Library	No
Number of Concurrent Migration	50 per account
License	90 days Free
vSphere Standard Switch Support	Yes
vSphere Distributed Switch Support	Yes
Network(NSX) Policy Migration	No

5.3 VMware Licensing

The licensing for vSphere is based on a CPU metric and licensing for other products is based on the number of OS instances. The vCloud Suite license is a single perpetual license that includes vSphere Enterprise Plus and vRealize Suite. For vSphere environments, the license can be purchased separately for vSphere Enterprise Plus and vRealize Suite. Other components have their own separate licenses and are optional add-ons. Table 4 lists the standard and optional components that are provided with a vCloud Suite License or vRealize Suite License.

Table 4: VMware Licensing

License	Component	vCloud Standard	vCloud Advanced	vCloud Enterprise	vRealize Standard	vRealize Advanced	vRealize Enterprise
Base	vSphere	Enterprise Plus	Enterprise Plus	Enterprise Plus			
	vRealize Suite Lifecycle Mgr	Included	Included	Included	Included	Included	Included
	vRealize Automation	N/A	Advanced	Enterprise	N/A	Advanced	Enterprise
	vRealize Operations • vRealize Operations Mgr • vRealize Configuration Mgr • vRealize Infrastructure Nav • vRealize Hyperic	Advanced	Advanced	Enterprise	Advanced	Advanced	Enterprise
	vRealize Business for Cloud	Standard	Advanced	Advanced	Standard	Advanced	Advanced
	vSphere Replication	Included	Included	Included	N/A	N/A	N/A
	vSphere Data Protection	Included	Included	Included	N/A	N/A	N/A
	vSphere Big Data Extensions	Included	Included	Included	N/A	N/A	N/A
	vRealize Log Insight	Included	Included	Included	Included	Included	Included
	vCenter Site Recovery Mgr	Add On	Add On	Add On	Add On	Add On	Add On
Other licenses	vRealize Orchestrator	Included	Included	Included	Add On	Add On	Add On
	vRealize Automation Public Cloud Extension	N/A	Add On	Add On	Add On	Add On	Add On
	vRealize Operations Public Cloud Extension	N/A	Add On	Add On	Add On	Add On	Add On
	NSX Data Center for vSphere	Included	Included	Included	Add On	Add On	Add On
	VMware Integrated OpenStack	Add On	Add On	Add On	Add On	Add On	Add On

5.4 HyTrust Security

HyTrust provides a suite of security-oriented products for vSphere environment. These products are HyTrust KeyControl, DataControl, and CloudControl. Note that the HyTrust products are currently supported on ESXi 6.5 U2 and NSX 6.3.1.

5.4.1 HyTrust KeyControl

HyTrust KeyControl (HTKC) enables enterprises to easily manage all their encryption keys at scale, how often they rotate them, and how they are shared securely. HyTrust KeyControl capabilities include:

- VMWare Certified Key Manager Server (KMS) for:
  - vSphere 6.5 and 6.7
- Universal key management for KMIP-compatible encryption agents
- Enterprise scalability and performance
- KeyControl can run in an active-active, high availability cluster
- FIPS 140-2 Level 1 validation and FIPS 140-2 Level 3 hardware security module (HSM)

5.4.2 HyTrust DataControl

HyTrust DataControl (HTDC) secures multi-cloud workloads throughout their lifecycle. DataControl helps manage workloads and encryption keys from a central location to reduce complexity, comply with regulations such as the GDPR.

DataControl provides granular encryption for better multi-cloud security. The protection boundary does not stop at the hypervisor or at the data store; VMs are individually encrypted. Inside the VM, unique keys can be assigned to encrypt individual partitions, including the boot (OS) disk. Encryption and rekeying can be done on the fly and there is no need to take workloads off-line.

Table 5 compares the data encryption features of vSphere and HyTrust DataControl/KeyControl.

Table 5: Comparison of Encryption Features

Encryption	vSphere VM Encryption	HyTrust DataControl
Protection level	Data at rest and in motion	Data at rest
Encryption Approach	Hypervisor does the encryption	In Guest encryption
Components	KMS, vCenter, ESXi Host	KMS, HyTrust DataControl Agent
Encryption Cipher	AES-XTS-256	AES-XTS-512,AES-XTS-256, AES 128
Encrypted objects	Virtual machine files, virtual disk files, and ESXi core dump files	All data in the drives
Interface	vSphere Web Client, vSphere Web Services SDK	HyTrust DataControl UI in the Guest OS. HyTrust KeyControl UI to manage VM Set, VMs and users
Enabling Option	Per VM level through vSphere Encryption Storage Policy	Enabled within Guest OS
Access Control	Users with vSphere Cryptographic Operations Privileges	Guest OS User uses KeyControl admin user. Authorization can also be done by HyTrust CloudControl
Interoperability Limitations	vSphere Fault Tolerance, vSphere Replication, Content Library	N/A
Platform Support	All Guest OS running on the Hypervisor	Most Windows and Linux flavors and version running on vSphere, KVM, Hyper-V, or XenSever

5.4.3 HyTrust CloudControl

HyTrust CloudControl (HTCC) provides a variety of security and policy enhancements without impacting the existing GUI of vSphere, NSX and ESXi. CloudControl is deployed as a transparent proxy and mediates the

actions taken by administrators using familiar interfaces. CloudControl provides the following security features:

- **Role Based Access Control** (RBAC) to control which functions have access to what resources and allows a much closer alignment of access rights to governance and compliance requirements.
- **Policy Control including Two Man Rule** to define and more importantly enforce policy including requiring secondary approval for potentially disruptive actions, reducing potential impact of human error or intentional malevolent behaviour.
- **Access Control including Two Factor Authentication** to significantly enhance the overall security posture of an organization without the traditional weaknesses of using even strong passwords.
- **Forensic grade logs** to provide an in-depth perspective on what has happened as well as what has not happened in your virtual environment.

Table 6 compares the access control features of vCenter and HyTrust CloudControl.

Table 6: Comparison of access control features

Access Control Feature	vCenter	HyTrust CloudControl
vSphere Web Client Access	vCenter URL	Published IP (PIP) associated with vCenter
Authentication	vCenter SSO, IWA	vCenter SSO, IWA, HTCC Service Account, Two factor authentication with RSA Secure ID, RADIUS, or TACACS+
Authorization	Predefined permissions to access various vCenter components	Uses permissions defined in vCenter
vCenter Users	SSO users from multiple AD Domain and vSphere local domain. Predefined solution users for vSphere services.	Users from Single AD Domain which includes configured HTCC Service Account
vCenter User Access Setup	Directory users/group need to be added in vCenter SSO users/group	Directory users need to be added to respective HTCC directory group which is associated with HTCC role
User Groups	14 predefined SSO groups. Directory users/group is mapped to SSO groups.	16 predefined rules for vSphere. HTCC directory group is mapped to HTCC rule.
Role Based Access Control	14 predefined roles with respective privileges	16 predefined roles for vSphere with appropriate privileges
Custom Roles Creation	Supported	Supported
Secondary Approval	Not Available	Available for set of compute and network operations

Access Control Feature	vCenter	HyTrust CloudControl
Auditing	Integrated with vRealize Log Insight. Auditing dashboard is available based on the event type. User's session details can be monitored in vSphere web client.	Has its own Log Viewer and dashboard. Logs can be redirected to use vRealize Log Insight as syslog server.

5.4.4 Compliance Management

An important part of security is compliance management. VMware vRealize Configuration Manager has twenty built-in compliance templates and others can be added. HyTrust CloudControl (HTCC) supports customizing built-in compliance templates but does not provide any out of the box.

Table 7 compares the compliance management features of vRealize Configuration Manager and HTCC.

Table 7: Comparison of compliance management features

Compliance Management Feature	vRealize Configuration Manager	HTCC
ESXi Host Compliance	Yes	Yes
Guest Virtual Machine Compliance	Yes	Limited
NSX Manager Compliance	No	Yes
Patching assessment and Deployment	Yes	No
Active Directory Compliance	Yes	No
Software Asset Management	Yes	No
Integration with vRealize Operation Manager	Yes	No
Manage Virtual Machines	Yes	No

6 Operational model

This section describes the options for mapping the logical components of vCloud Suite onto Lenovo ThinkSystem servers, storage, and Lenovo network switches. Each subsection contains performance data, has recommendations for how to size for that particular hardware, and a pointer to the BOM configurations that are described in Section 8 on page 49.

6.1 Hardware components

The following section describes the hardware components in a vCloud deployment.

6.1.1 Rack Servers

You can use various rack-based Lenovo ThinkSystem server platforms to implement edge, management, or compute clusters.

Lenovo ThinkSystem SR630

Lenovo ThinkSystem SR630 (as shown in Figure 4) is an ideal 2-socket 1U rack server for small businesses up to large enterprises that need industry-leading reliability, management, and security, as well as maximizing performance and flexibility for future growth. The SR630 server is designed to handle a wide range of workloads, such as databases, virtualization and cloud computing, virtual desktop infrastructure (VDI), infrastructure security, systems management, enterprise applications, collaboration/email, streaming media, web, and HPC. The ThinkSystem SR630 offers up to twelve 2.5-inch hot-swappable SAS/SATA HDDs or SSDs together with up to 4 on-board NVMe PCIe ports that allow direct connections to the U.2 NVMe PCIe SSDs.

Figure 4: Lenovo ThinkSystem SR630

For more information, see this website: lenovopress.com/lp0643

Lenovo ThinkSystem SR650

Lenovo ThinkSystem SR650 (as shown in Figure 5) is similar to the SR630 but in a 2U form factor.

Figure 5: Lenovo ThinkSystem SR650

The key differences compared to the SR630 server are more expansion slots and chassis to support up to twenty-four 2.5-inch or fourteen 3.5-inch hot-swappable SAS/SATA HDDs or SSDs together with up to 8 on-board NVMe PCIe ports that allow direct connections to the U.2 NVMe PCIe SSDs. The ThinkSystem SR650 server also supports up to two NVIDIA GRID cards for graphics acceleration.

For more information, see this website: lenovopress.com/lp0644

6.1.2 Bladed Servers using Lenovo Flex System

Flex System is a Lenovo enterprise-class platform that is specifically created to meet the demands of a virtualized data center and help clients establish a highly secure, private cloud environment. The Flex System includes the following features:

- Greatest choice for clients in processor type and OS platform (all in the same chassis) that is managed from a single point of control.
- The Flex System networking delivers 50% latency improvement through node-to-node (east-west) traffic rather than routing everything through the top-of-rack (ToR) switch (north-south).

Flex Enterprise Chassis

Flex System is anchored by the Flex System Enterprise Chassis which enables high-speed performance with integrated servers and networking. Furthermore, its flexible design can meet the needs of varying workloads with independently scalable IT resource pools for higher utilization and lower cost per workload.

Figure 6 shows the front and rear of the Flex Chassis with fourteen ThinkSystem SN550 compute nodes in the front and two 10 Gbe switches in the rear with power supplies and fans.

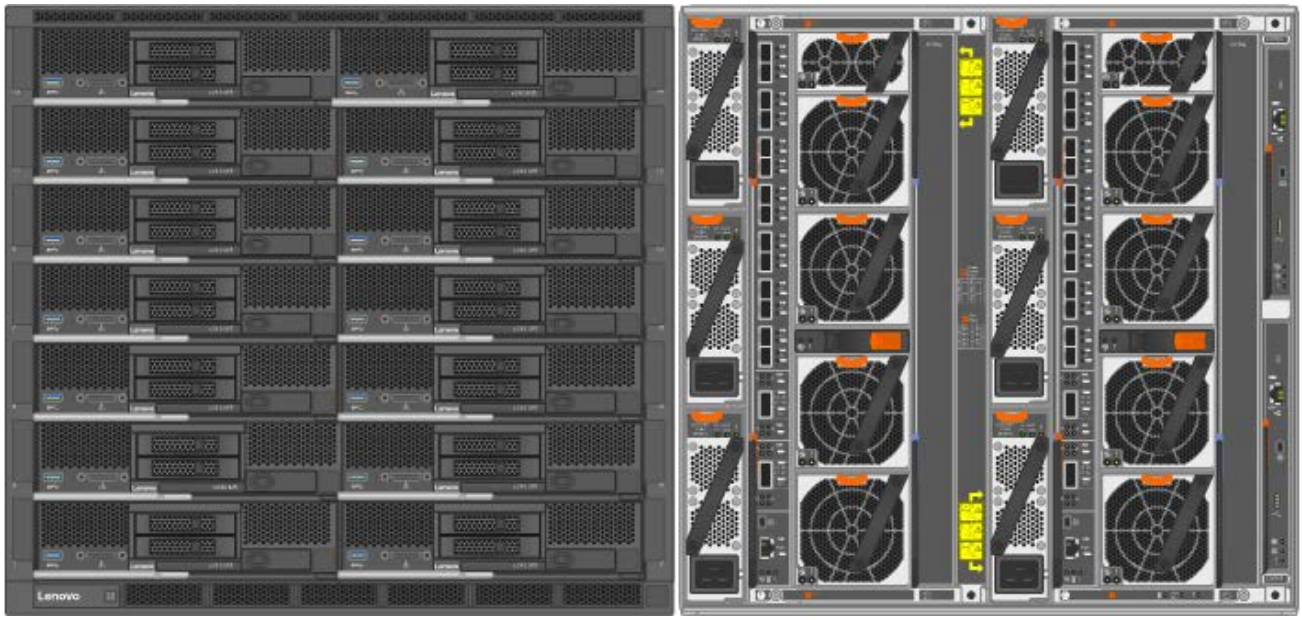

Figure 6: Flex System Enterprise Chassis front and rear

For more information, see the following websites:

- Lenovo Flex System Enterprise Chassis: lenovopress.com/tips0863

Lenovo ThinkSystem SN550 compute node

The Lenovo ThinkSystem SN550 compute node (as shown in Figure 7) is a high-performance server that offers enhanced security, efficiency, and reliability features to handle business-critical workloads. The blade server incorporates Intel Xeon Processor Scalable Family of processors. The processors feature up to 28 cores each and includes industry-leading two-socket memory capacity that features Lenovo TruDDR4™ memory up to 1.5 TB.

Figure 7: ThinkSystem SN550 Compute Node

For more information, see this website: lenovopress.com/lp0637

6.1.3 10 GbE networking

The standard network for vCloud Suite is 10 GbE. The following Lenovo 10GbE ToR switches are recommended:

- Lenovo RackSwitch G8272
- Lenovo Flex System Fabric SI4093

For more information about network switches, see this website:
shop.lenovo.com/us/en/systems/networking/ethernet-rackswitch

Lenovo RackSwitch G8272

The Lenovo RackSwitch G8272 that uses 10 Gb SFP+ and 40 Gb QSFP+ Ethernet technology is specifically designed for the data center. It is ideal for today's big data, cloud, and optimized workload solutions. It is an enterprise class Layer 2 and Layer 3 full featured switch that delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data center-grade buffers help keep traffic moving, while the hot-swap redundant power supplies and fans (along with numerous high-availability features) help provide high availability for business sensitive traffic.

