



Inside **OUT**

The ultimate, in-depth reference
Hundreds of timesaving solutions
Supremely organized, packed
with expert advice

Virtualizing Desktops & Apps with Windows Server 2012 R2

Virtualizing Desktops & Apps with Windows Server 2012 R2 Inside Out

Byron Wright
Brian Svidergol

PUBLISHED BY
Microsoft Press
A division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399
Copyright © 2015 by Microsoft Corporation

All rights reserved.

No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Control Number: 2014952206
ISBN: 978-0-7356-9721-8

Printed and bound in the United States of America.

First Printing

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Support at mspinput@microsoft.com. Please tell us what you think of this book at <http://aka.ms/tellpress>.

This book is provided "as-is" and expresses the author's views and opinions. The views, opinions and information expressed in this book, including URL and other Internet website references, may change without notice.

Some examples depicted herein are provided for illustration only and are fictitious. No real association or connection is intended or should be inferred.

Microsoft and the trademarks listed at <http://www.microsoft.com> on the "Trademarks" webpage are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

Acquisitions Editor: Alison Hirsch

Project Editor: Rosemary Caperton

Editorial Production: nSight, Inc.

Technical Reviewer: Todd Meister; Technical Review services provided by Content Master, a member of CM Group, Ltd.

Copyeditor: Ann Weaver

Indexer: Lucie Haskins

Cover: Twist Creative • Seattle



Contents at a glance

Chapter 1		Chapter 8	
Desktop and application virtualization	1	Planning and deploying session-based	
Chapter 2		virtual desktops	387
Planning and implementing user		Chapter 9	
state virtualization	31	Configuring RemoteApp programs and	
Chapter 3		client connectivity	451
Configuring Client Hyper-V	95	Chapter 10	
Chapter 4		Planning and implementing pooled and	
Planning and implementing App-V	157	personal virtual desktops	481
Chapter 5		Chapter 11	
Planning and deploying App-V clients	213	Implementing Remote Access for VDI	541
Chapter 6		Chapter 12	
Managing and administering		Performance and Health Monitoring of	
Application Virtualization	277	Virtual Desktop Infrastructure	565
Chapter 7		Index	591
Application sequencing	319		

This page intentionally left blank



Table of contents

	Introduction	xiii
	System requirements	xiii
	Acknowledgments	xiv
	Free ebooks from Microsoft Press	xiv
	Errata, updates, & book support	xiv
	We want to hear from you.	xv
	Stay in touch.	xv
Chapter 1	Desktop and application virtualization	1
	Overview of virtualization technologies	1
	User state virtualization.	3
	Application virtualization	4
	Desktop Virtualization.	6
	Storage virtualization.	11
	Usage scenarios for desktop virtualization.	12
	Meeting legal and security requirements	13
	Supporting desktop management tasks	14
	Improving application compatibility	15
	Implementing desktop as a service.	16
	Supporting the mobile user experience.	16
	Considerations for implementing virtualization	17
	User experience	17
	Network connectivity	18
	Infrastructure.	20
	Licensing requirements.	22
	Challenges for implementing desktop and application virtualization	24
	Identify virtualization technologies for business needs.	25
	Improve roaming experience for users	26
	Improve performance of apps for mobile users	26
	Provide remote access to apps and data	27

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can improve our books and learning resources for you. To participate in a brief survey, please visit:

<http://aka.ms/tellpress>

	Update apps efficiently	28
	Provide unique apps and improve security	28
Chapter 2	Planning and implementing user state virtualization	31
	Understanding user state and user profiles	31
	User profile creation	32
	User profile content	34
	Understanding and planning user state virtualization	35
	Assess user data requirements	38
	Assess user settings requirements	39
	Evaluate compatibility considerations	40
	Evaluate infrastructure and manageability requirements	40
	Evaluate usage scenario considerations	41
	Configuring user state virtualization technologies	44
	Configuring roaming user profiles	44
	Mandatory user profiles	52
	Configuring Folder Redirection	53
	Configuring Offline Files	59
	Using the Primary Computer setting	66
	Enabling user profile disks	70
	Configuring User Experience Virtualization	70
	UE-V architecture	72
	Comparing user state virtualization options	75
	Preparing to deploy UE-V	77
	Deploying the UE-V agent	78
	Managing the UE-V agent	81
	Managing default settings location templates	84
	Creating and managing custom settings location templates	87
Chapter 3	Configuring Client Hyper-V	95
	Understanding Hyper-V	95
	Client Hyper-V architecture	97
	Installing Client Hyper-V	100
	Hyper-V management tools	104
	Managing virtual switches	108
	Creating virtual machines	112
	Virtual machine settings	121
	Generation 2 virtual machines	124
	Controlling virtual machines	125
	Managing virtual machine files	130
	Processing	135
	Dynamic memory	137
	Integration services	138
	Managing virtual hard disks	141
	Virtual hard disk formats	141
	Fixed and dynamically expanding disks	143
	Differencing disks	145

	Creating virtual hard disks	147
	Editing virtual hard disks	148
	Storage Quality of Service	150
	Managing checkpoints	151
	How checkpoints are created	152
	Using checkpoints	153
	Considerations for using checkpoints	155
Chapter 4	Planning and implementing App-V	157
	Overview of App-V	157
	Benefits of App-V	157
	Differences between standard and virtualized applications	159
	Placing and functionality of the virtualization engine	160
	Application virtualization infrastructure	161
	App-V application life cycle	161
	App-V technologies	177
	App-V deployment models	184
	Planning App-V infrastructure	187
	App-V infrastructure requirements	187
	App-V deployment possibilities	189
	Service disruption impact	193
	Functional and physical placement	194
	Sizing and performance	195
	High availability for App-V	197
	Disaster recovery	199
	Deploying App-V infrastructure	200
	App-V infrastructure requirements	201
	Installing management databases	201
	App-V Management Server configuration	204
	App-V publishing server deployment and configuration	208
	App-V for Remote Desktop Services client	210
	Integrating App-V with System Center Configuration Manager	211
Chapter 5	Planning and deploying App-V clients	213
	Overview of App-V client configuration	213
	App-V desktop client	213
	How the App-V client accesses applications	215
	Storage locations for App-V client data	216
	Planning for App-V 5.0 shared content store	220
	Methods for deploying the App-V client	224
	App-V client for Remote Desktop	227
	Installing and configuring the App-V client	227
	Prerequisites for App-V client installation	228
	Installing the client by using Configuration Manager	230
	Installing the App-V for Remote Desktop Services client	247
	Configuring the App-V client for stand-alone mode	250
	App-V and Virtual Desktop Infrastructure (VDI) solution	252

	App-V 5.0 and third-party production integration	252
	Benefits of App-V and VDI integration.	253
	Building images with App-V.	254
	Managing App-V client properties	259
	Managing virtual applications	264
	Managing file type associations	265
	Managing server connections.	266
	Using Windows PowerShell to configure the App-V client.	267
	Using Group Policy to manage the App-V client	270
	Autoload	273
	Registry settings for the App-V client.	273
Chapter 6	Managing and administering Application Virtualization	277
	Using the Application Virtualization Management Console.	277
	Managing App-V administrators.	281
	Registering and unregistering servers	282
	Managing application packages	284
	Connection groups	287
	Managing management servers by using Windows PowerShell.	289
	Modifying and upgrading published applications	297
	Update an application	297
	Copy access and configuration	303
	Update a connection group	304
	Remove applications	305
	Edit the default configuration for a package.	306
	Exporting the configuration.	307
	Assignment of applications	308
	Naming conventions	309
	Enabling scripting for dynamic configuration.	310
	App-V reporting	312
	How App-V reporting works	312
	Data collected by App-V reporting.	314
	App-V client configuration for reporting.	315
	Generating App-V reports	317
Chapter 7	Application sequencing	319
	Overview of application sequencing	319
	App-V Sequencer	319
	The sequencing process	321
	Items to document in a recipe.	322
	Portions of a sequenced application	324
	Planning for application sequencing.	326
	Sequencer configuration.	326
	Best practices for application installation	330
	Best practices for package configuration.	331
	Applications that cannot be sequenced.	332
	Sequencing an application	333

	Types of applications that can be sequenced	333
	Preparing for sequencing	335
	Sequencing tasks	337
	Customizing the package	344
	Package editor	345
	Windows PowerShell	351
	Deploying Office 2013 by using App-V	351
	Advanced application sequencing	357
	Package accelerators	357
	Creating a package accelerator	358
	Options for updating packages	378
	Sequencing for connection groups	380
	Dynamic configuration and targeted scripting	381
Chapter 8	Planning and deploying session-based virtual desktops	387
	Understanding RDS	387
	Comparing RDS and the Remote Desktop feature	389
	RDS architecture	390
	Connecting to virtual desktops and RemoteApp programs	392
	RDS functionality that enhances the client experience	395
	RemoteFX	397
	Remote Desktop Connection configuration options	399
	RDS licensing	401
	Planning infrastructure for session-based desktops	403
	Assessing RDS infrastructure requirements	403
	Planning for the RD Session Host role service	406
	Planning for the RD Connection Broker role service	408
	Planning for the RD Web Access role service	410
	Planning for preserving user state	411
	Deploying session-based virtual desktops	415
	Understanding the session-based desktop deployment process	415
	Understanding session collections	422
	Configuring session collections	426
	Configuring RD Licensing servers	434
	Understanding high availability for RDS	438
	Understanding load balancing	439
	High availability for RD Session Host servers	441
	High availability for the RD Connection Broker role service	443
	High availability for the RD Web Access role service	447
	High availability for the RD Licensing role service	447
Chapter 9	Configuring RemoteApp programs and client connectivity	451
	Publishing and configuring RemoteApp programs	451
	Understanding RemoteApp programs	453
	Installing applications on RD Session Host servers	454
	Publishing RemoteApp programs	457
	Configuring RemoteApp programs	458

- Configuring and managing client connections to RDS 461
 - Configuring RemoteApp and Desktop Connections 462
 - Customizing RD Web Access 466
 - Understanding device redirection 467
 - Understanding printer redirection 469
 - Managing connections 470
- Configuring certificates and single sign-on 472
 - Understanding RDS certificates 473
 - Requesting and configuring RDS certificates 475
 - Understanding single sign-on 478

Chapter 10 Planning and implementing pooled and personal virtual desktops . 481

- Understanding pooled and personal virtual desktops 481
 - Using pooled virtual desktops 483
 - Using personal virtual desktops 484
 - Comparing virtual desktop options 484
 - High availability for pooled virtual desktops 485
 - High availability for personal virtual desktops 486
- Planning and creating virtual desktop templates 489
 - Selecting an operating system 490
 - Activating the operating system 491
 - Updating applications and the operating system 494
 - Eliminating the system partition 495
- Optimizing operating system configuration 497
 - Optimizing App-V 500
 - Optimizing antivirus software 500
 - Using Sysprep to prepare a virtual desktop template 501
- Planning storage for pooled and personal virtual desktops 503
 - Using local storage for pooled and personal virtual desktops 503
 - Using a SAN for pooled and personal virtual desktops 505
 - Using scale-out file servers for pooled and personal virtual desktops 506
 - Using additional Windows Server 2012 R2 storage technologies 508
- Capacity planning for pooled and personal virtual desktops 511
 - Capacity planning for storage 512
 - Capacity planning for memory 513
 - Capacity planning for networking 514
 - Capacity planning for processing 515
 - Capacity planning example 517
- Implementing pooled and personal virtual desktops 518
 - Deploying RD Virtualization Host servers 519
 - Understanding user profile disks for VM-based virtual desktops 523
 - Creating a virtual desktop collection 524
 - Updating pooled virtual desktops 533
 - Implementing RemoteApp for Hyper-V 535

Chapter 11 Implementing Remote Access for VDI 541

- Extending VDI outside the organization 541

	Why remote access is important for VDI	541
	Methods for securing remote access to VDI	542
	Network configuration for RD Gateway	543
	Configuration options for RD Gateway	548
	Controlling RD Gateway access	557
	Overview of controlling RD Gateway access	557
	RD CAPs	557
	RD RAPs	559
	Central RD CAP store	560
	Integrating Microsoft Azure Multi-Factor Authentication	561
Chapter 12	Performance and Health Monitoring of Virtual Desktop Infrastructure	565
	Monitoring desktop and application virtualization	565
	Understanding monitoring for desktop and application virtualization	565
	Event monitoring for desktop and application virtualization	566
	Performance monitoring tools for desktop and application virtualization	569
	Using Process Monitor to identify application issues	571
	Using Operations Manager for monitoring	572
	Parts of an Operations Manager implementation	573
	Understanding management packs and overrides	579
	Management packs for monitoring application and desktop virtualization	581
	Installing management packs	583
	Monitoring desktop virtualization infrastructure	585
	Understanding resource bottlenecks	585
	Considerations for monitoring desktop virtualization	586
	Monitoring RD Session Host server performance	587
	Optimizing RD Session Host server performance	588
	Index	591

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can improve our books and learning resources for you. To participate in a brief survey, please visit:

<http://aka.ms/tellpress>

This page intentionally left blank

Introduction

If you want to learn about using Microsoft technologies to implement application and desktop virtualization, then this book is for you. If you need to support users who roam among multiple computers and platforms, then this book will help you identify, select, and implement options for user state virtualizations. If you are implementing application virtualization by using Microsoft Application Virtualization (App-V), then this book will help you learn how to design and configure an App-V deployment. It also will teach you how to sequence applications for deployment. If you want to implement virtual desktops to simplify hardware upgrades and support mobile users, then this book will teach you about both session-based virtual desktops and virtual machine–based (VM-based) virtual desktops. It also will teach you how to secure remote access to virtual desktops by using Remote Desktop Gateway (RD Gateway). Finally, to ensure that your application and desktop virtualization meets the performance requirements of your users, you will learn about monitoring application and desktop virtualization.

This book assumes that you have a working knowledge of Windows client and server administration or have access to that information. So, we assume that you understand basic facts like how to connect a computer to a domain, how to create and apply Group Policy, and how to work with management tools. This book focuses on implementing the application and desktop virtualization technologies that layer on top of Microsoft Windows Server 2012 R2. Each of the technologies covered in the book is examined from a real-life perspective. This book provides examples and recommendations for implementation.

System requirements

The following are the recommended system requirements necessary to implement the technologies described in this book in a development environment:

- A processor with Second Level Address Translation (SLAT) support
- Windows 8.1 (Enterprise edition recommended, Professional edition minimum)
- Windows Server 2012 R2
- Microsoft Desktop Optimization Pack (MDOP)
- At least 8 GB of random access memory (RAM)
- A hard drive with at least 128 GB of free space
- A CD-ROM or DVD-ROM drive
- A mouse or other pointing device
- A 1024 x 768 or greater monitor display

Acknowledgments

We'd like to thank Alison Hirsch at Microsoft Press for helping us get started on this book and ensuring that we were on the right track. We'd also like to thank Rosemary Caperton at Microsoft Press for getting us to the finish line. Finally, we'd like to thank Michael McMann, Brad Joseph, Phil Helsel, and Tony Jamieson from Microsoft Learning Experiences for working with us on a new development process.

Byron would personally like to thank Tracey, Samantha, and Michelle for tolerating the occasional late night required to get a large project like this done.

Brian would personally like to thank his wife Lindsay, his son Jack, and his daughter Leah for supporting him in another project. He would also like to thank Elias Mereb, Charles Pluta, and Bob Clements for their continued support and expertise.

Free ebooks from Microsoft Press

From technical overviews to in-depth information on special topics, the free ebooks from Microsoft Press cover a wide range of topics. These ebooks are available in PDF, EPUB, and Mobi for Kindle formats, ready for you to download at:

<http://aka.ms/mspressfree>

Check back often to see what is new!

Errata, updates, & book support

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates to this book—in the form of a list of submitted errata and their related corrections—at:

<http://aka.ms/VirtApps/errata>

If you discover an error that is not already listed, please submit it to us at the same page.

If you need additional support, email Microsoft Press Book Support at *mspinput@microsoft.com*.

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to *<http://support.microsoft.com>*.

We want to hear from you

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

<http://aka.ms/tellpress>

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

Stay in touch

Let's keep the conversation going! We're on Twitter: *<http://twitter.com/MicrosoftPress>*.

This page intentionally left blank



Overview of App-V	157	Planning App-V infrastructure	187
Application virtualization infrastructure	161	Deploying App-V infrastructure	200

Virtualization usage has been expanding throughout the last several years. Although many IT administrators are now familiar with desktop and server virtualization, application virtualization is a lesser-known technology that is expanding rapidly. Simply put, application virtualization is a technology that delivers applications to computing devices in isolated containers without the need to perform a standard application installation on each computing device.

App-V provides a new way to deliver applications to users, a new way to centrally manage applications, and additional application capabilities for complex environments. Virtualizing applications, when appropriately planned and implemented, is an effective supplement to desktop and server virtualization and can help organizations move closer to cloud computing.

Overview of App-V

Microsoft Application Virtualization (App-V) is Microsoft's application virtualization technology. It is part of the Microsoft Desktop Optimization Pack (MDOP). MDOP is available to Microsoft Software Assurance customers and to Microsoft MSDN subscribers. It is a suite of desktop optimization applications that help IT administrators deploy, manage, and troubleshoot desktop computing environments. The newest version, at the time of this writing, is App-V 5.0 Service Pack 3.

Benefits of App-V

You know that App-V is Microsoft's application virtualization technology, and you have an idea of how to access it. But, to fully appreciate it, you should have a good understanding of its capabilities. The core capabilities that you can deliver with App-V are the following:

- **Run multiple versions of applications at the same time** You can use App-V to run different versions of applications concurrently on the same computer. For example, it's possible to run Internet Explorer 9, Internet Explorer 10, and Internet Explorer 11 concurrently if they all are set up as App-V applications; otherwise, you can't run multiple versions of Internet Explorer side by side on the same computer. It's also possible to

use App-V in conjunction with Remote Desktop Services (RDS). This allows users to run applications side by side on Remote Desktop Session Host (RD Session Host) servers.

- **Minimize application conflicts** Sometimes two or more applications conflict with one another because of dynamic-link library (DLL) or application programming interface (API) conflicts. When you deliver these applications as App-V applications, however, there isn't a conflict because each App-V application runs in its own isolated environment.
- **Simplify application removal** App-V applications do not install locally, which makes it easier to remove them. Clean removal isn't always possible with applications that install directly on Windows-based clients, even if an application has been designed to remove all files and settings when it uninstalls. Virtualized applications are removed easily after a user signs out from a computer.
- **Simplify application upgrades** Instead of upgrading a locally installed application on all computers in your organization with a hotfix, service pack, or new version, the modular nature of virtualized applications means that you can replace one version of an application with an updated version with less effort. You only need to update an application on the App-V server(s), and clients receive the latest version of the application from the publishing server the next time they launch the application.
- **Minimize license-compliance risks** App-V has application metering functionality that enables you to detect every use of a virtualized application to ensure license compliance.
- **Scale infrastructure** You can add publishing servers to an App-V deployment as necessary to ensure that service levels are maintained as demand grows.
- **Take advantage of client hardware resources** Even though App-V applications do not install locally, they can use a local computer's processor and RAM resources. In environments where client computers have adequate hardware resources, this can create a better experience for users than running applications on an RD Session Host server would.
- **Allow users to use roaming applications** If applications stream with App-V rather than install locally, users can sign in to any computer that has an App-V client and quickly access their applications. You also can configure App-V to work with user virtualization to allow users to have application settings and data for App-V applications that roam across client computers.
- **Give users quick access to their applications** Imagine a scenario in which a user needs to have an application installed. In many environments, a routine application installation requires several steps, such as copying the installation files to the client

computer, manually performing the installation, manually performing application updates, and cleaning up the installation files. With App-V, you can just add the users to a group, and they can immediately run the application through App-V.

Differences between standard and virtualized applications

In a traditional IT environment, applications are installed on each computing device. For this discussion, we'll refer to the traditional installed application as a standard application. During a standard application installation, the computing device is modified to meet the application's requirements. For example, a common installation of a standard application would include the following steps:

The application's files are copied During the installation, the installer copies the necessary application files to a specified folder. During this process, the files often are uncompressed first.

The registry is modified to support the application Most applications add to the registry to support the installation and running of the application. In addition, it's common for existing registry keys and values to be modified to support automatic application startup on boot and to set any required dependencies.

DLL files are registered Often, .dll files must be registered during the installation.

Permissions are configured You often need to set permissions to configure which user(s) can run the application, which profiles the shortcuts are added to, and which users can modify configuration files.

Shortcuts are added to the device Many applications add a shortcut to the Start screen, desktop, and taskbar. Additionally, it is common for applications to add a tray icon for quick user access.

App-V applications, which are virtualized, aren't installed onto computing devices. Instead, they are packaged for deployment. In App-V, packaging an application for deployment is known as sequencing an application. Later in this chapter, we'll discuss sequencing further. For now, let's examine the differences between an App-V application deployed to a computing device and a standard application installed locally on a computing device. For applications deployed with App-V, the following characteristics highlight the different methods by which applications interact with the computing device:

App-V applications run in their own isolated environment This reduces application conflicts and application crashes impacting other applications or the operating system and provides the foundation to allow multiple versions of applications to run at the same time.

App-V applications use a virtual registry This reduces installation difficulties and application conflicts and improves the stability of a client computer because applications aren't sharing a single registry.

App-V applications use virtual file systems This reduces conflicts by ensuring that applications don't overwrite shared files.

App-V applications use virtual services This also reduces conflicts because virtual services do not have dependencies on other non-virtual services and provide isolation from other virtual services.

App-V applications are installed by being packaged and delivered to computing devices. App-V applications only need to be updated once, by updating the application's package on the App-V server. By reducing the number of times you have to install and update an application, you can greatly reduce the number of hours required to manage your application infrastructure.

Placing and functionality of the virtualization engine

Virtual applications require access to resources on a host computer. Access typically uses a system request that a virtualization engine needs to intercept. The engine provides functions for capturing an application's system call and manipulating it where needed.

Interception in user mode

Applications run in user mode and perform operations on system services that reside in kernel mode. When an interception occurs in user mode, the virtualization engine must be placed over the native API layer. Some applications call the functions of the native API directly instead of by using Windows API.

One advantage of the user mode strategy is that the virtualization engine doesn't have to filter between system calls of different applications. The virtualization engine only works with an application that it built for the virtual environment.

Interception in kernel mode

In contrast, if the virtualization engine operates in kernel mode, it can intercept all system calls before they reach Windows executive services. The advantage is that it is easy to loosen the isolation of applications and let them share the same virtual environment. This enables you to create dependencies between packages, similar to working with middleware or plug-ins. However, the agent that builds the virtual environment must install natively on a computer to gain the required privileges for an interception in kernel mode. In this case, all applications that run on the host machine are affected. Furthermore, the virtualization layer captures system calls from all applications and must filter the processes to respond to every call correctly.

App-V implements a hybrid approach by identifying the locations of resources that the application requires. While the application is running, the virtualization engine ensures that function calls are modified only if they request a path inside the virtual file system or a key inside the virtual registry. When the engine runs in kernel mode, it also must check which package to consult to find the rules, because it will receive function calls from several running applications that belong to different packages.

The virtualization layer must intercept and redirect requests to the file system and registry to virtual counterparts that contain the files and keys that belong to a certain application. To decide which function calls to intercept and which to handle as usual, the virtualization layer needs rules. These rules generate when the App-V Sequencer, which collects all the files and registry keys that the installer creates or modifies, monitors an application's installation. Then, at the application's runtime, the virtualization engine ensures that function calls are modified only if they request a path inside the virtual file system or a key inside the virtual registry.

Application virtualization infrastructure

The infrastructure technologies of an App-V deployment are extensive, based on the deployment model that you choose. They work together to provide the complete suite of App-V technologies. It is important to familiarize yourself with all of the technologies, the typical life cycle of a virtual application, the deployment models, and some of the characteristics of packages and content packages.

App-V application life cycle

To effectively manage your virtual application infrastructure, you need to plan for the life cycle of your virtual applications. Without an effective life cycle, you may end up with application sprawl—a situation in which you have too many applications to manage. To avoid this, you should spend ample planning time designing and documenting an operational framework for your virtual application life cycle. In this section, we'll discuss the four phases of the virtual application life cycle: sequencing, publishing and deployment, updating, and termination.

Application sequencing

The App-V Sequencer is one of the primary applications of an App-V deployment. You use it to create virtual application packages. Then, you deploy the packages to your App-V clients. It is important to consider the following before you begin deploying the sequencer and sequencing applications:

- **Prerequisites** If the computer that runs the sequencer isn't running Windows 8 or newer or Windows Server 2012 or newer, then it must have the following software installed prior to installing the sequencer. Note that the App-V client installation automatically will install the Visual C++ prerequisites.
 - Visual C++ Redistributable Package for Visual Studio 2013
 - Visual C++ 2005 Redistributable
 - Microsoft .NET Framework 4
 - Windows PowerShell 3.0
 - Microsoft KB2533623 hotfix
- **Windows 8 or newer or Windows Server 2012 or newer** If the computer that runs the sequencer is running Windows 8 or newer or Windows Server 2012 or newer, it already has the prerequisite software.
- **Match the hardware and software** The computer that runs the sequencer should have a hardware and software configuration that matches the App-V client computers. For example, if all of your App-V client computers run Windows 8.1, you should install the sequencer on similar computer hardware that runs Windows 8.1.
- **Use a virtual machine** When possible, use a virtual machine (VM) as the computer that runs the sequencer. This allows you to take a snapshot of the VM prior to sequencing an application. Then, after you finish sequencing an application, you should revert the VM to the snapshot. This allows you to sequence an application with the same baseline configuration, which minimizes issues. Although the sequencer will allow you to sequence multiple applications without reverting the sequencer to a baseline configuration, it will warn you that you may encounter issues.
- **Multiple sequencers** If you have multiple operating system versions running the App-V client, you seriously should consider having multiple sequencers. This enables you to sequence applications on the same operating system to which you will deploy or stream the applications.

Application publishing and deployment

After you sequence an application, you need to publish it and deliver it to the clients. Publishing a virtual application makes the application available to App-V clients. Before the actual publishing process, you need to be aware of the different methods that you can use based on the type of App-V deployment you have.

Stand-alone deployment model

In the stand-alone model, you need to add the App-V package of the application to clients. You can do this by using the `Add-AppVClientPackage` Windows PowerShell cmdlet. For example, if you have an App-V package named `7-Zip.appv` located at `\\tt-util-01\share\7-Zip.appv`, you can run the following command to add the package:

```
Add-AppVClientPackage -Path \\tt-util-01\share\7-Zip.appv
```

After running the command, the output will show the details of the package added. In fact, the output is the same as if you were to run the `Get-AppVClientPackage -Name 7-Zip` Windows PowerShell command. The output is shown in Figure 4-1.

```
PackageId      : 30252a91-5cc5-419a-bcfb-645c942df001
VersionId     : 5c57388f-d692-49f0-a609-25dcb6e3e690
Name          : 7-Zip
Version       : 0.0.0.1
Path          : \\tt-util-01\share\7-Zip.appv
IsPublishedToUser : False
UserPending   : False
IsPublishedGlobally : False
GlobalPending : False
InUse         : False
InUseByCurrentUser : False
PackageSize   : 4872031
PercentLoaded : 10
IsLoading     : False
HasAssetIntelligence : False
```

Figure 4-1 Adding a package

In the output, notice that the `IsPublishedToUser` property is set to `False`. This is an important detail because while it is set to `False`, the user won't see or be able to use the virtual application. After you've added the package, the client will begin receiving the files that make up the package. The data will be stored locally on the client. You can look at the `%ProgramData%\App-V` directory to see data from the package.

The next step is to publish the application. However, you can't publish an application until the application has been added to the client. Once you are ready to publish, you can use the `Publish-AppVClientPackage` Windows PowerShell cmdlet. For example, if you added a package named `7-Zip`, you can publish it to the client by running the following Windows PowerShell command:

```
Publish-AppVClientPackage -Name 7-Zip
```

Once you run that, the output will be similar to when you added the package. The key difference is that the `IsPublishedToUser` property will be updated to a value of `True`, as shown in Figure 4-2.

```

PackageId      : 30252a91-5cc5-419a-bcfb-645c942df001
VersionId     : 5c57388f-d692-49f0-a609-25dcb6e3e690
Name          : 7-Zip
Version       : 0.0.0.1
Path          : \\tt-util-01\share\7-Zip.appv
IsPublishedToUser : True
UserPending   : False
IsPublishedGlobally : False
GlobalPending : False
InUse         : False
InUseByCurrentUser : False
PackageSize   : 4872031
PercentLoaded : 20
IsLoading     : False
HasAssetIntelligence : True

```

Figure 4-2 Publishing a package

After you publish the package, the application becomes available to the user. If shortcuts are configured for the package, they will begin to be displayed after publishing the package.

Full infrastructure model

In the full infrastructure model, the publishing process is a bit more automated, especially when you are dealing with a large number of virtual applications. The high-level process to publish an application is shown below. Note that the first step involving Group Policy is a one-time step in a new full infrastructure model and would not need to be performed for each application that you want to publish.

1. Create a new Group Policy Object (GPO) and modify the App-V–related GPO settings for your environment. Link it to the computers that have the App-V client software. At a minimum, you should configure an App-V publishing server in the GPO so that App-V clients will automatically be configured for a publishing server.
2. Add the application package on the App-V management portal. Configure the settings based on your environment. At a minimum, you need to ensure that the users have access to the application and that the application is published. You can configure access by right-clicking the application and then clicking Edit Active Directory Access.
3. Publish the application. To publish an application from the management portal, right-click it and then click Publish. You also can publish an application by using Windows PowerShell. For example, to publish an App-V package named 7-Zip, you can run the `Publish-AppvServerPackage -Name 7-Zip` command.
4. Sync the clients or wait for the next automatic sync. To immediately sync a client with an App-V publishing server named TT-UTIL-01, run the `Sync-AppvPublishingServer -Name TT-UTIL-01` Windows PowerShell command.

Configuration Manager model

If you use App-V and Configuration Manager to manage and deliver applications, then the steps to publish an application are different from other App-V models. The following high-level steps describe the process of publishing by using Configuration Manager.

1. In the Configuration Manager console, create a new application.
2. On the General page of the Create New Application Wizard, configure the application type to be Microsoft Application Virtualization (App-V) Client 5.0. Then, browse to the location of the .appv package that you want to publish. Also, ensure that the user and deployment .xml configuration files are in the same location as the .appv file. By default, the name of the user configuration file is <app>_UserConfig.xml. For example, if the name of the .appv file is 7-Zip.appv, then the name of the user configuration file is 7-Zip_UserConfig.xml. By default, the name of the deployment configuration file is <app>_DeploymentConfig.xml. For example, if the name of the .appv file is 7-Zip.appv, then the name of the deployment configuration file is 7-Zip_DeploymentConfig.xml.
3. Complete the Create New Application Wizard by specifying application details or maintaining the default values.
4. Distribute the application to Configuration Manager distribution points so that clients can obtain the application from the nearest distribution point.
5. Deploy the application to clients. You can deploy the application as a streaming application if you have a full infrastructure App-V deployment model. Otherwise, you can opt for the download and execute method. Each method has pros and cons. See <http://technet.microsoft.com/en-us/library/jj822982.aspx> for more information on the two deployment methods.

Application update

One ongoing maintenance task that you'll need to perform is updating applications. An application update, sometimes called an application upgrade, occurs when a software company releases a newer version of the application. Often, companies release newer versions of software to fix security issues or provide new or enhanced functionality. In a standard application deployment, in which applications are installed on every computer, you need to perform the update on every computer. In a virtualized application deployment, in which applications are packaged and delivered by App-V, you only need to update the packaged application on the App-V Sequencer and then update the distribution method with the updated package.

To update an existing App-V application with the App-V Sequencer, perform the following steps:

1. Run the App-V Sequencer on the client computer that you use for sequencing.

2. Click the Modify An Existing Virtual Application Package option, shown in Figure 4-3.

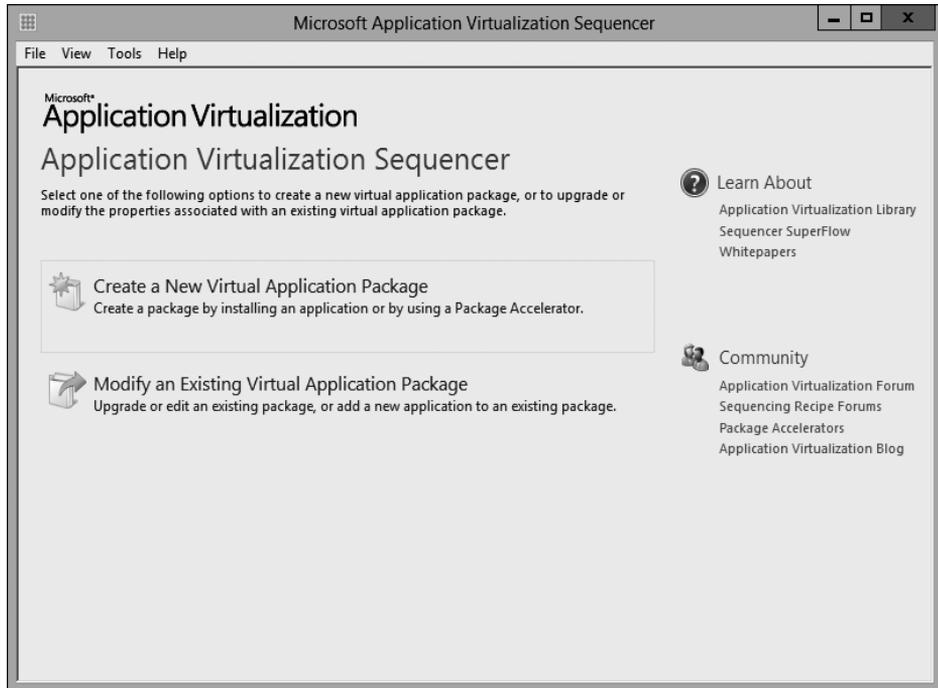


Figure 4-3 App-V Sequencer

3. On the Select Task page, shown in Figure 4-4, keep the Update Application In Existing Package option selected and then click Next.

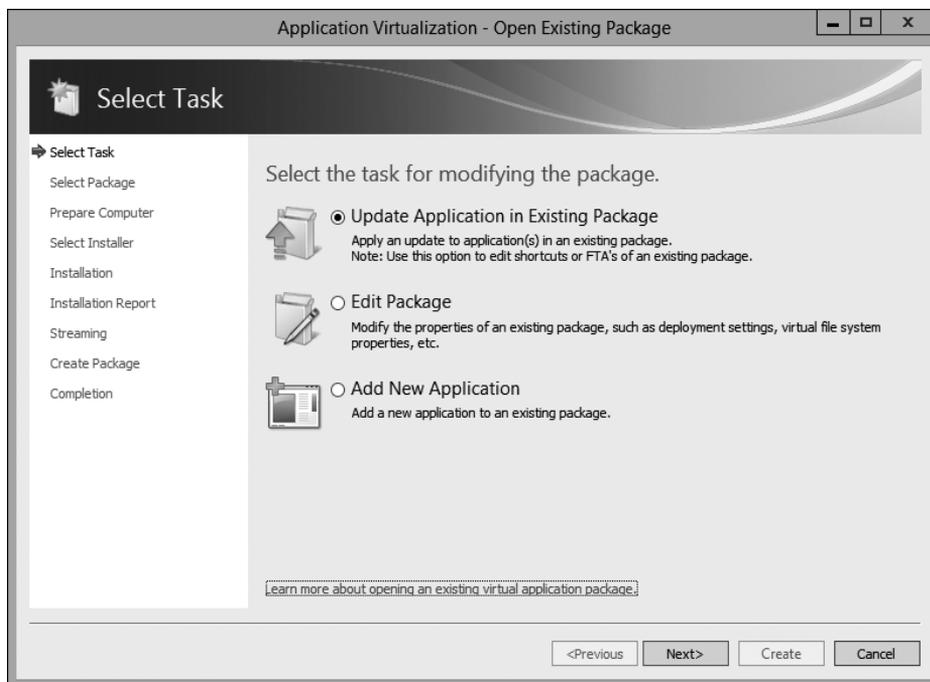


Figure 4-4 App-V Sequencer Select Task page

4. On the Select Package page, shown in Figure 4-5, click Browse and navigate to the existing App-V package file (.appv), click the file, click Open, and then click Next.

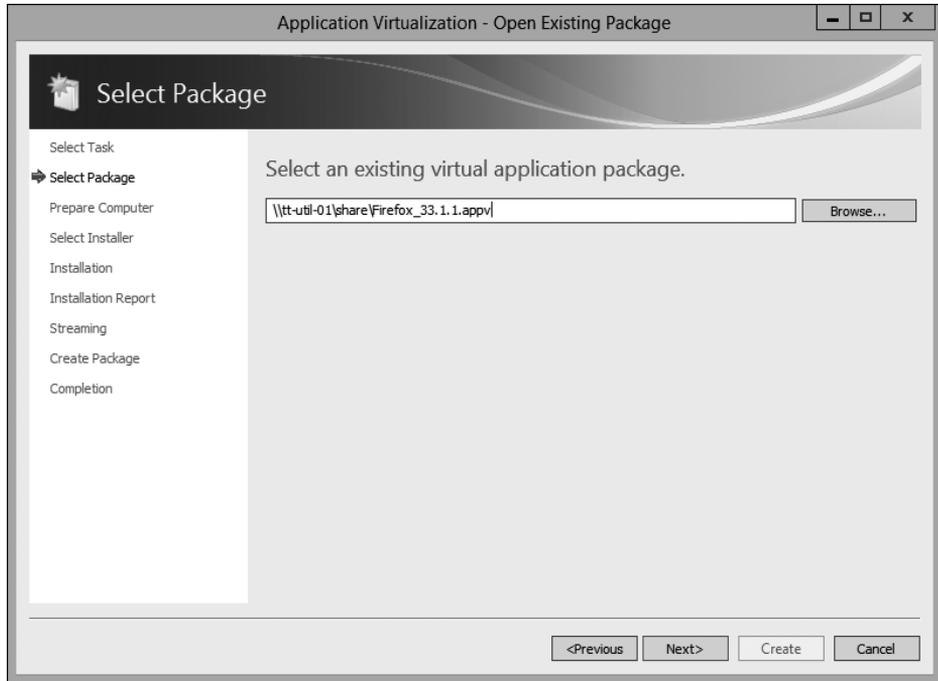


Figure 4-5 App-V Sequencer Select Package page

5. On the Prepare Computer page, shown in Figure 4-6, if the computer is ready to create a package, click Next. If issues are listed, remediate the issues if necessary and then click Refresh until the computer is ready to create a package. Click Next.

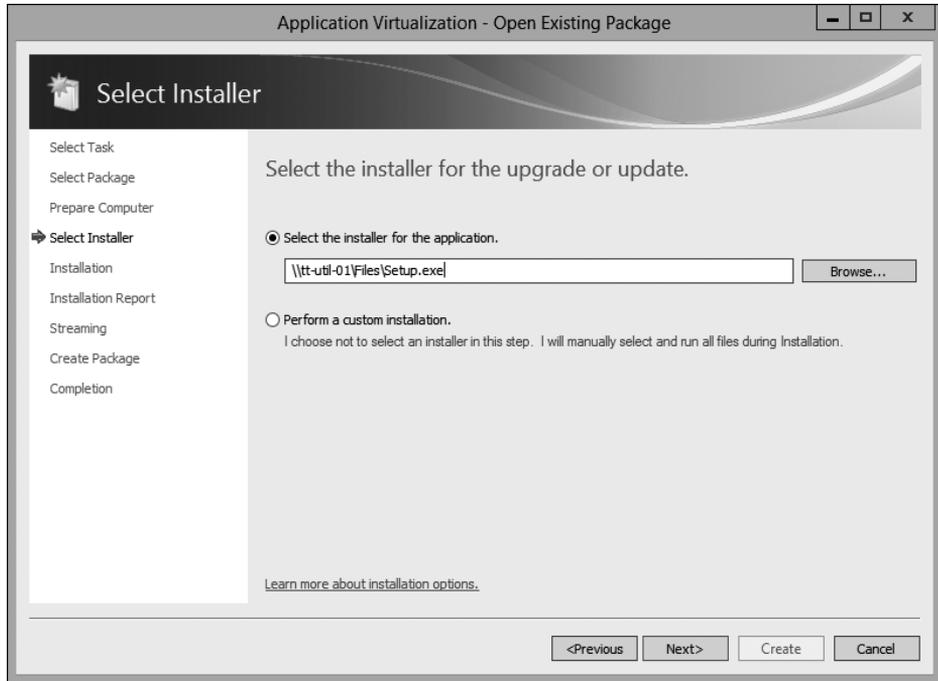


Figure 4-7 App-V Sequencer Select Installer page

7. The application update installation will begin. Update the application based on the installation program. When finished, select the I Am Finished Installing check box, as shown in Figure 4-8, and then click Next.

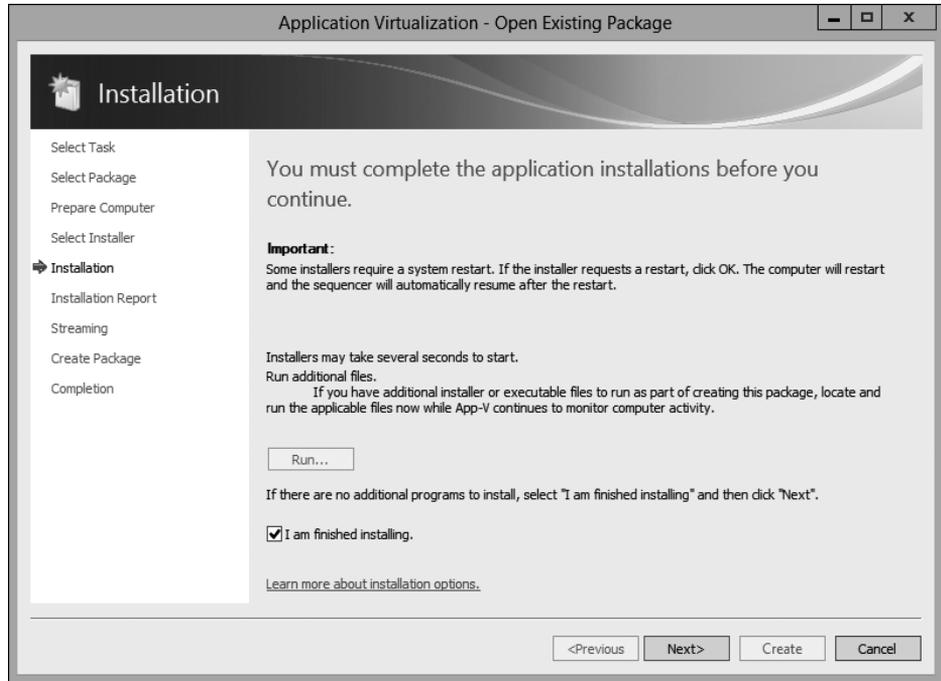


Figure 4-8 App-V Sequencer Installation page

8. On the Installation Report page, shown in Figure 4-9, if the update was successful, the App-V wizard should report that there weren't any issues detected. Click Next.

