View Architecture Planning

VMware Horizon 6
Version 6.2

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see http://www.vmware.com/support/pubs.

EN-001924-02

vmware®

You can find the most up-to-date technical documentation on the VMware Web site at:

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

View Architecture Planning

View Architecture Planning provides an introduction to VMware Horizon™ 6, including a description of its major features and deployment options and an overview of how the components are typically set up in a production environment.

This guide answers the following questions:

- Does the product solve the problems you need it to solve?
- Would it be feasible and cost-effective to implement this solution in your enterprise?

Not all features and capabilities of VMware Horizon 6 are available in all editions. For a comparison of feature sets in each edition, see http://www.vmware.com/files/pdf/products/horizon-view/VMware-Horizon-View-Pricing-Licensing-FAQ.pdf.

To help you protect your installation, this guide also provides a discussion of security features.

Intended Audience

This information is intended for IT decision makers, architects, administrators, and others who need to familiarize themselves with the components and capabilities of this product. With this information, architects and planners can determine whether View satisfies the requirements of their enterprise for efficiently and securely delivering Windows desktops and applications to their end users. The example architecture helps planners understand the hardware requirements and setup effort required for a large-scale deployment.

1 Introduction to View

With View, IT departments can run remote desktops and applications in the datacenter and deliver these desktops and applications to employees as a managed service. End users gain a familiar, personalized environment that they can access from any number of devices anywhere throughout the enterprise or from home. Administrators gain centralized control, efficiency, and security by having desktop data in the datacenter.

This chapter includes the following topics:

Advantages of Using View

When you manage enterprise desktops with View, the benefits include increased reliability, security, hardware independence, and convenience.

Reliability and Security

Desktops and applications can be centralized by integrating with VMware vSphere® and virtualizing server, storage, and networking resources. Placing desktop operating systems and applications on a server in the datacenter provides the following advantages:

- Access to data can easily be restricted. Sensitive data can be prevented from being copied onto a remote employee's home computer.
- RADIUS support provides flexibility when choosing among two-factor authentication vendors. Supported vendors include RSA SecureID, VASCO DIGIPASS, SMS Passcode, and SafeNet, among others.
- Integration with VMware Workspace™ Portal means that end users have on-demand access to remote desktops through the same Web-based application catalog they use to access SaaS, Web, and Windows applications. Inside a remote desktop, users can also use this custom app store to access applications.
- The ability to provision remote desktops with pre-created Active Directory accounts addresses the requirements of locked-down Active Directory environments that have read-only access policies.
- Data backups can be scheduled without considering when end users' systems might be turned off.

- Remote desktops and applications that are hosted in a datacenter experience little or no downtime. Virtual machines can reside on high-availability clusters of VMware servers.

Virtual desktops can also connect to back-end physical systems and Microsoft Remote Desktop Services (RDS) hosts.

Convenience

The unified management console is built for scalability so that even the largest View deployments can be efficiently managed from a single management interface. Wizards and dashboards enhance the workflow and facilitate drilling down to see details or change settings. Figure 1-1 provides an example of the browser-based user interface for View Administrator.

Figure 1-1. Administrative Console Showing the Dashboard View

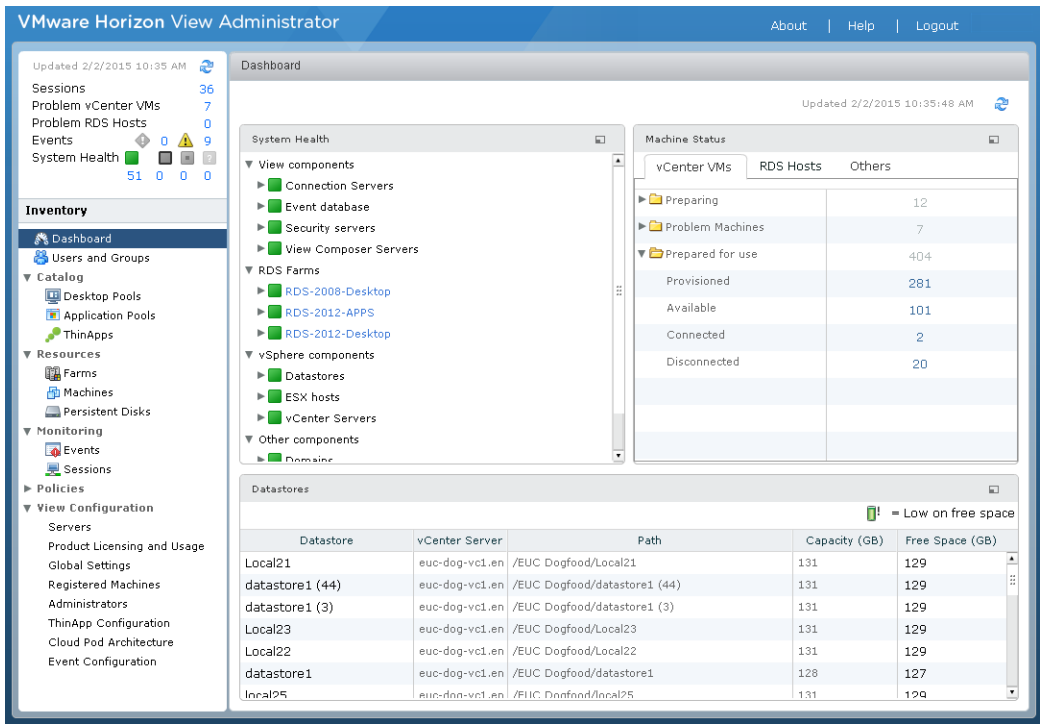

Another feature that increases convenience is the VMware remote display protocol, PCoIP. The PCoIP (PC-over-IP) display protocol delivers an end-user experience equal to the current experience of using a physical PC:

- On LANs, the display is faster and smoother than traditional remote displays.
- On WANs, the display protocol can compensate for an increase in latency or a reduction in bandwidth, ensuring that end users can remain productive regardless of network conditions.

Manageability

Provisioning desktops and applications for end users is a quick process. No one is required to install applications one by one on each end user's physical PC. End users connect to a remote application or a remote desktop complete with applications. End users can access their same remote desktop or application from various devices at various locations.

Using VMware vSphere to host virtual desktops and RDS host servers provides the following benefits:

- Administration tasks and management chores are reduced. Administrators can patch and upgrade applications and operating systems without touching a user's physical PC.

- Integration with Workspace Portal means that IT managers can use the Web-based Workspace Portal administration interface to monitor user and group entitlements to remote desktops.
- With View Persona Management, physical and virtual desktops can be centrally managed, including user profiles, application entitlement, policies, performance, and other settings. Deploy View Persona Management to physical desktop users prior to converting to virtual desktops.
- Storage management is simplified. Using VMware vSphere, you can virtualize volumes and file systems to avoid managing separate storage devices.
- With vSphere 6.0 or a later release, you can use Virtual Volumes (VVols). This feature maps virtual disks and their derivatives, clones, snapshots, and replicas, directly to objects, called virtual volumes, on a storage system. This mapping allows vSphere to offload intensive storage operations such as snapshoting, cloning, and replication to the storage system. The result, for example, is that a cloning operation that previously took an hour might now take just a few minutes using Virtual Volumes.
- With vSphere 5.5 Update 1 or a later release, you can use Virtual SAN, which virtualizes the local physical solid-state disks and hard disk drives available on ESXi™ hosts into a single datastore shared by all hosts in a cluster. You specify only one datastore when creating a desktop pool, and the various components, such as virtual machine files, replicas, user data, and operating system files, are placed on either SSD disks or hard drive disks, as appropriate.

  You manage virtual machine storage requirements, such as capacity, performance, and availability, in the form of default storage policy profiles, which get created automatically when you create a desktop pool.
- With the View storage accelerator, the IOPS storage load is dramatically reduced, supporting end-user logins at larger scales without requiring any special storage array technology.
- If remote desktops use the space-efficient disk format available with vSphere 5.1 and later, stale or deleted data within a guest operating system is automatically reclaimed with a wipe and shrink process.

Hardware Independence

Remote desktops and applications are hardware-independent. For example, because a remote desktop runs on a server in the datacenter and is only accessed from a client device, a remote desktop can use an operating system that might not be compatible with the hardware of the client device.

For example, although Windows 8 can run only on Windows 8-enabled devices, you can install Windows 8 in a virtual machine and use that virtual machine on a PC that is not Windows 8-enabled.

Remote desktops run on PCs, Macs, thin clients, PCs that have been repurposed as thin clients, tablets, and phones. Remote applications run on a subset of these devices. New device support is added quarterly.

If you use the HTML Access feature, end users can open a remote desktop inside a browser, without having to install any client application on the client system or device.

View Features

Features included in View support usability, security, centralized control, and scalability.

The following features provide a familiar experience for the end user:

- On Microsoft Windows client devices, print from a virtual desktop to any local or networked printer that is defined on the Windows client device. This virtual printer feature solves compatibility issues and does not require you to install additional print drivers in a virtual machine.
- On most client devices, use the location-based printing feature to map to printers that are physically near the client system. Location-based printing does require that you install print drivers in the virtual machine.

- Use multiple monitors. With PCoIP multiple-monitor support, you can adjust the display resolution and rotation separately for each monitor.
- Access USB devices and other peripherals that are connected to the local device that displays your virtual desktop.

  You can specify which types of USB devices end users are allowed to connect to. For composite devices that contain multiple types of devices, such as a video input device and a storage device, you can split the device so that one device (for example, the video input device) is allowed but the other device (for example, the storage device) is not.
- Use View Persona Management to retain user settings and data between sessions even after the desktop has been refreshed or recomposed. View Persona Management has the ability to replicate user profiles to a remote profile store (CIFS share) at configurable intervals.

  You can also use a standalone version of View Persona Management on physical computers and virtual machines that are not managed by View.

View offers the following security features, among others:

- Use two-factor authentication, such as RSA SecurID or RADIUS (Remote Authentication Dial-In User Service), or smart cards to log in.
- Use pre-created Active Directory accounts when provisioning remote desktops and applications in environments that have read-only access policies for Active Directory.
- Use SSL tunneling to ensure that all connections are completely encrypted.
- Use VMware High Availability to ensure automatic failover.

Scalability features depend on the VMware virtualization platform to manage both desktops and servers:

- Integrate with VMware vSphere to achieve cost-effective densities, high levels of availability, and advanced resource allocation control for your remote desktops and applications.
- Use the View storage accelerator feature to support end-user logins at larger scales with the same storage resources. This storage accelerator uses features in the vSphere 5 platform to create a host memory cache of common block reads.
- Configure View Connection Server to broker connections between end users and the remote desktops and applications that they are authorized to access.
- Use View Composer to quickly create desktop images that share virtual disks with a master image. Using linked clones in this way conserves disk space and simplifies the management of patches and updates to the operating system.

The following features provide centralized administration and management:

- Use Microsoft Active Directory to manage access to remote desktops and applications and to manage policies.
- Use View Persona Management to simplify and streamline migration from physical to virtual desktops.
- Use the Web-based administrative console to manage remote desktops and applications from any location.
- Use View Administrator to distribute and manage applications packaged with VMware ThinApp™.
- Use a template, or master image, to quickly create and provision pools of desktops.
- Send updates and patches to virtual desktops without affecting user settings, data, or preferences.
- Integrate with Workspace Portal so that end users can access remote desktops through the user portal on the Web, as well as use Workspace Portal from a browser inside a remote desktop.

- Integrate with Mirage™ and Horizon FLEX™ to manage locally installed virtual machine desktops and to deploy and update applications on dedicated full-clone remote desktops without overwriting user-installed applications.

How the Components Fit Together

End users start Horizon Client to log in to View Connection Server. This server, which integrates with Windows Active Directory, provides access to remote desktops hosted on a VMware vSphere server, a physical PC, or a Microsoft RDS host. Horizon Client also provides access to remote applications on a Microsoft RDS host.

NOTE View supports the following Active Directory Domain Services (AD DS) domain functional levels:

- Windows Server 2003
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2

View does not support Novell DSFW (Domain Services For Windows).

Figure 1-2 shows the relationship between the major components of a View deployment.

Figure 1-2. High-Level Example of a View Environment

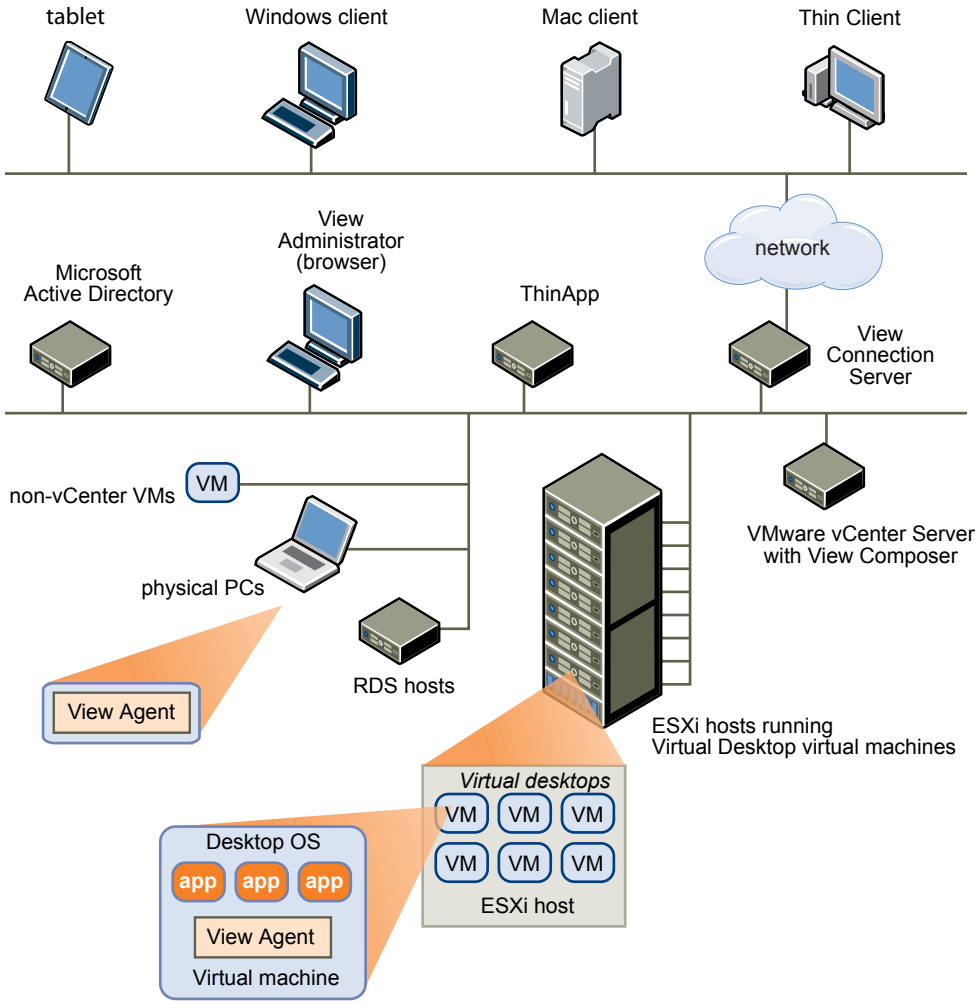

Client Devices

A major advantage of using View is that remote desktops and applications follow the end user regardless of device or location. Users can access their personalized virtual desktop or remote application from a company laptop, their home PC, a thin client device, a Mac, or a tablet or phone.

End users open Horizon Client to display their remote desktops and applications. Thin client devices use View thin client software and can be configured so that the only application that users can launch directly on the device is View Thin Client. Repurposing a legacy PC into a thin client desktop can extend the life of the hardware by three to five years. For example, by using View on a thin desktop, you can use a newer operating system such as Windows 8.x on older desktop hardware.

If you use the HTML Access feature, end users can open a remote desktop inside a browser, without having to install any client application on the client system or device.

View Connection Server

This software service acts as a broker for client connections. View Connection Server authenticates users through Windows Active Directory and directs the request to the appropriate virtual machine, physical PC, or Microsoft RDS host.

View Connection Server provides the following management capabilities:

- Authenticating users
- Entitling users to specific desktops and pools
- Assigning applications packaged with VMware ThinApp to specific desktops and pools
- Managing remote desktop and application sessions
- Establishing secure connections between users and remote desktops and applications
- Enabling single sign-on
- Setting and applying policies

Inside the corporate firewall, you install and configure a group of two or more View Connection Server instances. Their configuration data is stored in an embedded LDAP directory and is replicated among members of the group.

Outside the corporate firewall, in the DMZ, you can install and configure View Connection Server as a security server, or you can install an Access Point appliance. Security servers and Access Point appliances in the DMZ communicate with View Connection Servers inside the corporate firewall. Security servers and Access Point appliances ensure that the only remote desktop and application traffic that can enter the corporate data center is traffic on behalf of a strongly authenticated user. Users can access only the resources that they are authorized to access.

Security servers offer a subset of functionality and are not required to be in an Active Directory domain. You install View Connection Server in a Windows Server 2008 R2 or Windows Server 2012 R2 server, preferably on a VMware virtual machine. For more information about Access Point appliances, see *Deploying and Configuring Access Point*.

Important It is possible to create a View setup that does not use View Connection Server. If you install the View Agent Direct Connect Plugin in a remote virtual machine desktop, the client can connect directly to the virtual machine. All the remote desktop features, including PCoIP, HTML Access, RDP, USB redirection, and session management work in the same way, as if the user had connected through View Connection Server. For more information, see *View Agent Direct-Connection Plugin Administration*.

Horizon Client

The client software for accessing remote desktops and applications can run on a tablet, a phone, a Windows, Linux, or Mac PC or laptop, a thin client, and more.

After logging in, users select from a list of remote desktops and applications that they are authorized to use. Authorization can require Active Directory credentials, a UPN, a smart card PIN, or an RSA SecurID or other two-factor authentication token.

An administrator can configure Horizon Client to allow end users to select a display protocol. Protocols include PCoIP and Microsoft RDP for remote desktops. The speed and display quality of PCoIP rival that of a physical PC.

Features differ according to which Horizon Client you use. This guide focuses on Horizon Client for Windows. The following types of clients are not described in detail in this guide:

- Details about Horizon Client for tablets, Linux clients, and Mac clients. See the Horizon Client documentation at https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.

- Details about the HTML Access Web client, which allows you to open a remote desktop inside a browser. No Horizon Client application is installed on the client system or device. See the Horizon Client documentation at https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.
- Various third-party thin clients and zero clients, available only through certified partners.
- View Open Client, which supports the VMware partner certification program. View Open Client is not an official client application and is not supported as such.

VMware Horizon User Web Portal

From a Web browser on a client device, end users can connect to remote desktops and applications through the browser, automatically start Horizon Client if it is installed, or download the Horizon Client installer.

When you open a browser and enter the URL of a View Connection Server instance, the Web page that appears contains links to the VMware Downloads site for downloading Horizon Client. The links on the Web page are configurable, however. For example, you can configure the links to point to an internal Web server, or you can limit which client versions are available on your own View Connection Server.

If you use the HTML Access feature, the Web page also displays a link for accessing remote desktops inside a supported browser. With this feature, no Horizon Client application is installed on the client system or device. For more information, see the Horizon Client documentation at https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.

View Agent

You install the View Agent service on all virtual machines, physical systems, and Microsoft RDS hosts that you use as sources for remote desktops and applications. On virtual machines, this agent communicates with Horizon Client to provide features such as connection monitoring, virtual printing, View Persona Management, and access to locally connected USB devices.

If the desktop source is a virtual machine, you first install the View Agent service on that virtual machine and then use the virtual machine as a template or as a parent of linked clones. When you create a pool from this virtual machine, the agent is automatically installed on every remote desktop.

You can install the agent with an option for single sign-on. With single sign-on, users are prompted to log in only when they connect to View Connection Server and are not prompted a second time to connect to a remote desktop or application.

View Administrator

This Web-based application allows administrators to configure View Connection Server, deploy and manage remote desktops and applications, control user authentication, and troubleshoot end user issues.

When you install a View Connection Server instance, the View Administrator application is also installed. This application allows administrators to manage View Connection Server instances from anywhere without having to install an application on their local computer.

View Composer

You can install this software service on a vCenter Server instance that manages virtual machines or on a separate server. View Composer can then create a pool of linked clones from a specified parent virtual machine. This strategy reduces storage costs by up to 90 percent.

Each linked clone acts like an independent desktop, with a unique host name and IP address, yet the linked clone requires significantly less storage because it shares a base image with the parent. Because linked-clone desktop pools share a base image, you can quickly deploy updates and patches by updating only the parent virtual machine. End users' settings, data, and applications are not affected.

You can also use View Composer to create automated farms of linked-clone Microsoft RDS hosts, which provide hosted applications to end users.

Although you can install View Composer on its own server host, a View Composer service can operate with only one vCenter Server instance. Similarly, a vCenter Server instance can be associated with only one View Composer service.

vCenter Server

This service acts as a central administrator for VMware ESXi servers that are connected on a network. vCenter Server provides the central point for configuring, provisioning, and managing virtual machines in the datacenter.

In addition to using these virtual machines as sources for virtual machine desktop pools, you can use virtual machines to host the server components of View, including View Connection Server instances, Active Directory servers, Microsoft RDS hosts, and vCenter Server instances.

You can install View Composer on the same server as vCenter Server or on a different server. vCenter Server then manages the assignment of the virtual machines to physical servers and storage and manages the assignment of CPU and memory resources to virtual machines.

You can install vCenter Server either as a VMware virtual appliance or install vCenter Server in a Windows Server 2008 R2 server or a Windows Server 2012 R2 server, preferably on a VMware virtual machine.

Integrating and Customizing View

To enhance the effectiveness of View in your organization, you can use several interfaces to integrate View with external applications or to create administration scripts that you can run from the command line or in batch mode.

Integrating with Other Components

VMware Workspace Portal

You can integrate Workspace Portal with View to provide the following benefits to IT managers and end users:

- End users have on-demand access to remote desktops and applications through the same user portal on the Web that they use to access SaaS, Web, and Windows applications, with the same single sign-on convenience.
- End users can access Workspace Portal on the Web from inside a remote desktop for applications they need.
- If you also use HTML Access, end users can open a remote desktop inside a browser, without having to install any client application on the client system or device.
- IT managers can use the browser-based administration console of Workspace Portal to monitor user and group entitlements to remote desktops.

VMware Mirage and Horizon FLEX

You can use Mirage and Horizon FLEX to deploy and update applications on dedicated full-clone remote desktops without overwriting user-installed applications or data.

Mirage provides a better offline virtual desktop solution than the Local Mode feature that was previously included with View. Mirage includes the following security and management features for offline desktops:

- Encrypts the locally installed virtual machine and prevents a user from modifying virtual machine settings that affect the integrity of the secure container.
- Provides policies, including expiration, available in VMware Fusion™ Professional and VMware® Player Plus™, that are comparable to the polices provided with the previous Local Mode feature. Fusion Pro and Player Plus are included with Mirage.
- Eliminates the need for users to check in or check out their desktops to receive updates.
- Enables administrators to utilize the Mirage layering capability, backup features, and file portal.

VMware Horizon vRealize Orchestrator plug-in

The VMware Horizon vRealize Orchestrator plug-in allows interaction between vCenter Orchestrator and VMware Horizon 6. You can use this plug-in to expand the settings and methods for provisioning remote desktops and applications.

The plug-in contains a set of standard workflows that enable automation, self-service by request and approval, and scalable delegated administration across multi-tenant or highly distributed environments. You can also use these pre-defined workflows to create custom workflows.

Integrating with Popular Video Conferencing Software

Flash URL Redirection

Streaming Flash content directly from Adobe Media Server to client endpoints lowers the load on the datacenter ESXi host, removes the extra routing through the datacenter, and reduces the bandwidth required to simultaneously stream live video events to multiple client endpoints.

The Flash URL redirection feature uses a JavaScript that is embedded inside a Web page by the Web page administrator. Whenever a virtual desktop user clicks on the designated URL link from within a Web page, the JavaScript intercepts and redirects the ShockWave File (SWF) from the virtual desktop session to the client endpoint. The endpoint then opens a local VMware Flash Projector outside of the virtual desktop session and plays the media stream locally.

NOTE With Flash URL Redirection, the multicast or unicast stream is redirected to client devices that might be outside your organization's firewall. Your clients must have access to the Adobe Web server that hosts the ShockWave Flash (SWF) file that initiates the multicast or unicast streaming. If needed, configure your firewall to open the appropriate ports to allow client devices to access this server.

This feature is available only on some types of clients. To find out whether this feature is supported on a particular type of client, see the feature support matrix included in the "Using VMware Horizon Client" document for the specific type of desktop or mobile client device. Go to https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.

Microsoft Lync

You can use a Microsoft Lync 2013 client on remote desktops to participate in Unified Communications (UC) VoIP (voice over IP) and video chat calls with Lync certified USB audio and video devices. A dedicated IP phone is no longer required.

This architecture requires the installation of a Microsoft Lync 2013 client on the remote desktop and a Microsoft Lync VDI plug-in on the Windows 7 or 8 client endpoint. Customers can use the Microsoft Lync 2013 client for presence, instant messaging, Web conferencing, and Microsoft Office functionality.