The RackSwitch G8272 (shown in Figure 8), is ideal for latency sensitive applications, such as high-performance computing clusters and financial applications. In addition to the 10 Gb Ethernet (GbE) and 40 GbE connections, the G8272 can use 1 GbE connections. The G8272 supports the newest protocols, including Data Center Bridging/Converged Enhanced Ethernet (DCB/CEE) for Fibre Channel over Ethernet (FCoE), iSCSI and network-attached storage (NAS).

Figure 8: Lenovo RackSwitch G8272

For more information, see this website: lenovopress.com/tips1267

Lenovo Flex System Fabric SI4093 system interconnect module

The Lenovo® Flex System Fabric SI4093 System Interconnect Module (as shown in Figure 9) enables simplified integration of Flex System™ into networking infrastructure and provides the capability of building simple connectivity for points of delivery (PODs) or clusters up to 126 nodes.

Figure 9: Lenovo Flex System Fabric SI4093

The base switch configuration comes standard with 24x 10 GbE port licenses that can be assigned to internal connections or external SFP+ or QSFP+ ports using a flexible port mapping. For example, by using this feature customers can trade off four 10 GbE ports for one 40 GbE port (or vice versa) or trade off one external 10 GbE SFP+ port for one internal 10 GbE port (or vice versa). Customers then have the flexibility of turning on more ports when needed by using Features on Demand upgrade licensing capabilities that provide “pay as you grow” scalability without the need to buy more hardware.

The SI4093 provides transparent Flex System connectivity to top of rack (ToR) switches. The SI4093 aggregates compute node ports by appearing as a simple pass-thru device. The upstream network sees a “large pipe” of server traffic coming to and from the chassis, with the main difference being that intra-chassis switching is supported. For more information, see this website: lenovopress.com/tips1292

6.1.4 1 Gbe networking

The following Lenovo 1GbE ToR switch is recommended for use with vCloud Suite:

- Lenovo RackSwitch G8052

Lenovo RackSwitch G8052

The Lenovo System Networking RackSwitch G8052 (as shown in Figure 10) is an Ethernet switch that is designed for the data center and provides a virtualized, cooler, and simpler network solution. The Lenovo RackSwitch G8052 offers up to 48 1 GbE ports and up to 4 10 GbE ports in a 1U footprint. The G8052 switch is always available for business-sensitive traffic by using redundant power supplies, fans, and numerous high-availability features. For more information, see this website: lenovopress.com/tips0813

Figure 10: Lenovo RackSwitch G8052

6.1.5 Shared storage

Shared storage can be used as the primary storage for a VMware vCloud deployment.

Lenovo ThinkSystem DM5000F storage array

Lenovo ThinkSystem DM5000F is a unified, all flash entry-level storage system that is designed to provide performance, simplicity, capacity, security, and high availability for medium to large businesses. Powered by the ONTAP software, ThinkSystem DM5000F delivers enterprise-class storage management capabilities with a wide choice of host connectivity options and enhanced data management features. The ThinkSystem DM5000F is a perfect fit for a wide range of enterprise workloads, including big data and analytics, artificial intelligence, engineering and design, enterprise applications, and other storage I/O-intensive applications.

ThinkSystem DM5000F models (as shown in Figure 11) are 2U rack-mount controller enclosures that include two controllers, 64 GB RAM and 8 GB battery-backed NVRAM (32 GB RAM and 4 GB NVRAM per controller), and 24 SFF hot-swap drive bays (2U24 form factor). Controllers provide universal 1/10 GbE NAS/iSCSI or 8/16 Gb Fibre Channel (FC) ports for host connectivity.

A single ThinkSystem DM5000F Storage Array scales up to 144 solid-state drives (SSDs) with the attachment of Lenovo ThinkSystem DM240S 2U24 SFF Expansion Enclosures. Up to 12 DM5000F Storage Arrays can be combined into a clustered system in a NAS environment, or up to 6 DM5000F Storage Arrays can be combined into a clustered system in a SAN environment.

Figure 11: Lenovo DM5000F storage array

The ThinkSystem DM5000F offers the following key features and benefits:

- All-flash array capabilities to meet the demand for higher speed storage and provide higher IOPs and bandwidth with lower power usage and total cost of ownership than hybrid or HDD-based solutions.
- Improved performance and data protection with RAID-DP and RAID-TEC, as well as support for traditional RAID 4.
- Flexible host connectivity to match diverse client needs with support for unified NAS and SAN storage protocols, including 1/10 GbE NAS and iSCSI, and 8/16 Gb Fibre Channel connectivity.
- 12 Gb SAS drive-side connectivity with multipathing with up to 24x 2.5-inch small form factor (SFF) drives in the 2U24 SFF enclosures.
- Rich set of standard storage management functions available at no extra cost, including snapshots, volume copy, quality of service, thin provisioning, compression, deduplication, encryption, disk-based

backup, application- and virtual machine-aware backup, quick data recovery, and asynchronous mirroring.

- Optional WORM (write once, read many) (SnapLock) licensed function to reinforce permanence and integrity of stored data and to ensure compliance with applicable regulations.
- Intuitive, web-based GUI for easy system setup and management.
- Lenovo XClarity support for centralized systems management of Lenovo x86 servers, switches, and storage, which provides automated agent-less discovery, inventory, monitoring, and additional platform-specific functions across multiple systems.
- Designed for high availability with redundant hot-swap components, including controllers and I/O modules, power supplies, and non-disruptive firmware upgrades.

For more information, see this website: lenovopress.com/lp0911.

6.1.6 SAN switches

The following Lenovo SAN switch for Fibre Channel can be used with this solution.

Lenovo ThinkSystem DB610S

The Lenovo ThinkSystem DB610S FC SAN Switch provides exceptional price/performance value by delivering market-leading 32 Gb Gen 6 Fibre Channel technology and combining flexibility, simplicity, and enterprise-class functionality to meet the demands of growing flash-based storage environments.

Designed to enable maximum flexibility and reliability, the ThinkSystem DB610S is a compact, 1U rack-mount FC switch that offers low-cost access to industry-leading Storage Area Network (SAN) technology while providing “pay-as-you-grow” scalability to meet the needs of an evolving storage environment.

The DB610S FC SAN Switch (as shown in Figure 12) offers 24x SFP+ ports that support 4/8/16/32 Gbps speeds. The DB610S FC SAN switch provides easy integration into the existing SAN environments while realizing the benefits of Gen 6 Fibre Channel connectivity, and the switch offers a rich set of standard features with the options to expand its capabilities as needed.

The DB610S FC SAN Switch features the EZSwitch Setup wizard and can be configured in Access Gateway Mode to simplify deployment. The switch provides full non-blocking performance with Ports On Demand scalability to support SAN expansion and enable long-term investment protection.

Figure 12: Lenovo ThinkSystem DB610S SAN Switch

For more information, see this website: lenovopress.com/lp0582.

6.2 Shared edge and compute cluster servers

The shared edge and compute cluster uses its own dedicated vCenter server.

6.2.1 Edge and Infrastructure Services VMs

The VMs used for infrastructure services such as Active Directory, DNS/DHCP, firewalls, proxy and anti-virus are deployed in the shared edge and compute cluster. Table 8 lists each infrastructure service VM with the recommended sizes in terms of virtual CPUs, RAM, storage, and networking.

Table 8: Infrastructure services VMs

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
AD, DHCP, DNS server	2	4	70	1 GbE	clustered
http proxy server	2	4	30	1 GbE	clustered

6.2.2 Hybrid cloud VMs

Table 9 lists the cloud connectivity VMs with the recommended sizes in terms of virtual CPUs, RAM, storage, networking, and location. Note that these VMs do not have options for high availability.

Table 9: Cloud connectivity VMs

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	Location
AWS SMS Connector for vCenter	2	4	300	1 GbE	On-Premise

6.2.3 Server configuration

Since the shared cluster hosts compute workloads and edge services, the servers need to be sized appropriately. Refer to “Compute cluster servers” on page 24 for sizing guidance.

The following configuration is recommended:

- 2 x Intel® Xeon® Gold 6230 Processor (2.10 GHz 20 cores)
- 384 GB of system memory
- 2 x Dual port 10 GbE network card
- Dual M.2 boot drives for ESXi
- ESXi 6.7u1

The shared edge and compute cluster should have a minimum of two hosts for high availability. For more information, see “Server BOM” on page 49.

6.2.4 Load balancing and protection

An essential part of the infrastructure is load balancing of the server VMs and recognizing when a server is down and failing over to a second server.

For the shared edge and compute cluster connected to the Internet, it is also important to provide a firewall and protection against external threats. There are many ways to solve these problems such as using a F5 Big-IP edge gateway device or virtual machine. Using F5 protection and load balancing is outside the scope of this document.

6.3 Management cluster servers

The number of VMware vCloud Suite components in the management cluster increases as capabilities are added. This section addresses the management components that could be used. Third party add-ons must be sized separately.

6.3.1 Management cluster VMs

There are several considerations that contribute to an end-to-end sizing of an entire VMware vCloud environment including Lenovo software for systems management. This section is intended to provide some high-level guidance for management cluster configuration sizing. The recommended number of virtual CPUs, memory size, storage size, and network bandwidth is given for each VM and the VMs are grouped by each major component or appliance.

An essential part of the infrastructure is load balancing of the server VMs and recognizing when a server is down and failing over to another server. The following cases are available for VMs in the management cluster:

- vSphere HA: vCenter automatically restarts the VM on another server, but there is some downtime while the VM starts up.
- Microsoft SQL server clustering: The SQL server cluster automatically handles failover.
- Clustering within component to provide built-in high availability.
- Load balancing: An external load balancer such as a Big-IP switch from F5

Table 10 lists each management cluster VM for vSphere with its recommended size in terms of virtual CPUs, RAM, storage, and networking.

Table 10: Management cluster VMs for vSphere

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
vCenter Server(1) Management Cluster	8	24	50	1 GbE	load balancer
vCenter Server(2) Edge and Compute Cluster	8	24	50	1 GbE	load balancer
vCenter Server Database (MS SQL)	4	8	200	1 GbE	SQL AlwaysOn Availability Group
Platform Service Controller (1) Management Cluster	2	4	50	1 GbE	load balancer
Platform Service Controller (2) Edge and Compute Cluster	2	4	50	1 GbE	load balancer
vSphere Replication	2	4	20	1 GbE	not required
vSphere Data Protection	4	4	1600	1 GbE	not required
vRealize Orchestrator Appliance	2	3	12	1 GbE	Clustered

Table 11 lists each management cluster VM for vRealize Automation with its size in terms of virtual CPUs, RAM, storage, and networking.

Table 11: Management cluster VMs for vRealize Automation

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
vRealize Suite Lifecycle Manager	4	16	135	1 GbE	N/A
vRealize Automation Appliance	4	16	30	1 GbE	load balancer
IaaS Database (MS SQL)	8	16	100	1 GbE	SQL AlwaysOn Availability Group
Infrastructure Web Server	2	4	40	1 GbE	load balancer
Infrastructure Manager Server	2	4	40	1 GbE	load balancer
Distributed Execution Manager (DEM)	2	6	40	1 GbE	load balancer
vSphere Proxy Agent	2	4	40	1 GbE	load balancer
vRealize Application Services	8	16	50	1 GbE	vSphere HA

Table 12 lists each management cluster VM for vRealize Operations Manager with its size in terms of virtual CPUs, RAM, storage, and networking.

Table 12: Management cluster VMs for vRealize Operations Manager

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
vRealize Operations Manager – Master	4	16	500	1 GbE	clustered
vRealize Operations Manager – Data	4	16	500	1 GbE	not required
vRealize Configuration Manager – Collector	4	16	150	1 GbE	load balancer
vRealize Configuration Manager Database (MS SQL)	4	16	1000	1 GbE	SQL AlwaysOn Availability Group
vRealize Hyperic Server	8	12	16	1 GbE	load balancer
vRealize Hyperic Server - Postgres DB	8	12	75	1 GbE	load balancer
vRealize Infrastructure Navigator	2	4	24	1 GbE	not required

Table 13 lists each of the remaining management cluster VMs.

Table 13: Other Management cluster VMs

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High Availability
vRealize Business for Cloud	2	4	50	1 GbE	vSphere HA
Site Recovery Manager	4	4	20	1 GbE	not required
Site Recovery Manager Database (MS SQL)	2	4	100	1 GbE	SQL AlwaysOn Availability Group
vRealize Log Insight	8	16	100	1 GbE	Cluster of 3 nodes

Table 14 lists each management cluster VM for HyTrust with its size in terms of virtual CPUs, RAM, storage, and networking.

Table 14: Management cluster VMs for HyTrust

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
HyTrust CloudControl	4	16	70	1 GbE	Clustered
HyTrust KeyControl	2	8	20	1 GbE	Clustered

Table 15 lists the VMs that are needed for Lenovo software for systems management.

Table 15: Lenovo System Management VMs

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
Lenovo XClarity Administrator	2	4	64	1 GbE	not required
Lenovo XClarity Integrator (Windows OS)	1	2	30	1 GbE	not required

6.3.2 Server configuration

The management cluster should have a minimum of four hosts for high availability. Because of the large number of management VMs that can be used in the management cluster, the following configuration is recommended for each server:

- 2 x Intel® Xeon® Gold 6230 Processor (2.10 GHz 20 cores)
- 384 GB of system memory
- 2 x Dual port 10 GbE network card
- Dual M.2 boot drives for ESXi
- ESXi 6.7u1

For more information about the server configuration, see “Server BOM” on page 49.

6.4 Compute cluster servers

A detailed sizing for the compute cluster requires more information that is not in scope for this reference architecture; that is, it depends entirely on the virtualized applications.

To provide some general guidance, this section describes the expected performance for VM consolidation, and the Yahoo Cloud System Benchmark (YCSB). Lenovo makes no representation or warranty that an individual user can achieve results equivalent to the stated levels.

6.4.1 VM consolidation and estimation

To consolidate VMs onto ESXi servers, it is important to consider overcommitment ratios for CPU and system memory. Overcommitment makes sense because, some VMs often are lightly loaded while others are more heavily loaded, and relative activity levels vary over time. I/O overcommitment in terms of IOPs and storage capacity is not recommended at this time.

The CPU allocation ratio indicates the number of virtual cores that can be assigned to a node for each physical core. A 6:1 ratio is a balanced choice for performance and cost effectiveness on servers with Intel Xeon processors, as shown in the following equation:

```
Virtual CPU (vCPU) = Physical Cores * CPU allocation ratio
```

To improve memory utilization, the ESXi host transfers memory from idle VMs to VMs that need more memory. The Reservation or Shares parameter can be used to preferentially allocate memory to important VMs. ESXi implements various mechanisms, such as ballooning, memory sharing, memory compression, and swapping to provide reasonable performance, even if the host is not heavily memory overcommitted.

ESXi memory compression is enabled by default. When a server's memory becomes overcommitted, ESXi compresses virtual pages and stores them in memory. Because accessing compressed memory is faster than accessing memory that is swapped to disk, memory compression in ESXi allows memory overcommitment without significantly hindering performance. When a virtual page must be swapped, ESXi first attempts to compress the page. Pages that can be compressed to 2 KB or smaller are stored in the VM's compression cache, increasing the capacity of the host.