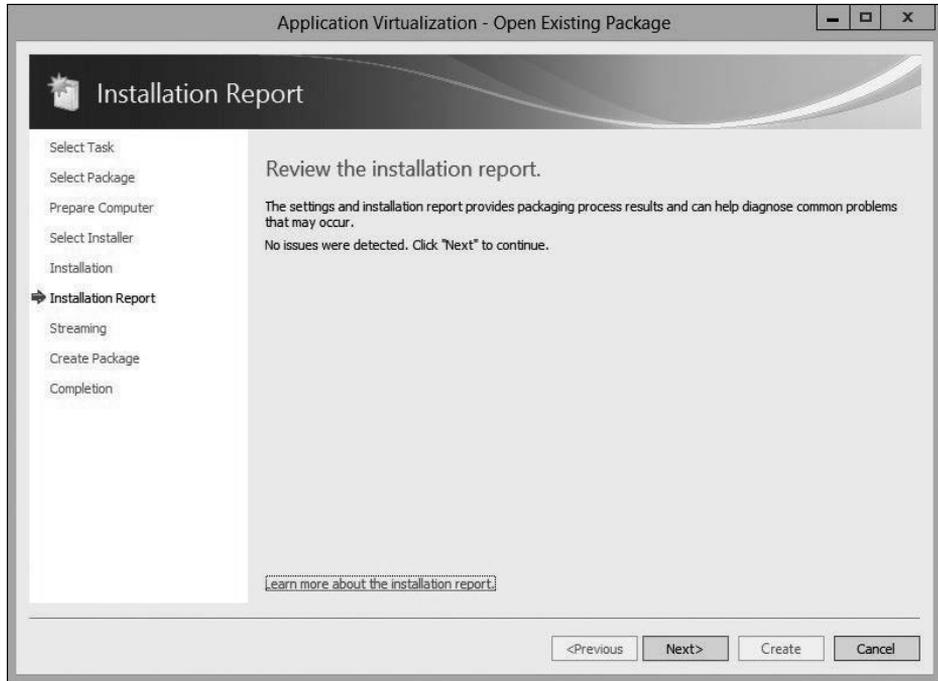


Figure 4-9 App-V Sequencer Installation Report page

9. On the Prepare For Streaming page, shown in Figure 4-10, click Run All. Then, perform any needed first-run application configurations. When finished, close the updated application and then click Next in the App-V wizard.

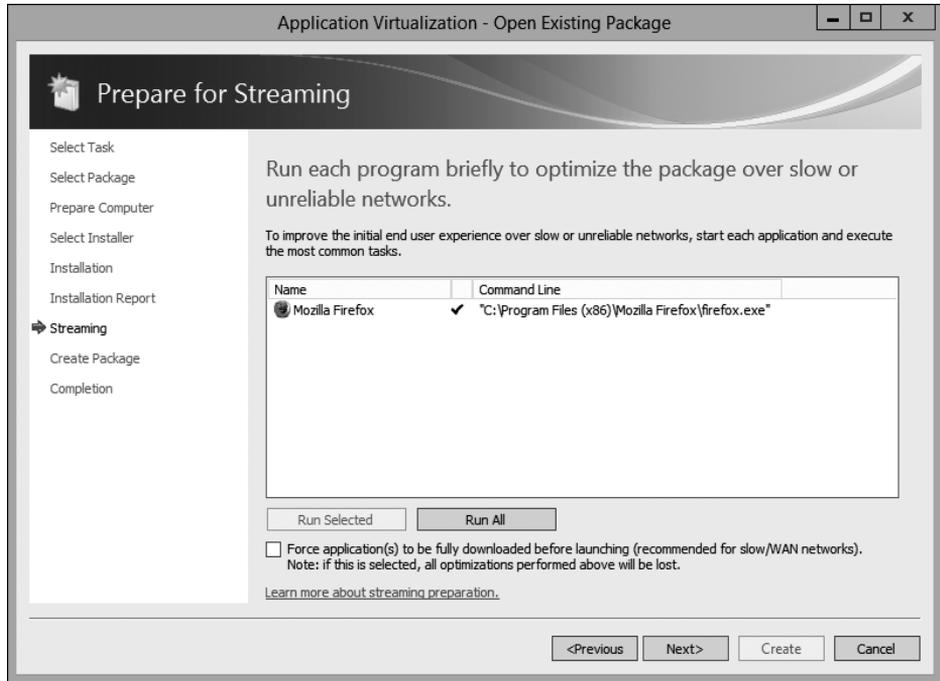


Figure 4-10 App-V Sequencer Prepare For Streaming page

- On the Create Package page, shown in Figure 4-11, enter a location to save the updated application package and then click Create.

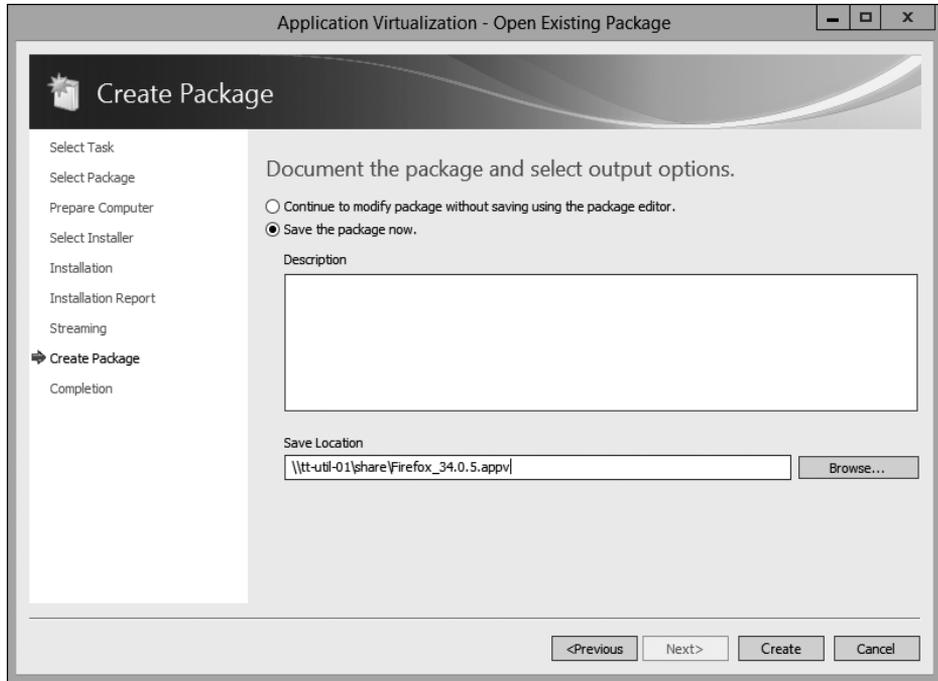


Figure 4-11 App-V Sequencer Create Package page

11. After the package creation completes, click Close. Next, you need to make the updated .appv package available to publishing servers, App-V clients, or Configuration Manager, depending on your App-V deployment model.

Application termination

As part of your routine application maintenance, you'll need to temporarily and permanently remove applications from your environment. Sometimes, you'll need to remove applications from specific users. Other times, you'll need to permanently remove an application from all of your users. You should understand the options available to you for handling these tasks. Let's look at a few scenarios and walk through the process.

Removing an application from a small number of users

There are many situations in which you will need to remove an application from one user or a small group of users. For example, if you were running end user pilot testing for a new application, you may want to remove users after they've completed their testing. There are multiple methods for doing this, but we'll focus on one of the most common methods in an environment with the App-V full infrastructure deployment model. In this model, you should be assigning application access based on Active Directory Domain Services (AD DS) security

groups. When it is time to remove one or more users, the simplest method is to remove those users from the security group. The virtual application remains available in App-V, but only for users with the appropriate access rights. In situations in which you are completely and permanently removing an application from all users, you have a few options:

In the App-V Management Console, you can delete the application by right-clicking it and then clicking Delete, as shown in Figure 4-12.

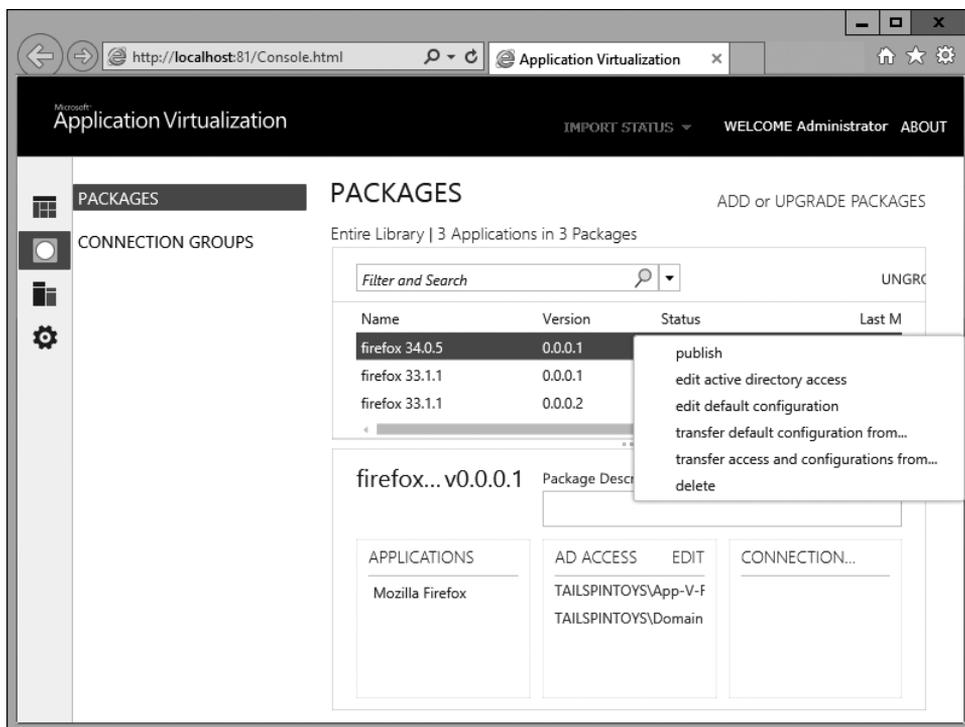


Figure 4-12 App-V Management Console showing Packages workspace with two published applications

From that point on, new App-V clients can't launch the deleted application. Users who already have the application or have used the application will still see it and be able to launch it. You should be aware of a couple of things that cause this. You delete apps from the App-V Management Server. The publishing server, by default, refreshes the list of published applications every 10 minutes. Clients get the list of applications from the publishing server. This 10-minute refresh occurs even if the management server and publishing server run on the same server! Even if you sync the App-V client with the publishing server immediately following an application deletion, it won't matter because the publishing server isn't yet aware that the application was deleted. In most production environments, this 10-minute refresh cycle is

okay. However, if you need to reduce the amount of time between refreshes or manually perform an immediate refresh, you have the following options:

On the publishing server, stop the AppVPublishing application pool, shown in Figure 4-13, in Internet Information Server (IIS).

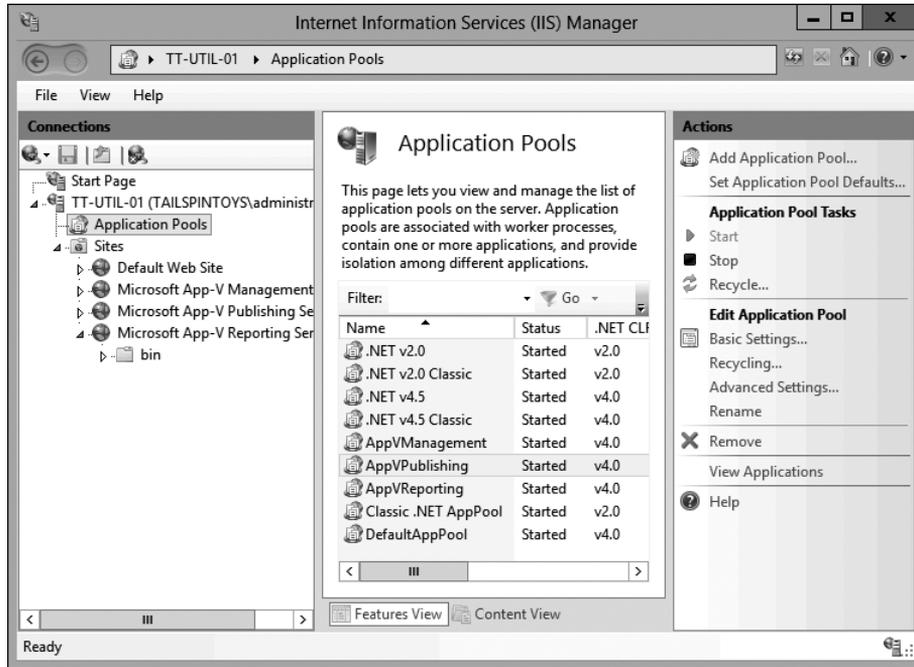


Figure 4-13 IIS application pools for App-V

Then, start the AppVPublishing application pool. This kicks off an immediate refresh of the published apps. Thereafter, you should perform a client sync to complete the removal process. Otherwise, you must wait until the next client sync.

Alternatively, you can adjust the refresh intervals in the registry on the publishing server. To do this, go to the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\AppV\Server\PublishingService key and reduce the value of the PUBLISHING_MGT_SERVER_REFRESH_INTERVAL entry. The default is 600 seconds. You can reduce it to something much smaller, such as 10 seconds, as shown in Figure 4-14.

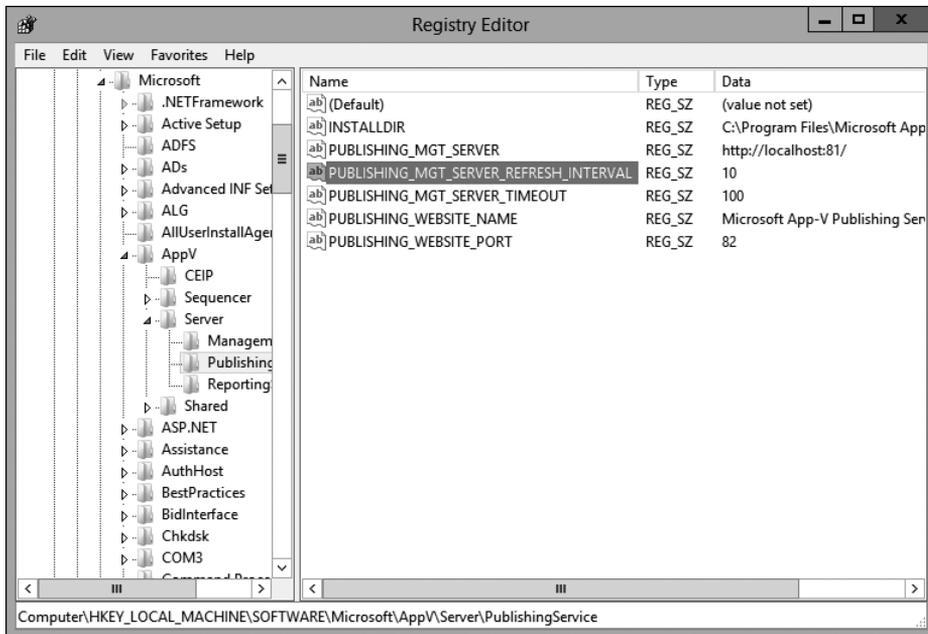


Figure 4-14 App-v publishing registry values

App-V technologies

An App-V 5.0 deployment includes a number of technologies, some of which must be present in all App-V 5.0 models, and some of which are used by only specific App-V deployments. These technologies are the management server, publishing server, management server database, reporting server, reporting server database, Sequencer, and App-V client.

From the Feature Selection page of the Microsoft Application Virtualization (App-V) Server 5.0 Setup Wizard, you can select which servers and databases you want to install, as shown in Figure 4-15.

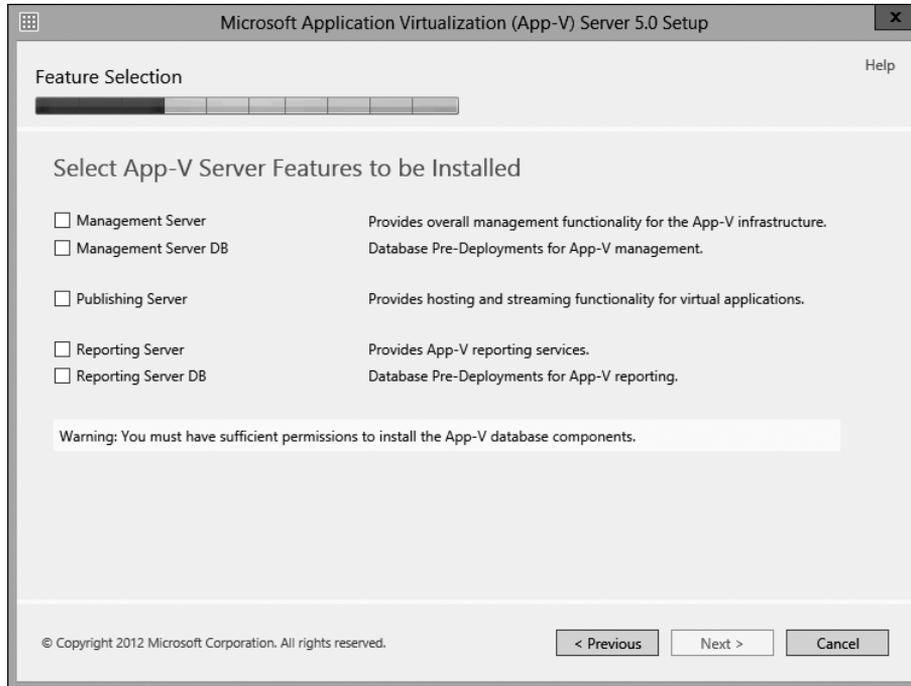


Figure 4-15 The Feature Selection page showing the available App-V server features

Management server and management server database

The management server hosts a Microsoft Silverlight–enabled web application that App-V administrators use to manage an App-V infrastructure. When you use App-V 5.0, all management occurs through the web console or Windows PowerShell, unlike previous versions of App-V. A computer must have Silverlight installed to access the console.

You can use the management server to work with the following objects:

- Packages** You can import packages in the App-V file format, which uses the .appv file extension. You then can publish the packages to App-V publishing servers. You also can use the console to configure package security. In Figure 4-16, the management console shows the Packages workspace with a couple of published applications.
- Connection groups** Connection groups make it possible for virtual applications to interact with one another. You also can use the console to configure security for connection groups.

- **Publishing servers** You can authorize publishing servers. You perform this task from the Servers node. You must specify publishing servers in the *domain\computername* format.
- **Administrators** You can add and manage App-V administrators. Administrators are able to import and publish packages, configure connection groups, and add publishing servers.

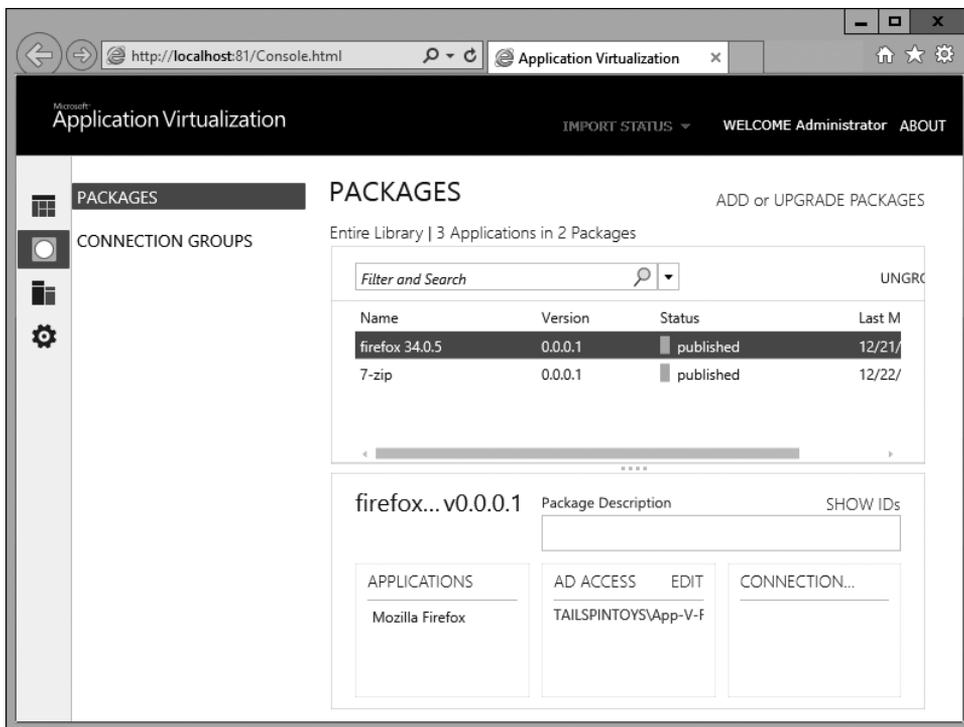


Figure 4-16 App-V Management Console showing Packages workspace with two published applications

The management server database stores the App-V configuration and data settings. The database is hosted on a Microsoft SQL Server instance that is running SQL Server 2008 Service Pack 2 (SP2), SQL Server 2008 R2, or SQL Server 2012. You should install the management server database separately from the management server, and if you do, you need to deploy the database first and then specify its location when you deploy the management server.

Publishing server, sequencer, and App-V client

The publishing server is a web server, built on IIS, which hosts and deploys applications for App-V clients. In previous versions of App-V, a publishing server was called a streaming server.

You deploy App-V applications to a publishing server in the App-V full infrastructure model by using the App-V Management Server console or by using Windows PowerShell.

You can deploy a management server and a publishing server on the same computer. You also can deploy a publishing server on a separate computer as long as you already have deployed an existing management server. This is different from previous versions of App-V, in which it was possible to deploy App-V streaming servers without having deployed a management server.

Previous versions of App-V streamed applications by using the Real-Time Streaming Protocol (RTSP). App-V 5.0 applications stream from a publishing server by using HTTP, HTTPS, or SMB. When you configure a publishing server, you specify a TCP/IP port that is used to stream applications. You subsequently use this port address when you configure Group Policy for use with App-V clients so that clients can access published App-V applications.

Publishing servers are useful in environments that have multiple geographic locations and are connected by wide area network (WAN) links. If you have the App-V full infrastructure model with a management server, publishing server, and database server at the headquarters office, clients outside the headquarters office likely would have a degraded experience compared to having App-V publishing servers locally at their office. In this case, you should deploy publishing servers at branch offices so that clients at those branch offices would be able to receive applications directly from the publishing server in the branch office, rather than across a WAN link from a publishing server in another office.

You can configure a client with the addresses of up to five publishing servers when you use Group Policy, as shown in Figure 4-17.

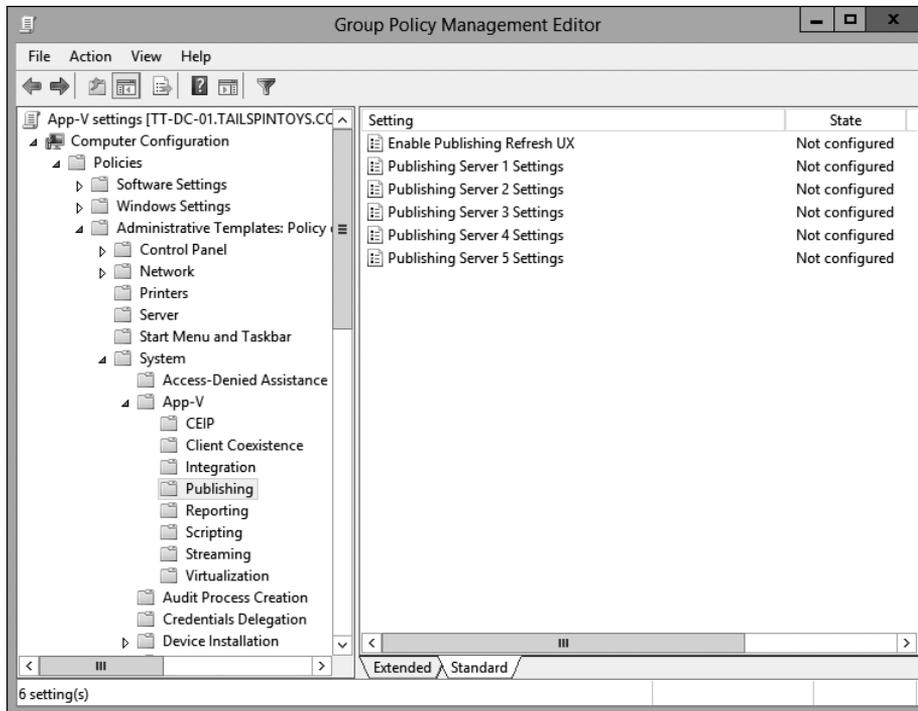


Figure 4-17 GPO settings for publishing

When configuring Group Policy to support clients in branch offices, provide the address of both the local publishing server and a second publishing server to use in case the local publishing server fails. The diagram in Figure 4-18 represents an App-V environment with a headquarters office and two branch offices. Each App-V client is configured to use the local App-V publishing server in its local office and also is configured with a secondary App-V publishing server in an alternate site.

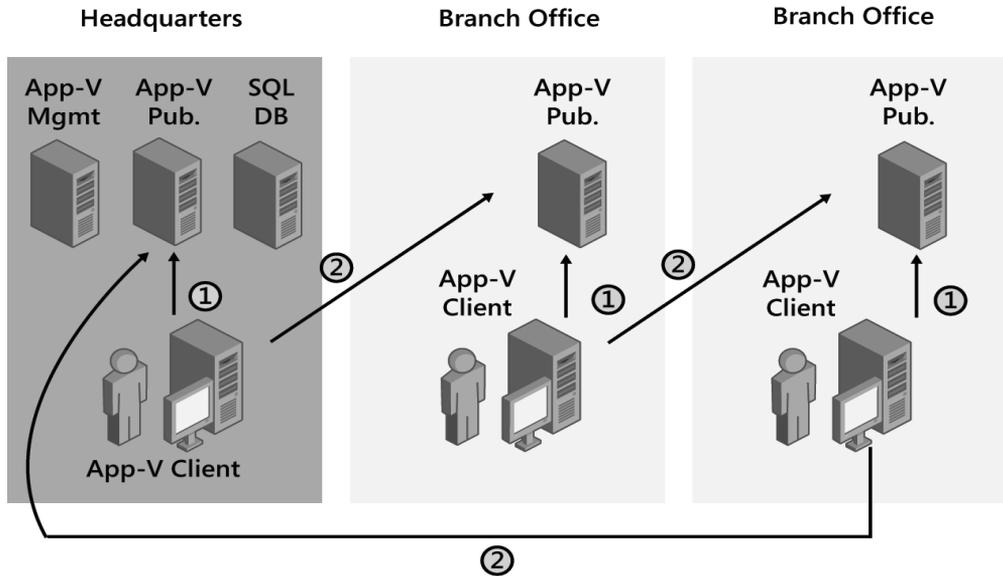


Figure 4-18 App-V environment with multiple publishing servers

You can use the Sequencer to create applications that can run under the App-V client. You should sequence an application on a Sequencer that is configured on the operating system on which the application will run. This helps minimize compatibility issues for your virtualized applications.

The App-V 5.0 Sequencer can install only on computers that run Windows 7 Service Pack 1 (SP1), Windows Server 2008 R2 SP1, Windows 8, Windows 8.1, Windows Server 2012, or Windows Server 2012 R2. In most cases, you should configure a client operating system as a sequencer. This is because most of your virtualized apps likely will be on a client operating system.

The App-V Sequencer has the following requirements:

- Windows 7 Service Pack 1 (SP1), Windows Server 2008 R2 SP1, Windows 8, Windows 8.1, Windows Server 2012, or Windows Server 2012 R2
- Microsoft .NET Framework 4 or 4.5
- Windows PowerShell 3.0 or newer
- Update for Windows KB2533623

The App-V client must be installed on computers before those computers can run virtualized applications. The App-V client supports both x86 and x64 operating system architectures:

- On x86 operating system architecture, only x86 applications are supported.
- On x64 operating system architecture, both x86 and x64 applications are supported.

You need to deploy a separate App-V client on RD Session Host servers. We look at that client in an upcoming section in this chapter titled “App-V for Remote Desktop Services client.”

Reporting server and reporting server database

The App-V reporting server, an optional feature, is the built-in reporting feature that you can use for reporting on virtual application usage. The reporting server records the following information:

- Application use, including launch status, startup times, and shutdown times
- Client information such as the host name, client version, operating system version and type, processor architecture, and operating system service pack level
- Package information, such as the package name, version, source, and the percentage cached

You configure the address of the reporting server when you use App-V Group Policy settings, as shown in Figure 4-19. Clients forward data to this address, which the reporting server then forwards to the reporting server database. You can install the reporting server separately from the reporting server database, though if you do this, you already must have deployed the database on another server.

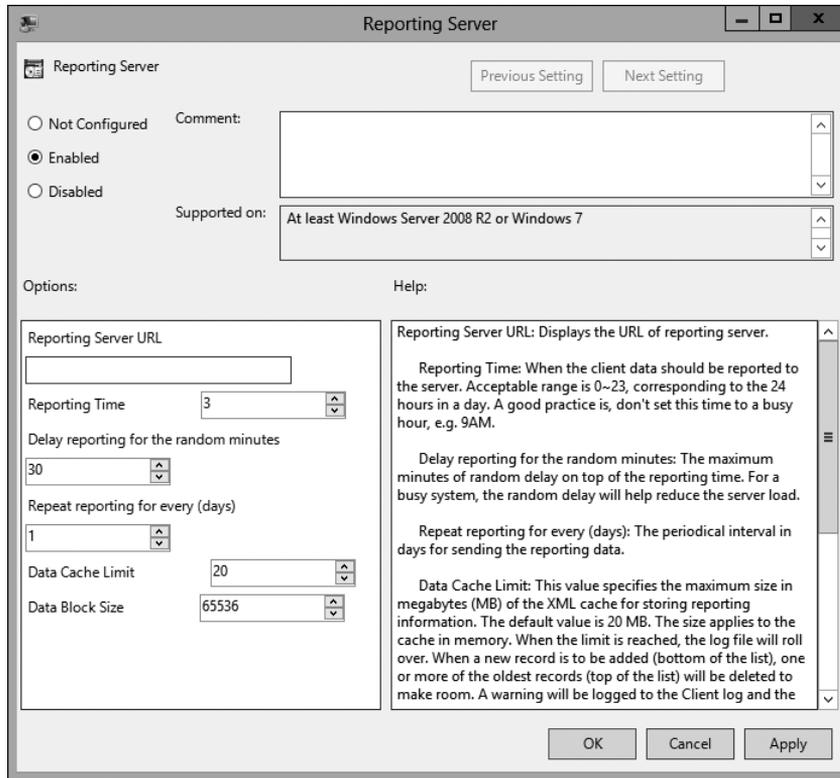


Figure 4-19 GPO settings for reporting

The reporting server database stores all the information that is forwarded to the reporting server. The instance that hosts the reporting server database must meet the same requirements as the instance that hosts the management server database. You can host both databases on the same server. You don't have to install SQL Server Reporting Services to deploy an App-V reporting server, but it is helpful because the App-V product doesn't include report generation.

App-V deployment models

There are three deployment models that you can use to deploy App-V. Each model and its characteristics are described below:

- Full infrastructure model** This is the most complete deployment of App-V and also is the most commonly deployed model. It offers the most services and functionality and has the largest footprint. The key technologies of the full infrastructure deployment are the App-V Management Server, the App-V Publishing Server, the App-V client, and an

App-V Sequencer. SQL is required on the back end, and the solution ties into AD DS. Optional technologies are the reporting server and associated reporting database. This model offers streaming of applications without requiring a System Center Configuration Manager environment. An example of a typical full infrastructure model deployment is shown in Figure 4-20.

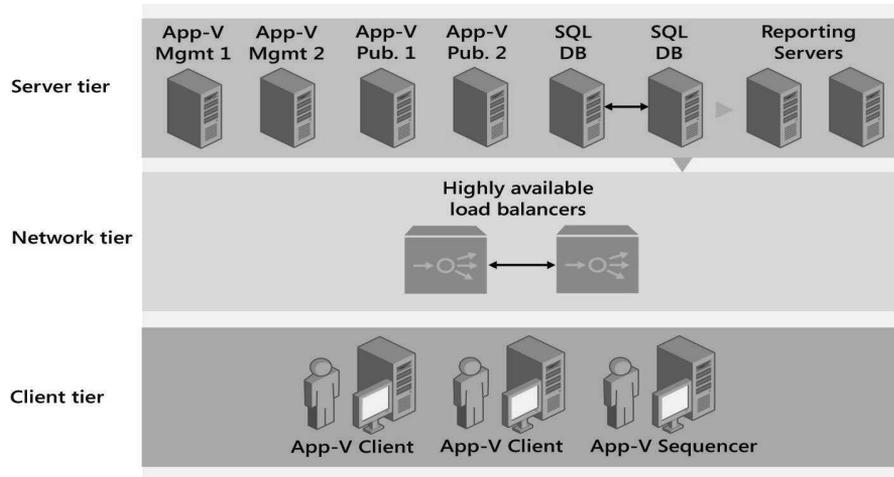


Figure 4-20 Technologies of a typical full infrastructure model

- App-V integration with Configuration Manager model** This is an implementation of App-V that, at a minimum, includes App-V clients, an App-V Sequencer, and System Center Configuration Manager 2012 or later. The App-V Management Server isn't part of this deployment model, the App-V Publishing Server isn't part of this deployment model, there isn't a desktop configuration service, and you can't track licensing or perform metering. The desktop configuration service is used by App-V clients to find out information about available virtual applications. This deployment model doesn't require SQL or AD DS, but it requires System Center Configuration Manager, which enables you to stream App-V applications from Configuration Manager distribution points and automatically deploy the App-V client to computers.
- Stand-alone deployment model** This is a small implementation of App-V. This deployment model has only App-V clients and an App-V Sequencer, which enables you to create .appv files or .msi files for delivery via a separate application delivery solution such as Group Policy or network file shares. The .appv files and the .msi files can be run only by the App-V client. This model isn't used often but can be valuable for a test environment or an environment with a lot of users who do not routinely connect to the network.

App-V packages

An App-V package comprises several files that have specific functionality. The primary package file is the .appv file that contains the captured assets and state information. Additional files provide custom integration information for publishing applications, detailed sequencing reporting, and, optionally, sequencing templates and package accelerators. You can use the following files to provide custom integration:

- The .appv file contains the captured files and state from the sequencing process in a single file. This file includes the architecture of the package file, publishing information, and registry settings in a tokenized form that can reapply to a machine and to a specific user on delivery.
- The .msi file is used in stand-alone deployments or, optionally, when deploying by using Configuration Manager or other deployment platforms.
- The _DeploymentConfig.xml file contains default publishing parameters for all applications in the package, and it can be modified to support customization.
- The _UserConfig.xml file allows customization of publishing parameters for specific domain users. You can customize items such as shortcuts and file associations with this config file.
- The Report.xml file contains diagnostic information, how sequencing is done, and what files are excluded from a package. It also contains the messages that are displayed in the Sequencer after you finish sequencing an application.
- The .cab file is an optional package accelerator file that speeds up the creation of sequenced virtual application packages.
- The .appvt file is an optional Sequencer template file that retains commonly reused Sequencer settings.

Contents of an .appv package

An .appv file is a compressed file that contains the contents of a virtual application package. It is based on the Open Packaging Conventions standard. It is used to store a combination of XML and non-XML files in a single entity. You can view .appv file contents by renaming the file with a .zip extension and exploring its contents.

The following list describes the primary .appv file contents:

- **StreamMap.xml** Contains Feature Block 0, also named the Publishing Feature Block.
- **PackageHistory.xml** Contains information about the origin of a package, for example, which user sequenced the package, on which machine, and at what time.

- **FilesystemMetadata.xml** Contains a list of the files that are part of the application and were captured during the sequencing.
- **AppxManifest.xml** Metadata for a package that contains everything that is needed to publish.
- **Registry.dat** A mountable .dat file containing the registry that was captured as part of the package.
- **Root** Contains the file system for the virtualized application that was captured during sequencing.

Planning App-V infrastructure

When introducing a new technology, such as App-V, planning and designing the infrastructure are fundamental to a successful implementation. As we'll discuss in this section, a reliable App-V environment depends heavily on the design and infrastructure. The process for implementing application virtualization is flexible and scalable, with larger deployments requiring more planning and different technologies.

Some key areas of interest when planning your App-V infrastructure include the following:

- The App-V infrastructure requirements
- The various App-V deployment models
- Sizing and performance
- High availability and disaster recovery

App-V infrastructure requirements

Before deploying App-V in your environment, you must ensure that the supporting infrastructure is in place and configured. App-V 5.0 has the following infrastructure requirements:

- **Active Directory Domain Services** AD DS is required for authentication and authorization of applications and connection groups. AD DS is needed only if you plan to deploy an App-V server, such as in a full infrastructure deployment model.
- **Installation service account** A service account in AD DS is required for the initial installation of the App-V server, presuming that your deployment is a full infrastructure model. This account needs Read permission to query AD DS and local Administrators group access on the server on which you perform the App-V installation. Following the installation of the management server, you can transition this to a security group in

AD DS, allowing you to easily add users who require administrative access to the management console.

- **Package repository** This is the location where package files will be stored for delivery to App-V clients.

The servers in an App-V environment have the following requirements::

- **Management server** Supported on Windows Server 2008 R2 with SP1 and newer. It requires the following technologies:
 - Microsoft .NET Framework 4.0 or newer
 - Windows PowerShell 3.0
 - Microsoft Visual C++ 2010 SP1 Redistributable Package (x86/x64)
 - Microsoft SQL Server 2008 Standard, Datacenter, or Developer edition (32-bit or 64-bit) or newer
- **Reporting server** Supported on Windows Server 2008 and newer. It requires the following:
 - Microsoft .NET Framework 4.0 or newer
 - Microsoft Visual C++ 2010 SP1 Redistributable Package (x86/x64)
 - Windows Web Server with the IIS role installed
 - Common HTTP Features (static content and default document)
 - Application Development (ASP.NET, .NET Extensibility, ISAPI Extensions, and ISAPI Filters)
 - Security (Windows Authentication, Request Filtering)
 - Management Tools (IIS Management Console)
 - 64-bit ASP.NET
 - Microsoft SQL Server 2008 Standard, Datacenter, or Developer edition (32-bit or 64-bit) or newer
- **Publishing server** Supported on Windows Server 2008 and newer. It requires the following:
 - Microsoft .NET Framework 4.0 or newer
 - Microsoft Visual C++ 2010 SP1 Redistributable Package (x86/x64)
 - Windows Web Server with the IIS role installed
 - Common HTTP Features (static content and default document)

- Application Development (ASP.NET, .NET Extensibility, ISAPI Extensions, and ISAPI Filters)
- Security (Windows Authentication, Request Filtering)
- Management Tools (IIS Management Console)
- 64-bit ASP.NET

Although the design of an App-V environment is very flexible, certain scenarios are not supported:

- Installation on domain controllers isn't supported for any App-V server technology.
- Installation isn't supported on Server Core installations of Windows Server.
- App-V 5.0 can't be installed on a system that has a previous version of the App-V Management Server.
- Microsoft SQL Server Express as a database engine isn't supported.

App-V deployment possibilities

Distributing virtual applications requires the App-V client software on the target computer. As you design your server infrastructure, you'll need to review the four main deployment models that we introduced earlier in this chapter. Each model has its own strengths, and the model you choose will determine which type of server infrastructure you deploy.

App-V full infrastructure model

The full infrastructure model provides all of the management server capabilities that App-V offers, including application streaming, authentication, security, licensing, and metering.

When planning for the full infrastructure model, you'll need AD DS and Microsoft SQL Server. The App-V Management Server should be on the same LAN segment hosting the database. Publishing servers are used in this model to publish content from a file server share to a distributed environment's remote locations by providing streaming capabilities close to the clients that are using the applications. This reduces latency and improves the end-user experience.

System Center 2012 Configuration Manager–integrated model

If you have an existing System Center 2012 Configuration Manager infrastructure, or you are looking to implement one, you can leverage Configuration Manager to distribute virtual applications in the same way that you distribute traditional application packages. You can add virtual applications to a Configuration Manager environment by using the same Create Application Wizard, as shown in Figure 4-21.

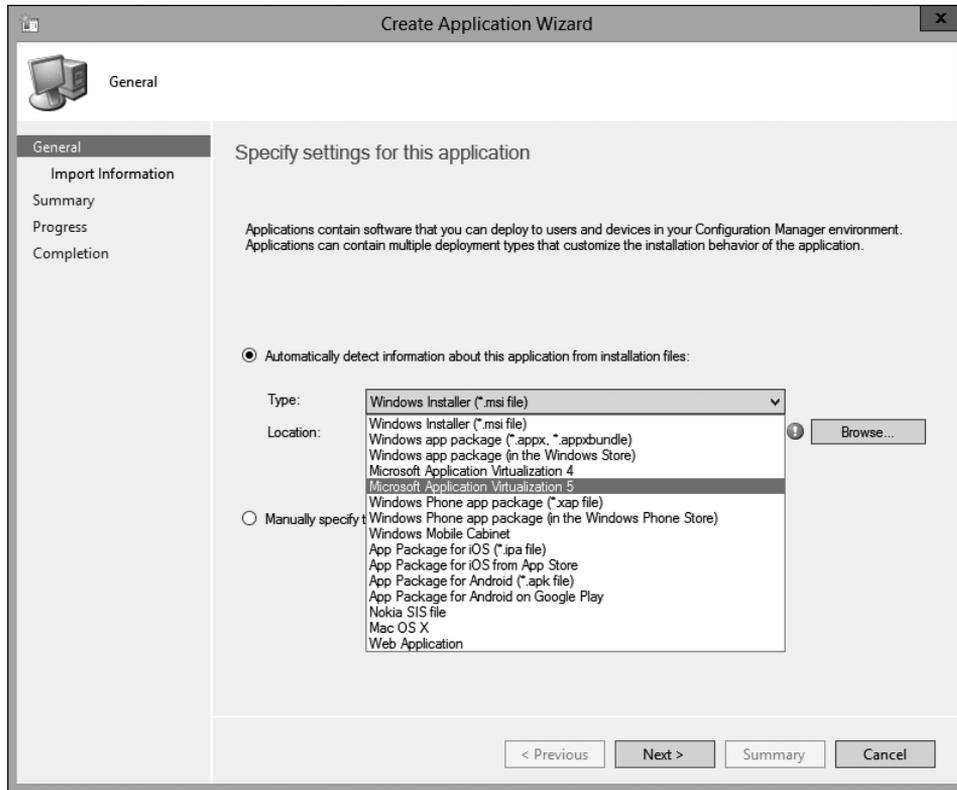


Figure 4-21 Configuration Manager Create Application Wizard

Many of the advanced capabilities that are available for managing a traditional application—such as using task sequences and building queries in collections to define which devices are targeted—also are available for a virtual application. You can target both users and computers to deliver an application in a more intelligent way, expanding on capabilities of the App-V full infrastructure model. For example, when you use a primary device as one of the possible rule requirements, you can identify which deployment type is used based on whether the user is working on his or her primary device.