Whenever a Lync VoIP or video chat call occurs, the Lync VDI plug-in offloads all the media processing from the datacenter server to the client endpoint, and encodes all media into Lync-optimized audio and video codecs. This optimized architecture is highly scalable, results in lower network bandwidth used, and provides point-to-point media delivery with support for high-quality real-time VoIP and video. For more information, see the white paper about VMware Horizon 6 and Microsoft Lync 2013, at http://www.vmware.com/files/pdf/techpaper/vmware-horizon-view-microsoft-lync-install-configure.pdf.

NOTE Recording audio is not yet supported. This integration is supported only with the PCoIP display protocol.

Integrating View with Business Intelligence Software

You can configure View Connection Server to record events to a Microsoft SQL Server or Oracle database.

- End-user actions such as logging in and starting a desktop session.
- Administrator actions such as adding entitlements and creating desktop pools.
- Alerts that report system failures and errors.
- Statistical sampling such as recording the maximum number of users over a 24-hour period.

You can use business intelligence reporting engines such as Crystal Reports, IBM Cognos, MicroStrategy 9, and Oracle Enterprise Performance Management System to access and analyze the event database.

For more information, see the *View Integration* document.

You can alternatively generate View events in Syslog format so that the event data can be accessible to analytics software. If you enable file-based logging of events, events are accumulated in a local log file. If you specify a file share, the log files are moved to that share. For more information, see the *View Installation* document.

Using View PowerCLI to Create Administration Scripts

Windows PowerShell is a command-line and scripting environment that is designed for Microsoft Windows. PowerShell uses the .NET object model and provides administrators with management and automation capabilities. As with any other console environment, you work with PowerShell by running commands, which are called cmdlets in PowerShell.

The View PowerCLI provides an easy-to-use PowerShell interface to View. You can use the View PowerCLI cmdlets to perform various administration tasks on View components.

- Create and update desktop pools.
- Configure multiple network labels to greatly expand the number of IP addresses assigned to virtual machines in a pool.
- Add datacenter resources to a full virtual machine or linked-clone pool.
- Perform rebalance, refresh, or recompose operations on linked-clone desktops.
- Sample the usage of specific desktops or desktop pools over time.
- Query the event database.
- Query the state of services.

You can use the cmdlets in conjunction with the vSphere PowerCLI cmdlets, which provide an administrative interface to the VMware vSphere product.

For more information, see the *View Integration* document.

Modifying LDAP Configuration Data in View

When you use View Administrator to modify the configuration of View, the appropriate LDAP data in the repository is updated. View Connection Server stores its configuration information in an LDAP compatible repository. For example, if you add a desktop pool, View Connection Server stores information about users, user groups, and entitlements in LDAP.

You can use VMware and Microsoft command-line tools to export and import LDAP configuration data in LDAP Data Interchange Format (LDIF) files from and into View. These commands are for advanced administrators who want to use scripts to update configuration data without using View Administrator or View PowerCLI.

You can use LDIF files to perform a number of tasks.

- Transfer configuration data between View Connection Server instances.
- Define a large number of View objects, such as desktop pools, and add these to your View Connection Server instances without using View Administrator or View PowerCLI.
- Back up a configuration so that you can restore the state of a View Connection Server instance.

For more information, see the *View Integration* document.

Using SCOM to Monitor View Components

You can use Microsoft System Center Operations Manager (SCOM) to monitor the state and performance of View components, including View Connection Server instances and security servers and the services running on these hosts.

For more information, see the *View Integration* document.

Using the vdmadmin Command

You can use the `vdmadmin` command line interface to perform a variety of administration tasks on a View Connection Server instance. You can use `vdmadmin` to perform administration tasks that are not possible from within the View Administrator user interface or that need to run automatically from scripts.

For more information, see the *View Administration* document.

Planning a Rich User Experience 2

View provides the familiar, personalized desktop environment that end users expect. For example, on some client systems, end users can access USB and other devices connected to their local computer, send documents to any printer that their local computer can detect, authenticate with smart cards, and use multiple display monitors.

View includes many features that you might want to make available to your end users. Before you decide which features to use, you must understand the limitations and restrictions of each feature.

This chapter includes the following topics:

Feature Support Matrix for View Agent

When planning which display protocol and features to make available to your end users, use the following information to determine which agent (remote desktop and application) operating systems support the feature.

The types and editions of the supported guest operating system depend on the Windows version.

Table 2-1. Operating Systems for Linked-Clone and Full-Clone Remote Desktops

Guest Operating System	Version	Edition	Service Pack
Windows 10	64-bit and 32-bit	Enterprise	None
Windows 8.1	64-bit and 32-bit	Enterprise and Professional	Latest update

Table 2-1. Operating Systems for Linked-Clone and Full-Clone Remote Desktops (Continued)

Guest Operating System	Version	Edition	Service Pack
Windows 8	64-bit and 32-bit	Enterprise and Professional	None
Windows 7	64-bit and 32-bit	Enterprise and Professional	SP1
Windows Server 2012 R2	64-bit	Datacenter	None
Windows Server 2008 R2	64-bit	Datacenter	SP1

Table 2-2. Operating Systems for RDS Hosts, Providing Remote Desktops or Applications

Guest Operating System	Edition	Service Pack
Windows Server 2008 R2	Standard, Enterprise, and Datacenter	SP1
Windows Server 2012	Standard and Datacenter	None
Windows Server 2012 R2	Standard and Datacenter	Latest update

Table 2-3. Features Supported on Windows Operating Systems Where View Agent Is Installed

Feature	Windows 7 Desktop	Windows 8.x Desktop	Windows 10 Desktop	Windows Server 2008/2012 R2 Desktop	Microsoft RDS-Hosted Desktops and Apps
USB redirection	X	X	X	X	USB flash drives and hard disks on Windows Server 2012 R2 RDS hosts
Client drive redirection	X	X	X	X	X
Real-Time Audio-Video (RTAV)	X	X	X	X	
Scanner redirection	X	X	X	X	X
Serial port redirection	X	X	X	X	
RDP display protocol	X	X	X	X	Session-based desktops only
PCoIP display protocol	X	X	X	X	X
Persona Management	X	X	X	X	
Wyse MMR					
Windows Media MMR	X	X			X
Location-based printing	X	X	X	X	X
Virtual printing	X	X	X	X	X
Smart cards	X	X	X	X	X
RSA SecurID or RADIUS	X	X	X	X	X

Table 2-3. Features Supported on Windows Operating Systems Where View Agent Is Installed (Continued)

Feature	Windows 7 Desktop	Windows 8.x Desktop	Windows 10 Desktop	Windows Server 2008/2012 R2 Desktop	Microsoft RDS-Hosted Desktops and Apps
Single sign-on	X	X	X	X	X
Multiple monitors	X	X	X	X	X

NOTE For information about which features are supported on the various types of client devices, see the Horizon Client documentation at https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.

In addition, several VMware partners offer thin and zero client devices for View deployments. The features that are available for each thin or zero client device are determined by the vendor and model and the configuration that an enterprise chooses to use. For information about the vendors and models for thin and zero client devices, see the VMware Compatibility Guide, available on the VMware Web site.

Choosing a Display Protocol

A display protocol provides end users with a graphical interface to a remote desktop or application that resides in the datacenter. Depending on which type of client device you have, you can choose between PCoIP (PC-over-IP), which VMware provides, or Microsoft RDP (Remote Desktop Protocol).

You can set policies to control which protocol is used or to allow end users to choose the protocol when they log in to a desktop.

NOTE For some types of clients, neither the PCoIP nor the RDP remote display protocol is used. For example, if you use the HTML Access client, available with the HTML Access feature, the Blast protocol is used, rather than PCoIP or RDP.

PCoIP

PCoIP (PC over IP) provides an optimized desktop experience for the delivery of a remote application or an entire remote desktop environment, including applications, images, audio, and video content for a wide range of users on the LAN or across the WAN. PCoIP can compensate for an increase in latency or a reduction in bandwidth, to ensure that end users can remain productive regardless of network conditions.

PCoIP is supported as the display protocol for remote applications and for remote desktops that use virtual machines, physical machines that contain Teradici host cards, or shared session desktops on an RDS host.

PCoIP Features

Key features of PCoIP include the following:

- Users outside the corporate firewall can use this protocol with your company's virtual private network (VPN), or users can make secure, encrypted connections to a security server or Access Point appliance in the corporate DMZ.
- Advanced Encryption Standard (AES) 128-bit encryption is supported and is turned on by default. You can, however, change the encryption key cipher to AES-192 or AES-256.
- Connections from all types of client devices.
- Optimization controls for reducing bandwidth usage on the LAN and WAN.
- 32-bit color is supported for virtual displays.
- ClearType fonts are supported.

- Audio redirection with dynamic audio quality adjustment for LAN and WAN.
- Real-Time Audio-Video for using webcams and microphones on some client types.
- Copy and paste of text and, on some clients, images between the client operating system and a remote application or desktop. For other client types, only copy and paste of plain text is supported. You cannot copy and paste system objects such as folders and files between systems.
- Multiple monitors are supported for some client types. On some clients, you can use up to 4 monitors with a resolution of up to 2560 x 1600 per display or up to 3 monitors with a resolution of 4K (3840 x 2160) for Windows 7 remote desktops with Aero disabled. Pivot display and autofit are also supported.

  When the 3D feature is enabled, up to 2 monitors are supported with a resolution of up to 1920 x 1200, or one monitor with a resolution of 4K (3840 x 2160).
- USB redirection is supported for some client types.
- MMR redirection is supported for some Windows client operating systems and some remote desktop operating systems (with View Agent-installed).

For information about which desktop operating systems support specific PCoIP features, see "Feature Support Matrix for View Agent," on page 19.

For information about which client devices support specific PCoIP features, go to https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.

Recommended Guest Operating System Settings

1GB of RAM or more and a dual CPU is recommended for playing in high-definition, full screen mode, or 720p or higher formatted video. To use Virtual Dedicated Graphics Acceleration for graphics-intensive applications such as CAD applications, 4GB of RAM is required.

Video Quality Requirements

480p-formatted video — You can play video at 480p or lower at native resolutions when the remote desktop has a single virtual CPU. If you want to play the video in high-definition Flash or in full screen mode, the desktop requires a dual virtual CPU. Even with a dual virtual CPU desktop, as low as 360p-formatted video played in full screen mode can lag behind audio, particularly on Windows clients.

720p-formatted video — You can play video at 720p at native resolutions if the remote desktop has a dual virtual CPU. Performance might be affected if you play videos at 720p in high definition or in full screen mode.

1080p-formatted video — If the remote desktop has a dual virtual CPU, you can play 1080p formatted video, although the media player might need to be adjusted to a smaller window size.

3D rendering — You can configure remote desktops to use software- or hardware-accelerated graphics. The software-accelerated graphics feature enables you to run DirectX 9 and OpenGL 2.1 applications without requiring a physical graphics processing unit (GPU). The hardware-accelerated graphics features enable virtual machines to either share the physical GPUs (graphical processing unit) on a vSphere host or dedicate a physical GPU to a single virtual machine desktop.

For 3D applications, up to 2 monitors are supported, and the maximum screen resolution is 1920 x 1200. The guest operating system on the remote desktops must be Windows 7 or later.

For more information about 3D features, see "Using 3D Graphics Applications," on page 27.

Hardware Requirements for Client Systems

For information about processor and memory requirements, see the "Using VMware Horizon Client" document for the specific type of desktop or mobile client device. Go to https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.

Microsoft RDP

Remote Desktop Protocol is the same multichannel protocol many people already use to access their work computer from their home computer. Microsoft Remote Desktop Connection (RDC) uses RDP to transmit data.

Microsoft RDP is a supported display protocol for remote desktops that use virtual machines, physical machines, or shared session desktops on an RDS host. (Only the PCoIP display protocol is supported for remote applications.) Microsoft RDP provides the following features:

- RDP 7 has true multiple monitor support, for up to 16 monitors.
- You can copy and paste text and system objects such as folders and files between the local system and the remote desktop.
- 32-bit color is supported for virtual displays.
- RDP supports 128-bit encryption.
- Users outside the corporate firewall can use this protocol with your company's virtual private network (VPN), or users can make secure, encrypted connections to a View security server in the corporate DMZ.

Hardware Requirements for Client Systems

For information about processor and memory requirements, see the "Using VMware Horizon Client" document for the specific type of client system. Go to https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.

NOTE Mobile client devices use only the PCoIP display protocol.

Using Hosted Applications

You can use Horizon Client to securely access remote Windows-based applications, in addition to remote desktops.

With this feature, after launching Horizon Client and logging in to a View server, users see all the remote applications they are entitled to use, in addition to remote desktops. Selecting an application opens a window for that application on the local client device, and the application looks and behaves as if it were locally installed.

For example, on a Windows client computer, if you minimize the application window, an item for that application remains in the Taskbar and looks identical to the way it would look if it were installed on the local Windows computer. You can also create a shortcut for the application that will appear on your client desktop, just like shortcuts for locally installed applications.

Deploying remote applications in this way might be preferable to deploying complete remote desktops under the following conditions:

- If an application is set up with a multi-tiered architecture, where the components work better if they are located geographically near each other, using remote, hosted applications is a good solution.

For example, when a user must access a database remotely, if large amounts of data must be transmitted over the WAN, performance is usually affected. With hosted applications, all parts of the application can be located in the same data center as the database, so that traffic is isolated and only the screen updates are sent across the WAN.

- From a mobile device, accessing an individual application is easier than opening a remote Windows desktop and then navigating to the application.

To use this feature, you install applications on a Microsoft RDS host. In this respect, View hosted applications work similarly to other application remoting solutions. View hosted applications are delivered using the PCoIP display protocol, for an optimized user experience.

Using View Persona Management to Retain User Data and Settings

You can use View Persona Management with remote desktops and with physical computers and virtual machines that are not managed by View. View Persona Management retains changes that users make to their profiles. User profiles comprise a variety of user-generated information.

- User-specific data and desktop settings, which allow the desktop appearance to be the same regard less of which desktop a user logs in to.
- Application data and settings. For example, these settings allow applications to remember toolbar positions and preferences.
- Windows registry entries configured by user applications.

To facilitate these abilities, View Persona Management requires storage on a CIFS share equal or greater than the size of the user's local profile.

Minimizing Logon and Logoff Times

View Persona Management minimizes the time it takes to log on to and off of desktops.

- View takes recent changes in the profile on the remote desktop and copies them to the remote repository at regular intervals. The default is every 10 minutes. In contrast, Windows roaming profiles wait until logoff time and copy all changes to the server at logoff.
- During logon, by default, View downloads only the files that Windows requires, such as user registry files. Other files are copied to the remote desktop when the user or an application opens them from the profile folder in the remote desktop.

  You can configure View to download specified files when the user logs in and download other files in the background.
- With View Persona Management, during logoff, only files that were updated since the last replication are copied to the remote repository.

With View Persona Management, you can avoid making any changes to Active Directory in order to have a managed profile. To configure Persona Management, you specify a central repository, without changing the user's properties in Active Directory. With this central repository, you can manage a user's profile in one environment without affecting the physical machines that users might also log on to.

With View Persona Management, if you provision desktops with VMware ThinApp applications, the ThinApp sandbox data can also be stored in the user profile. This data can roam with the user but does not significantly affect logon times. This strategy provides better protection against data loss or corruption.

Configuration Options

You can configure View personas at several levels: a single remote desktop, a desktop pool, an OU, or all remote desktops in your deployment. You can also use a standalone version of View Persona Management on physical computers and virtual machines that are not managed by View.

By setting group policies (GPOs), you have granular control of the files and folders to include in a persona:

- Specify whether to include the local settings folder. This policy affects the `AppData\Local` folder.
- Specify which files and folders to load at login time. For example: `Application Data\Microsoft\Certificates`. Within a folder, you can also specify files to exclude.
- Specify which files and folders to download in the background after a user logs in to the desktop. Within a folder, you can also specify files to exclude.
- Specify which files and folders within a user's persona to manage with Windows roaming profiles functionality instead of View Persona Management. Within a folder, you can also specify files to exclude.

As with Windows roaming profiles, you can configure folder redirection. You can redirect the following folders to a network share.

Contacts	My Documents	Save Games
Cookies	My Music	Searches
Desktop	My Pictures	Start Menu
Downloads	My Videos	Startup Items
Favorites	Network Neighborhood	Templates
History	Printer Neighborhood	Temporary Internet Files
Links	Recent Items	

To configure a remote repository to store personas, you can use either a network share or an existing Active Directory user profile path that you configured for Windows roaming profiles. The network share can be a folder on a server, a network-attached storage (NAS) device, or a network server. To support a large View deployment, you can configure separate repositories for different desktop pools.

You can install a standalone version of View Persona Management on physical computers and virtual machines that are not managed by View, allowing you to accomplish these goals:

- Share and manage profiles across standalone systems and remote desktops.
- Migrate user profiles from physical systems to remote desktops.
- Perform a staged migration from physical systems to remote desktops.
- Support up-to-date profiles when users go offline.

Limitations

View Persona Management has the following limitations and restrictions:

- You must have a View license that includes the View Personal Management component.
- View Persona Management requires a CIFS (Common Internet File System) share.

Using USB Devices with Remote Desktops and Applications

Administrators can configure the ability to use USB devices, such as thumb flash drives, cameras, VoIP (voice-over-IP) devices, and printers, from a remote desktop. This feature is called USB redirection, and it supports using either the RDP or the PCoIP display protocol. A remote desktop can accommodate up to 128 USB devices.

You can also redirect locally connected USB thumb flash drives and hard disks for use in RDS desktops and applications. Other types of USB devices, including other types of storage devices, are not supported in RDS desktops and applications.

When you use this feature in desktop pools that are deployed on single-user machines, most USB devices that are attached to the local client system become available in the remote desktop. You can even connect to and manage an iPad from a remote desktop. For example, you can sync your iPad with iTunes installed in your remote desktop. On some client devices, such as Windows and Mac OS X computers, the USB devices are listed in a menu in Horizon Client. You use the menu to connect and disconnect the devices.

In most cases, you cannot use a USB device in your client system and in your remote desktop or application at the same time. Only a few types of USB devices can be shared between a remote desktop and the local computer. These devices include smart card readers and human interface devices such as keyboards and pointing devices.

Administrators can specify which types of USB devices end users are allowed to connect to. For composite devices that contain multiple types of devices, such as a video input device and a storage device, on some client systems, administrators can split the device so that one device (for example, the video input device) is allowed but the other device (for example, the storage device) is not.

The USB redirection feature is available only on some types of clients. To find out whether this feature is supported on a particular type of client, see the feature support matrix included in the "Using VMware Horizon Client" document for the specific type of desktop or mobile client device. Go to https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.

Using the Real-Time Audio-Video Feature for Webcams and Microphones

With the Real-Time Audio-Video feature, you can use your local computer's webcam or microphone on your remote desktop. Real-Time Audio-Video is compatible with standard conferencing applications and browser-based video applications, and supports standard webcams, audio USB devices, and analog audio input.

End users can run Skype, Webex, Google Hangouts, and other online conferencing applications on their virtual desktops. This feature redirects video and audio data to the remote desktop with a significantly lower bandwidth than can be achieved by using USB redirection. With Real-Time Audio-Video, webcam images and audio input are encoded on the client and then sent to the remote desktop. On the remote desktop the stream is decoded and played by a virtual webcam and virtual microphone, which can be used by the third-party application.

No special configuration is necessary, although administrators can set GPOs and registry keys for the remote desktop to configure frame rate and image resolution, or to turn the feature off altogether. By default the resolution is 320 by 240 pixels at 15 frames per second. Administrators can also use client-side configuration settings to set a preferred webcam or audio device if needed.

NOTE This feature is available only on some types of clients. To find out whether this feature is supported on a particular type of client, see the feature support matrix included in the "Using VMware Horizon Client" document for the specific type of desktop or mobile client device. Go to https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.

Using 3D Graphics Applications

The software- and hardware-accelerated graphics features available with the PCoIP display protocol enable remote desktop users to run 3D applications ranging from Google Earth to CAD and other graphics-intensive applications.

NVIDIA GRID vGPU (shared GPU hardware acceleration)	Available with vSphere 6.0 and later, this feature allows a physical GPU (graphical processing unit) on an ESXi host to be shared among virtual machines. Use this feature if you require high-end, hardware-accelerated workstation graphics.
Virtual Dedicated Graphics Acceleration (vDGA)	Available with vSphere 5.5 and later, this feature dedicates a single physical GPU on an ESXi host to a single virtual machine. Use this feature if you require high-end, hardware-accelerated workstation graphics.
Virtual Shared Graphics Acceleration (vSGA)	Available with vSphere 5.1 and later, this feature allows multiple virtual machines to share the physical GPUs on ESXi hosts. You can use 3D applications for design, modeling, and multimedia.
Soft 3D	Software-accelerated graphics, available with vSphere 5.0 and later, allows you to run DirectX 9 and OpenGL 2.1 applications without requiring a physical GPU. Use this feature for less demanding 3D applications such as Windows Aero themes, Microsoft Office 2010, and Google Earth.

For these features, up to 2 monitors are supported, and the maximum screen resolution is 1920 x 1200.

With Horizon 6 version 6.2, NVIDIA GRID vGPU and vDGA are now also supported in remote applications running on Microsoft RDS hosts.

Important For more information on the various choices and requirements for 3D rendering, see the VMware white paper about graphics acceleration and the NVIDIA GRID vGPU Deployment Guide for VMware Horizon 6.1.

Streaming Multimedia to a Remote Desktop

The Windows Media MMR (multimedia redirection) feature, for Windows 7 and Windows 8/8.1 desktops and clients, enables full-fidelity playback on Windows client computers when multimedia files are streamed to a remote desktop.

With MMR, the multimedia stream is processed, that is, decoded, on the Windows client system. The client system plays the media content, thereby offloading the demand on the ESXi host. Media formats that are supported on Windows Media Player are supported; for example: M4V; MOV; MP4; WMP; MPEG-4 Part 2; WMV 7, 8, and 9; WMA; AVI; ACE; MP3; WAV.

Note You must add the MMR port as an exception to your firewall software. The default port for MMR is 9427.

Printing from a Remote Desktop

The virtual printing feature allows end users on some client systems to use local or network printers from a remote desktop without requiring that additional print drivers be installed in the remote desktop operating system. The location-based printing feature allows you to map remote desktops to the printer that is closest to the endpoint client device.

With virtual printing, after a printer is added on a local client computer, that printer is automatically added to the list of available printers on the remote desktop. No further configuration is required. For each printer available through this feature, you can set preferences for data compression, print quality, double-sided printing, color, and so on. Users who have administrator privileges can still install printer drivers on the remote desktop without creating a conflict with the virtual printing component.

To send print jobs to a USB printer, you can either use the USB redirection feature or use the virtual printing feature.

Location-based printing allows IT organizations to map remote desktops to the printer that is closest to the endpoint client device. For example, as a doctor moves from room to room in a hospital, each time the doctor prints a document, the print job is sent to the nearest printer. Using this feature does require that the correct printer drivers be installed in the remote desktop.

NOTE These printing features are available only on some types of clients. To find out whether a printing feature is supported on a particular type of client, see the feature support matrix included in the "Using VMware Horizon Client" document for the specific type of desktop or mobile client device. Go to https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.

Using Single Sign-On for Logging In to a Remote Desktop

The single-sign-on (SSO) feature allows end users to supply login credentials only once.

If you do not use the single-sign-on feature, end users must log in twice. They are first prompted to log in to View Connection Server and then are prompted log in to their remote desktop. If smart cards are also used, end users must sign in three times because users must also log in when the smart card reader prompts them for a PIN.

For remote desktops, this feature includes a credential provider dynamic-link library.

Using Multiple Monitors

Regardless of the display protocol, you can use multiple monitors with a remote desktop.

If you are using All Monitors display mode and click the Minimize button, if you then maximize the window, the window will go back to All Monitors mode. Similarly, if you are using Fullscreen mode and minimize the window, when you maximize the window, the window will go back to Fullscreen mode on one monitor.

Using All Monitors for Horizon Client

If you have Horizon Client use all monitors, if you maximize an application window, the window expands to the full screen of only the monitor that contains it.

Horizon Client supports the following monitor configurations:

- If you use 2 monitors, the monitors are not required to be in the same mode. For example, if you are using a laptop connected to an external monitor, the external monitor can be in portrait mode or landscape mode.

- If you use more than 2 monitors, the monitors must be in the same mode and have the same screen resolution. That is, if you use 3 monitors, all 3 monitors must be in either portrait mode or landscape mode and must use the same screen resolution.
- Monitors can be placed side by side, stacked 2 by 2, or vertically stacked only if you are using 2 monitors and the total height is less than 4096 pixels.
- To use the 3D rendering feature, you must use the PCoIP display protocol. You can use up to 2 monitors, with a resolution of up to 1920 X 1200. For a resolution of 4K (3840 X 2160), only one monitor is supported.
- With Horizon Client 3.4 or earlier and PCoIP, the maximum number of monitors that you can use to display a remote desktop is 4, with a resolution of up to 2560 X 1600 if you have enough video RAM.
- With Horizon Client 3.5 and PCoIP, a remote desktop screen resolution of 4K (3840 x 2160) is supported. The number of 4K displays that are supported depends on the hardware version of the desktop virtual machine and the Windows version.