The RAM allocation ratio is used to allocate virtual resources in excess of what is physically available on a host through ESXi memory compression technology. The RAM allocation ratio can be 1 for no compression and 2 for extreme compression. A reasonable ratio for generalized workloads is 1.25, as shown in the following equation:

```
Virtual Memory (vRAM) = (Physical Memory – Reserved Memory) * RAM allocation
ratio
```

The reserved memory for the ESXi hypervisor often is 4 GB. By using these formulas, the usable virtual resources can be calculated from two example compute node hardware configurations, as shown in Table 16.

Table 16: Virtual resource

Example Configuration	Cores	vCPU	System Memory	vRAM
SR630 "medium"	2 x 16	192	384 GB	475 GB (384-4)*1.25
SR630 "large"	2 x 20	240	768 GB	955 GB (768-4)*1.25

Table 17 lists the assumptions for some representative VM workload characteristics.

Table 17: Representative VM workload characteristics

Workload	vCPU	vRAM	Storage	IOPS
Small	1	2 GB	20 GB	15
Medium	2	6 GB	60 GB	35
Large	6	12 GB	200 GB	100

The three standard workloads of small, medium, and large as listed in Table 17 can be mixed in various proportions to create an "averaged" workload or a completely new workload could be defined, as in Table 18.

Table 18: Example mixed workload

Workload	vCPU	vRAM	Storage	IOPS
mixed	2	8 GB	80 GB	100

By using the mixed workload requirement from Table 18, it is easy to calculate the expected VM density. The smallest number determines the number of VMs that can be supported per compute server.

Table 19 shows that the "medium" configuration is memory bound and supports only a maximum of 59 VMs and that the "large" configuration is more balanced and supports a maximum of 119 VMs of the example workload.

Table 19: Resource-based VM density for mixed workload

Example Configuration	CPU	Memory
SR630 “medium”	192/2 = 96 VMs	475/8 = 59 VMs
SR630 “large”	240/2 = 120 VMs	955/8 = 119 VMs

In any deployment scenario, workloads are unlikely to have the same characteristics, and workloads might not be balanced between hosts. For better resource utilization, consider putting similar workloads on the same guest OS type in the same host to use memory compression.

6.4.2 Yahoo Cloud System Benchmark (YCSB)

This benchmark is used to give an indication of relative storage performance for compute cluster workloads

YCSB is a NoSQL benchmark framework. It contains different workloads to perform read, insert, update, delete, and scan operations to stress NoSQL databases and the I/O performance of storage systems. It supports customizing the workloads with different read and write ratios. The source code can be downloaded from this website: github.com/brianfrankcooper/YCSB.

The YCSB benchmark was run in Ubuntu 16.04.4 LTS VMs with mongodb v3.6.3 with the “WiredTiger” storage engine. Each VM was configured with 2 vCPUs, 4 GB memory, and 50 GB storage.

The benchmark was performed with 96 VMs. In addition, four other Ubuntu 16.04.4 LTS VMs with the YCSB clients were used to simulate the load for the mongodb VMs.

By default, the YCSB benchmark uses a data size of 1 KB records (10 fields, 100 bytes each, plus key). To stress the I/O system, 125 simultaneous threads were used to attempt a target throughput of 5000 database operations/second.

Three different workloads were tested each with a different ratio of reads and/writes. The following workloads and command lines were used:

- Workload A – Load (100% Write) – Create 10 million records

```
python ./bin/ycsb load mongodb -P workloads/workloada_load -p
host=172.29.65.1 -p threads=125 -target 5000 -s >>
TestResults/24VMs_load.dat 2>&1 &
```

- Workload A – Run (50% Read, 50% Write) – Update 8 million records

```
python ./bin/ycsb run mongodb -P workloads/workloada_run -p
host=172.29.65.1 -p threads=125 -target 5000 -s >>
TestResults/24VMs_run50R.dat 2>&1 &
```

- Workload A- Run (80% Read, 20% Write) – Update 8 million records

```
python ./bin/ycsb run mongodb -P workloads/workloada_run80R -p
host=172.29.65.1 -p threads=125 -target 5000 -s >>
TestResults/24VMs_run80R.dat 2>&1 &
```

The YCSB benchmark was run on four Lenovo ThinkSystem SR630 servers, each with two Gold 6230 processors (20 cores @2.1GHz) and 384 GB of system memory and a ThinkSystem DM5000F unified All

Flash storage array with protection level RAID-DP. The VMs were evenly distributed across the four servers and for the YCSB benchmark, 96 VMs use 10.7TB.

The ThinkSystem DM5000F all flash storage array is with two controllers and configured with forty-eight 1.8 TB SSDs equally assigned to one aggregate per controller. The DM5000Fstorage is with FC connectivity to the four servers and the virtual machines on each host are load balanced across two controllers. Table 20 lists the performance test results for the YCSB scenarios.

Table 20: YCSB performance test results

Workload	**Throughput (operations/sec)**	**DM5000F All Flash**
		96 VMs
LOAD (100% Write)	Total	150195
	Mean per VM	1565
RUN (50% Read, 50% Write)	Total	39528
	Mean per VM	412
RUN (80% Read, 20% Write)	Total	52880
	Mean per VM	551

The following three diagrams show the throughput (operation/second) curves for the YCSB scenarios in Table 20. The performance and throughput pattern for all 3 scenarios are consistent throughout the test.

- For 100% writes, the performance is steady with 150000 operations/second
- For 50% read/50% write, the performance fluctuates from 25000-50000 operations/second
- For 80% read/20% write, the performance fluctuates from 45000-55000 operations/second

Figure 13 shows the throughput curve for the load workload with 100% writes

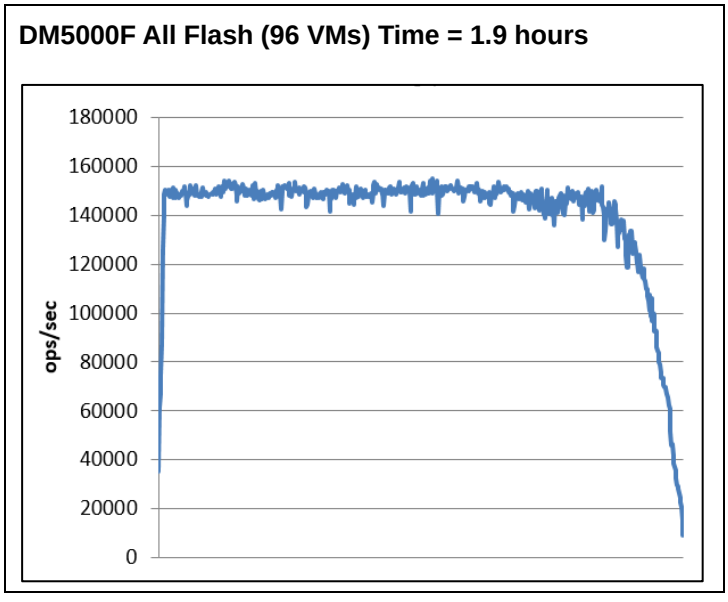

Figure 13: YCSB performance for load workload

Figure 14 shows the throughput curve for the run workload with 50% reads and 50% writes.

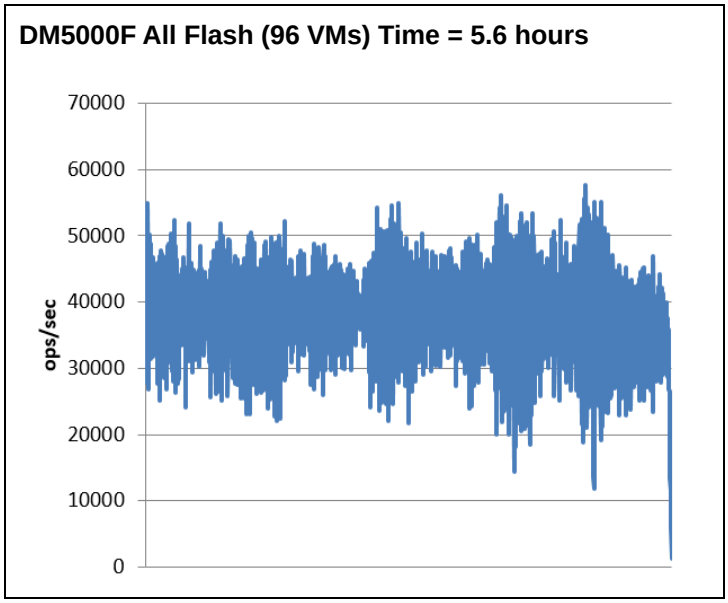

Figure 14: YCSB performance for run workload with 50% reads and 50% writes

Figure 15 shows the throughput curve for the run workload with 80% reads and 20% writes.

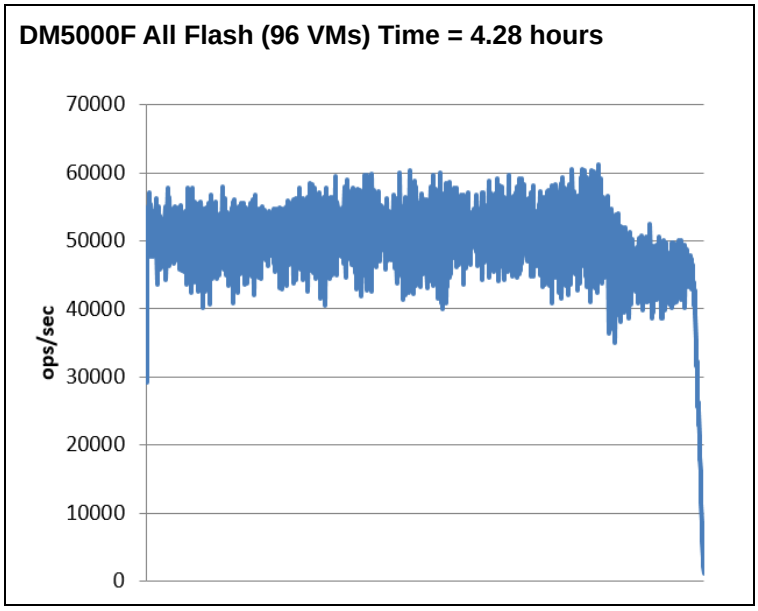

Figure 15: YCSB performance for run workload with 80% reads and 20% writes

6.5 Shared storage

This section describes the shared storage that is required by VMware vCloud. Storage for workloads is outside the scope of this document.

The total capacity that is required for the management cluster VMs is 9 TB, assuming two VMs for active-active or active-passive high availability. Adding 50% capacity for potential growth means that 13 TB is needed. However the actual storage required depends on the RAID level chosen. Much of this capacity is required by the SQL server VMs and it might be best to store these VMs in a separate volume. All of the VMs are considered to be persistent and must be backed up on a regular basis.

This reference architecture describes the following shared storage solution:

- Block I/O to Lenovo ThinkSystem DM5000F using Fibre Channel (FC)

For more information, see “Shared storage BOM” on page 52.

The storage configurations that are presented in this section include conservative assumptions about the VM size, changes to the VM, and user data sizes to ensure that the configurations can cope with the most demanding user scenarios.

6.5.1 ThinkSystem DM5000F Unified All Flash storage array

The recommended configuration is 12 SSDs of 3.84TB capacity with active-active dual controllers. Two RAID-DP aggregates of 6 drives each are created on each controller for a total capacity of 46 TB usable storage. The management VMs need to be load balanced across controllers. Large drives or additional expansion enclosures can be added to provide more storage capacity for backups or workload VMs.

6.6 Physical networking

This section describes the physical networking topology and includes design guidance to correctly configure the network environment for redundancy and failover.

6.6.1 Key physical networking concepts and terminology

This section describes the following basic networking concepts and terminology that are used throughout the next sections:

- Inter-Switch Link (ISL)

  An ISL is a physical network connection from a physical network port on one switch to a physical network port on another switch that enables communication between the two switches. This reference architecture uses two or four 10 GbE connections between the two networking switches, which are aggregated by using a trunk group.

- Link Aggregation Control Protocol (LACP)

  LACP is an IEEE 802.3ad standard for grouping several physical ports into one logical port (known as a dynamic trunk group) with any device that supports the standard. The 802.3ad standard allows standard Ethernet links to form a single Layer 2 link that uses LACP. Link aggregation is a method of grouping physical link segments of the same media type and speed in full duplex and treating them as if they were part of a single, logical link segment. If a link in a LACP trunk group fails, traffic is reassigned dynamically to the remaining links of the dynamic trunk group.

- Trunk Group

  A trunk group creates a virtual link between two switches that operates with aggregated throughput of the physical ports that are used. Most networking switches support two trunk types: static trunk groups and dynamic LACP trunk groups. Lenovo’s recommendation (and the method that is used in this reference architecture) is to use dynamic trunk groups when available. Figure 16 shows a dynamic trunk group aggregating two ports from each switch to form an ISL.

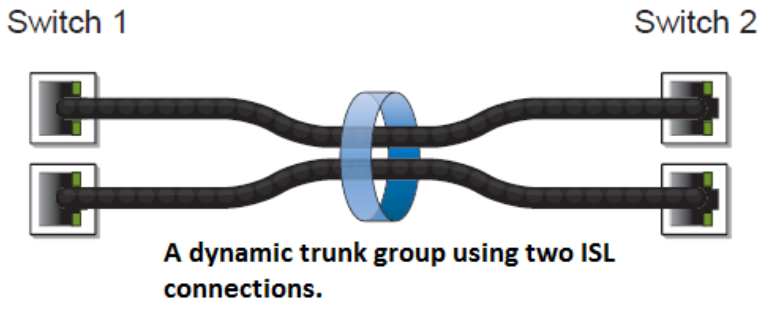

Figure 16: A dynamic trunk group aggregating two ISL connections between two switches

- Virtual Link Aggregation Group (vLAG)

  A switch or server in the access layer can be connected to more than one switch in the aggregation layer to provide for network redundancy. Typically, Spanning Tree Protocol (STP) is used to prevent broadcast loops, which blocks redundant uplink paths. Therefore, there is the unwanted consequence of reducing the available bandwidth between the layers by as much as 50%. In addition, STP can be slow to resolve topology changes that occur during a link failure and can result in considerable MAC address flooding.

  By using vLAGs, the redundant uplinks remain active and use all available bandwidth. To maintain maximum bandwidth over the multiple connections, vLAG is enabled on the LACP teams in this reference architecture.

- Virtual LAN (VLAN)

  VLANs are a way to logically segment networks to increase network flexibility without changing the physical network topology. With network segmentation, each switch port connects to a segment that is a single broadcast domain. When a switch port is configured to be a member of a VLAN, it is added to a group of ports that belong to one broadcast domain. Each VLAN is identified by a VLAN identifier (VID) in the range of 1 - 4094 (values of 0 and 4095 have a special meaning).

- Tagged Port

  A tagged port is a port on a physical switch that is configured as a member of one or more specific VLANs by using IEEE 802.1Q standard. Frames that have a valid VLAN tag matching this port are accepted and frames without the correct VLAN tag are ignored. When an untagged frame exits the switch through a tagged member port, the frame header is modified to include the 32-bit tag that is associated with the port VLAN ID (PVID). When a tagged frame exits the switch through a tagged member port, the frame header remains unchanged (original VID remains).

- Untagged Port

  An untagged port is also known as an access port on a physical switch. The access port is configured as a member of the native VLAN. All Ethernet frames that ingress on this port are retagged with the native VLAN to identify the frame VLAN. An Ethernet frame that egresses through an access port has its native VLAN tag stripped before sending to the next device.

