Appendix E

VCP6-DCV to VCP6.5-DCV Exam Updates

This appendix is intended to provide you with updated information, as VMware has released a minor update to the exam upon which this book is based. When VMware releases an entirely new exam (for example, v7), the changes are usually too extensive to provide in a simple update appendix. In that case, you might need to consult the new edition of the book for the updated content.

This appendix attempts to fill the void that occurs within any print book. In particular, it does the following:

- Mentions technical items that were not mentioned elsewhere in the book
- Covers new topics that VMware has added to the minor update of the exam
- Provides a way to get up-to-the-minute current information about content for the exam

Always Get the Latest at the Companion Website

The main purpose of this appendix is to be a living, changing document; it is important that you look for the latest version online at the book's companion website.

To do so, follow these steps:

Step 1. Browse to http://www.vmwarepress.com/title/9780789756480.

Step 2. Select the Updates tab.

Step 3. Download the latest "Appendix E Exam Update" document.

NOTE The downloaded document has a version number. Be sure to compare the version of the Appendix E Exam Update (version 1.0) with the latest online version of this appendix.

Technical Content

The version for this appendix is version 1.0. The version history is as follows:

- **Version 1.0**: Added content to cover VCP6.5-DCV objectives.

Changes to Exam Blueprint

The following table summarizes the changes in exam blueprint in VCP6.5-DCV versus VCP6.0-DCV. It covers blueprint changes, blueprint additions, and associated vSphere 6.5 technology changes.

The Blueprint Changes column identifies items that changed in the 6.5 blueprint from the 6.0 blueprint, including the following:

- Small changes in the wording of an objective
- Items that were moved or removed
- New objectives and large objectives that are now split into separate objectives

The Blueprint Additions column identifies items that are new in the 6.5 blueprint. Many of these items relate to new features in vSphere 6.5, but some relate to unchanged vSphere details that are now covered in the exam. The last column identifies new features and technology changes in vSphere 6.5 that are related to the exam objective but are not reflected by an exam change or addition.

You can ignore the Blueprint Changes column and focus on the last two columns, which are covered in detail in this appendix.

Blueprint Objective	Blueprint Changes	Blueprint Additions	Associated v6.5 Technology Changes

1.1	Changed objective from: VMware Directory ServiceTo: VMware Identity Sources		
1.2	The Securing VMs items moved to section 1.4.	[lb] Configure encrypted vMotion [lb] Describe (ESXi host) Secure Boot	
1.3	Replaced SSO with PSC in these items: [lb] Describe PSC architecture and components [lb] Perform a multisite PSC installation [lb] Upgrade a single/complex PSC installation Changed item from: Configure/Manage Active Directory Authentication To: Identity Sources		
1.4	Contains items moved from section 1.2: [lb] Harden VM access [lb] Harden a VM against DoS attacks Contains repeats from section 1.2: [lb] Describe Secure Boot [lb] Configure encrypted vMotion	[lb] Enable/disable VM encryption [lb] Describe (VM) Secure Boot	
2.1		[lb] Configure multiple VMkernel default gateways [lb] Configure ERSPAN [lb] Create and configure custom TCP/IP stacks [lb] Configure NetFlow	
2.2			
3.1	Changed objective from: Manage vSphere Storage Virtualization To: Manage vSphere Integration with Physical Storage	[lb] Perform NFS v3 and v4.1 configuration	LUN scalability
3.2		[lb] Create vSAN cluster [lb] Create disk groups [lb] Monitor vSAN [lb] Describe vVOLs [lb] Understand a vSAN iSCSI target	

Security

Configure Encrypted vMotion

Encrypted vMotion is a new feature that encrypts vMotion data as it travels over the network from the source to the target ESXi host. Previously, in order to protect vMotion data, you needed to place it on a protected network. Encrypted vMotion enables more flexibility and easier implementation. A 256-bit random key and a 64-bit nonce, used only once for this vMotion migration, are generated. The nonce is used to generate a unique counter for every packet sent over the network. This prevents replay attacks and enables the encryption of 264 128-bit blocks of data. The key and the nonce are packaged into a vSphere vMotion migration specification. The migration specification is sent to both systems in the cluster via the existing encrypted management connections between the vCenter Server instance and the ESXi hosts.

The vSphere vMotion traffic begins with every packet being encrypted with the key and the nonce on the source host.. Each uniquely encrypted packet is decrypted on the receiving host, completing the vSphere vMotion migration.

Starting with vSphere 6.5, vMotion always uses encryption when migrating encrypted virtual machines. You cannot turn off encrypted vMotion for encrypted virtual machines. For non-encrypted virtual machines, you can set Encrypted vMotion per virtual machine to *disabled*, *opportunistic* or *required*. Disabled means to not use Encrypted vMotion. Opportunistic means to use Encrypted vMotion if the source and target hosts support it, which requires ESXi 6.5 or later. Required means to disallow vMotion if the source and target do not support it.

To enable Encrypted vMotion for a virtual machine, follow these steps:

1. Right-click the VM and select **Edit Settings**.
2. Select **VM Options**.
3. Click **Encryption** and choose an option (Disabled, Opportunistic or Required) in the **Encrypted VMotion** drop-down menu.

Describe ESXi Host Secure Boot

vSphere 6.5 introduces Secure Boot support for ESXi hypervisor. Unified Extensible Firmware Interface (UEFI) Secure Boot is a mechanism that ensures that only trusted code is loaded by the Extensible Firmware Interface (EFI) firmware prior to OS handoff. When Secure Boot is enabled, the UEFI firmware validates the digitally signed kernel of an OS against a digital certificate stored in the UEFI firmware. For ESXi 6.5, this capability is further leveraged by the ESXi kernel, adding cryptographic assurance of ESXi components.

ESXi is already composed of digitally signed packages called vSphere installation bundles (VIBs). These packages are never broken open. At boot time, the ESXi file system maps to the content of those packages. By leveraging the same digital certificate in the host UEFI firmware used to validate the signed ESXi kernel, the kernel then validates each VIB using the Secure Boot verifier against the firmware-based certificate, ensuring a cryptographically “clean” boot.

When Secure Boot is enabled, ESXi will prevent the installation of unsigned code on ESXi. To install unsigned code such as beta drivers, you must disable Secure Boot. When Secure Boot is enabled, the Secure Boot verifier will run, detect the unsigned VIB, and crash the system, which produces the Purple

Screen of Death (PSOD) event that identifies the VIB that must be removed. To remediate, boot the ESXi host with Secure Boot disabled, remove the VIB, and reboot with Secure Boot enabled.