The Configuration Manager–integrated model requires both the App-V client and the Configuration Manager client on each managed system. It doesn't use any server technologies of the full App-V infrastructure to deliver virtual applications; instead, it uses existing Configuration Manager distribution points to deliver the virtual application to client devices. Note that some reporting capabilities aren't available in the integrated model when compared to the full infrastructure model. For example, if you use local delivery where clients download and execute the application, you only can report if the application has been used and the

last application usage time. In the full infrastructure model with reporting, you can report the number of times an application has been used.

Application delivery to a Configuration Manager client works differently from the App-V full infrastructure scenario. In the full infrastructure model, the App-V client manages its own content, and it can refresh instantly against the publishing server. In the Configuration Manager–integrated scenario, the Configuration Manager client manages the App-V client.

Configuration Manager supports two types of delivery methods for virtual applications:

- **Streaming delivery** You can enable streaming delivery on Configuration Manager distribution points. This option streams a virtual application to a client through HTTP or HTTPS.
- **Local delivery** This delivery first uses the Configuration Manager client to download all the files needed for the application through Background Intelligent Transfer Service (BITS). After downloading the files, the package fully loads into the App-V client cache.

Electronic software distribution model

The electronic software distribution (ESD) model is ideal for environments in which you prefer to leverage an existing software distribution solution. In this case, most distribution systems can use the virtual .msi file produced by the App-V Sequencer for delivery with an .appv package.

Planning considerations for the ESD model include the following:

- **Existing software distribution system** An existing software distribution system that can recognize and distribute .msi packages to client devices.
- **App-V Sequencer** A system deployed in your environment with the App-V Sequencer installed for building and managing virtual applications.
- **Windows PowerShell** The ability to deploy a script that contains the App-V client module for Windows PowerShell cmdlets. This provides the ability to add and publish packages in ESD mode.
- **Connection groups** Designating connection groups (grouping one or more App-V packages to enable interaction with one another) requires manually creating a connection group XML file and deploying it by using a custom Windows PowerShell script.
- **Group Policy** Having Group Policy available simplifies the task of configuring the App-V client. Alternatively, a manual or scripted configuration is possible through the Windows Registry. In Figure 4-22, a GPO named App-V settings provides several App-V settings to computers.

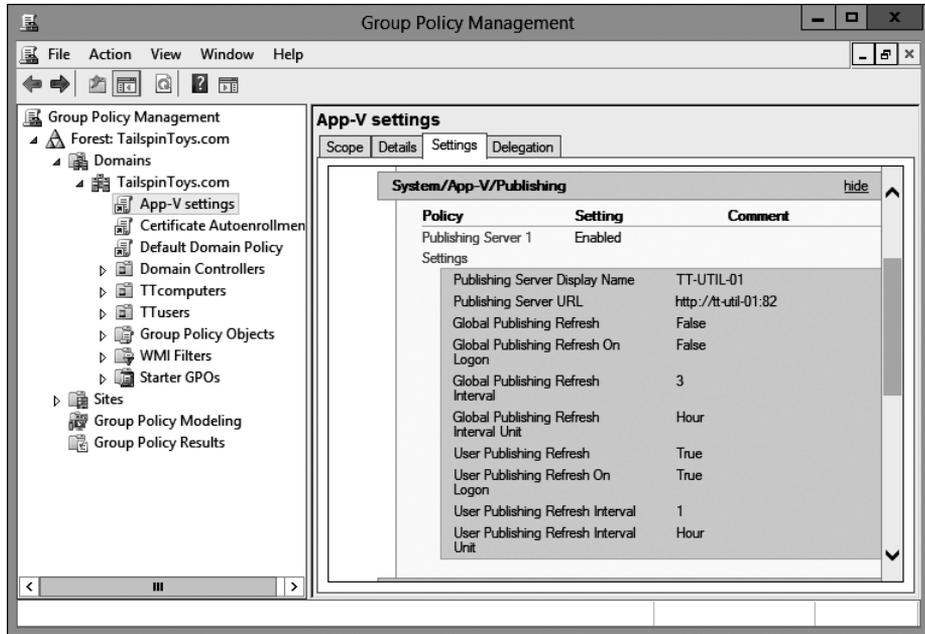


Figure 4-22 App-V GPO settings

Stand-alone deployment model

The App-V stand-alone model consists of the App-V Sequencer and an App-V client, and it requires no additional App-V infrastructure. The Sequencer has an option to create a virtual .msi file during the sequencing process. The virtual .msi file invokes Windows PowerShell commands and then publishes and loads the application to the App-V client cache.

The App-V Sequencer packages publication information, shortcuts, and the installation routines into an .appv file package, and the Sequencer generates virtual .msi files that you can execute manually. When executed, the installer adds the virtual application package to the App-V client and configures publication information to load applications from a local location rather than stream them across a WAN.

Stand-alone deployments require an App-V client on the computers, which allows a virtual .msi file to publish and load virtual applications or enables management through Windows PowerShell. You don't configure an App-V client to connect to any App-V server.

The stand-alone delivery scenario enables an organization to deploy virtual applications in situations where no servers are available to support other deployment methods for virtual applications. Use stand-alone deployments in the following scenarios:

- There are remote users who can't connect to an App-V infrastructure.
- Software management systems, such as Configuration Manager or another electronic software distribution system, already are in place.
- Network bandwidth limitations prevent electronic software distribution. In this case, you can use virtual application delivery on physical media.

Because the stand-alone model employs an .msi file, you can distribute the file if you use an existing software distribution infrastructure, such as GPOs, shared folders, optical media such as CDs and DVDs, USB flash drives, or others.

Service disruption impact

One of the common design steps in implementing App-V is to make the infrastructure highly scalable, which limits the impact of service disruption. The App-V infrastructure is highly dependent on AD DS. Therefore, it is recommended that you carefully plan your AD DS architecture to avoid unwanted service disruptions.

It's important to point out that from a client perspective, once an application is loaded on a computer, that device can run the application independently from the server. A previously published package can have different states on client computers:

- **Not Available** In this state, the package isn't registered or isn't available on the client.
- **Registered** In this state, the package is registered to the computer, but it still is not registered for the user.
- **Published** In this state, the application is registered and published on the client, and the user can start using it.
- **Partially Loaded** In this state, the application can be started because the client already has downloaded the initial feature block. Depending on which portion is missing, the rest of the files can download over the network, so the file server repository is the critical technology that provides that functionality.
- **Fully Loaded** In this state, the application downloads and extracts entirely onto a client machine, and it can be used in an offline scenario.

The following areas will be a concern if the virtual applications aren't configured to fully load on the client machine or aren't already published on the publishing server:

- **File server repository** The most critical technology that influences an application's functionality will be the file server repository. Storage availability and AD DS will need extra planning considerations in this scenario.
- **Management server** If the management server or the management database is down, adding new packages, updating existing packages, or managing connection groups won't be possible.
- **Publishing server** Publishing server failure affects the ability to make changes to the publishing list that clients previously received, which is for non-persistent Virtual Desktop Infrastructure (VDI) and RDS scenarios.
- **Reporting server** Reporting server or reporting database failure isn't critical to running App-V applications; the only functionality that might not work is reports that a client sends about usage statistics for virtual applications, which are stored in the reporting database.

Functional and physical placement

Organizations might plan different App-V infrastructure deployments based on their needs. When you start to plan for your App-V environment, you should try to answer the following questions to help with your design and implementation:

- Are there requirements that all roles must live on a single server? Decide whether you want to combine or cohost functionality.
- Do you need centralized or decentralized roles in a distributed environment?
- What are the requirements for high availability?
- Are the virtualized application users located in all of your office locations?
- Is your virtualization environment able to virtualize your entire App-V deployment?

Based on the answers to these questions, there are several design scenarios:

- **Small and midsize deployments** For small and midsize scenarios, which commonly address an environment with a small number of users and few packages in a single geographical site, you might cohost all of the roles on a single server.
- **Midsize and large deployments** For midsize and large deployments, which commonly require a flexible and scalable environment, you might consider a more complex design in which all services implement individually. In this scenario, every connection addresses a virtual IP address and machine name, and no services cohost on any given computer.

- **Distributed deployments** For distributed deployments, which commonly need to support a large number of users in different locations with many different requirements, you should implement a scenario that can address locations with no datacenter and weak Internet connectivity. For this type of design, all configuration data that is stored in a management database should be located in a major datacenter. Because the management server communicates with the management database, it should be located close to the SQL server because SQL communication is time-sensitive and network-sensitive. The file server repository share that holds the application should be located close to clients, and interval refreshes can be adjusted according to the actual network capabilities.
- **High availability deployments** In this scenario, you must have two identical machines (physical or virtual) that are configured in NLB mode (or behind a third-party hardware load balancer), where the following services are installed:
 - App-V management
 - App-V publishing
 - App-V reporting

Even if you don't start with a highly available environment, you should consider using load balancing. It can simplify scaling out later and provide some additional capabilities such as drainstopping a server.

It is recommended that you host the SQL Server database separately from the App-V services. This consideration is made for performance, security, and scalability. For highly available designs, you should consider implementing a SQL Server cluster.

Sizing and performance

Actual sizing and performance planning depends on multiple factors, such as scaling an App-V infrastructure properly to lower the round-trip response time and providing proper package optimization for streaming across slow networks.

Round-trip response time on a publishing server is the time that is needed for the publishing server to receive a successful package metadata update from the management server. Round-trip response time on a client is the time the App-V client computer takes to receive a successful notification from the publishing server.

If you have increased internal demand, you can implement an additional management server behind your load balancers.

Often, users might demand external scalability based on the location you must support. A design should include a content repository in each location to provide conveniently located packages to clients. Additionally, you might consider implementing a publishing server and a management server to lower the round-trip response time on clients. Capacity planning

should be included to evaluate future demands in planned growth to meet expected performance levels.

A few factors influence round-trip response time on a publishing server. Some of these include the number of:

- Publishing servers that make simultaneous requests.
- Connection groups that are configured on a management server.
- Access groups that are configured on a management server.
- A single management server can simultaneously respond to up to 320 publishing servers with a round-trip response time of approximately 40 seconds; a single management server with fewer than 50 publishing servers results in a round-trip response time of less than 5 seconds.
- The number of connection groups starts to influence round-trip response time after more than 400 are created.
- The number of access groups increases the round-trip response time as it grows.

The number of publishing servers that simultaneously connect to a management server does not influence central processing unit (CPU) utilization and SQL database transactions per second; batches per second are identical, regardless of the number of publishing servers.

For App-V, reporting server capacity planning should focus on the number of clients that simultaneously send reporting information to a reporting server. Round-trip response time increases linearly with an increased number of clients. For example, round-trip response time is 2.6 seconds with 500 clients and 5.2 seconds with 1,000 clients.

Capacity planning for the publishing server influences the round-trip response time on an App-V 5.0 client computer to send a publishing refresh request and to receive a response.

The following are the main factors that influence capacity planning of an App-V publishing server:

- The number of clients that simultaneously connect to a single publishing server.
- The number of packages in each refreshThe network bandwidth between clients and the publishing server.
- A publishing server with a dual-core processor can respond to up to 5,000 clients that simultaneously request refreshes. From 5,000 through 10,000 clients, a publishing server should have a quad-core processor at minimum. Increasing the number of packages increases response time by 40 percent, and network bandwidth has a major influence on

response time. For example, clients that run on slow networks—less than 1.5 megabits per second—will have a significantly slower response time than the same number of clients that run on LAN networks.

High availability for App-V

You should plan for a highly available App-V infrastructure in organizations where App-V is important. The high availability strategy for the App-V infrastructure depends on the App-V deployment model, because different procedures and settings for high availability are needed for different App-V deployment models.

Stand-alone deployment model

The stand-alone deployment model only requires an App-V Sequencer and client computers that have an installed App-V client. In the stand-alone deployment model, the App-V Sequencer is used only when a new application needs to be sequenced. Because the App-V Sequencer installs only when a new application needs to be sequenced, it isn't necessary to make the Sequencer highly available. If you stream from a central share, this share can deploy on a clustered file system or on an NLB web farm. In cases where you require access to sequencing, even in a disaster recovery (DR) scenario, you can deploy multiple sequencing computers.

App-V full infrastructure model

From a planning perspective, the App-V full infrastructure model requires the most attention. Because there are a multitude of technologies with differing high availability models, you should spend time looking at the options available and decide which one makes the most sense for your environment. The following are some questions that you should answer to help you plan for high availability:

- Does the reporting service need to be highly available?
- Does the sequencing computer need to be highly available?
- Are there infrastructure technologies outside App-V that may impact the high availability of App-V, such as load balancers, switches, virtualization servers, or storage?
- Which secondary site should each office use in the case of a publishing server failure at the office?

The answers to the above questions will help you plan the services, the number of servers required, and the high-level design of your highly available environment.

The App-V full infrastructure model stores all configuration and application information in the management server database and stores all utilization data in the reporting server database,

so each of these databases is a single point of failure in this model. Therefore, when you are configuring the App-V full infrastructure model to be highly available, you need to ensure that the management server database and the reporting server database remain accessible. You can do this by deploying these two databases on an instance that is installed on a highly available VM or on a clustered SQL Server instance.

App-V 5.0 supports multiple management, reporting, and publishing servers. You can configure the App-V 5.0 management and reporting server databases to work with multiple management servers and reporting servers by using a security group when specifying the computer account location during setup. At any time, you can add publishing servers to an existing App-V full infrastructure model deployment.

Consider Figure 4-23, which shows what can happen when the application team provides high availability for its services but another team isn't engaged in the project and is unaware of potential impacts.

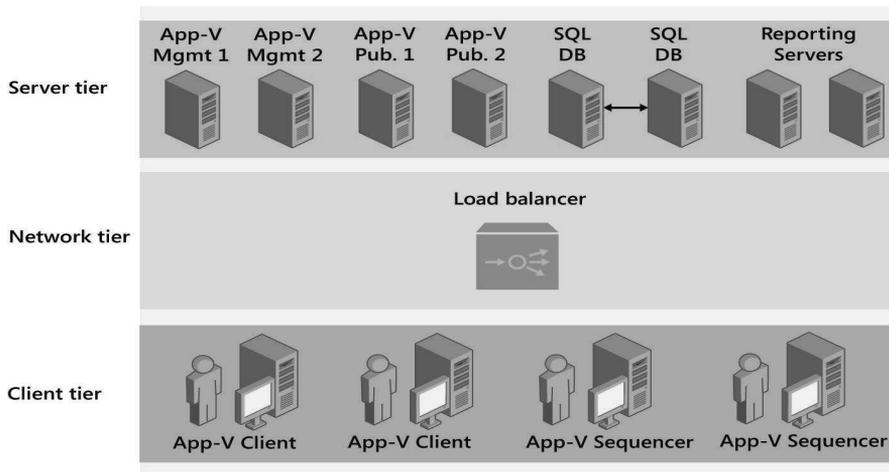


Figure 4-23 A diagram of an App-V full infrastructure deployment model

In the diagram above, although all of the App-V technologies are highly available, there is only one load balancer. It represents a single point of failure. Instead of this scenario, the availability of all services on which App-V is dependent should match the App-V availability.

Integrated Configuration Manager model

For the integrated Configuration Manager model to be made highly available, you should look at all of the options to meet your high availability requirements and figure out if any infrastructure changes are required:

- Are highly available VMs available in your environment? If so, you can use one or more for your Configuration Manager servers.
- Is your existing SQL environment highly available?
- Is the storage that the virtual environment uses highly available?
- Do you have distribution points in all of the locations to which you plan to deliver virtual applications? If so, you need to plan for scenarios in which a distribution point becomes unavailable. In environments with an existing Configuration Manager deployment, it isn't unusual to have sites with a single distribution point. You should consider multiple distribution points if your requirements include immediate access to virtual applications, especially for new App-V clients.
- You also may want to look at the overall high availability of the Configuration Manager environment. Is there an existing hierarchy with a central administration site?

Similar to other App-V deployment models, you should look at all aspects of your environment and ensure that all of the services involved are configured for high availability.

- For more information about high availability in Configuration Manager, see the “Planning for High Availability with Configuration Manager” website at <http://technet.microsoft.com/en-us/library/hh846246.aspx>.

Disaster recovery

A disaster recovery plan should include a proper backup of critical technologies to respond to a service outage. In addition, regular testing of restore operations will help ensure that the backups are functional and that the operational procedures are adequate to recover from a disaster. At a minimum, an App-V infrastructure backup should protect the management database and the package repository that contains the App-V packages. Outside App-V, you should ensure that the services on which App-V depends also are backed up and restored first after a disaster. For example, you need to ensure that AD DS, SQL, Server virtualization, and the core networking services are up and running before you can successfully recover App-V.

In an App-V backup and recovery scenario, each role has different requirements:

- A management server does not hold any unique data other than the registry configuration of the database source, so you can easily re-create this role in case of a disaster.
- A publishing server contains a registry value that indicates the host name of the management server to contact and a cached copy of the latest publishing data, which you need to back up.
- The reporting server has a registry value that indicates the name of the reporting database.

There are two different scenarios for recovery procedures:

- If a server that contains all of the roles fails, administrators should perform standard image recovery procedures that are defined by organizational policy. This often means restoring the VM or physical server to the most recent backup.
- If an App-V service fails, administrators should perform a recovery by installing the App-V technologies and prerequisites, such as installing and configuring Internet Information Services (IIS) and installing SQL Server on the database server.

When you restore the management server, it enables this service to become operational as soon as the service can contact the App-V database. When you complete the restoration of publishing servers, client requests will be serviced as soon as the service contacts the publishing servers (from 1 to 10 minutes). After the App-V services are restored, the reporting server starts accepting client connections.

Deploying App-V infrastructure

After you determine which application virtualization model to use, you need to deploy the appropriate roles to support that model. In most production environments, you should host the management server database and the reporting server database separately from the management server and the reporting server. It isn't uncommon to see the management database and reporting database also separated, depending on the database infrastructure you have. Before a deployment, you need to understand the order in which to deploy App V technologies. In this section, we will explain the following:

- App-V infrastructure requirements
- Installing management databases
- App-V Management Server configuration
- App-V Publishing Server deployment and configuration
- App-V for Remote Desktop Services clients
- Integrating App-V with System Center Configuration Manager

App-V infrastructure requirements

The technologies of App-V should be deployed in a specific order. If you deploy all roles on the same server when you deploy the App-V full infrastructure model, the installation wizard automatically deploys them in the correct order. If you deploy the roles on more than one server, you should deploy them in the following order:

1. Management server database
2. Management server
3. Publishing server
4. Reporting server database (optional)
5. Reporting server (optional)

When you configure management server database settings, you need to specify the security account of the computer that will access the database. This can be a security group account or the computer account of the server that will function as the management server. When you install a management server separately, you need to specify the instance location and credentials that will be used to access the management server database.

When you deploy a publishing server, you need to specify the network address of the management server. You can't deploy a publishing server without having a management server already deployed, unless you deploy all of the roles at the same time.

You must deploy a reporting server database before deploying a reporting server. The reporting server doesn't have dependencies on any other services in the App-V full infrastructure model.

You don't need to install the App-V Sequencer when deploying the App-V full infrastructure model. However, it is a good idea to deploy a sequencer as soon as possible to begin testing the deployment. The App-V client usually is the last application to be deployed.

Installing management databases

The management server stores all of the configuration data in an App-V management database, which includes all application metadata, the deployment configuration, the relationships, and the security assignments. The management server only communicates with the management database, and it is the first technology that should be installed in the App-V full infrastructure model. When you add management servers for a scalable deployment, you only need to allow Read and Write permissions to the database. You don't have to provide additional configurations.

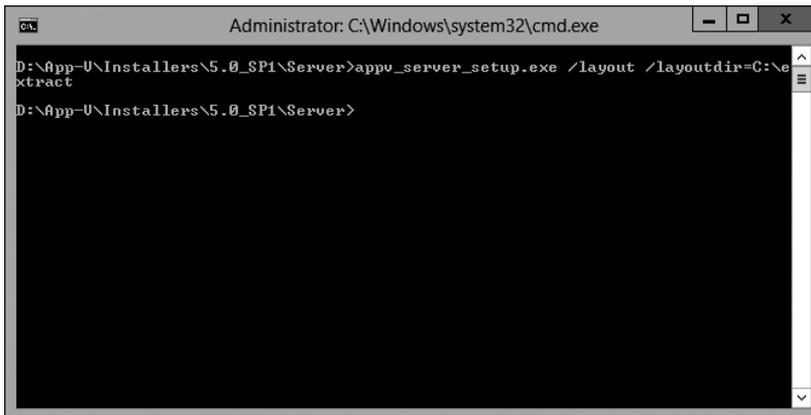
The minimum supported database platform is Microsoft SQL Server 2008 R2 Standard, Enterprise, Datacenter, or Developer edition (32-bit or 64-bit). The Developer edition should not be used in a production environment. Additional prerequisites include the installation of the following:

- .NET Framework 4 (Full Package)
- Microsoft Visual C++ 2010 SP1 redistributable package (x86)
- Windows Server 2008 or newer

In the deployment scenario in which you install all of the technologies on the same computer, the App-V server setup GUI-based installation first installs the App-V management database and then installs the management server and the publishing server. Finally, if selected, the reporting database and reporting server are installed.

If you are implementing scalable deployment, you should run the GUI installation on a server that hosts the management database because remote SQL database creation isn't supported in the installer.

As an alternative, you can install the SQL database when you execute SQL Server scripts that are extracted from the server setup, as shown in Figure 4-24. SQL Server scripts extract from the setup with the following command: `appv_server_setup.exe /layout /layoutdir=c:\extract`.



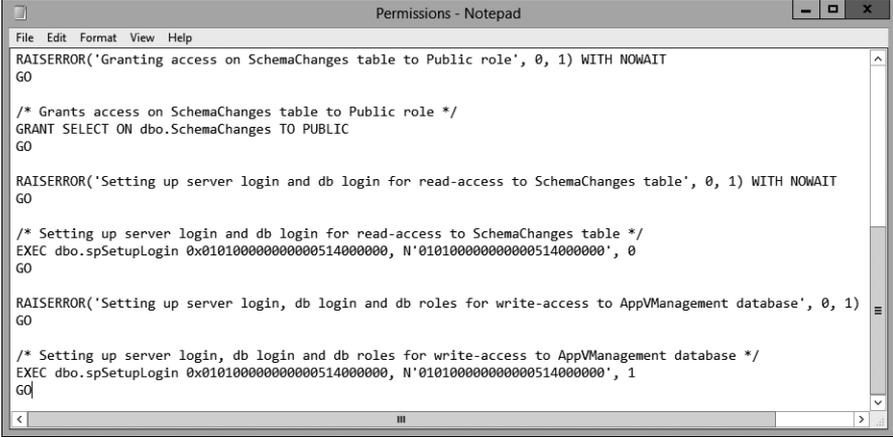
The screenshot shows a command prompt window titled "Administrator: C:\Windows\system32\cmd.exe". The prompt is at the directory `D:\App-V\Installers\5.0_SP1\Server`. The command `appv_server_setup.exe /layout /layoutdir=c:\extract` has been entered and executed. The output shows the command being processed and the directory path being used.

```
Administrator: C:\Windows\system32\cmd.exe
D:\App-V\Installers\5.0_SP1\Server>appv_server_setup.exe /layout /layoutdir=c:\e
xtract
D:\App-V\Installers\5.0_SP1\Server>
```

Figure 4-24 Example `appv_server_setup.exe` command

In the destination folder, you must modify two of the six scripts to provide the appropriate Read and Write permissions for the domain accounts or domain groups that you need to manage an App-V infrastructure. The first modification must be done to the `Permissions.sql`

script to replace the entry for [ManagementDBWriteAccessAccountSid] and for [ManagementDBWriteAccessAccountName] with the security identifier (SID) and the name for the domain group that requires Write permissions to the database. This group should include the App-V administrator account and all management servers in the environment. If you use the same account for installation and App-V administration, then you should use the same entries for [ManagementDBPublishAccessAccountSid] and [ManagementDBPublishAccessAccountName]. Otherwise, you should enter the correct SID and name for the installation account. In Figure 4-25, the Permissions.sql file has been modified and is ready for use.



```

File Edit Format View Help
RAISERROR('Granting access on SchemaChanges table to Public role', 0, 1) WITH NOWAIT
GO

/* Grants access on SchemaChanges table to Public role */
GRANT SELECT ON dbo.SchemaChanges TO PUBLIC
GO

RAISERROR('Setting up server login and db login for read-access to SchemaChanges table', 0, 1) WITH NOWAIT
GO

/* Setting up server login and db login for read-access to SchemaChanges table */
EXEC dbo.spSetupLogin 0x010100000000000514000000, N'010100000000000514000000', 0
GO

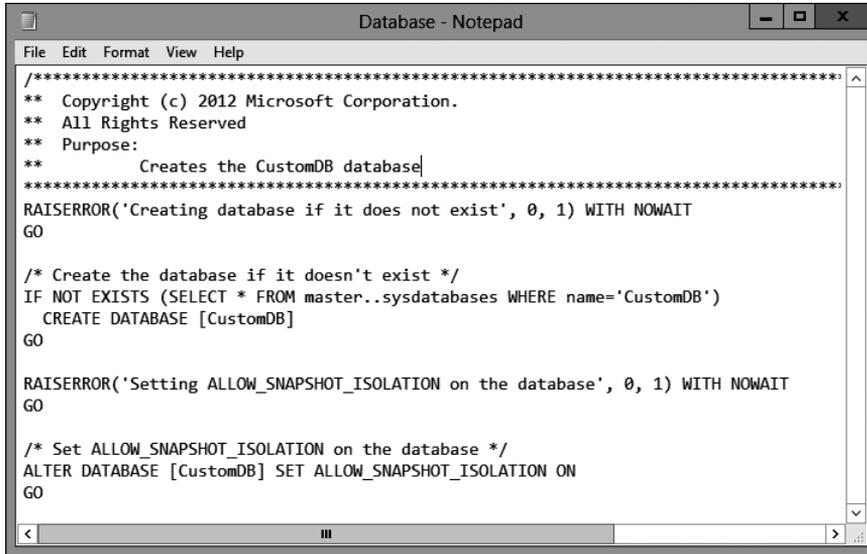
RAISERROR('Setting up server login, db login and db roles for write-access to AppVManagement database', 0, 1)
GO

/* Setting up server login, db login and db roles for write-access to AppVManagement database */
EXEC dbo.spSetupLogin 0x010100000000000514000000, N'010100000000000514000000', 1
GO

```

Figure 4-25 An example Permissions.sql file

Modification of the second script, Database.sql, is optional and has to be done only if you plan to replace the default database name AppVManagement with a unique name. In Figure 4-26, the Database.sql file has been updated to create a database named CustomDB.



```

Database - Notepad
File Edit Format View Help
/******
** Copyright (c) 2012 Microsoft Corporation.
** All Rights Reserved
** Purpose:
**     Creates the CustomDB database
*****
RAISERROR('Creating database if it does not exist', 0, 1) WITH NOWAIT
GO

/* Create the database if it doesn't exist */
IF NOT EXISTS (SELECT * FROM master..sysdatabases WHERE name='CustomDB')
CREATE DATABASE [CustomDB]
GO

RAISERROR('Setting ALLOW_SNAPSHOT_ISOLATION on the database', 0, 1) WITH NOWAIT
GO

/* Set ALLOW_SNAPSHOT_ISOLATION on the database */
ALTER DATABASE [CustomDB] SET ALLOW_SNAPSHOT_ISOLATION ON
GO

```

Figure 4-26 An example Database.sql file

A SQL Server administrator must run prepared SQL scripts against a computer that is running SQL Server that will host the database. SQL Sysadmin permissions are required. You can run the script if you first open the SQL Server Management Studio console and run it as a query, but you need select the proper database. The second method that you can use is the **OSQL** command-line application. “The switches /E, /i, and /d are case-sensitive. Reporting database setup is identical to a management database and can be done with the App-V setup installer, or it can be pre-created with SQL scripts. The following commands can be used:

```
OSQL -E -i database.SQL
```

```
OSQL -E -d MS_Appv5_Management -i CreateTablees.sql
```

```
OSQL -E -d MS_Appv5_Management -i CreateStoredProcs.sql
```

```
OSQL -E -d MS_Appv5_Management -i UpdateTables.sql
```

```
OSQL -E -d MS_Appv5_Management -i insertversionInfo.sql
```

```
OSQL -E -d MS_Appv5_Management -i Permissions.sql
```

App-V Management Server configuration

An App-V Management Server provides a centralized location to manage an App-V 5.0 infrastructure for delivering virtual applications to both an App-V client and an RDS (formerly Terminal Services) client. Unlike previous versions of App-V, a web application that runs on Silverlight manages the App-V 5.0 infrastructure. You configure this web application’s address

when you install a management server. The installation of an App-V Management Server creates a dedicated IIS website, for which you can specify the name during the installation setup. By default, it is called the Microsoft App-V Management Service. The App-V Management Service will be configured to listen on a dedicated port number, which can be provided during setup.

App-V server features can install on multiple servers to provide scalability and high availability; however, all App-V server features would need a common way to be accessed, such as by using a load balancer. Each management server node needs connectivity to the database on the computer that is running SQL Server. If a single server hosts multiple technologies, they can use different ports, or you can configure them to share a single port.

Preinstallation tasks include configuring appropriate user and administrative groups that can install and administer a management server. A management server requires that an IIS server is installed and configured to be trusted for delegation. If you plan to support Secure Sockets Layer (SSL) for connectivity to a management server, you also need a server certificate that is issued from either an internal or a public certification authority.

Installing an App-V Management Server at the command line requires elevated privileges. You can display the installation parameters, shown in Figure 4-27, by running the following command:

```
appv_server_setup.exe /?
```

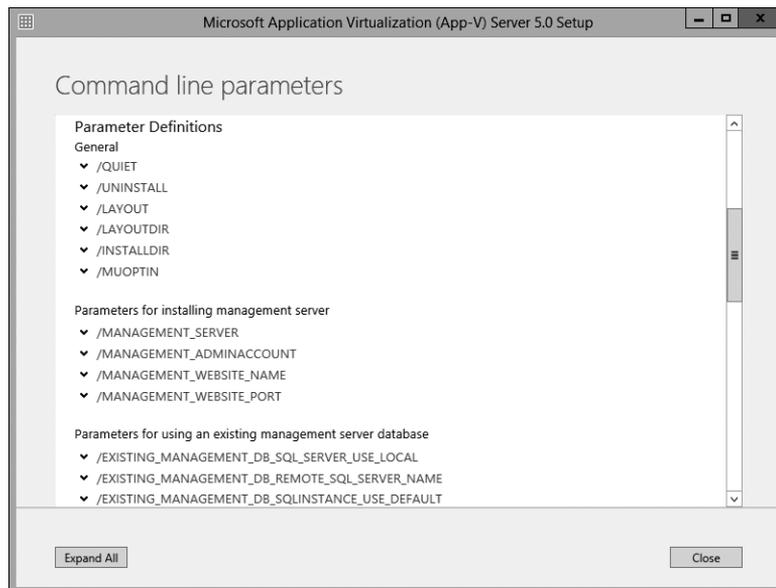


Figure 4-27 App-V setup command-line parameters

The following commands provide an example of an App-V Management Server installation, as shown in Figure 4-28. You can verify the output of the installation in the log file `appv_server_datatime.log` in the `%temp%` directory.

```
appv_server_setup.exe /quiet /management_server /MANAGEMENT_ADMINACCOUNT="Adatum
\AppVAdmin" /MANAGEMENT_WEBSITE_NAME="Microsoft App-V Management
service" /MANAGEMENT_WEBSITE_PORT="80" /EXISTING_MANAGEMENT_DB_REMOTE_SQL_SERVER
_NAME="SQLSRV.adatum.com " /EXISTING_MANAGEMENT_DB_SQLIN STANCE_USE_DEFAULT
/EXISTING_MANAGEMENT_DB_NAME="AppVManagement"
```

Figure 4-28 Command-line setup of an App-V server

Common postinstallation tasks include sharing the content folder that is used to store the App-V package. It also is common to enable firewall rule exceptions.

You can install a management server by using a very basic GUI, but for enterprise deployment, we recommend script-based installation. Management servers and publishing servers have a dependency on IIS with the following features:

- Common HTTP features: static content and default document
- Application development features: Microsoft ASP.NET, Microsoft .NET Extensibility, and Internet Server API (ISAPI) extensions and filters
- Security features: Windows authentication and request filtering
- Management tools features: IIS Manager

One common method for proper installation of IIS and all required services is to use the Deployment Image Servicing and Management (DISM) tool, which you can use to create a script by saving the following commands in a text editor with the `.cmd` extension: `dism /Online /Enable-Feature /FeatureName:IIS-ApplicationDevelopment ^`

```

/FeatureName:IIS-ASPNET /FeatureName:HS-commonHttpFeatures ^
/FeatureName:Iis-DefaultDocument /FeatureName:Iis-DirectoryBrowsing ^
/FeatureName:Iis-HealthAndDiagnostics ^
/FeatureName:HS-Httpcompressionstatic ^
/FeatureName:HS-HttpErrors /FeatureName:HS-HttpLogging ^
/FeatureName:HS-HttpTracing /FeatureName:HS-ISAPIExtensions ^
/FeatureName:HS-ISAPIFilter ^
/FeatureName:Iis-LoggingLibraries /FeatureName:HS-ManagementConsole ^
/FeatureName:Iis-ManagementService /FeatureName:HS-NetFxExtensibility ^
/FeatureName:IIS-Performance /FeatureName:ns-RequestFiltering ^
/FeatureName:HS-RequestMonitor /FeatureName:HS-Security ^
/FeatureName:Iis-staticcontent /FeatureName:Iis-webserver ^
/FeatureName:HS-webserverManagementTools ^
/FeatureName:HS-webserverRole /FeatureName:Iis-windowsAuthentication ^
/FeatureName:WAS-ConfigurationAPI /FeatureName:WAS-NetFxEnvironment ^
/FeatureName:WAS-ProcessModel /FeatureName:WAS-WindowsActivationService ^

```

A management server has the following requirements:

- 1-gigahertz (GHz) or faster x64 processor; two cores Intel Xeon 2.0 GHz or faster recommended
- 2 gigabytes (GB) or more of RAM; 4 GB of RAM recommended
- 200 megabytes (MB) of free disk space (does not include content); 40 GB recommended
- Windows Server 2008 R2 SP1 or newer.NET Framework 4 Extended
- .NET Framework 3.5.1 Features (or 4.5)
- Visual C++ 2010 SP1 Redistributable Package (64-bit)
- Visual C++ 2010 SP1 Redistributable Package (32-bit)
- Silverlight
- Windows PowerShell 3.0

App-V publishing server deployment and configuration

When you deploy an App-V publishing server, you must specify the location of an existing App-V Management Server. This is different from previous versions of App-V, in which it was possible to deploy a stand-alone streaming server without having to configure a management server.

Publishing servers function as distribution points for virtualized applications when you use the App-V full infrastructure model. Applications stream from these servers to clients. The entire application doesn't need to stream before a user can start interacting with it; therefore, you won't need as much bandwidth as you would with other deployment methods. Nonetheless, you still need to provision adequate bandwidth for the connection between a publishing server and the client.

To install a publishing server by using a GUI installer, you must follow the same steps as installing the management server. You have to point to an existing management server, and if these two roles coexist on the same computer, you must choose a different port for the website.

When you perform a command-line installation, you can use the help that the installer provides, which presents examples and definitions that construct the following command:

```
appv_server_setup.exe /?
```

You can use the following commands to perform publishing server installation at the command line:

```
appv_server_setup.exe /quiet /publishing_server
/PUBLISHING_MGT_SERVER=http://lon-svr1.adatum.com
/PUBLISHING_WEBSITE_NAME="Microsoft Appv Publishing service"/PUBLISHING_WEBSITE_PORT="80"
/EXISTING_MANAGEMENT_DB_REMOTE_SQL_SERVER_NAME="SQLSRV.adatum.com" /EXISTING_MANAGEMENT_DB_SQLIN_STANCE_USE_DEFAULT /EXISTING_MANAGEMENT_DB_NAME="AppVManagement"
```

As a best practice, when you install on the same server as the management server, use port 80 for the publishing server and an alternate port for the management server.

After you deploy both the management server and the publishing server, you need to configure them with appropriate firewall rules to provide management and client connectivity.

A publishing server is a web application that is hosted on IIS, and any configuration changes can be done through IIS Manager or by using the HKLM\Software\Microsoft\APPV\Server\PublishingService\PUBLISHING_MGT_SERVER registry settings to point to the protocol and port that establish connectivity with the App-V Management Server. The PUBLISHING_MGT_SERVER_REFRESH_INTERVAL registry setting specifies how often a publishing server queries a management server for packages. The default value is 600 seconds (10 minutes), and for testing purposes, you can shorten the interval to propagate changes to

clients more quickly. Any registry changes will become effective when the application pool restarts or IIS restarts.

You also can configure management server settings in IIS and the registry, as shown in Figure 4-29.

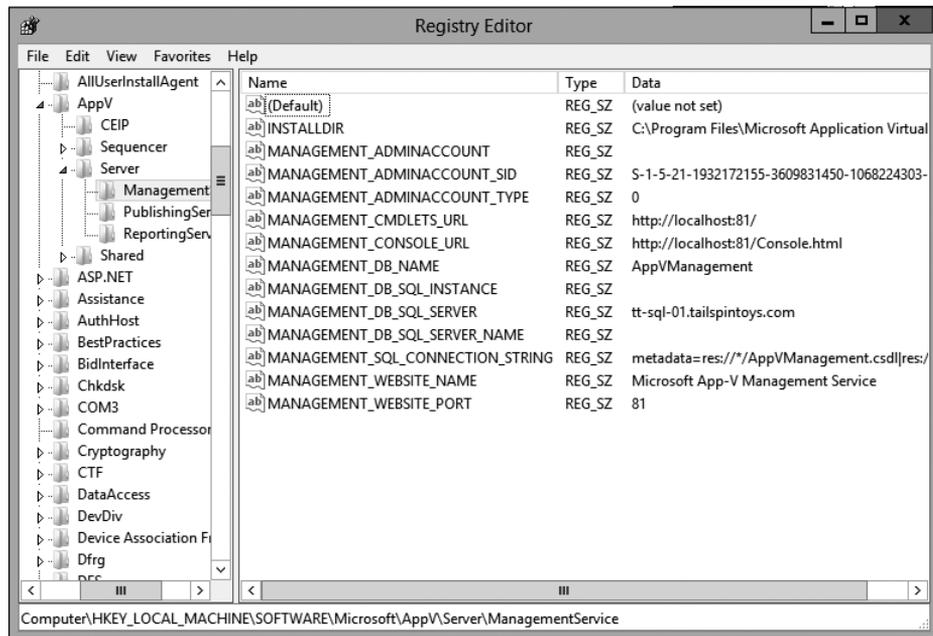


Figure 4-29 App-V registry settings for the management server

HKLM\Software\Microsoft\APPV\Server\ManagementService contains configuration data for a management server. From this registry location, you can identify or change a connection string to the management database (*MANAGEMENT_SQL_CONNECTION_STRING*) or identify the port and name for the management website.

In addition to the registry, some configuration settings are stored in the files in the *INSTALLDIR*. The *AdminGroup.xml* file contains information to recover access to an App-V console when you remove the last administrator from the console.

A publishing server has the following requirements:

- 1 GHz or faster x64 processor; two cores Intel Xeon 2.0 GHz or faster recommended
- 2 GB or more of RAM; 4 GB of RAM recommended
- 200 MB of free disk space (does not include content); 40 GB recommended

- Windows Server 2008 R2 SP1 or newer.NET Framework 4 Extended
- Visual C++ 2010 SP1 Redistributable Package (32-bit)
- Windows PowerShell 3.0
- The Web Server role with the following features:
 - Common HTTP features: static content and default document
 - Application development features: ASP.NET, .NET Extensibility, ISAPI extensions and filters
 - Security features: Windows authentication and request filtering
 - Management tools features: IIS Manager

App-V for Remote Desktop Services client

App-V 5.0 has a separate, special client that makes it possible to run virtualized applications on RD Session Host servers. With this client, you can run applications on RD Session Host servers that might not otherwise run on an RD Session Host server.

The App-V for RDS client has the following system requirements:

- 1.4 GHz or faster x86 or x64 processor
- Windows Server 2008 R2 SP1 or Windows Server 2012
- .NET Framework 3.51 and 4 (Full)
- Windows PowerShell 3.0
- Microsoft KB2533623 (Windows Server 2008 R2)
- Visual C++ 2008 redistributable (if installing by using an executable file)

You must configure Windows Server 2008 R2 or Windows Server 2012 as an RD Session Host server before you install the App-V for RDS client.

You can use the App-V for RDS client with the App-V full infrastructure, stand-alone, and Configuration Manager-integrated models. The App-V for RDS client uses the same Group Policy settings as the normal App-V client.

Integrating App-V with System Center Configuration Manager

The Configuration Manager–integrated model requires that you have an existing Configuration Manager or newer deployment. This model allows you to deploy sequenced App-V applications as one of many different application deployment types.

Before deploying sequenced App-V applications, you should configure App-V client software as an application that you can deploy. You then can specify the App-V client as a requirement when deploying any sequenced App-V application.

You can create the App-V client as an application by performing the following procedure:

1. Copy the App-V client installation file, `corecli_amd64.msi` or `corecli_i386.msi`, to a shared folder. In the Configuration Manager console, in the Software Library workspace, under the Application Management node, click Applications.
2. On the ribbon, click Create Application.
3. On the General page of the Create Application Wizard, set the type to Windows Installer (*.msi file) and then click Browse.
4. Browse to the shared folder where you copied the App-V client installation file. Finish the wizard and then click Close.

To create an App-V application in Configuration Manager, perform the following procedure:

1. In the Configuration Manager console, in the Software Library workspace, under the Application Management node, click Applications.
2. On the ribbon, click Create Application.
3. On the General page of the Create Application Wizard, set the type to Microsoft Application Virtualization 5 and then click Browse to go to the network location that hosts the file in .appv file format.
4. Finish the wizard and then click Close.

This page intentionally left blank

This page intentionally left blank



Planning and deploying session-based virtual desktops

Understanding RDS	387	Deploying session-based virtual desktops	415
Planning infrastructure for session-based desktops	403	Understanding high availability for RDS	438

Session-based virtual desktops are widely used by organizations to provide remote access to data and applications in a centralized and controlled environment. In Windows Server 2012 R2, Remote Desktop Services (RDS) provides the infrastructure to implement session-based virtual desktops and virtual machine (VM)-based virtual desktops.

In older versions of Windows Server, session-based desktops were provided by a feature named Terminal Services. Terminal Services had the same basic functionality for session-based desktops as RDS, but RDS has been extended with additional functionality to improve the user experience and manageability.

Understanding RDS

RDS is a Windows Server role that provides much more than just remote desktops. RDS includes six role services that enable you to create a scalable and fault-tolerant RDS deployment. You can manage an RDS deployment centrally and in the same way, regardless of the number of servers in an RDS deployment. This makes RDS very scalable.