Hardware Version	Windows Version	Number of 4K Displays Supported
10 (ESXi 5.5.x compatible)	7, 8, 8.x, 10	1
11 (ESXi 6.0 compatible)	7 (3D rendering feature disabled; Windows Aero disabled)	3
11	7 (3D rendering feature enabled)	1
11	8, 8.x, 10	1

NOTE When the remote desktop screen resolution is set to 3840 x 2160 (4K), items on the screen might appear smaller, and you might not be able to use the Screen Resolution dialog box in the remote desktop to make text and other items larger.

- If you use Microsoft RDP 7, the maximum number of monitors that you can use to display a remote desktop is 16.
- If you use Microsoft RDP display protocol, you must have Microsoft Remote Desktop Connection (RDC) 6.0 or higher installed in the remote desktop.

Using One Monitor in a Multiple-Monitor Setup

If you have multiple monitors but want Horizon Client to use only one of them, after client installation, you can select to have a desktop window launch in any mode other than All Monitors. By default, the window is launched on the primary monitor.

Managing Desktop and Application Pools from a Central Location

3

You can create pools that include one or hundreds or thousands of remote desktops. As a desktop source, you can use virtual machines, physical machines, and Windows Remote Desktop Services (RDS) hosts. Create one virtual machine as a base image, and View can generate a pool of remote desktops from that image. You can also create pools of applications that give users remote access to applications.

This chapter includes the following topics:

Advantages of Desktop Pools

View offers the ability to create and provision pools of desktops as its basis of centralized management.

You create a remote desktop pool from one of the following sources:

- A physical system such as a physical desktop PC or an RDS host
- A virtual machine that is hosted on an ESXi host and managed by vCenter Server
- A virtual machine that runs on a virtualization platform other than vCenter Server that supports View Agent

If you use a vSphere virtual machine as a desktop source, you can automate the process of making as many identical virtual desktops as you need. You can set a minimum and maximum number of virtual desktops to be generated for the pool. Setting these parameters ensures that you always have enough remote desktops available for immediate use but not so many that you overuse available resources.

Using pools to manage desktops allows you to apply settings or deploy applications to all remote desktops in a pool. The following examples show some of the settings available:

- Specify which remote display protocol to use as the default for the remote desktop and whether to let end users override the default.
- If using a virtual machine, specify whether to power off the virtual machine when it is not in use and whether to delete it altogether.
- Specify whether to use a Microsoft Sysprep customization specification or QuickPrep from VMware. Sysprep generates a unique SID and GUID for each virtual machine in the pool.

In addition, using desktop pools provides many conveniences.

Dedicated-assignment pools	Each user is assigned a particular remote desktop and returns to the same desktop at each login. Users can personalize their desktops, install applications, and store data.
Floating-assignment pools	The remote desktop is optionally deleted and re-created after each use, offering a highly controlled environment. A floating-assignment desktop is like a computer lab or kiosk environment where each desktop is loaded with the necessary applications and all desktops have access to necessary data. Using floating-assignment pools also allows you to create a pool of desktops that can be used by shifts of users. For example, a pool of 100 desktops could be used by 300 users if they worked in shifts of 100 users at a time.

Advantages of Application Pools

With application pools, you give users access to applications that run on servers in a data center instead of on their personal computers or devices.

Application pools offer several important benefits:

- Accessibility

  Users can access applications from anywhere on the network. You can also configure secure network access.

- Device independence

  With application pools, you can support a range of client devices, such as smart phones, tablets, laptops, thin clients, and personal computers. The client devices can run various operating systems, such as Windows, iOS, Mac OS, or Android.

- Access control

  You can easily and quickly grant or remove access to applications for one user or a group of users.

- Accelerated deployment

  With application pools, deploying applications can be accelerated because you only deploy applications on servers in a data center and each server can support multiple users.

- Manageability

  Managing software that is deployed on client computers and devices typically requires significant resources. Management tasks include deployment, configuration, maintenance, support, and upgrades. With application pools, you can simplify software management in an enterprise because the software runs on servers in a data center, which requires fewer installed copies.

- Security and regulatory compliance

  With application pools, you can improve security because applications and their associated data are centrally located in a data center. Centralized data can address security concerns and regulatory compliance issues.

- Reduced cost

  Depending on software license agreements, hosting applications in a data center can be more cost-effective. Other factors, including accelerated deployment and improved manageability, can also reduce the cost of software in an enterprise.

Reducing and Managing Storage Requirements

Deploying desktops on virtual machines that are managed by vCenter Server provides all the storage efficiencies that were previously available only for virtualized servers. Using View Composer increases the storage savings because all virtual machines in a pool share a virtual disk with a base image.

- Managing Storage with vSphere on page 33

  vSphere lets you virtualize disk volumes and file systems so that you can manage and configure storage without having to consider where the data is physically stored.

- Using Virtual SAN for High-Performance Storage and Policy-Based Management on page 34

  VMware Virtual SAN is a software-defined storage tier, available with vSphere 5.5 Update 1 or a later release, that virtualizes the local physical storage disks available on a cluster of vSphere hosts. You specify only one datastore when creating an automated desktop pool or an automated farm, and the various components, such as virtual machine files, replicas, user data, and operating system files, are placed on the appropriate solid-state drive (SSD) disks or direct-attached hard disks (HDDs).

- Using Virtual Volumes for Virtual-Machine-Centric Storage and Policy-Based Management on page 36

  With Virtual Volumes (VVols), available with vSphere 6.0 or a later release, an individual virtual machine, not the datastore, becomes a unit of storage management. The storage hardware gains control over virtual disk content, layout, and management.

- Reducing Storage Requirements with View Composer on page 37

  Because View Composer creates desktop images that share virtual disks with a base image, you can reduce the required storage capacity by 50 to 90 percent.

Managing Storage with vSphere

vSphere lets you virtualize disk volumes and file systems so that you can manage and configure storage without having to consider where the data is physically stored.

Fibre Channel SAN arrays, iSCSI SAN arrays, and NAS arrays are widely used storage technologies supported by vSphere to meet different datacenter storage needs. The storage arrays are connected to and shared between groups of servers through storage area networks. This arrangement allows aggregation of the storage resources and provides more flexibility in provisioning them to virtual machines.

Compatible vSphere 5.0 and 5.1 or Later Features

With vSphere 5.0 or a later release, you can use the following features:

- With the View storage accelerator feature, you can configure ESXi hosts to cache virtual machine disk data.

  Using this content-based read cache (CBRC) can reduce IOPS and improve performance during boot storms, when many machines start up and run anti-virus scans at the same time. Instead of reading the entire OS from the storage system over and over, a host can read common data blocks from cache.

- If remote desktops use the space-efficient disk format available with vSphere 5.1 and later, stale or deleted data within a guest operating system is automatically reclaimed with a wipe and shrink process.

- You can deploy a desktop pool on a cluster that contains up to 32 ESXi hosts, with certain restrictions.

  Replica disks must be stored on VMFS5 or later datastores or NFS datastores. If you store replicas on a VMFS version earlier than VMFS5, a cluster can have at most eight hosts. OS disks and persistent disks can be stored on NFS or VMFS datastores.

Compatible vSphere 5.5 Update 1 or Later Features

With vSphere 5.5 Update 1 or a later release, you can use Virtual SAN, which virtualizes the local physical solid-state disks and hard disk drives available on ESXi hosts into a single datastore shared by all hosts in a cluster. Virtual SAN provides high-performance storage with policy-based management, so that you specify only one datastore when creating a desktop pool, and the various components, such as virtual machine files, replicas, user data, and operating system files, are placed on the appropriate solid-state drive (SSD) disks or direct-attached hard disks (HDDs).

Virtual SAN also lets you manage virtual machine storage and performance by using storage policy profiles. If the policy becomes noncompliant because of a host, disk, or network failure, or workload changes, Virtual SAN reconfigures the data of the affected virtual machines and optimizes the use of resources across the cluster. You can deploy a desktop pool on a cluster that contains up to 20 ESXi hosts.

While supporting VMware features that require shared storage, such as HA, vMotion, and DRS, Virtual SAN eliminates the need for an external shared storage and simplifies storage configuration and virtual machine provisioning activities.

Important The Virtual SAN feature available with vSphere 6.0 and later releases contains many performance improvements over the feature that was available with vSphere 5.5 Update 1. With vSphere 6.0 this feature also has broader HCL (hardware compatibility) support. For more information about Virtual SAN in vSphere 6 or later, see the *Administering VMware Virtual SAN* document.

Note Virtual SAN is compatible with the View storage accelerator feature but not with the space-efficient disk format feature, which reclaims disk space by wiping and shrinking disks.

Compatible vSphere 6.0 or Later Features

With vSphere 6.0 or a later release, you can use Virtual Volumes (VVols). This feature maps virtual disks and their derivatives, clones, snapshots, and replicas, directly to objects, called virtual volumes, on a storage system. This mapping allows vSphere to offload intensive storage operations such as snapshoting, cloning, and replication to the storage system.

Virtual Volumes also lets you manage virtual machine storage and performance by using storage policy profiles in vSphere. These storage policy profiles dictate storage services on a per-virtual-machine basis. This type of granular provisioning increases capacity utilization. You can deploy a desktop pool on a cluster that contains up to 32 ESXi hosts.

Note Virtual Volumes is compatible with the View storage accelerator feature but not with the space-efficient disk format feature, which reclaims disk space by wiping and shrinking disks.

Using Virtual SAN for High-Performance Storage and Policy-Based Management

VMware Virtual SAN is a software-defined storage tier, available with vSphere 5.5 Update 1 or a later release, that virtualizes the local physical storage disks available on a cluster of vSphere hosts. You specify only one datastore when creating an automated desktop pool or an automated farm, and the various components, such as virtual machine files, replicas, user data, and operating system files, are placed on the appropriate solid-state drive (SSD) disks or direct-attached hard disks (HDDs).

Virtual SAN implements a policy-based approach to storage management. When you use Virtual SAN, View defines virtual machine storage requirements, such as capacity, performance, and availability, in the form of default storage policy profiles, which you can modify. Storage is provisioned and automatically configured according to the assigned policies. You can use Virtual SAN for linked-clone desktop pools, full-clone desktop pools, or an automated farm.

Each virtual machine maintains its policy regardless of its physical location in the cluster. If the policy becomes noncompliant because of a host, disk, or network failure, or workload changes, Virtual SAN reconfigures the data of the affected virtual machines and load-balances to meet the policies of each virtual machine.

While supporting VMware features that require shared storage, such as HA, vMotion, and DRS, Virtual SAN eliminates the need for an external shared storage infrastructure and simplifies storage configuration and virtual machine provisioning activities.

Important The Virtual SAN feature available with vSphere 6.0 and later releases contains many performance improvements over the feature that was available with vSphere 5.5 Update 1. With vSphere 6.0 this feature also has broader HCL (hardware compatibility) support. Also, VMware Virtual SAN 6.0 supports an all-flash architecture that uses flash-based devices for both caching and persistent storage.

Requirements and Limitations

The Virtual SAN feature has the following limitations when used in a View deployment:

- This release does not support using the View space-efficient disk format feature, which reclaims disk space by wiping and shrinking disks.
- Virtual SAN does not support the View Composer Array Integration (VAAI) feature because Virtual SAN does not use NAS devices.
- Virtual SAN datastores are not compatible with Virtual Volumes datastores for this release.

Note Virtual SAN is compatible with the View Storage Accelerator feature. Virtual SAN provides a caching layer on SSD disks, and the View Storage Accelerator feature provides a content-based cache that reduces IOPS and improves performance during boot storms.

The Virtual SAN feature has the following requirements:

- vSphere 5.5 Update 1 or a later release.
- Appropriate hardware. For example, VMware recommends a 10GB NIC and at least one SSD and one HDD for each capacity-contributing node. For specifics, see the VMware Compatibility Guide.
- A cluster of at least three ESXi hosts. You need enough ESXi hosts to accommodate your setup. For more information, see the *vSphere Configuration Maximums* document, available from https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html.
- SSD capacity that is at least 10 percent of HDD capacity.
- Enough HDDs to accommodate your setup. Do not exceed more than 75% utilization on a magnetic disk.

For more information about Virtual SAN requirements, see "Working with Virtual SAN" in the *vSphere 5.5 Update 1 Storage* document. For vSphere 6 or later, see the *Administering VMware Virtual SAN* document. For guidance on sizing and designing the key components of View virtual desktop infrastructures for VMware Virtual SAN, see the white paper at http://www.vmware.com/files/pdf/products/vsan/VMW-TMD-Virt-SAN-Dsn-Szing-Guid-Horizon-View.pdf.

Using Virtual Volumes for Virtual-Machine-Centric Storage and Policy-Based Management

With Virtual Volumes (VVols), available with vSphere 6.0 or a later release, an individual virtual machine, not the datastore, becomes a unit of storage management. The storage hardware gains control over virtual disk content, layout, and management.

With Virtual Volumes, abstract storage containers replace traditional storage volumes based on LUNs or NFS shares. Virtual Volumes maps virtual disks and their derivatives, clones, snapshots, and replicas, directly to objects, called virtual volumes, on a storage system. This mapping allows vSphere to offload intensive storage operations such as snapshoting, cloning, and replication to the storage system. The result, for example, is that a cloning operation that previously took an hour might now take just a few minutes using Virtual Volumes.

IMPORTANT Although one of the key benefits of Virtual Volumes is the ability to use Software Policy-Based Management (SPBM), for this release of View, no default granular storage policies are created by View, as they are when you use the Virtual SAN feature. Instead, you can set a global default storage policy in vCenter Server that will apply to all Virtual Volume datastores.

Virtual Volumes has the following benefits:

- Virtual Volumes supports offloading a number of operations to storage hardware. These operations include snapshotting, cloning, and Storage DRS.
- With Virtual Volumes, you can use advanced storage services that include replication, encryption, deduplication, and compression on individual virtual disks.
- Virtual Volumes supports such vSphere features as vMotion, Storage vMotion, snapshots, linked clones, Flash Read Cache, and DRS.
- You can use Virtual Volumes with storage arrays that support vSphere APIs for Array Integration (VAAI).

Requirements and Limitations

The Virtual Volumes feature has the following limitations when used in a View deployment:

- This release does not support using the View space-efficient disk format feature, which reclaims disk space by wiping and shrinking disks.
- Virtual Volumes does not support using View Composer Array Integration (VAAI).
- Virtual Volumes datastores are not compatible with Virtual SAN datastores for this release.

NOTE Virtual Volumes is compatible with the View Storage Accelerator feature. Virtual SAN provides a caching layer on SSD disks, and the View Storage Accelerator feature provides a content-based cache that reduces IOPS and improves performance during boot storms.

The Virtual Volumes feature has the following requirements:

- vSphere 6.0 or a later release.
- Appropriate hardware. Certain storage vendors are responsible for supplying storage providers that can integrate with vSphere and provide support for Virtual Volumes. Every storage provider must be certified by VMware and properly deployed.
- All virtual disks that you provision on a virtual datastore must be an even multiple of 1 MB.

Virtual Volumes is a vSphere 6.0 feature. For more information about the requirements, functionality, background, and setup requirements, see the topics about Virtual Volumes in the *vSphere Storage* document.

Reducing Storage Requirements with View Composer

Because View Composer creates desktop images that share virtual disks with a base image, you can reduce the required storage capacity by 50 to 90 percent.

View Composer uses a base image, or parent virtual machine, and creates a pool of up to 2,000 linked-clone virtual machines. Each linked clone acts like an independent desktop, with a unique host name and IP address, yet the linked clone requires significantly less storage.

Replica and Linked Clones on the Same Datastore

When you create a linked-clone desktop pool or farm of Microsoft RDS hosts, a full clone is first made from the parent virtual machine. The full clone, or replica, and the clones linked to it can be placed on the same data store, or LUN (logical unit number). If necessary, you can use the rebalance feature to move the replica and linked-clone desktop pools from one LUN to another or to move linked-clone desktop pools to a Virtual SAN datastore or from a Virtual SAN datastore to a LUN.

Replica and Linked Clones on Different Datastores

Alternatively, you can place View Composer replicas and linked clones on separate datastores with different performance characteristics. For example, you can store the replica virtual machines on a solid-state drive (SSD). Solid-state drives have low storage capacity and high read performance, typically supporting tens of thousands of I/Os per second (IOPS). You can store linked clones on traditional, spinning media-backed datastores. These disks provide lower performance, but are less expensive and provide higher storage capacity, which makes them suited for storing the many linked clones in a large pool. Tiered storage configurations can be used to cost-effectively handle intensive I/O scenarios such as simultaneous rebooting of many virtual machines or running scheduled antivirus scans.

For more information, see the best-practices guide called *Storage Considerations for VMware View.*

If you use Virtual SAN datastores or Virtual Volumes datastores, you cannot manually select different datastores for replicas and linked clones. Because the Virtual SAN and Virtual Volumes features automatically place objects on the appropriate type of disk and cache of all I/O operations, there is no need to use replica tiering for Virtual SAN and Virtual Volumes datastores.

Disposable Disks for Paging and Temp Files

When you create a linked-clone pool or farm, you can also optionally configure a separate, disposable virtual disk to store the guest operating system's paging and temp files that are generated during user sessions. When the virtual machine is powered off, the disposable disk is deleted. Using disposable disks can save storage space by slowing the growth of linked clones and reducing the space used by powered off virtual machines.

Persistent Disks for Dedicated Desktops

When you create dedicated-assignment desktop pools, View Composer can also optionally create a separate persistent virtual disk for each virtual desktop. The end user's Windows profile and application data are saved on the persistent disk. When a linked clone is refreshed, recomposed, or rebalanced, the contents of the persistent virtual disk are preserved. VMware recommends that you keep View Composer persistent disks on a separate datastore. You can then back up the whole LUN that holds persistent disks.

Local Datastores for Floating, Stateless Desktops

Linked-clone desktops can be stored on local datastores, which are internal spare disks on ESXi hosts. Local storage offers advantages such as inexpensive hardware, fast virtual-machine provisioning, high-performance power operations, and simple management. However, using local storage limits the vSphere infrastructure configuration options that are available to you. Using local storage is beneficial in certain environments but not appropriate in others.

NOTE The limitations described in this section do not apply to Virtual SAN datastores, which also use local storage disks but require specific hardware, as described in the preceding section about Virtual SAN.

Using local datastores is most likely to work well if the remote desktops in your environment are stateless. For example, you might use local datastores if you deploy stateless kiosks or classroom and training stations.

If you intend to take advantage of the benefits of local storage, you must carefully consider the following limitations:

- You cannot use VMotion, VMware High Availability (HA), or vSphere Distributed Resource Scheduler (DRS).
- You cannot use the View Composer rebalance operation to load-balance virtual machines across a resource pool.
- You cannot store a View Composer replica and linked clones on separate datastores, and, in fact, VMware recommends storing them on the same volume.

If you manage local disk usage by controlling the number of virtual machines and their disk growth, and if you use floating assignments and perform regular refresh and delete operations, you can successfully deploy linked clones to local datastores.

For more information, see the chapter about creating desktop pools in the *ViewAdministration* document.

Application Provisioning

With View, you have several options regarding application provisioning: You can use traditional application provisioning techniques, you can provide remote applications rather than a remote desktop, you can distribute application packages created with VMware ThinApp, or you can deploy applications as part of a View Composer base image.

- Deploying Individual Applications Using an RDS Host on page 39

  You might choose to provide end users with remote applications rather than remote desktops. Individual remote applications might be easier to navigate on a small mobile device.

- Deploying Applications and System Updates with View Composer on page 39

  Because linked-clone desktop pools share a base image, you can quickly deploy updates and patches by updating the parent virtual machine.

- Managing VMware ThinApp Applications in View Administrator on page 39

  VMware ThinApp™ lets you package an application into a single file that runs in a virtualized application sandbox. This strategy results in flexible, conflict-free application provisioning.

- Using Existing Processes or VMware Mirage for Application Provisioning on page 40

  With View, you can continue to use the application provisioning techniques that your company currently uses, and you can use Mirage. Two additional considerations include managing server CPU usage and storage I/O and determining whether users are permitted to install applications.

Deploying Individual Applications Using an RDS Host

You might choose to provide end users with remote applications rather than remote desktops. Individual remote applications might be easier to navigate on a small mobile device.

End users can access remote Windows-based applications by using the same Horizon Client that they previously used for accessing remote desktops, and they use the same PCoIP display protocol.

To provide a remote application, you install the application on a Microsoft Remote Desktop Session (RDS) host. One or more RDS hosts make up a farm, and from that farm administrators create application pools in a similar manner to creating desktop pools. A farm can contain up to 200 RDS hosts. A View pod can support up to 200 farms.

Using this strategy simplifies adding, removing, and updating applications; adding or removing user entitlements to applications; and providing access from any device or network to centrally or distributed application farms.

Deploying Applications and System Updates with View Composer

Because linked-clone desktop pools share a base image, you can quickly deploy updates and patches by updating the parent virtual machine.

The recompose feature allows you to make changes to the parent virtual machine, take a snapshot of the new state, and push the new version of the image to all, or a subset of, users and desktops. You can use this feature for the following tasks:

- Applying operating system and software patches and upgrades
- Applying service packs
- Adding applications
- Adding virtual devices
- Changing other virtual machine settings, such as available memory

NOTE Because you can also use View Composer to create farms of linked-clone Microsoft RDS hosts, the recompose feature lets you update the guest operating system and applications on RDS hosts.

You can create a View Composer persistent disk that contains user settings and other user-generated data. This persistent disk is not affected by a recompose operation. When a linked clone is deleted, you can preserve the user data. When an employee leaves the company, another employee can access the departing employee's user data. A user who has multiple desktops can consolidate the user data on a single desktop.

If you want to disallow users from adding or removing software or changing settings, you can use the refresh feature to bring the desktop back to its default values. This feature also reduces the size of linked clones, which tend to grow over time.

Managing VMware ThinApp Applications in View Administrator

VMware ThinApp™ lets you package an application into a single file that runs in a virtualized application sandbox. This strategy results in flexible, conflict-free application provisioning.

VMware ThinApp provides application virtualization by decoupling an application from the underlying operating system and its libraries and framework and bundling the application into a single executable file called an application package. You can use View Administrator to distribute VMware ThinApp applications to desktops and pools.

IMPORTANT If, instead of distributing ThinApps by assigning them to desktops and pools, you would rather assign ThinApps to Active Directory users and groups, you can use VMware Workspace Portal.

After you create a virtualized application with VMware ThinApp, you can choose to either stream the application from a shared file server or install the application on the virtual desktops. If you configure the virtualized application for streaming, you must address the following architectural considerations:

- Access for specific user groups to specific application repositories, where the application package is stored
- Storage configuration for the application repository
- Network traffic generated by streaming, which depends largely on the type of application

For streamed applications, users launch the applications by using a desktop shortcut.

If you assign a ThinApp package so that it is installed on a virtual desktop, the architectural considerations are similar to those that you address when you use traditional MSI-based software provisioning. Storage configuration for the application repository is a consideration both for streamed applications and for ThinApp packages installed in remote desktops.

Using Existing Processes or VMware Mirage for Application Provisioning

With View, you can continue to use the application provisioning techniques that your company currently uses, and you can use Mirage. Two additional considerations include managing server CPU usage and storage I/O and determining whether users are permitted to install applications.

If you push applications out to large numbers of remote desktops at exactly the same time, you might see significant spikes in CPU usage and storage I/O. These peak workloads can have noticeable effects on desktop performance. As a best practice, schedule application updates to occur during off-peak hours and stagger updates to desktops if possible. You must also verify that your storage solution is designed to support such workloads.

If your company allows users to install applications, you can continue your current policies, but you cannot take advantage of View Composer features such as refreshing and recomposing the desktop. With View Composer, if an application is not virtualized or otherwise included in the user's profile or data settings, that application is discarded whenever a View Composer refresh, recompose, or rebalance operation occurs. In many cases, this ability to tightly control which applications are installed is a benefit. View Composer desktops are easy to support because they are kept close to a known good configuration.

If users have firm requirements for installing their own applications and having those applications persist for the lifetime of the remote desktop, instead of using View Composer for application provisioning, VMware recommends that you create full-clone dedicated desktops, allow users to install applications, and then use Mirage to manage and update the desktops without overwriting user-installed applications.

Important Also use Mirage to manage locally installed offline desktops and their applications. For more information, see the Mirage Documentation page.

Using Active Directory GPOs to Manage Users and Desktops

View includes many Group Policy administrative (ADM) templates for centralizing the management and configuration of View components and remote desktops.

After you import these templates into Active Directory, you can use them to set policies that apply to the following groups and components:

- All systems regardless of which user logs in
- All users regardless of the system they log in to
- View Connection Server configuration
- Horizon Client configuration
- View Agent configuration

After a GPO is applied, properties are stored in the local Windows registry of the specified component.

You can use GPOs to set all the policies that are available from the View Administrator user interface (UI). You can also use GPOs to set policies that are not available from the UI. For a complete list and description of the settings available through ADM templates, see *Setting Up Desktop and Application Pools in View*.

Architecture Design Elements and Planning Guidelines for Remote Desktop Deployments

4

A typical View architecture design uses a pod strategy that consists of components that support up to 10,000 remote desktops using a vSphere 5.1 or later infrastructure. Pod definitions can vary, based on hardware configuration, View and vSphere software versions used, and other environment-specific design factors.