6.6.2 Physical network topology

This reference architecture uses two ultra-low latency, high-performance Lenovo RackSwitch G8272 10 GbE network switches to provide primary data communication services. Lenovo recommends the use of Emulex Virtual Fabric Adapter 5.2 (VFA5.2) network adapters for servers. These adapters also include built-in support for FCoE or iSCSI networking via a feature on demand (FoD).

Lenovo recommends that clusters that use shared storage should have two dual port network adapters to provide both connectivity and failover.

Figure 17 shows an example of the physical networking topology with the shared edge and compute cluster, management cluster, and an additional compute cluster. For clarity a representative server from each cluster is shown; that is, only 12 connections are shown to the two network switches rather than the 48 that are actually needed.

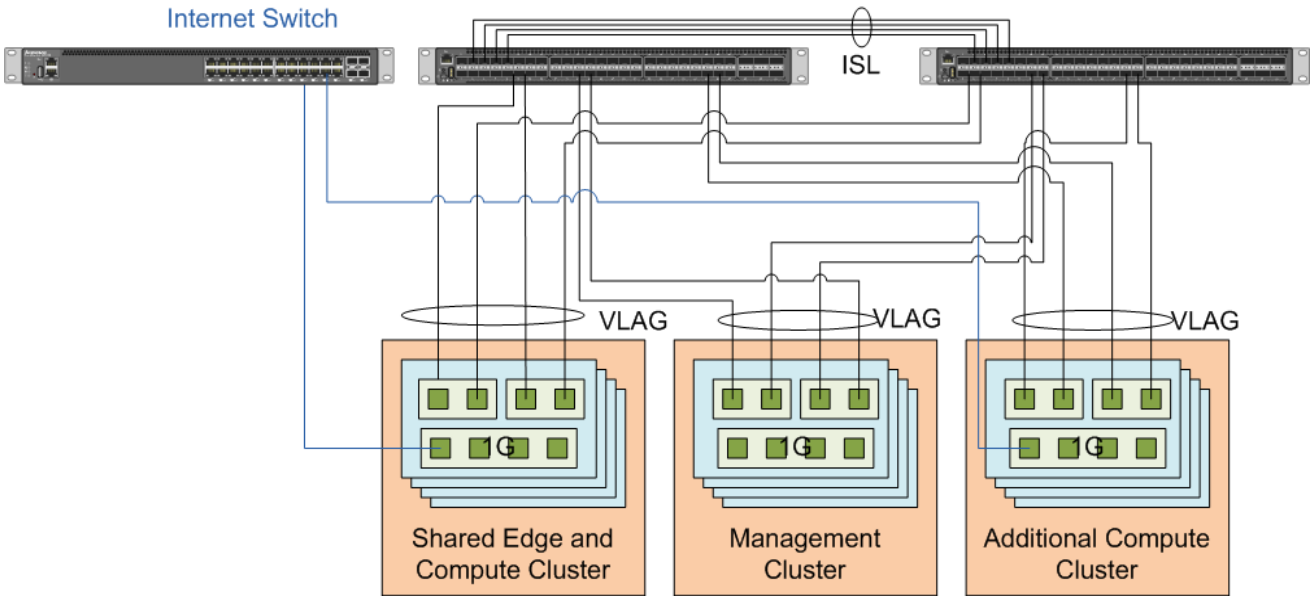

Figure 17: Physical Networking Topology

The two G8272 network switches have four physical 10 GbE connections that use ISL in a single trunk group. There are also 40 GbE ports on this switch that can be connected with a single 40 GbE cable. Some ports may be used for upstream connectivity. For all clusters, the network ports are grouped into a single trunk group configured to use VLAG.

Also shown are four 1G ports in each server. The servers in the shared edge and compute cluster and the additional compute cluster are connected to a 1G switch using one of these four 1G ports. This switch in turn is connected to the internet via a gateway and firewall (not shown).

Figure 17 does not show the out-of-band management network, which is used to manage the servers, switches, and other hardware that uses XClarity.

6.7 Virtual networking

This section describes the virtual networking topology used in a vSphere based virtualized environment that is layered on top of the physical networking topology.

6.7.1 Key virtual networking concepts and terminology

The following terminology is used in this section:

- vSphere Standard Switch

  The vSphere Standard Switch (VSS) is a software construct in ESXi that allows networking of VMs. A VSS can be connected to one or more physical Network Interface Cards (NICs) or none at all.

- Port Group

  Each VM interfaces with the VSS via a port. A VSS can consist of one or more port groups, which describe how the virtual switch should route traffic between the virtual network and the VMs that are connected to the specified ports. Administrators can use port groups to configure traffic shaping and bandwidth limitations, NIC failover, and other settings.

- vSphere Distributed Switch

  A vSphere Distributed Switch (VDS) builds on the idea of a vSwitch and is an abstract representation of multiple servers defining the same name, network policy, and port group to allow connectivity between the servers. A VDS consists of a control plane and a data plane. The control plane is at vCenter server and is responsible for managing and configuring the VDS, It also is used as a configuration template to configure the settings on each host that are connected to the VDS. The data plane (or I/O plane) is VSS that is on each host and is responsible for forwarding the data packets to the relevant uplinks. Because of this data plane, network communication continues to work even if vCenter is down.

- Distributed virtual port group

  A distributed virtual port group (DV port group) is the VDS equivalent of a port group.

- dvUplink ports

  dvUplink ports connect the VDS to physical NICs on associated hosts.

- dvUpLink port group

  The dvUpLink port group contains the dvUplink ports available in the host.

- Underlay Network

  The underlay network is another term for the physical and virtual networking infrastructure.

- Overlay Network (also known as software defined network)

  An overlay network is effectively a logical representation of a physical network. For example, Routers, Load Balancers, Firewalls, are all virtual appliances that can be configured and used in the overlay network without requiring physical devices.

6.7.2 VLANs

VLANs are used to logically separate different kinds of network traffic. For vCloud and vRealize, the following standard VLANs should be used:

- Management — Used for all management traffic for vCenter server, ESXi, and all of the vRealize management appliances (for example, Automation, Operations, and Business Suite)
- vSphere vMotion — Used to move VMs from one server to another.
- Fault Tolerance — Used to support the fault tolerance (FT) feature of vSphere.
- Storage — Used for iSCSI, NFS, or FCoE (VLAN 1002) storage traffic.

In addition, each workload application might require one or more VLANs for its logical networks, such as compute, data, and management.

It is recommended to use different VLANs in each of the clusters (shared edge and compute, management, and compute). These different VLANs can be connected together by leveraging Layer 3 routing either in the G8272 physical switch or in upstream switches.

In larger networks with many applications, it is easy to run out of unique VLANs. In this case, VXLANs and VMware NSX can be used. NSX is outside the scope of this document.

6.7.3 Virtual network topology

The details of the virtual network topology depend on the customer environment. The example that is presented in this section assumes the following VLANs are used:

- management
- vMotion
- application 1
- application 2

Figure 18 shows the VDS configuration in vCenter server. It is recommended that each DV port groups has a different VLAN. Each of the VMs is configured with one or more virtual network ports (vnic) that connects to a DV port group in the DVS. Standard kernel functions like vMotion should be configured so their vmknic connects to a DV port group.

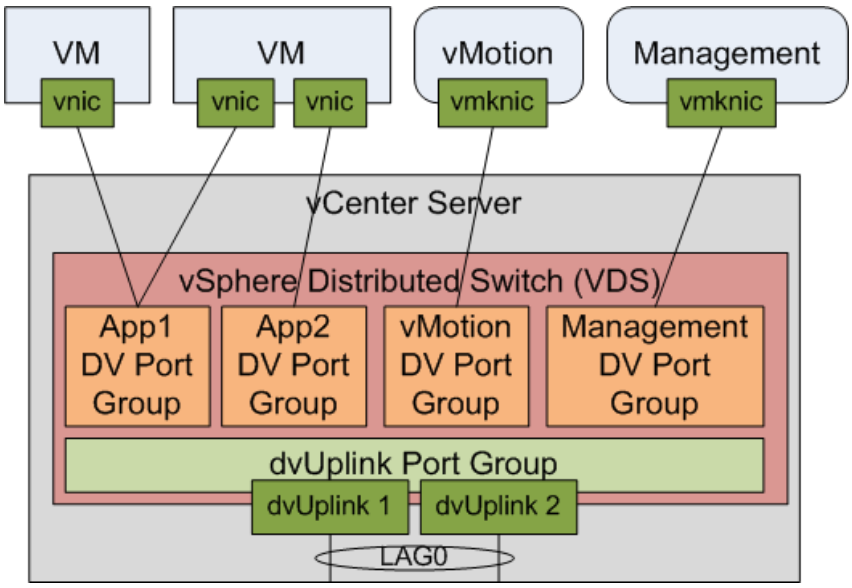

Figure 18: Example template of a virtual networking topology

Figure 19 shows the VDS template implemented on the ESXi hosts for two servers. The DV port groups are connected to dvUplinks which in turn are connected to the two physical NICs (vmnic) and the physical switches. The VLAN traffic flows through this infrastructure into the physical switches.

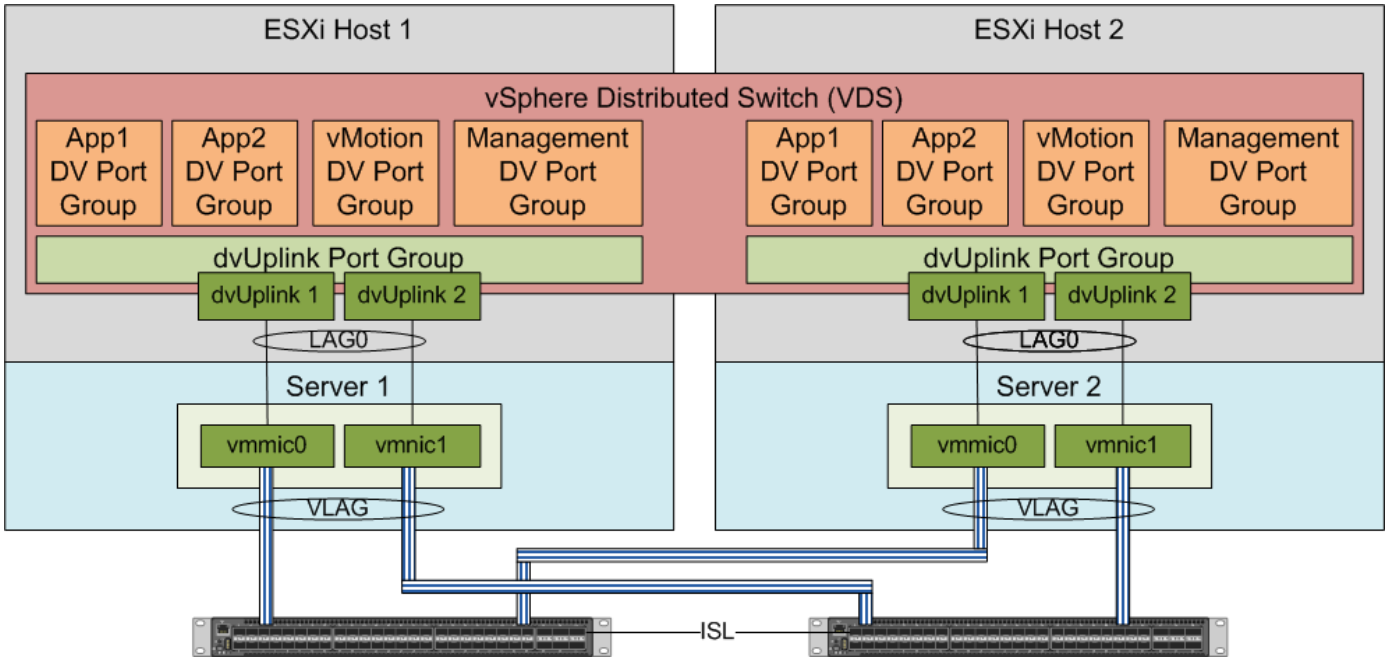

Figure 19: Deployment of the example virtual networking topology

6.8 Hybrid networking to public clouds

This section contains deployment considerations for hybrid networking from an on-premise cloud to public clouds such as AWS.

6.8.1 AWS networking

Table 21 shows the required connectivity for each component used for AWS.

Table 21: AWS component connectivity

Virtual Machine	DvSwitch	http proxy support
AWS SMS Connector for vCenter	Edge-Compute	Yes

6.8.2 Best Practices

Amazon Web Services (AWS) provide internet and direct connect accessibility options to connect from the vSphere private cloud environments. The AWS connectors leverage an http proxy to establish connectivity to the AWS cloud.

The server hardware for an on-premise vSphere cloud may not use the same processor family. In this case Lenovo recommends enabling Enhanced vMotion Compatibility (EVC) on the cluster to avoid compatibility errors during vMotion.

AWS Connectors are deployed in an on-premise vSphere environment and they do not provide capabilities for bi-directional migration. The on-premise virtual machines can be migrated to AWS cloud but there is no option in the connectors to migrate a virtual machine back to vSphere environment. Instead the VMs need to be exported to OVA format and stored in S3. Then the image can be imported or downloaded to the on-premise vSphere environment.

6.9 Systems management

Lenovo XClarity™ Administrator is a centralized resource management solution that reduces complexity, speeds up response, and enhances the availability of Lenovo® server systems and solutions.

The Lenovo XClarity Administrator provides agent-free hardware management for Lenovo's ThinkSystem® rack servers, System x® rack servers, and Flex System™ compute nodes and components, including the Chassis Management Module (CMM) and Flex System I/O modules. Figure 20 shows the Lenovo XClarity administrator interface, in which Flex System components and rack servers are managed and are seen on the dashboard. Lenovo XClarity Administrator is a virtual appliance that is quickly imported into a virtualized environment server configuration.

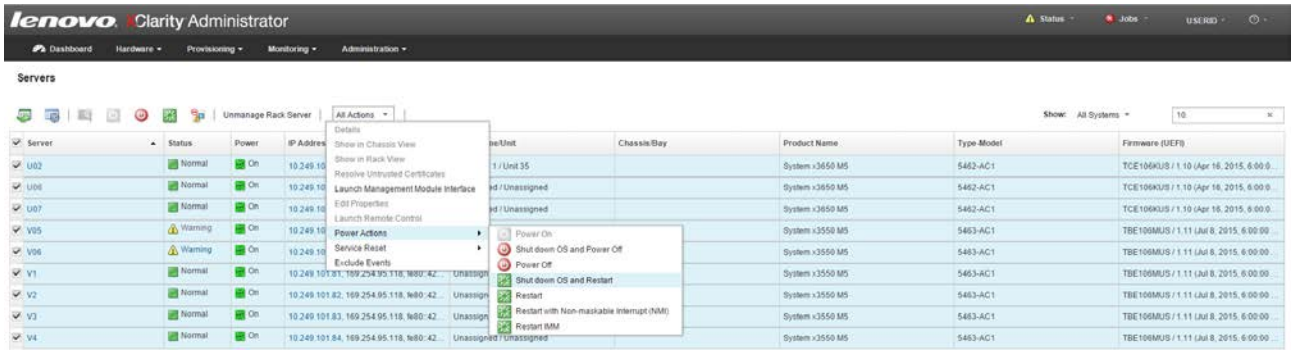

Figure 20: XClarity Administrator interface

6.9.1 Lenovo XClarity integration

Lenovo also provides XClarity integration modules for VMware vCenter, VMware vRealize Orchestrator and VMware vRealize Log Insight. For more information, see this website: http://www3.lenovo.com/us/en/data-center/software/systems-management/c/systems-management.

