
PlateSpin® Migrate 12.1

User Guide

May 2016

Legal Notice

For information about legal notices, trademarks, disclaimers, warranties, export and other use restrictions, U.S. Government rights, patent policy, and FIPS compliance, see <https://www.netiq.com/company/legal/>.

Copyright © 2016 NetIQ Corporation. All Rights Reserved.

Third-Party Software

Please refer to the *PlateSpin Third-Party License Usage and Copyright* (https://www.netiq.com/documentation/platespin_licensing/platespin_licensing_qs/data/platespin_licensing_qs.html) page for information about third party software used in PlateSpin Migrate.

Contents

About NetIQ Corporation	9
About This Guide	11
1 PlateSpin Migrate Feature Overview	13
1.1 About Workload Portability	13
1.1.1 Business Applications for PlateSpin Migrate	14
1.1.2 Workload Migration Tasks	15
1.1.3 Migration Tasks Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface	15
1.2 Supported Configurations	17
1.2.1 Supported Source Workloads	17
1.2.2 Supported Target Virtualization Platforms	20
1.3 Supported Transfer Methods	21
1.3.1 File-Level Transfer (Live)	21
1.3.2 Block-Level Transfer (Live)	22
1.3.3 Offline Transfer with Temporary Boot Environment	22
1.4 Security and Privacy	23
1.4.1 Security of Workload Data in Transmission	23
1.4.2 Security of Client-Server Communications	24
1.4.3 Security of Credentials	24
1.4.4 User Authorization and Authentication	24
1.5 Performance	24
1.5.1 About Product Performance Characteristics	24
1.5.2 Data Compression	25
1.5.3 Bandwidth Throttling	25
1.5.4 Scalability	25
2 Post-Installation Configuration	27
2.1 PlateSpin Migrate Product Licensing	27
2.1.1 Activating Your Product License	27
2.1.2 How Licensing Works	29
2.1.3 Using the License Manager	30
2.1.4 License Splitting	31
2.2 Setting Up User Authorization and Authentication	31
2.2.1 PlateSpin Migrate Roles	32
2.2.2 Assigning PlateSpin Migrate Roles to Windows Users	34
2.2.3 Setting Up PlateSpin Migrate Multitenancy on VMware	34
2.2.4 Setting Up User Activity Logging	39
2.3 Access and Communication Requirements across your Migration Network	40
2.3.1 Requirements for Discovery	40
2.3.2 Requirements for Migration	41
2.3.3 Migrations Across Public and Private Networks through NAT	42
2.4 Configuring PlateSpin Migrate Default Options By Using the PlateSpin Migrate Client	43
2.4.1 General Options	44
2.4.2 Notification Service	45
2.4.3 User Activity Logging	46
2.4.4 Default Job Values	47
2.4.5 Source Service Defaults	48
2.4.6 Target Service Defaults	49
2.5 Configuring PlateSpin Server Behavior through Configuration Parameters	50

2.5.1	Changing the Adapter Type Used during the Target Take Control Process of Workload Migration to a Target VM on a Hyper-V Host	50
2.5.2	Increasing the Size Limit on Post-Migration Actions Uploaded to the PlateSpin Server	51
2.5.3	Optimizing Data Transfer over WAN Connections	51
2.5.4	Other Use Cases for Making Modifications in the Web Config Tool (Advanced)	55

3 Getting Started with PlateSpin Migrate 57

3.1	Connecting to a PlateSpin Server	57
3.2	About the PlateSpin Migrate Client User Interface	59
3.2.1	Servers View	60
3.2.2	Jobs View	63
3.2.3	Tasks Pane	64
3.2.4	Status Bar	64
3.3	Discovering Source Workloads and Migration Targets	64
3.3.1	Pre-populating the Servers View with Windows Computers in the Network	65
3.3.2	Discovering the Details of Source Workloads and Targets	65
3.3.3	Discovering All Machines in a Windows Domain	66
3.3.4	Refreshing Machine Details	66
3.3.5	Resetting the Mechanism Used to Authenticate Access to ESX Hosts	67
3.3.6	Discovering Target Physical Machines	67
3.3.7	Discovering, Registering and Preparing a VM for a Server Sync Job	69
3.3.8	Discovery Guidelines for Machine Types and Credentials	70
3.4	Undiscovering Source Workloads and Migration Targets	71
3.5	Analyzing Machines with PlateSpin Analyzer	71
3.6	Working with Device Drivers	72
3.6.1	Packaging Device Drivers for Windows Systems	73
3.6.2	Packaging Device Drivers for Linux Systems	73
3.6.3	Uploading Drivers to the PlateSpin Migrate Device Driver Database	74
3.6.4	Using the Plug and Play (PnP) ID Translator Feature	75
3.7	Managing Custom Actions	80
3.7.1	Managing Post-migration Actions (Windows and Linux)	81
3.7.2	Freeze and Thaw Scripting Capabilities (Linux Block-Level Migrations)	82
3.8	Setting Up, Executing, and Managing Jobs	83
3.8.1	Setting a Migration Job	83
3.8.2	Pre-Migration Job Validation	85
3.8.3	Saving Migration Jobs	85
3.8.4	Starting a Saved Migration Job in the Jobs View	85
3.8.5	Changing the Start Time of a Scheduled Migration Job	86
3.8.6	Viewing the Parameters of a Completed or In-Progress Migration Job	86
3.8.7	Canceling an In-Progress Job	86
3.8.8	Controlling the State of Source Workloads in the Jobs View	86
3.8.9	Generating Job Reports	86
3.8.10	Job Diagnostics	87

4 Working with the PlateSpin Migrate Web Interface 89

4.1	Planning For Migrating Workloads to Azure	90
4.1.1	Prerequisites for Migrating Workloads to Azure	91
4.2	Enabling an Azure Subscription to Deploy the Replication Environment VM	91
4.3	Azure Networking Guidelines	92
4.3.1	Windows Workloads with Multiple NICs	93
4.3.2	Private and Public IP Addresses for Workloads Connected on an Azure VPN	93
4.4	Accessing the PlateSpin Migrate Web Interface	93
4.5	Navigating the Web Interface	94
4.6	Dashboard	95
4.7	Discovering Source Workloads and Targets	95

4.8	Workloads	96
4.8.1	Adding or Discovering a Source Workload	97
4.8.2	Modifying the Listing of the Workloads	97
4.8.3	Understanding the Migration Tasks	97
4.8.4	Configuring the Workload for Migration	98
4.8.5	Preparing the Migration	107
4.8.6	Running the Migration	107
4.9	Targets	108
4.10	Tasks	108
4.11	Generating Reports Using the PlateSpin Migrate Web Interface	109
4.11.1	Generating Workload and Workload Migration Reports	109
4.11.2	Generating Licensing Report	109
4.12	Configuring the Settings For the Web Interface	110
4.12.1	Using Tags to Help Sort Workloads	110
4.12.2	Managing PlateSpin Migrate Security Groups and Workload Permissions	111
4.12.3	Setting Up Automatic Email Notifications of Events and Reports	112
4.12.4	Configuring Product Licenses	114
4.13	Configuring the Refresh Rates for the Web Interface	115
5	Workload Portability with Peer-to-Peer Workload Migrations	117
5.1	Migrating a Workload to a VM Host (P2V, V2V)	117
5.2	Converting a Workload to a Physical Machine (P2P, V2P)	118
5.2.1	Best Practices (X2P)	119
5.3	Synchronizing Workloads with Server Sync	119
5.3.1	Server Sync to a Virtual Target	119
5.3.2	Server Sync to a Physical Target	122
5.4	Semi-Automated Workload Virtualization Using the X2P Workflow	122
5.4.1	Migrating a Workload to Microsoft Hyper-V	123
5.4.2	Migrating a Workload to Citrix XenServer	125
5.4.3	Migrating a Workload to RHEL KVM	127
5.4.4	Migrating a Workload to the SLES 11 SP3 Xen Hypervisor	129
5.5	Migrating Windows Clusters	131
5.5.1	Planning Your Cluster Migration	132
5.5.2	Enabling or Disabling Windows Cluster Discovery	134
5.5.3	Adding Resource Name Search Values	134
5.5.4	Setting the Quorum Arbitration Timeout	135
5.5.5	Setting Local Volume Serial Numbers	135
5.5.6	PlateSpin CutOver	135
5.5.7	PlateSpin Cluster Migration	136
5.5.8	Migrating Windows Clusters with the Web Interface	136
5.5.9	Migrating Windows Clusters Using the Migrate Client	136
6	Windows Workload Portability with a PlateSpin Image	137
6.1	About PlateSpin Images	137
6.2	Designating a PlateSpin Image Server	137
6.3	Capturing a Workload to a PlateSpin Image	139
6.4	Deploying a PlateSpin Image	140
6.5	Managing PlateSpin Images	141
6.5.1	Moving Images from One PlateSpin Image Server to Another	142
6.5.2	Automating Image Operations	142
6.5.3	Browsing and Extracting Image Files	142
7	Essentials of Workload Migration Jobs	145
7.1	Selecting a License Key for a Migration Job	146

7.2	Configuring Automatic Email Alerts of Job Status and Progress	146
7.3	Specifying Conversion Options	147
7.4	Specifying End States for Source and Target Workloads.	147
7.5	Specifying Network Options	147
7.6	Specifying Credentials to Access Sources and Targets	148
7.7	Managing a Workload's Network Identity	149
7.7.1	Managing the Identity of Windows Workloads	149
7.7.2	Managing the Network Identity of Linux Workloads.	150
7.8	Job Scheduling.	150
7.9	Setting Up Migration Networking	151
7.9.1	Temporary (Take Control) Network Settings	151
7.9.2	Target Post-Migration Networking	153
7.9.3	TCP/IP and Advanced Network Settings	155
7.10	Configuring Target Virtual Machines	157
7.10.1	Virtualization Platform-Specific Virtual Machine Configuration Options.	158
7.11	Handling Operating System Services and Applications	164
7.11.1	Handling the Startup Mode of Services (Windows Targets).	165
7.11.2	Handling Source Workload Services or Daemons During Live Transfer (Windows and Linux)	166
7.11.3	Viewing Windows System Files Selected for Replacement During a Migration	167
7.11.4	Handling the Run Level of Daemons (Linux Targets)	168
7.11.5	Handling Virtualization Enhancement Software.	168
7.12	Handling a Workload's Storage Media and Volumes.	169
7.12.1	Storage Layout and Volume Configuration	171
7.12.2	Volume Configuration in Server Sync	177
7.13	Including a Custom Post-migration Action in a Migration Job	178
8	MigrateAgent Utility	181
A	Linux Distributions Supported by Migrate	185
A.1	Analyzing Your Linux Workload	185
A.1.1	Determining the Release String.	185
A.1.2	Determining the Architecture	185
A.2	Pre-compiled "blkwatch" Driver Support (Linux Distros)	186
B	Using the PlateSpin Migrate Command Line Interface	195
B.1	Where Is the Tool Located?	195
B.2	Before You Use the Tool.	195
B.2.1	Pre-configuring the Migrate Server Values for CLI	196
B.2.2	Becoming Familiar with the Commands	196
B.3	Configurable .ini Files (Jobs) You Can Use with the Tool.	198
B.3.1	Conversion Jobs	199
B.3.2	ServerSync Jobs	206
B.3.3	Imaging Jobs	210
C	Rebranding the PlateSpin Migrate Web Interface	221
C.1	Rebranding the Interface Using Configuration Parameters	221
C.1.1	PlateSpin Migrate Web Interface with Configurable Elements Labeled	222
C.1.2	PlateSpin Server Configuration Settings	222
C.2	Rebranding the Product Name in the Windows Registry	224

D	Synchronizing Serial Numbers on Cluster Node Local Storage	227
E	Frequently Asked Questions	229
F	Troubleshooting PlateSpin Migrate	231
F.1	Discovery	231
F.2	Peer-to-Peer Migrations (Windows)	232
F.3	Using Images	233
F.4	Post-migration Cleanup of Source Workloads	234
F.4.1	Cleaning Up Windows Workloads	234
F.4.2	Cleaning Up Linux Workloads	235
F.5	Shrinking the PlateSpin Migrate Databases	236
G	Reference of Tables	237
	Glossary	239

About NetIQ Corporation

NetIQ, an Attachmate business, is a global leader in systems and security management. With more than 12,000 customers in over 60 countries, NetIQ solutions maximize technology investments and enable IT process improvements to achieve measurable cost savings. The company's portfolio includes award-winning management products for IT Process Automation, Systems Management, Security Management, Configuration Audit and Control, Enterprise Administration, and Unified Communications Management. For more information, please visit www.netiq.com.

Contacting Sales Support

For questions about products, pricing, and capabilities, please contact your local partner. If you cannot contact your partner, please contact our Sales Support team.

Worldwide:	www.netiq.com/about_netiq/officelocations.asp
United States and Canada:	1-888-323-6768
Email:	info@netiq.com
Web Site:	www.netiq.com

Contacting Technical Support

For specific product issues, contact our Technical Support team.

Worldwide:	www.netiq.com/support/contactinfo.asp
North and South America:	1-713-418-5555
Europe, Middle East, and Africa:	+353 (0) 91-782 677
Email:	support@netiq.com
Web Site:	www.netiq.com/support

Contacting Documentation Support

Our goal is to provide documentation that meets your needs. If you have suggestions for improvements, click **comment on this topic** at the bottom of any page in the HTML versions of the documentation posted at www.netiq.com/documentation. You can also email Documentation-Feedback@netiq.com. We value your input and look forward to hearing from you.

Contacting the Online User Community

Qmunity, the NetIQ online community, is a collaborative network connecting you to your peers and NetIQ experts. By providing more immediate information, useful links to helpful resources, and access to NetIQ experts, Qmunity helps ensure you are mastering the knowledge you need to realize the full potential of IT investments upon which you rely. For more information, please visit <http://community.netiq.com>.

About This Guide

This guide provides information about using PlateSpin Migrate.

- ◆ Chapter 1, “PlateSpin Migrate Feature Overview,” on page 13
- ◆ Chapter 2, “Post-Installation Configuration,” on page 27
- ◆ Chapter 3, “Getting Started with PlateSpin Migrate,” on page 57
- ◆ Chapter 4, “Working with the PlateSpin Migrate Web Interface,” on page 89
- ◆ Chapter 5, “Workload Portability with Peer-to-Peer Workload Migrations,” on page 117
- ◆ Chapter 6, “Windows Workload Portability with a PlateSpin Image,” on page 137
- ◆ Chapter 7, “Essentials of Workload Migration Jobs,” on page 145
- ◆ Chapter 8, “MigrateAgent Utility,” on page 181
- ◆ Appendix A, “Linux Distributions Supported by Migrate,” on page 185
- ◆ Appendix B, “Using the PlateSpin Migrate Command Line Interface,” on page 195
- ◆ Appendix C, “Rebranding the PlateSpin Migrate Web Interface,” on page 221
- ◆ Appendix D, “Synchronizing Serial Numbers on Cluster Node Local Storage,” on page 227
- ◆ Appendix E, “Frequently Asked Questions,” on page 229
- ◆ Appendix F, “Troubleshooting PlateSpin Migrate,” on page 231
- ◆ Appendix G, “Reference of Tables,” on page 237
- ◆ “Glossary” on page 239

Audience

This guide is intended for IT staff, such as data center administrators and operators, who use PlateSpin Migrate in their ongoing workload migration projects.