The examples in this document illustrate a scalable design that you can adapt to your enterprise environment and special requirements. This chapter includes key details about requirements for memory, CPU, storage capacity, network components, and hardware to give IT architects and planners a practical understanding of what is involved in deploying a View solution.

IMPORTANT This chapter does not cover the following topics:

Architecture design for hosted applications	A View pod can support up to 200 farms of Microsoft RDS hosts, and each farm can contain up to 200 RDS hosts. Supported operating systems for RDS hosts include Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2. For more information, see the *Setting Up Desktop and Application Pools in View*. If you plan to use virtual machines for RDS hosts, also see "RDS Host Virtual Machine Configuration," on page 53.
Architecture design for View Agent Direct Connect Plugin	With this plugin running on a remote virtual machine desktop, the client can connect directly to the virtual machine. All the remote desktop features, including PCoIP, HTML Access, RDP, USB redirection, and session management work in the same way, as if the user had connected through View Connection Server. For more information, see *View Agent Direct-Connection Plugin Administration*.

This chapter includes the following topics:

Virtual Machine Requirements for Remote Desktops

When you plan the specifications for remote desktops, the choices that you make regarding RAM, CPU, and disk space have a significant effect on your choices for server and storage hardware and expenditures.

- Planning Based on Types of Workers on page 44

  For many configuration elements, including RAM, CPU, and storage sizing, requirements depend largely on the type of worker who uses the virtual desktop and on the applications that must be installed.

- Estimating Memory Requirements for Virtual Machine Desktops on page 45

  RAM costs more for servers than it does for PCs. Because the cost of RAM is a high percentage of overall server hardware costs and total storage capacity needed, determining the correct memory allocation is crucial to planning your desktop deployment.

- Estimating CPU Requirements for Virtual Machine Desktops on page 47

  When estimating CPU, you must gather information about the average CPU utilization for various types of workers in your enterprise.

- Choosing the Appropriate System Disk Size on page 48

  When allocating disk space, provide only enough space for the operating system, applications, and additional content that users might install or generate. Usually this amount is smaller than the size of the disk that is included on a physical PC.

Planning Based on Types of Workers

For many configuration elements, including RAM, CPU, and storage sizing, requirements depend largely on the type of worker who uses the virtual desktop and on the applications that must be installed.

For architecture planning, workers can be categorized into several types.

Task workers	Task workers and administrative workers perform repetitive tasks within a small set of applications, usually at a stationary computer. The applications are usually not as CPU- and memory-intensive as the applications used by knowledge workers. Task workers who work specific shifts might all log in to their virtual desktops at the same time. Task workers include call center analysts, retail employees, warehouse workers, and so on.
Knowledge workers	Knowledge workers' daily tasks include accessing the Internet, using email, and creating complex documents, presentations, and spreadsheets. Knowledge workers include accountants, sales managers, marketing research analysts, and so on.
Power users	Power users include application developers and people who use graphics-intensive applications.
Kiosk users	These users need to share a desktop that is located in a public place. Examples of kiosk users include students using a shared computer in a classroom, nurses at nursing stations, and computers used for job placement and recruiting. These desktops require automatic login. Authentication can be done through certain applications if necessary.

Estimating Memory Requirements for Virtual Machine Desktops

RAM costs more for servers than it does for PCs. Because the cost of RAM is a high percentage of overall server hardware costs and total storage capacity needed, determining the correct memory allocation is crucial to planning your desktop deployment.

If the RAM allocation is too low, storage I/O can be negatively affected because too much Windows paging occurs. If the RAM allocation is too high, storage capacity can be negatively affected because the paging file in the guest operating system and the swap and suspend files for each virtual machine grow too large.

RAM Sizing Impact on Performance

When allocating RAM, avoid choosing an overly conservative setting. Take the following considerations into account:

- Insufficient RAM allocations can cause excessive Windows paging, which can generate I/O that causes significant performance degradations and increases storage I/O load.
- VMware ESXi supports sophisticated memory resource management algorithms such as transparent page sharing and memory ballooning, which can significantly reduce the physical RAM needed to support a given guest RAM allocation. For example, even though 2GB might be allocated to a virtual desktop, only a fraction of that number is consumed in physical RAM.
- Because virtual desktop performance is sensitive to response times, on the ESXi host, set nonzero values for RAM reservation settings. Reserving some RAM guarantees that idle but in-use desktops are never completely swapped out to disk. It can also reduce storage space consumed by ESXi swap files. However, higher reservation settings affect your ability to overcommit memory on an ESXi host and might affect VMotion maintenance operations.

RAM Sizing Impact on Storage

The amount of RAM that you allocate to a virtual machine is directly related to the size of the certain files that the virtual machine uses. To access the files in the following list, use the Windows guest operating system to locate the Windows page and hibernate files, and use the ESXi host's file system to locate the ESXi swap and suspend files.

Windows page file — By default, this file is sized at 150 percent of guest RAM. This file, which is by default located at `C:\pagefile.sys`, causes thin-provisioned storage to grow because it is accessed frequently. On linked-clone virtual machines, the page file and temporary files can be redirected to a separate virtual disk that is deleted when the virtual machines are powered off. Disposable page-file redirection saves storage, slowing the growth of linked clones and also can improve performance. Although you can adjust the size from within Windows, doing so might have a negative effect on application performance.

Windows hibernate file for laptops — This file can equal 100 percent of guest RAM. You can safely delete this file because it is not needed in View deployments.

ESXi swap file — This file, which has a `.vswp` extension, is created if you reserve less than 100 percent of a virtual machine's RAM. The size of the swap file is equal to the unreserved portion of guest RAM. For example, if 50 percent of guest RAM is reserved and guest RAM is 2GB, the ESXi swap file is 1GB. This file can be stored on the local data store on the ESXi host or cluster.

ESXi suspend file — This file, which has a `.vmss` extension, is created if you set the desktop pool logoff policy so that the virtual desktop is suspended when the end user logs off. The size of this file is equal to the size of guest RAM.

RAM Sizing for Specific Monitor Configurations When Using PCoIP

In addition to system memory, a virtual machine also requires a small amount of RAM on the ESXi host for video overhead. This VRAM size requirement depends in on the display resolution and number of monitors configured for end users. Table 4-1 lists the amount of overhead RAM required for various configurations. The amounts of memory listed in the columns are in addition to the amount of memory required for other PCoIP functionality.

Table 4-1. PCoIP Client Display Overhead

Display Resolution Standard	Width, in Pixels	Height, in Pixels	1-Monitor Overhead	2-Monitor Overhead	3-Monitor Overhead	4-Monitor Overhead
VGA	640	480	1.20MB	3.20MB	4.80MB	5.60MB
WXGA	1280	800	4.00MB	12.50MB	18.75MB	25.00MB
1080p	1920	1080	8.00MB	25.40MB	38.00MB	50.60MB
WQXGA	2560	1600	16.00MB	60.00MB	84.80MB	109.60MB
UHD (4K)	3840	2160	32.00MB	78.00MB	124.00MB	Not supported

For calculating system requirements, the VRAM values are in addition to the base system RAM for the virtual machine. Overhead memory is automatically calculated and configured when you specify the maximum number of monitors and select the display resolution in View Administrator.

If you use the 3D rendering feature and select Soft3D or vSGA, you can recalculate using the additional VRAM values in a View Administrator control for configuring VRAM for 3D guests. Alternatively, you can specify the exact amount of VRAM if you elect to manage VRAM by using vSphere Client. With the 3D rendering feature, you can select from the following options:

- The Soft3D (software-accelerated graphics) feature, available with vSphere 5.0 and later, allows you to use 3D applications such as Windows Aero themes or Google Earth. By default, the amount of VRAM set for this feature is 64MB. The maximum number of monitors is 2 and the maximum resolution is 1920 x 1200.
- The Virtual Shared Graphics Acceleration (vSGA) feature, available with vSphere 5.1 and later, allows multiple virtual machines to share the physical GPUs on the ESXi hosts. You can use 3D applications for design, modeling, and multimedia. By default, the amount of VRAM set for this feature is 96MB. The maximum number of monitors is 2 and the maximum resolution is 1920 x 1200.
- The Virtual Dedicated Graphics Acceleration (vDGA) feature, available with vSphere 5.5 and later, dedicates a single physical GPU (graphical processing unit) on an ESXi host to a single virtual machine. This feature provides high-end, hardware accelerated workstation graphics. When you create the virtual machine in vSphere, you are prompted to reserve all memory. For information about supported display resolutions, see the vendor's documentation. For a list of vendors, see the VMware Compatibility Guide.
- The NVIDIA GRID vGPU (shared GPU hardware acceleration) feature, available with vSphere 6.0 and later, allows a physical GPU on an ESXi host to be shared among virtual machines. This feature provides high-end, hardware accelerated workstation graphics. When you create the virtual machine in vSphere, you are prompted to reserve all memory. For information about supported display resolutions, see NVIDIA GRID vGPU Deployment Guide for VMware Horizon 6.1.

By default, the multiple-monitor configuration matches the host topology. There is extra overhead precalcuated for more than 2 monitors to accommodate additional topology schemes. If you encounter a black screen when starting a remote desktop session, verify that the values for the number of monitors and the display resolution, which are set in View Administrator, match the host system, or manually adjust the amount of memory by using selecting **Manage using vSphere Client** in View Administrator and then set the total video memory value to maximum of 128MB.

RAM Sizing for Specific Workloads and Operating Systems

Because the amount of RAM required can vary widely, depending on the type of worker, many companies conduct a pilot phase to determine the correct setting for various pools of workers in their enterprise.

A good starting point is to allocate 1GB for 32-bit Windows 7 or later desktops and 2GB for 64-bit Windows 7 or later desktops. If you want to use one of the hardware accelerated graphics features for 3D workloads, VMware recommends 2 virtual CPUs and 4GB of RAM. During a pilot, monitor the performance and disk space used with various types of workers and make adjustments until you find the optimal setting for each pool of workers.

Estimating CPU Requirements for Virtual Machine Desktops

When estimating CPU, you must gather information about the average CPU utilization for various types of workers in your enterprise.

CPU requirements vary by worker type. During your pilot phase, use a performance monitoring tool, such as Perfmon in the virtual machine, `esxtop` in ESXi, or vCenter Server performance monitoring tools, to understand both the average and peak CPU use levels for these groups of workers. Also use the following guidelines:

- Software developers or other power uses with high-performance needs might have much higher CPU requirements than knowledge workers and task workers. Dual or Quad virtual CPUs are recommended for 64-bit Windows 7 virtual machines running compute-intensive tasks such as using CAD applications, playing HD videos, or driving 4K display resolutions.
- Single virtual CPUs are generally recommended for other cases.

Because many virtual machines run on one server, CPU can spike if agents such as antivirus agents all check for updates at exactly the same time. Determine which agents and how many agents could cause performance issues and adopt a strategy for addressing these issues. For example, the following strategies might be helpful in your enterprise:

- Use View Composer to update images rather than having software management agents download software updates to each individual virtual desktop.
- Schedule antivirus and software updates to run at nonpeak hours, when few users are likely to be logged in.
- Stagger or randomize when updates occur.
- Use an antivirus product that is compatible with the VMware vShield API. For example, this API has been integrated into VMware vCloud® Networking and Security 5.1 and later.

As an informal initial sizing approach, to start, assume that each virtual machine requires 1/8 to 1/10 of a CPU core as the minimum guaranteed compute power. That is, plan a pilot that uses 8 to 10 virtual machines per core. For example, if you assume 8 virtual machines per core and have a 2-socket 8-core ESXi host, you can host 128 virtual machines on the server during the pilot. Monitor the overall CPU usage on the host during this period and ensure that it rarely exceeds a safety margin such as 80 percent to give enough headroom for spikes.

Choosing the Appropriate System Disk Size

When allocating disk space, provide only enough space for the operating system, applications, and additional content that users might install or generate. Usually this amount is smaller than the size of the disk that is included on a physical PC.

Because datacenter disk space usually costs more per gigabyte than desktop or laptop disk space in a traditional PC deployment, optimize the operating system image size. The following suggestions might help optimize image size:

- Remove unnecessary files. For example, reduce the quotas on temporary Internet files.
- Turn off Windows services such as the indexer service, the defragmenter service, and restore points. For details, see the topics "Optimize Windows Guest Operating System Performance," "Optimize Windows 7 and Windows 8 Guest Operating System Performance," and "Overview of Windows 7 and Windows 8 Services and Tasks That Cause Linked-Clone Growth," in *Setting Up Desktop and Application Pools in View*.
- Choose a virtual disk size that is sufficient to allow for future growth, but is not unrealistically large.
- Use centralized file shares or a View Composer persistent disk for user-generated content and user-installed applications.
- If you are using vSphere 5.1 or later, enable space reclamation for vCenter Server and for the linked-clone desktop pools.

  If virtual machine desktops use the space-efficient disk format available with vSphere 5.1 or later, stale or deleted data within a guest operating system is automatically reclaimed with a wipe and shrink process.

The amount of storage space required must take into account the following files for each virtual desktop:

- The ESXi suspend file is equivalent to the amount of RAM allocated to the virtual machine.
- By default, the Windows page file is equivalent to 150 percent of RAM.
- Log files can take up as much as 100MB for each virtual machine.
- The virtual disk, or `.vmdk` file, must accommodate the operating system, applications, and future applications and software updates. The virtual disk must also accommodate local user data and user-installed applications if they are located on the virtual desktop rather than on file shares.

  If you use View Composer, the `.vmdk` files grow over time, but you can control the amount of growth by scheduling View Composer refresh operations, setting a storage over-commit policy for virtual machine desktop pools, and redirecting Windows page and temporary files to a separate, nonpersistent disk.

You can also add 15 percent to this estimate to be sure that users do not run out of disk space.

View ESXi Node

A node is a single VMware ESXi host that hosts virtual machine desktops in a View deployment.

View is most cost-effective when you maximize the consolidation ratio, which is the number of desktops hosted on an ESXi host. Although many factors affect server selection, if you are optimizing strictly for acquisition price, you must find server configurations that have an appropriate balance of processing power and memory.

There is no substitute for measuring performance under actual, real world scenarios, such as in a pilot, to determine an appropriate consolidation ratio for your environment and hardware configuration. Consolidation ratios can vary significantly, based on usage patterns and environmental factors. Use the following guidelines:

- As a general framework, consider compute capacity in terms of 8 to 10 virtual desktops per CPU core. For information about calculating CPU requirements for each virtual machine, see "Estimating CPU Requirements for Virtual Machine Desktops," on page 47.
- Think of memory capacity in terms of virtual desktop RAM, host RAM, and overcommit ratio. Although you can have between 8 and 10 virtual desktops per CPU core, if virtual desktops have 1GB or more of RAM, you must also carefully consider physical RAM requirements. For information about calculating the amount of RAM required per virtual machine, see "Estimating Memory Requirements for Virtual Machine Desktops," on page 45.

  Note that physical RAM costs are not linear and that in some situations, it can be cost-effective to purchase more smaller servers that do not use expensive DIMM chips. In other cases, rack density, storage connectivity, manageability and other considerations can make minimizing the number of servers in a deployment a better choice.
- Note that in View 5.2 and later, the View Storage Accelerator feature is turned on by default, which allows ESXi 5.0 and later hosts to cache common virtual machine disk data. View Storage Accelerator can improve performance and reduce the need for extra storage I/O bandwidth to manage boot storms and anti-virus scanning I/O storms. This feature requires 1GB of RAM per ESXi host.
- Finally, consider cluster requirements and any failover requirements. For more information, see "Determining Requirements for High Availability," on page 57.

For information about specifications of ESXi hosts in vSphere, see the *VMware vSphere Configuration Maximums* document.

Desktop Pools for Specific Types of Workers

View provides many features to help you conserve storage and reduce the amount of processing power required for various use cases. Many of these features are available as pool settings.

The most fundamental question to consider is whether a certain type of user needs a stateful desktop image or a stateless desktop image. Users who need a stateful desktop image have data in the operating system image itself that must be preserved, maintained, and backed up. For example, these users install some of their own applications or have data that cannot be saved outside of the virtual machine itself, such as on a file server or in an application database.

Stateless desktop images — Stateless architectures have many advantages, such as being easier to support and having lower storage costs. Other benefits include a limited need to back up the linked-clone virtual machines and easier, less expensive disaster recovery and business continuity options.

Stateful desktop images — These images might require traditional image management techniques. Stateful images can have low storage costs in conjunction with certain storage system technologies. Backup and recovery technologies such as VMware Consolidated Backup and VMware Site Recovery Manager are important when considering strategies for backup, disaster recovery, and business continuity.

You create stateless desktop images by using View Composer and creating floating-assignment pools of linked-clone virtual machines.

You create stateful desktop images by creating dedicated-assignment pools of either linked-clone virtual machines or full virtual machines. If you use linked-clone virtual machines, you can configure View Composer persistent disks and folder redirection. Some storage vendors have cost-effective storage solutions for stateful desktop images. These vendors often have their own best practices and provisioning utilities. Using one of these vendors might require that you create a manual dedicated-assignment pool.

- Pools for Task Workers on page 50

  You can standardize on stateless desktop images for task workers so that the image is always in a well-known, easily supportable configuration and so that workers can log in to any available desktop.

- Pools for Knowledge Workers and Power Users on page 51

  Knowledge workers must be able to create complex documents and have them persist on the desktop. Power users must be able to install their own applications and have them persist. Depending on the nature and amount of personal data that must be retained, the desktop can be stateful or stateless.

- Pools for Kiosk Users on page 52

  Kiosk users might include customers at airline check-in stations, students in classrooms or libraries, medical personnel at medical data entry workstations, or customers at self-service points. Accounts associated with client devices rather than users are entitled to use these desktop pools because users do not need to log in to use the client device or the remote desktop. Users can still be required to provide authentication credentials for some applications.

Pools for Task Workers

You can standardize on stateless desktop images for task workers so that the image is always in a well-known, easily supportable configuration and so that workers can log in to any available desktop.

Because task workers perform repetitive tasks within a small set of applications, you can create stateless desktop images, which help conserve storage space and processing requirements. Use the following pool settings:

- Create an automated pool so that desktops can be created when the pool is created or can be generated on demand based on pool usage.
- Use floating assignment so that users log in to any available desktop. This setting reduces the number of desktops required if everyone does not need to be logged in at the same time.
- Create View Composer linked-clone desktops so that desktops share the same base image and use less storage space in the datacenter than full virtual machines.
- Determine what action, if any, to take when users log off. Disks grow over time. You can conserve disk space by refreshing the desktop to its original state when users log off. You can also set a schedule for periodically refreshing desktops. For example, you can schedule desktops to refresh daily, weekly, or monthly.
- If applicable, consider storing desktops on local ESXi datastores. This strategy can offer advantages such as inexpensive hardware, fast virtual-machine provisioning, high-performance power operations, and simple management. For a list of the limitations, see "Local Datastores for Floating, Stateless Desktops," on page 38.

NOTE For information about other types of storage options, see "Reducing and Managing Storage Requirements," on page 33.

- Use the Persona Management feature so that users always have their preferred desktop appearance and application settings, as with Windows user profiles. If you do not have the desktops set to be refreshed or deleted at logoff, you can configure the persona to be removed at logoff.

IMPORTANT View Persona Management facilitates implementing a floating-assignment pool for those users who want to retain settings between sessions. Previously, one of the limitations of floating-assignment desktops was that when end users logged off, they lost all their configuration settings and any data stored in the remote desktop.

Each time end users logged on, their desktop background was set to the default wallpaper, and they would have to configure each application's preferences again. With View Persona Management, an end user of a floating-assignment desktop cannot tell the difference between their session and a session on a dedicated-assignment desktop.

Pools for Knowledge Workers and Power Users

Knowledge workers must be able to create complex documents and have them persist on the desktop. Power users must be able to install their own applications and have them persist. Depending on the nature and amount of personal data that must be retained, the desktop can be stateful or stateless.

Because power users and knowledge workers, such as accountants, sales managers, marketing research analysts, must be able to create and retain documents and settings, you create dedicated-assignment desktops for them. For knowledge workers who do not need user-installed applications except for temporary use, you can create stateless desktop images and save all their personal data outside of the virtual machine, on a file server or in an application database. For other knowledge workers and for power users, you can create stateful desktop images. Use the following pool settings:

- Use dedicated assignment pools so that each knowledge worker or power user logs in to the same desktop every time.
- Use the Persona Management feature so that users always have their preferred desktop appearance and application settings, as with Windows user profiles.
- Use vStorage thin provisioning so that at first, each desktop uses only as much storage space as the disk needs for its initial operation.
- For power users and knowledge workers who must install their own applications, which adds data to the operating system disk, create full virtual machine desktops. Use Mirage to deploy and update applications without overwriting user-installed applications.
- If knowledge workers do not require user-installed applications except for temporary use, you can create View Composer linked-clone desktops. The desktop images share the same base image and use less storage space than full virtual machines.
- If you use View Composer with vSphere 5.1 or later virtual desktops, enable the space reclamation feature for vCenter Server and for the desktop pool. With the space reclamation feature, stale or deleted data within a guest operating system is automatically reclaimed with a wipe and shrink process.
- If you use View Composer linked-clone desktops, implement View Persona Management, roaming profiles, or another profile management solution.

  Configure persistent disks so that you can refresh and recompose the linked-clone OS disks while keeping a copy of the user profile on the persistent disks.

Pools for Kiosk Users

Kiosk users might include customers at airline check-in stations, students in classrooms or libraries, medical personnel at medical data entry workstations, or customers at self-service points. Accounts associated with client devices rather than users are entitled to use these desktop pools because users do not need to log in to use the client device or the remote desktop. Users can still be required to provide authentication credentials for some applications.

Virtual machine desktops that are set to run in kiosk mode use stateless desktop images because user data does not need to be preserved in the operating system disk. Kiosk mode desktops are used with thin client devices or locked-down PCs. You must ensure that the desktop application implements authentication mechanisms for secure transactions, that the physical network is secure against tampering and snooping, and that all devices connected to the network are trusted.

As a best practice, use dedicated View Connection Server instances to handle clients in kiosk mode, and create dedicated organizational units and groups in Active Directory for the accounts of these clients. This practice not only partitions these systems against unwarranted intrusion, but also makes it easier to configure and administer the clients.

To set up kiosk mode, you must use the `vdmadmin` command-line interface and perform several procedures documented in the topics about kiosk mode in the *View Administration* document. As part of this setup, you can use the following pool settings.

- Create an automated pool so that desktops can be created when the pool is created or can be generated on demand based on pool usage.
- Use floating assignment so that users can access any available desktop in the pool.
- Create View Composer linked-clone desktops so that desktops share the same base image and use less storage space in the datacenter than full virtual machines.
- Institute a refresh policy so that the desktop is refreshed frequently, such as at every user logoff.
- If applicable, consider storing desktops on local ESXi datastores. This strategy can offer advantages such as inexpensive hardware, fast virtual-machine provisioning, high-performance power operations, and simple management. For a list of the limitations, see "Local Datastores for Floating, Stateless Desktops," on page 38.

  NOTE For information about other types of storage options, see "Reducing and Managing Storage Requirements," on page 33.

- Use an Active Directory GPO (group policy object) to configure location-based printing, so that the desktop uses the nearest printer. For a complete list and description of the settings available through Group Policy administrative (ADM) templates, see *Setting Up Desktop and Application Pools in View*.
- Use a GPO if you want to override the default policy that enables connecting local USB devices to the desktop when the desktop is launched or when USB devices are plugged in to the client computer.

Desktop Virtual Machine Configuration

The example settings for items such as memory, number of virtual processors, and disk space are View-specific.

The amount of system disk space required depends on the number of applications required in the base image. VMware has validated a setup that included 8GB of disk space. Applications included Microsoft Word, Excel, PowerPoint, Adobe Reader, Internet Explorer, McAfee Antivirus, and PKZIP.

The amount of disk space required for user data depends on the role of the end user and organizational policies for data storage. If you use View Composer, this data is kept on a persistent disk.

The guidelines listed in the following table are for a standard Windows 7 or later virtual machine desktop.

Table 4-2. Desktop Virtual Machine Example for Windows 7 or Windows 8

Item	Example
Operating system	32-bit or 64-bit Windows 7 or later (with the latest service pack)
RAM	1GB (4GB if users must have hardware-accelerated graphics for 3D rendering)
Virtual CPU	1 (2 for 64-bit systems or if users must play high-definition or full screen video)
System disk capacity	24GB (slightly less than standard)
User data capacity (as a persistent disk)	5GB (starting point)
Virtual SCSI adapter type	LSI Logic SAS (the default)
Virtual network adapter	VMXNET 3

IMPORTANT Horizon 6 version 6.1 and later releases do not support Windows XP and Windows Vista desktops. View Agent 6.0.2 is the last View release that supports these guest operating systems. Customers who have an extended support agreement with Microsoft for Windows XP and Vista, and an extended support agreement with VMware for these guest operating systems, can deploy the View Agent 6.0.2 version of their Windows XP and Vista desktops with View Connection Server 6.1.

RDS Host Virtual Machine Configuration

Use RDS (Remote Desktop Services) hosts for providing hosted applications and session-based remote desktops to end users.

An RDS host can be a physical machine or a virtual machine. This example uses a virtual machine with the specifications listed in the following table. The ESXi host for this virtual machine can be part of a VMware HA cluster to guard against physical server failures.