One of the most common uses for RDS is the deployment of session-based virtual desktops. In a session-based virtual desktop, all processing is performed on a Remote Desktop Session Host (RD Session Host) server, and the results are displayed on a Remote Desktop client. The communication between the client and the RD Session Host server uses Remote Desktop Protocol (RDP).

RDP is a very efficient protocol and sends a limited amount of data over the network. This makes it possible to use RDS to provide desktops and applications for users over a LAN, from branch locations over a WAN, or over the Internet.

RDS includes the following functionality:

- **Provides users with a full desktop** Whether you use session-based virtual desktops or VM-based virtual desktops, you can provide users with access to a full remote

desktop that you can access from almost anywhere if you configure the necessary infrastructure.

- **Provides users with access to applications** You can use RemoteApp to provide users with access to applications running on an RD Session Host server. These applications run in a window just as regular applications do on users' desktops. From the user's perspective, applications delivered by RemoteApp function as if they are installed locally.
- **Allows secure remote access without using a virtual private network** The Remote Desktop Gateway (RD Gateway) role service is used as a proxy for accessing session-based virtual desktops or VM-based virtual desktops. This is suitable for securing access from the Internet.

The Terminal Services functionality found in older versions of Windows Server had only session-based virtual desktops and applications. In Windows Server 2012 R2, you also can use RDS to deploy VM-based virtual desktops. Connectivity to the VMs is done by using RDP, just as in a session-based deployment.

Some benefits of using RDS for virtual desktops and applications include the following:

- **Easier application deployment and updates** A typical application deployment requires you to install and update the application on each client computer. In all but the smallest environments, this requires you to implement some type of automated deployment tools for applications. With RDS, you only need to install and update applications on the central servers. This is significantly less work than installing and updating applications on individual client computers.
- **Simplified access to data and applications** When you implement RDS, applications and their data can be accessed from anywhere. You can allow users to use applications from a computer in the office, a home computer, and mobile devices.
- **Faster access to remote data** Access to data over a virtual private network (VPN) or WAN links often results in poor application performance. For example, an application that requires access to a SQL server may be very slow if the connectivity to the SQL server has high latency. When you use RDS, you place the central servers with the application installed close to the application data, and network latency is removed as a performance problem.
- **Higher data security for mobile users** Without RDS, mobile users copy data onto a mobile computer and take it with them. Or, in some cases, they use a VPN to access data remotely while offsite. In both cases, there is a risk that the mobile computer could be lost or stolen and the data accessed by unauthorized users. When you use RDS for remote access to data, there is no need to copy data to the remote device. This mitigates the risk that your organization will lose control of the data.

- Simplified client hardware management** Using RDS to provide virtual desktops reduces the effort to manage client device computers because the devices are performing much less work. Computers used to access virtual desktops become essentially disposable because the only configuration information they contain is the connection information to the remote desktop. In some cases, you may be able to extend hardware life because the client device is performing very little work.

Comparing RDS and the Remote Desktop feature

Remote Desktop is a feature in Windows 8.1 and Windows Server 2012 R2 that enables you to connect to a computer remotely and to view its desktop, just as when you sign in to that computer locally. The primary intention of the Remote Desktop feature is remote administration. That is why, when you enable the feature, by default only the administrator who enables it can connect to the remote desktop. Other users can connect to the remote desktop only if you grant them permission.

RDS is a Windows Server role that is available only in the Windows Server operating system. To deploy RDS, you need to install at least three role services and perform an additional configuration. RDS provides a similar experience to the Remote Desktop feature, but the primary intention of RDS is to enable users to have a standard remote environment that is available from any device and to use remote resources while integrating remote applications on the local user desktop. Table 8-1 compares RDS and the Remote Desktop feature.

Table 8-1 Comparing RDS and the Remote Desktop feature

RDS	Remote Desktop Feature
Can support many simultaneous users.	Desktop operating systems are limited to one simultaneous user. Server operating systems are limited to two simultaneous users.
Proper licensing must be purchased and configured.	No additional licensing is required.
Used to access a full remote desktop or remote applications (RemoteApp).	Used to access the full remote desktop.
Supports advanced features such as RemoteFX USB Redirection and multimedia redirection.	Does not support advanced features.
Requires an infrastructure of multiple servers that has been properly planned and deployed.	Is enabled on a single computer.

RDS architecture

There are six RDS service roles that can be included in an RDS deployment. At minimum, you need to have the Remote Desktop Connection Broker (RD Connection Broker) role service, the Remote Desktop Web Access (RD Web Access) role service, and either the RD Session Host or Remote Desktop Virtualization Host (RD Virtualization Host) role service. You can install individual RDS role services, but you won't be able to manage them unless they are part of an RDS deployment. Depending on your implementation goals, an RDS deployment can include additional RDS role services, and RDS role services can be installed on multiple servers for scalability and high availability.

Windows Server 2012 R2 RDS includes the following role services:

- **RD Session Host** This role service configures a server to provide session-based desktops and applications. Users can connect to an RD Session Host server and then run applications and use the network resources that the RD Session Host offers. RD Session Host is a required role service in a session-based desktop deployment of RDS.
- **RD Virtualization Host** This role service integrates with the Hyper-V role in Windows Server 2012 R2 to provide VMs that can be used as virtual desktops. The RD Virtualization Host role service also monitors and reports on established client sessions to the RD Connection Broker role service. This role service is responsible for managing the VMs that function as pooled and personal virtual desktops. If VMs are in a saved state, the RD Virtualization Host role service starts the VMs to prepare them for a user connection. For pooled virtual desktops, the RD Virtualization Host role service reverts the VMs to their initial state when users sign out. The RD Virtualization Host role service is required in a VM-based deployment of RDS.
- **RD Connection Broker** This role service manages connections to RemoteApp programs and virtual desktops, and it directs client connection requests to an appropriate endpoint. The RD Connection Broker role service also provides session reconnection and session load balancing. For example, when a user disconnects from a session and later establishes a connection, the RD Connection Broker role service ensures that the user reconnects to his or her existing session. This role service is mandatory in all RDS deployments, but it does not require large amounts of server resources.
- **RD Web Access** This role service provides a web-based interface to RemoteApp programs, session-based virtual desktops, or VM-based virtual desktops. A webpage provides each user with a customized view of all RDS resources that have been published to that user. This role service supports organizing resources in folders, which enables administrators to group remote applications in a logical manner. It also publishes

available RDS resources in an RDWeb feed, which can integrate with the Start screen on client devices. RD Web Access is a mandatory role service for each RDS deployment.

- **Remote Desktop Licensing (RD Licensing)** This role service manages RDS client access licenses (RDS CALs) that are required for each device or user to connect to an RD Session Host server. You use RD Licensing to install, issue, and track RDS CAL availability on an RD Licensing server. You are not required to install this role service during an initial RDS deployment, but an RDS deployment without proper licensing ceases to function after 120 days.
- **RD Gateway** This role service allows authorized remote users to connect securely to RemoteApp programs and virtual desktops from outside the organization over the Internet. An RD Gateway server acts as a proxy for external users to connect to internal RDS resources. To increase compatibility with firewalls in public locations such as hotels, RDP traffic is encapsulated in Hypertext Transfer Protocol Secure (HTTPS) packets. Access is controlled by configuring Remote Desktop connection authorization policies (RD CAPs) and Remote Desktop resource authorization policies (RD RAPs). An RD CAP specifies who is authorized to make a connection, and an RD RAP specifies to which resources authorized users may connect.

All deployment and management of RDS is done by using Server Manager, as shown in Figure 8-1. Server Manager provides an overview of all servers in an RDS deployment and a management interface for each server. RDS in Server Manager uses a discovery process to detect the role services that are installed on each machine that is added to Server Manager.

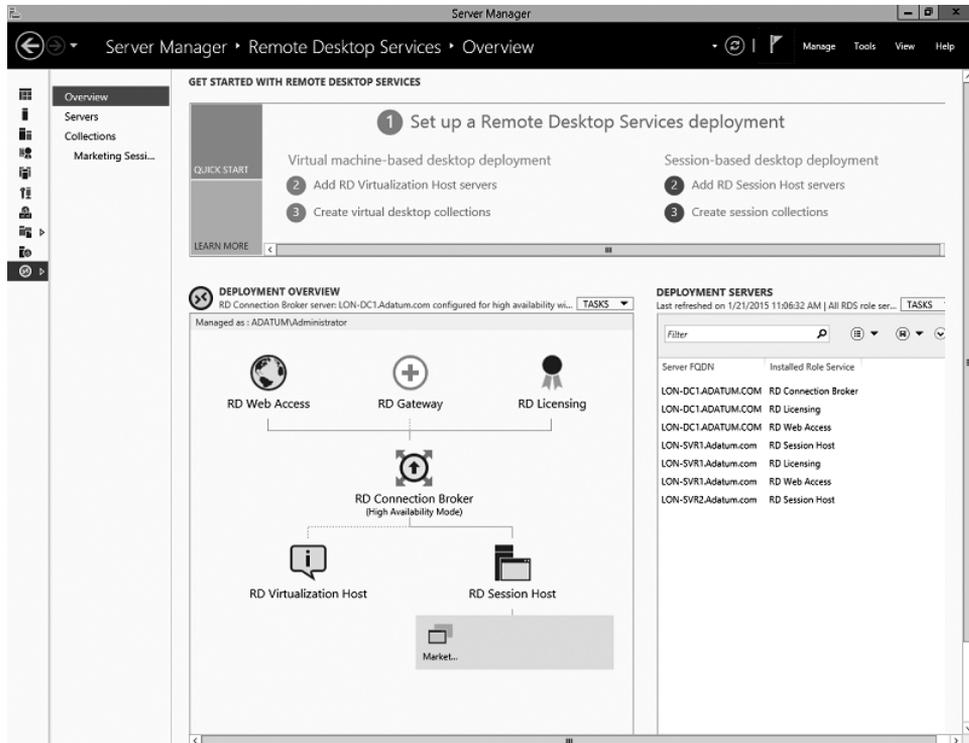


Figure 8-1 RDS configuration in Server Manager

NOTE

Legacy Remote Desktop administration tools such as Remote Desktop Services Manager and RD Session Host Configuration, which were used for configuring and administering RDS in Windows Server 2008 R2, are replaced with RDS in Server Manager in Windows Server 2012 and Windows Server 2012 R2.

Connecting to virtual desktops and RemoteApp programs

Windows client operating systems include Remote Desktop Connection (RDC), which is used to connect to virtual desktops and applications. Microsoft also provides Microsoft Remote Desktop for iOS and Android devices. All of these applications use RDP to connect to virtual desktops and RemoteApp programs.

When you use RDC to access a computer with the Remote Desktop feature enabled, you enter the IP address or DNS name of the remote computer, as shown in Figure 8-2. This type of direct connectivity doesn't work when connecting to RDS because you are connecting through

the RD Connection Broker and need to be directed to a specific collection for the RemoteApp program or virtual desktop.



Figure 8-2 Remote Desktop Connection (RDC)

After you implement servers for the RDS infrastructure, you need to create collections that define what the clients are connecting to and how it is configured. There are two types of collections:

- **Virtual desktop collections** This type of collection contains VMs hosted on RD Virtualization Host servers.
- **Session collections** This type of collection contains RD Session Host servers that provide session-based virtual desktops or RemoteApp programs.

To connect to collections in RDS, you need to have an .rdp file with the correct connectivity information for the RD Connection Broker and the collection to which you are connecting. RDC uses the connectivity information in the .rdp file.

You can create an .rdp file manually and make it available to users. When the user opens the .rdp file, RDC launches and connects to the RD Connection Broker. This method is functional but relatively complex because you need to learn the syntax for creating .rdp files and need to update them if your infrastructure changes.

The simplest way to provide user connectivity to RDS is by using RD Web Access, shown in Figure 8-3. When users connect to RD Web Access, they are provided with a list of collections to which they have access. When they click the appropriate collection, an .rdp file with the correct configuration information is generated, and RDC launches using the information in the .rdp file. This provides a consistent access method even if the RDS deployment is modified.



Figure 8-3 RD Web Access

The following process, shown in Figure 8-4, is used when clients connect to a session collection by using RD Web Access:

1. Users connect to the RD Web Access portal and identify the RDS resource to which they want to connect.
2. Users click the link on the RD Web Access portal for the RDS resource they want to access. This downloads the .rdp file, which contains information about the resource to which the user wants to connect.
3. RDC is launched, and it uses the information in the .rdp file to initiate a connection with the RD Connection Broker role service. After users authenticate to the RD Connection Broker role service, the RDC passes the request about the RDS resource to which the user wants to connect.
4. The RD Connection Broker role service examines the request to find an available RD Session Host server in the desired collection and sends the connection information back to the RDC client. If the request matches a session that already is established for the associated user, RD Connection Broker redirects the client to the server in the collection where the session was established. If the user doesn't have an existing session in the collection, the client redirects to the server that is most appropriate for the user.

connection, based on the RD Connection Broker load balancing algorithm—for example, weight factor, fewest connections, and least utilized.

5. The RDC client establishes a session with the RD Session Host server that RD Connection Broker provided.

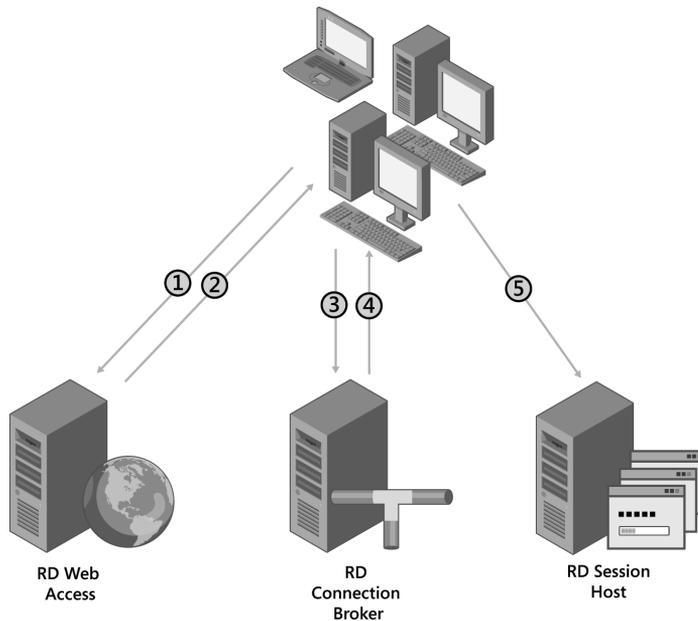


Figure 8-4 Connectivity for session collections

RDS functionality that enhances the client experience

RDC uses the RDP protocol to connect to RDS servers. The following are some of the specific features available that enhance the client experience:

- **Bandwidth reduction features** When an RDP connection is established, various methods to reduce network bandwidth are used, such as data compression and caching. Caching enables an adaptive user experience over LANs and WANs. Clients can detect available bandwidth and adjust the level of graphic detail that is used.
- **Full desktop or application window only** When a client connects to RDS, it can display either a full remote desktop or only the window of a remotely running application (RemoteApp program). With full desktops, users can perform remote administration or run multiple applications. However, the user must deal with two desktops: local and

remote. RemoteApp programs integrate with local desktops, but they still require network connectivity to RDS.

- **RemoteApp programs that look and feel like locally installed applications** The window displayed when you connect to a RemoteApp program looks like a locally installed application. Links to RemoteApp programs can be added to a client's Start screen. RemoteApp program icons support pinning, tabbed windows, live thumbnails, and overlay icons. RemoteApp windows can be transparent, and the content of a RemoteApp window displays while you are moving it.
- **Reconnection to existing sessions** If a user disconnects from a remote desktop session, the user can reconnect to the session and continue to work from the point at which he or she disconnected. The user can connect from the same device or from a different client device. If a session disconnects for a different reason, for example, because network connectivity is lost, the user automatically reconnects to the disconnected session when network connectivity is restored.
- **Redirection of local resources** Client resources such as drives, printers, the Clipboard, smart card readers, and USB devices can redirect to a remote desktop session. This enables you to use locally attached devices while working on RDS and to use the Clipboard to copy content between a local and remote desktop. You even can redirect USB devices that you plug in when the remote desktop connection already is established.
- **Windows media redirection** This feature provides high-quality multimedia by redirecting Windows media files and streams from RDS to a client. When Windows Media Player is used in a session-based virtual desktop, the multimedia file is not rendered on the RD Session Host. Instead, the multimedia stream is redirected to the RDC client and is rendered on the client. This reduces load on the RD Session Host and provides higher quality audio and video playback on the client. If the RDC client does not have the necessary codec for the multimedia content, then the content is rendered on the RD Session Host.
- **Multi-monitor support** This feature enables support for up to 16 monitors of any size, resolution, and layout. Applications function just as they do when you run them locally in multi-monitor configurations.

NOTE

Multi-monitor support requires RDC version 7.0 or later. This software is included with Windows 7 and later. If you are connecting to a computer running Windows 7, multi-monitor support is available only for the Ultimate and Enterprise editions. If you are connecting to a computer running Windows 8 or Windows 8.1, multi-monitor support is available only for the Professional and Enterprise editions.

- **Single sign-on (SSO)** When users connect to RDS, they have to provide their credentials again. With SSO, a user can connect to a remote desktop or start a RemoteApp program as the user who signed in to the local computer, without reentering credentials.
- **CPU, disk, and network Fair Share** Fair Share features are enabled by default on RD Session Host servers to ensure even resource distribution among users. One user can't monopolize resources or negatively affect the performance of other users' sessions. Fair Share can distribute network, disk, and CPU resources dynamically among user sessions on the same RD Session Host server. You can control Fair Share settings through Group Policy.

RemoteFX

RemoteFX introduces a set of enhancements to RDP that enables rich graphics and video capabilities within a remote desktop session, regardless of whether you are connecting to a session-based virtual desktop, running a RemoteApp program, or connecting to a VM-based virtual desktop. In all three cases, the user experience is almost identical to using a local physical desktop. RemoteFX is included in RDS, and you don't need to enable it explicitly unless you want to use the RemoteFX virtual graphics processing unit (vGPU) on a VM-based virtual desktop. In that case, you must add hardware to the VMs that are used for the virtual desktop.

The following is a list of some RemoteFX features:

- **RemoteFX for WAN** This feature delivers an improved user experience over lower-speed networks, such as at a branch office, on a wireless device, or working from home over a WAN connection. RemoteFX for WAN combines the RemoteFX Adaptive Graphics feature with intelligent WAN-aware transports. TCP and UDP can be used for remote desktop connections. The protocol that is better suited for the current connection is selected automatically, and automatic detection of network conditions to adjust the encoding of content is available.
- **RemoteFX Adaptive Graphics** This feature dynamically adapts to changing network conditions and optimizes encodings based on the content delivered. RemoteFX Adaptive Graphics use multiple codecs, which are optimized for different types of content, such as text, images, and video.
- **RemoteFX Media Streaming** This feature provides redirection of multimedia content. When a user attempts to play multimedia content in a remote session, the content is intercepted and redirected to the client. The client receives the compressed content, decodes the content, and plays it back locally.
- **RemoteFX Multi-Touch** This feature extends the Windows 8.1 touch experience to devices on which multi-touch is the primary means of user interaction. Windows 8.1 users are able to interact with remote desktop sessions in the same way as

a local desktop, including support for multi-touch gestures and the ability to navigate between local and remote sessions by using touch.

- **RemoteFX USB Redirection** This feature enables devices to redirect at the USB level. Because of this, no device drivers are required on the client computer, and any USB device—including audio, storage, all-in-one printers, and scanners—can be redirected.

Inside OUT

vGPU for VM-Based Virtual Desktops

An RD Virtualization Host server can suffer from high processor utilization due to graphics processing in the VMs it's hosting. This occurs because the processors in the RD Virtualization Host server do the graphics processing for each VM. This can limit the scalability of the RD Virtualization Host server.

One solution for reducing processor load on the RD Virtualization Host server is to use the vGPU functionality that is available in Windows Server 2012 R2. When you use the vGPU functionality, the VMs can use a dedicated graphics processor in the RD Virtualization Host server for graphics processing. From a performance perspective, this is like putting a more advanced video card in a desktop computer.

To use vGPU on an RD Virtualization Host server, you must meet the following requirements:

- **Second Level Address Translation (SLAT) support** The processor in the RD Virtualization Host server must include support for SLAT. For Intel processors, this is called Extended Page Tables. For AMD processors, this is called Nested Page Tables.
- **Supported video adapter** The video adapter in the RD Virtualization Host server must be DirectX 11-capable with a Windows Display Driver Model (WDDM) 1.2-compatible driver. WDDM 1.2 was introduced with Windows 8.
- **Windows 7 Enterprise with SP1 or Windows 8 Enterprise** Only the Enterprise editions of Windows 7, Windows 8, and Windows 8.1 support the use of vGPU in the VM.
- **Generation 1 VMs** The vGPU functionality isn't supported for generation 2 VMs in Hyper-V. You must configure VMs as generation 1 VMs.

If you meet the requirements for using vGPU, then you can add a RemoteFX 3D Video Adapter to VMs. In the configuration for the RemoteFX 3D Video Adapter, you can configure a maximum number of monitors and maximum monitor resolution.

Remote Desktop Connection configuration options

When you connect to a virtual desktop through RDS, RDC is configured automatically by using an RDP file that is provided by the RD Web Access server. When you use RDC to connect to a server or client with the Remote Desktop feature enabled, you can configure the connectivity settings manually. The configuration options are grouped on several different tabs. Microsoft Remote Desktop for iOS and Android have similar configuration options but different user interfaces.

On the General tab, you can specify the computer to which you want to connect by using RDC and user credentials. You also can save RDC settings in a text file with an .rdp file name extension to initiate a connection later without configuring RDC settings again.

The Display tab is shown in Figure 8-5. On this tab, you can choose the size of the remote desktop window, including the option to run the remote desktop in full-screen mode. You can select to use all local monitors for a remote session, select color depth, and enable a connection bar when the remote desktop is running in full-screen mode.

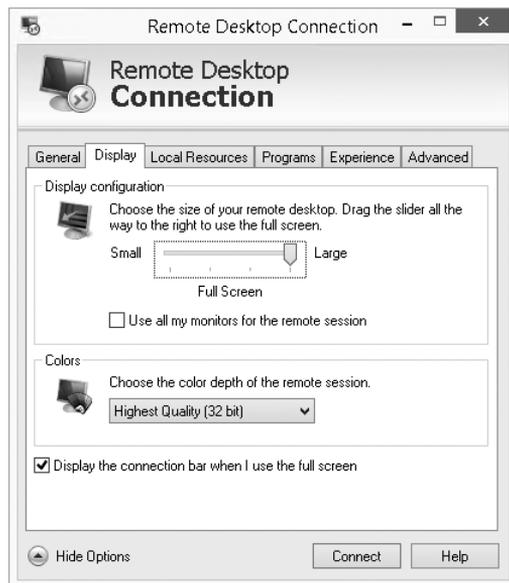


Figure 8-5 Remote Desktop Connection, Display tab

The Local Resources tab is shown in Figure 8-6. On this tab, you can set remote audio settings, such as whether you want to enable remote audio playback and recording. You also can specify a location where Windows key combinations, such as Alt+Tab, are applied and whether local devices and resources in remote sessions are available. For example, you can enable the

option to make the Clipboard, local drive, printers, and devices that you plug in later available in a remote session.



Figure 8-6 Remote Desktop Connection, Local Resources tab

On the Programs tab, you can specify a program that starts automatically in a remote desktop session when you connect to a remote computer. If you configure this option, when you close the program, your session is signed out automatically.

On the Experience tab, you can select a connection speed to optimize performance. You can enable different features, such as the following:

- Desktop background
- Font smoothing or visual styles in RDC
- Show window contents while dragging

By default, RDC automatically detects connection quality and configures connection quality–dependent features accordingly. On this tab, you also can configure persistent bitmap caching and automatic reconnection if a connection drops.

RDC displays the bandwidth with an icon on the connection bar (top of the window) that is similar to a signal strength meter. The meter is based only on bandwidth and does not take

latency into account. The number of bars in the icon identify the bandwidth, as show in Table 8-2.

Table 8-2 RDC bandwidth values

Icon	Bandwidth
4 bars	10 megabits per second (Mbps) and higher
3 bars	2000–9999 kilobits per second (Kbps)
2 bars	512 Kbps - 19999 Kbps
1 bar	Less than 512 Kbps
No icon shown	No bandwidth detected or older remote desktop host

On the Advanced tab, you can configure server authentication and Connect From Anywhere settings. The server authentication options allow you to define what should be done if the certificate provided by the server during authentication isn't valid. By default, a warning is displayed and you have the option to continue. If desired, you can configure this setting to connect without warning or prevent connections.

The Connect From Anywhere settings allow you to configure connectivity through an RD Gateway server. You can configure the alternate credentials for authentication to the RD Gateway server and the location of the RD Gateway server.

RDS licensing

If you want to use RDS, you need to purchase additional RDS CALs for each user or device that uses RDS. This is in addition to the typical licensing that is required for desktop computers. For example, in an environment where users have desktop computers and some applications are delivered by RemoteApp, you would need the following licenses:

- Operating system license for the desktop computer
- Server licenses for the Windows-based servers that deliver the RemoteApp programs
- Windows CALs for each user or computer that accesses the Windows servers
- RDS CALs for each user or desktop that uses RemoteApp programs
- Application licenses for each user or desktop that uses RemoteApp programs

RDS CALs provide users with access to session-based virtual desktops or RemoteApp programs. Licensing for VM-based virtual desktops is slightly more complex because the operating system for the VM also needs to be licensed. If you connect to a VM-based virtual desktop from a device that is covered by a Microsoft Software Assurance agreement, then the license

includes rights to use that same operating system in a VM-based virtual desktop. If the device isn't covered by a Microsoft Software Assurance agreement, then you need to purchase Windows Virtual Desktop Access (Windows VDA) licenses.

- **For more information about licensing VM-based virtual desktops, see Chapter 10, "Planning and implementing pooled and personal virtual desktops."**

When a client attempts to connect to an RDS deployment, the server that accepts the connection determines if an RDS CAL is needed. If an RDS CAL is required, then the server requests the RDS CAL on behalf of the client that is attempting the connection. If an appropriate RDS CAL is available, it is issued to the client, and then the client can connect to RDS.

RD Licensing manages the RDS CALs that are required for each device or user to connect to an RD Session Host server. You use RD Licensing to install, issue, and track the availability of RDS CALs on an RD Licensing server. At least one RD Licensing server must be deployed in the environment. The role service can be installed on any server, but for large deployments, the role service should not be installed on an RD Session Host server.

After an RDS installation, there is an initial grace period of 120 days. This grace period begins after the RD Session Host accepts the first client connection. If you have not installed valid licenses by the time the grace period expires, clients will not be able to sign in to the RD Session Host.

Inside Out

Licensing modes

Each RD Session Host server is configured with a licensing mode. The licensing mode determines the type of RDS CALs that an RD Session Host server requests from an RD Licensing server on behalf of a client that is connecting to an RD Session Host server. There are two licensing modes:

- **Per User** This gives one user the right to access any RD Session Host server in an RDS deployment from an unlimited number of client computers or devices. You should use RDS Per User CALs when the same user connects to RDS from many devices.
- **Per Device** This gives any user the right to connect to any RD Session Host server in an RDS deployment from a specific device. When a client connects to an RD Session Host server for the first time, a temporary license is issued. When the client computer or device connects to an RD Session Host server for the second time, if the license server is activated and enough RDS Per Device CALs are available, the license server issues the client computer or device a permanent RDS Per Device CAL. You should consider RDS Per Device CALs when multiple users use the same device for connecting to RDS, for example, a point-of-sale device that is used by different clerks.

A single RDS deployment can be configured with only one licensing mode. If you need a mix of Per User and Per Device RDS CALs, then you need to implement two RDS deployments.

NOTE

A permanent RDS Per Device CAL is valid for a randomly selected number of days between 52 and 89. If the RDS Per Device CAL isn't renewed, then it is returned to the available licenses on the RD Licensing server. You can revoke RDS Per Device CALs before they expire and return them immediately to the available licenses. You are limited to revoking a maximum of 20 percent of the RDS Per Devices CALs.

If you need to provide access to RDS for multiple external users who are not employees of your organization, then you should consider using an RDS External Connector License. An RDS External Connector License allows an unlimited number of nonemployees to connect to a specific RD Session Host. If you have multiple RD Session Host servers, you need multiple RDS External Connector Licenses in addition to any required Windows Server External Connector Licenses.

Planning infrastructure for session-based desktops

The planning for implementing RDS for session-based desktops can be fairly complex compared to other Windows-based role services. Most Windows-based role services require only one server. RDS requires at least three role services and, in most cases, the role services are spread across multiple servers. You should be aware of the functionality that each role service provides. You also should be aware of how an RDS deployment uses each role service. You need to know role service requirements and which hardware resources are most critical for each role service.

Assessing RDS infrastructure requirements

Before you implement RDS, you must determine your organization's requirements. To do so, you first must evaluate if RDS is the appropriate solution for your needs, and then you must choose between session-based and VM-based desktop deployments. If necessary, an RDS deployment can include both session-based and VM-based desktop deployments. You also must evaluate the existing server infrastructure and estimate the required server hardware, network bandwidth, client types and requirements, and connectivity needs for a successful RDS deployment.

Determine your RDS needs

To determine if RDS is an appropriate solution for your needs, you should assess and analyze the types of users, hardware, and applications in your organization. Areas of consideration include the following:

- **User types** Do you have users in remote locations, single-task users, contractors, and other types of users who would benefit from remote applications or virtual desktops?
- **Hardware** What client hardware currently is deployed in your organization? Would it be beneficial to move from traditional desktops to thin clients for some users? Do you allow users to bring their own devices into the organization's network? Do users wish to use mobile devices to run certain applications?
- **Application compatibility** Can the applications run in a multiuser environment? If not, will the applications run in a virtual environment?
- **Application performance** How do the applications perform in a remote or virtual environment? Keep in mind that many applications perform better as RemoteApp programs on RDS because processing takes place on a server.
- **Application support** Do vendors support the applications in a virtual or multiuser environment? Do vendors provide support to multiple users?
- **Licensing** Can the applications be licensed for a virtual or multiuser environment?
- **Business benefits** Are there justifiable business reasons to implement this solution? Potential benefits include cost savings, reduced deployment time, centralized management, and reduced administration costs.
- **Legal requirements** Because of financial and legal requirements, some organizations mandate that applications and data remain on-premises. RDS enables users to connect to a standard virtual desktop to use familiar applications and to work with data from almost any device, while organizational data stays in the data center.

Choosing between session-based and VM-based desktop deployments

RDS has two deployment types:

- **Session-based virtual desktop deployment** This provides users the ability to connect to an RD Session Host and use a full desktop or run remote applications and present them on a client as if they were installed locally.
- **VM-based virtual desktop deployment** This provides users with access to a full Windows client operating system that runs on a VM, for example, Windows 7 or Windows 8.1.

You need to decide which RDS deployment type is best for your environment based on various requirements. For example, you must consider if users must be completely isolated or if they must have administrative access. You should consider whether the applications work properly in a multiuser environment. In addition, you must consider whether you can install and run

applications on Windows Server. Remember that a VM-based virtual desktop deployment typically requires a more powerful server infrastructure and more disk storage than a session-based virtual desktop deployment for the same number of users. For some applications, VM-based virtual desktops might be the only viable solution.

Generally, you should choose session-based virtual desktops if possible. Session-based virtual desktops support a larger number of users than VM-based virtual desktops on the same hardware.

Determine server hardware and network resource requirements

Once you determine the RDS deployment benefits for your organization, you must consider the hardware requirements to support your users, including the following:

- **Number of users** How many users will use RDS, and where are they located?
- **User types** How many users run CPU-intensive and bandwidth-intensive applications? Will you have to provide more bandwidth and server hardware to support expected usage?
- **Connection characteristics** How many concurrent connections do you expect? Can your server and bandwidth resources handle peak usage times?
- **Application silos** Will you have to create multiple server collections to support different applications that might not be able to run on the same server?
- **Load balancing** Will you have to include multiple servers in a collection to spread the load among the servers? This increases available resources and provides redundancy.
- **High availability** What is the organization's tolerance for downtime? Do you need close to zero downtime, or could your organization tolerate the time it would take to restore from backups?
- **Expansion considerations** What are the growth expectations? At what point will new resources need to be brought online?

Determine user requirements

Another aspect to consider is user requirements. A large organization with multiple locations might have a number of mitigating factors to consider, such as the following:

- **Languages** Organizations with a global presence need to support multiple languages. You might need to install language packs on all of your RDS servers.

- **Profile management** How will you store user states? Do users require the same user state when they sign in locally and to an RDS session? Which type of Windows user state virtualization will be used?
- **Printing** Will existing printers function properly in a remote desktop environment? Will there be problems finding printer drivers to support existing printers? Is there a budget to replace older printer models?

Determine how clients access RDS

Clients can connect to RDS in various ways. You probably will need to provide different access methods for different groups of users. Areas to consider include the following:

- Will you allow users to connect over the Internet from remote locations? If so, you will need to set up an RD Gateway and obtain certificates.
- How will you handle Secure Sockets Layer (SSL) certificates—by using certificates from non-Microsoft certification authorities (CAs) or by using certificates that an internal CA issues?

Based on your assessment results, start designing your RDS deployment. You should identify RDS role services that are required and that you will deploy. You also should determine the number and hardware configuration of servers that are required, in addition to planning required storage, connectivity, and firewall configuration.

Planning for the RD Session Host role service

The RD Session Host role service provides Windows-based apps or full Windows desktops for RDS clients. This role service is mandatory for every RDS deployment that provides users with session-based desktops or RemoteApp programs. An RD Session Host server accepts incoming RDP requests, and after a client authenticates, it provides a desktop-based or application-based session to the client. An RD Session Host server is the central location where remote applications are installed, accessed, and maintained.

To plan the deployment of an RD Session Host server, you must consider the number of installed applications, the type of applications, resource use, the number of connected clients, and the type of user interaction. While connected to one RD Session Host, users might run a simple application that has low resource utilization and rarely runs, for example, an old data entry application. On another RD Session Host, users often might run a resource-intensive graphical application that requires many CPU resources, a considerable amount of RAM, intensive disk I/O operations, and that causes a lot of network traffic. If the hardware configuration on both of the RD Session Hosts is the same, the second server is considerably more utilized and can accept fewer user connections.

RD Session Host planning focuses on the number of concurrent users and the workload they generate. A server with a particular hardware configuration might support many simultaneous users or only a few, depending on their usage patterns and the applications that they are running on the RD Session Host.

The following are the main resources that you should consider when estimating RD Session Host utilization:

- **CPU** Each remote application that users start runs on an RD Session Host and utilizes CPU resources on the RD Session Host. In an environment where many users are connected to the same host, CPU and memory typically are the most critical resources.
- **Memory** Additional memory must be allocated to an RD Session Host for each user who connects to the RD Session Host, whether connecting to a full Windows desktop or running a RemoteApp program.
- **Disk** Because user state typically isn't stored on an RD Session Host, disk storage usually isn't a critical resource. However, many applications run simultaneously on an RD Session Host, and the disk subsystem should be able to meet their disk I/O needs.
- **Network** The network should provide enough bandwidth for connected users and for the applications that they run. For example, applications that use a SQL database use the network to connect to that SQL database. Also remember to consider the network bandwidth required to support the user connectivity to the RD Session Host server.
- **GPU** Applications that are graphically intensive, especially those that include three-dimensional graphics, might require vGPU support and RemoteFX to perform well. Without such support, graphics render on the server's CPU and may limit the number of users on the RD Session Host to a relatively small number.

When estimating the required resources for an RD Session Host, you can use one of the following methods:

- **Pilot deployment** This is a common and a simple approach. You first need to deploy RDS in a test environment and capture its initial performance. After that, you start increasing server load by increasing the number of users and monitoring response times and user feedback. You can find out how many users can connect to an RD Session Host and still have an acceptable user experience based on the number of users and the system response time. Based on the findings, you can estimate the number of servers that are needed for a production environment. This approach is reliable and simple, but it requires initial investments for the pilot deployment.
- **Load simulation** This method also uses an initial RDS deployment in a test environment. You need to gather information about applications that users operate and how

users interact with the applications. After that, you can use load simulator tools to generate various levels of typical user loads against an RDS deployment. When a load simulator tool runs, you need to monitor server utilization and responsiveness. This method is similar to the pilot deployment method, but it uses a load simulation tool instead of real users to generate user load. It also requires an initial investment, and its results depend on the initial estimation of actual user usage.

- **Projection based on single-user systems** This method uses data that is collected from a single-user system for projecting expected utilization on an RD Session Host with multiple user sessions. This method requires detailed knowledge of applications that are used, and it usually is not very reliable because a single-user system has a different overhead than a multiuser system.

It is critical that you plan for future scalability of an RDS deployment. User needs for applications will change over time, and you need to be ready to expand your RDS deployment to meet those needs. In some cases, you may be able to scale up the capacity of the individual servers with additional processors or additional memory. Scaling up by using more powerful servers tends to be expensive. Scaling out by adding servers generally is less expensive.

Fortunately, you can scale out an RDS deployment for session-based virtual desktops and RemoteApp programs by adding RD Session Host servers. For example, if you have an RDS deployment for session-based virtual desktops that uses two RD Session Host servers, and those two servers are experiencing frequent peaks of 100 percent CPU utilization, you can add a third RD Session Host server. The RD Connection Broker then automatically load balances the connections across three servers instead of two and reduces the CPU utilization on the two existing servers.

Planning for the RD Connection Broker role service

During RDS deployment planning, you must designate a server on which to install the RD Connection Broker role service. The RD Connection Broker role service is required in each RDS deployment. It provides users with access to RemoteApp programs, session-based virtual desktops, and VM-based virtual desktops. The RD Connection Broker role service manages all aspects of session connectivity. Functions performed by the RD Connection Broker role service include the following:

- **Routes connection requests** Determining the most appropriate RD Session Host or virtual desktop to which to send a connection request, based on a user's identity and the current load on RD Session Host or RD Virtualization Host servers.
- **Stores information about connections to VMs and sessions** By default, connection information is stored in the Windows Internal Database (WID) on an RD Connection Broker server. By storing this information, the RD Connection Broker role

service can reconnect users to the same session in an RDS deployment with multiple RD Session Host servers.

- **Configures RDS servers in the same group (collection)** You configure settings—for example, session settings or certificates—once, and RD Connection Broker applies the settings to servers in the collection.
- **Manages VM creation and deletion** In VM-based desktop deployments, RD Connection Broker manages VM creation and deletion for managed collections, and it assigns personal virtual desktops to users.
- **Provides information to RD Web Access servers** The RD Connection Broker role service gathers collection information about RemoteApp programs, session-based virtual desktops, and VM-based virtual desktops.

When a user initiates a session, the session request is received by the RD Connection Broker role service, which queries the database to determine if there is an existing disconnected session for that user. If so, the user is directed to the disconnected session. If not, the RD Connection Broker role service determines the server in the collection that is best able to handle the new connection, based on the load-balancing algorithm.

Inside OUT

RD Connection Broker scalability

Performance for an RD Connection Broker server depends on the number of requests it receives in a given time frame. Users generate a load on the RD Connection Broker server only when they perform the initial connection to a RemoteApp program or a session-based virtual desktop. This makes it different from an RD Session Host server, for which performance depends on the number of simultaneous users.

A single RD Connection Broker server with four processor cores and four gigabytes (GB) of RAM can process approximately 10 connections per second with a maximum processing time of 1 second. If you are willing to increase the maximum processing time to 12 seconds, that same server can process approximately 60 connections per second. For more detailed information about RD Session Broker scalability, see the white paper “RD Connection Broker Performance and Scalability” at <http://go.microsoft.com/fwlink/?LinkID=510038&clcid=0x409>.

When you consider the number of connections per second to which your RD Connection Broker server will be subject, you need to consider peak utilization times. First thing in the morning when users arrive and after lunch likely are peak times when users are connecting.

A single RD Connection Broker server can handle a large number of connection requests, and for performance, your RDS deployment may require only one. A more critical consideration for the RD Connection Broker role service is availability.

The RD Connection Broker role service is an entry point to an RDS deployment, and it is critical that it is available all the time. If the RD Connection Broker role service isn't available, then clients can't connect to RemoteApp programs or virtual desktops, but existing connections to RemoteApp programs and virtual desktops continue to function properly. When an RDS deployment only has one RD Connection Broker server, the server represents a single point of failure. To make the RD Connection Broker role service highly available or to increase scalability, you can add RD Connection Broker servers.

- **Configuring high availability for the RD Connection Broker role service is covered in more detail later in this chapter in the section titled "Understanding high availability for RDS."**

Planning for the RD Web Access role service

The RD Web Access role service is a mandatory part of each RDS deployment, and it installs the Web server role, Internet Information Services (IIS), as its prerequisite. The benefits of RD Web Access include the following:

- From almost anywhere, authorized users quickly can access a list of available RemoteApp programs, remote desktops, and virtual desktops on a webpage.
 - A list of available RDS resources publishes automatically via an RDWeb feed, and it can integrate with the Start screen on the client.
 - Changes in available RDS resources update automatically on clients that have subscriptions to an RDWeb feed.
 - Users can launch the RDC client from the RD Web Access portal, which enables them to connect remotely to the desktop of any computer on which they have Remote Desktop access.
 - RD Web Access and RDWeb feeds are personalized and show only RDS resources for which users have permissions.
 - Administrators can customize an RD Web Access portal without programming.
- **More information about customizing RD Web Access is provided in Chapter 9, "Configuring RemoteApp programs and client connectivity."**

It's important to remember that the RD Web Access role service only provides a link to launch RemoteApp programs or to connect to a Remote Desktop session. The RD Web Access role

service doesn't proxy client requests. When a user connects to a RemoteApp program or a virtual desktop, the client establishes a direct connection to the target server.

Performance considerations for an RD Web Access server are similar to those for an RD Connection Broker server because the RD Web Access role service provides only initial connectivity to RemoteApp programs and virtual desktops. After users are connected to requested resources, the RD Web Access role service is no longer used. Therefore, RD Web Access server performance needs to be designed to accommodate usage at peak times like morning arrivals and after lunch. If required for high availability or scalability, you can implement multiple RD Web Access servers and load balance them.

Planning for preserving user state

In a session collection with multiple RD Session Host servers, the connections from clients are load balanced across the RD Session Host servers by the RD Connection Broker server. By default, when a user connects to a specific RD Session Host server, a local profile is created for that user on the RD Session Host server. The next time a user connects, the RD Connection Broker may direct the client to a different RD Session Host server, where a different local profile is created. Each time users sign in, they may be using a different profile on a different RD Session Host server. This means that user state information such as application configuration, Desktop configuration, Favorites, and Documents are not the same across sessions. To provide a consistent user experience, you should preserve user state across multiple RD Session Host servers.

If users have desktop computers and session-based virtual desktops, you also need to consider whether you want user state to be preserved between desktop computers and the virtual desktops. This can be complicated by the fact that session-based virtual desktops may not have the same configuration as the desktop computers, and, consequently, it may not make sense to synchronize all of the user state information. For example, synchronizing Desktops may result in desktop shortcuts appearing that point to applications that are not available on the RD Session Host servers.