By using the Lenovo XClarity Integrator for VMware vCenter, administrators can consolidate physical resource management in VMware vCenter, which reduces the time that is required for routine system administration.

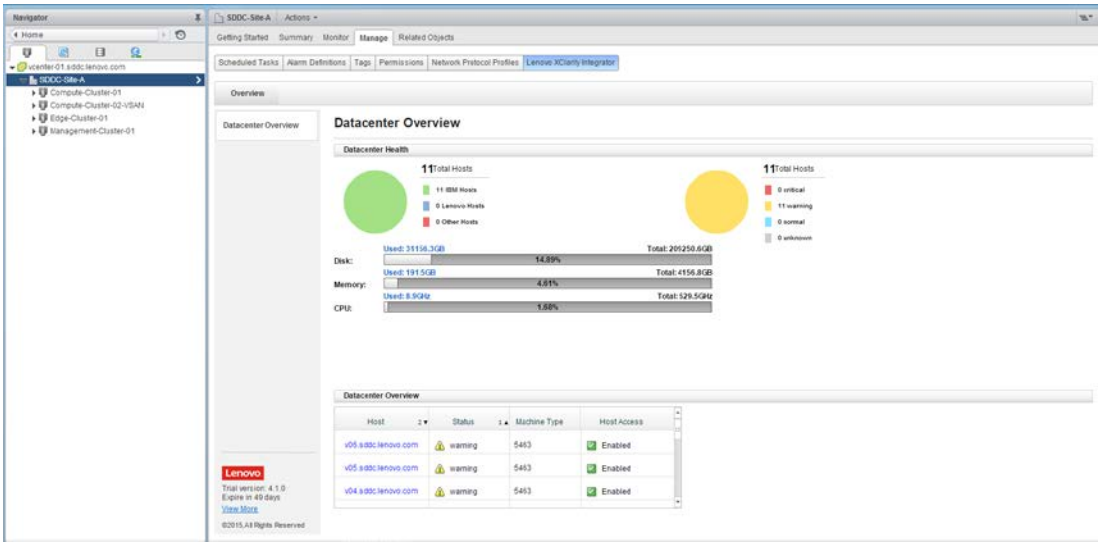

Figure 21: Lenovo XClarity Integrator for VMware vCenter

The Lenovo XClarity Integrator for VMware vCenter provides the following features and benefits:

- Extends Lenovo XClarity Administrator features to the virtualization management console
- Enables management of legacy infrastructure from the virtualization management console
- Reduces workload downtime by dynamically triggering workload migration in clustered environments during rolling server reboots or firmware updates, and predicted hardware failures

The Lenovo XClarity Integrator for VMware vRealize Orchestrator provides IT administrators with the ability to coordinate physical server provisioning features of Lenovo XClarity Pro with broader vRealize Orchestrator workflows. Lenovo XClarity Integrator for VMware vRealize Orchestrator provides a library of simple yet robust and customizable workflow routines and actions designed to automate complex, repetitive IT infrastructure tasks such as system discovery and configuration, hypervisor installation, and addition of new hosts to vCenter.

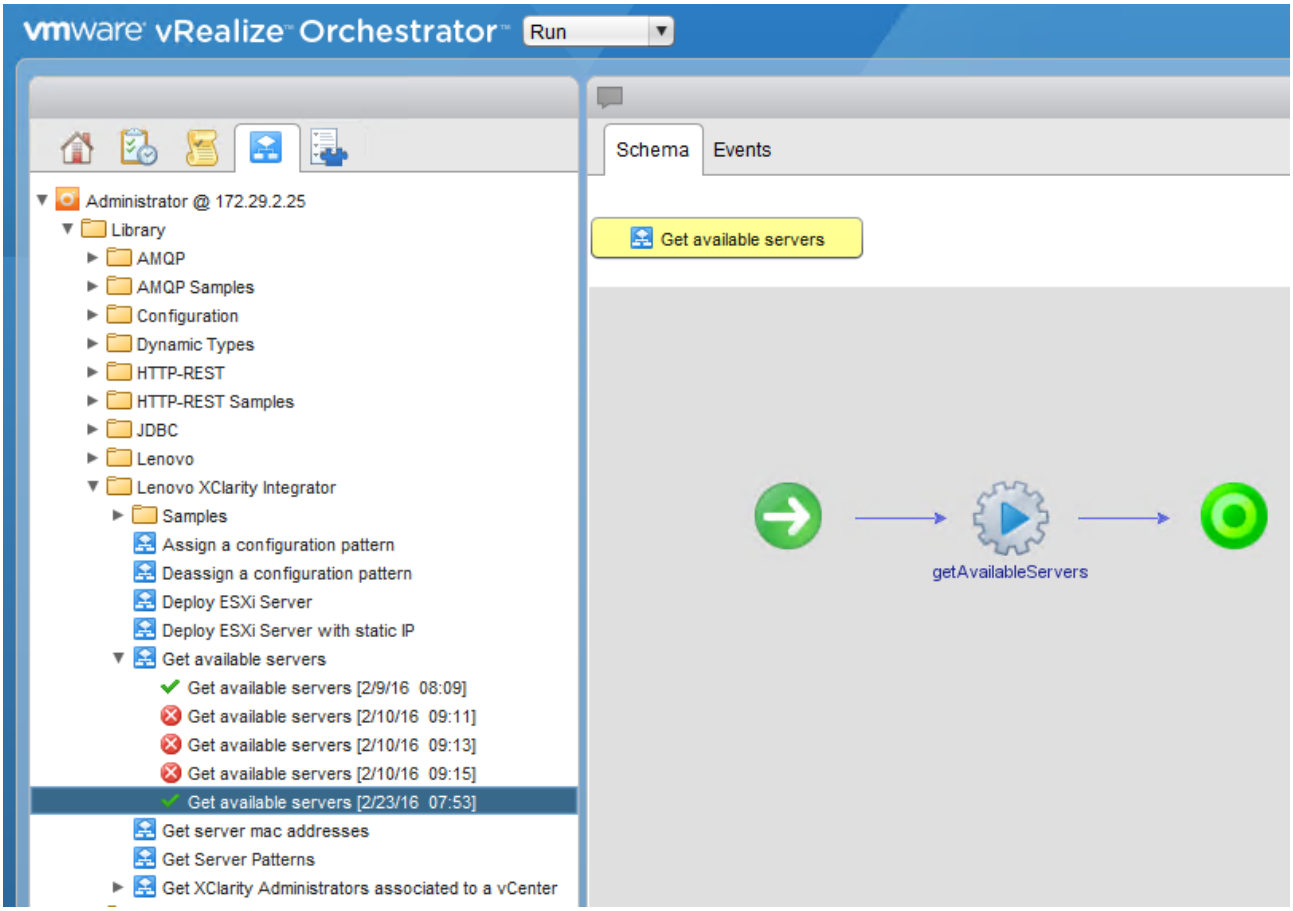

Figure 22: Lenovo XClarity Integrator for VMware vRealize Orchestrator interface

The Lenovo XClarity Administrator Content Pack for VMware vRealize Log Insight simplifies the collection and forwarding of Lenovo XClarity Administrator logs to VMware vRealize Log Insight for powerful processing and analytics, and displaying insightful information in an intuitive format.

The VMs for VMware vCenter, vRealize Orchestrator, Lenovo XClarity Administrator and Lenovo XClarity Administrator Integrator should have access to the management network used for managing servers, storage and networking.

Lenovo XClarity Integrator for vRealize Automation provides a set of blueprints to provision infrastructure services based on Lenovo servers, network switches and vSphere. This eases provisioning a new Lenovo server with vSphere installed, network isolation parameters configured on the Lenovo switches, apply vSphere distributed switch configurations and adding the server to the existing or new vSphere Cluster. These services leverage the workflows defined in the Lenovo vRealize SoftBundle for vRealize Orchestrator, Lenovo XClarity Integrator for vCenter, Lenovo XClarity Integrator for vRealize Orchestrator, and Lenovo Networking Integration plugin for vRealize Orchestrator.

The Lenovo vRealize SoftBundle package for vRealize Automation needs to imported into vRealize Orchestrator and then the Blueprints package is imported using the vRealize Cloud Client command line utility by Tenant Administrators and it creates catalog items automatically. The catalog items are created under Lenovo Servers, Lenovo Network, and Lenovo Virtualization services.

Figure 23 shows the workflows available in the Lenovo SoftBundle for vRealize Orchestrator.

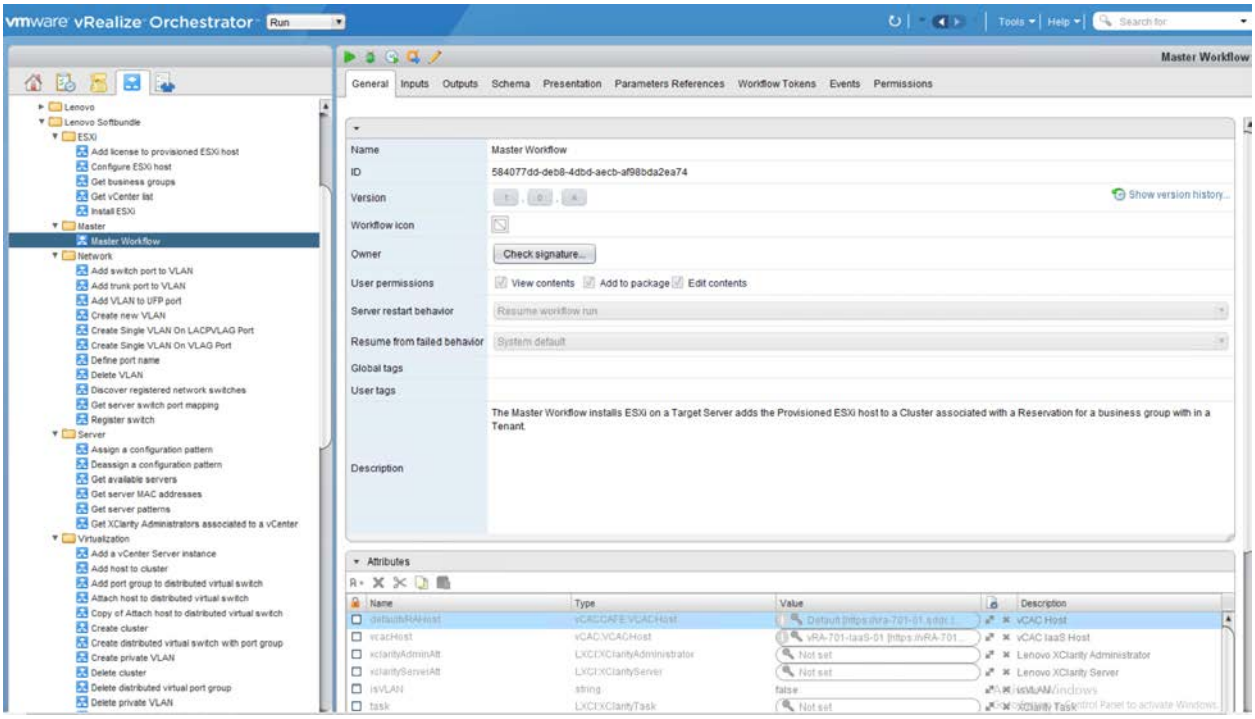

Figure 23: Lenovo SoftBundle workflows for vRealize Orchestrator

Figure 24 shows Lenovo XClarity Integrator catalog items for vRealize Automation.

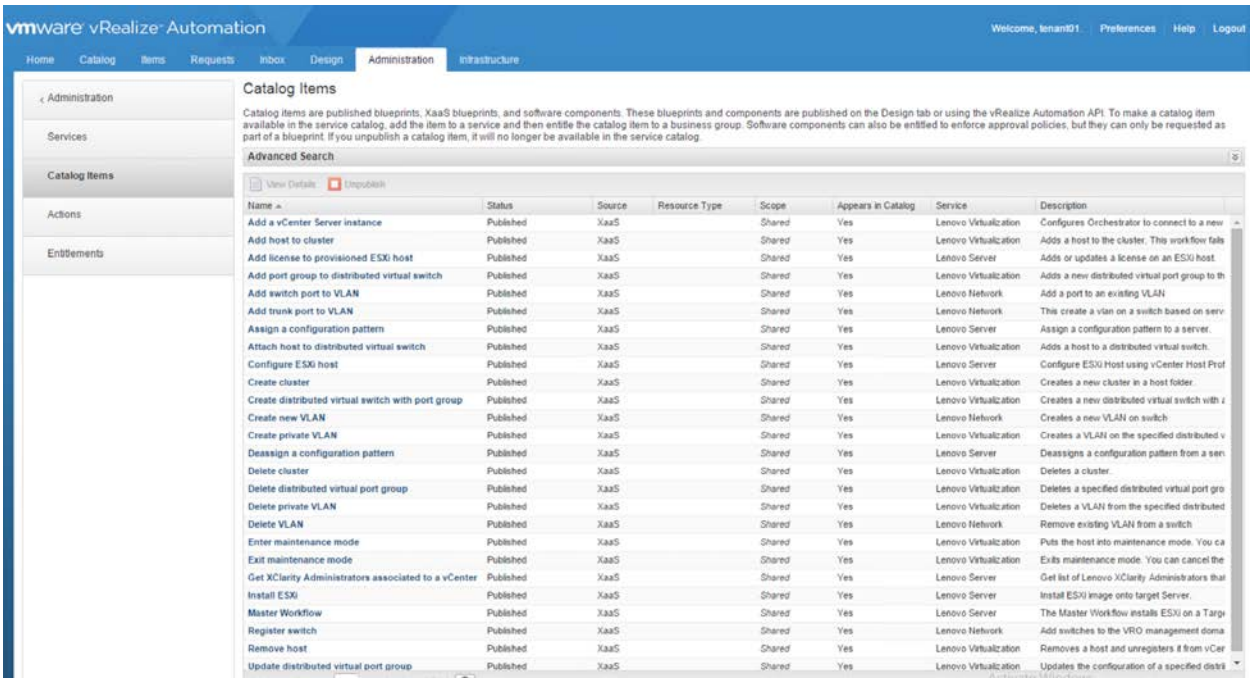

Figure 24: Lenovo XClarity Integrator for vRealize Automation Catalog Items

Figure 25 shows Lenovo XClarity Integrator services available for vRealize Automation.