Enable/Disable VM Encryption

vSphere 6.5 introduces Secure Boot support of virtual machines. VM Encryption is a VM-agnostic method (guest OS–agnostic method) of encryption for virtual machines that is scalable, easy to implement, and easy to manage. It provides security to virtual machine disk (VMDK) data by encrypting I/O from a virtual machine before it gets stored in the VMDK Advantages. Here are some specifics:

- Encryption works with any guest OS, because encryption occurs at the hypervisor level.
- You can apply and manage encryption via policy, which uses a framework that leverages vSphere Storage Policy Based Management (SPBM).
- Encryption keys and configuration are not contained in the VM guest OS.
- Key management is based on the industry standard Key Management Interoperability Protocol (KMIP). The vCenter Server is a KMIP client, which works with many KMIP key managers, providing flexibility and separation of duties between key usage and key management. In large environments, the Security Team performs the key management duty and vCenter performs key usage.
- VM Encryption leverages the latest CPU hardware advances in AES-NI encryption. The Advanced Encryption Standard instruction set is an extension to the x86 instruction set and provides accelerated encryption and decryption functions on a per-core basis in the CPU.

VM encryption supports the encryption of virtual machine files, virtual disk files, and core dump files. Some of the files associated with a virtual machine, such as log files, VM configuration files, and virtual disk descriptor files, are not encrypted. Core dumps on an ESXi host that has encryption mode enabled are always encrypted. Core dumps on the vCenter Server are not encrypted.

VM encryption uses vSphere APIs for I/O filtering (VAIO), which is typically called the IOFilter. The IOFilter is an ESXi framework that allows for the interception of VM I/Os in the virtual SCSI emulation (VSCSI) layer, which is just below the VM and above the VMFS file system. It enables VMware and third-party developers to develop services using VM I/O, such as encryption, caching, and replication. It is implemented entirely in user space, which isolates it cleanly from the core architecture and core functionality of the hypervisor. In case of any failure, only the VM in question would be affected. Multiple filters can be enabled for a particular VM or a VMDK, which are typically chained in a manner so that I/Os are processed by each of these filters serially, one after the other, and then finally either passed down to VMFS or completed within one of the filters.

The vCenter Server requests keys from the Key Management System (KMS). The keys used are the key encryption key (KEK) and AES-256 keys. vCenter Server stores only the ID of each KEK, but not the key itself.

When you create an encrypted virtual machine from the vSphere Web Client, all virtual disks are encrypted. You can later add disks and set their encryption policies. You cannot add an encrypted disk to a virtual machine that is not encrypted, and you cannot encrypt a disk if the virtual machine is not encrypted.

VM encryption supports the encryption of virtual machine files, virtual disk files, and core dump files. Some of the files associated with a virtual machine, such as log files, VM configuration files, and virtual disk descriptor files, are not encrypted.

The main prerequisite for VM Encryption is that the ESXi host must be configured with encryption mode enabled. Only administrators with encryption privileges can perform encryption and decryption tasks.

The default Administrator system role includes all Cryptographic Operations privileges. A new default role, the No Cryptography Administrator, supports all Administrator privileges except for the Cryptographic Operations privileges. You can create a custom role that contains granular Cryptographic Operations privileges such as **Cryptographic operations > Encrypt** (allows a user to encrypt a virtual machine or virtual disk) and **Cryptographic operations > Add disk** (allows a user to add a disk to an encrypted virtual machine).

The vSphere Web Client can be used to encrypt and decrypt virtual machines. The vSphere Web Services SDK can also be used to encrypt and decrypt virtual machines as well as to perform a deep recrypt (using a different DEK) or a shallow recrypt (using a different KEK) of a virtual machine. The **crypto-util** command line utility can be used to decrypt core dumps, check for file encryption, and perform management tasks on the ESXi host.

The main prerequisites for encrypting a virtual machine are to enable encryption on the ESXi host and to create an encryption storage policy.

NOTE You cannot encrypt a virtual disk of an unencrypted virtual machine.

Create an Encryption Storage Policy

You can use the following steps to create an encryption storage policy:

1. In the vSphere Web Client, drill down to **Home** > **Policies and Profiles** > **VM Storage Policies**.
2. Click **Create VM Storage Policy**.
3. Use the wizard to do the following:
    a. Provide a policy name.
    b. Select the **Use common rules in VM storage policy** check box.
    c. Click **Add component** and select **Encryption > Default Encryption Properties**.
    d. Deselect the **Use rule-sets in the storage policy** check box.
    e. Ensure that the **Storage Compatibility** > **Compatible** check box is selected.

VM Encryption: Prerequisites

You must meet the following prerequisites before you can encrypt a virtual machine:

- Establish a trusted connection with the KMS and select a default KMS.
- Create an encryption storage policy.
- Power off the virtual machine.
- Ensure that **Security Profile > Encryption Mode** is enabled on the ESXi host.

- Ensure you have the required privilege (**Cryptographic operations > Encrypt new**).

Enable VM Encryption: Procedure

You can use the following steps to encrypt a virtual machine:

1. In the vSphere Web Client, right-click the virtual machine and select **VM Policies** > **Edit VM Storage Policies**.
2. In the drop-down menu, select the encryption storage policy and choose either **Apply to all** (to apply to the virtual machine and all of its virtual disks) or **Apply** (to apply to just the virtual machine but not to the virtual disks).

You can use the following steps to encrypt a virtual disk:

1. In the vSphere Web Client, right-click the virtual machine and select **VM Policies** > **Edit Settings**.
2. Drill down to and select the virtual disk and choose the encrypted storage policy in the **VM Storage Policy** pull-down menu.
3. Click **OK**.

Describe VM Secure Boot

Virtual machines must be booted from the EFI firmware to enable Secure Boot. EFI firmware supports Windows, Linux, and nested ESXi. For Secure Boot to work, the guest OS must also support Secure Boot. Examples include Windows 8 and Windows Server 2012 and newer, VMware Photon OS, RHEL/Centos 7.0, Ubuntu 14.04, and ESXi 6.5.

When the virtual machine boots, only components with valid signatures are allowed. The boot process stops with an error if it encounters a component with a missing or invalid signature.

For Linux virtual machines, VMware Host-Guest Filesystem is not supported in Secure Boot mode. Remove VMware Host-Guest Filesystem from VMware Tools before you enable Secure Boot.

VM Secure Boot: Prerequisites

You must meet the following prerequisites before you can enable Secure Boot for a virtual machine:

- The virtual machine's **Boot Options** > **Firmware** is set to **EFI**.
- The virtual machine uses virtual hardware version 13 or later.
- The guest OS supports UEFI Secure Boot.
- VMware Tools version 10.1 is installed.

Enable VM Secure Boot: Procedure

To enable Secure Boot for a virtual machine, use the vSphere Web Client to edit the virtual machine settings, ensure **Boot Options** > **Firmware** is set to **EFI**, and select the **VM Options** > **Secure Boot** check box.