Roaming profiles

Roaming user profiles can be used to synchronize user state, but they synchronize entire user profiles. This typically is not desired for session-based desktops because not all user state information needs to be synchronized between desktop computers and RD Session Host servers. If you use roaming profiles for the desktop computers in your organization and you want to ensure that roaming profiles are not used on the RD Session Host servers, then you can configure the `msDS-PrimaryComputer` attribute for users and enable the Download Roaming Profiles On Primary Computers Only Group Policy setting.

You also can set user properties for roaming user profiles that are specific to RD Session Host servers, as shown in Figure 8-7. If you configure the Profile Path, then a user connecting to

a session-based virtual desktop uses the specified profile path rather than a roaming profile configured on the Profile tab. Effectively, the RDS user profile becomes a roaming profile used only when connected to an RD Session Host server.

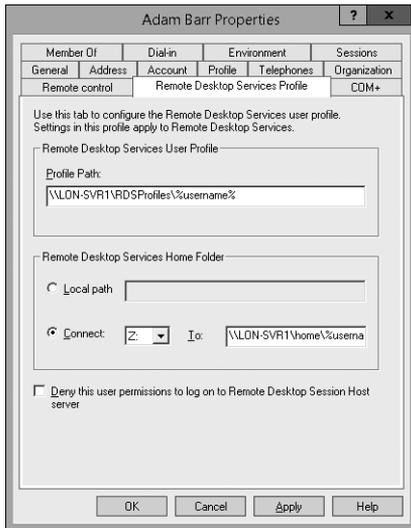


Figure 8-7 User Properties, Remote Desktop Services Profile tab

Instead of configuring individual user accounts with RDS-specific profiles, you can use Group Policy. In a Group Policy object that applies to the RD Session Host servers, you can configure settings in Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Profiles. There are two relevant settings:

- **Set Path For Remote Desktop Services Roaming User Profile** Specify a UNC path for storing all user profiles. A subfolder for each user is created automatically.
- **Use Mandatory Profiles On The RD Session Host Server** Indicates that the path specified in the Set Path For Remote Desktop Services Roaming User Profile setting contains a mandatory profile that can't be modified. When this setting is enabled, the UNC path for profiles does not contain subfolders for each user.

Folder redirection

Folder redirection also is an option for users with session-based virtual desktops. You can redirect only the folders that are suitable for use on the virtual desktops and desktop computers. Commonly redirected folders include Documents, Favorites, and AppData\Roaming.

If you use folder redirection for desktop computers and don't want folder redirection used when users sign in to the RD Session Host servers, you can use the `msDS-PrimaryComputer` attribute in user accounts just as you can for roaming profiles. In addition to configuring the attribute, you need to enable the `Redirect Folders On Primary Computers Only` Group Policy setting.

User profile disks

RDS in Windows Server 2012 and newer offers the option to configure user profile disks to preserve user state across sessions. A user profile disk is a VHDX file that is mounted to the user's profile path at `C:\Users\%username%` on the RD Session Host. The user profile disk is mounted during sign in. During a user's session, all changes to the profile write in his or her VHDX file, and when the user signs out, his or her profile disk is unmounted. The administrator specifies the maximum size of user profile disks and can limit which folders in a user profile are included in or excluded from a user profile disk.

User profile disks are configured individually for each session collection and can't be shared among collections. A share is specified in the collection configuration to store the user profile disks. All RD Session Host servers in the collection have access to the user profile disks in the share. This provides users with consistent user state from any RD Session Host server in the collection.

User profile disks can be used in conjunction with folder redirection and roaming user profiles. Folder redirection will reduce the size of user profile disks and allow the redirected folders to be accessed from desktop computers. Roaming user profiles are synchronized with the user profile disk.

From a server management perspective, one benefit of user profile disks is controlling the amount of data stored on the C drive of RD Session Host servers. Large user profiles stored on RD Session Host servers can cause the C drive to run low on space and cause performance issues. Because user profile disks are stored on a network share and mounted in `C:\Users`, the C drive never is used to store profile data.

The primary consideration when planning user profile disks is ensuring that the necessary disk space is available for network storage. To ensure that network storage is sufficient, you need to determine the average user profile size. The amount of storage that you need to allocate for user profile disks is the average user profile size times the number of users plus an allowance for growth in both the number of users and the average profile size.

User profile disks are dynamically expanding VHDX files. By default, the maximum size of a user profile disk is 20 GB, but you can set this to be larger or smaller depending on the needs of your users.

When you configure the share for user profile disks, all RD Session Host servers need to have Full Control permissions. This allows the RD Session Host servers to create and manage the user profile disks. When you configure a collection with user profile disks, these permissions are assigned automatically.

Inside OUT

Infrastructure testing prior to rollout

After you assess RDS infrastructure requirements and familiarize yourself with RDS and its role services, you should perform a proof-of-concept (POC) deployment. POC deployment is critical for a successful RDS deployment. It enables you to evaluate whether all the requirements are met and to perform a load simulation, which simulates typical user actions and validates your estimates for capacity, application workloads, and usage patterns by performing a test run in a controlled environment. During testing, you should find answers to the following questions:

- **How many users can connect, and what is an average response time?** Can POC deployment support the expected number of RD users, and is the response time acceptable? How utilized are servers, how long do user sign-in and sign-out take, and is the user experience as expected?
- **How does the application consume system resources?** Does it do so in accordance with documented estimates? If the application uses hardware as expected, the rest of the deployment can continue based on initial estimates. If it doesn't use hardware as expected, you must recalculate capacity requirements to ensure accurate estimates.
- **Are all of the potential user environment scenarios being tested?** You should test the application by accessing it in all the ways a user might use it. If there are access methods that you can't replicate in a POC environment, these access methods should be implemented in a controlled manner when performing the final deployment.
- **Are the applications and hardware running as expected?** Is additional performance tuning required? Do you need to perform any additional application configuration to run as expected in an RDS environment? Also, confirm that hardware performance is within estimated parameters.
- **Are there any unexpected changes in usage or access?** If any part of the presentation virtualization POC deployment does not reflect your production environment, alter the POC deployment so that it is as similar as possible to your final, planned infrastructure.

Using testing to eliminate errors in a deployment is important because problems with a presentation virtualization environment are much easier to resolve during testing than during full deployment.

Windows Server 2012 R2 includes a Best Practices Analyzer (BPA) for the Remote Desktop Services server role. BPA for RDS can analyze an RDS environment and check for changes that need to be made for RDS to perform optimally. You can access BPA in Server Manager or by running the `Invoke-BpaModel` cmdlet.

Deploying session-based virtual desktops

RDS includes multiple role services. If you use Server Manager for RDS deployment, you should be aware that if you use role-based or feature-based installation, you can install individual RDS role services. However, if you install an RDS role service in this way, you can't manage it. If you want to manage RDS, a deployment must have at least three role services: RD Connection Broker, RD Web Access, and either RD Session Host or RD Virtualization Host. Individual RDS role services can't be managed if they are not part of an RDS deployment.

Understanding the session-based desktop deployment process

You can deploy RDS by using Server Manager or Windows PowerShell. Server Manager has the ability to install the necessary server roles, role services, and features on multiple servers that are part of an RDS deployment. All management of RDS also can be done from Server Manager.

Inside Out

Adding servers to Server Manager for RDS deployment

Server Manager can be used to manage the local server and remote servers, but you need to add the remote servers manually before they can be managed. A typical RDS deployment has multiple servers, and you should add each of the servers to Server Manager before you begin the deployment process.

To add a server to Server Manager, perform the following steps:

1. In Server Manager, click **Manage** and click **Add Servers**.
2. In the **Add Servers** window, on the **Active Directory** tab, in the **Name (CN)** box, type the name of the server and click **Find Now**.
3. Double-click the server you want add and then click **OK**.

The high-level steps for deploying session-based virtual desktops are as follows:

1. **Start the RDS installation** In Server Manager, use the Add Roles And Features Wizard to select the Remote Desktop Services Installation option, shown in Figure 8-8. This option configures the wizard to collect the information necessary to perform a deployment of RDS across multiple servers.

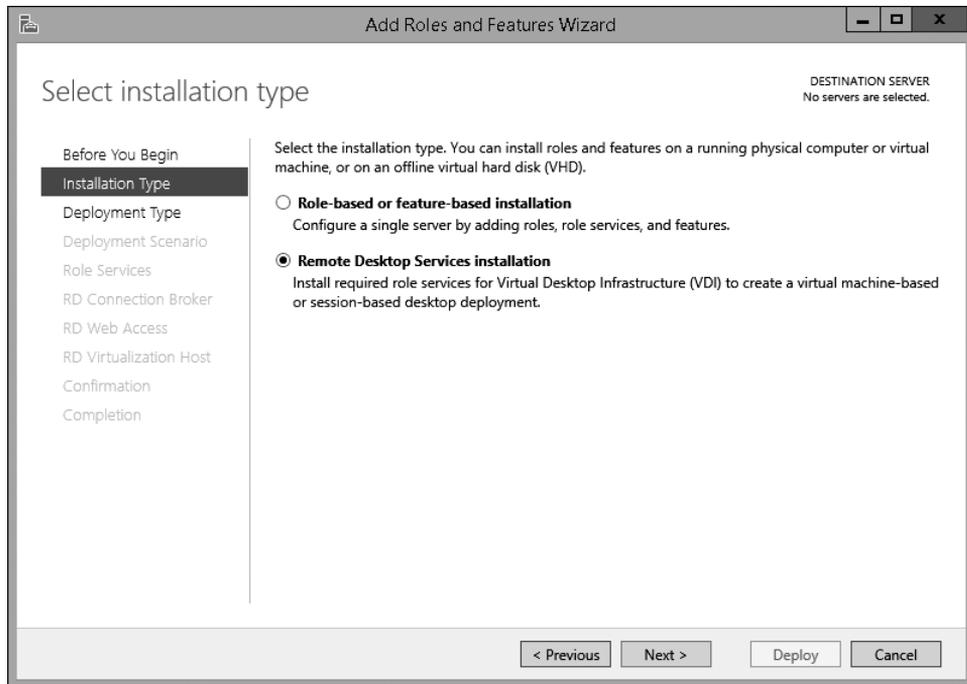


Figure 8-8 Add Roles And Features Wizard, Select Installation Type page

2. **Select the RDS deployment type** On the Select Deployment Type page, shown in Figure 8-9, select the appropriate deployment type. The Quick Start option installs the required role services on a single server and creates a session collection with several sample RemoteApp programs (Calculator, Paint, and WordPad). You only should use the Quick Start option for testing. In most cases, you want to select the Standard Deployment option because this allows you to customize the deployment for your environment.

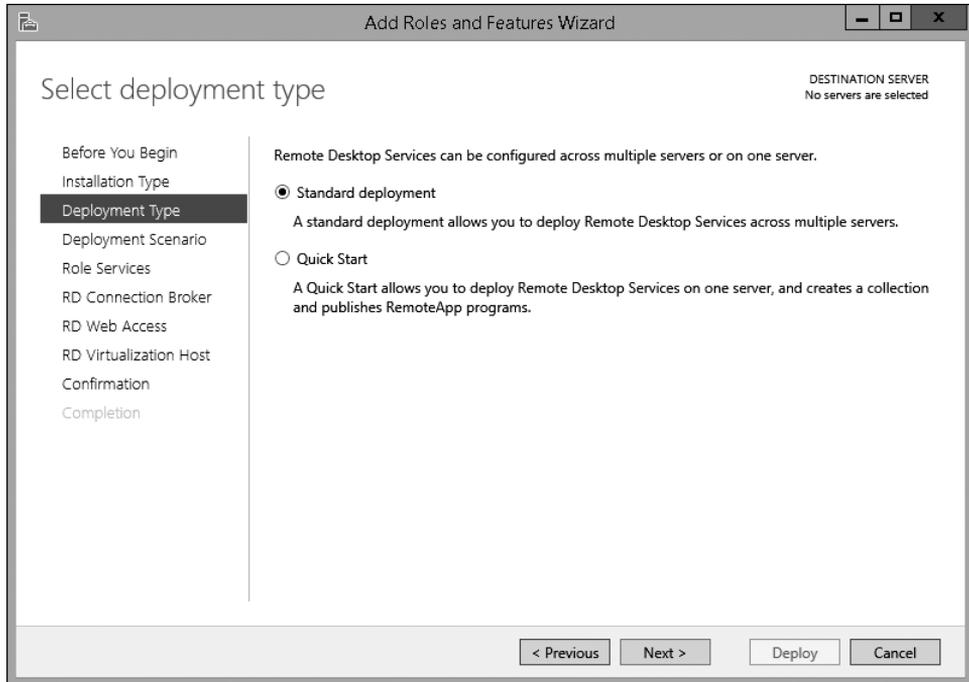


Figure 8-9 Add Roles And Features Wizard, Select Deployment Type page

3. **Select the RD deployment scenario** On the Select Deployment Scenario page, shown in Figure 8-10, select the Virtual Machine–Based Desktop Deployment option or the Session-Based Desktop Deployment option. A VM-based desktop deployment is used to deploy personal and pooled virtual desktops on computers running Hyper-V. A session-based desktop deployment uses RD Session Hosts.

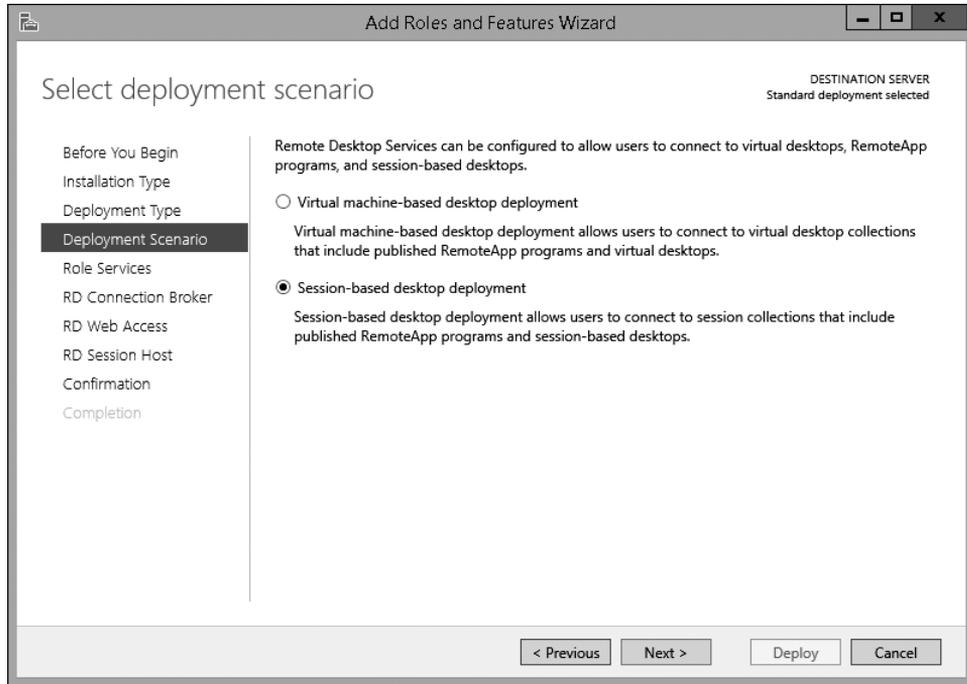


Figure 8-10 Add Roles And Features Wizard, Select Deployment Scenario page

4. **Select servers for RDS role services** In the Add Roles And Features Wizard, select the servers on which you want to install the RD Connection Broker, RD Web Access, and RD Session Host role services. As part of making RDS highly available, you can install each role service on multiple servers. In most RDS deployments, the RD Session Host role service isn't combined with other role services. The RD Connection Broker and RD Web Access role services can be combined in smaller RDS deployments.

During the deployment, the servers on which you installed the RD Session Host role are restarted. After the installation, you can perform initial configuration of the RDS deployment. You also can add servers to the deployment. At minimum, you should add RD Licensing, because you can't connect to an RD Session Host without valid RDS CALs after the initial grace period of 120 days expires. You also should consider installing multiple instances of the RDS role services for high availability.

Inside OUT

Using Windows PowerShell to deploy RDS

You can use Windows PowerShell to deploy RDS on Windows Server 2012 or newer. In Windows Server 2012, use the `New-SessionDeployment` cmdlet. In Windows Server 2012 R2, use the `New-RDSessionDeployment` cmdlet.

- ▶ You can use Windows PowerShell to manage all aspects of deploying and managing RDS. For more information about cmdlets available for deploying and managing RDS in Windows Server 2012 R2, see the Remote Desktop Cmdlets in Windows PowerShell page on TechNet at <http://technet.microsoft.com/en-us/library/jj215451.aspx>.

To install a session-based deployment of RDS, perform the following steps:

1. In Server Manager, click Manage and then click Add Roles And Features.
2. In the Add Roles And Features Wizard, on the Before You Begin page, click Next.
3. On the Select Installation Type page, click Remote Desktop Services Installation and click Next.
4. On the Select Deployment Type page, click Standard Deployment and click Next.
5. On the Select Deployment Scenario page, click Session-Based Desktop Deployment and click Next.
6. On the Review Role Services page, click Next. This page provides a brief description of each role service, but there is nothing to configure. The currently logged-on account is being used to create the deployment and is displayed here as a reminder.
7. On the Specify RD Connection Broker Server page, shown in Figure 8-11, in the Server Pool box, double-click the server on which you wish to install the RD Connection Broker role service and click Next.

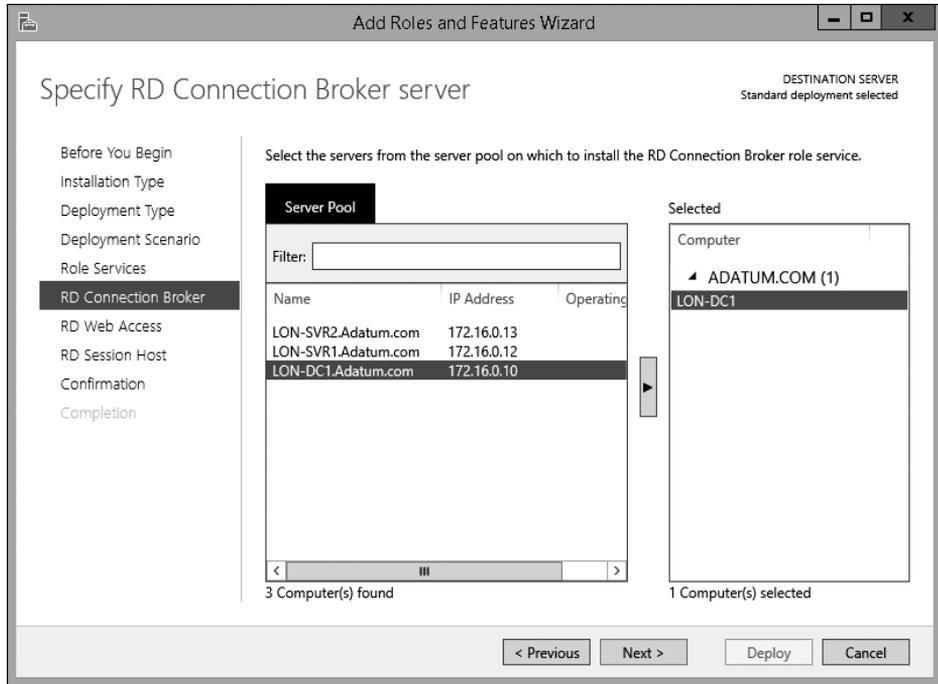


Figure 8-11 Add Roles And Features Wizard, Specify RD Connection Broker Server page

8. On the Specify RD Web Access Server page, shown in Figure 8-12, select the Install The RD Web Access Role Service On The RD Connection Broker Server check box and click Next. Alternatively, you can select another server on which to install the RD Web Access role service.

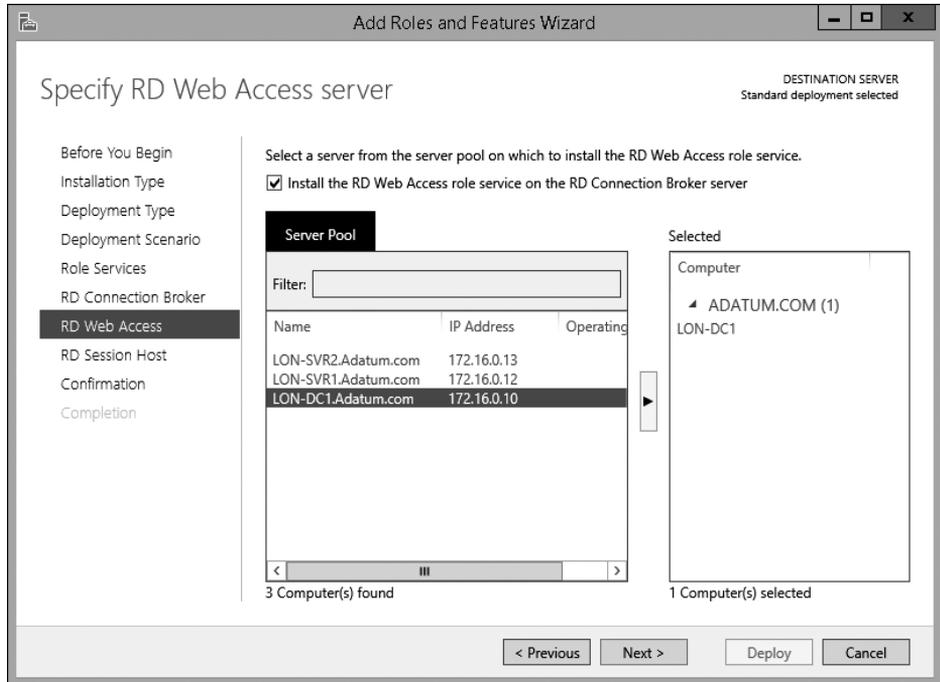


Figure 8-12 Add Roles And Features Wizard, Specify RD Web Access Server page

9. On the Specify RD Session Host Servers page, shown in Figure 8-13, double-click the server on which you wish to install the RD Session Host role service and click Next.

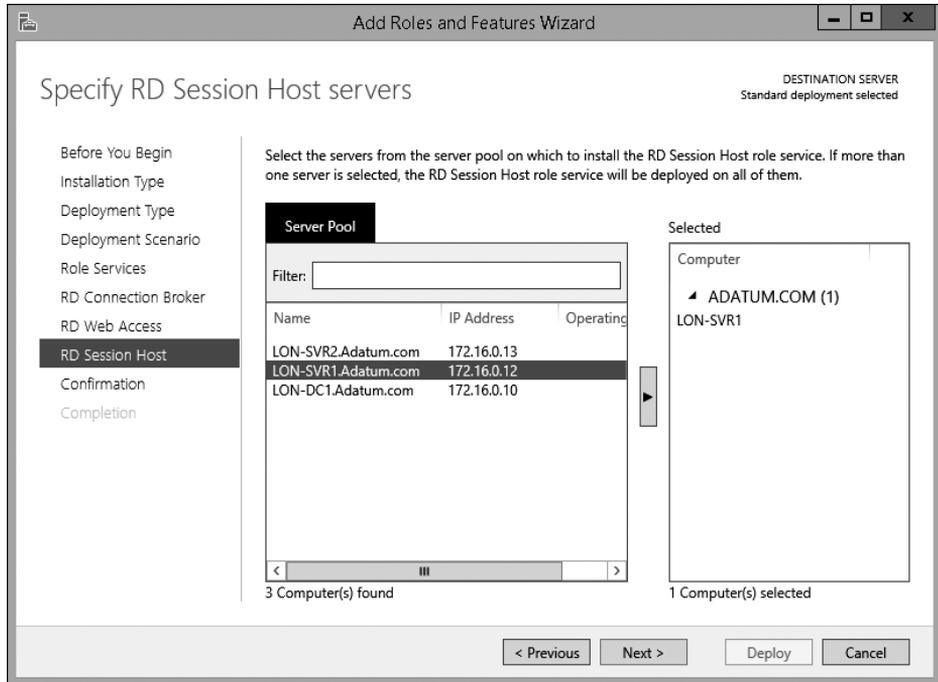


Figure 8-13 Add Roles And Features Wizard, Specify RD Session Host Servers page

10. On the Confirmation page, review the selected servers, select the Restart Destination Server Automatically If Required check box, and click Deploy.
11. On the Completion page, wait for the installation of the RDS role services to complete and click Close. If you are installing roles on the server from which you started the installation, the server may restart and require you to sign in again.

Understanding session collections

Session collections enable you to organize and control user connectivity to RDS. Each session collection contains either RD Session Host servers for session-based virtual desktops or VMs on Hyper-V for pooled or personal virtual desktops.

Collections simplify the administration process by enabling you to manage all collection members as a unit instead of managing them individually. For example, after you configure a collection with session settings, those settings automatically apply to all the servers in the collection. If you add a server to a collection, session settings also automatically apply to the added server.

When you add multiple RD Session Host servers to a collection, connections automatically are load balanced among them. The RD Connection Broker server uses the collection configuration information to identify that there are multiple RD Session Host servers and connects an equal number of clients to each. If an RD Session Host server in a collection fails, the RD Connection Broker connects all users to the remaining RD Session Host servers in the collection.

When there are multiple RD Session Host servers in a collection, they need to be configured with identical applications. Users expect the same applications to be available each time they sign in. If RD Session Host servers have different applications installed, it will appear to users that applications are randomly appearing and disappearing with each connection.

To create a session collection, perform the following steps:

1. In Server Manager, in the navigation pane, click Remote Desktop Services.
2. In Remote Desktop Services > Overview, click Create Session Collections.
3. In the Create Collection Wizard, on the Before You Begin page, click Next.
4. On the Name The Collection page, in the Name box, type the name of the collection and click Next. Make the name something that accurately describes how the collection will be used. You also can type in a more detailed description.
5. On the Specify RD Session Host Servers page, shown in Figure 8-14, double-click the RD Session Host server you want to add to the collection and click Next. Only RD Session Host servers already added to the RDS deployment appear in the Server Pool box. An RD Session Host server can be added to only one collection.

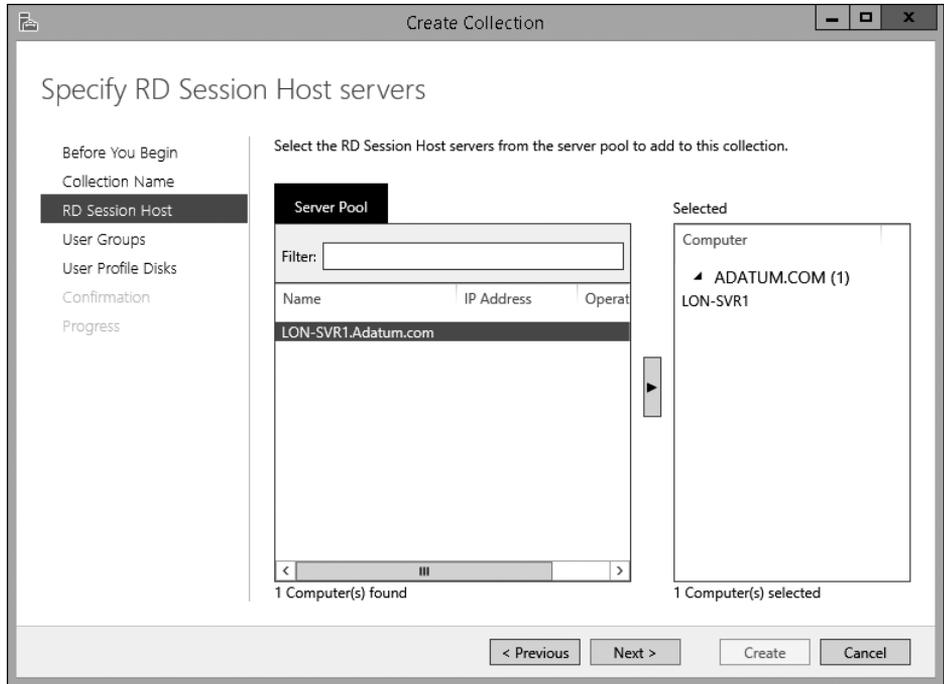


Figure 8-14 Create Collection Wizard, Specify RD Session Host Servers page

6. On the Specify User Groups page, shown in Figure 8-15, remove the Domain Users group, add the groups you want to have access to the collection, and then click Next. The Domain Users group is listed by default and would allow any user in your organization to access the collection. In most cases, you want to restrict collection access to a specific group of users.

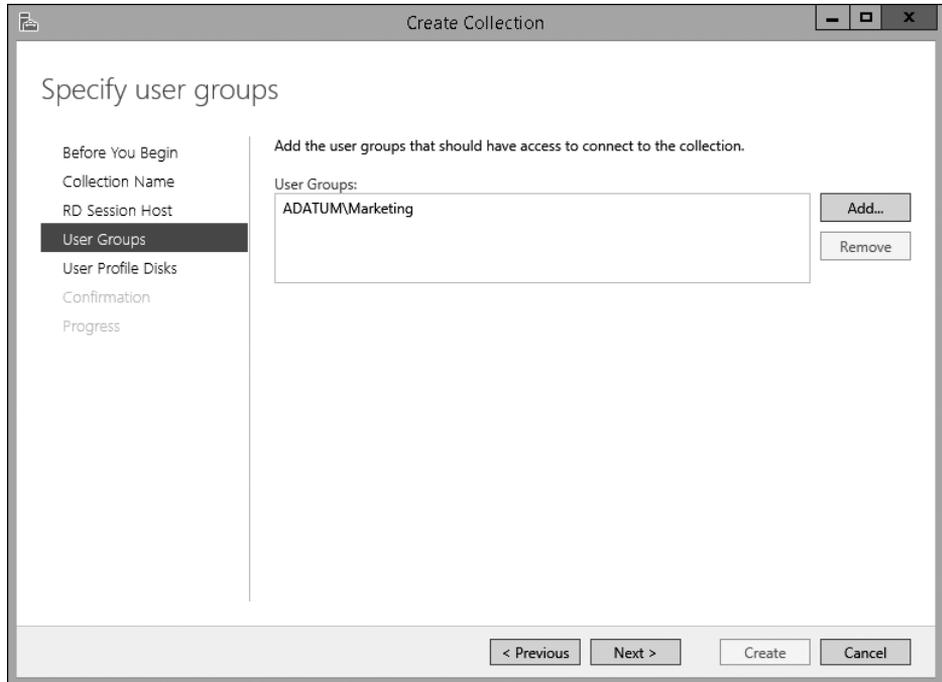


Figure 8-15 Create Collection wizard, Specify User Groups page

7. On the Specify User Profile Disks page, shown in Figure 8-16, select the Enable User Profile Disks check box if you have decided to implement user profile disks for users. If you select this option, you need to enter the UNC path where the user profile disks will be stored in the Location Of User Profile Disks box. You also need to specify a size in the Maximum Size (In GB) box.

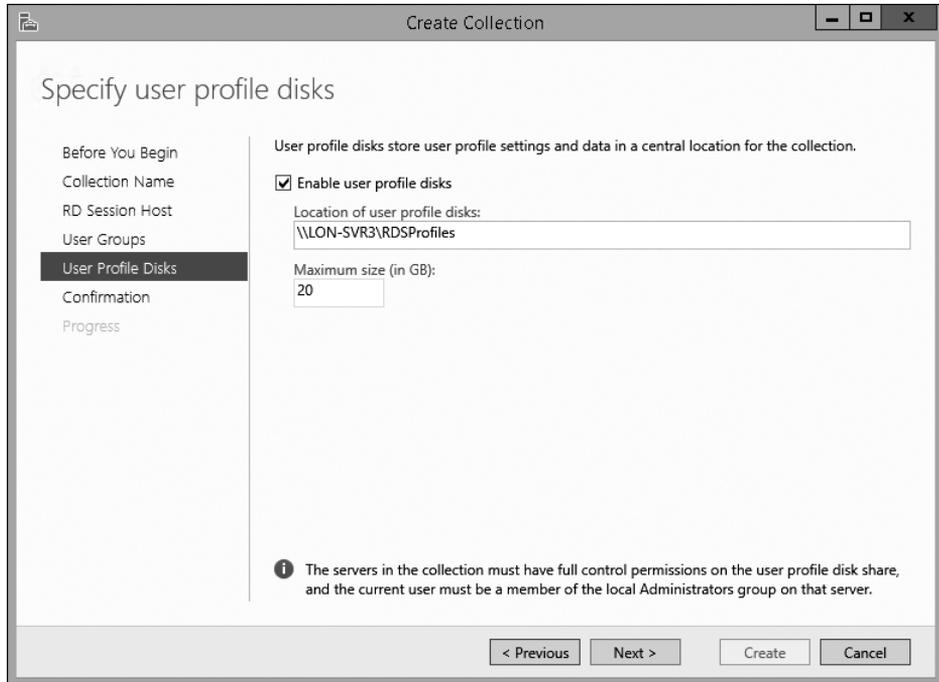


Figure 8-16 Create Collection Wizard, Specify User Profile Disks page

8. On the Confirm Selections page, click Create.
9. On the View Progress page, wait until all tasks are complete and then click Close.

Configuring session collections

The user interface for creating a session collection allows you to configure only a few of the configuration options for a session collection. After the session collection is created, you can edit the session collection and configure many more options.

To edit a session collection, perform the following steps:

1. In Server Manager, in the navigation pane, click Remote Desktop Services.
2. In Remote Desktop Services, in the navigation pane, click the collection you want to edit.
3. While viewing the collection, next to the Properties box, click Tasks and click Edit Properties.
4. In the *CollectionName* Properties window, edit the properties as required and click OK.

When you are editing the properties of a session collection, the editing window is divided into pages with groups of related options. The General page, shown in Figure 8-17, has the Name and Description that you entered during creation. The Show The Session Collection In RD Web Access check box was not available during creation. It is selected by default. Consider disabling this option during scheduled outages when you are performing maintenance on a session collection, for example, when you are upgrading an application on the RD Session Hosts in the collection.

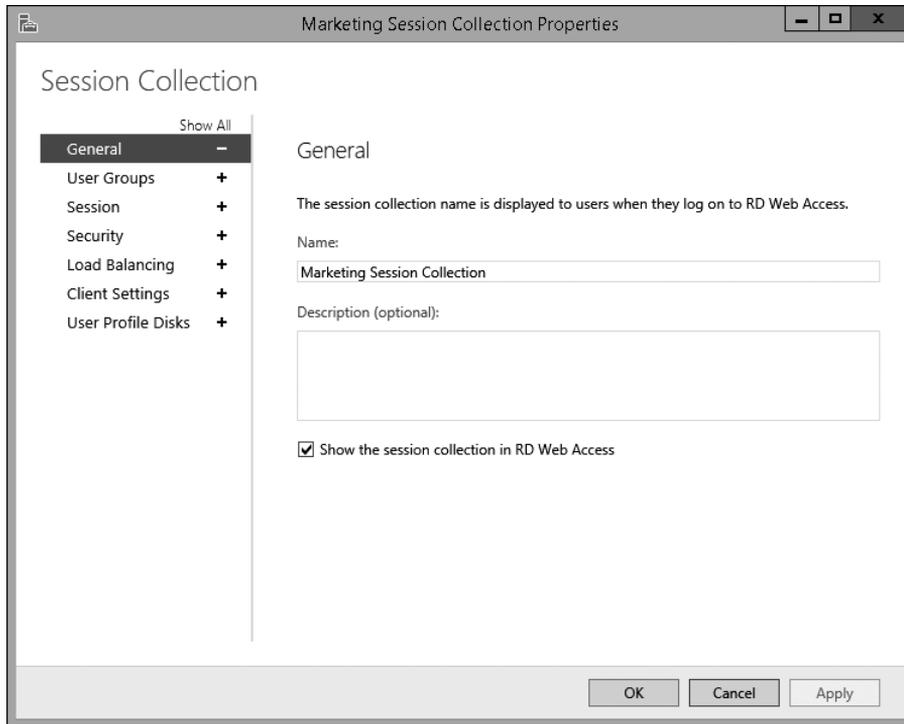


Figure 8-17 Properties of a session collection, General page

The User Groups page in the properties of a session collection allows you to configure which groups of users can connect to the session collection. This is the same as the user groups configured during creation.

The Session page, shown in Figure 8-18, has a number of settings that control session limits and temporary folders.

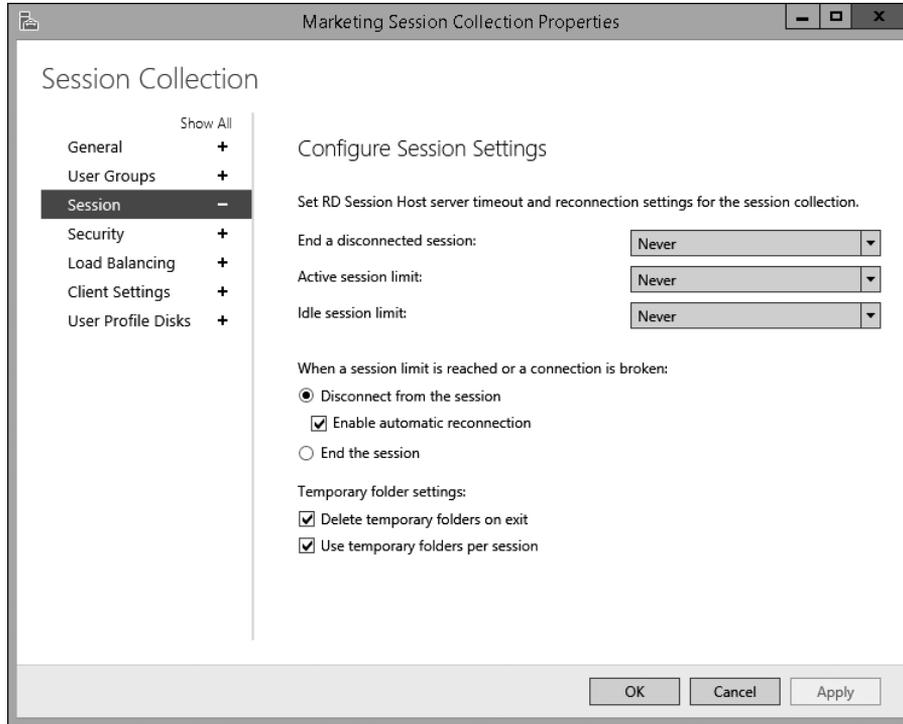


Figure 8-18 Properties of a session collection, Session page

Table 8-3 describes the session settings available on the Session page.

Table 8-3 Session settings for a session collection

Session setting	Description
End A Disconnected Session	Controls when a session is ended after a user disconnects. You can select to never end disconnected sessions or select a value ranging from one minute to five days. The default value is Never. A session is disconnected when there is a network connectivity issue or when a user closes the Remote Desktop Connection window without signing out. A disconnected session keeps all of the applications open and continues to use memory on the RD Session Host. Users can reconnect to the session and resume working where they left off, but if there are too many disconnected sessions, the RD Session Host may not have enough memory.

Active Session Limit	<p>Controls how long an active session can be before it is disconnected or ended. To allow active sessions with no limit, select Never. To limit active sessions, select a time value ranging from one minute to five days. The default value is Never.</p> <p>Users receive a warning two minutes before the active session limit is reached. This provides users with time to save their work. An active session is one in which the user is performing a task. An active session is identified by mouse movement or keyboard input. There is seldom a need to limit active sessions, but you could limit them if you are concerned that unauthorized users are accessing a session. This will force the user to reconnect and provide authentication credentials when the limit is reached.</p>
Idle Session Limit	<p>Controls how long an idle session can be idle before it is disconnected or ended. To allow idle sessions with no limit, select Never. To limit idle sessions, select a time value ranging from one minute to five days. The default value is Never.</p> <p>Users receive a warning two minutes before the idle session limit is reached. This provides users with an opportunity to move the mouse or press a key to make the session active and avoid the idle session limit.</p> <p>An idle session is one in which the user isn't performing a task. An idle session is identified by a lack of mouse movement or keyboard input. Most organizations configure an idle session limit. This has a similar effect on security as having the screen lock on a desktop computer. If a session is connected but unused, it may mean that the user has left his or her connection unattended.</p>
When A Session Limit Is Reached Or A Connection Is Broken	<p>Controls the action that is taken when the active session limit is reached, the idle session limit is reached, or a network problem disconnects a client. You can choose Disconnect From The Session or End The Session.</p> <p>In most cases, you will select Disconnect From The Session to prevent users from losing their work when they are disconnected. When you select this option, you also can select Enable Automatic Reconnection. This allows the RDC client to reconnect automatically after short network interruptions. If you do not select this option, the users must provide authentication credentials to reconnect to their disconnected session.</p> <p>By default, Disconnect From The Session and Enable Automatic Reconnection are selected.</p>
Delete Temporary Folders On Exit	<p>Configures temporary folders to be deleted when a session ends. This ensures that temporary files do not consume unnecessary disk space. This option is enabled by default.</p>

Use Temporary Folders Per Session	Configures each session for a user to have separate temporary folders on RD Session Host servers where a single user account is allowed to have multiple simultaneous sessions. This option is enabled by default to ensure that multiple sessions on an RD Session Host server do not conflict. However, it isn't relevant in most deployments because users typically are limited to a single session.
-----------------------------------	--

The Configure Security Settings page, shown in Figure 8-19, allows you to configure the Security Layer and the Encryption Level to use for the session. The Security Layer defines encryption methods that are used to encrypt communication between the RDC client and the RD Session Host. The available options for security layer are as follows:

- **RDP Security Layer** This is the weakest option for the security layer. It is available to support older RDP clients. This security layer does not support the use of Network Level Authentication.
- **SSL (TLS 1.0)** This is the strongest security layer. This security layer supports the use of network-level authentication. When this security layer is used, a certificate on the RD Session Host is used to establish the encryption channel. If the name on the certificate does not match the name used when connecting to the RD Session Host, then a warning is displayed on the client. This is supported by Windows XP SP3 and newer operating systems.
- **Negotiate** This is the default selection for security layer. SSL (TLS 1.0) is used if available on the server and client. If SSL (TLS 1.0) can't be used, then RDP Security Layer is used.