Additional Documentation

This guide is part of the PlateSpin Migrate documentation set. For a complete list of publications supporting this release, visit the [PlateSpin Migrate Online Documentation Web Site](#).

Documentation Updates

The most recent version of this guide can be found at the [PlateSpin Migrate Online Documentation Web Site](#).

1 PlateSpin Migrate Feature Overview

PlateSpin Migrate enables you to migrate heterogeneous workloads across x86-based physical, virtual, image, and cloud infrastructures in your data center. It decouples the workload infrastructure from its software (operating system, applications, and data) to allow any-to-any migrations. Migrate provides tools to easily discover workloads and hosts in your environment. You can efficiently plan, configure, execute, and test migrations from a single client or web-based interface. A visual dashboard tracks project health and metrics. With Migrate, you can dramatically increase project predictability, migration speed, and success ratios, which help reduce the costs for your migration projects.

- ◆ [Section 1.1, “About Workload Portability,” on page 13](#)
- ◆ [Section 1.2, “Supported Configurations,” on page 17](#)
- ◆ [Section 1.3, “Supported Transfer Methods,” on page 21](#)
- ◆ [Section 1.4, “Security and Privacy,” on page 23](#)
- ◆ [Section 1.5, “Performance,” on page 24](#)

1.1 About Workload Portability

PlateSpin Migrate automates the migration of workloads among physical, virtual machine, volume imaging, and cloud. PlateSpin Migrate 12.1 introduces support for migration of workloads to Microsoft Azure.

Figure 1-1 Workload Portability

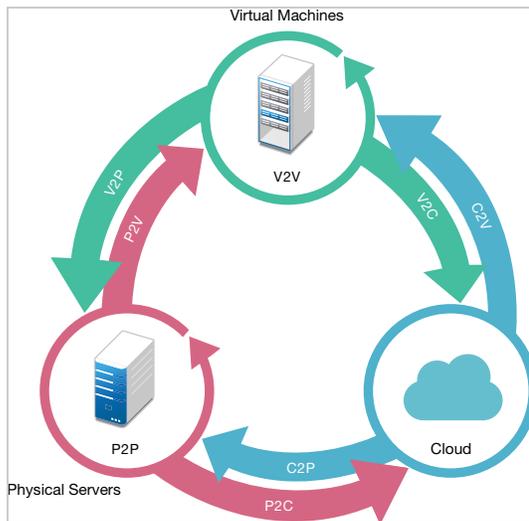


Table 1-1 Workload Portability Operations

Category of Operation	Migration Infrastructures
Workload Portability: Peer-to-peer	<ul style="list-style-type: none">◆ Physical to Virtual (P2V)◆ Virtual to Virtual (V2V)◆ Virtual to Physical (V2P)◆ Physical to Physical (P2P)
Workload Portability: Imaging	<ul style="list-style-type: none">◆ Physical to Image (P2I)◆ Virtual to Image (V2I)◆ Image to Virtual (I2V)◆ Image to Physical (I2P)
Workload Portability: Cloud	<ul style="list-style-type: none">◆ Physical to Microsoft Azure (P2C)◆ Virtual to Microsoft Azure (V2C)

PlateSpin Migrate supports multiple workload types and virtualization platforms. Imaging is supported for workloads with Microsoft Windows operating systems. For a more detailed list of supported workloads and infrastructures, see [“Supported Configurations” on page 17](#).

PlateSpin Migrate also provides a Web Interface that supports large scale migration of workloads to VMware containers. Do not use the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface interchangeably to perform the migration tasks during the migration process of a workload. NetIQ Corporation strongly recommends that you either use the PlateSpin Migrate Client or the PlateSpin Migrate Web Interface throughout the migration cycle of a workload. For information about the tasks you can perform using the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface, see [Section 1.1.3, “Migration Tasks Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface,” on page 15](#).

PlateSpin Migrate 12.1 enhances the Web Interface to let you migrate Windows and Linux workloads to Microsoft Azure. For information about the Windows and Linux workloads that you can migrate to Microsoft Azure, see [“Supported Windows and Linux Workloads For Migration to Microsoft Azure” on page 20](#)

- ◆ [Section 1.1.1, “Business Applications for PlateSpin Migrate,” on page 14](#)
- ◆ [Section 1.1.2, “Workload Migration Tasks,” on page 15](#)
- ◆ [Section 1.1.3, “Migration Tasks Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface,” on page 15](#)

1.1.1 Business Applications for PlateSpin Migrate

PlateSpin Migrate is designed to be used for the following scenarios:

- ◆ **Consolidation.** Automating large-scale migrations of physical machines to virtual machines, accelerating consolidation projects, and reducing administrative effort and errors.
- ◆ **Migration.** Moving fully configured workloads from old hardware to new hardware without rebuilding the entire software stack.
- ◆ **Test Lab Deployment.** Consolidating test lab workloads by running multiple virtual machines on a single VM host, quickly deploying virtual test lab environments with ease, and replicating an entire production environment in matter of hours or days.

- ♦ **Maintenance and Support Agreement Integrity.** De-virtualizing workloads along with the applications installed on them and moving them back to physical machines over the network so that the support agreements can remain valid.
- ♦ **Machine Provisioning.** Easily capturing an entire library of hardware-independent PlateSpin Images and deploying them to new infrastructures over the network without manually configuring the hardware, drivers, and so on.
- ♦ **Continuous Workload Optimization.** Moving workloads to and from any geographical location, onto any platform, in any direction. Workloads can be virtualized or de-virtualized during ongoing and continuous optimization of resources.

1.1.2 Workload Migration Tasks

PlateSpin Migrate enables you to define, save, schedule, execute, and monitor the following migration tasks.

Table 1-2 PlateSpin Migrate Workload Migration Tasks

Task	Description
Copy Workload	Results in a virtual or physical duplicate of a selected physical or virtual workload, except that the new workload is assigned a new network identity. Use this migration task when you intend to keep the source workload operational.
Move Workload	Results in an exact virtual or physical duplicate of a selected physical or virtual workload. Use this migration task when you intend to retire or repurpose the original infrastructure.
Server Sync	Synchronizes a virtual or physical workload with another virtual or physical workload without transferring the entire source volume data over the network.
Capture Image	Creates an image of a physical or virtual workload as a single entity, in PlateSpin Image format.
Deploy Image	Converts a PlateSpin Image into a booted or bootable workload on a physical or virtual machine.

1.1.3 Migration Tasks Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface

To migrate a workload, you should either use the PlateSpin Migrate Client or the PlateSpin Migrate Web Interface throughout the migration cycle of the workload.

The following table lists the tasks that you can perform using the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface:

Tasks	PlateSpin Migrate Client	PlateSpin Migrate web Interface
Monitor workload migration workflow	x	✓
Discover Windows standalone workloads	✓	✓
Discover Windows cluster workloads	✓	✓

Tasks	PlateSpin Migrate Client	PlateSpin Migrate web Interface
Discover Linux standalone workloads	✓	✓
Discover Linux cluster workloads	x	x
Discover target VMware hosts	✓	✓
Discover target non-VMware hosts	✓	x
Migrate to physical machines	✓	x
Migrate to VMware hosts	✓	✓
Migrate to non-VMware hosts	✓	x
Migrate to Azure Cloud	x	✓
Migrate to image	✓	x
Migrate Windows workloads with block-based transfer	x	✓
Migrate Linux workloads with block-based transfer	✓	✓
Migrate Windows workloads with file-based transfer	✓	✓
Migrate Linux workloads with file-based transfer	✓	x
Migrate Windows clusters with block-based transfer	✓	✓
Migrate workloads using live transfer	✓	✓
Migrate workloads using offline transfer	✓	x
Schedule incremental replication	x	✓
Migrate staged workloads using imaging	✓	x
Support post migration scripts	✓	x
Add new disks during migration	✓	x
Change disk volume mapping for target workload	✓	x
Migrate a VM to a vCenter folder	x	✓
Move a VM to a resource pool	✓	✓
Set compression level	✓	✓
Throttle bandwidth	✓	✓
Set encryption	✓	✓
Create tags	x	✓
View workload migration report	✓	✓
View workload migration status reports	x	✓
Add or remove licenses	✓	✓

Tasks	PlateSpin Migrate Client	PlateSpin Migrate web Interface
Check licenses status	✓	✓
Use security groups	x	✓

1.2 Supported Configurations

- ◆ [Section 1.2.1, “Supported Source Workloads,” on page 17](#)
- ◆ [Section 1.2.2, “Supported Target Virtualization Platforms,” on page 20](#)

1.2.1 Supported Source Workloads

PlateSpin Migrate supports the following operating system families:

- ◆ Microsoft Windows
- ◆ Linux

The following topics provide more details:

- ◆ [“Supported Microsoft Windows Workloads” on page 17](#)
- ◆ [“Supported Linux Workloads” on page 19](#)
- ◆ [“Supported Windows and Linux Workloads For Migration to Microsoft Azure” on page 20](#)

NOTE: Not all workloads are supported on all target VM platforms. For details, see [KB Article 7012976 \(https://www.netiq.com/support/kb/doc.php?id=7012976\)](https://www.netiq.com/support/kb/doc.php?id=7012976).

Supported Microsoft Windows Workloads

For the Windows platform, PlateSpin Migrate supports the following portability features:

- ◆ Workload portability: Peer-to-peer migrations (P2V, V2V, V2P, P2P).
- ◆ Peer-to-peer workload synchronization with Server Sync (P2V, V2V, P2P, V2P).
- ◆ Workload portability: Migrations to Microsoft Azure (P2C, V2C).

Table 1-3 Supported Microsoft Windows Workloads

Operating System	Remarks
Windows Server 2012 and 2012 R2	
Windows Server 2008 and 2008 R2	Including domain controller (DC) systems and Small Business Server (SBS) editions Migration of Windows Server 2008 R2 SP0 to Hyper-V is not supported because Microsoft no longer supports it. See Microsoft TechNet Website .
Windows Server 2003 R2 Windows Server 2003 SP 1 and later	

Operating System	Remarks
Windows 2000 Server SP 4 with Update Rollup 1	
Windows Server 2012 R2 Cluster Windows Server 2008 R2 Cluster Windows Server 2003 R2 Cluster	You can use the Web Interface or the Migrate Client to migrate supported Windows Server clusters. See Section 5.5, "Migrating Windows Clusters," on page 131 . PlateSpin Migrate does not support migration of Windows Server clusters to the following target infrastructures: <ul style="list-style-type: none"> ◆ Images ◆ Microsoft Azure
Windows 8 and 8.1	
Windows 7	
Windows Vista SP 1 and later	Supports only the Business/Enterprise and Ultimate versions.
Windows XP SP 1 and later	

NOTE: Only NTFS files systems are supported.

PlateSpin Migrate supports the following localized versions of source Windows workloads:

- ◆ English
- ◆ French
- ◆ German
- ◆ Japanese
- ◆ Chinese (traditional and simplified)

Workload Firmware (UEFI and BIOS) Support: PlateSpin Migrate mirrors the Microsoft support of UEFI or BIOS-based Windows workloads. It transfers workloads (both File and Block-based transfers are supported) from source to target while enforcing the supported firmware for the respective source and target operating systems. When any migration between UEFI and BIOS systems are initiated, Migrate analyzes the transition and alerts you about its validity.

NOTE: If you are migrating UEFI-based workload onto vSphere target container and you want to continue using the same firmware boot mode, you need to target a vSphere 5.0 container or newer.

The following are examples of Migrate behavior when doing conversion between UEFI and BIOS-based systems:

- ◆ When transferring a UEFI-based workload to a VMware vSphere 4.x container (which does not support UEFI), Migrate transitions the workload's UEFI firmware to BIOS firmware.
- ◆ When migrating a UEFI-based source on a BIOS-based target, Migrate converts the UEFI system's boot disks, which were GPT, to MBR disks. When migrating BIOS workload on a UEFI-based target, Migrate converts the BIOS system's boot disks, which are MBR, to GPT disks.

Supported Linux Workloads

For the Linux platform, PlateSpin Migrate the following portability features:

- ◆ Live peer-to-peer and Offline workload portability support (P2P, P2V, V2P, V2V), including workload synchronization with Server Sync.
- ◆ Support for EXT2, EXT3, EXT4, REISERFS, XFS Linux file systems.

IMPORTANT

- ◆ Workload imaging is not supported in Linux workloads.
- ◆ Migration of UEFI-based Linux workloads to Hyper-V target container is not supported.
- ◆ Conversion between UEFI and BIOS based Linux systems is not supported.

Some of the supported Linux versions require that you compile the PlateSpin `blkwatch` module for your specific kernel. Those workloads are called out explicitly.

For information about the pre-compiled versions of the `blkwatch` driver for many non-debug Linux distributions (32-bit and 64-bit) that PlateSpin Migrate includes, see [Appendix A, “Linux Distributions Supported by Migrate,” on page 185](#).

Table 1-4 Supported Linux Workloads

Distribution	Remarks
<ul style="list-style-type: none">◆ CentOS 4.x, 5.x, 6.x, 7 (For ESX)◆ CentOS 5.x, 6.x, 7 (For Hyper-V)	
Red Hat Enterprise Linux (RHEL) AS/ES/WS 4, 5, 5.1-5.11, 6, 6.1- 6.7, 7, 7.1, 7.2	For RHEL 6.7, 7, 7.1, and 7.2: Only BIOS-based workloads are supported.
SUSE Linux Enterprise Server (SLES) 9, 10, 11 (SP1, SP2, SP3, SP4)	The SLES 11 SP2 (32-bit) with kernel 3.0.13-0.27-pae is not supported. The kernel for this version of SLES must be upgraded to 3.0.51-0.7.9-pae so that conversion works. For SLES 11 SP4: Only BIOS-based workloads are supported.
Oracle Enterprise Linux (OEL)	<ul style="list-style-type: none">◆ Same level of support as that for workloads running RHEL.◆ Workloads using the Unbreakable Enterprise Kernel are not supported.
Red Hat Linux 8	32-bit systems only

NOTE

- ◆ Linux-based source workloads must be running a Secure Shell (SSH) server.
 - ◆ Live transfer is supported for these Linux workloads:
 - ◆ RHEL 4, 5, 6.x, 7.x
 - ◆ SLES 9, 10, 11
 - ◆ Migration of encrypted volumes is not supported.
-

Supported Windows and Linux Workloads For Migration to Microsoft Azure

PlateSpin Migrate 12.1 enhances the Web Interface to let you migrate the following Windows and Linux workloads to Microsoft Azure:

Windows:

- ◆ Microsoft Windows Server 2012 R2
- ◆ Microsoft Windows Server 2012
- ◆ Microsoft Windows Server 2008 R2

Linux:

- ◆ Red Hat Enterprise Linux (RHEL) 7.1
- ◆ Red Hat Enterprise Linux (RHEL) 6.7
- ◆ SUSE Linux Enterprise Server (SLES) 11 SP4
- ◆ SUSE Linux Enterprise Server (SLES) 11 SP3

NOTE: Migration of UEFI workloads to Microsoft Azure is not supported.