Networking

Configure Multiple VMkernel Default Gateways

vSphere 6.0 introduced the ability to create custom TCP/IP stacks on a host. This enables you to create multiple stacks and assign each VMkernel adapter to a separate stack with a separate default gateway. The only services supporting this feature are vMotion and Provisioning (NFC). Each TCP/IP stack on a host can have only one default gateway. This default gateway is part of the routing table, and all services that operate on the TCP/IP stack use it. In vSphere 6.5, a new feature is introduced where you can override the default gateway for a VMkernel adapter. This allows multiple VMkernel adapters to use unique gateways, even if they share the same TCP/IP stack. Again, only vMotion and Provisioning (NFC) support the ability to configure L3 connectivity for services in 6.5. Thus, these two services can override the default gateway for the host, and have their own individual gateway address.

To change the gateway for a VMkernel virtual adapter, use the vSphere Web Client to edit the settings of the adapter, select the **IPv4** > **Use static IPv4 settings** > **Override default gateway for this adapter** check box, and enter a gateway address.

Configure ERSPAN

Encapsulated Remote Switched Port Analyzer (ERSPAN) enables you to monitor traffic on multiple network interfaces or virtual local area networks (VLANs) and then send the monitored traffic to one or more destinations. To use ERSPAN in vSphere 6.5, create a port-mirroring session on the distributed virtual switch and set the type to Encapsulated Remote Mirroring (L3) source.

Port Mirroring

Port mirroring allows you to mirror a distributed port's traffic to other distributed ports or specific physical switch ports. It is typically used to analyze and debug data or diagnose errors on a network. You can create port mirror sessions of various types:

- **Distributed port mirroring**: Mirrors packets from a set of distributed ports to another set of distributed ports on the same ESXi host.
- **Remote mirroring source**: Mirrors packets from a set of distributed ports to a set of uplink ports on the same ESXi host.
- **Remote mirroring destination**: Mirrors packets from a set of VLANs to a set of distributed ports.
- **Encapsulated remote mirroring (L3) source**: Mirrors packets from a set of distributed ports to the IP address of a remote agent through an IP tunnel.
- **Distributed port mirroring (legacy)**: Mirrors packets from a set of distributed ports to a set of distributed and uplink ports on the same host.

You can create a port-mirroring session using the vSphere Web Client by following this procedure:

1. Select a distributed virtual switch and click **Configure** > **Settings**.
2. Select **Port Mirroring** and click **New**.

3. Select a session type, such as **Encapsulated Remote Mirroring (L3) Source**.
4. Set the session properties, which are dependent on the session type. These properties include, Name, Status (enable or disable), Normal I/O on ports, Mirrored packet length (bytes), sampling rate, and description.
5. Select the port-mirroring sources, which are dependent on the session type. You can select ports from a list, add ports by port number, set traffic direction, and specify a source VLAN (for remote mirror destination only).
6. Select the port-mirroring destinations, which are dependent on the session type. You can select ports from a list, add ports by port number, select an uplink, or specify an IP address (ERSPAN).

Create and Configure Custom TCP/IP Stacks

In vSphere 5.1 and earlier versions, a single TCP/IP stack is used for all system traffic, which limits your ability to segregate network traffic. In vSphere 5.5, custom TCP/IP stacks may be created. In vSphere 6.0, VMware added the vMotion TCP/IP stack and the Provisioning TCP/IP stack. The Default TCP/IP stack can be used for all management, migration, and provisioning activities. If the vMotion TCP/IP stack is used, vMotion traffic is automatically disabled on the Default TCP/IP stack. If the Provisioning TCP/IP stack is used, traffic for cold virtual machine migration, snapshotting, and cloning uses this stack instead of the Default TCP/IP stack.

In vSphere 5.5 and later, you can create a custom TCP/IP stack using the following command:

```
esxcli network ip netstack add -N="stack_name"
```

After creating a new custom TCP/IP stack using this command, you can use the vSphere Web Client to configure the stack, as follows:

1. Select the ESXi host and select **Configure** > **Networking**.
2. Select **TCP/IP configuration** and select a TCP/IP stack.
3. Change the name, DNS settings, and the routing or advanced settings of the stack.
4. Click **OK**.

Configure NetFlow

A network flow is a set of network packets sharing the same source IP address, target IP address, source TCP/UDP port, destination TCP/UDP port, and other settings. NetFlow is a feature introduced on Cisco routers that collects network flows, which can be used for network analysis, such as to identify the network usage by application components and to identify the root cause of bottlenecks.

In vSphere, you can configure a distributed virtual switch to collect network flows and send data to a NetFlow collector. Follow these steps to configure NetFlow using the vSphere Web Client:

1. Select a distributed virtual switch and select **Actions** > **Settings** > **Edit NetFlow**.
2. Provide the IP address and port used by the NetFlow collector.
3. Set the Observation Domain ID (to identify data related to the virtual switch).
4. Optionally, configure these other settings:

- Switch IP address (to facilitate viewing the data under a single network device instead of under separate devices for each host)
- Active and idle flow export timeouts
- Sampling rate
- Process internal flows only (to only collect data on network activity between virtual machines on the same ESXi host)

5. Click **OK**.

Storage

Perform NFS v3 and v4.1 Configuration

Review these sections in the exam guide:

- Chapter 4: Storage Protocols > NFS Protocol
- Chapter 4: Create an NFS Share for Use with vSphere
- Chapter 6: Differentiate VMware File System Technologies
- Chapter 6: Compare and Contrast VMFS and NFS Datastore Properties
- Chapter 6: Differentiate NFS 3.x and 4.1 Capabilities
- Chapter 6: Connect an NFS 4.1 Datastore Using Kerberos
- Chapter 6: Mount/Unmount an NFS Datastore

NFS 4.1, which has been supported since vSphere 6.0, has stronger cryptographic algorithms with Kerberos authentication using Microsoft Active Directory (AD). vSphere 6.5 introduces Kerberos integrity check (SEC_KRB5i) along with Kerberos authentication. There is also support for IPv6 with Kerberos. Host Profiles in vSphere 6.5 includes support for NFS 4.1. These enhancements provide better security for customers.

If you use NFS 4.1 storage with Kerberos, you must add each ESXi host to an Active Directory domain and enable Kerberos authentication. Kerberos integrates with Active Directory to enable single sign-on and provides an extra layer of security when used across an insecure network connection.

When multiple ESXi hosts share the NFS 4.1 datastore, you must use the same Active Directory credentials for all hosts that access the shared datastore. You can automate the assignment process by setting the user in Host Profiles and applying the profile to all ESXi hosts. Configure the Network Time Protocol (NTP) settings on these hosts to ensure their time is properly synchronized.

NOTE Enable AES256-CTS-HMAC-SHA1-96 or AES128-CTS-HMAC-SHA1-96 encryption modes on AD. The NFS 4.1 client does not support the DES-CBC-MD5 encryption mode.

When you use NFS 4.1 with Kerberos, you must change the DNS settings on ESXi hosts. The settings must point to the DNS server that is configured to hand out DNS records for the Kerberos Key

Distribution Center (KDC). For example, use the Active Directory server address if AD is used as a DNS server.

