Getting Started with XenDesktop® 7.x

Deliver desktops and applications to your end users, anywhere, anytime, with XenDesktop® 7.x

Craig Thomas Ellrod

[PACKT] PUBLISHING enterprise professional expertise distilled

www.allitebooks.com

Getting Started with XenDesktop® 7.x

Deliver desktops and applications to your end users, anywhere, anytime, with XenDesktop® 7.x

Craig Thomas Ellrod

[PACKT] PUBLISHING enterprise professional expertise distilled

BIRMINGHAM - MUMBAI

www.allitebooks.com

Getting Started with XenDesktop® 7.x

Copyright © 2014 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the author, nor Packt Publishing, and its dealers and distributors will be held liable for any damages caused or alleged to be caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

First Published: April 2014

Production Reference: 1150414

Published by Packt Publishing Ltd.
Livery Place
35 Livery Street
Birmingham B3 2PB, UK.

ISBN 978-1-84968-976-2

www.packtpub.com

Cover Image by Artie Ng (artherng@yahoo.com.au)

www.allitebooks.com

Credits

Author
Craig Thomas Ellrod

Reviewers
Jack Cobben
Lars Flaskager
Tom Franken
Govardhan Gunnala
Jan Hendrik Meier
Joseph Muniz
Peter Nap
Puthiyavan.Udayakumar
Florian Zoller

Acquisition Editor
Joanne Fitzpatrick

Content Development Editor
Sweny Sukumaran

Technical Editor
Mrunal Chavan

Copy Editors
Dipti Kapadia
Aditya Nair
Kirti Pai

Project Coordinator
Akash Poojary

Proofreaders
Maria Gould
Paul Hindle

Indexer
Hemangini Bari

Production Coordinator
Komal Ramchandani

Cover Work
Komal Ramchandani

www.allitebooks.com

Notice

The statements made and opinions expressed herein belong exclusively to the author/s and reviewer/s of this publication, and are not shared by or represent the viewpoint of Citrix Systems®, Inc. This publication does not constitute an endorsement of any product, service or point of view. Citrix® makes no representations, warranties or assurances of any kind, express or implied, as to the completeness, accuracy, reliability, suitability, availability or currency of the content contained in this publication or any material related to this publication. Any reliance you place on such content is strictly at your own risk. In no event shall Citrix®, its agents, officers, employees, licensees or affiliates be liable for any damages whatsoever (including, without limitation, damages for loss of profits, business information, loss of information) arising out of the information or statements contained in the publication, even if Citrix® has been advised of the possibility of such loss or damages.

Citrix®, Citrix Systems®, XenApp®, XenDesktop®, and CloudPortal™ are trademarks of Citrix Systems®, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries.

www.allitebooks.com

About the Author

Craig Thomas Ellrod has more than 25 years of experience in the computer industry and holds a Bachelor's degree in Computer Science from California State University, Chico, and a Master's degree in Business Administration from Pepperdine University. He has held many positions in the computer industry, including software programmer, support engineer, field and corporate systems engineer, technical marketing manager, product marketing manager, and product manager. He has worked for companies such as Celerity Computing, Emulex, Pinnacle Micro, Sync Research, Cisco Systems, Extreme Networks, and smaller startup ventures. Craig currently works for Citrix Systems® as a sales engineer and system architect in the Rockies region of USA. He has authored patent applications and designs and has received an award for innovation while working at Extreme Networks. Craig is a top blogger at Citrix®, has written many deployment guides, and is well-versed with all the Citrix® products. Craig has also written a book, *Technical Marketing*, *Stratequest*, and has authored a video series, *XenApp® 6.5*, *Packt Publishing*.

www.allitebooks.com

About the Reviewers

Jack Cobben is no stranger to the challenges that enterprises experience when managing large deployments of Windows systems and Citrix® implementations as he holds over 13 years of experience in systems management. Jack writes in his free time for his own blog, `www.jackcobben.nl`, and is active on the Citrix® support forums. He loves to test new software and share his knowledge in any way he can. You can follow him on Twitter via `@jackcobben`. While he works for Citrix®, Citrix® didn't help with or support this book in any way or form.

A great thanks to my wife and my twins for letting me have the time to review this book.

Lars Flaskager has a combined experience of 26 years in electronics and IT systems, with a successful track record of achieving first class results in IT design and implementation. His main focus and interests are in Citrix® products and solutions, and for more than 10 years, he has strived to be at the forefront when it comes to knowledge about Citrix® technology and how it can support businesses. Lars has worked for SimCorp for 12 years, where he gained all his knowledge about Citrix® solutions. He now works for Conecto, which is the only Citrix®-dedicated consultancy company in Denmark.

I would like to thank my former colleague, Torben Mæhle, and Citrix®, Denmark, for sharing their knowledge with me.

Tom Franken has 11 years of experience with virtualization technologies. He has built and manages XenDesktop® 5.5 environments, several vSphere systems, a Hyper-V cluster farm, and a vCloud implementation.

www.allitebooks.com

Govardhan Gunnala is a technical architect with a blend of cross-platform technologies, understanding and applying them to complex business requirements. He is a Microsoft- and Citrix®-certified professional specializing in server and application virtualization technologies. He is a skilled IT network security analyst and is highly regarded for sophisticated Perl and PowerShell scripting.

He has designed and delivered various cloud software solutions based on web, Citrix®, and VMware technologies. He maintains the delivered solutions along with their operational auditing, automation, and simplification. He is also responsible for the data center architecture and network security administration. He earlier worked as a senior systems engineer and as a member of the IT systems.

He is also a technical blogger and a corporate and institutional trainer with more than 8 years of experience in the IT software industry. You can follow his blog at `http://gunnalag.com/` and can get in touch with him on `http://www.linkedin.com/in/gunnalag`.

I would like to thank my intern students who have joined me to learn XenDesktop® and reiterated all the basic concerns and questions about the XenDesktop® technology from their perspective.

Jan Hendrik Meier has more than 10 years of experience in the IT industry. He started as a trainee for an IT specialist company. During this time, he had his first contact with products from Microsoft and Citrix®. Now, he is an expert for infrastructure and virtualization solutions. In the Citrix® area, he started work with an early XenDesktop® (then XenApp®) version—MetaFrame XP. He deepened his knowledge in Citrix products such as Presentation Server®, XenApp®, and XenDesktop®, and started to extend them with knowledge about various other Citrix® products such as Provisioning Services™, NetScaler®, and XenMobile®.

After staying for half a year in Australia, he picked up a job as a consultant in a mid-sized company, where he helped customers with his big stock of knowledge and a deep understanding of technical coherences.

Furthermore, he writes books and professional articles on different IT technologies. If he finds interesting problems at work, he writes their description and solutions for them on his blog at `http://www.jhmeier.de`.

I wish my new born daughter, Evi, an awesome and wonderful life. May all her wishes be fulfilled.

---

www.allitebooks.com

Joseph Muniz is a CSE at Cisco Systems and a security researcher as well. He started his career in software development, and later, managed networks as a contracted technical resource. Joseph moved into consulting and found a passion for security while meeting a variety of customers. He has been involved with the design and implementation of multiple projects ranging from Fortune 500 corporations to large federal networks.

Joseph runs the `TheSecurityBlogger.com` website, a popular resource for security and product implementation. You can also find Joseph speaking at live events, and he is involved with other publications too. Recent projects include being a speaker for Social Media Deception at the 2013 ASIS International conference; author of *Web Penetration Testing with Kali Linux*, *Packt Publishing*, September 2013; and an article on *Compromising Passwords* in *PenTest Magazine - Backtrack Compendium*, July 2013.

Outside of work, he can be found behind turntables scratching classic vinyls or on the soccer pitch, hacking away at the local club teams.

I would not have been able to contribute to this book without the support of my charming wife, Ning, and creative inspirations from my daughter, Raylin. I credit my passion for learning, to my brother, Alex, who has raised me along with my loving parents, Irene and Ray. I would like to give a final thank you to all my friends, family, and colleagues who have supported me over the years.

Peter Nap is an experienced Microsoft and Citrix® specialist with 14 years of experience. Mostly interested in server-based computing environments, his main areas of expertise are XenApp®, XenDesktop®, Microsoft Windows Server deployments, and the virtualization of applications, servers, and operating systems.

In his free time, he maintains his own website, `http://napplications.nl`, with free tools for ICT professionals because programming in C# is his passion. Currently, he is working for CGI as an infrastructure architect.

Peter Nap has also reviewed *Getting Started with XenApp 6.5*, *XenDesktop 5.6 Cookbook*, *XenDesktop 5 Starter*, and *Citrix® XenApp® 6.5 Expert Cookbook*, all of which have been published by Packt Publishing.

www.allitebooks.com

Puthiyavan.Udayakumar has more than 6 years of IT experience with an expertise in Citrix®, VMware, Microsoft products, and Apache CloudStack. He has extensive experience in designing and implementing virtualization solutions using various Citrix® products, VMware products, and Microsoft products. He is an IBM-certified solution architect and Citrix®-certified enterprise engineer, with more than 15 certifications in infrastructure products. He is the author of the book, *Getting Started with Citrix® CloudPortal™, Packt Publishing*. He holds a Master's degree in Science with a specialization in system software from Birla Institute of Technology and Science, Pilani, a Bachelor's degree in Engineering through SKR Engineering College from Anna University, and has received a national award from the Indian Society for Technical Education. He has presented various research papers at more than 15 national and international conferences including IADIS (held in Dublin, Ireland) followed by the IEEE pattern.

I would like to thank Packt Publishing for giving me the opportunity to review this book. This book is well-written by the author and the project is well-coordinated by the project coordinator.

Florian Zoller works as a lead IT architect for a consulting company based in Germany. He has several years of experience in designing and implementing Citrix® infrastructures for mid-sized and large deployments. Besides his expertise on XenApp®/XenDesktop®, XenMobile®, and NetScaler®, he focuses on software distribution and automation technologies such as FrontRange Desktop and Server Management. He is one of the few Immidio Valued Professionals (IVP).

www.allitebooks.com

www.PacktPub.com

Support files, eBooks, discount offers, and more

You might want to visit `www.PacktPub.com` for support files and downloads related to your book.

Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at `www.PacktPub.com` and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at `service@packtpub.com` for more details.

At `www.PacktPub.com`, you can also read a collection of free technical articles, sign up for a range of free newsletters and receive exclusive discounts and offers on Packt books and eBooks.

`http://PacktLib.PacktPub.com`

Do you need instant solutions to your IT questions? PacktLib is Packt's online digital book library. Here, you can access, read and search across Packt's entire library of books.

Why subscribe?

- Fully searchable across every book published by Packt
- Copy and paste, print and bookmark content
- On demand and accessible via web browser

Free access for Packt account holders

If you have an account with Packt at `www.PacktPub.com`, you can use this to access PacktLib today and view nine entirely free books. Simply use your login credentials for immediate access.

Instant updates on New Packt books

Get notified! Find out when new books are published by following `@PacktEnterprise` on Twitter, or the *Packt Enterprise* Facebook page.

Table of Contents

Preface 1

Chapter 1: Designing a XenDesktop® Site 11

The core components of a XenDesktop® Site 11

Terminology and concepts 13

Server side 13

Hypervisor 13

Database 13

Delivery Controller 13

Studio 14

Director 14

StoreFront 14

Virtual machines 14

The Virtual Desktop Agent 14

Server OS machines 14

Desktop OS machines 15

Active Directory 15

Desktop 15

XenApp® 15

Edgesight® 15

FlexCast® 16

Storage 17

The client side 17

Receiver 17

System requirements 18

Receiver 18

StoreFront 2.1 18

Databases 19

Studio 20

Delivery Controller 20

Director 21

The Virtual Delivery Agent (VDA) 21

Server host 22
Active Directory 23
Designing a basic XenDesktop® Site 23
Scenario 23
Common Citrix® communication ports 24
Summary 27
Chapter 2: Installing XenDesktop® 29
Planning the XenDesktop® installation 30
Step 1 – installing the controller (XD1) 32
Installing the components on XD1 33
Configuring a Site 38
Step 2 – installing StoreFront (XD2) 41
Installing the components on XD2 41
Creating a server certificate and adding a Site binding 41
Installing StoreFront 46
Step 3 – installing Director (XD3) 47
Installing the components on XD3 47
Step 4 – creating the virtual desktop and application delivery master images 48
Step 5 – installing the Virtual Delivery Agent on the master images 49
Creating the desktop master images 50
Creating the application master images 53
Step 6 – configuring the StoreFront server 56
Step 7 – creating the machine catalogs 58
Creating desktops 58
Creating the application servers 64
Step 8 – creating the delivery groups 67
Creating desktop delivery groups 68
Creating the application delivery groups 69
Installation checkpoint 72
Step 9 – installing Citrix Receiver™ on the client devices 75
Step 10 – testing the connection 77
Testing the desktops 77
Testing the applications 77
Summary 79
Chapter 3: Managing Machine Catalogs, Hosts, and Personal vDisks 81
Machine catalogs 81
Prerequisites 82
Creating the master images 82

Adding and configuring the virtual machines 83
Creating the computer accounts 83
Creating a machine catalog 84
Operating systems and hardware 84
Machine management 86
User experience 87
Managing the machine catalogs 88
Taking a snapshot of the master image 89
Updating the master image 90
Reverting to a previous master image 92
Managing the Active Directory computer accounts 93
Adding machines to a machine catalog 94
Modifying a machine catalog 95
Renaming a machine catalog 95
Deleting a machine catalog 96
Managing the hosts 96
Managing Personal vDisks 100
Updating Personal vDisks used by the master images 102
Adjusting the space available for applications 103
Disabling automatic resizing 104
Reallocating user profiles 104
Summary 104
Chapter 4: Managing Delivery Groups 105
Managing the delivery groups 105
Creating a delivery group 106
Editing a delivery group 107
Managing desktop sessions 109
Logging off or disconnecting sessions 109
Sending messages to users 109
Managing the delivery group resources 110
Adding and reallocating desktops 110
Locating desktops, sessions, and delivery groups 111
Shutting down and restarting desktops 112
Removing desktops from delivery groups 113
Deleting desktops from delivery groups 114
Restricting access to desktops 114
Securing the ICA® protocol communications 116
Managing power settings for desktops 117
Importing and exporting user data 118
Enabling and disabling the maintenance mode 119
Managing the server load 120
Managing the hosted applications 122
Application desktop delivery groups 123
Application sharing 123
Publishing applications to multiple desktop groups 124
Content redirection 124

Creating an application 125
Managing application sessions 127
Modifying the applications 127
Managing the Delivery Controller environment 129
Controller discovery 129
Adding, moving, or removing Delivery Controllers 132
Moving a Virtual Delivery Agent (VDA) to another Site 134
Active Directory OU-based controller discovery 134
Using SSL on controllers 136
Changing the default HTTP and HTTPS ports 136
Summary 137
Chapter 5: Managing Policies 139
XenDesktop® Studio versus Microsoft Group Policy Editor 140
Administrative roles 140
Working with policies 141
Navigating policies 141
Accessing policies 142
Searching policies 143
Creating policies 143
Creating a policy in Studio 144
Creating a policy in Microsoft Group Policy Editor 144
Configuring policies 146
Configuring policy settings 146
Best practices for designing policy settings 147
Applying policies 147
Using default values 147
Using filters 148
Implementing multiple policies 151
Implementing priorities 151
Implementing exceptions 153
The resulting set of policies 154
Running the Citrix® Group Policy Modeling Wizard 154
Running the Microsoft Group Policy Results tool 155
Troubleshooting policy scenarios 156
Comparing policies 157
Implementing policies with NetScaler Gateway™ 158
Implementing NetScaler Gateway™ policy filters 158
Summary 160
Chapter 6: Managing Printing 161
How printing works 161
Using locally attached printers 162
Using network attached printers 162

Using default printing, preferences, and drivers 163
Setting printing preferences 164
Printing policies 165
Universal Print Server and Driver 166
Autocreation of printers 169
Mapping printers and drivers 171
Optimization of printing 174
Summary 179
Chapter 7: Virtualizing USB Support 181
USB devices in virtualization 181
How XenDesktop® uses USB redirection 183
Enabling USB support 184
Preventing the mapping of USB devices 187
Using USB mass storage 187
USB redirection with XenApp® versus XenDesktop® 189
Using USB automatic redirection 189
Using voice and video 189
Summary 191
Chapter 8: Virtualizing Storage and Backup 193
XenDesktop® storage considerations 194
Desktop storage 194
High Availability 195
Performance 196
IOPS 196
Personal vDisk 196
XenDesktop® storage requirements 198
Virtual desktop storage requirements – dedicated desktop model 199
Virtual desktop storage requirements – dedicated shared desktop model 201
Virtual desktop storage requirements – shared hosted desktop model 203
Backup and restore 203
Backing up a SQL Server 203
Restoring a SQL Server 205
Backing up and restoring VMs and user data 206
USB mass storage 207
Summary 207

Chapter 9: High Definition Experience (HDX™) 209
Introducing high definition experience 210
HDX™ system requirements 210
The reality of HDX™ 212
Aero redirection 213
Configuring Aero redirection or desktop composition redirection 213
Windows Media 215
Configuring Windows Media client-side fetching 215
Configuring real-time Windows Media multimedia transcoding 217
Flash Media 219
Configuring Flash redirection on a server 220
Configuring Flash redirection on the client 221
HDX™ 3D 223
GPU versus vGPU 224
GPU 224
vGPU 225
HDX™ 3D requirements 225
Client 225
Server 225
HDX™ GPU sharing 226
HDX™ 3D – how it works 226
Installing and configuring HDX™ 3D 227
Upgrading HDX™ 3D 229
Configuring monitors for HDX™ 3D 229
Configuring image quality 230
Configuring audio 230
Configuring webcams 231
Configuring color compression 231
Configuring network priorities 232
Adaptive display 233
Summary 234
Chapter 10: Application Delivery 235
Delivering applications 236
Differences between XenApp® and XenDesktop® 236
What's new? 236
What's gone? 237
What's changed? 238
What hasn't changed? 238
Application Delivery Controllers 238
Application Delivery Networks 240
Summary 242

Chapter 11: Working with the XenDesktop® SDK 243
Microsoft Windows PowerShell 244
PowerShell snap-ins and cmdlets for XenDesktop® 244
Using the XenDesktop® SDK 247
Creating an SDK script 248
Troubleshooting using the XD PowerShell SDK 249
Useful desktop cmdlets 249
Useful controller cmdlets 250
Site debugging tools 252
Citrix Ready® 252
Summary 252
Chapter 12: Working with Citrix Receiver™ and Plugins 253
Understanding Receiver 254
Changing the Receiver settings 255
Pushing the Receiver settings from the server 255
Changing the Receiver settings from the client's desktop 256
Using plugins 258
The online plugin 258
Using workspace control 258
Changing the resolution of the virtual desktop 259
Moving the toolbar 259
Controlling local file access 260
Accessing devices 261
Accessing USB devices 261
Accessing local microphones and webcams 262
Redirecting Flash to a local device 263
Switching between virtual desktops 264
Logging off virtual desktops 265
Disconnecting from virtual desktops 265
Restarting a virtual desktop 266
Using Desktop Lock 267
Printing in virtual desktops 268
Understanding the keyboard input 268
The offline plugin 269
The CloudBridge™ plugin 269
Running Receiver on Microsoft Windows 270
Running Receiver on Apple 270
Running Receiver on other devices 271
Summary 272
Chapter 13: Securing XenDesktop® 273
DMZ and DMZ² 274
Securing XenDesktop® with NetScaler Gateway™ 275
Importing NetScaler VPX™ into XenServer® 276

Installing a NetScaler® license 276
Installing an SSL certificate 279
Creating a NetScaler Gateway™ virtual server 279
Configuring NetScaler Gateway™ for StoreFront 284
Configuring NetScaler® for an ICA proxy 286
Configuring a StoreFront connection to NetScaler Gateway™ 288
Exporting the StoreFront certificate 291
Importing the StoreFront certificate into NetScaler Gateway™ 294
Secure Ticket Authority 297
Securing the ICA/HDX protocols 297
Securing StoreFront 298
Securing Receiver 299
Securing controller 299
IIS 299
Non-IIS 299
Changing the controller port to HTTPS 300
Securing Studio and Director 300
IIS 300
Securing the XenDesktop® to XenServer® communications 300
Using smart cards 302
Summary 302
Chapter 14: Managing and Monitoring XenDesktop® 303
Using Studio to manage the XenDesktop® Site 304
Using Director to monitor the XenDesktop® Site 305
Using HDX Insight™ 310
Troubleshooting XenDesktop® 315
Troubleshooting users 316
Troubleshooting applications 316
Troubleshooting desktops 317
Troubleshooting sessions 317
Troubleshooting HDX™ 317
Troubleshooting Personal vDisks 318
Third-party tools 318
Summary 319
Chapter 15: VDI in the Cloud 321
Understanding virtualization in the cloud 321
Private cloud 322
Public cloud 323

Hybrid cloud 323
Personal cloud 324
Your cloud 325
Summary 325
Appendix A: Creating a Domain Certificate Authority 327
Appendix B: XenDesktop® Policy Settings Reference 331
Audio policies 333
Bandwidth policies 334
Redirection policies 335
Desktop UI policies 338
Graphics and multimedia policies 338
Caching policies 342
Multistream traffic policies 342
Printing policies 342
ICA® policies 345
Keep alive policies 346
Autoreconnection policies 346
Mobility policies 347
Session policies 347
Time zone policies 349
Load management policies 349
Delivery Agent policies 350
HDX™ 3D policies 351
Appendix C: Creating Self-signed Certificates for NetScaler Gateway™ 353
Enabling SSL on NetScaler Gateway™ 353
Creating a self-signed root CA certificate 354
Creating a public-facing server certificate 357
Installing the root CA and public certificates 359
Linking the public and root CA certificates 361
Viewing the root CA and server certificate bindings 362
Binding the certificates to the NetScaler Gateway™ VIP 362
Testing the certificates 364
Testing the NetScaler Gateway™ connection 365
Testing NetScaler Gateway™ with a Windows client 365
Appendix D: Using Public CA-signed SSL Wildcard Certificates on NetScaler Gateway™ 373
Enabling SSL on NetScaler Gateway™ 374
Creating a certificate request 374

www.allitebooks.com

Submitting the request to the public CA **378**
Installing the public-signed wildcard certificate **380**
Binding the public-signed certificate to the NetScaler Gateway™ VIP **382**
Testing NetScaler Gateway™ and certificates **383**
Index **385**

Preface

Citrix® XenDesktop® is a desktop virtualization and VDI solution that delivers a Windows desktop experience as an on-demand service to any user, anytime, anywhere. It suits all types of workers such as task workers, knowledge workers, or mobile workshifting workers. XenDesktop® quickly and securely delivers complete desktops or applications while providing a high-definition user experience.

XenDesktop® is a desktop virtualization solution that optimizes the delivery of desktops, applications, and data to end users. It includes all of the capabilities to deliver desktops, applications, and data securely to every type of user in an enterprise. Instead of managing thousands of static desktop images, you can manage and update the desktop OS and applications once, from one location.

Getting Started with XenDesktop® 7.x provides comprehensive details on how to design, implement, and maintain a desktop delivery Site using XenDesktop®. Along the way, you will also learn about management, policies, printing, USB support, storage and backup, High Definition User Experience (HDX™), application delivery, the XenDesktop® SDK, Citrix Receiver™, and about running XenDesktop® from the cloud.

If you are reading this book, you have most likely heard of the concept of desktop virtualization. You may have done some basic research on the topic or have installed a previous version of XenDesktop®. In any case, XenDesktop® 7 is different from the previous versions. So, if you are a desktop virtualization veteran or are new to the game and starting your Proof of Concept, this book will be helpful. In this book, we will walk you through the implementation of Citrix® XenDesktop® for a small deployment to help you understand not only how to install the product, but also how the desktop and application technology works.

Getting started with Hypervisors

Before you get started, you need to understand what a Hypervisor is. A Hypervisor is an operating system that hosts multiple instances of disparate operating systems. It can also be defined as a software that can create and run virtual machines. The Hypervisor software runs on a server hardware that has been enabled for virtualization. Once this is installed, you can then install several instances of different operating systems onto the Hypervisor. The Hypervisor was the game changer because instead of running one operating system per server, you could now run X number of operating systems on one server, and thus save space and money.

There are several vendors that make Hypervisors, such as Citrix® XenServer®, VMware ESX, Microsoft Hyper-V, and KVM. There are Type 1 Hypervisors that run directly on the server hardware; these are also known as bare-metal Hypervisors. There are Type 2 Hypervisors that run on top of an operating system, which then runs on the server. As you can imagine, Type 1 Hypervisors have been touted to have better performance as they interact directly with the server hardware resources.

Citrix® XenServer® is a Type 1 Hypervisor. Citrix® XenDesktop® runs on Citrix® XenServer®. It can also run on VMware ESX and Microsoft Hyper-V. This book will focus on the use of XenDesktop® running on XenServer®.

The following diagram gives you a visual idea of the differences between the types of Hypervisors as compared to traditional servers and how the interaction between these components contend for hardware resources, which ultimately affects the performance and sizing of hardware resources:

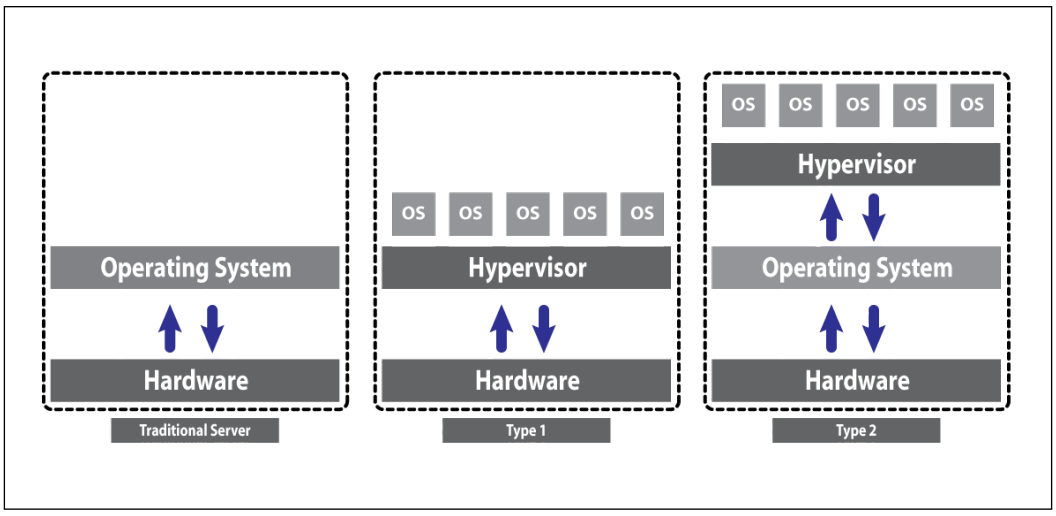

What this book covers

Chapter 1, *Designing a XenDesktop® Site*, starts by defining the pieces or components that make up a XenDesktop® Site along with the terminology and concepts involved. We then set out to design a basic XenDesktop® architecture, ending with a network diagram that we will use as a roadmap for the remainder of the book.

Chapter 2, *Installing XenDesktop®*, explains the installation of XenDesktop® as you now know what it looks like via a network diagram and what it sounds like from the components, terminology, and concepts learned. This chapter discusses how to use the plan that is built in the previous chapter and then execute the plan to start deploying the XenDesktop® Site.

Chapter 3, *Managing Machine Catalogs, Hosts, and Personal vDisks*, discusses how to use machine catalogs, hosts, and Personal vDisks for XenDesktop®. After you create a XenDesktop® Site with the initial desktops and applications, you may want to expand the Site. Machine catalogs contain a group of computers or desktops that define the hosting infrastructure for desktops and applications.

Chapter 4, *Managing Delivery Groups*, discusses in detail how to manage delivery groups for desktops and applications. Delivery groups are collections of machines that deliver desktops and applications to users.

Chapter 5, *Managing Policies*, explains that Citrix® policies are the best way to control connections, security, and other settings in XenDesktop®. Everything is done with policies, at least when it comes to giving users access and managing sessions.

Chapter 6, *Managing Printing*, explains that printing in XenDesktop® is handled the same way it is handled in XenApp. You can print using printers that are connected locally or networked; so, we discuss how to do this. We also talk about the installed printer drivers and controlling printers with policies.

Chapter 7, *Virtualizing USB Support*, discusses how USB support allows virtual desktops to access the local USB resources connected to the user/client device. XenDesktop® also provides direct connectivity support for some devices, such as keyboards, mice, and smart cards. Think about it; if you use a virtual desktop, you won't have a physical USB port to plug in to on that virtual machine, so we have to use the USB port on our client device and somehow map this to the virtual desktop.

Chapter 8, *Virtualizing Storage and Backup*, discusses the storage and backup requirements for XenDesktop®. You need storage for the XenDesktop® Site and the individual virtual desktops. A virtual desktop deployment is very dynamic, and the storage infrastructure needs to be able to accommodate it.

Chapter 9, *High Definition Experience (HDX™)*, explains that delivering HDX™ to any device, anywhere, has some requirements, especially with regard to which end is doing the processing, the server or client. The high definition experience is a broad set of technologies that provide a high-definition user experience to any device.

Chapter 10, *Application Delivery*, discusses application delivery in the context of XenDesktop®. You have your virtual desktop, so where are the applications? Applications are delivered from XenDesktop® by a VM running the app called a VM hosted app.

Chapter 11, *Working with the XenDesktop® SDK*, talks about the XenDesktop® SDK and how to use it. The XenDesktop® SDK is based on PowerShell Version 3.0 snap-ins and is a powerful tool for third-party vendors who wish to integrate their products with XenDesktop®. Later in the chapter, we will look at how to identify third-party vendors who have been certified to work with Citrix® XenDesktop® through the Citrix Ready® program.

Chapter 12, *Working with Citrix Receiver® and Plugins*, talks about the client side of the equation, specifically using Citrix® Receiver to receive and run the virtual desktop on the client device. Citrix® Receiver is device agnostic, so we discuss Receiver for the many different platforms that a client might use, including thin clients and mobile devices. Citrix® also uses plugins that plug in to Receiver, so we address these briefly.

Chapter 13, *Securing XenDesktop®*, explains that XenDesktop® is not secure by itself, but you can make it secure by following some simple guidelines. XenDesktop® and XenApp® have, for a long time, had a feature called the Secure Ticket Authority (STA); however, this doesn't provide complete security. In this chapter, we discuss how to secure XenDesktop® with SSL.

Chapter 14, *Managing and Monitoring XenDesktop®*, discusses monitoring XenDesktop® using Director and other tools. What is seemingly an afterthought is actually very important. If you can't see it, you can't manage it. In this chapter, we discuss how to manage a XenDesktop® Site. XenDesktop® Director is a web-based tool that enables the IT and support teams to monitor a XenDesktop® environment and perform troubleshooting.

Chapter 15, *VDI in the Cloud*, explains that since you can now deliver desktops and applications from anywhere to any device, where are you going to deliver them from? When we talk about the cloud, it means XenDesktop® can be installed in your data center (private cloud), or a hosting service provider (public cloud), or a combination of both (hybrid). We will look at the advantages and disadvantages of each.

Appendix A, *Creating a Domain Certificate Authority*, walks you through creating a domain certificate authority in Microsoft Windows Server, which you can use in the book's examples and in your own deployment.

Appendix B, XenDesktop® Policy Settings Reference, lists all of the policies and potential settings for use with XenDesktop®. Everything in XenDesktop® is done through policies and there are a lot of them.

Appendix C, Creating Self-signed Certificates for NetScaler Gateway™, walks you through the creation of a NetScaler® self-signed Certificate Authority (CA) certificate and a NetScaler® self-signed server certificate. This is perfect for getting started and for Proof of Concepts because it doesn't cost you anything to get SSL configured and running on NetScaler®.

Appendix D, Using Public CA-signed SSL Wildcard Certificates on NetScaler Gateway™, walks you through the process of obtaining a valid public Certificate Authority (CA) signed server certificate for use on NetScaler®. You only need the server certificate because the CA certificate is already populated in all of the browsers out there. There is a cost associated with obtaining a server certificate.

What you need for this book

The following are the software requirements for this book:

- Microsoft Windows Server 2012 R2
- Citrix® XenDesktop® 7.x
- A Hypervisor (Citrix® XenServer® 6.x.x, VMware vSphere (ESX 5.x), and Microsoft System Center Virtual Machine Manager 2012 Rollup 1)

The following are the license requirements for this book:

- Microsoft Windows Server 2012 R2
- Microsoft Windows 8
- Microsoft Terminal Services
- Citrix® XenDesktop®

The following are the hardware requirements for this book:

- Hypervisor host server
- Network infrastructure
- Client devices

The following are the Citrix XenDesktop components used in this book:

- Microsoft Active Directory
- XenDesktop® server 1
    - Delivery Controller
    - Studio
    - License server
    - Microsoft SQL Server
- XenDesktop® server 2
    - StoreFront
- XenDesktop® server 3
    - Director
- XenDesktop® server 4
    - Application server
- NetScaler Gateway™
    - StoreFront frontend
    - CloudBridge™
    - CloudBridge™ connector
    - CloudBridge™ WAN optimization

Who this book is for

If you are a system administrator, consultant, or beginner who wants to implement and administer Citrix® XenDesktop® Sites, then this book is for you. This book will help both new and experienced XenDesktop® professionals to deliver desktops and applications using the new version of XenDesktop® to any user on any device, anywhere, any time.

This book is based on XenDesktop® 7.x, which is a major architectural change from XenDesktop® 5.6 and XenApp® 6.5.

Conventions

In this book, you will find a number of styles of text that distinguish between different kinds of information. Here are some examples of these styles, and an explanation of their meaning.

Code words in text, database table names, folder names, filenames, file extensions, pathnames, dummy URLs, user input, and Twitter handles are shown as follows: "Type `Desktop` or the name of the desktop group you created".

A block of code is set as follows:

```
$grp = Get-XdDesktopGroup 'example'
C:\PS&gt;$grp.Desktops.Add( New-XdVirtualDesktop machine4 )
C:\PS&gt;Set-XdDesktopGroup $grp
```

When we wish to draw your attention to a particular block of code, the relevant lines or items will be in bold print as follows:

```
$grp = Get-XdDesktopGroup 'example'
C:\PS&gt;$grp.Desktops.Add( New-XdVirtualDesktop machine4 )
C:\PS&gt;Set-XdDesktopGroup $grp
```

Any command-line input or output is written as follows:

```
PS C:\Windows\system32> Set-ExecutionPolicy Unrestricted
```

New terms and **important words** are shown in bold. Words that you see on the screen, in menus or dialog boxes for example, appear in the text like this: "Click on the **Servers** link and select your XenServer version".

Warnings or important notes appear in a box like this.

Tips and tricks appear like this.

Reader feedback

Feedback from our readers is always welcome. Let us know what you think about this book — what you liked or may have disliked. Reader feedback is important for us to develop titles that you really get the most out of.

To send us general feedback, simply send an e-mail to `feedback@packtpub.com`, and mention the book title via the subject of your message.

If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, see our author guide on `www.packtpub.com/authors`.

Customer support

Now that you are the proud owner of a Packt book, we have a number of things to help you to get the most from your purchase.

Errata

Although we have taken every care to ensure the accuracy of our content, mistakes do happen. If you find a mistake in one of our books — maybe a mistake in the text or the code — we would be grateful if you would report this to us. By doing so, you can save other readers from frustration and help us improve subsequent versions of this book. If you find any errata, please report them by visiting `http://www.packtpub.com/support`, selecting your book, clicking on the **errata submission form** link, and entering the details of your errata. Once your errata are verified, your submission will be accepted and the errata will be uploaded on our website, or added to any list of existing errata, under the Errata section of that title. Any existing errata can be viewed by selecting your title from `http://www.packtpub.com/support`.

Piracy

Piracy of copyright material on the Internet is an ongoing problem across all media. At Packt, we take the protection of our copyright and licenses very seriously. If you come across any illegal copies of our works, in any form, on the Internet, please provide us with the location address or website name immediately so that we can pursue a remedy.

Please contact us at `copyright@packtpub.com` with a link to the suspected pirated material.

We appreciate your help in protecting our authors, and our ability to bring you valuable content.

Questions

You can contact us at `questions@packtpub.com` if you are having a problem with any aspect of the book, and we will do our best to address it.

www.allitebooks.com

1
Designing a XenDesktop® Site

In this chapter, we start with defining the pieces or components that make up a XenDesktop Site along with the terminology and concepts involved. We then set out to design a basic XenDesktop architecture, which ends with a network diagram that we will use as a roadmap for the remainder of the book. In this chapter, we will cover the following topics:

- The components of XenDesktop
- Terminology and concepts
- System requirements
- Designing a basic XenDesktop Site
- Common ports used in network communication

The core components of a XenDesktop® Site

Before we get started with the designing of the XenDesktop Site, we need to understand the core components that go into building it. XenDesktop can support all types of workers—from task workers who run Microsoft Office applications to knowledge users who host business applications, to mobile workshifting users, and to high-end 3D application users. It scales from small businesses that support five to ten users to large enterprises that support thousands of users.

Please follow the steps in the guide in the order in which they are presented; do not skip steps or topics for a successful implementation of XenDesktop.

The following is a simple diagram to illustrate the components that make up the XenDesktop architecture:

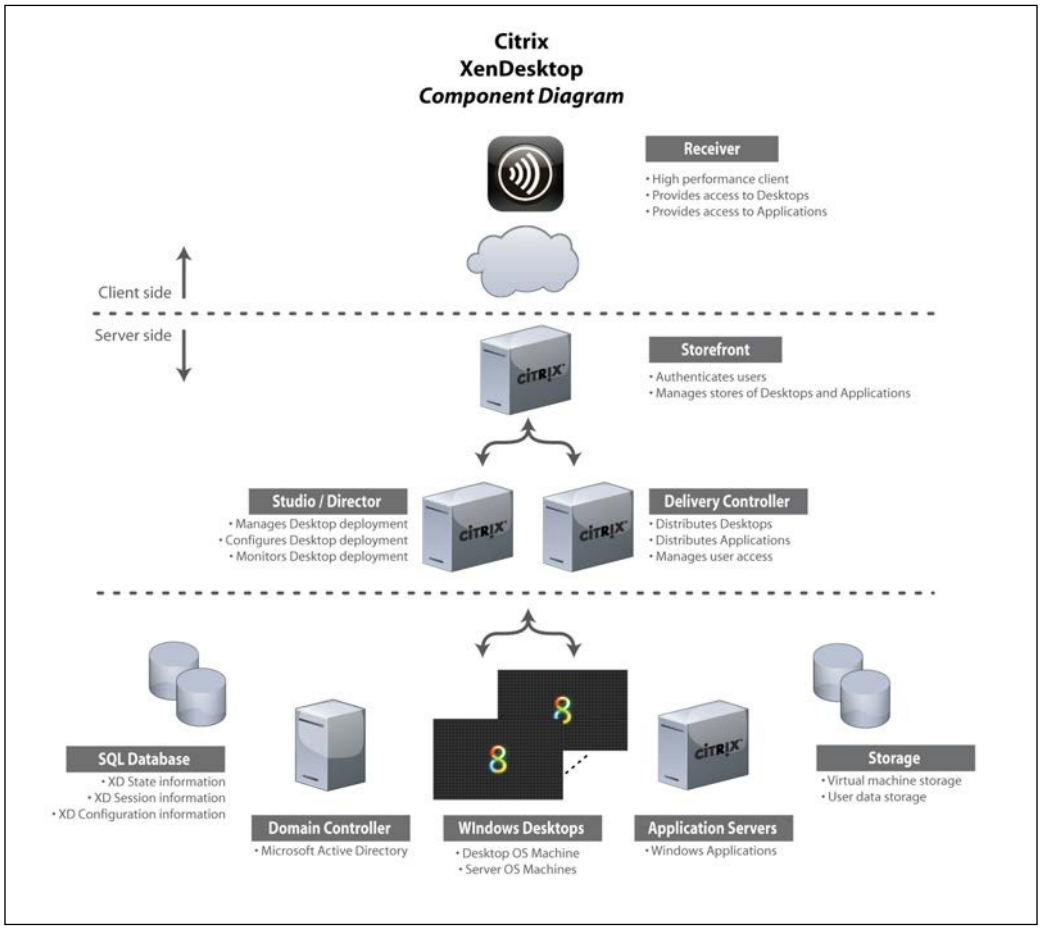

If you have the experience of using XenDesktop and XenApp, you will be pleased to learn that XenDesktop and XenApp now share management and delivery components to give you a unified management experience.

Now that you have a visual of how a simple Site will look when it is completed, let's take a look at each individual component so that you can understand their roles.

Terminology and concepts

In this section, we will cover some commonly used terminology and concepts used with XenDesktop.

Server side

It is important to understand the terminology and concepts as they apply to the server side of the XenDesktop architecture, so we will cover that in this section.

Hypervisor

As mentioned in the *Preface* of this book, a Hypervisor is an operating system that hosts multiple instances of other operating systems. XenDesktop is supported by three Hypervisors—Citrix XenServer, VMware ESX, and Microsoft Hyper-V.

Database

In XenDesktop, we use the Microsoft SQL Server. The database is sometimes referred to as the data store. Almost everything in XenDesktop is database driven, and the SQL database holds all state information in addition to the session and configuration information. The XenDesktop Site is only available if the database is available.

If the database server fails, existing connections to virtual desktops will continue to function until the user either logs off or disconnects from their virtual desktop; new connections cannot be established if the database server is unavailable. There is no caching in XenDesktop 7.x, so Citrix recommends that you implement SQL mirroring and clustering for High Availability.

The IMA data store is no longer used, and everything is now done in the SQL database for both session and configuration information. The data collector is shared evenly across XenDesktop controllers.

Delivery Controller

The Delivery Controller distributes desktops and applications, manages user access, and optimizes connections to applications. Each Site has one or more Delivery Controllers.

Studio

Studio is the management console that enables you to configure and manage your XenDesktop and XenApp deployment, eliminating the need for two separate management consoles to manage the delivery of desktops and applications. Studio provides you with various wizards to guide you through the process of setting up your environment, creating your workloads to host and assign applications and desktops, and assigning applications and desktops to users.

Citrix Studio replaces the Delivery Services Console and the Citrix AppCenter from previous XenDesktop versions.

Director

Director is used to monitor and troubleshoot the XenDesktop deployment.

StoreFront

StoreFront authenticates users to Site(s) hosting the XenApp and XenDesktop resources and manages the stores of desktops and applications that users access.

Virtual machines

A **virtual machine** (**VM**) is a software-implemented version of the hardware. For example, Windows Server 2012 R2 is installed as a virtual machine running in XenServer. In fact, every server and desktop in this book's examples will be installed as a VM with the exception of the Hypervisor, which obviously needs to be installed on the server hardware before we can install any VMs.

The Virtual Desktop Agent

The **Virtual Desktop Agent** (**VDA**) has to be installed on the VM to which users will connect. It enables the machines to register with controllers and manages the ICA/HDX connection between the machines and the user devices. The VDA is installed on the desktop operating system VM, such as Windows 7 or Windows 8, which is served to the client. The VDA maintains a heartbeat with the Delivery Controller, updates policies, and registers the controllers with the Delivery Controller.

Server OS machines

VMs or physical machines based on the Windows Server operating system are used to deliver applications or host shared desktops to users.

Desktop OS machines

VMs or physical machines based on the Windows desktop operating system are used to deliver personalized desktops to users or applications from desktop operating systems.

Active Directory

Microsoft Active Directory is required for authentication and authorization. Active Directory can also be used for controller discovery by desktops to discover the controllers within a Site. Desktops determine which controllers are available by referring to information that controllers publish in Active Directory.

Active Directory's built-in security infrastructure is used by desktops to verify whether communication between controllers comes from authorized controllers in the appropriate Site. Active Directory's security infrastructure also ensures that the data exchanged between desktops and controllers is confidential.

Installing XenDesktop or SQL Server on the domain controller is not supported; in fact, it is not even possible.

Desktop

A desktop is the instantiation of a complete Windows operating system, typically Windows 7 or Windows 8. In XenDesktop, we install the Windows 7 or Windows 8 desktop in a VM and add the VDA to it so that it can work with XenDesktop and can be delivered to clients. This will be the end user's virtual desktop.

XenApp®

Citrix XenApp is an on-demand application delivery solution that enables any Windows application to be virtualized, centralized, and managed in the data center and instantly delivered as a service. Prior to XenDesktop 7.x, XenApp delivered applications and XenDesktop delivered desktops. Now, with the release of XenDesktop 7.x, XenApp delivers both desktops and applications.

Edgesight®

Citrix Edgesight is a performance and availability management solution for XenDesktop, XenApp, and endpoint systems. Edgesight monitors applications, devices, sessions, license usage, and the network in real time. Edgesight will be phased out as a product.

FlexCast®

Don't let the term FlexCast confuse you. FlexCast is just a marketing term designed to encompass all of the different architectures that XenDesktop can be deployed in. FlexCast allows you to deliver virtual desktops and applications according to the needs of diverse performance, security, and flexibility requirements of every type of user in your organization. FlexCast is a way of describing the different ways to deploy XenDesktop. For example, task workers who use low-end thin clients in remote offices will use a different FlexCast model than a group of HDX 3D high-end graphics users. The following table lists the FlexCast models you may want to consider; these are available at `http://flexcast.citrix.com`:

FlexCast model	Use case	Citrix products used
Local VM	Local VM desktops extend the benefit of a centralized, single-instance management to mobile workers who need to use their laptops offline. Changes to the OS, apps, and data are synchronized when they connect to the network.	XenClient
Streamed VHD	Streamed VHDs leverage the local processing power of rich clients, which provides a centralized, single-image management of the desktop. It is an easy, low-cost way to get started with desktop virtualization (rarely used).	Receiver XenApp
Hosted VDI	Hosted VDI desktops offer a personalized Windows desktop experience typically required by office workers, which can be delivered to any device. This combines the central management of the desktop with complete user personalization. The user's desktop runs in a virtual machine. Users get the same high-definition experience that they had with a local PC but with a centralized management. The VDI approach provides the best combination of security and customization. Personalization is stored in the Personal vDisk. VDI desktops can be accessed from any device, such as thin clients, laptops, PCs, and mobile devices (most common).	Receiver XenDesktop Personal vDisk

FlexCast model	Use case	Citrix products used
Hosted shared	Hosted shared desktops provide a locked-down, streamlined, and standardized environment with a core set of applications. This is ideal for task workers where personalization is not required. All the users share a single desktop image. These desktops cannot be modified, except by the IT personnel. It is not appropriate for mobile workers or workers who need personalization, but it is appropriate for task workers who use thin clients.	Receiver XenDesktop
On-demand applications	This allows any Windows application to be centralized and managed in the data center, which is hosted on either multiuser terminal servers or virtual machines, and delivered as a service to physical and virtual desktops.	Receiver XenApp and XenDesktop App Edition

Storage

All of the XenDesktop components use storage. Storage is managed by the Hypervisor, such as Citrix XenServer. There is a personalization feature to store personal data from virtual desktops called the **Personal vDisk** (**PvD**).

The client side

For a complete end-to-end solution, an important part of the architecture that needs to be mentioned is the end user device or client. There isn't much to consider here; however, the client devices can range from a high-powered Windows desktop to low-end thin clients and to mobile devices.

Receiver

Citrix Receiver is a universal software client that provides a secure, high-performance delivery of virtual desktops and applications to any device anywhere. Receiver is platform agnostic. The Citrix Receiver is device agnostic, meaning that there is a Receiver for just about every device out there, from Windows to Linux-based thin clients and to mobile devices including iOS and Android. In fact, some thin-client vendors have performed a close integration with the Citrix Ready program to embed the Citrix Receiver code directly into their homegrown operating system for seamless operation with XenDesktop.

The Citrix Receiver must be installed on the end user client device in order to receive the desktop and applications from XenDesktop. It must also be installed on the virtual desktop in order to receive applications from the application servers (XenApp or XenDesktop), and this is taken care of for you automatically when you install the VDA on the virtual desktop machine.

System requirements

Each component has its requirements in terms of operating system and licensing. You will need to build these operating systems on VMs before installing each component. For help in creating VMs, look at the relevant Hypervisor documentation; in this book, we have used Citrix XenServer as the Hypervisor.

Receiver

The Citrix Receiver is a universal software client that provides a secure, high-performance delivery of virtual desktops and applications. The Receiver is available for Windows, Mac, mobile devices such as iOS and Android, HTML5, Chromebook, and Java 10.1.

You will need to install the Citrix Receiver twice for a complete end-to-end connection to be made.

Once on the end user's client device—there are many supported devices including iOS and Android—and once on the Windows virtual desktop (for Windows) that you will serve your users. This is done automatically when you install the **Virtual Desktop Agent** (**VDA**) on the Windows virtual desktop.

You need this Receiver to access the applications that are running on a separate application server (XenApp or XenDesktop).

StoreFront 2.1

StoreFront replaces the web interface. StoreFront 2.1 can also be used with XenApp and XenDesktop 5.5 and above. The operating systems that are supported are as follows:

- Windows Server 2012 R2, Standard or Data center
- Windows Server 2012, Standard or Data center
- Windows Server 2008 R2 SP1, Standard or Enterprise

System requirements are as follows:

- RAM: 2 GB
- Microsoft **Internet Information Services** (**IIS**)
- Microsoft Internet Information Services Manager
- .NET Framework 4.0

Firewall ports – external:

As StoreFront is the gateway to the Site, you will need to open specific ports on the firewall to allow connections in, mentioned as follows:

- Ports: `80` (http) and `443` (https)

Firewall ports – internal:

By default, StoreFront communicates with the internal XenDesktop Delivery Controller servers using the following ports:

- `80` (for StoreFront servers) and `8080` (for HTML5 clients)

You can specify different ports.

For more information on StoreFront and how to plug it into the architecture, refer to `http://support.citrix.com/article/CTX136547`.

Databases

The supported Microsoft SQL Server versions are as follows:

- SQL Server 2012 SP1, Express, Standard, and Enterprise Edition
- SQL Server 2008 R2 SP2, Express, Standard, Enterprise, and Data center Edition

The installer deploys this automatically. It can also be found on the XenDesktop installation media in the `Support` folder.

The following databases are also supported:

- SQL Server clustered instances
- SQL Server Mirroring
- SQL Server 2012, AlwaysOn Availability Groups

www.allitebooks.com

Studio

The operating systems that are supported are as follows:

- Windows 8.1, Pro and Enterprise
- Windows 8, Pro and Enterprise
- Windows 7, Pro, Enterprise, and Ultimate
- Windows Server 2012 R2, Standard, and Data center
- Windows Server 2012, Standard and Data center
- Windows Server 2008 R2 SP1, Standard, Enterprise, and Data center

System requirements are as follows:

- Disk space: 75 MB
- Microsoft .NET Framework 3.5 SP1 (Windows 2008 R2 only)
- Microsoft Management Console 3.0
- Windows PowerShell 2.0 (Windows 7 and Windows 2008 R2) or PowerShell 3.0 (Windows 8.1, Windows 8, Windows 2012 R2, and Windows 2012)

MMC 3.0 and PowerShell are included in the Windows Server.

Delivery Controller

The operating systems that are supported are as follows:

- Windows Server 2012 R2, Standard or Data center Edition
- Windows Server 2012, Standard or Data center Edition
- Windows Server 2008 R2, Standard or Enterprise Edition

System requirements are as follows:

- Disk space: 100 MB
- Microsoft .NET Framework 3.5 SP1 (Windows 2008 R2 only)
- Microsoft .NET 4.0
- Windows PowerShell 2.0 (included with Windows 2008 R2) or PowerShell 3.0 (included with Windows 2012 R2)
- Visual C++ 2005, 2008 SP1, and 2010 Redistributable Package

The installer installs the mentioned software automatically for you. It is also available on the XenDesktop installation media in the `Support` folder.

Director

The operating systems that are supported are as follows:

- Windows Server 2012 R2, Standard or Data center
- Windows Server 2012, Standard or Data center
- Windows Server 2008 R2 SP1, Standard or Data center

System requirements are as follows:

- Disk space: 50 MB
- Microsoft .NET Framework 4.0

The installer deploys this framework automatically for you.

- Microsoft **Internet Information Services** (**IIS**) 7.0 and ASP.NET 2.0

The supported browsers to view Director are as follows:

- Internet Explorer 11, 10, and 9 (IE 10 compatibility mode is not supported)
- Firefox
- Chrome

The Virtual Delivery Agent (VDA)

The VDA has also been referred to as the **Delivery Agent** (**DA**) in this book. It is available for both Windows desktop OSes as well as for Windows Server OSes.

The supported operating systems are as follows:

- Windows 8.1, Pro or Enterprise
- Windows 8, Pro or Enterprise
- Windows 7 SP1, Pro, Enterprise, or Ultimate
- Windows Server 2008 R2 SP1, Data center, Enterprise, or Standard

- Windows Server 2012 R2, Standard or Data center
- Windows Server 2012, Standard or Data center
- Windows Server 2008 R2 SP1, Standard, Enterprise, or Data center

The installer automatically deploys the support components such as the Microsoft .NET Framework and the Visual C++ Runtime Library. The Visual C++ components are also available on the XenDesktop installation media in the `Support` folder.

Multimedia acceleration features for HDX require Microsoft Media Foundation to be installed prior to installing the VDA on the machine.

To use a Windows XP or Vista machine in XenDesktop 7, you will need to install an earlier version of the Citrix VDA, which can be downloaded from the `Citrix.com` downloads website.

Server host

XenDesktop runs operating systems in VMs. These VMs exist on Hypervisors that run on top of the server hardware.

The supported Hypervisor operating systems are as follows:

- Citrix XenServer 6.0.2, 6.1, and 6.2
- VMware vSphere 5.0 update 2 and vSphere 5.1 update 1
- Microsoft System Center Virtual Machine Manager 2012 R2, 2012 SP1, or 2012

To see a list of server hardware that is compatible with XenServer, go to `http://hcl.xensource.com`. Click on the **Servers** link and select your XenServer version. I purchased a compatible server on eBay at a cheap price. For creating a production environment and to do anything with HDX 3D, you should purchase a new system with support.

A more exhaustive list of the supported Hypervisors can be found at `http://support.citrix.com/article/CTX131239`.

Active Directory

The supported operating system is as follows:

- Windows Server 2003 or higher

Designing a basic XenDesktop® Site

We are just about to get started with installing XenDesktop, but before we do, we need to do some initial assessment of the design. We need to think about what the XenDesktop Site will look like when we are finished, taking into account the number of users we want to service. The resulting design will tell us how much server, hardware, and storage capacity we will need, which FlexCast model to deploy, and which user groups to start with. We will also end up with an architecture diagram of the complete solution which will show how all the components fit together.

Ultimately, you can navigate to the Citrix Project Accelerator that has a handy tool to help you to quickly assess, design, and deploy your XenDesktop Site. It is located at `http://project.citrix.com`.

The Project Accelerator can be complex and confusing. Don't get caught up in it too much but use it as a general guideline.

Scenario

To help guide you through the process, I have created a fictitious company called `Xenpipe.com`. For now, there is just one type of user at Xenpipe—normal users who require access to Microsoft Office applications. In future, we can add heavy bandwidth users who require access to design applications (HDX 3D), mobile users who require remote access, and task workers who don't require any personalization, just a locked-down desktop. After plugging this information into the Citrix Project Accelerator, we came up with the following table to help us size our deployment. We chose to implement a Hosted VDI solution because it provides the most common form of virtual computing to any device, such as thin clients, PCs, laptops, or mobile devices.

User group	FlexCast	Users	Servers	Hardware	Storage
Main HQ	Hosted VDI	10	1 physical 18 virtual	21 cores 72 GB RAM	723 GB HD 400 IOPS

The resulting architecture will look as follows:

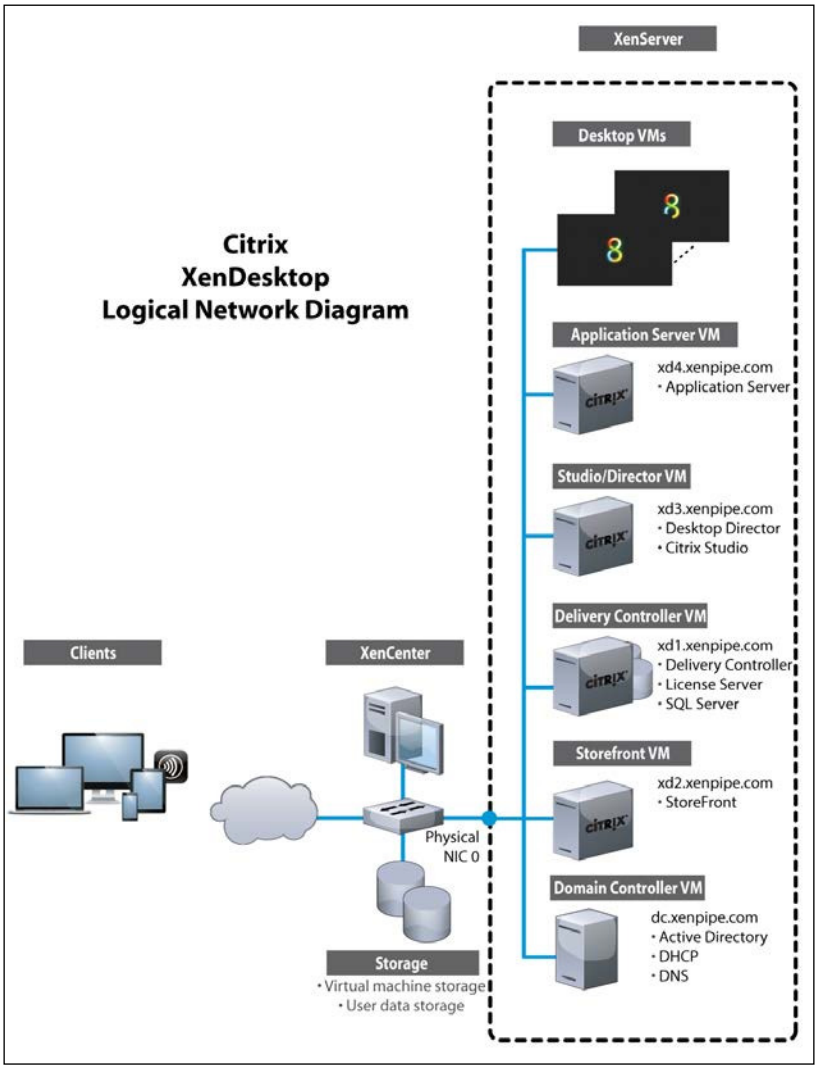

Common Citrix® communication ports

As you are building your infrastructure, it's important to know what type of protocols will run across your network. Sometimes, system administrators separate devices with network routers, switches, and firewalls that can block the XenDesktop implementation from working. The following is a list of protocols that you should allow through the routers, switches, and firewalls. All the Citrix protocols can be found in *CTX Article 101810* at `http://support.citrix.com/article/CTX101810`.

Citrix product	Protocol	Port(s)	Description
Citrix license server			
License Manager Daemon	TCP	27000	Handles license requests
Citrix Vendor Daemon	TCP	7279	Check-in and check-out of licenses
License Management Console	TCP	8082	Browser-based administration console
Common communication ports			
Citrix Receiver	TCP	80, 443	Communication with StoreFront or the NetScaler gateway
ICA, HDX	TCP	1494	Desktops and applications flow over this protocol
Session Reliability	TCP	2598	Session Reliability for ICA, HDX
Management Console	TCP	2513	Citrix Management Consoles
XML Server	TCP	80, 8080, 443	Desktop and application requests
STA	TCP	80, 8080, 443	Secure Ticket Authority embedded into XML service requests
Citrix XenDesktop			
Citrix XenServer	TCP	80, 443	Communication with XenServer
Microsoft Hyper-V	TCP	8100	SCVMM Administrator Console
VMware vSphere	TCP	443	VMware Web Services communication
Broker	TCP	80, 443	Used for communication with VDA, SDK, and XML service
Active Directory Identity Service	TCP	80	Used for Active Directory communications
Configuration Service	TCP	80	Used by the configuration service

Citrix product	Protocol	Port(s)	Description
Host Service	TCP	80	Used by the host service
Machine Creation Service	TCP	80	Used by machine creation services
Machine Identity Service	TCP	80	Used by machine identity services
License Configuration Service	TCP	80	Used by the licensing service
Desktop Director	TCP	80, 443	Used by Desktop Director
Virtual Desktop Agent	TCP	80	Communication with the Desktop Delivery Controller
	TCP	135, 3389	Communication with the Desktop Delivery Controller for remote assistance
	UDP	16500~16509	HDX audio
	TCP	80, 5985	Communication with Desktop Director
Citrix Desktop Service	TCP	80	Used by the workstation agent to communicate with the Broker
Database	TCP	1433, 1434	Microsoft SQL Server
Citrix XenServer			
XenCenter	TCP	22	SSH
	TCP	443	Management using XenAPI
	TCP	5900	VNC for Linux guests
	TCP	3389	RDP for Windows guests
Resource Pool	TCP	22	SSH
	TCP	443	Management using XenAPI

Citrix product	Protocol	Port(s)	Description
Infrastructure	TCP/ UDP	123	Network Time Protocol
	TCP/ UDP	53	DNS
	TCP	389	Active Directory
	TCP/ UDP	139	ISO Store: NetBIOS Session Service
	TCP/ UDP	445	ISO Store: Microsoft-DS
Storage	TCP	3260	iSCSI storage
	TCP	2049	NFS storage
	TCP	21605	SOAP over HTTP StorageLink

Summary

Now you should have a good grasp of the components, system requirements, and terminology used in Citrix XenDesktop. This chapter also serves as a good reference to look back on as you move forward. Remember to use the Internet to search for XenDesktop sizing guides and best practices, and don't forget to try out the Citrix Project Accelerator at `http://project.citrix.com`.

Now that you have an understanding of what the XenDesktop Site will look like from the network diagram, components, terminology, and concepts, we will install XenDesktop. The next chapter discusses how to plan and execute the installation.

2
Installing XenDesktop®

The move to **Virtual Desktop Infrastructure** (**VDI**) will be dominant for many reasons. Once you have built the foundation with a network diagram that contains the components and you understand the terminology and concepts, it will be easier to move forward. Understanding the details of the implementation takes time as any experienced person will tell you. That is why we started with building the foundation in the first chapter. When you are done with this chapter, you will have a complete XenDesktop Site installed and running. In this chapter, we will walk you through the following topics:

- Planning the XenDesktop installation
- Installing the controller (XD1)
- Installing StoreFront (XD2)
- Installing Director (XD3)
- Creating the virtual desktop and application delivery master images
- Installing the Virtual Delivery Agent on the master images
- Configuring the StoreFront server
- Creating the machine catalogs
- Creating the delivery groups
- Installing the Citrix Receiver on the client devices
- Testing the connection

www.allitebooks.com

Planning the XenDesktop® installation

There are some specific tasks that should be performed during the installation process. They are outlined in the following task list, which you can use as a guide to check off as you go along:

The installation task list		
Step	**Description**	**Completed**
1	Install the Delivery Controller, Studio, license server, and SQL database. Configure a Site.	
2	Install StoreFront.	
3	Install Director.	
4	Create the master images.	
5	Install the Delivery Agent on the master images.	
6	Configure StoreFront.	
7	Create the machine catalogs.	
8	Create the delivery groups.	
	The installation checkpoint	
9	Install and configure the Citrix Receiver on the client device.	
10	Test the connection.	

Before you get started, the following is a checklist of the items you will need to have on hand in order to start building and installing your XenDesktop Site. For our example, we have downloaded the `.iso` files of the installation DVDs from the Microsoft and Citrix download websites:

The installation inventory checklist		
Item	**Description**	**Check**
1	Physical server for hosting Hypervisor and virtual machines	
2	Hypervisor installation media and license	
3	Windows Server 2012 R2 with license	
4	Windows 7 installation media and license	
5	Windows 8 installation media and license	
6	XenDesktop installation media and license	

The first step is to create physical and virtual hosting environments. In our examples, we will use one physical server that runs the XenServer Hypervisor.

How to install XenServer onto the physical server?

We use a tool called `http://pendrivelinux.com` to install XenServer onto a bootable USB thumb drive, boot the server from that thumb drive, and proceed with the installation.

All the servers used for the XenDesktop VMs in this demo will run on Windows Server 2012 R2. Windows Server 2008 R2 SP2 will work as well.

It is assumed that your physical host (XenServer) is already installed and is ready to begin the installation of XenDesktop VMs, in addition to an Active Directory domain controller which is already available either as a VM or a physical server. In our example, we have used a domain controller as a VM. This same domain controller will be used as a DHCP and DNS server. After installing the Hypervisor, you need to install the Hypervisor management console on a separate physical machine such as a laptop or desktop computer. An example screenshot of our XenCenter console that manages the XenServer can be seen in the next screenshot that follows.

Install the domain controller if you don't already have one, and install a Windows Server 2012 R2 virtual machine to host the first XenDesktop VM with a **Fully Qualified Domain Name** (**FQDN**); in our example, it is `xd1.xenpipe.com`.

To create a storage repository for your installation media `.iso` files (DVDs), follow the video tip found at `http://www.citrix.com/tv/#videos/7933`.

To mount and install Windows Server 2012 R2 `.iso` (DVD) in XenServer and then install XenTools, follow the video tip found at `http://www.citrix.com/tv/#videos/7934`. If you are using XenServer as your Hypervisor, be sure to install the XenServer tools on each virtual machine you create.

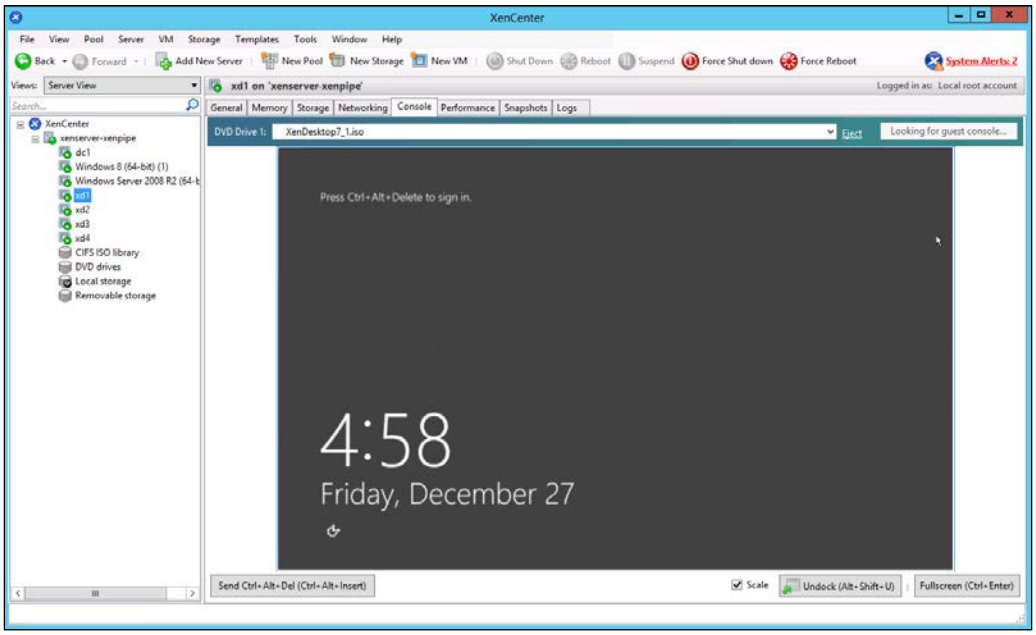

Step 1 – installing the controller (XD1)

The first machine you will install is the XenDesktop controller. If you refer back to our network diagram in *Chapter 1*, *Designing a XenDesktop® Site*, this is labeled as `xd1.xenpipe.com` or XD1 for short.

Installing the components on XD1

We are ready to begin our installation of the first XenDesktop server in the XenDesktop Site. You need to work from your laptop or desktop, typically the machine that is running XenCenter to perform this work. Perform the following steps to install the components on XD1:

You can use the console screen of the Windows Server in XenCenter; however, it will be much easier for you to connect directly to the machine using RDP. From a command-line prompt, type `c:\> mstsc /admin`. Enter the IP address and log in to the server.

1. Mount the XenDesktop `.iso` installation media in XenCenter as shown in the following screenshot:

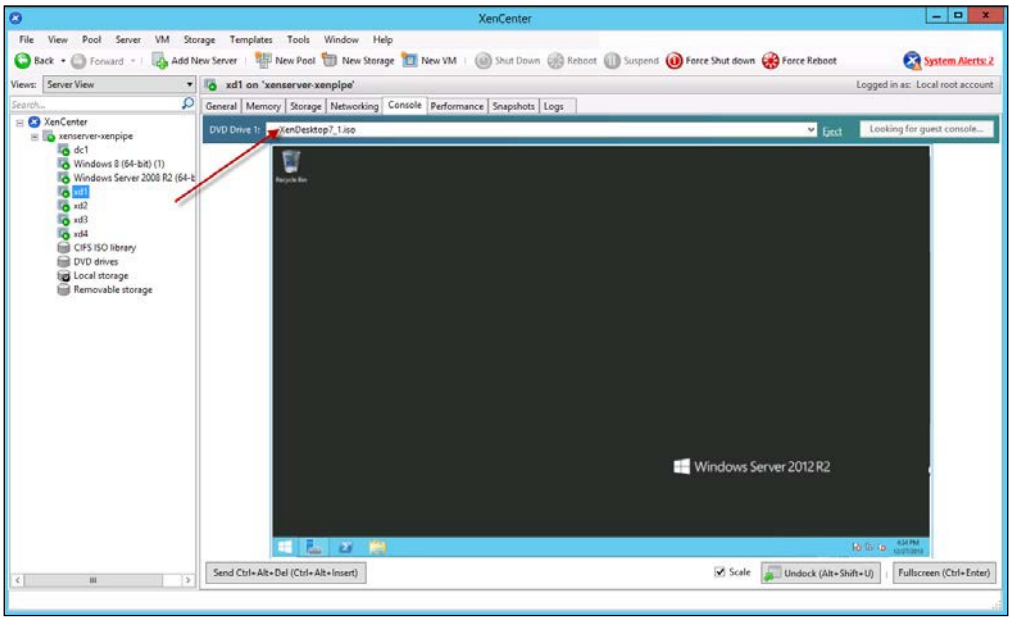

2. Log in to the server using a domain administrator account. Your account must have local administrator privileges, and you should be a member of the Domain Admins group in Active Directory.

From now on, you should always log in to the servers with a domain account for installation. Logging in to the local machine and installing XenDesktop will not work; you will have to start over on the server if you do. To log in with a domain account, you must use the `<domain>\Administrator` format in the **User** field of the login box; in our example, `xenpipe\administrator`.

3. Launch the installation from the CD drive of the virtual machine.
4. On the welcome page, click on **Start** as shown in the following screenshot:

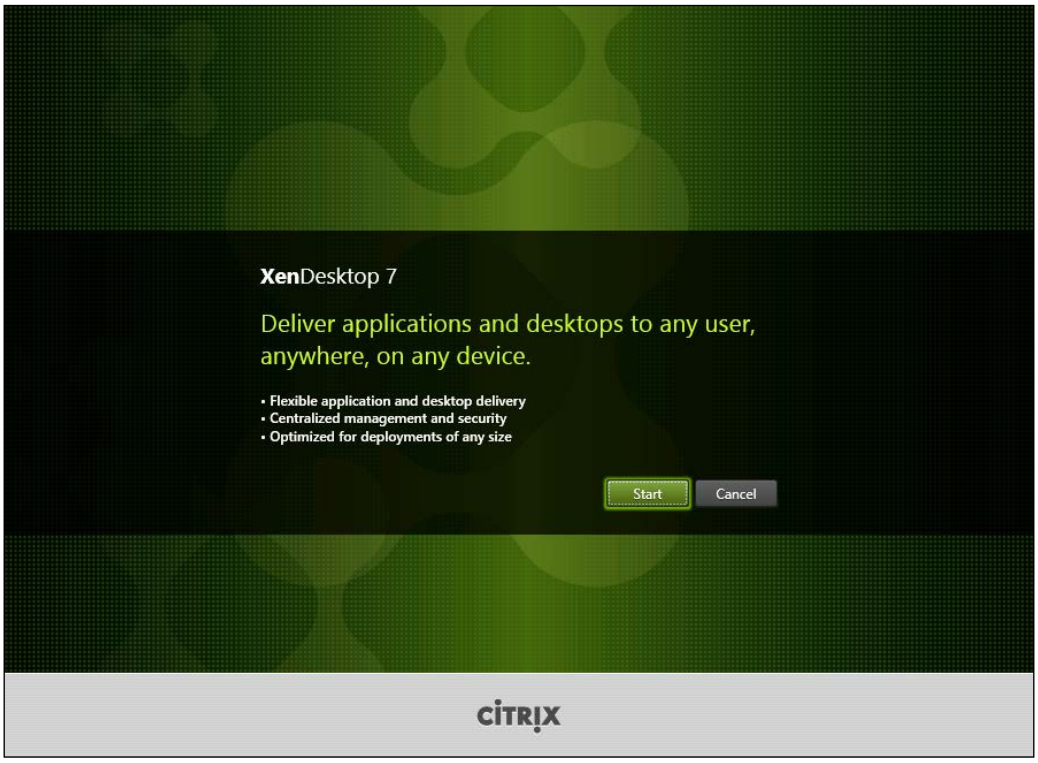

5. Select **Delivery Controller** to install XenDesktop and accept the license agreement as shown in the following screenshot:

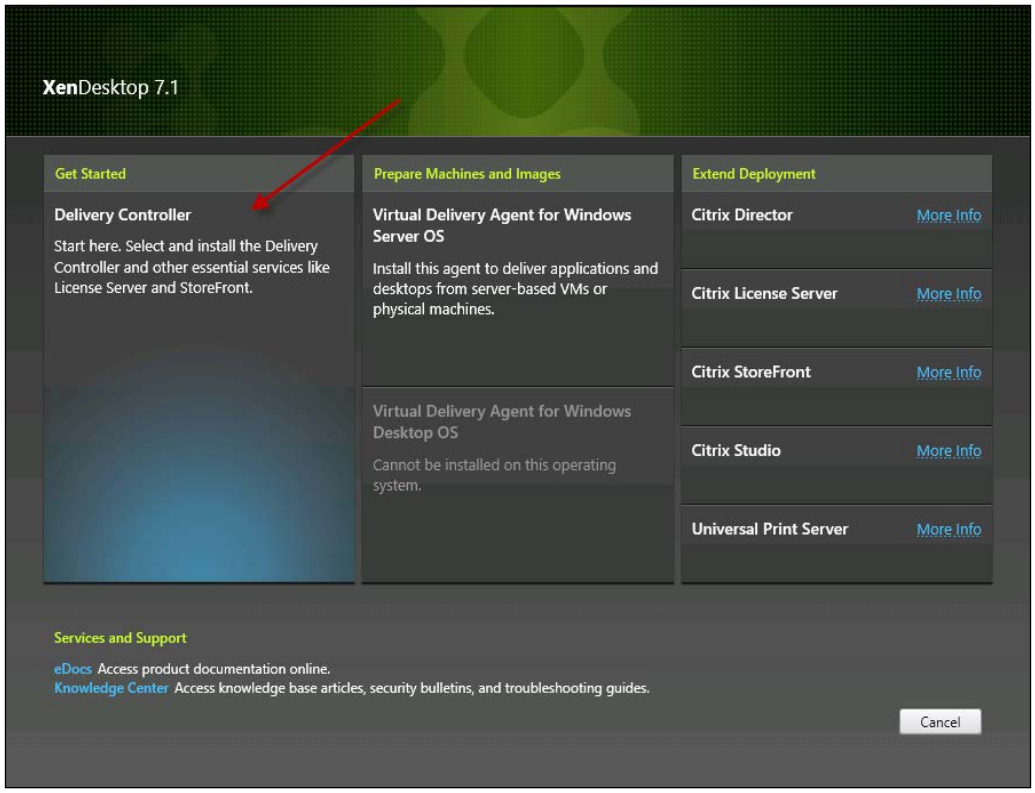

6. Select the **Delivery Controller**, **Studio**, and **License Server** components as shown in the following screenshot:

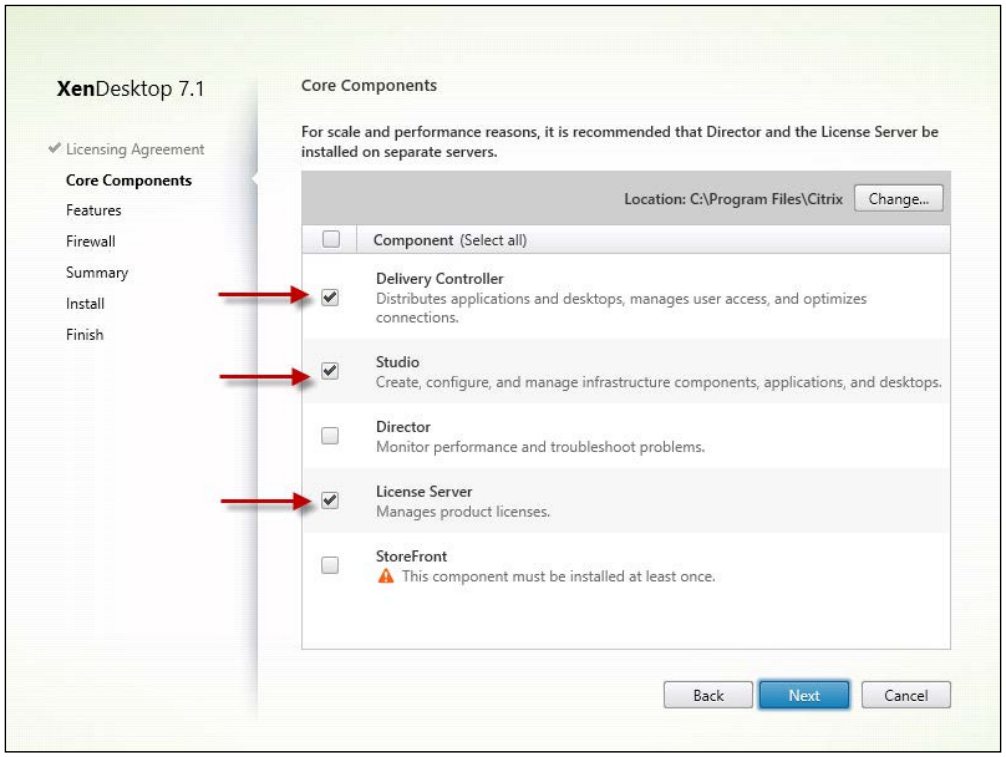

7. Install **Microsoft SQL Server 2012 Express**. Set up the Windows firewall by selecting the automatic option. Select **Install**.

8. Select **Finish**, as shown in the following screenshot:

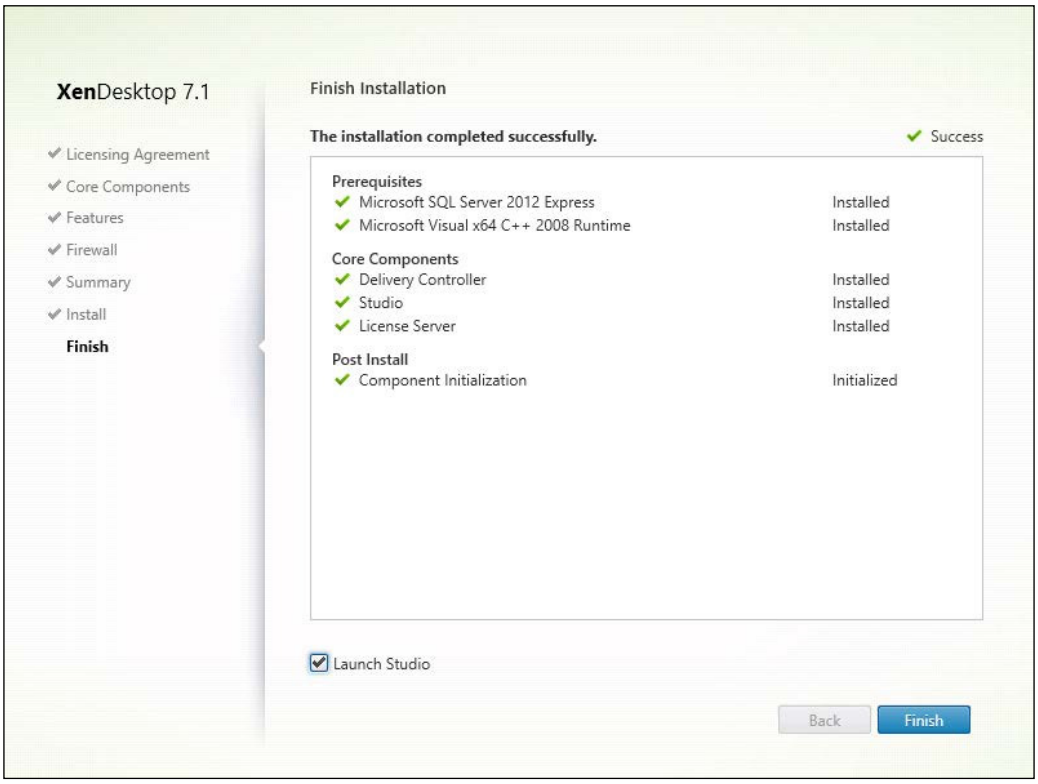

Configuring a Site

Now that you have the components installed on the first XenDesktop server, you can license the Site and set up the database and connection details using the following steps:

1. Launch Studio, and on the welcome page under the **Full deployment** section, select **Get started! Create a Site,** as shown in the following screenshot:

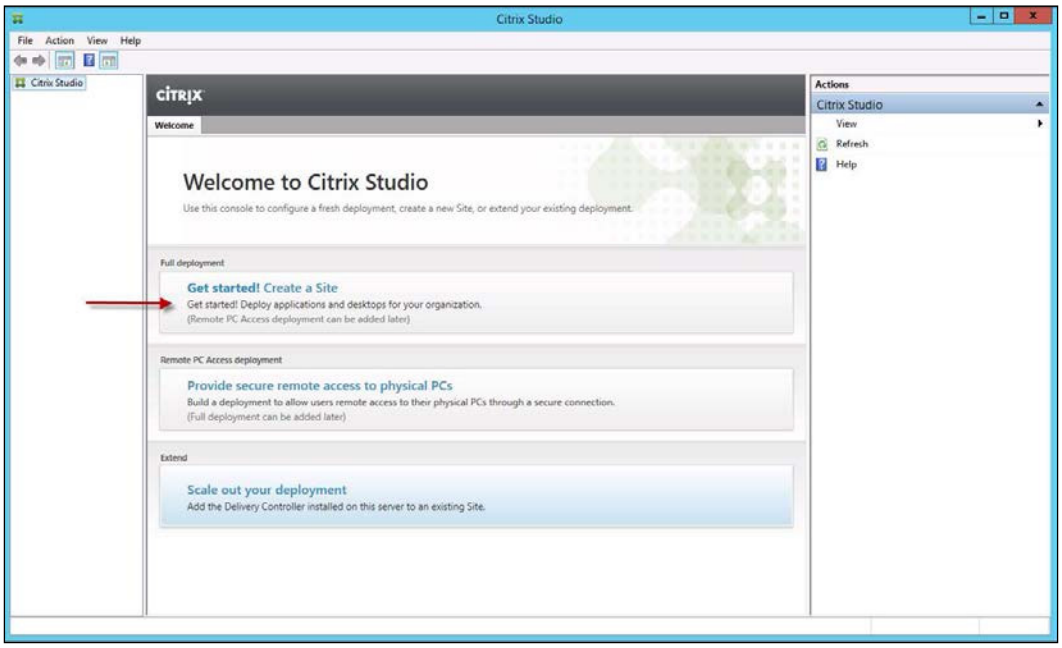

2. Select **Configure the Site and start delivering applications and desktops to users** and name your Site.

3. For the database location, enter `.\SQLEXPRESS` and select **Test connection**. To create a database automatically, click on **OK**. When the test connection is passed, select **Next** as shown in the following screenshot:

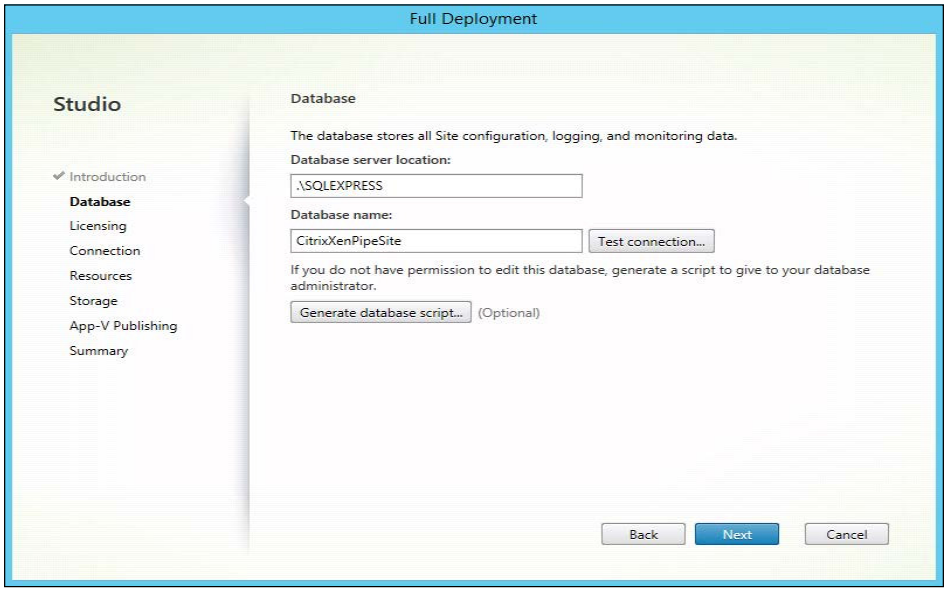

4. On the **Licensing** page, select **Use the free 30-day trial** or **Use an existing license** as shown in the following screenshot:

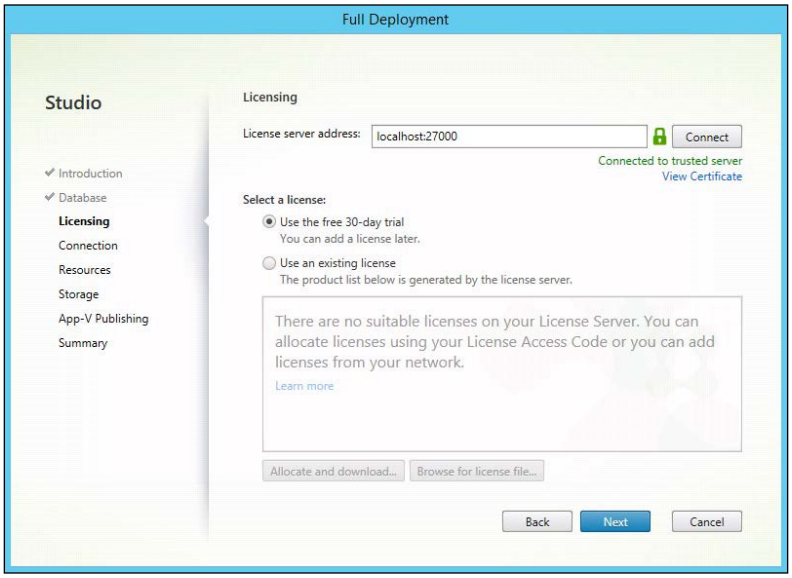

www.allitebooks.com

5. On the **Connection** page, enter the details for your Hypervisor—**Host type**, **Address**, **Username**, **Password**, and **Connection name**. Select **Studio tools (Machine Creation Services)** as shown in the following screenshot:

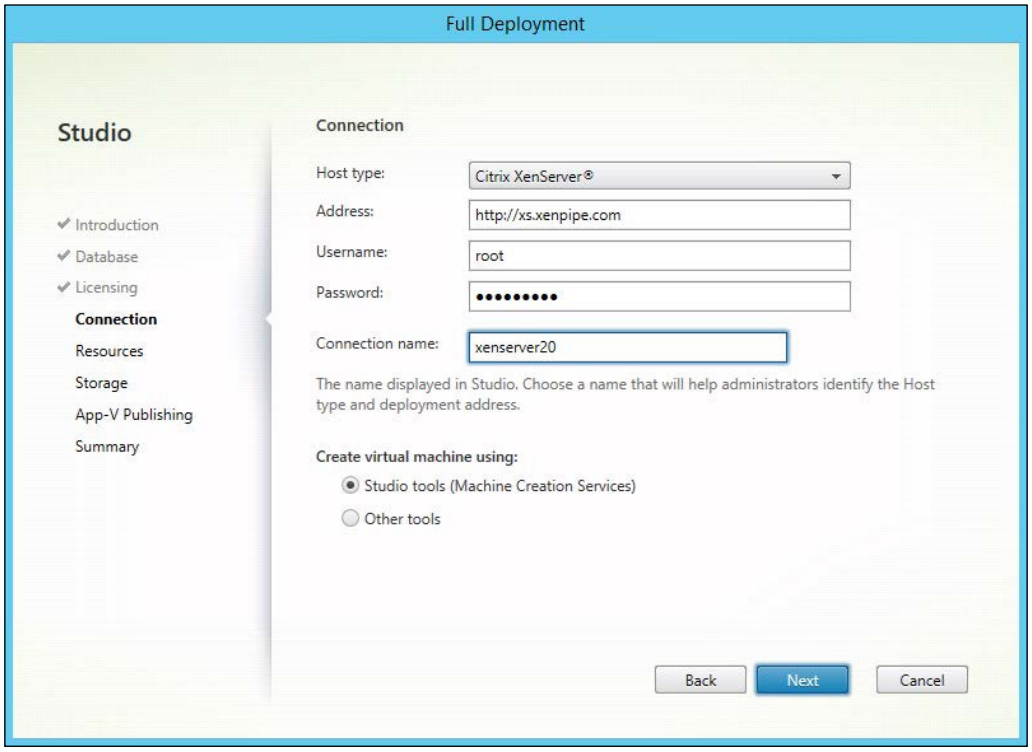

6. For the **Network** settings, type a name, select the network interface, and click on **Next**.
7. For **Storage**, select your storage device(s) and select **Use same storage for Virtual Machines and personal vDisks**. Then, click on **Next**.
8. Select **No** for **App-V Publishing** and then click on **Next** followed by **Finish**.

Step 2 – installing StoreFront (XD2)

The second server that we will install is StoreFront. If you refer back to the network diagram in *Chapter 1*, *Designing a XenDesktop® Site*, this is the server labeled as `xd2.xenpipe.com`, or XD2.

Installing the components on XD2

We will need a server certificate signed by the domain certificate authority to install StoreFront.

If you don't have a domain certificate authority, see *Appendix A*, *Creating a Domain Certificate Authority*, for tips on creating one for your domain.

Make sure that the **Internet Information Services** (**IIS**) role is installed on the StoreFront server XD2.

Creating a server certificate and adding a Site binding

The first step for creating the StoreFront server is to create a server certificate and add a Site. The steps can be summarized as follows:

1. Log in to the second XenDesktop server, or XD2 in our example, using a domain administrator account.
2. Navigate to **Start** | **Administrative Tools** | **Internet Information Services (IIS) Manager**.

3. In the **Connections** pane, select the second server or **XD2** in this case, as shown in the following screenshot:

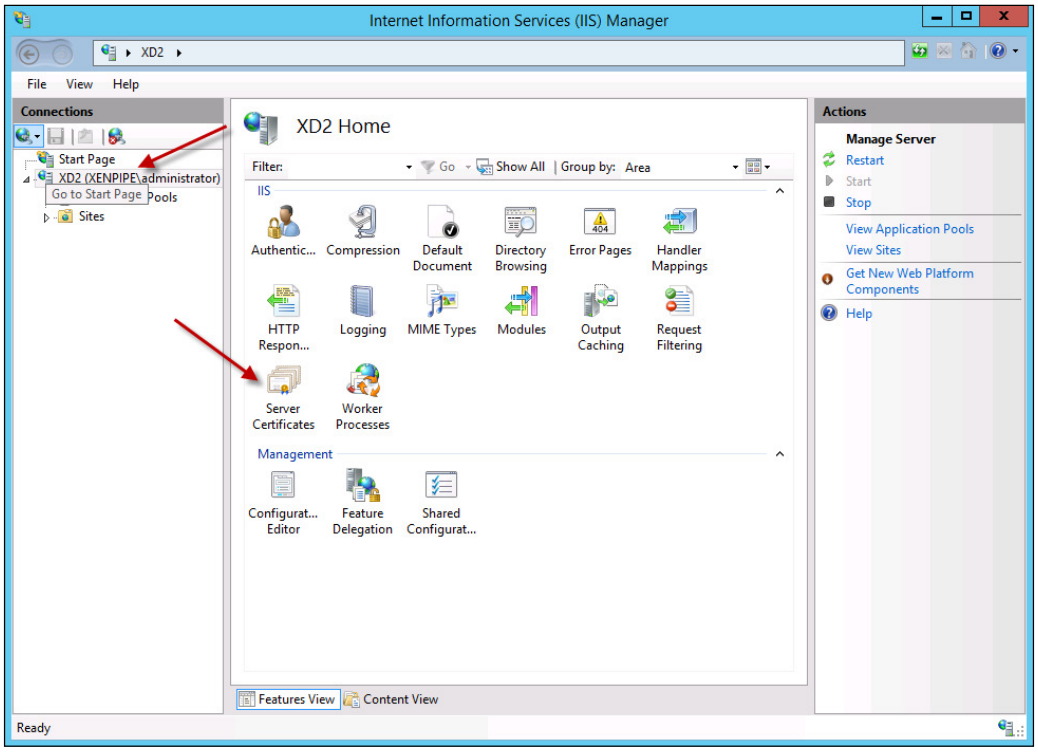

4. Open **Server Certificates** and then select **Create Domain Certificate**. Fill out the fields as shown in the following screenshot.

Make sure that you create a wildcard certificate with an asterisk placed before the domain name in the **Common name** field so that we can change the base URL of the StoreFront server as required. Trust me; it is better to do this now.

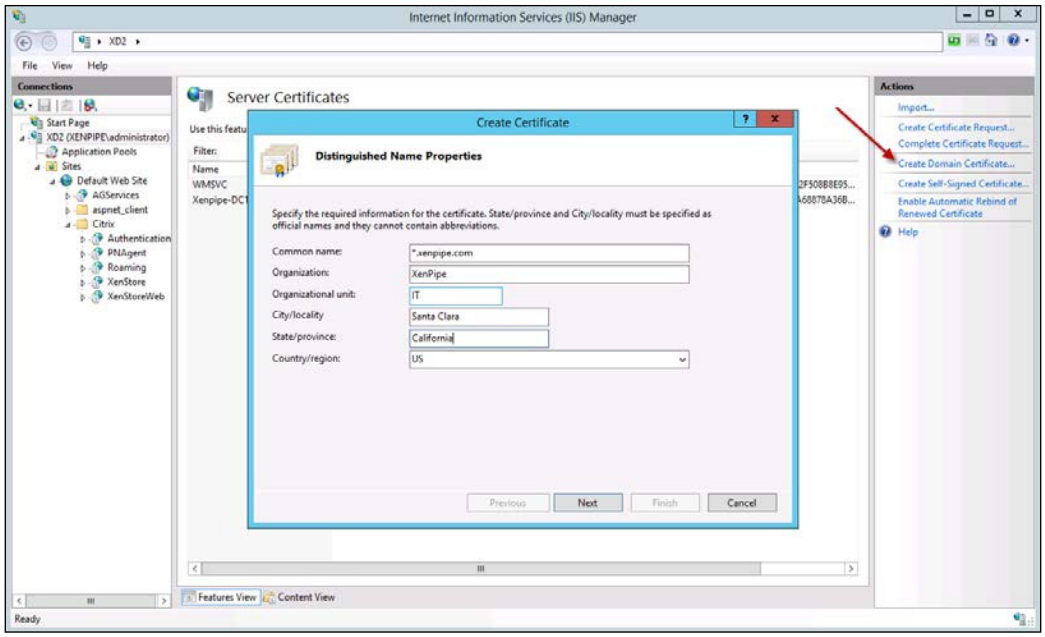

5. Type the certificate authority name in the `<CertificateAuthorityName>\<ServerName>` format and a friendly name. In our example, it is `xenpipe-DC1-CA\dc1` and `Xenpipe-DC1-CA`, respectively.

If you receive an error message, it is most likely because you did not select **Enterprise Certificate Authority** when installing the **Certificate Authority Role** on the domain controller.

6. Go back to the **Connections** pane and select the **Default Web Site** option for server two or XD2. Select and open **Bindings…**, as shown in the following screenshot:

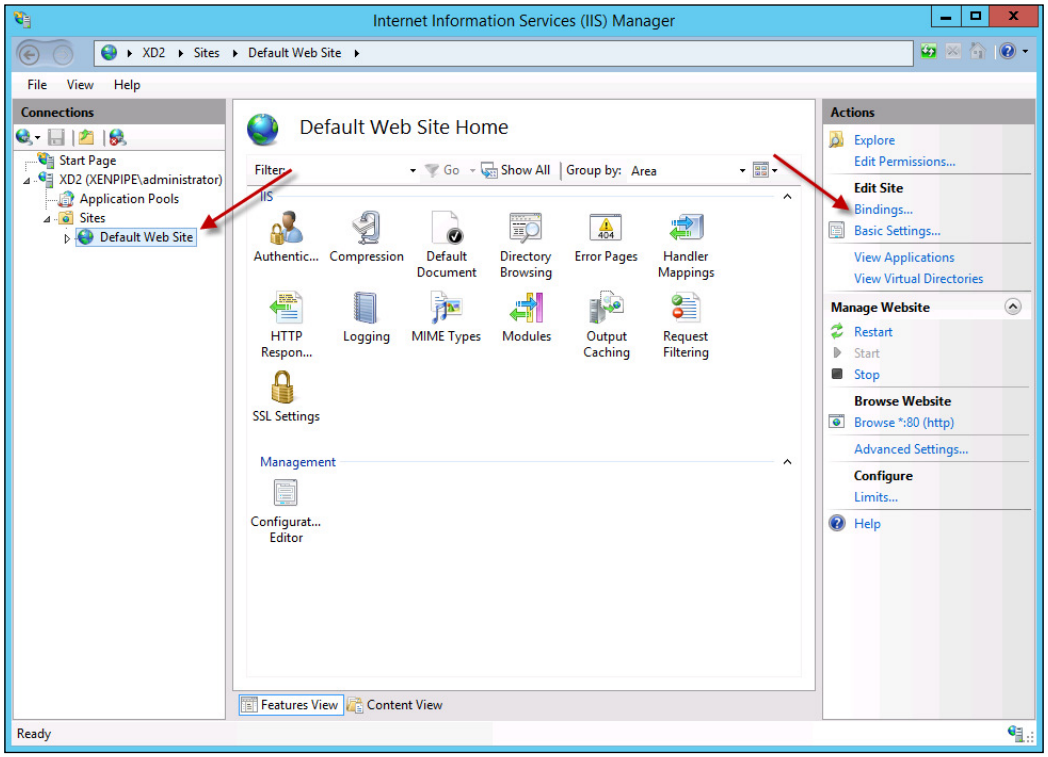

7. In **Site Bindings**, select **Add**. Choose **https**, **Port 443**, and select the certificate we just created. Click on **OK** and then on **Close**, as shown in the following screenshot:

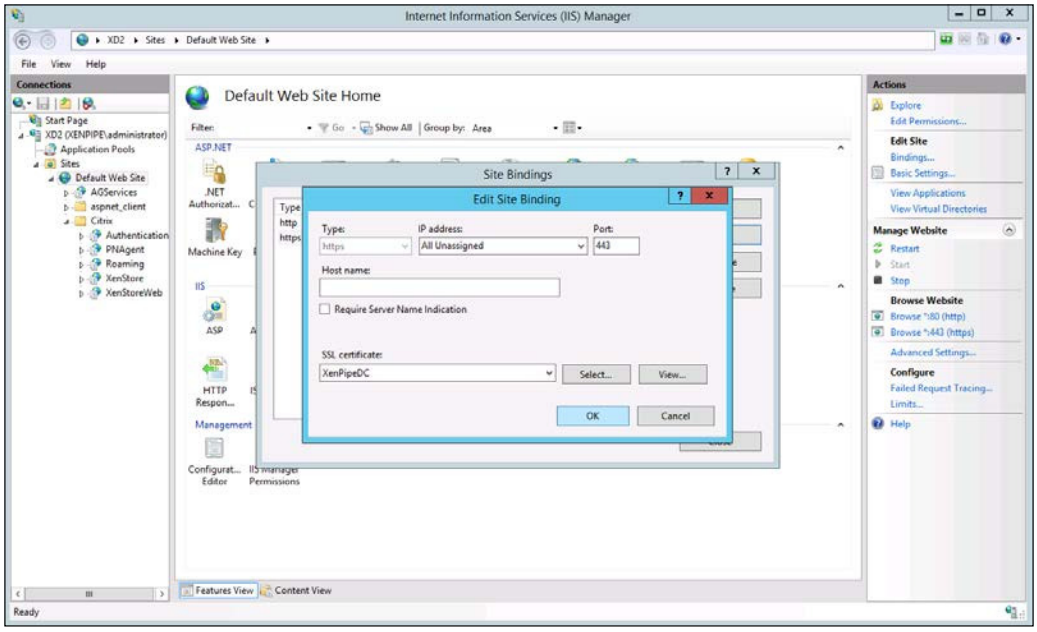

Installing StoreFront

Now that the basic server is built and ready for StoreFront, we can go about installing the StoreFront software using the following steps:

1. Mount the XenDesktop installation media in the DVD drive for the second server, XD2. You can use the XenCenter console for this from the XenCenter administrator machine, as illustrated in the following screenshot:

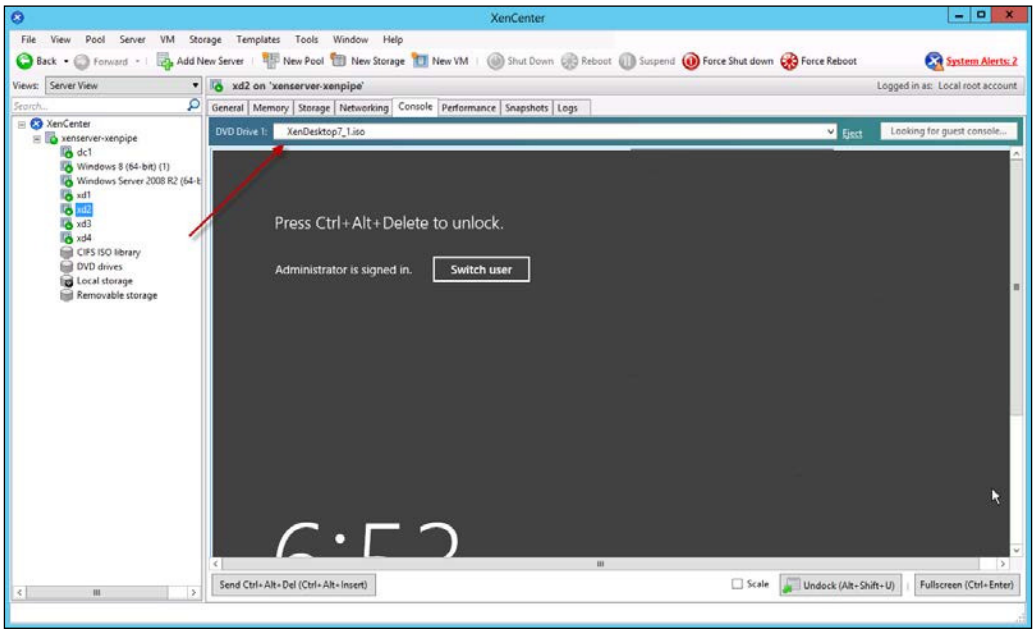

2. Log in to the second XenDesktop server, or XD2 in our example, using a domain administrator account. Launch the XenDesktop installation media and click on **Start**. In the installation options, select **Delivery Controller**.