Table 4-3. RDS Host Virtual Machine Example

Item	Example
Operating system	64-bit Windows Server 2008 R2 or Windows Server 2012 R2
RAM	24GB
Virtual CPU	4
System disk capacity	40GB
Virtual SCSI adapter type	LSI Logic SAS (the default for Windows Server 2008)
Virtual network adapter	VMXNET 3
1 NIC	1 Gigabit
Maximum number of client connections total (including session-based remote desktop connections and remote application connections)	50

For more information about RDS host configuration and tested workloads, see the *VMware Horizon 6 Reference Architecture* white paper at http://www.vmware.com/files/pdf/techpaper/VMware-Reference-Architecture-Horizon-6-View-Mirage-Workspace.pdf.

vCenter Server and View Composer Virtual Machine Configuration

You can install vCenter Server and View Composer on the same virtual machine or on separate servers. These servers require much more memory and processing power than a desktop virtual machine.

VMware tested having View Composer create and provision 2,000 desktops per pool using vSphere 5.1 or later. VMware also tested having View Composer perform a recompose operation on 2,000 desktops at a time. For these tests, vCenter Server and View Composer were installed on separate virtual machines.

Desktop pool size is limited by the following factors:

- Each desktop pool can contain only one vSphere cluster.
- With some setups, clusters can contain up to 32 hosts. With other setups, clusters are limited to 8 hosts. For more information, see "vSphere Clusters," on page 57.
- Each CPU core has compute capacity for 8 to 10 virtual desktops.
- The number of IP addresses available for the subnet limits the number of desktops in the pool. For example, if your network is set up so that the subnet for the pool contains only 256 usable IP addresses, the pool size is limited to 256 desktops. You can, however, configure multiple network labels to greatly expand the number of IP addresses assigned to virtual machines in a pool.

Although you can install vCenter Server and View Composer on a physical machine, this example uses separate virtual machines with the specifications listed in the following tables. The ESXi host for these virtual machines can be part of a VMware HA cluster to guard against physical server failures.

This example assumes that you are using View with vSphere 5.1 or later and vCenter Server 5.1 or later.

IMPORTANT This example also assumes that View Composer and vCenter Server are installed on separate virtual machines.

Table 4-4. vCenter Server Virtual Machine Example

Item	Example for a vCenter Server That Manages 10,000 Desktops	Example for a vCenter Server That Manages 2,000 Desktops
Operating system	64-bit Windows Server 2008 R2 Enterprise	64-bit Windows Server 2008 R2 Enterprise
RAM	48GB	4GB
Virtual CPU	16	2
System disk capacity	180GB	40GB
Virtual SCSI adapter type	LSI Logic SAS (the default for Windows Server 2008)	LSI Logic SAS (the default for Windows Server 2008)
Virtual network adapter	E1000 (the default)	VMXNET 3 (though E1000, the default, is fine too)
Maximum concurrent vCenter provisioning operations	20	20
Maximum concurrent power operations	50	50

Table 4-5. View Composer Virtual Machine Example

Item	Example for a View Composer That Manages 10,000 Desktops	Example for a View Composer That Manages 2,000 Desktops
Operating system	64-bit Windows Server 2008 R2 Enterprise	64-bit Windows Server 2008 R2 Enterprise
RAM	10 GB	4GB

Table 4-5. View Composer Virtual Machine Example (Continued)

Item	Example for a View Composer That Manages 10,000 Desktops	Example for a View Composer That Manages 2,000 Desktops
Virtual CPU	4	2
System disk capacity	50GB	40GB
Virtual SCSI adapter type	LSI Logic SAS (the default for Windows Server 2008)	LSI Logic SAS (the default for Windows Server 2008)
Virtual network adapter	VMXNET 3	VMXNET 3
Maximum View Composer pool size	2,000 desktops	1,000 desktops
Maximum concurrent View Composer maintenance operations	12	12
Maximum concurrent View Composer provisioning operations	8	8

IMPORTANT VMware recommends that you place the database to which vCenter Server and View Composer connect on a separate virtual machine.

View Connection Server Maximums and Virtual Machine Configuration

When you install View Connection Server, the View Administrator user interface is also installed.

View Connection Server Configuration

Although you can install View Connection Server on a physical machine, this example uses a virtual machine with the specifications listed in Table 4-6. The ESXi host for this virtual machine can be part of a VMware HA cluster to guard against physical server failures.

Table 4-6. Connection Server Virtual Machine Example

Item	Example
Operating system	64-bit Windows Server 2008 R2 or Windows Server 2012 R2
RAM	10GB
Virtual CPU	4
System disk capacity	70GB
Virtual SCSI adapter type	LSI Logic SAS (the default for Windows Server 2008)
Virtual network adapter	VMXNET 3
Network adapter	1Gbps NIC

View Connection Server Cluster Design Considerations

You can deploy multiple replicated View Connection Server instances in a group to support load balancing and high availability. Groups of replicated instances are designed to support clustering within a LAN-connected single-datacenter environment.

IMPORTANT To use a group of replicated View Connection Server instances across a WAN, MAN (metropolitan area network), or other non-LAN, in scenarios where a View deployment needs to span datacenters, you must use the Cloud Pod Architecture feature. You can link together four View pods to provide a single large desktop brokering and management environment for two geographically distant sites and manage up to 20,000 remote desktops. For more information, see *Administering View Cloud Pod Architecture*.

Maximum Connections for View Connection Server

Table 4-7 provides information about the tested limits regarding the number of simultaneous connections that a View deployment can accommodate.

This example assumes that View Connection Server is running on a 64-bit Windows Server 2008 R2 Enterprise operating system.

Table 4-7. Remote Desktop Connections

Connection Servers per Deployment	Connection Type	Maximum Simultaneous Connections
1 Connection Server	Direct connection, RDP or PCoIP: Tunneled connection, RDP: PCoIP Secure Gateway connection:	2,000 (tested limit) 2,000 (hard limit) 2,000 (hard limit)
7 Connection Servers	Direct connection, RDP or PCoIP	10,000 (tested, and therefore supported, limit)
1 Connection Server	Unified Access to physical PCs	2,000
1 Connection Server	Unified Access to RDS hosts	2,000
1 Connection Server	Blast Secure Gateway connections to remote desktops using HTML Access	2,000 (tested limit)

PCoIP Secure Gateway connections are required if you use security servers or Access Point appliances for PCoIP connections from outside the corporate network. Tunneled connections are required if you use security servers or Access Point appliances for RDP connections from outside the corporate network and for USB and multimedia redirection (MMR) acceleration with a PCoIP Secure Gateway connection. You can pair multiple security servers to a single View Connection Server instance.

Although the maximum number of simultaneous connections to security servers is also 2,000, instead of using just one security server per View Connection Server instance (with 2,000 sessions), you might choose to use 2 or 4. Monitoring of the security server might indicate that the PCoIP activity for 2,000 users is too great. The required amount of memory and CPU usage might dictate that you add more security servers per View Connection Server instance to spread the load. For example, you might use 2 security servers, with each one handling 1,000 connections, or you might use 4 security servers, with each one handling 500 connections. The ratio of security servers to View Connection Server instances depends on the requirements of the particular environment.

The number of connections per Access Point appliance is similar to those for security servers. For more information about Access Point appliances, see *Deploying and Configuring Access Point*.

NOTE In this example, although 5 View Connection Server instances could handle 10,000 connections, the number 7 is shown in the table for availability planning purposes, to accommodate connections coming from both inside and outside of the corporate network.

For example, if you had 10,000 users, with 8,000 of them inside the corporate network, you would need 5 View Connection Server instances inside the corporate network. That way, if one of the instances became unavailable, the 4 remaining instances could handle the load. Similarly, for the 2,000 connections coming from outside the corporate network, you would use 2 View Connection Server instances so that if one became unavailable, you would still have one instance left that could handle the load.

.

---

vSphere Clusters

View deployments can use VMware HA clusters to guard against physical server failures. Depending on your setup, clusters can contain up to 32 nodes.

vSphere and vCenter Server provide a rich set of features for managing clusters of servers that host virtual machine desktops. The cluster configuration is also important because each virtual machine desktop pool must be associated with a vCenter Server resource pool. Therefore, the maximum number of desktops per pool is related to the number of servers and virtual machines that you plan to run per cluster.

In very large View deployments, vCenter Server performance and responsiveness can be improved by having only one cluster object per datacenter object, which is not the default behavior. By default, vCenter Server creates new clusters within the same datacenter object.

Under the following conditions, vSphere clusters can contain up to 32 ESXi hosts, or nodes:

- vSphere 5.1 and later, with View Composer linked-clone pools, and store replica disks on NFS datastores or VMFS5 or later datastores
-
- vSphere 6.0 and later, and store pools on Virtual Volumes datastores

If you have vSphere 5.5 Update 1 and later, and store pools on Virtual SAN datastores, the vSphere clusters can contain up to 20 ESXi hosts.

If you store View Composer replicas on a VMFS version earlier than VMFS5, a cluster can have at most eight hosts. OS disks and persistent disks can be stored on NFS or VMFS datastores.

For more information, see the chapter about creating desktop pools, in the *Setting Up Desktop and Application Pools in View*. Networking requirements depend on the type of server, the number of network adapters, and the way in which VMotion is configured.

Determining Requirements for High Availability

vSphere, through its efficiency and resource management, lets you achieve industry-leading levels of virtual machines per server. But achieving a higher density of virtual machines per server means that more users are affected if a server fails.

Requirements for high availability can differ substantially based on the purpose of the desktop pool. For example, a stateless desktop image (floating-assignment) pool might have different recovery point objective (RPO) requirements than a stateful desktop image (dedicated-assignment) pool. For a floating-assignment pool, an acceptable solution might be to have users log in to a different desktop if the desktop they are using becomes unavailable.

In cases where availability requirements are high, proper configuration of VMware HA is essential. If you use VMware HA and are planning for a fixed number of desktops per server, run each server at a reduced capacity. If a server fails, the capacity of desktops per server is not exceeded when the desktops are restarted on a different host.

For example, in an 8-host cluster, where each host is capable of running 128 desktops, and the goal is to tolerate a single server failure, make sure that no more than 128 * (8 - 1) = 896 desktops are running on that cluster. You can also use VMware DRS (Distributed Resource Scheduler) to help balance the desktops among all 8 hosts. You get full use of the extra server capacity without letting any hot-spare resources sit idle. Additionally, DRS can help rebalance the cluster after a failed server is restored to service.

You must also make sure that storage is properly configured to support the I/O load that results from many virtual machines restarting at once in response to a server failure. Storage IOPS has the most effect on how quickly desktops recover from a server failure.

Example: Cluster Configuration Examples

The settings listed in the following tables are View-specific. For information about limits of HA clusters in vSphere, see the *VMware vSphere Configuration Maximums* document.

The following infrastructure example was tested with View 5.2 and vSphere 5.1.

Table 4-8. View Infrastructure Cluster Example

Item	Example
Virtual machines	vCenter Server instances, Active Directory, SQL database server, View Composer, View Connection Server instances, security servers, parent virtual machines to use as desktop pool sources
Nodes (ESXi hosts)	6 Dell PowerEdge R720 servers (16 cores * 2 GHz; and 192GB RAM on each host)
SSD storage	Virtual machines for vCenter Server, View Composer, SQL database server, and the parent virtual machines
Non-SSD storage	Virtual machines for Active Directory, View Connection Server, and security server
Cluster type	DRS (Distributed Resource Scheduler)/HA

Table 4-9. Virtual Machine Desktop Cluster Example

Item	Example
Number of clusters	5
Number of desktops and pools per cluster	1 pool of 2,000 desktops (virtual machines) per cluster
Nodes (ESXi hosts)	Following are examples of various servers that could be used for each cluster: ■ 12 Dell PowerEdge R720 (16 cores * 2 GHz; and 192GB RAM on each host) ■ 16 Dell PowerEdge R710 (12 cores * 2.526 GHz; and 144GB RAM on each host) ■ 8 Dell PowerEdge R810 (24 cores * 2 GHz; and 256GB RAM on each host) ■ 6 Dell PowerEdge R810 + 3 PowerEdge R720
SSD storage	Replica virtual machines
Non-SSD storage	32 Non-SSD datastores for clones (450 GB per datastore)
Cluster type	DRS (Distributed Resource Scheduler)/HA

Storage and Bandwidth Requirements

Several considerations go into planning for shared storage of virtual machine desktops, planning for storage bandwidth requirements with regard to I/O storms, and planning network bandwidth needs.

Details about the storage and networking components used in a test setup at VMware are provided in these related topics.

- Shared Storage Example on page 59

  For a View 5.2 test environment, View Composer replica virtual machines were placed on high-read-performance solid-state drives (SSD), which support tens of thousands of I/Os per second (IOPS). Linked clones were placed on traditional, lower-performance spinning media-backed datastores, which are less expensive and provide higher storage capacity.

- Storage Bandwidth Considerations on page 62

  In a View environment, logon storms are the main consideration when determining bandwidth requirements.

- Network Bandwidth Considerations on page 62

  Certain virtual and physical networking components are required to accommodate a typical workload.

- View Composer Performance Test Results on page 64

  These test results describe a View 5.2 setup with 10,000-desktops, in which one vCenter Server 5.1 instance managed 5 pools of 2,000 virtual machine desktops each. Only one maintenance period was required for provisioning a new pool or for recomposing, refreshing, or rebalancing an existing pool of 2,000 virtual machines. A logon storm of 10,000 users was also tested.

- WAN Support and PCoIP on page 66

  For wide-area networks (WANs), you must consider bandwidth constraints and latency issues. The PCoIP display protocol provided by VMware adapts to varying latency and bandwidth conditions.

Shared Storage Example

For a View 5.2 test environment, View Composer replica virtual machines were placed on high-read-performance solid-state drives (SSD), which support tens of thousands of I/Os per second (IOPS). Linked clones were placed on traditional, lower-performance spinning media-backed datastores, which are less expensive and provide higher storage capacity.

Storage design considerations are one of the most important elements of a successful View architecture. The decision that has the greatest architectural impact is whether to use View Composer desktops, which use linked-clone technology. The ESXi binaries, virtual machine swap files, and View Composer replicas of parent virtual machines are stored on the shared storage system.

The external storage system that vSphere uses can be a Fibre Channel or iSCSI SAN (storage area network), or an NFS (Network File System) NAS (network-attached storage). With the Virtual SAN feature, available with vSphere 5.5 Update 1 or later, the storage system can also be aggregated local server-attached storage.

The following example describes the tiered storage strategy used in a View 5.2 test setup in which one vCenter Server managed 10,000 desktops.

NOTE This example was used in a View 5.2 setup, which was carried out prior to the release of VMware Virtual SAN. For guidance on sizing and designing the key components of View virtual desktop infrastructures for VMware Virtual SAN, see the white paper at http://www.vmware.com/files/pdf/products/vsan/VMW-TMD-Virt-SAN-Dsn-Szing-Guid-Horizon-View.pdf.

The Virtual SAN feature available with vSphere 6.0 and later releases contains many performance improvements over the feature that was available with vSphere 5.5 Update 1. With vSphere 6.0 this feature also has broader HCL (hardware compatibility) support. For more information about Virtual SAN in vSphere 6 or later, see the *Administering VMware Virtual SAN* document.

Physical storage

- EMC VNX7500-block only
- 1.8TB Fast Cache (SSD)
- Eight 10Gbit FCoE front end connections (4 per controller).

SSD storage tier

A single RAID5 storage pool:

- 12 * 200GB EFD
- 250GB LUN for parent images
- 500GB LUN for infrastructure
- 75GB LUNs for replica stores (1 per desktop pool cluster)

Virtual machine desktop storage tier

Two RAID 1/0 storage pools:

For pool 1:

- 360 15K 300GB HDD (47TB usable)
- 97 450GB LUNs for desktops

For pool 2:

- 296 15K 300GB HDD (39TB usable)
- 7 450GB LUNs for infrastructure
- 85 450GB LUNs for desktops

This storage strategy is illustrated in the following figure.

Figure 4-1. Tiered Storage Example for a Large Desktop Pool

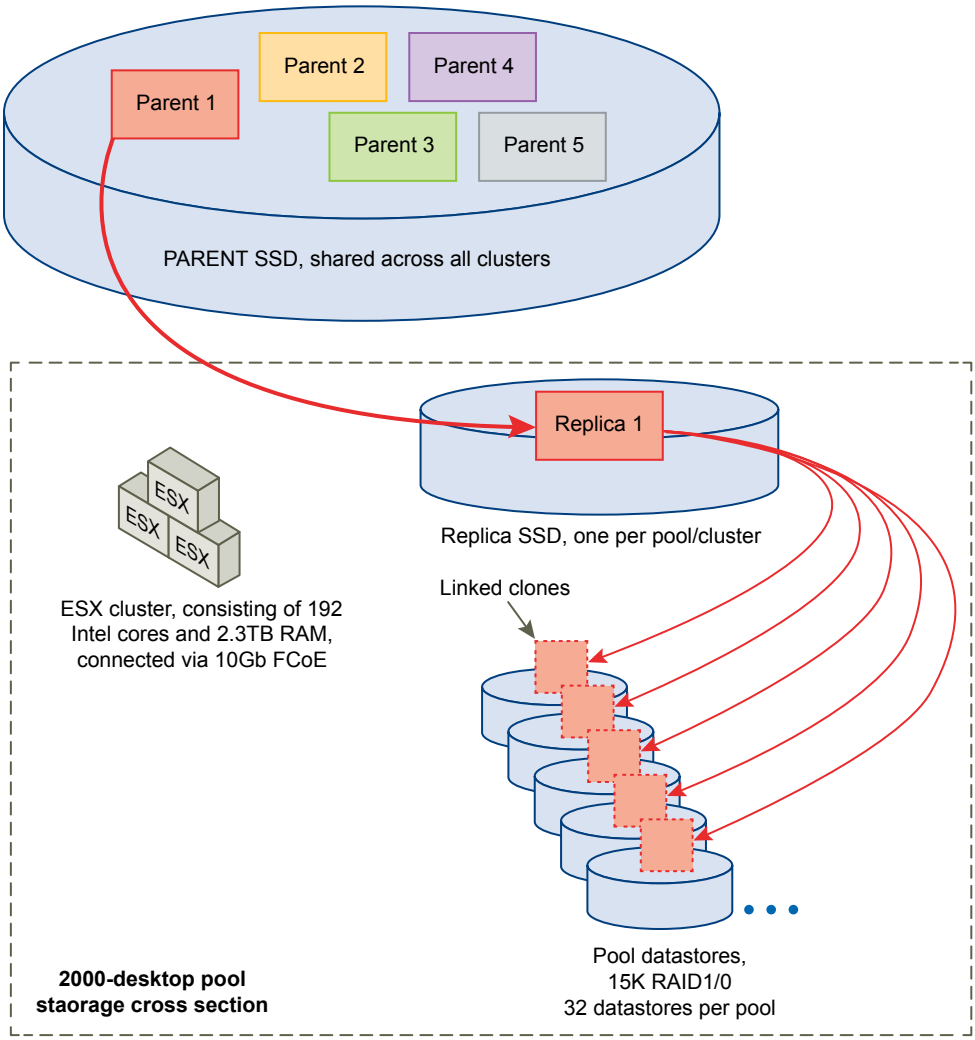

From an architectural perspective, View Composer creates desktop images that share a base image, which can reduce storage requirements by 50 percent or more. You can further reduce storage requirements by setting a refresh policy that periodically returns the desktop to its original state and reclaims space that is used to track changes since the last refresh operation.

If you use View Composer with vSphere 5.1 or later virtual machine desktops, you can use the space reclamation feature. With this feature, stale or deleted data within a guest operating system is automatically reclaimed with a wipe and shrink process when the amount of unused disk space reaches a certain threshold. Note that the space reclamation feature is not supported if you use a Virtual SAN datastore.

You can also reduce operating system disk space by using View Composer persistent disks or a shared file server as the primary repository for the user profile and user documents. Because View Composer lets you separate user data from the operating system, you might find that only the persistent disk needs to be backed up or replicated, which further reduces storage requirements. For more information, see "Reducing Storage Requirements with View Composer," on page 37.

NOTE Decisions regarding dedicated storage components can best be made during a pilot phase. The main consideration is I/Os per second (IOPS). You might experiment with a tiered-storage strategy or Virtual SAN storage to maximize performance and cost savings.

For more information, see the best-practices guide called *Storage Considerations for VMware View*.

Storage Bandwidth Considerations

In a View environment, logon storms are the main consideration when determining bandwidth requirements.

Although many elements are important to designing a storage system that supports a View environment, from a server configuration perspective, planning for proper storage bandwidth is essential. You must also consider the effects of port consolidation hardware.

View environments can occasionally experience I/O storm loads, during which all virtual machines undertake an activity at the same time. I/O storms can be triggered by guest-based agents such as antivirus software or software-update agents. I/O storms can also be triggered by human behavior, such as when all employees log in at nearly the same time in the morning. VMware has tested a logon storm scenario for 10,000 desktops. For more information, see "View Composer Performance Test Results," on page 64.

You can minimize these storm workloads through operational best practices, such as staggering updates to different virtual machines. You can also test various log-off policies during a pilot phase to determine whether suspending or powering off virtual machines when users log off causes an I/O storm. By storing View Composer replicas on separate, high-performance datastores, you can speed up intensive, concurrent read operations to contend with I/O storm loads. For example, you can use one of the following storage strategies:

- Manually configure the pool settings so that replicas are stored on separate, high-performance datastores.
- Use Virtual SAN, available with vSphere 5.5 Update 1 or later, which uses Software Policy-Based Management to determine which kinds of disks to use for replicas.
- Use Virtual Volumes, available with vSphere 6.0 or later, which uses Software Policy-Based Management to determine which kinds of disks to use for replicas.

In addition to determining best practices, VMware recommends that you provide bandwidth of 1Gbps per 100 virtual machines, even though average bandwidth might be 10 times less than that. Such conservative planning guarantees sufficient storage connectivity for peak loads.

Network Bandwidth Considerations

Certain virtual and physical networking components are required to accommodate a typical workload.

For display traffic, many elements can affect network bandwidth, such as protocol used, monitor resolution and configuration, and the amount of multimedia content in the workload. Concurrent launches of streamed applications can also cause usage spikes.

Because the effects of these issues can vary widely, many companies monitor bandwidth consumption as part of a pilot project. As a starting point for a pilot, plan for 150 to 200Kbps of capacity for a typical knowledge worker.

With the PCoIP display protocol, if you have an enterprise LAN with 100Mb or a 1Gb switched network, your end users can expect excellent performance under the following conditions:

- Two monitors (1920 x 1080)
- Heavy use of Microsoft Office applications
- Heavy use of Flash-embedded Web browsing
- Frequent use of multimedia with limited use of full screen mode
- Frequent use of USB-based peripherals
- Network-based printing

For more information, see the information guide called *PCoIP Display Protocol: Information and Scenario-Based Network Sizing Guide*.

Optimization Controls Available with PCoIP

If you use the PCoIP display protocol from VMware, you can adjust several elements that affect bandwidth usage.

- You can configure the image quality level and frame rate used during periods of network congestion. The quality level setting allows you to limit the initial quality of the changed regions of the display image. Unchanged regions of the image progressively build to a lossless (perfect) quality. You can adjust the frame rate from 1 to 120 frames per second.

  This control works well for static screen content that does not need to be updated or in situations where only a portion needs to be refreshed.
- You can also turn off the build-to-lossless feature altogether if instead of progressively building to perfect quality (lossless), you choose to build to perceptual lossless.
- You can control which encryption algorithms are advertised by the PCoIP endpoint during session negotiation. By default, both Salsa20-256round12 and AES-128-GCM algorithms are available.
- With regard to session bandwidth, you can configure the maximum bandwidth, in kilobits per second, to correspond to the type of network connection, such as a 4Mbit/s Internet connection. The bandwidth includes all imaging, audio, virtual channel, USB, and control PCoIP traffic.

  You can also configure a lower limit, in kilobits per second, for the bandwidth that is reserved for the session, so that a user does not have to wait for bandwidth to become available. You can specify the Maximum Transmission Unit (MTU) size for UDP packets for a PCoIP session, from 500 to 1500 bytes.
- You can specify the maximum bandwidth that can be used for audio (sound playback) in a PCoIP session.

In addition, on most client systems, PCoIP client-side image caching stores image content on the client to avoid retransmission. By default, the cache is 90MB if the client version is 2.0 or later.

Network Configuration Example

In a View 5.2 test pod in which one vCenter Server 5.1 instance managed 5 pools of 2,000 virtual machines in each pool, each ESXi host had the following hardware and software for networking requirements.

NOTE This example was used in a View 5.2 setup, which was carried out prior to the release of VMware Virtual SAN. For guidance on sizing and designing the key components of View virtual desktop infrastructures for VMware Virtual SAN, see the white paper at http://www.vmware.com/files/pdf/products/vsan/VMW-TMD-Virt-SAN-Dsn-Szing-Guid-Horizon-View.pdf.