LUN Scalability

In vSphere 6.0, the maximum number of LUNs is 256 and number of paths is 1,024. These limits pose challenges for customers in the following scenarios:

- When the infrastructure has eight paths to each LUN, the effective maximum number of LUNs per host is 128.
- Customers who tend to use small LUNs to support full LUN backup and restore may exhaust the LUN limits.
- Active-active clusters are effectively limited to 128 LUNs (half the normal maximum).

In vSphere 6.5, these scenarios are addressed due to support for 512 LUNs and 2,000 paths.

VSAN

Although the VCP6-DCV exam blueprint already covered VSAN, the 6.5 exam blueprint explicitly added several VSAN-specific items.

Create a VSAN Cluster

To create a VSAN cluster, apply following characteristics, prerequisites, and procedure:

Characteristics:

- You can have multiple vSAN clusters for each vCenter Server instance. You can use a single vCenter Server to manage more than one vSAN cluster.
- vSAN consumes all devices, including flash cache and capacity devices, and does not share devices with other features.
- vSAN clusters can include hosts with or without capacity devices. The minimum requirement is three hosts with capacity devices. For best results, create a vSAN cluster with uniformly configured hosts.
- If a host contributes capacity, it must have at least one flash cache device and one capacity device.
- In hybrid clusters, the magnetic disks are used for capacity and flash devices for read and write cache. vSAN allocates 70% of all available cache for read cache and 30% of available cache for the write buffer. In a hybrid configuration, the flash devices serve as a read cache and a write buffer.
- In all-flash clusters, one designated flash device is used as a write cache and additional flash devices are used for capacity. In all-flash clusters, all read requests come directly from the flash pool capacity.
- Only local or direct-attached capacity devices can participate in a vSAN cluster. vSAN cannot consume other external storage, such as SAN or NAS, attached to cluster.

Prerequisites:

To use the full set of vSAN capabilities, the ESXi hosts that participate in vSAN clusters must be version 6.5 or later. During the vSAN upgrade from previous versions, you can keep the current on-disk format version, but you cannot use many of the new features. vSAN 6.6 and later software supports all on-disk formats. Prior to creating the VSAN cluster, verify the following:

- ESXi hosts
  - Use ESXi 6.5 or later. (Alternatively, if you do not need the latest VSAN features, ESXi 5.5 Update 1 or later hosts can join the VSAN cluster. All hosts in the VSAN cluster must have the same on-disk format.)
  - Prepare at least three hosts, and preferably four or more hosts.
  - A host does not have to contribute storage to the VSAN cluster. However, to be able to access a VSAN datastore, a host must be a member of the VSAN cluster.
- Memory
  - Configure each host with at least 8GB.
  - Configure each host for 32GB if you need larger configurations and better performance.
- Storage controller
  - Verify that the storage I/O controllers, drivers, and firmware versions are VSAN certified.
  - Configure the controller for passthrough or RAID 0 mode.
  - Disable the controller cache and advanced features. Alternatively, set the controller cache to 100%.
  - Use a controller with queue depth of at least 256.
- Cache and capacity
  - Ensure that each host that contributes storage to the VSAN cluster has at least one cache drive and one capacity drive. These devices must be dedicated to VSAN and not used for other purposes, such as Virtual Flash, VMFS, or boot partitions.
  - For best results, create the VSAN cluster with uniformly configured ESXi hosts.
- Network connectivity
  - Configure each host with at least one network adapter.
  - For hybrid configurations, ensure at least 1GbE is dedicated to VSAN
  - For all flash configurations, ensure at least 10GbE is dedicated to VSAN.
- vCenter Server
  - Use vCenter Server 6.5 or later.
- License key
  - Use a valid VSAN license key that supports the features you require, such as all flash, stretched clusters, deduplication, and compression.

- Ensure the license capacity is equal to (or greater than) the total number of CPUs that participate in the VSAN cluster.

Procedure:

The following procedure can be used to create a VSAN cluster using the vSphere Web Client:

1. Configure a VMkernel network for VSAN.
   a. On each host that will participate in the cluster, create a VMkernel network adapter.
   b. In the settings of the VMkernel network adapter, select the **vSAN traffic** check box.
2. Create a VSAN cluster.
   a. Right-click a data center and select **New Cluster**.
   b. Provide a name for the cluster.
   c. Select **VSAN Turn ON** check box and click **OK**.
   d. Add hosts to the cluster by dragging and dropping existing hosts or by right-clicking the cluster and choosing **Add Host**.
   e. Select the vCenter Server, click the **Configure** tab, and then click **Storage Providers**. Verify that each host has a VSAN storage provider, but only one is enabled.
3. Configure a VSAN cluster.
   a. Select the VSAN cluster and click the **Configure** tab.
   b. Select the desired VSAN capabilities, such as deduplication, encryption, and fault tolerance mode.
   c. Click **Next**.
   d. Use the **Claim disks** page to select the cache and capacity disks to be used by the cluster. Click **Next**.
   e. If you selected **Configure two host vSAN cluster**, then choose a witness host and claim disks for the witness host.
   f. If you selected **Configure stretched cluster**, define fault domains for the cluster, choose a witness host, and claim disks for the witness host.
   g. If you selected **Configure fault domains**, define fault domains for the cluster.
   h. On the **Ready to Complete** page, click **Finish**.

Create Disk Groups

When you create disk groups, you must specify each host and each device to be used for the vSAN datastore. You organize cache and capacity devices into disk groups. To create a disk group, you define the disk group and individually select devices to include in the disk group. Each disk group contains one flash cache device and one or more capacity devices.

The vSAN cluster initially contains a single vSAN datastore with zero bytes consumed. As you create disk groups on each host and add cache and capacity devices, the size of the datastore increases according to the amount of physical capacity added by those devices. vSAN creates a single distributed vSAN datastore using the local empty capacity available from the hosts added to the cluster.

If the cluster requires multiple flash cache devices, you must create multiple disk groups, because a maximum of one flash cache device is allowed per disk group.

NOTE When you add an ESXi host to a vSAN cluster, the local storage from that host is not added to the vSAN datastore automatically. You have to create a disk group using storage from the new ESXi host.

The following procedure can be used to create a disk group on a VSAN host using the vSphere Web Client:

1. Select the VSAN cluster and click the **Configure** tab.
2. Under **VSAN**, click **Disk Management**.
3. Select the host and click the **Create a new disk group** icon.
4. Select a flash device to be used for cache.
5. In the **Capacity type** drop-down menu, select the type of capacity disks to be used (HDD or Flash).
6. Select the capacity drives.
7. Click **OK**.

Alternatively, you claim storage devices for a VSAN cluster and allow VSAN to organize the devices into default disk groups. To do this, select the cluster and then select **Configure** > **Disk Management** > **Claim Disks**. For each host, select any available desired device and click **Claim for cache tier** or click **Claim for capacity tier**.