3. Accept the licensing agreement, and on the **Core Components** page, only select **StoreFront**. Then, click on **Next**, **Install**, and finally **Finish** as shown in the following screenshot:

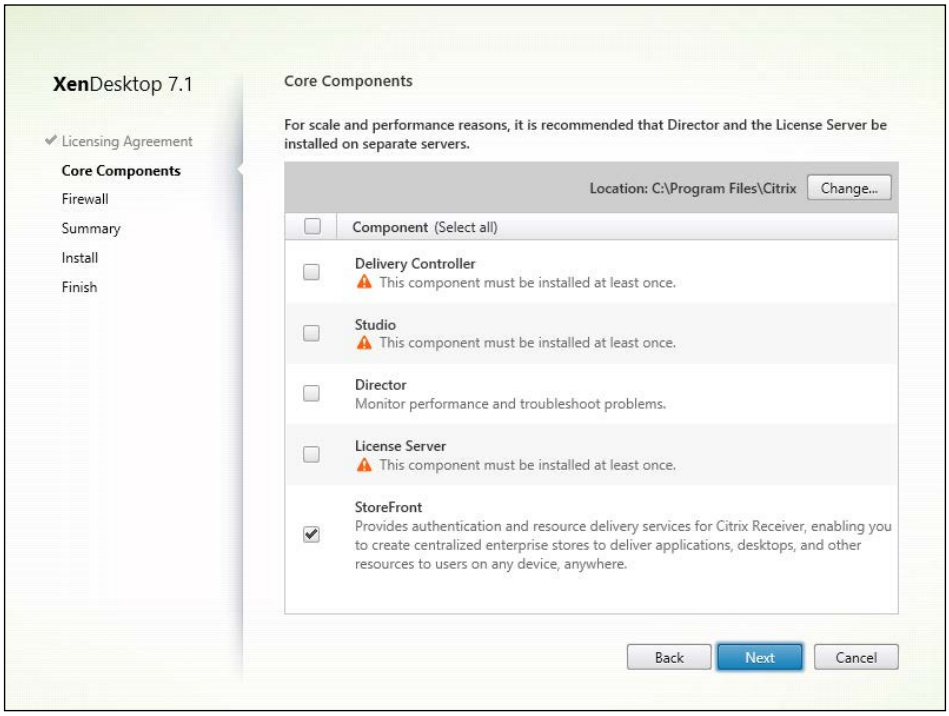

Step 3 – installing Director (XD3)

In this step, we will install the XenDesktop Director. Referring back to the network diagram, this is server `xd3.xenpipe.com` or XD3.

Installing the components on XD3

We will create and use the XD3 server to host the XenDesktop Director using the following steps:

1. Mount the XenDesktop installation media on the XD3 server in the same way as you did for XD2 in the previous example.
2. Log in to the third XenDesktop server, or XD3 in our example, using a domain administrator account.

3. Launch the XenDesktop installation CD. Click on **Start,** and under the installation options, select **Delivery Controller**.
4. Accept the license agreement; only select **Director** and click on **Next**.
5. Enter the FQDN of server 1 or XD1, where the Delivery Controller is located. In our example, `xd1.xenpipe.com`. Then, test the connection.
6. Automatically enable the remote assistance and firewall rules. Click on **Next**, **Install**, and then **Finish**.

Step 4 – creating the virtual desktop and application delivery master images

In this task, you will install a master image of the virtual machine that will be delivered as a desktop to Windows 8 users. You will also create a master image for the application server. At the end of this task, you should have two master images as listed in the following table:

Master images		
No.	**Master image**	**Use case**
1	Windows 8 master	Delivering Windows 8 desktops
2	Windows Server 2012 master	Delivering applications

Make sure that these machines are configured to use DHCP from the domain controller and join them to the domain using a domain administrator account.

In this example, we have used Windows 8 and Windows 7 for desktop master images, but this can also be done easily with Windows XP using the following steps:

1. From XenCenter, navigate to **New VM** | **Windows 8 64 bit**.
2. Mount the Windows 8 installation media in the DVD drive and click on **Create**.
3. Repeat the preceding two steps for Windows Server 2012.

4. After creating the Windows Server 2012 VM, install your applications on it. In our example, we will install Microsoft Office Professional SP1. Be sure to log in with a domain administrator account.

To watch a demo on how to create a Windows desktop in XenServer, follow the video tip at `http://www.citrix.com/tv/#videos/7974`.

When navigating around Windows 8, you might want to look up shortcuts for Windows 8 to help you navigate as the start menu is no longer there. I found *Alt* + *F4* and Windows key + *X* useful as well as the Windows key + *D*, which will move between the desktop and tiles.

5. If Windows Server 2012 doesn't get installed with a GUI command line, you can turn on the GUI by performing the following steps:

    1. Type in the following in the command prompt:

        ```
        PowerShell, <Enter>.

        Install-WindowsFeature Server-Gui-Shell,
        Server-Gui-Mgmt-Infra, <Enter>.

        shutdown -r -t 0
        ```

    2. Move the mouse to one of the two right-hand hot spots in order to get the start menu.
    3. Right-click anywhere on the screen and you will be able to see all the apps.

Step 5 – installing the Virtual Delivery Agent on the master images

Now that you have built a master image of the desktop, you may want to serve it to your users. For this, you need to install the XenDesktop Delivery Agent and Receiver onto that image.

www.allitebooks.com

Creating the desktop master images

In this task, we will install the **Virtual Delivery Agent** (**VDA**) and Citrix Receiver on the master image for desktops. VDA is the piece of software that allows the desktop VM to communicate with the XenDesktop servers. The Citrix Receiver allows applications to launch seamlessly from the application servers.
The following are the steps to create the desktop master images:

1. Mount the XenDesktop installation media in the DVD drive in XenCenter, then log in to the master image, and launch the install media. Select **Virtual Delivery Agent for Windows Desktop OS** as shown in the following screenshot:

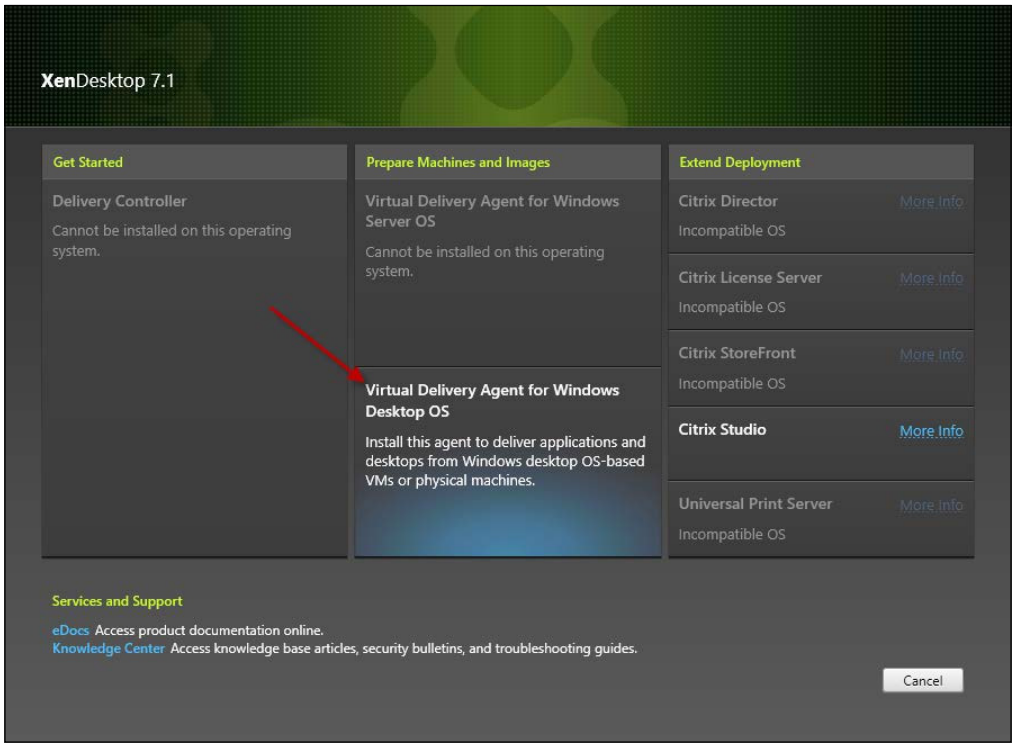

2. Select **Create a Master Image** as shown in the following screenshot:

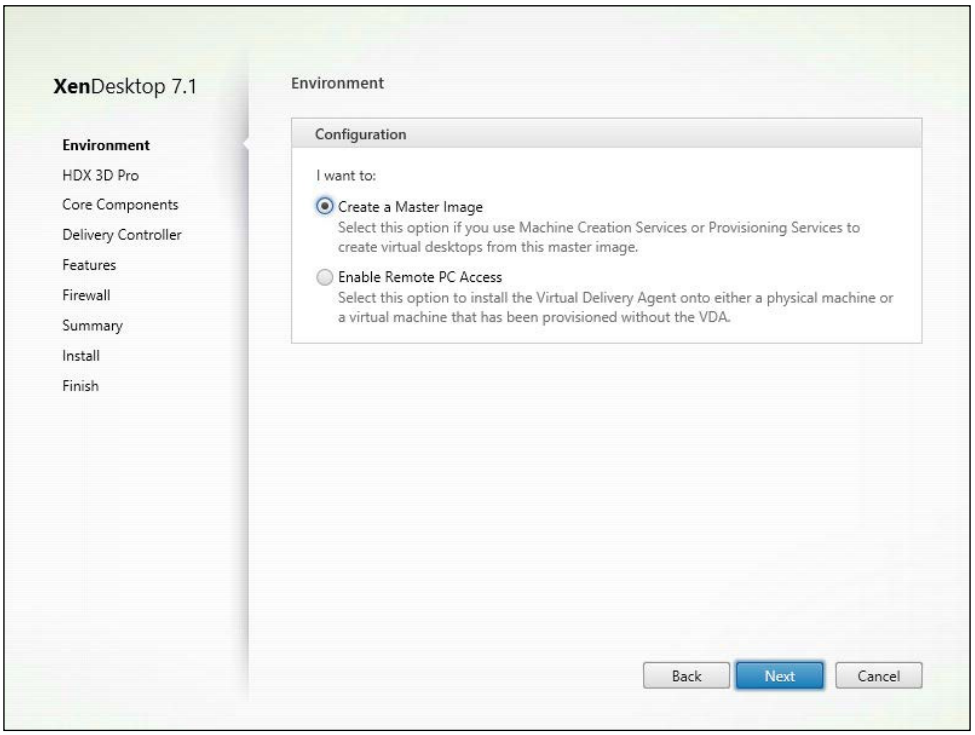

3. Install the standard VDA.
4. Install the Citrix Receiver so that users can access their documents and applications.

5. Type the FQDN in **Delivery Controller**. In our example, it is the first XenDesktop server that we built, `xd1.xenpipe.com`. Be sure to test the connection. Click on **Add** and then **Next** as shown in the following screenshot:

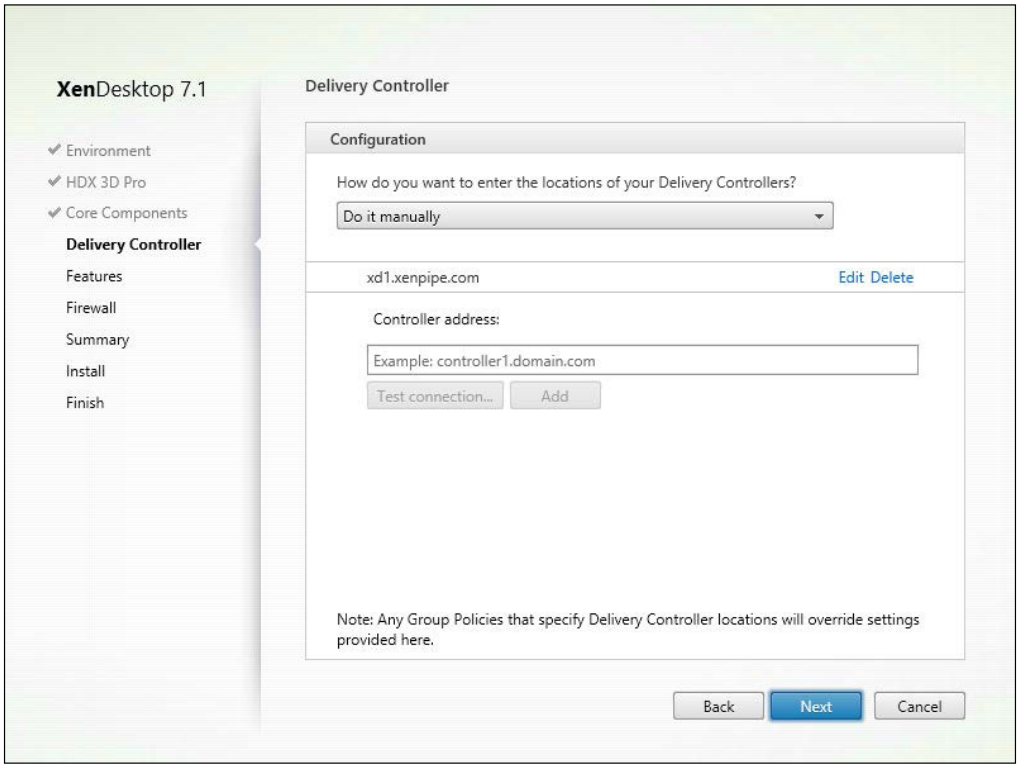

6. Select the features you want to install. We are going to enable **Personal vDisk** so that users can save their personal settings on the desktop when they log out. Select **Automatically create rules in Windows Firewall** and then click on **Next**.
7. On the install **Summary** page, review the components and select **Install**. You will need to restart the machine.
8. When the desktop has restarted, log in again to make sure that the installation has been completed.
9. When the desktop has rebooted, in preparation for creating the machine catalog, shut down the VM from the XenCenter console.

Creating the application master images

Make sure that you've successfully built the master image for the application server you want to serve applications from and the applications are installed on that server. In our example, this will be the `xd4.xenpipe.com` server. Now, you need to install the XenDesktop Delivery Agent and Citrix Receiver onto that image using the following steps:

1. Mount the XenDesktop installation media in the DVD drive in XenCenter, then log in to the master image, and launch the install media. Select **Virtual Delivery Agent for Windows Server OS** as shown in the following screenshot:

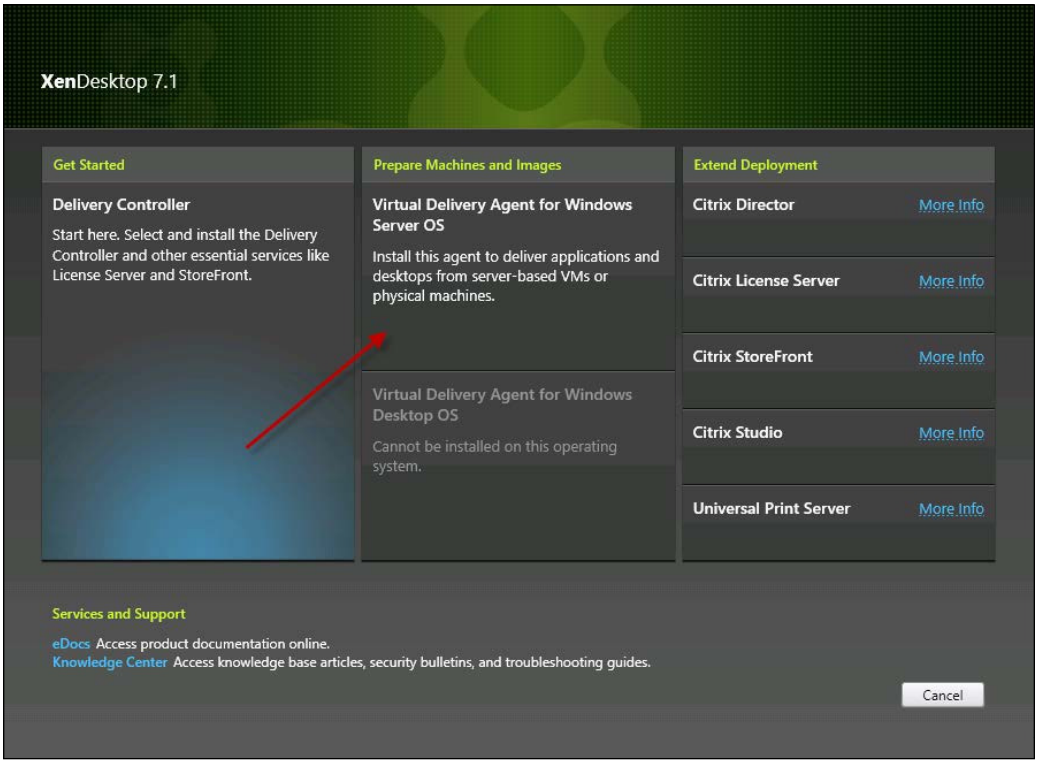

2. Select **Create a Master Image** as shown in the following screenshot:

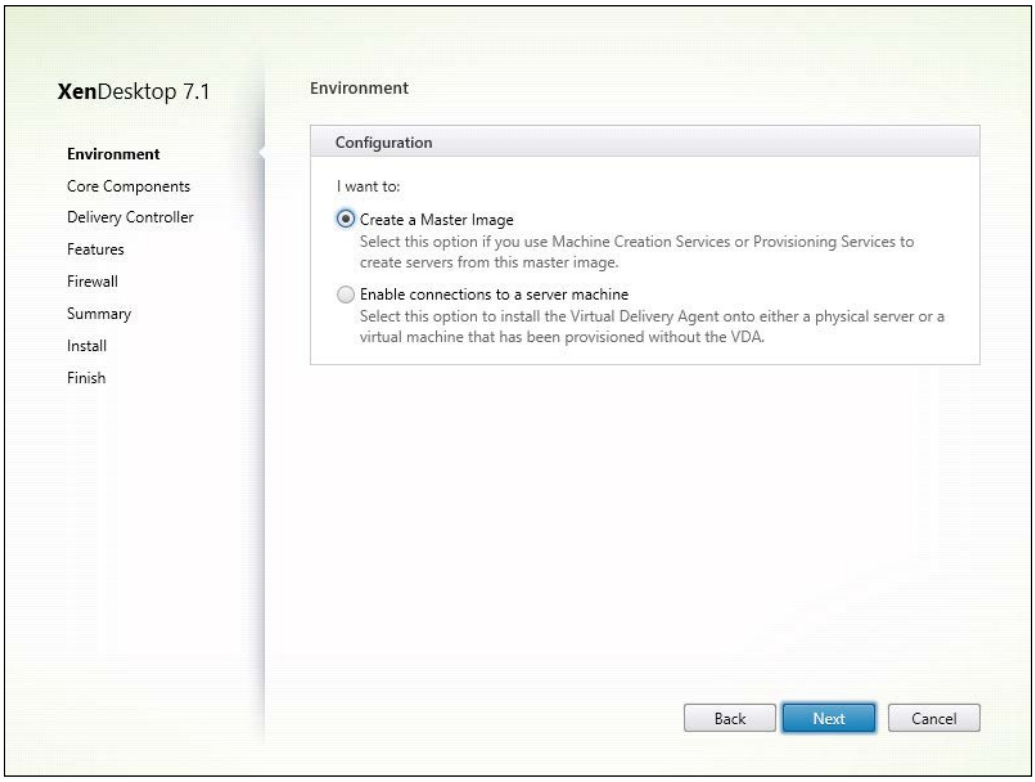

3. Install the Citrix Receiver so that users can access their documents and applications.

4. Type the FQDN in **Delivery Controller**. Be sure to test the connection. Click on **Add** and then click on **Next**, as shown in the following screenshot:

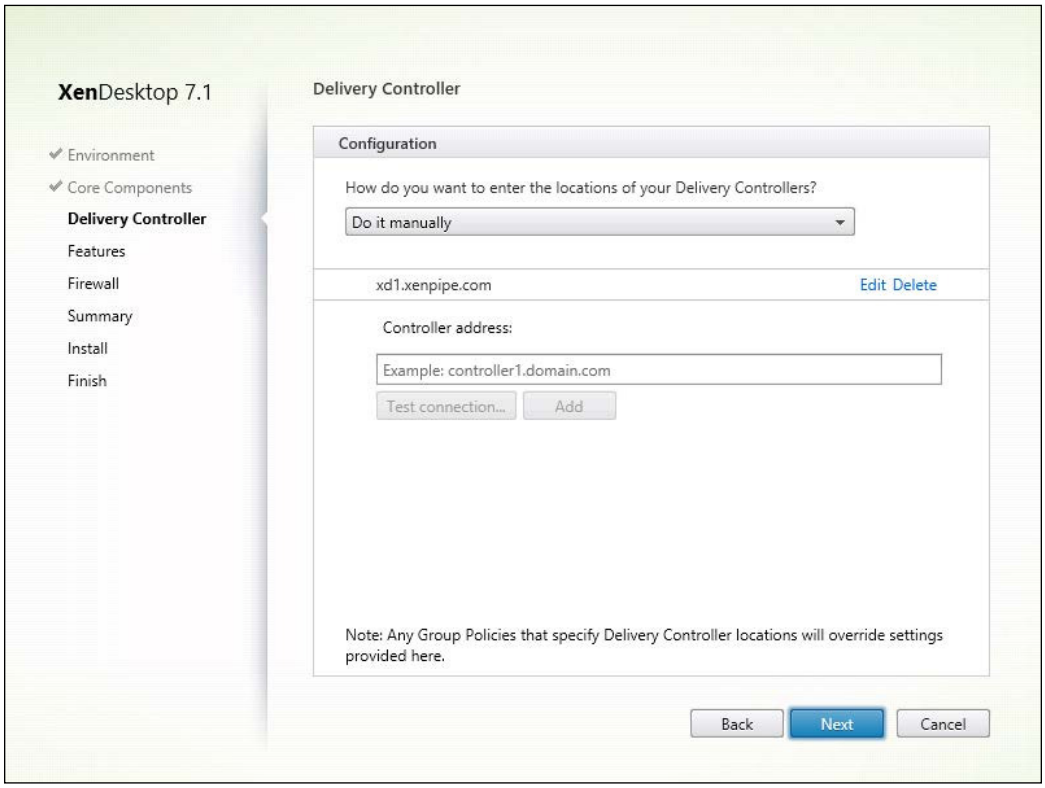

5. Select the features you want to install.
6. Select **Automatically create rules in Windows Firewall** and click on **Next**.
7. On the install **Summary** page, review the components and select **Install**. You will need to restart the machine.
8. When the desktop server has restarted, log in again to make sure that the installation is completed.
9. When the desktop server has rebooted, in preparation for creating the machine catalog, shut down the VM from the XenCenter console.

Step 6 – configuring the StoreFront server

Now that the StoreFront server, Delivery Controller, and database have been installed, we need to configure the StoreFront server to communicate with the Delivery Controller and database. The following are the steps to perform this action:

1. Log in to the second server, or `xd2.xenpipe.com` in our example, using a domain administrator account.
2. If the Citrix StoreFront management console is not already running, navigate to **Start** | **Apps** | **Citrix** | **StoreFront**.
3. Select **Create a new deployment** as shown in the following screenshot:

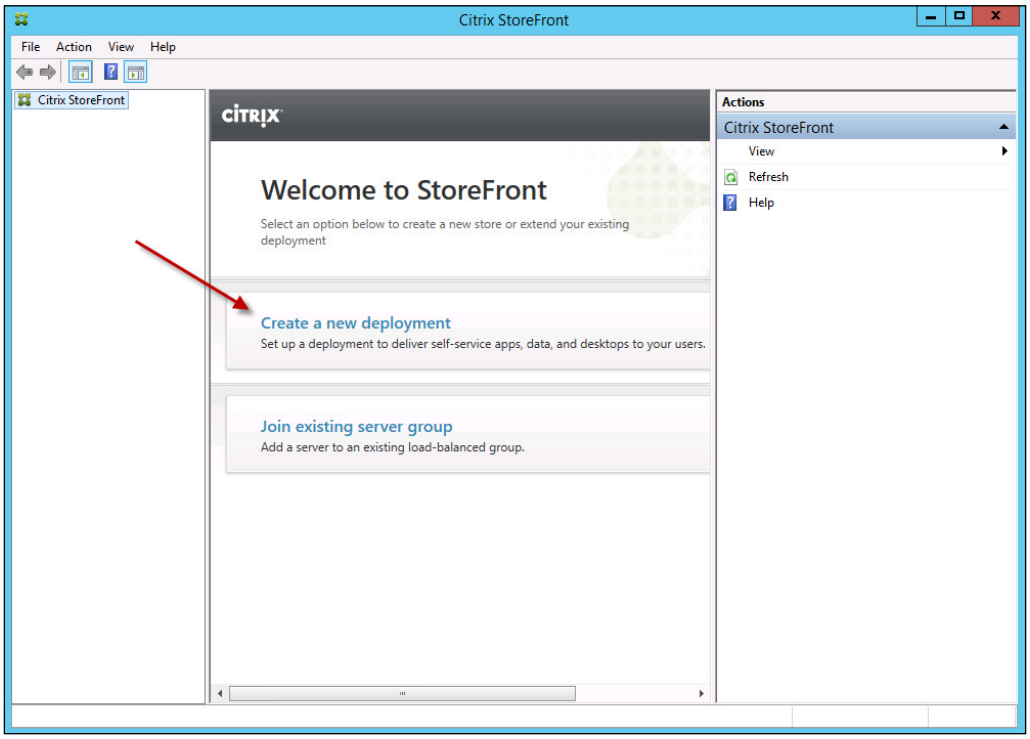

4. On the **Create a new deployment** page, select the default URL and click on **Next**. In our example, it is `https://xd2.xenpipe.com`.

You can change the base URL that users connect to with the Receiver later by navigating to **StoreFront | Server Group | Change Base URL**. For example, you may want your users to use `https://sf.xenpipe.com` in order to resemble a StoreFront vanity URL.

5. On the **Create Store** page, enter your store's name. In our example, we will use `XenStore`. Click on **Next**.
6. On the **Delivery Controllers** page, select **Add**. Type in a name for the controller and then select **XenDesktop**. For servers, click on **Add**. Enter the name or IP address of the first server, or `xd1.xenpipe.com` in our example. Select **Port 80** or **HTTP** for **Transport type**. Click on **OK** and then on **Next**, as shown in the following screenshot.

If you use **HTTPS**, be sure to install a server certificate on the Delivery Controller.

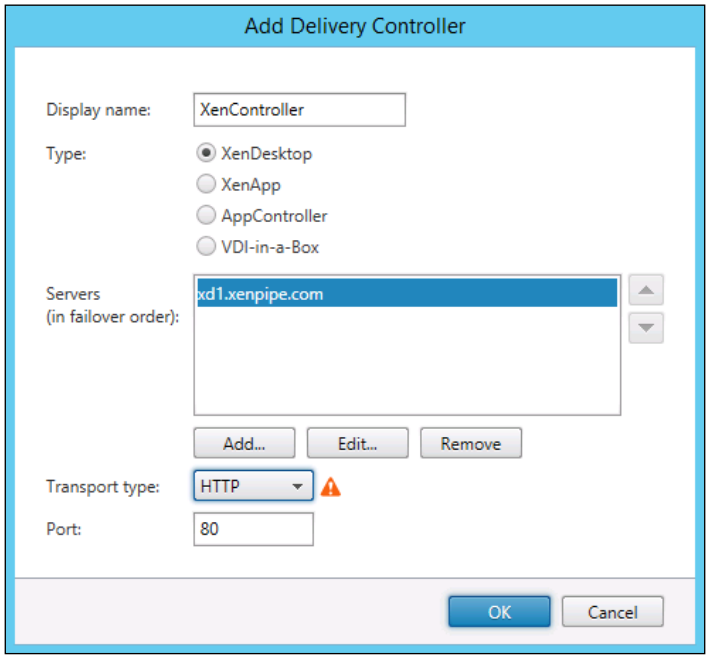

7. On the **Remote Access** page, select **None**.

8. Select **Create** and then click on **Finish**. Make a note of the website created for your store, `https://xenpipe.com/Citrix/XenStoreWeb` in our example, as shown in the following screenshot:

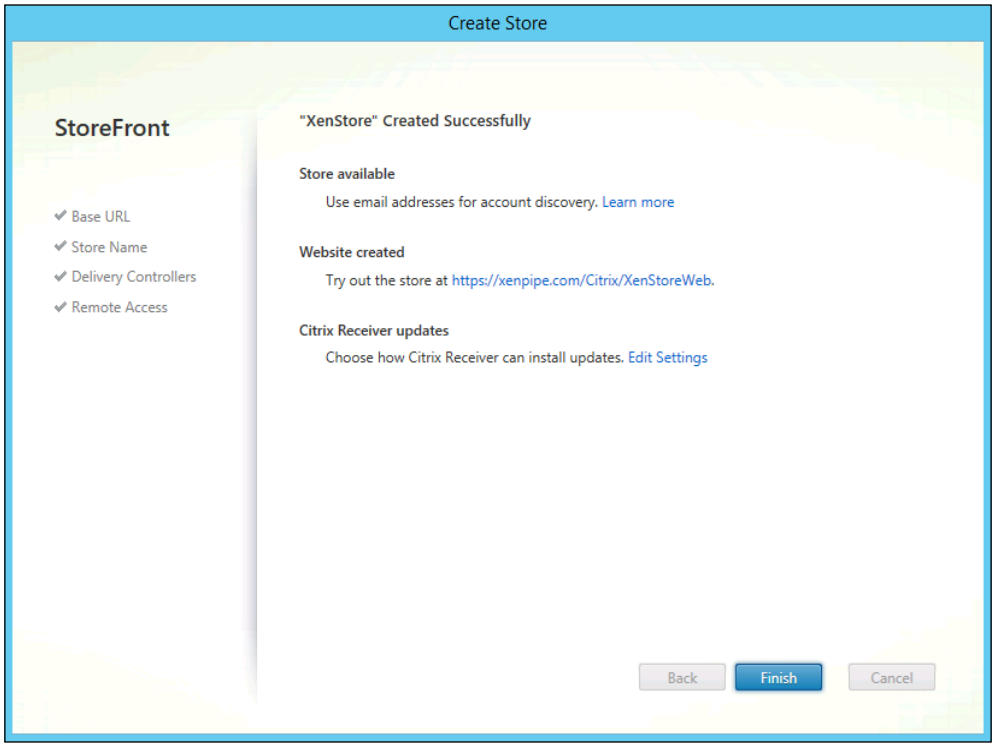

Step 7 – creating the machine catalogs

Collections of desktops or physical computers are managed as a single entity called a machine catalog. To deliver desktops and applications to users, the machine administrator creates a catalog of machines, and the assignment administrator allocates the machines from the machine catalog to the users by creating delivery groups.

Creating desktops

For this deployment example, note the following conventions:

- Machines are based on the Windows operating system.
- Desktops are delivered from virtual machines.

- Use the account naming scheme `XenUser###`. The `###` (hash) marks are placeholders for the sequential numbering of your virtual desktop machines. Using this format, your virtual machines will be named `XenUser001` through `XenUserXXX` and so on.

Until now, we have been creating the foundation of a VDI deployment that delivers virtual desktops and applications. Now, we will go about the task of creating the virtual desktops that will be delivered to the end users. The virtual desktops can be created using the following steps:

1. Make sure that the master images have been shut down.
2. Log in to server 1, or `xd1.xenpipe.com` in this example, using a domain administrator account.
3. Launch Citrix Studio by navigating to **Start** | **All Programs** | **Citrix** | **Studio**.
4. The entire deployment checklist appears on the screen. Select **Create Catalog**. When the **Create Machine Catalog** dialog box appears, click on **Next** to skip the splash screen.
5. For **Operating System and Hardware**, select **Windows Desktop OS** and then click on **Next** as shown in the following screenshot:

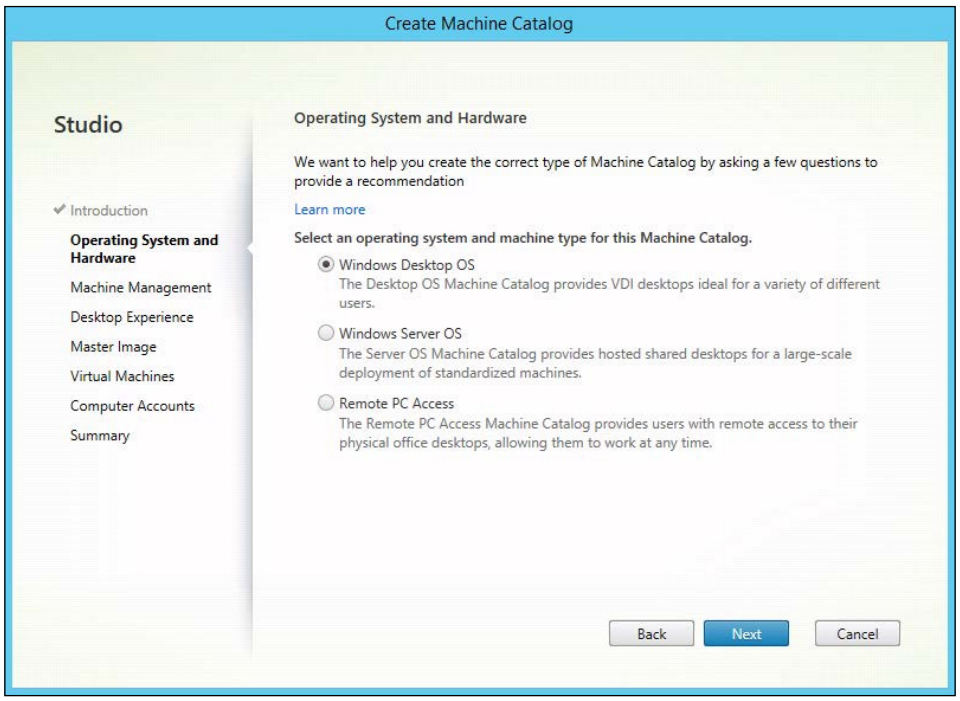

6. In **Machine Management**, select **Virtual machines** and **Machine Creation Services (MCS)** and then click on **Next** as shown in the following screenshot:

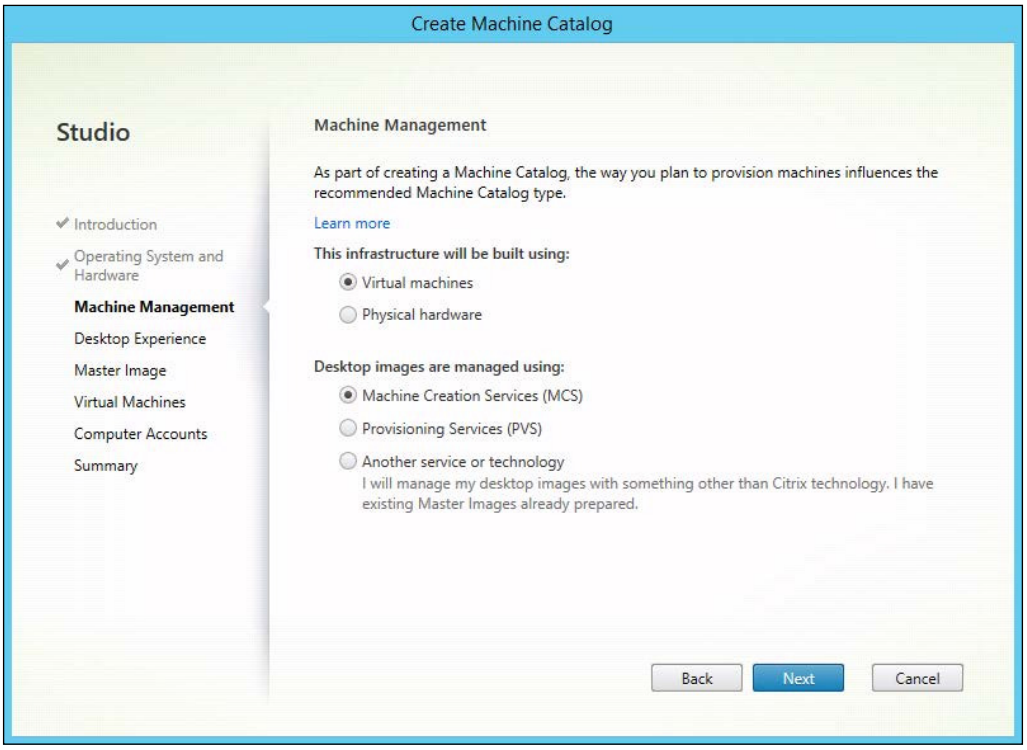

7. For **Desktop Experience**, select **I want users to connect to the same (static) desktop each time they log on** and **Yes, save changes on a separate Personal vDisk** as shown in the following screenshot:

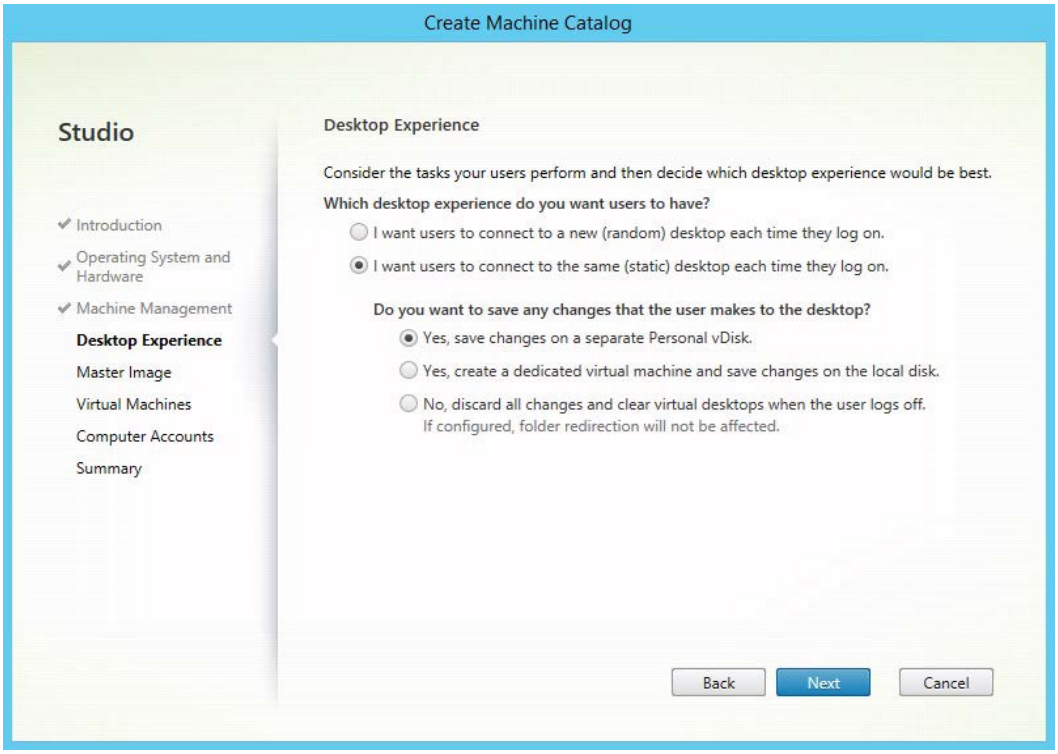

8. On the **Master Image** page, select the machine you previously created to be the master image. In our example, it is **Windows 8**.

9. On the add and configure **Virtual Machines** page, select the number of virtual machines (desktops) that have to be created, the vCPUs, memory, and disk storage for each desktop, as shown in the following screenshot:

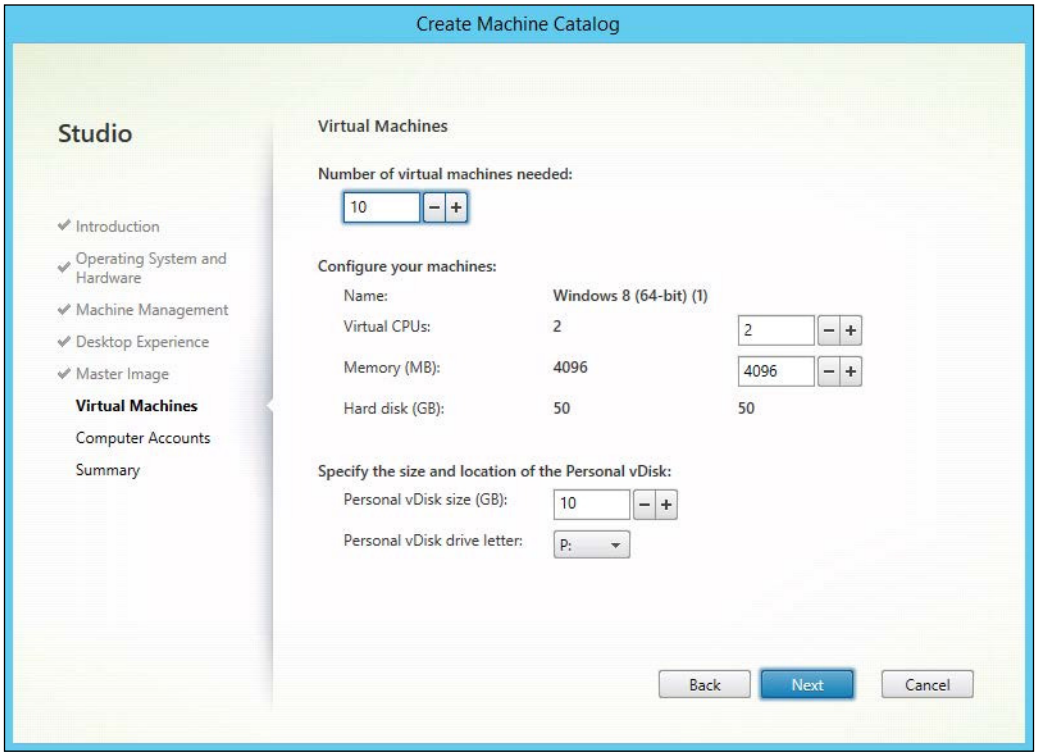

10. On the **Active Directory Computer Accounts** page, select **Create new Active Directory accounts**, enter the correct domain (**xenpipe.com** in our example), then select the **Default OU** option as the location, and enter `XenUser###` in the **Account naming scheme** field, as shown in the following screenshot:

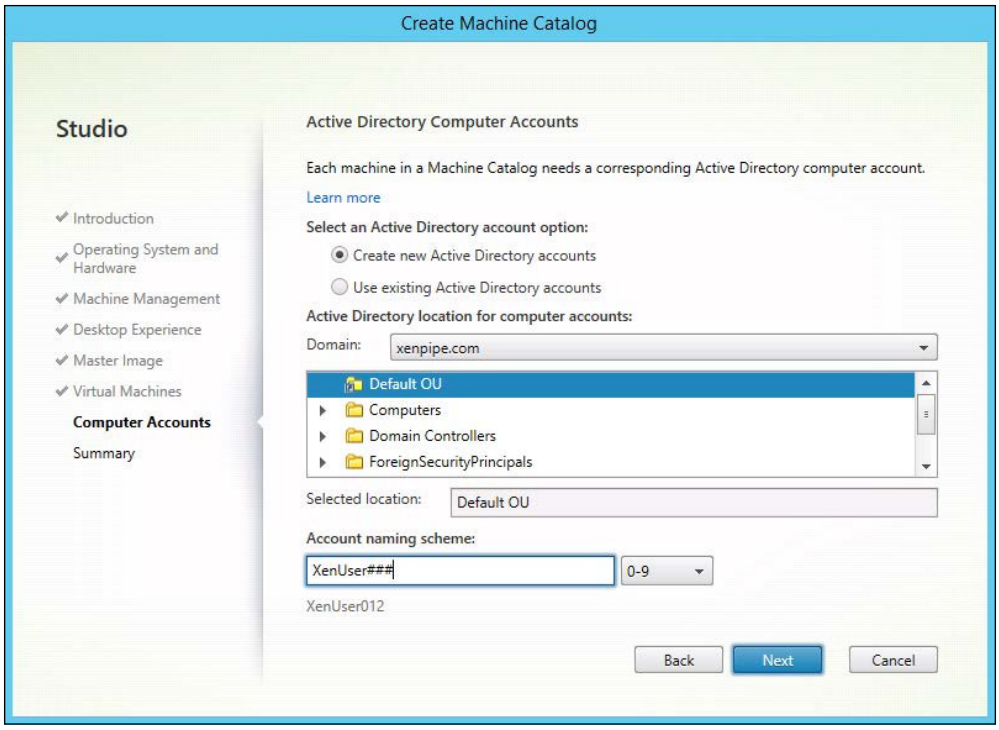

11. For **Machine Catalog Name**, enter a name and description.
12. Go through the **Summary** page and then select **Finish**. When complete, the machine catalog will be created, accounts will be created in Active Directory, and the desktop machines will be created in XenCenter. Please note that this might take a few minutes.

Check to make sure that the virtual machines have been started in XenCenter. If not, start them manually in XenCenter before moving on.

Creating the application servers

For this deployment example, please note the following conventions:

- Machines are based on the Windows operating system
- Applications are delivered from virtual machines

The VDI foundation is built and the virtual desktops are ready; now we need to add the virtual applications that the users will use. For this, perform the following steps:

1. Make sure that the master images have been shut down.
2. Log in to server 1, or `xd1.xenpipe.com` in this example, using a domain administrator account.
3. Launch Citrix Studio by navigating to **Start** | **All Programs** | **Citrix** | **Studio**.
4. In the left-hand navigation menu, select **Machine Catalogs**. Select **Create Machine Catalog**. When the **Create Machine Catalog** dialog box appears, click on **Next** to skip the splash screen.
5. For **Operating System and Hardware**, select **Windows Server OS** and then click on **Next** as shown in the following screenshot:

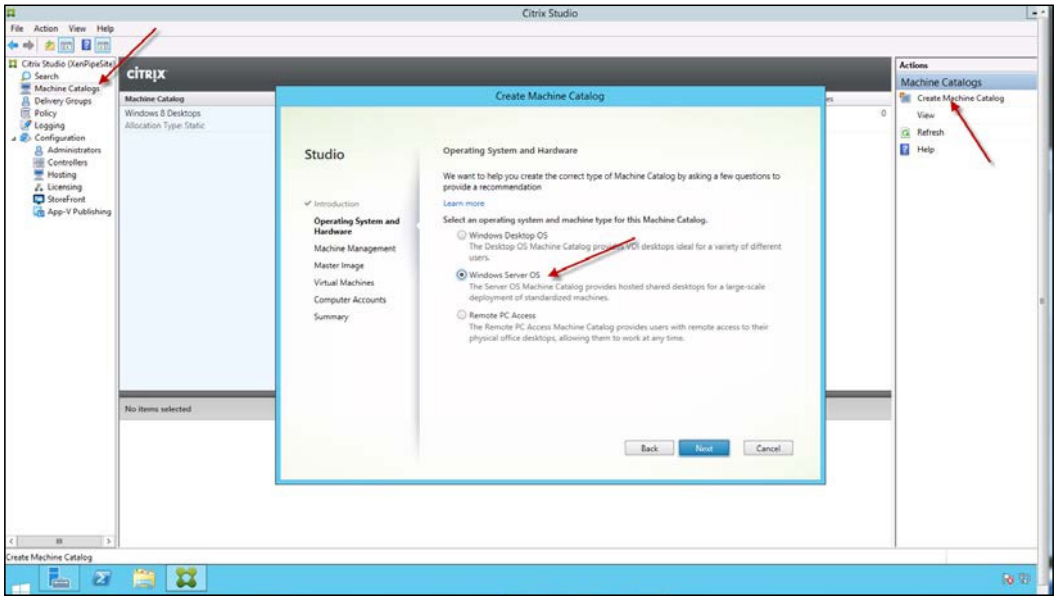

6. In **Machine Management**, select **Virtual machines** and **Machine Creation Services (MCS)** and then click on **Next** as shown in the following screenshot:

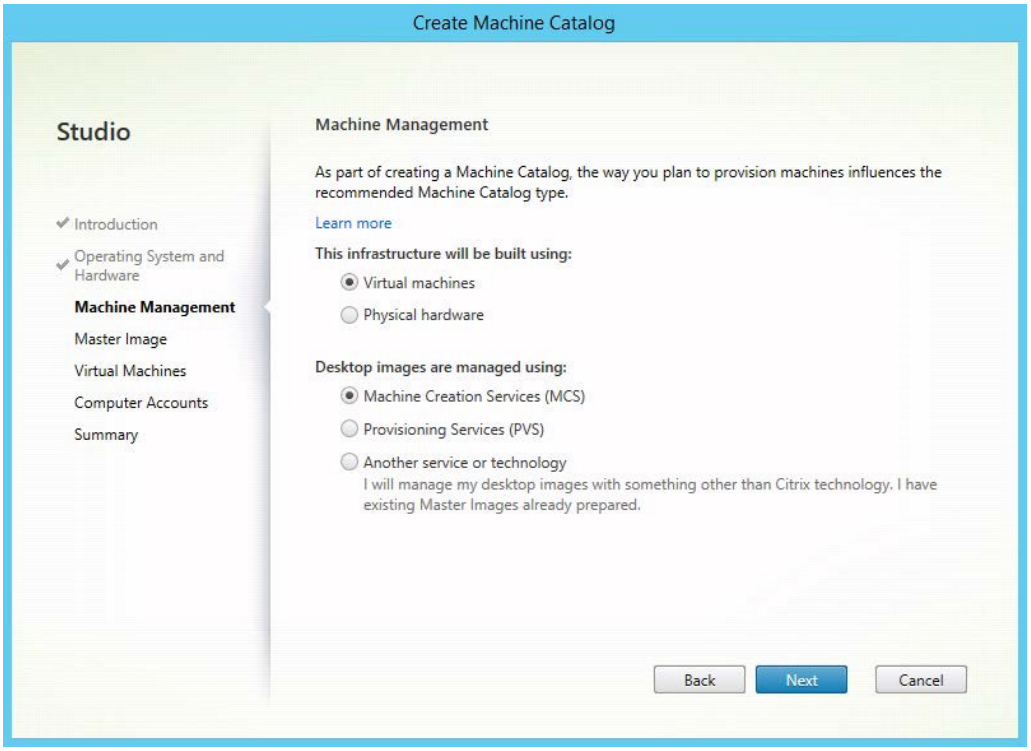

7. On the **Master Image** page, select the machine you created to be the master image for the applications.

8. On the add and configure **Virtual Machines** page, select the number of virtual machines (application servers) to be created, the vCPUs, memory, and disk storage for each. We only need one server to deliver apps at the moment, as shown in the following screenshot:

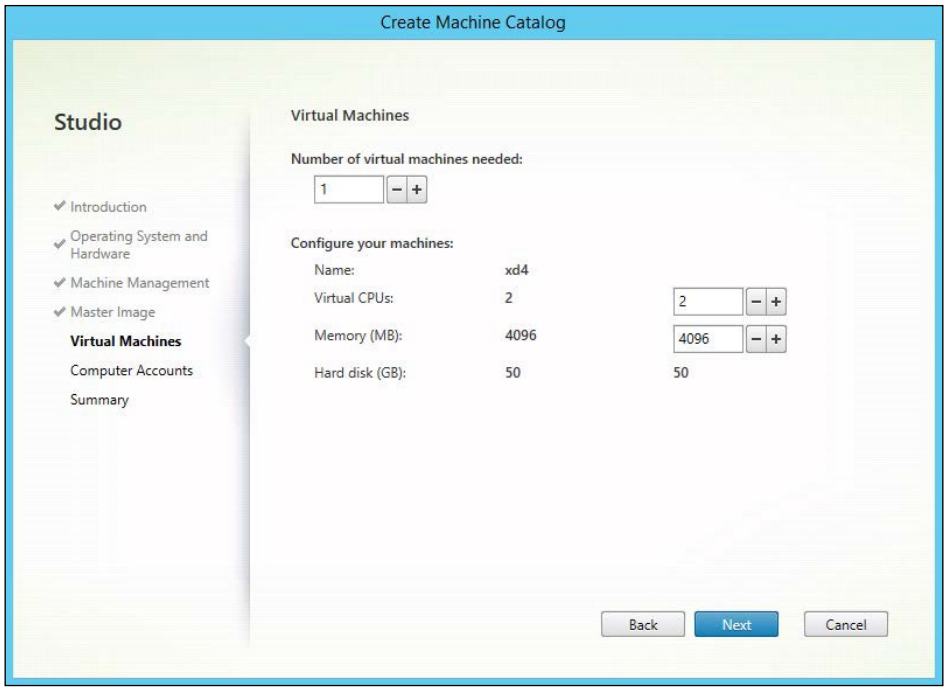

9. On the **Active Directory Computer Accounts** page, select **Create new Active Directory accounts**, enter the correct domain (**xenpipe.com** in our example), then select the **Default OU** option as the location, and enter `AppServer###` in the **Account naming scheme** field, as shown in the following screenshot:

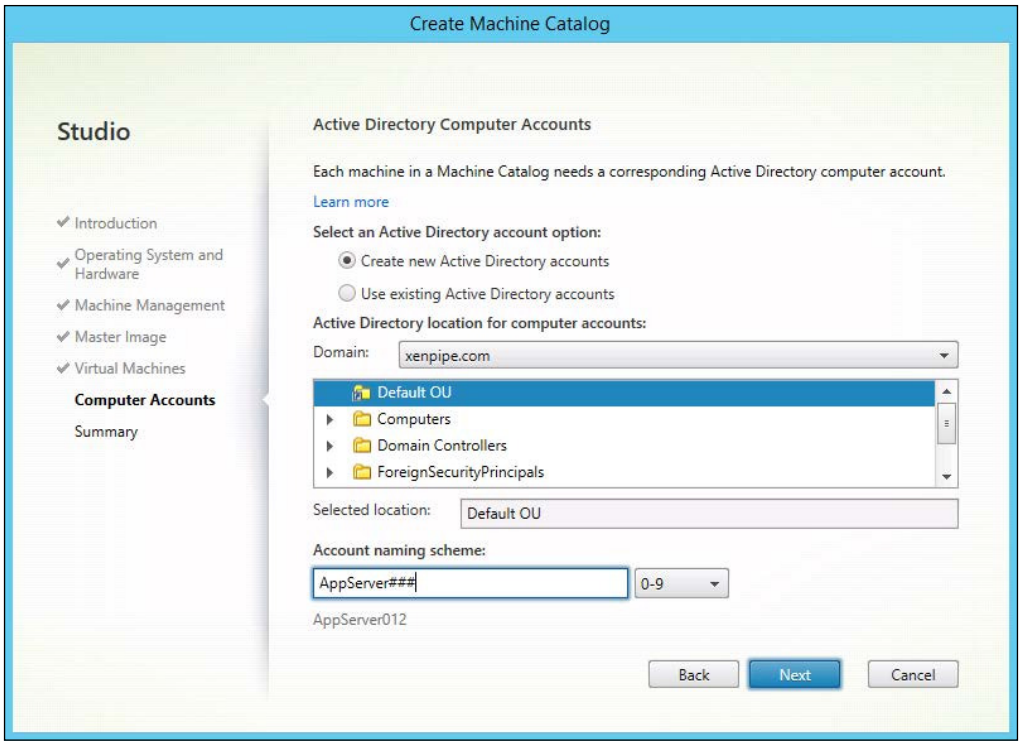

10. For **Machine Catalog Name**, enter a name and description.
11. Go through the **Summary** page and then select **Finish**. When complete, the machine catalog will be created, accounts will be created in Active Directory, and the application machine(s) will be created in XenCenter. Please note that this might take a few minutes.

Check to make sure that the virtual machines have been started in XenCenter. If not, start them manually in XenCenter before moving on.

Step 8 – creating the delivery groups

Now that you have configured StoreFront and created a machine catalog, you can create the delivery groups.

Creating desktop delivery groups

In this task, we will assign users to a delivery group so that they can receive their desktop. The following steps are used to create the desktop delivery groups:

1. Log in to the Delivery Controller, or `xd1.xenpipe.com` in our example, using a domain administrator account.
2. Launch Studio and navigate to **Delivery Groups** in the left-hand side navigation pane. Select **Create Delivery Group**.
3. On the **Machines** page, select the catalog that we created earlier for Windows 8 desktops. We enter five machines to be assigned, as shown in the following screenshot.

The add machines value is the number of machines that the delivery group consumes. This value should be less than or equal to the number of available machines. For example, we created 10 virtual desktops in **Machine Creation Services**, but we only assigned five to this delivery group.

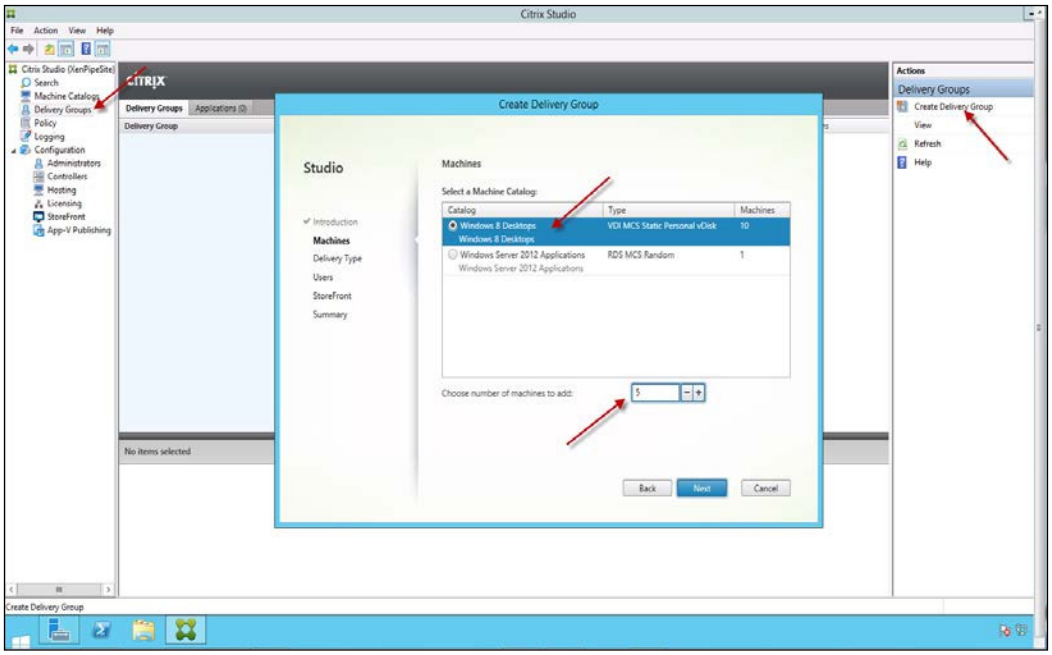

4. For delivery mode, select **Desktops** and then click on **Next**.

5. Add users.

Here, we use the **Domain Users** group in Active Directory; that is, `XENPIPE\Domain Users`.

6. Add the StoreFront server as shown in the following screenshot:

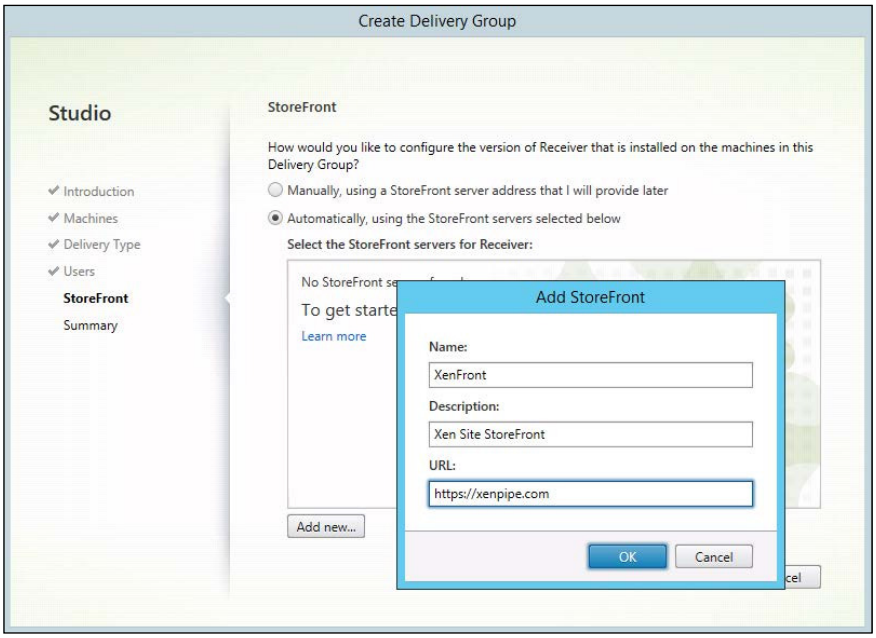

7. Enter a delivery group name and description and then click on **Finish**.

Creating the application delivery groups

In the previous section, we created a delivery group to deliver desktops to users. In the following steps, we will create a delivery group that provides hosted applications to users:

1. Log in to the Delivery Controller, `xd1.citrix.com`, using a domain administrator account.
2. Launch Studio.

3. Select **Delivery Groups** in the left-hand navigation menu, and then select **Create Delivery Group** in the right-hand navigation menu, as shown in the following screenshot:

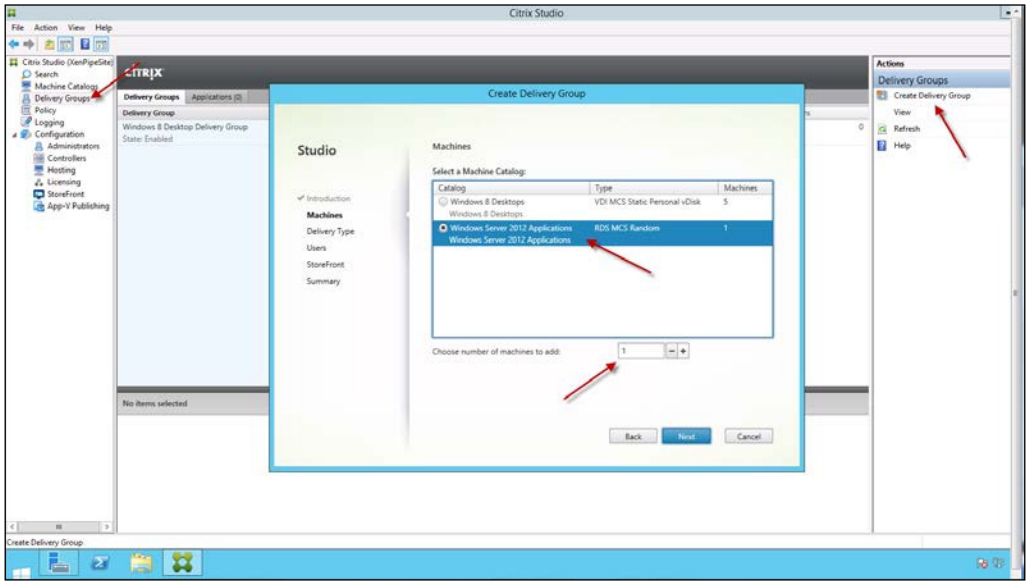

4. Choose the number of machines to be added from the total available machines. In our case, we will add one machine.

5. Select **Applications** to deliver applications to users. Then, click on **Next**, as shown in the following screenshot:

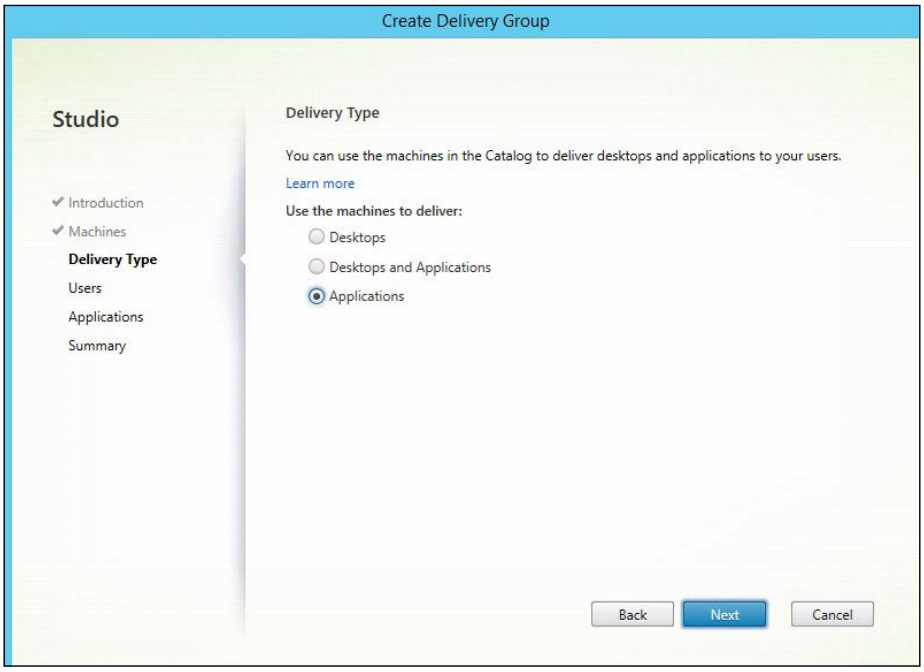

6. Add users. In our example, it is `XENPIPE\Domain Users`.
7. Select the applications that you want to deliver to your end users.

You may need to wait for the machine to start before selecting the application's launch location on the virtual machine.

8. You can select the applications by browsing, or you can add them manually. Select **OK** and then click on **Next**, as shown in the following screenshot:

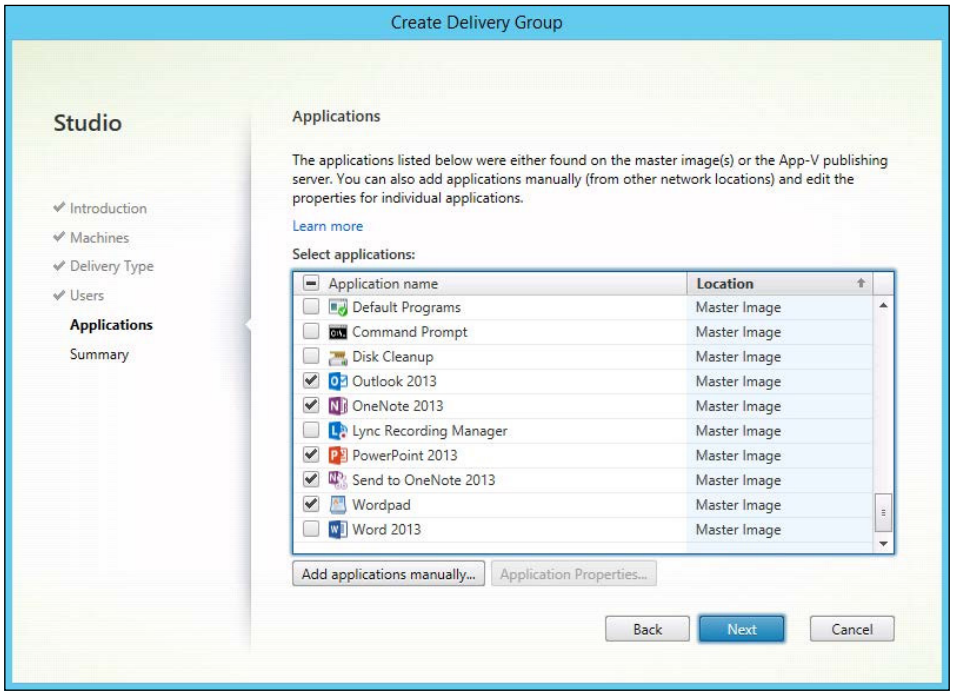

9. Enter a delivery group name and description and then click on **Finish**.

Installation checkpoint

All of the server-side stuff is done. This is a good time for an installation checkpoint. The following steps will help:

1. You should now have two machine catalogs—one for desktops and one for applications—as shown in the following screenshot:

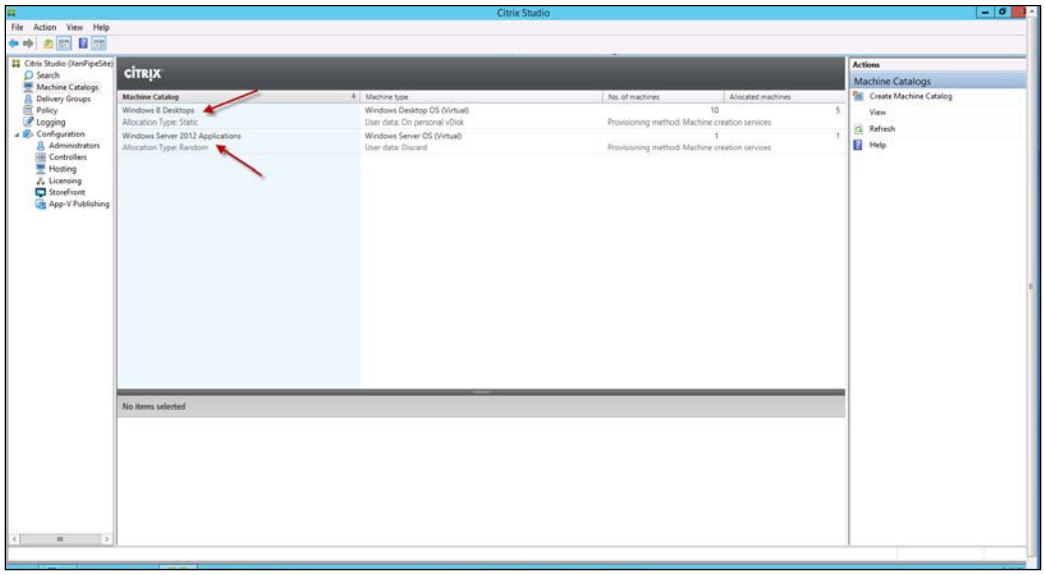

2. You should have two delivery groups—one for desktops and one for applications—as shown in the following screenshot:

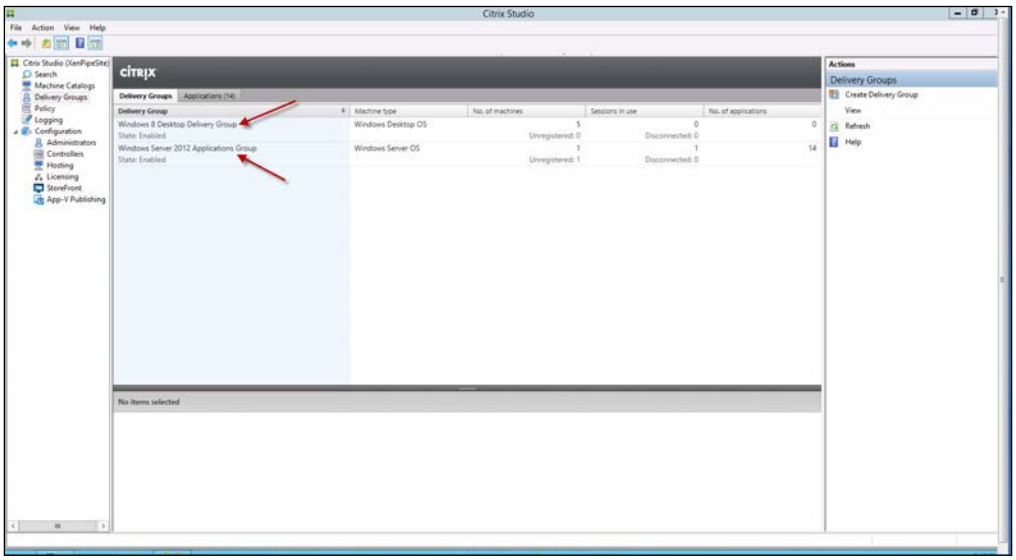

3. You should see the applications listed under the **Applications** tab as shown in the following screenshot:

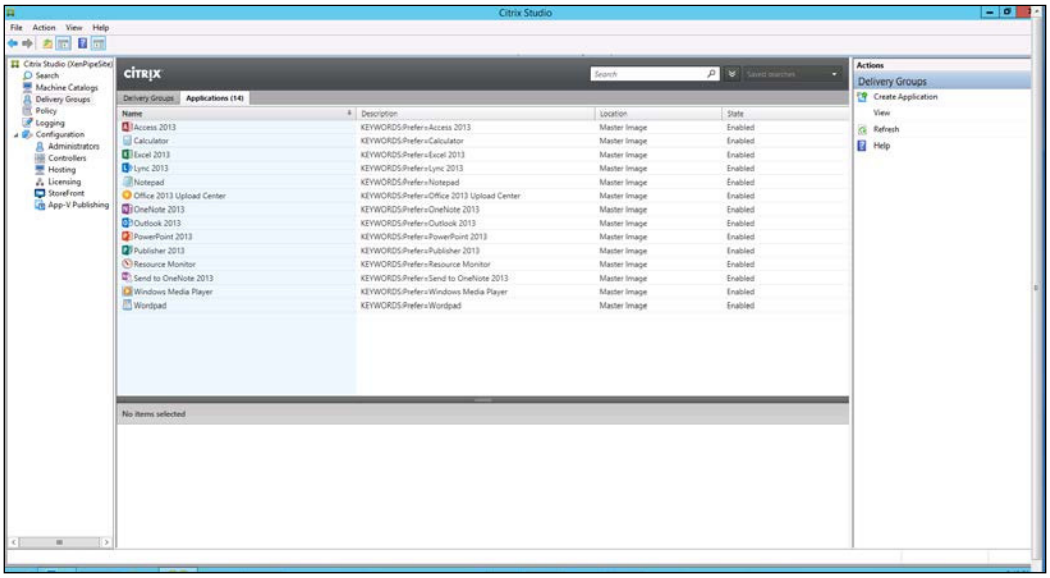

4. In the XenCenter (Hypervisor) management console, you should see the following machines:

Installation checkpoint

No.	Virtual machine	Use case
1	dc	Domain controller, AD, DHCP, and DNS
2	XD1	XenDesktop controller
3	XD2	XenDesktop StoreFront
4	XD3	XenDesktop Director
5	AppServer001	The application's server
6	XenUser 001 to 00X	Windows desktops
7	Win2012-AppsMaster	Office 2013 application server master
8	Win8-Master	Windows 8 desktop master

The virtual machines present in the XenCenter (Hypervisor) management console are shown in the following screenshot:

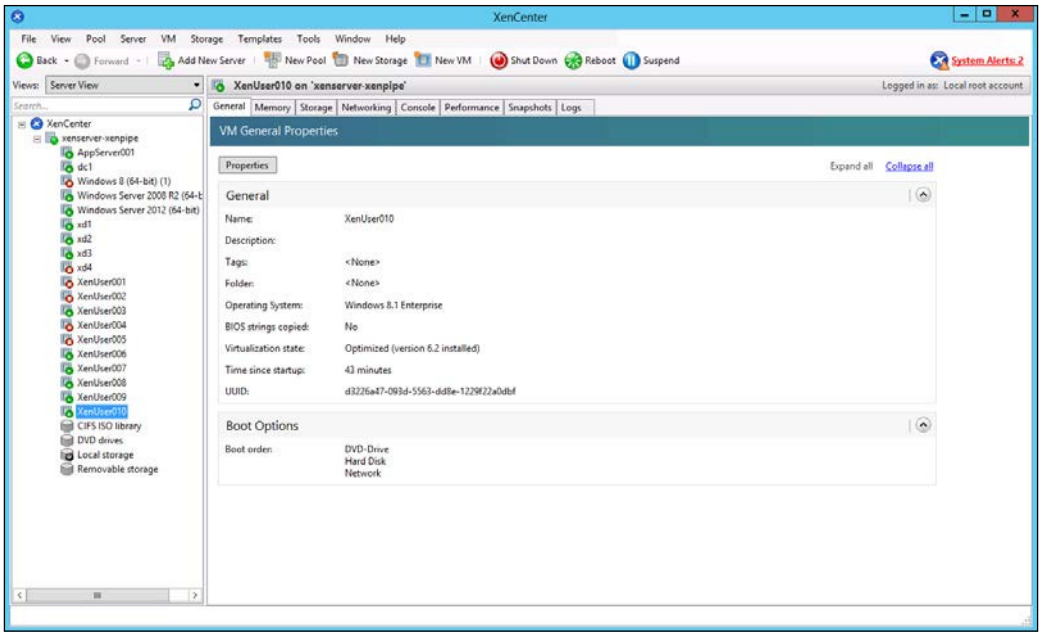

Step 9 – installing Citrix Receiver™ on the client devices

Now that the virtual desktop and application's infrastructure have been built, you need to access these from the StoreFront server using the Citrix Receiver on the client. This can be done using the following steps:

1. Log in to the client device with administrator privileges.

This is the end user device, such as a laptop, workstation, thin client, or even a mobile device. This is *not* the virtual desktop or master images we created in the previous steps.

2. Mount the XenDesktop installation media.

You can use XenDesktop `.iso` to burn a DVD to mount in the client drive; however, I find it is easier to mount the `.iso` image on the machine using Virtual CloneDrive, `http://www.slysoft.com/en/virtual-clonedrive.html`. It is also a common practice to configure client deployment from the StoreFront console and not the installation media—something you may want to try after this chapter.

3. Navigate to the Citrix Receiver and plugins directory.
4. Select your platform; for example, Windows.
5. Navigate to the Receiver directory and launch `CitrixReceiver.exe`. Select **Install** as shown in the following screenshot:

6. Click on **Add Account** and type in the URL to the StoreFront server. In our example, it is `https://xd2.xenpipe.com`.

Remember that in step 6, you can change the base URL that the users connect to in the XenDesktop StoreFront server in the StoreFront application. For example, if you want your users to connect to a vanity URL such as `https://sf.xenpipe.com`, just make sure your certificates match the URL, which is why we use a wildcard certificate, `*.xenpipe.com`.

7. Provide user credentials and connect.

Step 10 – testing the connection

You can test the environment by connecting to the StoreFront server with a web browser or by logging in to the client device and launching the Receiver.

Testing the desktops

Connect to the StoreFront server using a web browser and log in. You should see a list of desktops that are available to you in the **Desktops** section as shown in the following screenshot:

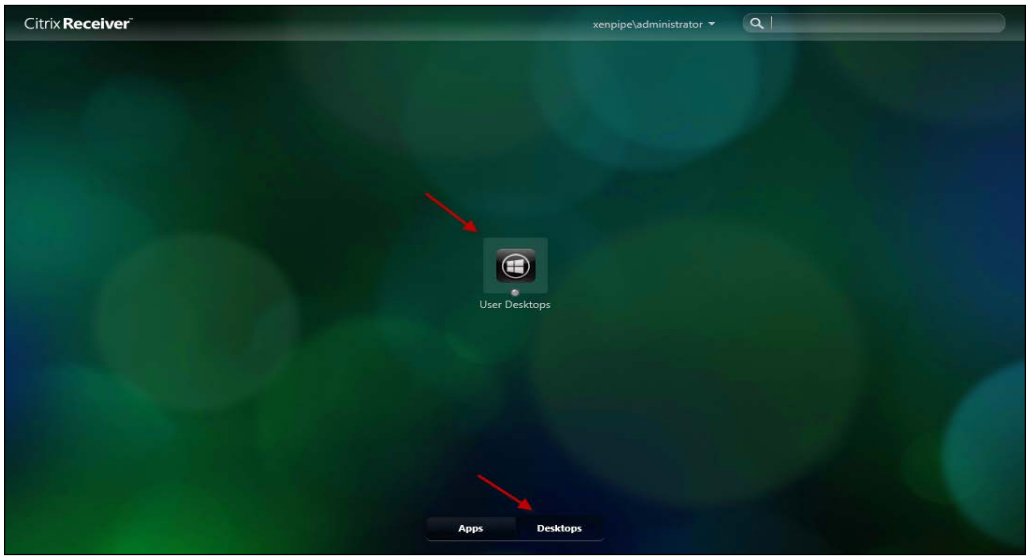

Testing the applications

Connect to the StoreFront server using a web browser and log in. You should see a list of the applications available to you in the **Apps** section. Click on the **+** sign on the left-hand side of the window, and add the applications that you want to use to your dashboard.

Click on the **Application** icon; it is automatically launched and run from the server and is delivered to your device over the ICA protocol using the Citrix Receiver, as shown in the following screenshot:

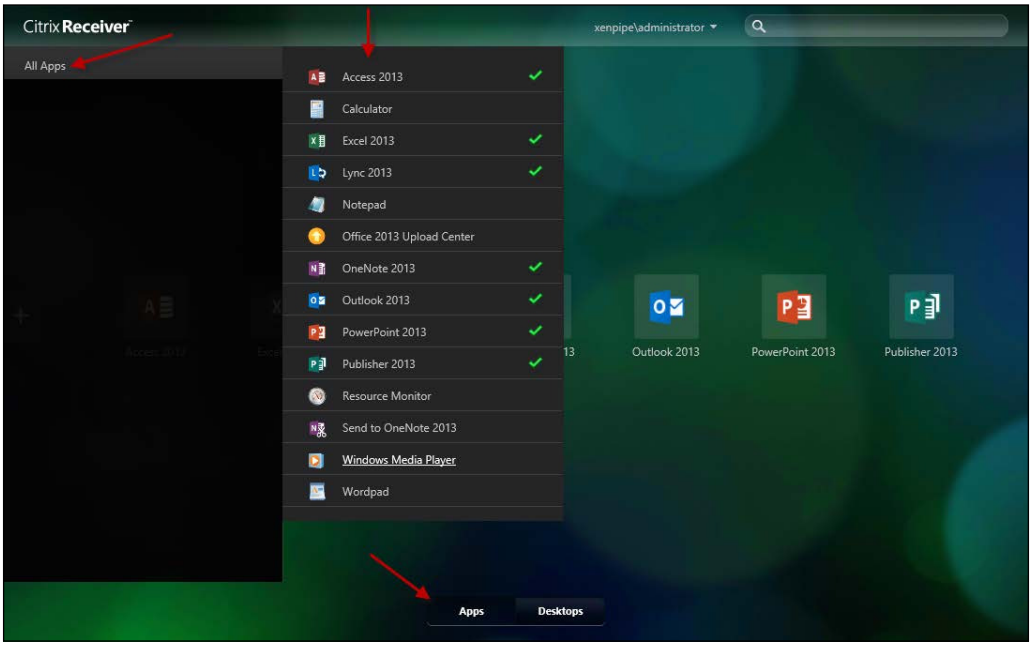

If you launch the Receiver and log in but still don't see your desktop, there is a little trick that you can do. In the search field, in the upper-right corner of the screen, type `Desktop` or the name of the desktop group you created. Then, you can add it to your Receiver home screen as shown in the following screenshot:

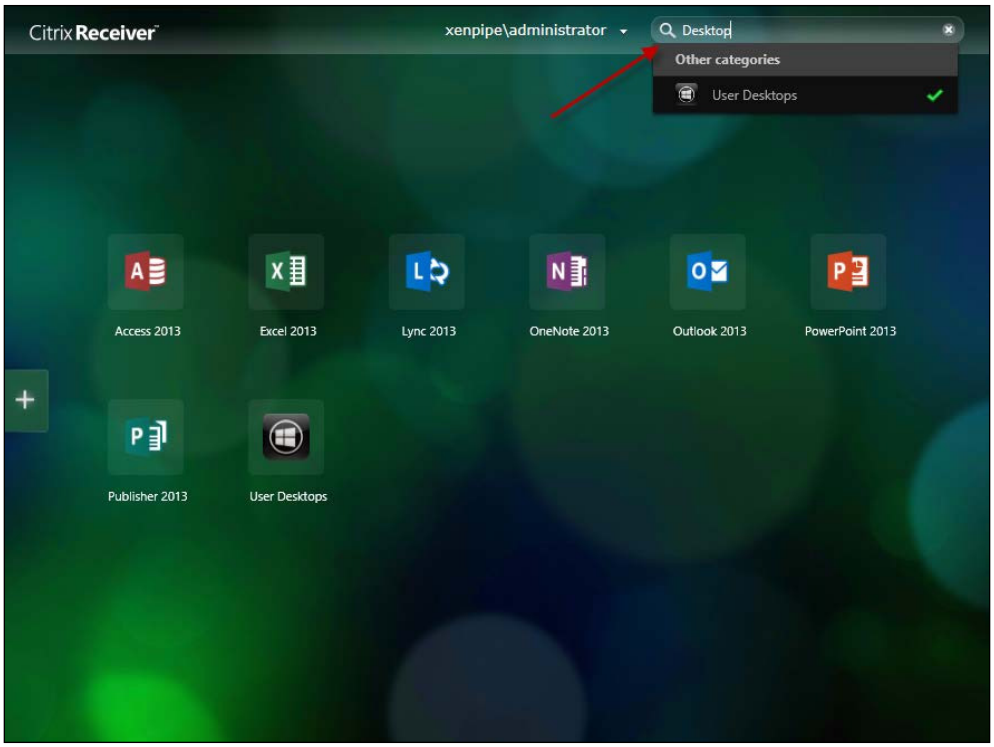

To see a working demo of the preceding interface, you can view the video at `http://www.citrix.com/tv/#videos/8041`.

Summary

I hope that you found this chapter useful as a roadmap in order to get your first XenDesktop deployment installed and running. I'm sure that by following this chapter, you were able to install XenDesktop in a much faster timeframe than if you did not have the guidance.

Now that you have gone through the steps to create a basic desktop and application delivery Site, you need to understand how to manage them. You are now familiar with the machine catalog and desktop delivery group wizards. In the next chapter, we will expand on the options available in these tools.

3

Managing Machine Catalogs, Hosts, and Personal vDisks

Now that you have gone through the steps to create a basic desktop and application delivery Site, you need to understand how to manage it. In the previous chapter, you have already used the tools needed to manage XenDesktop. In fact, you are now familiar with the machine catalog and desktop delivery group wizards. We will expand on the options available in these tools so that you are able to understand a wider range of options to create your XenDesktop Site. In this chapter, we will cover the following management topics of XenDesktop:

- Creating machine catalogs
- Managing machine catalogs
- Managing hosts (Hypervisors)
- Managing Personal vDisks

In XenDesktop 7.x, groups of virtual or physical desktops can be managed as an entity called a machine catalog. The administrator delivers the desktops to users by creating a machine catalog and then allocating the machines in the catalog to the users through the delivery groups. We will discuss delivery groups in the next chapter.

Machine catalogs

A machine catalog is a group of computers or desktops that define the hosting infrastructure for those desktops and applications and the level of control that the users have over their environments.

Prerequisites

The following are the prerequisites to create a machine catalog:

- To create a machine catalog, you first need to prepare the master images similar to what we did in *Chapter 2*, *Installing XenDesktop®*.
- The master image should contain the elements that are common to all users, such as antivirus software, Citrix Receiver and plugins, and other default programs.
- You also need the Active Directory computer accounts to assign to each machine. The machine catalog creation wizard creates these for you.

Creating the master images

To prepare desktops and applications for machines in Windows Server or Windows desktop machine catalogs, you must prepare the master image that is used as a base or template to create user desktops and applications.

You can create these master images the same way you would install any other Windows VM. Keep the following checklist in mind:

The master image checklist		
Item	**Description**	**Check**
1	Install all service packs and updates.	
2	Allocate enough hard disk space required for the user's desktops and applications—this cannot be changed later.	
3	Install the Hypervisor tools such as XenServer Tools, Hyper-V Integration Services, or VMware tools.	
4	Install the XenDesktop Virtual Delivery Agent and optimize the desktop. The Citrix Receiver is automatically installed during this process.	
5	Install third-party tools such as antivirus software.	
6	Configure the Windows updates.	
7	Install applications that won't be virtualized (delivered from a XenDesktop or XenApp server through the Receiver).	
8	Join the computer to the domain.	
9	Create a snapshot of the master image.	
10	Shut down the master image.	

Adding and configuring the virtual machines

This is where you provide details on how to use the master image that you created previously in order to create many VMs that will be used as VMs for desktops or applications. For example, if you want to create desktops for 10 users, you should specify the following attributes:

The virtual machine checklist		
Item	**Description**	**Check**
1	**Number of virtual machines needed**	`10`
2	**Virtual CPUs**	`2`
3	**Memory (MB)**	`4096`
4	**Hard disk (GB)**	`50`
5	**Personal vDisk size (GB)**	`10`
6	**Drive letter**	`P:`

Creating the computer accounts

As in the physical computer world, in the virtual computer world too, you need to create computer accounts for the VMs created previously. Every machine in a machine catalog must have an Active Directory computer account. If you are going to use static or random machine catalogs and have access to the Active Directory domain administrator account, you can make XenDesktop create new accounts when a machine catalog is created. If you don't have the required permissions, then you need to ensure that you have an adequate number of unused Active Directory computer accounts available before you begin creating machine catalogs.

You can use the existing computer accounts by browsing Active Directory when you create a machine catalog or you can import a list of computer account names in the `.csv` file format. XenDesktop uses the following format for computer accounts imported from the `.csv` files:

```
[ADComputerAccount]
ADcomputeraccountname.domain
...
```

For machine catalogs that already exist, you simply select or import the computer accounts and assign each virtual or physical machine to an Active Directory computer account and an Active Directory user account. If you are using streamed machines, then you need to use the Provisioning Services and existing Microsoft Active Directory tools to manage the devices and machines.

Creating a machine catalog

Machine catalogs allow the administrator to manage the users' desktops and applications. This is where you store your master images and tie them to Active Directory.

To create a machine catalog, use the following steps:

1. Log in to the machine that has Studio installed on it.
2. Launch Studio.
3. Select **Machine Catalogs** in the left-hand side navigation menu.
4. Select **Create Machine Catalog** in the right-hand side navigation menu as shown in the following screenshot:

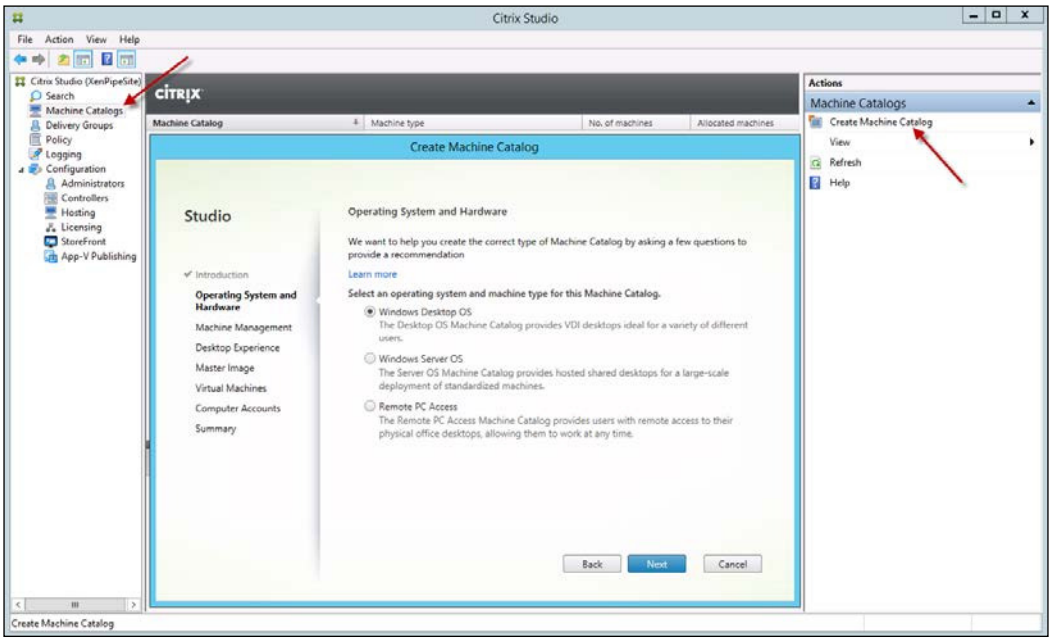

Operating systems and hardware

When you set up machine catalogs in Studio, you can create them for Windows Server and/or Windows desktops.

Windows desktop

You would use the Windows desktop type of machine to allocate shared desktops to users when they first log on. This machine type allows personalization, uses the Microsoft Windows desktop operating systems such as Windows 7 or 8, and has the following characteristics:

- Used for hosted VDI—dedicated or pooled—where each user gets their own desktop
- Used for delivering static or pooled desktops
- Users can personalize their desktops and applications and store data to a Personal vDisk
- Windows desktop can deliver desktops or applications

This is good for performance users who need their own desktops and the ability to save personal preferences. This is typical of a knowledge worker or someone who would use a desktop and applications extensively. Users with **mobile workstyles** will benefit from this deployment model as they can access the same desktop from a wide mixture of device types. It is also useful for users who need to install their own applications. You can standardize the configurations of the users' desktops by using a common template—a master image. Desktops are delivered to multiple device types regardless of the hardware compatibility, which reduces the cost because you provide users with a single personalized desktop experience. This has huge implications for administrators as it eases management and troubleshooting.

Windows Server

You would use the Windows Server type of machine to deliver a server desktop or the applications hosted on a server. There are no personalization options for this machine type. These machine catalogs reside on the machines with Microsoft Windows Server operating systems, such as Server 2012, and have the following characteristics:

- Used for the hosted shared delivery of server desktops or server-hosted applications that multiple users connect to and share
- Mostly used to deliver applications
- Can be used to deliver a shared server desktop
- Assigns machines to users at logon, based on load
- Assigned to users on a first-come, first-serve basis
- Standardized machines for large-scale sharing
- Delivers desktops or applications hosted on a server

This type of machine has many uses. It is good for task workers who perform specific tasks and don't need personalization or applications installed on their desktops. It is also useful for power users if they need a lot of personalization and a lot of applications at their disposal. This machine type also provides strict security control as it can restrict users from installing anything or making permanent changes, reducing cost by providing a locked-down standard machine for all the users in the catalog.

Remote PC Access

New in XenDesktop 7.x is the ability to use the `GoToMyPC` functionality in order to allow users to remotely connect to their physical office desktops.

You can read more about this feature at `http://support.citrix.com/proddocs/topic/xendesktop-71/cds-remotepc-plan.html`.

Machine management

There are different types of machines that can be defined as follows:

- **Virtual**: The virtual machine type is for machine catalogs that are based on machines that are power-managed through XenDesktop.
- **Physical**: The physical machine type is for machine catalogs that are based on machines that are *not* power-managed through XenDesktop.
- **Machine Creation Services (MCS)**: These use a master VM within XenDesktop to manage VMs, which enables you to easily manage and update target devices through one master image. This is the most common method used in **Virtual Desktop Infrastructure** (**VDI**).

MCS makes copies of a master image in order to assign these to individual users.

- **Provisioning Services (PVS)**: These manage the target devices as a device collection. Desktops and applications are delivered from a PVS vDisk.

PVS creates one image and share it with multiple users.

- **Another service or technology**: This manages and delivers the desktops and applications that you have already migrated to the VMs in the data center. This option provides you with the capability to use the third-party **electronic software distribution** (**ESD**) tools.

User experience

User experience is where you define the user types based on the tasks they perform and the environment they work in. The following are the user types available in Studio:

- **Random**: This user type assigns standardized desktops and applications with no personalization. These machine catalogs reside on the machines with the Microsoft Windows Server operating systems such as Server 2012. Some of their characteristics are as follows:
    - Randomly assigns machines to users at logon
    - Standardized machines for large-scale sharing
    - Delivers desktops or applications
    - Good for task workers who perform specific tasks and don't need personalization or users who don't need to install applications themselves on their desktops
- **Static**: This user type assigns shared desktops to users when they first log on and allows personalization. These machine catalogs reside on the machines with the Microsoft Windows desktop operating systems such as Windows 7 and 8. They are characterized by the following:
    - Used for static, fresh machine assignment
    - Private desktops for each user
    - Users can personalize their desktops and applications and store data to a Personal vDisk
    - Delivers desktops or applications
    - Good for performance and knowledge users who need their own desktop and the ability to save personal preferences

The following screenshot shows the user types available in Citrix Studio:

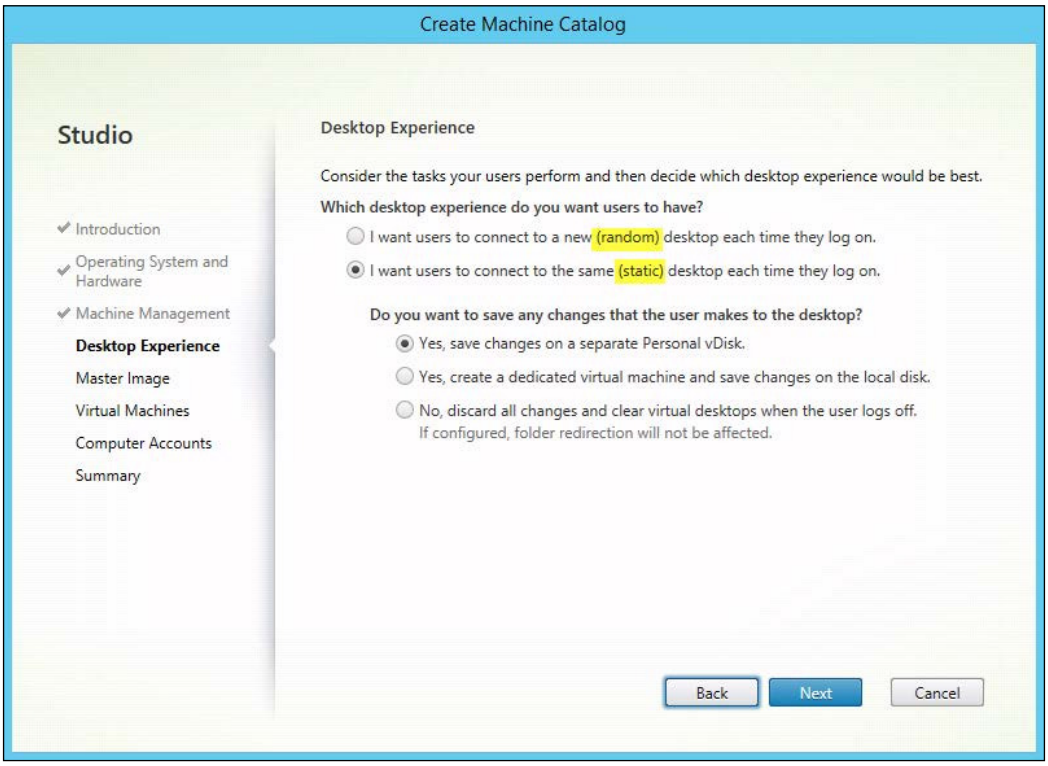

Managing the machine catalogs

Once you have created a machine catalog, you will undoubtedly find the need to manage it.

To manage a machine catalog, the following steps can be used:

1. Launch Studio.
2. Select **Machine Catalogs** in the left-hand side navigation menu.

3. Right-click on the machine catalog to view the additional options as shown in the following screenshot:

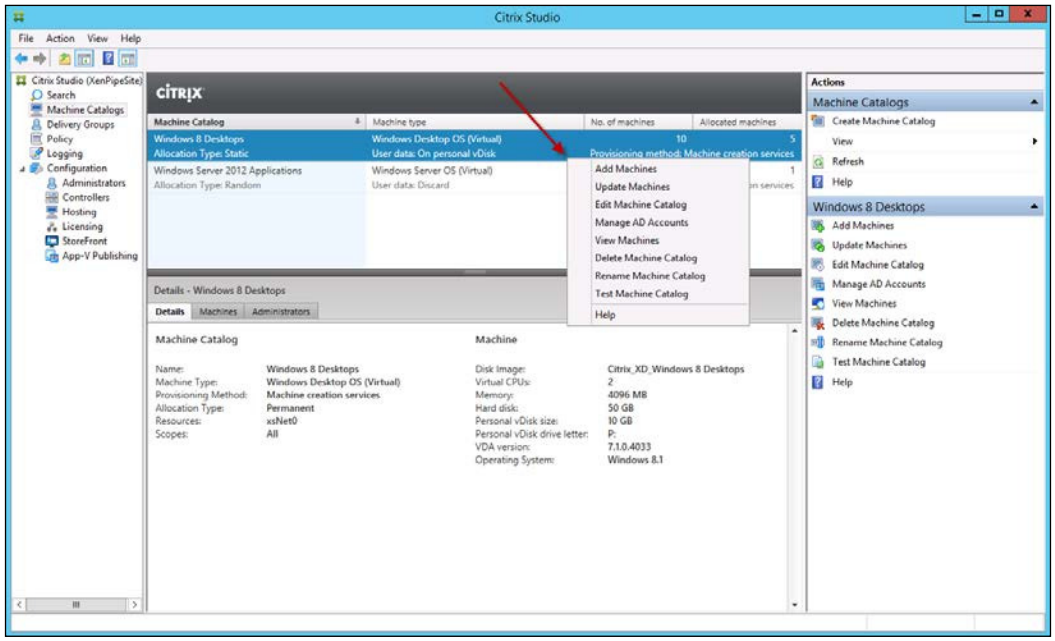

Taking a snapshot of the master image

Citrix urges you to take snapshots of the master images before you make any updates, in case you need to rollback.

To take a snapshot of a master image, perform the following steps:

1. Launch the XenCenter console (or another Hypervisor management console).
2. In the left-hand side pane, select your master image.
3. Select the **Snapshots** tab.

4. In the middle frame, select **Take Snapshot…,** as shown in the following screenshot:

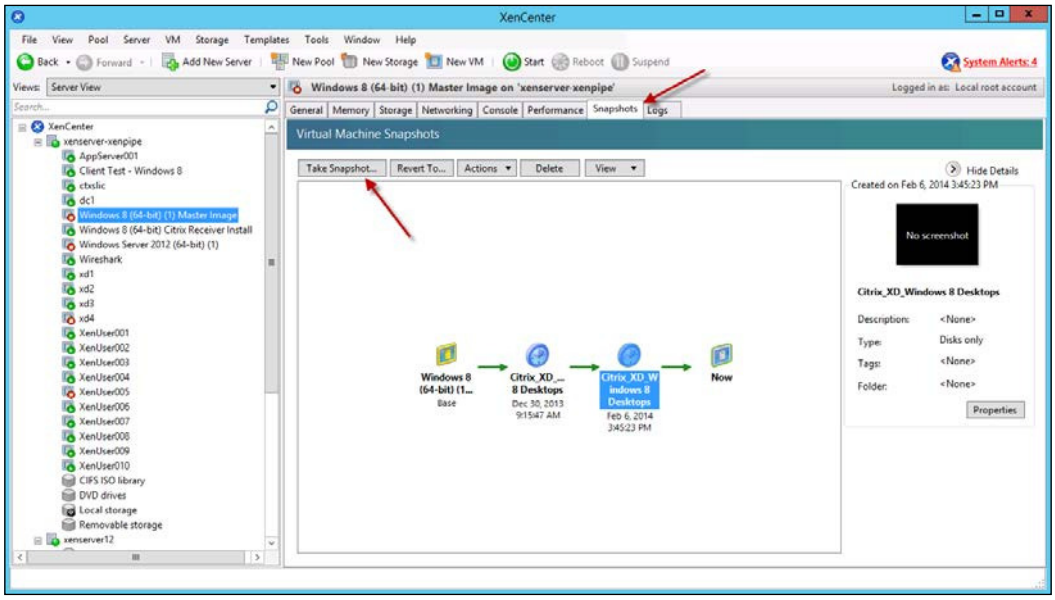

Updating the master image

Maintaining users' desktops is easier than you think. Random machine catalogs contain user desktops that you can maintain by applying global updates to the master image; for example, the Windows system updates the patches and antivirus software. After that, you simply modify the machine catalog to use the updated master image. The users will receive the updated desktop the next time they log on. This gives you incredible leverage to make substantial changes to the user desktops for a large number of users in a short amount of time. You can even do complete operating system upgrades. Static and physical machine catalogs are a little different and not as easy. These user types must be updated outside the XenDesktop management environment. You can do it manually on your own—machine by machine—or you can use the third-party electronic software distribution tools. If you are using streamed machines, you can update the user desktops by updating the vDisk through the PVS.

As mentioned, random machine catalogs can be updated once you have created and tested a new or updated master image. Desktops are updated with the new master image the next time the users log off, and the users receive the new or updated desktop the next time they log on.

When you create your snapshots, take two snapshots. One snapshot is for historical purposes. Rename the other snapshot of the master image and then use this to apply updates to the master image. The XenDesktop datastore or database keeps a historical record of the master images used with each machine catalog. If you do not delete, move, or rename the old master images, you can easily and instantly rollback a machine catalog to the previous version of the master image in the event that your users have problems with the updates you just deployed.