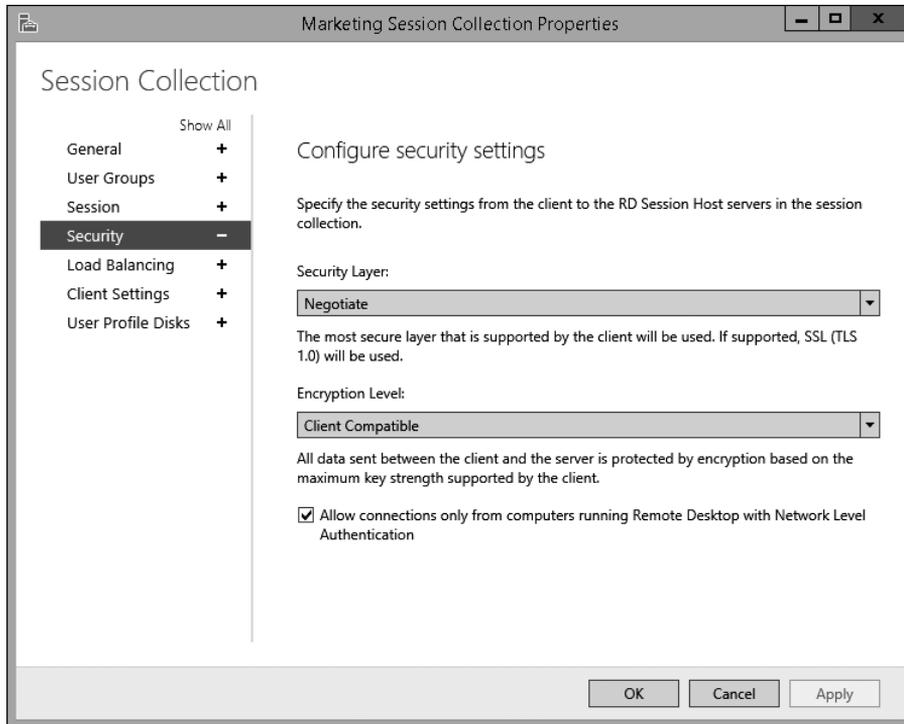


Figure 8-19 Properties of a session collection, Configure Security Settings page

Network Level Authentication is an authentication method that requires clients to enter authentication credentials before they are connected to the RD Session Host server. The credentials are passed by the RDC client to the RD Session Host server, and if the credentials are valid, the sign-in process is performed. When Network Level Authentication isn't used, clients can connect to the RD Session Host server and interact with the sign-in screen on the RD Session Host before they are authenticated. This is a security risk because it is possible for unauthenticated clients that have access to RD Session Host servers to see recently used user names and the operating system version.

You can force all clients to use Network Level Authentication by selecting the Allow Connections Only From Computers Running Remote Desktop With Network Level Authentication check box. This is enabled by default.

The Encryption Level setting allows you to configure the number of bits used for encryption. This setting applies for both security layers, and more bits provide stronger encryption. The options for Encryption Level are as follows:

- **Low** Uses 56-bit encryption for data sent from the client to the server. Data sent from the server to the client isn't encrypted. This option is provided to support older clients and typically isn't required.
 - **High** Uses 128-bit encryption for all data sent between the client and server. This option can be used by Windows XP and newer operating systems. This is the preferred option.
 - **FIPS Compliant** Uses encryption algorithms that are FIPS 140-1 or FIPS 140-2 compliant for all data sent between the client and server. Federal Information Processing Standards (FIPS) is a United States government standard for data encryption. This option typically isn't used unless requested specifically by an organization that needs to meet FIPS requirements.
 - **Client Compatible** Negotiates the highest level of encryption supported by the client and uses that. This is the default configuration, but it can be considered a security risk because it allows 56-bit encryption for clients that request it. Use this option only if you need to support clients that can't use 128-bit encryption.
- The **Configure Load Balancing Settings** page in the properties of a session collection is covered later in this chapter in the section titled "High availability for RD Session Host servers."

The **Configure Client Settings** page has settings for device redirection and monitors. By default, redirection is enabled for all available options. If desired, you can select to enable or disable redirection for the following:

- Audio And Video Playback
 - Audio Recording
 - Smart Cards
 - Plug And Play Devices
 - Drives
 - Clipboard
 - Printers
- You can find more information about the client settings in Chapter 9.

The User Profile Disks page, shown in Figure 8-20, allows you to configure all of the information entered during collection creation and to define what data is stored on the user profile disks. There are two options for user profile disks data settings:

- **Store All User Settings And Data On The User Profile Disk** Specifies that the complete user profile is stored on the user profile disk. You can add specific folders and files within the profile to exclude.
- **Store Only The Following Folders On The User Profile Disk** Specifies that only selected folders in the user profile are stored on the user profile disk. The folders available for selection are Contacts, Desktop, Documents, Downloads, Links, Music, Pictures, Roaming User Profile Data, and User Registry Data. You also can add specific files and folders within the profile to include.

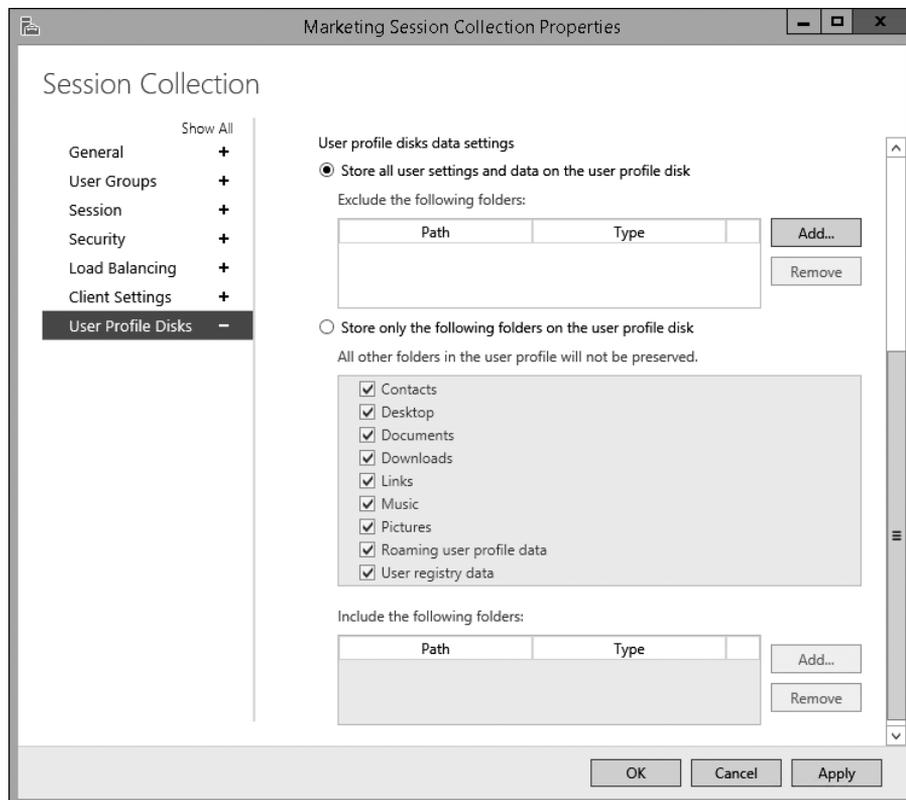


Figure 8-20 Properties of a session collection, User Profile Disks page

Configuring RD Licensing servers

The initial configuration of RDS doesn't configure licensing. However, a functional RDS deployment in production must have licensing properly configured to ensure that users can connect. To configure licensing for RDS, you need to complete the following tasks:

- Set the licensing mode
- Install an RD Licensing server
- Activate an RD Licensing server
- Install and activate CALs

To install an RD Licensing server, perform the following steps:

1. In Server Manager, in the navigation pane, click Remote Desktop Services.
2. On the Overview page, in the Deployment Overview area, click RD Licensing.
3. In the Add RD Licensing Servers Wizard, on the Select A Server page, double-click the server you want to configure as an RD Licensing server and click Next.
4. On the Confirmation page, click Add.
5. Wait until the installation is complete and click Close.

To set the licensing mode for an RDS deployment, perform the following steps:

1. In Server Manager, in the navigation pane, click Remote Desktop Services.
2. On the Overview page, in the Deployment Overview area, click Tasks and click Edit Deployment Properties.
3. In the Deployment Properties window, in the navigation pane, click RD Licensing.
4. On the RD Licensing page, select Per Device or Per User and click OK.

The Microsoft Clearinghouse is the service that is used to activate RD Licensing servers and RDS CALs. When you install an RD Licensing server, you need to activate it before it can begin servicing clients. To do this, you use Remote Desktop Licensing Manager (RD Licensing Manager), shown in Figure 8-21. RD Licensing Manager is installed on each RD Licensing server.

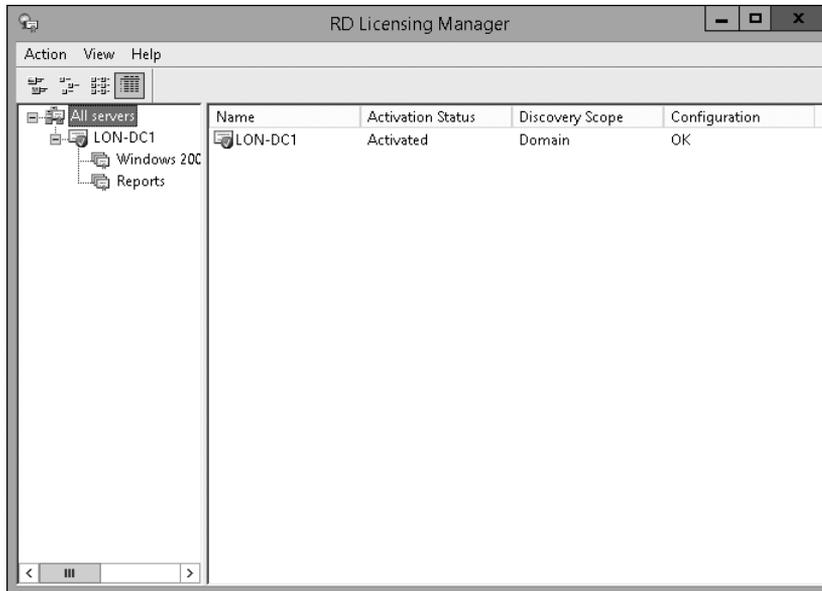


Figure 8-21 RD Licensing Manager

When you activate an RD Licensing server, you need to exchange information with the Microsoft Clearinghouse. You can choose from the following connection methods:

- **Automatic Connection (Recommended)** Transfers the necessary information between the RD Licensing server and the Microsoft Clearinghouse over the Internet. The RD Licensing server must have connectivity to the Internet.
- **Web Browser** Requires you to enter a Product ID at the website <https://activate.microsoft.com>. Then, you type the license server ID provided by the website into the Activate Server Wizard. Use this connection method if the RD Licensing server does not have access to the Internet.
- **Telephone** Requires you to phone the Microsoft Clearinghouse and provide the Product ID for your server. You are then give a license server ID, which you need to enter into the Activate Server Wizard. Use this connection method if you have no access to the Internet.

Installing RDS CALs is a similar process to activating an RD Licensing server. The same connectivity methods to the Microsoft Clearinghouse are supported. The installation process automatically uses the method that you used when activating the server. You can change the connectivity method in the Properties of the server if required.

To activate an RD Licensing server over the Internet, perform the following steps:

1. In Server Manager, click Tools, point to Terminal Services, and click Remote Desktop Licensing Manager.
2. Right-click the licensing server and click Activate Server.
3. In the Activate Server Wizard, on the Welcome To The Activate Server Wizard page, click Next.
4. On the Connection Method page, in the Connection Method box, select Automatic Connection (Recommended) and click Next.
5. If you have not already configured the company information for your server, you are prompted to do so. On the Company Information page, enter the required company information and click Next.
6. On the next Company Information page, if desired, enter the optional information and click Next.
7. On the Completing The Activate Server Wizard page, deselect the Start Install Licenses Wizard Now check box and click Finish.
8. Right-click the server and click Review Configuration.
9. In the Server Configuration window, shown in Figure 8-22, click Add To Group.

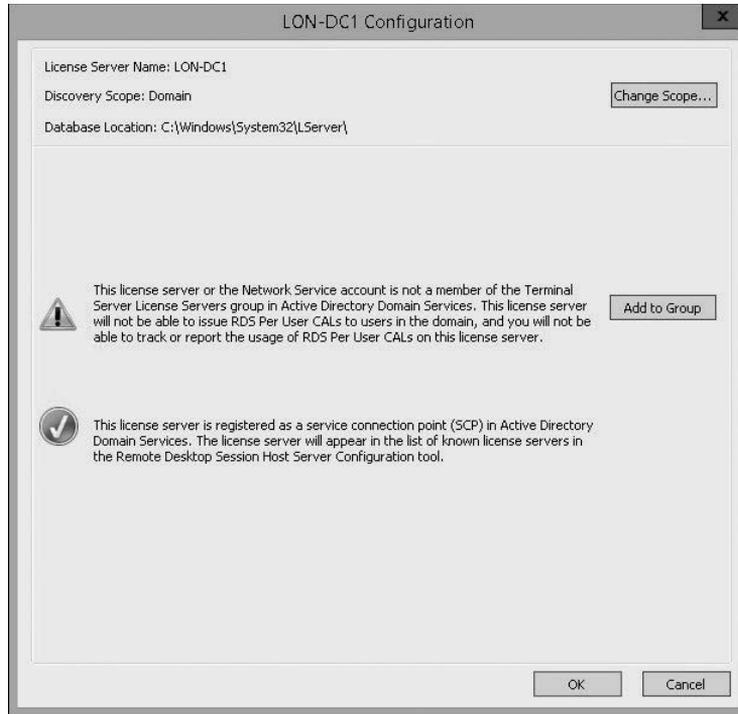


Figure 8-22 License server configuration

10. In the RD Licensing Manager window, click Continue to acknowledge the warning about requiring Domain Admin privileges.
11. In the RD Licensing Manager dialog box, click OK to acknowledge that the server has been added to the Terminal Server License Servers Group.
12. In the Server Configuration window, click OK.

To install RDS CALs over the Internet, perform the following steps:

1. In Server Manager, click Tools, point to Terminal Services, and click Remote Desktop Licensing Manager.
2. Right-click the license server and click Install Licenses.
3. In the Install Licenses Wizard, on the Welcome To The Install Licenses Wizard page, click Next.

4. On the License Program page, select the license program used to purchase your RDS CALs and click Next. Available license programs include Open License, Enterprise Agreement, Campus Agreement, and more.
5. Enter the information requested and click Next. The information requested varies depending on the licensing program used, but it will include either a license code or an agreement number.
6. On the Product Version And License Type page, enter the product version, license type, and number of RDS CALs based on your license and click Next.
7. Wait while the Microsoft Clearinghouse processes the request and the RDS CALs are installed and then click Finish.

Understanding high availability for RDS

A highly available service is one that is available almost all of the time. High availability often is expressed numerically as the percentage of time that a service is available. For example, a requirement for 99.9 percent availability allows 8.75 hours of downtime per year, or approximately 40 minutes of downtime every four weeks. With 99.999 percent uptime, the allowed service downtime is reduced to only 5 minutes per year.

To achieve high availability for a service such as RDS, you need to identify single points of failure in the infrastructure and work to eliminate them. For example, if you have only one RD Session Host server, that is a single point of failure. There will be a service outage if that server has hardware problems or is taken offline for maintenance.

To make infrastructure highly available, you need to make it redundant. For example, within a server, using mirroring (RAID 1) for disks ensures that that failure of a single hard disk does not cause the server to fail. This principle also can be applied at other levels of infrastructure such as networking, network services, and data center power. To make RDS highly available, you need multiple servers for each of the RDS role services, but not all RDS role services automatically become highly available just because you add more servers running the role service. Some RDS role services require you to implement load balancing.

Inside OUT

Redundant hardware vs. redundant servers

Within a server, two common hardware components for high availability are dual power supplies and hardware RAID cards. Combined, these options can add \$1,000 to the cost of a server.

When you have highly available infrastructure with multiple servers, you can balance the high availability of components within a server with the high availability provided by having multiple servers. Instead of providing hardware redundancy in servers, you can provide the redundancy with additional servers. This can be a cost-effective way to provide high availability.

In the previous example, you could save \$1,000 per server by not putting in dual power supplies and hardware RAID cards. In a large RDS deployment, if you are implementing eight RD Session Host servers, this would save \$8,000 in hardware costs, which you could use to purchase two or three additional servers to provide high availability.

Most administrators want their individual servers to be highly available, but it's really the services that need to be highly available. The key to making this strategy work is providing enough redundant servers and monitoring so that you can resolve issues quickly when a server fails

Understanding load balancing

Load balancing is a technology that you can use to achieve high availability and scalability primarily for stateless services. The term *stateless* refers to workloads that respond to each request independently from previous requests and without keeping client state. For example, when a client requests a webpage, a Web server gathers all of the necessary information from the request and then returns a generated webpage to the client. When the client requests another webpage, it might request the webpage from the same Web server or from any other identically configured Web server in a load-balancing cluster.

The servers that are part of a load balancing cluster are referred to as nodes. Each node is configured with the same software so that clients can connect with any node and obtain the intended service. For example, each node would have the same website configured.

Clients connect to a virtual IP address that is used to access all nodes in the cluster. To the clients, the virtual IP address behaves the same way that an IP address configured on a physical server would. Only one cluster node responds to each client request.

If a node in the cluster fails, the remaining nodes continue to service clients. This makes a load balancing cluster highly available. Adding nodes to the cluster increases the capacity of the cluster. Adding nodes is scaling out.

Windows Network Load Balancing

Windows Network Load Balancing (NLB) is a software solution for load balancing that is included in Windows Server operating systems. NLB creates a virtual IP address that all of the

nodes in the load balancing cluster share. When a request comes in to virtual IP, the request is received by all nodes, but only one node responds. The nodes determine the appropriate node to respond based on an algorithm that they all use.

The most common reason organizations consider using NLB is the cost. Because it is included in Windows Server operating systems, it is effectively free. However, there are a few drawbacks to using NLB:

- **It is not service aware** NLB is capable of identifying when a server is no longer responding but not when a service is no longer responding. This means that some types of failures result in clients being directed to a nonresponsive service on a partially functional node.
- **Scalability is limited** NLB supports up to 32 nodes in a cluster, but performance peaks at 8 nodes.
- **Network hardware configuration may be required** Some network switches need additional configuration to work with NLB. This is required because multiple devices are sharing the same virtual IP address but are connected to different switch ports.

Hardware-based load balancing

Most large organizations use specialized hardware load balancers instead of using NLB. Hardware load balancers are more scalable than NLB, but they also are significantly more expensive. The least expensive hardware load balancers are about \$2,000, and they can cost more than \$40,000.

The configuration of a hardware load balancer varies depending on the vendor, but all of them provide the same basic functionality. The virtual IP address for the load balancing cluster is assigned to the load balancer, and the load balancer receives requests from the clients. The load balancer then forwards each request to a single node.

The load balancer is responsible for identifying failed nodes. Node failure can be identified at the node level, as NLB does, or at the service level. If service-level failure is used, the load balancer monitors the service on each node and stops sending client requests if the service stops responding.

Inside OUT

DNS round robin

An alternate method for load balancing is DNS round robin. This method is implemented entirely by using DNS records and does not require any additional software or hardware. The main concern with DNS round robin is that it can be less reliable than NLB or hardware-based load balancing.

In DNS, a host (A) record is used to identify the IP address to which a name should resolve. For example, a host (A) record identifies the IP address to which `www.microsoft.com` resolves. If you create multiple host (A) records for a name, then it resolves to multiple IP addresses. When a client resolves the name, the DNS server provides IP addresses from all of the host (A) records for that name. This configuration is called DNS round robin.

When a client receives multiple IP addresses from a DNS server, the typical behavior is to contact the first IP address in the list. If the first IP address in the list doesn't respond, then the client contacts the second IP address in the list. This process continues until the client successfully connects to an IP address or the list is exhausted.

The main drawback to DNS round robin is that it's unpredictable. The client is in control of the process for managing server failure and the list of IP addresses. If the software developer for the client software does not manage server failure well, then users experience poor performance. For client software that is designed to use DNS round robin, it is a simple and effective load balancing mechanism.

High availability for RD Session Host servers

When an RDS deployment has a single RD Session Host server, that server becomes a single point of failure. When it fails, the failure will affect all users who are connected to it and who run RemoteApp programs on that server. You must consider the possibility of failure or the lack of RD Session Host server availability in your disaster recovery plan.

You can take several steps to improve RD Session Host availability. You can use reliable and redundant hardware from respected vendors to minimize the probability of hardware failure. You also should make sure that the network is reliable and that there are multiple network paths to an RD Session Host server. You should be aware, however, that failures are unavoidable, and no single server can always be available without downtime. For example, after you install Windows updates, computer restart is often required, which causes server downtime.

To make the RD Session Host server role highly available, you should have multiple RD Session Host servers in each collection. The RD Connection Broker role service automatically load balances connections to the RD Session Host servers. If the RD Connection Broker role service identifies that an RD Session Host server is unavailable, clients are not directed to the failed RD Session Host server. Clients are directed only to the remaining functional RD Session Host server.

As a best practice, all RD Session Host servers should have a similar hardware configuration. This ensures that all RD Session Host servers have similar performance and can handle a similar

number of clients. The default configuration of load balancing for a collection, shown in Figure 8-23, is best suited for this scenario.

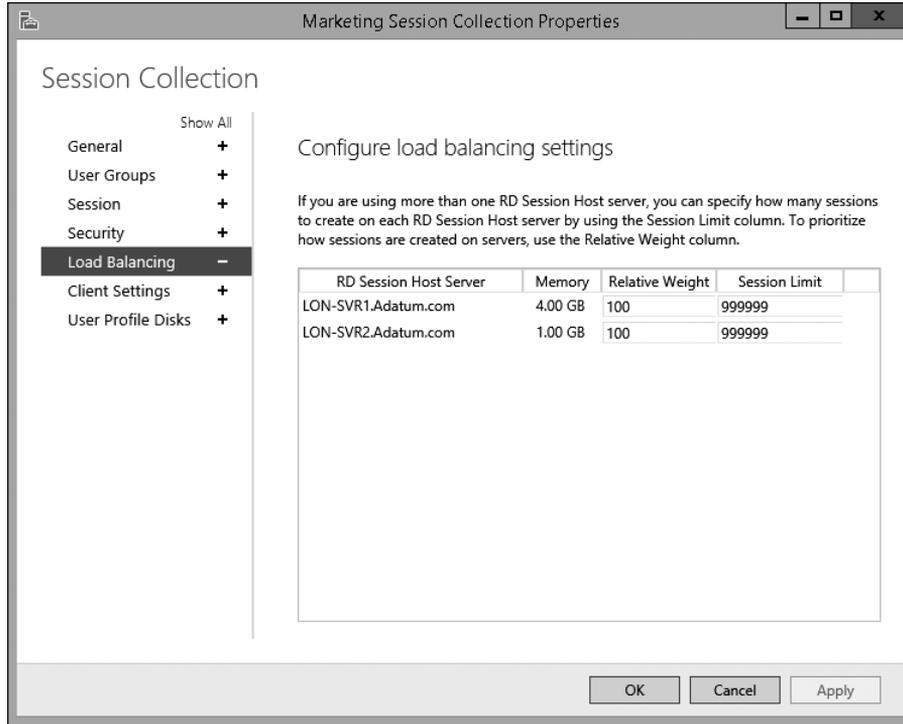


Figure 8-23 Session Collection Properties page, Configure Load Balancing Settings page

You can adjust the ratio of sessions allocated to an RD Session Host server by adjusting the Relative Weight value for that server. The default value for all servers is 100. When all servers have the same value, they all receive the same number of sessions. If you have one server with significantly better hardware and give that server a Relative Weight of 200, then it will receive twice as many sessions as a server with a Relative Weight of 100.

You also can set a Session Limit for each server. The default value for the Session Limit is 999,999, which effectively is unlimited. If you have determined that users experience performance issues when more than 80 clients are connected, then you can set a session limit of 80 to ensure that performance is satisfactory for all users.

To add a second RD Session Broker server to an RDS deployment, perform the following steps:

1. In Server Manager, in the navigation pane, click Remote Desktop Services.

2. On the Overview page, click Add RD Session Host Servers.
3. In the Add RD Session Host Server window, on the Select A Server page, in the Server Pool box, double-click the server you want to configure as an RD Session Host server and click Next.
4. On the Confirm Selections page, select the Restart Remote Computers As Needed check box and click Add.
5. Wait until the RD Session Host role service is installed on the server and click Close.

To add an RD Session Host server to a session collection, perform the following steps:

1. In Server Manager, in the navigation pane, click Remote Desktop Services.
2. In the navigation pane, click the session collection.
3. Scroll down to the Host Servers box, click Tasks, and click Add RD Session Host Servers.
4. In the Add Servers To Collection Wizard, on the Specify RD Session Host Servers page, double-click the RD Session Host server that you want to add to the session collection and click Next.
5. On the Confirm Selections page, click Add.
6. Wait until the task is complete and then click Close.

High availability for the RD Connection Broker role service

The RD Connection Broker role service is responsible for directing clients to an available RD Session Host server. If the RD Connection Broker role service is unavailable, then users are not able to access session-based virtual desktops. Having a single RD Connection Broker server creates a single point of failure.

To make the RD Connection Broker role service highly available, you need to have multiple RD Connection Broker servers. The RD Connection Broker role service uses a SQL Server database to track sessions that have been allocated to RD Session Host servers. For multiple RD Connection Brokers servers to work together, they need to share a single SQL Server database.

NOTE

High availability for the RD Connection Broker role service in Windows Server 2012 and newer is active/active. This means that multiple RD Connection Broker servers can respond to client requests at the same time. Older implementations of the RD Connection Broker and Terminal Services (TS) Connection Broker high availability

were active/passive. A second server was used only when the first server failed. This provided high availability but not scalability.

To prepare the RD Connection Broker role service for high availability, you need to do the following:

- Configure a server running Microsoft SQL Server 2008 R2 or newer. The RD Connection Broker servers must have permission to create a database on the server.
- Install the SQL Server Native Client on all RD Connection Broker servers. The RD Connection Broker servers use this to connect to the SQL database.
- Configure a static IP address on all RD Connection Broker servers. This is required to implement DNS round robin for load balancing.
- Configure a DNS round robin record for the RD Connection Broker servers. Select a name that is meaningful, such as rds.adatum.com.

Inside OUT

Configuring SQL permissions for RD Session Broker high availability

To assign the necessary permissions for RD Session Broker servers on the SQL server, perform the following steps:

1. Create a security group in Active Directory Domain Services (AD DS) and add the computer accounts for the RD Connection Broker servers.
2. Restart the RD Connection Broker server so that the new group membership takes effect.
3. On the SQL server, open SQL Server Management Studio.
4. In the Connect To Server window, verify that the correct instance of SQL server is listed and click Connect.
5. In SQL Server Management Studio, in Object Explorer, expand Security, and expand Logins.
6. Right-click Logins and click New Login.
7. In the Login – New window, on the General page, in the Login Name box, type Domain/GroupName.
8. Click the Server Roles page, select the Dbcreator check box, and click OK.
9. Close SQL Server Management Studio.

When you configure the RD Connection Broker role service for high availability, its database moves from a local WID to a computer that is running SQL Server. Even when an RDS deployment has multiple RD Connection Broker servers, SQL Server still can be a single point of failure. You should make sure that SQL Server is highly available by running it in a failover cluster.

When you configure high availability for the RD Connection Broker role service, you need to provide a Database Connection String that the RD Connection Broker servers use to connect to the SQL server. The Database Connection String has the following format:

```
DRIVER=SQL Server Native Client 11.0;SERVER=LON-SQL.Adatum.com;Trusted_
Connection=Yes;APP=Remote Desktop Services Connection Broker;Database=RDS-DB
```

There are several things to note about the Database Connection String:

- **A SQL native client version is specified** In this example, the SQL native client version is 11.0. This is used when your SQL server is SQL Server 2012. If your SQL server is SQL Server 2008 R2, then the SQL native client version is 10.0.
- **A server name is specified** In this example, the server name is LON-SQL.Adatum.com. In your deployment, you should specify the name of the SQL server that will be hosting the database for the RD Connection Broker servers.
- **A database name is specified** In this example, the database name is RDS-DB. This is the name of the database that will be created for the RD Connection Broker servers to use. You can select an alternate name, but it should be a meaningful name to make it easy to identify.

To configure the RD Connection Broker role service for high availability, perform the following steps:

1. In Server Manager, in the navigation pane, click Remote Desktop Services.
2. On the Overview page, in the Deployment Overview area, right-click RD Connection Broker and click Configure High Availability.
3. In the Configure RD Connection Broker For High Availability Wizard, on the Before You Begin page, click Next.
4. On the Configure RD Connection Broker For High Availability page, shown in Figure 8-24, in the Database Connection String box, type the appropriate Database Connection String for your environment.

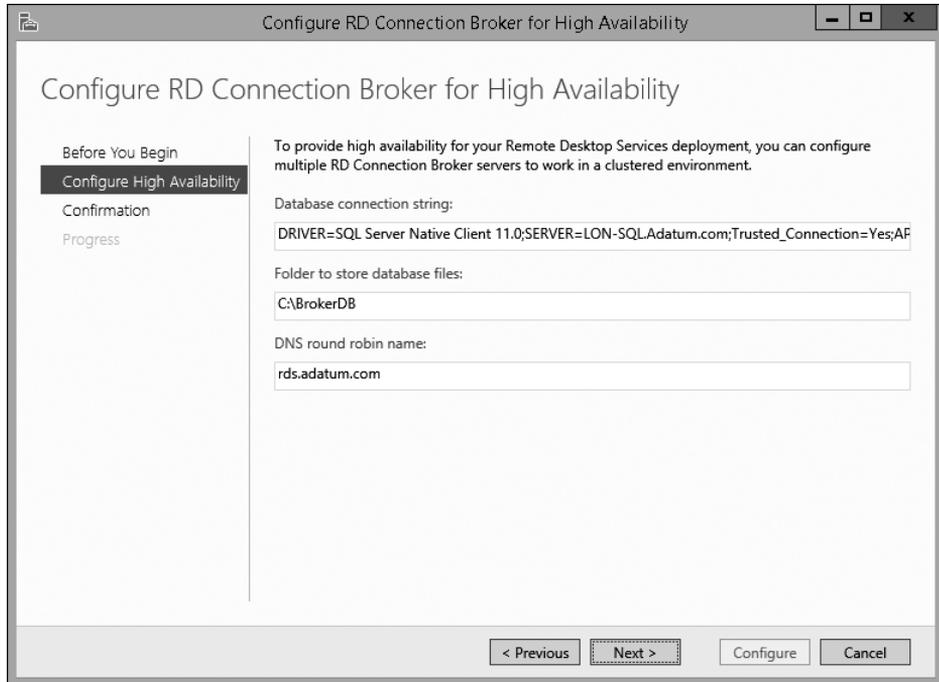


Figure 8-24 Configure RD Connection Broker For High Availability Wizard, Configure RD Connection Broker For High Availability page

5. In the Folder To Store Database Files box, type the path for the database on the SQL server. The database will be created in this location. This folder must already exist.
6. In the DNS Round Robin Name box, type the name of the DNS round robin record that you created for the RD Connection Broker servers and then click Next.
7. On the Confirmation page, click Configure.
8. On the Progress page, click Close.

After you have configured high availability for the RD Connection Broker role service, the RD Connection Broker icon in the Deployment Overview area is updated with the text (High Availability Mode). Now you can add another RD Connection Broker server by right-clicking the RD Connection Broker icon and clicking Add RD Connection Broker server. The new RD Connection Broker server will use the central SQL database that you have configured.

High availability for the RD Web Access role service

RD Web Access servers are a critical part of an RDS deployment. The RD Web Access servers are responsible for providing clients with an .rdp file that contains connectivity information to collections. If RD Web Access isn't available, then clients can't obtain the necessary configuration information to connect to session-based virtual desktops. You should make RD Web Access servers highly available.

Load balancing is used to make RD Web Access servers highly available. You can use NLB, hardware-based load balancing, or DNS round robin. If you are using NLB or hardware-based load balancing, you'll need to create a DNS record for the virtual IP address used by the load balancing cluster. For example, you could create a host (A) record for RDWeb.adatum.com that resolves to the virtual IP address. If you are using DNS round robin, then you need to create multiple host (A) records for RDWeb.adatum.com that resolve to the IP addresses of the RD Web Access servers.

To add an RD Web Access server, perform the following steps:

1. In Server Manager, in the navigation pane, click Remote Desktop Services.
2. On the Overview page, right-click RD Web Access and click Add RD Web Access Servers.
3. In the Add RD Web Access Servers Wizard, on the Select A Server page, double-click the server you want to configure as an RD Web Access server and click Next.
4. On the Confirmation page, click Add.
5. Wait until the installation is complete and click Close.
6. Configure your load-balancing solution with the IP address of the new RD Web Access server.

High availability for the RD Licensing role service

The effect of the RD Licensing role service for an RDS deployment varies, depending on the licensing mode that has been selected. When an RDS deployment is configured for Per User licensing, the RD Session Host servers contact an RD Licensing server each time a client connects. If an RD Licensing server isn't available, then users can't connect.

Per User licensing isn't enforced by RD Licensing servers. If an RD Session Host server can contact an RD Licensing server, that is sufficient to allow a connection. You are responsible for ensuring that you are in compliance with licensing requirements, but they are not enforced.

To make an RDS deployment with Per User licensing highly available, you need to install multiple RD Licensing servers. If the first RD Licensing server is unavailable, then the second is contacted.

Allocation of RDS User CALs among the RD Licensing servers does not matter because they are not enforced. To simplify license management, you can install and activate all RDS User CALs on a single RD Licensing server.

High availability for an RDS deployment configured for Per Device licensing also requires multiple RD Licensing servers, but configuration is more complex because RDS Device CAL usage is enforced. If an RDS Device CAL isn't available, then connectivity can be blocked. Because of this, you need to consider how CALs are allocated among the RD Licensing servers.

RDS client behavior for Per Device licensing varies, depending on the state of the client:

- **First connection** The first time a device connects, it is issued a temporary CAL that can be used only once. If an RD Licensing server is unavailable, the temporary CAL can't be issued, and new devices are unable to connect. The temporary CAL can be issued by any RD Licensing server even if no RDS Device CALs are available on that server.
- **Temporary license** The second time a device connects, it is issued a permanent RDS Device CAL. For a device to be issued a permanent RDS Device CAL, an RD Licensing server with unallocated Per Device CALs must be available. If an RD Licensing server with unallocated Per Device CALs isn't available, then the temporary CAL remains valid for 90 days.
- **Permanent CAL** Devices with a permanent CAL can connect to an RD Session Host when no RD Licensing server is available. Permanent RDS Device CALs are valid for 52 to 89 days and can't be renewed if no RD Licensing server is available.
- **Permanent CAL expired** If the permanent CAL has expired and an RD Licensing server isn't available, the connection is blocked. An RD Licensing server with unused Per Device CALs must be available to issue a new permanent CAL.

The simplest way to configure high availability for the RD Licensing role service when using Per Device licensing is to put all RDS Device CALs on a single RD Licensing server. The second RD Licensing server has no CALs installed and issues only temporary licenses. In this configuration, failure of the RD Licensing server with CALs has no effect on devices with a permanent or temporary license, which typically are the majority of devices. Devices connecting for the first time are issued a temporary license from the remaining RD Licensing server without CALs. The only clients unable to connect are devices with an expired license, which should be a small number of devices.

A slightly more complex way to configure high availability for the RD Licensing role service when using Per Device licensing is to split RDS Device CALs among RD Licensing servers. Most CALs are installed on the primary RD Licensing server, but some are installed on a secondary RD Licensing server. This configuration is better because if the primary RD Licensing server fails, then CALs can be issued by the secondary RD Licensing server, and no devices should be prevented from connecting.

Splitting CALs between two RD Licensing servers is slightly more expensive because you need to purchase additional CALs for the secondary RD Licensing server. In a large deployment of RDS, this likely is worth the additional cost to avoid outages. In a small deployment of RDS, it may not be worth the cost because very few users would be affected.

When you have multiple RD Licensing servers, it is critical that you configure the RDS deployment to use the RD Licensing server you have configured with the CALs as the primary RD Licensing server.

To configure the order of RD Licensing servers, perform the following steps:

1. In Server Manager, in the navigation pane, click Remote Desktop Services.
2. On the Overview page, in the Deployment Overview area, click Edit Deployment Properties.
3. In the Deployment Properties window, in the navigation area, click RD Licensing.
4. Select the server you want to be primary, click Move Up until it is at the top of the list, and click OK.

To add an RD Licensing server, perform the following steps:

1. In Server Manager, in the navigation pane, click Remote Desktop Services.
2. On the Overview page, in the Deployment Overview area, right-click RD Licensing and click Add RD Licensing Servers.
3. In the Add RD Licensing Servers Wizard, on the Select A Server page, double-click the server you want to configure as an RD Licensing server and click Next.
4. On the Confirmation page, click Add.
5. Wait until the installation is complete and click Close.

This page intentionally left blank



Index

Numbers and Symbols

- 32-bit operating systems, 491
- 64-bit operating systems, 491
- 1670 error code, 255
- 1671 error code, 255
- 1672 error code, 255
- 1673 error code, 255
- 1674 error code, 255
- 1677 error code, 255

A

- A (host) records, 441, 447
- Access RemoteApp And Desktops Wizard, 464
- Activate Server Wizard, 436
- Active Directory
 - activating operating systems, 492–493
 - adding users or groups, 278
 - configuring Credential Roaming, 69
 - editing access, 164, 285, 289
 - managing checkpoints, 151
 - Primary Computer user attribute, 76
 - VM-based desktop deployments, 522
- Active Directory Administrative Center, 48–49, 66–67
- Active Directory Domain Services (AD DS)
 - App-V client and, 228
 - App-V infrastructure requirements, 187–188
 - dynamic configuration, 310, 382
 - high availability and, 21
 - personal virtual desktops, 492
 - pooled virtual desktops, 492
 - Primary Computer feature and, 66
 - RD CAPs, 557
 - RD Gateway configuration scenarios, 544, 546
 - RD RAPs, 559
 - security groups and, 174–175, 308–309
- Active Directory Users and Computers, 48
- ActiveX controls, 380
- AD DS (Active Directory Domain Services)
 - App-V client and, 228
 - App-V infrastructure requirements, 187–188
 - dynamic configuration, 310, 382
 - high availability and, 21
 - personal virtual desktops, 492
 - pooled virtual desktops, 492
 - Primary Computer feature and, 66
 - RD CAPs, 557
 - RD Gateway configuration scenarios, 544, 546
 - RD RAPs, 559
 - security groups and, 174–175, 308–309, 444
- Add-AppvClientConnectionGroup cmdlet, 295, 297
- Add-AppvClientPackage cmdlet
 - adding packages, 163, 250–251, 269, 357
 - described, 295
 - elevated prompts, 297
- Add-AppvPublishingServer cmdlet, 266–267, 295, 297
- add-ons, 322–323, 334, 380
- Add RD Licensing Servers Wizard, 434, 449
- Add RD Web Access Servers Wizard, 447
- Add Roles and Features Wizard
 - Confirm Selections page, 521
 - Confirmation page, 422
 - Microsoft Virtual Machine-Based Desktop Deployment scenario, 519–523
 - Review Role Services page, 419, 520
 - Select Deployment Scenario page, 417–419, 520
 - Select Deployment Type page, 416–417, 419, 519
 - Select Installation Type page, 416, 419, 519
 - Specify RD Connection Broker Server page, 419–420, 520
 - Specify RD Session Host Servers page, 421–422
 - Specify RD Virtualization Host Server page, 521
 - Specify RD Web Access Server page, 420–421, 521
- Add Servers To Collection Wizard, 443
- Admin event log, 256, 568
- AdminGroup.xml file, 209
- Administrative Events custom view, 566–567
- .ADML files, 78
- .ADMX files, 78
- Advanced Group Policy Management (AGPM), 2
- agents (Operations Manager), 574
- AgentSetup32.msi file, 79
- AgentSetup64.msi file, 79–80
- AgentSetup.exe utility, 79–80
- AGPM (Advanced Group Policy Management), 2
- alerts (Operations Manager), 25, 572–575, 579–582
- always roaming user scenario, 43–44
- antivirus software, 500–501
- App-V (Application Virtualization)
 - application compatibility, 5, 7, 15
 - application life cycle, 161–177, 297–303, 305–306, 312
 - benefits of, 157–159
 - building images, 254–259
 - deploying Office 2013, 351–357

- deployment models, 184–187, 189–193
- disaster recovery, 199–200
- event logs, 568
- functional and physical placement, 194–195
- Group Policy support, 272–273, 384–385
- high availability, 197–199
- infrastructure requirements, 187–189, 200–201
- installation error codes, 255
- installing management databases, 201–204
- integrating with System Center Configuration Manager, 211
- licensing requirements, 24
- management packs, 582
- MDOP and, 2, 4–5
- modifying and upgrading published applications, 297–312
- optimizing, 500
- Remote Desktop Services client and, 210
- service disruption impact, 193–194
- shared content store, 220–224
- sizing and performance, 195–197
- standard versus virtualized applications, 159–160
- technologies supported, 177–184
- third-party production integration, 252–253
- UE-V Generator and, 88
- unsupported applications, 324
- VDI solution, 252–254
- virtualization engine and, 160–161

App-V client

- accessing applications, 215–216
- autoload mechanism, 273
- building images, 254–259
- configuring for stand-alone mode, 250–252
- data storage locations, 216–220
- deployment methods, 224–227
- described, 179–183, 213
- .INI files and, 380–381
- installation error codes, 255
- installation prerequisites, 228–230
- installing using Configuration Manager, 230–246
- isolating applications, 214
- managing file type associations, 265–266
- managing properties, 259–266
- managing server connections, 266–275
- managing using Group Policy, 180–181, 191, 210, 221–222, 225–228, 270–273
- managing virtual applications, 264
- registry settings, 273–275
- for Remote Desktop, 227, 247–250
- reporting process overview, 312–318
- requesting and executing applications, 213–214
- shared content store, 220–224
- shortcut support, 217–220
- third-party production integration, 252–253
- troubleshooting installation, 255–259
- VDI solution, 252–254
- Windows PowerShell support, 267–269, 294–296

App-V client setup window

- Customer Experience Improvement Program page, 248–249
- described, 247
- Software License Terms page, 247–248

App-V client window

- App Connection Groups tab, 262–263
- described, 260
- displaying Windows PowerShell cmdlets, 263–264
- Overview tab, 260–261, 264
- Virtual Apps tab, 261–262, 264

App-V Management Console

- adding packages, 303–304
- Administrators tab, 278, 281–282
- described, 277–280
- granting access, 308–309
- importing new versions of packages, 308–309
- managing administrators, 281–282
- managing application packages, 284–287
- managing connection groups, 287–288
- Overview tab, 277
- Packages tab, 175, 278, 285–288, 304
- registering and unregistering servers, 282–284
- Servers tab, 278, 282–284

App-V management servers

- accessing packages, 284
- adding packages, 284
- configuring, 204–207
- described, 178–179, 282
- disaster recovery and, 199–200
- event logs, 568
- IIS support, 205–206, 278
- infrastructure requirements, 188
- installing management databases, 201–204
- managing by using Windows PowerShell, 289–297
- publishing packages, 285
- refreshed list of published applications and, 175
- round-trip response time, 195–196
- service disruption impact, 194
- Windows PowerShell support, 294, 296

App-V publishing servers

- App-V client and, 215, 295–296
- capacity planning, 196–197
- deploying and configuring, 208–210
- described, 179–183
- disaster recovery and, 199–200
- displaying associated, 296
- IIS support, 176, 179, 278
- infrastructure requirements, 188–189
- monitoring performance, 566
- optimizing, 500–501
- refreshed list of published applications and, 175–176
- registering and unregistering, 282–284
- round-trip response time, 195–196
- service disruption impact, 193–194

App-V reporting servers

- App-V client configuration, 315–317
- data collection, 314
- described, 183–184
- disaster recovery and, 199–200
- generating reports, 317–318
- IIS support, 278
- infrastructure requirements, 188
- reporting process overview, 312–313
- service disruption impact, 194