For information about using the Web Interface to migrate workloads to Microsoft Azure, see [Chapter 4, “Working with the PlateSpin Migrate Web Interface,”](#) on page 89.

1.2.2 Supported Target Virtualization Platforms

The following is a list of supported virtualization platforms. For more details on supported configurations, as well as the most up-to-date list, see [KB Article 7012976 \(https://www.netiq.com/support/kb/doc.php?id=7012976\)](https://www.netiq.com/support/kb/doc.php?id=7012976).

IMPORTANT: You need an OS license for the migrated target workload. For Azure target workloads, you must provide Azure with the license information or you will be charged for the OS license.

Table 1-5 Supported Target Virtualization Platforms

Platform	Notes
VMware vCenter 4.1, including Updates 1, 2, and 3	For creating the target VM disk using Raw Device Mapping (RDM), use Semi-Automated Workload Virtualization Using the X2P Workflow .
VMware vCenter 5.0, including Updates 1, 2, and 3	
VMware vCenter 5.1, including Updates 1 and 2	
VMware vCenter 5.5, including Updates 1 and 2	
VMware vCenter 6.0	

Platform	Notes
VMware ESX 4.1, including Updates 1, 2, and 3	All ESXi versions must have a paid license; migration is unsupported with these systems if they are operating with a free license. For creating the target VM disk using Raw Device Mapping (RDM), use Semi-Automated Workload Virtualization Using the X2P Workflow .
VMware ESXi 4.1, including Updates 1, 2, and 3	
VMware ESXi 5.0, including Updates 1, 2, and 3	
VMware ESXi 5.1, including Updates 1 and 2	
VMware ESXi 5.5, including Updates 1 and 2	
VMware ESXi 6.0	
Microsoft Windows Server 2012 Hyper-V	Migration of UEFI-based Linux workloads to Hyper-V target container is not supported.
Microsoft Windows Server 2012 R2 Hyper-V	
Microsoft Windows Server 2008 R2 Hyper-V	Supported through Semi-Automated Workload Virtualization Using the X2P Workflow . Migration of UEFI-based Linux workloads to Hyper-V target container is not supported.
Citrix XenServer 6, 6.1, 6.2, 6.5	Supported through Semi-Automated Workload Virtualization Using the X2P Workflow .
SUSE Linux Enterprise Server (SLES) 11 SP3 XEN, SLES 11 SP3 KVM	<ul style="list-style-type: none"> ◆ Fully virtualized guests are supported. ◆ Supported through Semi-Automated Workload Virtualization Using the X2P Workflow.
Red hat Enterprise Linux (RHEL) 6.4 - 7 KVM	Supported through Semi-Automated Workload Virtualization Using the X2P Workflow .

1.3 Supported Transfer Methods

Depending on the selected workload and the migration type, PlateSpin Migrate enables you to select different methods for transferring workload data from the source to the target.

For information on how to select a transfer method, see [“Specifying Conversion Options” on page 147](#).

- ◆ [Section 1.3.1, “File-Level Transfer \(Live\),” on page 21](#)
- ◆ [Section 1.3.2, “Block-Level Transfer \(Live\),” on page 22](#)
- ◆ [Section 1.3.3, “Offline Transfer with Temporary Boot Environment,” on page 22](#)

1.3.1 File-Level Transfer (Live)

The File-Based Live Transfer method, available for Windows workloads, copies data and replicates changes at the file level.

To ensure data consistency, this method leverages the Microsoft Volume Shadow Copy Service (VSS) if available. Many enterprise apps are integrated with VSS; for those which are not, PlateSpin Migrate provides the capability to briefly pause services while the VSS snapshot is captured, to ensure that the data of those applications is captured in a consistent state.

If VSS unavailable (for example, in workloads running Windows Server 2003 with no service packs or Windows Server 2000), PlateSpin Migrate monitors source volumes for changes while transferring data. When the initial transfer is complete, migrate re-sends any files that have changed. If the rate of file system changes is consistently high, data transfer is stopped and a job progress warning is shown.

You can configure your migration job to stop high-transaction services, such as Microsoft SQL Server or Microsoft Exchange Server, during the transfer (see [“Handling Source Workload Services or Daemons During Live Transfer \(Windows and Linux\)” on page 166](#)). This has two benefits:

- ♦ It ensures that the databases of these applications are transferred in a more consistent state.
- ♦ It reduces the rate of file system changes so that PlateSpin Migrate is able to keep up with them and complete the transfer.

This method might be appropriate for moderately active systems and it provides you with the capability to resize your volumes on the target workload.

1.3.2 Block-Level Transfer (Live)

The Block-Based Live Transfer method, available for both Windows and Linux workloads, enables PlateSpin Migrate to transfer data at the block level, providing an exact copy of the source workload.

For Windows workloads, PlateSpin Migrate leverages the Microsoft Volume Snapshot Service (VSS) (Windows 2003 SP1 and later) with applications and services that support VSS.

For Linux workloads, Migrate supports only block-based data transfer with a blkwatch driver. The Migrate distribution includes precompiled blkwatch drivers for workloads running the standard, non-debug kernels of supported Linux distributions. See [Section A.2, “Pre-compiled “blkwatch” Driver Support \(Linux Distros\),” on page 186](#).

If your workloads have a non-standard, customized, or newer kernel, you can build a custom blkwatch driver for your specific kernel. See [Knowledgebase Article 7005873 How to Build a Custom Block-Based Linux Kernel Driver \(https://www.netiq.com/support/kb/doc.php?id=7005873\)](https://www.netiq.com/support/kb/doc.php?id=7005873).

NOTE: Deployment or removal of the blkwatch driver is transparent, has no continuity impact, and requires no intervention and no reboot.

The blkwatch driver leverages LVM snapshots if they are available. Copying data from the snapshot helps avoid potential open file conflicts. See [Knowledgebase Article 7005872 Using LVM Snapshots for Migrating and Protecting Linux Workloads \(https://www.netiq.com/support/kb/doc.php?id=7005872\)](https://www.netiq.com/support/kb/doc.php?id=7005872). If LVM snapshots are not available, Migrate locks and releases each block in turn for data transfer.

The Block-Based Live Transfer method is the preferred data transfer method for both Windows and Linux workloads.

1.3.3 Offline Transfer with Temporary Boot Environment

This method enables PlateSpin Migrate to boot your source machine into a temporary pre-execution environment and transfer the data while the source is offline. This method is not applicable with the PlateSpin Migrate Web Interface.

NOTE: The Offline Transfer method lets you migrate the following workloads:

- ♦ Windows Server 2000 SP0 (or prior to SP4)
- ♦ Windows Server 2003 SP0

Before you use the Offline Transfer method to migrate a Windows Server 2003 workload, you must do the following:

1. Edit the `boot.ini` file on the workload to set the `/noexecute` parameter to `alwaysoff`.
2. Restart the workload.

The pre-execution environment underlying the Offline transfer method makes use of a Linux Ramdisk, which contains a minimal set of system files, drivers, and executables, sufficient for an initial, temporary boot. To ensure that the source operating system properly loads and operates in the temporary pre-execution environment, PlateSpin Migrate temporarily modifies its boot files and restores them to their original state after the pre-execution environment has successfully loaded.

The Ramdisk is also used to temporarily boot target physical machines in X2P migrations, as well as to boot target VMs in semi-automated migrations.

See also, [Section 3.3.6, “Discovering Target Physical Machines,” on page 67](#), and [“Semi-Automated Workload Virtualization Using the X2P Workflow” on page 122](#).

1.4 Security and Privacy

PlateSpin Migrate provides several features to help you safeguard your data and increase security.

- ♦ [Section 1.4.1, “Security of Workload Data in Transmission,” on page 23](#)
- ♦ [Section 1.4.2, “Security of Client-Server Communications,” on page 24](#)
- ♦ [Section 1.4.3, “Security of Credentials,” on page 24](#)
- ♦ [Section 1.4.4, “User Authorization and Authentication,” on page 24](#)

1.4.1 Security of Workload Data in Transmission

To make the transfer of your workload data more secure, you can configure your migration jobs to encrypt the data in transit to the target. When encryption is enabled, over-the-network data transfer from the source to the target is encrypted by using AES (Advanced Encryption Standard).

You can configure your PlateSpin Server to use a data encryption algorithm that is compliant with FIPS (Federal Information Processing Standards, Publication 140-2). If compliance with FIPS is required, it must be set up on your system prior to the PlateSpin Server installation. See [“Enabling Support for FIPS-Compliant Data Encryption Algorithms \(Optional\)”](#) in your *Installation Guide*.

For information on how to enable replication for a migration job, see [“Specifying Conversion Options” on page 147](#).

If FIPS is enabled in the source workload, you must do the following before you migrate the workload:

- 1 In the address bar of a web browser, type the following URL to launch the PlateSpin Server Configuration Settings page:

```
https://Your_PlateSpin_Server/platespinconfiguration/
```

- 2 Log in as Administrator.
- 3 Locate the `EnforceFIPSCompliance` parameter and click **Edit** to change its value to `true`.

- 4 Click Save.

After you modify the settings in the configuration tool, it might take up to 30 seconds for the change to take reflect on the interface. You need not reboot or restart the services.

- 5 Discover the FIPS enabled source workload.

1.4.2 Security of Client-Server Communications

Data transmission between the PlateSpin Server and the PlateSpin Migrate Client can be configured to use either HTTP (default) or HTTPS (Secure Hypertext Transfer Protocol). To secure data transmission between the client and the server, enable SSL on your PlateSpin Server host and use HTTPS when specifying the server URL. See [“Connecting to a PlateSpin Server” on page 57](#).

1.4.3 Security of Credentials

Credentials that you use to access sources and targets in workload migration jobs are:

- ◆ Cached, encrypted, and securely stored by the PlateSpin Migrate Client, by using operating system APIs.
- ◆ Stored in the PlateSpin Migrate database and are therefore covered by the same security safeguards that you have in place for PlateSpin Server hosts.
- ◆ Included within diagnostics, which are accessible to accredited users. You should ensure workload migration projects are handled by authorized staff.

1.4.4 User Authorization and Authentication

PlateSpin Migrate provides a role-based user authorization and authentication mechanism. See [Section 2.2, “Setting Up User Authorization and Authentication,” on page 31](#).

NOTE: If you have installed a PlateSpin Migrate Server localized for one language and a PlateSpin Migrate Client localized for a different language, do not use authorization credentials that include any language-specific characters. Using such characters in the login credentials causes miscommunication between the client and the server: the credentials are rejected as invalid.

1.5 Performance

- ◆ [Section 1.5.1, “About Product Performance Characteristics,” on page 24](#)
- ◆ [Section 1.5.2, “Data Compression,” on page 25](#)
- ◆ [Section 1.5.3, “Bandwidth Throttling,” on page 25](#)
- ◆ [Section 1.5.4, “Scalability,” on page 25](#)

1.5.1 About Product Performance Characteristics

The performance characteristics of your PlateSpin Migrate product depend on a number of factors, including:

- ◆ Hardware and software profiles of your source and target
- ◆ Hardware and software profiles of your PlateSpin Server host
- ◆ The specifics of your network bandwidth, configuration, and conditions

- ◆ The number of your source workloads' volumes and their sizes
- ◆ File density (number of files per unit of capacity) on your source workloads' volumes
- ◆ Source I/O levels (how busy your workloads are)
- ◆ The number of concurrent migrations and the number and type of the targets
- ◆ Whether data encryption is enabled or disabled
- ◆ Whether data compression is enabled or disabled

For planning large-scale workload migrations, you should perform a test migration of an average workload and use the result as a benchmark, fine-tuning your metrics regularly throughout the project. In addition to the data transfer process, also consider the other phases that a migration job goes through, as applicable to your project:

- ◆ Preparation and network setup
- ◆ Source workload and target machine discovery
- ◆ Target configuration

1.5.2 Data Compression

If necessary, PlateSpin Migrate can compress the workload data before transferring it over the network. This enables you to reduce the overall amount of data transferred during a workload migration job.

Compression ratios depend on the type of files on a source workload's volumes, and might vary from approximately 0.9 (100MB of data compressed to 90 MB) to approximately 0.5 (100MB compressed to 50MB).

NOTE: Data compression utilizes the source workload's processor power.

Data Compression can be configured per migration job. You can also specify a default compression value to be applied globally (see [“Default Job Values” on page 47](#)).

1.5.3 Bandwidth Throttling

PlateSpin Migrate enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration; you can specify a throughput rate for each migration job. This provides a way to prevent migration traffic from congesting your production network and reduces the overall load of your PlateSpin Server.

Bandwidth throttling is a parameter of a workload migration job's configuration properties. See [“Specifying Network Options” on page 147](#).

1.5.4 Scalability

You can set up multiple workload migrations and run them simultaneously. Scalability testing performed with VMware ESX hosts suggests the following benchmark recommendations:

- ◆ Multiple migrations to a single VM server: no more than 10
- ◆ Multiple migrations against multiple VM servers (ESX): no more than 40

For a Migrate Server with 8GB RAM and 4 CPUs, you can use the PlateSpin Migrate Web Interface to discover 200 workloads.

2 Post-Installation Configuration

This section provides information on typical, usually one-time configuration tasks following product installation. For installation information, see your *Installation Guide*.

- ◆ [Section 2.1, “PlateSpin Migrate Product Licensing,” on page 27](#)
- ◆ [Section 2.2, “Setting Up User Authorization and Authentication,” on page 31](#)
- ◆ [Section 2.3, “Access and Communication Requirements across your Migration Network,” on page 40](#)
- ◆ [Section 2.4, “Configuring PlateSpin Migrate Default Options By Using the PlateSpin Migrate Client,” on page 43](#)
- ◆ [Section 2.5, “Configuring PlateSpin Server Behavior through Configuration Parameters,” on page 50](#)

2.1 PlateSpin Migrate Product Licensing

This section provides information about licensing and activating your PlateSpin Migrate product, as well as managing your license keys using the PlateSpin Migrate Client.