To update a master image, perform the following steps:

1. Start XenDesktop Studio.
2. In the left-hand side pane, select **Machines** and then select a machine catalog. Click on **Update Machines** as shown in the following screenshot:

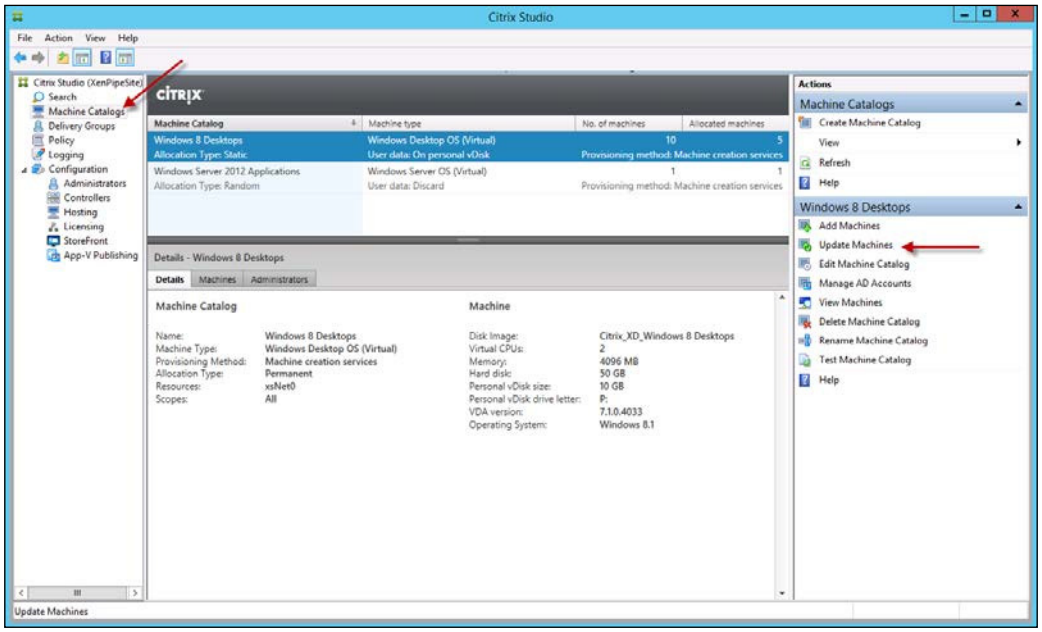

3. On the **Master Image** page, select the host and the new or updated master image.
4. On the **Strategy** page, specify how the new or updated master image will be applied to the user desktops.
5. Select **None** to apply the changes when the users log off.
6. Select **Send Message** to inform the users of the nonurgent update.
7. Select **Restart Immediately** to automatically log users off and restart their desktops.

8. Select **Send Message** and then **Restart after delay** to send a message, give the users some time to log off, and then log them off and restart their desktops.

Reverting to a previous master image

XenDesktop keeps a historical record of the master images for each machine catalog. This gives you the ability to instantly revert a machine catalog back to the previous master image in the event that your users run into problems with the updated image. Desktops revert back the next time when the users log off.

Don't delete, move, or rename any old master images including the snapshots in the chains to master images; otherwise, you will not be able to revert the machine catalog to the master image. If you are unable to locate the previous master image, you can browse for a different one to use as the update for the desktops.

To revert back to a previous master image, perform the following steps:

1. Launch Studio.
2. In the left-hand side navigation pane, select **Machines**.
3. Select a machine catalog and click on **Rollback machine update** as shown in the following screenshot:

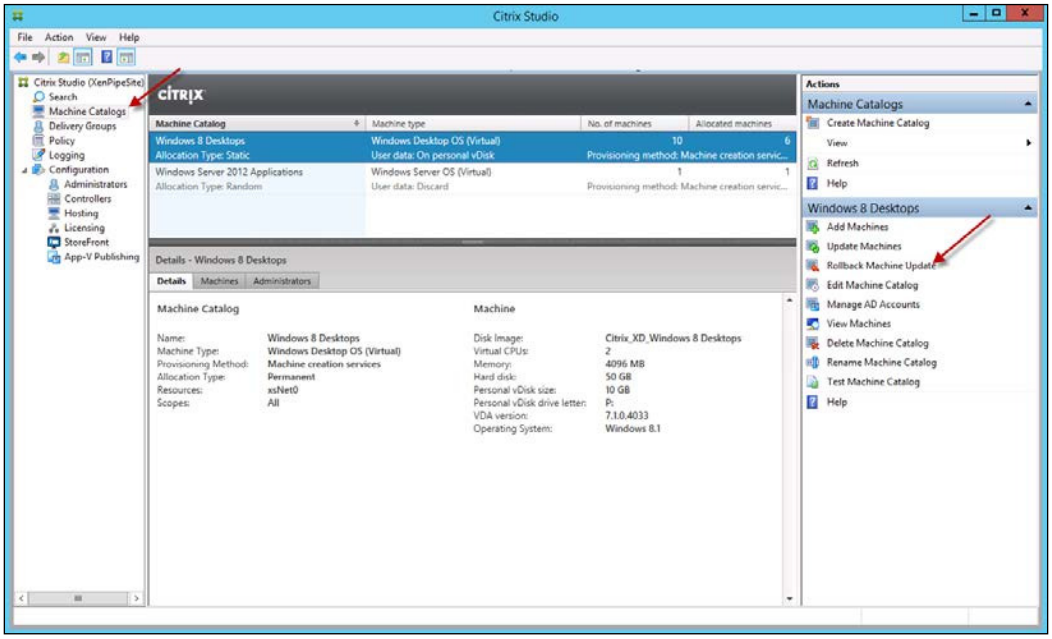

4. On the **Strategy** page, specify how the new or updated master image will be applied to users' desktops.
5. Select **None** to apply the changes when the users log off.
6. Select **Send Message** to inform the users of the nonurgent update.
7. Select **Restart Immediately** to automatically log the users off and restart their desktops.
8. Select **Send Message** and then **Restart** after delay to send a message, give the users some time, and then log them off and restart their desktops.

The rollback strategy is only applied to desktops that need to be reverted. You can only rollback if an update has been made previously.

Managing the Active Directory computer accounts

You can remove the Active Directory computer accounts from machine catalogs to free up the unused accounts in case you want to use them in other machine catalogs. Similarly, you can attach additional computer accounts to a machine catalog when more machines are added.

To add and remove the Active Directory computer accounts, use the following steps:

1. Launch Studio.
2. Select **Machine Catalogs** in the left-hand side navigation pane.
3. Select the machine catalog you want to operate on in the center navigation pane.
4. In the right-hand side navigation pane, select **View Machines**. This lists the Active Directory computer accounts for the machine catalog.
5. To remove computer accounts from the machine catalog and XenDesktop, select the accounts and click on **Delete**. You can also delete the accounts from Active Directory.
6. To add computer accounts to the machine catalog, select the machine catalog, click on **Add Machines** in the right-hand side navigation pane, and select the computer account in Active Directory.

You can also choose to reset the password for the accounts.

Adding machines to a machine catalog

Once you create a machine catalog, you can add machines for new users to that machine catalog. For Windows desktop and Windows Server machine catalogs, you can simply create more machines and Active Directory computer accounts and then add them to the machine catalog.

How to create machines is covered in the *Step 7 – creating the machine catalogs* section in *Chapter 2*, *Installing XenDesktop®*.

If you use streamed machine catalogs, you need to add more machines by joining the new target devices using the PVS. You can also create new device collections in the PVS and then add the collections to an existing machine catalog.

Make sure that the virtualization infrastructure (Hypervisor host) that hosts the master image for the machine catalog has sufficient memory, storage, and processors to host the additional machines. Ensure that there are adequate unused Active Directory computer accounts for the additional machines.

To add machines to a machine catalog, use the following steps:

1. Launch Studio.
2. Select **Machine Catalogs** and click on **Add Machines** as shown in the following screenshot:

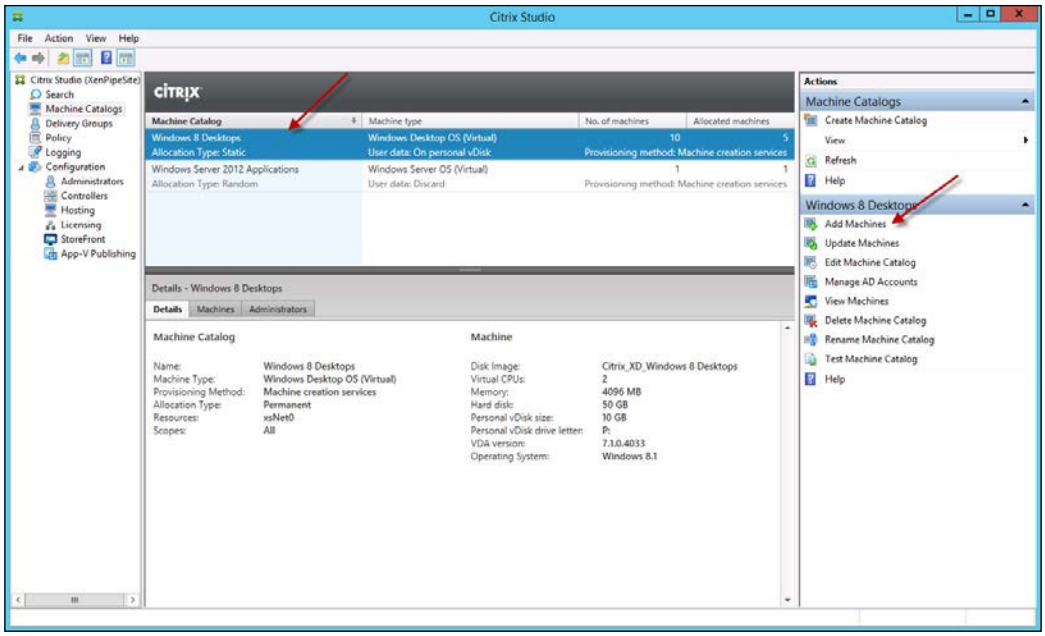

3. On the **Machines and Users** page, click on **Add computers** and select the computers from the Active Directory domain.
4. To create additional computer accounts or use the existing computer accounts, provide the Active Directory domain and **Organizational Unit** (**OU**) that contains the computer accounts. To use the existing accounts, select **Browse** or **Import**.

Modifying a machine catalog

Using Studio, you can change the description of a catalog.

To edit a machine catalog, use the following steps:

1. Launch Studio.
2. Select the machine catalog and then click on **Edit Machine Catalog**.

Renaming a machine catalog

You can rename a machine catalog. To rename a machine catalog, use the following steps:

1. Launch Studio.
2. Select the machine catalog and then click on **Rename Machine Catalog**.

Deleting a machine catalog

You can delete a machine catalog. If you do, it will also remove the machines and Active Directory computer accounts from XenDesktop. You can also delete the machines and computer accounts from the Hypervisor host.

Before deleting a machine catalog, make sure that all users have logged off from the desktops in the machine catalog, no disconnected user sessions are still running, all of the machines in the catalog are in maintenance mode, and all machines in the catalog are powered off.

To delete a machine catalog, use the following steps:

1. Launch Studio.
2. Select the machine catalog and then click on **Delete Machine Catalog**.
3. You can also delete the machines hosting the users' desktops.

Managing the hosts

You can use Studio to create hosts, add storage to hosts, rename hosts, update connection details, rename connections, configure high availability, configure throttling, enable/disable maintenance mode, view details of VMs, manage VMs, delete hosts, and delete connections.

You need full administrator rights to carry out the tasks mentioned in the preceding section.

To create a new host, use the following steps:

1. Launch Studio.
2. Navigate to **Configuration** | **Hosting**.
3. Select **Add Connection and Resources** to a new host.
4. Specify the type, address, and credentials. It is recommended to use HTTPS to secure communication between XenDesktop and XenServer. You can also enable high availability if using XenServer.

Select **None** for the host type if you are running desktops on dedicated blade PCs.

5. Name the connection.
6. Either use XenDesktop to create VMs or create them manually; then select **Next**.

To create a host by means of existing connection details, use the following steps:

1. Launch Studio.
2. Navigate to **Configuration** | **Hosting**.
3. Select **Add Connection and Resources**.
4. Select **Use an existing Connection**, select the connection, and then click on **Next**.
5. Select the storage and network for the VMs; then select **Next**.
6. Type a new name for the new host and click on **Finish**. These steps are summarized in the following screenshot:

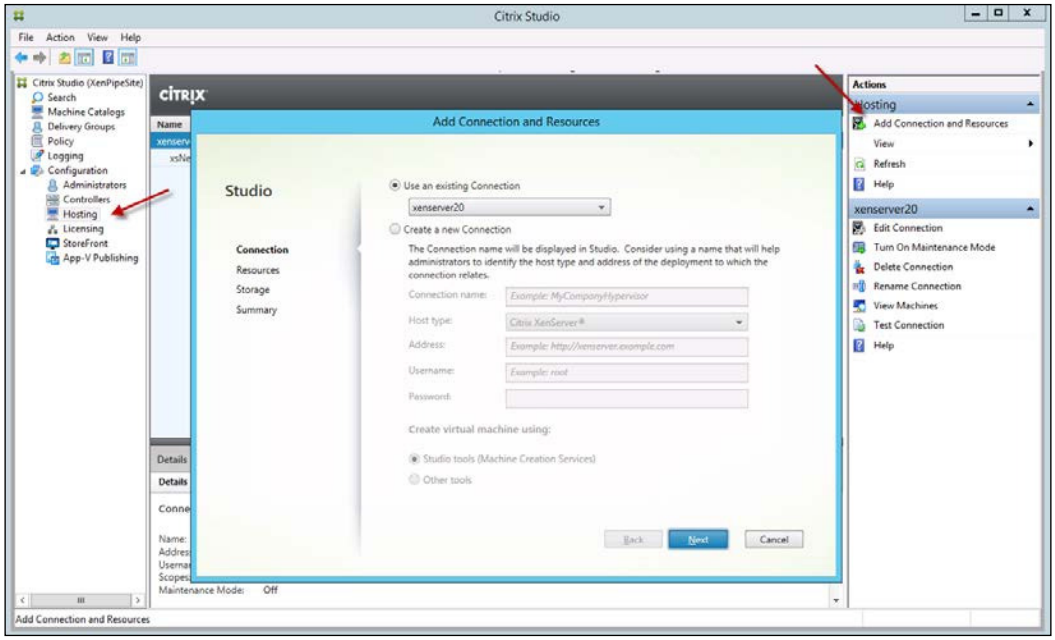

To edit a host or connection, use the following steps:

1. Launch Studio.
2. Navigate to **Configuration** | **Hosting**.
3. Select the host and click on **Edit Connection**.
4. Select **Change Details**.
5. You can update the connection's address and credentials or rename it.
6. You can configure high availability under **Edit**.
7. You can configure Hypervisor throttling under **Advanced**.

To prevent more than a specific number of operations or actions running at one time, modify the **Max active actions** property. To limit concurrent actions to a percentage of the total number of VMs configured, modify the max power as a percentage of desktops. The lower of the two is used. You can also limit the number of new actions by modifying the **Max new actions per minute** property.

The steps to edit a host or connection are displayed in the following screenshot:

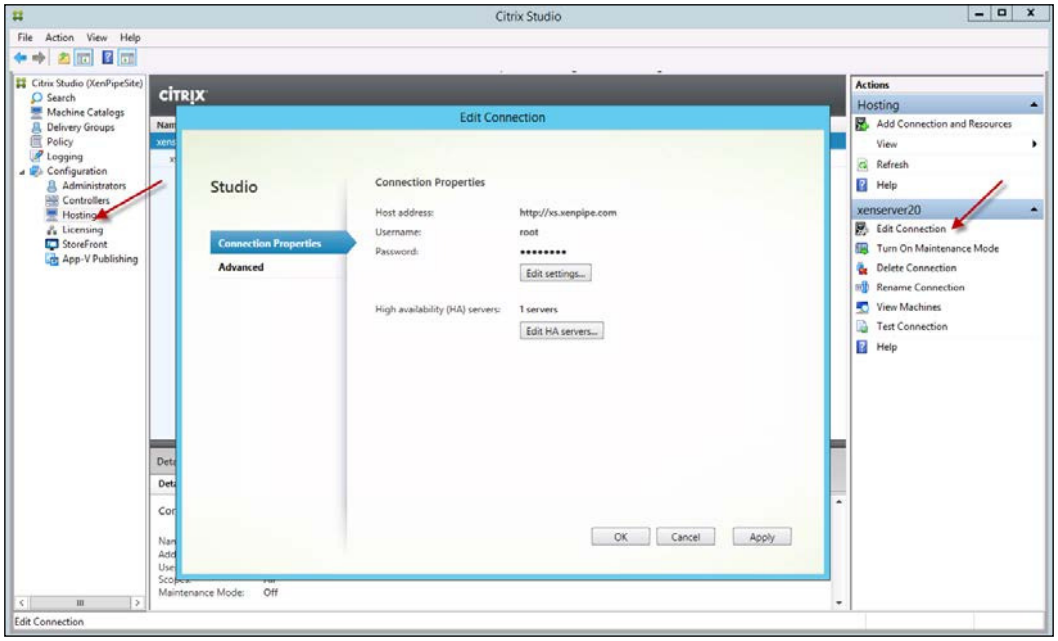

To place a connection into maintenance mode, use the following steps:

1. Launch Studio.
2. Navigate to **Configuration | Hosting**.
3. Navigate to **Connection | Turn On Maintenance Mode**.
4. Click on **Disable** to take a connection out of maintenance mode.

To manage machines accessed through a specific connection, use the following steps:

1. Launch Studio.
2. Navigate to **Configuration | Hosting**.
3. Select the connection.
4. Select **View Machines**; then select the machine(s).
5. Select an action from the following options:
    - **Start**: This starts the machine if it is powered off or suspended
    - **Suspend**: This pauses the desktop without shutting it down
    - **Shut down**: This requests the desktop's operating system to shut down
    - **Force Shut down**: This forcibly powers off the desktop
    - **Restart**: This requests the desktop's operating system to shut down and then start again
    - **Turn On Maintenance Mode**: This temporarily halts connections to a machine so that maintenance can be carried out

Users cannot connect to a machine in maintenance mode. If a user is connected, the maintenance mode takes effect when the user logs off.

    - **Remove from Delivery Group**: Removing a machine deletes it from a delivery group, but does not delete it from the catalog
    - **Delete**: Deleting a machine removes it from the catalog

Make sure all data is backed up before deleting a machine.

To delete a host or connection, use the following steps:

1. Launch Studio.
2. Navigate to **Configuration** | **Hosting**.
3. Select the connection to delete and then select **Delete Connection**.

A catalog will become unusable when a connection is deleted that is associated with a host that is referenced by that catalog. Before you delete a catalog, make sure it is not supported by other hosts.

Managing Personal vDisks

At this point, you are probably wondering how to harness the economics of a shared desktop with the personalization capability of a dedicated desktop. Citrix made this possible with the acquisition of RingCube, and called the feature **Personal vDisk** (**PvD**). The PvD feature allows you to manage pooled and streamed master image with the addition of allowing people to make their desktops personal by installing applications and saving their desktop preference settings. This is a key difference from traditional VDI deployments with pooled desktops, where users would lose their customized preferences and personal applications after the administrator modifies the base VM or master image. In other words, administrators can quickly and centrally manage their base VMs and provide users with a customized and personalized desktop experience.

Personal vDisks provide this separation of pooled/shared desktops from personalization by sending all changes made on the user's VM to a separate disk, called the Personal vDisk, which XenDesktop attaches to the user's VM. The content from the Personal vDisk is merged with the content from their base VM when the user logs on to their virtual desktop. This happens seamlessly from the user's perspective. Ultimately, users access applications provisioned by their administrator in the base VM while getting the look and feel they are used to seeing in their personal desktops.

Personal vDisks contain the following two pieces, which are the same size by default:

- The first is the path located at `C:\Users` (in Windows 7 and 8) or `C:\Documents and Settings` (in Windows XP). This location contains the user's profile, data, and documents. By default, this is assigned the drive letter `P:`. Yet, you can change it to a different drive letter when you use Studio to create a machine catalog that uses Personal vDisks.
- The second uses a virtual hard disk file (a `.vhd` file). This contains all the other stuff, such as applications installed in `C:\Program Files`. This piece is hidden from users and is not displayed in Windows Explorer.

An algorithm automatically adjusts the sizes of the two pieces depending on how the vDisks are used. So, if a user installs some big applications on a Personal vDisk such that space becomes limited, the application space is increased to accommodate user data. The complete size of the Personal vDisk does not change. This resizing feature is configurable. As for the physical location, the Personal vDisk does not need to be stored with the dedicated pool VM. This is significant because it allows you to use high-speed disks for VM storage and less expensive disks for Personal vDisk storage.

When you create a host, you define the storage locations for disks that are used by the VMs. You have the option to separate Personal vDisks from the disks used for the VMs. If you use local storage for both VM and PvD, they need to be accessible from the same Hypervisor. Citrix Studio will only offer up such compatible storage locations when you create a host.

You can add Personal vDisks to new hosts when you create a new XenDesktop Site. You can also add PvDs and storage for them to existing hosts (but not machine catalogs).

To add Personal vDisks to existing hosts, use the following steps:

1. Launch Studio.
2. Navigate to **Configuration** | **Hosting** and select a host.
3. Click on **Add Connection and Resources**.

4. In the **Storage** section, under **Personal vDisk storage (Desktop OS only)**, select one of the options to either use the same or different Personal vDisk storage, and specify the storage location, as shown in the following screenshot:

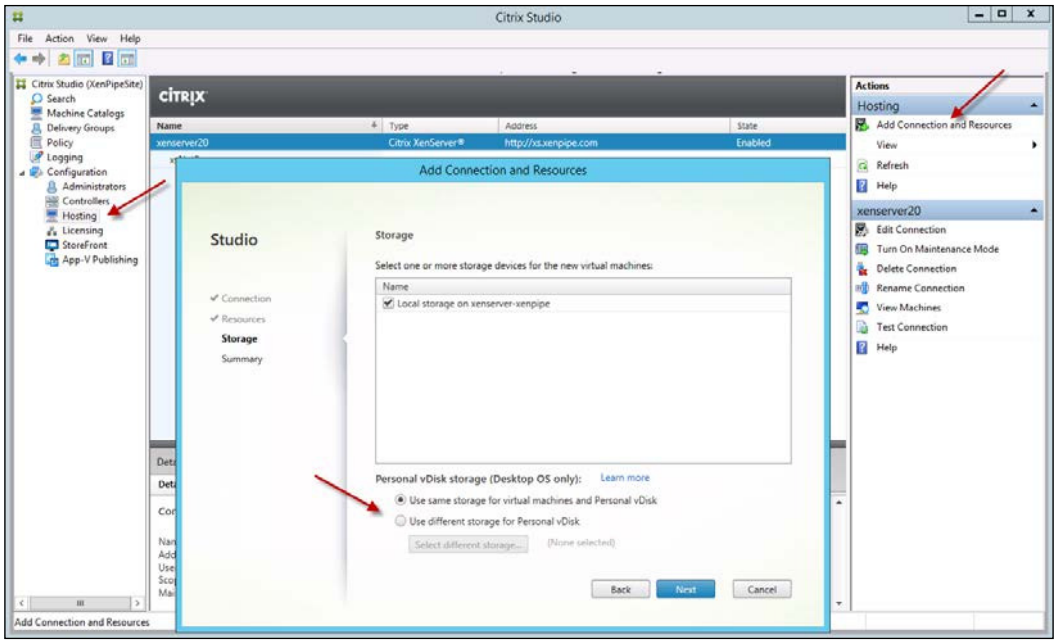

Updating Personal vDisks used by the master images

You can enable the Personal vDisk feature for use with a master image when you install the Delivery Agent using the following steps:

1. Update the master image with application or operating system updates.

You can prevent prompts for unsigned drivers in the **Control Panel** of the master image by selecting the **Ignore warnings** option under **System** | **Hardware** | **Driver Signing**.

2. Shut down the machine.
3. Select **Update Inventory** in the Personal vDisk dialog box.

If you interrupt the shutdown, you must restart the machine, shut it down, and update the inventory again.

4. It is recommended to take a snapshot of the master image when the inventory operation shuts down the machine.

Adjusting the space available for applications

It is possible to manually adjust the automatic resizing algorithm that controls the size of the VHD relative to the `P:`. XenDesktop manages this automatically; however, there might be a case where a user has installed an application that doesn't fit on the VHD even after it has been resized with the algorithm. Currently, the way you adjust the space is by editing the primary size of the VHD in the registry.

It is best practice to make the adjustment on a machine catalog's master image; however, you can adjust the size of the VHD on a virtual desktop for an affected user. One use case is to create a big enough VHD to store large antivirus definition files.

The following steps will guide you in adjusting the space available for applications:

1. On the master image or desktop, launch the Registry Editor.
2. Find the location for `HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\personal vDisk\Config`.
3. Set **MinimumVHDSizeMB** to the new primary size for the VHD (in MB). It needs to be greater than the current size and less than the size of the physical disk minus the **PvDReservedSpaceMB** value.
4. Make sure that **PercentOfPvDForApps** is set to `50`. This sets the default allocation of space on the Personal vDisk to 50 percent. If any other value is used, the dynamic resizing algorithm gets disabled.
5. Activate the algorithm by setting **EnableDynamicResizeOfAppContainer** to a value of `1`.
6. If you are using profile management or Citrix profile management, ensure that **EnableUserProfileRedirection** is set to a value of `0`. This makes sure all of the space on `P:` is allocated for applications.
7. If you are performing this operation on a virtual desktop rather than a machine catalog, resizing takes place when the desktop is restarted.

Disabling automatic resizing

To disable the automatic resizing algorithm, use the following process:

- On the master image or desktop in the Registry Editor, set **EnableDynamicResizeOfAppContainer** to a value of `0`. This is located at `HKEY_LOCAL_MACHINE\SOFTWARE\Citrix\personal vDisk\Config`.

Editing the registry may create problems if not done correctly. Use the Registry Editor at your own risk.

Reallocating user profiles

In XenDesktop, user profiles are stored on the Personal vDisk (`P:`) by default and not on the virtual desktop (`C:`). If you would rather have Citrix profile management process the profiles instead of the Personal vDisk, you can change this default value when installing the Virtual Desktop Agent by modifying the registry on the master image that is used for the new catalog. In this instance, the catalog is new and no users have logged on; therefore, there are no profiles to store on the `P:`.

On the other hand, if you activate profile management on machines in existing machine catalogs with Personal vDisks and the machine catalog is already in use, the logons will already have taken place and profiles will exist on the `P:`. The profiles will stay on the drive even after you modify the registry, so you should adjust the default.

Summary

In this chapter, we covered machine catalogs, hosts, and Personal vDisks; important concepts with regard to XenDesktop; and delivering desktops and applications. After reading this chapter, you now have the resources to understand how to create and deliver desktops and applications as well as how to manage them. In the next chapter, we will look at delivery groups and how to manage them.

4
Managing Delivery Groups

You've built a XenDesktop Site, learned about machine catalogs, hosts, and Personal vDisks. You have also been exposed to the delivery group wizard. In this chapter, we will expand on delivery groups including the following topics:

- Managing delivery groups
- Managing hosted applications
- Managing Delivery Controllers

As you learned in the previous chapter, collections of virtual desktops or physical computers are managed as a single entity called a machine catalog. As an administrator, you deliver desktops to users by creating a catalog of machines and then allocate the machines in the catalog to users by creating delivery groups. You can also deliver applications to users by creating application delivery groups. You can change the delivery groups at any time, meaning that you can add desktops or applications to a delivery group at any time.

Managing the delivery groups

Delivery groups are collections of machines that deliver desktops and/or applications to users. Delivery groups specify which users can access the desktops or applications and are usually based on user characteristics or the types of users. We created delivery groups for desktops and applications in *Chapter 2*, *Installing XenDesktop®*. Delivery groups can include the following machines:

- **Windows desktop machine**: This option provides secure, easily controllable access for users such as external contractors or third-party collaborators in addition to internal users. This provides a secure and centralized resource management while providing individual desktop and application environments for each user.

Use this option to deliver desktops.

- **Windows Server machine**: This option supports a high volume of users while providing a high-definition user experience when accessing desktops or applications. This is useful for task users who perform a set of well-defined tasks and do not require personalization.

Use this option to deliver applications or shared server desktops.

Creating a delivery group

Before you create a delivery group, please note the following prerequisites:

- You can only create a delivery group if at least one desktop remains unused in the machine catalog.
- You cannot use a desktop in more than one delivery group.
- You can create delivery groups from multiple machine catalogs with the same desktop characteristics.
- You cannot create mixed delivery groups from machine catalogs with different desktop types. The machine catalog characteristics must match if you want to put the desktops in a single group.

The following steps will help you to create a delivery group:

1. Launch Studio.
2. In the right-hand side pane, click on **Create Delivery Group**.
3. Select the **Machine Catalogs** section and enter the number of machines to be added from the machine catalog.
4. On the **Delivery Type** page, select what the desktops will deliver to your users:
    - Virtual desktops
    - Virtual applications

This is shown in the following screenshot:

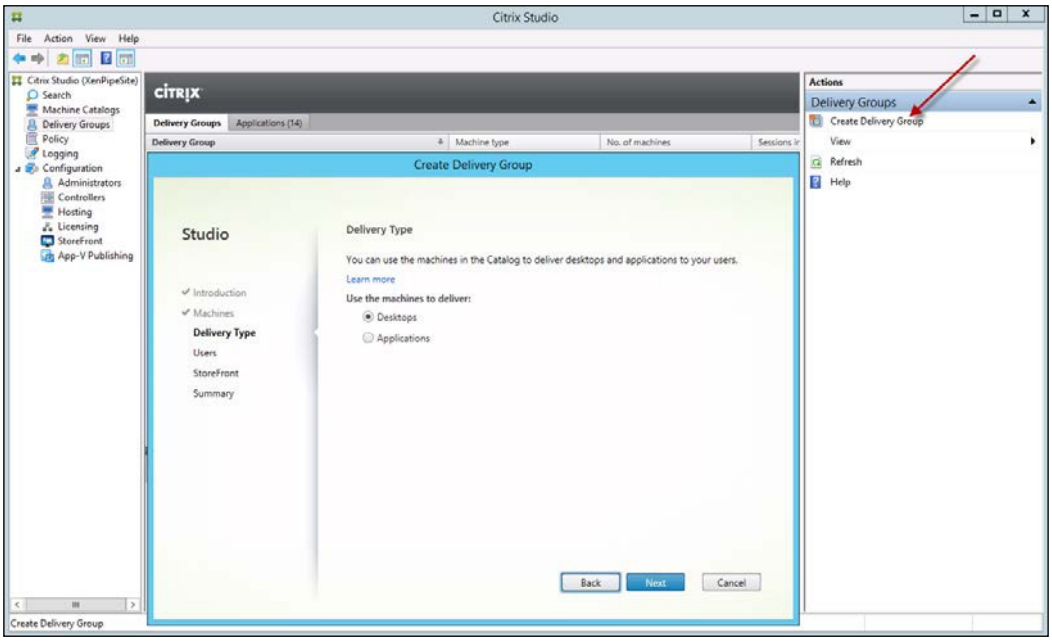

5. On the **Users** page, add the user or user groups that can access the desktops. You can select the user groups by browsing or entering a list of Active Directory users and groups, each separated by a semicolon. For Windows desktop machine delivery groups, you can import user data from a file after you create the group.
6. Select the StoreFront URLs to be pushed to the Citrix Receiver so that the Receiver can connect to a store without user intervention. You can choose to connect automatically or manually.
7. On the **Summary** page, check all the details. Enter a display name that the users and administrators can see and a descriptive delivery group name that only the administrators can see.

Editing a delivery group

You can change a delivery group by editing it. To edit a delivery group, perform the following steps:

1. Launch Studio.

2. Select **Delivery Groups** in the left-hand side pane.
3. Select the **Delivery Group** tab in the center pane and click on **Edit Delivery Group**.
4. Add the users or groups that can access the desktops. You can browse or enter a list of Active Directory users and groups, each separated by a semicolon.
5. You can change the **Delivery Type** option to **Desktops** or **Applications**.
6. On the **End User Settings** page, you can change the following fields:
    - The delivery group's **Display name** and **Description**
    - The number of hosted **Desktops per user**
    - **Color depth**
    - **Time zone**
    - The **Enabled** option for the delivery group
    - **Enable Secure ICA**
7. On the **StoreFront** page, select the store URLs to be pushed to the Citrix Receiver so that it can connect to a store without user intervention.
8. On **Power Management**, you can dictate when the machines will be powered on/off and disconnected/logged off as shown in the following screenshot:

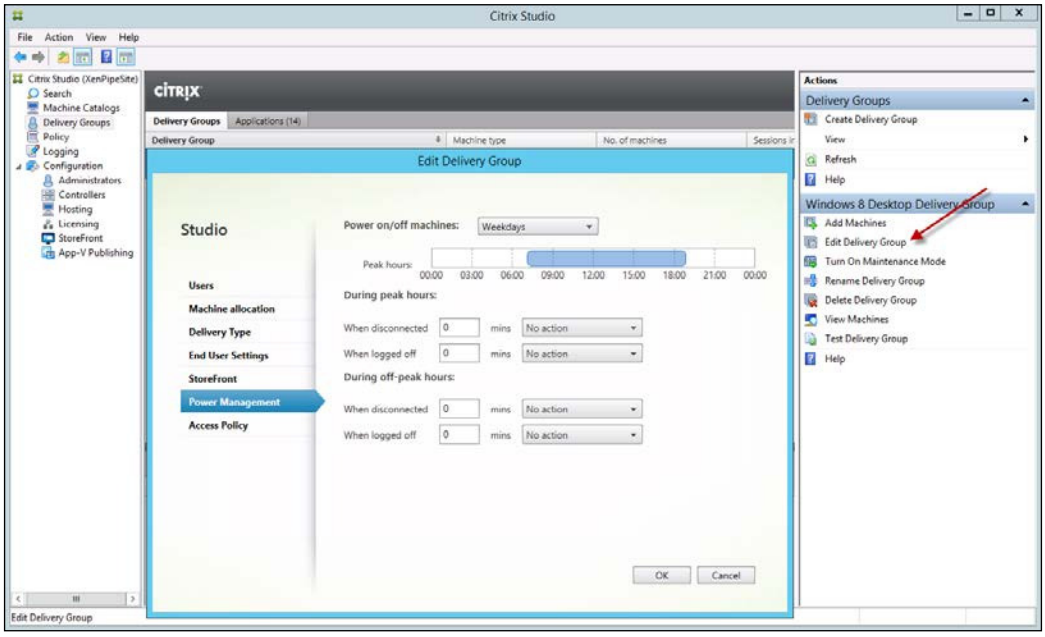

This is a really important setting. XenDesktop keeps a certain number of machines powered on and ready for use. In the previous versions, it was 10 percent of the total pool; so, for example, out of 50 desktops, 5 would be powered on by default.

9. On the **Access Policy** page, you can select the delivery group connections to go through the NetScaler Gateway.

Managing desktop sessions

The VDA registers the machine with the **Delivery Controller** (**DC**). When a user logs on to a desktop, the VDA registers the machine as being in use with the DC, and the user logs on to their desktop. When you want to carry out maintenance or assist users, you can control sessions by logging users off from the sessions, disconnecting sessions, sending messages to users, as well as searching and locating sessions, users, and desktops.

Logging off or disconnecting sessions

To log off or disconnect sessions, perform the following steps:

1. Launch Studio.
2. Use **Search** to locate the session, or select a desktop group and click on **View Machines**.
3. Select the session or desktop and click on **Log off**.

Logging a session off closes the session and frees the desktop for other users, unless it is allocated to a specific user. Disconnecting a session keeps the user's applications running, and the desktop remains allocated to that user. When the user reconnects, the same desktop is allocated.

You can configure the power state timers through policies to automatically process unused sessions, freeing up desktops and saving power.

Sending messages to users

In order to send messages to users to inform them about desktop maintenance, perform the following steps:

1. Launch Studio.

2. Use **Search** to locate the session, desktop, or user. You can also select a desktop group and click on **View Machines**.
3. Select the session, desktop, or user and click on **Send Message** as shown in the following screenshot:

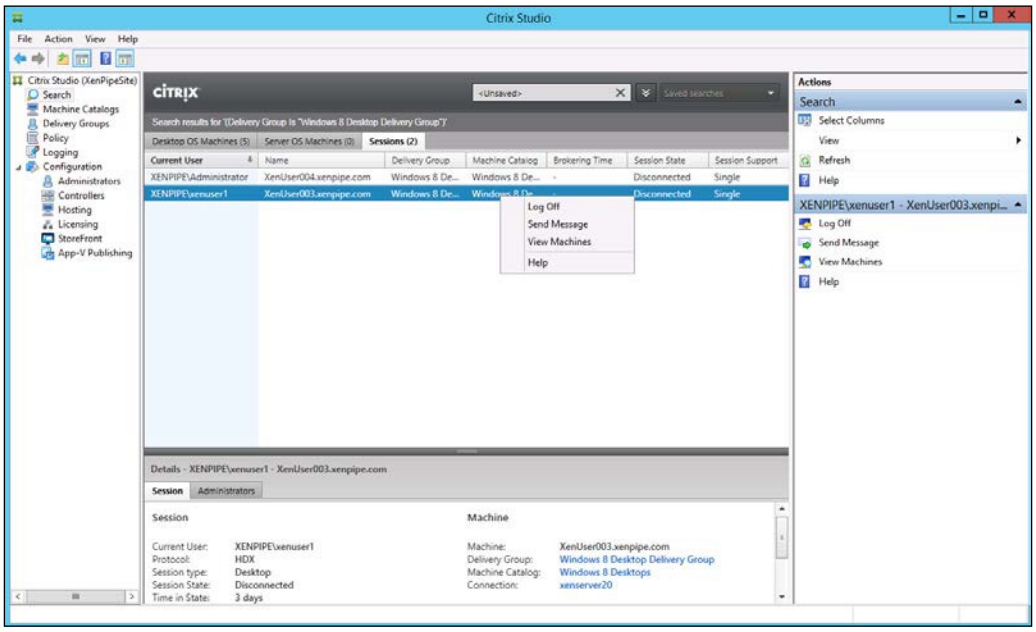

Managing the delivery group resources

You can add and reallocate desktops from the delivery groups.

Adding and reallocating desktops

To add or reallocate desktops from a delivery group, perform the following steps:

1. Launch Studio.

2. Select the **Delivery Groups** option.
3. Select **Edit Delivery Group**.
4. On the **End User Settings** page, set the number of desktops to be used, as displayed in the following screenshot:

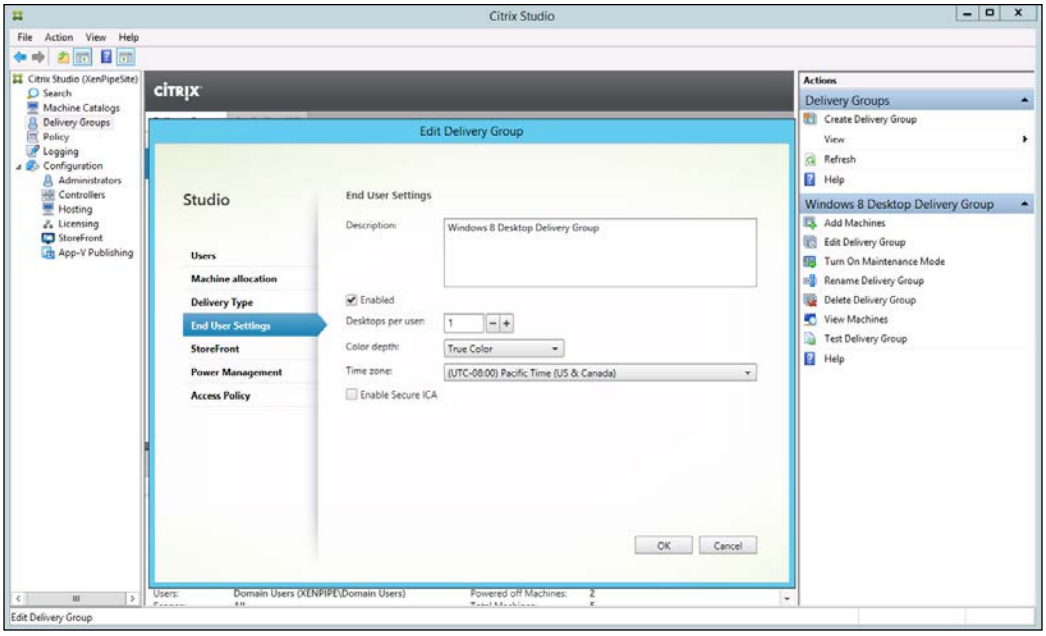

Locating desktops, sessions, and delivery groups

To locate desktops, sessions, and delivery groups, perform the following steps:

1. Launch Studio.
2. Select the **Search** node.

3. Enter the name or a part of the name of the desktop, session, or delivery group, as shown in the following screenshot:

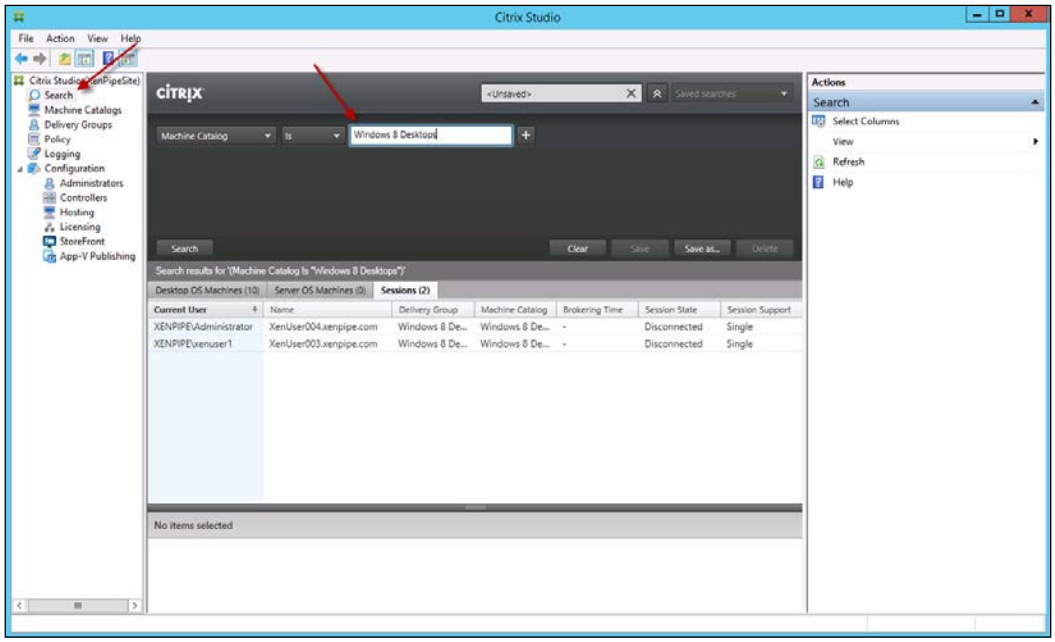

You can use the **Advanced search** option to find an endpoint using IP addresses, connected sessions, and delivery groups.

Shutting down and restarting desktops

To shut down or restart a desktop, perform the following steps:

1. Launch Studio.
2. Select the **Search** node to find the desktop you want to shut down or restart.
3. You have the following options, as shown in the following screenshot.
    - **Force Shutdown**: This option forces the desktop to shut down.
    - **Restart**: This option sends a request to the operating system to shut down and then start again.
    - **Force Restart**: This option sends a `RESET` command to the Hypervisor.

- **Suspend**: This option pauses the desktop.
- **Shut Down**: This option sends a request to the operating system to shut down. A machine that is not placed in maintenance can be powered on again.

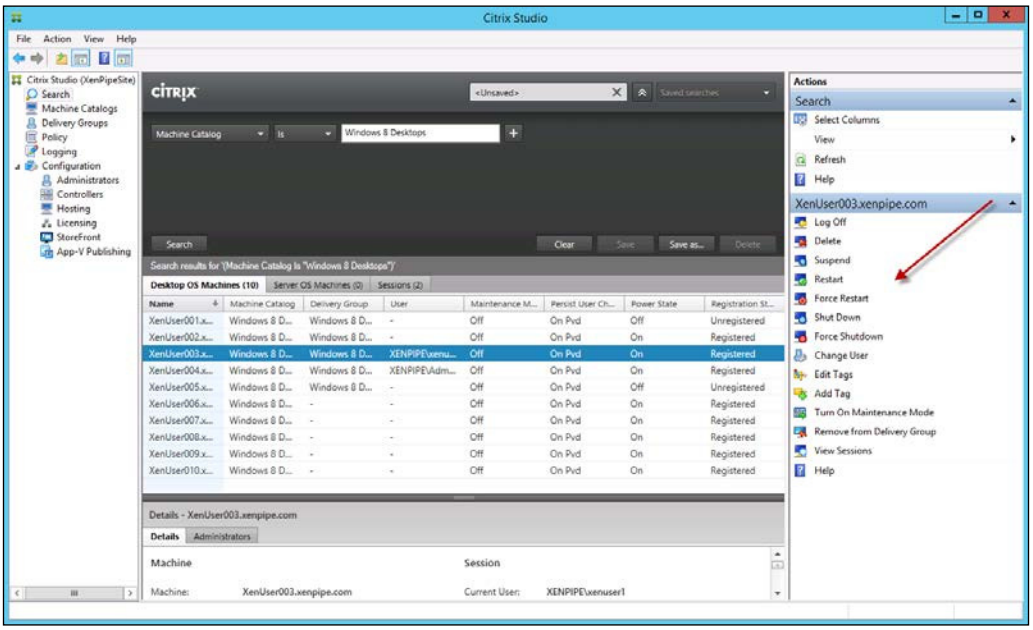

Removing desktops from delivery groups

To remove a desktop from a delivery group, perform the following steps:

1. Launch Studio.
2. Select the **Search** node to find the desktop you want to remove.
3. Put the desktop in the maintenance mode.
4. Make sure that the desktop has been shut down.
5. Select **Remove from Delivery Group**.

Placing the desktop in the maintenance mode prevents users from connecting to it. Also, desktops may contain personal data, so you may need to reimage the virtual machine.

Deleting desktops from delivery groups

To delete a desktop from a delivery group, perform the following steps:

1. Launch Studio.
2. Select the **Search** node to find the desktop you want to delete.
3. Put the desktop in the maintenance mode.
4. Make sure that the desktop has been shut down.
5. Select **Delete**.

If you want to delete a desktop but keep the virtual machine, you must remove it instead of deleting it.

Restricting access to desktops

You can restrict access to a delivery group's desktops by using **Scopes**, **Smart Access** strings, and **Exclusion filters**. You can use Scopes for administrator access. You can use the Smart Access strings and Exclusion filters for user access. Smart Access strings allow users to securely access desktops and applications using any device at any location. Exclusion filters are used to restrict access to machines.

Using Smart Access

To restrict user access using the Smart Access strings, perform the following steps:

1. Launch Studio.
2. Select the **Delivery Groups** option.
3. Select **Edit Delivery Group** and click on **Access Policy**.
4. Select **Connections through NetScaler Gateway**.
5. Add, edit, or remove the Smart Access strings that allow users to access scenarios for the delivery group, as shown in the following screenshot:

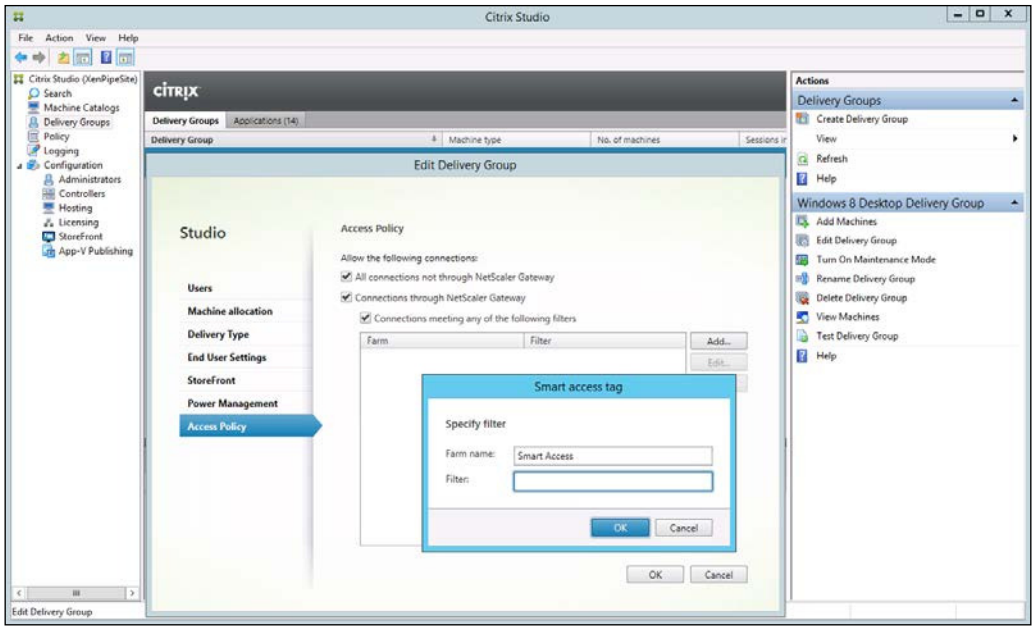

Smart Access filters must be used with a NetScaler Gateway. When you use the Smart Access filters, only connections through the NetScaler Gateway are allowed.

Using Exclusion filters

Exclusion filters are used through the XenDesktop **Software Development Kit** (**SDK**).

For example, to prevent access to a certain delivery group regardless of who is using the desktops in the lab, you would enter the following `SDK` command:

```
Set-BrokerAccessPolicy -Name VPDesktops_Direct
 -ExcludedClientIPFilterEnabled $True -
```

You can also use the asterisk (*) key as a wildcard to match all the tags that start with the same string. For example, `VPDesktops_*` applies the filter to all desktops.

Securing the ICA® protocol communications

You can secure all the communications to and from desktops in any delivery group using the **SecureICA** feature. SecureICA encrypts the `ICA` protocol. SecureICA is disabled by default. If you enable it, it defaults to a 128-bit encryption. You can change the level of encryption using the SDK.

Don't use SecureICA as the only encryption method as it doesn't provide authentication and doesn't check message integrity. You should use the `SSL/TLS` encryption in addition to SecureICA. Using the NetScaler Gateway to frontend the XenDesktop Site is an effective way to provide `SSL/TLS` to the XenDesktop Site.

To turn on SecureICA for desktops, perform the following steps:

1. Launch Studio.
2. Select the **Delivery Groups** option.
3. Select **Edit Delivery Group** and click on **Users**.
4. In **End User Settings**, select **Enable Secure ICA** as shown in the following screenshot:

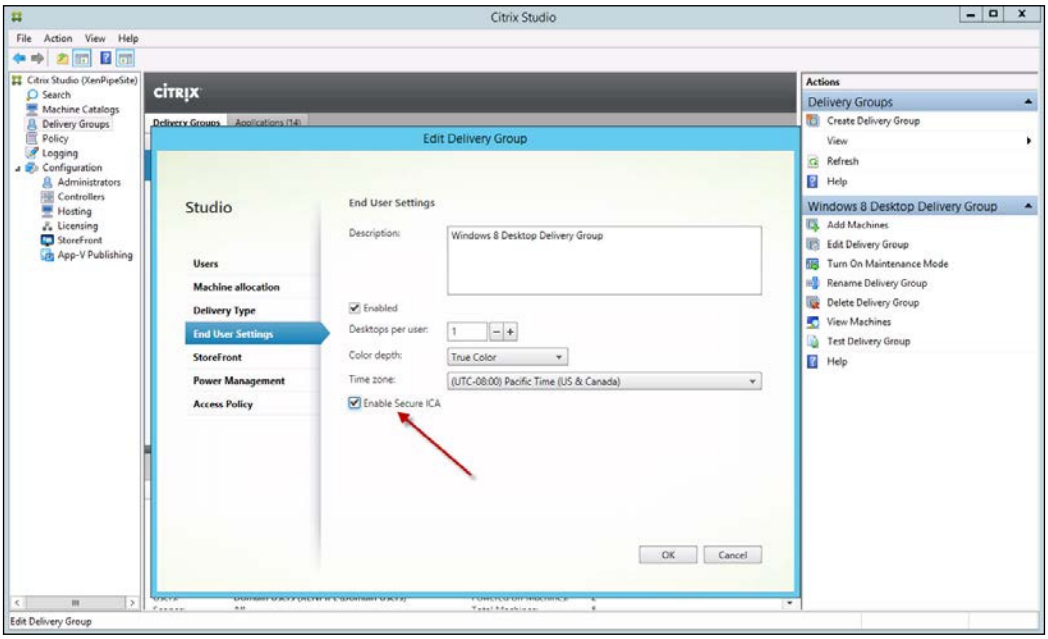

Managing power settings for desktops

You can manage partial or full power settings for desktop delivery groups. This applies to the virtual machines only and not to the physical machines. Permanently allocated desktops can only be power managed partially. Static delivery groups can contain both permanently allocated and unallocated desktops; however, you can only power manage the unallocated desktops.

Pools and buffers

In random delivery groups and for unallocated desktops in static delivery groups, a **pool** is a set of unallocated desktops that are kept in a powered-on state and are ready for users to connect to. The pool size is configurable.

When users log on, they are immediately presented with a desktop. You may want to make the pool larger during peak usage times.

A **buffer** is an extra set of unallocated desktops that are turned on and ready for users to connect to. For random delivery groups and unallocated desktops in static delivery groups, the desktops in the buffer are turned on when the number of desktops in the pool drops below the threshold buffer size. By default, this is 10 percent. You can adjust the buffer size using the SDK.

Power state timers

You can use power state timers to suspend desktops after users have disconnected for a defined period of time. Random or streamed desktops are always automatically shut down when users log off. This can be changed using the `ShutdownDesktopsAfterUse` property in the SDK. You can set the timers for peak and off-peak usage. You can manage the power state timers using the SDK.

Partial power management

With permanently allocated desktops, you can set the power state timers but not the pools or buffers. XenDesktop will turn on the desktops at the start of the peak period and turn them off in the off-peak period.

You don't have control over the buffer size or number of desktops that become available to compensate for the desktops that are used.

To change the power state timers, perform the following steps:

1. Launch Studio.
2. Select the **Delivery Groups** option.
3. Select **Edit Delivery Group** and click on **Power Management**.
4. Select **Weekdays**.
5. For random delivery groups, edit and specify the pool size.
6. In peak hours, set the peak and off-peak hours.
7. Set power state timers for the following options:
    - **When disconnected**: Specify the delay (in minutes) before the disconnected machines should be suspended and then select **Suspend**.
    - **When logged off**: Specify the delay before the logged-off machines in the delivery group should be turned off and then select **Shut Down**. This is not available for random desktops.

Importing and exporting user data

You can allocate desktops and applications to users by importing data from a file. You can import to a delivery group that is based on the existing or physical machines if you have correct permissions to access the file and delivery group. You can also export user data to a file.

The import file can contain data from previous XenDesktop versions. You can only use this type of file to update the delivery groups on physical machines.

The import and export file format requirements are as follows:

- A comma-separated value (`.csv`) file format
- The first line must contain the column headings as follows:

  ```
  [ADComputerAccount],[AssignedUser],[VirtualMachine],[HostId]
  ```

- The `ADComputerAccount` values can be any of the following options:
    - A common name (`VirtualComputer01`)
    - An IP address (`x.x.x.x`)
    - A distinguished name (`VirtualComputer01.domain.com`)

- Domain and computer name pairs (for example, `domain\VirtualComputer01`)

To import or export data, perform the following steps:

1. Launch Studio.
2. Select the **Delivery Groups** option.
3. Select **Edit Delivery Group**.
4. Under **Machine allocation**, select **Import list…** or **Export list…,** as shown in the following screenshot:

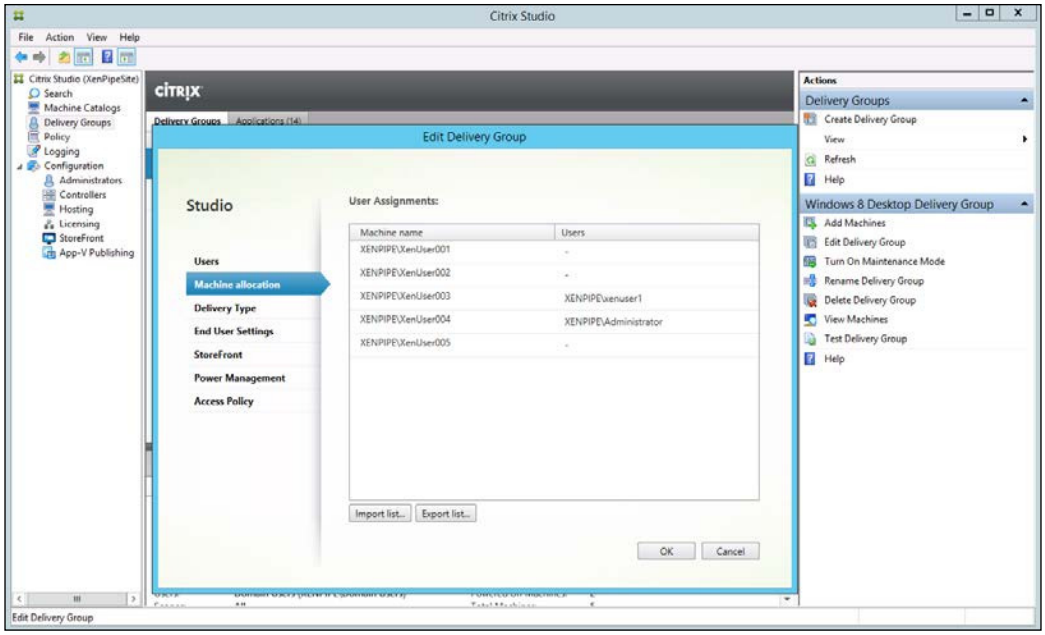

Enabling and disabling the maintenance mode

To stop connections to the desktops or perform maintenance tasks, you can put the desktops into the maintenance mode. You can perform administrative tasks while the desktop is in the maintenance mode, such as applying patches and upgrades.

You can perform maintenance tasks on delivery groups or individual desktops.

User connectivity is affected as follows:

- Users can connect to the existing sessions on the Windows Server machine desktops but cannot start new sessions.
- Users cannot connect or reconnect to the Windows desktop machines. If they are already connected, they stay connected until they disconnect or log off.

XenDesktop regains control of the desktops when they are taken out of the maintenance mode.

To enable or disable the maintenance mode, perform the following steps:

1. Launch Studio.
2. Select the delivery group or select individual machines.
3. Click on **Turn On Maintenance Mode** or **Disable Maintenance Mode**.

Managing the server load

You can manage the server load for the Windows Server machines. When a user logs in to a Windows Server machine, load management assigns the desktop to the server that is best suited to handle the request.

Server selection is based on the following attributes:

- **Server maintenance mode status**
- **Server load index**: This is represented by the aggregated load based on the CPU utilization, memory utilization, and number of sessions

Load indices are calculated using the load evaluator formula.

- **Concurrent logon tolerance setting** or **allowed number of concurrent logon requests**

Windows Server machines will not be considered for load balancing under the following conditions:

- The maintenance mode is on
- RDC is set to **Don't allow connections to this computer**

- RDC is *not* set to **Don't allow connections to this computer** and the **Remote Host Configuration User Logon** setting is one of the following:
    - **Allow reconnections, but prevent new logons**
    - **Allow reconnections, but prevent new logons until the server is restarted**

The server load index

A server's load index determines how likely it is that a server is delivering the Windows Server machines to receive connections. It is a combination of the following factors:

- The number of sessions
- Performance metric settings for CPU, disk, and memory use

To monitor the server load, perform the following steps:

1. Launch Studio.
2. Use the **Search** feature and select a machine.
3. Scroll to the right-hand side of the window to display the **Server Load Index** column.

The **Server Load Index** column may be hidden.

The concurrent logon tolerance setting

You can control the number of pending logons that a server delivering Windows Server OS machines can concurrently accept with this setting. If all the servers have a setting higher than the **Concurrent logons tolerance** setting, the next logon request is assigned to the server with the lowest pending logons. If more than one server has low pending logons, then the lowest load index is considered. Load management settings are managed using policies.

Use the **Concurrent logons tolerance** setting to avoid server performance problems. An example is shown in the screenshot that follows.

To change load management settings, perform the following steps:

1. Launch Studio.
2. Select **Policy**.
3. Select **Create Policy** using the **New Policy** wizard.
4. Under **Categories**, select **Load Management**.
5. Select **Concurrent logons tolerance**.
6. You can select the default value or change the settings as shown in the following screenshot:

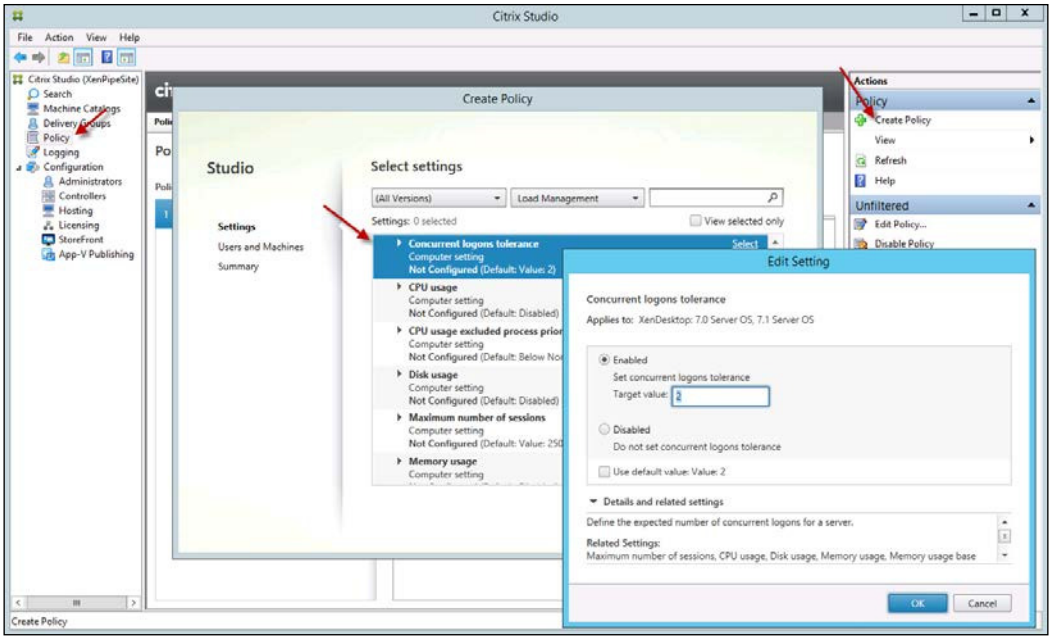

Managing the hosted applications

You can manage the hosted applications with the XenDesktop infrastructure. You can create and manage application desktop groups and the applications they host, manage the XenDesktop controller environment, configure hosts and connections, enable the use of smart cards, control user access and sessions, and monitor the applications. These settings can be controlled using Studio and the SDK.

Hosted applications are the applications that are installed on and accessed from the server, where the processing takes place. This is the traditional application publishing model. For many organizations, this provides the lowest cost of ownership for IT resources because it provides the highest scalability.

Published applications publish specific applications and deliver only these applications to users. This option provides greater administrative control and is used frequently. Applications are typically published using XenApp.

Hosted applications do not support thin clients. For thin clients, you need to publish applications using XenApp.

Application desktop delivery groups

When you create an application, you assign desktop groups to deliver the application to users. All the desktops in a desktop group publish the same application or set of applications. Desktops that host applications can be virtual or physical machines.

Application sharing

The Windows desktop machine's desktop groups don't support application sharing. If a user accesses an application from a desktop, the applications on that desktop are not available to other users. Other users must access the applications published by the desktop group from other desktops in the group if they are available.

A Windows Server machine's desktop groups allow application sharing. Applications published in the same desktop delivery group are shared in the same session. If session sharing is not supported, applications are launched in separate sessions.

Session sharing requires the following settings to be the same for all applications:

- **Color depth**
- **Encryption**
- **Audio quality**
- **Domain name**
- **Username**
- **Site name**

- **Special folder redirection**
- **Virtual COM port mapping**
- **Display size**
- **Client printer port mapping**
- **Client printer spooling**
- **EnableSessionSharing**
- **TWIDisableSessionSharing**

To determine whether the applications are compatible with each other for session sharing, use the `Get-BrokerSessionSharing IncompatibleApplication` cmdlet in the SDK.

Publishing applications to multiple desktop groups

By publishing the same applications to different types of desktop groups that contain different machine types, you can provide a different user experience for the application. For example, you might want some users to have the ability to customize the application and retain their changes while other users' changes can be discarded.

An application can be published from a private desktop group and a shared desktop group. If a user who has access to the application from both the desktop groups accesses the application, the application is launched from the private desktop group provided that there is a desktop available in that group. If a desktop is not available in the private desktop group, the application is launched from the shared desktop group.

Content redirection

An application can be configured to redirect content from the user device to the desktop hosting the application by associating file types with the application. When a user opens a file on their device, XenDesktop launches the application on a desktop hosting the application based on the file type association.

File types and their associated applications are stored in the XenDesktop Site's database. You can update the list of file types while configuring the content redirection for the application by importing the file types from the desktops in the desktop group assigned to an application. To update the file types, a desktop must be in the maintenance mode.

You can also associate the file types using the SDK.

Creating an application

Just as you have delivery groups to deliver desktops, you need a delivery group to deliver applications to the end users. The following steps will create an application:

1. Launch Studio.
2. Select **Delivery Groups**.
3. Click on the **Applications** tab and select **Create Application**.
4. Use the **Create Application** wizard to create the application as shown in the following screenshot:

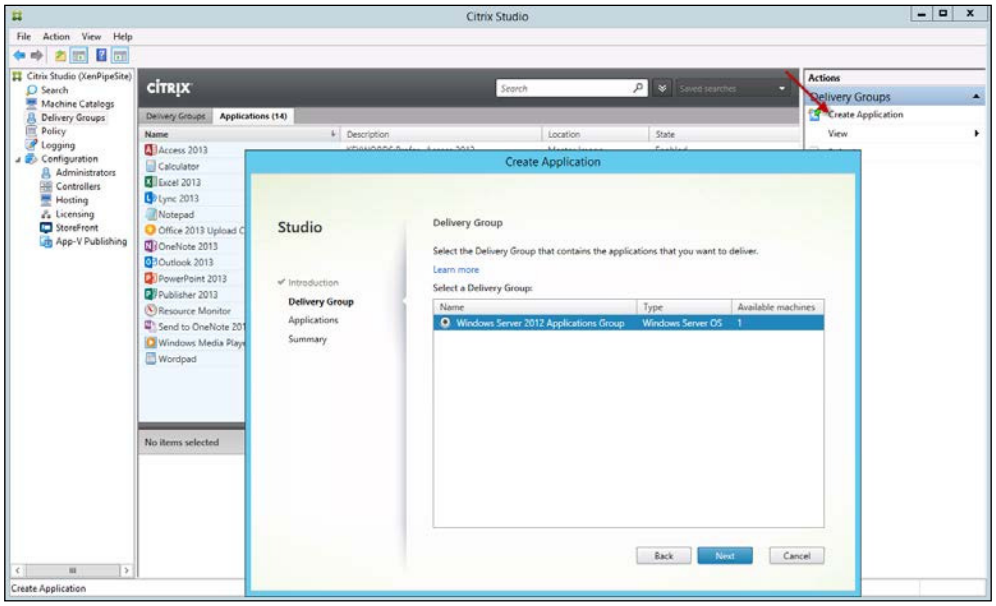

Managing and creating application desktop delivery groups

You can perform similiar functions on an application desktop delivery group to those you can perform on a desktop delivery group. For example, you can enable/disable the maintenance mode, find and search for applications using the search function, power manage application machines, shut down/restart, reallocate, import/export user data, remove application machines, and delete application machines from the catalog.

To create an application desktop delivery group, perform the following steps:

1. Launch Studio.
2. Select **Create Delivery Group**.
3. Use the **Create Delivery Group** wizard to create a desktop group for applications.

The **Users** page appears if the desktop group is based on shared or static existing machines or physical machines and the machine has not been allocated to other user accounts.

4. To give users access to the applications hosted on the desktops in a dedicated desktop group, give them access to the desktop group. On the **Users** page, add the users or user groups that can access the desktops and enter the number of desktops available to each user.
5. Select the **Applications** option to deliver applications with this delivery group.
6. Enter a name on the **Summary** page. The steps are displayed in the following screenshot:

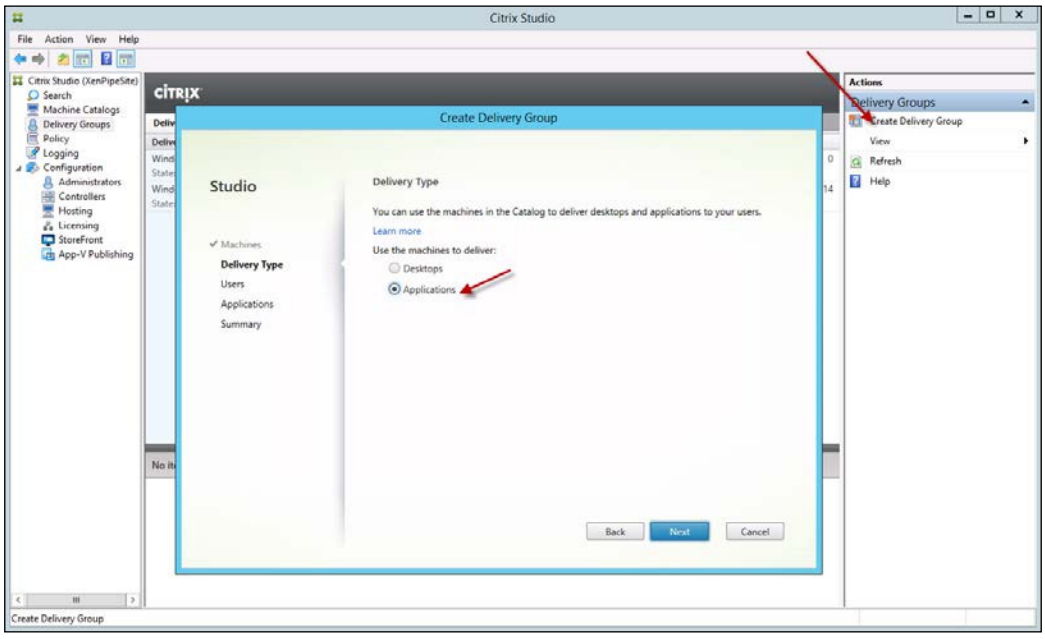

Managing application sessions

When a user requests for a hosted application, the user device links to the VDA on the desktop and establishes a session. You can control the session by logging users off, disconnecting sessions, and sending messages to users.

To log off or disconnect sessions, perform the following steps:

1. Launch Studio.
2. Select **Delivery Groups** and click on the **Applications** tab. Then, select the application.
3. Under **Sessions**, select the session or machine and click on **Logoff** or **Disconnect**.

To send messages to users, perform the following steps:

1. Launch Studio.
2. Select **Delivery Groups** and then click on the **Applications** tab. Select the application.
3. Click on **Sessions** and then select the session.
4. Click on **Send Message**.
5. Compose the message and click on **OK**.

Modifying the applications

You can modify the application properties, add or remove desktop delivery groups that host applications, add or remove users who can access applications, change the application name, and remove applications from desktop delivery groups.

To modify application properties, perform the following steps:

1. Launch Studio.
2. Select **Delivery Groups** and click on the **Applications** tab. Then, select **Properties**.
3. Use the wizard to modify the application.

To add or remove desktop delivery groups that host the application, perform the following steps:

1. Launch Studio.
2. Select **Delivery Groups** and click on the **Applications** tab. Then, select the application.
3. Click on **Edit Desktop Groups**.
4. Click on **Add** or **Remove**.

To add or remove users who can access the application, perform the following steps:

1. Launch Studio.
2. Select **Delivery Groups** and click on the **Applications** tab. Then, select the application.
3. Select **Properties**.
4. Click on **Limit Visibility**.
5. Select **Add** or **Remove**.

To modify the application name displayed to users, perform the following steps:

1. Launch Studio.
2. Select **Delivery Groups** and click on the **Applications** tab. Then, select the application.
3. Select **Properties**.
4. Choose **Identification**.

To remove applications from a desktop group, perform the following steps:

1. Launch Studio.
2. Select **Delivery Groups** and click on the **Applications** tab. Then, select the application.
3. Select **Properties**.
4. Click on **Delete**. The steps are summarized in the following screenshot:

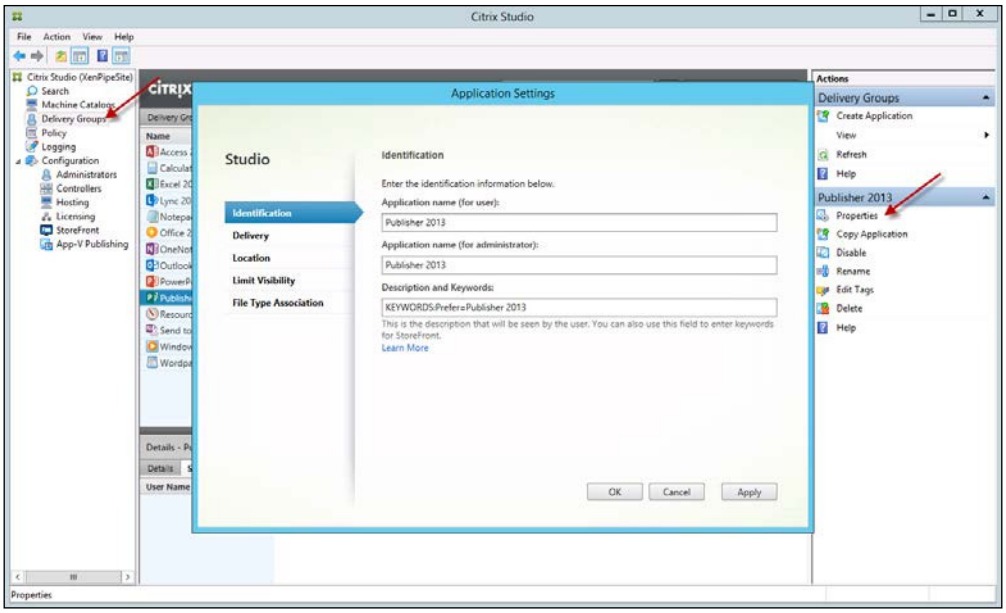

Managing the Delivery Controller environment

The Delivery Controller is the component that distributes desktops, manages user access, and optimizes connections. A XenDesktop Site can have one or more controllers depending on the size of the deployment.

Controller discovery

Before a desktop can be used, the **Virtual Delivery Agent** (**VDA**) must establish a connection with a Delivery Controller. The VDA finds a controller by checking the `ListOfDDCs` registry key. The `ListOfDDCs` registry key contains the DNS entries or IP addresses that point to the controllers for the XenDesktop Site. For load balancing, the VDA automatically distributes connections across all the controllers in the list.

The `ListOfSIDs` registry key indicates which machine **Security IDs** (**SIDs**) the VDA will allow to contact it as a controller. The `ListOfSIDs` registry key can be used to decrease the load on Active Directory or to avoid possible security threats from a compromised DNS server.

You need to make sure that the `ListOfDDCs` and `ListOfSIDs` registry keys are up to date because controllers are sometimes added and removed from the XenDesktop Site. Failing to do so could result in rejected session launches and launch delays. You can keep the list current by using the auto update feature, which is enabled by default, or by manually updating the policy or registry settings.

The VDA looks for the `ListOfDDCs` and `ListofSIDs` registry keys in the following locations and uses the first one it finds:

- When auto update is enabled:
    - A persistent storage location maintained for the auto update feature
- When auto update is disabled:
    - Policy settings (controllers; controller SIDs).
    - The controller information under the Virtual Desktop Agent in the registry. The VDA installer automatically populates these based on the controller information you specify when you install the DA.
    - OU-based controller discovery. This is a legacy method for backward compatibility.
    - The `Personality.ini` file created by **Machine Creation Services** (**MCS**).

If `ListOfDDCs` specifies more than one controller, the DA attempts to connect to them in a random order. The `ListOfDDCs` registry key can also contain controller groups. The DA attempts to connect to each controller in a group before moving to other entries in the `ListOfDDCs` registry key. These DDCs cannot be DNS aliases as XenDesktop checks if there is a corresponding Active Directory account.

Auto update is enabled by default.

The following types of deployments cannot use auto update and must use self management:

- Deployments that use controller groups
- Deployments that use a `ListOfSIDs` registry keys for security reasons
- Deployments that use the Provisioning Services without a write-back disk
- Deployments that use the controllers or controller SIDs policy setting

To enable or disable auto update, perform the following steps:

1. Launch Studio.
2. Navigate to **Policy | Create Policy**.
3. Set the **Enable auto update of controllers** policy to **Allowed** (enabled by default) or select the **Disable auto update of controllers** policy, as shown in the following screenshot:

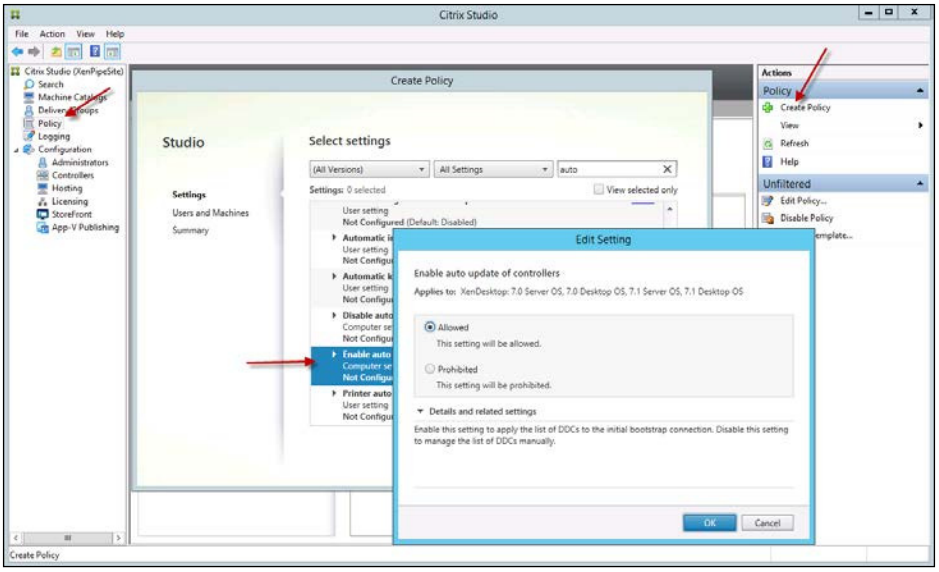

To self manage the controller connections using Studio, perform the following steps:

1. Launch Studio.
2. Navigate to the **Policy** section and locate the **Virtual Desktop Agent** section.
3. Find the controller's machine policy setting and update the **FQDN** values.
4. If you also use the `ListOfSIDs` registry key in your deployment, update the **SID** values.

To self manage the controller connections using registry settings, perform the steps that follow. The registry settings should be modified on each virtual desktop image after you add, move, or remove the Delivery Controllers in the Site:

1. Open the Registry Editor.

2. Update the `ListOfDDCs` registry key, which lists all the FQDNs of the controllers in the Site, using the following path. Separate multiple values with spaces and surround the controller groups with brackets:

   ```
   HKEY_LOCAL_MACHINE\Software\Citrix\VirtualDesktopAgent\ListOfDDCs
   (REG_SZ)
   ```

3. If a Site's OU was specified during the VDA installation, then you can also modify the `FarmGUID` keys.
4. You can also update the `ListOfSIDs` key (optional) using the following path:

   ```
   HKEY_LOCAL_MACHINE\Software\Citrix\VirtualDesktopAgent\
   ListOfSIDs (REG_SZ)
   ```

Use the Registry Editor at your own risk; we are not responsible for problems resulting from the improper use of the Registry Editor. Back up your Registry before editing.

Adding, moving, or removing Delivery Controllers

To add, move, or remove Delivery Controllers from the Site, you need the following roles and permissions:

- **Sysadmin or dbcreator database server role**: If you don't have these roles, you need the `CreateAnyDatabase` and `AlterAnyDatabase` server permissions
- **The db_owner or db_datawriter database user role**: If you do not have either of these roles, you need the `Insert, Delete, and Update user` permission

Before you add, move, or remove controllers and if you use database mirroring, make sure that both the principal and mirrored databases are running. Also, if you are executing scripts using the SQL Server Management Studio, enable the `SQLCMD` mode.

After adding, moving, or removing a Delivery Controller, if auto update is enabled, the DAs will receive an updated list of controllers within 90 minutes.

After adding, moving, or removing a Delivery Controller, if using self manage, ensure that the controller policy settings or the `ListOfDDCs` registry key is updated for all virtual desktops.

If you move a controller from one Site to another, you will need to update the registry settings on both the Sites.

To add a Delivery Controller, perform the following steps:

1. On the server that you want to add the Delivery Controller to, run the XenDesktop installer and select **Delivery Controller** along with the other components.
2. Launch Studio.
3. Select **Join Existing Deployment** and enter the Site address.

To move a Delivery Controller to a different XenDesktop Site, perform the following steps:

1. On a controller in the old Site, launch Studio.
2. Navigate to **Configuration** | **Controllers** and then select the controller that you want to move.
3. Select **Remove Controller** under **Actions**.
4. On the controller that's being moved, launch Studio.
5. Reset services when prompted.
6. Join the existing Site.

To remove a Delivery Controller, perform the following steps:

1. Launch Studio.
2. Navigate to **Configuration** | **Controllers** and then select the controller that you want to remove.
3. Select **Remove Controller** under **Actions**.

Adding, moving, and removing controllers is not backward compatible with the previous versions of XenDesktop. A XenDesktop Site requires at least one controller; so, you cannot move or remove the last controller in the Site.

Moving a Virtual Delivery Agent (VDA) to another Site

You might need to move a VDA from one Site to another when upgrading or when moving a DA image from a test Site to a production Site.

For moving a VDA with the installer, perform the following steps:

1. Run the XenDesktop installer on the image.
2. Specify a valid DNS entry or the IP address of a Delivery Controller for the new Site.

For moving a VDA with policies, perform the following steps:

1. Create a policy in the old Site with the following settings:
    - **Controllers**: This contains the DNS entries or IP addresses of one or more controllers in the new Site
    - Set **Enable Auto update of controllers** to **Disabled**
2. Apply the policy to the desktop delivery group to stage a migration.
3. In 90 minutes, every VDA in the desktop group is alerted of the new policy, but the DA ignores the list because **Enable auto update of controllers** is disabled. Yet, it selects one of the controllers in the policy, which is in the new Site.
4. When the VDA registers with the controller in the new Site, it receives the new Site's `ListOfDDCs` and policies, which have **Enable auto update of controllers** set to **Enabled**.

Active Directory OU-based controller discovery

This discovery method is primarily for backward compatibility and is valid only for DAs for Windows desktop machines and not for DAs for Windows Server machines. Active Directory OU-based discovery requires that all the computers in a Site be members of a domain, with mutual trust relationships between the domain used by the controller and the domain(s) used by desktops. If you use this method, you must configure the GUID and OU in each desktop registry.

To perform OU-based controller discovery, run the following PowerShell script on the controller located at `$Env:ProgramFiles\Citrix\Broker\Service\Setup Scripts`. You must have the `CreateChild` permissions on a parent OU plus full administrator rights.

```
Set-ADControllerDiscovery.ps1
```

When you create a Site, a corresponding OU must be created in Active Directory if you want desktops to discover the controllers through Active Directory.

The PowerShell script creates several objects that are essential for the operation of the farm. It is not necessary to extend the schema. The objects that are created by the PowerShell script along with their description are listed in the following table:

Active Directory objects	
Object	**Description**
Controllers security group	The computer account of all controllers in the Site must be a member of this security group. Desktops in a Site accept data from a controller only if they are members of this security group. Be sure to give the **Access this computer from the Network** privilege to all the virtual desktops that run the DA.
Service Connection Point (**SCP**)	This contains information about the Site such as the Site's name. Enable the advanced features in the **Active Directory Users** and **Computers** listing to see the SCP objects.
RegistrationServices container	This contains one SCP object for each controller in the Site, which is validated and updated each time the controller starts. This object is created in the Site's OU.

Administrators need to create and delete the `children` permissions on the `RegistrationServices` container and write permissions on the controllers' security group. These are automatically granted by running the PowerShell script `Set-ADControllerDiscovery.ps1`.

When using a Site's OU, information is written to Active Directory when installing or uninstalling XenDesktop. It is also written when a controller starts and needs to update information in its SCP. By default, the `Set-ADControllerDiscovery.ps1` script creates permissions on the objects in the Site's OU, providing each controller a write access to its SCP. The contents of the objects in the Site's OU are used to establish trust between the desktops and controllers.

Try to ensure that only authorized administrators can add or remove computers from the controllers' security group by using the security group's **Access Control List** (**ACL**). Also, make sure that the respective controller can change the information in the controller's SCP.

Also note that replication can cause delays. Changes made to Active Directory when initially creating the OU, installing or removing controllers, and so on, may not be visible to desktops until the information is replicated to the domain controller.

To move a controller using OU-based controller discovery, perform the following steps:

1. On a controller in the old Site, launch Studio.
2. Navigate to **Configuration** | **Controllers** and then select the controller you want to move.
3. Select **Remove Controller** under **Actions**.
4. Run the following PowerShell script on the controller, located at `$Env:ProgramFiles\Citrix\Broker\Service\Setup Scripts`. You must have the `CreateChild` permissions on a parent OU plus full administration rights.

   ```
   Set-ADControllerDiscovery.ps1
   ```

Using SSL on controllers

The XML service runs on the Delivery Controller, and it supports both the HTTP and HTTPS protocols. The XML service supports **Secure Sockets Layer** (**SSL**) using server certificates. You must obtain and use a valid server certificate on all the controllers.

Changing the default HTTP and HTTPS ports

By default, the XML service listens on port `80` (HTTP) and port `443` (HTTPS). You can use different ports for HTTP and HTTPS communication; however, beware of the security risks of exposing the controller to untrusted networks.

To change the default HTTP and HTTPS ports, run the following command on the Delivery Controller:

```
BrokerService.exe –WIPORT <http port> –WISSLPORT <https port>
```

To ignore HTTP and HTTPS on the default ports, set the registry keys `HKLM\Software\Citrix\DesktopServer\XmlServicesEnableNonSsl` and `XmlServicesEnableSsl` to `0`.

Summary

In this chapter, we covered delivery groups, an important concept with regard to XenDesktop and delivering desktops and applications. After reading these last two chapters, you now have the resources to understand how to create and deliver desktops and applications and how to manage them. In the next chapter, we will look at polices and how to manage them, considering that sessions are managed according to policies.

5
Managing Policies

Everything in XenDesktop is done with policies, at least when it comes to giving users access and managing sessions. At this point in the book, we have provided everything you need to know when it comes to giving users access and managing sessions. Now, we need to manage the policies created in the previous chapters. Citrix policies are the best way to control connections, security, and other settings in XenDesktop, and we will discuss them in this chapter.

XenDesktop policies are very flexible; you can create policies for users, groups of users, specific devices, or types of connections. A policy can include several settings; for example, you might want to create policies to perform the following tasks:

- Monitor CPU usage
- Monitor **Independent Computing Architecture** (**ICA**) latency
- Monitor profiles
- Permit user's access to the documents on their local devices
- Permit/block users from saving data to their hard drives
- Permit/block users from accessing clipboard and USB drives

XenDesktop® Studio versus Microsoft Group Policy Editor

There are two ways to manage policies in XenDesktop. You can use XenDesktop Studio or the Microsoft Group Policy Editor. Which tool you should use depends on your Site. If you have Active Directory installed and the appropriate permissions to manage **Group Policy Objects** (**GPOs**), you might want to use the Microsoft Group Policy Editor. If you are a Citrix administrator and don't have permissions to manage group policies, you might want to use XenDesktop Studio to create policies. If you use Studio to create policies, they get stored in the XenDesktop database and updates are applied to the desktop when the desktop gets registered or when a user connects to the desktop.

If your Site has Active Directory installed and you have permissions to manage the group policies, you should use the Microsoft Group Policy Editor. The settings that you configure will affect the GPOs you select in the Group Policy Editor. The policies created in the Group Policy Editor are stored on the domain controller, and the updates are sent out to the virtual desktops at standard intervals, which are set as part of the GPO refresh policy.

In an Active Directory environment, the Active Directory GPOs take precedence over the Citrix Site policy settings. Policy changes do not affect users who are already connected to their desktops. Policy updates take place when the users log on or reconnect to their desktops.

Administrative roles

There are two types of XenDesktop administrator roles as follows:

- **Full admin**: This role has all the administration permissions with the authority to control all the facets of policy administration, including policy creation, management, editing, and modeling.
- **Read-only admin**: This role allows you to view all the aspects of policy administration but does not have the authority to change any settings. A read-only administrator can run the **Policy Modeling** wizard.

The afore mentioned roles also apply when you use PowerShell to configure policies using the command-line interface.

Working with policies

It is likely that you will create and enable more than one policy in your implementation. If you do, make sure that you prioritize the policies so that it is clear which policy should be executed first.

Policies can be configured using the following simple steps:

1. Create the policy.
2. Configure the policy settings.
3. Apply filters to the policy.
4. Set a policy priority.

Use priorities with your policies to make sure that they are executed correctly. Also, note that unfiltered policies take precedence over filtered policies.

Navigating policies

In XenDesktop Studio, policy settings are grouped into two main categories as follows:

- **Machine policy settings**: These settings are used to control the behavior of virtual desktops and are applied when a virtual desktop starts. These settings are applied even if there are no active user sessions on the desktop.
- **User policy settings**: These settings are used to control the user experience when connecting to desktops using the `ICA/HDX` protocol. User policies are applied whenever a user connects or reconnects to a desktop using `ICA/HDX`. Keep in mind that if a user connects to a desktop directly at the console or using RDP, the user policies are not enforced.

Policies and settings are similarly mapped in Active Directory, and if you use Microsoft Group Policy Editor, you will notice that they are organized into the categories of **Computer Configuration (Machine Policy)** and **User Configuration (User Policy)**.

Note that the top-level names for policies are different in Studio (machine versus user) and the Group Policy Editor (computer versus user); however, the individual policy settings are the same.

Accessing policies

Policies and their settings can be found in Studio by clicking on the **Policy** node in the left-hand side navigation pane. You then click on **Create Policy** in the right-hand side navigation pane, and select the relevant policies in the **Settings** page. You can assign them to the user and machine objects or to all the objects in a Site, as shown in the following screenshot:

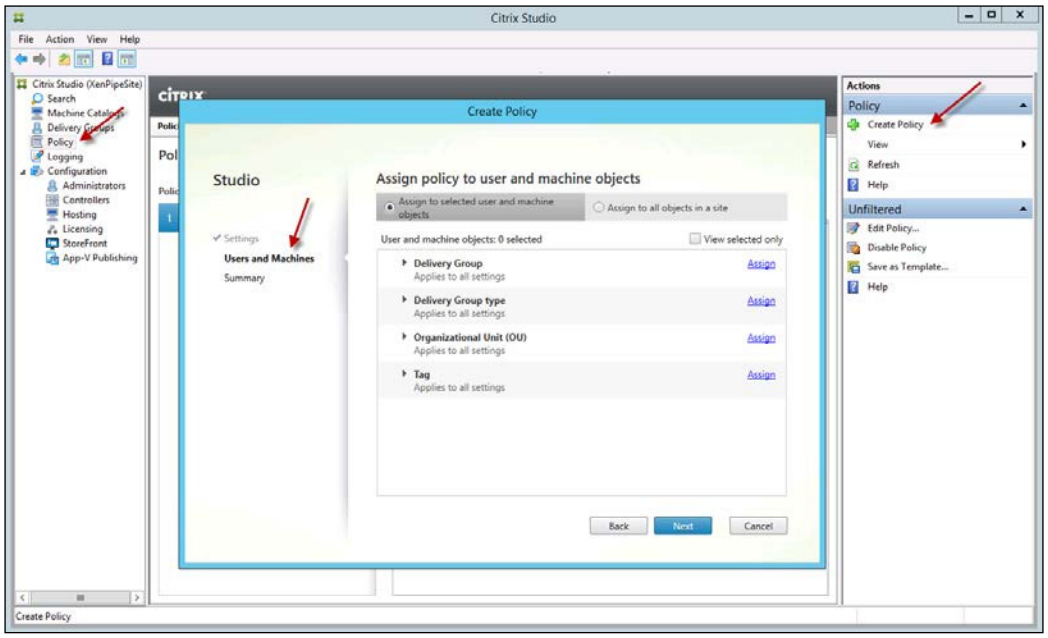

The Microsoft Group Policy Editor can be used to configure policies by selecting the **Citrix Policies** node under **Computer Configuration** or **User Configuration**.

Launch the Microsoft Group Policy Editor by searching for and running `gpedit.msc` on the Desktop Delivery Controller server.

The **Citrix Policies** node displays a list of the policies available for the Site. You can display the policy settings with **Summary**, **Settings**, and **Filters** as shown in the following screenshot:

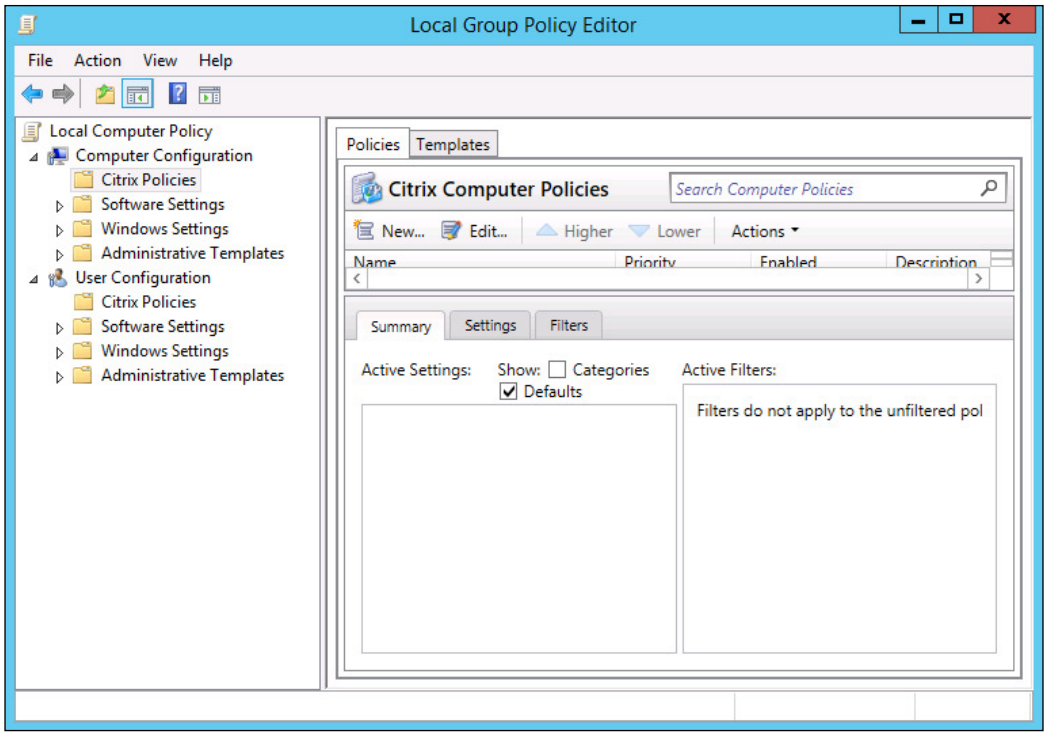

Searching policies

From each of these consoles, Studio or Group Policy Editor, you can search through the policies. You will find the **Search** tool in Studio in the upper-right corner on the **Settings** page of the **Policy** wizard. In the Group Policy Editor, the search tool appears in the upper-right corner of the window and also when **Citrix Policies** is selected in the left-hand side navigation pane. In the Group Policy Editor, you can extend your search using the **Settings** and **Filters** tabs.

Creating policies

Before you create a policy, it is good practice to determine on which group of users or devices you want to apply the policy. For example, you might want to create policies that are based on the end users' job function, device, connection type, or location. In fact, you might even want to use the same methodology that has been used for your Windows Active Directory installation. It is good practice to minimize the number of policies you create by using a few policies that apply to a group as opposed to creating a new policy just to enable or disable a specific setting.

Creating a policy in Studio

To create a policy in Studio, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node in the left-hand side navigation pane.
3. Click on **Create Policy** in the right-hand side navigation pane.
4. Select the policies that you want to configure.
5. Assign the policies to selected users and machines or to all the objects.
6. Give a name to the policy and click on **Finish**, as shown in the following screenshot:

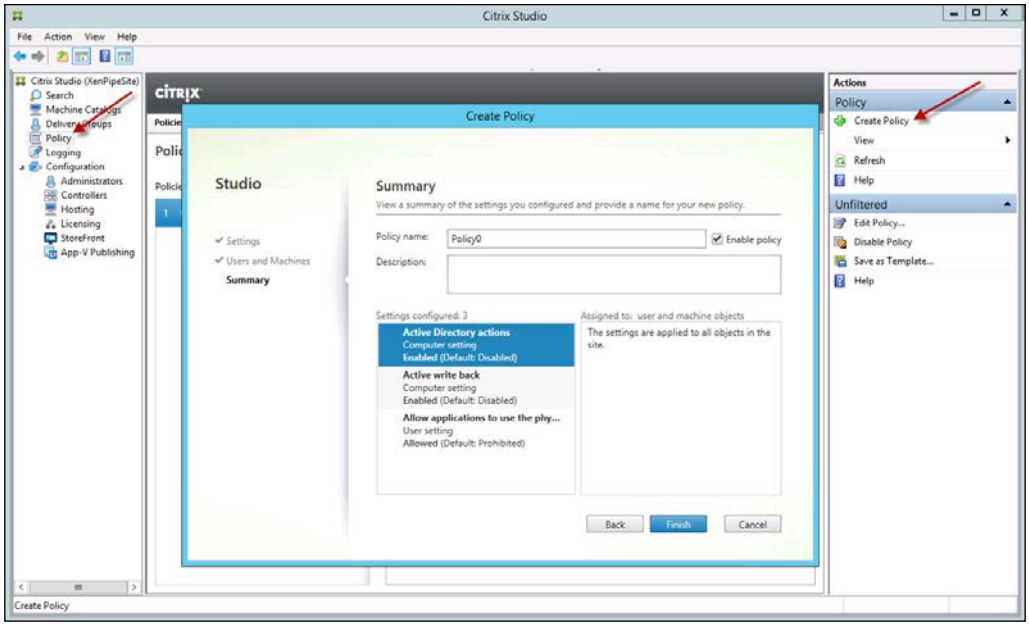

Creating a policy in Microsoft Group Policy Editor

To use **Citrix Policies** in Active Directory, you need to install the Citrix Group Policy Management component on one of your group policy management servers. To install the Group Policy Management component, locate the appropriate `CitrixGroupPolicyManagement` MSI in the installation media under the `\x86\Citrix Policy or \x64\Citrix Policy` directories. Install `CitrixGroupPolicyManagement_x86.MSI` or `CitrixGroupPolicyManagement_x64.MSI` for 32-bit or 64-bit computers, respectively.

The ADM and ADMX files are located on the installation media in the `\x86\ProfileManagement` or `\x64\Profile Management` directories, respectively. The ADM files can be loaded on a local domain controller or in an Active Directory partition for AD replication.

To create a policy in Microsoft Group Policy Editor, perform the following steps:

1. In Windows 2012, you can use the search bar on the desktop to search for `gpedit.msc`, or you can manually run the **Microsoft Management Console** (**MMC**) and add the Group Policy Editor as a snap-in.
2. Select **Citrix Policies** under **Computer Configuration** or **User Configuration**.
3. Select **New**.
4. Follow the **New Policy** wizard to perform the following steps:
    1. Choose the settings you want to configure.
    2. Set the filters that determine to whom/what the policy should be applied.
    3. Mark the checkbox to enable or disable the policy.

    The steps are summarized in the following screenshot:

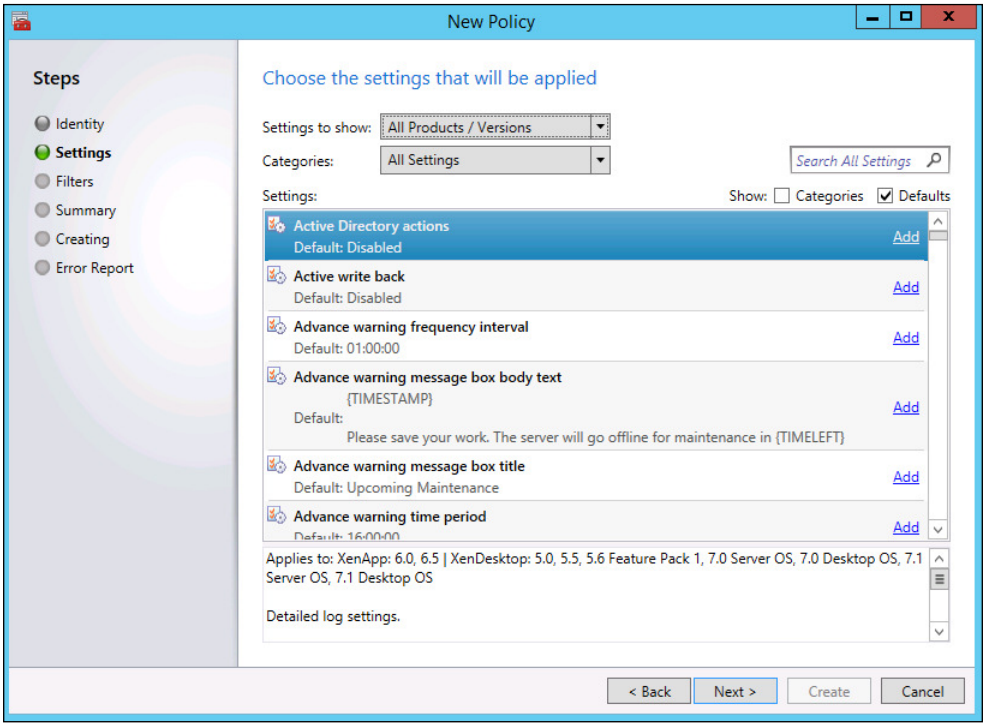

Enabling the policy allows it to be immediately applied to the users who are logging in to their virtual desktops. You may want to add some priorities to the policies so that they are executed in the correct order or you may want to fine-tune the settings at a later time, in which case you should disable the policy temporarily.

Configuring policies

Connections are controlled with the policies that contain settings. Policy settings have three states—not configured, enabled, or disabled. By default, policies and settings are not configured, which is why you will notice that the policies need to be selected the first time you create a policy in the wizard. To put it in another way, settings are not added to a policy automatically. Settings can only be applied after they have been added to a policy.

Configuring policy settings

Policy settings can have two states—allowed/enabled or prohibited/disabled. If a policy setting is set to **Allowed**, it will be applied. If a policy setting is set to **Prohibited**, it will not be applied. The same is true in the case that if you enable a policy setting, it becomes active; and if you disable a policy setting, it is deactivated. You will find a mix of allowed/enabled and prohibited/disabled in the Citrix policy terminology.

The typical use case for policy settings is to allow or prevent users from doing the action specified in the setting. For example, if the USB drive redirection setting is set to **Allowed**, then users will be able to access their USB drives on their local client.

Some policy settings depend upon other policy settings. For instance, the **Client USB drive redirection** setting dictates whether users can access the drives on their local devices. For users to access their network drives, both the **Client USB drive redirection** setting and the **Client network drives** setting must be enabled and added to the policy you create. If the **Client USB drive redirection** setting is disabled, users cannot access their network drives, even if another policy setting, **Client network drives**, is enabled.

Best practices for designing policy settings

The following are some basic best practices to follow when designing policy settings:

- If you have multiple XenDesktop Sites, you might want to have the basic or common policies set within the Microsoft Group Policy Objects and then have more Site-specific policies set within XenDesktop Studio.
- Assign policies to groups and not individual users. By using groups, policies are updated for the users in that group automatically when you add or remove users from the group.
- Avoid conflicting settings in the Remote Desktop Session Host Configuration because it specifies a similar functionality to the Citrix policy settings.
- Keep policy settings consistent for easy troubleshooting.
- Disable unused policies because they create unnecessary processing.