Physical components for each host

- Brocade 1860 Fabric Adapter utilizing 10Gig Ethernet and FCoE for network and storage traffic, respectively.
- Connection to a Brocade VCS Ethernet fabric consisting of 6 VDX6720-60 switches. The switches uplinked to the rest of the network with two 1GB connections to a Juniper J6350 router.

vLAN summary

- One 10Gb vLAN per desktop pool (5 pools)
- One 1Gb vLAN for the management network
- One 1Gb vLAN for the VMotion network
- One 10Gb vLAN for the infrastructure network

Virtual VMotion-dvswitch (1 uplink per host)	This switch was used by the ESXi hosts of infrastructure, parent, and desktop virtual machines. ■ Jumbo Frame (9000 MTU) ■ 1 Ephemeral Distributed Port Group ■ Private VLAN and 192.168.x.x addressing
Infra-dvswitch (2 uplink per host)	This switch was used by the ESXi hosts of infrastructure virtual machines. ■ Jumbo frame (9000 MTU) ■ 1 Ephemeral distributed port group ■ Infrastructure VLAN /24 (256 addresses)
Desktop-dvswitch (2 uplink per host)	This switch was used by the ESXi hosts of parent, and desktop virtual machines. ■ Jumbo frame (9000 MTU) ■ 6 Ephemeral distributed port groups ■ 5 Desktop port groups (1 per pool) ■ Each network was /21, 2048 addresses

View Composer Performance Test Results

These test results describe a View 5.2 setup with 10,000-desktops, in which one vCenter Server 5.1 instance managed 5 pools of 2,000 virtual machine desktops each. Only one maintenance period was required for provisioning a new pool or for recomposing, refreshing, or rebalancing an existing pool of 2,000 virtual machines. A logon storm of 10,000 users was also tested.

The test results provided here were accomplished with the software, hardware, and configuration settings described in the following topics:

- Desktop and pool configurations described in "View Connection Server Maximums and Virtual Machine Configuration," on page 55
- Tiered-storage components described in "Shared Storage Example," on page 59
- Networking components described in "Network Bandwidth Considerations," on page 62

Capacity for an Hour-Long Logon Storm of 10,000 Users

NOTE This example was used in a View 5.2 setup, which was carried out prior to the release of VMware Virtual SAN. For guidance on sizing and designing the key components of View virtual desktop infrastructures for VMware Virtual SAN, see the white paper at http://www.vmware.com/files/pdf/products/vsan/VMW-TMD-Virt-SAN-Dsn-Szing-Guid-Horizon-View.pdf. For test results with various workloads and View operations when using Virtual SAN, see the reference architecture white paper at http://www.vmware.com/files/pdf/techpaper/vmware-horizon-view-virtual-san-reference-architecture.pdf.

The Virtual SAN feature available with vSphere 6.0 and later releases contains many performance improvements over the feature that was available with vSphere 5.5 Update 1. With vSphere 6.0 this feature also has broader HCL (hardware compatibility) support. For more information about Virtual SAN in vSphere 6 or later, see the *Administering VMware Virtual SAN* document.

In a test setup, the following desktop and pool configurations were used for a logon storm scenario for 10,000 desktops. The power policy for desktops was set to Always On.

For 10,000 desktops the logon storm occurred over a 60-minute period, using a normal distribution of logon times. The virtual machines were powered on and were available before the logon storm began. After logon, a workload started, which included the following applications: Adobe Reader, Microsoft Outlook, Internet Explorer, Microsoft Word, and Notepad.

Following are additional details of the logon storm that was sustained during testing:

- 95% of logons occurred within +/- 2 standard deviation window (40 minutes).
- 68% of logons occurred within +/- 1 standard deviation window (20 minutes).
- Peak logon rate was 400/min, or 6.67/second.

Time Required for Provisioning a Pool

Pools are provisioned either up front, when you create the pool, or on demand, as users are assigned to them. Provisioning means creating the virtual machine and configuring it to use the correct operating system image and network settings.

In a test setup already containing 4 pools of 2,000 virtual machines in each pool, provisioning a fifth pool that contained 2,000 virtual machines took 4 hours. All virtual machines were provisioned up front.

Time Required for Recomposing a Pool

You can use a recompose operation to provide operating system patches, install or update applications, or modify the desktop hardware settings of virtual machines in a pool. Before recomposing a pool, you take a snapshot of a virtual machine that has new configuration. The recompose operation uses that snapshot to update all virtual machines in the pool.

In a test setup of 5 pools of 2,000 virtual machines in each pool, a recompose of one pool of 2,000 virtual machines took 6 hours and 40 minutes. All virtual machines were powered on and available before the recompose operation began.

Time Required for Refreshing a Pool

Because disks grow over time, you can conserve disk space by refreshing a desktop to its original state when users log off, or you can set a schedule for periodically refreshing desktops. For example, you can schedule desktops to refresh daily, weekly, or monthly.

In a test setup of 5 pools of 2,000 virtual machines in each pool, a refresh of one pool of 2,000 virtual machines took 2 hours and 40 minutes. All virtual machines were powered on and available before the refresh operation began.

Time Required for Rebalancing a Pool

A desktop rebalance operation evenly redistributes linked-clone desktops among available logical drives. A rebalance operation saves storage space on overloaded drives and ensures that no drives are underused. You can also use a rebalance operation to migrate all virtual machines in a desktop pool to or from a Virtual SAN datastore.

In a test pod that contained 5 pools of 2,000 virtual machines in each pool, 2 datastores were added to the pod for one test. For another test, 2 datastores were removed from the pod. After the datastores were added or removed, a rebalance operation was performed on one of the pools. A rebalance of one pool of 2,000 virtual machines took 9 hours. All virtual machines were powered on and available before the rebalance operation began.

WAN Support and PCoIP

For wide-area networks (WANs), you must consider bandwidth constraints and latency issues. The PCoIP display protocol provided by VMware adapts to varying latency and bandwidth conditions.

If you use the RDP display protocol, you must have a WAN optimization product to accelerate applications for users in branch offices or small offices. With PCoIP, many WAN optimization techniques are built into the base protocol.

- WAN optimization is valuable for TCP-based protocols such as RDP because these protocols require many handshakes between client and server. The latency of these handshakes can be quite large. WAN accelerators spoof replies to handshakes so that the latency of the network is hidden from the protocol. Because PCoIP is UDP-based, this form of WAN acceleration is unnecessary.
- WAN accelerators also compress network traffic between client and server, but this compression is usually limited to 2:1 compression ratios. PCoIP is able to provide compression ratios of up to 100:1 for images and audio.

For information about the controls introduced with View 5 that you can use to adjust the way PCoIP consumes bandwidth, see "Optimization Controls Available with PCoIP," on page 63.

Bandwidth Requirements for Various Types of Users

When determining minimum bandwidth requirements for PCoIP, plan with the following estimates:

- 100 to 150Kbps average bandwidth for a basic office productivity desktop: typical office applications with no video, no 3D graphics, and the default Windows and View settings.
- 50 to 100Kbps average bandwidth for an optimized office productivity desktop: typical office applications with no video, no 3D graphics, with Windows desktop settings optimized and View optimized.
- 400 to 600Kbps average bandwidth for virtual desktops utilizing multiple monitors, 3D, Aero, and Office 2010.
- 500Kbps to 1Mbps minimum peak bandwidth to provide headroom for bursts of display changes. In general, size your network using the average bandwidth, but consider peak bandwidth to accommodate bursts of imaging traffic associated with large screen changes.
- 2Mbps per simultaneous user running 480p video, depending upon the configured frame rate limit and the video type.

NOTE The estimate of 50 to 150Kbps per typical user is based on the assumption that all users are operating continuously and performing similar tasks over an 8- to 10- hour day. The 50Kbps bandwidth usage figure is from View Planner testing on a LAN with the Build-to-Lossless feature disabled. Situations may vary in that some users may be fairly inactive and consume almost no bandwidth, allowing more users per link. Therefore, these guidelines are intended to provide a starting point for more detailed bandwidth planning and testing.

The following example shows how to calculate the number of concurrent users at a branch or remote office that has a 1.5Mbps T1 line.

Branch or Remote Office Scenario

- Users have basic Microsoft Office productivity applications, no video, no 3D graphics, and USB keyboards and mouse devices.
- The bandwidth required per typical office user on View is from 50-150Kbps.
- The T1 network capacity is 1.5Mbps.

- Bandwidth utilization is 80 percent (.8 utilization factor).

Formula for Determining the Number of Users Supported

- In the worst case, users require 150Kbps: (1.5Mbps*.8)/150Kbps = (1500*.8)/150 = 8 users
- In the best case, users require 50Kbps: (1.5Mbps*.8)/50Kbps = (1500*.8)/50 = 24 users

Result

This remote office can support between 8 and 24 concurrent users per T1 line with 1.5Mbps capacity.

IMPORTANT You might require optimization of both View and Windows desktop settings to achieve this user density.

This information was excerpted from the information guide called *VMware View 5 with PCoIP: Network Optimization Guide*.

View Building Blocks

A building block consists of physical servers, a vSphere infrastructure, View servers, shared storage, and virtual machine desktops for end users. A building block is a logical construct and should not be sized for more than 2,000 View desktops. Customers usually include up to five building blocks in a View pod, although in theory you can use more blocks than that, as long as the pod does not go above 10,000 sessions and 7 View Connection Server instances.

Table 4-10. Example of a LAN-Based View Building Block for 2,000 Virtual Machine Desktops

Item	Example
vSphere clusters	1 or more
80-port network switch	1
Shared storage system	1
vCenter Server with View Composer on the same host	1 (can be run in the block itself)
Database	MS SQL Server or Oracle database server (can be run in the block itself)
VLANs	3 (a 1Gbit Ethernet network for each: management network, storage network, and VMotion network)

Each vCenter Server can support up to 10,000 virtual machines. This support enables you to have building blocks that contain more than 2,000 virtual machine desktops. However, the actual block size is also subject to other View-specific limitations.

If you have only one building block in a pod, use two View Connection Server instances for redundancy.

View Pods

A pod is a unit of organization determined by View scalability limits.

Pod Example Using Five Building Blocks

A traditional View pod integrates five 2,000-user building blocks that you can manage as one entity.

Table 4-11. Example of a LAN-Based View Pod Constructed of 5 Building Blocks

Item	Number
Building blocks for a View pod	5
vCenter Server and View Composer	5 (1 virtual machine that hosts both in each building block)
Database server	5 (1 standalone database server in each building block) MS SQL Server or Oracle database server
View Connection Servers	7 (5 for connections from inside the corporate network and 2 for connections from outside)
vLANs	See Table 4-10.
10Gb Ethernet module	1
Modular networking switch	1

Each vCenter Server can support up to 10,000 virtual machines. This support enables you to have building blocks that contain more than 2,000 virtual machine desktops. However, the actual block size is also subject to other View-specific limitations.

For both examples described here, a network core can load balance incoming requests across View Connection Server instances. Support for a redundancy and failover mechanism, usually at the network level, can prevent the load balancer from becoming a single point of failure. For example, the Virtual Router Redundancy Protocol (VRRP) can communicate with a load balancer to add redundancy and failover capability.

If a View Connection Server instance fails or becomes unresponsive during an active session, users do not lose data. Desktop states are preserved in the virtual machine desktop so that users can connect to a different View Connection Server instance and their desktop session resumes from where it was when the failure occurred.

Figure 4-2. Pod Diagram for 10,000 Virtual Machine Desktops

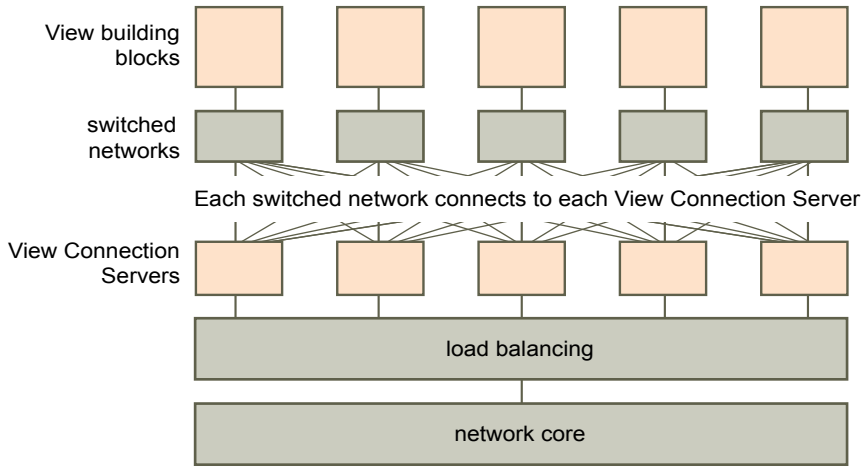

Pod Example Using One vCenter Server

In the previous section, the View pod consisted of multiple building blocks. Each building block supported 2,000 virtual machines with a single vCenter Server. VMware has received many requests from both customers and partners to use a single vCenter Server to manage a View pod. This request arises from the fact that a single instance of vCenter Server can support 10,000 virtual machines. With View 5.2 and later, customers have the ability to use a single vCenter Server to manage a 10,000-desktop environment. This topic illustrates an architecture based on using a single vCenter Serverto manage 10,000 desktops

Although using one vCenter Server and one View Composer for 10,000 desktops is possible, doing so creates a situation where there is a single point of failure. The loss of that single vCenter Server renders the entire desktop deployment unavailable for power, provisioning, and refit operations. For this reason, choose a deployment architecture that meets your requirements for overall component resiliency.

For this example, a 10,000-user pod consists of physical servers, a vSphere infrastructure, View servers, shared storage, and 5 clusters of 2,000 virtual desktops per cluster.

Table 4-12. Example of a LAN-Based View Pod with One vCenter Server

Item	Example
vSphere clusters	6 (5 clusters with one linked-clone pool per cluster, and 1 infrastructure cluster)
vCenter Server	1
View Composer	1 (standalone)
Database server	1 (standalone) MS SQL Server or Oracle database server
Active Directory server	1 or 2
View Connection Server instances	5
Security servers	5
vLANs	8 (5 for the desktop pool clusters, and 1 each for management, VMotion, and the infrastructure cluster)

Cloud Pod Architecture Overview

To use a group of replicated View Connection Server instances across a WAN, MAN (metropolitan area network), or other non-LAN, in scenarios where a View deployment needs to span datacenters, you must use the Cloud Pod Architecture feature.

This feature uses standard View components to provide cross-datacenter administration, global and flexible user-to-desktop mapping, high-availability desktops, and disaster recovery capabilities. You can link together four View pods to provide a single large desktop brokering and management environment for two geographically distant sites and manage up to 20,000 remote desktops.

The following diagram is an example of a basic Cloud Pod Architecture topology.

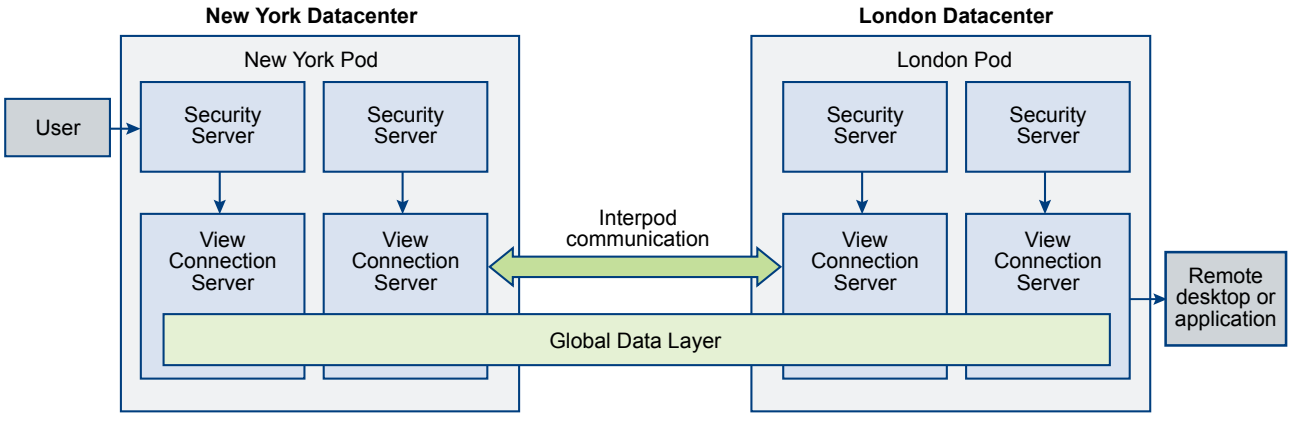

In the example topology, two previously standalone View pods in different datacenters are joined together to form a single pod federation. An end user in this environment can connect to a View Connection Server instance in the New York datacenter and receive a desktop or application in the London data center.

The Cloud Pod Architecture feature is not supported in an IPv6 environment.

For more information, see *Administering View Cloud Pod Architecture*.

Advantages of Using Multiple vCenter Servers in a Pod

When you create a design for a View production environment that accommodates more than 500 desktops, several considerations affect whether to use one vCenter Server instance rather than multiple instances.

Starting with View 5.2, VMware supports managing up to 10,000 desktop virtual machines within a single View pod with a single vCenter 5.1 or later server. Before you attempt to manage 10,000 virtual machines with a single vCenter Server instance, take the following considerations into account:

- Duration of your company's maintenance windows
- Capacity for tolerating View component failures
- Frequency of power, provisioning, and refit operations
- Simplicity of infrastructure

Duration of Maintenance Windows

Concurrency settings for virtual machine power, provisioning, and maintenance operations are determined per vCenter Server instance.

Pod designs with one vCenter Server instance	Concurrency settings determine how many operations can be queued up for an entire View pod at one time. For example, if you set concurrent provisioning operations to 20 and you have only one vCenter Server instance in a pod, a desktop pool larger than 20 will cause provisioning operations to be serialized. After queuing 20 concurrent operations simultaneously, one operation must complete before the next begins. In large-scale View deployments, this provisioning operation can take a long time.
Pod designs with multiple vCenter Server instances	Each instance can provision 20 virtual machines concurrently.

To ensure more operations are completed simultaneously within one maintenance window, you can add multiple vCenter Server instances (up to five) to your pod, and deploy multiple desktop pools in vSphere clusters managed by separate vCenter Server instances. A vSphere cluster can be managed by only one vCenter Server instance at one time. To achieve concurrency across vCenter Server instances, you must deploy your desktop pools accordingly.

Capacity for Tolerating Component Failures

The role of vCenter Server in View pods is to provide power, provisioning, and refit (refresh, recompose, and rebalance) operations. After a virtual machine desktop is deployed and powered on, View does not rely on vCenter Server for the normal course of operations.

Because each vSphere cluster must be managed by a single vCenter Server instance, this server represents a single point of failure in every View design. This risk is also true for each View Composer instance. (There is a one-to-one mapping between each View Composer instance and vCenter Server instance.) Using one of the following products can mitigate the impact of a vCenter Server or View Composer outage:

- VMware vSphere High Availability (HA)
- VMware vCenter Server Heartbeat™

- Compatible third-party failover products

IMPORTANT To use one of these failover strategies, the vCenter Server instance must not be installed in a virtual machine that is part of the cluster that the vCenter Server instance manages.

In addition to these automated options for vCenter Server failover, you can also choose to rebuild the failed server on a new virtual machine or physical server. Most key information is stored in the vCenter Server database.

Risk tolerance is an important factor in determining whether to use one or multiple vCenter Server instances in your pod design. If your operations require the ability to perform desktop management tasks such as power and refit of all desktops simultaneously, you should spread the impact of an outage across fewer desktops at a time by deploying multiple vCenter Server instances. If you can tolerate your desktop environment being unavailable for management or provisioning operations for a long period, or if you choose to use a manual rebuild process, you can deploy a single vCenter Server instance for your pod.

Frequency of Power, Provisioning, and Refit Operations

Certain virtual machine desktop power, provisioning, and refit operations are initiated only by administrator actions, are usually predictable and controllable, and can be confined to established maintenance windows. Other virtual machine desktop power and refit operations are triggered by user behavior, such as using the Refresh on Logoff or Suspend on Logoff settings, or by scripted action, such as using Distributed Power Management (DPM) during windows of user inactivity to power off idle ESXi hosts.

If your View design does not require user-triggered power and refit operations, a single vCenter Server instance can probably suit your needs. Without a high frequency of user-triggered power and refit operations, no long queue of operations can accumulate that might cause View Connection Server to time-out waiting for vCenter Server to complete the requested operations within the defined concurrency setting limits.

Many customers elect to deploy floating pools and use the Refresh on Logoff setting to consistently deliver desktops that are free of stale data from previous sessions. Examples of stale data include unclaimed memory pages in `pagefile.sys` or Windows `temp` files. Floating pools can also minimize the impact of malware by frequently resetting desktops to a known clean state.

Some customers are reducing electricity usage by configuring View to power off desktops not in use so that vSphere DRS (Distributed Resources Scheduler) can consolidate the running virtual machines onto a minimum number of ESXi hosts. VMware Distributed Power Management then powers off the idle hosts. In scenarios such as these, multiple vCenter Server instances can better accommodate the higher frequency of power and refit operations required to avoid operations time-outs.

Simplicity of Infrastructure

A single vCenter Server instance in a large-scale View design offers some compelling benefits, such as a single place to manage golden master images and parent virtual machines, a single vCenter Server view to match the View Administrator console view, and fewer production back-end databases and database servers. Disaster Recovery planning is simpler for one vCenter Server than it is for multiple instances. Make sure you weigh the advantages of multiple vCenter Server instances, such as duration of maintenance windows and frequency of power and refit operations, against the disadvantages, such as the additional administrative overhead of managing parent virtual machine images and the increased number of infrastructure components required.

Your design might benefit from a hybrid approach. You can choose to have very large and relatively static pools managed by one vCenter Server instance and have several smaller, more dynamic desktop pools managed by multiple vCenter Server instances. The best strategy for upgrading existing large-scale pods is to first upgrade the VMware software components of your existing pod. Before changing your pod design, gauge the impact of the improvements of the latest version's power, provisioning, and refit operations, and later experiment with increasing the size of your desktop pools to find the right balance of more large desktop pools on fewer vCenter Server instances.

Planning for Security Features 5

View offers strong network security to protect sensitive corporate data. For added security, you can integrate View with certain third-party user-authentication solutions, use a security server, and implement the restricted entitlements feature.

IMPORTANT With Horizon 6 version 6.2 and later releases, View can perform cryptographic operations using FIPS (Federal Information Processing Standard) 140-2 compliant algorithms. You can enable the use of these algorithms by installing View in FIPS mode. Not all View features are supported in FIPS mode. For more information, see the *View Installation* document.

This chapter includes the following topics:

Understanding Client Connections

Horizon Client and View Administrator communicate with a View Connection Server host over secure HTTPS connections. Information about the server certificate on View Connection Server is communicated to the client as part of the SSL handshake between client and server.

The initial Horizon Client connection, which is used for user authentication and remote desktop and application selection, is created when a user opens Horizon Client and provides a fully qualified domain name for the View Connection Server, security server, or Access Point host. The View Administrator connection is created when an administrator types the View Administrator URL into a Web browser.

A default SSL server certificate is generated during View Connection Server installation. By default, SSL clients are presented with this certificate when they visit a secure page such as View Administrator.

You can use the default certificate for testing, but you should replace it with your own certificate as soon as possible. The default certificate is not signed by a commercial Certificate Authority (CA). Use of noncertified certificates can allow untrusted parties to intercept traffic by masquerading as your server.

- Client Connections Using the PCoIP Secure Gateway on page 74

  When clients connect to a remote desktop or application with the PCoIP display protocol from VMware, Horizon Client can make a second connection to the PCoIP Secure Gateway component on a View Connection Server instance, security server, or Access Point appliance. This connection provides the required level of security and connectivity when accessing remote desktops and applications from the Internet.

- Tunneled Client Connections with Microsoft RDP on page 75

  When users connect to a remote desktop with the Microsoft RDP display protocol, Horizon Client can make a second HTTPS connection to the View Connection Server host. This connection is called the tunnel connection because it provides a tunnel for carrying RDP data.

- Direct Client Connections on page 75

  Administrators can configure View Connection Server settings so that remote desktop and application sessions are established directly between the client system and the remote application or desktop virtual machine, bypassing the View Connection Server host. This type of connection is called a direct client connection.

Client Connections Using the PCoIP Secure Gateway

When clients connect to a remote desktop or application with the PCoIP display protocol from VMware, Horizon Client can make a second connection to the PCoIP Secure Gateway component on a View Connection Server instance, security server, or Access Point appliance. This connection provides the required level of security and connectivity when accessing remote desktops and applications from the Internet.

Security servers and Access Point appliances include a PCoIP Secure Gateway component, which offers the following advantages:

- The only remote desktop and application traffic that can enter the corporate data center is traffic on behalf of a strongly authenticated user.
- Users can access only the resources that they are authorized to access.
- This connection supports PCoIP, which is an advanced remote display protocol that makes more efficient use of the network by encapsulating video display packets in UDP instead of TCP.
- PCoIP is secured by AES-128 encryption by default. You can, however, change the encryption cipher to AES-256.
- No VPN is required, as long as PCoIP is not blocked by any networking component. For example, someone trying to access their remote desktop or application from inside a hotel room might find that the proxy the hotel uses is not configured to pass PCoIP.

  For more information, see "Firewall Rules for DMZ-Based Security Servers," on page 84.