App-V Sequencer

- Add New Application option, 303
- application sequencing, 161–162
- application update, 165–174, 297–303
- best practices, 330–332
- Completion page, 302, 375–376
- Configure Software page, 320, 341–342, 373–374
- configuring, 326–330
- connection groups and, 380–381
- Create A New Virtual Application Package option, 337, 367
- Create Package Accelerator page, 361–363, 365–366
- Create Package page, 173–174, 302, 345, 372–373, 379
- creating package accelerators, 358–377
- creating packages, 284, 310
- Customize page, 344
- deployment models and, 192
- described, 161, 179–183, 319–321
- editing packages, 302–303, 306–307
- ESD model and, 191
- Guidance page, 369–370
- Installation Files page, 360–361
- Installation page, 170–171, 300–301, 341, 379
- Installation Report page, 171–172, 301, 343, 379
- installing, 336
- Local Installation page, 370–371
- Modify An Existing Virtual Application Package option, 166
- Package Name page, 338–339, 371–372
- Packaging Method page, 337, 367–368
- Prepare Computer page, 168–169, 299, 337, 378
- Prepare For Streaming page, 172–173, 301–302
- Run Each Program page, 374–375
- Select Guidance page, 364–365
- Select Installer page, 169–170, 299–300, 338, 378
- Select Package Accelerator page, 368–369
- Select Package page, 167–168, 299, 359–360
- Select Task page, 166–167, 298
- sequencing process overview, 321–322
- Streaming page, 379
- Target OS page, 344–345
- Update Application in Existing Package option, 166–167, 298–299
- updating packages, 378–381
- Verify Applications page, 363–364
- virtual machines as, 328–329
- Windows PowerShell support, 294–295, 351

App-V Sequencing Guide, 333

- App-V_SSRS_Reports.exe utility, 317
- AppData folder, 34–35, 53
- Append Data NTFS permission, 47, 77

application compatibility

- App-V, 5, 7, 15
- comparing virtual desktop options, 484–485
- improving, 15–16
- operating system differences, 491
- RD Session Host service, 455, 485
- RemoteApp, 6, 15, 453
- session-based desktops and, 9
- user state virtualization, 40

Application event log, 568**application files**

- launching virtual applications, 321
- in packages, 284
- standard application installation, 159
- streaming delivery, 216, 500

application life cycle (App-V)

- application deployment, 162
- application publishing, 162, 297–312
- application removal, 305–306
- application sequencing, 161–162
- application termination, 174
- application update, 165–175, 297–303
- Configuration Manager model, 165
- described, 161
- full infrastructure model, 164
- removing applications from users, 174–177, 305–306
- script execution support, 312, 383–384
- stand-alone deployment model, 163–164

application request process, 213–214**application sequencing. See also App-V Sequencer**

- application types supporting, 333–334
- applications incompatible with, 332–333
- best practices, 330–332
- for connection groups, 380–381
- customizing packages, 344–345
- deploying Office 2013 by using App-V, 351–357
- described, 161–162
- dynamic configuration and targeted scripting, 381–385
- items to document in recipes, 322–324
- package accelerators, 357–377
- package editor, 345–350
- planning, 326–333
- portions of sequenced applications, 324–325
- preparing for, 335–336
- process overview, 321–322
- sequencing tasks, 337–343
- updating packages, 378–385
- Windows PowerShell support, 351

application virtualization

- App-V. *See* App-V
- App-V client and, 214

- described, 1, 4–6
- implementation considerations, 17–25
- infrastructure, 161–187
- management packs, 581–583
- monitoring, 565–572
- RemoteApp. *See* RemoteApp
- standard applications versus, 159–160
- usage scenarios, 12–17

ApplySettingsTemplateCatalog.exe utility, 93

AppSense applications, 253

.APPV files

- contents of, 186–187
- described, 178, 186, 325
- extracting virtual registry, 275
- hosting on web servers, 217
- saving packages as, 320, 355–356
- sequenced applications, 321
- stand-alone deployment model, 185, 192
- version information, 298, 378

AppVClient module (PowerShell), 268–269, 294–296

appv_client_MSI_x64.msi file, 229

appv_client_MSI_x86.msi file, 229

appv_client_setup.exe command, 225, 228–229, 247

AppVManagement database, 281

AppVPackages directory, 355

AppVReporting database, 281

AppVSequencer module (PowerShell), 294, 296, 351

AppVServer module (PowerShell), 294–295

appv_server_setup.exe command, 202, 205–206, 208

.APPVT files, 186, 325

AppxManifest.xml file, 187, 310, 325, 381

asset folders, 219

assigned computer user scenario, 41–42

authentication

- AD DS and, 187
- App-V and, 188–189
- certificates and, 68
- Microsoft accounts and, 71, 76
- multi-factor, 557, 561–564
- Network Level Authentication, 431
- RD CAPs, 558
- RD Gateway, 549
- scenario with unauthorized users, 544
- Server Authentication, 552
- wireless networks and, 111

authorization

- AD DS, 187
- RD Gateway, 391, 549–550, 557

Autoload registry key, 273–274

Autoruns program, 590

Azure Multi-Factor Authentication, 557, 561–564

B

Background Intelligent Transfer Service (BITS), 498

bandwidth

- monitoring performance, 589

- RDC detection, 400–401

- RDS reduction, 395

Berson, Frank, 564

Best Practices Analyzer (BPA), 415

BitLocker Drive Encryption, 495

BITS (Background Intelligent Transfer Service), 498

Block-Level Backup Engine Services, 498

Bluetooth Support Service, 498

boot options (virtual machines), 124–125

BPA (Best Practices Analyzer), 415

branch office environment, 453

BranchCache feature, 63

Bring Your Own Device (BYOD), 17, 541

BYOD (Bring Your Own Device), 17, 541

C

.CAB files, 186, 325, 358, 376–377

CALs (client access licenses), 22–24, 401–403, 437–438, 448–449

capacity planning

- described, 195–197, 511–512

- high availability and, 21

- for memory, 513–514

- for networking, 514–515

- for processing, 515–516

- scenario for, 517–518

- for storage, 512–513

CEIP (Customer Experience Improvement Program), 81, 272

CEIP node (App-V policies), 272

central RD CAP store, 560–561

certificates

- common mistakes creating, 476

- configuring, 472–479

- Credential Roaming feature, 68–69

- load-balanced environment, 280

- RDS, 473–478

- SAN, 473

- self-signed, 476

- server-based, 280, 473

- SSL, 205, 278–280, 377, 475, 550–553

- UCC, 473

- wildcard, 279, 473

Certificates snap-in (MMC), 475

change user command

- /execute options, 250, 456

- /install option, 247, 456

- /query option, 456

Checkpoint Name window, 328

Checkpoint-VM cmdlet, 152, 329

checkpoints (snapshots)

- applying, 154

- creating, 152–153, 328–329, 336

- deleting, 154–155

- described, 151–152, 328

- exporting, 154
- pooled virtual desktops, 533
- renaming, 154
- reverting to previous state, 154
- SAN support, 505
- troubleshooting, 329
- usage considerations, 153–156
- child partitions, 98–99**
- Citrix applications, 253**
- Clear-FileStorageTier cmdlet, 509**
- client access licenses (CALs), 22–24, 401–403, 437–438, 448–449**
- Client Coexistence node (App-V policies), 272**
- Client Hyper-V**
 - architectural overview, 97–100
 - controlling virtual machines, 125–130
 - creating virtual machines, 112–121
 - described, 6–7, 95–97
 - dynamic memory, 137–138
 - Generation 2 virtual machines, 124–125
 - Hyper-V comparison, 99–100
 - installing, 100–104
 - integration services, 138–141
 - management tools, 104–107
 - managing checkpoints, 151–156
 - managing virtual hard disks, 141–151
 - managing virtual machine files, 130–135
 - managing virtual switches, 108–112
 - processing overview, 135–136
 - virtual machine settings, 121–124
- Client registry key, 274**
- cluster shared volume (CSV), 488, 507–509, 525**
- cmdlets. *See also* specific cmdlets**
 - described, 105–106
 - elevated prompts, 297
 - viewing, 263–264
- COM objects, 214, 219**
- compact VHD files, 149**
- Company Settings Center dialog box, 86**
- compatibility, application**
 - App-V, 5, 7, 15
 - comparing virtual desktop options, 484–485
 - improving, 15–16
 - operating system differences, 491
 - RD Session Host service, 455, 485
 - RemoteApp, 6, 15, 453
 - session-based desktops and, 9
 - user state virtualization, 40
- Computer And Device Management Wizard, 577–579**
- computer certificates, 553**
- computer images, deploying App-V client, 226–227**
- configuration files**
 - App-V client, 217
 - control settings, 311, 382–383
 - customizing, 295
 - dynamic, 310–312, 349, 382
 - file name, 165
 - importing, 308
 - MDT supported, 495
 - post-installation steps, 323
 - usage considerations, 310
 - Windows PowerShell support, 295
- Configuration Manager (System Center)**
 - App-V and, 165, 185, 189–191, 198–199, 211
 - App-V client and, 226, 228, 230–246, 312
 - installing App-V client using, 230–246
 - personal virtual desktops, 494
 - pooled virtual desktops, 494
 - UE-V agent and, 79, 81
 - usage considerations, 228
- Configuration.xml file, 352–353**
- configure function (ODT), 354**
- Configure RD Connection Broker For High Availability Wizard, 445–446**
- connection groups**
 - adding packages to, 291–292
 - application sequencing, 380–381
 - creating, 290–294
 - described, 178, 214, 287
 - ESD model and, 191
 - granting access, 292–293
 - managing, 287–288, 290–296
 - ordering packages in, 292
 - publishing, 293–294
 - troubleshooting issues, 288
 - updating, 304–305
 - version considerations, 304
- Contacts folder, 54**
- Content_Types.xml file, 310, 381**
- Control Panel, 456, 465**
- Convert-VHD cmdlet, 150**
- copy-on-write (CoW) roaming, 217**
- coreinfo tool, 102**
- cost of ownership, reducing, 254**
- CoW (copy-on-write) roaming, 217**
- Create A Settings Location Template Wizard**
 - Edit Templates page, 91–92
 - Finish page, 92
 - Review Locations page, 90–91
 - Specify Applications page, 89–90
- Create Application Wizard**
 - Application Catalog page, 232–233
 - Completion page, 245–246
 - Deployment Types page, 233–234, 243–244
 - described, 189–190
 - General Information page, 232
 - General page, 190, 231
 - Summary page, 244–245
- Create Collection Wizard**
 - automatic VM creation, 525

- Confirm Selections page, 426
- described, 524
- Name The Collection page, 423, 526
- Specify RD Session Host Servers page, 423–424
- Specify The Collection Type page, 526–527
- Specify The Unattended Installation Settings page, 528–529
- Specify The Virtual Desktop Settings page, 528
- Specify The Virtual Desktop Template page, 527
- Specify User Groups page, 424–425
- Specify User Profile Disks page, 425–426, 532
- Specify Users And User Groups page, 529–530
- Specify Virtual Desktop Allocations page, 530–531
- Specify Virtual Storage page, 531–532
- View Progress page, 426
- virtual desktop templates, 490

Create Deployment Type Wizard

- Completion page, 242–243
- Content page, 236–237
- Create Requirement window, 239–240
- Dependencies page, 240–241
- described, 234
- Detection Method page, 237–238
- Detection Rule window, 237–238
- General Information page, 235–236
- General page, 234–235
- Requirements page, 239–240
- Summary page, 241–242
- User Experience page, 238–239

Create Folders NTFS permission, 47, 77

Create New Application Wizard, 165

Create New Package Wizard

- Completion page, 375–376
- Configure Software page, 320, 341–342, 373–374
- Create Package page, 345, 372–373
- Customize page, 344
- Guidance page, 369–370
- Installation page, 341
- Installation Report page, 343
- Local Installation page, 370–371
- Package Name page, 338–339, 371–372
- Packaging Method page, 367–368
- Prepare Computer page, 337
- Run Each Program page, 374–375
- Select Installer page, 338
- Select Package Accelerator page, 368–369
- Target OS page, 344–345

Create Package Accelerator Wizard

- Create Package Accelerator page, 361–363, 365–366
- Installation Files page, 360–361
- Select Guidance page, 364–365
- Select Package page, 359–360
- Verify Applications page, 363–364

Create Requirement window, 239–240

Creator Owner account, 47, 77–78

Credential Roaming, 4, 68–69

- CSV cache, 508–509, 511, 513

- CSV (cluster shared volume), 488, 507–509, 525

- Customer Experience Improvement Program (CEIP), 81, 272

- customsettings.ini file, 495

D

- DaRTRT (Diagnostics and Recovery Toolset), 2

- Data Deduplication feature, 508, 510–511

- data redundancy, 504–505

- data warehouse database, 574

- Database Connection String, 445

- Database.sql file, 203–204

- debug logs, 257–259

- demonstrations, 7

- deployment methods (App-V client)

- deployment by imaging, 226

- GPO deployment, 225–226

- manual deployment, 225

- stand-alone mode, 221, 250–252

- deployment models (App-V)

- Configuration Manager model, 165, 185, 189–191, 198–199, 312

- described, 162

- electronic software distribution model, 191–192

- full infrastructure model, 164, 184–185, 187, 189, 197–198, 313

- stand-alone, 163–164, 185, 192–193, 197, 312

- DeploymentConfig.xml file

- App-V client, 214

- App-V packages, 186

- application sequencing, 325, 355

- Configuration Manager model, 165

- dynamic configuration, 310, 382–383

- Desktop folder, 35, 53**

- desktop virtualization**

- application compatibility, 15–16

- Client Hyper-V. *See* Client Hyper-V

- described, 2, 6

- desktop as a service, 16

- implementation considerations, 17–25

- legal and security requirements, 13

- management packs, 581–583

- management task support, 14

- mobile user experience, 16–17

- monitoring, 565–572, 585–590

- user profile disks, 70

- virtual desktops. *See* virtual desktops

- Detection Rule window, 237–238**

- device-based CALs, 22–24**

- device drivers, 323, 333**

- device emulation, 99, 112–113**

- device redirection, 467–469, 558, 590**

- Diagnostic Policy Service, 498**

- Diagnostics and Recovery Toolset (DaRTRT), 2**

- differencing disks, 145–147, 150, 152–153, 155**

- Disable-AppvClientConnectionGroup cmdlet, 295
- disaster recovery, 100, 199–200
- Disconnect option (connections), 471
- Disconnect-RDUser cmdlet, 471
- diskpart utility, 496
- Dism.exe utility, 103
- DLLs (dynamic-link libraries), 15, 158–159
- DNS (Domain Name System), 21, 441, 493, 550
- DNS round robin, 440–441, 444, 446, 486
- document invocation, 462
- Documents folder, 35, 39, 53, 58
- Domain Admin group, 282
- Domain Computers account, 78
- Domain Name System (DNS), 21, 441, 493, 550
- Domain Users group, 308
- download function (ODT), 354
- Downloads folder, 35, 54
- Dsmod.exe utility, 68
- dual booting, 96–97
- dynamic configuration files
 - disabling virtual services, 349
 - enabling scripting, 310–312
 - targeted scripting, 381–385
- Dynamic CPU Fair Share feature, 589
- Dynamic Disk Fair Share feature, 589
- dynamic DNS, 493
- dynamic-link libraries (DLLs), 15, 158–159
- dynamic memory, 137–138, 156, 513–514
- Dynamic Network Fair Share feature, 589
- dynamical expanding VHDs, 143–145

E

- Edit Virtual Hard Disk Wizard, 148–150
- electronic software distribution (ESD) model, 191–192
- emulated devices, 99, 112–113
- Enable-AppvClientConnectionGroup cmdlet, 295
- Enable-DedupVolume cmdlet, 510
- Enable-UevTemplate cmdlet, 75
- Enable-WindowsOptionalFeature cmdlet, 103
- EnableCpuQuota registry key, 589
- EnableFairShare registry key, 589
- EnablePackageScript registry key, 275
- Enhanced Session Mode, 129
- enterprise user state virtualization, 38
- environment variables, 50, 219
- error codes (App-V installations), 255
- ESD (electronic software distribution) model, 191–192
- event logs, 255–256, 282, 499, 566–568
- Event Viewer, 256–259, 555, 565–567
- Everyone account, 78
- Exchange Server, 42–44, 579
- Expand-AppvSequencerPackage cmdlet, 296
- Export Virtual Machine dialog box, 132
- exporting
 - checkpoints, 154

- package configuration, 287, 307–308
 - virtual machines, 132, 134
- Extended Page Tables, 398
- extension points, 219, 310, 381
- external virtual switches, 109–110

F

- failover clusters, 487–489, 506–507
- Fair Share feature, 397, 588–589
- Fault Streaming, 320–321
- fault tolerance, 12, 387
- Favorites folder, 35, 54
- feature block maps, 320–321, 324
- Fibre Channel, 100, 488, 506
- File Explorer
 - accessing Sysprep, 502
 - browsing ODT logging path, 353
 - browsing source path, 353, 355
 - configuring offline files and folders, 61, 64–66
 - creating Central Store, 271
- file server repository, 193–195
- file systems, 214
- file type associations (FTAs)
 - editing, 265–266, 306–307
 - extension points as, 219
 - package editor support, 350
 - post-installation steps, 323
 - RemoteApp programs, 454
 - removing, 306–307
 - viewing existing, 286–287
- FilesystemMetadata.xml file, 187, 325
- fixed-size VHDs, 143–145
- Folder Redirection feature
 - configuring, 53–58
 - described, 3, 53
 - preserving user state, 412–413
 - user profile disks and, 524
 - user state virtualization and, 36, 39, 41–43
- folders
 - application sequencing, 324–325
 - asset, 219
 - configuring offline, 61, 64–66
 - Folder Redirection. *See* Folder Redirection feature
 - known, 53–54
 - naming conventions, 309
 - roaming user profiles, 48
 - storing on user profile disks, 523
 - user-specific data, 34–35
- FQDN (fully qualified domain name), 473–475, 479, 552
- FreeBSD operating system, 125, 139
- FTAs (file type associations)
 - editing, 265–266, 306–307
 - extension points as, 219
 - package editor support, 350
 - post-installation steps, 323

RemoteApp programs, 454
 removing, 306–307
 viewing existing, 286–287

Full Control NTFS permission, 47–48, 77–78, 251

full infrastructure model

App-V infrastructure requirements, 187
 deployment possibilities, 189
 described, 164, 184–185
 high availability and, 197–198
 reporting process overview, 313

Fully Loaded state, 193

fully qualified domain name (FQDN), 473–475, 479, 552

G

gateway servers (Operations Manager), 575

generation 1 virtual machines, 121–124, 398

generation 2 virtual machines, 124–125

generic volume license key (GVLK), 492

Get-AppvClientApplication cmdlet, 295

Get-AppvClientConfiguration cmdlet, 295

Get-AppvClientConnectionGroup cmdlet, 215, 295

Get-AppvClientMode cmdlet, 295

Get-AppvClientPackage cmdlet, 215–216, 251, 295

Get-AppvClientPackage cmdlet, 163

Get-AppvPublishingServer cmdlet, 215, 296

Get-AppvServerConnectionGroup cmdlet, 292, 294

Get-AppvServerPackage cmdlet, 291, 294

Get-AppvServerPackageDeploymentConfiguration cmdlet, 294

Get-AppvServerPackageUserConfiguration cmdlet, 295

Get-AppvVirtualProcess cmdlet, 295, 306

Get-Cluster cmdlet, 509

Get-Command cmdlet, 105, 268, 461

Get-Help cmdlet, 106

Get-Module cmdlet, 267

Get-RDUserSession cmdlet, 471

Get-StorageTier cmdlet, 509

Get-UevAppxPackage cmdlet, 87

Get-UevConfiguration cmdlet, 83

Get-UevTemplate cmdlet, 83

Get-VMHost cmdlet, 106

Get-VMSnapshot cmdlet, 330

Get-VMSwitch cmdlet, 112

Get-WindowsOptionalFeature cmdlet, 103

GlobalEnabled registry key, 275

GlobalLogonRefresh registry key, 275

GlobalPeriodicRefreshInterval registry key, 275

GPOs (Group Policy Objects)

App-V and, 164, 191–192
 App-V client and, 221–222, 225–226, 228, 282
 Configure Background Sync setting, 60
 configuring Folder Redirection, 55–56, 58
 creating, 164
 Enable Package Script setting, 384
 managing user profiles, 50–51

modifying settings, 270–271
 publishing server settings, 181
 reporting server settings, 184
 shared content store settings, 221–222

Grant-AppvServerConnectionGroup cmdlet, 293, 295

Grant-AppvServerPackage cmdlet, 289–290, 295

Griffin, Kristin L., 564

Group Policy

App-V administrator access, 282
 App-V client settings, 180–181, 191, 210, 221–222, 225–228, 270–273

App-V settings, 165, 181, 183–185

Central Store for, 270–271

configuring scripting, 384

Credential Roaming settings, 68–69

Ease of Access settings, 75

ensuring consistent configuration, 315

ESD model and, 191

Fair Share settings, 589

Folder Redirection settings, 41, 55

network latency configuration, 60

Offline Files settings, 60

Primary Computer setting, 66–67

settings storage location, 73

shared content store setting, 221–222

UE-V agent deployment, 78–79

UE-V agent management, 81–84

UE-V settings, 77–78, 86

user profile management, 50–52, 412–413

user state virtualization implementation, 41

Windows 8 app settings, 87

Group Policy Management Editor, 272–273

Group Policy Objects (GPOs)

App-V and, 164, 191–192

App-V client and, 221–222, 225–226, 228, 282

Configure Background Sync setting, 60

configuring Folder Redirection, 55–56, 58

creating, 164

managing user profiles, 50–51

modifying settings, 270–271

publishing server settings, 181

reporting server settings, 184

shared content store settings, 221–222

GVLK (generic volume license key), 492

H

hardware-based load balancing, 440

heartbeat networks, 489

hiberfil.sys file, 499

hibernate mode, 99

high availability

App-V, 197–199

deployment considerations, 195

load balancing and, 439–441

personal virtual desktops, 486–489

- planning, 21–22
- pooled virtual desktops, 485–486
- RD Connection Broker service, 443–446
- RD Licensing role service, 447–449
- RD Session Host service, 21, 441–443
- RD Web Access role service, 447
- RDS support, 438–449
- host (A) records, 441, 447
- HTTP (Hypertext Transfer Protocol Service), 278
- HTTPS binding, 278–279
- Hyper-V**
 - Client Hyper-V. *See* Client Hyper-V
 - Data Deduplication feature, 510
 - described, 95
 - event logs, 568
 - integration services, 138–141
 - management pack, 583
 - Storage Spaces and, 12
 - troubleshooting checkpoints, 329
 - unique IDs, 133–134
 - virtual process support, 135
- Hyper-V Data Exchange Service**, 140
- Hyper-V Guest Service Interface**, 140
- Hyper-V Guest Shutdown Service**, 139
- Hyper-V GUI Management Tools feature**, 105
- Hyper-V Heartbeat Service**, 140
- Hyper-V-Hypervisor event log**, 568
- Hyper-V Manager**
 - creating checkpoints, 328–329
 - creating external virtual switches, 109–110
 - creating virtual hard disks, 147–148
 - creating virtual machines, 113–120
 - described, 104–106
 - viewing checkpoints, 153–154
- Hyper-V module (PowerShell)**, 105
- Hyper-V Remote Desktop Virtualization Service**, 140
- Hyper-V Server**, 507–508
- Hyper-V Time Synchronization Service**, 139
- Hyper-V-VMMS event log**, 568
- Hyper-V Volume Shadow Copy Requestor**, 140
- Hyper-V-Worker event log**, 568
- Hypertext Transfer Protocol Service (HTTP)**, 278

I

- I/O operations per second (IOPS)**, 150, 512–513
- ICS (Internet Connection Sharing)**, 112
- IIS (Internet Information Server)**
 - App-V management server and, 205–206, 278
 - App-V publishing server and, 176, 179, 278
 - App-V reporting server and, 278
 - RD Web Access role service and, 410, 466–467
 - reverse proxy servers, 548
- IIS Manager**, 466–467, 475
- images**
 - building with App-V, 254–259

- building with App-V client, 228
- reducing number of, 254
- Sequencer configuration and, 326
- Sysprep tool and, 503
- Import-AppvServerPackage cmdlet**, 289, 295
- Import Certificates window**, 551
- Import Management Packs window**, 584–585
- Import-Module cmdlet**, 268, 351, 356
- Import Virtual Machine Wizard**, 133–134
- importing**
 - configuration files, 308
 - management packs, 583–585
 - packages, 178, 303–304
 - virtual machines, 133–135
- inherited environments**, 384
- .INI files**, 380–381, 455
- Install Licenses Wizard**, 437–438
- Installation Wizard**, 331
- Integration node (App-V policies)**, 272
- integration services, virtual machines**, 138–141
- internal networks**, 111–112
- Internet Connection Sharing (ICS)**, 112
- Internet Explorer**, 333
- Internet Information Server (IIS)**
 - App-V management server and, 205–206, 278
 - App-V publishing server and, 176, 179, 278
 - App-V reporting server and, 278
 - RD Web Access role service and, 410, 466–467
 - reverse proxy servers, 548
- Internet SCSI (iSCSI)**, 488, 506, 515
- Invoke-RDUserLogoff cmdlet**, 471
- IOPS (I/O operations per second)**, 150, 512–513
- IP addresses**, 441
- iSCSI (Internet SCSI)**, 488, 506, 515
- IsPublishedToUser property**, 163

K

- kernel mode**, 98, 160–161
- Key Management Service (KMS)**, 492–493
- killing processes**, 306
- KMS (Key Management Service)**, 492–493
- known folders**, 53–54

L

- last access time stamp**, 499
- latency, network**, 18–19, 66
- LDIFDE.exe utility**, 68
- legal requirements**
 - RDS, 404
 - Shadow option, 471
 - usage scenarios, 13
- licensing**
 - App-V considerations, 158
 - KMS support, 493
 - post-installation step, 323

RD Session Host servers, 22, 402, 448, 455
 RDS, 401–403
 virtualized environment requirements, 22–24

Links folder, 54

Linux operating system

Citrix support, 253
 Client Hyper-V, 125
 integration services, 99, 139
 Operations Manager, 574, 577
 Secure Boot, 125

List Folder Contents NTFS permission, 78

List Folder NTFS permission, 47, 77

Live Migration, 487

load balancing

described, 439
 DNS round robin, 440–441
 hardware-based, 440, 447
 high availability for server roles, 486
 RD Connection Broker service, 20
 RD Web Access service, 447
 SSL certificates and, 280
 Windows Network Load Balancing, 439–440

load simulations, 511–512

Local folder (AppData), 34

LocalHelp setting, 466–467

LocalLow folder (AppData), 34–35

locally attached storage, 503–505

log files and logging

App-V, 282
 App-V client, 255–259
 monitoring performance, 566–568
 personal virtual desktops, 499
 pooled virtual desktops, 499
 RD Gateway, 554–555

Log Off option (connections), 471

Login VSI, 512

M

MAC (media access control) addresses, 111

Mac OS X operating system, 8, 17

machine catalog, 217–218

machine configuration files, 217, 311, 382–383

mailto extension, 310, 381

management databases, 201–204

management networks, 489

management operating system, 98

management packs (Operations Manager)

described, 579–581
 importing, 583–585
 installing, 583–585
 monitoring application virtualization, 581–583
 monitoring desktop virtualization, 581–583

management server database (App-V), 178–179

management servers (App-V)

accessing packages, 284

adding packages, 284

configuring, 204–207

described, 178–179, 282

disaster recovery and, 199–200

event logs, 568

IIS support, 205–206, 278

infrastructure requirements, 188

installing management databases, 201–204

managing by using Windows PowerShell, 289–297

publishing packages, 285

refreshed list of published applications and, 175

round-trip response time, 195–196

service disruption impact, 194

Windows PowerShell support, 294, 296

management servers (Operations Manager), 574

mandatory user profiles, 52–53

manifest file, 310, 358, 382

Manifest.xml file, 214

manual deployment, 225

MBAM (Microsoft BitLocker Administration and Monitoring), 2

MDOP (Microsoft Desktop Optimization Pack), 2, 24, 88, 157

MDT (Microsoft Deployment Toolkit), 79, 495, 533

MED-V (Microsoft Enterprise Desktop Virtualization), 2

media access control (MAC) addresses, 111

memory management

capacity planning, 513–514, 517–518

CSV cache, 508–509

operating system differences, 491

optimizing, 588

user mode, 98

virtual machines, 101, 137–138, 156

Memory\Available MBytes counter, 588

Memory\Pages/sec counter, 586, 588

Merge-VHD cmdlet, 150

metered networks, 65–66

Microsoft accounts, 71, 75–76

Microsoft Advanced Group Policy Management (AGPM), 2

Microsoft Application Virtualization. See App-V (Application Virtualization)

Microsoft Application Virtualization (App-V) Server 5.0 Setup Wizard, 177–178

Microsoft Azure Multi-Factor Authentication, 557, 561–564

Microsoft BitLocker Administration and Monitoring (MBAM), 2

Microsoft Clearinghouse, 434–435, 438

Microsoft Deployment Toolkit (MDT), 79, 495, 533

Microsoft Desktop Optimization Pack (MDOP), 2, 24, 88, 157

Microsoft Diagnostics and Recovery Toolset (DaRT), 2

Microsoft Download Center, 317

Microsoft Enterprise Desktop Virtualization (MED-V), 2

Microsoft Exchange Server, 42–44, 579

Microsoft Forefront Threat Management Gateway (TMG), 547

Microsoft Hyper-V. *See* Client Hyper-V; Hyper-V

Microsoft Intune, 79, 227

Microsoft Management Console (MMC), 475

Microsoft .NET Framework, 71, 229, 377

Microsoft Office 365, 16, 42–43, 351

Microsoft Office 2013, 85, 351–357, 493

Microsoft Outlook, user state virtualization and, 42–44

Microsoft SharePoint Server, 324

Microsoft Software Assurance (SA), 23

Microsoft SQL Server

App-V compatibility and, 324

configuring permissions, 444

databases supported, 281, 443

high availability for server roles, 486

monitoring, 579

scripting support, 202–204

Microsoft User Experience Virtualization. *See* UE-V (User Experience Virtualization)

middleware, 334, 380

mirrored spaces, 12

MMC (Microsoft Management Console), 475

mobile devices

compatibility considerations, 40

identifying business needs, 26–27

metered networks and, 65

RD client, 542

RDS support, 388

remote access from, 555–557

Remote Desktop support, 392, 399

VPNs and, 8, 16–17

monitoring performance

App-V considerations, 195–197

application virtualization, 565–572

desktop virtualization, 565–572, 585–590

high availability and, 21

network connectivity and, 19

Operations Manager support, 572–585

resource bottlenecks, 585–586

roaming user profiles and, 46

SPEC CPU2006 Results, 516

storage virtualization and, 12

virtual machine optimization, 156

virtual machine storage speed, 101

virtualization implementation and, 25

Mount-AppvClientConnectionGroup cmdlet, 296

Mount-AppvClientPackage cmdlet, 296

Move Wizard, 131–132, 153

Mozilla Firefox Setup Wizard, 340

msDS-PrimaryComputer attribute, 413

.MSI files

creating package accelerators, 358

described, 186, 325

extracting from executables, 229

GPO deployment, 225

saving packages as, 320

sequenced applications, 321

stand-alone deployment model, 185–186, 192–193, 251

UE-V agent and, 79–80

multi-factor authentication, 557, 561–564

Multi-Factor Authentication Server, 563

Music folder, 35, 54

N

Name registry key, 274

naming conventions, 309

NAS (network attached storage), 36, 224

native installation path, 331–332

Nested Page Tables, 398

.NET Framework, 71, 229, 377

network attached storage (NAS), 36, 224

network bridges, 110–111

Network Connections window, 110

network interface card (NIC), 515

Network Interface\Bytes Total/sec counter, 586

Network Level Authentication, 431

Network Load Balancing (NLB), 439–440, 447

Network Policy Server (NPS), 550, 557, 560–561

network shares

offline folders, 62–63

redirecting folders in, 55

roaming user profiles, 47–48

user profiles disks, 523

networks and networking

capacity planning, 514–515

connectivity in virtualized environments, 18–20

failover clusters and, 489

network latency and, 18–19, 66

Process Monitor support, 571–572

RD Gateway configuration scenarios, 543–548

New-AppvPackageAccelerator cmdlet, 296, 366–367

New-AppvSequencerPackage cmdlet, 296, 351, 377

New-AppvServerConnectionGroup cmdlet, 291, 295

New-RDSessionDeployment cmdlet, 419

New-SessionDeployment cmdlet, 419

New-VHD cmdlet, 106, 148

New Virtual Hard Disk Wizard

Choose Disk Format page, 142, 147

Choose Disk Type page, 144, 147

Configure Disk page, 145–148

Specify Name And Location page, 147

New Virtual Machine Wizard

Assign Memory page, 116–117

Configure Networking page, 117–118

Connect Virtual Hard Disk page, 118–119

Installation Options page, 119–120

Specify Generation page, 115–116

Specify Name And Location page, 114–115

New-VirtualDisk cmdlet, 510

New-VM cmdlet, 120–121
New-VMSwitch cmdlet, 112
NIC (network interface card), 515
NIC Teaming feature, 515
NLB (Network Load Balancing), 439–440, 447
Not Available state, 193–194

notifications

App-V client, 195
 IOPS value, 150
 Operations Manager, 573
 RD Session Host server, 590
 RemoteApp, 454
 Shell Hardware Detection, 498
 UE-V, 82, 86

NPS (Network Policy Server), 550, 557, 560–561

NTFS permissions

roaming user profiles, 47–48
 UE-V, 77–78

NTUSER.DAT file, 34, 46

NTUSER.MAN file, 52–53

O

Object Linking and Embedding Database (OLE DB), 380

occasionally roaming user scenario, 42–43

off state (VMs), 126

Office 365, 16, 42–43, 351

Office 2013, 85, 351–357, 493

Office Deployment Tool for Click-to-Run, 351–356

Offline Files dialog box, 64

Offline Files feature

configuring, 59–66
 operating modes, 60–62
 synchronization problems, 61–62
 user state virtualization and, 36, 39, 41–43

offsite replication, 505

OLE DB (Object Linking and Embedding Database), 380

OOBE (out-of-box experience), 326, 502

Open Existing Package Wizard

Completion page, 302
 Create Package page, 173–174, 302
 Installation page, 170–171, 300–301
 Installation Report page, 171–172, 301
 Prepare Computer page, 168–169, 299
 Prepare For Streaming page, 172–173, 301–302
 Select Installer page, 169–170, 299–300
 Select Package page, 167–168, 299
 Select Task page, 166–167, 298

Open Packaging Conventions standard, 186

open software description (OSD) files, 219–220

operating system image, 228

Operational event log, 256, 568

Operations console (Operations Manager)

Administration Overview page, 576–577
 described, 575–576
 Monitoring Overview page, 583–584

Operations Manager (System Center)

described, 25, 565, 572–573
 implementation components, 573–576
 installing, 576–579
 management packs, 579–585

Optimize-VHD cmdlet, 150

ordering packages in connection groups, 292

OSD (open software description) files, 219–220

OSQL command-line application, 204

out-of-box experience (OOBE), 326, 502

Outlook, user state virtualization and, 42–44

P

package accelerators

creating, 358–367
 creating packages from, 367–376
 described, 357–358

package editor

Advanced tab, 349–350
 Change History tab, 347
 Deployment tab, 346
 described, 345
 Package Files tab, 348
 Properties tab, 346
 Shortcuts and FTAs tab, 350
 Virtual Registry tab, 347–348
 Virtual Services tab, 348–349

package function (ODT), 354

package ID (applications), 216, 310

PackageHistory.xml file, 186, 325

PackageInstallationRoot registry key, 274

packages (App-V)

adding applications, 303, 379–380
 adding by using Windows PowerShell, 289–290
 adding to App-V Management Console, 303–304
 adding to connection groups, 291–292
 adding to management servers, 284
 assigning, 308–309
 best practices for configuring, 331–332
 contents of, 186–187
 creating, 284, 355–356
 creating from package accelerators, 367–376
 customizing, 344–345
 data collection about, 314
 described, 186
 dynamic configuration, 310–312
 editing, 379
 editing default configuration, 286
 editing existing, 302–303
 exporting configuration, 287, 307–308
 granting access, 284–286
 importing, 178, 303–304
 isolating, 287
 managing, 284–287, 294–296
 moving to repositories, 284

- naming conventions, 309
- ordering in connection groups, 292
- publishing, 285
- publishing by using Windows PowerShell, 289–290
- removing, 251–252
- shared content store and, 220–224
- supported states, 193
- unpublishing, 285
- updating applications in, 298–302, 378–379
- version considerations, 308–309
- parent partitions, 98–99**
- parity spaces, 12**
- Partially Loaded state, 193**
- partitions**
 - creating during manual install, 496
 - eliminating system partition, 495–496
 - modifying, 496–497
- pass-through disks, 143**
- PasswordChangeEnabled setting, 466**
- paused state (VMs), 126, 152**
- Performance Monitor, 570–571, 586**
- performance monitoring**
 - App-V considerations, 195–197
 - application virtualization, 565–572
 - desktop virtualization, 565–572, 585–590
 - high availability and, 21
 - network connectivity and, 19
 - Operations Manager support, 572–585
 - resource bottlenecks, 585–586
 - roaming user profiles and, 46
 - SPEC CPU2006 Results, 516
 - storage virtualization and, 12
 - virtual machine optimization, 156
 - virtual machine storage speed, 101
 - virtualization implementation and, 25
- Permissions.sql file, 202–203**
- personal virtual desktops**
 - capacity planning, 511–518
 - comparing options, 484–485
 - creating virtual desktop collections, 524–533
 - deploying RD Virtualization Host servers, 519–523
 - described, 10–11, 481–484
 - high availability, 486–489
 - implementing RemoteApp for Hyper-V, 535–539
 - optimizing antivirus software, 500–501
 - optimizing operating system configuration, 497–503
 - planning and creating templates, 489–497
 - planning storage, 503–511
 - updating, 494–495
 - user profile disks, 523–524
 - virtual desktop collections, 525–526
- .PFX files, 475**
- Physical Disk\Current Disk Queue Length counter, 586**
- Pictures folder, 35, 54**
- Ping command, 19**
- PKI (Public Key Infrastructure), 550–552**
- plug-ins, 322–323, 334, 380–381**
- PnP device redirection, 469**
- POC (proof-of-concept) deployment, 414**
- pooled virtual desktops**
 - App-V optimizations, 500
 - capacity planning, 511–518
 - comparing options, 484–485
 - creating virtual desktop collections, 524–533
 - deploying RD Virtualization Host servers, 519–523
 - described, 10–11, 481–483
 - high availability, 485–486
 - implementing RemoteApp for Hyper-V, 535–539
 - optimizing antivirus software, 500–501
 - optimizing operating system configuration, 497–503
 - planning and creating templates, 489–497
 - planning storage, 503–511
 - supporting management tasks, 14
 - updating, 494, 533–535
 - user profile disks, 523–524
 - user state, 25
 - virtual desktop collections, 525–533
- port configurations, RD Gateway servers, 548–553**
- PowerShell (Windows). See also specific cmdlets**
 - adding and publishing packages, 289–290
 - App-V client and, 215–216, 229, 263–264, 267–269, 315–317
 - AppVClient module, 268–269, 294–296
 - AppVSequencer module, 294, 296, 351
 - AppVServer module, 294–295
 - bulk changes and, 68
 - Client Hyper-V settings, 103
 - configuring App-V client, 267–269
 - configuring UE-V settings, 77–78
 - creating and managing connection groups, 290–294
 - creating package accelerators, 366–367
 - creating packages from package accelerators, 377
 - deploying RDS, 419
 - deploying UE-V agent, 81
 - elevated prompts, 297
 - ESD model and, 191
 - Hyper-V module, 105
 - Integrated Scripting Environment, 106–107
 - managing connections, 471
 - managing management servers, 289–297
 - managing UE-V agent, 83–84
 - RemoteApp programs, 461
 - removing applications, 305–306
 - UE-V requirements, 71
- pre-publishing applications, 254**
- Primary Computer feature, 66–68, 413**
- Primary Feature Block, 321**
- printer redirection, 469–470**
- private key, 553**
- private networks, 111**

PrivateModeSessionTimeoutInMinutes setting, 466

process ID, 306

Process Monitor, 571–572

processing

application requests, 213–214

capacity planning, 515–518

identifying startup processes, 590

killing processes, 306

Processor Information\%Processor Time counter, 586

Process\Working Set counter, 588

ProfileImagePath registry key, 34

ProfileList registry key, 34

proof-of-concept (POC) deployment, 414

Properties dialog box

certificates, 476

Documents folder, 54–55, 58

Folder Redirection settings, 56–57

RD Gateway server, 548–557

RDS deployment, 477

RemoteApp programs, 458–461

session collections, 426–433, 442, 468–469

shared content store, 223–224

shared folders, 62

user accounts, 48–49

virtual hard disks, 146

ProPlusVolume_en-us_x86.appv file, 355

Public Key Infrastructure (PKI), 550–552

PublicModeSessionTimeoutInMinutes setting, 466

Publish-AppvClientPackage cmdlet

described, 296

-Global setting, 269, 357

-Name setting, 163, 251, 269

Publish-AppvServerConnectionGroup cmdlet, 293, 295

Publish-AppvServerPackage cmdlet, 164, 290, 295

Publish RemoteApp Programs Wizard

Select RemoteApp Programs page, 457–458, 537–538

Select Virtual Desktop page, 536–537

published applications

assignment of, 308–309

automatically loading, 273

copy access and configuration, 303–304

editing default configuration for packages, 306–307

enabling scripting for dynamic configuration, 310–312

exporting configuration, 307–308

global settings, 325

naming conventions, 309

removing, 305–306

updating, 297–303

updating connection groups, 304–305

viewing, 175–176, 178–179

Published state, 193

Publishing Feature Block, 321

Publishing node (App-V policies), 272

publishing servers (App-V)

App-V client and, 215, 295–296

capacity planning, 196–197

deploying and configuring, 208–210

described, 179–183

disaster recovery and, 199–200

displaying associated, 296

IIS support, 176, 179, 278

infrastructure requirements, 188–189

monitoring performance, 566

optimizing, 500

refreshed list of published applications and, 175–176

registering and unregistering, 282–284

round-trip response time, 195–196

service disruption impact, 193–194

PUBLISHING_MGT_SERVER_REFRESH_INTERVAL registry

setting, 176, 208

PublishingService registry key, 176

PXE boot, 124

Q

QoS (Quality of Service), 20, 150–151

Quality of Service (QoS), 20, 150–151

R

RADC (RemoteApp and Desktop Connections),

462–465, 479

RAID (Redundant Array of Independent Disks), 12, 504–505

RD CAPs (Remote Desktop connection authorization policies), 391, 557–558, 560–561