Applying policies

Policies go into effect in the following circumstances:

- The virtual desktop starts
- The user logs on

Also, if you are using Active Directory in your Site, then policy settings will go into effect when Active Directory re-evaluates policies at the regular 90 minute intervals.

You can force policies to be updated by issuing the following command:

```
gpupdate /force
```

Using default values

For many of the policy settings, a default value exists. When you select **Use default value**, it disables the configuration of the setting and specifies that only the default value will be used when the policy is applied.

For a complete list of all the Citrix policy settings and their default values, please refer to *Appendix B*, *XenDesktop® Policy Settings Reference*.

Using filters

In Studio, filters are applied in the **Users and Machines** page of the **Policy** wizard. In Microsoft's **Group Policy** wizard, filters are applied on the **Filters** page. When you add a filter to a policy, the policy's settings are applied only to the connections that match the criteria. If no filter is created, then the policy is applied to all connections.

You can add as many filters as you want to a policy; however, only certain filters are available depending on whether it is a machine policy or a user policy.

The following is a list of the available filters:

The policy filters

Name	Description	Scope
Access control	This applies a policy based on the access control conditions through which a client connects.	User policies
CloudBridge	This applies a policy based on whether or not the user session is launched through a Citrix CloudBridge for WAN optimization.	User policies
Client IP address	This applies a policy based on the IP address (IPv4 or IPv6) of the user device connecting to the session. Examples for IPv4 are as follows: • `10.1.0.0` • `10.1.0.*` • `10.1.0.1-10.1.0.80` • `10.1.0.20/24` Examples for IPv6 are as follows: • `2002:0ec9:4d5e:002a:0:0:bead:ab21` • `2002:0ec9:4d5e:`	User policies
Client name	This applies a policy based on the name of the user device connecting to the session.	User policies
Desktop group	This applies a policy based on the desktop group membership of the desktop session.	Machine and user policies
Desktop type	This applies a policy based on type of the desktop session.	Machine and user policies

The policy filters		
Name	**Description**	**Scope**
Organizational unit	This applies a policy based on **Organizational Unit** (**OU**) of the desktop session.	Machine and user policies
Tag	This applies a policy based on the tags applied to the desktop session.	Machine and user policies
User or group	This applies a policy based on the user or group membership of the user connecting to the session.	User policies

The preceding table can be found at `http://support.citrix.com/proddocs/topic/xendesktop-71/cds-policies-applying-rho.html`.

XenDesktop finds policies that match the filters for a connection when a user logs on. XenDesktop then classifies the policies according to their priority, compares policy settings, and deploys them according to the priority.

If you are using Active Directory, user policy settings are pushed out when Active Directory re-evaluates the policies at regular 90 minute intervals and also when a user logs on.

A disabled policy setting takes precedence over a low-ranking policy setting that is enabled. Policy settings that are not configured are ignored.

There is a known issue where the Group Policy Management Console filtering mechanism may not work with groups when you apply machine policies. The workaround is to use Citrix Studio. You can read more about this at `http://support.citrix.com/article/CTX127461`.

Unfiltered policies

You will notice after the initial installation that there is an unfiltered policy for both machine and user policy settings. Any policy setting added to this policy affects all the connections.

If you use Citrix Studio to manage Citrix policies, the settings you add to the unfiltered policy get applied to all the virtual desktops and connections.

If you use Microsoft Group Policy Editor to manage your policies, the policy settings that are added to an unfiltered policy are deployed to all the Sites and connections that are part of the GPOs that contain the policy. For example, the sales OU contains a GPO named Sales-US, which includes all the members of the US sales team. The Sales-OU GPO contains an unfiltered policy that contains several user policy settings. When a member of Sales-OU logs into the Site, the settings in the unfiltered policy are applied automatically to the session because the user is a member of the Sales-US GPO.

Filter modes

When you assign filters, you will notice that the mode can be set to **Allow** or **Deny**. The mode governs how the policy is applied or not applied to the connections that match the filter criteria. If the mode is set to **Allow** (default), the policy is applied to connections that match the filter criteria. However, if the mode is set to **Deny**, the policy is applied if the connection does not match the filter criteria. This can be confusing; so, the following sections will provide some example use cases.

Using the same filters with different modes

Let's take an example. In a policy where the filters are the same but have different modes, **Allow** and **Deny**, the filter set to **Deny** takes precedence.

A detailed example is as follows:

- In policy 1, filter 1 is a user filter for the accounting group. Its mode is set to **Allow**.
- In policy 1, filter 2 is a user filter for the accounting group's manager account. Its mode is set to **Deny**. Because filter 2 in policy 1 is set to **Deny**, the policy is not applied, even if the accounting manager logs on. This denies access to the accounting manager's resources for everyone. This is because **Deny** takes precedence. You would have to explicitly allow the account of the accounting manager to give them access. In other words, you would have to have a policy X with a filter for the accounting manager, which is set to **Allow the connection**.

Using different filters with similar modes

On the other hand, in a policy where the filters are different but have the same modes, **Allow** and **Allow**, both the filters must match for the policy to be applied.

A detailed example is as follows:

- In policy 2, filter 3 is a user filter for the accounting group. Its mode is set to **Allow**.

- In Policy 2, filter 4 is a client IP address filter for the subnet `192.168.10.*` with the mode set to **Allow also**. This policy is applied when anyone in the accounting group logs in from the private network `192.168.10.*` because both filters evaluate to true.

Implementing multiple policies

You can use many different policies to adapt XenDesktop to users' requirements based on their job function, physical location, or connection type. For example, you may want to put restrictions on user groups that work with sensitive data. You can create policies that prevent users from saving files to their local media. However, if users in the same group need access to their local media, you can create a different policy just for those specific users. You would then configure priorities on the two policies to control which one takes precedence. When implementing multiple policies, you will need to ascertain how to prioritize them, how to create exceptions, and how to view the resultant policy when there is a conflict. It is a good idea to create an environment in a Proof-of-Concept lab setting to test your policies before rolling them into production.

Also, Citrix policies interact with Windows operating system policies and some Windows policies take precedence, especially with regard to security.

Implementing priorities

The prioritization of policies gives you the ability to specify how the policies will apply or take precedence when they contain conflicting settings. The process for evaluating policies is as follows:

1. Suppose a user logs in to their virtual desktop and policies for which a connection has been identified. XenDesktop goes through a process of sorting the policies in the order of priority and then compares and applies the settings according to the priority number.
2. The prioritization of policies is done by assigning them different priority numbers. You assign the priority numbers and XenDesktop evaluates them. New policies are automatically assigned the lowest priority. If a conflict arises between the policy settings, obviously a setting with a higher priority will take precedence over a policy with a lower priority.
3. XenDesktop then goes through a merge process to sort the settings according to their priority and the settings' condition. Disabled settings (**Deny**) always override settings that are enabled if the enabled setting has a lower priority. Otherwise, settings that are not configured are ignored and do not override any other settings.

Prioritization starts with the lowest number having the highest priority, and a larger number having a lower priority. For example, a priority of `1` is the highest priority you can assign to a policy setting. A priority of `100` would have a lower priority.

To change the priority using Studio, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Select the policy you want to modify.
4. Click on **Higher Priority** or **Lower Priority** in the right-hand side navigation pane, as shown in the following screenshot:

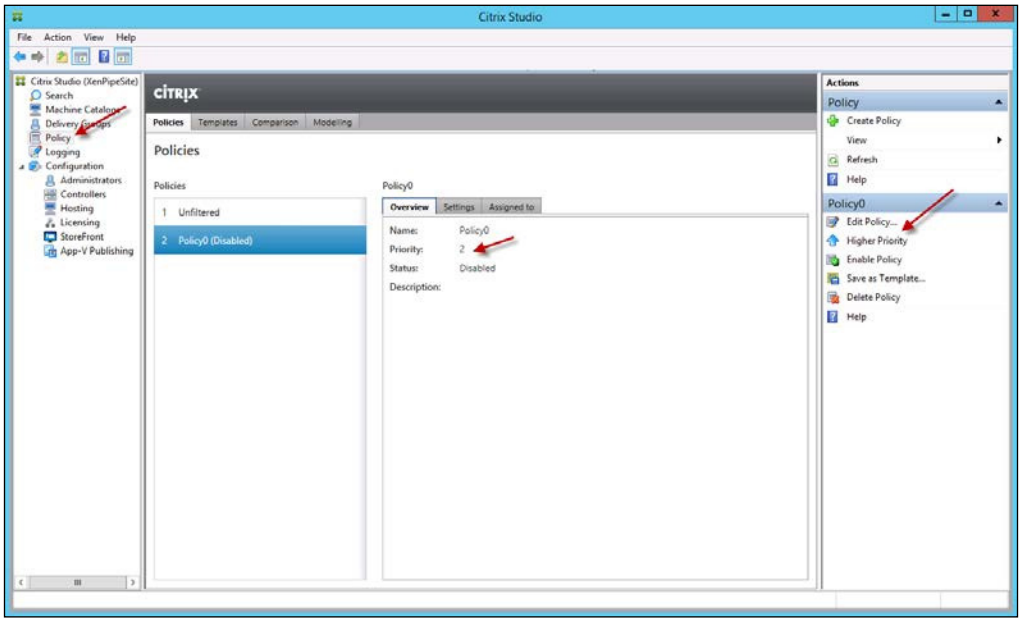

To change the priority using Microsoft Group Policy Editor, perform the following steps:

1. Launch Microsoft Group Policy Editor (`gpedit.msc`).
2. Highlight the Citrix policy.

3. Right-click to change the priority as shown in the following screenshot:

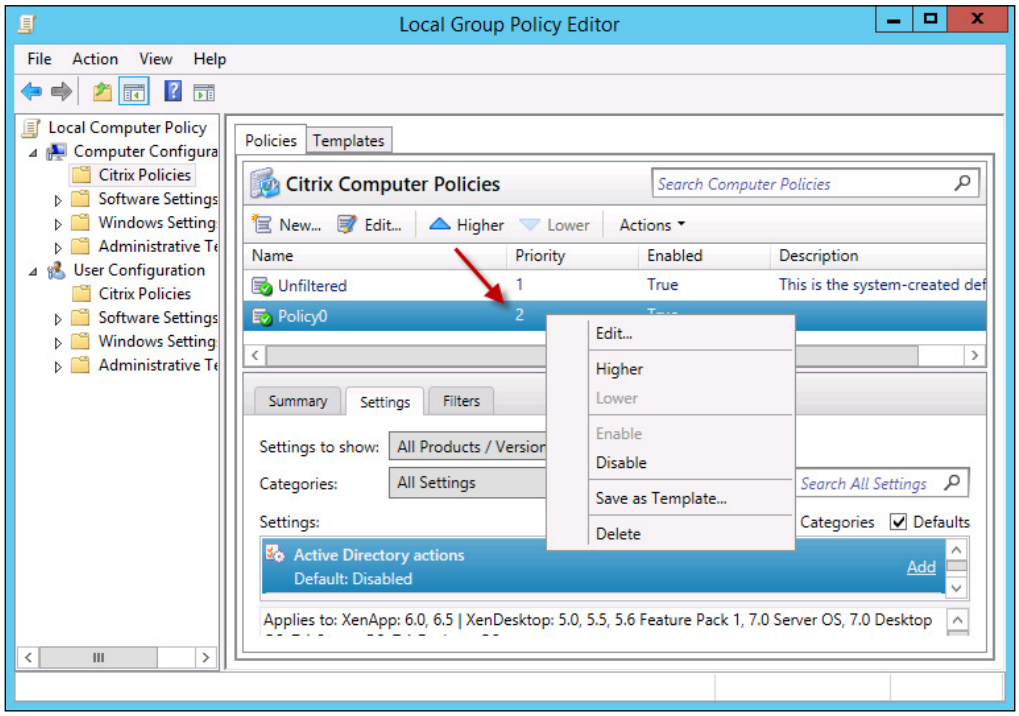

Implementing exceptions

When you create policies, you may need to create exceptions for specific use cases, as was the case in the previous example where we wanted only the accounting managers to have access to their account. You can create exceptions by performing the following steps:

1. Create an exception group; for example, accounting managers. Create a policy for the exception group. Then, rank this policy higher than the policies for the rest of the group.
2. Use an **Allow** filter on a policy for this group.

This will apply the policy to the accounting managers only. Alternatively, you could create an exception group for all the accounting employees and use a **Deny** filter.

A detailed example is as follows:

- In policy 1, filter 1 is a client IP address filter for the subnet `12.1.2.*`, set to **Allow**
- In policy 1, filter 2 is a user filter for the accounting managers, set to **Allow**

Policy 1 is applied to users who log in from the `12.1.2.*` subnet and are accounting managers.

Another detailed example is as follows:

- In policy 2, filter 1 is a client IP address filter for the subnet `12.1.2.*`, set to **Allow**
- In policy 2, filter 2 is a user filter for all the accounting employees, set to **Deny**

Policy 2 is not applied to the regular accounting users who log in from `12.1.2.*`.

The resulting set of policies

If you've been around Windows and policies for any length of time, you will be familiar with the resultant set of policies. Sometimes, unexpected things happen when multiple policies are in effect with different priorities. It can be very confusing and sometimes difficult to keep track of even on a whiteboard. There is a tool that helps you out with this.

You can ascertain ultimate policy settings by evaluating the Resultant Set of Policies.

Running the Citrix® Group Policy Modeling Wizard

The **Citrix Group Policy Modeling Wizard** is typically used to model a set of conditions for specific connection scenarios. For example, you may want to model how users who log in remotely over a slow connection connect and the policies that get applied to them. The wizard produces a report with the policies that will be applied to these users.

If you are logged into the XenDesktop controller as well as authenticated against Active Directory and are a domain user, the **Citrix Group Policy Modeling Wizard** will calculate the Resultant Set of Policies using the XenDesktop Site policy settings and the Active Directory GPOs together.

If you are logged on to the XenDesktop controller as simply a local user and run the **Citrix Group Policy Modeling Wizard** from within Citrix Studio, the wizard evaluates the Resultant Set of Policies using only the XenDesktop Site policy settings.

To run the **Citrix Group Policy Modeling Wizard**, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node in the left-hand side navigation pane.
3. Select the **Modelling** tab.
4. Click on **Launch Modeling Wizard** in the right-hand side navigation pane.
5. When the modeling is complete, click on **View Modeling Report**, as shown in the following screenshot:

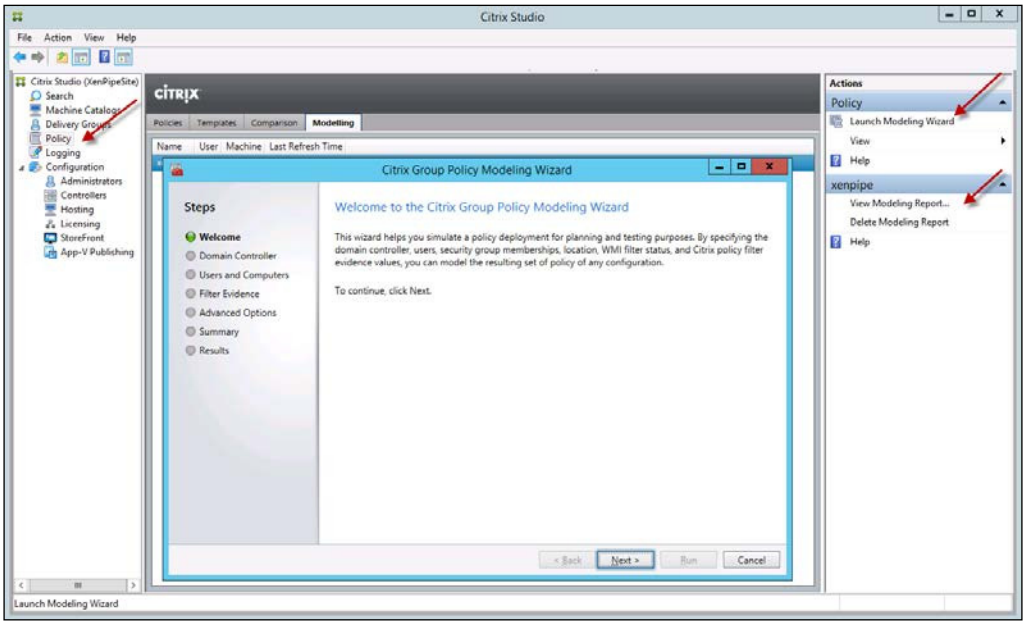

Running the Microsoft Group Policy Results tool

The Microsoft Group Policy Results tool (`GPResult.exe`) can also help you to evaluate the state of the GPOs in your Site. It can generate a report that shows how policy objects are being applied, including Citrix policies.

If you run the **Citrix Group Policy Modeling Wizard** or Microsoft Group Policy Results tool from the Microsoft Group Policy Management Console, the XenDesktop Site policy settings created using Studio are not included in the Resultant Set of Policies. Therefore, it is best practice to run the **Citrix Group Policy Modeling Wizard** from within XenDesktop Studio.

To run the Microsoft Group Policy Results tool, perform the following steps:

1. Open a command prompt or PowerShell window.
2. Run the `GPResult.exe` command with the `/R` switch as shown in the following screenshot:

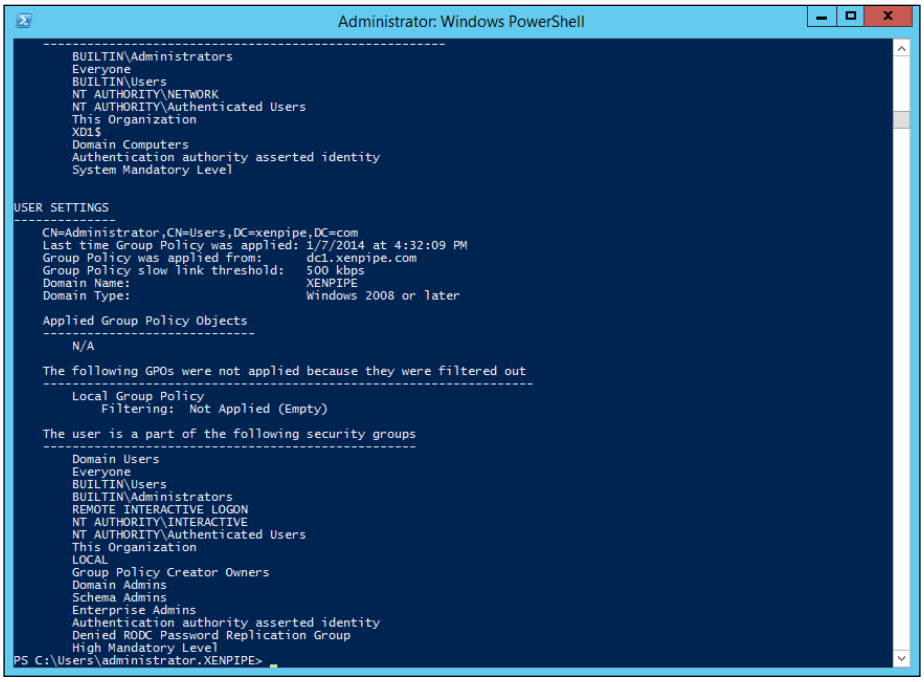

Troubleshooting policy scenarios

As you know, policies can become complex and sometimes conflict, which causes them to behave unexpectedly. When multiple policies have to be applied to a session, XenDesktop uses the resultant policy.

If you run the **Citrix Group Policy Modeling Wizard** or the Microsoft Group Policy Results tool and there are no reported policy settings, then the users' connection to virtual desktops will not be affected by any policies.

Users are not affected by policies under the following conditions:

- No policies have filters that match the policy evaluation criteria
- Policies with matching filters do not have any settings configured
- Polices with matching filters are disabled

When you apply policies to sessions, keep the following factors in mind:

- Ensure that the policies are enabled for those session connections
- Ensure that the policies contain valid and appropriate settings

Comparing policies

In Citrix Studio, you can compare policies side by side. This is a great tool that helps you to visualize policy settings and how they might affect the outcome of your deployment before implementing them.

To compare policies, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node in the left-hand side navigation pane.
3. Click on the **Comparison** tab.
4. Select the policies and run the comparison, as shown in the following screenshot:

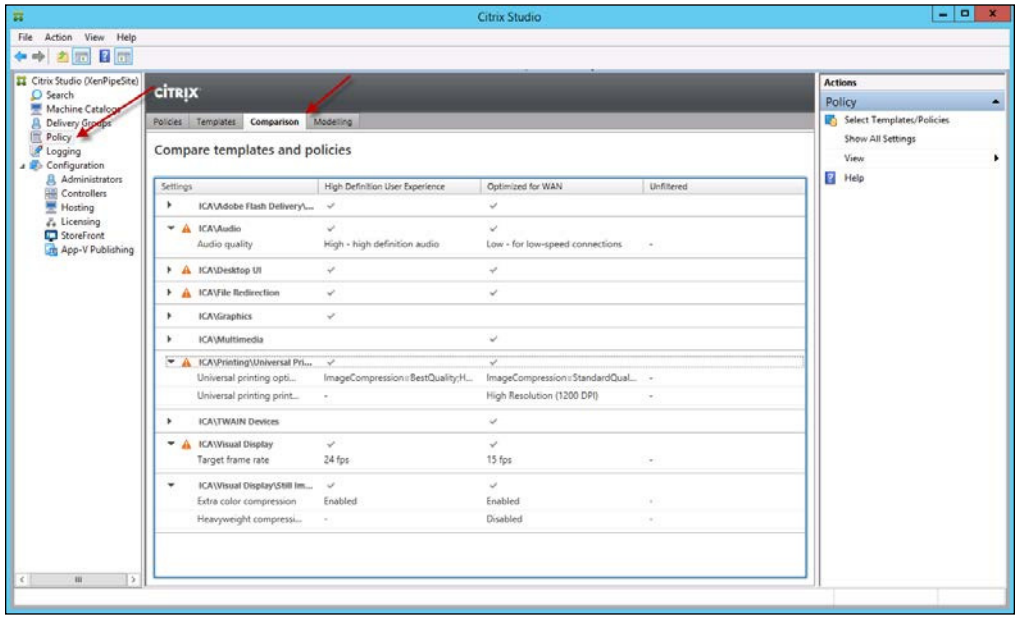

There is a tool that creates a Resultant Set of Policies, and it can be found at `http://support.citrix.com/article/CTX138533`.

Implementing policies with NetScaler Gateway™

The Citrix NetScaler Gateway is quite commonly used as a frontend to XenDesktop and provides the necessary security controls. Thus, applying policies for the NetScaler Gateway is a logical point of enforcement for XenDesktop sessions. You can create policies that apply to the NetScaler Gateway connections.

To do this, you would create Citrix policies that can be applied to different access scenarios. You can base them on elements such as authentication strength, logon location, and user device, including endpoint analysis for the software updates and antivirus. This is also an excellent point to apply client-side drive mapping controls, cutting and pasting controls, and local printing controls.

Implementing NetScaler Gateway™ policy filters

For XenDesktop to apply policies with filters on the NetScaler Gateway connections, you must complete the following prerequisites:

Component	Prerequisite
NetScaler Gateway	• Create at least one connection policy filter to define specific requirements for user logon. • You should be using NetScaler Gateway Version 10.1 to create filters that work with XenDesktop 7.x.
StoreFront	• Specify the point of authentication as **Users must log on at NetScaler Gateway** in the Citrix Receiver. • Make sure that the XenDesktop controllers for the Site are configured to trust requests sent to the Citrix XML service.

Component	Prerequisite
XenDesktop	• Make sure that any access policy configured on the XenDesktop controllers for a Site allows connections to virtual desktops through the NetScaler Gateway. • Create a user policy that includes a NetScaler Gateway filter.

To apply policies for NetScaler Gateway sessions using Citrix Studio, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the policy and click on **Next**.
5. On the **Access Control** tab, select **Assign**.
6. Enter the **NetScaler Gateway farm name** and **Access Condition** fields, as shown in the following screenshot:

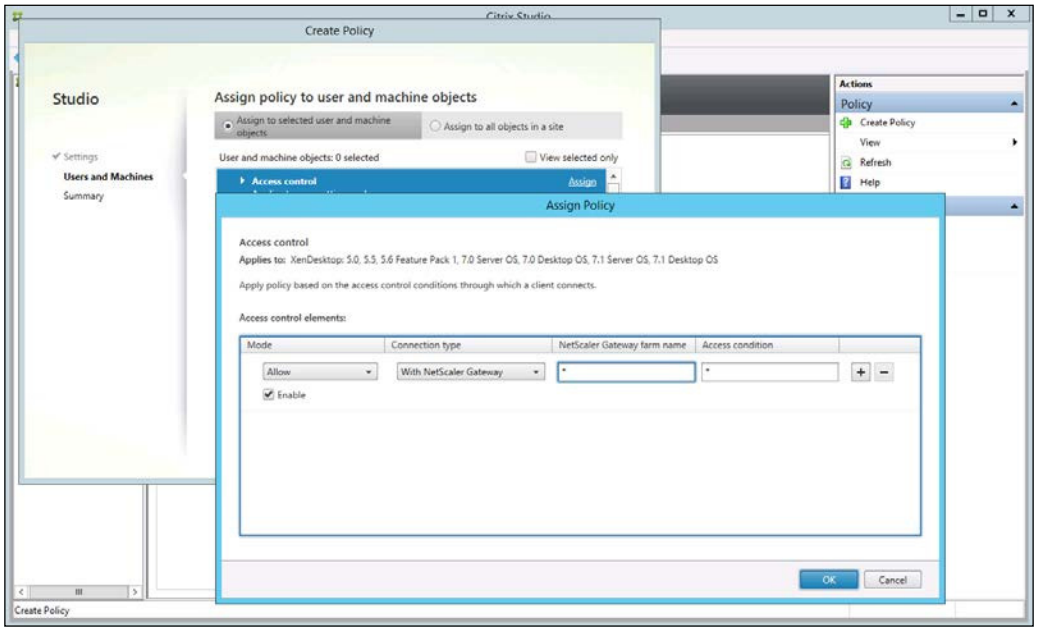

To apply policies for NetScaler Gateway sessions using Microsoft Group Policy Editor, perform the following steps:

1. Launch Microsoft Group Policy Editor.
2. Click on **New** to launch the **New Policy** wizard.
3. Select the policy or policies and click on **Next**.
4. Under **Filters**, click on **Add**.
5. Fill in the **AG farm name** and **Access condition** fields, as shown in the following screenshot:

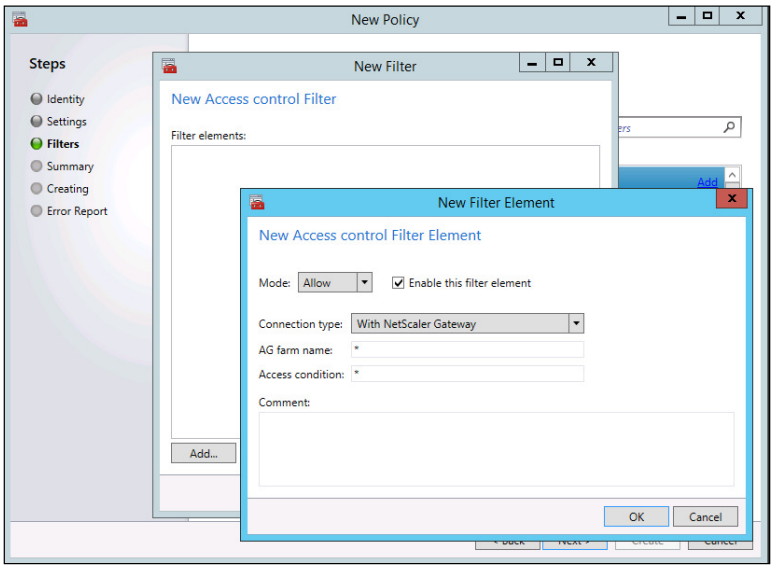

XenDesktop does not verify the Site, server, and filter names, so you need to validate this information in the NetScaler Gateway.

Summary

Citrix policies are the best way to control user sessions in XenDesktop. In this chapter, you learned the similarities and differences between the policy editor in Citrix Studio versus Microsoft Group Policy Editor. Now you know how policies work and how to apply them to specific connection types using filters. Towards the end, we included a quick troubleshooting section because policies can be tricky and you may need to figure out what the Resultant Set of Policies is for your Site. In the next chapter, we will discuss how printing works in XenDesktop and how it applies to networked and locally attached printers.

6
Managing Printing

Printing in XenDesktop is the same as printing in XenApp. I am sure you already know how printing works in Windows networking, so now you just need to understand how it works in XenDesktop. It's fairly simple; there are two ways to print something from XenDesktop: a locally attached printer on the user's USB port or a network attached printer.

In this chapter, we will cover the following topics:

- How printing works in a virtual desktop when they are connected to USB ports and networks
- The Citrix Universal Print Driver
- Printers that get autocreated
- The mapping of printers
- Optimization

How printing works

In a XenDesktop environment, printing is initiated by the user (the client side) on the machines that host the applications (the XenDesktop server side). Print jobs are redirected through the network print server or user device to the printer. Every time a user starts a desktop session, XenDesktop builds a user's workspace and provisions the printers that the user can use; this is known as printer autocreation. XenDesktop workspaces are not persistent, so they are rebuilt every time a user starts a session.

Printing works as follows:

1. A user starts a desktop session, logs in, or reconnects.
2. XenDesktop uses policies to build the list of available printers, which is known as provisioning.

3. Printers are dynamically adjusted based on policy, user location, and network changes. This is a key benefit of the Citrix printing methodology.
4. The user prints using a locally attached printer or a network printer.

Citrix also recommends using the Universal Print Driver as it eliminates administrative overhead and print driver issues.

Using locally attached printers

When you are using a locally attached printer (the client printing pathway), XenDesktop routes the print job to the locally attached printer from the server, through the client, and to the printer over the **Independent Computing Architecture** (**ICA**) protocol using the printer virtual channel. The ICA protocol optimizes and compresses the print job traffic over the printer virtual channel as shown in the following diagram:

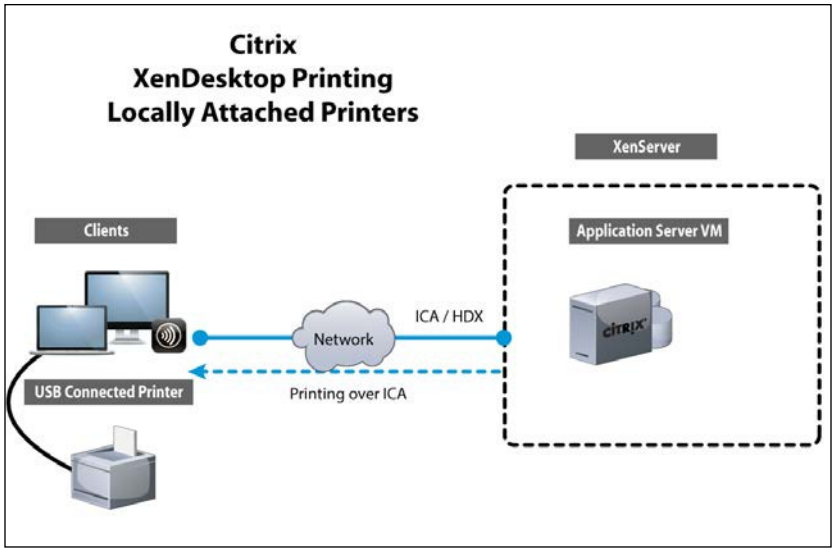

Using network attached printers

When you are using a network attached printer (the network printing pathway), XenDesktop routes the print job from the server through the network to the print server as shown in the following diagram.

If the virtual desktop or application cannot contact the print server or the native printer driver is not available on the server, the print job is routed over the ICA connection instead.

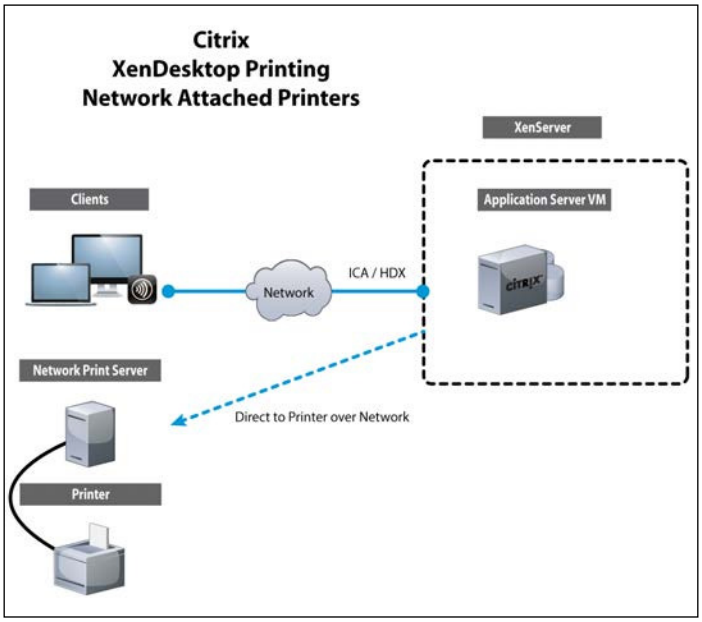

Using default printing, preferences, and drivers

If you do not configure any printing policies, the following is the process flow of how printing will work in XenDesktop:

1. The user starts their session and logs on or reconnects, and printers are autocreated. This is the equivalent of configuring the auto create client printers policy with the **Auto-create all client printers** option. We will discuss more about printer autocreation later in the chapter.
2. XenDesktop routes print jobs to locally attached printers over the ICA virtual channel (the ICA protocol) to the user device.
3. XenDesktop routes print jobs to network attached printers directly over the network. If the printer is not contactable, XenDesktop routes the print job over the ICA virtual channel instead, which is equivalent to having the **Direct connection to print server** policy configured.

4. XenDesktop stores printing preferences on the user device. If the user device doesn't support it, XenDesktop stores the printing preference in the user profiles on the server. This is equivalent to having the **Printer properties retention** policy with the **Held in profile only if not saved on client** option.
5. XenDesktop uses the Windows print driver if it is available on the server or attempts to install the driver. If the Windows print driver is not available, it uses the Citrix Universal Print Driver. This is equivalent to having the **Automatic installation of in-box printer drivers** policy with the **Universal printing** option and the **Use universal printing only if requested driver is unavailable** setting.

Citrix recommends using the Universal Print Driver to avoid installing a large number of printer drivers on the server.

Setting printing preferences

In XenDesktop, users can modify their printing preferences, and these can be stored in one of the following ways:

- **User device**: Preferences can be stored on the user's virtual desktop and can be changed by navigating to **Control Panel | View devices and printers**. Right-click on the printer to select **Printing preferences**, as shown in the following screenshot:

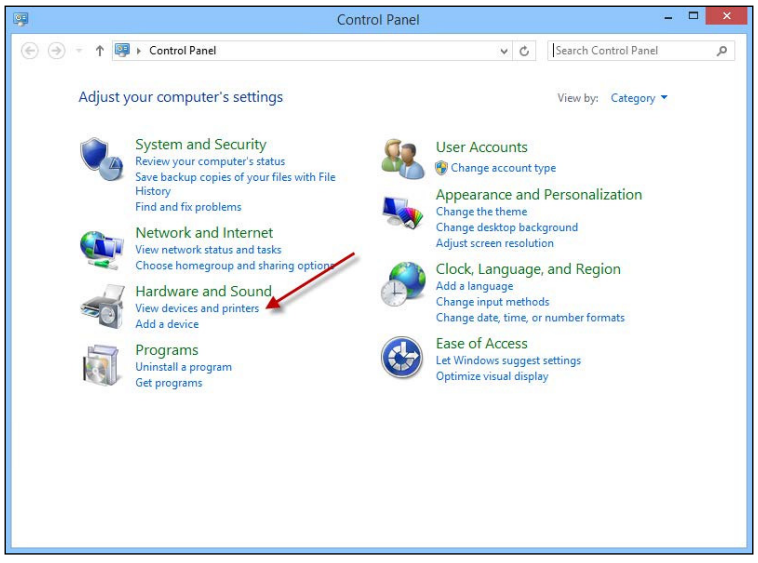

- **Document**: Preferences can be stored inside documents that are running as virtual applications from the Citrix Receiver. For example, launch Microsoft Word and select **Print**. You can access the printing preferences as follows:

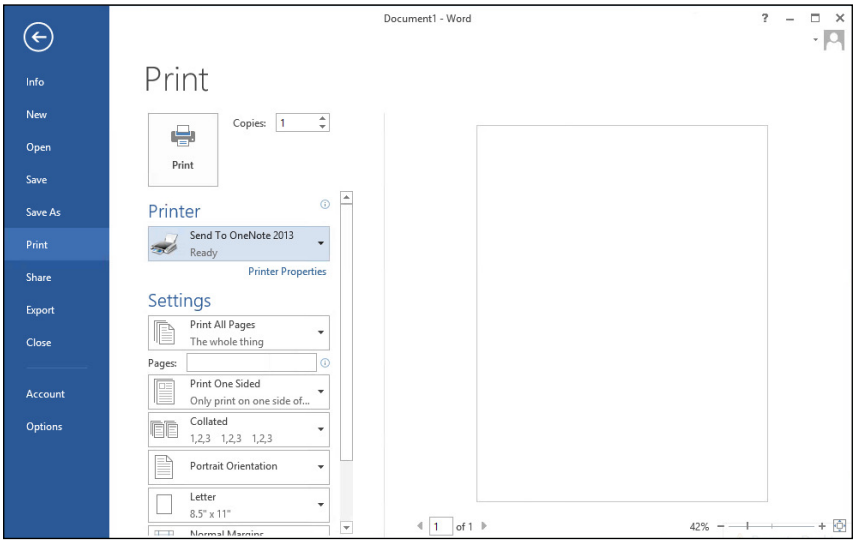

- **Session**: Preferences can be stored for an autocreated printer if the change was made in the **Printing preferences** section of the **Control Panel** through the desktop session. In other words, for changes made on the user device through the **Control Panel**, printer preferences will reflect on the autocreated printer.
- **Server**: These are default settings associated with a printer driver stored on the server. You can access these by navigating to **Control Panel** | **View devices and printers**. Select the printer, right-click on it, and select preferences.

Printing policies

You can use printing policies to control how users access printers in XenDesktop. You can control how printers are provisioned, how print jobs are routed, as well as how and which print drivers are used. You can have different printing policies for different types of users.

For a complete list of all the XenDesktop printing policy settings, please refer to the *Printing* section of *Appendix B, XenDesktop® Policy Settings Reference*.

Universal Print Server and Driver

The Citrix Universal Print Server extends the XenApp and XenDesktop Universal printing support to network printing. This feature eliminates the need to install numerous network printer drivers on XenApp and XenDesktop hosts and enables more efficient network utilization.

Many printer manufacturers have certified their printers and drivers to work with XenApp and XenDesktop; however, most of the printing can be done with the built-in Universal Print Server and Universal Print Driver. The Universal Print Server and Driver support network printing.

The Universal Print Server optimizes and compresses the print data, improving network performance and providing a better user experience. There are two components of the Universal Print Server—UPClient and UPServer—both of which can be found on the server. The Universal Print Server provides image and font caching, advanced compression, optimization, and QoS support over and above the basic Windows print server features. The Universal Print Server works with other driver manufacturers too, in case you have settings you need in the driver that are specific to the printer manufacturer. In this case, you can have the benefits of the Universal Printer Server combined with the specialized printer functionality.

It is a management decision whether you should use the Universal Print Server or vendor-specific drivers. It is far easier to support and troubleshoot printer issues if you only have a small number of drivers to support.

The Universal Print Server and Universal Print Driver are installed automatically with XenDesktop, and you can do all of your printing to different types of printers using the Universal Print Driver. The Universal Print Driver is a driver that is independent of devices and supports any print device. This simplifies the administration by reducing the number of drivers required for your Site. The Universal Print Driver also supports advanced printer functionality, such as stapling, sorting, and color depth.

You enable the Universal Print Server by enabling the **Universal Print Server enable** policy (disabled by default).

To enable the Universal Print Server, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Search for universal or print in **Settings**.

5. Select the **Universal Print Server enable** policy.
6. Enable the policy and click on **Next**.
7. Assign the policy to delivery group, group type, OU, or tagged machines.
8. Give a name to the policy and enable it.
9. Click on **OK**, as shown in the following screenshot:

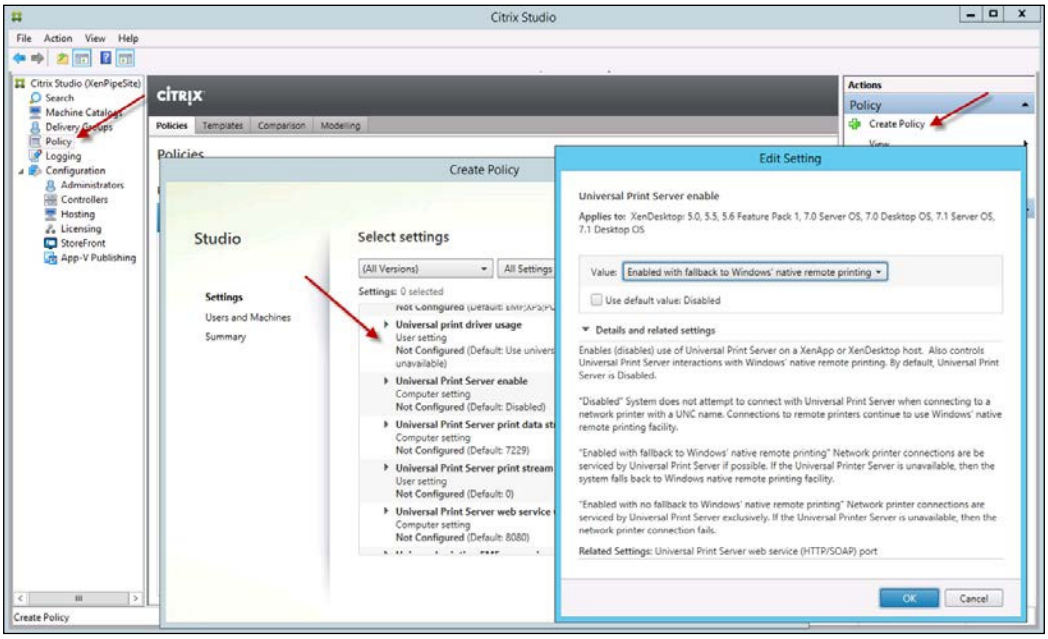

You specify the printer driver through the **Universal print driver usage** policy.

To enable the policy, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Search for universal or print in **Settings**.
5. Select the **Universal print driver usage** policy.
6. Enable the policy value and click on **OK**. Then, click on **Next**.
7. Assign the policy to the users and machines.
8. Give a name to the policy and enable it.

9. Click on **OK**, as shown in the following screenshot:

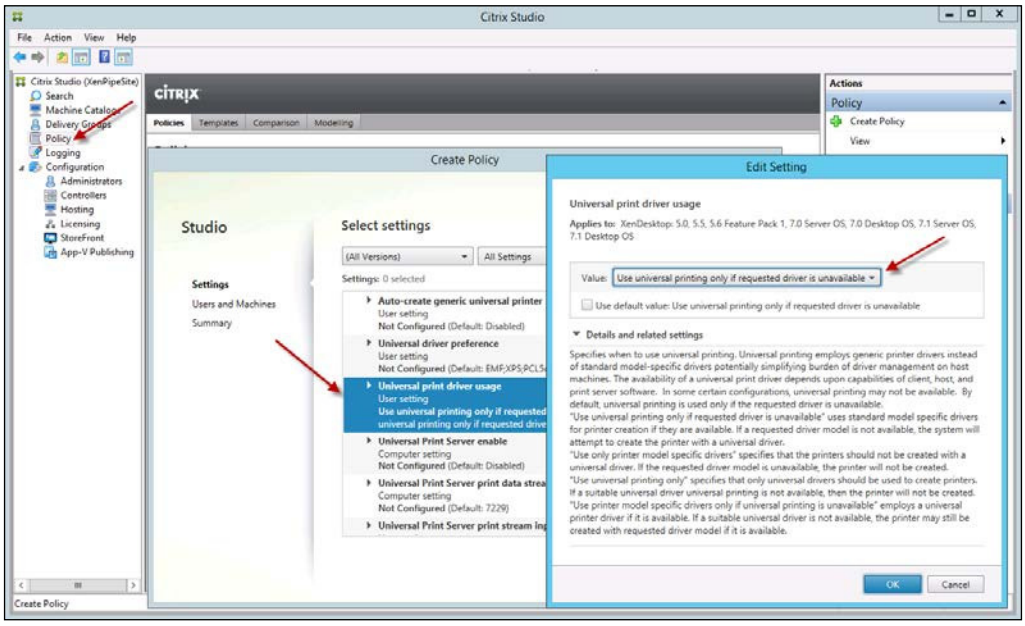

The other policies that can be enabled for use with the Universal Print Server are as follows:

- **Client printer redirection**
- **Auto-create client printers**
- **Session printers**
- **Direct connections to print server**
- **UPD preference**
- **Universal Print Server print data stream (CGP) port (default 7229)**
- **Universal Print Server web service (HTTP/SOAP) port (default 8080)**
- **Universal Print Server print stream input bandwidth limit (Kbps)**

The descriptions of these policies can be found in *Appendix B, XenDesktop® Policy Settings Reference*.

In order to enable these policies, you can follow steps similar to those used for the previous two policies. If you want to enable specific policies, you can locate these by searching. To search and locate a specific policy, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Enter some keywords for the policy you are searching, as shown in the following screenshot:

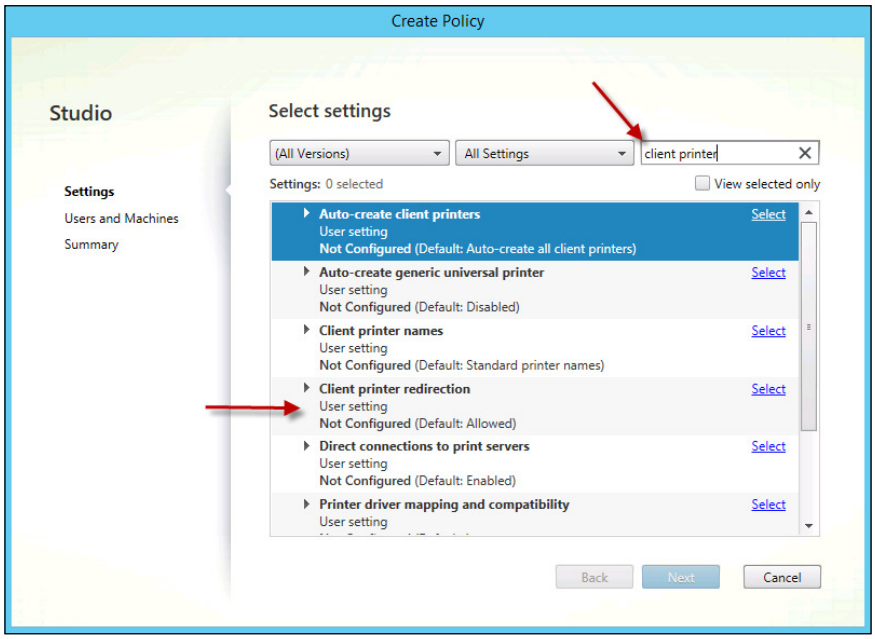

Autocreation of printers

The autocreation of printers is where XenDesktop automatically creates printers at the beginning of each user session. Both network printers and locally attached client printers can be autocreated. The autocreation of printers is dependent on the **Session printers** policy and the **Printer driver mapping and compatibility** policy.

If XenDesktop detects a new local printer connected to a user device during autocreation, it checks the server OS machine for the required printer driver. If a Windows-native driver is not available, XenDesktop uses the Universal Print Driver. This is the default setting.

Printers assigned through autocreation are based on the printers that are installed on the user device and the printers assigned through policies.

Use the **Auto-create client printers** policy to control autocreation. You can specify that all printers visible to the user device, including network and locally attached printers, are created automatically at the start of each session (default). All the local printers physically attached to the user device are created automatically. Only the default printer is created automatically if autocreation is disabled.

The **Auto-create client printers** policy needs to have the **Client printer redirection** policy set to **Allowed**.

To enable the **Auto-create client printers** policy, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Auto-create client printers** policy.
5. Select **Auto-create all client printers** and click on **OK**. Then, click on **Next**.
6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Then, click on **OK**, as shown in the following screenshot:

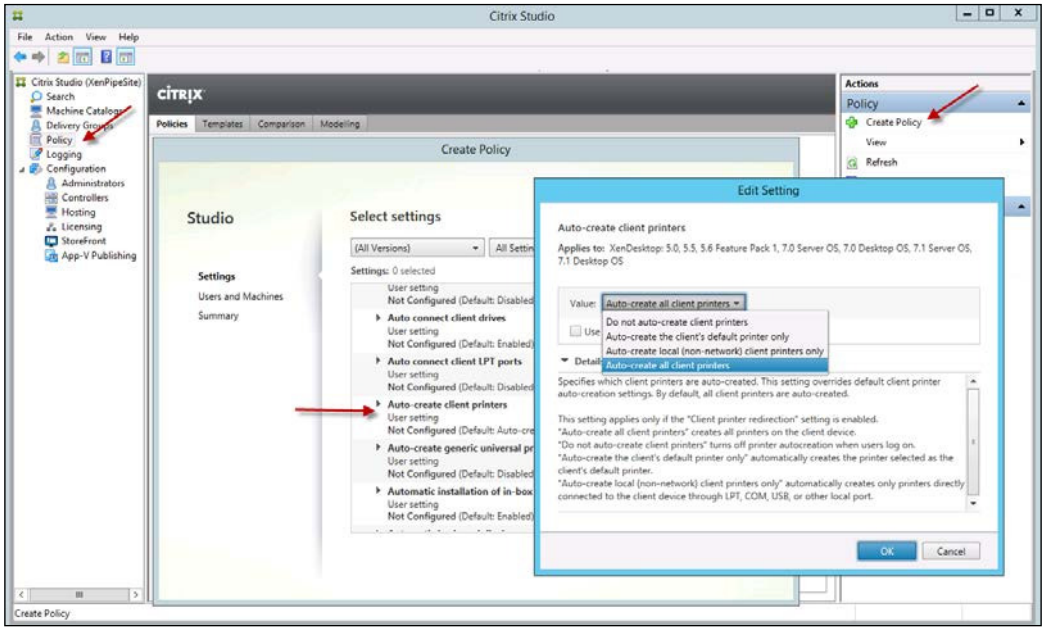

To enable the **Client printer redirection** policy, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Client printer redirection** policy.
5. Set the policy to **Allowed** and click on **OK**. Then, click on **Next**.
6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **OK**, as shown in the following screenshot:

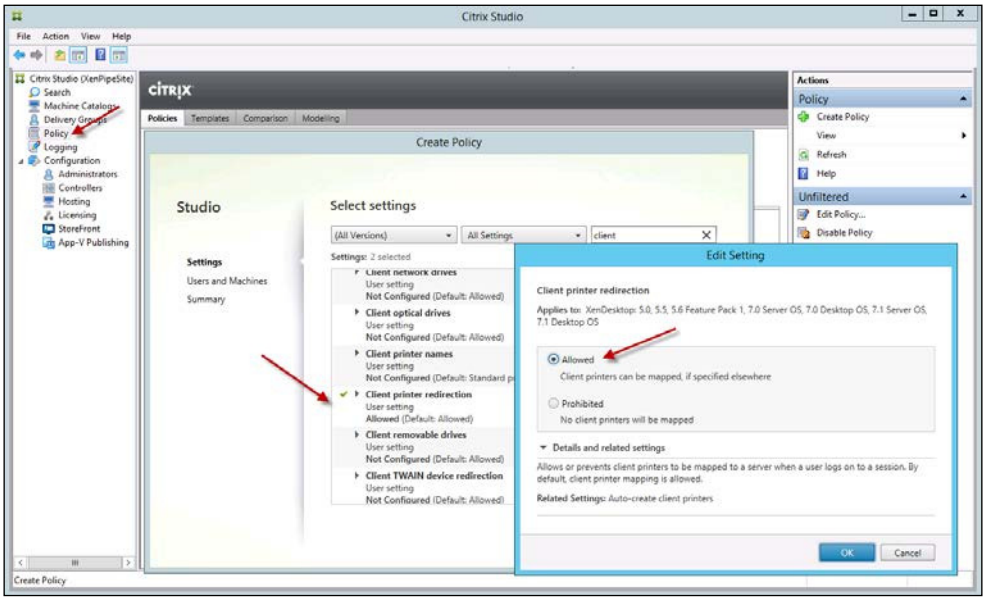

Mapping printers and drivers

You can map network printers and drivers for your users who are using policies. You can configure specific printers in the **Session printers** policy setting using the `\\servername\printername` format.

You can configure default printers in the **Default printer** policy setting; however, Citrix recommends that you do not modify the users' default printer. Instead, use the terminal services or the Windows user profile setting stored on the server for the printer.

Citrix also recommends using the Universal Print Driver as it eliminates administrative overhead and print driver issues.

If autocreation fails for some reason, XenDesktop will use a Windows native print driver, and if that is not available for some reason, XenDesktop will use the Universal Print Driver.

You can map and control the printer drivers you use with the **Automatic installation of in-box printer drivers** policy and the **Printer driver mapping and compatibility** policy.

To enable the **Session printers** policy, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Session printers** policy.
5. Enter the UNC path to the printer and click on **OK**. Then, click on **Next**.
6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **OK**, as shown in the following screenshot:

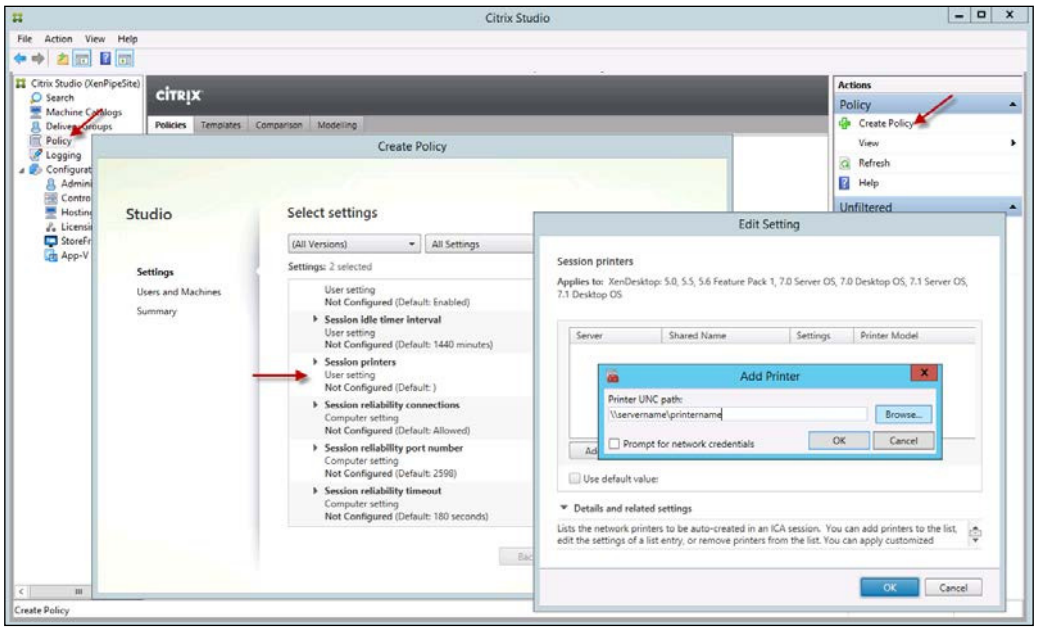

To modify the users' **Default printer** policy, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Default printer** policy.
5. Enter the UNC path to the printer and click on **OK**. Then, click on **Next**.
6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **OK**, as shown in the following screenshot:

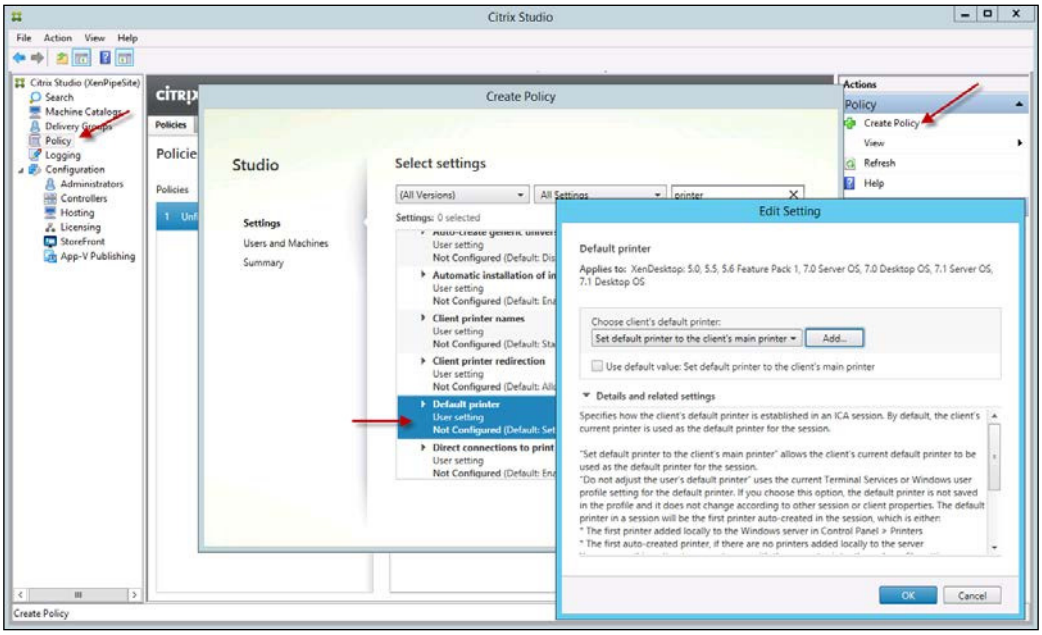

You can also use a policy to determine how the users' default printer value will be stored.

To modify the **Printer properties retention** policy, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.

4. Select the **Printer properties retention** policy.
5. Select **Value** and click on **OK**. Then, click on **Next**, as shown in the following screenshot:

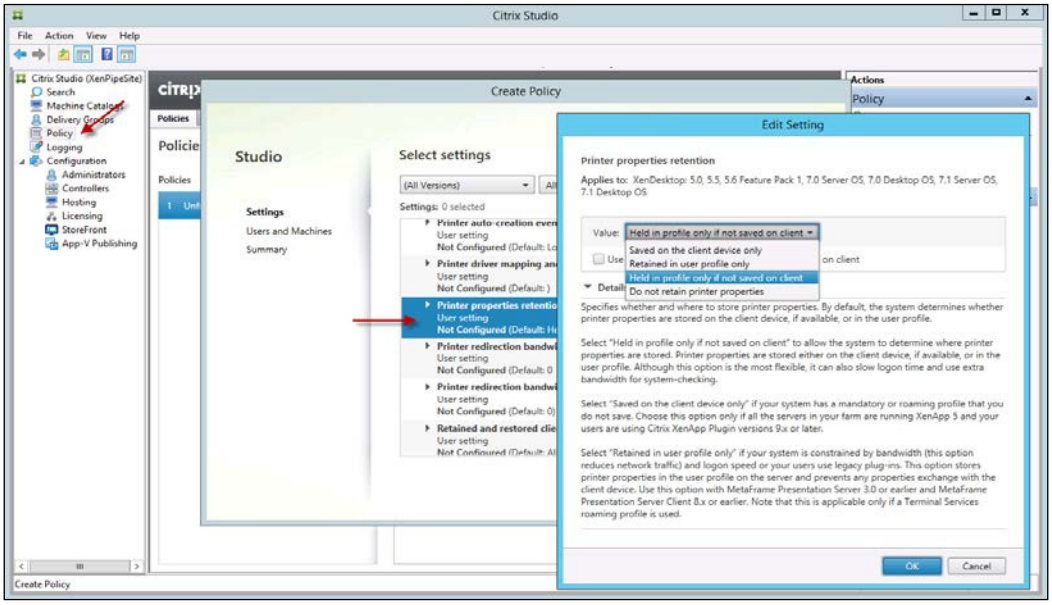

6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **Finish**.

Optimization of printing

The performance of VDI depends a lot on the performance of the network. This applies to printing as well because printing data has to start at the server, go across the network, and come out at the printer. Printing might be competing with other XenDesktop virtual channels for network resources, such as video processing, keystrokes, and mouse data, which can all be complicated further by using a WAN.

To optimize printing performance, you would want to use the **Universal printing optimization defaults** policy with the **Desired image quality**, **Enable heavyweight compression**, **Image and Font caching**, and **Allow non-administrators to modify these settings** options enabled. You should also use the **Universal printing image compression limit** and **Universal printing print quality limit** policies.

To optimize printing in the **Universal printing optimization** policy, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Universal printing optimization defaults** policy.
5. Select the desired image quality and caching options and click on **OK**. Then, click on **Next**, as shown in the following screenshot:

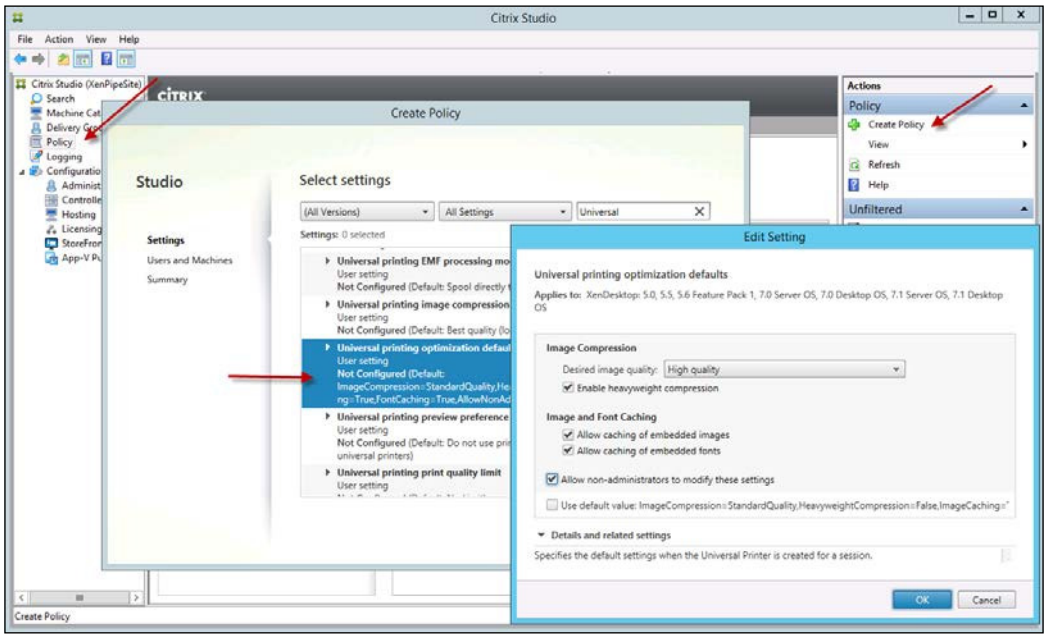

6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **OK**.

If you are facing latency or limited bandwidth situations, such as over a WAN, you may experience decreased performance. To improve performance, you might consider sending all the printing data over the ICA protocol instead of directly to the network printer, and you can do this by disabling the **Direct connections to print server** policy. Data sent over ICA is compressed for better performance.

To send all the printing data over ICA, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Direct connections to print servers** policy.
5. Disable the policy and click on **OK**. Then, click on **Next**.
6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **OK**, as shown in the following screenshot:

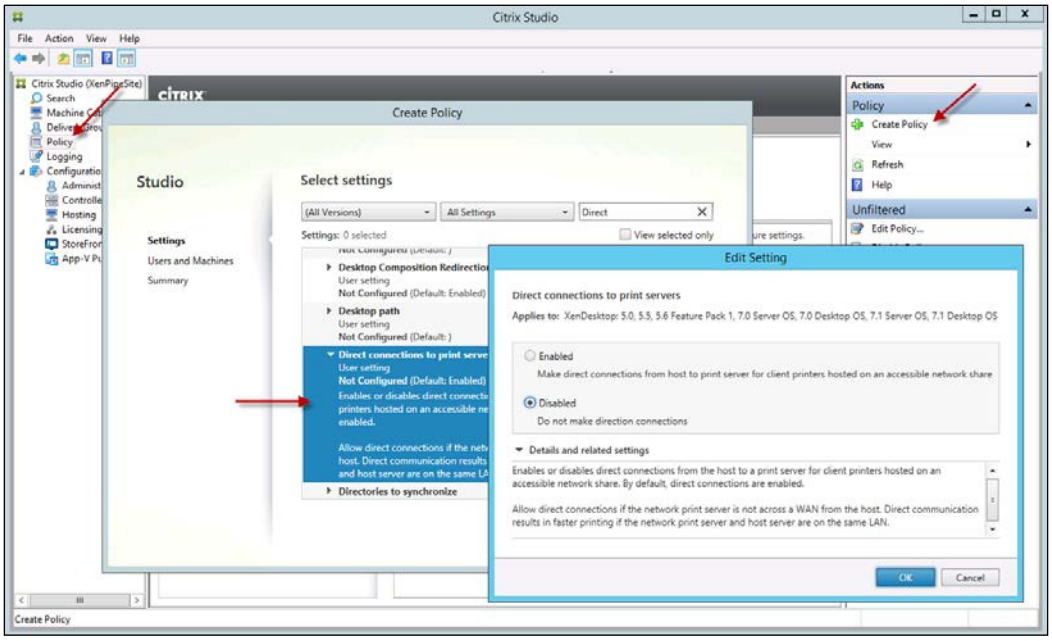

You can also limit the bandwidth used by printing so that it doesn't interfere with other data streams, such as video, keystroke, and mouse data. Additionally, you can prioritize different ICA channels if you are using CloudBridge for WAN optimization. To set bandwidth limits for printing, you can use the **Printer redirection bandwidth limit** and **Printer redirection bandwidth limit percent** policies along with the **Overall session bandwidth limit** policy.

To set a printing bandwidth limit in Kbps, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Printer redirection bandwidth limit** policy.
5. Enter a value in the **Value (Kbps)** field and click on **OK**. Then, click on **Next**, as shown in the following screenshot:

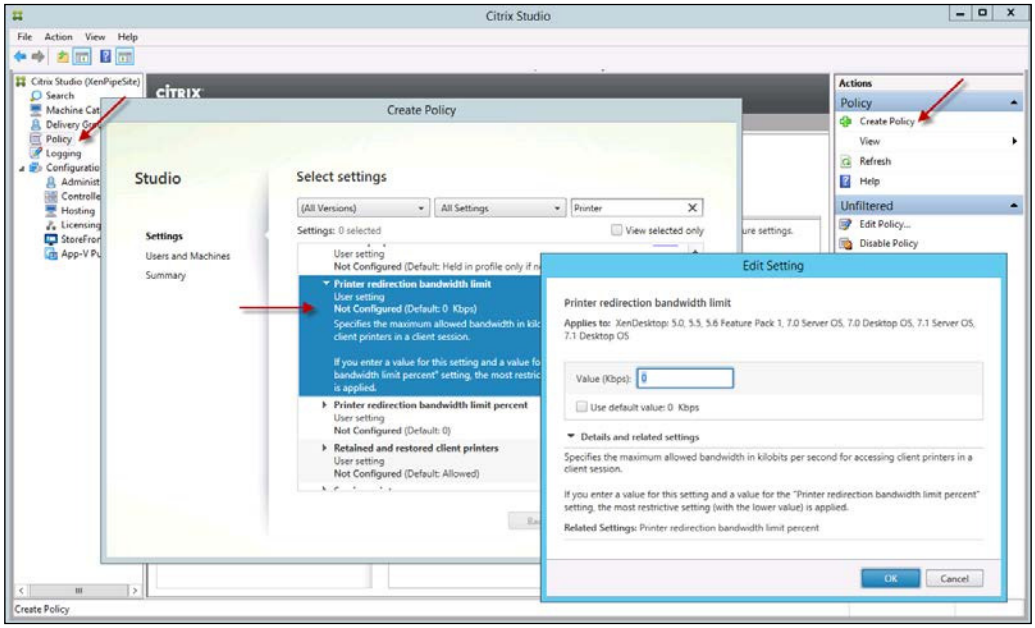

6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **OK**.

To set a printing bandwidth limit percent, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.

4. Select the **Printer redirection bandwidth limit percent** policy.
5. Enter a value for the percent limit and click on **OK**. Then, click on **Next,** as shown in the following screenshot:

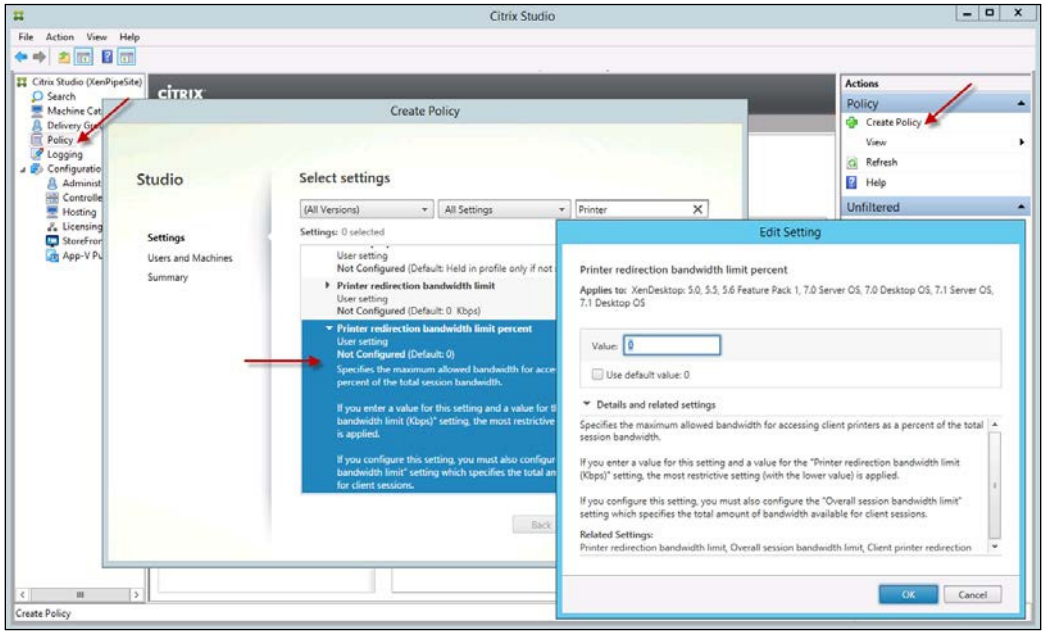

6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **OK**.

To set an overall printing bandwidth limit in Kbps, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Overall session bandwidth limit** policy.

5. Enter the value for Kbps and click on **OK**. Then, click on **Next**, as shown in the following screenshot:

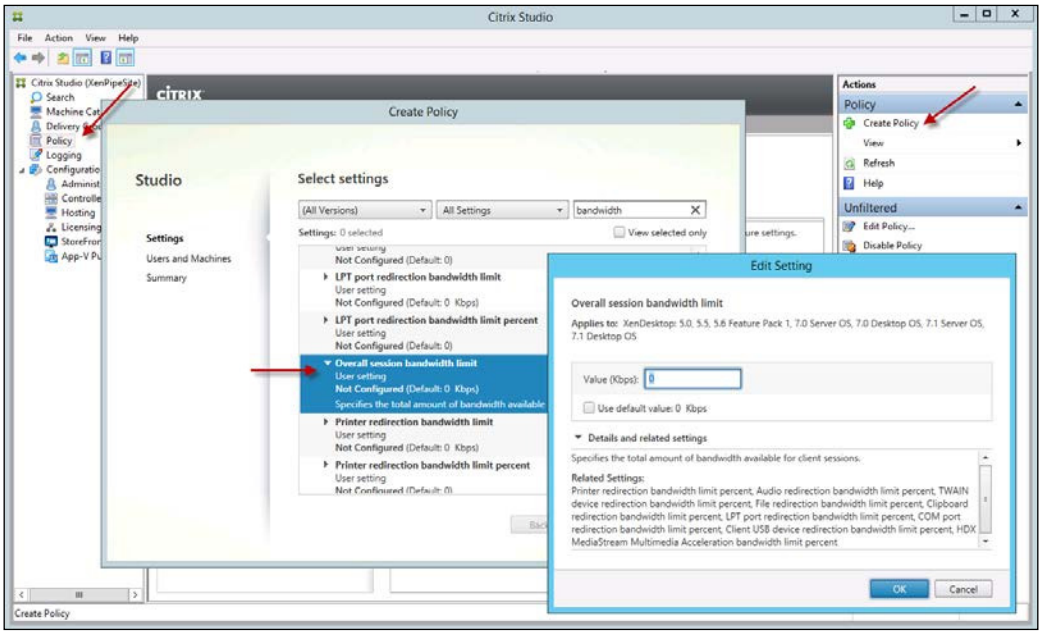

6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **OK**.

Summary

There are two ways to print something from XenDesktop: using a locally attached printer on the user's USB port or a network attached printer. Additional options about how printing is used and optimized are controlled through policy settings. In this chapter, we covered how printing works in a virtual environment, which includes locally attached and network attached printers. We showed you how to set and save printing preferences both from the server and the end user device. We covered printer autocreation, how to map printers, and the optimization of printing. Now that you understand how printing works, we will discuss how a USB works in XenDesktop in the next chapter.

7
Virtualizing USB Support

Think about it; if you use a virtual desktop, you won't have a physical USB port to plug in to that virtual machine. We need to use the USB port on our client device and somehow map the USB port on the client device to the virtual desktop. USB support allows virtual desktops to access the local USB resources connected to the user/client device. XenDesktop also provides direct connectivity support for some devices such as keyboards, mice, and smart cards.

In this chapter, we will cover the following topics:

- How USB redirection works in XenDesktop
- How to enable USB support in XenDesktop
- USB mass storage
- USB voice and video
- USB autoredirection

USB devices in virtualization

We all know how USB devices work when they are directly connected to our computer. You plug it in, Windows automatically detects the device and installs a driver, and then you get a popup that says the device is ready to use. Easy, right? Virtualization breaks this process because you can still plug the USB device into your client device, laptop, or workstation, but the operating system is absent. The operating system has been removed, and it now runs on a server in the data center or cloud. To further complicate matters, your desktop is now delivered from the server over a network to your client device. Your client device or thin client must now support USB on a virtual channel.

Now that the USB connection model is broken, you need a way to fix it. How do you connect the USB device at the client device to the operating system running on the server in the cloud? The answer is by using the Citrix USB virtual channel.

The Citrix USB virtual channel creates a virtual connection between the desktop running on the server and the end user's client device. When you plug a USB device in to the client device, the connection information is routed through the USB virtual channel to the VDI desktop, and it provides the same user experience as that of having a USB device connected to a local operating system.

All is well and good, isn't it? Actually, you would be surprised at how many client devices don't work with the USB virtual channel.

For a complete list of USB devices that have been certified to work with Citrix, visit the Citrix Ready Site at `http://citrixready.com`. For a complete list of USB devices that are supported and tested with XenDesktop, refer to the CTX article at `http://support.citrix.com/article/ctx123569`.

The following illustration shows a logical representation of the USB virtual channel that XenDesktop uses. The channel is used for printing as well, but you can see how a USB virtual channel is connected to the user's virtual desktop in XenDesktop all the way out to the end user's device. The end user device contains a USB port and USB-connected device, such as a printer or storage device.

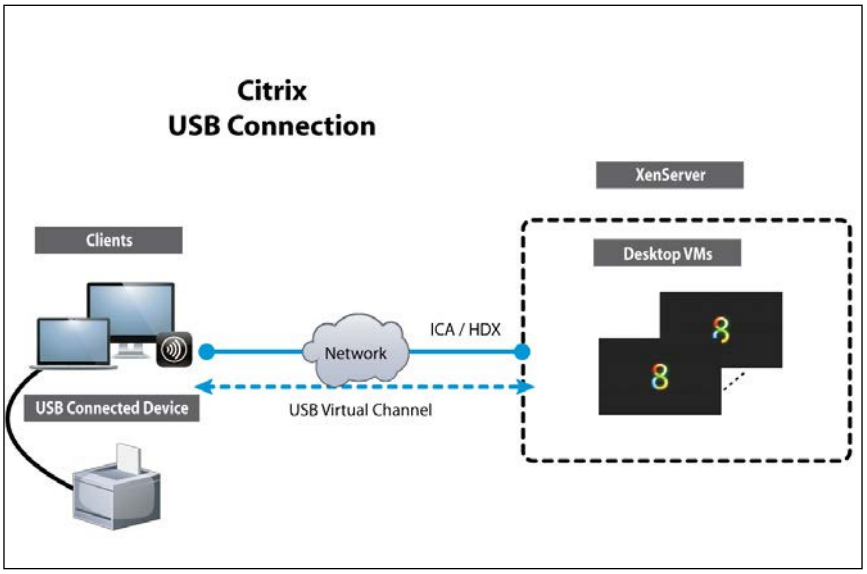

How XenDesktop® uses USB redirection

XenDesktop allows users to interact with a wide range of USB devices during their XenDesktop session. To the end user, the experience of using a USB device with a virtual desktop is the same as using a USB device with a local computer, with some caveats. Because XenDesktop uses a USB virtual channel to connect the device through the network, there are some limitations.

The level of support depends on the client installed on the user device. At the time of this writing, USB redirection is supported on the Windows and Mac Receiver clients.

Citrix refers to certain types of USB devices as **Isochronous**. Isochronous USB devices such as webcams, microphones, speakers, and headsets are supported as long as the network latency is low and has the throughput of a high-speed LAN. If these requirements are not met, the performance can't be guaranteed.

Isochronous devices require timing coordination to work correctly, such as voice or video transmission. In other words, the data flow to the desktop needs to be the same or close to the rate of the data flow out of the USB device. This is different from **asynchronous** devices, which allow processes to run independently until they are interrupted. This is also different from **synchronous** devices which require processes to wait for the completion of an event in another process before continuing.

XenDesktop uses optimized virtual channels to redirect some USB devices while providing optimal performance and bandwidth efficiency. Some devices that are supported directly in a XenDesktop session are as follows:

- Keyboards
- Mice
- Smart cards
- Scanners
- Audio devices
- Webcams

By default, USB redirection is allowed for certain classes of USB devices and denied for others. Refer to the Receiver documentation for a list.

If a device does not have an optimized virtual channel, it can still be supported by the generic USB virtual channel that provides raw USB redirection.

Generic USB redirection is supported for the following devices:

- Bloomberg keyboards
- Smartphones
- 3D mice
- **Human Interface Devices** (**HIDs**) such as USB phones, fingerprint readers, and software dongles

Some devices are not supported for USB redirection because they aren't useful as USB devices, such as network interface cards, Bluetooth dongles, USB hubs, USB graphics adapters, and USB storage devices for ReadyBoost.

If you are using thin clients, consult the manufacturer for the details of USB support and configuration.

USB devices connected to a USB hub are supported; however, the USB hub itself cannot be redirected over the USB virtual channel.

Enabling USB support

Enabling USB support is easy, but it needs to be done at both ends, that is, on the client and the server sides. First, you enable USB redirection through policies in the **USB Devices Policy Settings** section of the ICA policy settings.

To enable USB support on XenDesktop, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Client USB device redirection** policy.
5. Set it to **Allowed** and click on **OK**. Then, click on **Next**.
6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **Finish**, as shown in the following screenshot:

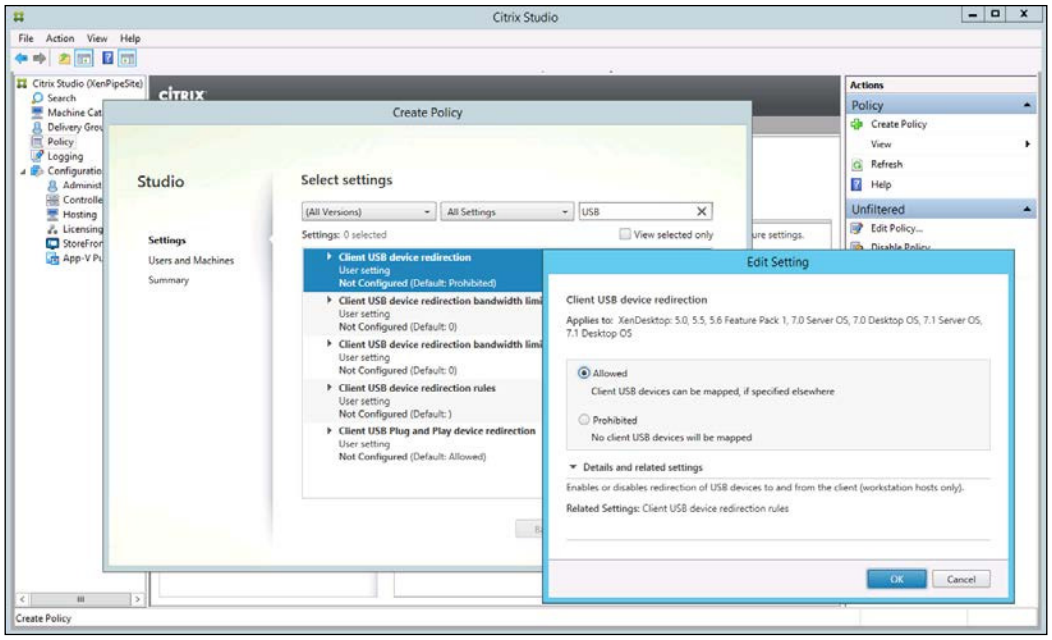

Second, you enable USB support on the end user's device in the Citrix Receiver settings or automatically enable it when you install the Receiver client on user devices.

To enable USB support in Citrix Receiver on the user device, perform the following steps:

1. Insert the XenDesktop installation DVD.
2. Navigate to the **Citrix Receiver and Plug-ins** directory.
3. Navigate to the `Receiver` folder.
4. Launch the `CitrixReceiver.exe` file and click on **Install**.
5. Add an account, connect to StoreFront, and click on **Finish**.
6. Launch Receiver if it isn't running already.
7. In the system tray, right-click on the Receiver icon and click on **About**.
8. Click on **Connection Center**.

9. Set USB access to **Ask Permission** or **Full Access** and click on **Close**, as shown in the following screenshot:

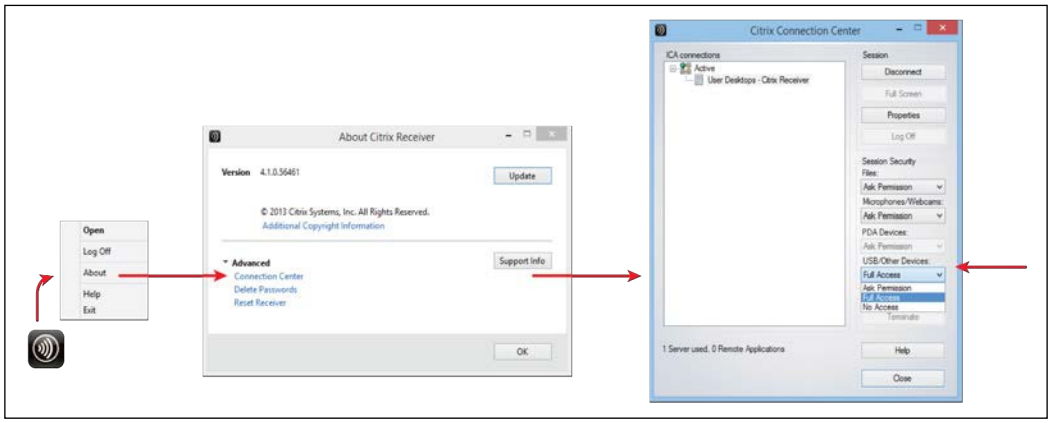

Optionally, you can specify the USB redirection rules in the **Client USB device redirection rules** policy.

The following table illustrates some of the rules and settings that apply to USB policy settings:

USB policy settings		
Name	**Description**	**Default**
Client USB device redirection	This controls whether the USB devices are available in a session or not.	**Prohibited**
Client USB device redirection rules	This specifies redirection rules for USB devices. Refer to the CTX article at `http://support.citrix.com/article/ctx119722`	None specified
Client USB Plug and Play device redirection	This permits the Plug and Play of PTP, MTP, and POS devices in a user session.	**Allowed**
Client USB device redirection bandwidth limit	This specifies the maximum allowed bandwidth, in kilobits per second, for the redirection of USB devices to and from the client.	**0** to no maximum is set
Client USB device redirection bandwidth limit percent	This specifies the maximum allowed bandwidth, in percent, for the redirection of USB devices to and from the client.	**0** to no maximum is set

Preventing the mapping of USB devices

By default, USB devices are automatically redirected when USB support is enabled, and the USB preferences settings are set to automatically connect USB devices. There might be cases where you may not want to automatically redirect all USB devices or limit the USB devices that can be used. Desktop users can explicitly redirect devices that are not automatically redirected by selecting them from the USB device list. It is also possible to prevent USB devices from ever being listed on or redirected to either the client endpoint or the **Virtual Desktop Agent** (**VDA**).

To limit or disable the automatic redirection of a USB device, refer to the CTX articles at `http://support.citrix.com/article/CTX123015` and `http://support.citrix.com/article/CTX132716`.

Using USB mass storage

I think everyone has used a USB drive at some point. You can configure USB-connected drives to be accessible to the virtual desktop through what are called client drive mappings. Client drive mappings can be set up to automatically connect USB drives to the virtual desktop. To configure client drive mapping, use the **Client removable drives** setting in the **File Redirection Policy Settings** section of the ICA policy settings.

The difference between having a USB drive show up as a connected drive or as a USB device is determined by policy, as shown in the following table:

Remote drive policy		
Feature	**Client drive mapping**	**USB mapping**
Enabled by default	Yes	No
Read-only access configurable	Yes	No
Safe to remove device	No	Yes, based on OS recommendations for safe removal

If the client drive mapping and the USB rule are enabled and a mass storage device is inserted before a session starts, it will be redirected using the client drive mapping. After this, the mass storage device is considered for redirection through USB mapping. If it is inserted after a session has started, the reverse happens. It will be considered for USB mapping before client drive mapping. Client drives are connected automatically unless disabled.

To enable or disable auto connecting of the client drives, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Auto connect client drives** policy.
5. Set it to **Enabled** or **Disabled**. Click on **OK** and then click on **Next**.
6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **Finish**, as shown in the following screenshot:

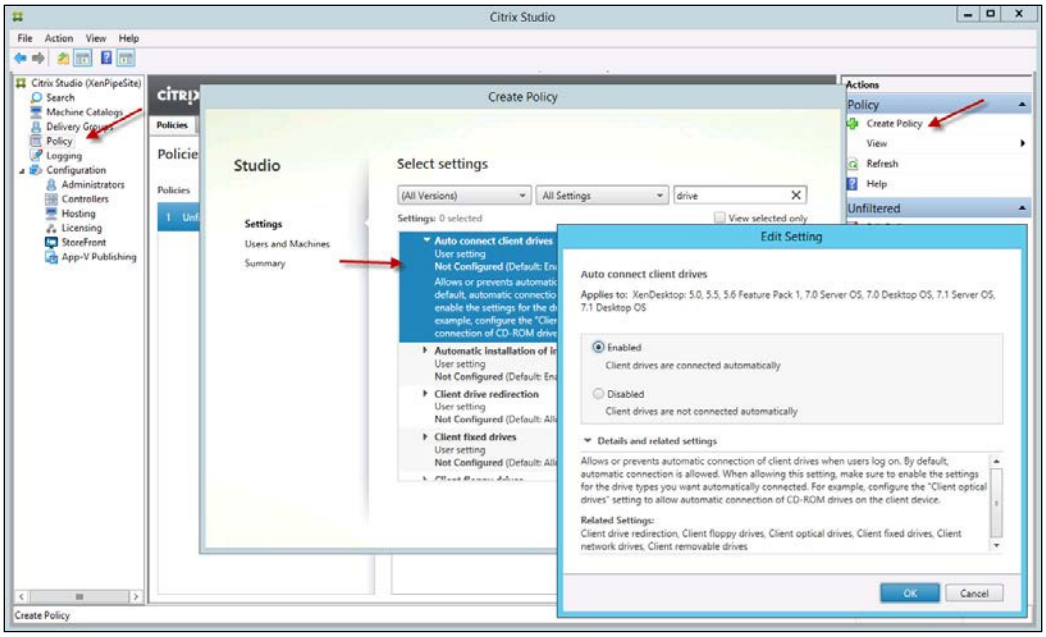

Be sure to set the **Client USB device redirection** policy to **Allowed**.

USB redirection with XenApp® versus XenDesktop®

XenApp is essentially the new name for Citrix Presentation Server, which was formerly Citrix MetaFrame. What this means is that XenApp (CPS or CM) has been around for a lot longer than XenDesktop. Therefore, USB redirection was handled differently between XenApp and XenDesktop. Moving forward, in Versions 7.x, USB redirection will be handled in the same way between XenApp and XenDesktop.

To further complicate matters, the support within each product is different depending on the version of the software. XenApp 4.5 and 5.0 support device-level USB redirection. XenApp 6.0 supports Plug and Play USB redirection. Starting with XenDesktop 4.0, isochronous USB devices are supported and these use a special service known as the Citrix USB service.

To read more about how USB is used differently between XenApp and XenDesktop, refer to the CTX article at `http://support.citrix.com/article/CTX124956`.

Using USB automatic redirection

When USB support is enabled and the USB user preferences are set to automatically connect, USB devices are automatically redirected. These specific USB devices are also automatically redirected when operating in a desktop appliance mode or with **Virtual Machine** (**VM**) hosted applications.

To configure the automatic redirection of USB devices, refer to the CTX article at `http://support.citrix.com/article/CTX123015`.

Using voice and video

It's clumsy, it's complicated, but it's used a lot. I'm talking about Microsoft Lync, formerly called Microsoft Communicator. Lync is used as a collaboration tool with instant messaging, video conferencing, and online meetings. It even supports telephony and integrates with Skype. The marketing spin says that Microsoft Lync is an "enterprise-ready unified communications platform. Lync connects people everywhere as a part of their everyday productivity experience." With so many organizations using Lync, you can imagine that Lync management servers will be one of the resources that gets virtualized.

There is a great deployment guide on frontending Lync with NetScaler, and now that NetScaler is a VM, you should be able to virtualize both NetScaler and Lync. The guide is available at `http://blogs.citrix.com/2011/03/30/load-balancing-lync/`.

The following CTX article is useful for using Microsoft Lync with XenDesktop: `http://support.citrix.com/article/CTX124124`.

If you are going to use Microsoft Lync 2010 in a XenDesktop environment, then you will want to use the optimization pack. The Citrix HDX RealTime Optimization Pack for Microsoft Lync 2010 offers clear, crisp, high-definition video calls in conjunction with Microsoft Lync 2010. Currently, there is no optimization pack for Lync 2013. The following are some usability points for consideration when using Lync with XenDesktop.

The video usability points to be considered are as follows:

- **Webcam**: If you are using a webcam, Lync must execute directly on the VDA platform or the client desktop. This ensures the best audio quality. It uses the CTXMM virtual channel that utilizes less network latency and less bandwidth, which is good for WAN connections.
- **Presence**: For presence information to work, Lync and Outlook must both be installed on the VDA platform or the client desktop.
- **FPS**: The default frames per second is 24 fps, but you can lower it to 15 or 20 fps and still have good video quality.
- **Compression**: Adaptive display dynamically determines the available bandwidth and adjusts traffic through quality and frame rate.

The voice usability points to be considered are as follows:

- **USB Headset**: Use the generic USB redirection (CTXGUSB virtual channel) as both signaling and audio are involved.
- **Headset**: Use the bidirectional audio (CTXCAM virtual channel) instead of the generic USB redirection virtual channel. This minimizes bandwidth consumption. It is important to use a good quality headset with noise and echo cancellation.
- **Codec**: Configure the bidirectional audio virtual channel to use the optimized-for-speech audio codec, also known as medium quality, where the bandwidth is 56 kilobits per second or 28 kilobits in both directions.

Make sure that you use Citrix Receiver 3.0 or higher.

One more note on CPU utilization is that voice and video are data intensive, and they place a high load on the CPU of the client desktop and VDA. Assigning two virtual CPUs to each virtual machine will help reduce the thread switching latency. You can also reduce the video display frame rate using the **ICA/Visual Display** policy.

The optimization pack for Microsoft Lync is located at `http://support.citrix.com/proddocs/topic/technologies/hdx-realtime-optimization-pack-wrapper.html`.

Summary

In this chapter, we learned how USB works in a virtual desktop; or how it doesn't work natively, but how you can configure it to work the way you want. Remember that XenDesktop uses a USB virtual channel between the remote end user and the virtual desktop. Think of it as a long USB cable connected to the remote end user through the network to the XenDesktop server. USB redirection works for many devices including printers and storage devices. In this chapter, you learned how to enable and disable access to USB devices. In the next chapter, we will discuss storage.

8
Virtualizing Storage and Backup

Virtualization changes the storage paradigm. In the old days, end users would have a hard disk inside their laptop or workstation, where the operating system was installed and where they also stored all of their data. If you use a virtual desktop, you won't have a local storage any more. You might have local storage; however, the **Information Technology** (**IT**) department will require user data to be stored close to the virtual machines somewhere in the data center or cloud. If you refer back to the network diagram presented in *Chapter 1*, *Designing a XenDesktop® Site*, and also to the diagram that follows, you will see two types of storage devices: one for a SQL Server and the other for VDI storage or virtual machine storage. We will discuss the differences between a SQL Server and VDI storage in this chapter. Specifically, we will cover the following topics:

- Storage considerations
- Storage requirements
- Backing up

- Restoring

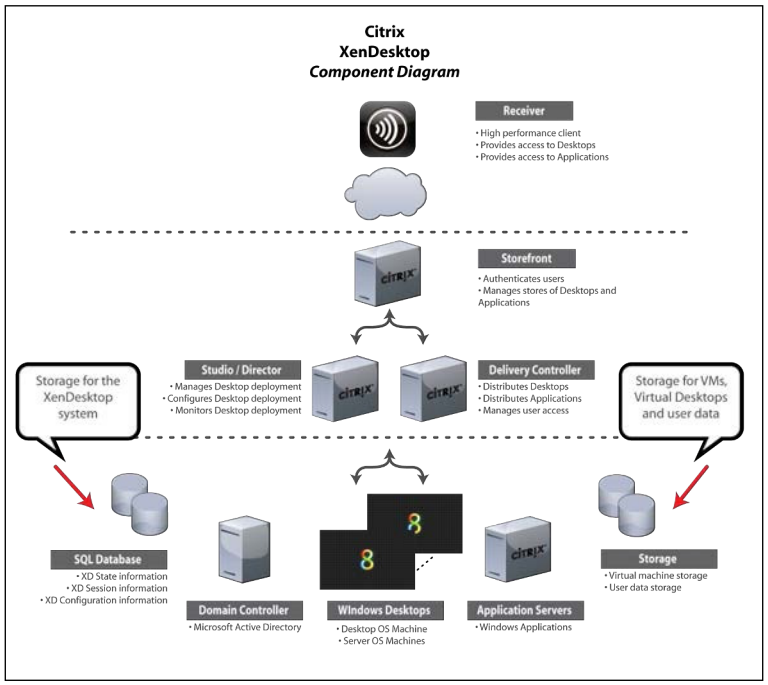

XenDesktop® storage considerations

If you think VDI is complicated, you are welcome to the subject of storage. There is a lot to consider; so, read the chapter and put some thought into it. You need storage for the virtual desktops in addition to the XenDesktop servers themselves.

Desktop storage

Storage is an important component of virtualization. You need storage, and lots of it, to store the virtual machines and user data. The only exception to this is if you are using the **Provisioning Server** (**PVS**) or XenServer with **thin provisioning**. The amount of storage available is determined by the type of architecture you implement, dedicated desktops versus shared desktops. As you can imagine, having dedicated desktops for each user will really start to eat up disk space. Shared desktops use a shared VDI image with the ability to store user data in a **Personal vDisk** (**PvD**) partition, which requires considerably less storage and keeps the costs down. Storage can be locally attached to the Hypervisor server; however, to scale and get better performance, you will likely use the **Network Attached Storage** (**NAS**) or **Storage Area Networking** (**SAN**) XenDesktop system storage.

A SQL Server is used by XenDesktop as a central location to keep track of the Site configuration. The SQL data store contains static information, such as Site policies, machine catalogs, delivery groups, and published applications and desktops. It also keeps track of dynamic information, such as who is connected to which resource, server load, and connection statuses for load balancing. In XenDesktop 7, the only option is to use a SQL Server; you cannot use an `Access` database any more. The SQL data store is not typically located on the same drive or medium as the virtual machine storage device. In our example, we installed a SQL Server on the XenDesktop Delivery Controller using local storage (as seen in the following diagram); however, you may want to consider using NAS or SAN for scalability and better performance. When you get into production, you will want to separate your SQL servers from the XenDesktop servers so that you can add high availability features to the database.

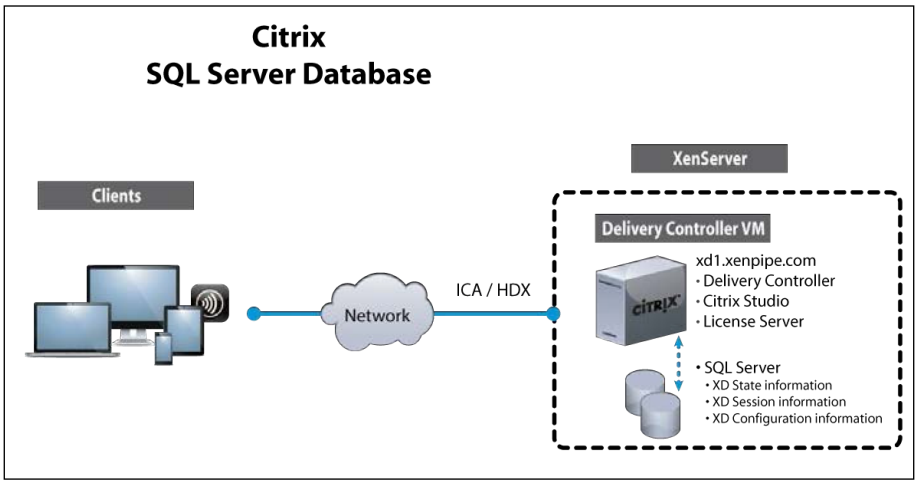

High Availability

High Availability (**HA**) refers to the ability of keeping the Site running in the event of an outage. In XenDesktop 7, if the database becomes unreachable, sessions will continue to run, but running new sessions and configuration changes will not be possible. It is recommended that you build some HA into your architecture through the use of SQL mirroring for automatic failover. It will also benefit you to do some SQL clustering for the failover of server tasks in addition to the failover of data. There are some new features in SQL Server 2012, such as AlwaysOn Availability Groups and Windows Server Failover Clustering, which are worth looking into. More information on AlwaysOn Availability can be found at `http://msdn.microsoft.com/en-us/library/ff877884.aspx`.

Backup your SQL Server and storage frequently in the event that your HA solution fails for some reason because HA is not a backup; it just keeps your Site up for users if a server fails.

Performance

Performance is mentioned here because one of the largest bottlenecks in VDI usability is storage and the performance of the I/O. You will hear people talk frequently about **Input Output operations Per Second** (**IOPS**). As IOPS has been identified as a key bottleneck in VDI performance, several companies have emerged to provide data deduplication technology to speed up the IOPS between a server and its storage. Some of these companies are Atlantis Computing, Nimble Storage, Nexenta, Fusion-io, and Tegile to name a few.

IOPS

IOPS stands for Input Output Operations Per Second processing. It is a server performance measurement, and we talked about it briefly under the *Performance* section. You will hear about it frequently when you are deploying VDI because as you can imagine, now that you are taking all the users and condensing them into a single data hardware paradigm, there will be bottlenecks. You might have heard rumors around the industry that VDI is unusable, and this can mostly be attributed to IOPS. There are some companies out there that are solving this problem and they have been mentioned. Investigating these technologies will be beneficial as they aim to solve this problem and can significantly improve the usability of VDI.

Personal vDisk

Personal vDisk is a result of the RingCube acquisition—and is a good one.
A Personal vDisk retains the image for shared (pooled) and streamed desktops, yet allows end users to install applications and change their desktop settings and preferences. Before this feature, in a shared desktop deployment, users would lose all of their customizations and personal applications when the administrator altered the base virtual machine.

All the changes made to a user's VM are stored on the Personal vDisk. The contents of a user's Personal vDisk are merged at runtime with the content from the base VM. This allows administrators to update base VMs with software updates and security patches while retaining the end user's own preferences and applications. Personal vDisks have the following two parts:

- Data containing user data, documents, and the user's profile in `c:\Users`, which is assigned to the `P:` drive letter (the drive letter can be changed).
- Data containing a virtual hard disk file (`.vhd`) that contains all the other data such as the applications installed in `c:\Program Files`. This drive is not visible to users.

Personal vDisks do not need to be stored in the same location as the Hypervisor or virtual desktop.

The following diagram provides a visual representation of how VDI uses shared images for desktops and then merges these images with each user's Personal vDisk to create a uniquely personal experience, similar to that of using a physical laptop or workstation:

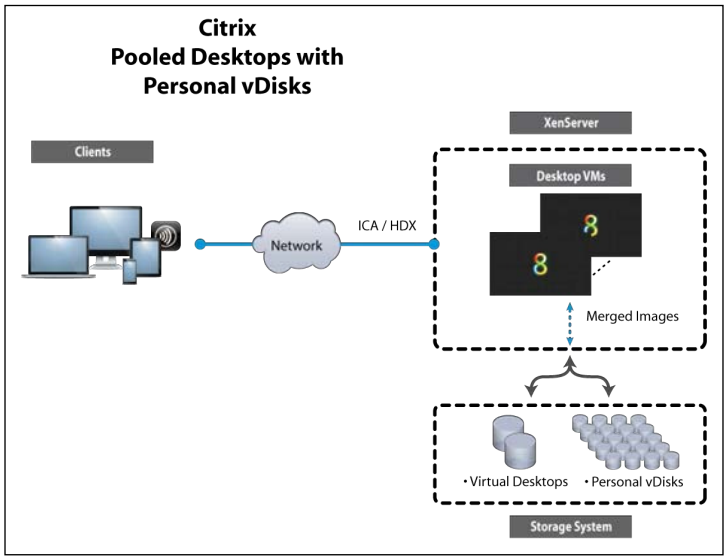

XenDesktop® storage requirements

You need a place to store the XenDesktop system database, and it uses a SQL Server. In a Proof-of-Concept or test environment such as ours, we can use the local disk of the XenDesktop Delivery Controller to install a SQL Server. A SQL Server is installed with the XenDesktop product; you just have to specify where you want it to live. Each component in XenDesktop has minimum storage requirements as explained in the following table:

The XenDesktop components' storage requirements

Component	Supported operating systems	Minimum storage
Hypervisor host	XenServer 6.2, 6.1, 6.0.2	16 GB minimum, 60 GB recommended
	vSphere 5.1 Update 1, 5.0 Update 2	
	Hyper-V 2012 SP1	
XenCenter	Windows XP, Vista, 7, 8	100 MB minimum
	Windows Server 2003, 2008, 2008 R2, 2012	
Delivery Controller	Windows Server 2012, Standard and Data center	100 MB
	Windows Server 2008 R2, Standard, Enterprise, Data center	
SQL database	SQL Server 2012 SP1, Express, Standard, Enterprise	None specified
	SQL Server 2008 R2 SP2, Express, Standard, Enterprise, DC	
Studio	Windows 7, Professional, Enterprise, Ultimate	75 MB
	Windows 8, Professional, Enterprise	
	Windows Server 2012, Standard, DC	
	Windows Server 2008 R2 SP1, Standard, Enterprise, DC	
Director	Windows Server 2012, Standard, DC	None specified
	Windows Server 2008 R2 SP1, Standard, Enterprise, DC	
StoreFront	Windows Server 2012, Standard, DC	None specified
	Windows Server 2008 R2 SP1, Standard, Enterprise	

The XenDesktop components' storage requirements		
Component	**Supported operating systems**	**Minimum storage**
License server	Windows Server 2008, 2008 R2, 2012 Windows 7 and 8, 32-bit and 64-bit	2.5 GB

Additional guest OS storage requirements can be found at `http://bit.ly/1auxjQp`.