You can also use the PlateSpin Migrate Web Interface to configure the product license. For more information about using the PlateSpin Migrate Web Interface for configuring the product license, see the *Working with PlateSpin Migrate Web Interface* section in the *PlateSpin User Guide* on the [Documentation Website](#).

NOTE: You cannot use the Licenses that you purchased for PlateSpin Migrate 9.3 and later versions with PlateSpin Migrate 9.2 and prior versions.

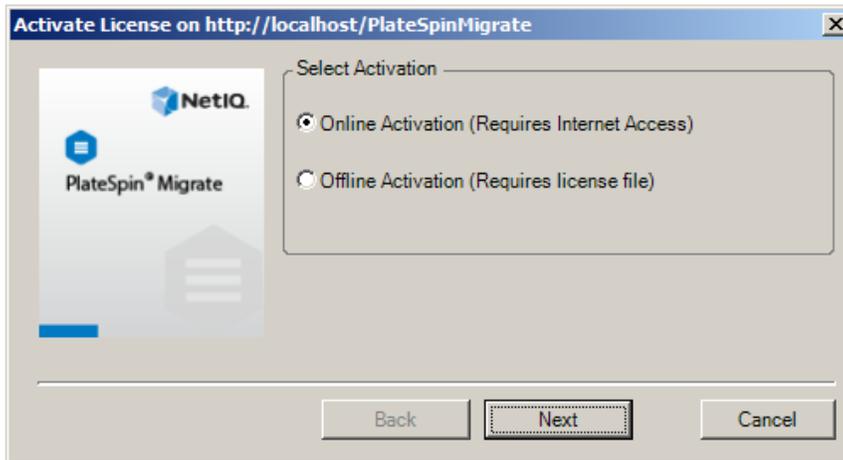
- ◆ [Section 2.1.1, “Activating Your Product License,” on page 27](#)
- ◆ [Section 2.1.2, “How Licensing Works,” on page 29](#)
- ◆ [Section 2.1.3, “Using the License Manager,” on page 30](#)
- ◆ [Section 2.1.4, “License Splitting,” on page 31](#)

2.1.1 Activating Your Product License

For product licensing, you must have a license activation code. If you do not have a license activation code, request one through the [Customer Center Web site \(https://www.netiq.com/center/\)](https://www.netiq.com/center/). A Micro Focus representative will contact you and provide the license activation code.

When you launch the PlateSpin Migrate Client for the first time after installation, the License Activation Wizard opens and prompts you to activate your product license.

Figure 2-1 License Activation Wizard



You have two options for activating your product license: online or offline.

- ◆ [“Online Activation” on page 28](#)
- ◆ [“Offline Activation” on page 28](#)

Online Activation

Online activation requires that your PlateSpin Migrate Client have Internet access.

NOTE: HTTP proxies might cause failures during online activation. If you are using an HTTP proxy server and are having problems with online activation, try the offline activation method.

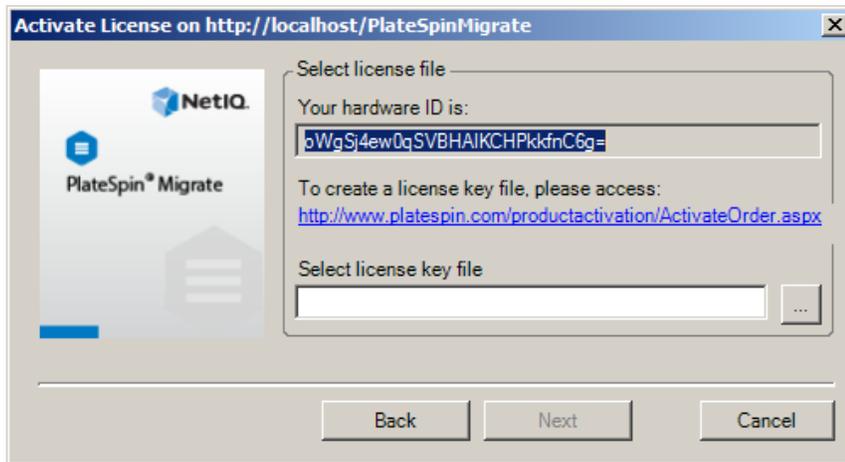
- 1 In the License Wizard, select the **Online Activation** option and click **Next**.
- 2 Enter the e-mail address that you provided when placing your order, and the activation code you received.

The PlateSpin Migrate Client obtains the required license over the Internet and activates the product.

Offline Activation

For offline activation, you obtain a license key over the Internet by using a machine that has Internet access.

- 1 In the License Wizard, select the **Offline Activation** option and click **Next**.
The Activate License dialog box is displayed:



- 2 Save your hardware ID for use in the next steps.
- 3 Use a computer with Internet access to obtain a license key through the [Web-based license activation utility \(http://www.platespin.com/productactivation/ActivateOrder.aspx\)](http://www.platespin.com/productactivation/ActivateOrder.aspx).

To obtain a license key, you must have a Novell account. If you are an existing PlateSpin customer and you don't have a Novell account, you must first create one. Use your existing PlateSpin username (a valid e-mail address registered with PlateSpin) as input for your Novell account username.

- 4 Save your new license key in a location accessible to your PlateSpin Migrate Client.
- 5 In the License Wizard, type the full path to, or browse to and select, the PlateSpin Migrate license file, then click **Next**.

The product is activated based on the selected license.

2.1.2 How Licensing Works

PlateSpin Migrate licenses are sold on a per-workload basis. A license entitles you to an unlimited number of migrations on a specific number of workloads. With every migration, a workload unit of the license is assigned to either the source or the target. The machine that has the workload unit assigned to it can subsequently be migrated an unlimited number of times. Each time a workload is assigned, the **Workloads remaining** number is decremented.

The following is a summary of workload assignment behavior by portability task.

Table 2-1 PlateSpin Migrate Workload License Assignment by Migration Type

Task	Workload Assignment Behavior
Copy Workload	A workload license remains with the source.
Move Workload	A workload license is transferred from the source to the target.
Server Sync	Not applicable
Capture Image	A workload license is assigned to the source and remains with it
Deploy Image	Not applicable

2.1.3 Using the License Manager

Use the License Manager to manage product licensing, add and delete license keys, generate licensing reports, view activation codes, and reset workload licensing. If you have multiple license keys and you want to select a specific, non-default license key for a particular migration job, see [Section 7.1, “Selecting a License Key for a Migration Job,”](#) on page 146.

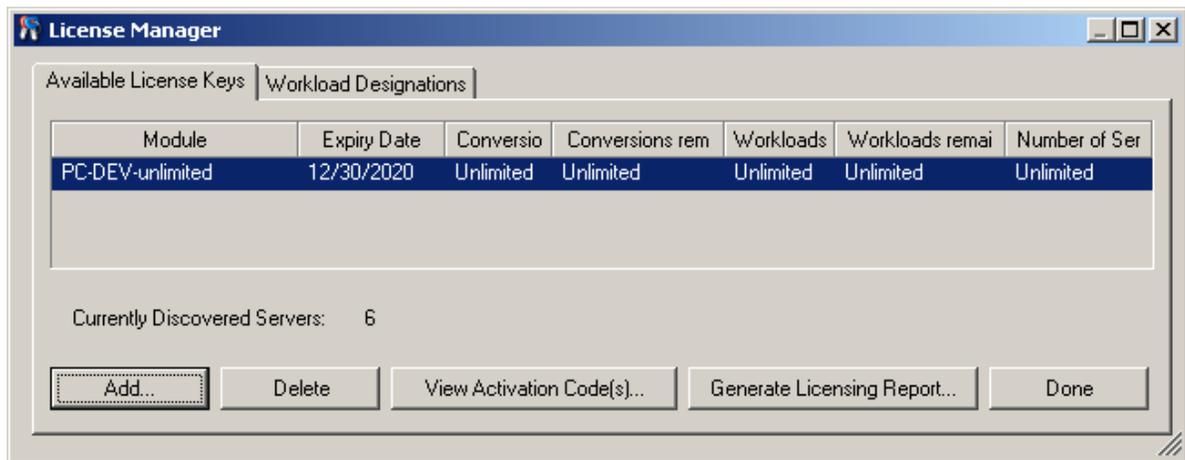
- ♦ “Managing Available License Keys” on page 30
- ♦ “Managing Workload Designations” on page 31

To access the License Manager, in your PlateSpin Migrate Client, click **Tools > License Manager**.

Managing Available License Keys

You can manage available license keys on the License Manager’s **Available License Keys** tab (in your PlateSpin Migrate Client, click **Tools > License Manager > Available License Keys**).

Figure 2-2 Available License Keys



The tab displays the license name (**Module**) along with its expiry date and entitlements. These depend on the license type. The **Number of Servers** column indicates the number of machines you can discover. This is generally the same as the number of machines that you can migrate. Use the buttons at the bottom for related license management tasks:

Table 2-2 License Manager Command Buttons

Command	Description
Add	Adds licenses.
Delete	Deletes expired licenses.
View Activation Code(s)	Select a license and click this button to see the activation code and the date it was activated.
Generate Licensing Report	Creates a *.ps1 file that is used by Technical Support to troubleshoot licensing issues.

Managing Workload Designations

You can manage workload licensing on the License Manager's **Workload Designations** tab (in your PlateSpin Migrate Client, click **Tools > License Manager > Workload Designations**).

The tab lists workloads with assigned licenses. In the PlateSpin Migrate Client Servers view, each of these servers has a key icon adjacent to it.

You can reset workload licensing so that a license is no longer assigned to a particular machine. For example, you might want to do this when decommissioning servers that are already in the inventory of the PlateSpin Server.

To reset workload licensing:

- 1 On the License Manager's **Workload Designations** tab, select the required workload and click **Transfer Selected Workload**.

The Transfer License dialog box is displayed.

- 2 Use the displayed **Workload Transfer Request** string to obtain a workload transfer code from the **License Entitlement Web portal** (<http://www.platespin.com/entitlementmgr/>). Log in with credentials associated with your purchase order.

You must have a Novell account. If you are an existing PlateSpin customer and you don't have a Novell account, you must first create one. Use your existing PlateSpin username (a valid e-mail address registered with PlateSpin) as input for your Novell account username.

- 3 Return to the License Manager and specify the newly obtained transfer code. Click **Next**.

PlateSpin Migrate resets the selected workload.

2.1.4 License Splitting

A license entitles you to one instance of PlateSpin Migrate per workload. Depending on the license you purchased, you can split a license either on a per-migration or a per-workload basis.

You can only split a license that has not yet been activated. For example, you can split a per-workload license of 1000 workloads into one license covering 400 workloads and another covering 600 workloads. You can split a per-migration license for 3000 migrations into one license for 1200 migrations and one license for 1800 migrations.

For assistance with multi-license scenarios, especially if you are uncertain how to utilize licenses across your network environment, see [KB Article 7920876](https://www.netiq.com/support/kb/doc.php?id=7920876) (<https://www.netiq.com/support/kb/doc.php?id=7920876>).

2.2 Setting Up User Authorization and Authentication

PlateSpin Migrate's user authorization and authentication mechanism is based on user roles, and controls application access and operations that users can perform. The mechanism is based on Integrated Windows Authentication (IWA) and its interaction with Internet Information Services (IIS).

NOTE: If you have installed a PlateSpin Migrate Server localized for one language and a PlateSpin Migrate Client localized for a different language, do not use authorization credentials that include any language-specific characters. Using such characters in the login credentials causes miscommunication between the client and the server: the credentials are rejected as invalid.

PlateSpin Migrate’s user auditing functionality is provided through the capability to log user actions (see “Setting Up User Activity Logging” on page 39).

- ◆ [Section 2.2.1, “PlateSpin Migrate Roles,” on page 32](#)
- ◆ [Section 2.2.2, “Assigning PlateSpin Migrate Roles to Windows Users,” on page 34](#)
- ◆ [Section 2.2.3, “Setting Up PlateSpin Migrate Multitenancy on VMware,” on page 34](#)
- ◆ [Section 2.2.4, “Setting Up User Activity Logging,” on page 39](#)

2.2.1 PlateSpin Migrate Roles

A PlateSpin Migrate role is a collection of PlateSpin Migrate privileges that entitle a particular user to perform specific actions. During installation, the PlateSpin Migrate installation program creates three local Windows groups on the PlateSpin Server host: PlateSpin Migrate Administrators, PlateSpin Migrate Power Users, and PlateSpin Migrate Operators. These groups map directly to the three PlateSpin Migrate roles that control user authorization and authentication:

Group for PlateSpin Migrate Client Users	Group for PlateSpin Migrate Web Interface Users	Description
PlateSpin Administrators	Workload Conversion Administrators	Have unlimited access to all features and functions of the application. A local administrator is implicitly part of this group.
PlateSpin Power Users	Workload Conversion Power Users	Have access to most features and functions of the application with some limitations, such as restrictions in the capability to modify system settings related to licensing and security.
PlateSpin Operators	Workload Conversion Operators	Have access to a limited subset of system features and functions, sufficient to maintain day-to-day operation.

When a user attempts to connect to a PlateSpin Server, the credentials provided through the PlateSpin Migrate Client are validated by IIS. If the user is not a member of one of the PlateSpin Migrate roles, connection is refused. If the user is a local administrator on the PlateSpin Server host, that account is implicitly regarded as a PlateSpin Migrate Administrator.

The Permission details for the PlateSpin Migrate roles depends on whether you use the PlateSpin Migrate Client or the PlateSpin Migrate Web Interface for migrating the workloads:

- ◆ For information on PlateSpin Migrate Roles and permission details when you use PlateSpin Migrate Client to perform the workload migration, see [Table 2-3 on page 32](#).
- ◆ For information on PlateSpin Migrate Roles and permission details when you use PlateSpin Migrate Web Interface to perform the workload migration, see [Table 2-4 on page 34](#).