Monitor VSAN

You can use the vSphere Web Client to monitor the following items:

- **The VSAN cluster**: Select the VSAN cluster and click **Monitor** > **vSAN**. Select **Physical Disks** to review hosts, cache devices, and capacity devices. Select **Health** to review VSAN health categories. Select **Configure** > **General** to check cluster status, Internet connectivity, and on-disk format.
- **VSAN capacity**: Select the VSAN cluster and click **Monitor** > **vSAN**. Select **Capacity** to review provisioned and used capacity. Here, you can view the percentage of capacity used by object type, such as virtual disks, swap objects, file system overhead, and deduplication/compression overhead.
- **Virtual devices in the VSAN cluster**: Select the VSAN cluster and click **Monitor** > **vSAN**. Select **Virtual Disks** to review the virtual disks in the VSAN cluster, their physical disk placement, and compliance failures.
- **Resynchronization tasks in the VSAN cluster**: Select the VSAN cluster and click **Monitor** > **vSAN**. Select **Resyncing Components** to track the progress of resynchronization of virtual machine objects and the number of remaining bytes.
- **Devices that participate in the VSAN datastore**: Select the VSAN cluster, click **Configure** > **Device Backing**, and select a disk group. View the devices in the **Disks** table.

- **VSAN health**: Select the VSAN cluster, click **Configure** > **VSAN** > **Health and Performance**, and then click the **Health Services > Edit Settings** button. You can turn on periodic vSAN health checks covering hardware compatibility, network configuration and operation, advanced vSAN configuration options, storage device health, and virtual machine objects.
- **VSAN performance**: Select the VSAN cluster and click **Configure** > **VSAN** > **Health and Performance**. Click **Edit** and select the **Turn on vSAN performance service** check box. This allows you to monitor the performance of your vSAN environment and investigate potential problems. With this setting, the cluster summary displays an overview of vSAN performance statistics, including IOPS (input/output operations per second), throughput, and latency.
- **VSAN default alarms**: You can examine the configuration of the VSAN default alarms and respond to these alarms whenever they are triggered. You cannot modify these alarms, but you can create custom VSAN alarms. To view the default VSAN alarms, select the cluster, select **Configure** > **Alarm Definitions**, and search for "vSAN."
- **Customer VSAN Alarms based on VMkernel Observations (VOBs)**: Select the vCenter Server and then select **Configure** > **Alarm Definitions** to create an event-based alarm. In the wizard, select **specific event occurring on this object** and use **Triggers** to add a vSAN event.

NOTE When a hardware device, host, or network fails, or if a host is placed into maintenance mode, vSAN initiates resynchronization in the vSAN cluster. The following events can trigger resynchronization: changing a virtual machine storage policy, restarting a host after a failure, recovering hosts from a failure, evacuating data by using the full data migration mode before placing a host in maintenance mode, and exceeding the capacity threshold (80% by default) of a capacity device.

Describe vVOLs

Review these sections in the exam guide:

- Chapter 5: Explain VSAN and VVOL Architectural Components
- Chapter 5: Determine the Role of Storage Providers in VVOLs
- Chapter 5: Create/Modify VMware Virtual Volumes (VVOLs)
- Chapter 5: Create Virtual Volumes Given the Workload and Availability Requirements

Understand a VSAN iSCSI Target

You can use the iSCSI target service to enable hosts and physical workloads that reside outside the vSAN cluster to access the vSAN datastore. This feature enables an iSCSI initiator on a remote host to transport block-level data to an iSCSI target on a storage device in the vSAN cluster.

After you configure the vSAN iSCSI target service, you can discover the vSAN iSCSI targets from a remote host. To discover vSAN iSCSI targets, use the IP address of any host in the vSAN cluster and the TCP port of the iSCSI target.

To enable the iSCSI target service, edit the settings of the cluster and select the **Enable vSAN iSCSI target service** check box. Configure the network settings, TCP port, and authentication method. Configure an iSCSI target, add one or more LUNs to the target, and configure an iSCSI initiator group.

NOTE The vSAN iSCSI target service does not support other vSphere or ESXi clients or initiators, third-party hypervisors, or migrations using raw device mapping (RDMs).

VMFS6

Configure VMFS6 Datastore

To create a VMFS6 datastore, use the New Datastore Wizard, just as you would to create other VMFS datastores, except select **VMFS6** as the datastore version. You can choose options such as **use all available partitions** and **use free space**.

VMFS6 uses a 1MB block size. You can choose **enable space reclamation granularity** and configure **space reclamation priority**.

After upgrading ESXi hosts to version 6.5, you can continue using existing VMFS5 datastores. You can migrate VMs from a VMFS5 datastore to a VMFS6 datastore, but you cannot upgrade a VMFS5 datastore to version 6. Storage DRS, VMFS5, and VMFS6 can coexist in the same datastore cluster.

Traditional 512n storage devices use a native 512-bytes sector size. 512e storage devices are an advanced format in which the physical sector size is 4,096 bytes but the logical sector size emulates 512-bytes sector size. Storage devices that use the 512e format can support legacy applications and guest operating systems. When you set up a datastore on a 512e storage device, VMFS6 is selected by default. For 512n storage devices, the default option is VMFS5, but you can select VMFS6.

Any new VMFS5 or VMFS6 datastore uses a GUID partition table (GPT) to format the storage device, which supports datastores larger than 2TB. If your VMFS5 datastore has been previously upgraded from VMFS3, it continues to use the master boot record (MBR) partition format, which limits the datastore size to 2TB.

Understand SIOC Metrics for Datastore Clusters and DRS

See Chapter 6, “Monitor SIOC,” in the exam guide.

Upgrading

VMware Tools

Version 10.1 is bundled with vSphere 6.5 for Windows Vista and later. Version 10.0.12 is bundled with vSphere 6.5 for Windows pre-Vista.

Migrate to VCSA

VMware vCenter Server 6.5 has many new and innovative features. The installer is now supported on Microsoft Windows, macOS, and Linux operating systems without the need for any plug-ins. With vSphere 6.5, the VMware vCenter Server Appliance (VCSA) has surpassed the Windows installable version. It includes the following exclusive features:

- Migration Tool
- Improved appliance management
- Native high availability
- Native backup and restore
- VMware vSphere Update Manager as part of the appliance

If you use a Windows-based vCenter Server, you can either upgrade it directly to version 6.5 or migrate it to VCSA 6.5.

Understand the Migration Paths to VCSA

You can migrate a Windows-based vCenter Server 5.5 or 6.0 to VCSA 6.5, regardless of whether the source vCenter Server deployment uses embedded or external single sign-on (SSO) 5.5 or embedded or external Platform Services Controller (PSC) 6.0. Migrations from external SSL 5.5 and external PSC 6.0 require two steps, where you have migrated the SSO or PSC components in the first step.

You can perform the migration whether the source vCenter Server uses an embedded or external database. In either case, the database is converted to an embedded PostgreSQL database on the VCSA.

You should use the VMware Migration Assistant to prepare for the migration. Use it to gather the required information on the source vCenter Server, SSO, and PSC instances.