Security servers with PCoIP support run on Windows Server 2008 R2 and Windows Server 2012 R2 operating systems and take full advantage of the 64-bit architecture. This security server can also take advantage of Intel processors that support AES New Instructions (AESNI) for highly optimized PCoIP encryption and decryption performance.

For more information about Access Point virtual appliances, see *Deploying and Configuring Access Point*.

Tunneled Client Connections with Microsoft RDP

When users connect to a remote desktop with the Microsoft RDP display protocol, Horizon Client can make a second HTTPS connection to the View Connection Server host. This connection is called the tunnel connection because it provides a tunnel for carrying RDP data.

The tunnel connection offers the following advantages:

- RDP data is tunneled through HTTPS and is encrypted using SSL. This powerful security protocol is consistent with the security provided by other secure Web sites, such as those that are used for online banking and credit card payments.
- A client can access multiple desktops over a single HTTPS connection, which reduces the overall protocol overhead.
- Because View manages the HTTPS connection, the reliability of the underlying protocols is significantly improved. If a user temporarily loses a network connection, the HTTP connection is reestablished after the network connection is restored and the RDP connection automatically resumes without requiring the user to reconnect and log in again.

In a standard deployment of View Connection Server instances, the HTTPS secure connection terminates at the View Connection Server. In a DMZ deployment, the HTTPS secure connection terminates at a security server or Access Point appliance. See "Preparing to Use a Security Server," on page 80 for information on DMZ deployments and security servers.

Clients that use the PCoIP display protocol can use the tunnel connection for USB redirection and multimedia redirection (MMR) acceleration, but for all other data, PCoIP uses the PCoIP Secure Gateway on a security server or Access Point appliance. For more information, see "Client Connections Using the PCoIP Secure Gateway," on page 74.

For more information about Access Point virtual appliances, see *Deploying and Configuring Access Point*.

Direct Client Connections

Administrators can configure View Connection Server settings so that remote desktop and application sessions are established directly between the client system and the remote application or desktop virtual machine, bypassing the View Connection Server host. This type of connection is called a direct client connection.

With direct client connections, an HTTPS connection is still made between the client and the View Connection Server host for users to authenticate and select remote desktops and applications, but the second HTTPS connection (the tunnel connection) is not used.

Direct PCoIP connections include the following built-in security features:

- PCoIP supports Advanced Encryption Standard (AES) encryption, which is turned on by default, and PCoIP uses IP Security (IPsec).
- PCoIP works with third-party VPN clients.

For clients that use the Microsoft RDP display protocol, direct client connections to remote desktops are appropriate only if your deployment is inside a corporate network. With direct client connections, RDP traffic is sent unencrypted over the connection between the client and the desktop virtual machine.

Choosing a User Authentication Method

View uses your existing Active Directory infrastructure for user authentication and management. For added security, you can integrate View with two-factor authentication solutions, such as RSA SecurID and RADIUS, and smart card authentication solutions.

- Active Directory Authentication on page 76

  Each View Connection Server instance is joined to an Active Directory domain, and users are authenticated against Active Directory for the joined domain. Users are also authenticated against any additional user domains with which a trust agreement exists.

- Using Two-Factor Authentication on page 77

  You can configure a View Connection Server instance so that users are required to use RSA SecurID authentication or RADIUS (Remote Authentication Dial-In User Service) authentication.

- Smart Card Authentication on page 77

  A smart card is a small plastic card that is embedded with a computer chip. Many government agencies and large enterprises use smart cards to authenticate users who access their computer networks. One type of smart card used by the United States Department of Defense is called a Common Access Card (CAC).

- Using the Log In as Current User Feature Available with Windows-Based Horizon Client on page 78

  With Horizon Client for Windows, when users select the **Log in as current user** check box, the credentials that they provided when logging in to the client system are used to authenticate to the View Connection Server instance and to the remote desktop. No further user authentication is required.

Active Directory Authentication

Each View Connection Server instance is joined to an Active Directory domain, and users are authenticated against Active Directory for the joined domain. Users are also authenticated against any additional user domains with which a trust agreement exists.

For example, if a View Connection Server instance is a member of Domain A and a trust agreement exists between Domain A and Domain B, users from both Domain A and Domain B can connect to the View Connection Server instance with Horizon Client.

Similarly, if a trust agreement exists between Domain A and an MIT Kerberos realm in a mixed domain environment, users from the Kerberos realm can select the Kerberos realm name when connecting to the View Connection Server instance with Horizon Client.

You can place users and groups in the following Active Directory domains:

- The View Connection Server domain
- A different domain that has a two-way trust relationship with the View Connection Server domain
- A domain in a different forest than the View Connection Server domain that is trusted by the View Connection Server domain in a one-way external or realm trust relationship
- A domain in a different forest than the View Connection Server domain that is trusted by the View Connection Server domain in a one-way or two-way transitive forest trust relationship

View Connection Server determines which domains are accessible by traversing trust relationships, starting with the domain in which the host resides. For a small, well-connected set of domains, View Connection Server can quickly determine a full list of domains, but the time that it takes increases as the number of domains increases or as the connectivity between the domains decreases. The list might also include domains that you would prefer not to offer to users when they log in to their remote desktops and applications.

Administrators can use the `vdmadmin` command-line interface to configure domain filtering, which limits the domains that a View Connection Server instance searches and that it displays to users. See the *View Administration* document for more information.

Policies, such as restricting permitted hours to log in and setting the expiration date for passwords, are also handled through existing Active Directory operational procedures.

Using Two-Factor Authentication

You can configure a View Connection Server instance so that users are required to use RSA SecurID authentication or RADIUS (Remote Authentication Dial-In User Service) authentication.

- RADIUS support offers a wide range of alternative two-factor token-based authentication options.
- View also provides an open standard extension interface to allow third-party solution providers to integrate advanced authentication extensions into View.

Because two-factor authentication solutions such as RSA SecurID and RADIUS work with authentication managers, installed on separate servers, you must have those servers configured and accessible to the View Connection Server host. For example, if you use RSA SecurID, the authentication manager would be RSA Authentication Manager. If you have RADIUS, the authentication manager would be a RADIUS server.

To use two-factor authentication, each user must have a token, such as an RSA SecurID token, that is registered with its authentication manager. A two-factor authentication token is a piece of hardware or software that generates an authentication code at fixed intervals. Often authentication requires knowledge of both a PIN and an authentication code.

If you have multiple View Connection Server instances, you can configure two-factor authentication on some instances and a different user authentication method on others. For example, you can configure two-factor authentication only for users who access remote desktops and applications from outside the corporate network, over the Internet.

View is certified through the RSA SecurID Ready program and supports the full range of SecurID capabilities, including New PIN Mode, Next Token Code Mode, RSA Authentication Manager, and load balancing.

Smart Card Authentication

A smart card is a small plastic card that is embedded with a computer chip. Many government agencies and large enterprises use smart cards to authenticate users who access their computer networks. One type of smart card used by the United States Department of Defense is called a Common Access Card (CAC).

Administrators can enable individual View Connection Server instances for smart card authentication. Enabling a View Connection Server instance to use smart card authentication typically involves adding your root certificate to a truststore file and then modifying View Connection Server settings.

All client connections, including client connections that use smart card authentication, are SSL enabled.

To use smart cards, client machines must have smart card middleware and a smart card reader. To install certificates on smart cards, you must set up a computer to act as an enrollment station. For information about whether a particular type of Horizon Client supports smart cards, see the Horizon Client documentation at https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.

Using the Log In as Current User Feature Available with Windows-Based Horizon Client

With Horizon Client for Windows, when users select the **Log in as current user** check box, the credentials that they provided when logging in to the client system are used to authenticate to the View Connection Server instance and to the remote desktop. No further user authentication is required.

To support this feature, user credentials are stored on both the View Connection Server instance and on the client system.

- On the View Connection Server instance, user credentials are encrypted and stored in the user session along with the username, domain, and optional UPN. The credentials are added when authentication occurs and are purged when the session object is destroyed. The session object is destroyed when the user logs out, the session times out, or authentication fails. The session object resides in volatile memory and is not stored in View LDAP or in a disk file.
- On the client system, user credentials are encrypted and stored in a table in the Authentication Package, which is a component of Horizon Client. The credentials are added to the table when the user logs in and are removed from the table when the user logs out. The table resides in volatile memory.

Administrators can use Horizon Client group policy settings to control the availability of the **Log in as current user** check box and to specify its default value. Administrators can also use group policy to specify which View Connection Server instances accept the user identity and credential information that is passed when users select the **Log in as current user** check box in Horizon Client.

The Log in as current user feature has the following limitations and requirements:

- When smart card authentication is set to Required on a View Connection Server instance, authentication fails for users who select the **Log in as current user** check box when they connect to the View Connection Server instance. These users must reauthenticate with their smart card and PIN when they log in to View Connection Server.
- The time on the system where the client logs in and the time on the View Connection Server host must be synchronized.
- If the default **Access this computer from the network** user-right assignments are modified on the client system, they must be modified as described in VMware Knowledge Base (KB) article 1025691.
- The client machine must be able to communicate with the corporate Active Directory server and not use cached credentials for authentication. For example, if users log in to their client machines from outside the corporate network, cached credentials are used for authentication. If the user then attempts to connect to a security server or a View Connection Server instance without first establishing a VPN connection, the user is prompted for credentials, and the Log in as Current User feature does not work.

Restricting Remote Desktop Access

You can use the restricted entitlements feature to restrict remote desktop access based on the View Connection Server instance that a user connects to.

With restricted entitlements, you assign one or more tags to a View Connection Server instance. Then, when configuring a desktop pool, you select the tags of the View Connection Server instances that you want to be able to access the desktop pool. When users log in through a tagged View Connection Server instance, they can access only those desktop pools that have at least one matching tag or no tags.

For example, your View deployment might include two View Connection Server instances. The first instance supports your internal users. The second instance is paired with a security server and supports your external users. To prevent external users from accessing certain desktops, you could set up restricted entitlements as follows:

- Assign the tag "Internal" to the View Connection Server instance that supports your internal users.

- Assign the tag "External" to the View Connection Server instance that is paired with the security server and supports your external users.
- Assign the "Internal" tag to the desktop pools that should be accessible only to internal users.
- Assign the "External" tag to the desktop pools that should be accessible only to external users.

External users cannot see the desktop pools tagged as Internal because they log in through the View Connection Server tagged as External, and internal users cannot see the desktop pools tagged as External because they log in through the View Connection Server tagged as Internal. Figure 5-1 illustrates this configuration.

Figure 5-1. Restricted Entitlements Example

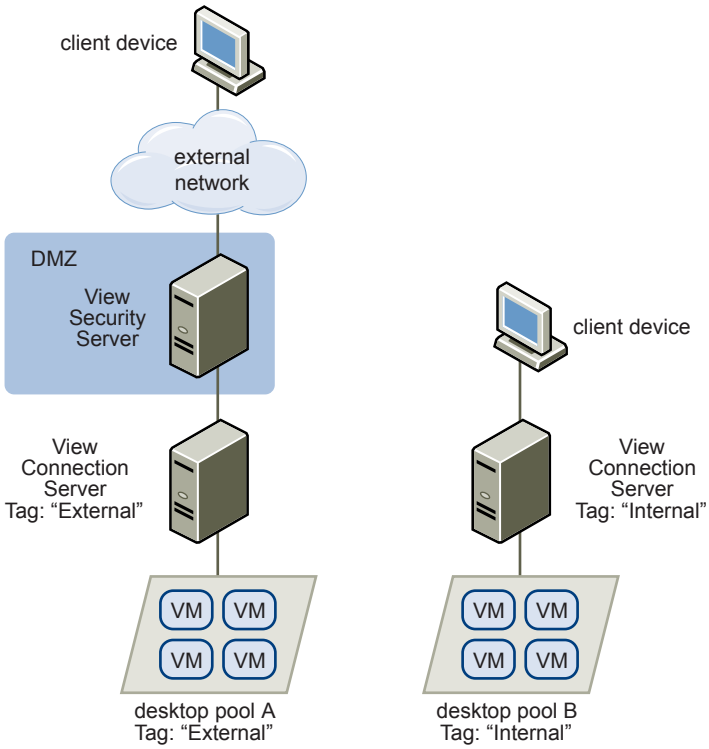

You can also use restricted entitlements to control desktop access based on the user-authentication method that you configure for a particular View Connection Server instance. For example, you can make certain desktop pools available only to users who have authenticated with a smart card.

The restricted entitlements feature only enforces tag matching. You must design your network topology to force certain clients to connect through a particular View Connection Server instance.

Using Group Policy Settings to Secure Remote Desktops and Applications

View includes Group Policy administrative (ADM) templates that contain security-related group policy settings that you can use to secure your remote desktops and applications.

For example, you can use group policy settings to perform the following tasks.

- Specify the View Connection Server instances that can accept user identity and credential information that is passed when a user selects the **Log in as current user** check box in Horizon Client for Windows.
- Enable single sign-on for smart card authentication in Horizon Client.
- Configure server SSL certificate checking in Horizon Client.

- Prevent users from providing credential information with Horizon Client command line options.
- Prevent non-Horizon Client systems from using RDP to connect to remote desktops. You can set this policy so that connections must be Horizon Client-managed, which means that users must use View to connect to remote desktops.

See the *Setting Up Desktop and Application Pools in View* for information on using remote desktop and Horizon Client group policy settings.

Implementing Best Practices to Secure Client Systems

You should implement best practices to secure client systems.

- Make sure that client systems are configured to go to sleep after a period of inactivity and require users to enter a password before the computer awakens.
- Require users to enter a username and password when starting client systems. Do not configure client systems to allow automatic logins.
- For Mac client systems, consider setting different passwords for the Keychain and the user account. When the passwords are different, users are prompted before the system enters any passwords on their behalf. Also consider turning on FileVault protection.

For a concise reference to all the security features View provides, see the *View Security* document.

Assigning Administrator Roles

A key management task in a View environment is to determine who can use View Administrator and what tasks those users are authorized to perform.

The authorization to perform tasks in View Administrator is governed by an access control system that consists of administrator roles and privileges. A role is a collection of privileges. Privileges grant the ability to perform specific actions, such as entitling a user to a desktop pool or changing a configuration setting. Privileges also control what an administrator can see in View Administrator.

An administrator can create folders to subdivide desktop pools and delegate the administration of specific desktop pools to different administrators in View Administrator. An administrator configures administrator access to the resources in a folder by assigning a role to a user on that folder. Administrators can only access the resources that reside in folders for which they have assigned roles. The role that an administrator has on a folder determines the level of access that the administrator has to the resources in that folder.

View Administrator includes a set of predefined roles. Administrators can also create custom roles by combining selected privileges.

Preparing to Use a Security Server

A security server is a special instance of View Connection Server that runs a subset of View Connection Server functions. You can use a security server to provide an additional layer of security between the Internet and your internal network.

IMPORTANT With Horizon 6 version 6.2 and later releases, you can use Access Point appliances in place of security servers. Access Point appliances are deployed as hardened virtual appliances, which are based on a Linux appliance that has been customized to provide secure access. For more information about Access Point virtual appliances, see *Deploying and Configuring Access Point*.

A security server resides within a DMZ and acts as a proxy host for connections inside your trusted network. Each security server is paired with an instance of View Connection Server and forwards all traffic to that instance. You can pair multiple security servers to a single connection server. This design provides an additional layer of security by shielding the View Connection Server instance from the public-facing Internet and by forcing all unprotected session requests through the security server.

A DMZ-based security server deployment requires a few ports to be opened on the firewall to allow clients to connect with security servers inside the DMZ. You must also configure ports for communication between security servers and the View Connection Server instances in the internal network. See "Firewall Rules for DMZ-Based Security Servers," on page 84 for information on specific ports.

Because users can connect directly with any View Connection Server instance from within their internal network, you do not need to implement a security server in a LAN-based deployment.

NOTE Security servers include a PCoIP Secure Gateway component so that clients that use the PCoIP display protocol can use a security server rather than a VPN.

For information about setting up VPNs for using PCoIP, see the VPN solution overviews, available in the Technology Partner Resources section of the Technical Resource Center at http://www.vmware.com/products/view/resources.html.

Best Practices for Security Server Deployments

You should follow best practice security policies and procedures when operating a security server in a DMZ.

The *DMZ Virtualization with VMware Infrastructure* white paper includes examples of best practices for a virtualized DMZ. Many of the recommendations in this white paper also apply to a physical DMZ.

To limit the scope of frame broadcasts, the View Connection Server instances that are paired with security servers should be deployed on an isolated network. This topology can help prevent a malicious user on the internal network from monitoring communication between the security servers and View Connection Server instances.

Alternatively, you might be able to use advanced security features on your network switch to prevent malicious monitoring of security server and View Connection Server communication and to guard against monitoring attacks such as ARP Cache Poisoning. See the administration documentation for your networking equipment for more information.

Security Server Topologies

You can implement several different security server topologies.

The topology illustrated in Figure 5-2 shows a high-availability environment that includes two load-balanced security servers in a DMZ. The security servers communicate with two View Connection Server instances inside the internal network.

Figure 5-2. Load-Balanced Security Servers in a DMZ

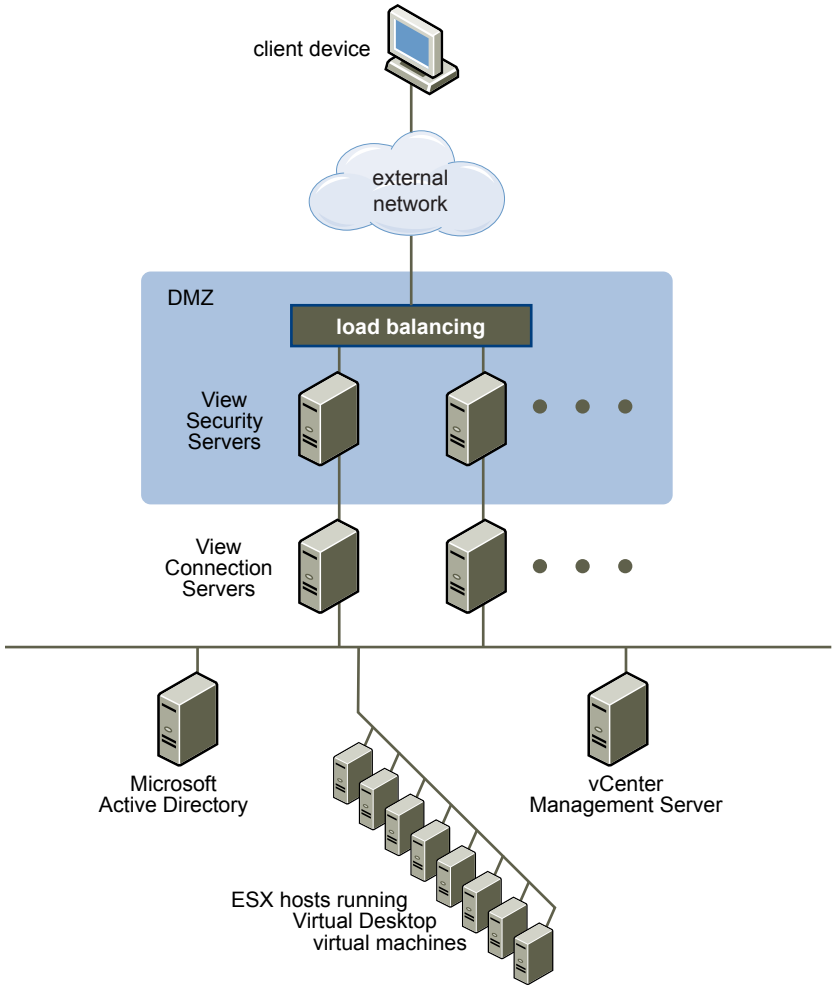

When users outside the corporate network connect to a security server, they must successfully authenticate before they can access remote desktops and applications. With appropriate firewall rules on both sides of the DMZ, this topology is suitable for accessing remote desktops and applications from client devices located on the Internet.

You can connect multiple security servers to each instance of View Connection Server. You can also combine a DMZ deployment with a standard deployment to offer access for internal users and external users.

The topology illustrated in Figure 5-3 shows an environment where four instances of View Connection Server act as one group. The instances in the internal network are dedicated to users of the internal network, and the instances in the external network are dedicated to users of the external network. If the View Connection Server instances paired with the security servers are enabled for RSA SecurID authentication, all external network users are required to authenticate by using RSA SecurID tokens.

Figure 5-3. Multiple Security Servers

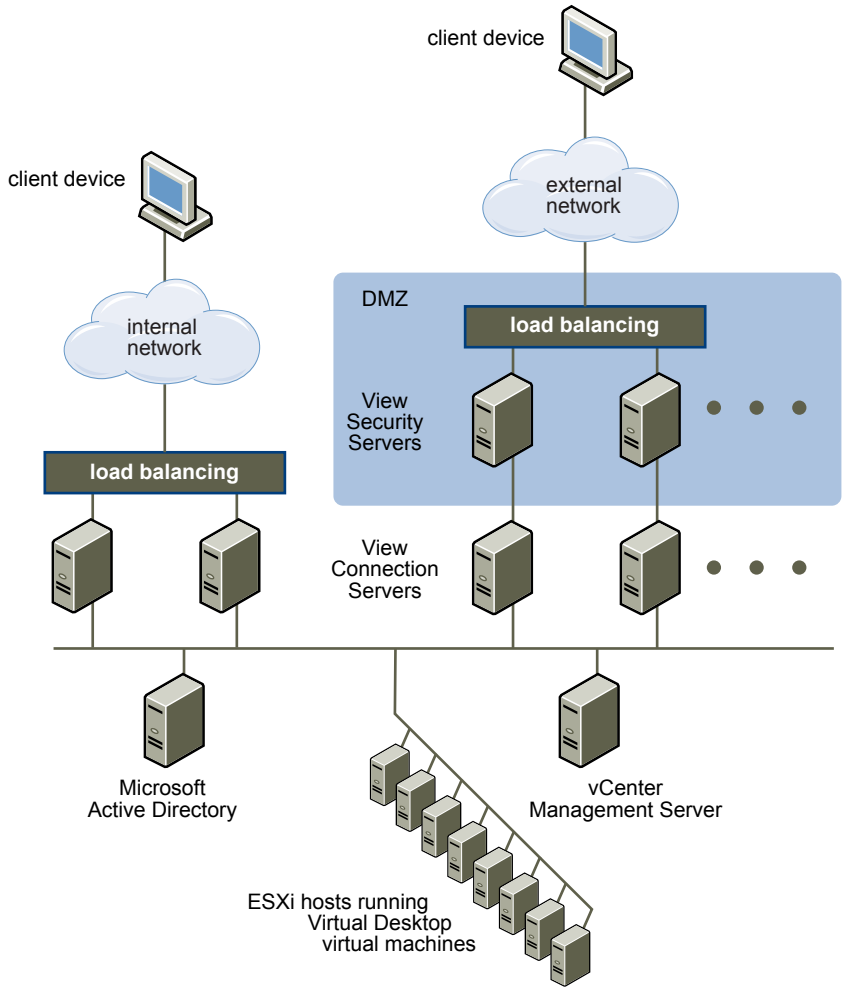

You must implement a hardware or software load balancing solution if you install more than one security server. View Connection Server does not provide its own load balancing functionality. View Connection Server works with standard third-party load balancing solutions.

Firewalls for DMZ-Based Security Servers

A DMZ-based security server deployment must include two firewalls.

- An external network-facing, front-end firewall is required to protect both the DMZ and the internal network. You configure this firewall to allow external network traffic to reach the DMZ.
- A back-end firewall, between the DMZ and the internal network, is required to provide a second tier of security. You configure this firewall to accept only traffic that originates from the services within the DMZ.

Firewall policy strictly controls inbound communications from DMZ services, which greatly reduces the risk of compromising your internal network.

Figure 5-4 shows an example of a configuration that includes front-end and back-end firewalls.

Figure 5-4. Dual Firewall Topology

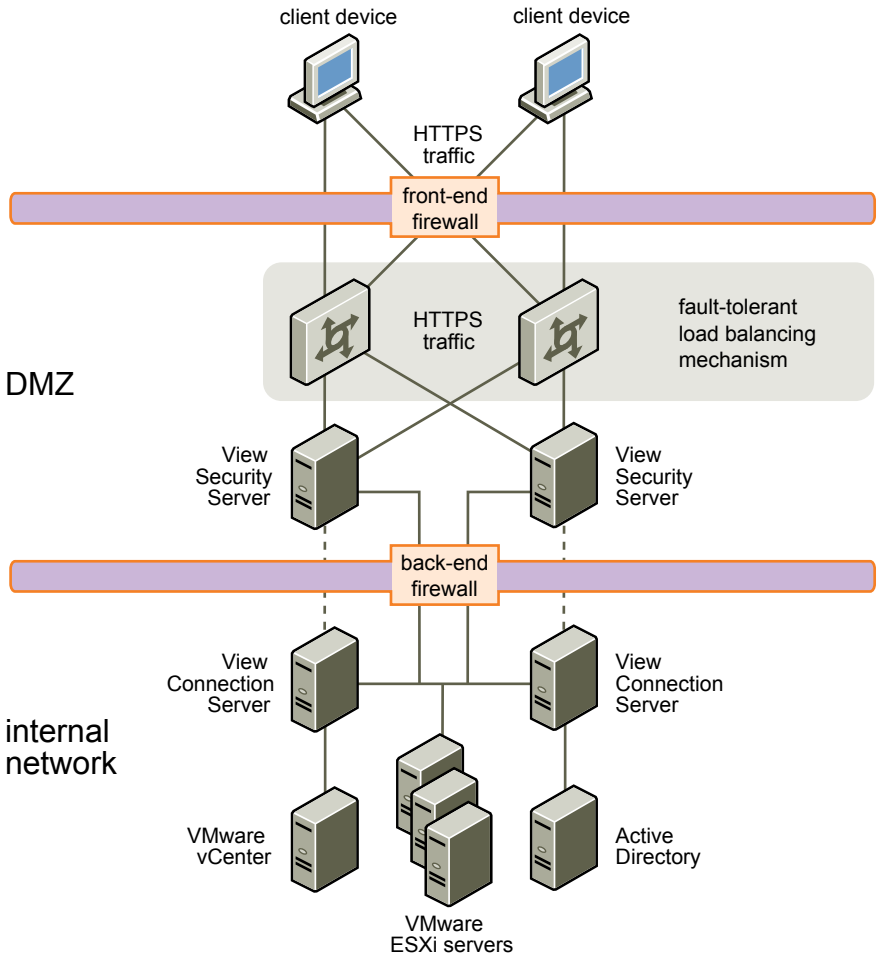

Firewall Rules for DMZ-Based Security Servers

DMZ-based security servers require certain firewall rules on the front-end and back-end firewalls. During installation, View services are set up to listen on certain network ports by default. If necessary, to comply with organization policies or to avoid contention, you can change which port numbers are used.