Table 2-3 *PlateSpin Migrate Roles and Permission Details For PlateSpin Migrate Client Users*

Role Details	Administrators	Power Users	Operators
Licensing: Add, delete licenses; transfer workload licenses	Yes	No	No
Machines: Discover, undiscover	Yes	Yes	No

Role Details	Administrators	Power Users	Operators
Machines: Delete virtual machines	Yes	Yes	No
Machines: View, refresh, export	Yes	Yes	Yes
Machines: Import	Yes	Yes	No
Machines: Export	Yes	Yes	Yes
PlateSpin Migrate Networks: Add, delete	Yes	No	No
Jobs: Create new job	Yes	Yes	No
Jobs: View, abort, change start time	Yes	Yes	Yes
Imaging: View, start synchronization in existing contracts	Yes	Yes	Yes
Imaging: Consolidate increments, apply increments to base, delete increments, install/delete image servers	Yes	Yes	No
Block-Based Transfer Components: Install, upgrade, remove	Yes	Yes	No
Device Drivers: View	Yes	Yes	Yes
Device Drivers: Upload, delete	Yes	Yes	No
PlateSpin Server access: View Web services, download client software	Yes	Yes	Yes
PlateSpin Server settings: Edit settings that control user activity logging and SMTP notifications	Yes	No	No
PlateSpin Server settings: Edit all server settings except those that control user activity logging and SMTP notifications	Yes	Yes	No
Run Diagnostics: Generate detailed diagnostic reports on jobs.	Yes	Yes	Yes
Post-conversion Actions: Add, update, delete	Yes	Yes	No

Table 2-4 PlateSpin Migrate Roles and Permission Details For PlateSpin Migrate Web Interface Users

Role Details	Administrators	Power Users	Operators
Add Workload	Yes	Yes	No
Remove Workload	Yes	Yes	No
Configure Migration	Yes	Yes	No
Prepare Migration	Yes	Yes	No
Run Full Replication	Yes	Yes	Yes
Run Incremental Replication	Yes	Yes	Yes
Pause/Resume Schedule	Yes	Yes	Yes
Test Cutover	Yes	Yes	Yes
Cutover	Yes	Yes	Yes
Abort	Yes	Yes	Yes
Settings (All)	Yes	No	No
Run Reports/Diagnostics	Yes	Yes	Yes

2.2.2 Assigning PlateSpin Migrate Roles to Windows Users

To allow specific Windows domain or local users to carry out specific PlateSpin Migrate operations according to designated role, add the required Windows domain or user account to the applicable Windows local group (PlateSpin Administrators, PlateSpin Power Users, or PlateSpin Operators) on the PlateSpin Server host. For more information, see your Windows documentation.

2.2.3 Setting Up PlateSpin Migrate Multitenancy on VMware

PlateSpin Migrate includes unique user roles (and a tool for creating them in a VMware datacenter) that make it possible non-administrative VMware users (or “enabled users”) to perform Migrate lifecycle operations in the VMware environment. These roles makes it possible for you, as a service provider, to segment your VMware cluster to allow multitenancy: where multiple Migrate containers are instantiated in your datacenter to accommodate Migrate customers or “tenants” who want to keep their data and evidence of their existence separate from and inaccessible to other customers who also use your datacenter.

This section includes the following information:

- ◆ [“Using Tools to Define VMware Roles” on page 35](#)
- ◆ [“Assigning Roles In vCenter” on page 36](#)

Using Tools to Define VMware Roles

PlateSpin Migrate requires certain privileges to access and perform tasks in the VMware Infrastructure (that is, VMware “containers”), making the Migrate workflow and functionality possible in that environment. Because there are many of these required privileges, NetIQ has created a file that defines the minimum required privileges and aggregates them respectively into three VMware custom roles:

- ◆ PlateSpin Virtual Machine Manager
- ◆ PlateSpin Infrastructure Manager
- ◆ PlateSpin User

This definition file, `PlateSpinRole.xml`, is included in the PlateSpin Migrate Server installation. An accompanying executable, `PlateSpin.VMwareRoleTool.exe`, accesses the file to enable the creation of these custom PlateSpin roles in a target vCenter environment.

This section includes the following information:

- ◆ [“Basic Command Line Syntax” on page 35](#)
- ◆ [“Additional Command Line Parameters and Flags” on page 35](#)
- ◆ [“Tool Usage Example” on page 36](#)
- ◆ [“\(Option\) Manually Defining the PlateSpin Roles in vCenter” on page 36](#)

Basic Command Line Syntax

From the location where the role tool was installed, run the tool from the command line, using this basic syntax:

```
PlateSpin.VMwareRoleTool.exe /host=[host name/IP] /user=[user name] /role=[the  
role definition file name and location] /create
```

NOTE: By default, the role definition file is located in the same folder with the role definition tool.

Additional Command Line Parameters and Flags

Apply the following parameters as needed when you use `PlateSpin.VMwareRoleTool.exe` to create or update roles in vCenter:

`/create` (mandatory) Creates the roles defined by the `/role` parameter

`/get_all_privileges` Display all server-defined privileges

Optional Flags

`/interactive` Run the tool with interactive options that allow you to choose to create individual roles, check role compatibility, or list all compatible roles.

`/password=[password]` Provide the VMware password (bypasses the password prompt)

`/verbose` Display detailed information

Tool Usage Example

Usage: `PlateSpin.VMwareRoleTool.exe /host=houston_sales /user=pedrom /role=PlateSpinRole.xml /create`

Resulting Actions:

1. The role definition tool runs on the `houston_sales` vCenter server, which has an administrator with the user name `pedrom`.
2. In the absence of the `/password` parameter, the tool prompts for the user password, which you enter.
3. The tool accesses the role definition file, `PlateSpinRole.xml`, which is located in the same directory as the tool executable (there was no need to further define its path).
4. The tool locates the definition file and is instructed (`/create`) to create the roles defined in the contents of that file in the vCenter environment.
5. The tool accesses the definition file and creates the new roles (including the appropriate minimum privileges for defined, limited access) inside vCenter.

The new custom roles are to be [assigned to users later in vCenter](#).

(Option) Manually Defining the PlateSpin Roles in vCenter

You use the vCenter client to manually create and assign the PlateSpin custom roles. This requires creating the roles with the enumerated privileges as defined in `PlateSpinRole.xml`. When you create manually, there is no restriction on the name of the role. The only restriction is that the role names you create as equivalents to those in the definition file have all of the appropriate minimum privileges from the definition file.

For more information about how to create custom roles in vCenter, see [Managing VMWare VirtualCenter Roles and Permissions](http://www.vmware.com/pdf/vi3_vc_roles.pdf) (http://www.vmware.com/pdf/vi3_vc_roles.pdf) in the VMware Technical Resource Center.

Assigning Roles In vCenter

As you set up a multitenancy environment, you need to provision a single Migrate server per customer or “tenant.” You assign this Migrate server an enabled user with special Migrate VMware roles. This enabled user creates the Migrate container. As service provider, you maintain this user’s credentials and do not disclose them to your tenant customer.

The following table lists the roles you need to define for the enabled user. It also includes more information about the purpose of the role:

vCenter Container for Role Assignment	Role Assignment Specifics	Propagate Instructions	More Information
Root of vCenter inventory tree.	Assign the enabled user the <i>PlateSpin Infrastructure Manager</i> (or equivalent) role.	For security reasons, define the permission as non-propagating.	This role is needed to monitor tasks being performed by the Migrate software and to end any stale VMware sessions.

vCenter Container for Role Assignment	Role Assignment Specifics	Propagate Instructions	More Information
All datacenter objects where the enabled user needs access	Assign the enabled user the <i>PlateSpin Infrastructure Manager</i> (or equivalent) role.	For security reasons, define the permission as non-propagating.	This role is needed to allow access to the datacenter's datastores for file upload/download. Define the permission as non-propagating.
Each cluster to be added to Migrate as a container, and each host contained in the cluster	Assign the enabled user the <i>PlateSpin Infrastructure Manager</i> (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	To assign to a host, propagate the permission from the cluster object or create an additional permission on each cluster host. If the role is assigned on the cluster object and is propagated, no further changes are necessary when you add a new host to the cluster. However, propagating this permission has security implications.
Each Resource Pool where the enabled user needs access.	Assign the enabled user the <i>PlateSpin Virtual Machine Manager</i> (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	Although you can assign access to any number of Resource Pools in any location in the tree, you must assign the enabled user this role on at least one Resource Pool.
Each VM folder where the enabled user needs access	Assign the enabled user the <i>PlateSpin Virtual Machine Manager</i> (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	Although you can assign access to any number of VM Folders in any location in the tree, you must assign the enabled user this role on at least one folder.
Each Network where the enabled user needs access. Distributed Virtual Networks with a dvSwitch and a dvPortgroup	Assign the enabled user the <i>PlateSpin Virtual Machine Manager</i> (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	Although you can assign access to any number of networks in any location in the tree, you must assign the enabled user this role on at least one folder. <ul style="list-style-type: none"> ◆ To assign the correct role to the dvSwitch, propagate the role on the Datacenter (resulting in an additional object receiving the role) or place the dvSwitch in a folder and assign the role on that folder. ◆ For a standard portgroup to be listed as an available network in the Migrate UI, create a definition for it on every host in the cluster.

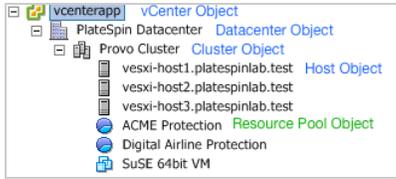
vCenter Container for Role Assignment	Role Assignment Specifics	Propagate Instructions	More Information
Each Datastore and Datastore Cluster where the enabled user needs access	Assign the enabled user the <i>PlateSpin Virtual Machine Manager</i> (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	The enabled user must have been assigned this role on at least one Datastore or Datastore Cluster. For Datastore Clusters, the permission must be propagated to the contained datastores. Not providing access to an individual member of the cluster causes both prepare and full replications to fail

The following table shows the role you can assign to the customer or tenant user.

vCenter Container for Role Assignment	Role Assignment Specifics	Propagate Instructions	More Information
Each resource pool(s) and folder(s) where the customer's VMs will be created.	Assign the tenant user the <i>PlateSpin User</i> (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	This tenant is a member of the PlateSpin Administrators group on the PlateSpin Migrate server and is also on the vCenter server. If the tenant will be granted the ability to change the resources used by the VM (that is, networks, ISO images, and so forth), grant this user the necessary permissions on those resources. For example, if you want to allow the customer to change the network where their VM is attached, this user should be assigned the Read-only role (or better) on all of the networks being made accessible to the customer.

The figure below illustrates a Virtual Infrastructure in the vCenter console. The objects labeled in blue are assigned the Infrastructure Manager role. The objects labeled in green are assigned the Virtual Machine Manager role. The tree does not show VM folders, Networks and Datastores. Those objects are assigned the *PlateSpin Virtual Machine Manager* role.

Figure 2-3 Roles assigned in vCenter



Security Implications of Assigning VMware Roles

PlateSpin software uses an enabled user only to perform protection lifecycle operations. From your perspective as a service provider, an end user never has access to the enabled user's credentials and is unable to access the same set of VMware resources. In an environment where multiple Migrate servers are configured to use the same vCenter environment, Migrate prevents possibilities for cross-client access. The major security implications include:

- With the *PlateSpin Infrastructure Manager* role assigned to the vCenter object, every enabled user can see (but not affect) the tasks performed by every other user.
- Because there is no way to set permissions on datastore folders/subfolders, all enabled users with permissions on a datastore have access to all other enabled users' disks stored on that datastore.
- With the *PlateSpin Infrastructure Manager* role assigned to the cluster object, every enabled user is able to turn off/on HA or DRS on the entire cluster
- With the *PlateSpin User* role assigned at the storage cluster object, every enabled user is able to turn off/on SDRS for the entire cluster
- Setting the *PlateSpin Infrastructure Manager Role* on the DRS Cluster object and propagating this role allows the enabled user to see all VMs placed in the default resource pool and/or default VM folder. Also, propagation requires the administrator to explicitly set the enabled user to have a "no-access" role on every resource pool/VM folder that he or she should not have access to.
- Setting the *PlateSpin Infrastructure Manager Role* on the vCenter object allows the enabled user to end sessions of any other user connected to the vCenter.

NOTE: Remember, in these scenarios, different enabled users are actually different instances of the PlateSpin software.

2.2.4 Setting Up User Activity Logging

By default, PlateSpin Migrate records all user activities in a log file, `PlateSpin.UserActivityLogging.log`, located on your PlateSpin Server host, in the following directory:

```
..\PlateSpin Migrate Server\logs.
```

The format of an individual log entry is:

```
date|Category|description|user|details1|details2
```

The `Category` element describes the functional area applicable to a particular action, such as `Security`, `Inventory` (discovery operations), `LicenseManagement`, or `Migration` (workload portability operations).

Elements `details1` and `details2` depend on the `Category` and provide additional information if applicable.

Below is an example of a log entry recording the login action of a user with the domain account MyDomain\John.Smith.

```
2008-09-02 14:14:47|Security|User logged in|MyDomain\John.Smith
```

When the size of a log file reaches a specified value, it is rolled over to a new file with a sequential number appended to the name:

```
PlateSpin.UserActivityLogging.log.1  
PlateSpin.UserActivityLogging.log.2  
PlateSpin.UserActivityLogging.log.3
```

When the number of log files reaches a specified value, the system starts overwriting the oldest file each time a rollover is performed.

To enable or disable user activity logging, and to specify log file size and rollover options:

- 1 In the PlateSpin Migrate Client, click **Tools > Options**.
- 2 Click the **Logging** tab.
- 3 Specify the required options, then click **OK**.

2.3 Access and Communication Requirements across your Migration Network

This section provides information about setting up user authorization and authentication, configuring your network environment, and managing your product's default settings and behavior.

- ♦ [Section 2.3.1, "Requirements for Discovery," on page 40](#)
- ♦ [Section 2.3.2, "Requirements for Migration," on page 41](#)
- ♦ [Section 2.3.3, "Migrations Across Public and Private Networks through NAT," on page 42](#)

2.3.1 Requirements for Discovery

The following table lists software, network, and firewall requirements that systems in your environment must meet for the discovery and inventory process. For information about the actual discovery procedures, see [Section 3.3, "Discovering Source Workloads and Migration Targets," on page 64](#).