To prepare for the VCSA migration, you should prepare the vCenter Server database for migration. This includes ensuring that passwords will not expire soon, reducing the database size, running cleanup scripts, and backing up the database.

NOTE Use the following command to run the cleanup script for a Microsoft SQL Server Express database:

```
sqlcmd -E -S localhost\VIM_SQLEXP -d VIM_VCDB -
i path/cleanup_orphaned_data_MSSQL.sql
```

To use the command line interface (CLI) to upgrade a vCenter Server appliance to VCSA 6.5, first build a JSON configuration file and then use the **vcsa-deploy** command. An example of the command is shown here:

```
vcsa-deploy upgrade --accept-eula --acknowledge-ceip optional_arguments
path_to_the_json_file
```

Prior to performing the upgrade, you can use the **vcsa-deploy** command with the **–precheck-only** argument to verify that you met prerequisites.

You can change your vCenter Server deployment type after upgrade or migration to version 6.5. The key is to repoint the vCenter Server to another external PSC. To repoint the vCenter Server to another PSC, use the **cmsso-util** command, as shown here:

```
cmsso-util repoint --repoint-psc psc_fqdn [--dc-port port_number]
```

vCenter Migration Tool

The specific steps for using the Migration Tool depend on the source vCenter Server configuration and other factors. Here is an example of how to use the GUI to migrate a vCenter Server instance with an embedded SSO or PSC to a VCSA with an embedded PSC:

1. Download and mount the VCSA Installer ISO.
2. Download and run the VMware Migration Assistant on the source Windows vCenter Server.
3. Use the Migration Assistant to gather the required information.
4. Deploy the OVA file for migrating to the target VCSA with the embedded PSC. Use the wizard to perform the following steps:
    a. Choose **Migrate**.
    b. Deploy a VCSA with the proper compute size (tiny, small, medium, large, or x-large) and storage size (default, large, or x-large) .
    c. Configure the network and other settings.
5. Set up the target VCSA. Connect to the source vCenter Server SSO and transfer data.

NOTE Your window of downtime does not begin until you begin to set up the target appliance (Step 5). You cannot cancel or interrupt the process until it completes with the shutdown of the source deployment. Your window of downtime ends when the target appliance starts.

Clusters

Understand Network DRS

In vSphere 6.5, DRS takes network utilization into account, in addition to all the metrics used in vSphere 6.0. It monitors the transmit and receive rates of the connected physical uplinks and avoids placing VMs on hosts that are more than 80% utilized. vSphere DRS does not reactively balance the hosts solely based on network utilization. Rather, it considers network utilization when determining whether a host is the best recipient for a VM. This input improves vSphere DRS placement decisions during powering on and load-balancing operations.

Differentiate Load-Balancing Policies

The standard deviation model that has always been used for vSphere DRS load balancing works well, but in larger clusters, where distribution patterns become normalized, the model may result in outliers. Outliers are hosts whose utilization is greater than the average of the cluster, but do not pose a significant impact to the standard deviation. In vSphere 6.5, DRS improvements allow it to do a better job of addressing outliers and balancing cluster resources. These improvements involve using pairwise calculations, where the difference between the most utilized and least utilized hosts is calculated and used when making migration recommendations.

In vSphere 6.5, DRS includes new options:

- **VM distribution**: Enables DRS to evenly distribute VMs among the hosts in the cluster, but continues to ensure that VM performance is the top priority and that resource demand is met.
- **Memory metric for load balancing**: Enables DRS to use consumed memory instead of active memory as part of the primary metric in calculating memory load.
- **CPU over commitment**: Enables DRS to enforce a virtual-CPU-to-physical-CPU ratio that you specify. After the cluster reaches this defined value, no additional VMs will be allowed to power on.

Describe Predictive DRS

Predictive DRS is a new feature that leverages the predictive analytics of vRealize Operations (vROps) Manager and vSphere DRS. Together, these two products can provide workload balancing prior to the occurrence of resource utilization spikes and resource contention. Nightly, vROps calculates dynamic thresholds, which are used to create forecasted metrics for the future utilization of VMs. The metrics are then passed to vSphere DRS to determine the best placement and balance of VMs before resource utilization spikes occur. Predictive DRS helps prevent resource contention on hosts that run VMs with predictable utilization patterns.

Prerequisites include the following:

- vCenter Server 6.5 or later.
- Predictive DRS must be configured and enabled in both vCenter Server and vROps.
- The vCenter Server and vROps clocks must be synchronized.

To configure Predictive DRS, use the vROps GUI to add a vCenter Server adapter instance, choose **Advanced Settings**, and select **True** in the **Provide data to vSphere Predictive DRS** drop-down menu

Backup, Recovery, and Replication

Configure VCSA File-Based Backup and Restore

Native backup and restore for the vCenter Server Appliance is new in version 6.5. It enables users to back up vCenter Server and Platform Services Controller appliances directly from the vCenter Server Appliance Management Interface (VAMI) or API. It streams a set of files to a designated storage device using SCP, HTTP, HTTPs, FTP, or FTPS. This backup fully supports vCenter Server Appliance instances with both embedded and external Platform Services Controller instances.

You can back up the VCSA using these steps:

1. Log on to the VAMI as root (https://*appliance-IP-address-or-FQDN*:5480)
2. Select **Summary** > **Backup**.
3. In the wizard, select the protocol (SCP, HTTP, HTTPs, FTP, or FTPS) and then provide the target server address and credentials.

4. Optionally, select **Encrypt Backup Data**, include historical data (stats, events, and tasks), and provide a description. Click **Next**.
5. Review the backup summary information and click **OK** to start the backup.

NOTE The backup operation for a vCenter High Availability cluster backs up only the active node.

NOTE If the non-active node in a vCenter HA cluster needs to be restored, you should rebuild the vCenter HA cluster from the active node instead of restoring from the VCSA backup. Likewise, if a PSC needs to be restored, you should simply deploy a new PSC to the SSO domain, if at least one healthy PSC is available in the domain. If the last remaining PSC in a SSO domain fails, then restore the PSC from backup.

You can use a two-stage process to perform a restore operation using the same installer from which the vCenter Server Appliance or Platform Services Controller instance was originally deployed or upgraded. In the first stage, launch the installer and choose **Restore**. Identify the backup type (SCP, HTTP, HTTPs, FTP, or FTPS) and the server from which to restore. Provide other details that are important when deploying a new VCSA appliance, such as size, password, and target vCenter Server and ESXi host. Click **Next** to continue to stage 2, where you start and monitor the data transfer to the new appliance. After completing the Stage 2 Wizard, you are redirected to the VCSA Getting Started page. If you restored a VCSA with an embedded PSC, then no post-restore recovery steps are needed. Otherwise, you need to run the following script in the VCSA shell. If you are restoring a VCSA with an external PSC, then run this command once in the VCSA shell. If you are restoring a PSC, then run this command once in the shell of each VCSA that shares the domain.

```
/usr/bin/vcenter-restore
```

VCSA restoration includes restoring the vCenter Server UUID and all configuration settings.