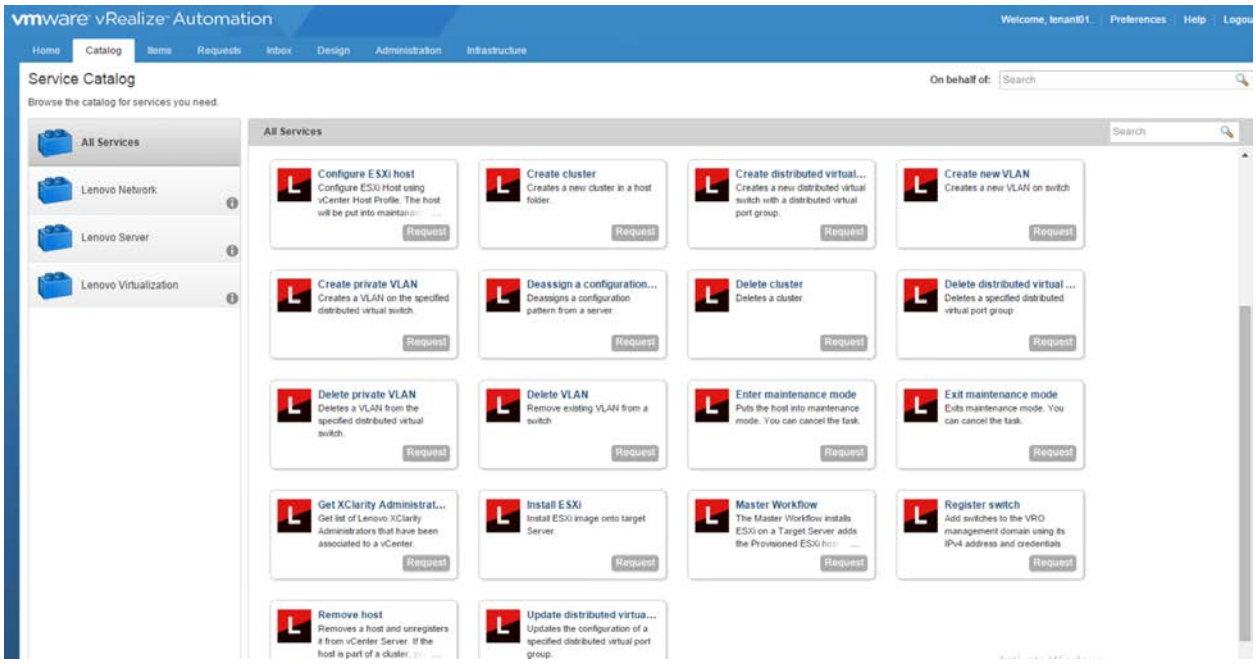

Figure 25: Lenovo XClarity Integrator for vRealize Automation Services Catalog

6.9.2 Lenovo network integration plug-ins

Lenovo also provides network integration plug-ins for VMware vRealize Orchestrator and vRealize Log Insight. For more information, see this website:

shop.lenovo.com/us/en/systems/software/systems-management/network-management

The Lenovo networking plug-In for VMware vRealize Orchestrator enables you to:

- Reduce new service delivery time on RackSwitch G8272, Flex System EN4093R, and Flex Systems Interconnect Fabric
- Leverage a comprehensive library of Lenovo Networking workflows and actions capturing network configuration best practices for rapid service orchestration
- Implement more Lenovo switch configuration activity through vRealize Orchestrator and vCenter with less reliance upon native switch interfaces.

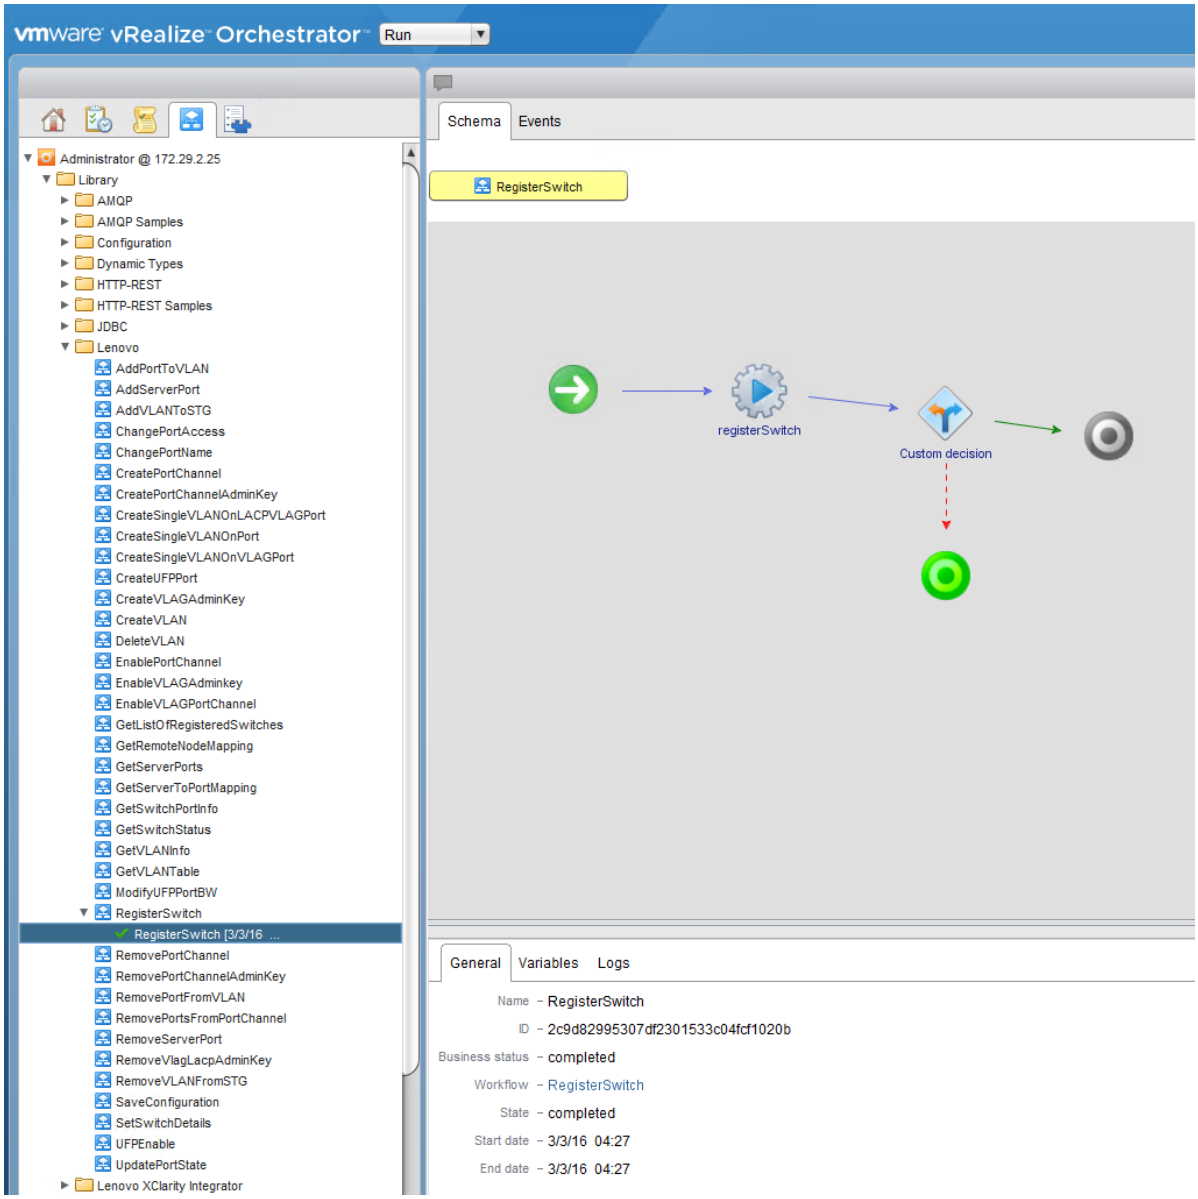

Figure 26: Lenovo networking plug-In for VMware vRealize Orchestrator interface

Lenovo Networking Content Pack for VMware vRealize Log Insight enables you to:

- Increase network reliability by allowing system or network administrators to monitor networks that feature Lenovo branded RackSwitch switches
- Gain access to extremely detailed switch log entries to facilitate deep insights into the network status
- Reduce initial provisioning time of Log Insight by using the 9 prebuilt dashboards and 12 predefined alarms featured in the pack

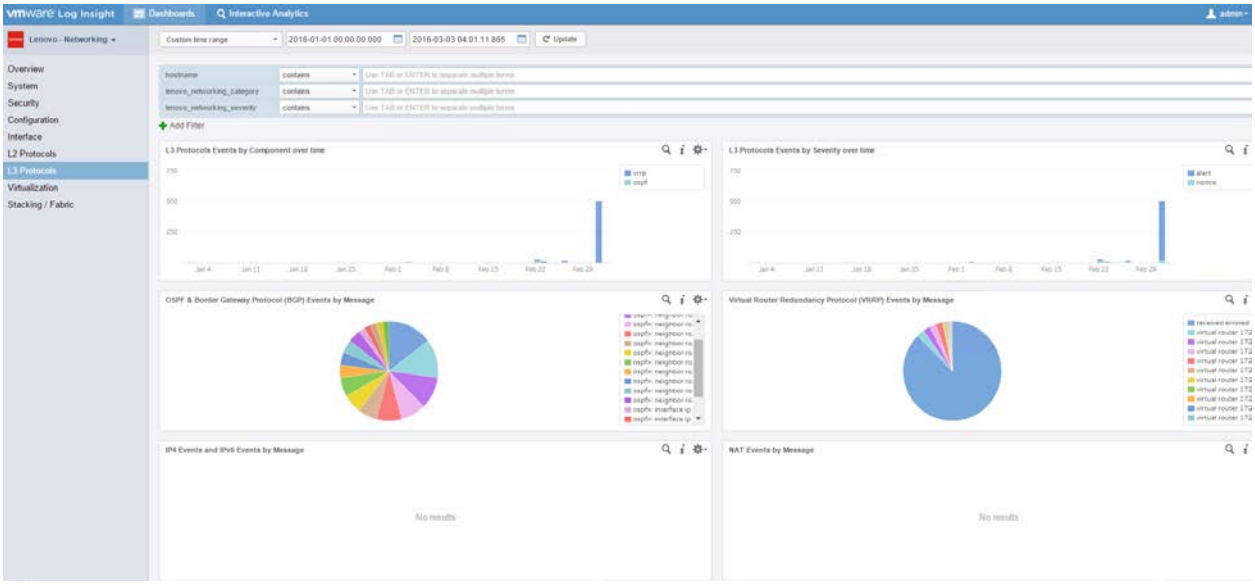

Figure 27: Lenovo Networking Content Pack for VMware vRealize Log Insight interface

6.9.3 Lenovo plug-ins compatibility

Table 22 below lists current versions of Lenovo integration plugins and the required or supported VMware vCenter and vRealize Suite products.

Table 22: Plug-in compatibility

Component Name	Version	Supported Product Versions
Lenovo XClarity Administrator(LXCA)	2.3.0	VMware vCenter 6.0U2/6.5/6.7, ESXi 6.0U2/6.5 U1/6.7u1
Lenovo XClarity Integrator(LXCI) for vCenter	6.0.0	Lenovo XClarity Administrator 1.4.x, 2.x VMware vCenter 5.x U1/U2/U3, 6.0 U1/U2/U3, 6.5 U1/U2,6.7(U1,U2)
Lenovo XClarity Administrator content pack for VMWare vRealize Log Insight	2.0	Lenovo XClarity Administrator 1.1 or higher VMware vRealize Log Insight 3.0 or higher
Lenovo XClarity Integrator for VMware vRealize Automation	1.0.1	VMware vRealize Automation 7.x
Lenovo XClarity Integrator for VMware vRealize Orchestrator	1.1.0	VMware vRealize Automation 7.0 VMware vRealize Orchestrator 6.0/7.0
Lenovo Network Plugin for VMware vRealize Orchestrator	1.4.0	VMware vRealize Orchestrator 7.4.x
Lenovo Networking Content Pack for VMware vRealize Log Insight	1.2	CNOS 10.3, ENOS 8.1, VMware vRealize Log Insight 4.3.x, 4.5.x

7 Deployment example

This section describes an example deployment of vRealize Suite 7.5. Four physical servers are used for each of the shared edge and compute, management, and additional compute clusters. Lenovo ThinkSystem SR630 servers are used for the shared edge and compute cluster and management cluster. The servers for the additional compute cluster are workload- dependent and Lenovo ThinkSystem SR650 servers are used in this case. Lenovo RackSwitch G8272 switches are used for the 10 GbE network. Lenovo ThinkSystem DM5000F storage array with SSDs is used for all of the VMs and is connected to the servers using Lenovo ThinkSystem DB610S SAN switch.

7.1 Hardware views

Figure 28 shows a view of the rack with the nine servers, shared storage, and switches.

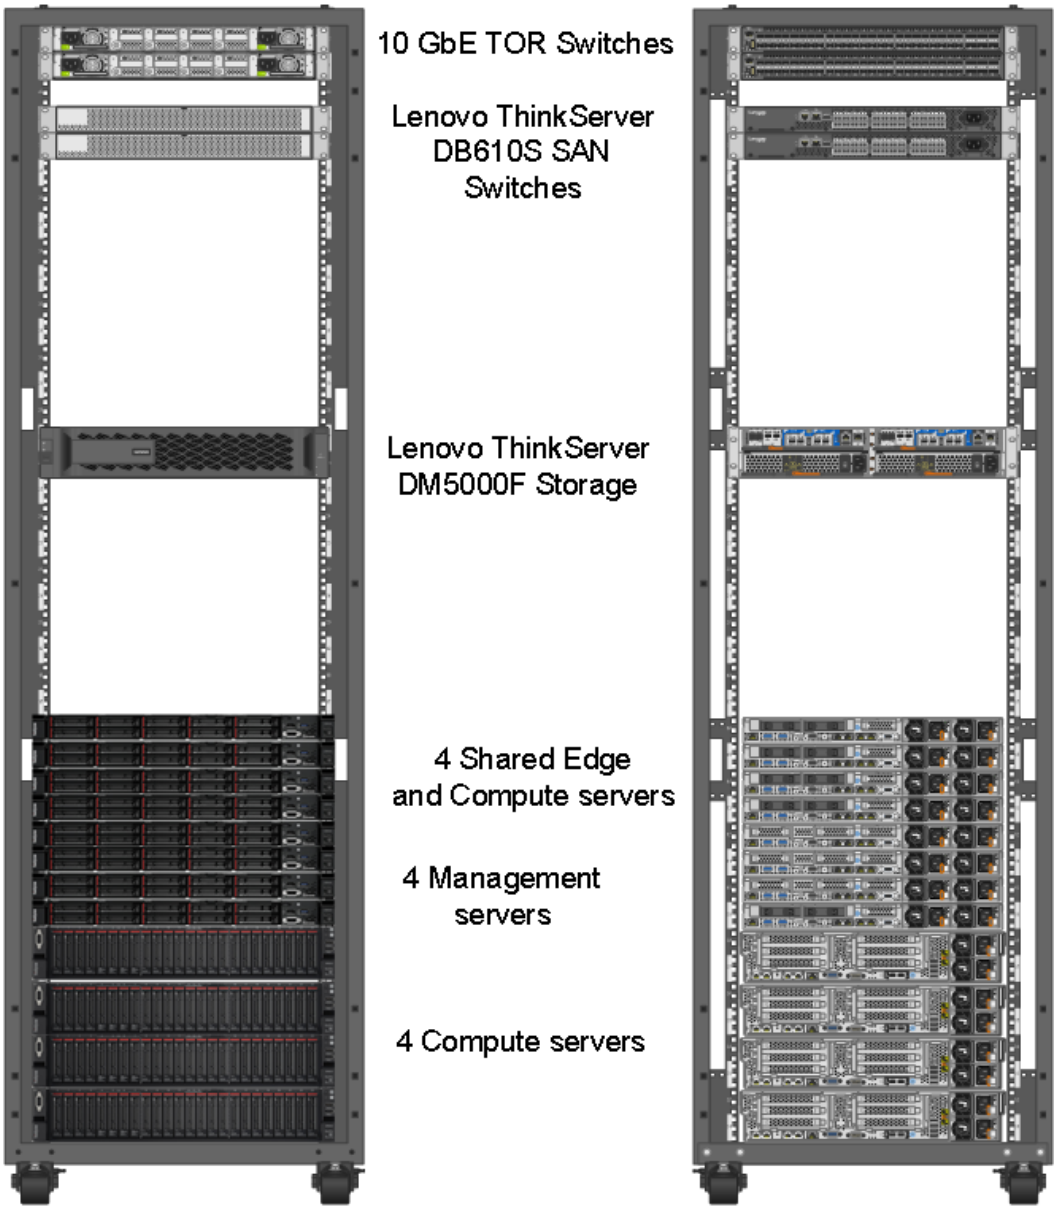

Figure 28: Rack Layout

Figure 29 shows a view of the physical 10 GbE network and connections to the external internet.