Virtual desktop storage requirements – dedicated desktop model

The purpose of XenDesktop is to deliver virtual desktops to the users. Just as you would have given each user a laptop with a local disk to store their operating system, programs, and user data, you need to provide storage for the virtual desktops to do the same.

Not only do you need to estimate how much storage you need for each copy of the operating system, for example, Windows 8; but you also need to estimate the number of desktops you will have. It would be nice if you could just say, "I want 100 desktops; now calculate the total storage and find me a product". It is usually a little more complicated than this, but we have simplified the process to get you moving in the right direction.

The good news is that you don't have to set aside space for each user's desktop. For example, 100 users requiring 40 GB of storage would require 4 Terabytes of storage. You can create a shared or pooled desktop and then add small amounts of storage for each user's Personal vDisk. This reduces the storage requirements and saves on the equipment cost.

Storage requirements for the vDisk (virtual desktop) are typically determined by the size of the operating system requirements (Windows 7 or 8) plus any applications and other software requirements. The vDisk is typically created as a master image, and it is recommended that you create several versions of the master image for software updates and backups.

To determine the amount of storage required, you will need to use some sizing formulas. The amount of storage for vDisks can be calculated if you have an idea of how many users you need to supply with virtual desktops. Once you have this, you can find a product in the market that will supply that amount of storage.

Storage required for dedicated desktops		
Variable	**Description**	**Value**
vUsers	This is the number of virtual desktops deployed to users	X
vMem	This is the amount of RAM allocated to each desktop	2 GB
vDisk	This is the amount of disk space allocated to the master OS	25 GB
WCache	This is the size of the write cache	1.5 GB
Storage capacity	This is the total capacity of the storage device	TBD

Based on the parameters in the previous table, we can approximately calculate how much storage we will need for our users' virtual desktops if we are using a dedicated VDI model.

For 1000 users, given the assumptions for the following parameters, we will need 33.7 Terabytes of storage. The formula to calculate the storage capacity is as follows:

*Storage capacity = ((vUsers * vDisk) + ((vUsers * (vMem + WCache)) * 1.25)) * 1.15*

- The assumptions for the parameters to calculate the storage capacity are as follows:
    - 2 GB RAM for each desktop
    - 25 GB disk for the Windows operating system
    - A dedicated desktop for each user
    - 1.5 GB write cache for each user
    - Disks get 15 percent storage headroom; 25 percent headroom for VM RAM
- For 1000 users, we calculate the storage capacity required to be equal to 33.7 Terabytes

For more information on how to size your storage, refer to the sizing guide for Citrix XenDesktop at `http://bit.ly/1auzB1O`.

Virtual desktop storage requirements – dedicated shared desktop model

Citrix Personal vDisks are an important part of the XenDesktop architecture because they provide an end user personalization of the desktop while maintaining the control of a centrally-managed desktop image. The use of Personal vDisks allows the IT administrators to scale into large deployments without using extremely large storage footprints. Personal vDisks give end users the ability to customize their virtual desktop by installing their own software and changing their desktop preferences.

Personal vDisks are deployed at a machine level during catalog creation or when a machine is added to an existing catalog. If you are using Machine Creation Services, Personal vDisks are deployed within Desktop Studio. You can also select **Pooled with Personal vDisk** as an option when you are creating machines in an MCS desktop catalog.

When creating the desktop catalog, you can change the default size of the Personal vDisk. Make sure that the Personal vDisk is large enough to handle profile and personalization operations. The recommended minimum size of a PvD is 10 GB by default; a user's Personal vDisk will be split between user personalization and application data with the default split being 50/50. If a user manages to fill their entire Personal vDisk, you can increase the size of their Personal vDisk after the virtual desktop has been assigned using Desktop Director.

Storage requirements – shared desktops versus Personal vDisks

Variable	Description	Value
vUsers	This is the number of virtual desktops deployed to users	X
vMem	This is the amount of RAM allocated to each desktop	2 GB
vDisk	This is the amount of disk space allocated to the master OS	25 GB

Storage requirements – shared desktops versus Personal vDisks		
Variable	**Description**	**Value**
PvDisk	This is the amount of Personal vDisk space allocated to each user for personal storage	10 GB
numvDisks	This is the number of vDisks	6
WCache	This is the size of the write cache	1.5 GB
Storage capacity	This is the total capacity of the storage device	TBD

Based on the parameters in the previous table, we can approximately calculate how much storage we will need for our users' virtual desktops if we are using a shared VDI model.

For 1000 users, given the assumptions for the following parameters, we will need 16.7 Terabytes of storage. The formula to calculate the storage capacity is as follows:

*Storage Capacity = ((vUsers * PvDisk) + ((vUsers * (vMem + WCache)) * 1.25) + (vDisk * numvDisks)) * 1.15*

- The assumptions for the parameters to calculate the storage capacity are as follows:
    - 2 GB RAM for each desktop
    - 25 GB disk for each operating system
    - 10 GB Personal vDisk space for each user
    - Six vDisks that will be shared among all users
    - 1.5 GB write cache for each user
    - 15 percent storage headroom; 25 percent headroom for VM RAM
- For 1000 users, we get the storage capacity required to be equal to 16.7 Terabytes

As you can see, using shared (pooled) desktops with Personal vDisk decreases the requirements on storage significantly: a 50 percent reduction in storage and cost from the 33.7 Terabytes required in the dedicated model.

For more information on Personal vDisks and sizing, visit `http://bit.ly/15nUHyn`.

A storage capacity calculation spreadsheet can be found at `http://bit.ly/1eiV7uF`.

Virtual desktop storage requirements – shared hosted desktop model

There is also a shared desktop model that is used very often in the field and is rather simple to implement. It is possible to create a single desktop that is not dedicated to any specific user yet everyone can use it. You build one master image and several users use this without the Personal vDisk feature. It is simple to manage this using Citrix Profile Management in order to manage user settings. This requires very little infrastructure and is easy to deploy and manage. However, it does rely on Microsoft Terminal Services licensing.

Backup and restore

The data in a XenDesktop environment resides in two different types of storages. If you refer back to our network diagram at the beginning of the chapter, data is stored for the XenDesktop infrastructure in the SQL database, while VMs and user data are housed on a separate storage. As mentioned earlier, it is a good practice to build HA and Disaster Recovery into the architecture. As an additional safeguard, you should get into the practice of backing up data, to avoid data loss in the event of a disaster.

Backing up a SQL Server

Consult the Microsoft SQL Server documentation for more information on how you should schedule regular backups. The following steps will help you to get a one-time manual backup done.

Don't allow any administrative changes while performing the backup.

To perform a one-time backup of a SQL Server, perform the following steps:

1. Launch PowerShell or access PowerShell from the Studio console, and run the `Get-BrokerDBConnection` command to get the name of the database server / data source and the database/initial catalog name.

   ```
   PS C:\> Get-BrokerDBConnection

   Server=sqlserver.xenpipe.com;Database=CitrixXenDesktopDB;
   Trusted_Connection=True
   ```

2. Log in to a SQL Server on the Delivery Controller and launch Microsoft SQL Server Management Studio.
3. In order to run the SQL Server Management Studio, we have to download SQL Server 2012 service pack 1 Express with tools and run the installation to add features to the existing installation, as shown in the following screenshot:

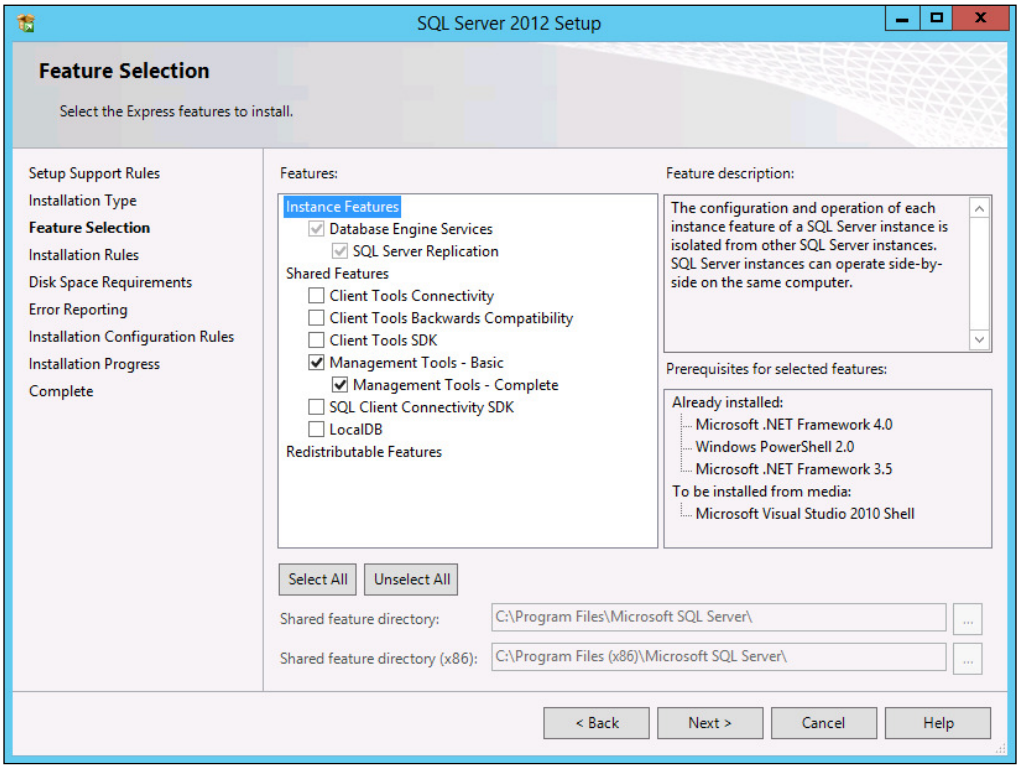

4. Connect to the database engine.
5. Expand the **Databases** node and navigate to the database identified in step 1.
6. Right-click on the database and select **Tasks**. Then, click on **Back Up…**.

7. Set **Backup type** to **Full** and choose to verify the backup, as shown in the following screenshot:

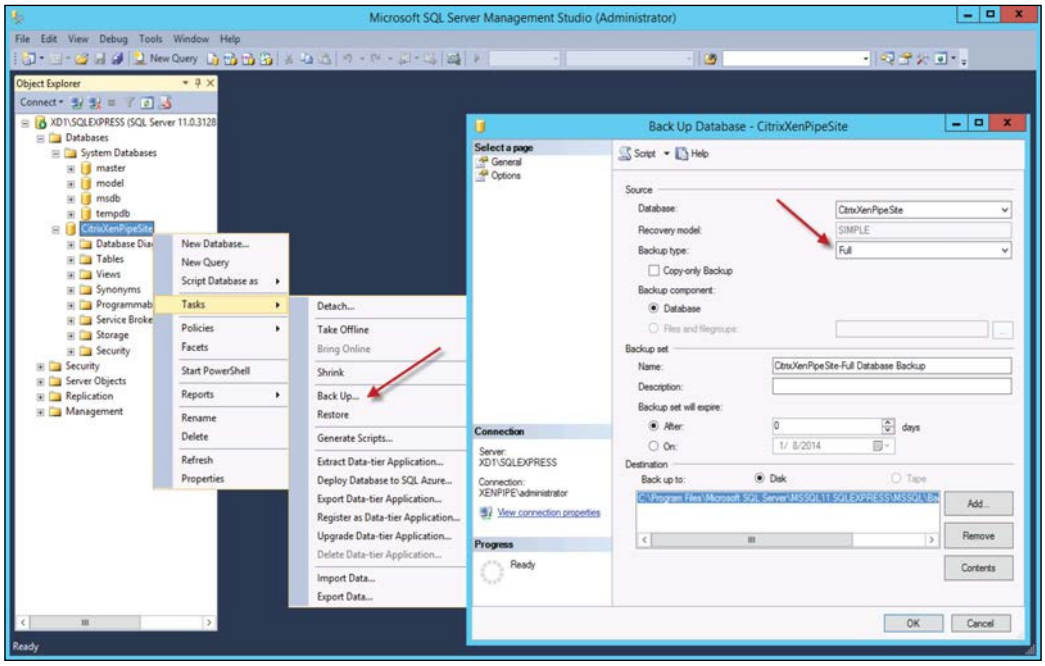

Restoring a SQL Server

Consult the Microsoft SQL Server documentation for more information on how you should restore SQL. The following steps will help you to restore the one-time backup created previously.

To perform a restore of the SQL database, perform the following steps:

1. Log in to a SQL Server on the Delivery Controller.
2. Launch Microsoft SQL Server Management Studio and then connect to the database engine.
3. Expand the **Databases** node and navigate to the database identified in make sure this is aligned correctly previous section.
4. Right-click on the database and select **Tasks**, then click on **Restore**, and finally **Database…**.

5. Select the backup set from which you want to restore. Several options are available, as shown in the following screenshot:

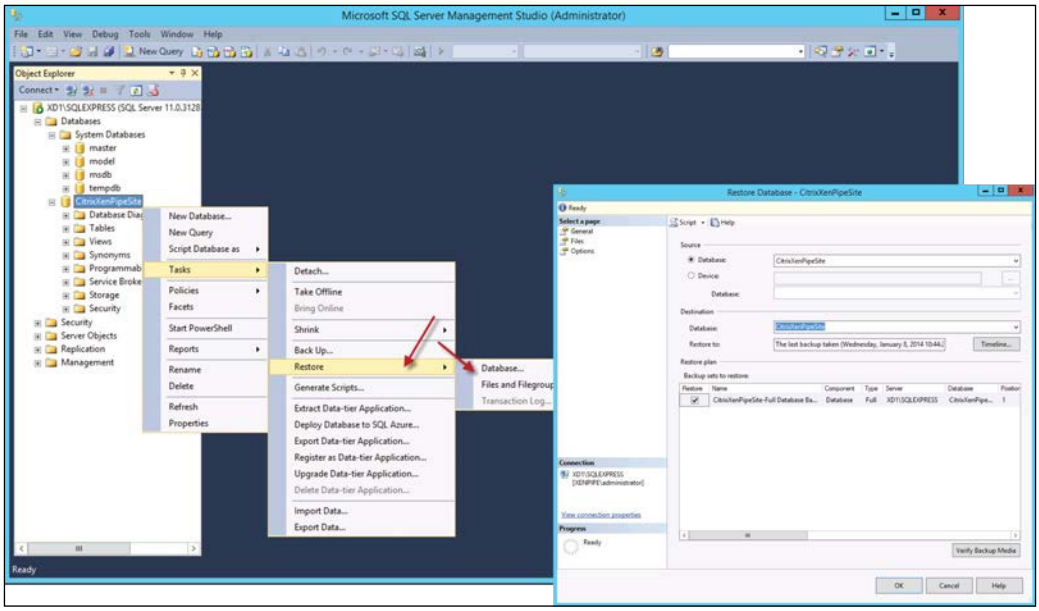

6. After a successful restore operation, you will need to restart each broker. Log in to each broker machine (Delivery Controller), launch PowerShell, and execute the following command:

```
PS C:\> Add-PSSnapin Citrix.* Get-AcctServiceStatus Get-
BrokerServiceStatus Get-ConfigServiceStatus Get-HypServiceStatus
Get-LicServiceStatus Get-ProvServiceStatus Get-PvsVmServiceStatus
```

7. All the preceding commands should return an **OK** status. If not, refer to the CTX article at `http://support.citrix.com/article/CTX128075` to troubleshoot.

Backing up and restoring VMs and user data

The next most important things to backup are the VMs and user data, basically everything on the storage device that may or may not include the SQL Server you just backed up in the previous section. The Hypervisor, for example, has features to backup and restore VMs and can be found in the relative administrator's reference. The only downside to using the XenCenter console is that you would be backing up each server or VM manually, and this can be time consuming.

What is probably a better practice is to make use of third-party backup and restore software. The third-party tools can provide more robust feature sets of agentless backup and restore of many VMs, dedupe technology, scheduling, and even file-level restores.

Citrix has invested in PHD Virtual (`www.phdvirtual.com`), so you can be sure that their backup/restore software is tightly integrated with Citrix Xen products, such as XenServer, XenApp, and XenDesktop.

USB mass storage

In the previous chapter, we talked about how to use an end user attached USB storage device through remoting over the virtual channel. This could be useful for an end user to access a USB storage device attached to their client device through a virtual desktop served from the data center.

You might get the idea that you can use a USB mass storage device as a primary storage for the XenDesktop Site, VMs, or user data. This is not a good idea because if you remove the drive, you could lose all of your data. So, don't try it. Don't entrust your data to a USB drive.

Summary

The storage paradigm is changing with large storage appliances becoming denser and faster along with the fact that storage is being merged more and more with the compute and networking resources into what is called converged infrastructure. In this chapter, you learned about the important components of storage with regard to XenDesktop. In the next chapter, we will learn about HDX and the high definition experience.

9
High Definition Experience (HDX™)

High definition experience (**HDX**) is a set of capabilities that delivers a high definition user experience. Around the time Citrix started to dig in its heels on protecting the intellectual property around the `ICA` protocol, it also figured out that there are other valuable features that can be used along with the `ICA` protocol. Whether you want to name these features HDX features or VDI features is a whole different topic of discussion. Either way, you are likely to come across the HDX feature set at some point during your implementation of VDI. HDX technology provides an extension to the `ICA` protocol by providing a set of unique, high definition experience features for the end user on a variety of devices.

In this chapter, we will cover the following topics:

- HDX
- Windows Aero
- Windows Media and Flash
- HDX 3D
- How to configure HDX

Introducing high definition experience

Before we dive in, I would like you to get a simple understanding of what HDX is. For the client device, HDX leverages the computing capability of user devices to enhance and optimize the user experience. HDX MediaStream technology ensures that users receive a smooth, seamless experience with multimedia content on their virtual desktop.

On the network, HDX uses advanced optimization and acceleration capabilities to deliver the best performance over any network, including low-bandwidth and high-latency connections.

In the data center, HDX leverages the processing power and scalability of servers to deliver advanced graphics performance, regardless of the capabilities of the client device.

In the network operations center, HDX Insight captures data about the ICA traffic and provides a dashboard view of real-time and historical data. HDX Insight requires additional software and is covered in *Chapter 10*, *Application Delivery*.

HDX supports the following two types of graphics technologies:

- **Direct X**: This is a collection of **application programming interfaces** (**APIs**) developed by Microsoft to handle multimedia tasks, such as games and videos on Microsoft platforms.
- **OpenGL**: This is an Open Graphics Library and is a multiplatform API for rendering 2D and 3D computer graphics. OpenGL was originally developed by **Silicon Graphics Inc.** (**SGI**) in 1992 and is used widely for CAD and other graphics applications.

HDX™ system requirements

To have a high definition experience, several requirements must be met on the client and server in order to have enough horsepower to deliver that experience. The HDX system requirements are as follows:

- **User device**: HDX features are historically only supported on Windows-based client devices; however, there is a new Linux client that supports HDX. With the new server-side codec, any client device would be able to use HDX as long as it is processed in the server. Either way, if you have a Windows device or thin client, it needs to support the following requirements:
    - DirectX 9

- Pixel Shader 2.0 in hardware
- 32 bits per pixel
- 1.5 GHz 32-bit or 64-bit processor
- 1 GB RAM
- 128 MB video memory

This is how HDX MediaStream works: HDX queries the Windows device to verify that it has the required GPU capabilities, processor speed, and RAM, and automatically reverts to server-side desktop composition if it doesn't. Devices that don't will be listed in the GPO group for excluded devices.

- **Server**: Most servers will meet the requirements if they support virtualization. The following are the basic requirements:
  - A PCI display card
  - Virtual desktop based on Windows 7 or Windows 8 or a Windows Server desktop OS for Aero features
  - Citrix XenServer, VMware ESX Hypervisor, and Hyper-V
  - Network card throughput of 1.5 Mbps to 5 Mbps
- **Windows Media**: The war on who gets to show you the video on the user device is a messy one. Support for the software that displays the video is required for each vendor. Naturally, HDX supports Windows Media Player under the HDX features of Windows Media client-side content fetching, Windows Media redirection, and real-time Windows multimedia transcoding. Windows Media for HDX is supported on the following operating systems:
  - Receiver for Windows
  - Receiver for iOS
  - Receiver for Linux
- **Flash Media**: Flash redirection is not supported on Windows 8 or Windows Server 2012 and is disabled by default for those operating systems. Flash redirection is supported on earlier Windows 7 and Windows 2008 operating systems.

The reality of HDX™

HDX was a great marketing concept. It evolved from some very real-world use cases around some capabilities that customers were asking for. I think it was also an opportunity for Citrix to create some intellectual property around delivering desktops in a virtualized environment, now that processing power and the end user device have been decoupled.

You might hear about the reference to the `HDX` protocol. There is no such thing, well, almost. The `HDX` protocol is just another marketing term, rebrand if you wish, of the `ICA` protocol. It is the `ICA` protocol that Citrix uses to deliver desktops and applications, sometimes referred to as the `ICA/HDX` protocol. It is also an attempt to move away from the competitors who have engineered solutions for the `ICA` protocol. If a vendor is a valid Citrix Ready partner and has been given the legal rights to develop for the `ICA/HDX` protocol, then you can be confident that their solution adheres to the protocol's specifications. Competitors who claim that they work with `ICA/HDX`, in fact don't work, because they haven't licensed the `ICA/HDX` protocol from Citrix and have somehow reverse engineered their solution to work, but there is no guarantee or support for anything that goes wrong.

The original thinking around HDX was that a powerful client device can use its own processing power to crunch multimedia, relieving the server in the data center or cloud from having to do it. This works great if you have a high-powered workstation, laptop, or PC with a high-end graphics card. Many thin-client vendors claim to support the HDX features; however, when they actually tested the features, they found that they can't support them because their hardware is underpowered. Some high-end, thin clients, meaning expensive, do support HDX multimedia on the device itself, but even these have been hit and miss. Citrix has had a hard time getting thin-client vendors to build hardware that works. Combine this with the reality that the end user computing market is moving towards mobile devices that contain no processing power whatsoever. This is why, in XenDesktop 7, there are new features that perform all of the processing on the server and deliver the bits over the long **keyboard, video, and mouse** (**KVM**) cable, or the `ICA/HDX` protocol, to the end user device.

Thin clients, mobile devices, server-bound processing, and slow networks are all bottlenecks that render HDX unusable, and for that matter, even VDI too. So, before you commit to delivering HDX, dig deep to find out the gotchas and how to really make this work for your environment. A good case in point, a customer found while evaluating XenDesktop that the server was so CPU-and memory-bound that hundreds of users were getting poor performance and were having a bad experience, so the customer chose to buy blade PCs for each user, housed centrally in the data center, and a dedicated virtual desktop for each user, thus guaranteeing a fully powered machine-like experience to the end user. This approach is expensive, so it probably won't fit everyone's budget, but it shows there is a trade-off in HDX and VDI: cost versus performance.

Aero redirection

HDX delivers a Windows Aero or Windows 8 experience to any user, regardless of the client, by rendering the graphics content on the server. Aero redirection leverages the **graphics processing unit** (**GPU**) or integrated graphics processor on Windows client devices to provide users with a fluid Windows 7 Aero or Windows 8 desktop experience.

The Windows Aero options that become available in a virtual desktop are as follows:

- The taskbar preview when a user hovers over an image in the taskbar
- Windows Peek when a user hovers over the preview image in the taskbar
- Flip when a user presses *Alt* + *Tab*
- Flip 3D when a user presses the Windows key + *Alt*

In XenDesktop 7.x, Citrix changed the name of the feature from Aero redirection to desktop composition redirection.

Configuring Aero redirection or desktop composition redirection

The end user has to enable the Aero theme on the virtual desktop, and then the server-side graphics are automatically available when the virtual desktop is configured for Aero redirection.

Desktop composition redirection (Aero redirection) and its quality are enabled using policies.

To configure desktop composition redirection (Aero redirection), perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Desktop Composition Redirection** policy.
5. Set it to **Enabled**, click on **OK,** and then click on **Next**.
6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **Finish**, as shown in the following screenshot:

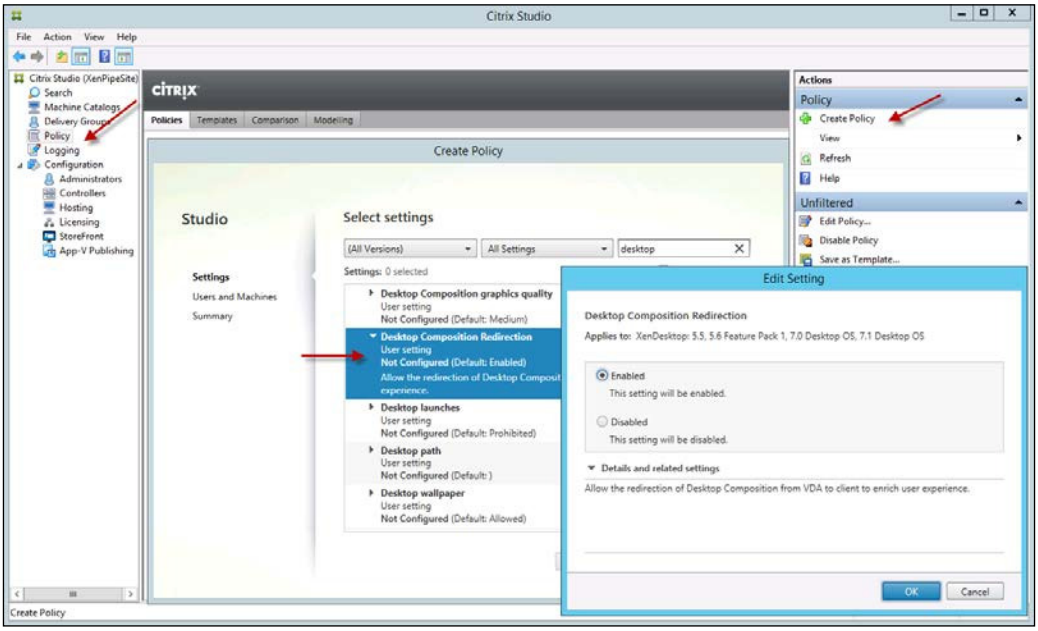

Disabling desktop composition redirection (Aero redirection) reduces the bandwidth consumption.

If you enable the **Desktop Composition Redirection** policy, be sure to enable the **Desktop Composition graphics quality** policy, as shown in the following screenshot:

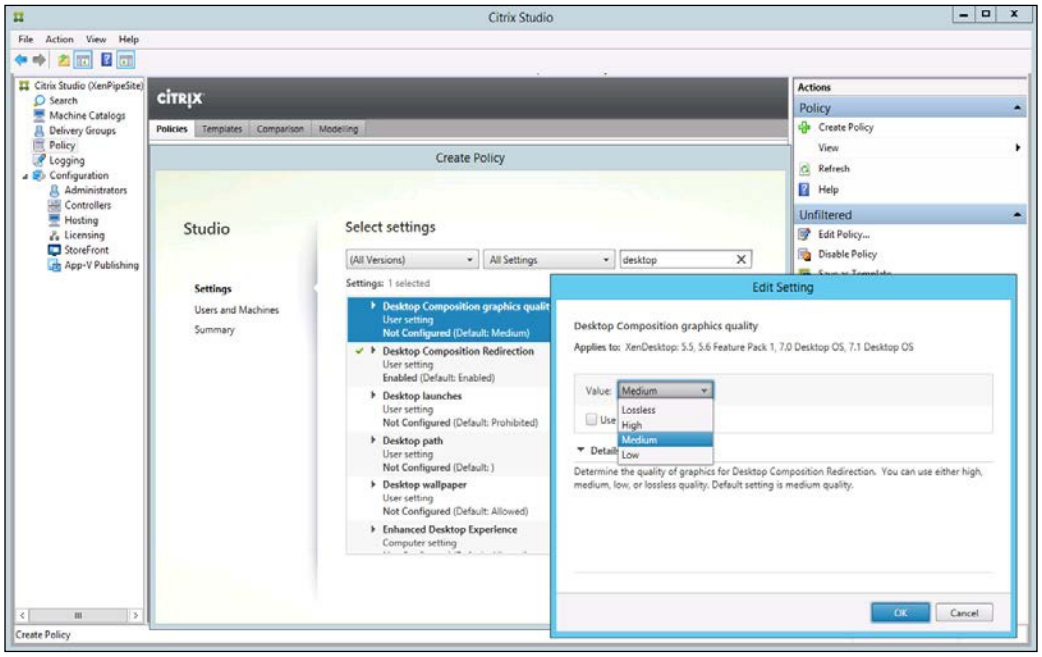

Windows Media

Windows Media client-side fetching allows a client device to pull multimedia files directly from the Internet or Intranet without having to go through the host server. This is known as Windows Media client-side redirection. By streaming media directly to the client, the server is relieved of the processing and the network is relieved of the bandwidth consumption.

Configuring Windows Media client-side fetching

Windows Media redirection needs to be enabled using policies. Make sure that the client has access to the Internet or Intranet to play the content.

Make sure the URL to the content hasn't been previously blacklisted. When an attempt to access a URL fails, the URL is added to a blacklist for the duration of the user's session. You can get more information on the blacklist and how to locate it at `http://support.citrix.com/article/CTX126817`.

Make sure the content uses one of the following protocols:

- `HTTP`, `HTTPS`
- `MMS`, `MMSU`, `MMST`
- `RTSP`, `RTSPU`, `RTPST`

If the preceding conditions aren't met, the media delivery automatically falls back to Windows Media server-side redirection. When Windows Media redirection falls back to server-side rendering, the VDI host server captures the media in its native compressed format and redirects it to the client. The client device then recreates the media pipeline to decompress and render the media received from the VDI host. This works best on clients that support the Windows operating system, because they already have the multimedia framework installed to rebuild the content delivered from the server. Linux clients can use a similar open source media framework to deliver content as well.

Windows Media client-side fetching or Windows Media server-side redirection should be transparent to the end user. What actually happens depends on several factors such as the processing power of both the server and client, combined with a low-latency and high-performance network between the two.

To configure client-side Windows Media redirection, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Windows Media client side content fetching** policy.
5. Set it to **Allowed** and click on **OK**, and then on **Next** (**Allowed** by default)
6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.

8. Click on **Finish**, as shown in the following screenshot:

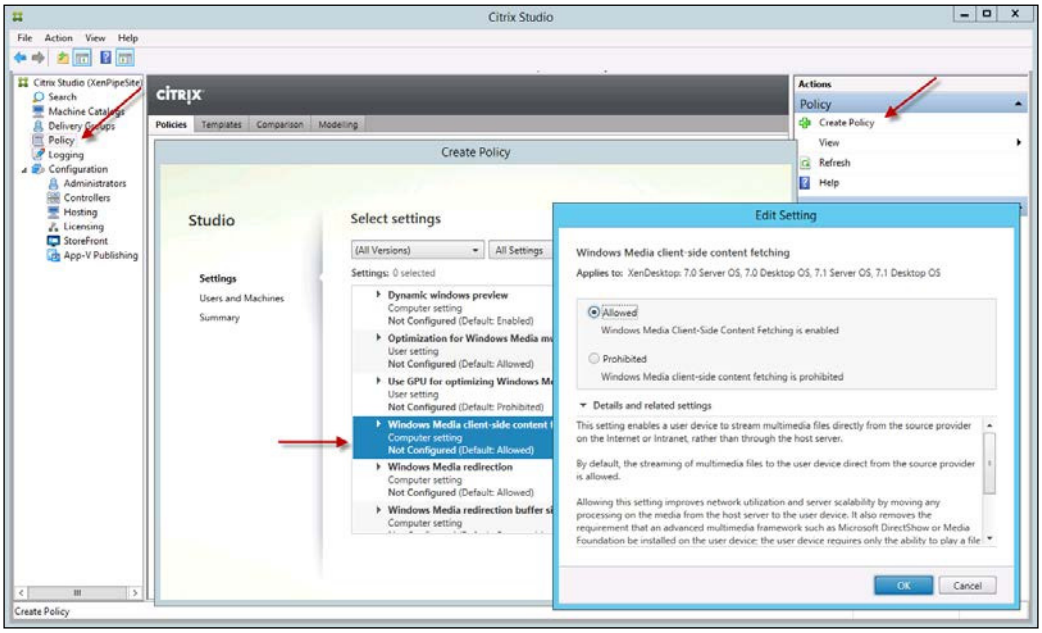

Configuring real-time Windows Media multimedia transcoding

Real-time multimedia transcoding enables audio and video media streaming to mobile devices, and by improving how Windows Media content is delivered over a WAN, it enhances the user experience. There are a couple of important terms used here. Transcoding transforms the media content from one encoding format to another. For example, compressing media to reduce the file size or converting data to a format supported on the target device. **Transrating** alters the bit rate of the media based on network conditions, for example, by decreasing the media's resolution or frame rate to achieve a lower bit rate.

When this feature is enabled, real-time multimedia transcoding is deployed automatically to enable media streaming, which provides a seamless user experience even in extreme network conditions. Transcoding occurs on the **Virtual Delivery Agent** (**VDA**) for the virtual desktop on the XenDesktop Delivery Controller. If the DA has a supported GPU for hardware acceleration, transcoding is done in the GPU; otherwise, it is performed in the server's CPU. The media is then transrated to achieve the target transmission bit rate and redirected to the client device, where it is decompressed and rendered.

Real-time multimedia transcoding enhances the user experience by converting media into a format that can be rendered locally on the client hardware by performing the following tasks:

- Eliminating the need for server-side rendering
- Lowering the bit rate of the media to match the bandwidth of the client for smooth playback of audio and video
- Providing policies to help administrators predict and manage network consumption for **Quality of Service** (**QoS**)
- Improving server utilization and scalability

To configure real-time multimedia transcoding, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Optimization for Windows Media multimedia redirection over WAN** policy.
5. Set it to **Allowed** click on **OK,** and then click on **Next** (**Allowed** by default).
6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **Finish**, as shown in the following screenshot:

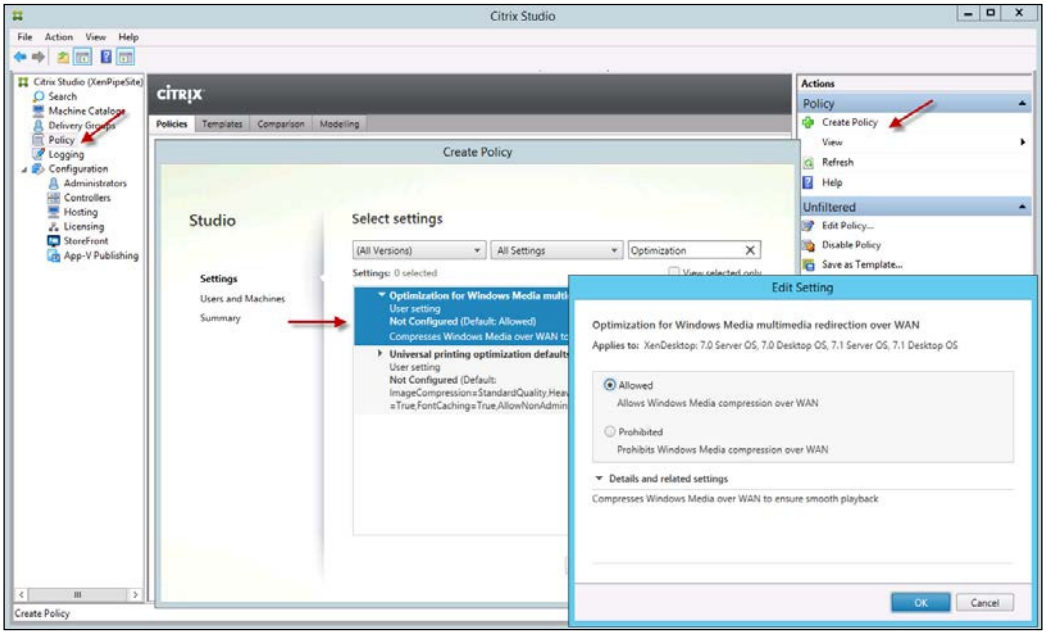

For GPU encoding, you must enable the registry key for 32-bit or 64-bit computers as follows:

- For 32-bit computers, open the registry and navigate to `HKLM\Software\Citrix\HDXMediaStream` and select **Add DWORD** `EnableNVidiaCompressor`, and set it to `1`
- For 64-bit computers, open the registry and navigate to `HKLM\Software\Wow6432Node\Citrix\HDXMediaStream` and select **Add DWORD** `EnableNVidiaCompressor`, and set it to `1`

Editing the registry incorrectly can cause serious harm to your computer and may require you to reinstall the operating system.

To turn off real-time multimedia transcoding, change the setting for **Optimization for Windows Media multimedia redirection Over WAN** to **Prohibited**.

You can also limit the video quality by setting the **Limit Video Quality** policy setting. High-quality video requires more processing and bandwidth and playing multiple videos simultaneously on the server can consume a large amount of resources.

Flash Media

Flash redirection offloads the processing of most Adobe Flash content including animations, videos, and applications to LAN- and WAN-connected hosted endpoints. By moving the processing of Flash to the user device, it reduces the server and network load resulting in greater scalability. Configuring Flash redirection requires both server-side and client-side configuration settings.

Early Flash redirection features are supported on the client-side only, while the latest Flash redirection is supported on both clients and servers. Flash redirection supports **intelligent fallback** to render content on the server when it is more efficient to do so. It also supports **Flash URL Compatibility List**, which controls whether URLs should be rendered on the client or server or should be blocked.

Flash redirection uses the Windows event log to log Flash events. You really need to check the event log to determine if Flash redirection is being used.

Configuring Flash redirection on a server

To configure Flash redirection on a server, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Flash default behavior** policy.
5. Set it to **Enable**, click on **OK,** and then on **Next**.

The options included in the **Flash default behavior** policy are as follows:

- **Block**: Flash redirection is not allowed
- **Disabled**: Flash redirection is not used and Flash content is available from the server side
- **Enabled**: Flash redirection is allowed and rendered on the client side

6. Assign the policy and click on **Next**.
7. Give a name to the policy and enable it.
8. Click on **Finish**, as shown in the following screenshot:

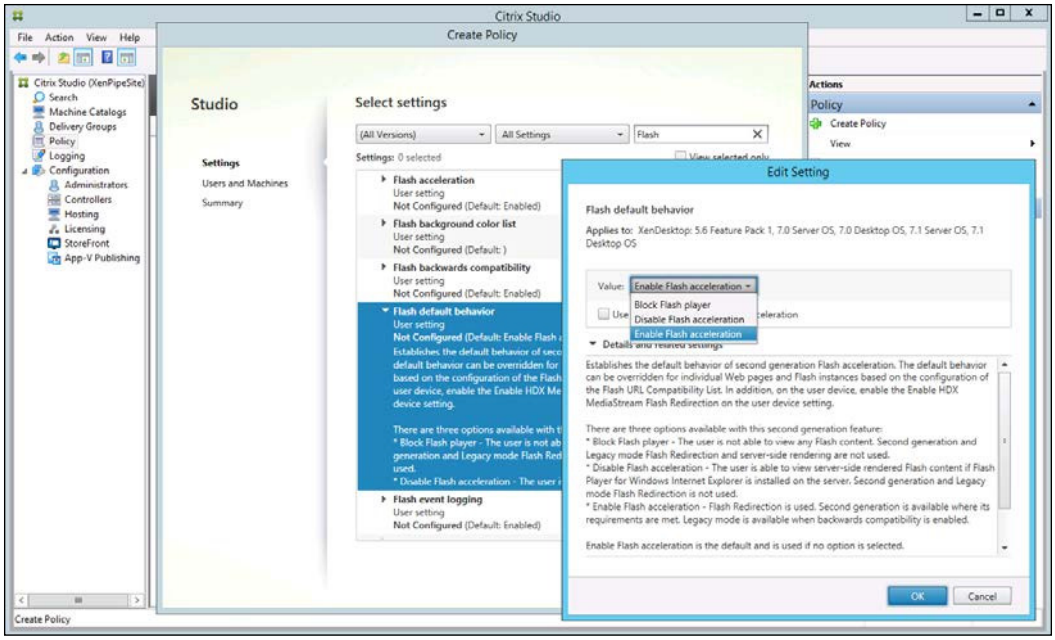

If you enable Flash redirection, make sure you also enable the following policies:

- **Enable** or **Disable the HDX MediaStream Flash Redirection** policy on the user device, depending on where you want it rendered.
- Ensure that **Use HDX MediaStream Flash redirection** is not set to **Never**.
- Enable the **Flash intelligent fallback** policy setting to allow small Flash movies used for advertisements to be rendered on the server.
- Select the **Flash server-side content fetching URL list** policy setting to specify websites whose Flash content you would rather want be rendered on the server. This is useful for Intranet sites and when the client doesn't have access to the Internet. This setting should also be used with the **Enable server-side content fetching** policy setting. You can use an * (asterisk) as a wildcard.
- Select the **Flash URL compatibility list** policy setting to specify whether you want the content rendered on the client, server, or have it blocked from rendering.
- Select the **Flash background color list** policy setting to match the colors of web pages with Flash instances.

Configuring Flash redirection on the client

Before you configure the client settings, make sure you have installed Citrix Receiver and Adobe Flash Player on the client device. Flash redirection on the client is enabled by default, so nothing else is required; however, you can change the default settings using Group Policy Objects in the Group Policy Editor.

To configure Flash redirection on the client, perform the following steps:

1. Log in to the client device.
2. Launch Microsoft Group Policy Management Editor (`gpedit.msc`).
3. Under **Action**, select **Add/Remove Templates**, click on **Add**.
4. Import and add the HDX MediaStream Flash redirection using the client administrative template `HdxFlash-Client.adm`.
5. For 32-bit computers, the template is found in `%Program Files%\Citrix\ICA Client\Configuration`.
6. For 64-bit computers, the template is found in `%Program Files (x86)%\Citrix\ICA Client\Configuration`.
7. Expand the administrative templates and select **HDX MediaStream Flash Redirection - Client**.

8. Enable or disable the **Enable HDX MediaStream Flash Redirection on the user device** policy setting. You can choose additional options beyond that, but mostly it is used with the default option, as shown in the following screenshot:

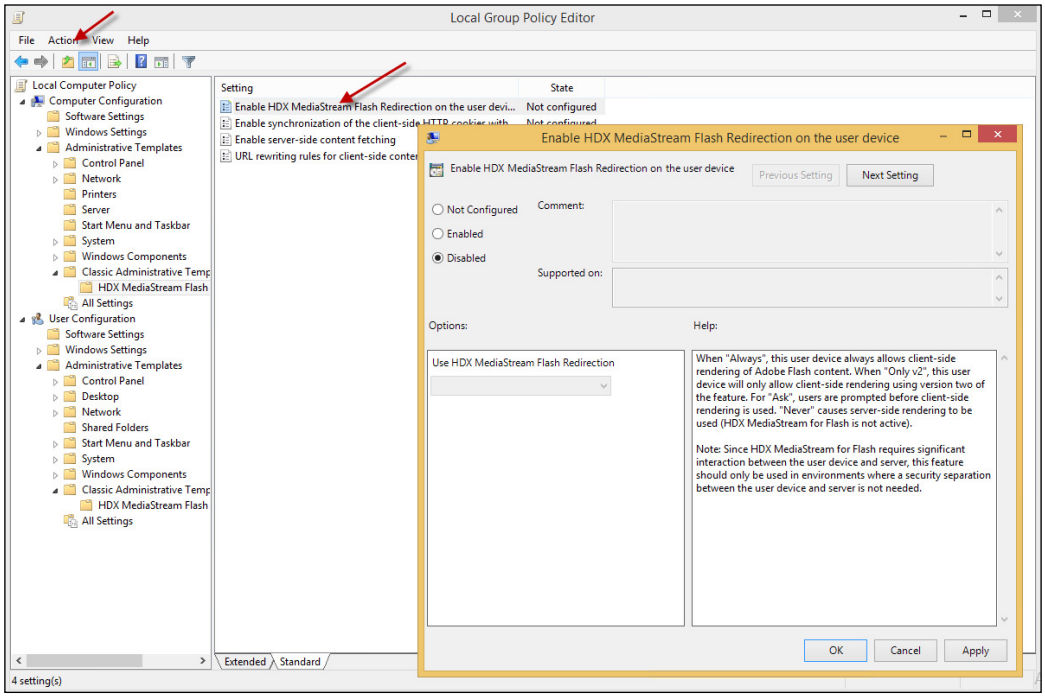

9. Additional options can be selected in the **Citrix Receiver - Desktop Viewer Preferences** window under the **Flash** tab on the client from within the client's virtual desktop session, as shown in the following screenshot:

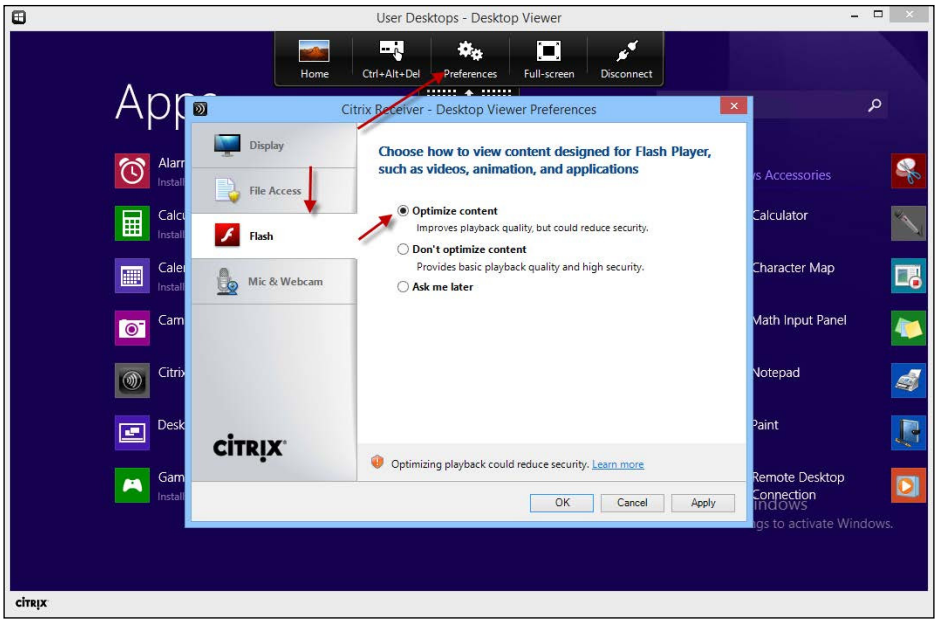

By default, Flash redirection downloads Adobe Flash content to the user device. You can have the server download the content by enabling server-side content fetching. This causes content to be downloaded to the server and then sent to the user device. You do this by enabling the **Enable server-side content fetching** policy setting found in the **HDX MediaStream Flash Redirection - Client** administrative template.

You can perform URL rewriting for Flash content using the **URL rewriting rules for client-side content fetching** policy setting. You provide two URL patterns using regular Perl expressions so that when a user requests content from a URL matching the first expression, it redirects it automatically to the URL with the second expression.

HDX™ 3D

HDX 3D is also referred to as HDX 3D Pro, although there is some internal debate at Citrix about whether to drop the Pro extension from the name, HDX 3D enables you to deliver desktops and applications that require a GPU for hardware acceleration. HDX 3D is particularly suitable for use with Direct X- and OpenGL-driven applications, and with rich media such as videos.

Ultimately, HDX 3D allows you to replace expensive workstations with simpler, low-end devices, moving the graphics processing to the server. HDX 3D is useful for **computer-aided design** (**CAD**), **computer-aided manufacturing** (**CAM**), **computer-aided engineering** (**CAE**), **geographical information software** (**GIS**), and **picture archiving and communication systems** (**PACS**) for medical imaging.

A small subset of the market is virtualizing their 3D applications. What these 3D customers and vendors are finding is that even though their solutions are very costly, they can still save some money with HDX. Running a 3D application on a server and delivering it to a virtual desktop and then out to a client, requires the presence of a high-performance 3D video card in the server that is running XenDesktop. Until recently, you could only have one user per GPU, which really limited the number of virtual desktops you could deliver. For example, if you had one card with two GPUs, you could only serve two desktops. A new feature developed by Citrix in XenServer and NVidia now allows GPU sharing, where you can now share a GPU among several HDX 3D users, which seems a little more practical. Whether you are buying NVidia's virtual desktop appliance that contains eight Kepler2 video cards or a server that supports multiple Kepler2 cards, you can now make 3D application delivery from a virtual desktop more economically feasible.

For a list of compatible GPU devices that work for XenDesktop, refer to the hardware compatibility list at `http://hcl.xensource.com/GPUPass-throughDeviceList.aspx`.

If you have questions about the effect of your 3D application supported in Citrix XenDesktop, the best thing to do is contact your 3D application vendor.

GPU versus vGPU

There is a difference between GPU and vGPU, so you should know the difference.

GPU

A GPU is used for HDX 3D pass-through capability, where the XenDesktop VDA passes through the graphics processing to the GPU installed on the server. A GPU is typically a high-end graphics processor installed as an add-on card to a workstation or server. The leader in this space is NVidia, and Citrix has been doing a lot of work to integrate their products. GPUs are very powerful. GPUs are typically **Single Instruction Multiple Data** (**SIMD**) processors, meaning that they can perform the same operation on multiple data units simultaneously. GPUs are excellent for highly intensive graphics applications.

For more detailed information on implementing GPUs, refer to the article at `http://www.citrix.com/skb/articles/RDY12010`.

vGPU

A **Virtual Graphics Processing Unit** (**vGPU**) contains multiple cores of GPUs on a single processor. Similar to the way Intel and AMD have multiple CPU cores on a single processor for virtualization of the CPU, NVidia has multiple vGPUs capable of being shared by VMs. vGPUs are **Multiple Instruction Multiple Data** (**MIMD**) processors, meaning they can execute many instructions on multiple units of data in parallel. The vGPU feature in XenDesktop allows multiple VMs to utilize the processing power of a single GPU. You will find these capabilities in the NVidia GRID cards, the **Kepler1** (**K1**) and **Kepler2** (**K2**) cards.

For more detailed information on implementing vGPUs, refer to the article at `http://www.citrix.com/skb/articles/RDY12202`.

HDX™ 3D requirements

XenDesktop with HDX 3D Pro is a desktop and app virtualization solution that supports high-end designers and engineers of 3D professional graphics applications and provides cost-effective support to viewers and editors of 3D data.

Client

Both servers and clients must meet the minimum requirements to support HDX 3D. To access desktops or applications delivered with HDX 3D, the user must install Citrix Receiver. All versions of Citrix Receiver are supported; however, you should be using the latest one. User devices do not need a dedicated GPU to access desktops or applications delivered with HDX 3D; now that XenServer supports GPU sharing, HDX 3D supports all monitor resolutions supported by the GPU on the server. Make sure your client device has enough processing power to handle the compression codec.

Server

The server needs to support a GPU pass-through or vGPU for HDX 3D, and a list of compatible ones can be found in the preceding link. The virtual desktop also needs a special VDA for HDX 3D, and this option can be selected while installing the VDA on the virtual desktop's master image.

HDX™ GPU sharing

HDX 3D takes graphics applications and renders them on the GPU, freeing up the CPU. With HDX 3D, multiple users can share graphics cards. The number of users/virtual desktops that can be supported on a single GPU card depends on the amount of video RAM and the GPU processor. 8-10 users have been found to be supported on NVidia cards running a specific program, and it is possible to achieve 32 concurrent users on high-end GPUs. Conduct the research, find out if the hardware and GPU vendor support VDI, GPU sharing, and how many users per GPU they support.

A list of compatible GPU pass-through devices can be found at `http://hcl.vmd.citrix.com/GPUPass-throughDeviceList.aspx`.

HDX™ 3D – how it works

HDX 3D supports physical computers as well as virtual computers with a GPU pass-through. This includes desktop, blade, racked workstations, and XenServer VMs. The XenServer GPU pass-through feature enables you to create VMs with exclusive access to dedicated graphics processing hardware. You can install multiple GPUs on the Hypervisor and assign VMs to each of these GPUs on a one-to-one basis. This is known as a GPU pass-through.

HDX 3D uses compression technologies that are different from the standard VDA to optimize the delivery of graphically intensive applications. If a compatible NVIDIA CUDA-enabled GPU is available on the Hypervisor, HDX 3D engages a codec that uses the GPU on the host to encode data. At the other end, user devices need the decoder for the codec to receive GPU-encoded data, but they do not need a dedicated GPU.

HDX 3D uses CPU-based compression as the default compression technique for encoding. GPU-based deep compression is used on systems with a low CPU capacity. If the user device does not have the decoder for the GPU codec and the CPU capacity is low, then HDX 3D uses CPU based-compression. In the case where pixel-perfect graphics are absolutely necessary, such as medical imaging applications, and lossless compression is required, the CPU codec is also used.

GPU-based deep compression makes the best use of the available bandwidth because you can deliver complex graphics over WAN connections with bandwidths as low as 2 Mbps. On LAN connections, the bandwidth consumed by graphically intensive applications can be reduced dramatically without compromising on the high definition user experience. CPU-based compression requires at least 3 Mbps of network bandwidth to deliver an interactive user experience. When lossless compression is enabled, this requirement increases to 10 Mbps.

Installing and configuring HDX™ 3D

To enable users for HDX 3D, you need to install the VDA for HDX 3D. It may be a good idea to create a separate desktop master image for HDX 3D users as they require special drivers. Then, you create a machine catalog and delivery group to assign the desktop or VM-hosted app to a user.

To deliver a 3D application in XenDesktop, perform the following steps:

1. Make sure XenDesktop is installed and you have a Site created.
2. Create the VM or physical machine that will host the graphical application and connect it to the domain.
3. Make sure you indicate that this VM can use either pass-through or vGPU.
4. Ensure that the NVidia drivers are installed along with the NVidia GPU card(s).
5. Install the graphics application.

By default, each virtual CPU that you allocate to a XenServer VM is assigned to a single core socket. If you use XenServer, you can assign multiple cores per virtual CPU. You can read about this in the article at `http://support.citrix.com/article/CTX126524`.

To install HDX 3D on a master image, perform the following steps (similar to the procedure in *Chapter 2*, *Installing XenDesktop®*):

1. Load the XenDesktop install media.
2. In the XenDesktop installation wizard, select **Virtual Delivery Agent for Windows Desktop OS**.

3. On the **HDX 3D Pro** screen, select **Yes, install the VDA for HDX 3D Pro**, and then select **Next**, as shown in the following screenshot:

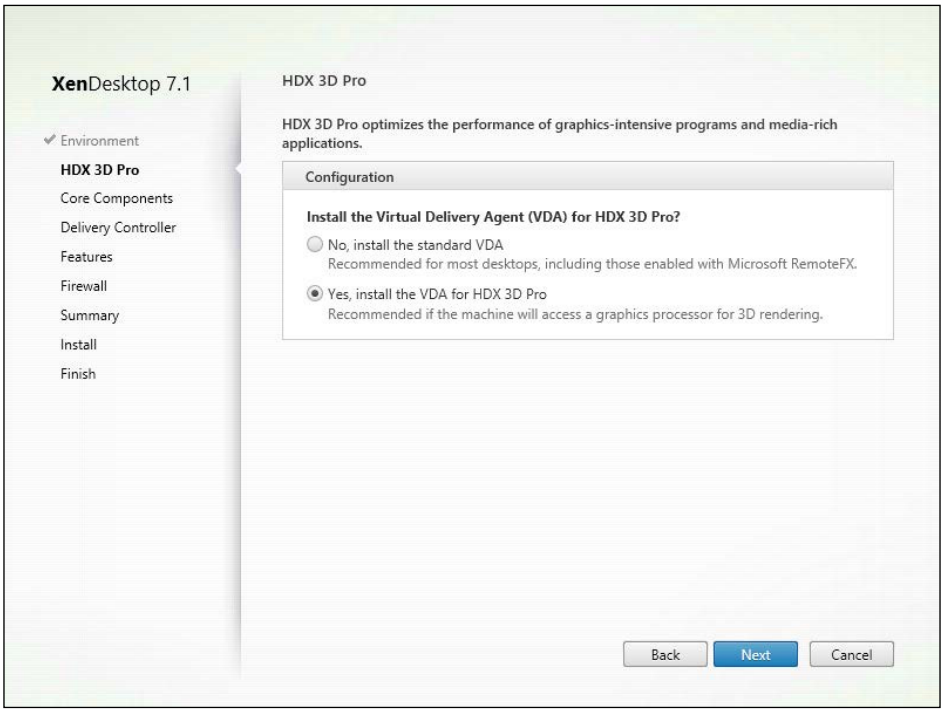

4. Install Citrix Receiver if you plan to deliver the applications from a XenApp server. You do not need Citrix Receiver if you will be delivering the graphics application as a VM-hosted app.

HDX 3D Pro is not available on an RDSH server.

5. Specify the controller location.
6. If you select shadowing and real-time monitoring, remember it takes a lot of processing power.
7. Make sure you select **Optimize Performance** in **Features**.
8. Configure firewall ports if you need to.
9. Click on **Install**.

After installation, log on to XenDesktop Studio, create a machine catalog, and add the computer hosting the graphical application to the catalog. Then, create a delivery group or an application delivery group using the machine catalog.

You can also install the VDA for HDX 3D from a command prompt by running `XenDesktopVdaSetup.exe`.

Upgrading HDX™ 3D

To upgrade to a new version of HDX 3D, uninstall the HDX 3D components and the HDX 3D VDA from the desktop image, and install the updated version. The same applies if you are upgrading from the regular VDA to the HDX 3D DA.

Configuring monitors for HDX™ 3D

In order for users to view graphical applications in a maximized window across multiple monitors, the computer hosting a 3D application or VM must be configured with at least the same number of monitors that users will require and with the same or higher resolution. You do this by attaching the monitors to the physical computer or physical server that hosts the VM before installing the DA for HDX 3D. You can also configure virtual monitors after the installation of the HDX 3D DA.

For example, you can configure a 1920 x 1200 pixel display as the primary monitor and a 1280 x 1024 pixel display as the secondary monitor. When a user uses the graphical application, they can use a 1920 x 1200 display or smaller. Users with larger screens and resolutions will not be able to see the application above 1920 x 1200 pixels. In a dual monitor setup, the window will span across both the monitors but will max out at the lowest monitor screen size. HDX 3D supports more than two monitors; however, some GPUs don't, so check with the GPU manufacturer.

To configure larger or additional monitors after installing the VDA for HDX 3D, you will need to create virtual monitors using the **Control Panel** of the GPU on the computer or VM hosting the graphical application. Obtain and attach the virtual monitor's **extended display identification data** (**EDID**) file to monitor resolutions that are the same or larger than those required by the users.

For more information on configuring monitors for HDX 3D, refer to the article at `http://support.citrix.com/article/CTX131501`.

Configuring image quality

Once you have installed the VDA for HDX 3D (HDX 3D Pro), you can configure the image quality configuration tool for your users. You use the image quality configuration tool to set policies for HDX 3D image quality. The policy settings it controls are of the **Enable lossless** and **HDX3DPro quality settings** policies.

A list of possible policy settings to use with HDX 3D are listed in the following table:

The HDX 3D policies			
Policy	**Function**	**Default**	**Use with policy**
Enable lossless	This specifies whether or not users can enable and disable lossless compression using the image quality configuration tool. Lossless compression is required to deliver pixel-perfect images.	**Disabled**	
Fixed quality	To ensure that all frames are lossless when interacting with an image, users must also select the **Fixed Quality** checkbox. When enabled, users can adjust the image quality with the slider.	**Disabled**	
HDX3DPro quality settings	This specifies the minimum and maximum values that define the range of image quality adjustment available to users in the image quality configuration tool.	**Not set**	**Fixed quality**

Configuring audio

Audio is controlled through policies in XenDesktop Studio. You control the settings for audio through the following user policy settings:

- **Audio quality**
- **Client audio redirection**
- **Client microphone redirection**
- **Audio redirection bandwidth limit**
- **Audio redirection bandwidth limit percent**

- **Audio over UDP Real-time Transport**
- **Audio UDP port range**

Refer to *Appendix B, XenDesktop® Policy Setting Reference*, for a list of policies and related settings.

In general, higher audio quality requires more bandwidth and CPU cycles because of the large amount of audio data that must be sent to users. You can use sound compression to balance audio quality with session performance. You may want to create separate audio policy groups for different types of users based on their LAN or WAN utilization.

Configuring webcams

HDX provides a webcam video compression option to improve bandwidth efficiency during video conferencing. The user can change the settings in the **Citrix Receiver – Desktop Viewer Preferences** dialog box.

To configure webcam video compression, perform the following steps:

1. Install Citrix Receiver on the client device.
2. Install the video conferencing application.
3. Make sure the user's hardware can produce sound.
4. Use the default camera settings.
5. Install drivers for the webcam, if needed.
6. Enable the following Citrix policies in Studio:
    - **Client audio redirection**
    - **Client microphone redirection**
    - **Windows media redirection**

Configuring color compression

HDX provides high-quality graphics experiences for Windows 7 and Windows 8 desktops. If you have low bandwidth, you can improve responsiveness by enabling extra color compression. Using color compression results in lower quality graphics. You can enable or disable extra color compression in Studio from the **Extra Color Compression** policy setting.

Configuring network priorities

The user experience can be controlled in the devices outside of XenDesktop's control, such as network routers and switches. You can assign priorities to network traffic across multiple connections for a session in routers and switches that support QoS. XenDesktop ICA/HDX traffic uses four TCP connections, and one UDP connection carries the ICA/HDX traffic between the user device and the server. Each virtual channel or the `ICA/HDX` protocol stream is coupled with a priority and transported on the corresponding connection. You can set the priorities for each ICA/HDX stream independently, based on the TCP port number used for the connection.

The four priorities are as follows:

- **Very high**: This is used for real-time connections such as webcams
- **High**: This is used for interactive elements such as keyboard, mouse, and video
- **Medium**: This is used for bulk processes such as client drive mapping
- **Low**: This is used for background activities such as printing

XenDesktop supports multiple channel streaming for DA on Windows 8 and Windows 7. QoS is supported only when multiple session reliability ports or the `CGP` protocol is configured. `CGP` also needs to be configured on the organization's routers.

It's important to note that `CGP` needs to be configured on all of the organization's routers to be effective, and not only on one router.

To enable QoS for multiple ICA/HDX streaming connections, you configure separate ports for the `ICA/HDX` protocol and then you configure the machine policies, **Multi-Port Policy** and **Multi-Stream**; and the user policy, **Multi-Stream**.

To assign priorities to network traffic, perform the following steps:

1. Launch Studio.
2. Select the **Policy** node.
3. Click on **Create Policy**.
4. Select the **Multi-Port Policy** policy and select a default value for `CGP`.
5. Type additional CGP ports in **CGP port 1**, **CGP port2**, and **CGP port3** as needed to identify priorities for each.

6. Select **OK** and click on **Next** (**Allowed** by default).
7. Assign the policy and click on **Next**.
8. Give a name to the policy and enable it.
9. Click on **Finish**, as shown in the following screenshot:

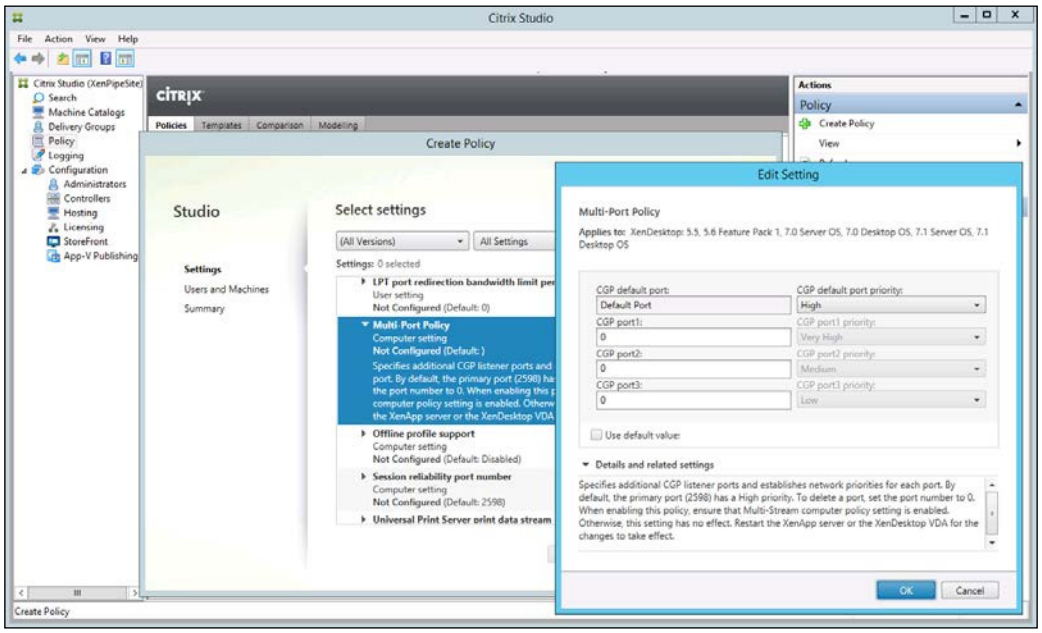

You will also need to enable the **Multi-Stream computer setting** and the **Multi-Stream user setting** policies.

Adaptive display

Adaptive display automatically adjusts the image quality of videos and moving graphics, such as transitions in slide shows, based on the available bandwidth. When adaptive display is enabled, users see smooth-running transitions in presentations with no reduction in quality.

Adaptive display generally doesn't require any configuration because it is enabled by default and is self-tuning. To change the settings, you need to configure the **Moving image compression** and **Lossy compression level** policies.

Summary

As the market moves more towards mobility, producing HDX on user devices will be a key component of delivering desktops with XenDesktop. Although using vGPUs is a little pricey today, I am sure all of the integration work will pay off. If you have 3D applications, you need to virtualize; this can be done, and you learned how to do it in this chapter. In the next chapter, we will learn about application delivery.

10
Application Delivery

Application delivery is the delivery of applications. This is not as simple as it sounds; the actual mechanics and details of delivering applications is wide ranged and complex. If you are familiar with Citrix XenApp, formerly Citrix Presentation Server that was formerly MetaFrame for Windows, then you have known about delivering applications for quite some time. In fact, not only can XenApp deliver applications, but it can also deliver, that's right, desktops. Desktop virtualization is not really all that new; it is just newly popular as the virtualization of data center resources is booming.

Citrix made some acquisitions, rebuilt some key pieces of technology, and released Citrix XenDesktop to appeal to this burgeoning marketplace, specifically for desktop delivery or **Virtual Desktop Infrastructure** (**VDI**). As the product line moved forward, Citrix made a major architectural change in its widely popular XenApp product and merged it with XenDesktop. A major architectural change is no understatement. Many of the familiar components of XenApp are simply discontinued. Many third-party vendors, who built applications around the earlier versions of XenApp now have to rebuild their applications.

In the new version of XenDesktop, Version 7.x, you will find that desktops and applications have been merged into a common architecture for hosted applications (XenApp) and desktops (XenDesktop) in order to provide a more unified management experience. Let me say this in another way: the capabilities previously available within XenApp are now delivered from XenDesktop. However, if you don't need desktops and just want applications, then there is a XenApp Version 7.x coming soon.

In this chapter, we will discuss the following topics:

- What delivering applications means in XenDesktop
- Differences between XenApp and XenDesktop
- Application Delivery Controllers
- Application Delivery Networks

Delivering applications

In the old days, before XenDesktop, we delivered applications by publishing them in XenApp. When you publish an app, you make it available to your users and they can connect and launch the application through a client; today, that client is called Citrix Receiver. In the prior versions of XenDesktop, you would build a XenApp farm (Site) and install Citrix Receiver on the XenDesktop virtual desktop images so that users could subscribe to and launch applications from a XenApp farm (Site) to the XenDesktop virtual desktop. This is done for scalability. This paradigm of delivering applications to virtual desktops has been carried forward in XenDesktop 7.x; that is, the applications still need to be hosted or run from separate application servers.

In XenDesktop 7.x, applications are delivered from VMs called VM hosted apps. In XenDesktop 7.x, you can now publish applications along with desktops. In *Chapter 2*, *Installing XenDesktop®*, in the *Step 4 – creating the virtual desktop and application delivery master images*, *Step 5 – installing the Virtual Delivery Agent on the master images*, *Step 7 – creating the machine catalogs*, and *Step 8 – creating the delivery groups* sections, we walked you through how to install XenDesktop for applications along with desktops. In the same way that you create desktops in Studio, you create and add applications in Studio by creating a Site, a machine catalog, and then a delivery group. In XenDesktop 7.x, the XenApp functionality is integrated into XenDesktop. However, if it is just applications that you want to deliver, there will be a XenApp Version 7.x available soon.

Differences between XenApp® and XenDesktop®

As mentioned previously, there are some major changes that occur as a result of merging XenApp into XenDesktop. If you have any history of managing XenApp farms, then these changes will resonate with you.

What's new?

The following are the new additions in XenDesktop:

- **StoreFront**: This was previously called the **Web Interface**, which is no longer available. Users connect to the StoreFront for the delivery of applications and desktops. However, at the time of this writing, I was told that the Web Interface would be brought back for XenApp.

- **Citrix Studio**: Even the **Delivery Services Console** (**DSC**) was a new term that was coined just recently. So, it is no surprise that it changes again with this update. DSC no longer exists. The delivery of applications and desktops occurs from one location, that is, Citrix Studio.
- **Administrative roles**: In XenDesktop 7.x, you can create custom administrators whose permissions are based on specific roles and scope. A role represents a job function and has defined permissions associated with it. A scope represents a collection of objects. A scope is used to group objects that map to your organization. There are several roles included in XenDesktop 7.x for help desk, applications, hosting, catalog, and administrator. Each role contains specific permissions.

What's gone?

The following have been discontinued or are no longer available in XenDesktop:

- **Terminal Services**: Terminal Services (Remote Desktop Services) is no longer required on the servers running the controller.
- **Zone master and data collector**: There is no dedicated zone master. In XenApp, there is a zone master or data collector responsible for user connection requests and communication with Hypervisors. However, in XenDesktop 7.x, this function is distributed evenly across all the controllers in the Site.
- **Shadow taskbar**: The shadow taskbar is gone. To view and interact with other users' sessions remotely, you can use the shadow feature from the Director console.
- **Farms**: Farms are now called Sites. This applies to both XenApp and XenDesktop.
- **XenApp Web and XenApp Services Sites**: These are gone; everything is published from Studio. StoreFront stores are equivalent to these Sites.
- **Online plugin and offline plugin**: The online and offline plugins were confusing at best. These are not completely gone; however, they are automatically administered through Citrix Receiver. Just in case you were wondering, the online plugin was used to connect to a XenApp Services Site, and a browser was used to connect to a XenApp Web Site. The offline plugin was used as a browser plugin. Before this, the client was called the **Program Neighborhood Agent** (**PNA**).

- **SmartAuditor**: The SmartAuditor feature in XenApp is gone, and now you have to look at third-party solutions such as ObserveIT (`http://www.observeit.com/`) for user monitoring and auditing.
- **Citrix streaming**: Streaming is no longer available in XenDesktop 7. In order to use this type of functionality, you need to use App-V.

What's changed?

The all too familiar **IMA data store** is gone. The central database that is used to store the configuration information for XenApp and XenDesktop is now Microsoft SQL Server. This database holds both configuration information and session information. IMA is now referred to as the **FlexCast Management Architecture** (**FMA**).

The localhost cache is no longer supported and you will need to look to Microsoft SQL clustering/mirroring and the Hypervisor for **High Availability** (**HA**) features.

HA can be achieved by using Microsoft SQL clustering or mirroring. You can also deploy Microsoft SQL Server as a VM and use the Hypervisor HA feature set.

Oracle is no longer a supported database.

What hasn't changed?

The XML broker continues to run on the Delivery Controller.

The previous sections talked enough about the components that have changed. The components and terminology for XenDesktop 7.x can be reviewed in *Chapter 1*, *Designing a XenDesktop® Site*.

Application Delivery Controllers

Any discussion about application delivery would not be complete unless we talk about the **Application Delivery Controllers** (**ADCs**). Citrix has a product called **NetScaler**, and its main functionality is application delivery. I know that we've been talking about application delivery from XenDesktop; however, there is actually an entire market just for ADCs outside of XenApp and XenDesktop. The term ADC is sort of a misnomer; that is, it doesn't actually deliver applications but rather sits in between XenDesktop and the user to optimize and secure the application delivery.

ADCs grew up as separate physical networking devices originally used for load balancing, HA, optimization, and security. ADCs are typically deployed in HA pairs to provide stateful failover in the event that one of them goes down. These days, the best option for HA is to run a pair of these devices in a HA pair for high availability and scalability. These devices can now be found as virtual appliances, such as the NetScaler VPX. Multiple NetScaler VPXs can be run on a NetScaler SDX. Clustering also supports up to 32 NetScalers in a cluster; so, it is really the best option for growth and scalability. As you might have guessed, NetScaler has a place in the XenDesktop architecture.

ADC as a network function has nothing to do with the XenDesktop Delivery Controller; they are completely separate. In fact, if you look at the following network diagram, you will see a component called **StoreFront**. StoreFront runs as a web service on Windows Server. The NetScaler ADC is used to frontend StoreFront because it does load balancing, scales better, and provides frontend high availability and stronger security. NetScaler provides data transport security through the use of encrypted communications with SSL. For security, scalability, and high availability, running StoreFront behind NetScaler is preferred. The network diagram is as follows:

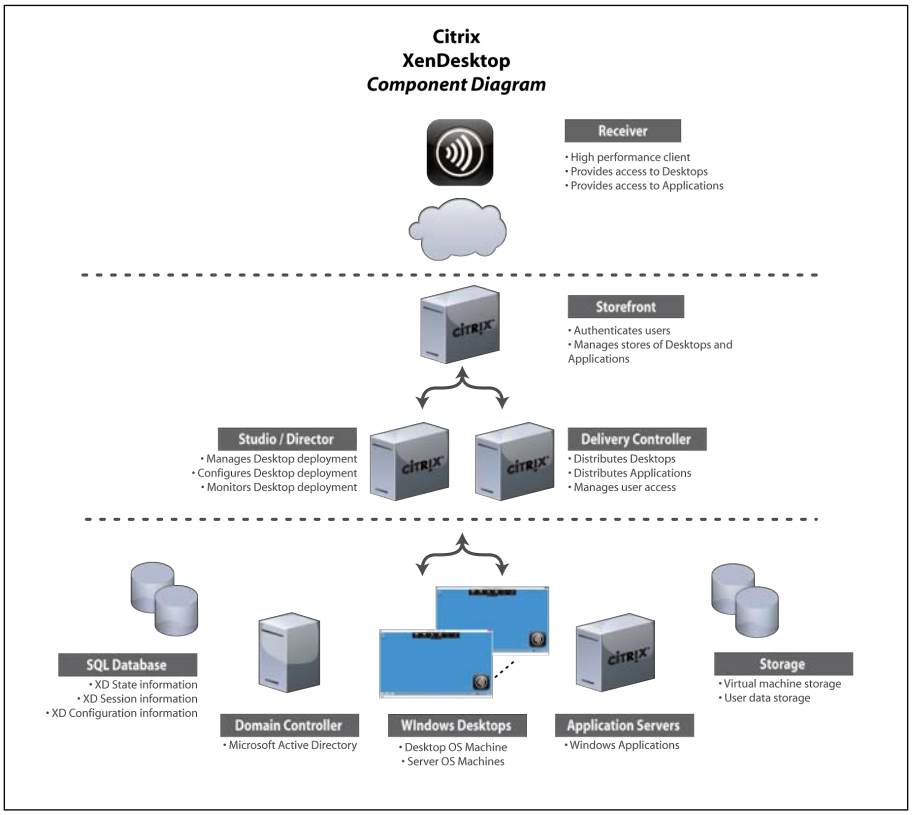

Application Delivery Networks

If you have built a XenApp or XenDesktop implementation, you have already built an **Application Delivery Network** (**ADN**). The **Application Centric Infrastructure** (**ACI**) is the same thing with a different name. There are a few more components to consider in order to complete the picture such as high availability, transport security, load balancing, and WAN optimization. An ADN is the complete end-to-end view of the application's delivery infrastructure, whereas ADC is only one piece. An ADN with all of the components included provides a complete system for application delivery, high availability, security, and optimization.

The two Citrix components that complete the XenDesktop ADN are Citrix NetScaler and Citrix CloudBridge, formerly **Branch Repeater** and **WANScaler**, respectively. NetScaler provides TCP multiplexing, load balancing, compression, caching, security, and high availability. CloudBridge provides data deduplication and the acceleration of data delivery. If you have ever seen CloudBridge in action, it is truly amazing; in some cases, it returns 90 percent of your bandwidth. Also, note that Citrix CloudBridge is the only technology that is authorized by Citrix to accelerate ICA/HDX traffic, which is the protocol that XenDesktop uses to deliver desktops and applications.

Both NetScaler and CloudBridge can be implemented as physical or virtual appliances. Where NetScaler and CloudBridge become useful is in the move to cloud computing or extending your data center into the cloud. If you have a cloud burst or disaster recovery Site sitting on hold in the cloud, the best way to keep it on hot standby is to extend your existing data center by front-ending it with NetScaler for HA and connecting it to the cloud across the CloudBridge.

Citrix NetScaler and CloudBridge can be used with the same Citrix License Server used in XenDesktop, yet they are managed separately through a web browser or Citrix Command Center.

You can read more about these technologies at the following websites:

- **NetScaler**: `http://www.citrix.com/netscaler`
- **CloudBridge**: `http://www.citrix.com/products/cloudbridge/overview.html`

If you choose to implement CloudBridge, make sure that you have enabled the CloudBridge settings in Citrix Receiver.

To implement CloudBridge in Citrix Receiver, perform the following steps:

1. Download and install the CloudBridge plugin on the client machine. It would be good to do this in the *Step 9 – installing the Citrix Receiver on the client devices* section of *Chapter 2*, *Installing XenDesktop®*, at the same time when you are installing the Citrix Receiver on the user client device(s).
2. Launch the CloudBridge plugin from the **Start** menu.
3. Right-click on the Citrix Receiver icon in the taskbar and select **About**.
4. Navigate to **Advanced** | **Accelerator Settings** | **Manage CloudBridge**.
5. Enter the signaling IP address of the CloudBridge appliance at the XenDesktop Site.
6. Click on **Apply**, as shown in the following screenshot:

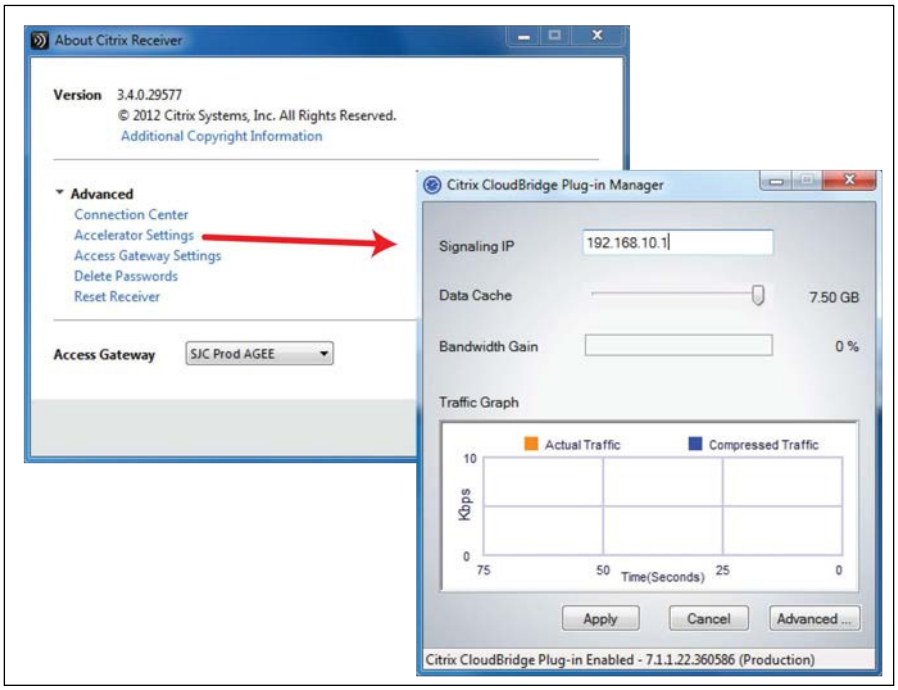

If you want an SSL acceleration, you will need to upload the CA certificate and client certificate into the accelerator.

The diagram that follows shows a hypothetical XenDesktop Site in the **Local Datacenter** front-ended by Citrix NetScaler, extended into the **Cloud** with the NetScaler **CloudBridge Connector VPN Tunnel**, and the acceleration of data using the **CloudBridge plug-in** for **Receiver** and CloudBridge VPX in the cloud. In reference to the XenDesktop diagram in *Chapter 1*, *Designing a XenDesktop® Site*, these are new components to that diagram, and the StoreFront server is placed in a load balancing service behind NetScaler to provide high availability for StoreFront.

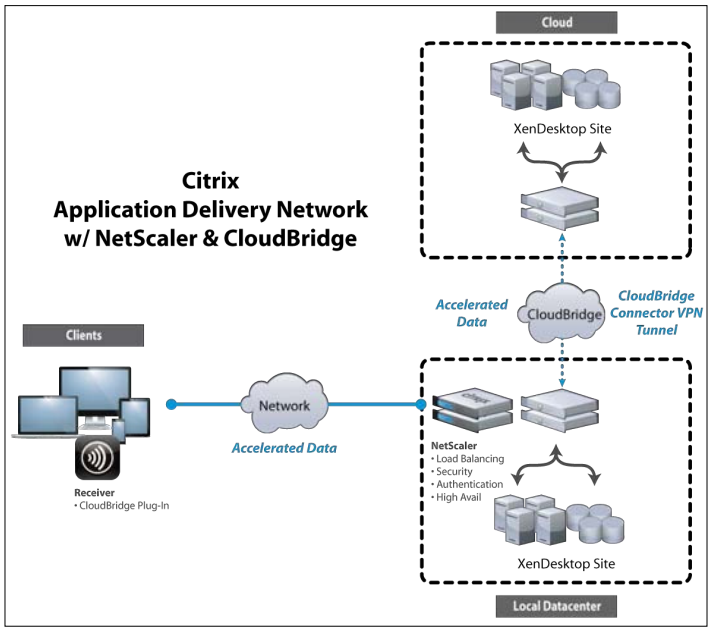

Summary

Application delivery is about receiving applications from anywhere on any device that can be anywhere. Application delivery is Citrix's sweet spot as they have been in this business since the start of the company in 1989. It started out as MetaFrame for Windows, evolved as Citrix Presentation Server, migrated to XenApp, and has now morphed into XenDesktop for the virtual delivery of both applications and desktops. With the popularization of virtualization for data center devices, such as network, compute, and storage, the virtualization of applications and now desktops is becoming more popular. Also, the options for the delivery of apps and desktops are abundant and powerful when using the complete vision of application delivery networks formulated with Citrix products. In this chapter, you learned about application delivery and how additional Citrix products complete the end-to-end desktop and application delivery vision. In the next chapter, we will discuss the XenDesktop SDK.

11
Working with the XenDesktop® SDK

Citrix has a thriving partner community, and developers will agree to this; it can be attributed to the **Software Development Kits** (**SDKs**) or **Application Programming Interfaces** (**APIs**). Citrix builds SDKs or APIs for almost all of their products so that partners and third-party developers can integrate their products with Citrix products. SDKs and APIs are the perfect way to provide integration and control with Citrix products without the need for Citrix to expose the important intellectual property.

XenDesktop has an SDK that is based on Microsoft Windows PowerShell Version 3.0. The XenDesktop SDK is installed automatically when the controller or Studio gets installed. Make sure that you create a separate administrator group to run the SDKs with limited special permissions on the Delivery Controller. Don't use the local administrator group as this group has too many privileges and thus is a security risk. The XenDesktop SDK is actually what is underneath the hood of XenDesktop, driving much of the activity. You can view the nuts and bolts of the XenDesktop SDK under the **PowerShell** tab in XenDesktop Studio.

These APIs or PowerShell snap-ins allow you or your partners to perform the same tasks that Citrix Studio would perform from a command-line interface or a script. The XenDesktop SDK is only compatible with XenDesktop 5 and above. Snap-ins that end in `.V1` are XenDesktop 7.x snap-ins, while snap-ins that end in `.V2` are XenDesktop 5 snap-ins. Both `.V1` and `.V2` snap-ins work with XenDesktop 7.x.

In this chapter, we will cover the following topics:

- Microsoft Windows PowerShell
- Using the XenDesktop SDK
- Creating an SDK script

- Troubleshooting using the PowerShell SDK
- The Citrix Ready program

Microsoft Windows PowerShell

PowerShell is Microsoft's task automation framework that has been designed to take the place of the command prompt in legacy Windows systems. PowerShell v3.0 is the latest version and gets installed automatically when you install the Windows operating system, such as Windows Server 2012 R2. PowerShell uses what is called a snap-in architecture that provides a .NET connection into the Windows machine. Snap-ins consist of cmdlets, which provide the common scripting environment for administrators and applications. Cmdlets are not executable, but are rather runtime modules.

You can launch PowerShell by navigating to the Windows **Start** menu and entering PowerShell in the execution window. In Windows Server 2012, you can launch PowerShell from the desktop menu bar directly, as seen in the following screenshot:

PowerShell snap-ins and cmdlets for XenDesktop®

PowerShell scripts use **cmdlets** (**command-lets**). The XenDesktop SDK contains over 300 cmdlets for PowerShell, which allow you to perform more actions with XenDesktop than you can perform using the Studio **User Interface** (**UI**). Cmdlets can be used for configuration or monitoring and return a Microsoft .NET Framework object. Cmdlets can accept parameters and variables. The XenDesktop SDK cmdlets can be executed from a remote environment. You can create PowerShell scripts using XenDesktop cmdlets to manage and control XenDesktop. Cmdlets are individual API calls and are made available through snap-ins. Each snap-in loads a set of cmdlets.

The XenDesktop PowerShell SDK follows the standard verb-item syntax for PowerShell. Some of the common verbs used in the SDK are as follows:

- `Get-`
- `Set-`
- `Add-`
- `Remove-`

- `Rename-`
- `Update-`
- `Reset-`

For example, the `Get-BrokerSite` cmdlet contains the verb `Get-` with the item `BrokerSite`. This cmdlet will return all the information about the Site's broker. It is a good idea to use parameters with the cmdlet to filter or reduce the amount of information that is returned. This is helpful for large deployments.

To execute cmdlets remotely, you will need to include the `-AdminAddress $ddcAddress` parameter, where `ddcAddress` is the controller in the `ddc.<domain.com>:80` form.

A complete reference to all of the XenDesktop PowerShell cmdlets can be found at `http://support.citrix.com/proddocs/topic/xendesktop-71/cds-sdk-cmdlet-help.html`.

The following diagram is a logical representation of how PowerShell snap-ins and cmdlets are used to interface directly with XenDesktop:

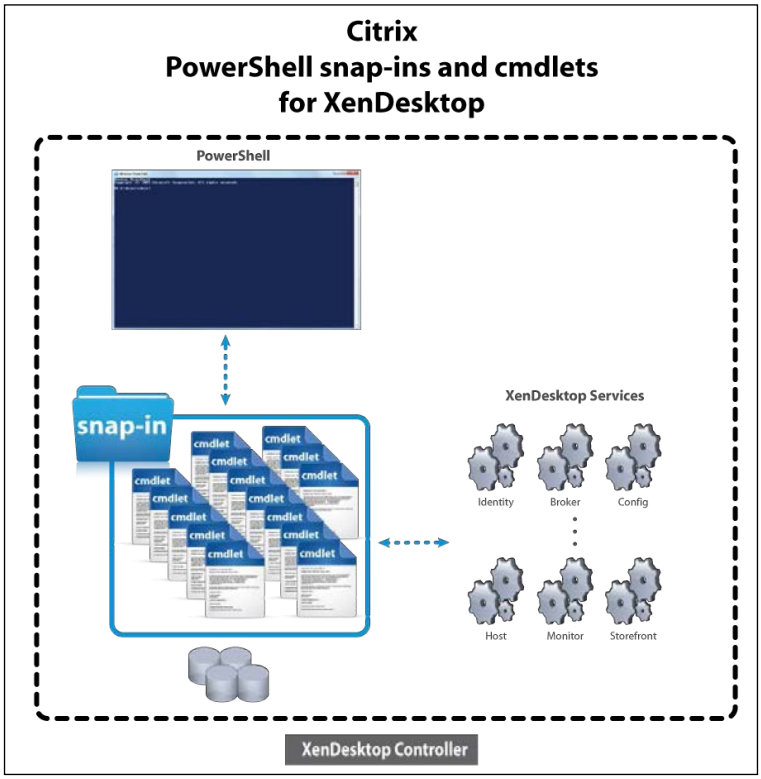

The following table illustrates the snap-ins that are available in the XenDesktop SDK:

Snap-in	Description	Use case
`Citrix.AdIdentity.AdminV2`	This is the Active Directory identity service	This provides administrative cmdlets for the Active Directory identity service
`Citrix.AppV.Admin.V1`	This is the App-V administrative snap-in	This provides cmdlets for App-V applications
`Citrix.Broker.Admin.V2`	This is the broker administrative snap-in	This provides cmdlets for many key functions in XenDesktop by communicating directly with the XML broker for things that range from publishing apps and desktops to creating delivery groups and policies
`Citrix.Configuration.Admin.V2`	This is the configuration administrative snap-in	This provides cmdlets for configuration services in the data store
`Citrix.ConfigurationLogging.Admin.V1`	This is the configuration logging administrative snap-in	This provides cmdlets for configuring database logging
`Citrix.DelegatedAdmin.Admin.V1`	This is the delegated administration service snap-in	This provides cmdlets for administrative roles
`Citrix.EnvTest.Admin.V1`	This is the Citrix environment test service snap-in	This provides cmdlets with tools to test XenDesktop, usually used for installation and configuration purposes
`Citrix.Host.Admin.V2`	This is the host service snap-in	This provides cmdlets for using Hypervisors, networking, and storage, which are consumed by Machine Creation Services to create virtual machines
`Citrix.MachineCreation.Admin.V2`	This is the Machine Creation Services snap-in	This provides cmdlets for Machine Creation Services used to create and manage virtual machines
`Citrix.Monitor.Admin.V1`	This is the monitor service snap-in	This provides cmdlets for monitoring
`Citrix.Storefront.Admin.V1`	This is the StoreFront service snap-in	This provides cmdlets for use with StoreFront

Using the XenDesktop® SDK

As mentioned, the XenDesktop SDK contains PowerShell snap-ins for XenDesktop that are installed automatically when the controller or Studio are installed.

To run a XenDesktop cmdlet, perform the following steps:

1. Log in to the controller with an administrator account.
2. There are two ways to launch PowerShell, as follows:
    - From the Windows start menu, navigate to **Start** | **PowerShell**
    - From Citrix Studio, navigate to **PowerShell** | **Launch PowerShell**
3. In the `PS c:\>` prompt, type `Set-ExecutionPolicy Unrestricted` and `Answer "Y" to the prompt.`
4. Load individual PowerShell snap-ins; for example, `Add-PSSnapin Citrix.AdIdentity.Admin.V2`.
5. Load all Citrix PowerShell snap-ins using the `Add-PSSnapin` alias, `asnp Citrix.*`, as shown in the following screenshot:

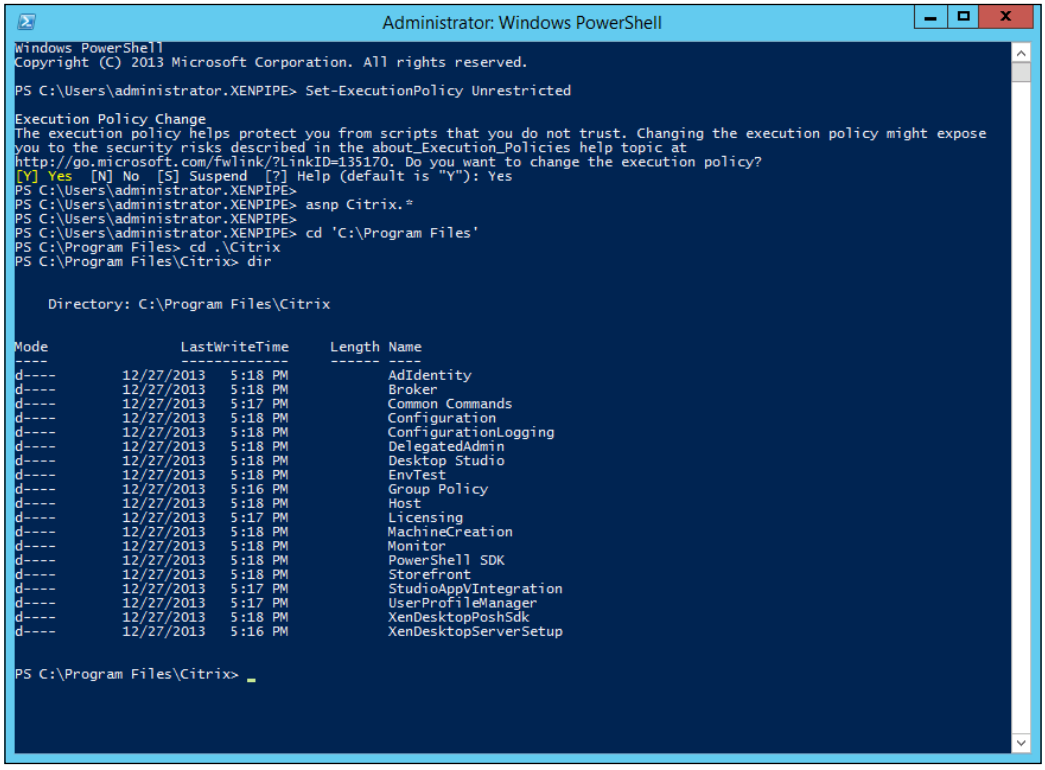

The SDK snap-ins are located on the Delivery Controller in the `\%Program Files%\Citrix` directory.

Creating an SDK script

After working with the XenDesktop SDK through a PowerShell window, you might want to build a script to execute the cmdlets. Scripts are more powerful than the command-line interface because you can have multiple cmdlets in a script and can execute them by running the script. If you are a third-party vendor integrating with XenDesktop, scripts are very important for you.

You can create scripts by reverse engineering; that is, using the ones that are used by XenDesktop.

To create scripts, perform the following steps:

1. Open the Citrix Studio console.
2. Perform the task that you want to replicate in a script in the Studio UI.
3. Retrieve the log of the SDK operations that was made by Studio to perform the task.
4. Reverse engineer the commands and understand the contents.
5. Create your own script using the log as an inspiration.

You can use variables in scripts. You can find out more about PowerShell scripting on the Microsoft website at `http://technet.microsoft.com/en-us/scriptcenter`.

When writing PowerShell scripts, you will want to use Windows Notepad or any other editor. You might also find some of the **Integrated Scripting Environment** (**ISE**) tools that are available out there useful. ISEs provide an advanced ability to create and run scripts more easily.

ISEs provide multiline editing, syntax coloring, tab completion, and variable and console panes. ISEs provide live debugging with breakpoints and step-in/step-out tools. Some ISEs provide IntelliSense, which provides context-specific coding, disambiguation from different cmdlets, and the reduction of types and syntax errors.

The Windows PowerShell ISE is free. Other ISEs include Jive PowerGUI and Idera PowerShell+. Sapien offers a paid ISE tool.

Troubleshooting using the XD PowerShell SDK

There are some common cmdlets that can be used for troubleshooting. The most useful cmdlet for troubleshooting a user's desktop is the `Get-BrokerDesktop` cmdlet. This shows the connection information for the desktop, such as connection time, deregistration time, and any errors that might have occurred.

Useful desktop cmdlets

To find desktops that are hung or not responding because they haven't been shut down after use, you can use the following command:

```
Get-BrokerDesktop -PowerActionPending $false
-PowerState On -SummaryState Available -WillShutdownAfterUse $true
```

To find and set the desktop group's settings, you can use the `Get-BrokerDesktopGroup` and `Set-BrokerDesktopGroup` cmdlets.

To change the protocol from ICA to RDP, you can use the following command:

The highlighted part of the code will change as per each administrator's requirement.

```
Set-BrokerDesktopGroup -Name $desktopGroupName -ProtocolPriority
RDP=$true
```

To create and configure a desktop group, you can use the following command lines:

```
New-BrokerDesktopGroup -DesktopKind $desktopKind -Name $desktopGroupName
-AdminAddress $ddcAddress

Add-BrokerMachinesToDesktopGroup -Catalog $catalogName -Count
$count -DesktopGroup $desktopGroupName -AdminAddress $ddcAddress

New-BrokerUser -Name '<domain>\<Username>'

New-BrokerEntitlementPolicyRule -Name $entitlementPolicyName -
DesktopGroupUid $desktopGroup.Uid -IncludedUsers $users

New-BrokerAccessPolicyRule -Name $accessPolicyName -
IncludedDesktopGroupFilterEnabled $true -IncludedDesktopGroups
$desktopGroupName

New-BrokerPowerTimeScheme -Name $powerTimeSchemeName -DisplayName
$powerTimeSchemeDispName -DesktopGroupUid $desktopGroup.
Uid -DaysOfWeek Weekdays
```

To create a lot of users at one time, you can use the following command line:

```
ForEach ($user In $users) { New-BrokerUser -Name $user -AdminAddress
$ddcAddress }
```

In the preceding command line, `$users` is an array.

To remove a desktop and delete a desktop group, you can use the following command line:

```
Remove-BrokerMachine -MachineName $machineSamName -DesktopGroup
$desktopGroupName
```

To remove all the desktops from a desktop group, you can use the following command lines:

```
$desktops = Get-BrokerDesktop -DesktopGroupName $desktopGroupName
$machineUids = $desktops | ForEach {$_.MachineUid}
Remove-BrokerMachine -InputObject $machineUids -DesktopGroupName
$desktopGroupName
Remove-BrokerEntitlementPolicyRule -Name $entitlementPolicyName
Remove-BrokerAccessPolicyRule -Name $accessPolicyName
Remove-BrokerPowerTimeScheme -Name $powerTimeSchemeName
Remove-BrokerDesktopGroup -Name $desktopGroupName
```

To monitor and manage desktops, use the following command line:

```
Get-BrokerDesktop | Select MachineName, PowerState, RegistrationState,
SummaryState
```

To power on a desktop, use the following command line:

```
New-BrokerHostingPowerAction -MachineName "<domain>\<machine>"
-Action TurnOn
```

To check a desktop's vital signs, you can use the following command line:

```
Get-WMIObject win32_service -ComputerName $name | Select Name, State
```

Useful controller cmdlets

There are some common cmdlets that can be used to discover the status of the controller and its services. These cmdlets are as follows:

- `Get-AcctServiceStatus`: This returns the status of the AD identity service

- `Get-BrokerServiceStatus`: This returns the status of the broker service
- `Get-ConfigServiceStatus`: This returns the status of the configuration service
- `Get-HypServiceStatus`: This returns the status of the host service
- `Get-ProvServiceStatus`: This returns the status of the Machine Creation Service
- `Get-PvsVmServiceStatus`: This returns the status of the Machine Identity Service

An easier way to check the status would be to get an object for each of the services and to pipe them to the `Test` service using the following command line:

```
Get-ConfigRegisteredServiceInstance | Test-ConfigServiceInstance
Availability
```

Some additional cmdlets that are useful for troubleshooting controllers are `Get-BrokerSite` and `Set-BrokerSite`. For example, if you want to use a secure ICA, you could change it for a session by issuing the following command:

```
Set-BrokerSite -SecureIcaRequired $true
```

To clear all the service instance registrations except for the licensing service, you can use the following command:

```
$instance | Unregister-ConfigRegisteredServiceinstance
```

To re-register the DDCs to the Site, you can use the following command lines:

```
Get-ServiceServiceInstance -AdminAddress $ddcAddress |
Register-ConfigServiceInstance
Get-ConfigRegisteredServiceInstance -AdminAddress $ddcAddress |
Reset-ServiceServiceGroupMembership
```

To join a controller to an existing Site, you can use the following command lines:

```
Get-ItemProperty -Path 'HKLM:\SOFTWARE\Citrix\DesktopServer\DataStore\
Connections\Controller' -Name ConnectionString
Get-ServiceDBSchema -DatabaseName $dbName -ScriptType Instance
Set-ServiceDBConnection -DBConnection $cs
Get-ServiceServiceInstance | Register-ConfigServiceInstance
Get-ConfigRegisteredServiceInstance | Reset-ServiceServiceGroupMembership
```

Site debugging tools

Site debugging tools are the Citrix tools that use the XenDesktop PowerShell SDK. They can be found in the Citrix Diagnostics Toolkit. Some of the diagnostic tools use the XenDesktop SDK to perform diagnostics on the XenDesktop Site and present them in a simple GUI. SiteDiag is another tool that is very helpful.

Citrix Diagnostic Toolkit for x64 can be found at `http://support.citrix.com/article/CTX135075`.

Citrix Diagnostic Toolkit for x86 can be found at `http://support.citrix.com/article/CTX134966`.

SiteDiag can be found at `http://desktopsandapps.com/2013/04/23/73/`.

Example XenDesktop SDK scripts can be found on the Internet. An example can be found at `http://blogs.citrix.com/2012/12/05/xendesktop-powershell-sdk-script-examples-part-1/`.

Citrix Ready®

The **Citrix Ready** program is tailored for the organizations that have demonstrated product and solution compatibility such that their technology is compatible with Citrix products and their product will complement your investment in Citrix technology. Having said that, there are a number of Citrix Ready validated partners who have integrated their products with XenDesktop using the SDK.

A list of these partners can be found on the Citrix Ready online catalog located at `www.citrixready.com`. You can narrow down your search by using the filters in the left-hand side navigation menu.

Summary

Whether you are just looking to use snap-ins and cmdlets for troubleshooting or you are a third-party vendor doing some major integration through scripting and automation, the XenDesktop SDK is powerful and will meet your needs. The XenDesktop SDK is a powerful way to plug in to the different XenDesktop services. In this chapter, you learned about the XenDesktop SDK. In the next chapter, you will learn about the Citrix Receiver.

12
Working with Citrix Receiver™ and Plugins

Throughout this book, we have mainly focused on the server side of XenDesktop. Now, we will take a moment to look at the other side of the equation: the client. Arguably, this chapter should come first, but having navigated your way through the complex server-side components, this final piece will be a downhill run for you.

Citrix Receiver provides users with secure, self-service access to virtual desktops and applications. A user can also use Receiver to access apps through the stores managed by Citrix StoreFront and the legacy web pages managed by the web interface.

Citrix Receiver is device agnostic, so you can conceivably run a Windows 8 desktop from XenDesktop using Receiver on a Linux, Android, or Apple device without having to install the Windows 8 operating system on it. What is more likely is that you will be running Windows 8 via the Receiver on a thin client to drive down the costs. As we move forward in the mobility revolution, Receiver will be a critical piece of desktop and application virtualization. Because Receiver is device agnostic, it runs on Windows to Mac OS and everything in between. In this chapter, we will cover the following topics:

- What is Citrix Receiver and what are plugins?
    - Online plugins
    - Offline plugins
    - CloudBridge plugins
- Receiver for Windows
- Receiver for Apple
    - Mac
    - iOS

- Receiver for Android
- Receiver for Linux
- Receiver for Java
- Receiver for HTML5
- Receiver for Blackberry
- Receiver for Chromebook

Citrix Receiver is free and can be downloaded from `http://receiver.citrix.com/`.