IMPORTANT For additional details and security recommendations, see the *View Security* document.

Front-End Firewall Rules

To allow external client devices to connect to a security server within the DMZ, the front-end firewall must allow traffic on certain TCP and UDP ports. Table 5-1 summarizes the front-end firewall rules.

Table 5-1. Front-End Firewall Rules

Source	Default Port	Protocol	Destination	Default Port	Notes
Horizon Client	TCP Any	HTTP	Security Server	TCP 80	(Optional) External client devices connect to a security server within the DMZ on TCP port 80 and are automatically directed to HTTPS. For information about the security considerations related to letting users connect with HTTP rather than HTTPS, see the *View Security* guide.
Horizon Client	TCP Any	HTTPS	Security server	TCP 443	External client devices connect to a security server within the DMZ on TCP port 443 to communicate with a Connection Server instance and remote desktops and applications.

Table 5-1. Front-End Firewall Rules (Continued)

Source	Default Port	Protocol	Destination	Default Port	Notes
Horizon Client	TCP Any UDP Any	PCoIP	Security server	TCP 4172 UDP 4172	External client devices connect to a security server within the DMZ on TCP port 4172 and UDP port 4172 to communicate with a remote desktop or application over PCoIP.
Security Server	UDP 4172	PCoIP	Horizon Client	UDP Any	Security servers send PCoIP data back to an external client device from UDP port 4172. The destination UDP port is the source port from the received UDP packets. Because these packets contain reply data, it is normally unnecessary to add an explicit firewall rule for this traffic.
Client Web browser	TCP Any	HTTPS	Security server	TCP 8443	If you use HTML Access, the external Web client connects to a security server within the DMZ on HTTPS port 8443 to communicate with remote desktops.

Back-End Firewall Rules

To allow a security server to communicate with each View Connection Server instance that resides within the internal network, the back-end firewall must allow inbound traffic on certain TCP ports. Behind the back-end firewall, internal firewalls must be similarly configured to allow remote desktops applications and View Connection Server instances to communicate with each other. Table 5-2 summarizes the back-end firewall rules.

Table 5-2. Back-End Firewall Rules

Source	Default Port	Protocol	Destination	Default Port	Notes
Security server	UDP 500	IPSec	Connection Server	UDP 500	Security servers negotiate IPSec with View Connection Server instances on UDP port 500.
Connection Server	UDP 500	IPSec	Security server	UDP 500	View Connection Server instances respond to security servers on UDP port 500.
Security Server	UDP 4500	NAT-T ISAKMP	Connection Server	UDP 4500	Required if NAT is used between a security server and its paired View Connection Server instance. Security servers use UDP port 4500 to traverse NATs and negotiate IPsec security.
Connection Server	UDP 4500	NAT-T ISAKMP	Security server	UDP 4500	View Connection Server instances respond to security servers on UDP port 4500 if NAT is used.
Security server	TCP Any	AJP13	Connection Server	TCP 8009	Security servers connect to View Connection Server instances on TCP port 8009 to forward Web traffic from external client devices. If you enable IPSec, AJP13 traffic does not use TCP port 8009 after pairing. Instead it flows over either NAT-T (UDP port 4500) or ESP.
Security server	TCP Any	JMS	Connection Server	TCP 4001	Security servers connect to View Connection Server instances on TCP port 4001 to exchange Java Message Service (JMS) traffic.
Security server	TCP Any	JMS	Connection Server	TCP 4002	Security servers connect to View Connection Server instances on TCP port 4002 to exchange secure Java Message Service (JMS) traffic.
Security server	TCP Any	RDP	Remote desktop	TCP 3389	Security servers connect to remote desktops on TCP port 3389 to exchange RDP traffic.
Security server	TCP Any	MMR	Remote desktop	TCP 9427	Security servers connect to remote desktops on TCP port 9427 to receive traffic relating to multimedia redirection (MMR) and client drive redirection.

Table 5-2. Back-End Firewall Rules (Continued)

Source	Default Port	Protocol	Destination	Default Port	Notes
Security server	TCP Any UDP 55000	PCoIP	Remote desktop or application	TCP 4172 UDP 4172	Security servers connect to remote desktops and applications on TCP port 4172 and UDP port 4172 to exchange PCoIP traffic.
Remote desktop or application	UDP 4172	PCoIP	Security server	UDP 55000	Remote desktops and applications send PCoIP data back to a security server from UDP port 4172 . The destination UDP port will be the source port from the received UDP packets and so as this is reply data, it is normally unnecessary to add an explicit firewall rule for this.
Security server	TCP Any	USB-R	Remote desktop	TCP 32111	Security servers connect to remote desktops on TCP port 32111 to exchange USB redirection traffic between an external client device and the remote desktop.
Security server	TCP Any	HTTPS	Remote desktop	TCP 22443	If you use HTML Access, security servers connect to remote desktops on HTTPS port 22443 to communicate with the Blast agent.
Security server		ESP	Connection Server		Encapsulated AJP13 traffic when NAT traversal is not required. ESP is IP protocol 50. Port numbers are not specified.
Connection Server		ESP	Security server		Encapsulated AJP13 traffic when NAT traversal is not required. ESP is IP protocol 50. Port numbers are not specified.

Understanding View Communications Protocols

View components exchange messages by using several different protocols.

Figure 5-5 illustrates the protocols that each component uses for communication when a security server is not configured. That is, the secure tunnel for RDP and the PCoIP secure gateway are not turned on. This configuration might be used in a typical LAN deployment.

Figure 5-5. View Components and Protocols Without a Security Server

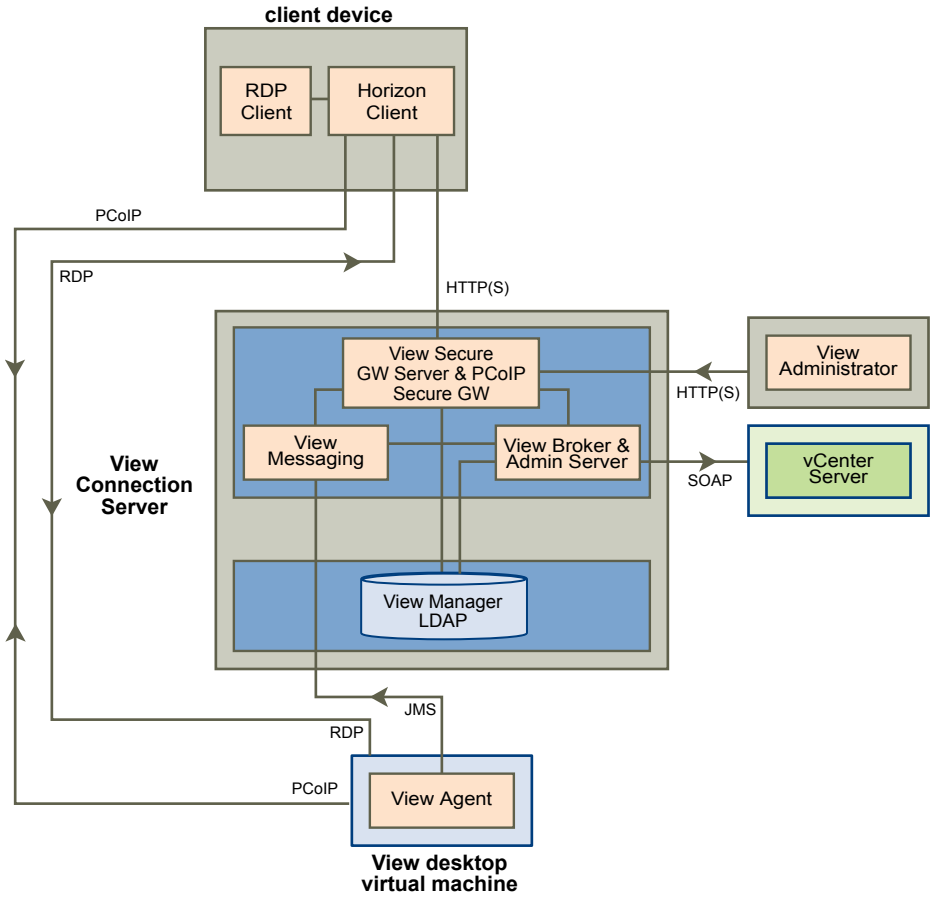

NOTE This figure shows direct connections for clients using either PCoIP or RDP. The default setting, however, is to have direct connections for PCoIP and tunnel connections for RDP.

See Table 5-3 for the default ports that are used for each protocol.

Figure 5-6 illustrates the protocols that each component uses for communication when a security server is configured. This configuration might be used in a typical WAN deployment.

Figure 5-6. View Components and Protocols with a Security Server

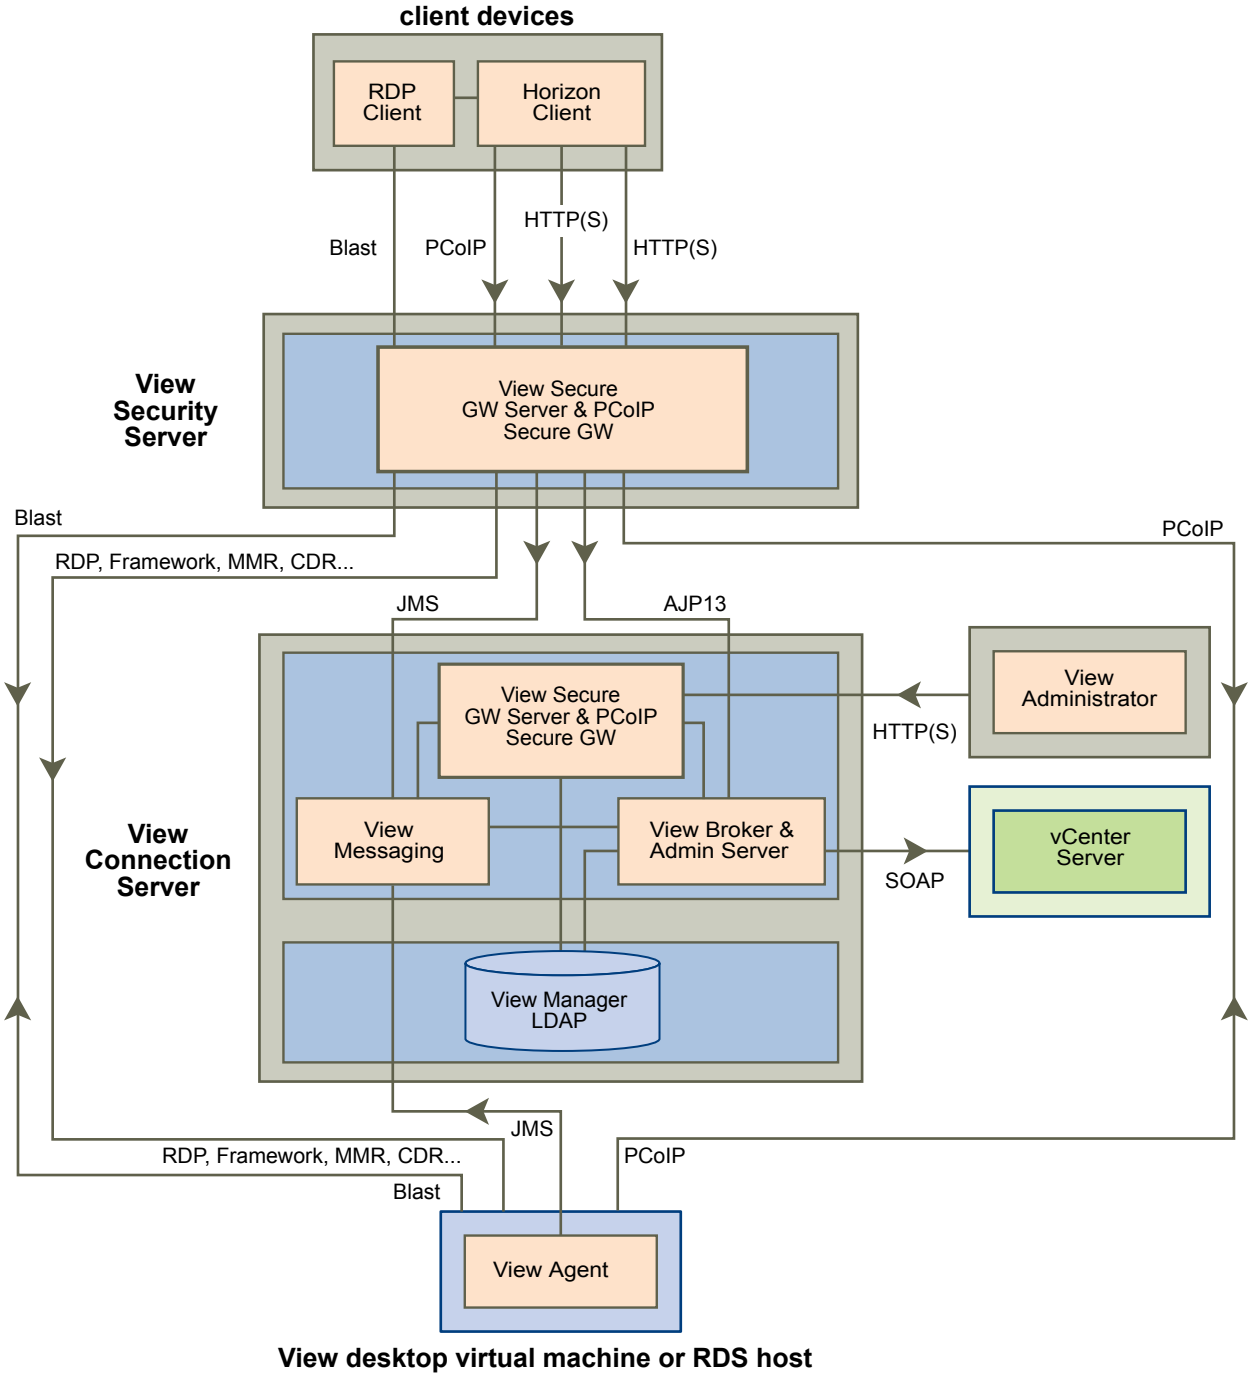

Table 5-3 lists the default ports that are used by each protocol. If necessary, to comply with organization policies or to avoid contention, you can change which port numbers are used.

Table 5-3. Default Ports

Protocol	Port
JMS	TCP port 4001 TCP port 4002
AJP13	TCP port 8009 NOTE AJP13 is used in a security server configuration only.

Table 5-3. Default Ports (Continued)

Protocol	Port
HTTP	TCP port 80
HTTPS	TCP port 443
MMR/CDR	For multimedia redirection and client drive redirection, TCP port 9427
RDP	TCP port 3389 **NOTE** If the View Connection Server instance is configured for direct client connections, these protocols connect directly from the client to the remote desktop and are not tunneled through the View Secure GW Server component.
SOAP	TCP port 80 or 443
PCoIP	Any TCP port from Horizon Client to port 4172 of the remote desktop or application. PCoIP also uses UDP port 50002 from Horizon Client (or UDP port 55000 from the PCoIP Secure Gateway) to port 4172 of the remote desktop or application.
USB redirection	TCP port 32111. This port is also used for time zone synchronization.
HTML Access	For the HTML Access Gateway on connection servers and security servers, TCP Port 8443 For View Agent connections, TCP Port 22443

TCP Ports for View Connection Server Intercommunication

View Connection Server instances in a group use additional TCP ports to communicate with each other. For example, View Connection Server instances use port 4100 or 4101 to transmit JMS inter-router (JMSIR) traffic to each other. Firewalls are generally not used between the View Connection Server instances in a group.

View Broker and Administration Server

The View Broker component, which is the core of View Connection Server, is responsible for all user interaction between clients and View Connection Server. View Broker also includes the Administration Server that is used by the View Administrator Web interface.

View Broker works closely with vCenter Server to provide advanced management of remote desktops, including virtual machine creation and power operations.

View Secure Gateway Server

View Secure Gateway Server is the server-side component for the secure HTTPS connection between client systems and a security server, Access Point appliance, or View Connection Server instance.

When you configure the tunnel connection for View Connection Server, RDP, USB, and Multimedia Redirection (MMR) traffic is tunneled through the View Secure Gateway component. When you configure direct client connections, these protocols connect directly from the client to the remote desktop and are not tunneled through the View Secure Gateway Server component.

NOTE Clients that use the PCoIP display protocol can use the tunnel connection for USB redirection and multimedia redirection (MMR) acceleration, but for all other data, PCoIP uses the PCoIP Secure Gateway on a security server or Access Point appliance.

View Secure Gateway Server is also responsible for forwarding other Web traffic, including user authentication and desktop and application selection traffic, from clients to the View Broker component. View Secure Gateway Server also passes View Administrator client Web traffic to the Administration Server component.

PCoIP Secure Gateway

Security servers and Access Point appliances include a PCoIP Secure Gateway component. When the PCoIP Secure Gateway is enabled, after authentication, clients that use PCoIP can make another secure connection to a security server or Access Point appliance. This connection allows clients to access remote desktops and applications from the Internet.

When you enable the PCoIP Secure Gateway component, PCoIP traffic is forwarded by a security server or Access Point appliance to remote desktops and applications. If clients that use PCoIP also use the USB redirection feature or multimedia redirection (MMR) acceleration, you can enable the View Secure Gateway component in order to forward that data.

When you configure direct client connections, PCoIP traffic and other traffic goes directly from a client to a remote desktop or application.

When end users such as home or mobile workers access desktops from the Internet, security servers or Access Point appliances provide the required level of security and connectivity so that a VPN connection is not necessary. The PCoIP Secure Gateway component ensures that the only remote traffic that can enter the corporate data center is traffic on behalf of a strongly authenticated user. End users can access only the resources that they are authorized to access.

View LDAP

View LDAP is an embedded LDAP directory in View Connection Server and is the configuration repository for all View configuration data.

View LDAP contains entries that represent each remote desktop and application, each accessible remote desktop, multiple remote desktops that are managed together, and View component configuration settings.

View LDAP also includes a set of View plug-in DLLs to provide automation and notification services for other View components.

View Messaging

The View Messaging component provides the messaging router for communication between View Connection Server components and between View Agent and View Connection Server.

This component supports the Java Message Service (JMS) API, which is used for messaging in View.

By default, RSA keys that are used for intercomponent message validation are 512 bits. The RSA key size can be increased to 1024 bits if you prefer stronger encryption.

NOTE When the message security mode is set to **Enhanced**, SSL is used to secure JMS connections rather than using per-message encryption.

If you want all keys to be 1024 bits, the RSA key size must be changed immediately after the first View Connection Server instance is installed and before additional servers and desktops are created. See VMware Knowledge Base (KB) article 1024431 for more information.

Firewall Rules for View Connection Server

Certain ports must be opened on the firewall for View Connection Server instances and security servers.

When you install View Connection Server, the installation program can optionally configure the required Windows Firewall rules for you. These rules open the ports that are used by default. If you change the default ports after installation, you must manually configure Windows Firewall to allow Horizon Client devices to connect to View through the updated ports.

If you choose to install HTML Access with View Connection Server, the installer configures the **VMware Horizon View Connection Server (Blast-In)** rule in Windows Firewall to open TCP port 8443, used by HTML Access.

The following table lists the default ports that can be opened automatically during installation. Ports are incoming unless otherwise noted.

Table 5-4. Ports Opened During View Connection Server Installation

Protocol	Ports	View Connection Server Instance Type
JMS	TCP 4001	Standard and replica
JMS	TCP 4002	Standard and replica
JMSIR	TCP 4100	Standard and replica
JMSIR	TCP 4101	Standard and replica
AJP13	TCP 8009	Standard and replica
HTTP	TCP 80	Standard, replica, and security server
HTTPS	TCP 443	Standard, replica, and security server
PCoIP	TCP 4172 in; UDP 4172 both directions	Standard, replica, and security server
HTTPS	TCP 8443	Standard, replica, and security server. After the initial connection to View is made, the Web browser on a client device connects to the Blast Secure Gateway on TCP port 8443. The Blast Secure Gateway must be enabled on a security server or View Connection Server instance to allow this second connection to take place.
HTTPS	TCP 8472	Standard and replica For the Cloud Pod Architecture feature: used for interpod communication.
HTTP	TCP 22389	Standard and replica For the Cloud Pod Architecture feature: used for global LDAP replication.
HTTPS	TCP 22636	Standard and replica For the Cloud Pod Architecture feature: used for secure global LDAP replication.

Firewall Rules for View Agent

The View Agent installation program opens certain TCP ports on the firewall. Ports are incoming unless otherwise noted.

Table 5-5. TCP Ports Opened During View Agent Installation

Protocol	Ports
RDP	3389
USB redirection	32111 (This port is also used for time zone synchronization.)
MMR (multimedia redirection) and CDR (client drive redirection)	9427
PCoIP	4172 (TCP and UDP)
HTML Access	22443

The View Agent installation program configures the local firewall rule for inbound RDP connections to match the current RDP port of the host operating system, which is typically 3389. If you change the RDP port number after installation, you must change the associated firewall rules.

If you instruct the View Agent installation program to not enable Remote Desktop support, it does not open ports 3389 and 32111, and you must open these ports manually.

If you use a virtual machine template as a desktop source, firewall exceptions carry over to deployed desktops only if the template is a member of the desktop domain. You can use Microsoft group policy settings to manage local firewall exceptions. See the Microsoft Knowledge Base (KB) article 875357 for more information.

Firewall Rules for Active Directory

If you have a firewall between your View environment and your Active Directory server, you must make sure that all of the necessary ports are opened.

For example, View Connection Server must be able to access the Active Directory Global Catalog and Lightweight Directory Access Protocol (LDAP) servers. If the Global Catalog and LDAP ports are blocked by your firewall software, administrators will have problems configuring user entitlements.

See the Microsoft documentation for your Active Directory server version for information about the ports that must be opened for Active Directory to function correctly through a firewall.

Overview of Steps to Setting Up a View Environment

6

Complete these high-level tasks to install View and configure an initial deployment.

Table 6-1. View Installation and Setup Check List

Step	Task
1	Set up the required administrator users and groups in Active Directory. Instructions: *View Installation* and vSphere documentation.
2	If you have not yet done so, install and set up ESXi hosts and vCenter Server. Instructions: VMware vSphere documentation.
3	If you are going to deploy linked-clone desktops, install View Composer, either on the vCenter Server system or on a separate server. Also install the View Composer database. Instructions: *View Installation* document.
4	Install and set up View Connection Server. Also install the Events database. Instructions: *View Installation* document.
5	Create one or more virtual machines that can be used as a template for full-clone desktop pools or as a parent for linked-clone desktop pools. Instructions: *Setting Up Desktop and Application Pools in View*.
6	(Optional) Set up an RDS host and install applications to be remoted to end users. Instructions: *Setting Up Desktop and Application Pools in View*.
7	Create desktop pools, application pools, or both. Instructions: *Setting Up Desktop and Application Pools in View*.
8	Control user access to desktops. Instructions: *Setting Up Desktop and Application Pools in View*.
9	Install Horizon Client on end users' machines and have end users access their remote desktops and applications. Instructions: Horizon Client documentation at https://www.vmware.com/support/viewclients/doc/viewclients_pubs.html.
10	(Optional) Create and configure additional administrators to allow different levels of access to specific inventory objects and settings. Instructions: *View Administration* document.
11	(Optional) Configure policies to control the behavior of View components, desktop and application pools, and end users. Instructions: *Setting Up Desktop and Application Pools in View*.

Table 6-1. View Installation and Setup Check List (Continued)

Step	Task
12	(Optional) Configure View Persona Management, which gives users access to personalized data and settings whenever they log in to a desktop. Instructions: *Setting Up Desktop and Application Pools in View*.
13	(Optional) For added security, integrate smart card authentication or a RADIUS two-factor authentication solution. Instructions: *View Administration* document.

Index

S

T

U

V

W