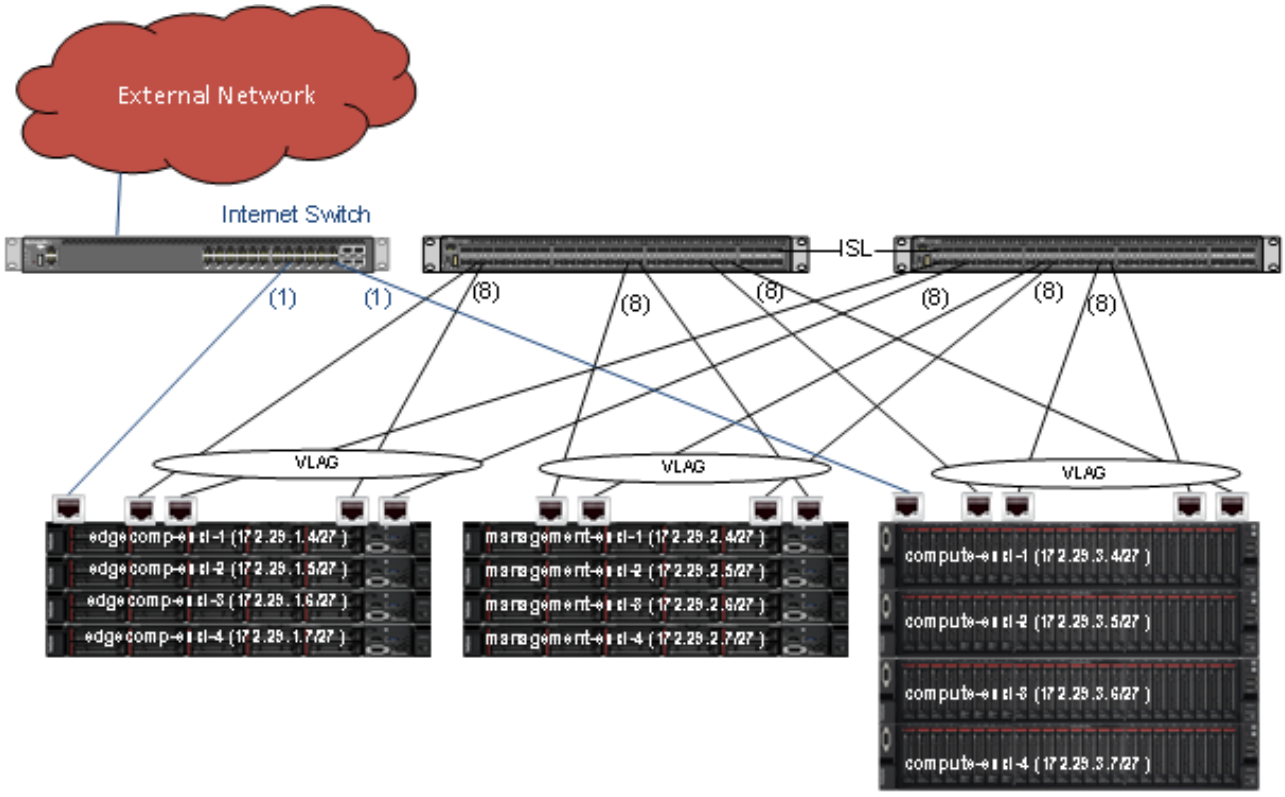

Figure 29: Networking Overview

For the shared edge and compute, management and additional compute clusters, the nodes use VLAG technology and as such are using a LAG configuration within the vSphere Distributed Switches. It is recommended to use VLAG for all the clusters connected to the same set of switches.

The servers in the shared edge and compute cluster and the additional compute cluster are connected to a 1G switch. This switch in turn is connected to the internet via a gateway and firewall (not shown).

7.2 IP/VLAN mapping

This example deployment uses the following five VLANs:

- Management
- vMotion
- FT
- vRA
- Comp

Table 23 lists example IP address ranges for the VLANs in each cluster where RID means Rack ID.

Table 23: Network Segments

Cluster	Shared Edge and Compute (RID 1)		Management (RID 2)		Compute (RID 3)	
	Subnet	VLAN	Subnet	VLAN	Subnet	VLAN
Manage	172.29.1.0/27	101	172.29.2.0/27	201	172.29.3.0/27	301
vMotion	172.29.1.32/27	102	172.29.2.32/27	202	172.29.3.32/27	302
FT	172.29.1.64/27	103	172.29.2.64/27	203	172.29.3.64/27	303
vRA	N/A	107	172.29.2.192/27	207	N/A	307
Comp	172.29.2.192/27	109	N/A	209	172.29.2.192/27	309

In this example, each cluster needs a minimum of five network segments within the 172.29.RID.x address range. Each segment does not require more than 30 IP addresses; therefore, a 255.255.255.224 (/27) netmask provides enough addresses. The same VLAN IDs can be used across racks with different IP segments. In this example, that option is not available because the switches and routers are shared across the three clusters.

7.2.1 Layer 3 Routing with VRRP

Virtual Routing Redundancy Protocol (VRRP) should be enabled in the switches for layer 3 routing. Each switch defaults to IP routing. VMs can use the respective routing IP to reach the switches. Routing occurs through either one of the switches and this causes intermittent routing failures when used with VLAG.

Layer 3 routing with VRRP removes this limitation by using a single routing IP address for both the switches. Each subnet requires 2 IPs reserved for normal layer 3 routers and one IP reserved for VRRP router.

Table 24 lists the example layer 3 routing IPs for each of the VLANs and the clusters.

Table 24: Layer 3 Example Routing

Cluster	VLAN	Subnet	L3 Interface IP		VRRP Router IP	
			G8272(1)	G8272(2)	G8272(1) Master	G8272(2) Backup
Shared Edge and Compute	Management	172.29.1.0/27	172.29.1.1	172.29.1.2	172.29.1.3	172.29.1.3
	vMotion	172.29.1.32/27	172.29.1.33	172.29.1.34	172.29.1.35	172.291.35
	FT	172.29.1.64/27	172.29.1.65	172.29.1.66	172.29.1.67	172.29.1.67
	Comp	172.29.1.192/27	172.29.1.193	172.29.1.194	172.29.1.195	172.29.1.195
Management	Management	172.29.2.0/27	172.29.2.1	172.29.2.2	172.29.2.3	172.29.2.3
	vMotion	172.29.2.32/27	172.29.2.33	172.29.2.34	172.29.2.35	172.29.2.35
	FT	172.29.2.64/27	172.29.2.65	172.29.2.66	172.29.2.67	172.29.2.67
	vRA	172.29.2.192/27	172.29.2.193	172.29.2.194	172.29.2.195	172.29.2.195
Additional Compute	Management	172.29.3.0/27	172.29.3.1	172.29.3.2	172.29.3.3	172.29.3.3
	vMotion	172.29.3.32/27	172.29.3.33	172.29.3.34	172.29.3.35	172.29.3.35
	FT	172.29.3.64/27	172.29.3.65	172.29.3.66	172.29.3.67	172.29.3.67
	Comp	172.29.3.192/27	172.29.3.193	172.29.3.194	172.29.3.195	172.29.3.195

7.3 Detailed cluster view

This section describes the detailed overlay virtualized networking for this example that is used for the clusters. The term vRA is used as a short-hand to denote all of the vRealize Automation VMs.

The Lenovo ESXi image is used for this example deployment. Lack of space in the ESXi boot bank may be a problem and it may be necessary to uninstall VIBs and reboot the host to make room for the VIBs that are required for NSX.

To use ESXi 6.7, download the Lenovo ThinkSystem custom image from the following website: my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/6_7#custom_iso.

To use ESXi 6.5 U1 or U2, download the Lenovo ThinkSystem custom image from the following website: my.vmware.com/web/vmware/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/6_5#custom_iso.

7.3.1 Management cluster

Although a single distributed switch can be deployed across various clusters, it is recommended to deploy unique distributed switches per cluster. Figure 30 shows the distributed switch for the management cluster.

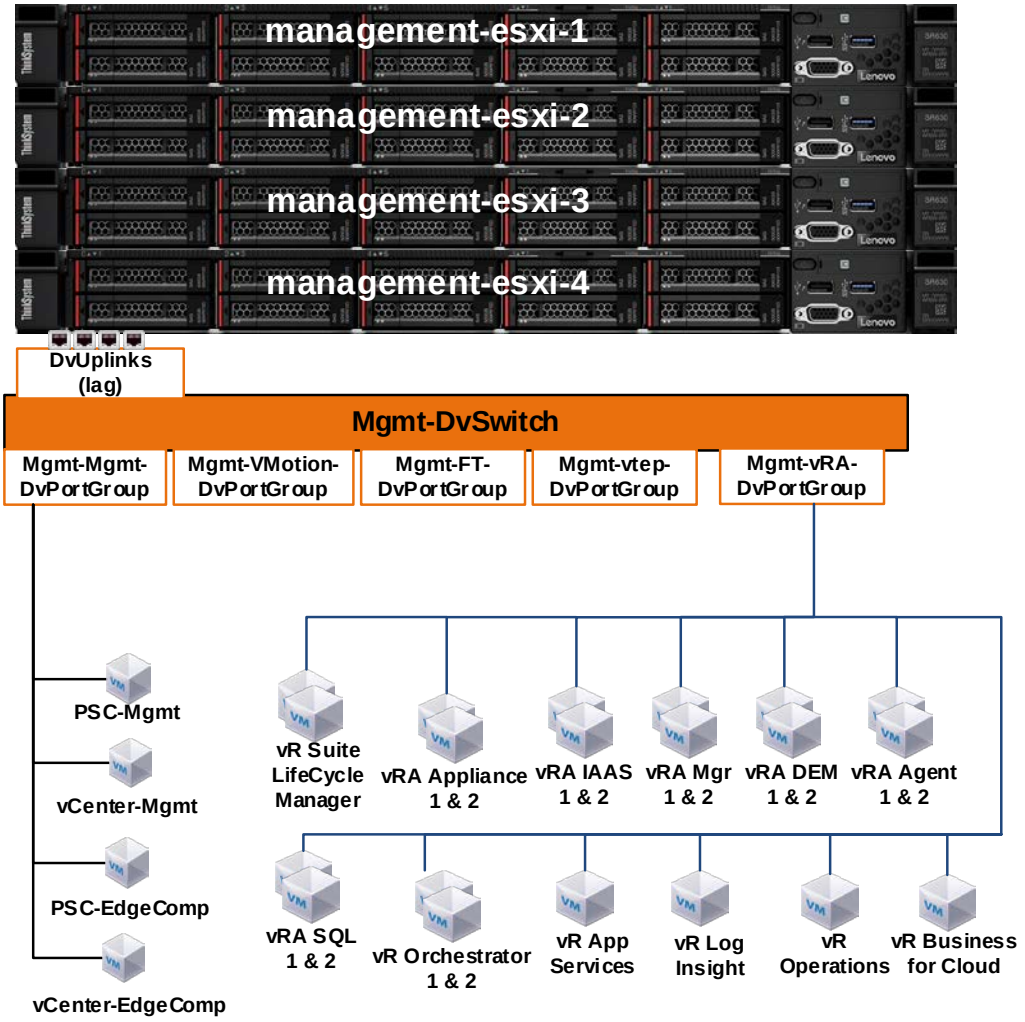

Figure 30: Management Cluster VDS

The infrastructure uses single SSO domain and all platform service controllers and vCenter servers are connected to this domain. All VMs in the management cluster can be configured on the same network segment. Separate vCenter and platform service controller instances are used for management cluster and shared edge and compute cluster. It is recommended that vCenter is deployed in a highly available configuration (depending on each customer's needs) which can result in multiple vCenters, multiple PSCs, and physical load balancers.

The vRealize VMs are deployed on a separate DvPortGroup (Mgmt-vRA) and subnet.

7.3.2 Shared Edge and Compute cluster

Figure 31 shows the distributed switches for the edge and compute cluster. One DVS is used for accessing the Internet via a 1GbE port on the server. The other DVS is used for all of the edge and compute network flows.

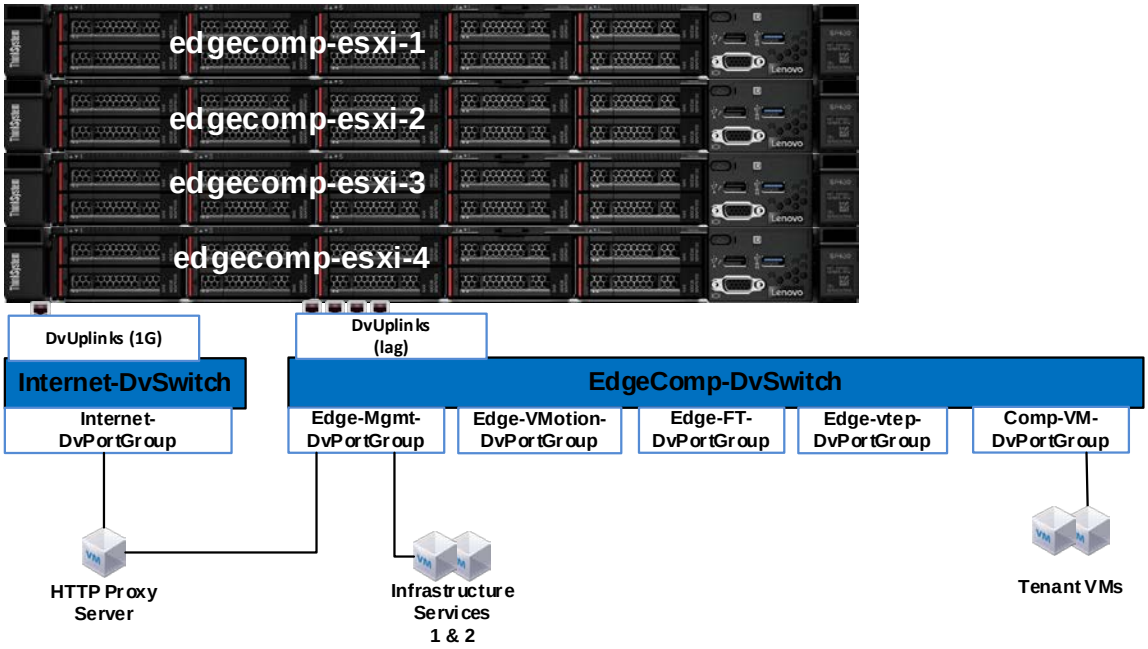

Figure 31: Shared Edge and Compute Cluster VDS

Infrastructure service VMs, such as Active Directory, DHCP, DNS, and NTP might exist in the customers' environment and these services are accessed through a clustered configuration, as shown in this example. However, if there is a requirement to virtualize these services within this environment, then they can be accessed through Edge-Mgmt-DvPortGroup, as shown in Figure 31.