Understanding Receiver

Citrix Receiver is an essential part for users to receive desktops and applications from XenDesktop because it is the other end of the connection. Remember when you were a kid and you would take two tin cans and put a piece of long string between them to talk to your friend at the other end? Think of XenDesktop as one side (server), Receiver as the other side (client), and the `ICA/HDX` protocol as the string in between them. Receiver is a critical component in the communication link between the client and the server. This setup is illustrated in the following figure:

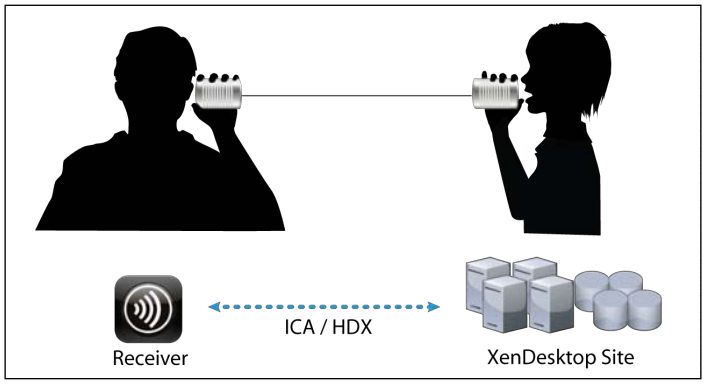

Citrix Receiver allows the virtual desktop to run on the client device. Technically, Citrix Receiver just displays the screen pixels on the client device, while the actual virtual desktop is running back in the data center on a server. The simplest way to describe the Receiver to ICA to XenDesktop connection is that Receiver is just a display device with a very long **Keyboard, Video, and Mouse** (**KVM**) cable and the `ICA/HDX` protocol connects it to the XenDesktop server.

Apart from simply connecting a client to their desktops and applications, Receiver performs many other functions as well. The many functions that Receiver performs include the following:

- Authentication
- Secure communications
- Performance enhancements for video
- Single sign-on
- WAN optimization
- Session reliability
- Smooth roaming between desktops
- Automatic StoreFront push and connection
- Multimonitor support and the redirection of resources (audio, video, USB, storage, and printers) from a virtual desktop to a client device

You can read more about Receiver at `http://support.citrix.com/proddocs/topic/receiver/rec-receiver-and-plugins.html`.

Changing the Receiver settings

The Receiver settings can be pushed down from the server or can be accessed and changed directly from the client by accessing the Receiver toolbar at the top of the virtual desktop window.

Pushing the Receiver settings from the server

Most of the settings that you require as an administrator are pushed down to the clients through policies, without the client having to do anything. You can set this up by automatically pushing Receiver down to the client by specifying the StoreFront URL to which the user will connect. The steps to perform this task are as follows:

1. Launch Studio.
2. Select the **Delivery Groups** node.
3. Select the desktop delivery group you wish to edit.
4. Click on **Edit Delivery Group**.

5. In the **StoreFront** section, enter the storefront URLs to push to the user.
6. Click on **OK**, as shown in the following screenshot:

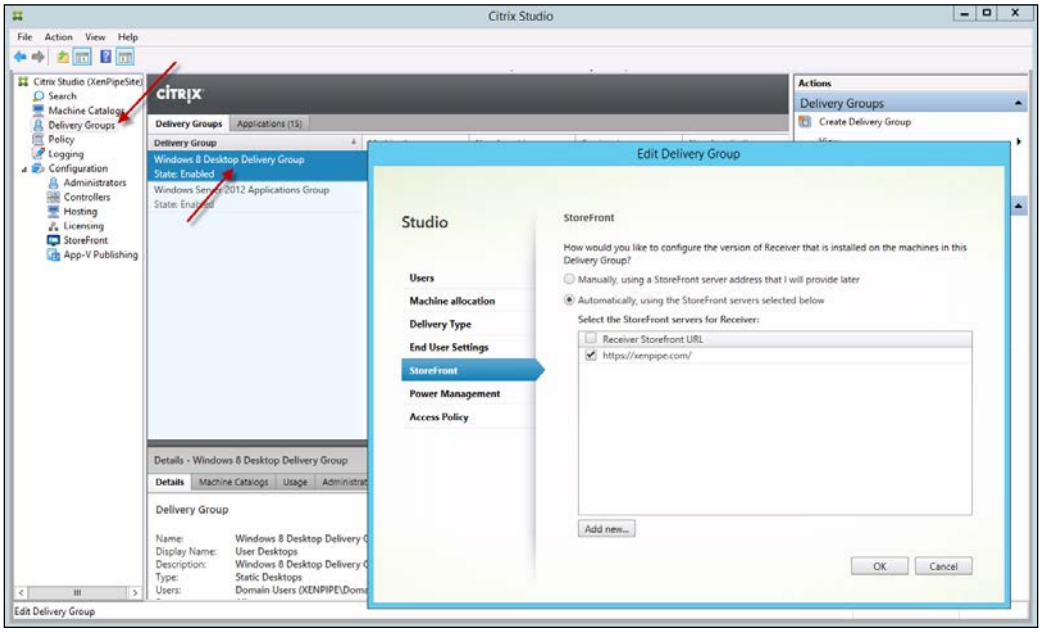

Changing the Receiver settings from the client's desktop

You can change a limited set of settings on Receiver on the client desktop by performing the following steps:

1. Launch the end user's virtual desktop by opening Receiver on the end user's client device and selecting **User Desktops,** as shown in the following screenshot:

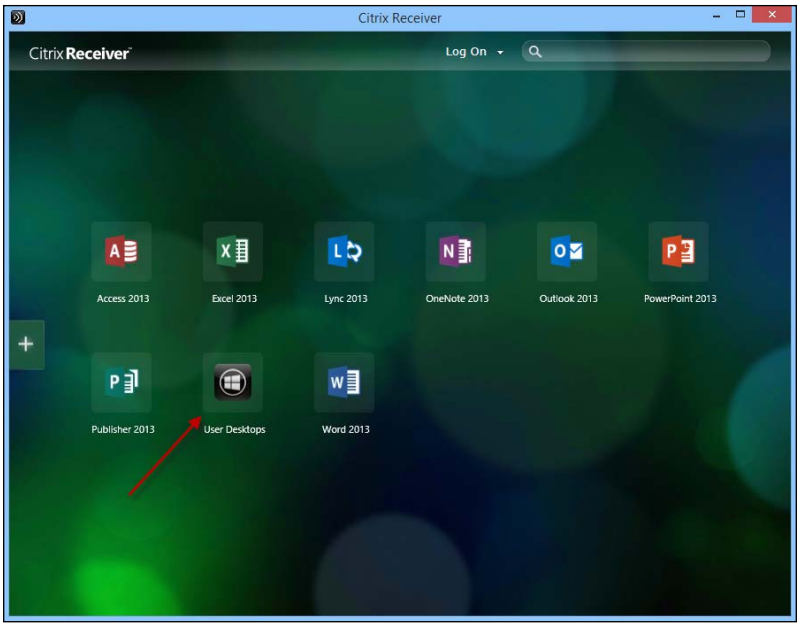

2. When the desktop appears, you will see a toolbar appear at the top of the screen.
3. Select **Preferences** to modify the Receiver settings. Make the required changes and click on **OK**, as shown in the following screenshot:

Using plugins

Citrix Receiver can be used with plugins to provide advanced features and capabilities. Several plugins already exist and can be downloaded from the Citrix downloads website.

The Citrix downloads website is located at `http://www.citrix.com/downloads.html`.

The online plugin

The online plugin is the most notable as it is required for users to access virtual desktops and applications. The origin of the online plugin goes all the way back to the initial founding of Citrix with the development of MetaFrame and Citrix Presentation Server. Later, when Citrix changed the name of the product to XenApp, they also changed the name of the Program Neighborhood Agent to the online plugin. The online plugin was used by XenApp to receive applications published in a XenApp Services Site, which could then be integrated directly with the user's desktops to appear in their Windows **Start** menu.

The functionality developed for XenApp has been carried forward to work with XenDesktop. The online plugin includes the Desktop Viewer software, which is the client-side software that supports the running of virtual desktops from XenDesktop.

The benefit of Receiver is that when you install it, it automatically installs the online plugin. When an update is available, Receiver can be configured to automatically download and install the update. When Receiver is uninstalled, the online plugin is also removed.

The following sections cover the important functions that the online plugin performs.

Using workspace control

Workspace control is enabled by default for your users' virtual desktops. Workspace control, also known as smooth roaming, allows the user to roam from one device to another while keeping the virtual desktop up and running. The user essentially disconnects from the virtual desktop running on one device and reconnects to the same virtual desktop from a different device. For example, running a virtual desktop on a Windows laptop and then switching over to an Apple iPad. If the user connects to the second device without disconnecting from the first, the session on the first device is automatically disconnected.

Changing the resolution of the virtual desktop

The Desktop Viewer can choose the best resolution to fit the entire virtual desktop in the client-side window; it can scale the virtual desktop to fit a specific window size or display the virtual desktop as its actual size. It can also work with multiple monitors by dragging the desktop window across all the monitors and then clicking on the **Maximize** button. The window will change resolution to match the monitor's resolution. The various actions that can be performed with the Desktop Viewer are as follows:

- You can have Desktop Viewer automatically select a resolution; navigate to **Preferences** | **Display** | **Best resolution** on the Receiver toolbar
- You can scale a virtual desktop into a window; navigate to **Preferences** | **Display** | **Scale to fit** on the Receiver toolbar
- You can display a virtual desktop as its actual size; navigate to **Preferences** | **Display** | **Actual size** on the Receiver toolbar

Using actual size will result in scroll bars appearing on the sides of the virtual desktop window as the window size is typically smaller than the actual desktop size. This can be seen in the preceding screenshot.

Moving the toolbar

You can move the toolbar to another location.

To move the toolbar, click and hold the grip of the toolbar and move it to a different location as shown in the following screenshot:

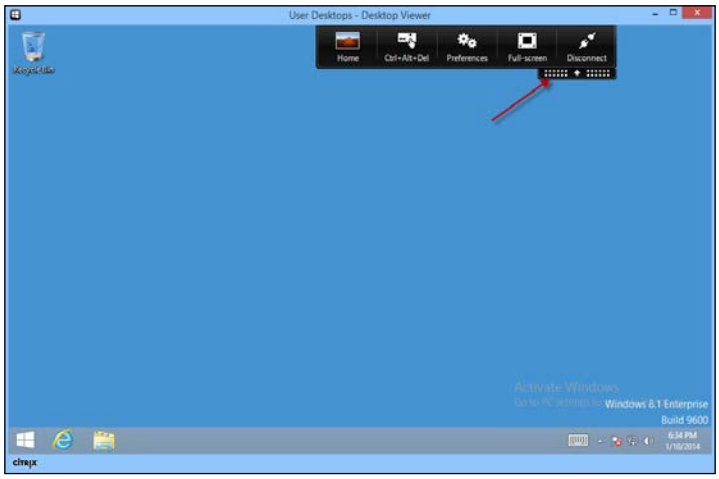

Controlling local file access

You can control how a virtual desktop accesses files on the local user device.

To change the local file access options, navigate to **Preferences | File Access** on the toolbar and choose one of the following options:

- **Read and write**: This allows the virtual desktop to read and write local files
- **Read only**: This allows the virtual desktop to read but not write files locally
- **No access**: This does not allow the virtual desktop to read or write local files
- **Ask me each time**: This prompts the user each time the files need to be accessed

An example of these options is shown in the following screenshot:

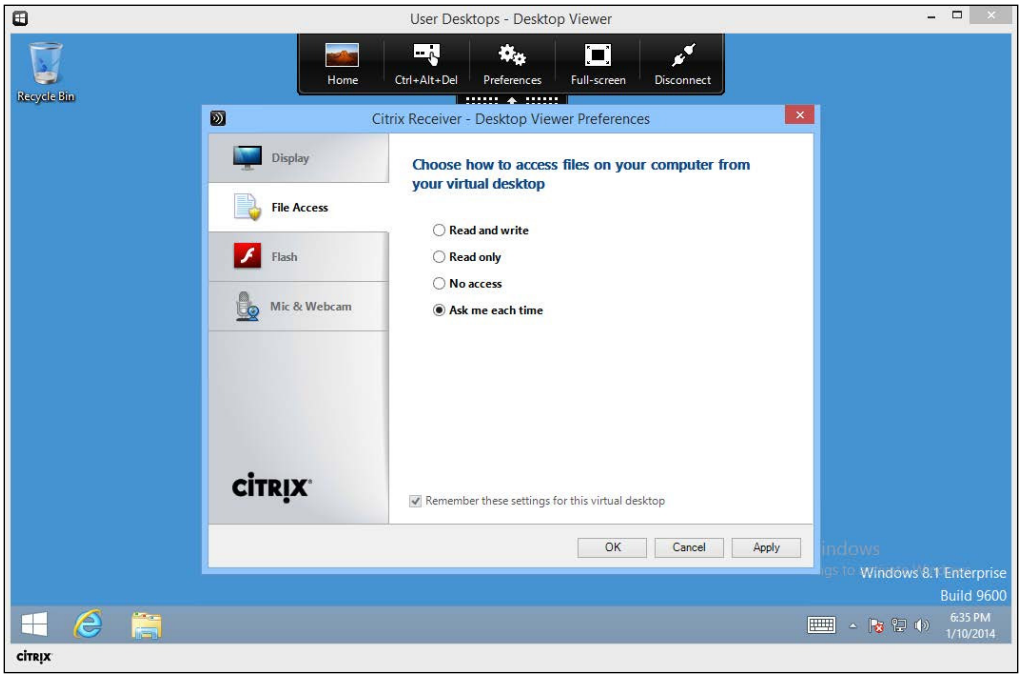

Accessing devices

You can access any local device through the virtual desktop. Make sure that you have enabled the policy settings for devices. We covered this in the *Accessing policies* section in *Chapter 5*, *Managing Policies*.

To access local devices, navigate to **Preferences** | **Devices** on the toolbar as shown in the following screenshot:

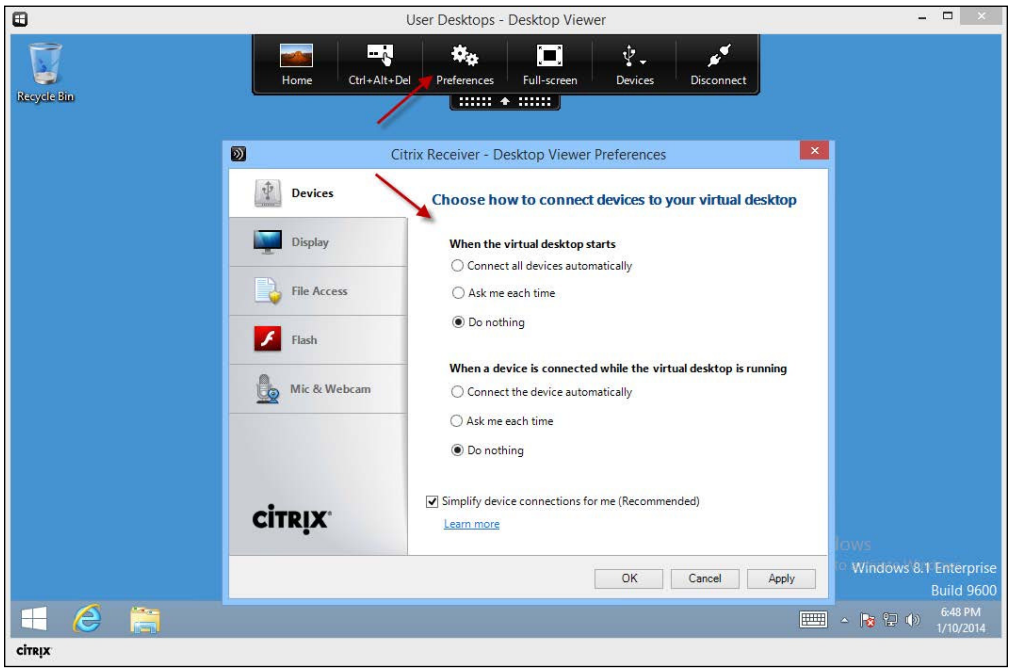

Accessing USB devices

If you have enabled USB support as we covered in the *Configuring policy settings* section in *Chapter 5*, *Managing Policies*, then you are presented with a list of USB devices each time a virtual desktop starts. You can customize how to use these USB devices. Make sure that you have the policies configured to allow USB access and have set the previous setting in **Preferences**.

To access USB devices, navigate to **Devices** and select the USB device on the toolbar, as shown in the following screenshot:

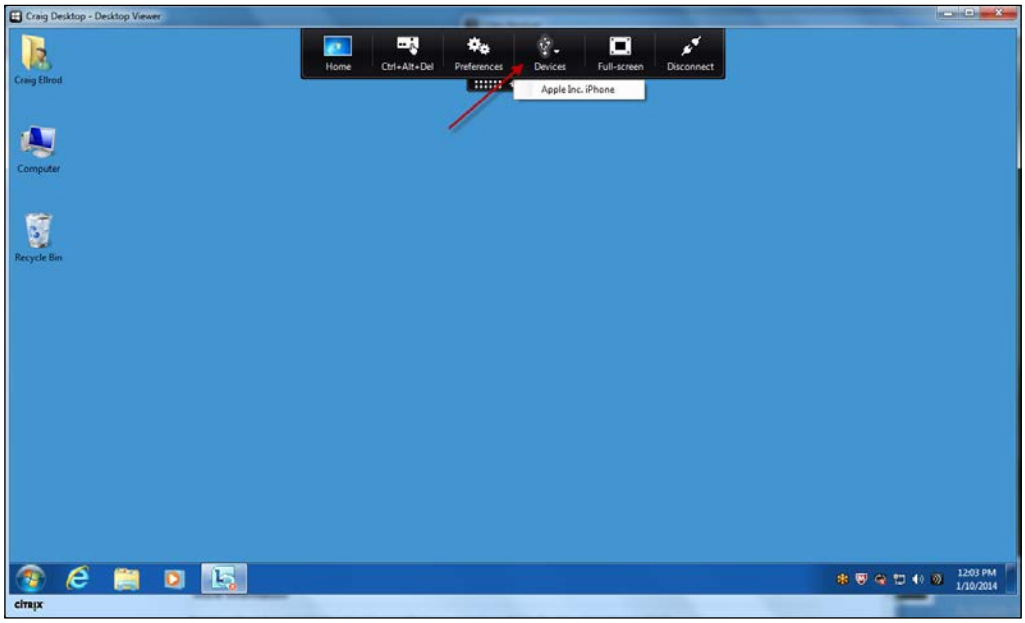

Accessing local microphones and webcams

You can change the way microphones and webcams are used.

To change the microphone and webcam settings, navigate to **Preferences | Mic & Webcam** on the toolbar and select one of the following options:

- **Use my microphone and webcam**: This allows the microphone or webcam to be used with the virtual desktop
- **Don't use my microphone or webcam**: This does not allow the microphone or webcam to be used
- **Ask me each time**: This prompts the user when the virtual desktop starts for permission to access the microphone or webcam

An example of the aforementioned tasks along with the options to change the microphone and webcam settings is shown in the following screenshot:

Redirecting Flash to a local device

You can use Flash redirection to play Adobe Flash content using local device resources instead of using the virtual desktop resources on the server. This results in better performance and higher quality images.

To redirect Flash, navigate to **Preferences** | **Flash** on the toolbar as shown in the following screenshot:

Switching between virtual desktops

If you configure users to have more than one virtual desktop, you can switch between these desktops using the toolbar.

To switch between virtual desktops, perform the following steps:

1. On the toolbar, select the name of the desktop you want to switch to.
2. To switch back, select **Home**.

Logging off virtual desktops

In the event that your users need to log off from virtual desktop, there is a way to do this. You can log off from a virtual desktop.

To log off from a virtual desktop, perform the following steps:

1. On the toolbar, click on **Ctrl+Alt+Del**.
2. Select the action you want to take, as shown in the following screenshot:

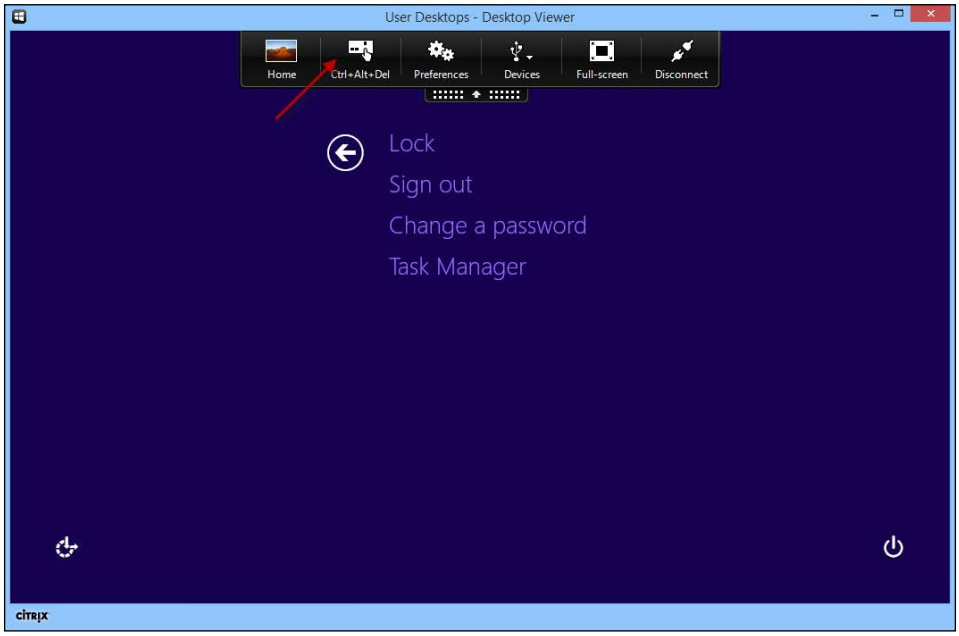

Disconnecting from virtual desktops

If you need to disconnect from a virtual desktop without logging off, you can do that, and it will keep your applications intact until you reconnect. You can disconnect from a virtual desktop.

To disconnect from a virtual desktop, click on **Disconnect** on the toolbar as shown in the following screenshot:

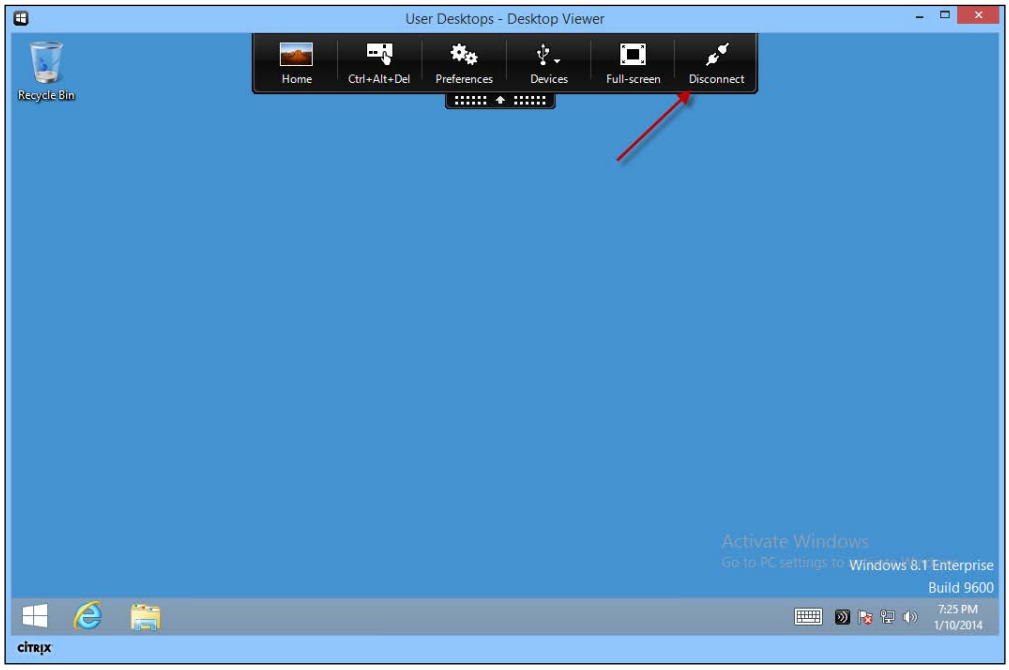

Restarting a virtual desktop

Just like restarting a laptop or PC, you can restart a virtual desktop. To restart a virtual desktop, perform the following step:

1. Right-click on the Windows **Start** menu and select **Restart,** as shown in the following screenshot:

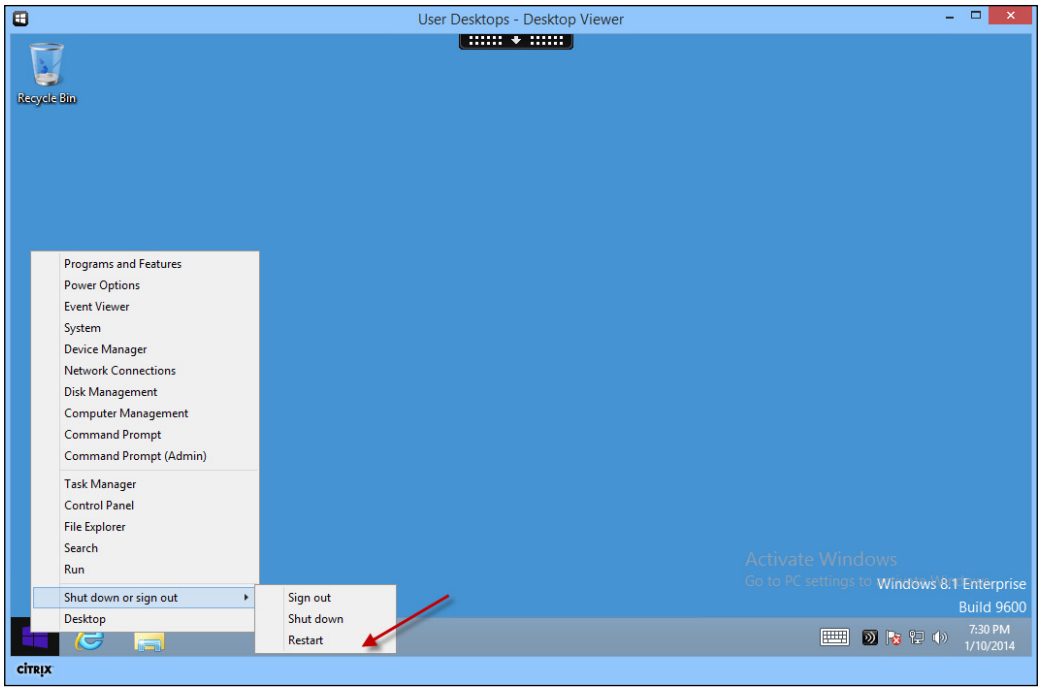

Using Desktop Lock

Desktop Lock is a feature that locks the user out of their local device when a virtual desktop is running. This gives the user the impression that the virtual desktop is the local desktop. This is really designed for low-end Windows devices and thin clients. When using Desktop Lock, the local device's **Start** menu is disabled and the Windows **Task Manager** is also replaced with the virtual desktop instance.

The standby and lock actions on the local computer do not affect the virtual desktop, but its behavior in response to a local logoff is configurable by the IT department.

It is not a good idea to use screensavers that automatically lock virtual desktops. If a screensaver has been activated on a virtual desktop, it might give the impression that the local computer is locked, which could pose a security risk. It is also not a good idea to hibernate virtual desktops.

The effects of the shutdown and lock actions on a virtual desktop with Desktop Lock are configurable and are summarized as follows:

Action	Description
Logoff	This option logs you off from the virtual desktop and the local computer and returns the user to the local login dialog box
Shutdown	This option shuts down the virtual desktop and local computer
Restart	This option restarts the virtual desktop and local computer and returns the user to the local login
Standby	This option puts the virtual desktop into standby mode
Lock computer	This option locks the virtual desktop and returns the user to the local computer's locked screen

For more information on installing Desktop Lock, refer to the documentation at `http://support.citrix.com/proddocs/topic/xendesktop-7/cds-desktop-lock-install.html`.

Printing in virtual desktops

You can view the printer and fax devices under the **Devices and Printers** section in the **Control Panel** of the virtual desktop. You can change your default printer, but when you disconnect the default printer, it reconnects to the local printer. We covered printing in *Chapter 6*, *Managing Printing*, and it is relevant to Receiver because printing policies are sent from the controller to the client's desktop via Receiver.

Understanding the keyboard input

By default, when you use a virtual desktop, all the keyboard inputs are directed to the virtual desktop with the following exceptions:

Key(s)	Action
Windows logo key + *L*	This directs you to the local computer
Ctrl + *Alt* + *Del*	This directs you to the local computer, unless you are using Desktop Lock
Sticky keys, filter keys, toggle keys (Microsoft)	These direct you to the local computer

Key(s)	Action
Ctrl + *Alt* + *Break*	This displays the Desktop Viewer toolbar buttons in a pop-up window
Ctrl + *Esc* and *Alt* + *Tab*	These direct you according to the selections done by the IT department
Ctrl + *F1*	This is a custom Citrix key combination that sends *Ctrl* + *Alt* + *Del* to the virtual desktop
Shift + *F2*	This is a custom Citrix key combination that switches virtual desktops between the fullscreen and window modes

The offline plugin

The offline plugin originated at a time when users would need to run their applications from XenApp directly on their client machine to harness the power of the client machine. This is known as streaming. Setting up streaming is a huge administrative task because you have to create profiles of the exact machine that you are going to stream the applications to. Citrix streaming is no longer available in XD7, and the offline plugin is not required.

The CloudBridge™ plugin

The CloudBridge plugin is supported on Citrix Receiver 3.0 and above and is administered by Citrix Receiver. It is truly amazing to see WAN optimization in action. Citrix Branch Repeater is now known as Citrix CloudBridge and was formerly known as Citrix WANScaler that Citrix acquired from Orbital Data Corporation in 2006.

For WAN optimization to work, you need two CloudBridge appliances at each end of the communication link so that it can sit in between the packet flow and perform its deduplication and caching operations. CloudBridge can be deployed as a virtual machine locally or in the cloud. What is really interesting is that there is a software client that you can run on your local machine, which takes the place of one of the CloudBridge appliances. So effectively, you can have WAN optimization between your client device and the CloudBridge device in the data center.

The CloudBridge software client is now available as a plugin for Citrix Receiver. I would highly recommend that you use this as it makes life so much more interesting by moving data at lightning speed. The user experience is amazing with CloudBridge.

Download the Citrix CloudBridge plugin from `www.citrix.com/downloads`. Make sure that you install it on your master image so that your users can take advantage of it.

If you use the CloudBridge plugin, make sure that you do the following two things:

- Use the transparent mode on the plugin
- Create an acceleration rule to accelerate all the traffic for your subnet at the headend CloudBridge appliance

Running Receiver on Microsoft Windows

Citrix Receiver runs on Microsoft Windows. This should come as no surprise as Windows has been the deployment model for Citrix ever since its inception as a company.

Citrix Receiver is compatible with the following Windows versions:

- Windows Server (2012 R2, 2012, 2008 R2, 2008)
- Windows desktop (Win 8.1, Win 8, Win 7, Vista, XP)
- Windows 8/RT
- Windows CE
- Windows Phone 8

Running Receiver on Apple

Citrix Receiver runs on Apple products. This goes to show how Citrix is delivering virtual desktops and applications to any device. This can be explained as follows:

- **Mac**: Citrix Receiver runs on Mac and provides the expected benefits. Citrix Receiver for Mac is installed with a `.dmg` file. Receiver for Mac works best when configured with NetScaler Gateway and StoreFront. It can also help simplify deployment by configuring **Email-based account discovery**, making it easier for end users to connect to the XenDesktop Site.
- **iOS**: As the iPhone and iPad have become popular, you can be sure that Citrix Receiver runs on iOS. As of this writing, it is compatible with iOS 7 and the iPhone 5c and 5s. It also works with the previous versions of iOS. Citrix Receiver for iOS can be downloaded from Apple's App Store, and you can download the legacy version on Citrix's download website.

Running Receiver on other devices

As you know, Citrix Receiver is device agnostic, meaning it can run on a wide variety of devices in the marketplace. The following are some examples:

- **Android**: Citrix Receiver runs on Android. The mobility market is heating up with a large part of the market running their devices on the Android operating system. Android is a product of Google and runs on many device types and manufacturers. Citrix Receiver supports Android Versions 2.3.3 or later.
- **Linux**: Citrix Receiver runs on Linux. Many thin-client vendors are using Linux or a customized version of Linux as their base operating system and are then preloading Citrix Receiver so that it connects to the XenDesktop Site with minimal configuration. Citrix Receiver is compatible with Linux kernel Version 2.6.29 or above, with glibc 2.7 or above, GTK 2.12.0 or above, libcap1 or libcap2, and udev support.
- **HTML5**: Citrix Receiver runs on HTML5-compatible web browsers. The virtual desktops and applications are accessed directly through the user's HTML5-compatible web browser. This is another option for device agnosticity because different browsers now support HTML5 such as Internet Explorer, Google Chrome, Mozilla Firefox, and Safari. It is best to use the HTML5 option with NetScaler Gateway as a frontend for SSL communication in front of the XenDesktop Site.
- **BlackBerry**: Citrix Receiver runs on BlackBerry devices including the BlackBerry Playbook. You can download Receiver for BlackBerry from the BlackBerry World app store.
- **Chromebook**: Citrix Receiver runs on Chromebook. Receiver for Chromebook provides access to XenDesktop through NetScaler Gateway.

A list of Citrix Receiver supported client devices and feature sets is located at `http://www.citrix.com/content/dam/citrix/en_us/documents/products-solutions/citrix-receiver-feature-matrix.pdf`.

Summary

Citrix Receiver is a vital piece of the communication equation between the client and server or user device and XenDesktop. Fundamentally, it provides the client-side `ICA/HDX` protocol connection to XenDesktop and has a rich feature set designed to deliver a high definition experience to the user, anywhere on any device.

In this chapter, we covered Citrix Receiver and the plugins that go with it. You learned how to customize user sessions from the client session in the Receiver settings. You also learned about all the different devices that Receiver runs on.

In the next chapter, we will learn about managing and monitoring XenDesktop.

13
Securing XenDesktop®

By itself, XenDesktop has a weak spot because the traffic is not totally secure, but you can make it secure by following the simple guidelines mentioned in this chapter. You must undoubtedly be familiar with SSL, with uses port `443` or HTTPS to encrypt data and to check the message integrity between the client and the server. XenDesktop and XenApp have, for a long time, had a feature called the **Secure Ticket Authority** (**STA**); however, this feature doesn't provide the complete message encryption security that SSL provides. In this chapter, we will discuss securing XenDesktop with SSL.

In this chapter, we will discuss the following topics:

- DMZ architectures
- Securing XenDesktop with NetScaler Gateway
- The STA
- Securing the `ICA/HDX` protocols
- Securing StoreFront
- Securing Receiver
- Securing the controller
- Securing Studio and Director
- Securing the XenDesktop to XenServer communications
- Smart cards

DMZ and DMZ²

The concept of a **Demilitarized Zone** (**DMZ**) in security architectures has been around for a long time. A DMZ essentially provides a separate network in front of a firewall to only allow connections into the web portal in the private network; in our case, the StoreFront server. When you set up a DMZ, you create an additional layer of security, or zone, that hackers have a hard time penetrating, and you only let in the specific connections on the port numbers that need to gain access, such as HTTP on port `80` or HTTPS on port `443`. When you set up a DMZ to secure XenDesktop, you will need to use HTTPS or port `443` and SSL certificates. The best way to do this is to install NetScaler running SSL and a load balancing service in the DMZ connected to StoreFront and the remaining XenDesktop components behind the second firewall. You open port `443` on the outside firewall, and only open the required ports for XenDesktop on the second firewall. Now, for an additional layer of security beyond this, you can install what is called a Double-Hop DMZ, where there are three firewalls and two DMZs. Some organizations that are overly zealous about security may opt for this type of configuration. The following are the network diagrams for Single-Hop DMZ and Double-Hop DMZ. The ports that need to be open on the inside firewall are the TCP ports `80`, `443`, `1494`, and `2598`.

The following diagram is an example of a Single-Hop DMZ architecture:

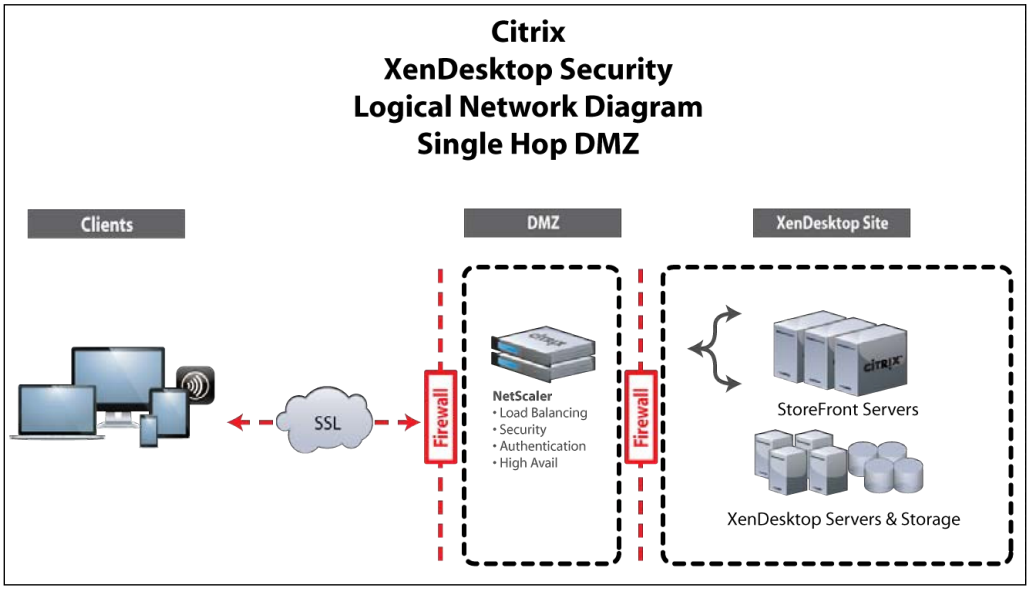

The following diagram is an example of a Double-Hop DMZ architecture:

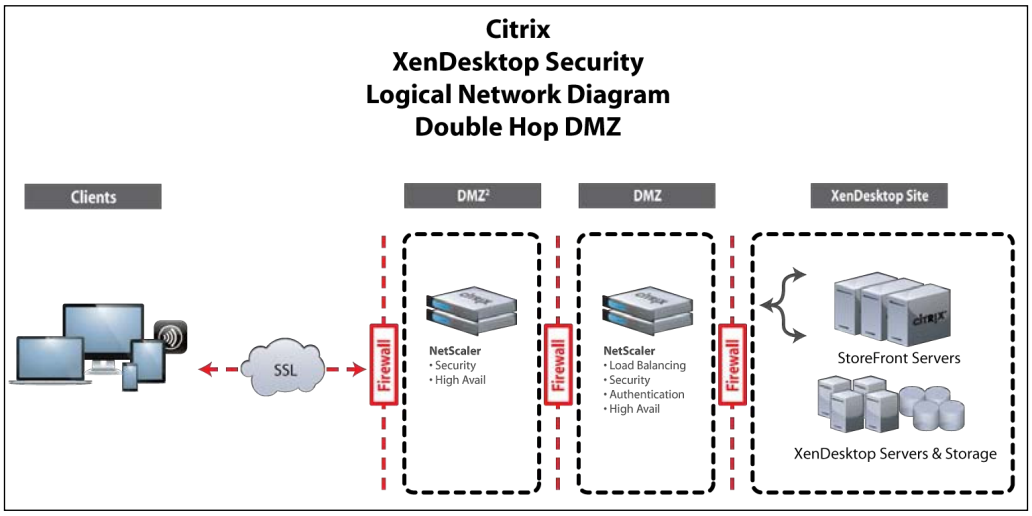

Securing XenDesktop® with NetScaler Gateway™

In this chapter, we will talk about all the ways to configure security around XenDesktop components. The most important security component is NetScaler Gateway, as it provides a completely secure frontend to the XenDesktop Site. If you can only do one thing to secure XenDesktop, make it NetScaler Gateway.

Citrix acquired a company called NetScaler, and the name hasn't changed since. NetScaler is a very powerful frontend device, and it is covered in more detail in *Chapter 10*, *Application Delivery*. No other application Delivery Controller works as well as NetScaler as it is engineered to work very closely with XenDesktop. NetScaler provides a special functionality for the `ICA/HDX` protocol.

Installing NetScaler Gateway involves many steps, so the following table is a checklist to help keep you on track:

The installation task list		
Step	**Description**	**Completed**
1	Import NetScaler VPX into XenServer (or an other Hypervisor)	
2	Configure the initial IP addresses and license	

The installation task list		
Step	**Description**	**Completed**
3	Install the SSL certificate	
4	Create the NetScaler Gateway virtual server	
5	Configure LDAP	
6	Configure the NetScaler to StoreFront connection	
7	Configure the StoreFront to NetScaler connection	
8	Test the connection	

Importing NetScaler VPX™ into XenServer®

The first step is to download NetScaler VPX from the Citrix website and to install or import it into XenServer or whichever Hypervisor you are using. When you download the VPX, there is a license key available, so download that as well. To provide additional security in this deployment, we have built a completely separate physical XenServer to run NetScaler Gateway on and will use a different subnet as this will be plugged in to the DMZ.

To import NetScaler VPX into XenServer, perform the following steps:

1. Launch the XenCenter console.
2. Go to **File** | **Import** and import NetScaler VPX.

Configure the **NetScaler IP management address** (**NSIP**) initially. You may want to put this IP address in the private subnet so that it isn't exposed externally. You may also want to run the setup wizard to assign a subnet IP address in the same subnet as the NSIP.

Installing a NetScaler® license

Once you have NetScaler installed and configured with the NSIP, you need to install a license. There are a couple of steps to get this done. First, you have to generate the license from the `citrix.com` website and then you need to install it.

To obtain the host ID (MAC address) of NetScaler, perform the following steps:

1. Log in to NetScaler from the console screen on the XenCenter console. You can also connect to the NSIP using SSH.

The default username and password are `nsroot` and `nsroot`, respectively.

2. Go to the command shell and type in the `shell` command.
3. Type in the `lmutil lmhostid -ether` command.
4. The MAC address that is returned is what you need to use as the host ID when obtaining a license from the `citrix.com` website, as shown in the following screenshot:

```
 Done
>

CLI session timed out at Thu Jan 30 17:00:17 2014
^D
Bye!

login: nsroot
Password:
Jan 30 18:22:04 <auth.notice> nsvpx1 login: ROOT LOGIN (nsroot) ON ttyv0
Copyright (c) 1992-2008 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
        The Regents of the University of California. All rights reserved.

 Done
> shell
Copyright (c) 1992-2008 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
        The Regents of the University of California. All rights reserved.

root@nsvpx1# lmutil lmhostid -ether
lmutil - Copyright (c) 1989-2007 Macrovision Europe Ltd. and/or Macrovision Corp
oration. All Rights Reserved.
The FLEXnet host ID of this machine is "5eb3ece3fcbe"
root@nsvpx1#
```

To generate a license for NetScaler, perform the following steps:

1. Log in to the `citrix.com` website.
2. Navigate to the license retrieval system.
3. Select the **Citrix NetScaler** license and click on **Continue**.
4. The **Configuration** page is displayed. Enter the following details and click on **Continue**:
    - **Host ID Type**: By default, the host ID type is selected as **Ethernet**; you should not change this.
    - **Host ID**: This refers to the MAC address on the NetScaler appliance retrieved in the previous section. Enter that MAC address here.

- **Quantity/Available**: Enter the quantity of license you would like to allocate; typically, you use the default of **1**.

5. Click on **Continue** and then on **Confirm** to complete the license allocation.
6. Download the license file, as shown in the following screenshot:

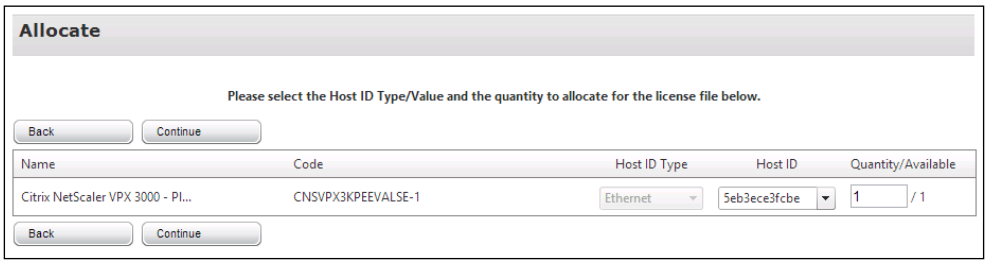

To install the license, perform the following steps:

1. Open a browser and connect to the NSIP.
2. Upon initial configuration, you can upload the license file. You can also install it by navigating to **System configuration** | **License** as shown in the following screenshot:

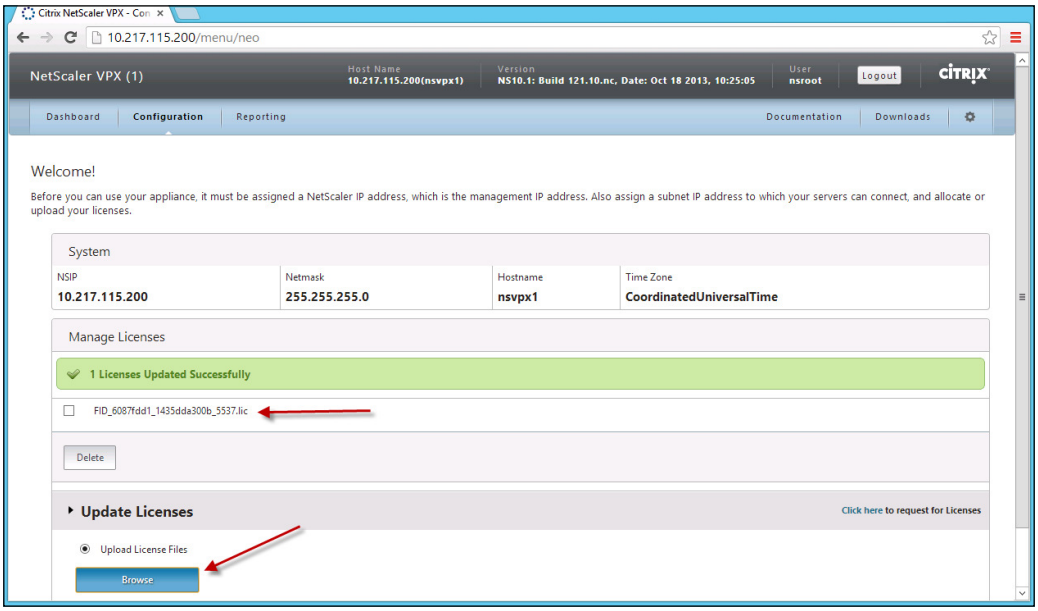

Installing an SSL certificate

You need to use SSL to secure XenDesktop deployment, and you can do this by installing an SSL certificate on NetScaler. There are a couple of ways to get an SSL certificate on NetScaler. NetScaler is capable of creating a self-signed test certificate, which is good for the **Proof of Concepts** (**POCs**). Another more desirable method is to purchase a valid SSL certificate from a public **Certificate Authority** (**CA**); you can get an inexpensive SSL certificate from GoDaddy, for instance, and then install these on NetScaler. In this section, we will use a test certificate.

Using a self-signed certificate will result in the users' browser reporting a security warning, and ultimately your connection to XenDesktop won't work.

For step-by-step instructions on how to generate and install a self-signed SSL certificate, refer to *Appendix C, Creating Self-signed Certificates for NetScaler Gateway™*.

For step-by-step instructions on how to generate and install a valid SSL certificate from a public CA, refer to *Appendix D, Using Public CA-signed SSL Wildcard Certificates on NetScaler Gateway™*.

Creating a NetScaler Gateway™ virtual server

Now, we need to create the public-facing virtual server for users to connect to from their devices. This is also known as the virtual IP or VIP server. This will also be used to frontend the **Virtual Private Network** (**VPN**) connection for the XenDesktop connections.

To create a NetScaler Gateway virtual server, perform the following steps:

1. Log in to NetScaler.
2. Navigate to **NetScaler Gateway**.

3. Right-click on **NetScaler Gateway** to select **Enable feature**, as shown in the following screenshot:

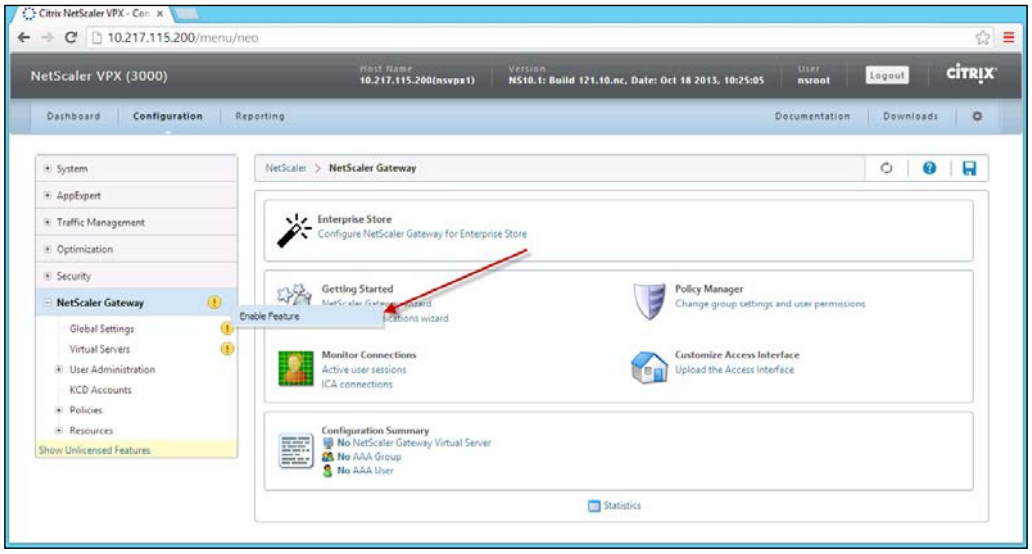

4. On the same screen, launch **NetScaler Gateway wizard**, as shown in the following screenshot:

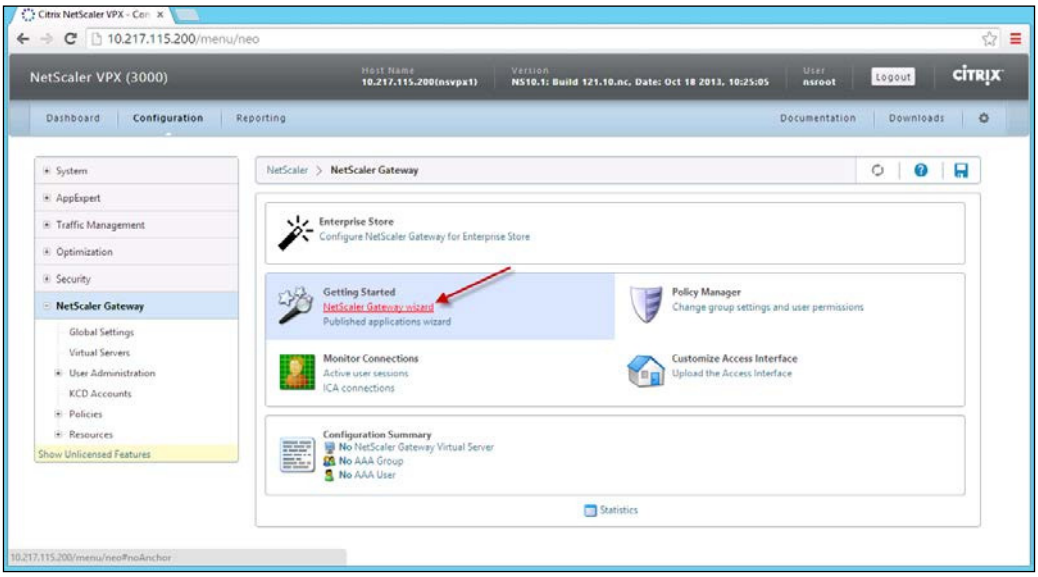

5. Create a virtual server by specifying an externally accessible IP address, as shown in the following screenshot:

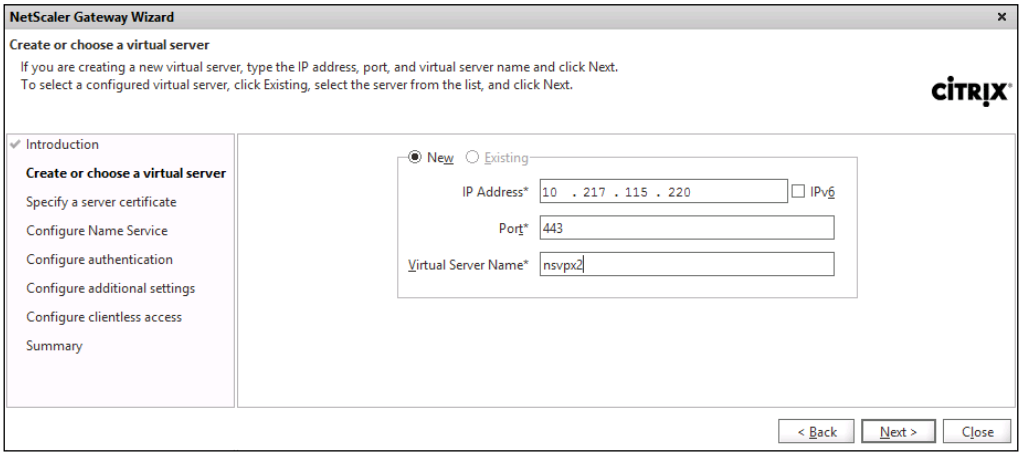

6. Specify the server certificate to be used. We will use an existing test certificate that we created earlier, as shown in the following screenshot:

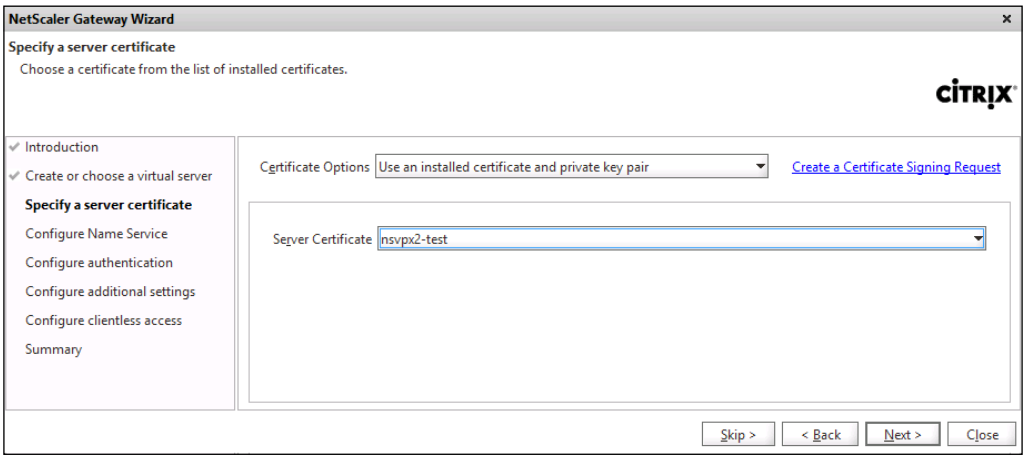

7. Configure the DNS server. Don't use WINS, as shown in the following screenshot:

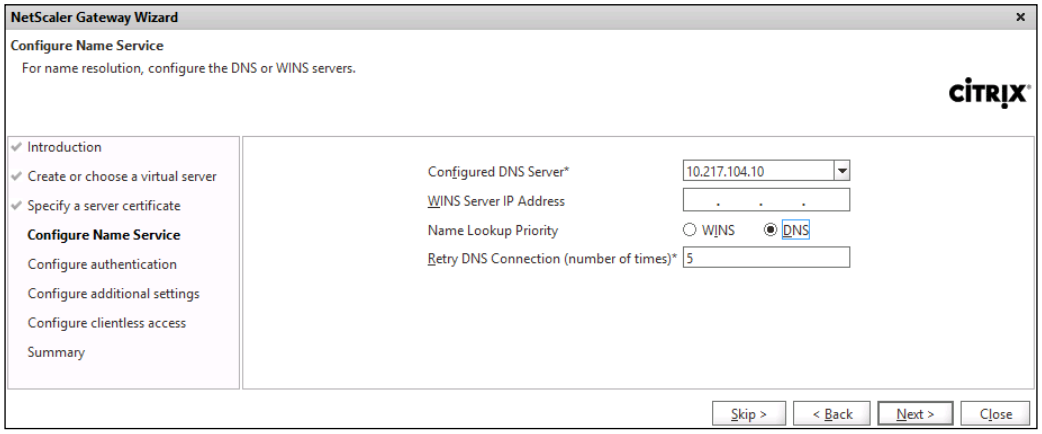

8. In the **Configure authentication** section, select **LDAP** as the authentication type, supply **Connection Settings**, and click on **Retrieve Attributes**, as shown in the following screenshot:

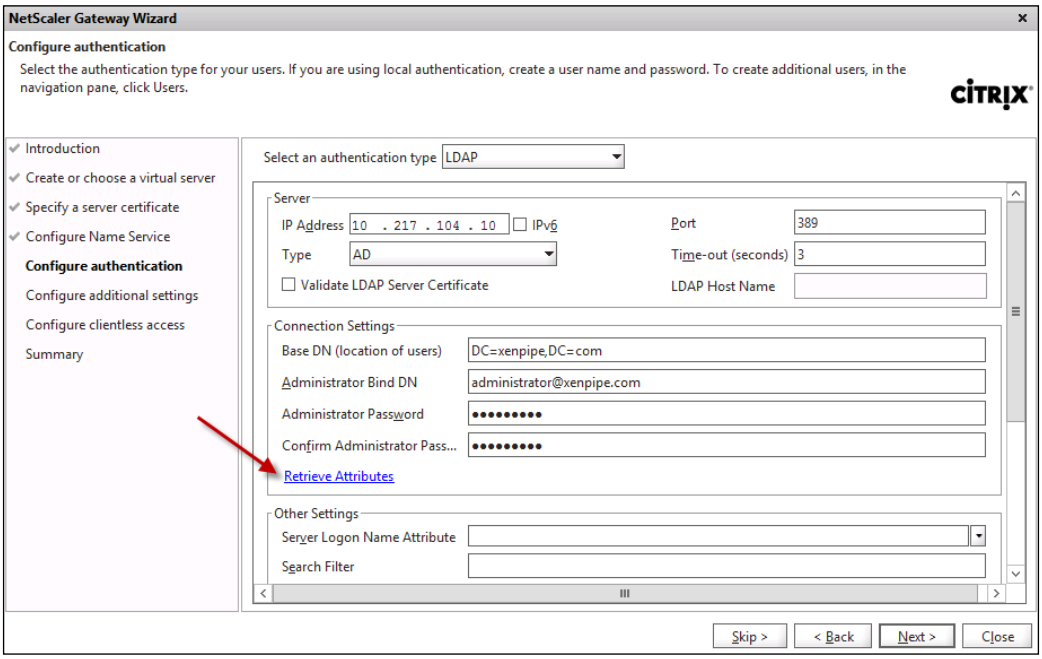

9. Select **Allow** to configure authorization. Also, sometimes people forget to specify HTTPS; so, in case they use HTTP, this redirects them to the HTTPS Site. The domain name must match the domain name used in the certificate, as shown in the following screenshot:

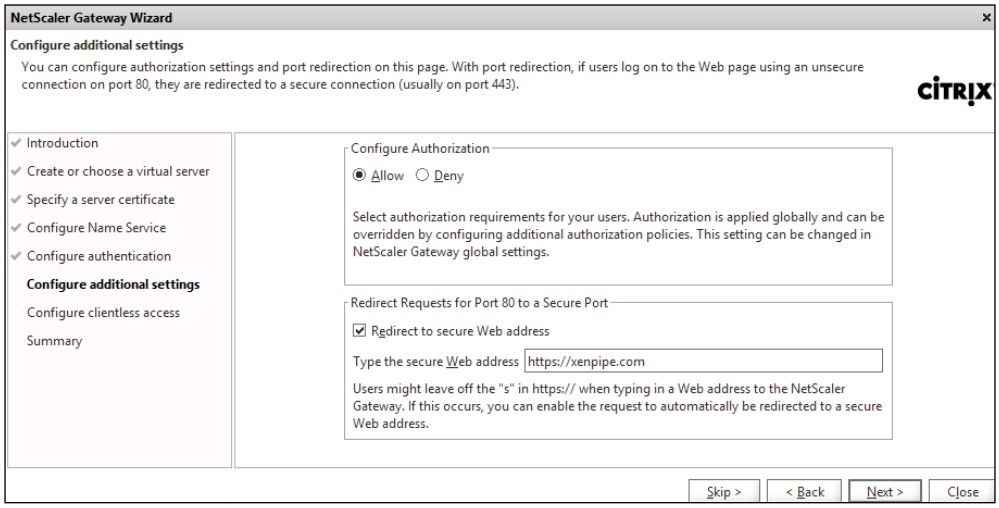

10. Configure **Clientless Access**. This is to allow users to connect with a web browser without having to download and install the plugin, as shown in the following screenshot:

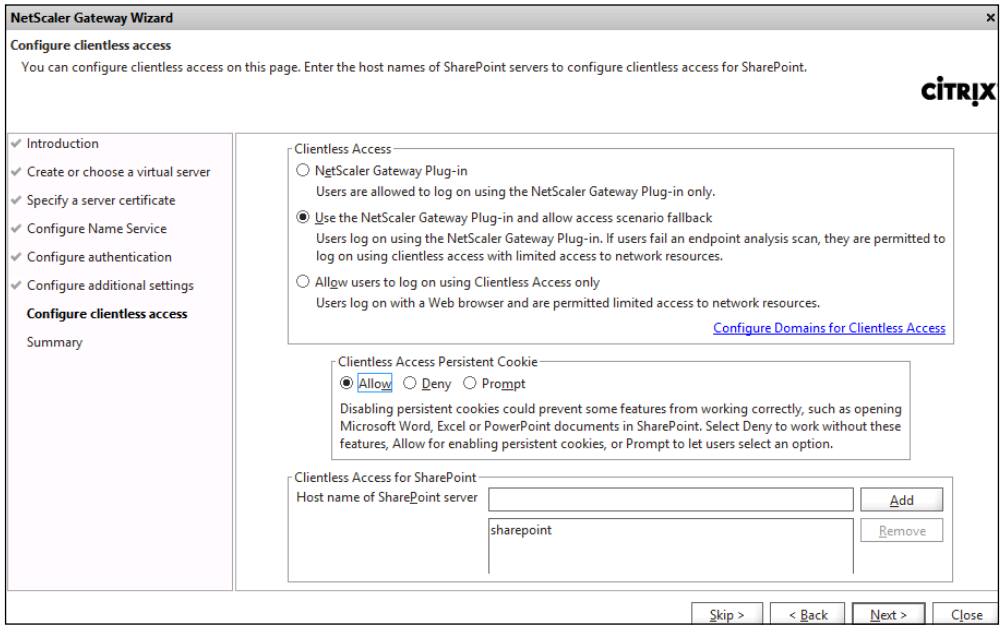

Configuring NetScaler Gateway™ for StoreFront

Next, we need to inform NetScaler Gateway to redirect connections to XenDesktop StoreFront using the following steps:

1. Log in to NetScaler.
2. Navigate to **NetScaler Gateway | Published applications wizard**, as shown in the following screenshot:

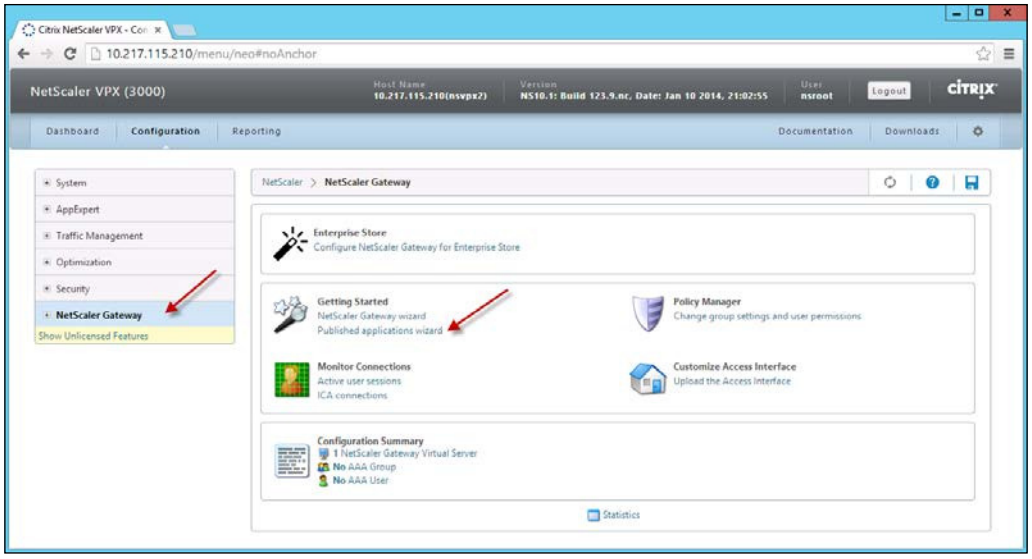

3. Skip the **Introduction** screen and select the virtual server we created earlier. Then, click on **Next**, as shown in the following screenshot:

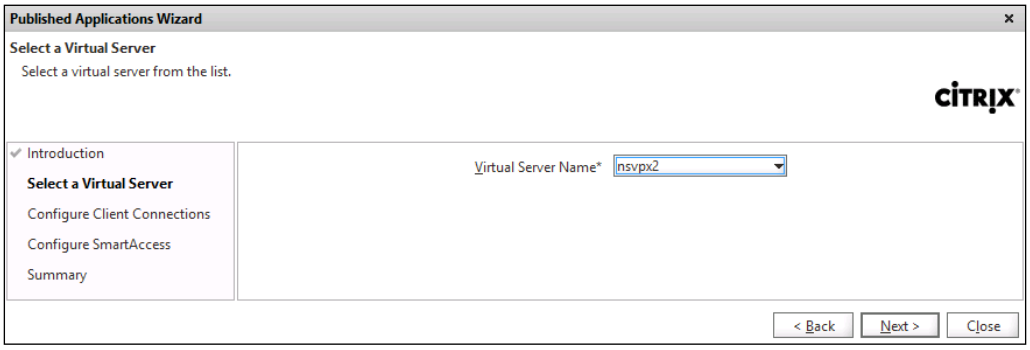

4. Enter the StoreFront URL in the **Web Interface Address** field and the authentication domain. Click on **Add** to add the STA in the `http(s)://<servername>/scripts/ctxsta.dll` format on port `80`. (The STA is located on the Delivery Controller.) Click on **Next**, as shown in the following screenshot.

If you decide to use HTTPS port `443` on the STA, make sure that you have an SSL certificate installed on the Delivery Controller. Either way, using HTTP or HTTPS, when you configure the STA, it should come up and the status should turn green in color.

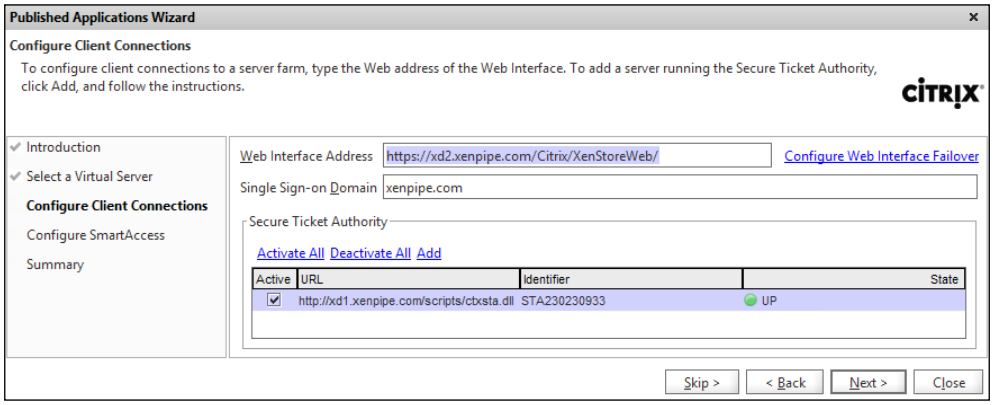

5. Enable the **SETVPNPARAMS_POL** policy. Click on **Next**, then select **Finish**, and then click on **Exit**, as shown in the following screenshot:

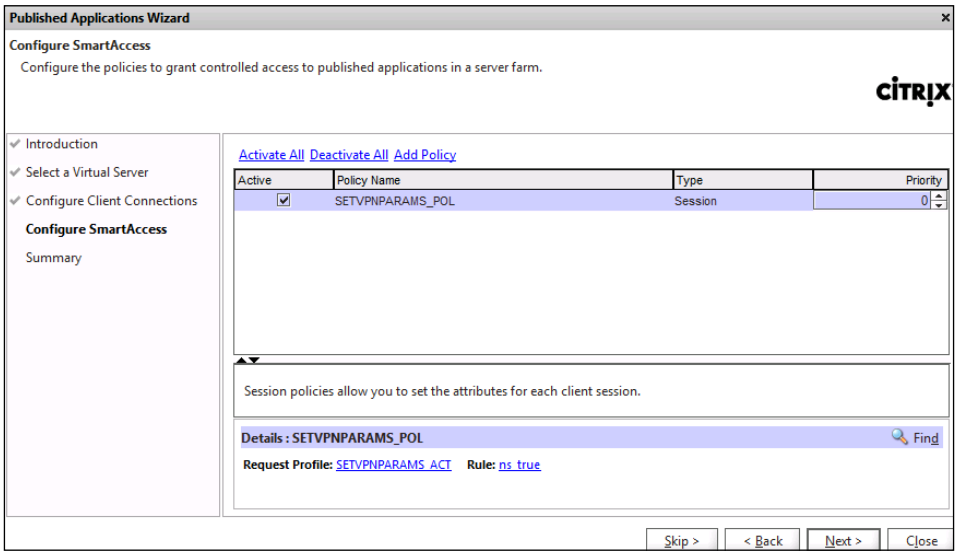

Configuring NetScaler® for an ICA proxy

Only Citrix NetScaler has the ability to frontend the `ICA/HDX` protocol. There might be other vendors who claim that they can do it, but they don't have the legal rights to and certainly don't have the technical specifications within the Citrix product to do it correctly. Frontending the `ICA/HDX` protocol is done by proxying the connections of XenDesktop through NetScaler. This can be done by performing the following steps:

1. Log in to NetScaler.
2. Navigate to **NetScaler Gateway | Global Settings | Change global settings**, as shown in the following screenshot:

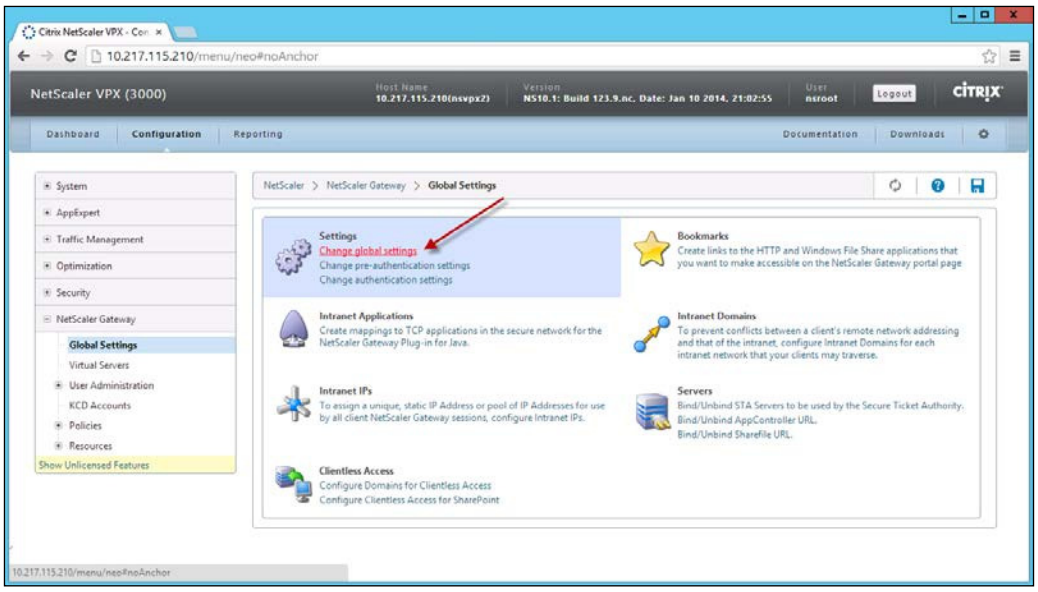

3. Select the **Client Experience** tab. Enable **Display Home Page**, configure **Split Tunnel** to **OFF** and **Clientless Access** to **On**, and change the **UI Theme** to **Green Bubble**, as shown in the following screenshot:

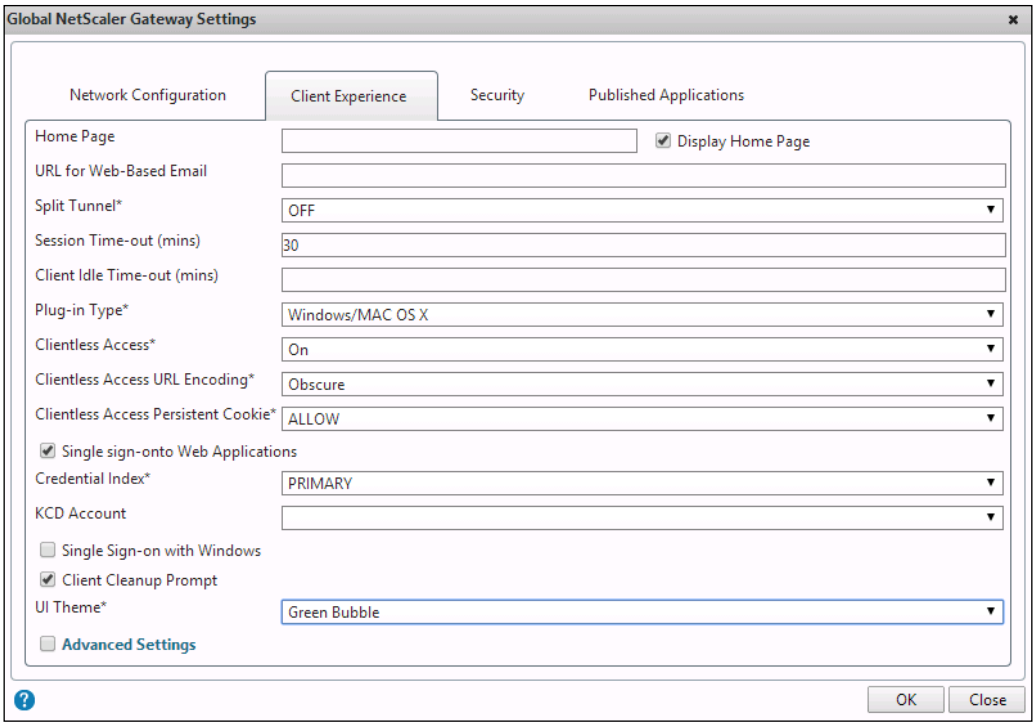

4. Select the **Published Applications** tab and make sure that **ICA Proxy** is turned **ON**, as shown in the following screenshot:

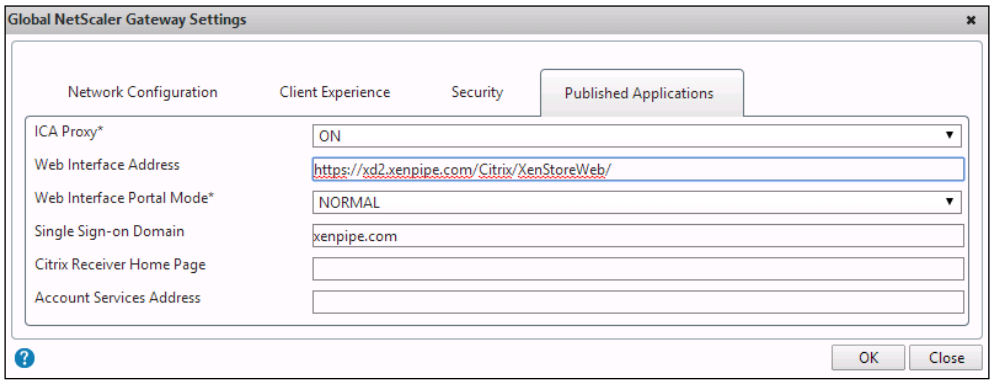

Configuring a StoreFront connection to NetScaler Gateway™

The last thing you need to do is configure the StoreFront server to communicate with the access gateway. This can be done by performing the following steps:

1. Log in to the StoreFront server and launch Citrix StoreFront.
2. Select **Authentication** in the left-hand side navigation menu.
3. Select **Add/Remove Methods** in the right-hand side navigation menu. Select **User name and password**, **Domain pass-through**, and **Pass-through from NetScaler Gateway**. Click on **OK**, as shown in the following screenshot:

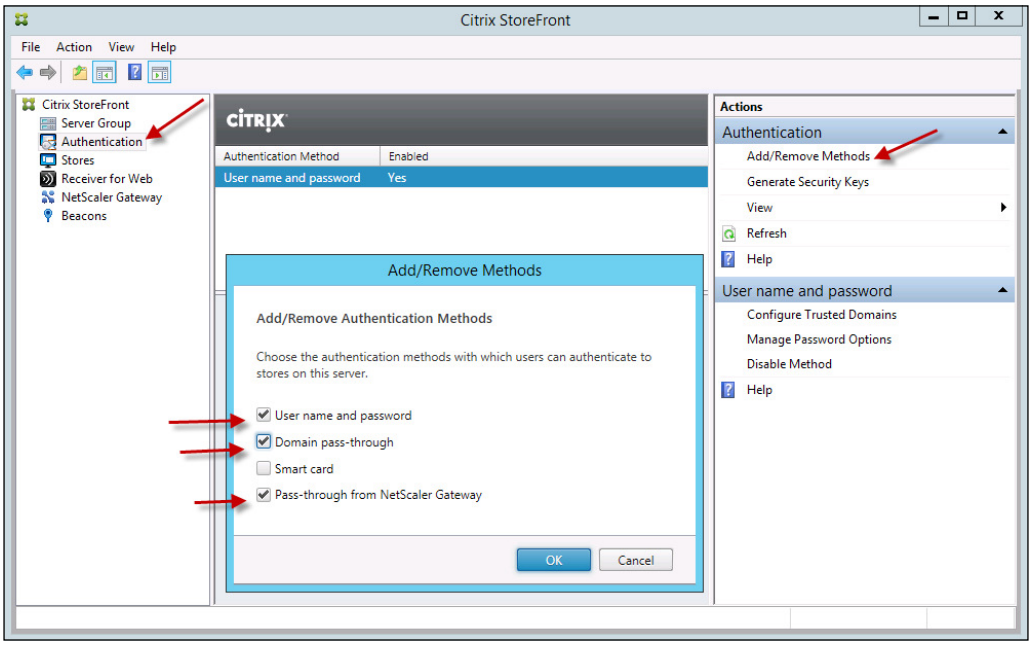

4. Select **NetScaler Gateway** in the left-hand side navigation menu.

5. Select **Add NetScaler Gateway Appliance** in the right-hand side navigation menu, as shown in the following screenshot:

6. Give a name to the gateway; this shows up in the console.
7. Enter the gateway URL that users will enter into their browsers in the `https://<NetScalerGatewayFQDN>/Citrix/<Storename>Web` format.

The **NetScalerGatewayFQDN** field must match the **Common Name** field in the SSL certificate. It is best to use a wildcard certificate so that you can give yourself the flexibility in naming a NetScaler Gateway connection. In our example, we created a certificate with the common name `*.xenpipe.com`. This way, we can name our NetScaler Gateway connection `https://ng.xenpipe.com` or `https://<sub-domain>.xenpipe.com`, where `sub-domain` can be anything we choose.

8. Enter the callback address for StoreFront to silently authenticate users on the inside network in the `https://<NetScalerGatewayFQDN>/CitrixAuthService/AuthService.asmx` form at, as shown in the following screenshot:

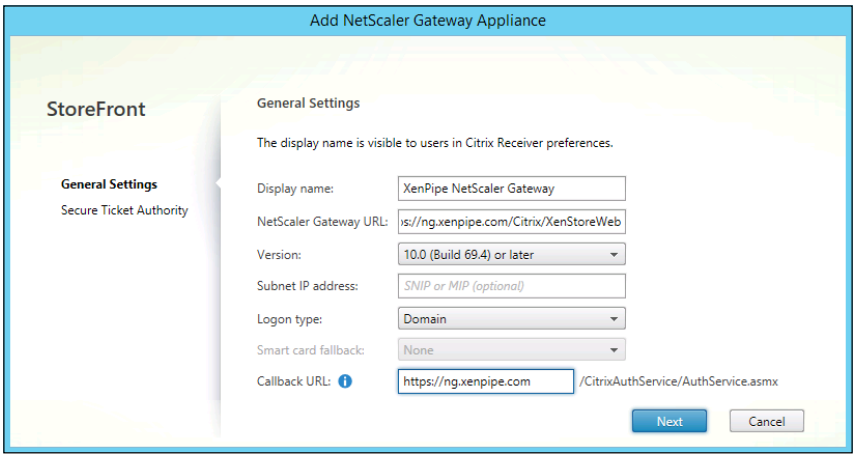

9. Click on **Add** to enter the URLs in the **Secure Ticket Authority URLs** section of your XenDesktop controller. In our example, this is `http://xd1.xenpipe.com`, as shown in the following screenshot.

This URL needs to be the same as what was configured on NetScaler Gateway.

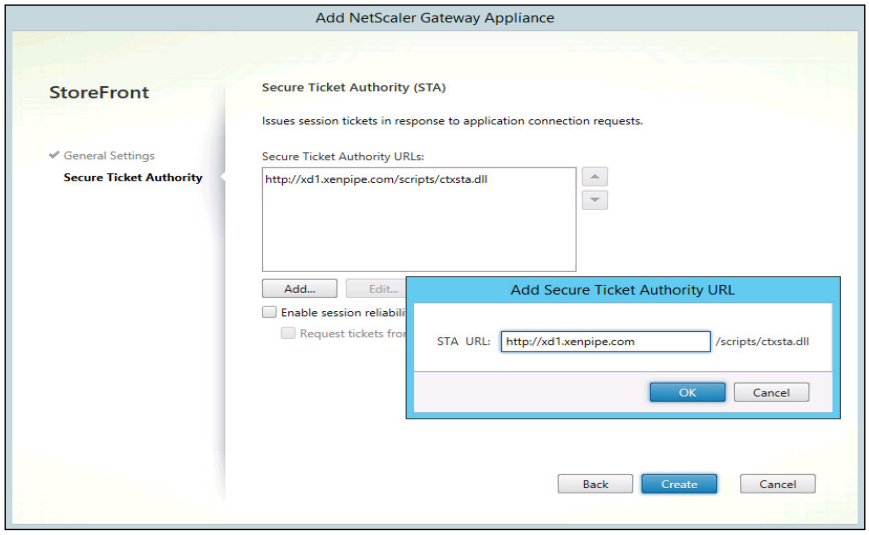

10. Click on **OK**, then select **Create**, and then click on **Finish**.
11. Select **Stores** in the left-hand side navigation menu.
12. Select the store in the middle navigation menu.
13. Click on **Enable Remote Access** in the right-hand side navigation menu.
14. Select **No VPN tunnel** and then click on **OK**, as shown in the following screenshot:

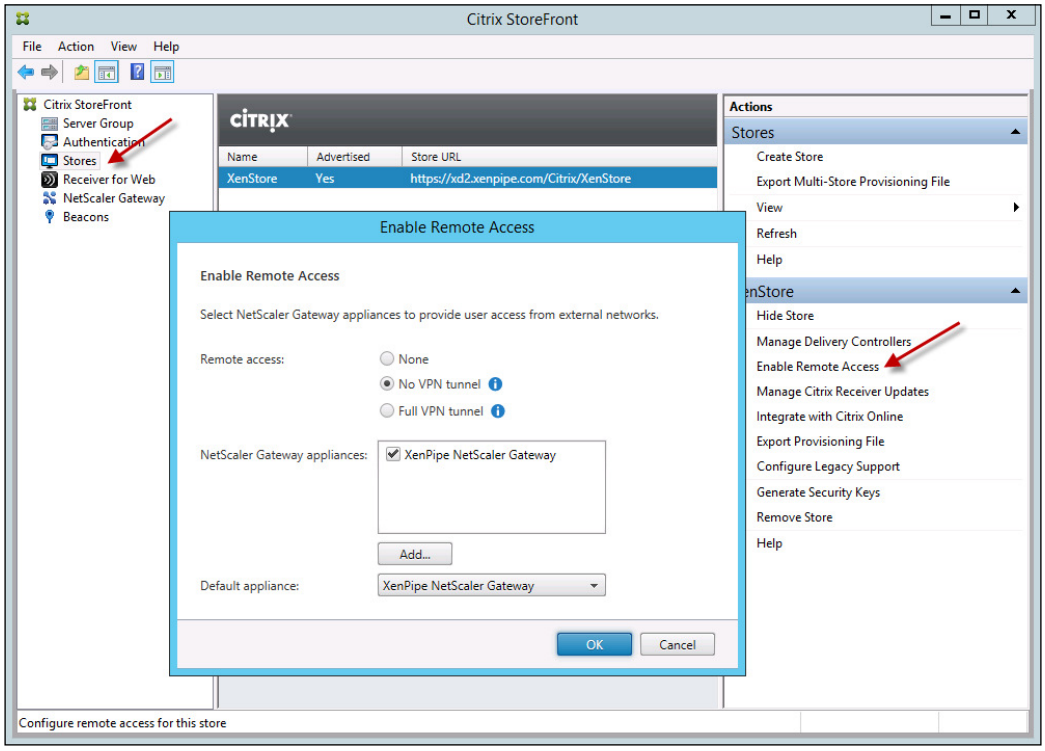

Exporting the StoreFront certificate

NetScaler needs to trust the StoreFront server when it makes a callback. You need to import the StoreFront certificate into NetScaler for this to happen. You need to export the key in two formats for the import: a `.pfx` format and a `.cer` format.

To export the StoreFront certificate in the `.pfx` format, perform the following steps:

1. Log in to the StoreFront server.
2. Launch the **Microsoft Management Console** (**MMC**).

3. Navigate to **Console Root** | **Certificates (Local Computer)** | **Personal** | **Certificates**.
4. Locate the server certificate from your Site and right-click on **Export…**, as shown in the following screenshot:

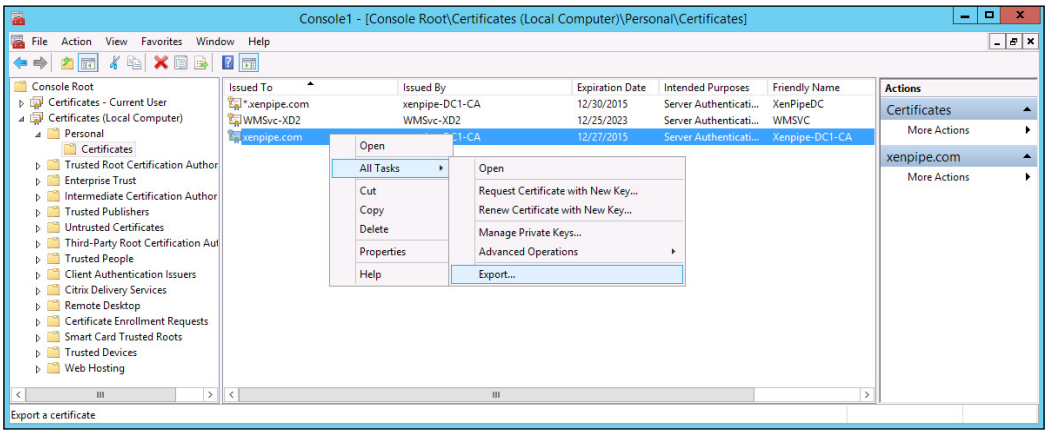

5. Select **Yes, export the private key** as shown in the following screenshot:

6. Use the **PKCS #12 (.PFX)** format. Select **Include all certificates in the certification path if possible** and **Export all extended properties**, as shown in the following screenshot:

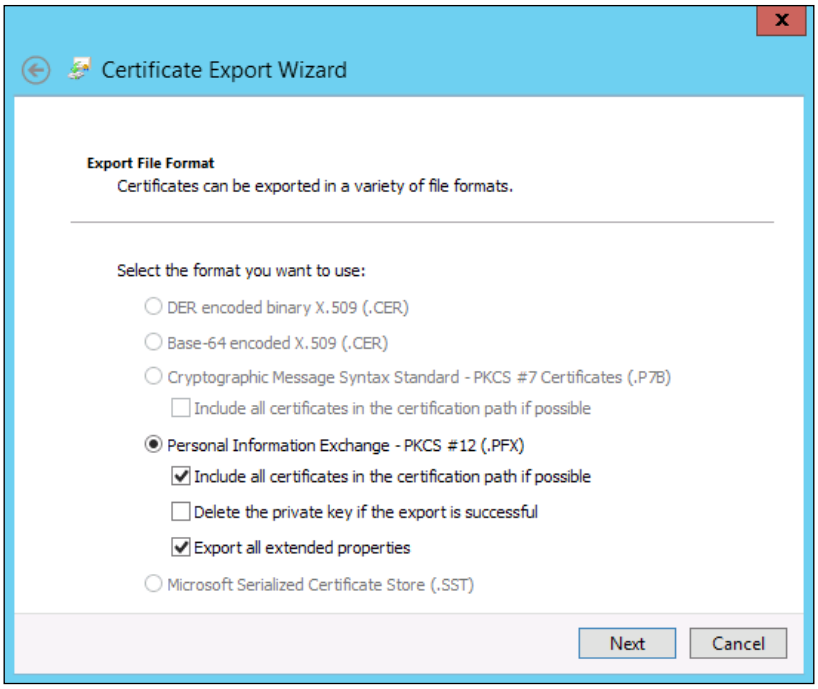

To export the StoreFront certificate in the `.cer` format, perform the following steps:

1. From the MMC, repeat the same export procedure.
2. Select **No, do not export the private key**, as shown in the following screenshot:

3. Select the **Base-64 encoded X.509 (.CER)** format as shown in the following screenshot:

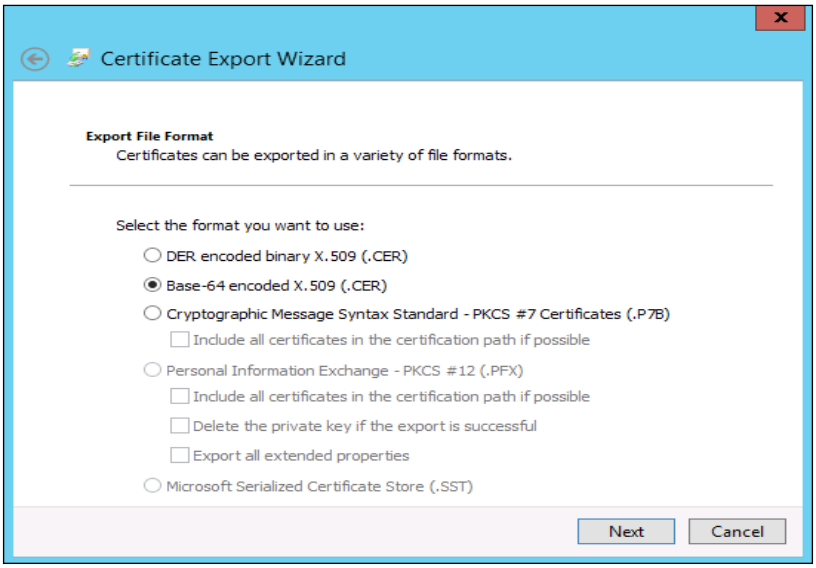

At this point, you should have two exported certificate files on the StoreFront server: one with a `.pfx` extension and the other with a `.cer` extension.

Importing the StoreFront certificate into NetScaler Gateway™

To import the StoreFront certificate into NetScaler Gateway, perform the following steps:

1. Log in to NetScaler Gateway.
2. Navigate to **Traffic Management | SSL**.
3. Under **Tools**, select **Import PKCS #12**.
4. Locate the certificate and import the `.pfx` file into the **PKCS12 File** field. Enter an output filename with a `.key` extension. Select **DES3** as the **Encoding Format** and specify the password. This will create a key file that we will use in a moment, as shown in the following screenshot:

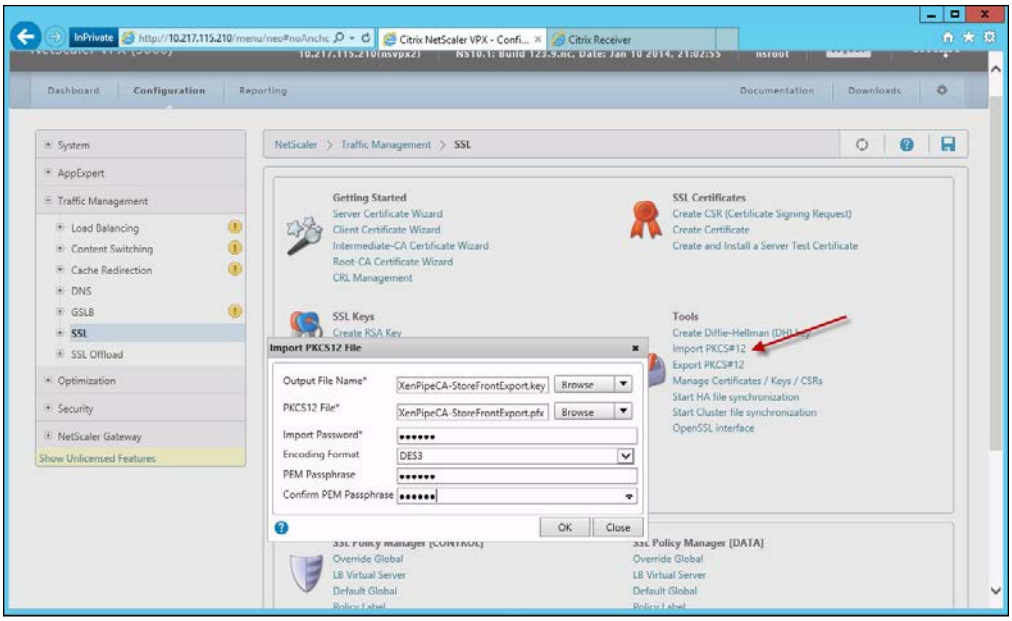

5. Under **Tools**, select **Manage Certificates / Keys / CSRs**. Upload the `.cer` certificate file, as shown in the following screenshot:

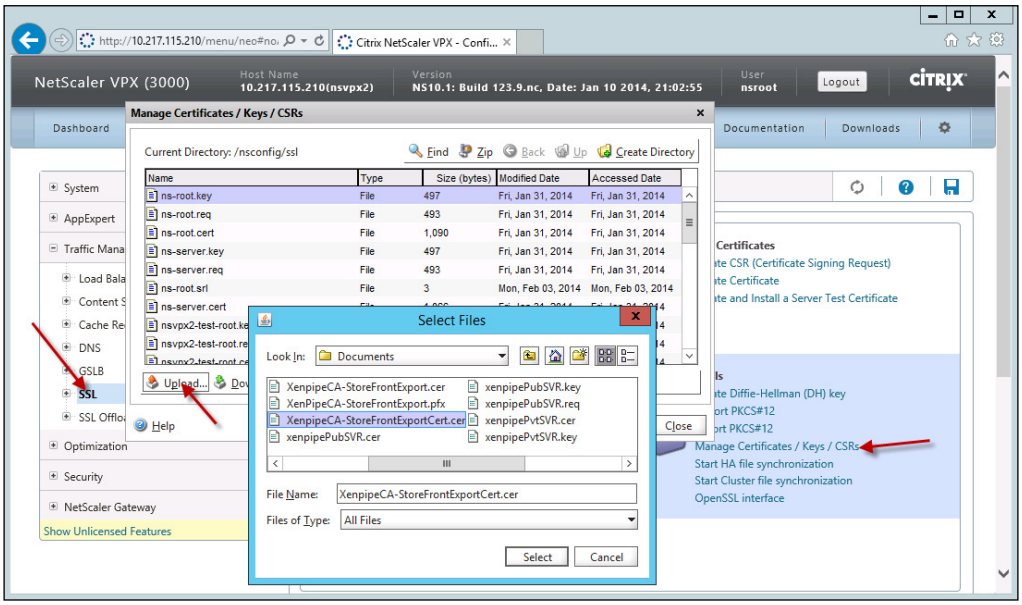

6. You should see three files on NetScaler Gateway, `.pfx`, `.cer`, and `.key` files, as displayed in the following screenshot:

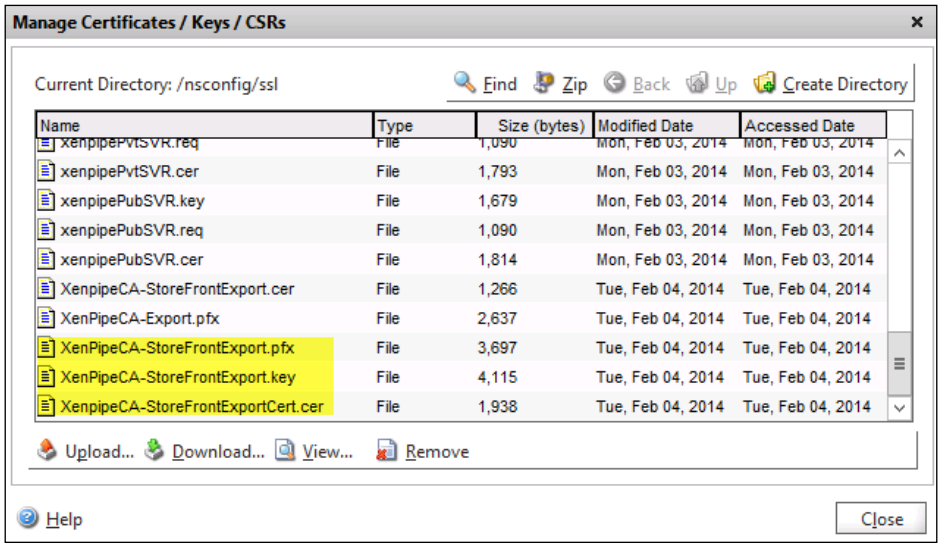

7. Navigate to **Traffic Management** | **SSL** | **Certificates**. Select **Install…**, as shown in the following screenshot:

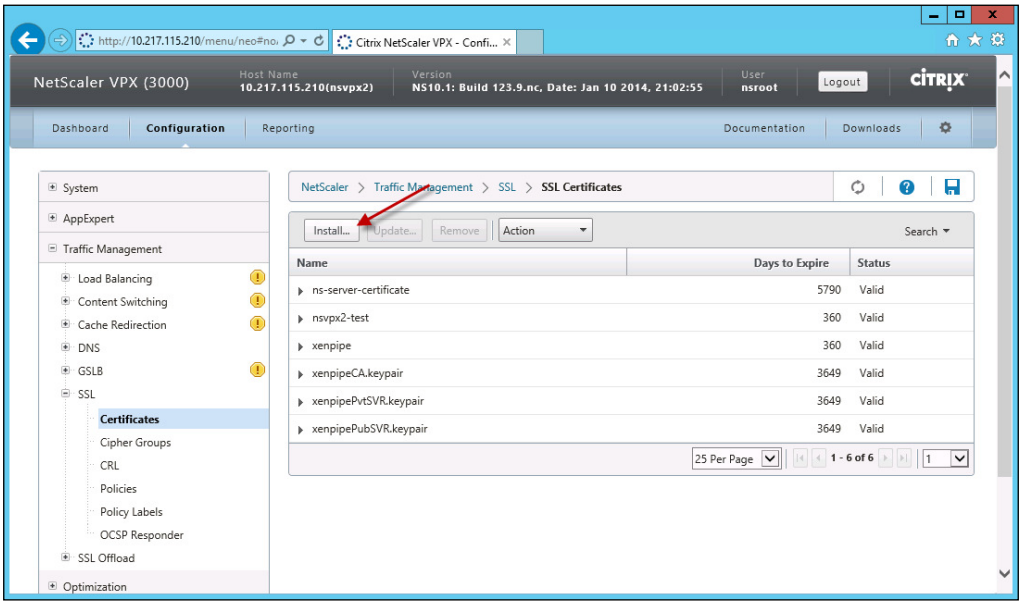

8. Create a certificate-key pair name of your own. Specify the `.cer` file and the `.key` file you just imported, as shown in the following screenshot:

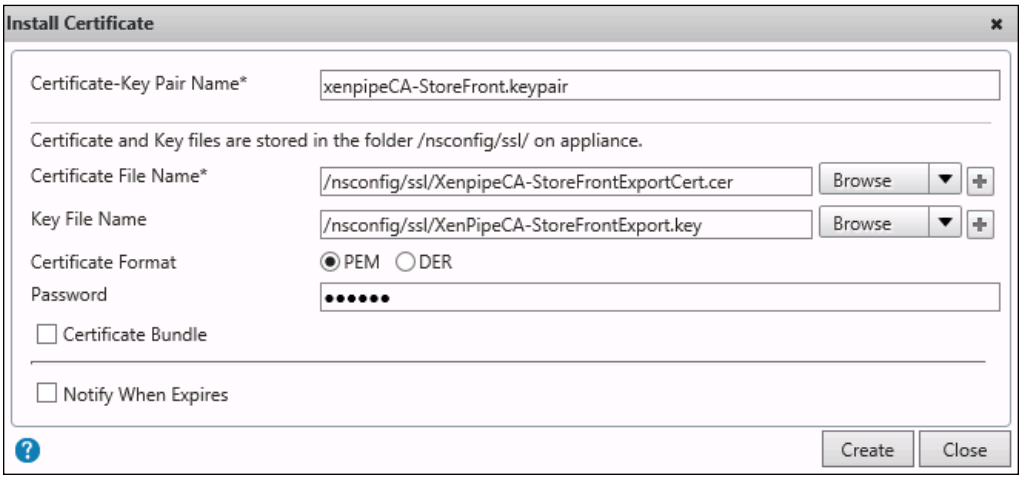

At this point, there is nothing left to do in order to configure NetScaler for securing XenDesktop.

Secure Ticket Authority

The STA runs on the controller and is embedded within the Citrix XML service. The STA issues session tickets in response to the connection requests for published resources. The STA is simply used to validate user sessions and session requests. It does not perform encryption or validate packet integrity. For these types of security, you need to use SSL.

Securing the ICA/HDX protocols

ICA is not secure by default. You can turn on a feature called SecureICA, which encrypts the `ICA` protocol; however, it does not perform data integrity checks, so you will eventually need to implement SSL for communications. We will go through SSL in the rest of this chapter.

To turn on the SecureICA feature, perform the following steps:

1. Launch Studio.
2. Select **Delivery Groups** in the left-hand side navigation menu.

3. Select the delivery group to be edited in the center pane.
4. In the right-hand side pane, click on **Edit Delivery Group** and then on **End User Settings**.
5. Select **Enable Secure ICA**, as shown in the following screenshot:

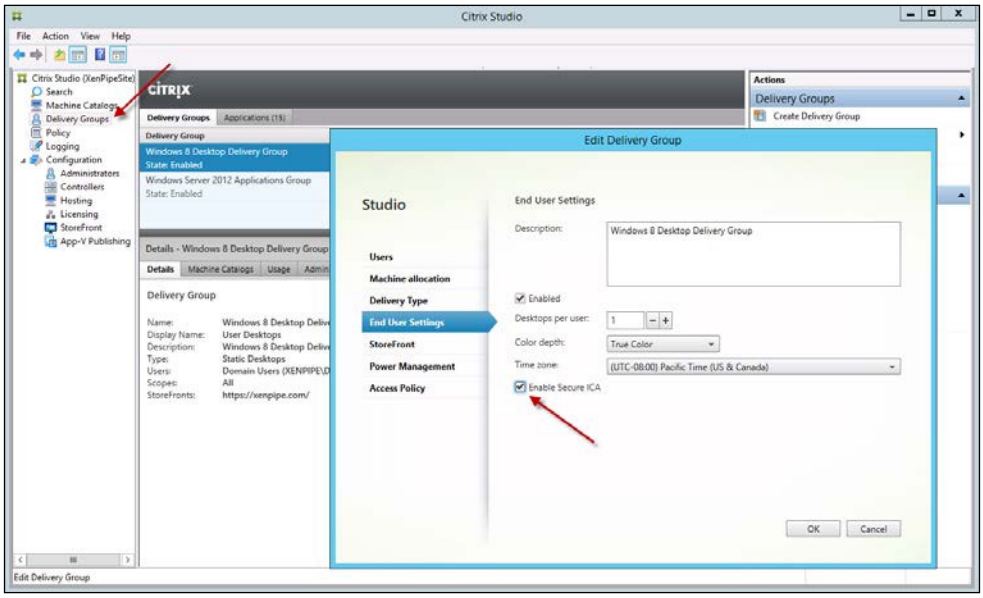

Securing StoreFront

Traditionally, StoreFront runs on Windows Server 2012 R2. Citrix uses the SSL standard to secure communications with XenDesktop. The most common deployment is to use a Single-Hop DMZ with NetScalers as the frontend to the XenDesktop deployment. The feature of NetScaler that is used for SSL security is called NetScaler Gateway. HTTPS or SSL is used on the external, public-facing side of NetScaler. HTTPS or SSL is optional on the internal, private-facing side of NetScaler. Some organizations that have very strict security policies also use SSL on the internal, private-facing side of NetScaler.