RD Client, 542

RD Connection Broker - Enable Single Sign On certificate, 474

RD Connection Broker - Publishing certificate, 473–474

RD Connection Broker role service

described, 390, 482

event log, 568

high availability, 443–446, 486

load balancing, 20

monitoring performance, 566, 583

planning, 408–410

pooled virtual desktops, 483

RDS considerations, 393, 522

scalability, 409

session collection, 394–395

RD Gateway Manager

described, 548

logging, 554–555

remote access from mobile devices, 555–556

simultaneous connections, 554

SSL bridging, 553

TCP and UDP ports, 548–553

RD Gateway role service

certificate support, 475, 550–553

controlling access, 388, 482, 557–564

deployment settings, 521–522

described, 391, 542–543

- monitoring performance, 566, 582
- network configuration, 543–548
- simultaneous connections, 554
- troubleshooting, 556–557
- RD Licensing Manager, 434–435, 437**
- RD Licensing role service**
 - configuring servers, 434–438
 - deployment settings, 522
 - described, 391
 - event logs, 568
 - high availability, 447–449
 - monitoring performance, 566, 582
- RD RAPs (Remote Desktop resource authorization policies), 557, 559–560**
- RD Session Host role service**
 - App-V client and, 183
 - App-V for RDS client and, 210
 - application compatibility, 485
 - application update support, 14
 - capacity planning, 21
 - described, 7–9, 387, 390
 - Fair Share feature, 397
 - high availability for, 21, 441–443
 - installing, 421–422
 - installing applications, 454–457
 - licensing considerations, 22, 402, 448, 455
 - load balancing and, 20, 411
 - monitoring performance, 565–566, 569, 571–572, 582, 585, 587–590
 - network connectivity, 18
 - planning, 406–408
 - preserving user state, 411–414
 - print driver management and, 470
 - RemoteApp support, 388, 451, 454–457, 535–536
 - session-based VDI, 252
 - session-based virtual desktops, 387–388, 404, 415
 - session collections, 422–432
 - UE-V support, 71
 - user state virtualization and, 25
- RD Virtualization Host service**
 - automating VM creation, 9
 - capacity planning, 511, 513–518
 - configuring storage, 503–504
 - deploying servers, 519–523
 - described, 482
 - event log, 568
 - failover clustering, 488
 - high availability, 486
 - monitoring performance, 566, 583, 587
 - network connectivity and, 18
 - optimizing operating system configuration, 497
 - saving virtual desktops, 533
 - video adapter support, 398
 - virtual desktop templates, 525
- RD Web Access role service**
 - certificate support, 474
 - customizing, 466–467
 - deployment settings, 522
 - described, 390–391, 482
 - high availability, 447, 486
 - monitoring performance, 583
 - planning, 410–411
 - RD Gateway server and, 543
 - user connectivity support, 393–395
 - web feed, 463
- RDC (Remote Desktop Connection)**
 - Advanced tab, 401
 - configuration options, 399–401
 - described, 392–396
 - Display tab, 399
 - Enhanced Session Mode, 129–130
 - Experience tab, 400–401
 - General tab, 399
 - Local Resources tab, 399–400
 - mobile devices, 555–556
 - port defaults, 559
 - Program tab, 400
 - RemoteApp support, 453
- .RDP files, 393**
- RDP (Remote Desktop Protocol), 387, 453, 514, 588**
- RDS (Remote Desktop Services)**
 - App-V and, 204, 210, 253
 - App-V client and, 227, 247–250, 253
 - architectural overview, 390–392
 - assessing infrastructure requirements, 403–406
 - certificate support, 473–478
 - configuring and managing client connections, 461–472
 - configuring certificates and single sign-on, 472–479
 - connecting to RemoteApp programs, 392–395
 - connecting to virtual desktops, 392–395
 - deploying, 404–405, 519–523
 - described, 387–389
 - determining needs, 403–404
 - determining resource requirements, 405
 - determining user requirements, 405–406
 - event logs, 568
 - functionality enhancing client experience, 395–397
 - high availability, 438–449
 - licensing, 401–403
 - management packs, 582
 - monitoring performance, 566
 - package configuration best practices, 332
 - RDC configuration options, 399–401
 - Remote Desktop comparison, 389
 - RemoteFX features, 397–398
 - session-based virtual desktops and, 9, 387
 - Storage Spaces and, 12
- Read Data NTFS permission, 47, 77**
- read-only domain controllers (RODCs), 546**

Real-Time Streaming Protocol (RTSP), 180

rebooting, 331

recipes, items to document in, 322–324

Recreate All Virtual Desktops Wizard, 534–535redirected folders. *See* Folder Redirection feature

redundancy, data, 504–505

Redundant Array of Independent Disks (RAID), 12, 504–505**Register-UevTemplate cmdlet, 83, 93**

Registered state, 193

registering publishing servers, 282–283

Registry Editor, 274–275

registry keys

App-V client, 214, 218, 273–275

App-V publishing server, 176

disabling Fair Share, 589

Process Monitor support, 571–572

UE-V settings, 92–93

user profile settings, 34–35

virtualized applications, 159

Registry.dat file, 187, 275, 284, 325**Relative Weight processor setting, 136****Remote Desktop connection authorization policies (RD CAPs), 391, 557–558, 560–561****Remote Desktop Connection Broker role service**

described, 390, 482

event log, 568

high availability, 443–446, 486

load balancing, 20

monitoring performance, 566, 583

planning, 408–410

pooled virtual desktops, 483

RDS considerations, 393, 522

scalability, 409

session collection, 394–395

Remote Desktop Connection (RDC)

Advanced tab, 401

configuration options, 399–401

described, 392–396

Display tab, 399

Enhanced Session Mode, 129–130

Experience tab, 400–401

General tab, 399

Local Resources tab, 399–400

mobile devices, 555–556

port defaults, 559

Program tab, 400

RemoteApp support, 453

Remote Desktop Easy Print driver, 470**Remote Desktop Gateway Manager**

described, 548

logging, 554–555

remote access from mobile devices, 555–556

simultaneous connections, 554

SSL bridging, 553

TCP and UDP ports, 548–553

Remote Desktop Gateway role service

certificate support, 475, 550–553

controlling access, 388, 482, 557–564

deployment settings, 521–522

described, 391, 542–543

monitoring performance, 566, 582

network configuration, 543–548

troubleshooting, 556–557

Remote Desktop Licensing Manager, 434–435, 437**Remote Desktop Licensing role service**

configuring servers, 434–438

deployment settings, 522

described, 391

event logs, 568

high availability, 447–449

monitoring performance, 566, 582

Remote-Desktop-Management-Service event log, 568**Remote Desktop Protocol (RDP), 387, 453, 588****Remote Desktop resource authorization policies (RD RAPs), 557, 559–560****Remote Desktop Services (RDS)**

App-V and, 204, 210, 253

App-V client and, 227, 247–250, 253

architectural overview, 390–392

assessing infrastructure requirements, 403–406

certificate support, 473–478

configuring and managing client connections, 461–472

configuring certificates and single sign-on, 472–479

deploying, 404–405, 519–523

described, 387–389

determining needs, 403–404

determining resource requirements, 405

determining user requirements, 405–406

event logs, 568

functionality enhancing client experience, 395–397

high availability, 438–449

licensing, 401–403

monitoring performance, 566

package configuration best practices, 332

RDC configuration options, 399–401

Remote Desktop comparison, 389

session-based virtual desktops and, 9

Storage Spaces and, 12

Remote Desktop Session Host role service

App-V client and, 183

App-V for RDS client and, 210

application compatibility, 485

application update support, 14

capacity planning, 21

described, 7–9, 387, 390

Fair Share feature, 397

high availability for, 21, 441–443

installing, 421–422

installing applications, 454–457

licensing considerations, 22, 402, 448, 455

- load balancing and, 20, 411
- monitoring performance, 582
- monitoring performance counters, 565–566, 569, 571–572, 585, 587–590
- network connectivity, 18
- planning, 406–408
- preserving user state, 411–414
- print driver management and, 470
- RemoteApp support, 388, 451, 454–457, 535–536
- session-based VDI, 252
- session-based virtual desktops, 387–388, 404, 415
- session collections, 422–432
- UE-V support, 71
- user state virtualization and, 25

Remote Desktop Virtualization Host service

- automating VM creation, 9
- capacity planning, 511, 513–518
- configuring storage, 503–504
- deploying servers, 519–523
- described, 482
- event log, 568
- failover clustering, 488
- high availability, 486
- monitoring performance, 566, 583
- network connectivity and, 18
- optimizing operating system configuration, 497
- video adapter support, 398
- virtual desktop templates, 525

Remote Desktop Web Access role service

- certificate support, 474
- customizing, 466–467
- deployment settings, 522
- described, 390–391, 482
- high availability, 447, 486
- monitoring performance, 583
- planning, 410–411
- RD Gateway server and, 543
- user connectivity support, 393–395
- web feed, 463

RemoteApp

- application compatibility, 6, 15
- configuring programs, 458–461
- configuring RADC, 462–465
- described, 5–6, 451–453, 490
- device redirection, 467–469
- managing connections, 470–472
- mobile devices, 556
- printer redirection, 469–470
- publishing programs, 457–458
- RD Session Host servers, 388, 451, 454–457
- RD Web Access customization, 466–467
- RDS connecting to, 392–396
- UE-V Generator and, 88
- understanding programs, 453–454

- RemoteApp and Desktop Connection Management event log, 568

- RemoteApp and Desktop Connections event log, 568

- RemoteApp and Desktop Connections (RADC), 462–465, 479

- RemoteApp for Hyper-V, 535–539

- RemoteDesktopServices-RdpCoreTS event log, 568

- RemoteDesktopServices-SessionServices event log, 568

- RemoteFX features, 100, 397–398, 490–491, 583

- Remove-AppvClientConnectionGroup cmdlet, 296–297

- Remove-AppvClientPackage cmdlet, 251–252, 296–297, 305

- Remove-AppvPublishingServer cmdlet, 296–297

- Remove-AppvServerConnectionGroup cmdlet, 295

- Remove-AppvServerPackage cmdlet, 295

- Remove-VMSwitch cmdlet, 112

- Rename-VMSwitch cmdlet, 112

- renaming checkpoints, 154

- Repair-AppvClientConnectionGroup cmdlet, 296

- Repair-AppvClientPackage cmdlet, 296

- Report Manager, 317–318

- Reporting node (App-V policies), 272

- reporting server database (App-V), 183–184

- reporting servers (App-V)

- App-V client configuration, 315–317

- data collection, 314

- described, 183–184

- disaster recovery and, 199–200

- generating reports, 317–318

- IIS support, 278

- infrastructure requirements, 188

- reporting process overview, 312–313

- service disruption impact, 194

- reporting servers (Operations Manager), 575

- Report.xml file, 186, 324

- Resize-VHD cmdlet, 150

- resource bottlenecks, 585–586

- Resource Monitor, 570

- Restore-UevUserSetting cmdlet, 75, 83

- reverse proxy servers, 547–548

- Roaming folder (AppData), 35, 53

- roaming user profiles

- configuring, 44–52

- configuring user object, 48–49

- described, 3

- folder names, 48

- managing by using Group Policy, 50–52

- managing user data requirements, 39

- mandatory, 52–53

- performance issues, 46

- preparing network shares, 47–48

- preserving user state, 411–412

- RemoteApp support, 453

- user profile disks and, 523–524

- user state virtualization and, 42–44, 75–76

- RODCs (read-only domain controllers), 546

Root (folder), 187, 284, 325, 380
 round-trip response time, 195–196
 RTSP (Real-Time Streaming Protocol), 180
 Run As profile, 580
 running state (VMs), 126

S

SA (Software Assurance), 23
 SAN (subject alternative names) certificates, 473
 SANs (storage area networks)
 capacity planning, 515
 Client Hyper-V, 100
 described, 505–506
 failover clustering, 488
 shared content store, 224
 SAS (shared serial attached SCSI), 488
 Saved Games folder, 54
 saved state (VMs), 126
 scale-out file servers, 488, 506–508, 515
 scripting
 enabling for dynamic configuration, 310–312
 optimizing Windows 8.1, 499
 SQL Server scripts, 202–204
 targeted, 381–385
 Scripting node (App-V policies), 272
 Searches folder, 54
 Second Level Address Translation (SLAT), 101–102, 398
 Secondary Feature Block, 321
 Secure Boot, 125
 Secure Sockets Layer (SSL)
 certificate support, 205, 278–280, 377, 550–553
 RD Gateway network configuration, 547, 550–554
 session collections and, 430
 security considerations
 HTTPS, 279
 identifying business needs, 28–29
 RDS, 388
 signing .CAB files, 376–377
 SSL inspection, 547
 usage scenarios, 13
 wildcards and FQDN, 479
 security groups
 App-V infrastructure requirements, 187–188
 assigning packages, 308–309
 creating, 444
 NTFS permissions, 77
 RD CAPs, 557
 RD RAPs, 559
 removing users from, 175
 roaming user profiles and, 51
 security identifier (SID), 203, 523
 Select Existing Certificate dialog box, 478
 Select Previous Version dialog box, 308
 Send-AppvClientReport cmdlet, 296–297
 Send Message option (connections), 471

Send-RDUserMessage cmdlet, 471

Sequencer (App-V)

Add New Application option, 303
 application sequencing, 161–162
 application update, 165–174, 297–303
 Completion page, 302, 375–376
 Configure Software page, 320, 341–342, 373–374
 configuring, 326–330
 connection groups and, 380–381
 Create Package Accelerator page, 361–363, 365–366
 Create Package page, 173–174, 302, 345, 372–373, 379
 creating packages, 284, 310
 Customize page, 344
 deployment models and, 192
 described, 161, 179–183
 editing packages, 302–303, 306–307
 ESD model and, 191
 Guidance page, 369–370
 Installation Files page, 360–361
 Installation page, 170–171, 300–301, 341, 379
 Installation Report page, 171–172, 301, 343, 379
 installing, 336
 Local Installation page, 370–371
 Modify An Existing Virtual Application Package option, 166
 Package Name page, 338–339, 371–372
 Packaging Method page, 337, 367–368
 Prepare Computer page, 168–169, 299, 337, 378
 Prepare For Streaming page, 172–173, 301–302
 Run Each Program page, 374–375
 Select Guidance page, 364–365
 Select Installer page, 169–170, 299–300, 338, 378
 Select Package Accelerator page, 368–369
 Select Package page, 167–168, 299, 359–360
 Select Task page, 166–167, 298
 sequencing process overview, 321–322
 Streaming page, 379
 Target OS page, 344–345
 Update Application in Existing Package option, 166–167, 298–299
 updating packages, 378–381
 Verify Applications page, 363–364
 virtual machines as, 328–329
 Windows PowerShell support, 294–295, 351
sequencing applications. See also App-V Sequencer
 application types supporting, 333–334
 applications incompatible with, 332–333
 best practices, 330–332
 for connection groups, 380–381
 customizing packages, 344–345
 deploying Office 2013 by using App-V, 351–357
 described, 161–162
 dynamic configuration and targeted scripting, 381–385
 items to document in recipes, 322–324
 package accelerators, 357–377
 package editor, 345–350

- planning, 326–333
- portions of sequenced applications, 324–325
- preparing for, 335–336
- process overview, 321–322
- sequencing tasks, 337–343
- updating packages, 378–385
 - Windows PowerShell support, 351
- Sequencing Wizard, 332**
- Server Authentication, 552**
- server-based certificates, 280, 473**
- Server Configuration window, 436–437**
- server connections (App-V client)**
 - Autoload mechanism, 273
 - configuring by using Windows PowerShell, 267–269
 - described, 266–267
 - managing by using Group Policy, 270–273
 - registry settings, 273–275
- Server Manager**
 - BPA support, 415
 - client access licenses, 437–438, 475
 - described, 391–392
 - monitoring performance, 565–566
 - pooled virtual desktops, 526, 534
 - publishing RemoteApp programs, 457
 - RD Connection Broker role service, 445
 - RD Connection Broker servers, 442–443
 - RD Gateway servers, 548
 - RD Licensing servers, 434–436, 449
 - RD Session Host servers, 443
 - RD Virtualization Host servers, 519
 - RD Web Access server, 447, 467
 - RDS deployment, 415–416, 419, 473, 475, 519, 523
 - RemoteApp for Hyper-V, 536
 - session collection management, 423, 426
 - viewing collections, 470
- Server Message Block (SMB), 488, 507**
- service resource (SRV) record, 493**
- session-based VDI, 12, 252, 587**
- session-based virtual desktops**
 - benefits of, 9
 - comparing options, 484–485
 - deploying, 404, 415–422
 - described, 7–8, 387
 - high availability for RDS, 438–449
 - planning infrastructure, 403–415
 - preserving user state, 411–415
 - RD Connection Broker role service, 408–410
 - RD Licensing servers, 434–438
 - RD Session Host role service, 406–408
 - RD Web Access role service, 410–411
 - RDS considerations, 387–403
 - selecting, 11
 - session collections, 422–434
 - supporting management tasks, 14–15
 - user profile disks, 70
 - user state, 25
- session collections**
 - configuring, 426–433
 - configuring load balancing settings, 442
 - connecting to, 394–395
 - creating, 393, 423–426
 - described, 422–423
 - device redirection, 468–469
 - printer redirection, 469–470
 - user profile disks, 413
- session ID, 471**
- Set-ADUser cmdlet, 68**
- Set-AppvClientConfiguration cmdlet**
 - Autoload setting, 273
 - described, 296–297
 - EnablePackageScripts setting, 312, 356, 384
 - ReportingServerURL setting, 316
 - SharedContentStoreMode setting, 221
- Set-AppvClientMode cmdlet, 296–297**
- Set-AppvClientPackage cmdlet, 296–297**
- Set-AppvPublishingServer cmdlet, 267, 296–297**
- Set-AppvServerConnectionGroup cmdlet, 291–292, 295**
- Set-AppvServerPackage cmdlet, 295**
- Set-ExecutionPolicy cmdlet, 268**
- Set-FileStorageTier cmdlet, 509**
- Set-UevConfiguration cmdlet, 74, 83**
- Set-VHD cmdlet, 150**
- Set-VMSSwitch cmdlet, 106, 112**
- settings location templates**
 - custom, 87–93
 - described, 72–74
 - managing default, 84–87
 - Office 2013, 85
- settings packages, 73**
- settings storage location, 72–73, 77, 80**
- settings template catalog, 72–73, 77–78, 93**
- Setup.exe utility, 352–353, 355**
- Shadow option (connections), 471–472**
- shared content store, 220–224, 254, 295, 500**
- shared SAS bus, 507**
- shared serial attached SCSI (SAS), 488**
- SharedContentStoreMode registry key, 274**
- SharePoint Server, 324**
- Shell Hardware Detection, 498**
- shortcuts (App-V), 217–220, 286, 350**
- ShowDesktops setting, 467**
- SID (security identifier), 203, 523**
- signtool.exe utility, 377**
- SIM (subscriber identity module) card, 65**
- single-root I/O virtualization (SR-IOV), 100**
- single sign-on (SSO), 397, 474, 478–479**
- SLAT (Second Level Address Translation), 101–102, 398**
- sleep mode, 99**
- SMB (Server Message Block), 488, 507**
- snapshots (checkpoints)**

- applying, 154
- creating, 152–153, 328–329, 336
- deleting, 154–155
- described, 151–152, 328
- exporting, 154
- pooled virtual desktops, 533
- renaming, 154
- reverting to previous state, 154
- SAN support, 505
- troubleshooting, 329
- usage considerations, 153–156
- Software Assurance (SA), 23**
- Software Client extension point, 219**
- solid-state drives (SSDs)**
 - Data Deduplication feature, 510
 - mobile devices, 101
 - performance considerations, 147, 156, 513, 518
 - storage tiers, 508–509, 513
 - write-back cache, 509
- SPEC CPU2006 Results, 516**
- SQL Server**
 - App-V compatibility and, 324
 - configuring permissions, 444
 - databases supported, 281, 443
 - high availability for server roles, 486
 - monitoring, 579
 - scripting support, 202–204
- SQL Server Management Studio, 444**
- SQL Server Reporting Services, 313, 317–318**
- SR-IOV (single-root I/O virtualization), 100**
- SRV (service resource) record, 493**
- SSDs (solid-state drives)**
 - Data Deduplication feature, 510
 - mobile devices, 101
 - performance considerations, 147, 156, 513, 518
 - storage tiers, 508–509, 513
 - write-back cache, 509
- SSL Bridging, 553**
- SSL inspection, 547**
- SSL (Secure Sockets Layer)**
 - certificate support, 205, 278–280, 377, 550–553
 - RD Gateway network configuration, 547, 550–554
 - session collections and, 430
- SSO (single sign-on), 397, 474, 478–479**
- stand-alone deployment model**
 - App-V client, 221, 250–252, 312
 - deployment possibilities, 192–193
 - described, 163–164, 185
 - high availability and, 197
- standard applications, 159–160, 334**
- Start-AppvVirtualProcess cmdlet, 295**
- Start Install Licenses Wizard, 436**
- Start Menu folder, 53**
- starting state (VMs), 126**
- startup processes, identifying, 590**
- startup sequence, 332–333**
- static memory, 137**
- Stokes, Jeff, 499**
- Stop-AppvClientConnectionGroup cmdlet, 296**
- Stop-AppvClientPackage cmdlet, 296**
- storage area networks (SANs)**
 - capacity planning, 515
 - Client Hyper-V, 100
 - described, 505–506
 - failover clustering, 488
 - shared content store, 224
- storage planning and performance**
 - antivirus software and, 500–501
 - capacity planning, 512–513, 517–518
 - described, 503
 - local storage, 503–505
 - networking, 489
 - optimizing operating system configuration, 499
 - personal virtual desktops, 503–511
 - pooled virtual desktops, 503–511
 - SAN considerations, 100, 224, 488, 505–506
 - scale-out file servers, 488–489, 506–508
 - storage virtualization, 2, 11–12
 - Windows Server, 508–511
- storage Quality of Service, 150–151**
- Storage Spaces**
 - described, 2, 11–12, 510
 - redundancy options, 505, 507
 - storage tiers, 508, 513
- storage tiers, 508–509, 513**
- storage virtualization, 2, 11–12**
- StoreMode registry key, 221**
- Streaming node (App-V policies), 272**
- streaming packages, 332**
- streaming server, 179**
- StreamMap.xml file, 186, 325**
- subject alternative names (SAN) certificates, 473**
- subscriber identity module (SIM) card, 65**
- Sync-AppvPublishingServer cmdlet, 164, 215, 296**
- Sync Controller Application task, 74**
- synchronization**
 - metered networks and, 65
 - Offline Files problems, 61–62
 - reducing frequency, 590
 - UE-V and, 70–75, 84–87
- Sysprep tool, 134, 490, 501–503, 525**
- System account, NTFS permissions, 47–48**
- System Center Configuration Manager**
 - App-V and, 165, 185, 189–191, 198–199, 211
 - App-V client and, 226, 228, 230–246, 312
 - installing App-V client using, 230–246
 - personal virtual desktops, 494
 - pooled virtual desktops, 494
 - UE-V agent and, 79, 81
 - usage considerations, 228

System Center Operations Manager

- described, 25, 565, 572–573
- implementation components, 573–576
- installing, 576–579
- management packs, 579–585

System Configuration, 590**System event log, 568****System Preparation Tool 3.14 window, 502****System Restore, 499****systeminfo utility, 102****T****Target Folder Location setting, 56****targeted scripting, 381–385****Task Manager, 568, 590****taskkill -PID command, 306****TCP ports, RD Gateway servers, 548–553****TechNet website**

- App-V deployment methods, 165
- BranchCache, 63
- CopyProfile, 33
- coreinfo tool, 102
- described, 376
- GVLKs, 492
- high availability with Configuration Manager, 199
- Hyper-V, 125, 135
- KMS, 493
- MDOP, 2
- MDT, 496
- optimization script, 499
- Process Monitor, 572
- RDS cmdlets, 419
- sequencing applications by using PowerShell, 351
- shadowing connections, 472
- Storage Spaces, 505
- UE-V agent, 80
- Windows SIM, 496

Template Auto Update task, 93**temporary files, 330****Terminal Services, 9, 204, 387–388, 443****TerminalServices-SessionBroker event log, 568****TerminalServices-TSV-VmHostAgent event log, 568****test environments, 7, 151****Test-VHD cmdlet, 150****third-party production integration, 252–253****Threat Management Gateway (TMG), 547****timeout settings, 558****TMG (Threat Management Gateway), 547****ToolsSetup.exe utility, 88****TPM (Trusted Platform Module), 495****troubleshooting**

- App-V client installation, 255–259
- checkpoints, 329
- connection groups, 288
- creating certificates, 476

- deleting user profiles, 33–34
 - enabling debug logs and, 259
 - examining virtual registry, 275
 - inherited environments and, 384
 - naming conventions and, 309
 - Offline Files synchronization problems, 61–62
 - performance issues, 571
 - RD Gateway servers, 556–557
 - removing applications, 306
- Trusted Platform Module (TPM), 495**

U**UAC (User Account Control), 326****UCC (unified communications certificate), 473****UDP ports, RD Gateway servers, 548–553****UE-V 2.0 template gallery, 88****UE-V agent**

- deploying, 78–81
- managing, 81–84
- preparing to deploy, 77–78

UE-V Generator, 73–74, 87–93**UE-V (User Experience Virtualization)**

- architectural overview, 72–75
- deploying UE-V agent, 78–81
- described, 4, 31, 70–72
- licensing requirements, 24
- managing UE-V agent, 81–84
- MDOP and, 2
- operating system requirements, 71
- preparing to deploy UE-V agent, 77–78

UEFI firmware, 125**UNC (universal naming convention), 48, 507–508****unified communications certificate (UCC), 473****universal naming convention (UNC), 48, 507–508****UNIX-based operating systems, 574, 577****Unpublish-AppvClientPackage cmdlet, 251, 296, 305****Unpublish-AppvServerConnectionGroup cmdlet, 295****Unpublish-AppvServerPackage cmdlet, 295****unpublishing packages, 285****unregistering publishing servers, 282–283****Update-AppvSequencerPackage cmdlet, 296****Update-AppvServerConnectionGroup cmdlet, 295****updating applications**

- App-V application life cycle, 165–175, 297–303
- App-V benefits, 158
- application sequencing and, 323
- best practices, 330
- identifying business needs, 28

updating connection groups, 304–305**updating integration services, 140–141****updating personal virtual desktops, 494–495****updating pooled virtual desktops, 494, 533–535****URL registry key, 275****usage scenarios**

- application virtualization, 12–17

- legal and security requirements, 13
 - management tasks, 14–15
 - user state virtualization, 41–44
 - virtualization technologies, 12–17
 - User Account Control (UAC), 326**
 - user-based CALs, 22–24**
 - user catalog, 217–218**
 - user certificates, 553**
 - user configuration files**
 - App-V client and, 217
 - control settings, 311, 383
 - customizing, 295
 - dynamic deployment, 310–311, 382
 - file name, 165
 - importing, 308
 - Windows PowerShell support, 295
 - user data**
 - assessing requirements, 38–39
 - backing up, 36
 - caching updates to, 37
 - configuring Folder Redirection, 53
 - configuring Offline Files, 59
 - described, 3
 - folders storing, 35
 - UE-V synchronization and, 71
 - User Experience Virtualization (UE-V)**
 - architectural overview, 72–75
 - deploying UE-V agent, 78–81
 - described, 4, 31, 70–72
 - licensing requirements, 24
 - managing UE-V agent, 81–84
 - MDOP and, 2
 - operating system requirements, 71
 - preparing to deploy UE-V agent, 77–78
 - user mode, 98, 160**
 - user profile disks, 70, 413–414, 491, 523–524**
 - user profiles**
 - creating, 32–34
 - deleting, 33–34
 - described, 32
 - managing by using Group Policy, 50–52
 - mandatory, 52–53
 - modifying default, 33
 - roaming. *See* roaming user profiles
 - user data requirements assessment, 38–39
 - user settings requirements assessment, 39
 - user settings**
 - assessing requirements, 39
 - backing up, 36
 - caching updates to, 37
 - described, 39
 - Folder Redirection and, 39
 - migrating, 36
 - personal virtual desktops and, 527
 - roaming user profiles and, 39, 43, 253
 - session-based virtual desktops and, 433
 - UE-V and, 39, 71, 75, 82–83
 - user state**
 - described, 3, 31
 - determining user requirements, 406
 - managing, 35–36
 - preserving, 411–415
 - user state virtualization**
 - assessing user data requirements, 38–39
 - assessing user settings requirements, 39
 - benefits of, 35–38
 - challenges implementing, 25, 36–37
 - comparing options, 75–76
 - configuring technologies, 44–70
 - Credential Roaming, 4, 68–69
 - described, 1–3
 - evaluating compatibility considerations, 40
 - evaluating infrastructure and manageability requirements, 40–41
 - evaluating usage scenario considerations, 41–44
 - Folder Redirection. *See* Folder Redirection feature
 - Offline Files. *See* Offline Files feature
 - UE-V. *See* UE-V
 - user profiles. *See* user profiles
 - _UserConfig.xml file**
 - App-V client, 214
 - App-V packages, 186
 - application sequencing, 325, 355
 - Configuration Manager model, 165
 - dynamic configuration, 310, 382–383
 - UVHD-Template.vhdx file, 524**
- ## V
- VDA (Virtual Desktop Access), 23, 402**
 - VDI (Virtual Desktop Infrastructure)**
 - App-V client, 252–254
 - described, 12
 - monitoring application virtualization, 565–572
 - monitoring desktop virtualization, 565–572, 585–590
 - Operations Manager monitoring support, 572–585
 - RD Gateway access control, 557–564
 - RD Gateway configuration, 541–557
 - remote access, 541–543
 - session-based, 12, 252
 - user state virtualization and, 37–38
 - VM-based, 10–12, 252
 - vGPU (virtual graphics processing unit), 397–398**
 - VHD format, 141–142, 145, 148–149, 506**
 - VHDs (virtual hard disks)**
 - compacting files, 149
 - converting, 149
 - creating, 147–148
 - described, 141
 - differencing disks, 145–147, 150, 152–153, 155
 - dynamical expanding, 143–145

- editing, 148–150
 - expanding, 149
 - fixed-size, 143–145
 - formats supported, 100, 141–143
 - merging, 150
 - shrinking, 149
 - storage Quality of Service, 150–151
 - user profile disks, 523–524
- VHDX format, 100, 142–145, 148, 413**
- Videos folder, 35, 54**
- Virtual Desktop Access (VDA), 23, 402**
- virtual desktop collections**
 - creating, 524–533
 - described, 393
 - save delay, 533
 - user profile disks and, 524
- Virtual Desktop Infrastructure (VDI)**
 - App-V client, 252–254
 - described, 12
 - monitoring application virtualization, 565–572
 - monitoring desktop virtualization, 565–572, 585–590
 - Operations Manager monitoring support, 572–585
 - RD Gateway access control, 557–564
 - RD Gateway configuration, 541–557
 - remote access, 541–543
 - session-based, 12, 252
 - user state virtualization and, 37–38
 - VM-based, 10–12, 252
- virtual desktop templates**
 - activating operating system, 491–494
 - capacity planning, 512
 - creating, 489
 - creating personal virtual desktops, 484
 - described, 489–490
 - eliminating system partition, 495–497
 - optimizing operating system configuration, 497–503
 - selecting operating systems, 490–491
 - Sysprep support, 501–503
 - updating, 534
 - updating applications, 494–495
 - updating operating system, 494–495
- virtual desktops**
 - comparing options, 484–485
 - monitoring performance, 587
 - network connectivity, 19–20
 - personal. *See* personal virtual desktops
 - pooled. *See* pooled virtual desktops
 - RDS connecting to, 392–395
 - selecting technology, 11
 - session-based. *See* session-based virtual desktops
 - terminology, 12
 - user experience, 17–18
 - VM-based. *See* VM-based virtual desktops
- virtual file systems, 160–161**
- virtual graphics processing unit (vGPU), 397–398**
- virtual hard disks (VHDs)**
 - compacting files, 149
 - converting, 149
 - creating, 147–148
 - described, 141
 - differencing disks, 145–147, 150, 152–153, 155
 - dynamical expanding, 143–145
 - editing, 148–150
 - expanding, 149
 - fixed-size, 143–145
 - formats supported, 100, 141–143
 - merging, 150
 - shrinking, 149
 - storage Quality of Service, 150–151
 - user profile disks, 523–524
- virtual machine bus (VMBus), 99**
- Virtual Machine Connection**
 - creating checkpoints, 152, 328–329
 - described, 105, 126–130
 - updating integration services, 140
- Virtual Machine Limit processor setting, 136**
- virtual machine networks, 489**
- Virtual Machine Reserve processor setting, 136**
- virtual machines (VMs)**
 - App-V and, 162
 - automatic creation, 525
 - checkpoint limitations, 155
 - Client Hyper-V and. *See* Client Hyper-V
 - configuration information, 113–114
 - controlling, 125–130
 - copying exported, 134
 - copying imported, 135
 - creating, 112–121
 - creating personal virtual desktops, 484
 - described, 1, 96
 - dynamic memory, 137–138
 - exporting, 132, 134
 - generation 1, 121–124, 398
 - generation 2, 124–125
 - importing, 133–135
 - integration services, 138–141
 - Live Migration, 487–488
 - managing files, 130–135
 - managing virtual hard disks, 141–151
 - memory considerations, 101, 137–138
 - migrating, 100
 - modifying settings, 121–124
 - moving files, 130–132
 - optimizing performance, 156
 - processing overview, 135–136
 - replicating, 100
 - as sequencers, 328–329
 - states supported, 126
 - storage considerations, 101, 253, 510
 - Sysprep tool, 503

- testing environments, 111
 - unique IDs, 133–134
 - virtual private networks (VPNs)**, 8, 388, 542
 - virtual processors**, 135–136
 - virtual registry**
 - App-V client support, 218, 275
 - App-V support, 160–161, 311
 - application sequencing, 347–348, 382
 - displaying keys, 347–348
 - registry.dat file, 187, 275, 284
 - troubleshooting, 275
 - Virtual Switch Manager**, 106
 - virtual switches**, 108–112
 - virtualization engine**, 160–161
 - Virtualization node (App-V policies)**, 273
 - virtualization technologies**
 - dual booting and, 97
 - identifying for business needs, 25–29
 - implementation considerations, 17–25
 - overview of, 1–12
 - usage scenarios for, 12–17
 - virtual machines. *See* virtual machines
 - VM-based VDI**, 10–12, 252
 - VM-based virtual desktops**. *See also* personal virtual desktops; pooled virtual desktops
 - deploying, 404, 522
 - described, 9–12
 - licensing considerations, 23
 - storage considerations, 510
 - user profile disks, 70, 523–524
 - virtual desktop collections, 393, 524–533
 - VMbus (virtual machine bus)**, 99
 - VMs (virtual machines)**
 - App-V and, 162
 - automatic creation, 525
 - checkpoint limitations, 155
 - Client Hyper-V and. *See* Client Hyper-V
 - configuration information, 113–114
 - controlling, 125–130
 - copying exported, 134
 - copying imported, 135
 - creating, 112–121
 - creating personal virtual desktops, 484
 - described, 1, 96
 - dynamic memory, 137–138
 - exporting, 132, 134
 - generation 1, 121–124, 398
 - generation 2, 124–125
 - importing, 133–135
 - integration services, 138–141
 - Live Migration, 487–488
 - managing files, 130–135
 - managing virtual hard disks, 141–151
 - memory considerations, 101, 137–138
 - migrating, 100
 - modifying settings, 121–124
 - moving files, 130–132
 - optimizing performance, 156
 - processing overview, 135–136
 - replicating, 100
 - as sequencers, 328–329
 - states supported, 126
 - storage considerations, 101, 253, 510
 - Sysprep tool, 503
 - testing environments, 111
 - unique IDs, 133–134
 - Volume Activation Services server role**, 492
 - Volume Activation Tools**, 492
 - Volume Shadow Copy Service (VSS)**, 498
 - VPNs (virtual private networks)**, 8, 16–17, 388, 542
 - VSS (Volume Shadow Copy Service)**, 498
- ## W
- WDDM (Windows Display Driver Model)**, 398
 - Web Server role**, 210
 - Wi-Fi connections, mobile devices and**, 65
 - wildcard certificates**, 279, 473
 - Windows 7 operating system**
 - App-V client, 213, 225, 229–230
 - App-V Sequencer, 182, 326
 - application sequencing, 326–327
 - Credential Roaming, 68
 - deleting user profiles, 33–34
 - integration services, 99
 - KMS support, 493
 - multi-monitor support, 396
 - Offline Files, 60–61
 - pooled and personal virtual desktops, 490–491
 - RADC feature, 6, 462
 - RemoteApp for Hyper-V, 539
 - roaming user profiles, 46
 - signing tool, 377
 - UE-V support, 68, 72
 - user state virtualization options, 75–76
 - vGPU support, 398
 - virtualization features, 490–491
 - VM-based virtual desktops, 404
 - Windows XP Mode, 97
 - Windows 8 operating system**
 - Active Directory-based activation, 493
 - App-V client, 213, 225, 229–230
 - application sequencing, 162, 182, 326–327
 - Client Hyper-V, 6, 100, 113, 125, 150
 - Group Policy settings, 87
 - modifying startup processes, 590
 - multi-monitor support, 396
 - network bridges, 111
 - Offline Files, 60–61, 64–66
 - optimizing configuration, 498
 - optimizing VM deployment, 503

- pooled and personal virtual desktops, 490–491
- Primary Computer setting, 66
- RemoteApp for Hyper-V, 539
- resource bottlenecks, 585
- SLAT support, 102
- Sysprep tool, 503
- UE-V and, 71, 83, 88
- user state virtualization options, 75–76
- vGPU support, 398
- VHDX format, 142
- virtualization features, 490–491
- Windows 8.1 operating system**
 - Active Directory-based activation, 493
 - App-V client, 162, 213, 225, 229–230
 - App-V Sequencer, 182
 - application sequencing, 327, 335, 344
 - application virtualization, 5
 - Client Hyper-V. *See* Client Hyper-V
 - evaluating compatibility considerations, 40
 - Group Policy support, 33, 60
 - GVLKs, 492
 - integration services, 140
 - managing checkpoints, 151, 155
 - metered networks, 65
 - Microsoft accounts and, 71, 76
 - mobile device connections, 65
 - modifying startup programs, 590
 - Offline Files, 64
 - optimizing configuration, 497–499
 - partitions, 495–496
 - RDS, 389, 396
 - RemoteFX, 397–398
 - shadowing connections, 472
 - Storage Spaces, 11
 - UE-V requirements, 71
 - user profile content, 34, 39
 - user state information, 31–32
 - user state virtualization, 44
 - virtual desktop templates, 491–492
 - virtualization features, 490–491
 - VM-based virtual desktops, 9, 404
 - Windows SIM, 495–496
- Windows Defender**, 327, 335–337
- Windows Display Driver Model (WDDM)**, 398
- Windows Event Log**, 255–256
- Windows Event Viewer**, 256–259, 555, 565–567
- Windows Features dialog box**, 102–103
- Windows Installer Service**, 456
- Windows Management Instrumentation (WMI)**, 75, 226
- Windows Media Player**, 336, 396
- Windows Network Load Balancing**, 439–440
- Windows PowerShell**. *See also* specific cmdlets
 - adding and publishing packages, 289–290
 - App-V client and, 215–216, 229, 263–264, 267–269, 315–317
 - application sequencing, 351
 - AppVClient module, 268–269, 294–296
 - AppVSequencer module, 294, 296, 351
 - AppVServer module, 294–295
 - bulk changes and, 68
 - Client Hyper-V settings, 103
 - configuring App-V client, 267–269
 - configuring UE-V settings, 77–78
 - creating and managing connection groups, 290–294
 - creating package accelerators, 366–367
 - creating packages from package accelerators, 377
 - deploying RDS, 419
 - deploying UE-V agent, 81
 - elevated prompts, 297
 - ESD model and, 191
 - Hyper-V module, 105
 - Integrated Scripting Environment, 106–107
 - managing checkpoints, 329–330
 - managing connections, 471
 - managing management servers, 289–297
 - managing UE-V agent, 83–84
 - RemoteApp programs, 461
 - removing applications, 305–306
 - UE-V requirements, 71
- Windows Registry**, 274–275
- Windows Search**, 327, 335–336, 498
- Windows Security dialog box**, 464
- Windows Server**
 - App-V, 188–189, 202, 207, 210
 - App-V client, 247
 - App-V Sequencer, 182
 - application sequencing, 327–328
 - application virtualization, 6, 162
 - Best Practices Analyzer, 415
 - Client Hyper-V, 95, 97, 99–101, 124–125, 329
 - Credential Roaming, 68
 - Data Deduplication feature, 510
 - described, 9
 - desktop virtualization, 7
 - failover clustering, 487, 489, 506
 - Folder Redirection, 57
 - high availability, 443, 486, 488
 - integration services, 99, 139–141
 - licensing requirements, 22, 401
 - load balancing, 439–440
 - management packs, 581–583
 - metered networks, 65
 - monitoring performance, 566, 568–571, 579, 588–590
 - NIC Teaming feature, 515
 - Offline Folders, 62
 - Primary Computer setting, 66, 68
 - RDS. *See* RDS (Remote Desktop Services)
 - Remote Desktop Easy Print, 470
 - RemoteApp, 5
 - resource bottlenecks, 585

- Server Message Block, 488, 507
- session-based desktops, 405
- shadowing connections, 472
- Storage Spaces, 2, 6, 9, 505
- storage technologies, 508–511
- storage virtualization, 11
 - UE-V, 71, 78
- use state virtualization, 44, 57, 62, 65
- user profile disks, 413
- user profiles, 48
- vGPU, 398
- virtual hard disks, 142, 147
- virtual machines, 113, 129, 135
 - Volume Activation Services server role, 492
- Windows Server Update Services (WSUS), 14, 141, 494**
- Windows Setup Screen, 496–497**
- Windows SIM (System Image Manager), 495–496**
- Windows Store apps, 70, 87**
- Windows Sysinternals tool suite, 571–572, 590**
- Windows System Configuration, 590**
- Windows System Image Manager, 495–496**
- Windows System Resource Manager (WSRM), 589**
- Windows Update, 327, 335–336**
- Windows Virtual Desktop Access (VDA), 23, 402**
- Windows Vista operating system**
 - integration services, 99
 - Offline Folders, 63
 - RemoteApp for Hyper-V, 539
 - user profiles, 48
- Windows XP operating system**
 - application compatibility, 5
 - Client Hyper-V, 96–97
 - Folder Redirection, 57
 - integration services, 99, 139
 - Offline Files, 61
 - Remote Desktop Easy Print, 470
 - RemoteApp for Hyper-V, 539
 - session collections, 430, 432
 - SSL support, 430
 - user profiles, 33–34, 46, 48
- wireless access points, 111**
- wireless network adapters, 100**
- WMI (Windows Management Instrumentation), 75, 226**
- write-back cache, 508–510**
- WSRM (Windows System Resource Manager), 589**
- WSUS (Windows Server Update Services), 14, 141, 494**

X

- xClipboard setting, 467**
- xDriveRedirection setting, 467**
- XenApp (Citrix), 253**
- XenDesktop (Citrix), 253**
- XML files**
 - connection groups, 305
 - described, 284
 - dynamic configuration and, 310
 - exporting package configuration to, 287, 307
- xPnPRedirection setting, 467**
- xPortRedirection setting, 467**
- xPrinterRedirection setting, 467**

Z

- .ZIP files, 275, 358**