Table 2-5 Network Communication Prerequisites for Discovery Operations

System	Prerequisites
All workloads	Ping (ICMP echo request and response) support
All Windows sources and Hyper-V hosts	<ul style="list-style-type: none">♦ Microsoft .NET Framework version 2.0 SP2 or 3.5 SP1♦ Credentials with local or domain admin privileges

System	Prerequisites
<ul style="list-style-type: none"> ◆ Windows 7 ◆ Windows 8, 8.1 ◆ Windows Server 2008 ◆ Windows Server 2008 R2 ◆ Windows Server 2012 ◆ Windows Server 2012 R2 ◆ Windows Server 2012 with Hyper-V ◆ Windows Server 2012 R2 with Hyper-V 	<ol style="list-style-type: none"> 1. Built-in Administrator or a domain account credentials (mere membership in the local Administrators group is insufficient). 2. The Windows Firewall configured to allow File and Printer Sharing. Use one of these options: <ul style="list-style-type: none"> ◆ Option 1, using Windows Firewall: Use the basic Windows Firewall Control Panel item (<code>firewall.cpl</code>) and select File and printer Sharing in the list of exceptions. - OR - ◆ Option 2, using Firewall with Advanced Security: Use the Windows Firewall with Advanced Security utility (<code>wf.msc</code>) with the following Inbound Rules enabled and set to <code>Allow</code>: <ul style="list-style-type: none"> ◆ File and Printer Sharing (Echo Request - ICMPv4In) ◆ File and Printer Sharing (Echo Request - ICMPv6In) ◆ File and Printer Sharing (NB-Datagram-In) ◆ File and Printer Sharing (NB-Name-In) ◆ File and Printer Sharing (NB-Session-In) ◆ File and Printer Sharing (SMB-In) ◆ File and Printer Sharing (Spooler Service - RPC) ◆ File and Printer Sharing (Spooler Service - RPC-EPMAP) 3. (Conditional) If the volumes are encrypted with the BitLocker disk encryption feature, they must be unlocked.
All Linux sources	<ul style="list-style-type: none"> ◆ Secure Shell (SSH) server
Citrix Xen Server	<ul style="list-style-type: none"> ◆ Open port 22 (TCP) ◆ Root-level access. For information on using an account other than root, see KB Article 7920711 (https://www.netiq.com/support/kb/doc.php?id=7920711). ◆ Custom SSH ports are supported; specify the port number during discovery: <code><hostname IP_address>:port_number</code>.
VMware ESX/ESXi Servers	<ul style="list-style-type: none"> ◆ VMware account with an Administrator role ◆ VMware Web services API and file management API (HTTPS / port 443 TCP)
VMware vCenter Servers	The user with access must be assigned the appropriate roles and permissions. Refer to the pertinent release of VMware documentation for more information.

2.3.2 Requirements for Migration

The following table lists firewall requirements that systems in your environment must meet for problem-free operation during workload migration jobs.

Table 2-6 Network Communication Prerequisites for Workload Portability

System	Open Port (Default)	Remarks
PlateSpin Server hosts	Either TCP 80 or TCP 443	<ul style="list-style-type: none"> Port 80 (TCP) is required for HTTP communication among the PlateSpin Server, sources, and targets. Port 443 (TCP) is required for HTTPS communication (if SSL is used) between the PlateSpin Server and the source or target machines.
All source workloads except those in image deployment jobs.	TCP 3725	Required for targets to initiate communication during file-level data transfer, except for I2X jobs, during which this port needs to be open on the migration target only. For Server Sync jobs, this port is required for both sources and targets.
All targets	TCP 3725	Required for: File-based Server Sync / Image Sync <ul style="list-style-type: none"> File-level Server Sync Image synchronization jobs
All Windows sources and targets	NetBIOS 137 - 139	Required for NetBIOS communications.
All sources	SMB (TCP 139, 445 and UDP 137, 138)	Required for communication and file-level data transfer during offline migration.
All Linux sources Citrix Xen Server	TCP 22	Required for communication during offline migration.
PlateSpin Server hosts; All Windows sources	TCP 135/445	For DCOM/RPC communication between PlateSpin Server and a source for taking control of and rebooting the workload through WMI. NOTE: WMI (RPC/DCOM) can use TCP ports 135 and 445 as well as random/dynamically assigned ports above 1024.

2.3.3 Migrations Across Public and Private Networks through NAT

In some cases, a source, a target, or PlateSpin Migrate itself, might be located in an internal (private) network behind a network address translator (NAT) device, unable to communicate with its counterpart during migration.

PlateSpin Migrate enables you to address this issue, depending on which of the following hosts is located behind the NAT device:

- PlateSpin Server:** In your server's *PlateSpin Server Configuration* tool, record the additional IP addresses assigned to that host:

- From any Web browser, open the `https://Your_PlateSpin_Server/platespinconfiguration/tool`.

2. Locate the **AlternateServerAddresses** server parameter, click **Edit**, then add additional IP addresses, delimited by a a semicolon (;), for example:

```
204.50.186.147;204.50.186.148
```

- ◆ **Source:** As part of that specific migration job, record the additional IP addresses assigned to that workload. See [“Specifying Network Options” on page 147](#).
- ◆ **Target:** When you are attempting to discover a target, such as VMware ESX, specify the public (or external) IP address in the discovery parameters.

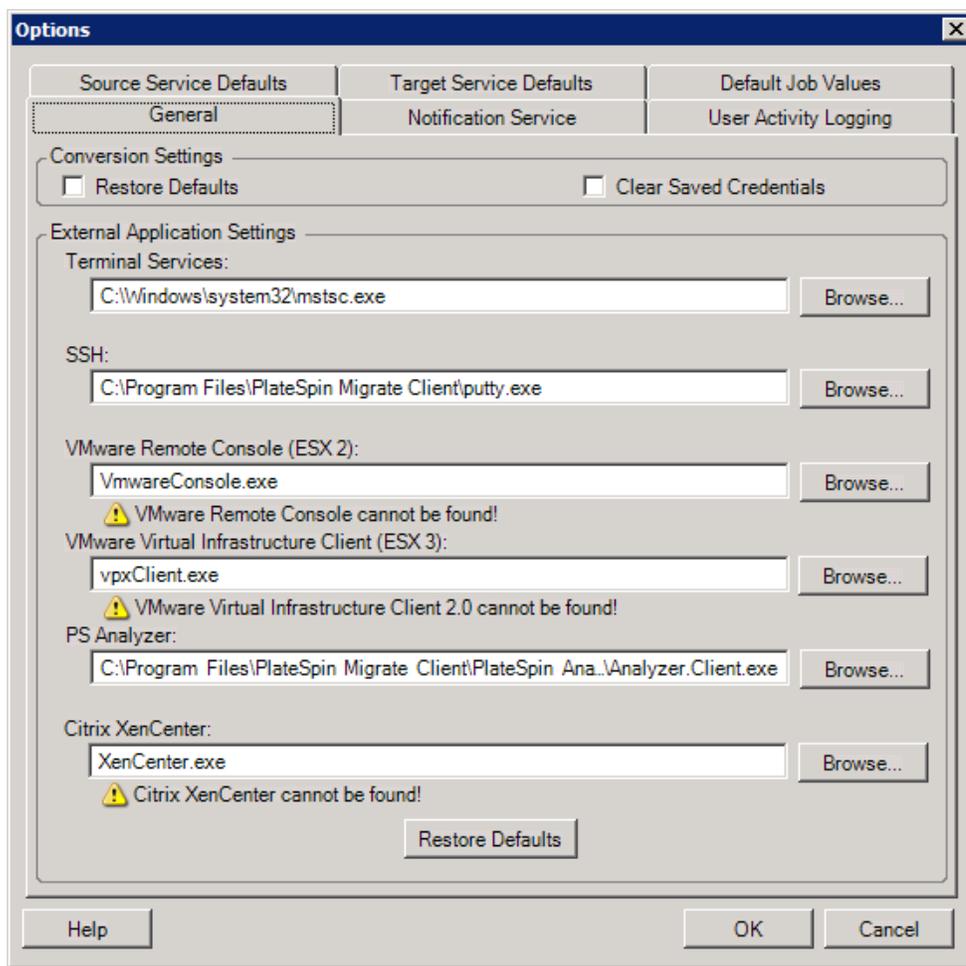
2.4 Configuring PlateSpin Migrate Default Options By Using the PlateSpin Migrate Client

Default options control PlateSpin Migrate’s global settings and its default behavior. To configure the default options following installation or to reflect changes in your environment, click **Tools > Options**. Review the following sections for more information:

- ◆ [Section 2.4.1, “General Options,” on page 44](#)
- ◆ [Section 2.4.2, “Notification Service,” on page 45](#)
- ◆ [Section 2.4.3, “User Activity Logging,” on page 46](#)
- ◆ [Section 2.4.4, “Default Job Values,” on page 47](#)
- ◆ [Section 2.4.5, “Source Service Defaults,” on page 48](#)
- ◆ [Section 2.4.6, “Target Service Defaults,” on page 49](#)

2.4.1 General Options

Use this tab to restore default settings, clear saved credentials, and to specify the locations of executable files for external applications you can launch from within the PlateSpin Migrate Client user interface. To access this dialog box in the PlateSpin Migrate Client, click **Tools > Options**.



Restore Defaults: When this option is selected, PlateSpin Migrate resets the job configuration method (launches the Actions dialog box after a drag-and-drop) and resumes checking for software updates on the Client startup.

Clear Saved Credentials: Removes stored usernames and passwords for source and target machines.

External Application Settings: Use the adjacent **Browse** buttons to locate application executables.

Restore Defaults: Resets the paths to their defaults.

2.4.2 Notification Service

Use this tab to specify Simple Mail Transfer Protocol (SMTP) server settings for event and job progress notifications. To access this dialog box in the PlateSpin Migrate Client, click **Tools > Options**.

The screenshot shows the 'Options' dialog box with the 'Notification Service' tab selected. The dialog is divided into three main sections: 'Source Service Defaults', 'Target Service Defaults', and 'Default Job Values'. The 'Target Service Defaults' section is further divided into 'General', 'Notification Service', and 'User Activity Logging'. The 'Notification Service' section contains two main groups of settings:

- SMTP Server Settings:**
 - Address: 192.168.1.128
 - Port: 25
 - Reply Address: support@platespin.com
- SMTP Account Credentials:**
 - User Name: John Smith
 - Password: [Redacted]
 - Confirm Password: [Redacted]

At the bottom of the dialog, there are three buttons: 'Help', 'OK', and 'Cancel'.

SMTP Server Settings: Specify your SMTP server's IP address, port, and a reply address for e-mail event and progress notifications.

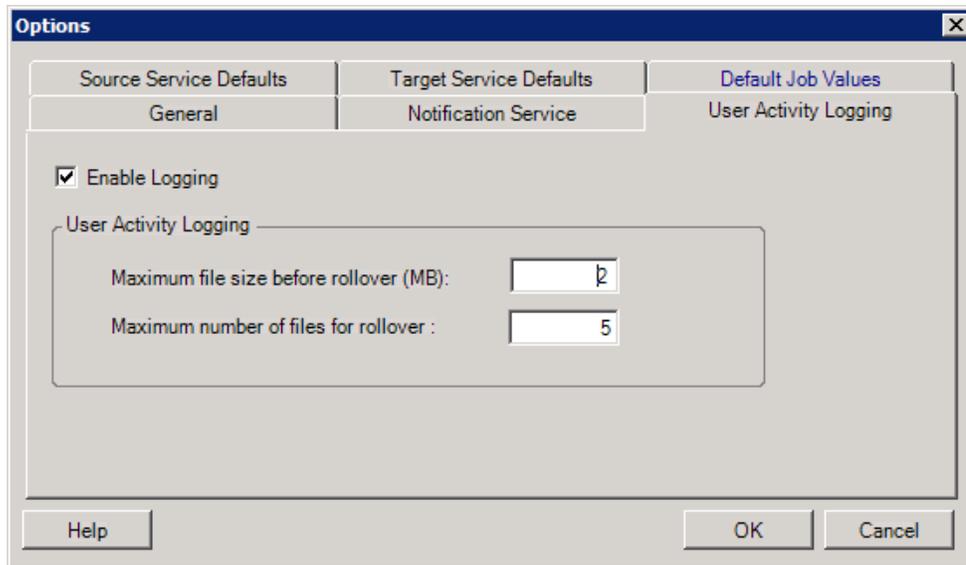
SMTP Account Credentials: Provide valid credentials if your SMTP server requires authentication.

You can also configure migration progress notifications on a per-migration basis. See [Section 7.2, "Configuring Automatic Email Alerts of Job Status and Progress,"](#) on page 146.

2.4.3 User Activity Logging

Use this tab to specify options related to user activity logging. See [“Setting Up User Activity Logging” on page 39](#).

To access this dialog box in the PlateSpin Migrate Client, click **Tools > Options**.



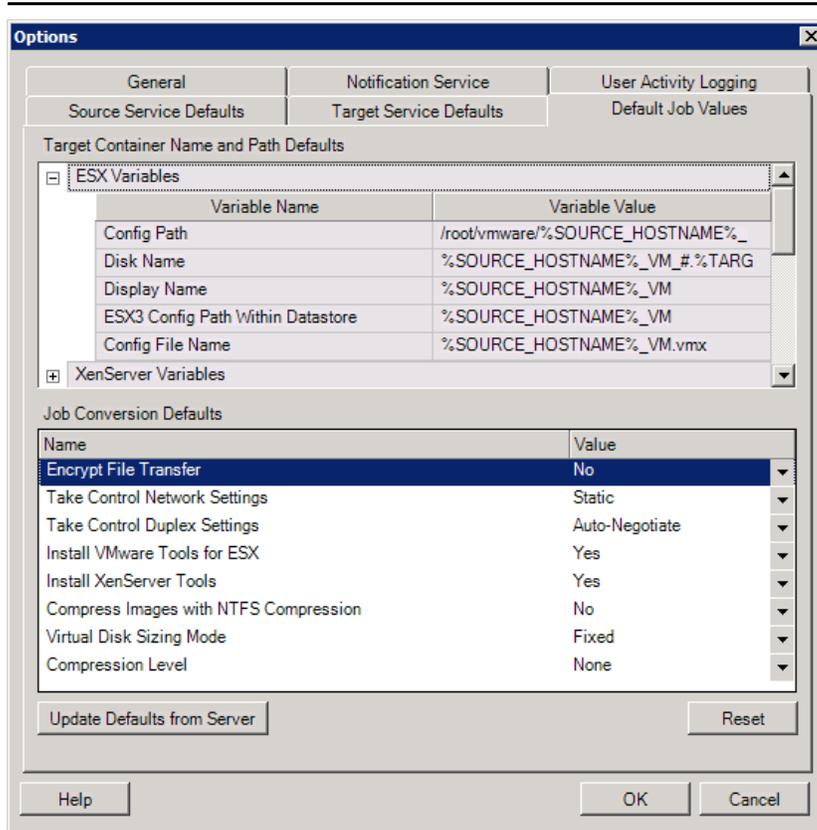
Enable Logging: When this option is selected, PlateSpin Migrate logs all user activities.

Maximum file size before rollover (MB): When the size of a log file reaches the specified value, it is rolled over to a new file with a sequential number appended to the name.

Maximum number of files for rollover: When the number of log files reaches the specified value, the system starts overwriting the oldest file each time a rollover is performed.

2.4.4 Default Job Values

Use this tab to specify default migration job values specific to the target virtualization platform. To access this dialog box in the PlateSpin Migrate Client, click **Tools > Options**.