NetScaler runs a load balancing service for StoreFront in addition to SSL. This is especially helpful in case you need to scale up or grow your deployment by adding more StoreFront servers without disrupting the existing operations. Also, if a StoreFront server goes down, NetScaler will load balance around the failed server, ensuring high availability for your users. Citrix makes setting this up really easy. In the NetScaler GUI, there is a wizard that walks you through setting up load balancing along with NetScaler Gateway for SSL communications, which we covered at the beginning of this chapter.

Securing Receiver

In order to complete the secure connection between the client and StoreFront, you need to configure SSL or VPN in Receiver under the NetScaler Gateway settings.

To configure SSL in Receiver, perform the following steps:

1. In the Windows notification area, right-click on the Receiver icon and choose **Preferences**.
2. Right-click on the **Online Plug-in** entry in the **Plug-in Status** entry and choose **Change Server**.
3. The **Change Server** screen displays the currently configured URL. Enter the server URL in the textbox in the `https://servername` format to encrypt the configuration data using SSL/TLS.
4. Click on **Update** to apply the change.

Securing controller

The XenDesktop controller runs on Microsoft Windows Server, so SSL is typically implemented on IIS. SSL is implemented using server certificates.

IIS

To install a server certificate on the controller running IIS, you need to create a certificate request, submit it to a CA, and then install the issued certificate in IIS.

To implement SSL on IIS, follow the Microsoft article at `http://support.microsoft.com/kb/299875`.

Non-IIS

If you are not running IIS on the controller, you can still use SSL. You can create a certificate request manually.

The steps for doing this manually are located at the following links:

- `http://blogs.technet.com/b/pki/archive/2009/08/05/how-to-create-a-web-server-ssl-certificate-manually.aspx`
- `http://msdn.microsoft.com/en-us/library/ms733791%28v=vs.110%29.aspx`

Changing the controller port to HTTPS

By default, the controller is installed to communicate with the XML service on HTTP port `80`. Once you install SSL, you can change the port to listen on `443` for additional security. This is optional.

To change the XML service port to `443`, perform the following steps:

1. Open a command window in the controller.
2. Run the following command:

```
BrokerService –WIPORT 80 –WISSLPORT 443
```

Securing Studio and Director

You can further bolster the security of your XenDesktop Site by installing SSL on both the Studio and the Director servers. Both Studio and Director are web-based interfaces, so they run on IIS. To secure these, you can install SSL on IIS.

IIS

To install a server certificate on the Studio and Director servers running IIS, you need to create a certificate request, submit it to a CA, and then install the issued certificate in IIS.

To implement SSL on IIS, follow the Microsoft article at `http://support.microsoft.com/kb/299875`.

Securing the XenDesktop® to XenServer® communications

When XenDesktop is installed on XenServer, you need to secure the communications between these two. You will need to replace the default SSL certificate on XenServer.

To replace the default certificate on XenServer, perform the following steps:

1. Log in to XenServer and get to a command prompt.

2. Modify `/etc/pki/tls/openssl.cnf` as follows:
    - Uncomment the following line:

```
req_extensions = v3_req
```

    - Modify the request section as follows:

```
[v3_req]
basicConstraints = CA:FALSE
keyUsage = keyEncipherment
extendedKeyUsage = serverAuth
```

3. Generate a certificate request as follows:

```
openssl genrsa -out [servername].private 2048
openssl req -new -outform PEM -out [servername].request -keyform
PEM -key [servername].private -days 1800
```

   Here, `[servername]` is the XenServer hostname.

4. Submit the request `[servername].request` to a CA and retrieve the CA-signed certificate.
5. Move the existing certificate as follows:

```
mv /etc/xensource/xapi-ssl.pem /etc/xensource/xapi-ssl.pem_orig
```

6. Add the new CA-signed certificate as follows:

```
cat [servername].public [servername].private > [servername].pem
install -m 0400 [servername].pem /etc/xensource/xapi-ssl.pem
```

7. Edit the `.xapissl` file as follows:

```
vi /etc/init.d/xapissl
PEMFILE="/etc/ssl/certs/[servername].pem"
```

8. Restart the XenServer SSL communications service as follows:

```
/etc/init.d/xapissl restart
```

To install the CA-signed certificate on the XenDesktop controller, perform the following steps:

1. On the controller, open Windows Explorer and locate the root certificate.
2. Right-click on the root certificate and select **Install Certificate**.
   The **Certificate Manager** wizard appears.

3. On the **Certificate Store** page, select **Place all certificates in the following store**.
4. Navigate to **Browse** | **Show Physical Stores**.
5. Expand **Trusted Root Certification Authorities** and select **Local Computer**.
6. Select the local computer and click on **OK**.
7. Complete the wizard and click on **Install**.

Using smart cards

Some organizations augment their security policy by requiring users to authenticate using smart cards. A smart card is a plastic card, much like your credit card or bank card, with a built-in microprocessor and is used for personal identification. The microprocessor replaces the magnetic strip you usually see on credit cards and bank cards. The data on magnetic strips can be easily read, written, or deleted with off-the-shelf equipment. Smart cards have memory and a microprocessor. They use a serial interface for data communications and receive power from a card reader. The card reader draws its power from the USB port it is plugged in to, typically a thin client. The smart card and card reader communicate with the server to authenticate a user. Setting up smart card authentication can be tricky. I don't have enough room in this book to cover the topic and it is not an exact science. Smart cards will have different nuances depending on which smart cards and card readers you use. Make sure that the smart card and card readers that you use have been validated by the manufacturer in the Citrix Ready program, `www.citrixready.com`.

Get started with setting up smart cards at `http://support.citrix.com/proddocs/topic/dws-storefront-21/dws-configure-smartcard.html`.

Summary

XenDesktop traffic is not completely secure. By following the guidelines in this chapter, you can make it secure. It is best to use SSL for security with a pair of NetScalers in the DMZ to encrypt the XenDesktop traffic. In this chapter, we looked at how to secure the traffic using NetScaler along with some other security implications for XenDesktop. Next, we will look at how to manage and monitor XenDesktop.

14
Managing and Monitoring XenDesktop®

Managing and monitoring is always important. If you can't see it, you can't manage it. If you can't manage it, you can't monitor it. XenDesktop Director is a web-based tool that enables IT and support teams to monitor a XenDesktop environment and perform troubleshooting.

In this chapter, we will discuss the following topics:

- How to manage XenDesktop
- How to monitor XenDesktop
- How to use HDX Insight
- Troubleshooting
- Third-party tools

Using Studio to manage the XenDesktop® Site

You've been learning how to manage XenDesktop all throughout this book using Studio. After successfully using the **Install** wizard, you learned how to manage machine catalogs, hosts, Personal vDisks, and delivery groups. You also learned how to manage hosted applications, Delivery Controllers, policies, printing, USB, storage, and HDX. You even had a chance to learn how the SDK works from within XenDesktop Studio.

Role-based administration is new in XenDesktop 7.x and uses roles that you can define in addition to the administrator role. The following table illustrates the roles:

Role	Capabilities
Full administrator	This role can perform all tasks.
Read-only administrator	This role can see all objects on a global level but cannot change anything.
Help desk administrator	This role can view delivery groups and manage the sessions and machines in the groups.
Machine catalog administrator	This role can create and manage machine catalogs and provide machines in these catalogs. This can also manage base images and install software.
Delivery group administrator	This role can configure applications, desktops, and machines for delivery along with sessions, policies, and power settings.
Host administrator	This role can manage host connections and resource settings. This cannot deliver machines, applications, or desktops to users.

To configure role-based administration in Studio, perform the following steps:

1. Launch Studio.
2. Navigate to **Configuration** | **Administrators**.

3. Select **Create Administrator**, as shown in the following screenshot:

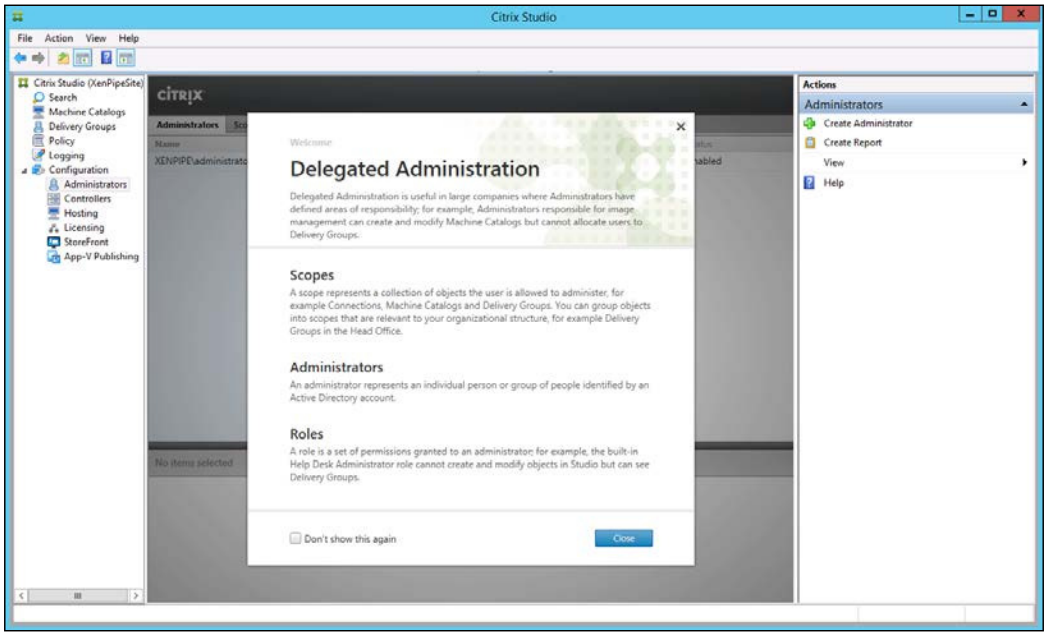

Using Director to monitor the XenDesktop® Site

XenDesktop is monitored using Director. Director is a web-based monitoring tool that ties in nicely with XenDesktop roles, which determine what can be displayed. With XenDesktop Director, you can view the status of sessions, users, and Site infrastructure. With XenDesktop Director, you can also view the users' applications and processes inside their virtual desktop and quickly help them out by ending the unresponsive applications or processes. You can even restart users' machines for them.

Director is installed by default as a website on the XenDesktop Delivery Controller. You can also install it on a separate server. In our network diagram, we have Studio and Director installed on separate servers.

As previously mentioned, role-based administration is new in XenDesktop 7.x and uses roles that you can define in addition to the administrator role. Several of the roles are displayed in the following table:

Role	Capabilities
Full administrator	This role can see all views and trends
Read-only administrator	This role can see all views on a global level but cannot change anything
Help desk administrator	This role can view only the help desk and user detail views for objects that the administrator can manage
Machine catalog administrator	This role cannot view anything
Delivery group administrator	This role can see all views and trends
Host administrator	This role cannot view anything

XenDesktop Director uses SSL to encrypt communications. You can run it with or without SSL.

To configure SSL on the Director server, perform the following steps:

1. Follow the steps provided in *Chapter 2*, *Installing XenDesktop®*.
2. Create a server certificate and add a Site binding. You can also import and bind the server certificate issued to the StoreFront server.

To disable SSL on the Director server, perform the following steps:

1. Log in to the Director server.
2. Launch the **Internet Information Server** (**IIS**).

3. Navigate to **Director** in the left-hand side navigation menu and select **Application Settings**, as shown in the following screenshot:

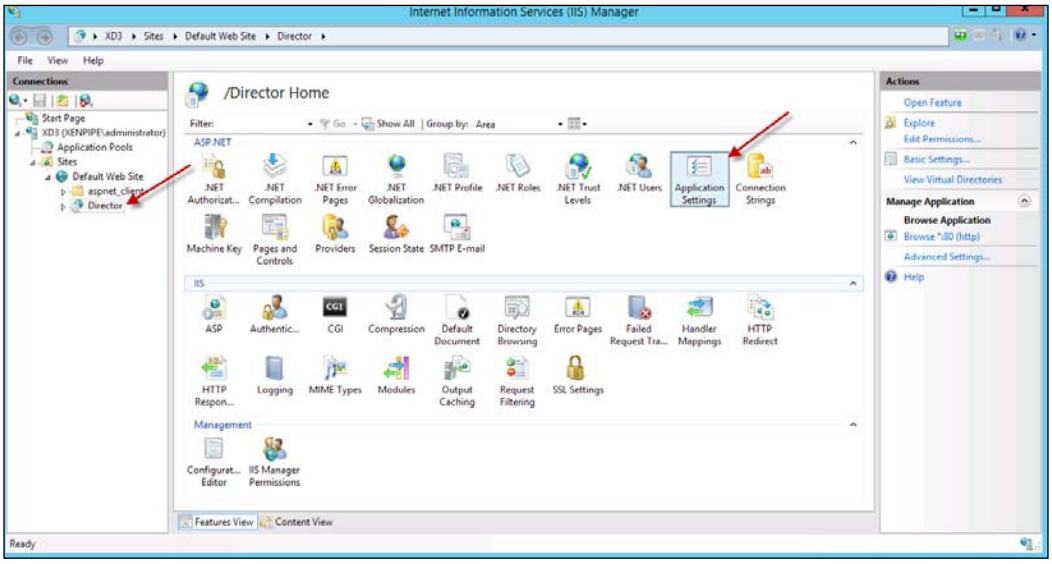

4. Change the value of **UI.EnableSslCheck** to `false`. Then, click on **OK**, as shown in the following screenshot:

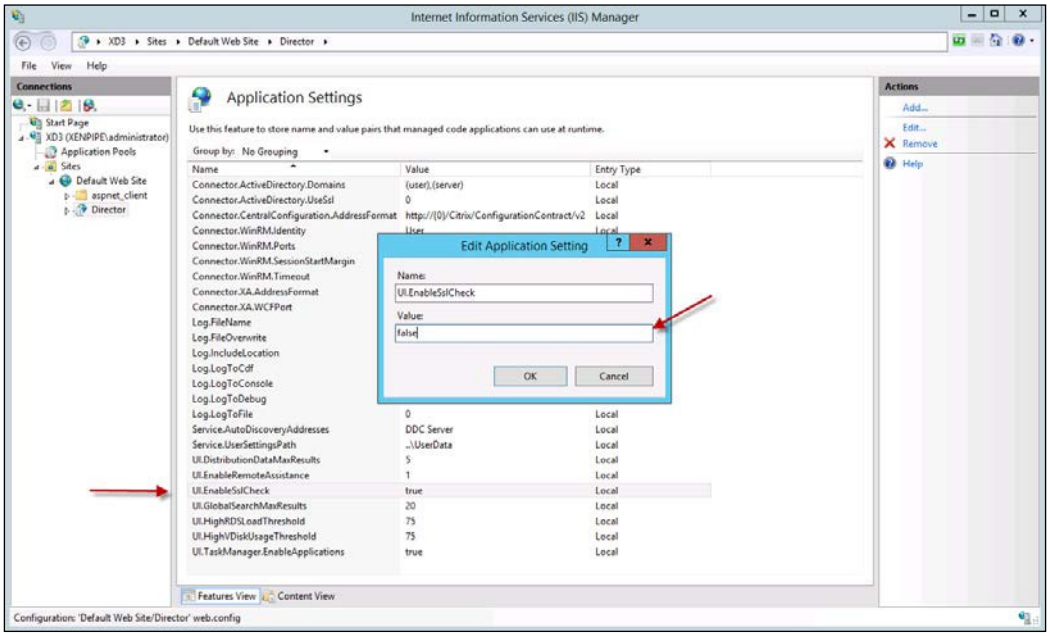

For desktops to be usable, they must be registered (that is, establish communication) with the correct controller or with any one of the controllers, if there are more than one. To change the discovery address, perform the following steps:

1. From the same location, change the value of **Service.AutoDiscoveryAddresses** to the name of the Delivery Controller. In our example, this is `xd1.xenpipe.com`, as shown in the following screenshot:

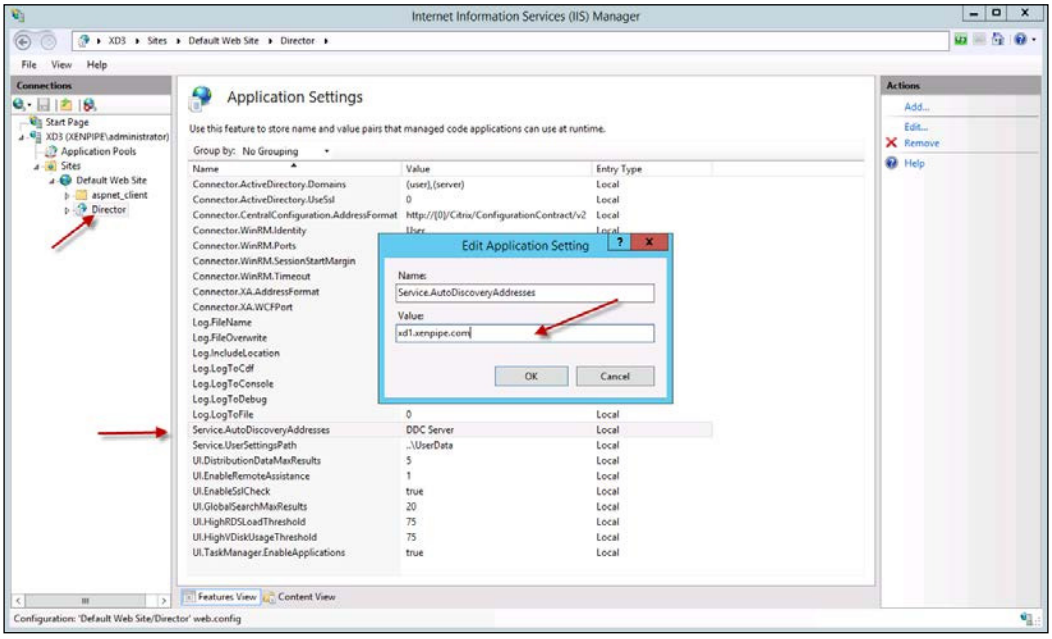

2. Click on **Restart** to restart IIS, as shown in the following screenshot:

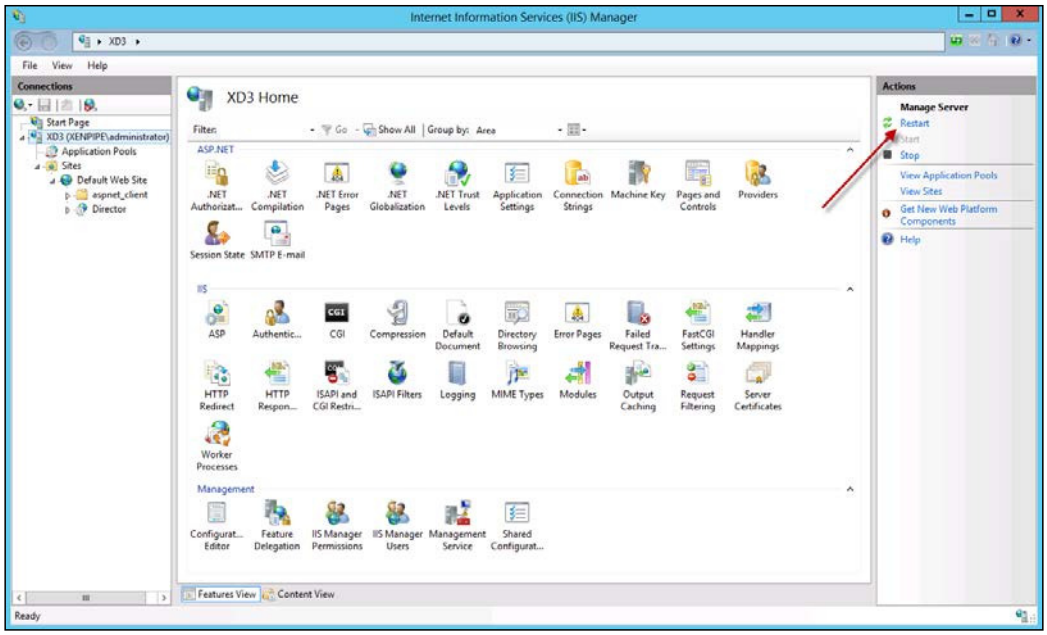

To run Director, perform the following steps:

1. Log in to the Director server. In our example, we use `xd3.xenpipe.com`.
2. Search for `Director` and launch it, as shown in the following screenshot:

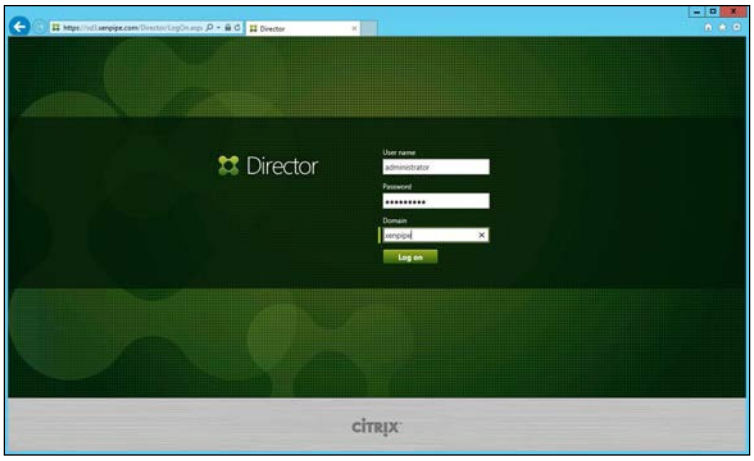

You can also run Director from a web browser.

Using HDX Insight™

NetScaler Insight Center is new and is the integration of the Edgesight network analysis and the Edgesight performance management functionality in Director. There are two main components of NetScaler Insight Center, as follows:

- Web Insight shows the performance of web applications
- HDX Insight shows the performance of XenDesktop and XenApp applications

HDX Insight provides the application and desktop views of the network. This provides an advanced analytics of the ICA traffic in the XenDesktop deployment. HDX Insight performance management provides the historical capacity and health trend reporting features. A key metric that is often looked at in Citrix XenDesktop deployments is the **ICA Session Round Trip Time** (**ICA RTT**), which is an indication of usability.

HDX Insight network analysis must be used with NetScaler Insight Center. NetScaler Insight Center must be installed and configured in Director. NetScaler Insight Center is a virtual machine (appliance) that can be downloaded from `http://citrix.com`. Director gathers information related to XenDesktop from HDX Insight. The analysis that it provides is a robust view of the Citrix `ICA/HDX` protocol from end to end, from the client all the way through to the backend infrastructure.

Once you are logged in to NetScaler Insight Center and have added a NetScaler appliance to the inventory, you will see a **Dashboard** tab that shows Web Insight and HDX Insight data and another tab called **Configuration**, where you add NetScalers and AppFlow connections.

To import NetScaler Insight Center, perform the following steps:

1. Log in to `http://www.citrix.com/downloads.html`.
2. Locate and download the NetScaler Insight Center virtual appliance for XenServer. It is located under the NetScaler download section in **Components**.
3. Log in to XenCenter and select the XenServer you want to install it on.
4. Navigate to **File** | **Import…**, as shown in the following screenshot:

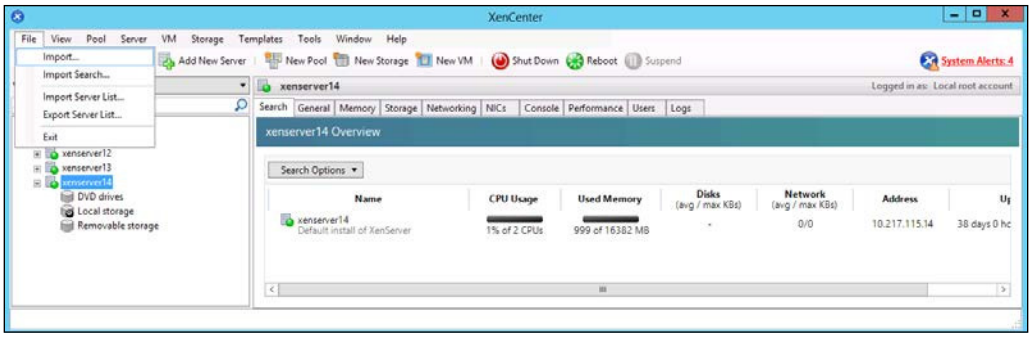

5. Browse for the `.xva` file and import it. Be sure to specify two network interfaces.
6. Once the virtual appliance boots up, in the **Console** screen, enter a basic networking configuration, as shown in the following screenshot:

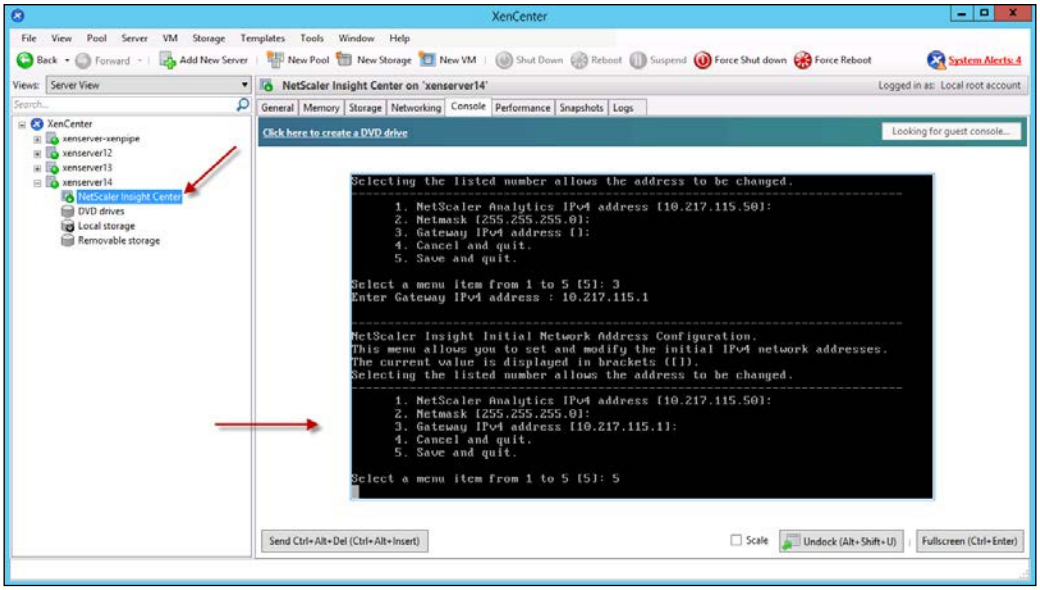

To get the NetScaler appliance ready, perform the following steps:

1. Log in to the NetScaler appliance and not the NetScaler Insight Center virtual machine that you just imported.
2. Navigate to **Network** | **IPs** and open the **NetScaler IP** (**NSIP**) address. Make sure that **GUI** is checked and uncheck **Secure Access Only**.

To access NetScaler Insight Center, perform the following step:

1. Open a browser and connect to the IP address that you configured in the previous section, as shown in the following screenshot:

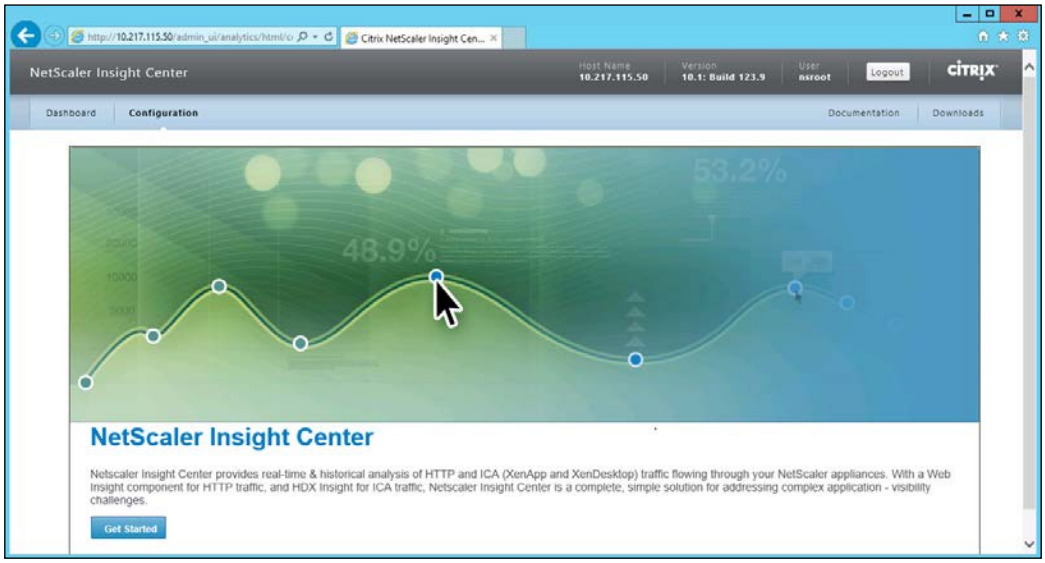

To add the NetScaler appliance to the inventory, perform the following step:

1. Either click on **Get Started** or on the **Configuration** tab and navigate to **Inventory** | **Add**. Enter the IP address that NetScaler uses to send its AppFlow data to NetScaler Insight Center, typically the NSIP, as shown in the following screenshot:

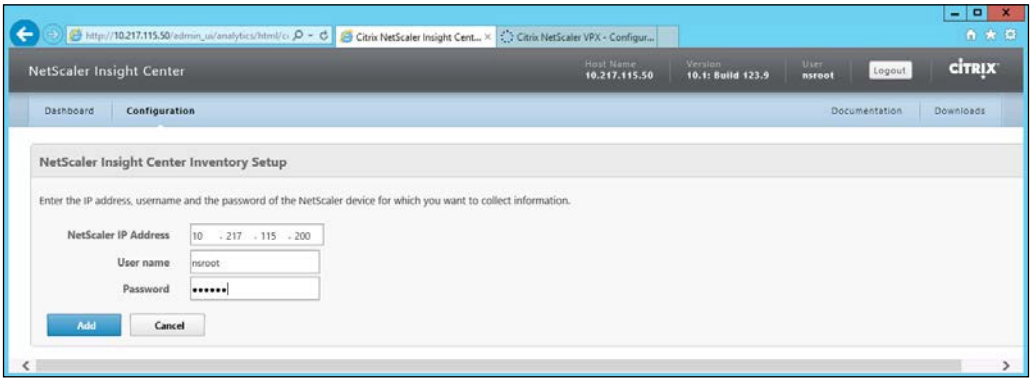

To enable AppFlow data collection on the NetScaler appliance, perform the following steps:

1. Log in to NetScaler Insight Center.
2. Select the NetScaler appliance's IP address, right-click on it, and select **Enable AppFlow**, as shown in the following screenshot:

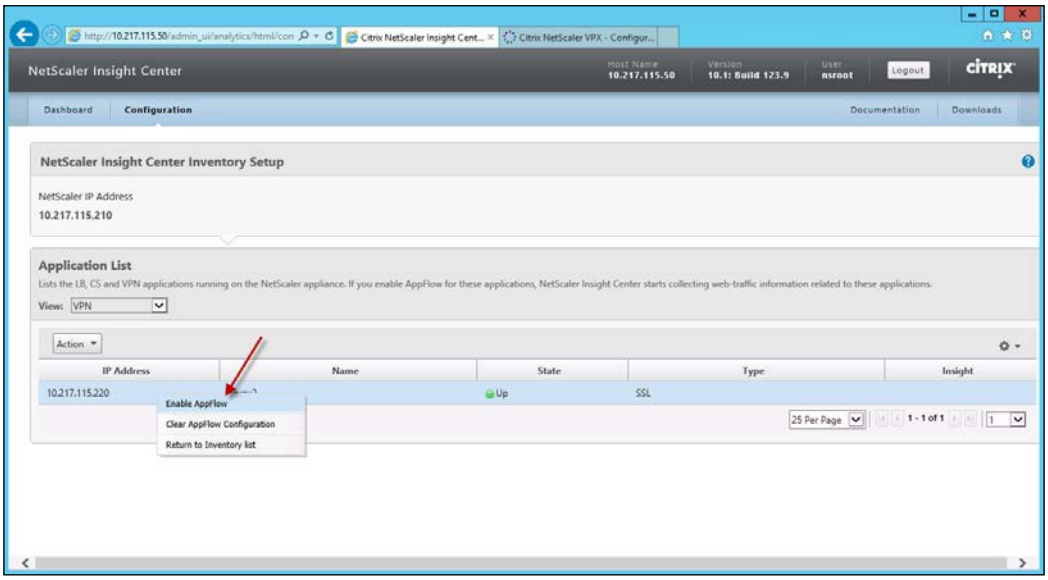

3. For now, we can enable AppFlow for all the traffic, as shown in the following screenshot:

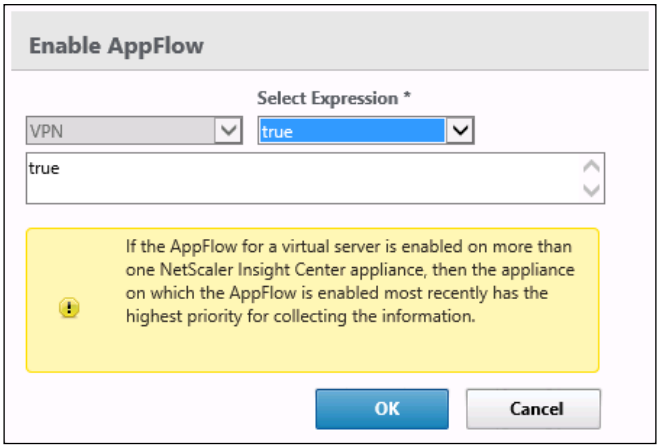

4. When you have successfully enabled AppFlow, the **Insight** column will show **ENABLED**, as shown in the following screenshot:

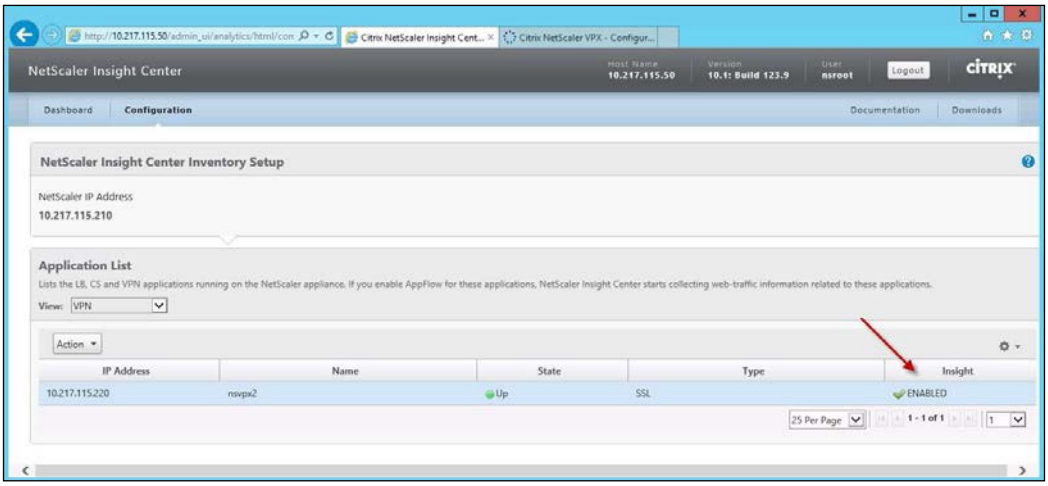

5. Now, you can generate some data on NetScaler and view it in the **Dashboard** tab.

To enable NetScaler Insight Center in Director, perform the following steps:

1. Log in to the Director server and run the `DirectorConfig` tool from a command prompt as follows:

```
cd \inetpub\wwwroot\Director\tools
.\DirectorConfig /configNetScaler
```

2. Enter the NetScaler Insight Center FQDN or IP address, username, password, and HTTP or HTTPS, as shown in the following screenshot:

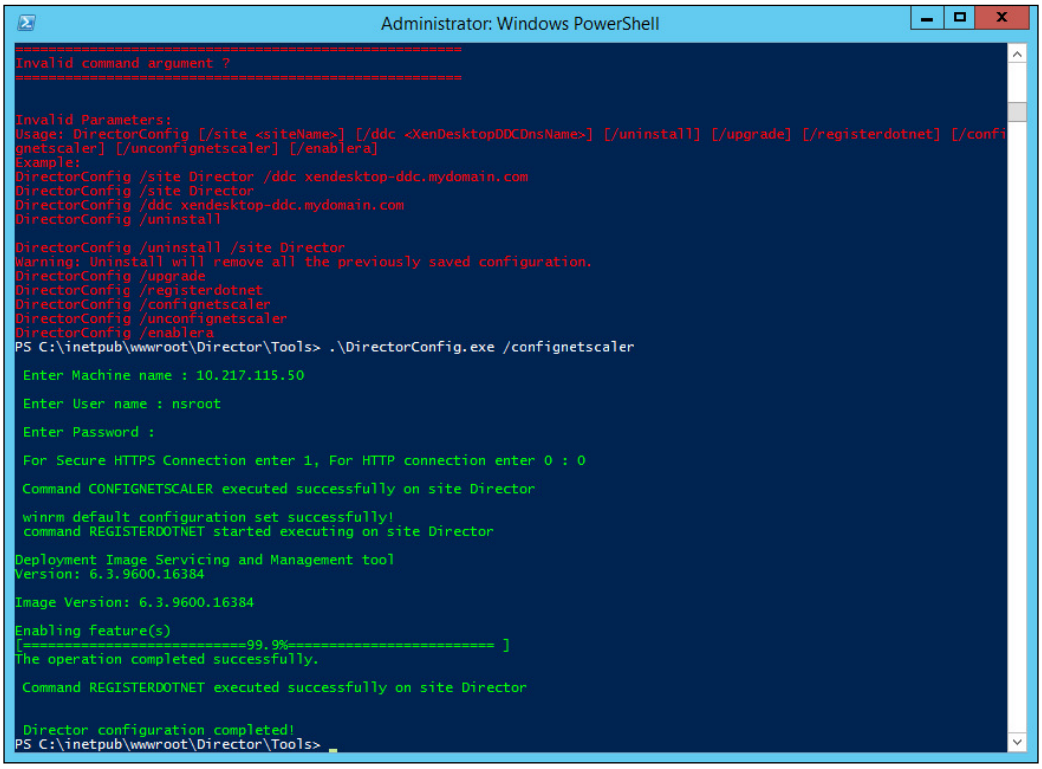

Troubleshooting XenDesktop®

XenDesktop Director can be used to troubleshoot issues. The most efficient way of troubleshooting XenDesktop is to use the activity manager in Director.

Troubleshooting users

If a user is unable to log on or it takes a really long time to log on, or even completely fails on repeated attempts, you might need to diagnose the logon issues.

To troubleshoot the logon issues, perform the following steps:

1. Launch Director.
2. Navigate to **User Details** | **Logon Duration**. Watch the user log off and log on time to view the process. Examine each phase to see where the problem lies.

To reset a user profile, perform the following steps:

1. Launch Director and navigate to **Help Desk**. Choose the machine where the user profile is located.
2. Select the user, click on **Reset Profile**, and then select **Reset**.

When a profile is reset, most of the user profile data is deleted.

To shadow users, perform the following steps:

1. Launch Director.
2. Navigate to **User Details** | **User Session** | **Session Details** | **Shadow**.
3. Instruct the user to click on **Yes**.

Troubleshooting applications

If an application becomes slow or doesn't respond, you may need to troubleshoot the application. To troubleshoot applications, perform the following steps:

1. Launch Director.
2. Navigate to **Activity Manager** | **Applications**. Here, you can terminate the applications and processes or restart a user's virtual desktop. Then, have the user start the application again.

Troubleshooting desktops

If desktop connections start to become problematic, you may need to troubleshoot them. To troubleshoot desktops, perform the following steps:

1. Launch Director.
2. Navigate to **User Details**. View the error that caused the failed connection. You can also restart a user's virtual desktop.

To troubleshoot desktop power states, perform the following steps:

1. Launch Director.
2. Select **Power Control** (only valid for desktop machines). From here, you can choose **Restart**, **Force Restart**, **Shutdown**, **Force Shutdown**, **Suspend**, **Resume**, and **Start Desktop VMs**.

To troubleshoot with Scout in the Citrix Diagnosis Facility, perform the following steps:

1. Log on to the Delivery Controller.
2. Navigate to **Start** | **Citrix** | **Citrix Scout**.
3. Run Scout and upload the information.

Troubleshooting sessions

Even if a session gets disconnected, it is still active and the applications continue to run. The user device no longer communicates with the server. To troubleshoot sessions, perform the following steps:

1. Launch Director.
2. Navigate to **User Details** | **Session Details**. You can terminate applications or processes. You can disconnect sessions and log off users from **Session Control**.

Troubleshooting HDX™

If the quality of the session degrades, you may need to troubleshoot HDX.
To troubleshoot HDX, perform the following steps:

1. Launch Director.
2. Navigate to **User Details** | **HDX**.

Troubleshooting Personal vDisks

You may need to troubleshoot Personal vDisks for the users.

To diagnose Personal vDisks, perform the following steps:

1. Log on to the desktop machine that you want to monitor.
2. From a command prompt, run the Personal vDisk `Diag` tool:

   ```
   cd \Program Files\Citrix\Personal vDisk\bin
   CtxPvdDiag.exe
   ```

3. Several reports can be selected. The reports show the changes that are made to the filesystem along with user data and applications.

To reset a Personal vDisk, perform the following steps:

1. Launch Director. Choose the desktop machine.
2. Navigate to **User Details** | **Personalization** | **Reset Personal vDisk** | **Reset**. If the reset is successful, the status will change to **Running**.

Resetting the Personal vDisk reverts the settings to their factory defaults and all the data is deleted including the applications. However, profile data is retained.

Third-party tools

XenDesktop Director provides monitoring; however, it doesn't do everything, so it is best complemented by third-party tools that specialize in management and monitoring. There are several vendors who provide robust reporting and monitoring for XenDesktop. Be sure to check the Citrix Ready catalog (`www.citrixready.com`) for Citrix-validated solution vendors.

Some third-party vendors who provide additional value to XenDesktop Director are as follows:

- **Lakeside Software**: `www.lakesidesoftware.com`
- **RES Software**: `www.ressoftware.com`
- **eG Innovations**: `www.eginnovations.com`
- **ca Technologies**: `www.ca.com`
- **Comtrade**: `www.comtrade.com`
- **Goliath Technologies**: `www.goliathtechnologies.com`

I will briefly mention some of the additional benefits of third-party solutions. After you have implemented XenDesktop, it will be critical to maintain a consistent high performance for desktops and applications. You will need a tool that characterizes the environment and usage patterns of users. With some of these tools, you can forecast the design and utilization of the virtual infrastructure. Some of these tools go as deep as predicting server compute, memory, storage, and IOPS based on usage patterns. It is all about the quality of the user experience in VDI; so, one vendor named Lakeside Software actually has an end user experience optimization tool that reports on how well the end user experience is evolving and what the areas of impact are.

Summary

Throughout the book, you have been managing the XenDesktop Site using Studio. In this chapter, you also learned additional topics, such as, how to manage, monitor, and troubleshoot XenDesktop. In the next chapter, we will discuss VDI in the cloud.

15
VDI in the Cloud

Now that I have made your understanding of virtualization cloudy, some terms have emerged to help guide you through the conceptualization of these networks into frameworks that can be easily referred to. They are the **private cloud**, **public cloud**, and **hybrid cloud**. Honestly, a cloud is just a rack of equipments. However, if this nomenclature helps us to better understand, talk about, and put things into nice, neat little boxes, then why not go with it.

In this chapter, we will discuss the following topics:

- Virtualization in the cloud
- Private cloud
- Public cloud
- Hybrid cloud
- Personal cloud
- Your cloud

Understanding virtualization in the cloud

Virtualization in the cloud is sort of an oxymoron because everything in the cloud is already virtualized, at least for the most part. If the cloud has an operating system, it is probably already running in the cloud. VDI in the cloud makes sense because users don't really care from where their virtual desktops or applications are being delivered as long as the usability is good.

Many of the Citrix products can be run as virtual machines on supported Hypervisors, such as XenServer, VMware, or Hyper-V. So, it makes sense to conceptualize that you can run many Citrix products in your private cloud or even in the public cloud, such as Amazon Web Services, Rackspace, or SoftLayer.

You can, and it has been done, build a complete XenDesktop desktop and application delivery solution using Citrix products in the public cloud. However, what you will find is that running a solution in the cloud can be costly in terms of money and performance.

Many of the largest public cloud service providers run Xen or a customized version of Xen. As you know, Xen (pronounced Zen) was acquired by Citrix, which they turned into XenServer, a commercial-grade enterprise product. Xen is still open source and supported by Citrix. Although around the time Citrix acquired Xen, a new open source Hypervisor emerged as part of the Linux distribution called KVM. Many of the cloud service providers utilize Hypervisors for their virtualization base operating system.

A cloud doesn't have to be all virtual. Some cloud providers such as SoftLayer, which is now an IBM company, will sell or rent you a rack of your own physical servers that you will most likely install a Hypervisor on. Service providers will also rent you a rack of the physical equipment in their facility known as a **colocation** or **COLO** arrangement. A couple of reasons why some companies might opt for this are better security and higher performance. Whether additional security and performance are perceived or actually true is a matter of debate, but it helps some managers to sleep better at night—a real tangible benefit.

Note that just because you can install something that was virtualized for use with XenServer doesn't mean that it will necessarily work in a public cloud provider that is running Xen. An example that applies to XenDesktop is that of Amazon Web Services or AWS. The Amazon Cloud is the biggest as of this writing and is built on a custom version of open source Xen. As a result, we just can't load XenDesktop up there and expect it to run. However, some integration work has been undertaken between Citrix and Amazon so that XenDesktop now runs in the Amazon Cloud. XenDesktop also runs in the Windows Azure cloud. For other cloud providers, you would likely need to rent your own physical rack of servers, a COLO, from them and install XenDesktop on the Hypervisor of your choice—XenServer perhaps.

Private cloud

A private cloud is your own data center, irrespective of whether it is located in your company's premises or in a COLO facility. You don't even need to have a virtualization software installed here, but my guess is you will. A private cloud is an infrastructure that you have a primary responsibility for, including the owning of and depreciation of assets. Keep in mind that you also need to have an IT staff that maintains this equipment. Private clouds are expensive, but they are more secure and you know exactly what is running in them.

Public cloud

A public cloud is any service provider that provides hosting or **Infrastructure as a Service** (**IaaS**). Public clouds were spurred on by the emergence of virtualization Hypervisors. You can rent virtual machines and virtual servers in the cloud for a fraction of the cost of buying and maintaining your own equipment. As more and more companies are realizing the benefits of a public cloud, they are moving some resources to the cloud. One of the caveats, of course, is that you don't really know who you are sharing your infrastructure with. Some companies have been slow to move to the public cloud as they are apprehensive about its security and performance.

Hybrid cloud

A hybrid cloud is a combination of the private and public cloud infrastructures, as shown in the following figure:

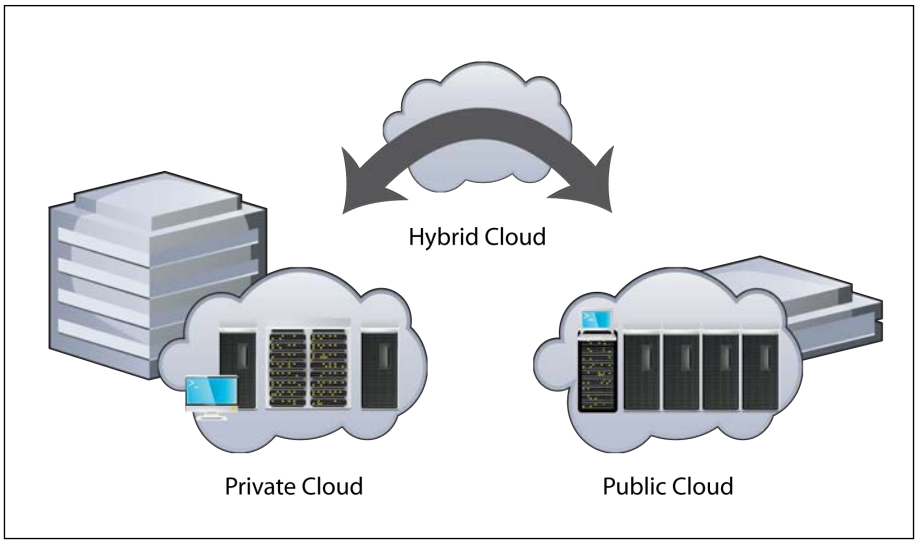

Many companies are choosing to keep a private cloud as the main infrastructure for security and performance reasons, while extending that into the public cloud using CloudBridge so that they have some sense of disaster recovery, failover, and extra capacity when they have a burst of data.

As an example, to implement a VDI solution in the cloud, it is possible to rent physical equipment from a colocation or COLO facility or directly from a cloud provider, such as SoftLayer or Rackspace. It is then possible to install XenServer and XenDesktop frontended by Citrix NetScaler Gateway and to use CloudBridge in order to bridge to a public cloud for high availability and disaster recovery. The network diagram for this type of solution would be strikingly similar to the one we reviewed in *Chapter 10*, *Application Delivery*. The network diagram to implement a VDI solution in the cloud is as follows:

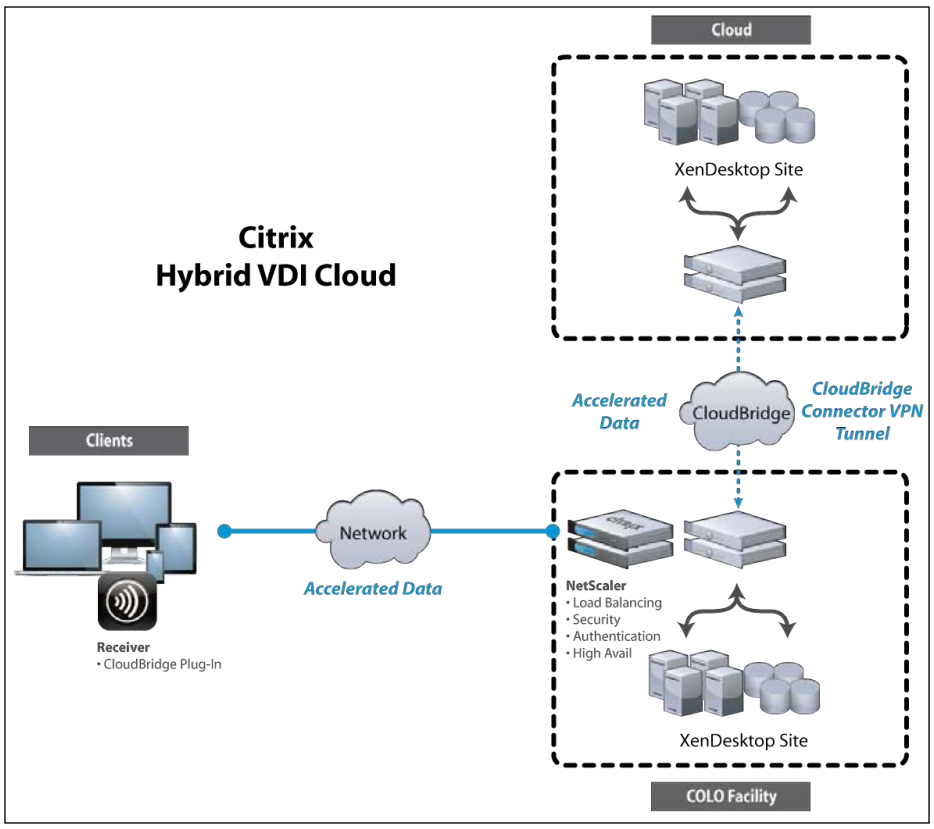

Personal cloud

The concept of a personal cloud is the one where you can run a Hypervisor on your personal device such as a laptop. There are other solutions that also use a dual-boot system, where you can boot onto the operating system of your choice. Running multiple OSes on a client machine is the concept of a personal cloud.

Your cloud

Don't let the industry movement persuade you into thinking that you must have a private, public, or hybrid cloud. Take a look at what is practical in terms of providing quality service to your users. Building a VDI solution for 100 users is very different from building a VDI solution for 10,000 users. Some or all of these may not work well in the cloud. Always do a Proof of Concept before you get ready to commit to a production environment. Although VDI is a compelling solution, it is still in its infancy, and improvements are continually being made. Follow what some influencers are saying about VDI including Brian Madden and Douglas Brown, as they know all about the pitfalls of VDI and aren't afraid to talk about them.

Brian Madden's views on VDI can be found at `http://www.brianmadden.com`.

Douglas Brown's views on VDI can be found at `http://www.dabcc.com/`.

Summary

In this chapter, we discussed virtualization as it relates to the cloud, including private, public, hybrid, and personal clouds.

Virtualization is an exciting concept. One thing I've learned in this business is that there is never a dull moment. Just when I thought that networking and data centers were old news and nothing much could be extracted in terms of value, along came virtualization, and this radically changed the game. As network, compute, and storage infrastructure continues to converge and miniaturize to become more powerful, virtualization and software machines will be able to perform more. Virtualization, **software-defined data center** (**SDDC**), and networking overlays are the next big waves of innovation for our industry, even though the last time I checked, there were still only seven layers to the OSI model, and everyone seems to need data processing.

As you look towards the future and plan your VDI deployment, take your time and make sure that you extract an appropriate amount of value for what you are going to put into it. I had a conversation recently with a virtualization storage startup guy and at the end of it, I told the guy, "Hey, why don't you just buy everyone a cheap laptop; it's cheaper and does the same job, if not better." In other words, we have a long way to go, but it is a road worth travelling.

A

Creating a Domain Certificate Authority

The domain certificate authority should be separate from the XenDesktop servers and should be running on Windows Server 2008 R2 SP2 or Windows Server 2012 R2. In our example in *Chapter 2*, *Installing XenDesktop®*, we used the Active Directory domain controller running on Windows Server 2012 R2 in our lab.

The steps to create a domain certificate authority are as follows:

1. Log in to your system, and from the **Start** menu, navigate to **Administrative Tools** | **Server Manager** as shown in the following screenshot:

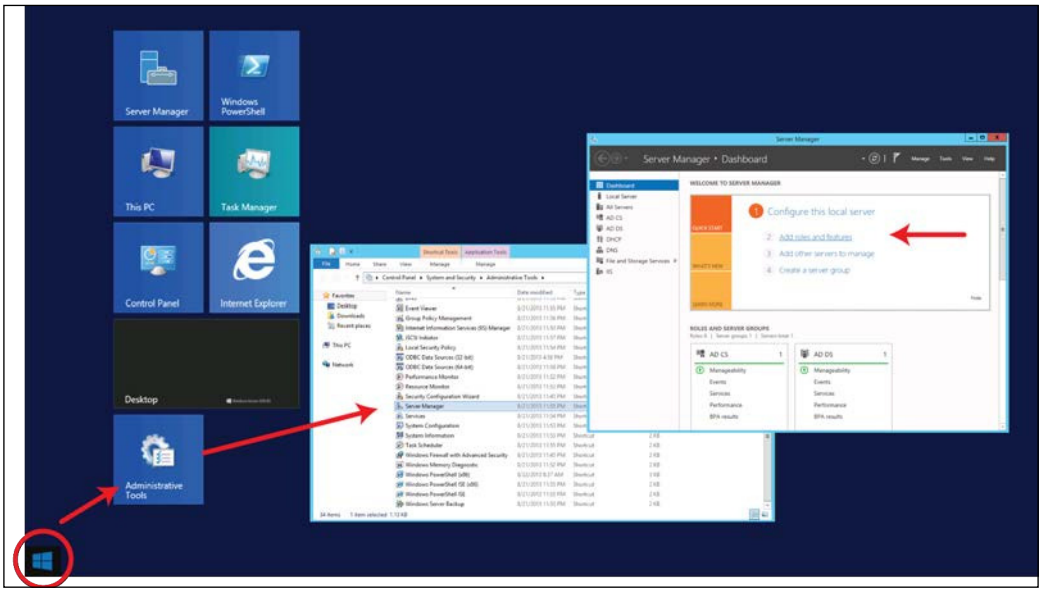

2. Navigate to **Roles | Add Roles**.
3. Select **Active Directory Certificate Services** and click on **Next**, as shown in the following screenshot:

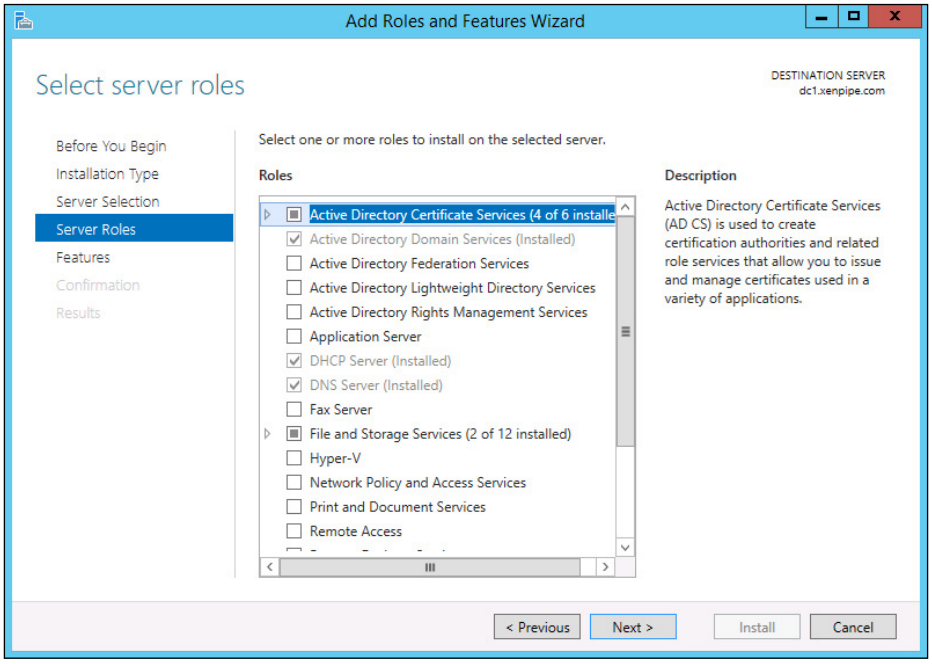

4. Select **Certificate Authority** and **Certificate Authority Web Enrollment** and then click on **Next**. Select **Add Required Role Services** and click on **Next**.
5. Select **Enterprise** and then click on **Next**.
6. Select **Root CA** and click on **Next**.

An enterprise root CA is a top-level CA in a certification hierarchy which requires the Active Directory domain service. It self signs its own CA certificate and uses a group policy to publish this certificate to the trusted root certification authorities store of all the servers and workstations in the domain.

7. Choose **Create a new private key** and click on **Next**.
8. Select **Set Fully Qualified Domain Name** and click on **Next**.
9. Next, select **Set validity period** and click on **Next**.

10. Then, choose **Install additional Role services for IIS** and click on **Next**.
11. Select **Install**.

You can see a demo of the aforementioned steps by following the video tip at `http://www.citrix.com/tv/#videos/7971`.

The preceding steps are detailed for you so that you can get your lab up and running quickly. Designing, building, and supporting a CA in production is a big undertaking. The following are some links to help you get started:

- `http://technet.microsoft.com/en-us/library/cc770827.aspx`
- `http://msdn.microsoft.com/en-us/library/ms755466(v=vs.85).aspx`

B

XenDesktop® Policy Settings Reference

Everything related to the control of sessions in XenDesktop is done through policies. Policies contain settings that are applied when the policy is active or allowed. You can configure the settings using Citrix Studio. You can also use the Active Directory Group Policy Management editor.

Many of these policies are merely references sourced from the Citrix XenDesktop documentation located at `http://support.citrix.com/proddocs/topic/xendesktop-71/cds-policies-rules-wrapper-rho.html`.

To configure XenDesktop policies using Citrix Studio, perform the following steps:

1. Log in to the Delivery Controller.
2. Move the mouse pointer to the lower-left corner of the screen.
3. Click on the **Start** menu and select the down arrow.

4. Under **Apps**, select **Citrix Studio**, as shown in the following screenshot:

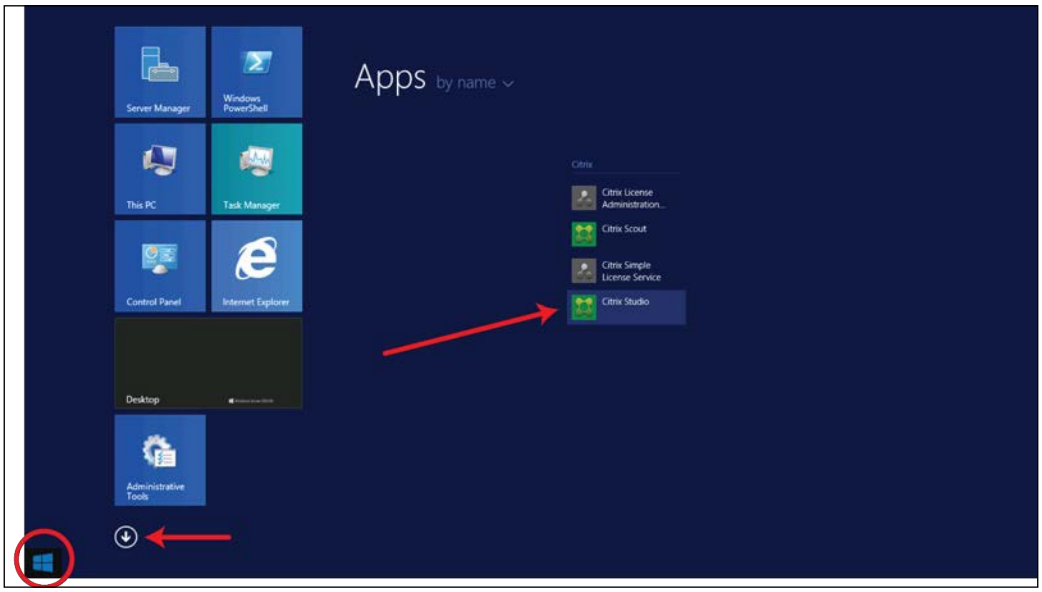

5. When Studio is launched, select the **Policy** node in the left-hand side navigation menu.
6. Policies will be listed in the center of the screen.
7. To create a new policy, click on **Create Policy** in the right-hand side navigation menu.

To configure XenDesktop policies using Microsoft Group Policy Management editor, perform the following steps:

1. Log in to the Active Directory domain controller.
2. Move the mouse pointer to the lower-left corner of the screen.
3. Click on the **Start** menu and select the down arrow.
4. In the **Search** window, search for `gpedit.msc` and launch it.
5. When GPEdit is launched, navigate to **User Configuration** | **Citrix Configuration** in the left-hand side navigation menu.
6. Policies will be listed in the center of the screen.

7. To create a new policy, click on **New** in the upper navigation menu.
8. Follow the **New Policy** wizard.

Audio policies

The following table lists the policies that apply to the audio settings:

Policy	Function	Default	Use with policy
Audio Plug-n -Play	This controls whether or not to allow the use of multiple audio devices to record and play sound.	**Allowed**	
Client microphone redirection	This controls whether or not to allow audio input from microphones on the user device.	**Prohibited**	
Audio quality	This controls the audio quality on the user device. The available options are **Low (16 Kbps)**, **Medium (64 Kbps)**, and **High (1.3 Mbps)**. If you record and play the audio, the bandwidth is doubled.	**High**	Audio redirection bandwidth limit
Client audio redirection	This allows or prevents the applications hosted on the server to play sounds through a sound device installed on the user device. This setting also allows or prevents the users from recording an audio input.	**Allowed**	Audio redirection bandwidth limit
Audio over UDP Real-time Transport	This allows or prevents the transmission and receipt of audio between the host and user device over RTP using UDP. When disabled, it uses TCP.	**Allowed**	
Audio over UDP	This allows or prevents audio over UDP on the server.	**Allowed**	
Audio UDP port range	This specifies the range of the ports used by the Delivery Agent to exchange audio packet data with the user device.	**16500-16509**	

Bandwidth policies

The following table lists the policies that apply to the bandwidth settings:

Policy	Function	Default	Use with policy
Audio redirection bandwidth limit	This specifies the maximum allowed bandwidth, in kilobits per second, for playing or recording audio in a user session	No maximum is set	Overall session bandwidth limit
Audio redirection bandwidth limit percent	This specifies the maximum allowed bandwidth limit for playing or recording audio as a percent of the total session bandwidth	No maximum is set	Overall session bandwidth limit
Clipboard redirection bandwidth limit or limit percent	This limits the bandwidth for a cut and paste operation and specifies the maximum allowed bandwidth, in kilobits per second, for data transfer or a percent of the total bandwidth between a session and the local clipboard	No maximum is set	Overall session bandwidth limit
COM port redirection limit or limit percent	This specifies the maximum allowed bandwidth, in kilobits per second or as a percent of the total bandwidth, for accessing a COM port in a client connection	No maximum is set	Overall session bandwidth limit
File redirection bandwidth limit or limit percent	This specifies the maximum allowed bandwidth, in kilobits per second or as a percent of the total bandwidth, for accessing a client drive in a user session	No maximum is set	Overall session bandwidth limit
HDX MediaStream Multimedia Acceleration bandwidth limit or percent	This specifies the maximum allowed bandwidth in order to deliver the streaming of an audio and video using HDX MediaStream multimedia acceleration	No maximum is set	Overall session bandwidth limit
LPT port redirection bandwidth limit or percent	This specifies the maximum allowed bandwidth, in kilobits per second or as a percent of total bandwidth, for print jobs using an LPT port in a single user session	No maximum is set	Overall session bandwidth limit

Policy	Function	Default	Use with policy
Overall session bandwidth limit	This specifies the total amount of bandwidth available in kilobits per second for the user sessions	No maximum is set	Overall session bandwidth limit
Printer redirection bandwidth limit or percent	This specifies the maximum allowed bandwidth, in kilobits per second or as a percent of the total bandwidth, for accessing the client printers in a user session	No maximum is set	Overall session bandwidth limit
TWAIN device redirection bandwidth limit or percent	This specifies the maximum allowed bandwidth, in kilobits per second or as a percent of the total bandwidth, for controlling the TWAIN imaging devices from the published applications	No maximum is set	Overall session bandwidth limit
Client USB device redirection bandwidth limit or percent	This specifies the maximum allowed bandwidth, in kilobits per second, for the redirection of USB devices to and from the client	No maximum is set	Overall session bandwidth limit

Redirection policies

The following table lists the policies that apply to the redirection settings:

Policy	Function	Default	Use with policy
Auto connect client drives	This allows or prevents the automatic connection of client drives when users log on.	**Allowed**	• Client drive redirection • Client floppy drives • Client optical drives • Client fixed drives • Client network drives • Client removable drives
Client clipboard redirection	This allows or prevents the clipboard on the user device to be mapped to the clipboard on the server.	**Allowed**	Clipboard redirection bandwidth limit

Policy	Function	Default	Use with policy
Client drive redirection	This enables or disables file redirection to and from the drives on the user device.	**Enabled**	• Client floppy drives • Client optical drives • Client fixed drives • Client network drives • Client removable drives
Client COM port redirection	This controls whether or not the user devices attached to the local COM ports are available in a session.	**Prohibited**	COM port redirection bandwidth limit and COM port redirection bandwidth limit percent
Client LPT port redirection	This controls whether or not the client printers attached to the local LPT ports are available in a session.	**Prohibited**	LPT port redirection bandwidth limit and LPT port redirection bandwidth limit percent
Client fixed drives and Client drive redirection	This allows or prevents users from accessing or saving files to fixed drives on the user device.	**Allowed**	Client drive redirection and Auto connect client drives
Client floppy drives and Client drive redirection	This allows or prevents users from accessing or saving files to floppy drives on the user device.	**Allowed**	Client drive redirection and Auto connect client drives
Client network drives and Client drive redirection	This allows or prevents users from accessing and saving files to network (remote) drives through the user device.	**Allowed**	Client drive redirection and Auto connect client drives
Client optical drives and Client drive redirection	This allows or prevents users from accessing or saving files to CD-ROM, DVD-ROM, and BD-ROM drives on the user device.	**Allowed**	Client drive redirection and Auto connect client drives

Policy	Function	Default	Use with policy
Client removable drives and Client drive redirection	This allows or prevents users from accessing or saving files to USB drives on the user device.	**Allowed**	Client drive redirection and Auto connect client drives
Client TWAIN device redirection and TWAIN compression redirection	This controls whether or not the users' TWAIN devices, such as scanners and cameras ,are available in a session and also controls the compression of image data transfers.	**Allowed**	• TWAIN compression level • TWAIN device redirection bandwidth limit • TWAIN device redirection bandwidth limit percent
TWAIN compression level	This specifies the level of compression of image transfers from a client to a server.	**Medium**	
Client USB device redirection	This controls whether or not USB devices are available in a session.	**Disabled**	Client USB device redirection rules
Client USB device redirection rules	This specifies the redirection rules for USB devices. Refer to CTX119722.	None specified	Client USB device redirection
Use asynchronous writes	This improves the speed of writing and copying files to a client disk over a WAN.	**Disabled**	Client drive redirection
Host to Client redirection	This enables or disables the file type associations for URLs and some media content to be opened on the user device. When disabled, content opens on the server.	**Disabled**	
Preserve client drive letters	This enables or disables the mapping of the client drives to the same drive letter in the session.	**Disabled**	Client drive redirection

Desktop UI policies

The following table lists the policies that apply to the desktop UI settings:

Policy	Function	Default	Use with policy
Aero Redirection	This redirects the processing of the Windows Aero interface to the **graphics processing unit** (**GPU**) of the user device rather than that of the server when enabled.	**Enabled**	
Aero Redirection Graphics Quality	This specifies the quality of the graphics used for Aero redirection. The options available are **Lossless**, **Low**, **Medium**, and **High**.	**High**	
Desktop wallpaper	This controls whether or not a desktop wallpaper is used in the users' sessions.	**Enabled**	
Menu Animation	This allows or prevents menu animation in user sessions.	**Allowed**	
View window contents while dragging	This controls whether or not a window's contents are viewable when being dragged.	**Allowed**	

Graphics and multimedia policies

The following table lists the policies that apply to the graphics and multimedia settings:

Policy	Function	Default	Use with policy
Target Frame Rate	This specifies the maximum number of frames per second that are sent to the user device from the virtual desktop.	**24 fps**	
Visual Quality	This controls the visual quality of the images that are displayed on the user device.	**Medium**	
Minimum image quality	This specifies the minimum acceptable image quality for Adaptive Display. The lesser the compression used, the higher the quality of the images displayed.	**Normal**	

Policy	Function	Default	Use with policy
Moving image compression	This specifies whether or not **Adaptive Display** is enabled. **Adaptive Display** automatically adjusts the image quality of videos and transitional slides in slide shows based on the available bandwidth. With **Adaptive Display** enabled, users should see smooth running presentations with no reduction in the quality.	**Enabled**	
Progressive compression level	This provides a less detailed but faster initial display of images.	**Disabled**	• Progressive compression threshold value • Lossy compression level • Progressive heavyweight compression
Progressive compression threshold value	This represents the maximum bandwidth, in kilobits per second, for a connection to which progressive compression is applied.	**2147483647 kilobits per second**	
Extra color compression	This enables or disables the use of extra color compression on images delivered over the client connections that are limited in bandwidth, thus improving the responsiveness by reducing the quality of the displayed images.	**Disabled**	
Extra color compression threshold	This represents the maximum bandwidth, in kilobits per second, for a connection below which extra color compression is applied. If the client connection bandwidth drops below the set value, extra color compression, if enabled, is applied.	**8192 kilobits per second**	
Heavyweight compression	This enables or disables the reduction of the bandwidth beyond progressive compression without losing image quality by using a more advanced, but more CPU-intensive, graphical algorithm.	**Disabled**	Progressive compression level

Policy	Function	Default	Use with policy
Lossy compression value	This controls the degree of lossy compression used on images delivered over the client connections that are limited in bandwidth. In such cases, displaying the images without compression can be slow.	**Medium**	Progressive compression level
Lossy compression threshold value	This represents the maximum bandwidth, in kilobits per second, for a connection to which lossy compression is applied.	**2147483647 kilobits per second**	
Flash default behavior	This controls whether or not Flash content can be rendered in sessions. **Enabled** indicates that Flash redirection is allowed on the client device. **Blocked** indicates that Flash redirection is not allowed. **Disabled** indicates that Flash redirection is rendered on the server.	**Enabled**	• Flash URL compatibility list • Enable HDX MediaStream for Flash on the user device
Flash server-side content fetching URL list and Flash URL compatibility list	This specifies the websites whose Flash content can be downloaded to the server and then transferred to the user device for rendering.	Not set	Enable server-side content fetching
Flash background color list	This enables you to set the key colors for the given URLs.	Not set	
Flash event logging	This enables or disables the recording of Flash events in the Windows application event log.	**Enabled**	
Flash intelligent fallback	This enables or disables the automatic attempts to employ the server-side rendering for Flash player instances where the client-side rendering is either unnecessary or provides a poor user experience.	**Enabled**	
Display memory limit	This specifies the maximum video buffer size in kilobytes for the session.	**32768 kilobytes**	Display mode degrade performance

Policy	Function	Default	Use with policy
Display mode degrade performance	This specifies the color depth or resolution that the session degrades to when the session display memory limit is reached.	**Color depth**	Notify user when display mode is degraded
Dynamic Windows preview	This enables or disables the display of seamless windows in flip, flip 3D, taskbar preview, and peek window preview modes.	**Enabled**	
Image caching	This enables or disables the caching of images in sessions. When needed, the images are retrieved in sections to make scrolling smoother.	**Enabled**	
Maximum allowed color depth	This specifies the maximum color depth allowed for a session.	**32 bits per pixel**	Display mode degrade performance
Notify user when display mode is degraded	This displays a brief explanation to the user when the color depth or resolution is degraded.	**Disabled**	
Queuing and tossing	This discards the queued images that are replaced by another image.	**Enabled**	
Multimedia conferencing	This allows or prevents support for video conferencing applications.	**Enabled**	Windows media redirection
Windows media redirection	Allowing this setting increases the quality of audio and video rendered from the server to a level that compares with the audio and video played locally on a client device. The server streams multimedia to the client in the original, compressed form and allows the client device to decompress and render the media.	**Allowed**	
Windows media redirection buffer size	This specifies a buffer size from 1 to 10 seconds for multimedia acceleration.	**5 secs**	
Windows media redirection buffer size use	This enables or disables using the buffer size that is specified in the **Windows media redirection buffer size** setting.	**Disabled**	

Caching policies

The following table shows the policy that applies to the caching settings:

Policy	Function	Default	Use with policy
Persistent cache threshold	This caches bitmaps on the hard drive of the user device. This enables the reuse of large, frequently-used images from previous sessions.	Bandwidth reaches 3,000,000 kilobits per second	

Multistream traffic policies

The following table lists the policies that apply to the multistream traffic settings:

Policy	Function	Default	Use with policy
Multi-Port Policy	This specifies the ports for the ICA traffic across multiple connections and establishes network priorities	**TCP port 2598**	Multi-Stream (Computer)
Multi-Stream (Computer)	This enables support for the multistream connections on the server	**Disabled**	
Multi-Stream (User)	This enables support for multistream on the user device	**Disabled**	Multi-Stream (Computer)

Printing policies

The following table lists the policies that apply to the printing settings:

Policy	Function	Default	Use with policy
Auto-create client printers and **Client printer redirection**	This specifies the client printers that are autocreated. This setting overrides the default **Client printer auto-creation** settings.	**Auto-created**	
Auto-create generic universal printer	This enables or disables the autocreation of the generic Citrix Universal Printer object for the sessions where a user device that is compatible with universal printing is in use.	Not autocreated	• Universal print driver usage • Universal driver preference

Policy	Function	Default	Use with policy
Universal print server enable	This enables or disables the Universal Print Server.	**Disabled**	
Universal print server print data stream (CGP) port	This specifies the TCP port number used by the Universal Print Server print data stream **Common Gateway Protocol** (**CGP**) listener.	**Port 7229**	
Universal print server web service (HTTP/ SOAP) port	This specifies the TCP port number used by the Universal Print Server Web service listener for the incoming HTTP/SOAP requests.	**Port 8080**	
Universal print server print stream input bandwidth limit (kbps)	This specifies the upper boundary (in kilobits per second) for the transfer rate of print data that is delivered from each XenApp or XenDesktop print job to the Universal Print Server using CGP.	**0**	
Universal print EMF processing mode	This controls the method of processing the EMF spool file on the Windows user device.	**Spooled directly to printer**	
Universal printing image compression limit	This specifies the maximum quality and the minimum compression level available for images that are printed with the Universal Printer Driver.	**Best quality (lossless compression)**	Universal printing optimization defaults
Universal printing optimization defaults	This specifies the default values for printing optimization when the Universal Printer driver is created for a session.	**Standard Quality**	
Universal printing preview performance	This specifies whether or not to use the **Print Preview** function for the autocreated or generic universal printers.	**Not used**	Universal print driver usage

Policy	Function	Default	Use with policy
Universal printing print quality limit	This specifies the maximum **dots per inch** (**dpi**) that are available for generating a printed output in the session.	**No limit**	
Client printer names	This selects the naming convention used for the autocreated client printers.	Standard printer names	Auto-create client printers
Printer properties retention	This controls the location where the printer properties are stored.		
Direct connections to print servers	This enables or disables the direct connections from the host to a print server for the client printers that are hosted on an accessible network share.	**Enabled**	
Printer driver mapping and compatibility	This specifies the driver substitution rules for the autocreated client printers.	None specified	
Printer properties retention	This specifies whether or not to store the printer properties and where to store them.	System determines	
Client printer redirection	This controls whether or not users can access the printers connected to their user devices.	**Allowed**	Auto-create client printers
Default Printer	This specifies how the default printer on the user device is established in a session.	Users' current printer	
Printer assignments	This provides an alternative to the **Default printer** and **Session printers** settings. Use the individual **Default Printer** and **Session Printers** settings to configure behaviors for a farm, Site, large group, or an organizational unit. Use the **Printer assignments** setting to assign a large group of printers to multiple users.	Users' current printer	

Policy	Function	Default	Use with policy
Printer auto-creation event log performance	This specifies the events that are logged during the printer autocreation process.	Errors and warnings	
Session printers	This specifies the network printers to be autocreated in a session.	None specified	
Automatic installation of in-box printer drivers	This enables or disables the automatic installation of the printer drivers from the Windows in-box driver set or from the driver packages staged on the host using `pnputil.exe /a`.	Installed as needed	
Automatic installation of in-box printer drivers	This controls the installation of the native Windows drivers when automatically creating the client and network printers.		
Universal print driver usage	This controls whether or not to use the Universal Printer Driver.	**Only used if requested driver is unavailable**	Auto-create generic universal printer

ICA® policies

The following table lists the policies that apply to the `ICA/HDX` protocol settings:

Policy	Function	Default	Use with policy
ICA listener connection timeout	This specifies the maximum wait time for a connection using the `ICA` protocol to be completed	**120,000 ms (2 mins)**	
ICA listener port number	This specifies the TCP/IP port number used by the `ICA` protocol	**Port 1494**	
ICA round trip calculation	This turns on the ICA round trip calculations on active connections	**Enabled**	
ICA round trip calculation interval	This specifies the frequency, in seconds, to calculate the ICA round trip	**15 secs**	
ICA round trip calculation for idle connections	This turns on the ICA round trip calculations for idle connections	**Disabled**	

Keep alive policies

The following table lists the policies that apply to the keep alive settings:

Policy	Function	Default	Use with policy
ICA keep alive timeout	This specifies the number of seconds between successive ICA keep-alive messages	**60 secs**	
ICA keep alives	This enables or disables the sending of the ICA keep-alive messages periodically	**Disabled**	Do not use with Session Reliability

Autoreconnection policies

The following table lists the policies that apply to the autoreconnection settings:

Policy	Function	Default	Use with policy
Auto client reconnect	This allows users to resume working from where they were when a connection was broken	**Allowed**	
Auto client reconnect authentication	This requires authentication for automatic client reconnects	**Disabled**	
Auto client reconnect logging	This enables the recording of the reconnections in the event log	**Disabled**	
Auto connect client COM ports	This enables or disables the automatic connection of the COM ports on the user devices when users log on to the Site	**Enabled**	
Auto connect client LPT ports	This enables or disables the automatic connection of the LPT ports on the user devices when users log on to the farm	**Enabled**	

Mobility policies

The following table lists the policies that apply to the mobility settings:

Policy	Function	Default	Use with policy
Allow applications to use the physical location of the client device	This determines whether or not the applications running in a session on a mobile device are allowed to use the physical location of the client device.	**Prohibited**	
Automatic keyboard display	This enables or disables the automatic display of the keyboard on mobile device screens.	**Disabled**	
Remote the combo box	This determines the types of combo boxes that you can display in the sessions on mobile devices. To display the device-native combo box control, set this policy to **Allowed**. When this setting is allowed, a user can change Receiver for an **iOS session** setting to use the Windows combo box.	**Prohibited**	

Session policies

The following table lists the policies that apply to the session settings:

Policy	Function	Default	Use with policy
Disconnected session timer	This determines how long a disconnected, locked workstation can remain locked before the session is logged off	**1440 mins (24 hours)**	
Session connection timer	This enables or disables a timer to determine the maximum duration of an uninterrupted connection between a user device and a workstation	**Disabled**	Session connection timer interval

Policy	Function	Default	Use with policy
Session connection timer interval	This determines, in minutes, the maximum duration of an uninterrupted connection between a user device and a workstation	**1440 mins (24 hours)**	Session connection timer
Session idle timer	This enables or disables a timer to determine how long an uninterrupted user device connection to a workstation will be maintained if there is no input from the user	**Enabled**	Session idle timer interval
Session idle timer interval	This determines, in minutes, how long an uninterrupted user device connection to a workstation will be maintained if there is no input from the user	**1440 mins (24 hours)**	Session idle timer
Session reliability connections	This allows or prevents sessions to remain open during a loss of network connectivity	**Allowed**	• Auto client reconnect authentication • Session reliability timeout
Session reliability port number	This specifies the TCP port number for the incoming session reliability connections	**Port 2598**	
Session reliability timeout	This specifies the length of time, in seconds, that the session reliability proxy waits for a client to reconnect before allowing the session to be disconnected	**180 secs (3 mins)**	• Auto client reconnect authentication • Session reliability timeout
Single Sign-On	This enables or disables the use of a single sign-on when users connect to XenApp servers or published applications	**Enabled**	
Single Sign-On central store	This specifies the UNC path of the single sign-on central store to which users are allowed to connect	Not specified	

Time zone policies

The following table lists the policies that apply to the time zone settings:

Policy	Function	Default	Use with policy
Estimate local time for legacy clients	This enables or disables estimating the local time zone of user devices that send inaccurate time zone information to the server	**Local time zone**	
Use local time of client	This determines the time zone setting of the user session	**User session time zone**	

Load management policies

The following table lists the policies that apply to the load management settings:

Policy	Function	Default	Use with policy
Concurrent logon tolerance	This specifies the maximum number of concurrent logons that a server can accept	**2**	
CPU usage	This specifies the level of CPU usage, as a percentage, at which the server reports a full load	**90%**	
Disk usage	This specifies the disk queue length at which the server reports a 75 percent full load	**8**, when enabled	
Maximum number of sessions	This specifies the maximum number of sessions that a server can host	**100**, when enabled	
Memory usage	This specifies the level of memory usage, as a percentage, at which the server reports a full load	**90%**	
Memory usage base load	This specifies an approximation of the base operating system's memory usage and defines, in MB, the memory usage below which a server is considered to have zero load	**768 MB**	

Delivery Agent policies

The following table lists the policies that apply to the **Virtual Delivery Agent** (**VDA**) settings:

Policy	Function	Default	Use with policy
Controllers	This specifies a space-separated list of controller **Fully Qualified Domain Names** (**FQDNs**) that the Delivery Agent uses to register with a controller when using registry-based registration.	Blank	• Controller SIDs • Enable Auto Update of Controllers when Disabled
Controller registration IPv6 netmask	This allows the administrators to restrict the Delivery Agent to only a preferred subnet (rather than a global IP, if one is registered). This setting specifies the IPv6 address and network where the Delivery Agent will be registered. The Delivery Agent will register only on the first address that matches the specified netmask.	Blank	Only use IPv6 controller registration
Controller registration port	This specifies the TCP/IP port number that the Delivery Agent uses to register with a controller when using a registry-based registration.	**80**	Enable Auto Update of Controllers when Disabled
Controller SIDs	This specifies a space-separated list of controller **Security Identifiers** (**SIDs**) that the Delivery Agent uses to register with a controller when using a registry-based registration.	Blank	Enable Auto Update of Controllers when Disabled
Enable Auto Update of Controllers	This enables the Delivery Agent to register with a controller automatically after installation.	**Enabled**	

Policy	Function	Default	Use with policy
Only use IPv6 controller registration	When enabled, the Delivery Agent registers with the controller using the machine's IPv6 address. When the Delivery Agent communicates with the controller, it uses the following address order: global IP address, **Unique Local Address** (**ULA**), link-local address (if no other IPv6 addresses are available). When disabled, the Delivery Agent registers and communicates with the controller using the machine's IPv4 address.	**Disabled**	
Set GUID	This specifies the **Globally Unique Identifier** (**GUID**) of the XenDesktop Site that the Delivery Agent uses to register with a controller when using an Active Directory-based registration.	Blank	Enable Auto Update of Controllers when Disabled

HDX™ 3D policies

The following table lists the policies that apply to the HDX 3D settings:

Policy	Function	Default	Use with policy
Enable lossless	This specifies whether or not the users can enable or disable lossless compression using the image quality configuration tool	**Disabled**	
HDX3DPro Quality Settings	This specifies the minimum and maximum values that define the range of image quality adjustment available to users in the image quality configuration tool	Not set	

C

Creating Self-signed Certificates for NetScaler Gateway™

The easiest and most inexpensive way to test a NetScaler Gateway SSL function is to create a self-signed **Certificate Authority** (**CA**) root certificate and a public-facing server certificate. In this Appendix, using the tools in NetScaler, you will use a self-signed root certificate to sign the public-facing server certificate. You can then import these root and public-facing server certificates' chain of trust into the client's browser for testing.

Keep in mind that self-signed certificates are only good for a Proof of Concept lab or testing. Self-signed certificates are not widely trusted by your client's browsers. When you move to production, make sure that you purchase a valid signed certificate.

To remain flexible in our deployment, we will also use a wildcard certificate, for example `*.xenpipe.com`, so that we can create and use different subdomains (for example, `ng.xenpipe.com`) and still have our certificates work.

Enabling SSL on NetScaler Gateway™

To enable SSL on NetScaler Gateway, perform the following steps:

1. Log in to NetScaler Gateway.
2. Navigate to **Traffic Management** | **SSL**.

3. Right-click and then click on **Enable Feature**, as shown in the following screenshot:

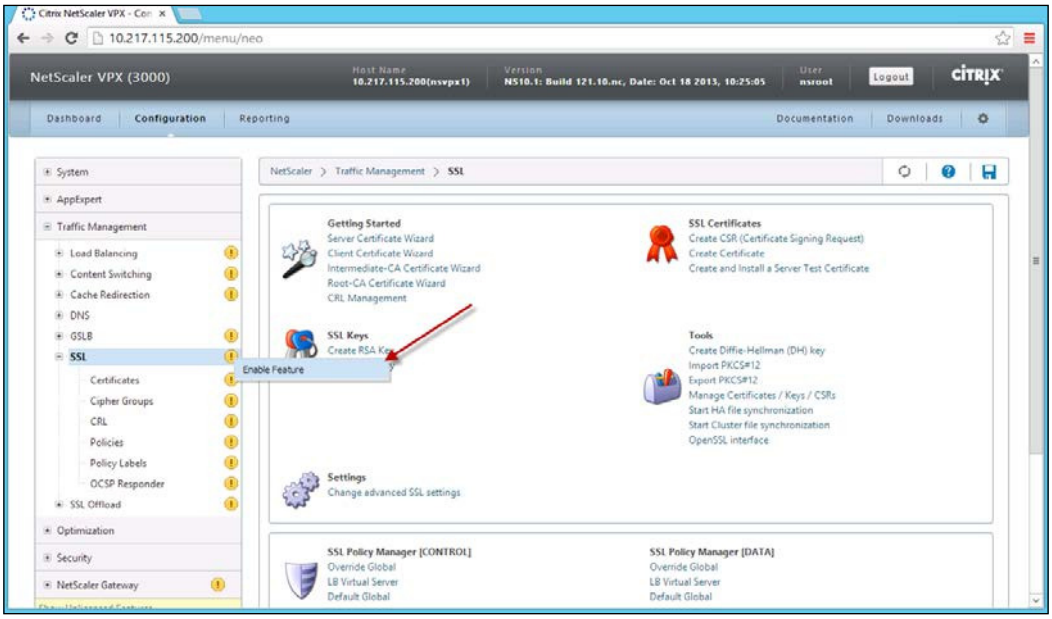

Creating a self-signed root CA certificate

First, you will create a self-signed root CA certificate to be used to sign the public-facing server certificate. To create a self-signed root CA certificate, perform the following steps:

1. Log in to NetScaler Gateway.

The default username and password are `nsroot` and `nsroot`, respectively.

2. Navigate to **Traffic Management | SSL** and select **Create RSA Key**, as shown in the following screenshot:

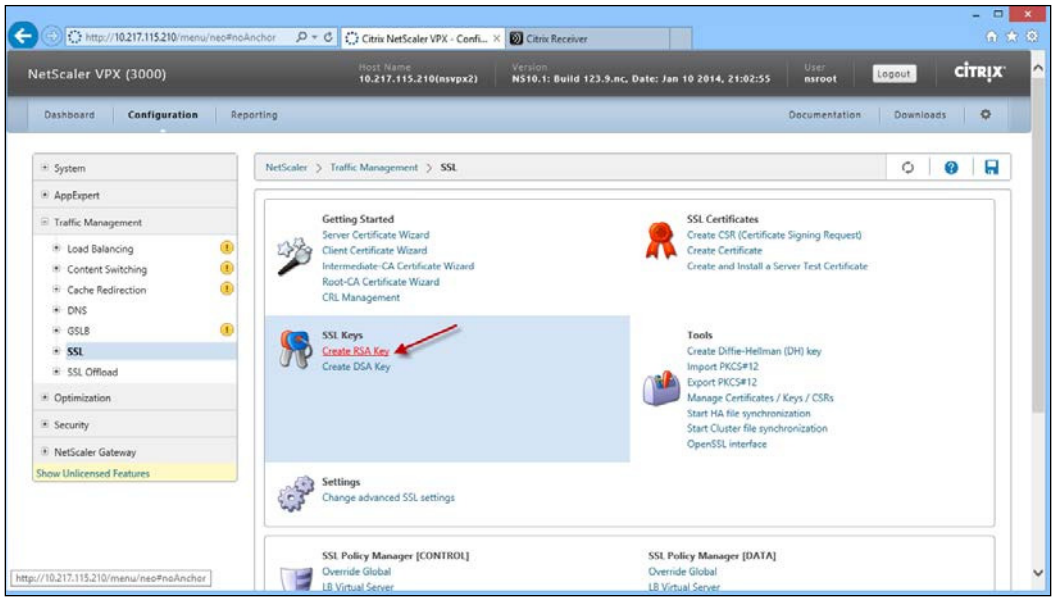

3. Create the RSA key by entering the **Key Filename** and **Key Size(bits)** fields. Set **Public Exponent Value** as **F4** and **Key Format** as **PEM**. Then, click on **OK**, as shown in the following screenshot:

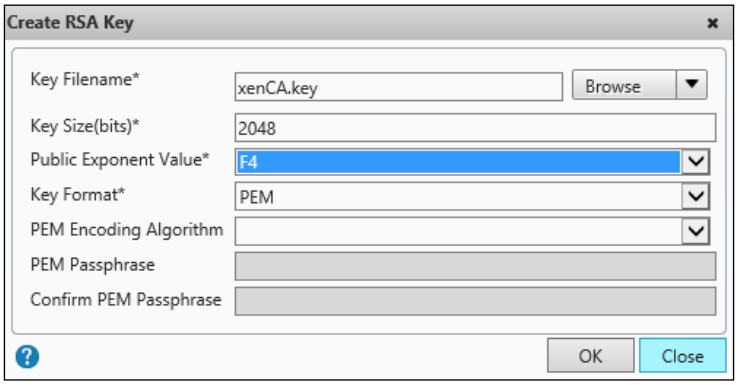

4. Under **SSL Certificates**, select **Create CSR (Certificate Signing Request)**. Enter the **Request File Name** and **Key Filename fields** that are created in the previous step. Choose **PEM** as the **Key Format**, create a password, enter the values in the **Distinguished Name Fields** section, and then click on **OK**, as shown in the following screenshot.

Common Name *must* be different from the server certificate common name in the next step.

5. Create the certificate. Under **SSL Certificates**, select **Create Certificate**. Enter a name in **Certificate File Name** and choose **Certificate Format** as **PEM**. Select **Root-CA** for **Certificate Type**, enter a name for **Certificate Request File Name**, and choose **PEM** as **Key Format**. Enter **Validity Period (Number of Days)** and **Key Filename** and then click on **OK**, as shown in the following screenshot:

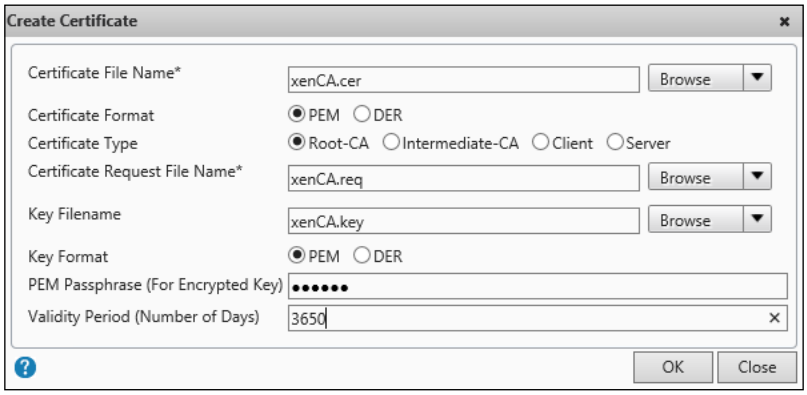

Creating a public-facing server certificate

Next, you will create a public-facing server certificate and will have it signed (self-signed, that is) by the root CA certificate that you have just created in the previous section. To create a public-facing server certificate, perform the following steps:

1. Navigate to **Traffic Management | SSL**. Under **SSL Keys**, select the **Create RSA Key** option, as shown in the following screenshot:

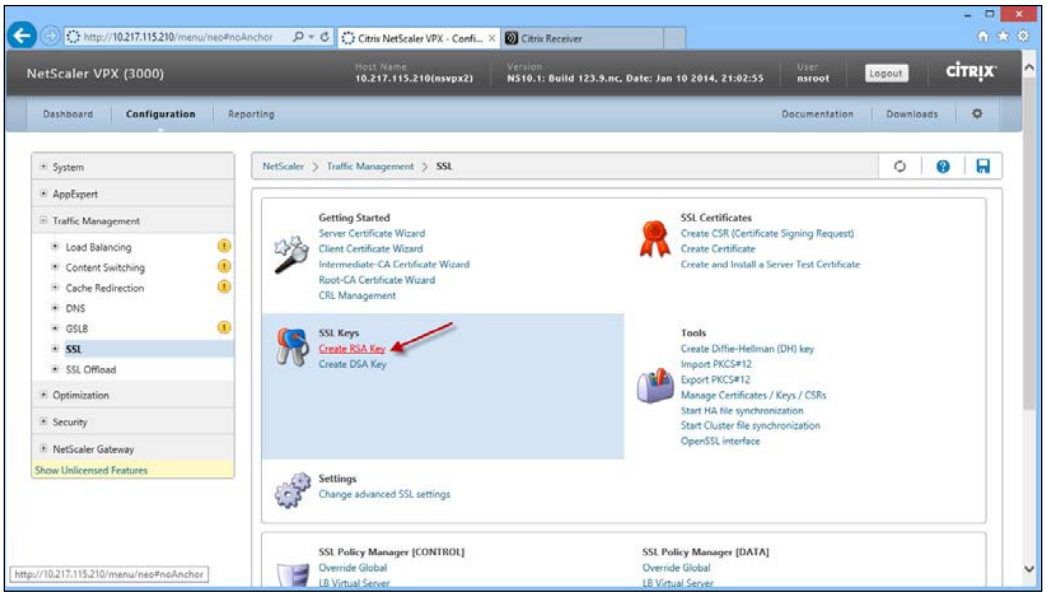

2. First, you create the RSA key by entering **Key Filename**, **Key Size(bits)**, **Public Exponent Value** as **F4**, and **Key Format** as **PEM** and then clicking on **OK**, as shown in the following screenshot:

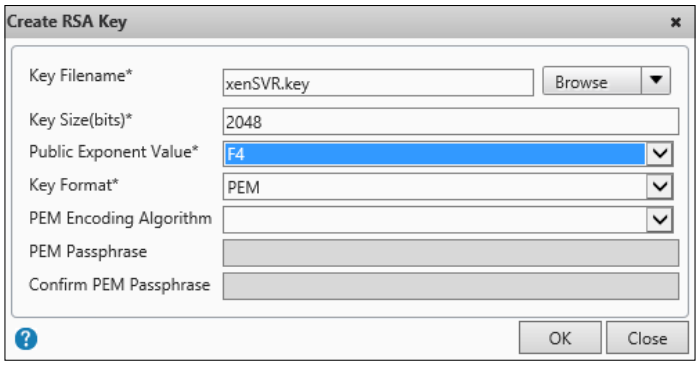

3. Under **Tools**, select **Create CSR (Certificate Signing Request)**. Enter a value in **Request File Name** and **Key Filename** that you have created in the previous step. Choose **PEM** as **Certificate Format**, create a password, and enter the values in the **Distinguished Name Fields** section. Then, click on **OK**, as shown in the following screenshot.

Common Name *must* match **Fully Qualified Domain Name** that is used to access the Site; for example, `ng.xenpipe.com` is what we will use to connect to NetScaler Gateway. In this example, we create a wildcard certificate that will work across all the subdomains; hence, you see `*.xenpipe.com` in the **Common Name** field. If you don't do this, the certificate will not work.

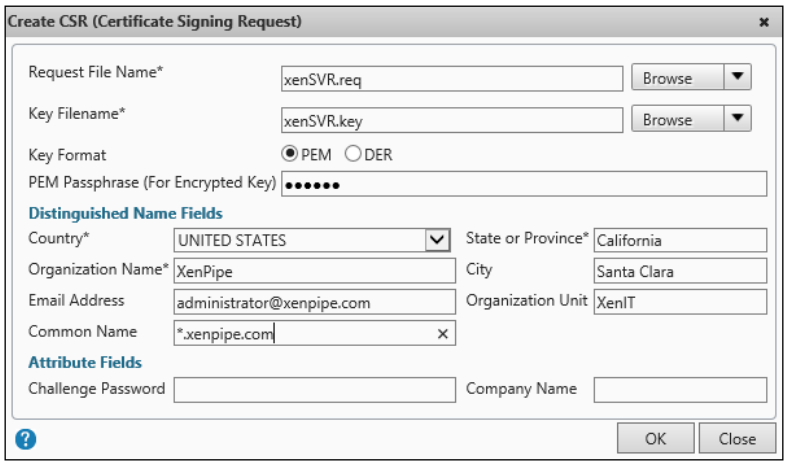

4. Next, we create the server certificate. Under **SSL Certificates**, select **Create Certificate**. Enter a name in **Certificate File Name** and select **Server**. Enter the **Certificate Request File Name**, choose **PEM** as **Certificate Format**, and enter the **Validity Period (Number of Days)**. Enter a name in **CA Certificate File Name**, choose **PEM** as **CA Certificate File format**, and enter a name in the **CA Key File Name** field. Enter the passphrase and select the serial number file from the appliance named `ns-root.srl`. These fields are provided so that the root CA that you have created earlier can sign and trust this server certificate. Fill them out and click on **OK**, as shown in the following screenshot.
The `ns-root.srl` serial number file resides on NetScaler.

This is where we create the public-facing server certificate and simultaneously sign it (self-signed, that is) with the root CA certificate.