The Comp-VM-DvPortGroup is used for the workload VMs.

7.3.3 Compute cluster

Figure 32 shows the distributed switch for the compute cluster. Workload-specific DvPortGroups and VMs are not shown because they are highly dependent on the specific workloads. More than one compute cluster can be used to add more workloads.

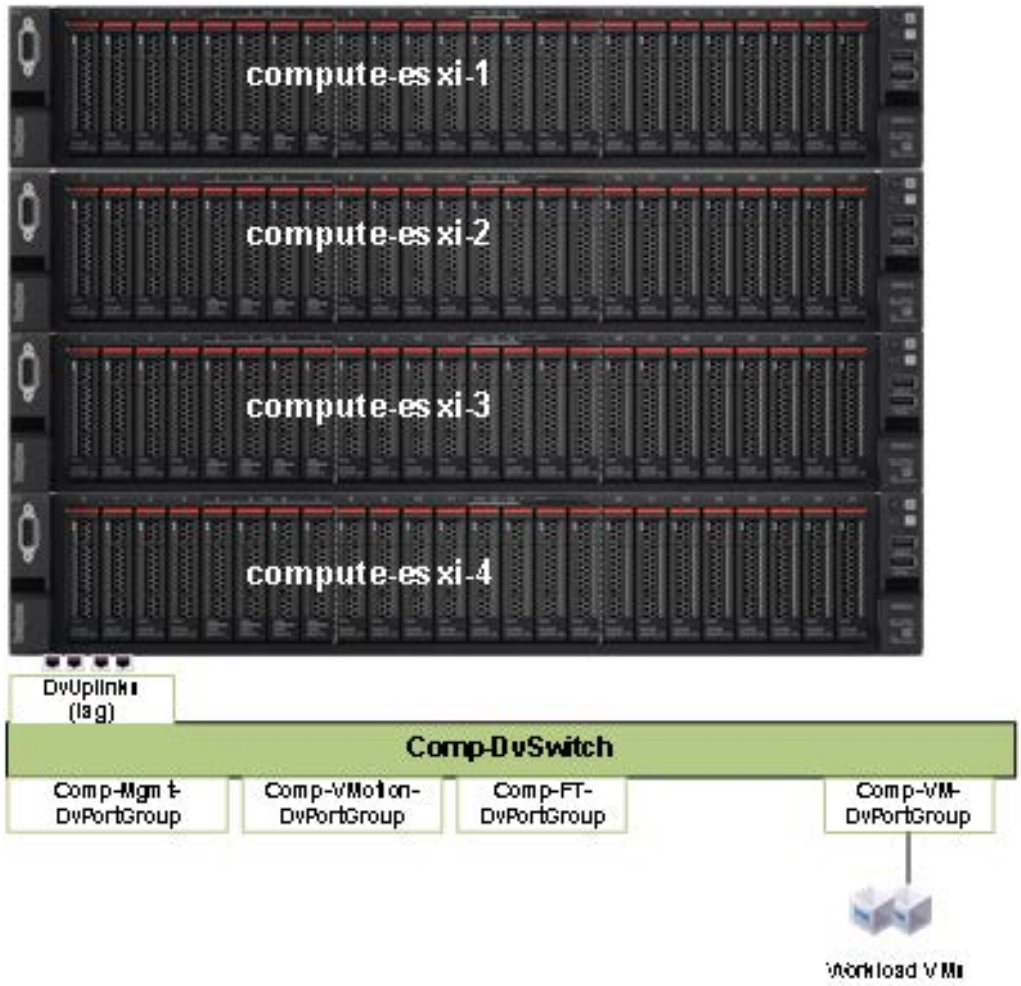

Figure 32: Compute Cluster VDS

8 Appendix: Bill of Materials

This appendix features the Bill of Materials (BOMs) for different configurations of hardware for vCloud Suite deployments. There are sections for edge servers, management servers, storage, and networking switches that are orderable from Lenovo.

The BOM lists in this appendix are not meant to be exhaustive and must always be confirmed with the configuration tools. Any description of pricing, support, and maintenance options is outside the scope of this document.

For connections between ToR switches and devices (servers, storage, and chassis), the connector cables are configured with the device. The ToR switch configuration includes only transceivers or other cabling that is needed for failover or redundancy.

8.1 Server BOM

The tables below list the BOM for a server (a minimum of two servers per cluster is needed). See Table 31 on page 52 for the Lenovo Flex System chassis BOM.

Table 25: ThinkSystem SR630

Code	Description	Quantity
7X02CTO1WW	ThinkSystem SR630 - 3yr Warranty	1
AUW3	Lenovo ThinkSystem Mainstream MB - 1U	1
AUW0	ThinkSystem MS 1U 8x2.5" Chassis	1
B4HJ	Intel Xeon Gold 6230 20C 125W 2.1GHz Processor	2
AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	12
AUWB	Lenovo ThinkSystem SR630 8x2.5" SATA/SAS Backplane	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 CV3 128GB SATA 6Gbps Non-Hot Swap SSD	2
AUNU	ThinkSystem QLogic QLE2742 PCIe 32Gb 2-Port SFP+ Fibre Channel	1
AUWC	ThinkSystem SR530/SR570/SR630 x8/x16 PCIe LP+LP Riser 1 Kit	1
AUWQ	Lenovo ThinkSystem 1U LP+LP BF Riser BKT	1
AUKK	ThinkSystem 10Gb 4-port SFP+ LOM	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AXCB	ThinkSystem Toolless Slide Rail Kit with 1U CMA	1
AVW9	ThinkSystem 750W (230V) Titanium Hot-Swap Power Supply	2
6311	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
5977	Select Storage devices - no configured RAID required	1
B4XA	VMware ESXi 6.7 U1 (factory installed)	1

Table 26: ThinkSystem SR650

Code	Description	Quantity
7X06CTO1WW	ThinkSystem SR650 - 3yr Warranty	1
AUQB	Lenovo ThinkSystem Mainstream MB - 2U	1
AUVX	ThinkSystem MS 2U 2.5" Chassis, up to 16	1
B4HJ	Intel Xeon Gold 6230 20C 125W 2.1GHz Processor	2
AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	12
AURA	ThinkSystem 2U/Twr 2.5" SATA/SAS 8-Bay Backplane	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 CV3 128GB SATA 6Gbps Non-Hot Swap SSD	2
AUNU	ThinkSystem QLogic QLE2742 PCIe 32Gb 2-Port SFP+ Fibre Channel	1
AUKK	ThinkSystem 10Gb 4-port SFP+ LOM	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AXCH	ThinkSystem Toolless Slide Rail Kit with 2U CMA	1
AUS8	ThinkSystem SR550/SR590/SR650 EIA Latch w/ VGA Upgrade Kit	1
AVWF	ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	2
6311	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
5977	Select Storage devices - no configured RAID required	1
B4XA	VMware ESXi 6.7 U1 (factory installed)	1

Table 27: ThinkSystem SD530

Code	Description	Quantity
7X21CTO1WW	ThinkSystem SD530 - 3yr Warranty	1
AUXN	ThinkSystem SD530 Computing Node	1
B4HJ	Intel Xeon Gold 6230 20C 125W 2.1GHz Processor	2
AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	12
AUYG	ThinkSystem SD530 3x2 SAS/SATA BP	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AUUL	ThinkSystem M.2 CV1 32GB SATA 6Gbps Non-Hot Swap SSD	2
B4XA	VMware ESXi 6.7 U1 (factory installed)	1

Table 28: ThinkSystem D2 Enclosure (for SD530)

Code	Description	Quantity
7X20CTO1WW	ThinkSystem D2 Enclosure -3yr Warranty	1
AUXM	ThinkSystem D2 Enclosure	1
AUY9	ThinkSystem D2 10Gb 8 port EIOM SFP+	1
AUYC	ThinkSystem D2 Slide Rail	1
AUYD	ThinkSystem D2 CMA (Cable Management Arm)	1
AUZ2	ThinkSystem D2 2000W Platinum PSU	2
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
AUNU	ThinkSystem QLogic QLE2742 PCIe 32Gb 2-Port SFP+ Fibre Channel	4
AVR1	ThinkSystem Single Ethernet Port SMM	1
AUY7	ThinkSystem D2 8-slot x8 Shuttle ASM	1

Table 29: ThinkSystem SN550

Code	Description	Quantity
7X16CTO1WW	ThinkSystem SN550 - 3yr Warranty	1
5977	Select Storage devices - no configured RAID required	1
B4HJ	Intel Xeon Gold 6230 20C 125W 2.1GHz Processor	2
AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	12
AUXP	Lenovo ThinkSystem SN550 Server	1
ATBS	Flex System CN4052S 4-port 10Gb Virtual Fabric Adapter	1
AVCV	ThinkSystem QLogic QML2692 Mezz 16Gb 2-Port Fibre Channel Adapter	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
AUUV	ThinkSystem M.2 CV3 128GB SATA 6Gbps Non-Hot Swap SSD	1
B4XA	VMware ESXi 6.7 U1 (factory installed)	1

8.2 Shared storage BOM

The following tables list the BOM for Lenovo ThinkSystem DM5000F storage array.

Table 30: Lenovo ThinkSystem DM5000F control enclosure

Code	Product Description	Quantity
7Y41CTO1WW	Controller : Lenovo ThinkSystem DM5000F All Flash Array	1
B38L	Lenovo ThinkSystem Storage 2U24 Chassis	1
B5RJ	DM Series Premium Offering	1
B39F	Lenovo ThinkSystem DM Series Cntr, 16Gb FC/10Gb Opt	2
B65R	Lenovo ThinkSystem23TB (6x 3.84TB, 2.5", SSD) Drive Pack for DM5000F	2
A1PJ	3m Passive DAC SFP+ Cable	2
B4BP	Lenovo ThinkSystem Storage USB Cable, Micro-USB	1
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
B6KC	Lenovo ThinkSystem DM Series ONTAP 9.5 SW, Encryption	1
	DM5000F Software Licenses	
B5AZ	DM Series SnapVault License	2
B4SJ	DM Series FCP Protocol License	2
B4SF	DM Series CIFS Protocol License	2
B4SU	TPM	2
B4SH	DM Series iSCSI Protocol License	2
B4SL	DM Series SnapRestore License	2
B4SP	DM Series SnapManager License	2
B4SM	DM Series FlexClone License	2
B4SG	DM Series NFS Protocol License	2
B4SK	DM Series SnapMirror License	2
B4SN	DM Series Software Encryption License	2

8.3 Chassis BOM

Table 31 lists the BOM for a Lenovo Flex System chassis.

Table 31: Chassis BOM

Code	Description	Quantity
8721HC2	Lenovo Flex System Enterprise Chassis w/ CMM2	1
ASUS	Lenovo Flex System Enterprise Chassis w/ CMM2	1
ASUV	Lenovo Flex System Fabric SI4093 System Interconnect Module	2
6292	2m, 16A/100-250V, C19 to IEC 320-C20 Rack Power Cable	6
A0UC	Flex System Enterprise Chassis 2500W Power Module Standard	2
A0UD	Flex System Enterprise Chassis 2500W Power Module	4
A0UA	Flex System Enterprise Chassis 80mm Fan Module	4

8.4 Networking BOM

Table 32 and Table 33 list the BOMs for the network switches for 1 GbE and 10 GbE connectivity respectively

Table 32: RackSwitch G8052

Code	Description	Quantity
7159G52	Lenovo System Networking RackSwitch G8052 (Rear to Front)	1
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
3802	1.5m Blue Cat5e Cable	3
A3KP	Lenovo System Networking Adjustable 19" 4 Post Rail Kit	1

Table 33: RackSwitch G8272

Code	Description	Quantity
7159CRW	Lenovo System Networking RackSwitch G8272 (Rear to Front)	1
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
A3KP	Lenovo System Networking Adjustable 19" 4 Post Rail Kit	1
A1DP	1m QSFP+-to-QSFP+ cable	1
A1DM	3m QSFP+ DAC Break Out Cable	0

8.5 SAN Networking BOM

Table 34 lists the BOM for the SAN network (Fibre Channel)

Table 34: ThinkSystem DB620S

Code	Description	Quantity
6415HC4	Lenovo ThinkSystem DB620S Gen6 FC Switch 24x32Gb SWL SFP (1yr)	1
B4QE	Lenovo ThinkSystem DB620S Gen6 FC Switch 24x32Gb SWL SFP (1yr)	1
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
AVFX	1.0m Green Cat6 Cable	1

8.6 Rack BOM

Table 35 lists the BOM for a rack.

Table 35: Rack BOM

Code	Description	Quantity
93634PX	42U 1100mm Enterprise V2 Dynamic Rack	1
39Y8941	DPI Single Phase C19 Enterprise PDU (without line cord)	6
40K9614	DPI 30a Line Cord (NEMA L6-30P)	6

Resources

For more information about the topics that are described in this document, see the following resources:

- vSphere Hypervisor (ESXi):

  vmware.com/products/vsphere-hypervisor

- vCenter Server:

  vmware.com/products/vcenter-server

- vCloud Suite:

  vmware.com/products/vcloud-suite

- vRealize Suite:

  vmware.com/products/vrealize-suite.html

- vRealize Automation:

  vmware.com/products/vrealize-automation

- vRealize Automation Reference Architecture:

  vmware.com/files/pdf/products/vCloud/vRealize-Automation-6x-Reference-Architecture.pdf

- vRealize Operations:

  vmware.com/products/vrealize-operations

- vRealize Business:

  vmware.com/products/vrealize-business

- vRealize Log Insight:

  vmware.com/products/vrealize-log-insight

Document history

Version 1.0	7 August 2017	• First version for Lenovo ThinkSystem servers
Version 1.1	15 August 2018	• Updated for vSphere 6.7 • Added HyTrust security products • Removed VMware Hybrid Cloud Manager and vCloud Air
Version 1.2	7 June 2019	• Updated for Intel Xeon Scalable Processor generation 2 • Updated for vSphere 6.7 u1 • Added ThinkSystem DM5000F All Flash storage • Removed ThinkSystem DS6200 storage array

Trademarks and special notices

References in this document to Lenovo products or services do not imply that Lenovo intends to make them available in every country.

Lenovo, the Lenovo logo, ThinkSystem, ThinkAgile, ThinkCentre, ThinkVision, ThinkVantage, ThinkPlus and Rescue and Recovery are trademarks of Lenovo.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo’s current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-Lenovo websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this Lenovo product and use of those websites is at your own risk.