Target Container Name and Path Details: These variables control naming conventions for target paths and containers. Expand a variable set and click a variable to edit its value. You can edit the following variables:

- ◆ **SOURCE_HOSTNAME:** Host name of the source computer.
- ◆ **IMAGESERVER_LOCATION:** Relevant with image server settings; specify this variable in the **Image Server Variables** section.
- ◆ **SOURCE_VOLUME_SERIAL_NUMBER:** Volume serial number of the source computer.

Job Conversion Defaults: Use this area to set defaults that affect all migration jobs. You can overwrite these defaults when configuring actual migration jobs.

- ◆ **Encrypt File Transfer:** See [Security and Privacy](#).
- ◆ **Install VMware Tools for ESX/XenServer Tools:** See [Handling Virtualization Enhancement Software](#).
- ◆ **Compress Images with NTFS Compression:** See [Capturing a Workload to a PlateSpin Image](#). Unrelated to data compression for over-the-network transfer.
- ◆ **Virtual Disk Sizing Mode:** (ESX only) **Fixed**—space is pre-allocated for the virtual disk; **Dynamic**—the virtual disk is assigned a minimum amount of space, which grows when needed.
- ◆ **Compression Level:** See [Data Compression](#).

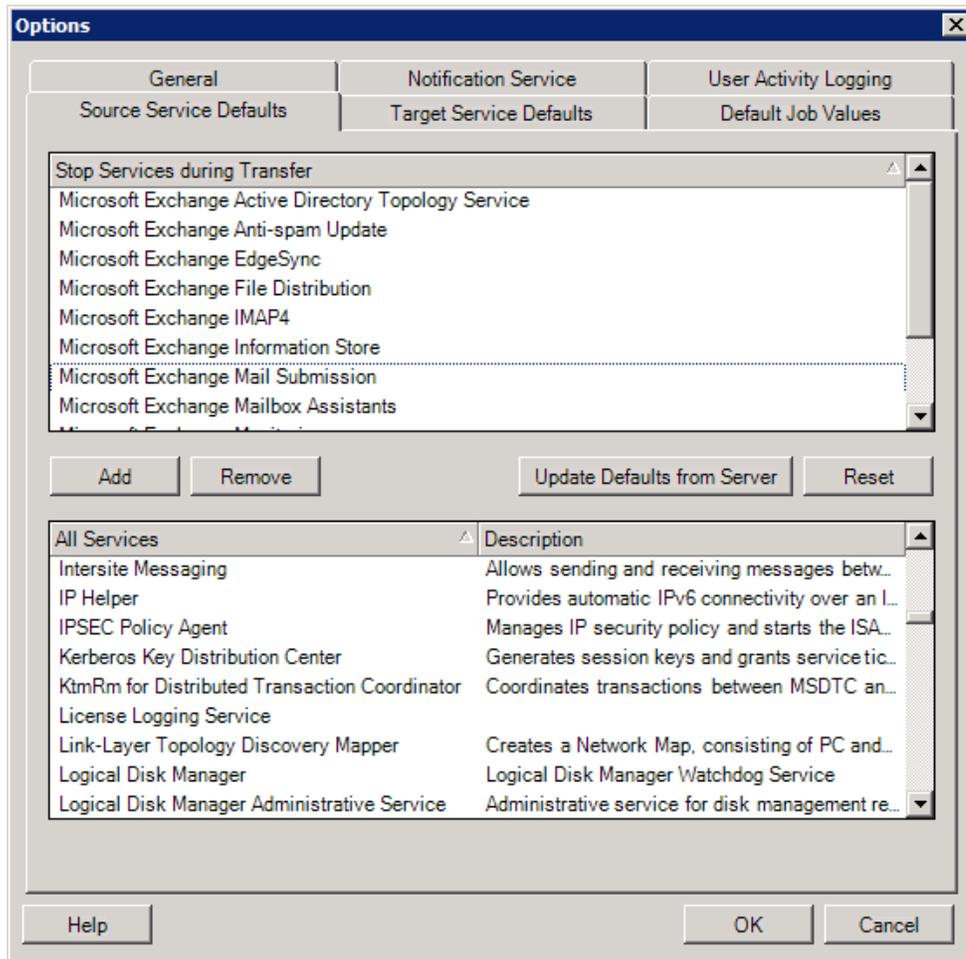
Reset: Restores default job values.

Update Defaults from Server: Click to retrieve defaults from the PlateSpin Server if available.

2.4.5 Source Service Defaults

Use this tab to select Windows services daemons to stop on the source workload during a Live Transfer migration. See [Handling Source Workload Services or Daemons During Live Transfer \(Windows and Linux\)](#).

To access this dialog box in the PlateSpin Migrate Client, click **Tools > Options**.



Stop Services during Transfer section: Lists services that are stopped by default. To stop a service during data transfer that uses a specific transfer method by default, select the corresponding check box. A deselected check box means the service remains active during Live Transfer.

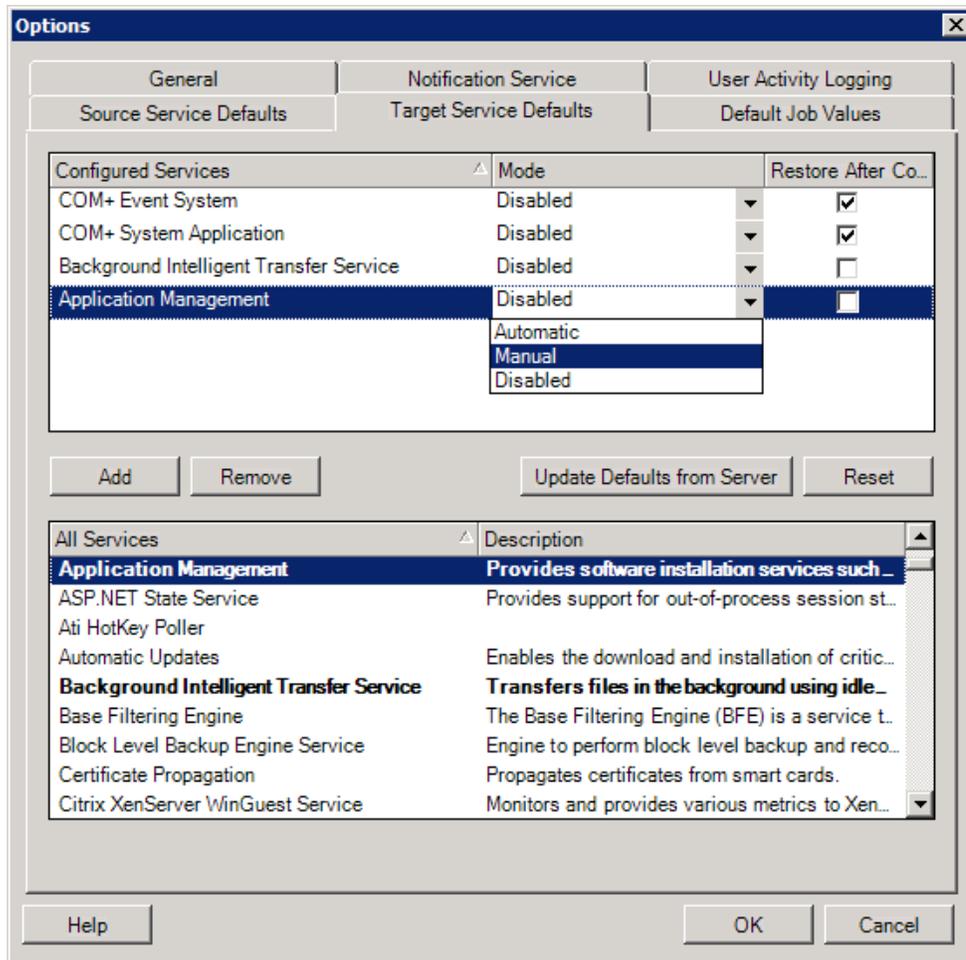
All Services section: Lists unique services on all discovered machines. Click **Add** to add a selected service from the lower section to the upper section and set it to stop during the migration.

Update Defaults from Server: Retrieves defaults from PlateSpin Server.

2.4.6 Target Service Defaults

Use this tab to select Windows services whose mode on the target is to be different from that of the source. See [Handling the Startup Mode of Services \(Windows Targets\)](#).

To access this dialog box in the PlateSpin Migrate Client, click **Tools > Options**.



Configure Services section: Lists services and their target startup modes. Select the **Restore After Conversion** check box to use the selected mode during the migration. The service is then restored to match the source after the migration is complete and the target machine is ready to run.

All Services section: Lists unique services on all discovered machines. Click **Add** to add a service to the upper section. Use the **Mode** drop-down list to select the service state for the target. This is set during the configuration step of the job.

Remove: Removes a service.

Reset: Clears the upper section. The modes of all services in the target will match those on the source.

2.5 Configuring PlateSpin Server Behavior through Configuration Parameters

Some aspects of your PlateSpin Server's behavior are controlled by configuration parameters that you set on a configuration Web page residing your PlateSpin Server host (https://Your_PlateSpin_Server/platespinconfiguration/).

Under normal circumstances you should not need to modify these settings unless you are advised to do so by PlateSpin Support. This section provides a number of common use cases along with information on the required procedure.

Use the following procedure for changing and applying any configuration parameters:

- 1 From any Web browser, open the https://Your_PlateSpin_Server/platespinconfiguration/tool.
- 2 Locate the required server parameter and change its value.
- 3 Save your settings and exit the page.

No reboot or restart of services is required after the change is made in the configuration tool.

The following topics provide information on specific situations, in which you might need to change product behavior using an XML configuration value.

- ♦ [Section 2.5.1, "Changing the Adapter Type Used during the Target Take Control Process of Workload Migration to a Target VM on a Hyper-V Host," on page 50](#)
- ♦ [Section 2.5.2, "Increasing the Size Limit on Post-Migration Actions Uploaded to the PlateSpin Server," on page 51](#)
- ♦ [Section 2.5.3, "Optimizing Data Transfer over WAN Connections," on page 51](#)
- ♦ [Section 2.5.4, "Other Use Cases for Making Modifications in the Web Config Tool \(Advanced\)," on page 55](#)

2.5.1 Changing the Adapter Type Used during the Target Take Control Process of Workload Migration to a Target VM on a Hyper-V Host

By default, PlateSpin Migrate selects the adapter type based on the Workload OS and Target Virtual Machine type. To edit the adapter type used during the Target Take Control process of workload migration to a Hyper-V target, do the following:

- 1 From any Web browser, open https://Your_PlateSpin_Server/platespinconfiguration/.
- 2 Locate the `PreferredHyperVTakeControlNetworkAdapter` parameter and edit its value as **synthetic** or **legacy** depending on your requirement.
- 3 Save your settings and exit the page.

2.5.2 Increasing the Size Limit on Post-Migration Actions Uploaded to the PlateSpin Server

By default, PlateSpin Migrate sets a 64 MB upload limit for each individual post-migration action, including its dependencies. See [Managing Custom Actions](#).

You can increase this limit by modifying the configuration setting on the PlateSpin Server Web configuration tool.

IMPORTANT: Decreasing the default size limit might have a negative impact on the stability of your PlateSpin Server.

- 1 From any Web browser, open `https://Your_PlateSpin_Server/platespinconfiguration/`.
- 2 Locate and edit the setting that defines the value of the `httpRuntime` element's `maxLength` attribute:
`maxLength=8192`
- 3 Replace the existing value with the required new value in kilobytes. For example, to increase the size to 16 MB, replace 8192 with 16384:
`maxLength=16384`
- 4 Save your settings and exit the page.

2.5.3 Optimizing Data Transfer over WAN Connections

You can optimize data transfer performance and fine tune it for WAN connections. You do this by modifying configuration parameters that the system reads from settings you make in a configuration tool residing on your PlateSpin Server host. For the generic procedure, see “[Configuring PlateSpin Server Behavior through Configuration Parameters](#)” on page 50.

- ♦ “[Tuning Parameters](#)” on page 51
- ♦ “[Tuning FileTransferSendReceiveBufferSize](#)” on page 53

Tuning Parameters

Use the file transfer configuration parameters settings to optimize data transfers across a WAN. These settings are global and affect all replications using the file-based and VSS replications.

NOTE: If these values are modified, replication times on high-speed networks, such as Gigabit Ethernet, might be negatively impacted. Before modifying any of these parameters, consider consulting PlateSpin Support first.

[Table 2-7](#) lists the configuration parameters on the PlateSpin Configuration page (`https://Your_PlateSpin_Server/platespinconfiguration/`) that control file transfer speeds with the defaults and maximum values. You can modify these values through trial-and-error testing in order to optimize operation in a high-latency WAN environment.

Table 2-7 Default and Optimized File Transfer Configuration Parameters

Parameter	Default Value	Maximum Value
AlwaysUseNonVSSFileTransferForWindows2003	False	
FileTransferCompressionThreadsCount	2	N/A
Controls the number of threads used for packet-level data compression. This setting is ignored if compression is disabled. Because the compression is CPU-bound, this setting might have a performance impact.		
FileTransferBufferThresholdPercentage	10	
Determines the minimum amount of data that must be buffered before creating and sending new network packets.		
FileTransferKeepAliveTimeOutMilliSec	120000	
Specifies how long to wait to start sending keep alive messages if TCP times out.		
FileTransferLongerThan24HoursSupport	True	
FileTransferLowMemoryThresholdInBytes	536870912	
Determines when the server considers itself to be in a low memory state, which causes augmentation of some networking behavior.		
FileTransferMaxBufferSizeForLowMemoryInBytes	5242880	
Specifies the internal buffer size used in a low memory state.		
FileTransferMaxBufferSizeInBytes	31457280	
Specifies internal buffer size for holding packet data.		
FileTransferMaxPacketSizeInBytes	1048576	
Determines the largest packets that will be sent.		
FileTransferMinCompressionLimit	0 (disabled)	max 65536 (64 KB)
Specifies the packet-level compression threshold in bytes.		
FileTransferPort	3725	

Parameter	Default Value	Maximum Value
FileTransferSendReceiveBufferSize	0 (8192 bytes)	max 5242880 (5 MB)
<p>Defines the maximum size (in bytes) of the send and receive buffers for TCP connections in the replication network. The buffer size affects the TCP Receive Window (RWIN) size, which sets the number of bytes that can be sent without TCP acknowledgement. This setting is relevant for both file-based and block-based transfers. Tuning the buffer size based on your network bandwidth and latency improves throughput and reduces CPU processing.</p> <p>When the value is set to zero (off), the default TCP window size is used (8 KB). For custom sizes, specify the size in bytes.</p> <p>Use the following formula to determine the proper value:</p> $((\text{LINK_SPEED in Mbps} / 8) * \text{DELAY in sec}) * 1000 * 1024$ <p>For example, for a 100 Mbps link with 10 ms latency, the proper buffer size would be:</p> $(100/8) * 0.01 * 1000 * 1024 = 128000 \text{ bytes}$ <p>For tuning information, see “Tuning FileTransferSendReceiveBufferSize” on page 53.</p>		
FileTransferSendReceiveBufferSizeLinux	0 (253952 bytes)	
<p>Specifies the TCP/IP Receive Window (RWIN) Size setting for file transfer connections for Linux. It controls the number of bytes sent without TCP acknowledgement, in bytes.</p> <p>When the value is set to zero (off), the TCP/IP window size value for Linux is automatically calculated from the FileTransferSendReceiveBufferSize setting. If both parameters are set to zero (off), the default value is 248 KB. For custom sizes, specify the size in bytes.</p> <p>NOTE: In previous release versions, you were required to set this parameter to 1/2 the desired value, but this is no longer required.</p>		
FileTransferShutDownTimeOutInMinutes	1090	
FileTransferTCPTimeOutMilliSec	30000	
Sets both the TCP Send and TCP Receive Timeout values.		
PostFileTransferActionsRequiredTimeInMinutes	60	

Tuning FileTransferSendReceiveBufferSize

The FileTransferSendReceiveBufferSize parameter defines the maximum size (in bytes) of the send and receive buffers for TCP connections in the replication network. The buffer size affects the TCP Receive Window (RWIN) size, which sets the number of bytes that can be sent without TCP

acknowledgement. This setting is relevant for both file-based and block-based transfers. Tuning the buffer size based on your network bandwidth and latency improves throughput and reduces CPU processing.