During the backup, you can select a check box and provide a password to encrypt the backup files via symmetric key encryption. The same password then must be used to decrypt the backup set during a restore procedure. If the password is lost, there is no way to restore the backup because the password is not stored with reversible encryption.

Define Supported VCSA Backup Targets

Supported VCSA backup targets are SCP, HTTP, HTTPs, FTP, and FTPS, using servers that are configured with sufficient storage space. You should dedicate a separate folder on the server for each file-based backup.

Deploy VMware Data Protection Agents

VDP supports granular guest-level backup and recovery support for Microsoft Exchange Servers, SQL Servers, and SharePoint Servers using VDP agents installed in the guest OS.

NOTE If you start a VDP client or plug-in installer without Administrator privileges on a computer with User Access Control (UAC) enabled, the software does not install correctly.

You can install the VDP agents in a VM guest OS by following these steps:

1. Right-click the **Command Prompt** icon and select **Run as Administrator**.
2. Use the **cd** command to change the working directory to the location of the installation package.
3. Start the installer that corresponds to the supported application. For example, to start the installer for the agent for the 64-bit version of SQL, use this command:

```
msiexec /i VMwareVDPSQL-windows-x86_64-<version>.msi
```

4. Complete the steps in the wizard, including the step to provide information about the VDP server.

Manage Snapshots on Recovered VMs (vSphere Replication)

Replication preserves snapshot history. If a snapshot was created and replicated, you can recover to the application-consistent snapshot.

If you enabled multiple point-in-time instances when you configured replication for the virtual machine, vSphere Replication presents the retained instances as standard snapshots after a successful recovery. You can select one of these snapshots to revert the virtual machine. vSphere Replication does not preserve the memory state when you revert to a snapshot.

After performing a successful recovery on the target vCenter Server site, you can perform a failback. To do so, you log in to the target site and manually configure replication in the reverse direction—from the target site to the source site. The disks on the source site are used as replication seeds. Before doing this, you must unregister the virtual machine from the inventory on the source site.

Troubleshooting

Understand the VCSA Monitoring Tool

VMware improved the VCSA monitoring and configuration tools, including the VAMI, which now allows you to monitor the embedded vPostgres database as well as compute usage and network usage. In the VAMI, you can select **Database** to view details on the disk space usage and the amount of storage used by alarms, events, tasks, and stats. You can click **CPU and Memory** to view charts showing CPU and memory usage and trends at various time intervals, such as one day or one month. Likewise, you can select **Networking** to view network usage charts.

To view the overall health status of the VCSA, use the VAMI to select **Summary** and view the **Health Status** pane and **Overall Status** badge. Green indicates all components are healthy. Red indicates that one or more components may be in an unusable state. Yellow indicates that one or more components might become overloaded soon. Orange indicates that one or more components may be degraded. For details on yellow orange and red badges, view the **Health Messages** pane.

You can view the health status of the VCSA hardware and other components by accessing the VCSA shell and using commands such as **mem.health.get**, **storage.health.get**, **swap.health.get**, and **softwarepackages.health.get**.

You can use the **vimtop** command in the VCSA shell, which is similar to **esxtop** for ESXi hosts. It provides details on CPU and memory usage for the processes and services running in VCSA. Like **esxtop**, **vimtop** is interactive. You can press **o** on your keyboard to view network performance details and press **K** to view disk performance details.

You can use the vSphere Web Client to view the status of vCenter Server nodes and services. To do so, log on to the vSphere Web Client in the **System Configuration Administrators** group (or as the SSO administrator). Navigate to **Administration** > **System Configuration**. You can select **Services** and then select a specific service, such as the License Service, and view details of the service in the **Summary** tab.

In VCSA 6.5, a new service life cycle framework called vMon is used that unifies the watchdog services in vCenter Server 6.0 to simplify managing and monitoring vCenter Server services. These watchdog services include vmware-watchdog, Java Service Wrapper, Likewise Service Wrapper, and Windows Service Control Manager. In VCSA 6.5, vMon keeps track of service dependencies, which may be complex. Some features, such as vCenter HA, leverage vMon to help determine when to fail over to another node.

Troubleshooting KMS Connectivity

You can use the vSphere Web Client to add KMS to your vCenter Server system, which creates a KMS cluster when you are adding the first KMS instance. You can add KMS instances from the same vendor to the cluster. If your environment supports KMS solutions from multiple vendors, you can create multiple KMS clusters. To do this, your user account must include the **Cryptographic operations** > **Manage key servers** privilege. In the vSphere Web Client, select the vCenter Server and click **Configure** > **Key Management Servers** > **Add KMS**. Provide the cluster name, address, and credentials.

Connecting to a KMS by using only an IPv6 address is not supported.

vCenter Certification Authority

You can view the certificates known to the vCenter Certificate Authority (VMCA) to see whether active certificates are about to expire, to check on expired certificates, and to see the status of the root certificate. To do so, you can log on to the vSphere Web Client as a user in the **CAAdmins vCenter Single Sign-On** group and use the following steps:

1. In the **Home** menu, select **Administration**.
2. In the inventory pane, expand **Nodes** and select the node.
3. Click **Manage** > **Certificate Authority**.
4. Click the appropriate certificate type: Active Certificates, Revoked Certificates, Expired Certificates, or Root Certificates.
5. Select the specific certificate and click **Show Certificate Details**.

The PSC web interface allows you to perform the following tasks:

- Add and remove certificate store entries.

- View the VMCA instance associated with the PSC.
- View certificates generate by the VMCA.
- Renew existing certificates or replace certificates.

Enhanced Logging

In vSphere 6.5, logging is enhanced to support auditing rather than just troubleshooting. Logs coming from vCenter Server via Syslog are now enriched to clearly show “before” and “after” setting changes. The enhanced logging covers VMs and all vSphere changes, such as changes to vCenter Server roles and permissions, datastore-browsing functions (including downloading a VM), and actions such as creating and modifying vCenter Server clusters and hosts.

VMware Update Manager Changes

In vSphere 6.5, VMware Update Manager (VUM) is still the preferred means for keeping ESXi hosts up to date; however, VUM is now integrated into the VCSA. The integration eliminates the need for an additional VM, OS license, and database server (VUM leverages the embedded vPostgres instance in the appliance but uses a database schema separate from vCenter Server). VUM is enabled by default and is ready to use immediately.

Deploy and Consolidate

Auto Deploy

Auto Deploy can be configured and managed using a graphical user interface (GUI) in vSphere 6.5. The PowerCLI method is still available, but the GUI provides an easier-to-use option. For the Auto Deploy GUI to be visible in the vSphere Web Client, both the Image Builder and Auto Deploy services must be running when you’re logging in to vCenter Server.

The Image Builder feature in the GUI enables you to download ESXi images from the VMware public repository or to upload ZIP files containing ESXi images or drivers. You can customize the images by adding or removing components and optionally export images to ISO or ZIP for use elsewhere. You can compare two images to see how their contents differ.