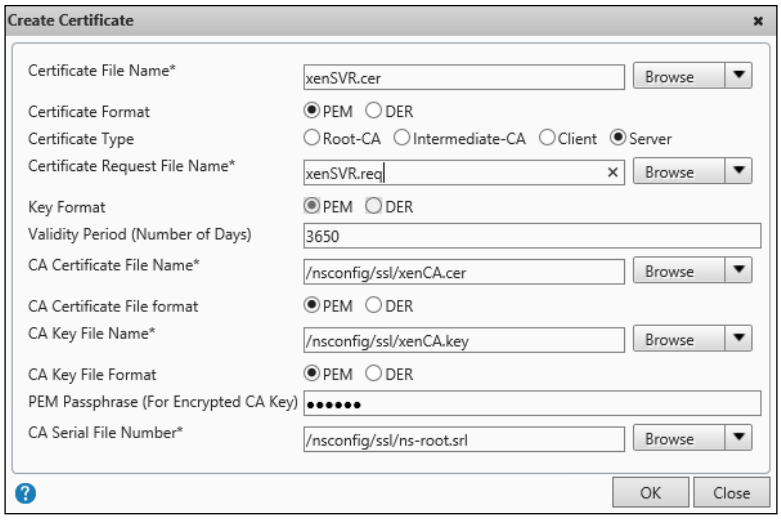

Installing the root CA and public certificates

To install the root CA and public certificates, perform the following steps:

1. You can install the root CA certificate by navigating to **Traffic Management | SSL | Certificates**. Select **Install…**, as shown in the following screenshot:

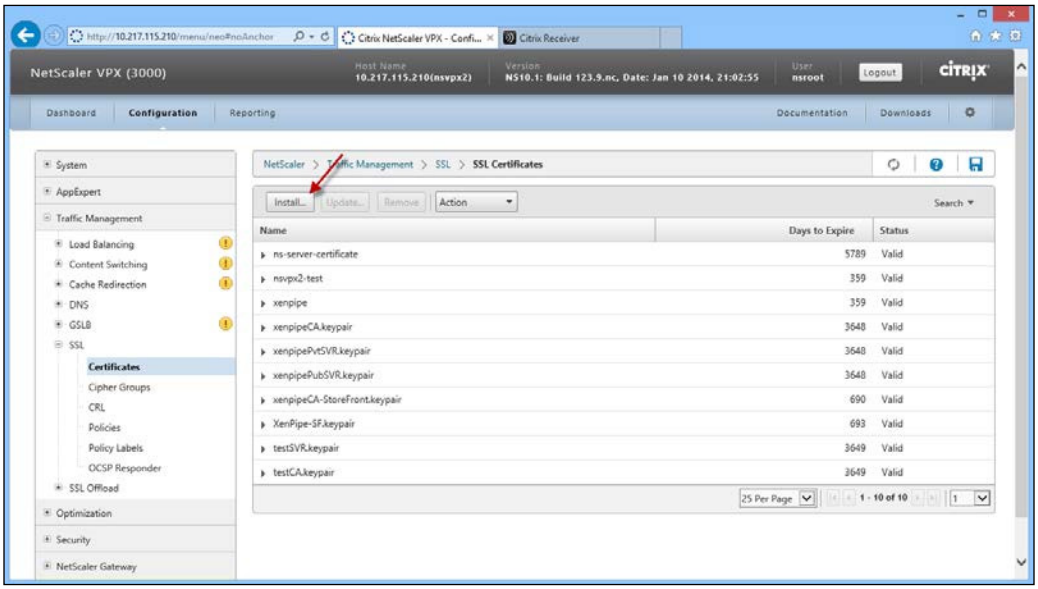

2. Enter a unique name for the root CA **Certificate-Key Pair Name** field. Enter a name in **Certificate File Name**. Click on **Create**, as shown in the following screenshot:

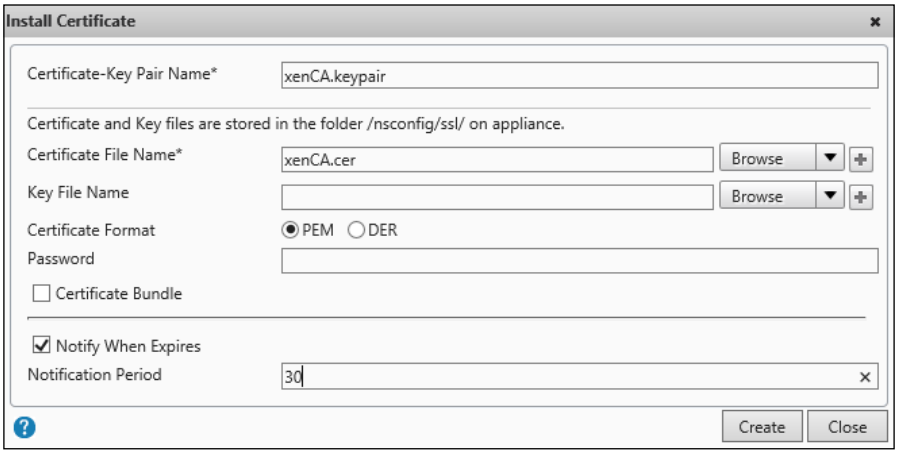

3. Select **Install…** again. Enter a unique name for the server **Certificate-Key Pair Name** field. Enter the name in the **Certificate File Name** and **Key File Name** fields, choose **PEM** as **Certificate Format**, and enter the **Password**. Click on **Create**, as shown in the following screenshot:

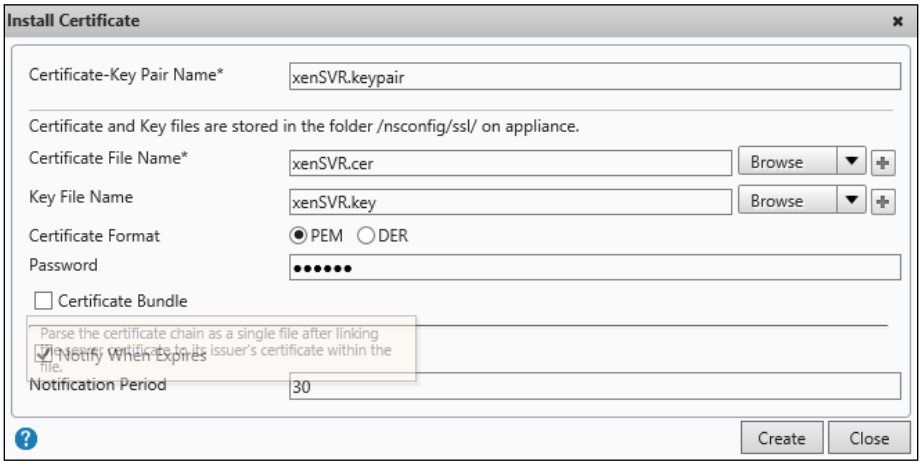

Linking the public and root CA certificates

For the certificates to work properly, the public server certificate must be linked to the root CA certificate.

To link the public and root CA certificates, perform the following step:

1. Navigate to **Traffic Management** | **SSL** | **Certificates**. Select the public certificate by clicking on the key-pair name. Navigate to **Actions** | **Link…**, and select the root CA certificate from the **CA Certificate Name** drop-down list, as shown in the following screenshot:

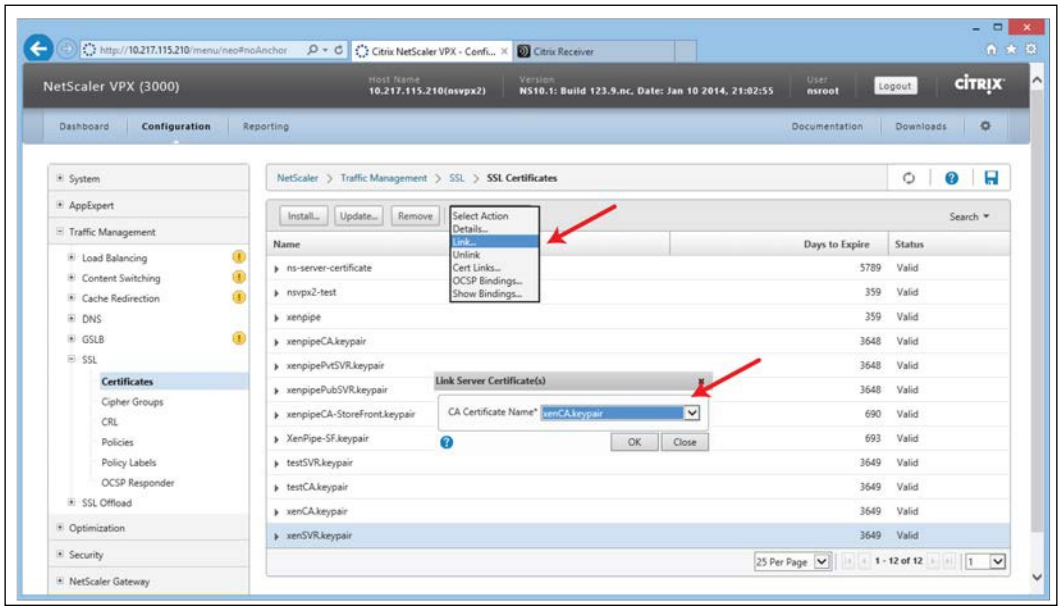

Viewing the root CA and server certificate bindings

To make sure that the certificates are linked, perform the following step:

1. Navigate to **Traffic Management | SSL | Certificates**. Select the root CA certificate. Navigate to **Action | Cert Links…**, as shown in the following screenshot:

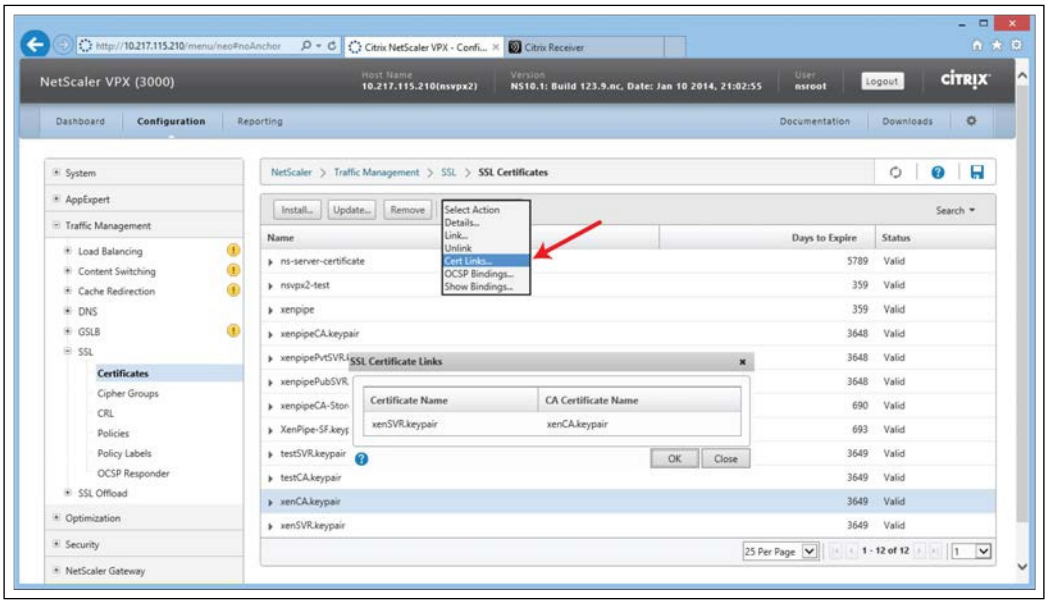

Binding the certificates to the NetScaler Gateway™ VIP

Once the certificates have been created and linked, you need to bind them to the NetScaler Gateway VIP.

To bind the certificates to the NetScaler Gateway VIP, perform the following steps:

1. Navigate to **NetScaler Gateway | Virtual Servers**. Select the VIP and select **Open…**, as shown in the following screenshot:

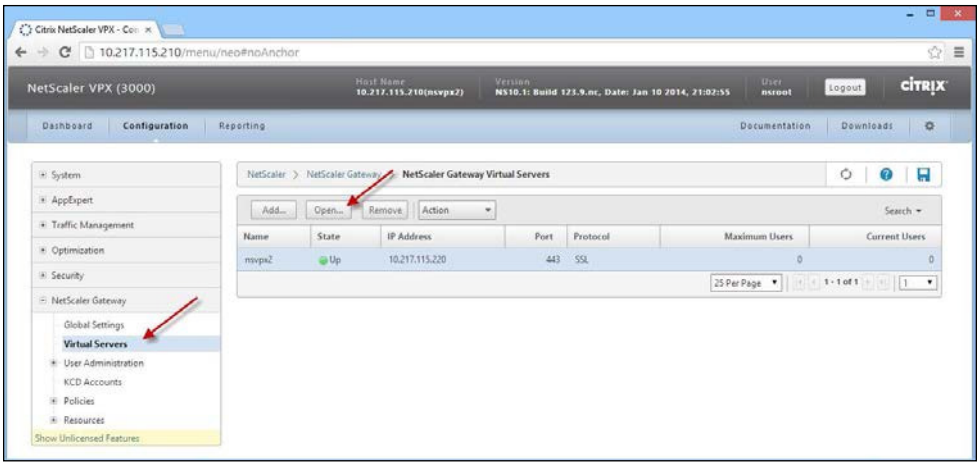

2. In the **Certificates** tab, select the server certificate and click on **Add**. Then, click on **OK,** as shown in the following screenshot.

If you receive an informational message about no useable ciphers, you can click on **OK** and continue.

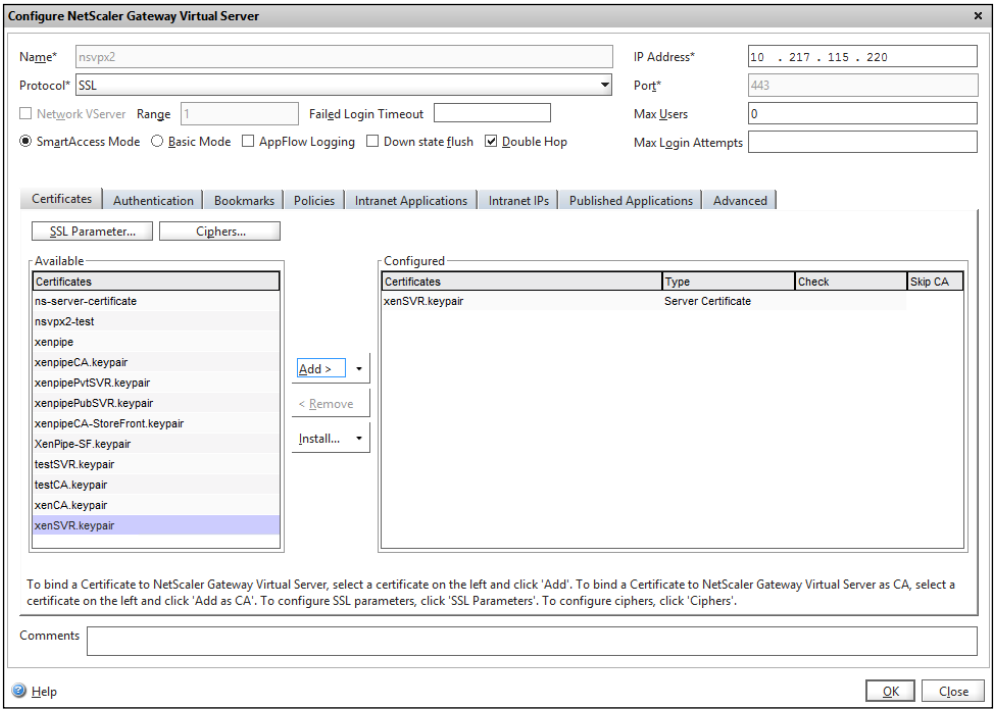

Testing the certificates

Now that the certificates are installed correctly, you should test them. To use a self-signed certificate, you will need to add the domain name to the `hosts` file on the client device unless the public DNS system is already directing to the correct IP address.

To test the certificates, perform the following steps:

1. Log in to the client device, for example, Windows. Run `Notepad.exe` with administrator privileges. Open the **hosts** file in `c:\Windows\System32\drivers\etc`, as shown in the following screenshot:

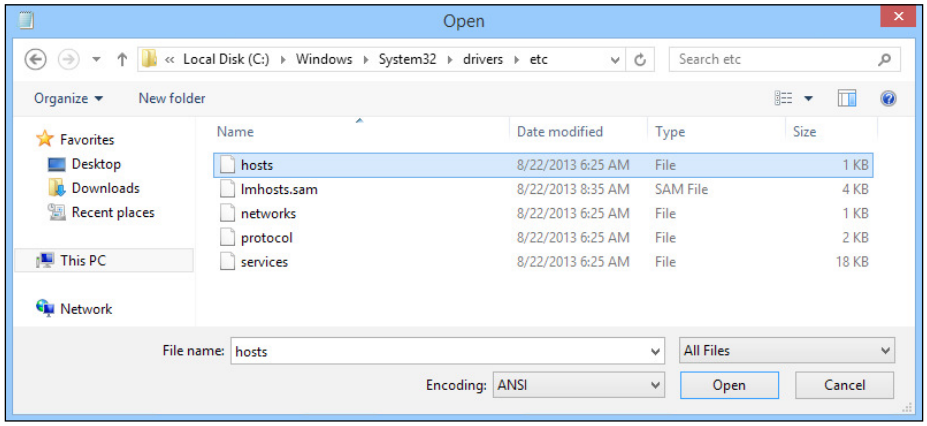

2. Add an entry for your Site. In our case, it will be `ng.xenpipe.com`, as shown in the following screenshot. It must match the **Common Name** field in the certificates that you have just created. Save the file.

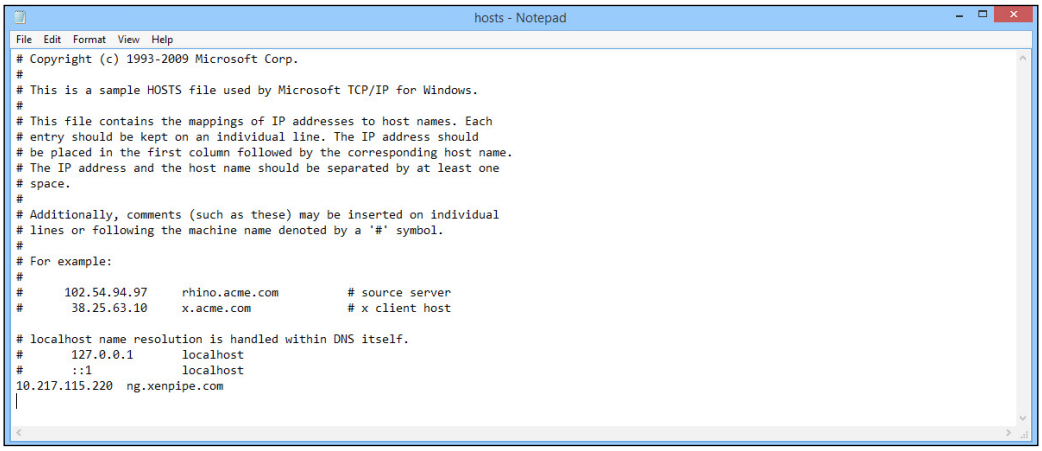

3. Open a command prompt and ping the entry to see if it is alive, as shown in the following screenshot:

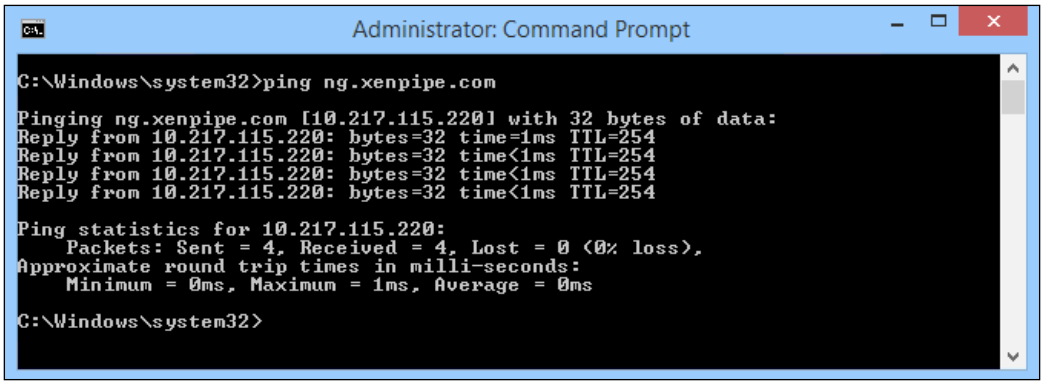

Testing the NetScaler Gateway™ connection

Once you have NetScaler Gateway built and configured, you need to test it. For it to work properly, the certificates must be valid. If you remember, we just created the self-signed test certificates. We now need to download these test certificates from NetScaler and import them into our client device.

Testing NetScaler Gateway™ with a Windows client

One quick way to test if the certificate chain is working on NetScaler is by downloading the root CA certificate, self-signing the public-facing server certificate from NetScaler, and importing these into the client browser. The following steps will guide you through how to test NetScaler Gateway with a Windows client:

1. Download the test certificate from NetScaler Gateway. You can do this using a secure FTP client such as `WinSCP.net`.
2. Download and install WinSCP from `http://winscp.net`.

3. Launch WinSCP and log in to NetScaler Gateway through the **NetScaler IP Address** (**NSIP**). Navigate to the `/nsconfig/ssl` directory and copy the `.cer`, `.req`, and `.key` files, as shown in the following screenshot.

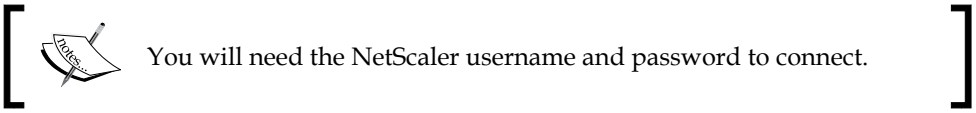

You will need the NetScaler username and password to connect.

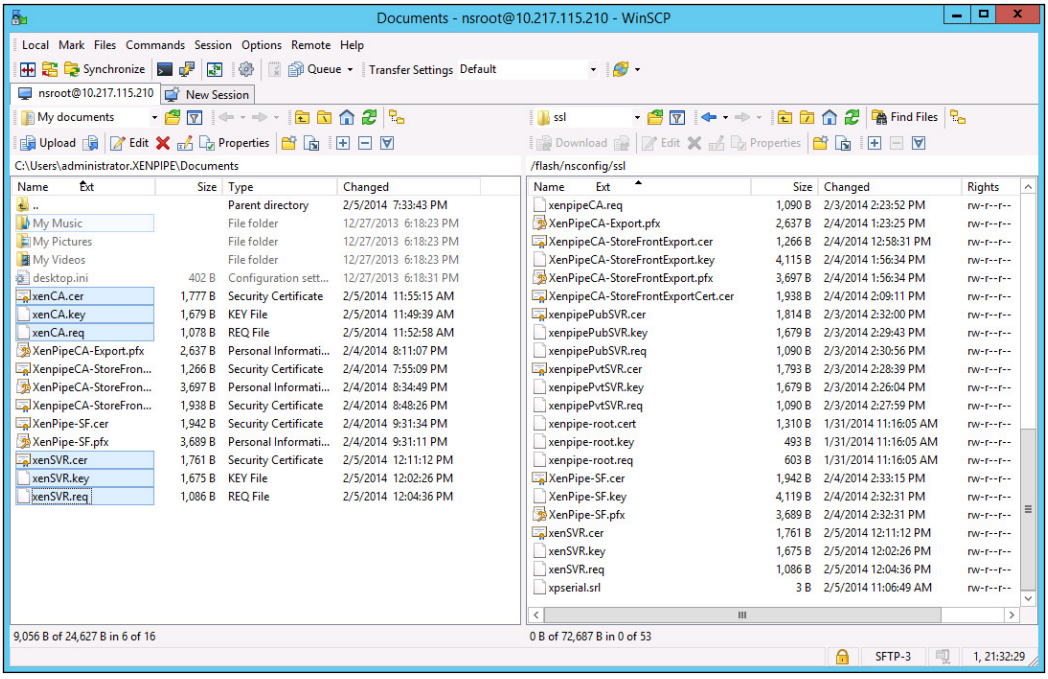

4. To launch **Microsoft Management Console** (**MMC**) from the **Start** menu, right-click and then choose **Run**. Then, enter `MMC`, as shown in the following screenshot:

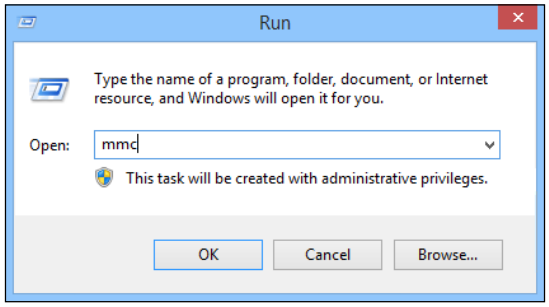

5. Load **Certificates snap-in** for **Computer account**, as shown in the following screenshot:

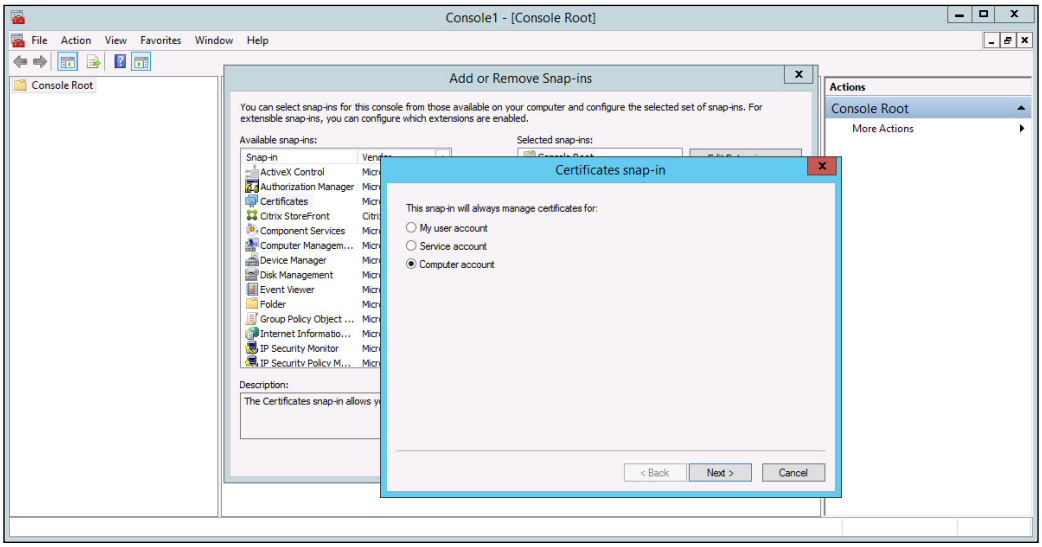

6. Expand the **Trusted Root Certification Authorities** tab and select **Certificates**, as shown in the following screenshot:

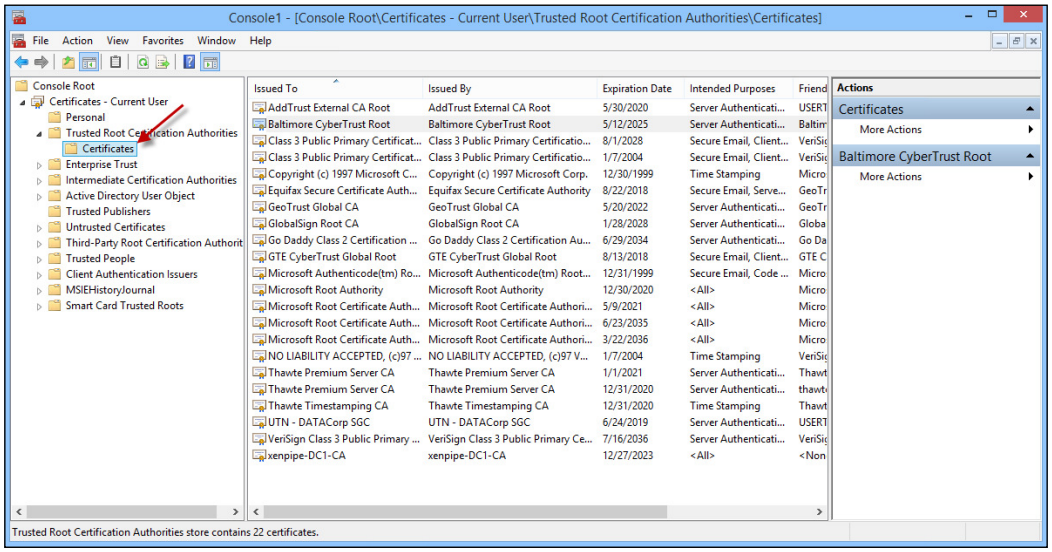

7. Navigate to **Action | All Tasks | Import**. This launches the **Certificate Import Wizard** window. Select **Local Machine** and then click on **Next**, as shown in the following screenshot:

8. Browse for the self-signed root CA certificate. Open it and then click on **Next**, as shown in the following screenshot:

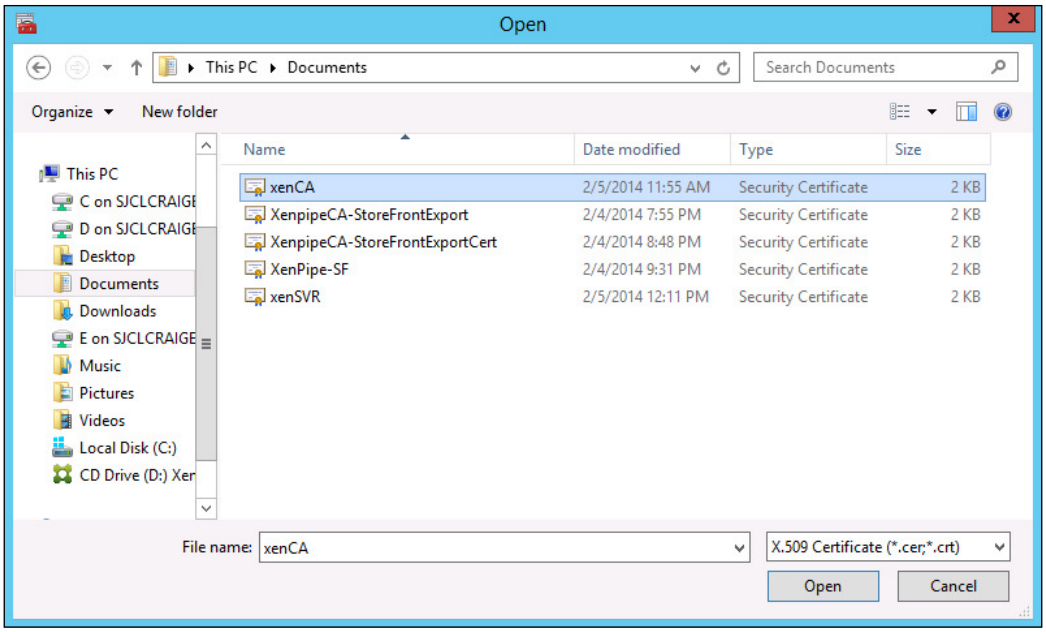

9. Install it in the **Trusted Root Certification Authorities** store. Click on **Next** and then click on **Finish**, as shown in the following screenshot:

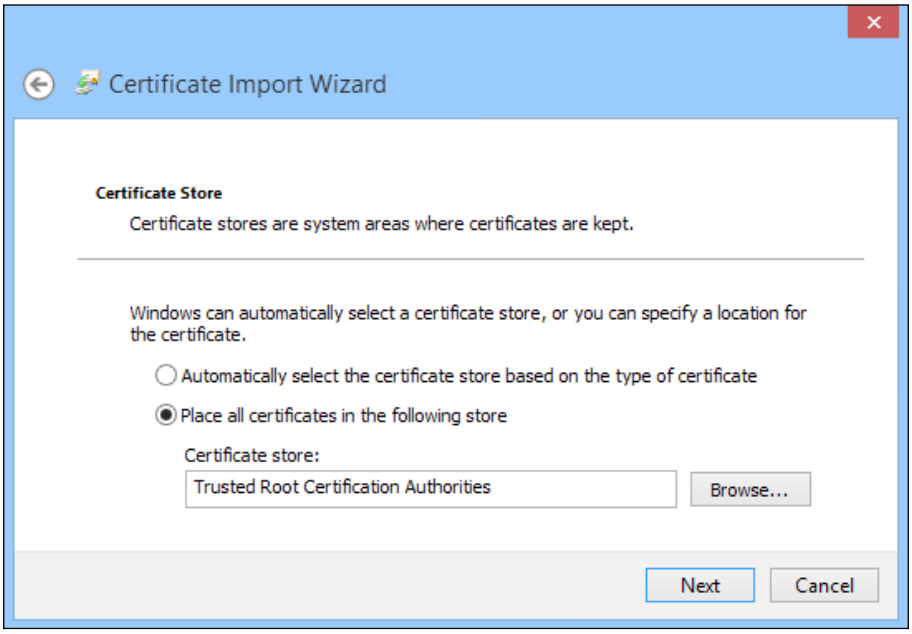

10. Navigate to **Personal** | **Certificates**, as shown in the following screenshot:

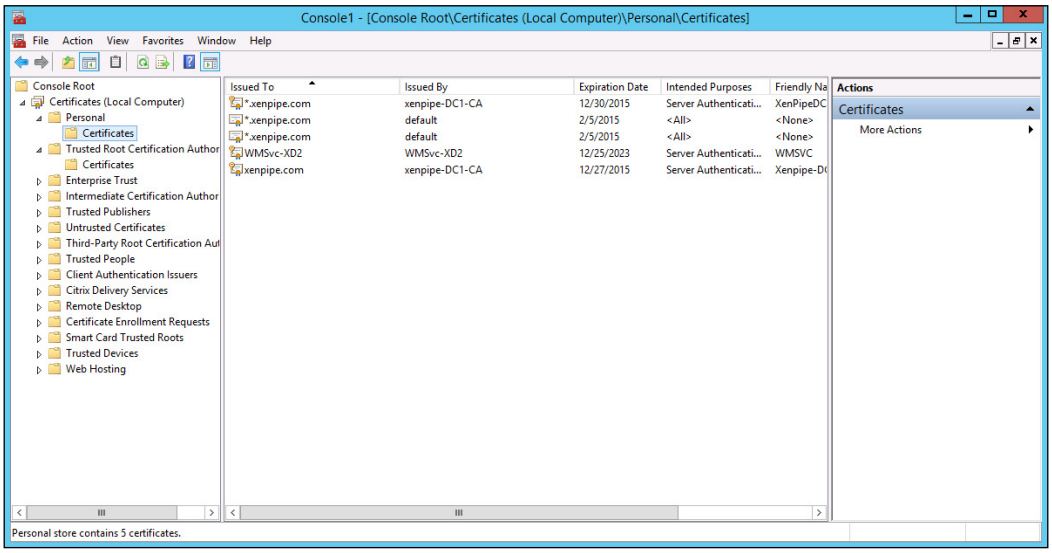

11. Navigate to **Action** | **All Tasks** | **Import**. This launches the **Certificate Import Wizard** window. Select **Local Machine** and then click on **Next**, as shown in the following screenshot:

12. Browse for the self-signed public server certificate. Open it and then click on **Next**, as shown in the following screenshot:

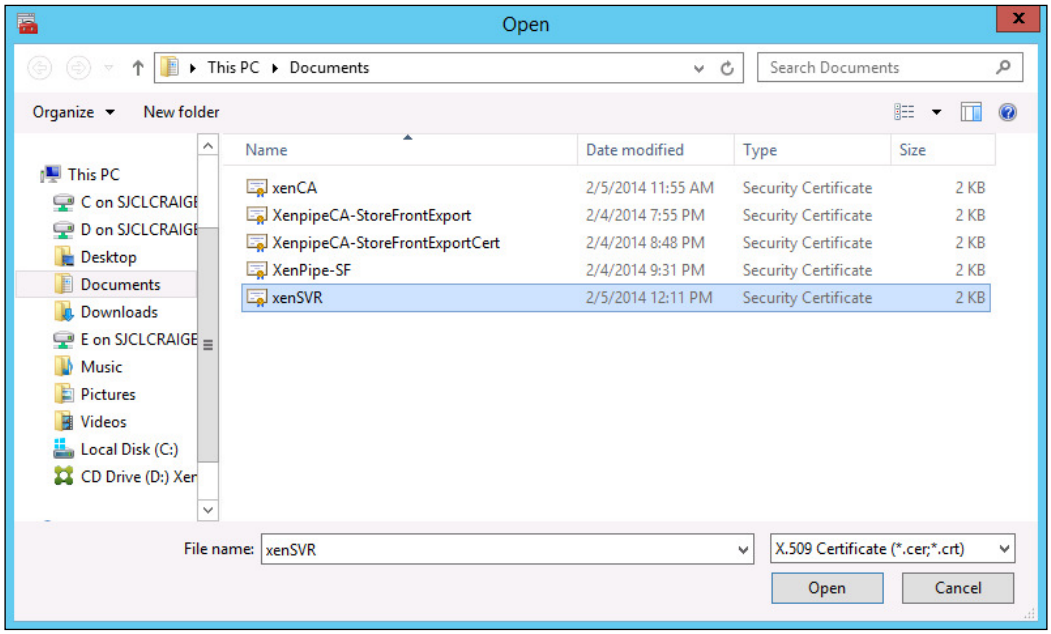

13. Install it in the **Personal** store. Click on **Next** and then on **Finish**, as shown in the following screenshot:

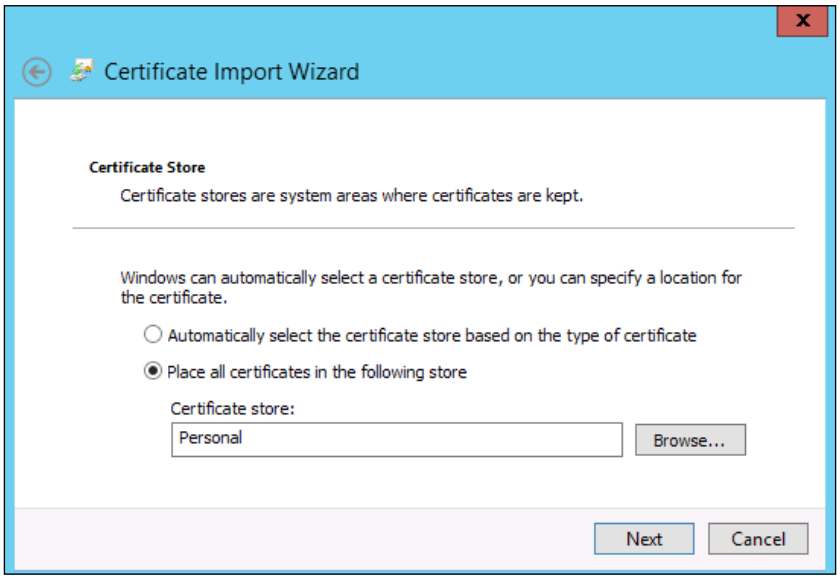

14. Launch a browser and connect to NetScaler Gateway using the entry that you have placed in the `hosts` file, for example, `ng.xenpipe.com`. You should be able to see the Citrix Receiver logon screen without any certificate issues. Select **Desktops** to launch a virtual desktop. Select **Apps** to launch a virtual application (you can also launch these apps from your virtual desktop), as shown in the following screenshot.

Keep in mind that the URL that you connect to and the entry in the `hosts` file must also match the **Common Name** field in the certificate for it to work. Because we created a wildcard certificate, `*.xenpipe.com`, the NetScaler Gateway URL, `https://ng.xenpipe.com`, will work as would any `https://<sub-domain>.xenpipe.com`.

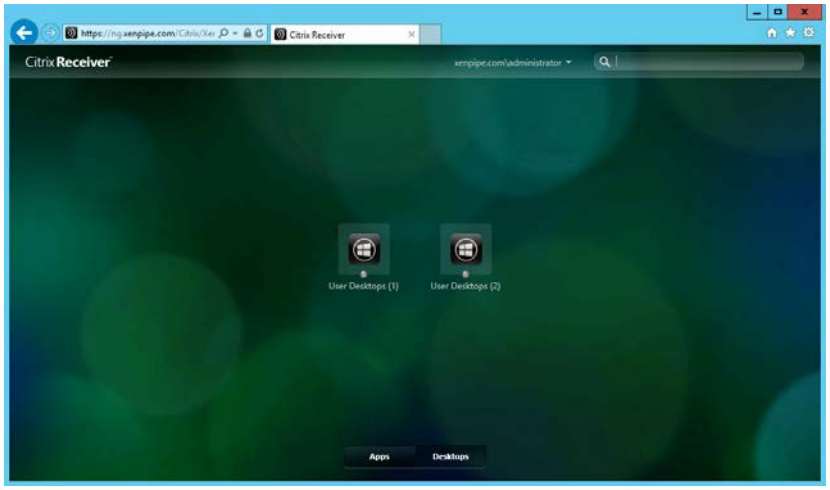

D

Using Public CA-signed SSL Wildcard Certificates on NetScaler Gateway™

When you get ready to move to production, you will need a valid certificate signed by a public certificate authority. There are many vendors to purchase these certificates from, ranging from expensive, well-known vendors to the cheap ones. Verisign is well-recognized, just as GoDaddy is. Do yourself a favor and shop around. The cheapest one I could find for a wildcard certificate was LuckyRegister. I will demonstrate the concept of using public CA-signed SSL wildcard certificates on NetScaler Gateway using this website.

Another option is to use the Class 2 Verification certificate option from `www.StartSSL.com`. With this, you can create unlimited certificates for unlimited domains including wildcard certificates.

To remain flexible in our deployment, we will also use a wildcard certificate, for example, `*.xenpipe.com`, so that we can create and use different subdomains (`ng.xenpipe.com`) and still have our certificates work.

The reason you want to use a certificate signed by a valid public **Certificate Authority** (**CA**) is because those vendors have trusted root certificates built into all of the browsers your clients use. You won't have to import any certificates and the certificate you purchase will be trusted by the clients connecting to NetScaler. Once you move from Proof of Concept to production, you really should use a public CA-signed certificate.

Enabling SSL on NetScaler Gateway™

To enable SSL on NetScaler Gateway, perform the following steps:

1. Log in to NetScaler Gateway.
2. Navigate to **Traffic Management | SSL**.
3. Right-click and select **Enable Feature,** as shown in the following screenshot:

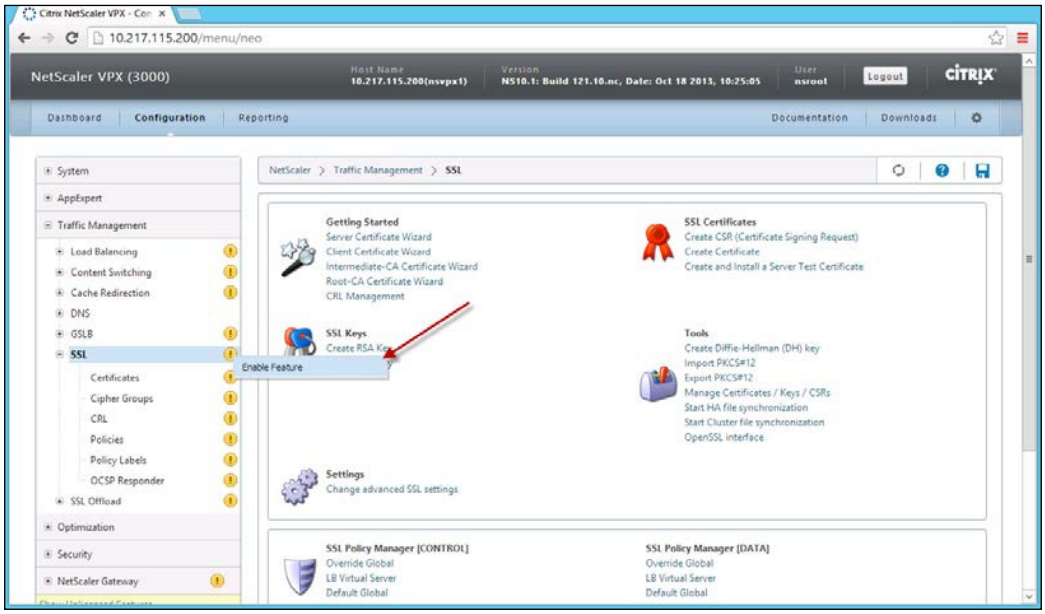

Creating a certificate request

To create a certificate request, perform the following steps:

1. Navigate to **Traffic Management | SSL**. Under **SSL Keys**, select the **Create RSA Key** option, as shown in the following screenshot:

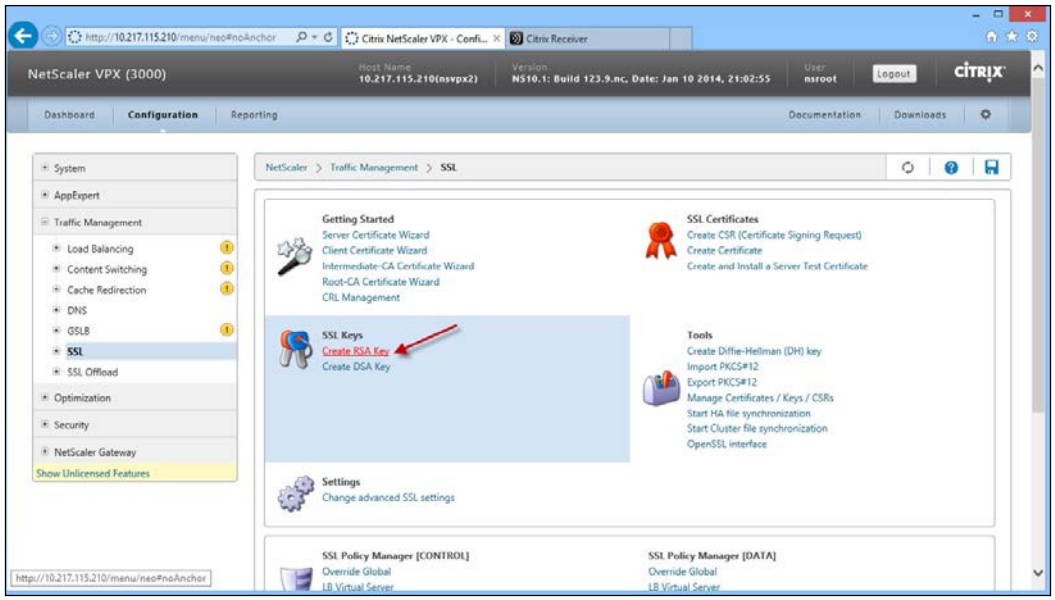

2. To create the RSA key, enter a **Key Filename** and **Key Size(bits)** value, set the **Public Exponent Value** as **F4**, choose **PEM** as **Key Format**, and click on **OK,** as shown in the following screenshot:

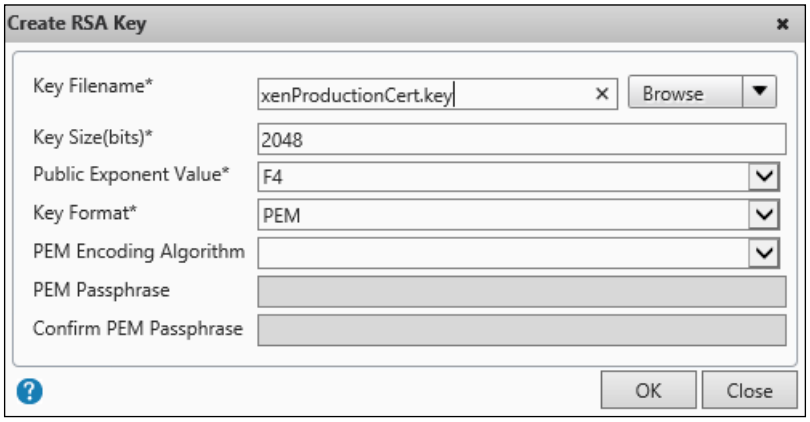

3. Under **Tools**, select **Create CSR (Certificate Signing Request)**. Enter the **Request File Name** and **Key Filename** you have created in the previous step. Choose **PEM** as **Key Format**, enter the password from the time the RSA key was created, enter the details in the **Distinguished Name Fields** section, and click on **OK,** as shown in the following screenshot.

Common Name *must* match the **Fully Qualified Domain Name** used to access the Site; for example, `ng.xenpipe.com` is what we will use to connect to NetScaler Gateway. So, in this example, as before, we create a wildcard certificate that will work across all subdomains; hence, you see `*.xenpipe.com` in the **Common Name** field. If you don't do this, the certificate will not work.

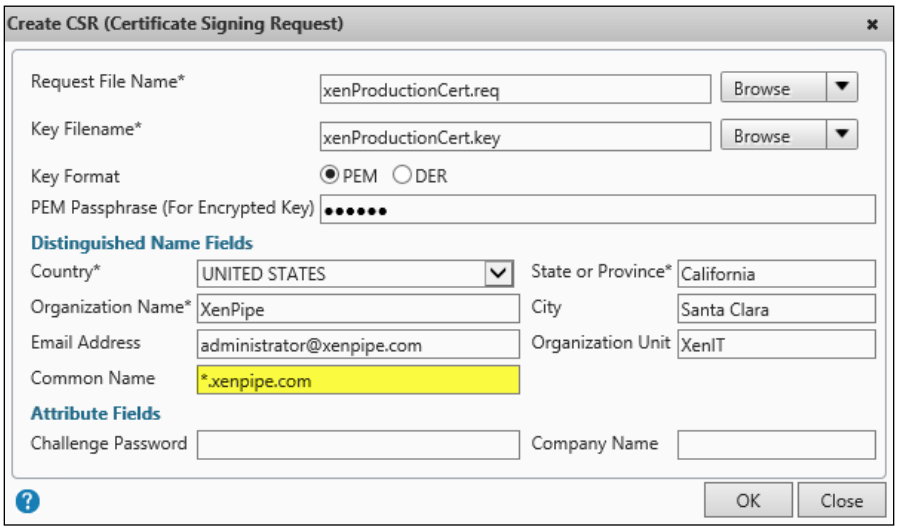

4. Download the Certificate Signing Request from NetScaler Gateway. You can do this using WinSCP as we did in *Appendix C, Creating Self-signed Certificates for NetScaler Gateway™*, or you can use the NetScaler tool. Navigate to **Traffic Management | SSL | Tools | Manage Certificates / Keys / CSRs** as shown in the following screenshot:

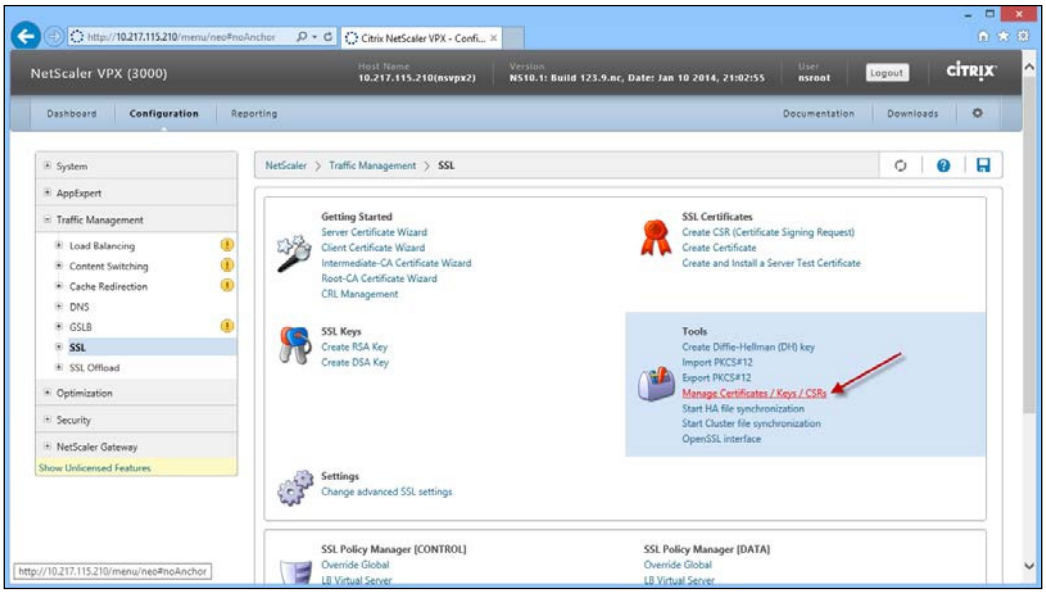

5. Select the `request` file you just created and click on **Download…,** as shown in the following screenshot:

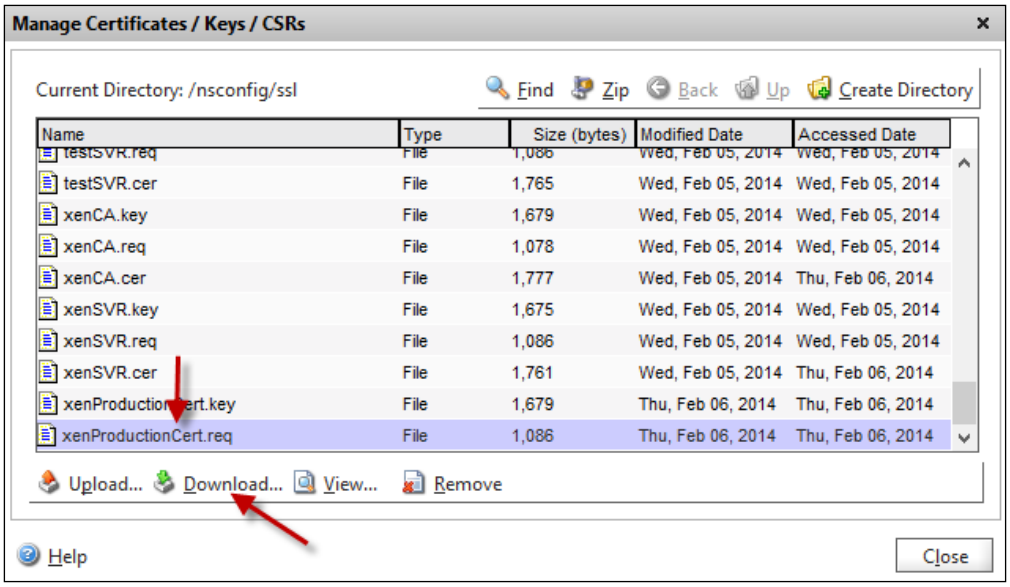

6. Open the `request` file using a text editor, for example, Notepad. Copy the text as shown in the following screenshot:

Certificate request files are encrypted and appear as a big block of garbled text. To submit a request to a public CA, you simply copy the text and paste it into the public CA's web form. We will see how to paste this text into the public CA's web form in the next section.

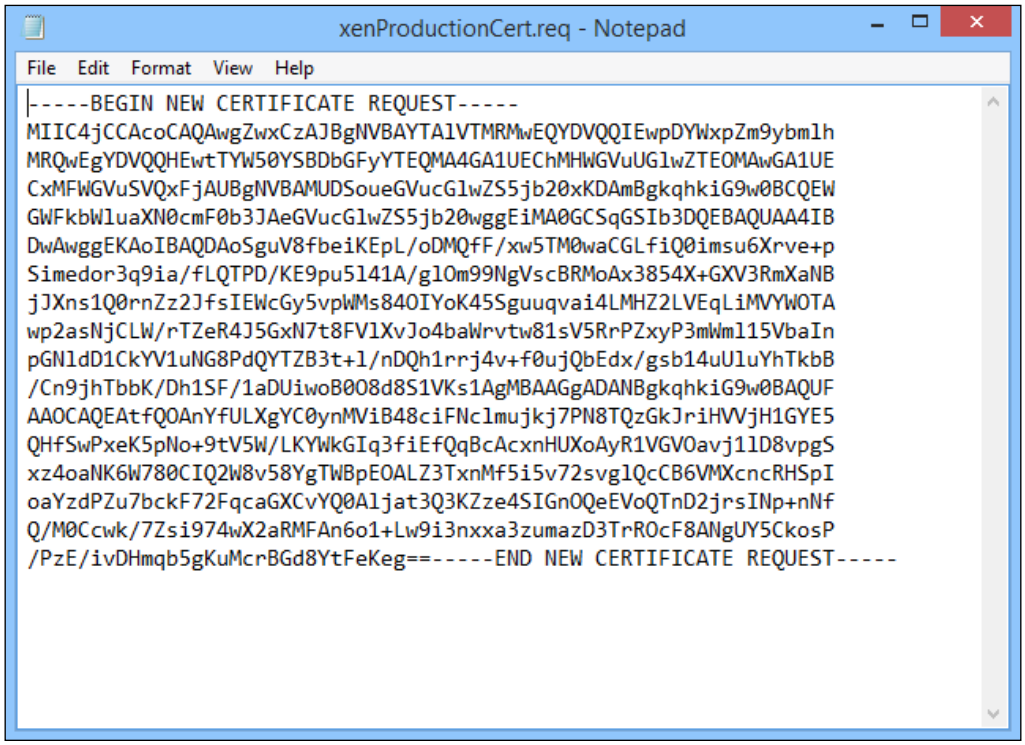

Submitting the request to the public CA

Now that you have the `request` file built, you need to submit it to a public CA to have it signed and trusted by every browser and device out there.

To submit the `request` file to the public CA, perform the following steps:

1. Open a browser and navigate to the public CA website. In our case, we are using `http://luckyregister.com`. Look for the wildcard certificate in the **SSL Certificates** section and add it to the cart and check out, as shown in the following screenshot:

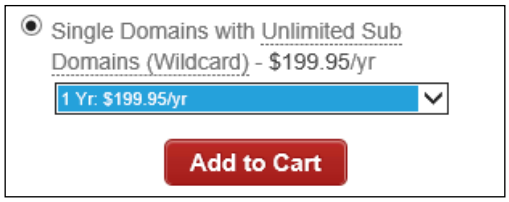

2. Once you've paid for the SSL certificate, navigate to **SSL Certificates** and click on **Set Up,** as shown in the following screenshot:

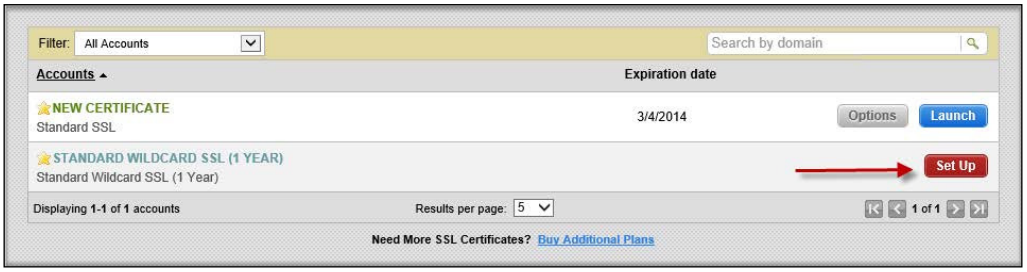

3. When the setup is done, launch the **NEW CERTIFICATE** tool, as shown in the following screenshot:

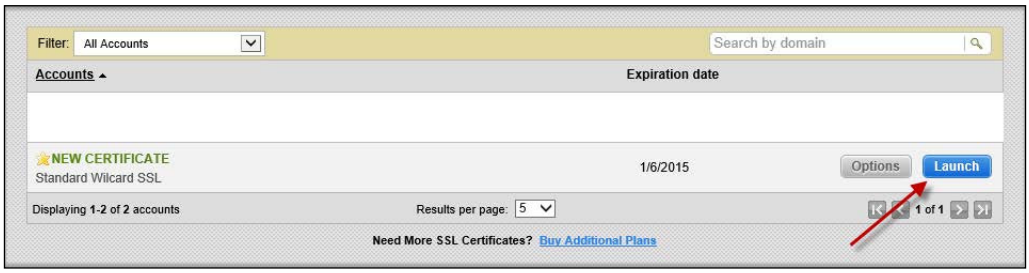

4. Insert or paste the Certificate Signing Request created in the previous section, as shown in the following screenshot:

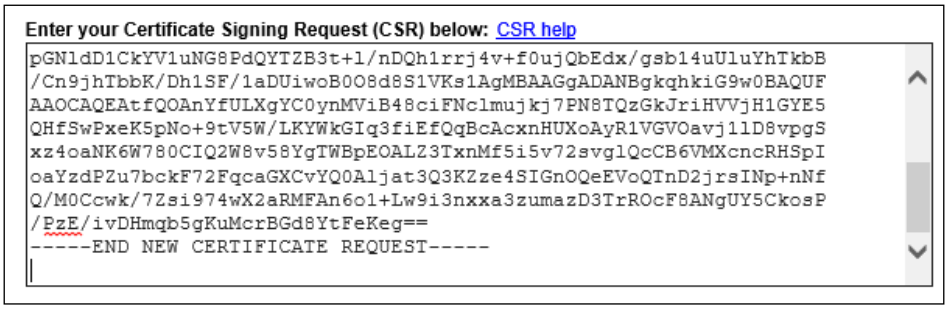

5. Click on **Next** through the prompts and wait for the certificate to be issued. Navigate to **SSL Certificates** and click on **Download** to download the issued certificate, as shown in the following screenshot:

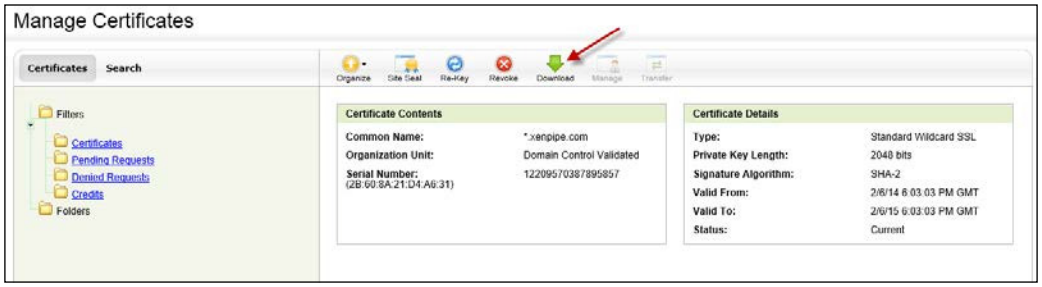

6. Use **IIS7** as the server type, as shown in the following screenshot:

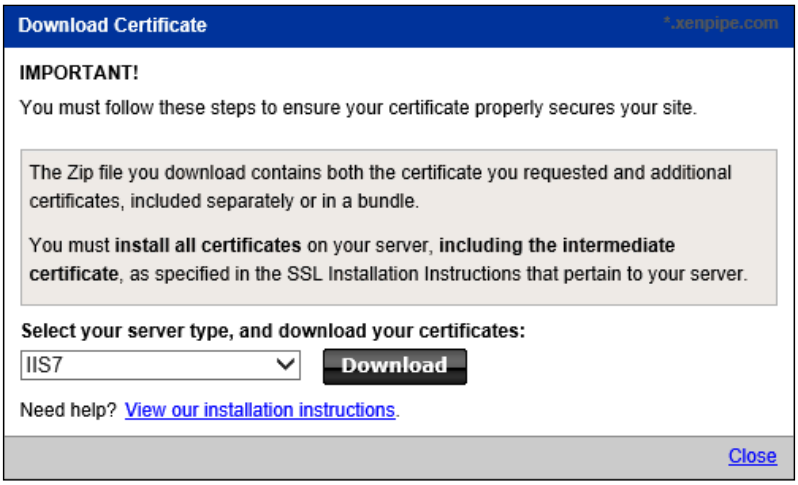

Installing the public-signed wildcard certificate

To install the public-signed wildcard certificate, perform the following steps:

1. Log in to NetScaler Gateway and navigate to **Traffic Management | SSL**. Navigate to **Manage Certificates / Keys / CSRs** and click on **Upload…** to upload the certificate to NetScaler Gateway, as shown in the following screenshot:

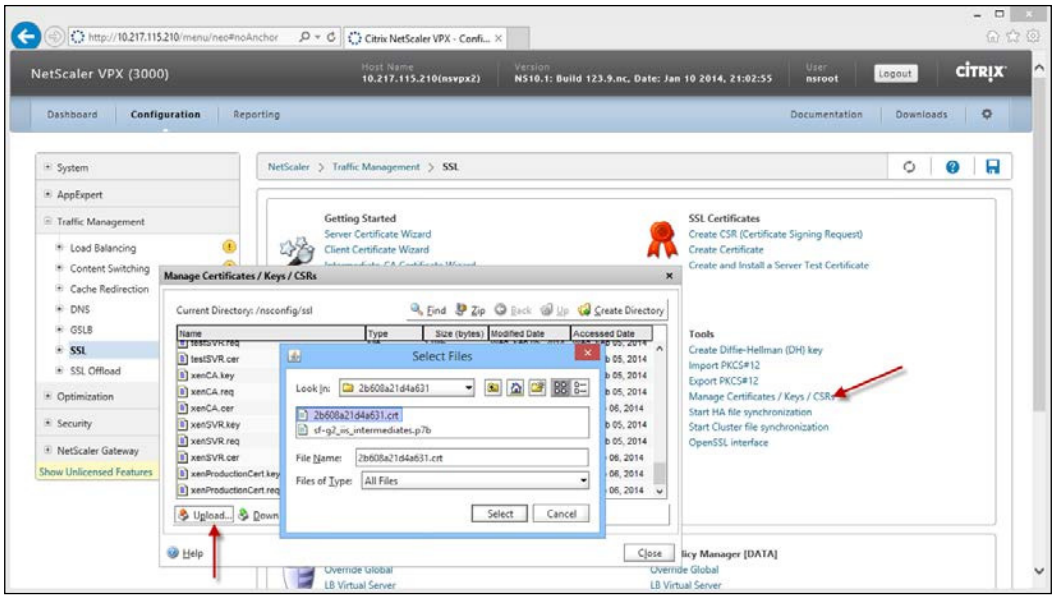

2. Navigate to **Traffic Management | SSL | Certificates**. Click on **Install…**.
3. Create a new **Certificate-Key Pair Name**. Use the newly uploaded certificate. Use the key file created for the certificate request. Use **PEM** as the **Certificate Format**, enter the key's password, and click on **Create,** as shown in the following screenshot:

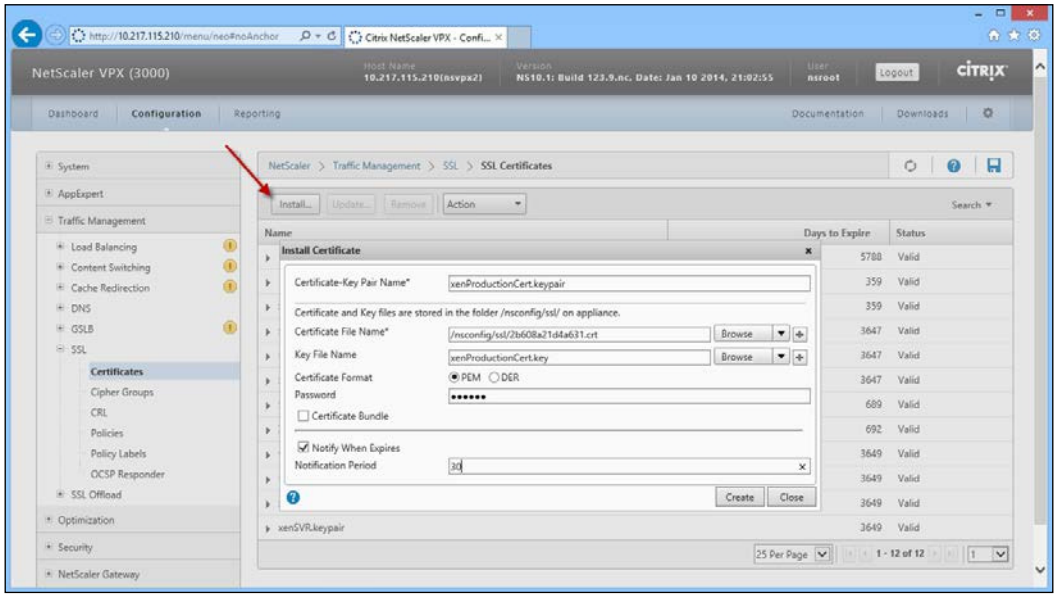

Binding the public-signed certificate to the NetScaler Gateway™ VIP

Once the public-signed certificate has been created and installed, you need to bind it to the NetScaler Gateway VIP.

To bind the certificate to the NetScaler Gateway VIP, perform the following steps:

1. Navigate to **NetScaler Gateway | Virtual Servers**. Select the VIP and click on **Open…,** as shown in the following screenshot:

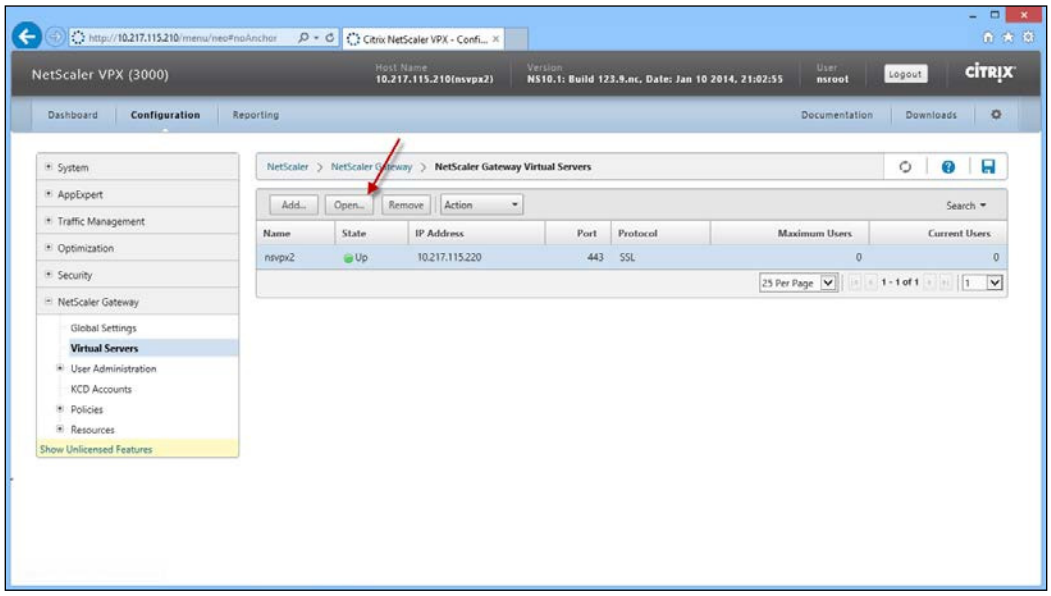

2. In the **Certificates** tab, select the new server certificate and add it. Then, click on **OK,** as shown in the following screenshot:

If you receive an informational message about no useable ciphers, you can click on **OK** and continue. Also keep in mind that you will need to unbind any certificates that are currently bound to the VIP before binding this new one.

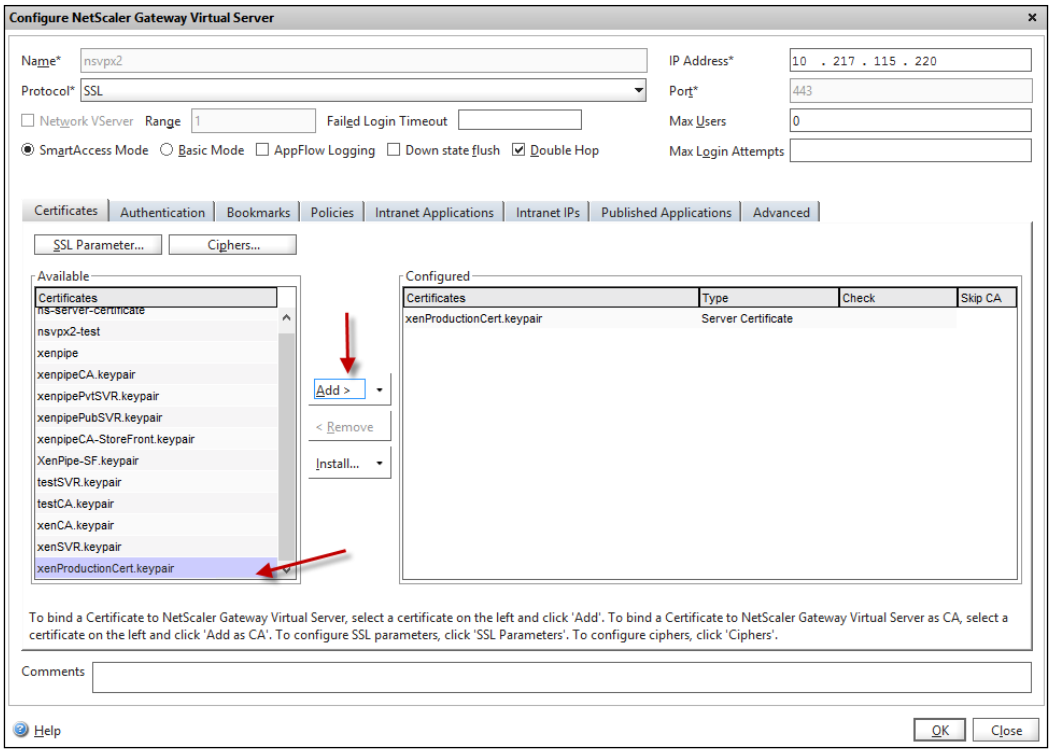

Testing NetScaler Gateway™ and certificates

Now that the certificates are installed correctly, you should test them. To use the new public-signed certificate, you can either have the public DNS system resolve to the IP address of your VIP or you can add the domain name to the `hosts` file on the client device. Either way, the certificate should work in any browser on any device without having to import any certificates on those devices.

To test the NetScaler Gateway VIP and certificates, perform the following steps:

1. Launch a browser and connect to NetScaler Gateway using the entry you created, for example, `ng.xenpipe.com`. You should be able to see the Citrix Receiver logon screen without any certificate issues. Select **Desktops** to launch a virtual desktop.

2. Select **Apps** to launch a virtual application (you can also launch these apps from your virtual desktop).

Keep in mind that the URL you connect to must also match the **Common Name** field in the certificate for it to work. Because we created a wildcard certificate, `*.xenpipe.com`, the NetScaler Gateway URL, `https://ng.xenpipe.com`, will work just as any `https://<sub-domain>.xenpipe.com` would.

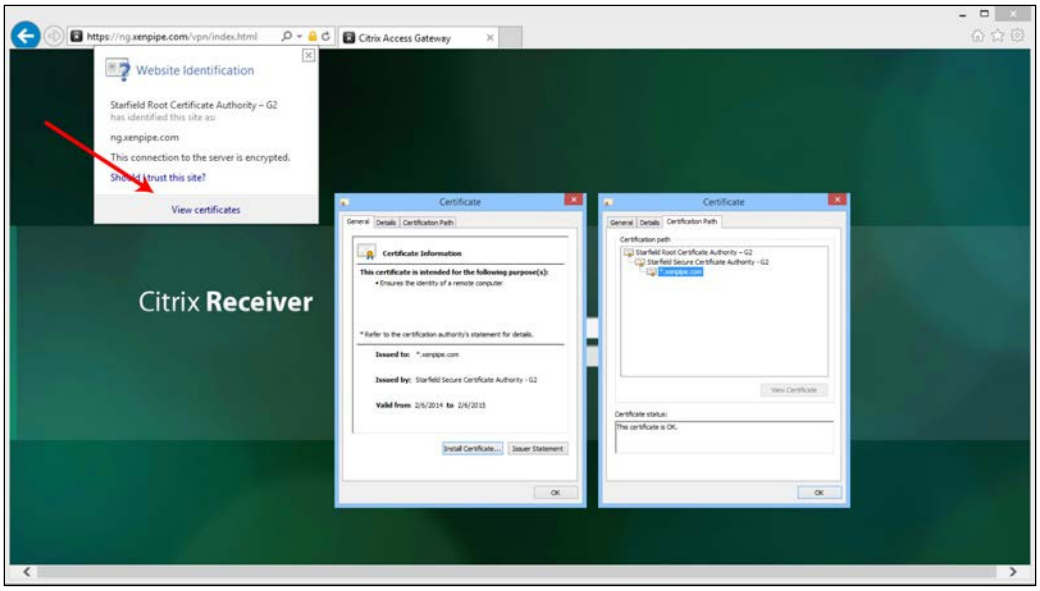

Index

A

Access Control List (ACL) 136
Active Directory
adaptive display, HDX 3D
administrative roles 237
administrator roles
  about 140
  full admin 140
  read-only admin 140
Aero redirection
  about 213
  configuring 213, 214
AlwaysOn Availability
  URL 195
Android
  Receiver, running on 271
Apple
  Receiver, running on 270
Application Centric Infrastructure (ACI) 240
application delivery 235
Application Delivery Controllers (ADCs) 238
Application Delivery Network (ADN) 240
application desktop delivery groups
  application, creating 125
  application sessions, managing 127
  application sharing 123
  applications, modifying 127, 128
  applications, publishing to multiple desktop groups 124
  content redirection 124
  creating 125, 126
  managing 125, 126
application master images
  creating 53-55
Application Programming Interfaces (APIs) 210, 243
applications
  delivering 236
  troubleshooting 316
audio, HDX 3D
  configuring 230
audio policies
  Audio over UDP 333
  Audio over UDP Real-time Transport 333
  Audio Plug-n-Play 333
  Audio quality 333
  Audio UDP port range 333
  Client audio redirection 333
  Client microphone redirection 333
Auto-create client printers policy
  about 170
  enabling 170
autoreconnection policies
  auto client reconnect 346
  auto client reconnect authentication 346
  auto client reconnect logging 346
  auto connect client COM ports 346
  auto connect client LPT ports 346

B

backup
  SQL Server, backing up 203, 204
  SQL Server, restoring 205, 206
bandwidth policies
  audio redirection bandwidth limit 334
  audio redirection bandwidth limit percent 334

client USB device redirection bandwidth limit or percent 335
clipboard redirection bandwidth limit or limit percent 334
COM port redirection limit or limit percent 334
file redirection bandwidth limit or limit percent 334
HDX™ MediaStream Multimedia Acceleration bandwidth limit or percent 334
LPT port redirection bandwidth limit or percent 334
overall session bandwidth limit 335
printer redirection bandwidth limit or percent 335
TWAIN device redirection bandwidth limit or percent 335
BlackBerry
Receiver, running on 271
Branch Repeater 240
buffer 117

C

caching policies
persistent cache threshold 342
certificate request
creating 374-378
submitting, to public CA 378-380
certificates
testing 383
Chromebook
Receiver, running on 271
Citrix.AdIdentity.AdminV2 246
Citrix.AppV.Admin.V1 246
Citrix.Broker.Admin.V2 246
Citrix® CloudBridge™
about 240
implementing, in Citrix Receiver™ 241
Citrix® communication ports
about 24
common communication ports 25
for Citrix® license server 25
for Citrix® XenDesktop® 25
for Citrix® XenServer® 26
Citrix.Configuration.Admin.V2 246
Citrix.ConfigurationLogging.Admin.V1 246
Citrix.DelegatedAdmin.Admin.V1 246
Citrix® Edgesight® 15
Citrix.EnvTest.Admin.V1 246
Citrix.Host.Admin.V2 246
Citrix.MachineCreation.Admin.V2 246
Citrix.Monitor.Admin.V1 246
Citrix® NetScaler® 240
Citrix® NetScaler Gateway™
about 158
used, for implementing policies 158
Citrix Ready®
about 252
URL 252
Citrix Receiver™. *See* **Receiver**
Citrix.Storefront.Admin.V1 246
Citrix® Studio 237
Citrix® VDI hybrid cloud 324
Citrix® XenApp® 15
Client printer redirection policy
enabling 171
client side, XenDesktop® architecture
about 17
Receiver 17
cloud
VDI 321
virtualization 321, 322
CloudBridge® plugin 269
cmdlets (command-lets)
about 244
executing 245
COLO arrangement 322
colocation 322
color compression, HDX 3D
configuring 231
computer-aided design (CAD) 224
computer-aided engineering (CAE) 224
computer-aided manufacturing (CAM) 224
controller cmdlets
Get-AcctServiceStatus 250
Get-BrokerServiceStatus 251
Get-ConfigServiceStatus 251
Get-HypServiceStatus 251
Get-ProvServiceStatus 251
Get-PvsVmServiceStatus 251
controller, securing
about 299

controller port, changing to HTTPS 300
IIS, installing 299
non-IIS, installing 299

D

database 13
Default printer policy
enabling 173
Delivery Agent policies
controller registration IPv6 netmask 350
controller registration port 350
controllers 350
controller SIDs 350
enable Auto Update of Controllers 350
only use IPv6 controller registration 351
set GUID 351
Delivery Controller
about 13, 20
managing 129
Delivery Controller environment
Active Directory OU-based controller discovery 134
auto update, enabling 130
controller connections, managing 131
controller discovery 129-131
default HTTP, changing 136
Delivery Agent (DA), moving to another Site 134
Delivery Controllers, adding 132, 133
Delivery Controllers, moving 133
Delivery Controllers, removing 133
HTTPS ports, changing 136
managing 129
SSL, using on controllers 136
delivery groups
about 105
application delivery groups, creating 69-72
creating 67, 106, 107
desktop delivery groups, creating 68, 69
desktop sessions, managing 109
editing 107-109
locating 111
managing 105
messages, sending to users 109
sessions, disconnecting 109
Windows desktop machine 105
Windows Server machine 106
delivery groups resources
desktop access, restricting 114
desktop access, restricting with Exclusion filters 115
desktop access, restricting with Smart Access 114
desktop power settings, managing 117
desktops, adding 110
desktops, deleting 114
desktops, locating 111
desktops, reallocating 110
desktops, removing 113
desktops, restarting 112, 113
desktops, shutting down 112
ICA® protocol communications, securing 116
maintenance mode, disabling 119, 120
maintenance mode, enabling 119, 120
managing 110
server load, managing 120
sessions, locating 111
user data, exporting 118
user data, importing 118
Delivery Services Console (DSC) 237
Demilitarized Zone (DMZ)
about 274
setting up 274
desktop cmdlets
Get-BrokerDesktopGroup cmdlet 249
Set-BrokerDesktopGroup cmdlet 249
desktop composition redirection. *See* **aero direction**
Desktop Lock 267
desktop master images
creating 50-52
desktop OS machines 15
desktops
about 15
troubleshooting 317
desktop UI policies
about 338
Aero Redirection 338
Aero Redirection Graphics Quality 338
desktop wallpaper 338
menu animation 338
view window contents while dragging 338

Desktop Viewer 259
Director
about 14, 305
SSL, configuring 306
SSL, disabling 306
used, for monitoring XenDesktop® Site 305-309
Director (XD3)
components, installing 47, 48
installing 47
Direct X 210
domain certificate authority
creating 327-329
Double-Hop DMZ architecture
about 274
diagrammatic representation 275

E

electronic software distribution (ESD) tools 87
Exclusion filters
used, for restricting user access 115
extended display identification data (EDID) file 229

F

filter modes, XenDesktop® policies
different filters, using with similar modes 150
filters, using with different modes 150
Flash Media
about 219
Flash redirection, configuring on client 221-223
Flash redirection, configuring on server 220, 221
Flash redirection
enabling 221
Flash URL Compatibility List 219
FlexCast® 16
FlexCast Management Architecture (FMA) 238
FlexCast® models
hosted shared 17
hosted VDI 16
Local VM 16
on-demand applications 17
Streamed VHDs 16
URL 16

G

geographical information software (GIS) 224
Get-AcctServiceStatus cmdlet 250
Get-BrokerDesktopGroup cmdlet 249
Get-BrokerServiceStatus cmdlet 251
Get-ConfigServiceStatus cmdlet 251
Get-HypServiceStatus cmdlet 251
Get-ProvServiceStatus cmdlet 251
Get-PvsVmServiceStatus cmdlet 251
GoDaddy 373
GoToMyPC functionality 86
GPU pass-through 226
graphics and multimedia policies
display memory limit 340
display mode degrade performance 341
dynamic windows preview 341
extra color compression 339
extra color compression threshold 339
flash background color list 340
flash default behavior 340
flash event logging 340
flash intelligent fallback 340
Flash server-side content fetching URL list and Flash URL compatibility list 340
heavyweight compression 339
image caching 341
lossy compression threshold value 340
lossy compression value 340
maximum allowed color depth 341
minimum image quality 338
moving image compression 339
multimedia conferencing 341
notify user when display mode is degraded 341
progressive compression level 339
progressive compression threshold value 339
queuing and tossing 341
target frame rate 338
visual quality 338
windows media redirection 341

windows media redirection buffer size 341
windows media redirection buffer size use 341
graphics processing unit (GPU) 213, 224

H

HDX™
about 209, 210
graphics technologies 210
reality 212
system requisites 210
troubleshooting 317
HDX™ 3D
about 223, 224
adaptive display 233
audio, configuring 230
color compression, configuring 231
configuring 229
GPU, versus vGPU 224
HDX™ GPU sharing 226
image quality, configuring 230
installing 227, 228
monitors, configuring 229
network priorities, configuring 232
requisites 225
upgrading 229
webcams, configuring 231
working 226, 227
HDX™ 3D policies
enable lossless 351
HDX3DPro quality settings 351
HDX™ GPU Sharing 226
HDX Insight™
about 310
using 310-314
HDX™ protocol 212
High Availability (HA) 195
high definition experience. *See* **HDX™**
hosted applications
application desktop delivery groups 123
managing 122
hosts
creating 96
deleting 100
editing 98
managing 96
HTML5 compatible web browsers
Receiver, running on 271
hybrid cloud 321-324
Hypervisor 13

I

ICA/HDX protocols
securing 297, 298
ICA® policies
ICA® listener connection timeout 345
ICA® listener port number 345
ICA® round trip calculation 345
ICA® round trip calculation for idle connections 345
ICA® round trip calculation interval 345
ICA® Session Round Trip Time (ICA RTT) 310
IMA data store 238
image quality, HDX™ 3D
configuring 230
Independent Computing Architecture (ICA®) protocol 162
Infrastructure as a Service (IaaS) 323
Input Output operations Per Second. *See* **IOPS**
installation checkpoint, XenDesktop 72-74
intelligent fallback 219
Internet Information Services (IIS) role 41
IOPS 196
ISE (Integrated Scripting Environment) 248
Isochronous 183

K

keep alive policies
ICA keep alives 346
ICA keep alive timeout 346

L

Linux
Receiver, running on 271
load management policies
concurrent logon tolerance 349
CPU usage 349
disk usage 349
maximum number of sessions 349

memory usage 349
memory usage base load 349
Local Data center 242
locally attached printers
using 162
LuckyRegister 373

M

machine catalog management
about 88
Active Directory computer accounts, managing 93
machines, adding 94
master image, reverting to 92, 93
master image snapshot, capturing 89
master image, updating 90-92
machine catalogs
about 81
application servers, creating 64-67
creating 58, 84
deleting 96
desktops, creating 58-63
machine management 86
managing 88
modifying 95
prerequisites 82
remote PC Access 86
renaming 95
user experience 87
Windows desktop, creating 85
Windows Server, creating 85
Machine Creation Services (MCS) 130
machine management
another service or technology option 87
Machine Creation Services (MCS) 86
physical machine type 86
virtual machine type 86
maintenance mode, delivery groups
enabling 119
Microsoft SQL Server versions 19
Microsoft Windows
Receiver, running on 270
Microsoft Windows PowerShell
about 244
cmdlets, for XenDesktop 244
launching 244
snap-ins, for XenDesktop 244
mobile workstyles 85
mobility policies
allow applications to use the physical location of the client device 347
automatic keyboard display 347
remote the combo box 347
Multiple Instruction Multiple Data (MIMD) processors 225
multiple policies
evaluating 151
exceptions, implementing 153, 154
implementing 151
priorities, implementing 151
priority, changing using Microsoft Group Policy Editor 152
priority, changing using Studio 152
multistream traffic policies
multi-port policy 342
multi-stream (computer) 342
multi-stream (user) 342

N

NetScaler® 238
NetScaler® ADC 239
NetScaler® CloudBridge™ Connector VPN Tunnel 242
NetScaler Gateway™
installation task list 275
installing 275
SSL, enabling 353, 374
used, for securing XenDesktop 275
NetScaler Gateway™ connection
testing 365
testing, with Windows client 365-372
NetScaler Gateway™
policy filters, implementing 158-160
used, for implementing policies 158
NetScaler Insight™ Center
about 310
accessing 312
enabling, in Director 314
HDX Insight™ 310
importing 310
Web Insight 310
NetScaler IP Address (NSIP) 366

network attached printers
using 162
Network Attached Storage (NAS) 194
network priorities, HDX 3D
configuring 232

O

ObserveIT
URL 238
offline plugin 269
online plugin 258
online plugin, functions
Desktop Lock, using 267, 268
devices, accessing 261
Flash redirection, using 263
keyboard input 268
local file access, controlling 260
microphones and webcams, accessing 262, 263
toolbar, moving 259
USB devices, accessing 261, 262
virtual desktop resolution, changing 259
virtual desktops, disconnecting from 265, 266
virtual desktops, logging off 265
virtual desktops, printing in 268
virtual desktops, restarting 266
virtual desktops, switching between 264
workspace control, using 258
OpenGL 210
optimizing
printing performance 174
Organizational Unit (OU) 95
OU-based controller discovery
performing 135
used, for moving controller 136

P

personal cloud 324
Personal vDisk (PvD)
about 17, 100, 196
adding, to hosts 101
automatic resizing, disabling 104
content 101
enabling, for using with master image 102
managing 100
parts 197
space available for applications, adjusting 103
troubleshooting 318
updating 102
user profiles, reallocating 104
Personal vDisk (PvD) partition 194
PHD Virtual
URL 207
picture archiving and communication system (PACS) 224
plugins
CloudBridge™ plugin 269
offline plugin 269
online plugin 258
using, with Receiver 258
policy filters
access control 148
Client IP address 148
client name 148
CloudBridge™ 148
desktop group 148
desktop type 148
Organizational Unit 149
tag 149
user or group 149
pool 117
power settings, for desktops
buffer 117
managing 117
partial power management 117
pool 117
power state timers 117
power state timers
changing 118
prerequisites, machine catalogs
about 82
computer accounts, creating 83
master images, creating 82
virtual machines, adding 83
virtual machines, configuring 83
printer drivers
mapping 172
Printer properties retention policy
modifying 173, 174
printers
autocreation 169

mapping 171
printing
about 161
default printing, using 163
locally attached printers, using 162
network attached printers, using 162
optimizing 174-179
policies 165
preferences, setting 164
working 161, 162
printing policies
auto-create client printers and Client printer redirection 342
auto-create generic universal printer 342
automatic installation of in-box printer drivers 345
client printer names 344
client printer redirection 344
Default Printer 344
direct connections to print servers 344
printer assignments 344
printer auto-creation event log performance 345
printer driver mapping and compatibility 344
printer properties retention 344
session printers 345
universal print driver usage 345
universal print EMF processing mode 343
universal printing image compression limit 343
universal printing optimization defaults 343
universal printing preview performance 343
universal printing print quality limit 344
universal print server enable 343
universal print server print data stream (CGP) port 343
universal print server print stream input bandwidth limit (kbps) 343
universal print server web service (HTTP/SOAP) port 343
printing preferences
document 165
server 165
session 165
setting 164
user device 164
private cloud 321, 322
Program Neighborhood Agent (PNA) 237
Provisioning Server (PVS) 194
public cloud 321, 323
public-facing server certificate
creating 357, 358
public server certificate
installing 359
linking, to root CA certificate 361
public-signed certificate
binding, to NetScaler Gateway™ VIP 382
installing 380, 381

Q

Quality of Service (QoS) 218

R

real-time multimedia transcoding
configuring 217-219
Receiver
about 17, 18, 253, 254
installing, on client devices 75, 76
plugins, using 258
running, on Android 271
running, on Apple 270
running, on BlackBerry devices 271
running, on Chromebook 271
running, on HTML5 compatible web browsers 271
running, on iOS 270
running, on Linux 271
running, on Mac 270
running, on Microsoft Windows 270
securing 299
Receiver settings
changing 255
changing, from client's desktop 256, 257
pushing, from server 255
redirection policies
auto connect client drives 335
client clipboard redirection 335
client COM port redirection 336

client drive redirection 336
client fixed drives and client drive redirection 336
client floppy drives and client drive redirection 336
client LPT port redirection 336
client network drives and client drive redirection 336
client optical drives and client drive redirection 336
client removable drives and client drive redirection 337
Client TWAIN device redirection and TWAIN compression redirection 337
client USB device redirection 337
client USB device redirection rules 337
host to client redirection 337
preserve client drive letters 337
TWAIN compression level 337
use asynchronous writes 337
requisites, HDX 3D
clients 225
server 225
Resultant Set of Policies
Citrix® Group Policy Modeling Wizard, running 154, 155
comparing 157
evaluating 154
Microsoft Group Policy Results tool, running 155
policy scenarios, troubleshooting 156, 157
root CA and server certificate bindings
viewing 362
root CA and server certificates
binding, to NetScaler Gateway™ VIP 362, 363
testing 364, 365
root CA certificate
installing 359, 360

S

SDK script
creating 248
SecureICA feature
about 116
enabling 116, 297
Secure Sockets Layer (SSL)
using, on controllers 136
Secure Ticket Authority (STA) 273, 297
Security IDs (SIDs) 129
security, XenDesktop™. *See* **XenDesktop™ security**
self-signed root CA certificate
creating 354-356
server host 22
server load, delivery groups
concurrent logon tolerance setting 121
managing 120
server load index 121
server OS machines 14
server side, XenDesktop architecture
about 13
Active Directory 15
database 13
Delivery Controller 13
desktop 15
desktop OS machines 15
Director 14
Edgesight® 15
FlexCast® 16
Hypervisor 13
server OS machines 14
storage 17
StoreFront 14
Studio 14
Virtual Desktop Agent (VDA) 14
virtual machine (VM) 14
XenApp® 15
session policies
disconnected session timer 347
session connection timer 347
session connection timer interval 348
session idle timer 348
session idle timer interval 348
session reliability connections 348
session reliability port number 348
session reliability timeout 348
single sign-on 348
single sign-on central store 348
session printers policy
enabling 172
sessions
troubleshooting 317

Set-BrokerDesktopGroup cmdlet 249
ShutdownDesktopsAfterUse property 117
Silicon Graphics Inc. (SGI) 210
Single-Hop DMZ architecture
about 274
diagrammatic representation 274
Single Instruction Multiple Data (SIMD) 224
Site debugging tools
troubleshooting 252
SiteDiag
about 252
URL 252
Smart Access
about 114
used, for restricting user access 114
smart cards
using, in security 302
snap-ins, XenDesktop® SDK
Citrix.AdIdentity.AdminV2 246
Citrix.AppV.Admin.V1 246
Citrix.Broker.Admin.V2 246
Citrix.Configuration. Citrix.Configuration-Logging.Admin.V1Admin.V2 246
Citrix.ConfigurationLogging.Admin.V1 246
Citrix.DelegatedAdmin.Admin.V1 246
Citrix.EnvTest.Admin.V1 246
Citrix.Host.Admin.V2 246
Citrix.MachineCreation.Admin.V2 246
Citrix.Monitor.Admin.V1 246
Citrix.Storefront.Admin.V1 246
Software Development Kits (SDKs) 243
SQL Server
backup, performing 203, 204
restore, performing 205
SSL
enabling, on NetScaler Gateway 353
Storage Area Networking (SAN) 194
StoreFront
about 14, 236, 239
securing 298
StoreFront 2.1 18, 19
StoreFront server
configuring 56-58
StoreFront (XD2)
components, installing 41
installing 41, 46
server certificate, creating 41, 42
Site binding, adding 41-45
Studio
about 14, 304
role-based administration, configuring 304
used, for managing XenDesktop® Site 304
Studio and Director, securing
about 300
IIS, installing 300
Studio User Interface (UI) 244
system requisites, HDX
about 210
Flash Media 211
server 211
user device 210
Windows Media 211
system requisites, XenDesktop Site
Active Directory 23
Citrix Receiver™ 18
databases 19
Delivery Controller 20
Director 21
server host 22
StoreFront 2.1 18
Studio 20
Virtual Delivery Agent (VDA) 21, 22

T

thin provisioning 194
third-party tools 318, 319
time zone policies
estimate local time for legacy clients 349
use local time of client 349
transcoding 217
transrating 217
troubleshooting
XD PowerShell SDK used 249
troubleshooting, with XD PowerShell SDK
controller cmdlets 250, 251
desktop cmdlets 250
Site debugging tools 252
troubleshooting, XenDesktop
applications 316
desktops 317
HDX™ 317

performing 315
Personal vDisks 318
sessions 317
users 316

U

unfiltered policies, XenDesktop® policies 149
Universal Print Driver
about 166
enabling 166
Universal Print Server
about 166
enabling 166
policy, enabling 167, 168
policy, searching 168
USB mass storage 207
USB Support virtualization
about 181
USB devices 181, 182
user data, delivery groups
exporting 119
importing 118
user device, printing preferences 164
user experience, machine catalogs
about 87
user types 87
users
troubleshooting 316
user types
random 87
static 87

V

VDI storage
about 193
visual representation 197
Verisign 373
vGPU 225
video usability points
compression 190
FPS 190
presence 190
webcam 190
Virtual CloneDrive 76
Virtual Delivery Agent (VDA) 21, 22, 50, 217
Virtual Desktop Agent (VDA) 14
Virtual Desktop Infrastructure (VDI)
about 86, 235
Provisioning Services™ (PVS) 86
Virtual desktop storage requirements
dedicated desktop model 199, 200
dedicated shared desktop model 201, 202
shared hosted desktop model 203
Virtual Graphics Processing Unit. *See* **vGPU**
virtualization 321
virtual machine (VMs)
about 14
backing up 206
restoring 206
voice usability points
about 190
codec 190
Headset 190
USB Headset 190

W

WANScaler 240
webcams, HDX 3D
configuring 231
Windows client
used, for testing NetScaler Gateway™ connection 365
Windows desktop, machine catalogs
creating 85
Windows Media
about 215
client-side fetching, configuring 215, 216
real-time multimedia transcoding, configuring 217-219
Windows Server, machine catalogs
creating 85

X

XD1 controller
components, installing 33-36
installing 32
Site, configuring 38-40
XenApp®
SmartAuditor feature 238
XenApp® and XenDesktop®

differences 236
XenDesktop®
administrative roles 237
administrator roles 140
backup 203
Citrix® Studio 237
discontinued features 237
domain certificate authority, creating 327-329
features 236
improvements 236
managing 303
modifications 238
monitoring 303
printing 161
StoreFront 236
third-party tools 318
troubleshooting 315
XenDesktop® architecture
client side 17
server side 13
XenDesktop® installation
application delivery master images, creating 48, 49
applications, testing 77, 78
Citrix Receiver™, installing on client devices 75, 76
delivery groups, creating 67
desktops, testing 77
Director (XD3), installing 47
installation checkpoint 72
inventory checklist 30
planning 30
StoreFront server, configuring 56-58
StoreFront (XD2), installing 41
task list 30
testing 77
Virtual Delivery Agent, installing on master images 49
virtual desktop, creating 48
XD1 controller, installing 32
XenDesktop® policies
about 139
accessing 142
applying 147
audio policies 333
autoreconnection policies 346
bandwidth policies 334
best practices, for designing policy settings 147
caching policies 342
configuring 146
configuring, Citrix® Studio used 331
configuring, Microsoft Group Policy Management editor used 332
creating 143
creating, in Microsoft Group Policy Editor 144
creating, in Studio 144
default values, using 147
Delivery Agent policies 350
desktop UI policies 338
filter modes 150
filters, using 148
graphics and multimedia policies 338
HDX™ 3D policies 351
ICA policies 345
implementing, NetScaler Gateway™ used 158
keep alive policies 346
load management policies 349
machine policy settings 141
mobility policies 347
multistream traffic policies 342
navigating 141
policy settings, configuring 146
printing policies 342
redirection policies 335
searching 143
session policies 347
time zone policies 349
unfiltered policies 149
user policy settings 141
working with 141
XenDesktop® SDK
cmdlet, running 247
snap-ins 246
using 247
XenDesktop® security
about 273
controller, securing 299
DMZ 274
ICA/HDX protocols, securing 297
NetScaler Gateway™, used 275

Receiver, securing 299
smart cards, using 302
STA 297
StoreFront, securing 298
Studio and Director, securing 300
XenServer® communications, securing 300
XenDesktop® security, with NetScaler Gateway™
NetScaler®, configuring for ICA® proxy 286, 287
NetScaler Gateway™, configuring for redirecting to StoreFront 284, 285
NetScaler Gateway™ virtual server, creating 279-283
NetScaler® license, installing 276-278
NetScaler VPX™, importing into XenServer 276
SSL certificate, installing 279
StoreFront certificate, exporting 291-294
StoreFront certificate, importing 294-297
StoreFront connection, configuring to NetScaler Gateway™ 288-291
XenDesktop® Site
concepts 13
core components 11
designing 23
diagrammatic representation 12
managing, Studio used 304
monitoring, Director used 305-308
scenario 23
system requisites 18
terminology 13
XenDesktop® Software Development Kit (SDK) 115
XenDesktop® storage considerations
about 194
desktop storage 194
High Availability (HA) 195
IOPS 196
performance 196
Personal vDisk 196
XenDesktop® storage requirements 198
XenDesktop® Studio
versus Microsoft Group Policy Editor 140
XenDesktop®, using USB redirection
about 183
USB automatic redirection, using 189
USB devices mapping, preventing 187
USB mass storage, using 187
USB support, enabling 184-186
voice and video, using 189, 190
XenApp, versus XenDesktop 189
Xenpipe scenario 23
XenServer® communications
securing 300, 301
XenServer® GPU pass-through 226

Thank you for buying

Getting Started with XenDesktop® 7.x

About Packt Publishing

Packt, pronounced 'packed', published its first book "Mastering phpMyAdmin for Effective MySQL Management" in April 2004 and subsequently continued to specialize in publishing highly focused books on specific technologies and solutions.

Our books and publications share the experiences of your fellow IT professionals in adapting and customizing today's systems, applications, and frameworks. Our solution based books give you the knowledge and power to customize the software and technologies you're using to get the job done. Packt books are more specific and less general than the IT books you have seen in the past. Our unique business model allows us to bring you more focused information, giving you more of what you need to know, and less of what you don't.

Packt is a modern, yet unique publishing company, which focuses on producing quality, cutting-edge books for communities of developers, administrators, and newbies alike. For more information, please visit our website: `www.packtpub.com`.

About Packt Enterprise

In 2010, Packt launched two new brands, Packt Enterprise and Packt Open Source, in order to continue its focus on specialization. This book is part of the Packt Enterprise brand, home to books published on enterprise software – software created by major vendors, including (but not limited to) IBM, Microsoft and Oracle, often for use in other corporations. Its titles will offer information relevant to a range of users of this software, including administrators, developers, architects, and end users.

Writing for Packt

We welcome all inquiries from people who are interested in authoring. Book proposals should be sent to `author@packtpub.com`. If your book idea is still at an early stage and you would like to discuss it first before writing a formal book proposal, contact us; one of our commissioning editors will get in touch with you.

We're not just looking for published authors; if you have strong technical skills but no writing experience, our experienced editors can help you develop a writing career, or simply get some additional reward for your expertise.

[PACKT] enterprise
professional expertise distilled
PUBLISHING

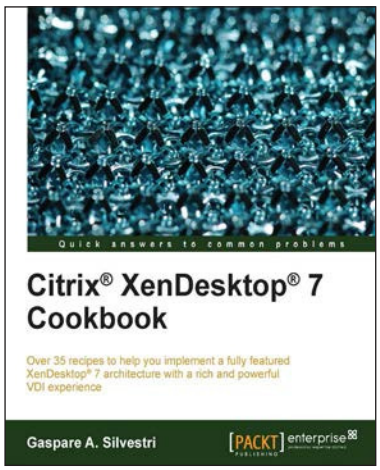

Citrix® XenDesktop® 7 Cookbook

ISBN: 978-1-78217-746-3 Paperback: 410 pages

Over 35 recipes to help you implement a fully featured XenDesktop® 7 architecture with a rich and powerful VDI experience

1. Implement the XenDesktop 7 architecture and its satellite components.
2. Learn how to publish desktops and applications to the end-user devices, optimizing their performance and increasing the general security.
3. Designed in a manner which will allow you to progress gradually from one chapter to another or to implement a single component only referring to the specific topic.

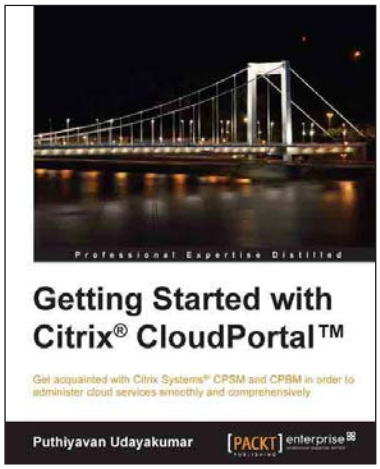

Getting Started with Citrix® CloudPortal™

ISBN: 978-1-78217-682-4 Paperback: 128 pages

Get acquainted with Citrix Systems® CPSM and CPBM in order to administer cloud services smoothly and comprehensively

1. Overview of CPSM and CPBM architectures, and planning CPSM and CPBM.
2. Become efficient in product management, workflow management, and billing and pricing management.
3. Provision services efficiently to cloud consumers and clients.

[PACKT] enterprise
professional expertise distilled
PUBLISHING

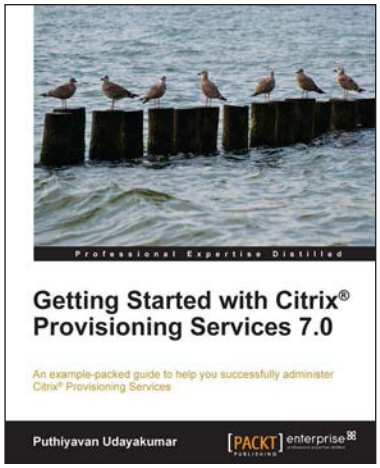

Getting Started with Citrix® Provisioning Services 7.0

ISBN: 978-1-78217-670-1 Paperback: 134 pages

An example-packed guide to help you successfully administer Citrix® Provisioning Services

1. Install and configure Citrix Provisioning Services quickly and efficiently.
2. Master the architecture of Citrix Provisioning Services.
3. Successfully manage and operate Citrix Provisioning Services.

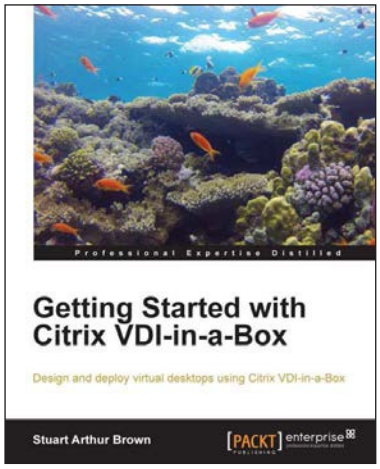

Getting Started with Citrix VDI-in-a-Box

ISBN: 978-1-78217-104-1 Paperback: 86 pages

Design and deploy virtual desktops using Citrix VDI-in-a-Box

1. Design a Citrix VDI-in-a-Box solution.
2. Get the budget for Citrix VDI-in-a-Box by building a case.
3. Implement a Citrix VDI-in-a-Box Proof of Concept and Citrix VDI-in-a-Box solution.