You can use the Deployed Hosts tab to view hosts that are provisioned with Auto Deploy and to perform tests and remediations.

Host Profiles

In vSphere 6.5, Host Profiles functionality is improved in the vSphere Web Client, which now provides an easy-to-use search function and the ability to mark elements as favorites. You can now create a hierarchy of host profiles by leveraging the new capability to copy settings from one profile to many profiles. In vSphere 6.5, it is now possible to manage settings for groups of hosts via a CSV file.

Compliance checks are more informative in vSphere 6.5, displaying a detailed, side-by-side comparison of values from a host with the values in a profile. Remediation is improved due to pre-checks that

determine if maintenance mode is required and speedy parallel remediation for changes that do not require maintenance mode.

Availability

Proactive HA

Proactive HA integrates with select hardware partners to detect degraded components and evacuate VMs from affected vSphere hosts *before* an incident causes a service interruption.

Hardware partners offer a vCenter Server plug-in to provide the health status of the system memory, local storage, power supplies, cooling fans, and network adapters. As hardware components become degraded, Proactive HA determines which hosts are at risk and places them into a new state, called *Quarantine Mode*. While in Quarantine Mode, VMs are migrated to healthy hosts, as long as affinity or anti-affinity rules are not violated and there is no impact to VM performance. In addition, the affected hosts are avoided when new VMs are added to the cluster.

vSphere HA Orchestrated Restarts

Beginning in vSphere 6.5, you can now configure rules to orchestrate the VM start order in a cluster. If a host fails, HA will automatically attempt to restart the VMs in the specified order, even when they will be spread among multiple hosts in the cluster. To configure this, select **Configure** > **VM/Host Groups** to create at least two VM groups and select **Configure** > **VM/Host Rules** to create a VM-to-VM rule that first restarts VMs in one group and then restarts VMs in the other group. You could create multiple rules that work together. For example, you can create a rule to start VM-Group-A before VM-Group-B and create another rule to start VM-Group-B before VM-Group-C. In this example, the effective boot order is VM-Group-A, then VM-Group-B, and finally VM-Group-C.

Fault Tolerance

In vSphere 6.5, the integration between vSphere Fault Tolerance (FT) and vSphere DRS is improved to enable better placement decisions by ranking the hosts based on the available network bandwidth and by recommending the datastore in which to place the secondary VMDK files. Also, the supported network latency between the primary and secondary VMs has been greatly decreased.

In vSphere 6.5, multiple port groups can now be used to increase the overall bandwidth available for vSphere FT logging traffic. This is similar to the multi-NIC feature of vSphere vMotion.

HA Admission Control

In vSphere 6.5, the default admission control setting is changed to *Cluster Resource Percentage*, which reserves a percentage of the total available CPU and memory resources in the cluster. For simplicity, the percentage is now calculated automatically by defining the number of host failures to tolerate (FTT). The percentage is dynamically changed as hosts are added or removed from the cluster.

Another new enhancement is the Performance Degradation VMs Tolerate setting, which controls the amount of performance reduction that is tolerated after a failure. A value of 0% indicates that no performance degradation is tolerated.

Understand and Describe the Architecture of VCSA HA

In vSphere 6.5, VCSA provides a native high availability solution. The solution consists of active, passive, and witness nodes that are cloned from an existing vCenter Server instance. It includes a maintenance mode that prevents planned maintenance from triggering a failover. It uses a native PostgreSQL replication for the database and a separate, asynchronous file system replication for data outside of the database.

In most scenarios, you can deploy the VCSA HA by using the basic method, where the nodes run in a single cluster. This approach is simple and automatically creates the passive and witness nodes. When you're deploying VCSA in a DRS-enabled cluster, the basic method creates an anti-affinity rule and uses DRS to place the nodes on separate hosts.

The alternative is to deploy VCSA HA using the advanced method, where the nodes can be placed in separate clusters, separate vCenter Servers, or even separate data centers. To use this method, you must manually create the passive and witness nodes by cloning the source vCenter Server instance and then migrate the nodes to the proper location.

VCSA HA supports both embedded and external PSCs. An external PSC instance is required when there are multiple vCenter Server instances in an Enhanced Linked Mode configuration. When you're using an external PSC with VCSA HA, an external load balancer is required to provide high availability to the PSC instances.

Only the active node has an active management interface (public IP). The three nodes communicate over a private vCenter HA network. The active node runs the active vCenter Server instance, replicates data to the passive node, and communicates with the witness node. The passive node constantly receives updates from the active node and automatically takes the role of the active node during a failover. The witness provides a quorum to protect against split-brain situations.

VCSA HA provides failover when a node is lost or when key services fail. For example, the failure of a host running the active node results in a failover. A recovery time objective (RTO) of 5 minutes is expected.

VCSA HA requires ESXi 5.5 or later, small-size (or larger) VCSA 6.5, the vCenter HA network on separate subnet with less than 10ms network latency, and a single vCenter Server, Standard license. VMware recommends you run the VCSA nodes on separate ESXi hosts in a DRS cluster, which may be managed by a separate vCenter Server that is version 5.5 or later. You can place the nodes in VMFS, NFS, or VSAN datastores.

Enable and Configure VCSA HA

You can use this procedure to deploy VCSA HA in the basic configuration:

1. Deploy the first VCSA, which will become the active node.
2. On each ESXi host in the cluster where the VCSA nodes will run, add a second network (port group) for vCenter HA traffic.
3. In the vSphere Web Client, right-click the VCSA and select **vCenter HA Settings** > **Configure**.
4. Start the VCSA HA configuration, select **Basic**, and then supply IP addresses and other information for the passive and witness nodes.

5. Verify the clone operations successfully create the witness and passive nodes and verify the VCSA HA is operational.

You can use this procedure to deploy VCSA HA in the advanced configuration:

1. Deploy the first VCSA, which will become the active node.
2. On each ESXi host in the cluster where the VCSA nodes will run, add a second network (port group) for vCenter HA traffic.
3. In the vSphere Web Client, right-click the VCSA and select **vCenter HA Settings** > **Configure**.
4. Start the VCSA HA configuration, select **Basic**, and then supply IP addresses and other information for the passive and witness nodes.
5. Verify the clone operations successfully create the witness and passive nodes and verify the VCSA HA is operational.

After VCSA HA is deployed, you can perform various management tasks, such as setting up SNMP traps, configuring custom certificates, initiating a VCSA HA failover, backing up the active node, and rebooting all vCenter HA nodes.

VMs

Content Library Improvements

In vSphere 6.5, you can now mount an ISO directly from the Content Library, apply a guest OS customization specification during VM deployment, and update existing templates. The Content Library performance is improved. The new Optimized HTTP Sync option stores content a compressed format, which reduces the synchronization time. The Content Library leverages new features in vCenter Server 6.5, including VCSA HA and backup/restore.