You can tune the `FileTransferSendReceiveBufferSize` parameter to optimize transfer of blocks or files from the source servers to the target servers in your replication environment. Set the parameter on the PlateSpin Configuration page (https://Your_PlateSpin_Server/platespinconfiguration/).

To calculate the optimum buffer size:

- 1 Determine the latency (delay) between the source server and target server.

The goal is to discover what the latency is for a packet size that approaches the MTU as closely as possible.

1a Log in to the source server as an Administrator user.

1b Enter the following at a command prompt:

```
# ping <target-server-ip-address> -f -l <MTU_minus_28> -n 10
```

Typically, the `-l` option for `ping` adds 28 bytes in headers of the specified payload for the `target-server-ip-address`. Thus, a size in bytes of `MTU minus 28` is a good initial value to try.

1c Iteratively modify the payload and re-enter the command in [Step 1b](#) until you get the following message:

```
The packet needs to be fragmented.
```

1d Note the latency in seconds.

For example, if the latency is 35 ms (milliseconds), then note 0.035 as the latency.

- 2 Calculate a byte value for your initial buffer size:

```
Buffer Size = (Bandwidth in Mbps / 8) * Latency in seconds * 1000 * 1024
```

Use binary values for the network bandwidth. That is, 10 Gbps = 10240 Mbps and 1 Gbps = 1024 Mbps.

For example, the calculation for a 10 Gbps network with a latency of 35 ms is:

```
Buffer Size = (10240 / 8) * 0.035 * 1000 * 1024 = 45875200 bytes
```

- 3 (Optional) Calculate an optimal buffer size by rounding up to a multiple of the Maximum Segment Size (MSS).

3a Determine the MSS:

```
MSS = MTU Size in bytes - (IP Header Size + TCP Header Size)
```

The IP header size is 20 bytes. The TCP header size is 20 bytes plus the bytes for options like timestamp.

For example, if your MTU size is 1470, then your MSS is typically 1430.

```
MSS = 1470 bytes - (20 bytes + 20 bytes) = 1430 bytes
```

3b Calculate the optimal buffer size:

```
Optimal Buffer Size = (roundup( Buffer Size / MSS )) * MSS
```

To continue the example:

```
Optimal Buffer Size = (roundup(45875200 / 1430)) * 1430  
= 32081 * 1430  
= 45875830
```

You round up instead of down, because rounding down gives a multiple of the MSS that is smaller than the Buffer Size of 45875200:

Non-optimal Buffer Size = 32080 * 1430 = 45874400

2.5.4 Other Use Cases for Making Modifications in the Web Config Tool (Advanced)

Below is a list of configuration values that might address various environmental or functional issues. Do not use them unless you are advised to do so by PlateSpin Support.

Table 2-8 List of Common Use Cases for changing settings in the Web configuration tool

Issue or Use Case	Value Shown in the Config Tool	Remarks
Discovery/Inventory issues	<pre><add key="UseServiceForCommonInventory" value="true" /> <add key="UseServiceForMigrateInventory" value="false" /> <add key="EarliestWindowsVersionForCommonInventory" value="5.2" /></pre>	
Target boot issues related to drivers	<pre><add key="TargetMachineRegistryCleanupLevel " value="None" /></pre>	
Controller installation issues on sources (mainly due to environmental restrains)	<pre><add key="InstallControllerUsingService " value="true" /> <add key="RunControllerUnderSameCredsAsRemoteAccess " value="false" /></pre>	
Issues related to database size growth	<pre><add key="PowerConvertDBSizeLimitInBytes " value="4294967296" /> <add key="PlateSpinDBCleanupThresholdPercent " value="80" /> <add key="OFXDBCleanupThresholdInBytes " value="4294967296" /></pre>	

3 Getting Started with PlateSpin Migrate

This section introduces the features of the PlateSpin Migrate product and how you use the PlateSpin Migrate Client to carry out workload migration and management tasks.

The PlateSpin Migrate Client connects to a specified PlateSpin Server and provides access to information in the PlateSpin Server database.

Most of your interaction with the server takes place through the client, including the discovery of source workloads and targets; setting up, executing, and monitoring jobs; managing license keys; and configuring the default behavior of the server.

IMPORTANT

- ◆ When you migrate workloads to a ESX server that belongs to a DRS cluster, the VMware DRS and VMware HA are set as disabled only for the target VM. However, the other VMs in the cluster are not be affected.
- ◆ If the ESX host is a part of the VMware DRS or HA enabled cluster, you must discover the ESX host via vCenter or vCenter Cluster.

-
- ◆ [Section 3.1, “Connecting to a PlateSpin Server,” on page 57](#)
 - ◆ [Section 3.2, “About the PlateSpin Migrate Client User Interface,” on page 59](#)
 - ◆ [Section 3.3, “Discovering Source Workloads and Migration Targets,” on page 64](#)
 - ◆ [Section 3.4, “Undiscovering Source Workloads and Migration Targets,” on page 71](#)
 - ◆ [Section 3.5, “Analyzing Machines with PlateSpin Analyzer,” on page 71](#)
 - ◆ [Section 3.6, “Working with Device Drivers,” on page 72](#)
 - ◆ [Section 3.7, “Managing Custom Actions,” on page 80](#)
 - ◆ [Section 3.8, “Setting Up, Executing, and Managing Jobs,” on page 83](#)

3.1 Connecting to a PlateSpin Server

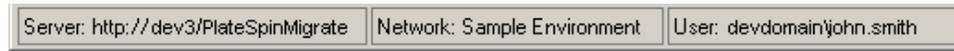
Every time you start the PlateSpin Migrate Client, it performs the following actions:

- ◆ Performs authentication of the specified user account with the PlateSpin Server.
See [Section 2.2, “Setting Up User Authorization and Authentication,” on page 31](#).
- ◆ Connects to a specified PlateSpin Server.
- ◆ Loads a specified PlateSpin Migrate Network, a collection of discovered source workloads and targets that you work with at one time.

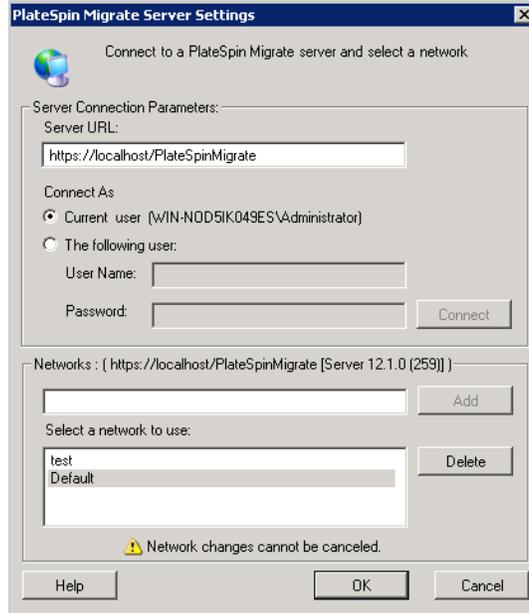
You specify your connection credentials, the PlateSpin Server instance, and the required PlateSpin Migrate Network in the PlateSpin Server settings.

- 1 In the PlateSpin Migrate Client, click **Tools > PlateSpin Server Settings**.
or

Double-click one of the following three areas in PlateSpin Migrate Client's status bar at the bottom: **Server**, **Network**, or **User**.



The PlateSpin Server Settings dialog box opens.



2 Specify the required PlateSpin Server URL, user, and network parameters as required:

Interface Element	Description
Server URL	Type the PlateSpin Server URL in the following format: <code>http://<i>server_host</i>/platespinmigrate</code> (If SSL is enabled on the PlateSpin Server host, replace <code>http</code> in the URL with <code>https</code>).
Connect As	To connect to a PlateSpin Server, you must have administrative access to the PlateSpin Server host or be a member of one of the PlateSpin Migrate roles. See Section 2.2, "Setting Up User Authorization and Authentication," on page 31 .
Networks	To familiarize yourself with PlateSpin Migrate features, use the Sample Environment network. To work with actual source workloads and targets, use the Default network or create your own. To add a network, type the name, then click Add . To remove a network, select it, then click Delete .

3 When you have finished, click **OK**.

3.2 About the PlateSpin Migrate Client User Interface

The PlateSpin Migrate Client window consists of the following elements:

- ♦ **Menu bar:** Reflects the current view and provides command groups for accessing program features and operations.
- ♦ **Toolbar:** Reflects the current view and provides visual shortcuts to program features and operations.
- ♦ **Current view:** The main work area of the interface; lists either machines (when in Servers view mode), or jobs (when in Jobs view mode).
- ♦ **Panes:** Vertically aligned at the left side of the window, panes facilitate the selection of the current view (View pane) or a migration job (Tasks pane). A Details pane reflects the current view and provides summary information about an item selected in the current view.
- ♦ **Status bar:** At the bottom of the PlateSpin Migrate Client window, the status bar displays the PlateSpin Server that the client is currently connected to, the PlateSpin Migrate Network you are currently working with, the name and role of the current user logged in, and the status of the Automatic Network Discovery feature.

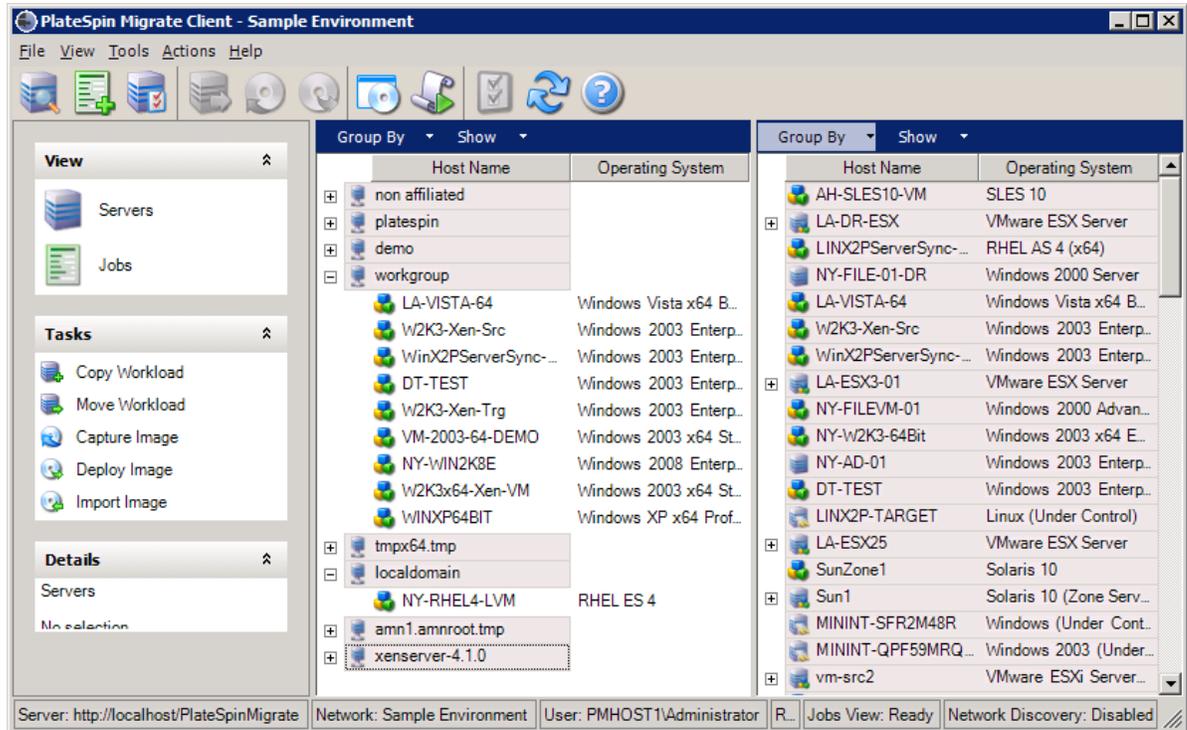
For additional information, see the following sections.

- ♦ [Section 3.2.1, “Servers View,” on page 60](#)
- ♦ [Section 3.2.2, “Jobs View,” on page 63](#)
- ♦ [Section 3.2.3, “Tasks Pane,” on page 64](#)
- ♦ [Section 3.2.4, “Status Bar,” on page 64](#)

3.2.1 Servers View

The Servers view is the main visual interface to your discovered source workloads and targets. Its main area consists of two split panes that you can customize to suit your needs.

Figure 3-1 PlateSpin Migrate Client's Servers View



The hierarchical display of items in the Servers view reflects the organization of items on their respective platforms; for example: VMs are shown nested beneath their VM hosts, and PlateSpin Images are beneath their image servers.

In addition, the **Group By** bar enables you to group machines by affiliation to a domain or to a vCenter Server (for VMware ESX server systems). See [“Organizing the Servers View” on page 60](#).

NOTE: The Servers view hierarchy does not reflect advanced VM resource management hierarchies and structures, such as membership in resource pools or affiliation with ESX Distributed Resource Scheduler (DRS) clusters. You can view such information in an item's properties. See [“Viewing the Properties of Source Workloads and Targets” on page 61](#).

- ◆ [“Organizing the Servers View” on page 60](#)
- ◆ [“Viewing the Properties of Source Workloads and Targets” on page 61](#)
- ◆ [“List of Machine-Specific Icons in the Servers View” on page 62](#)

Organizing the Servers View

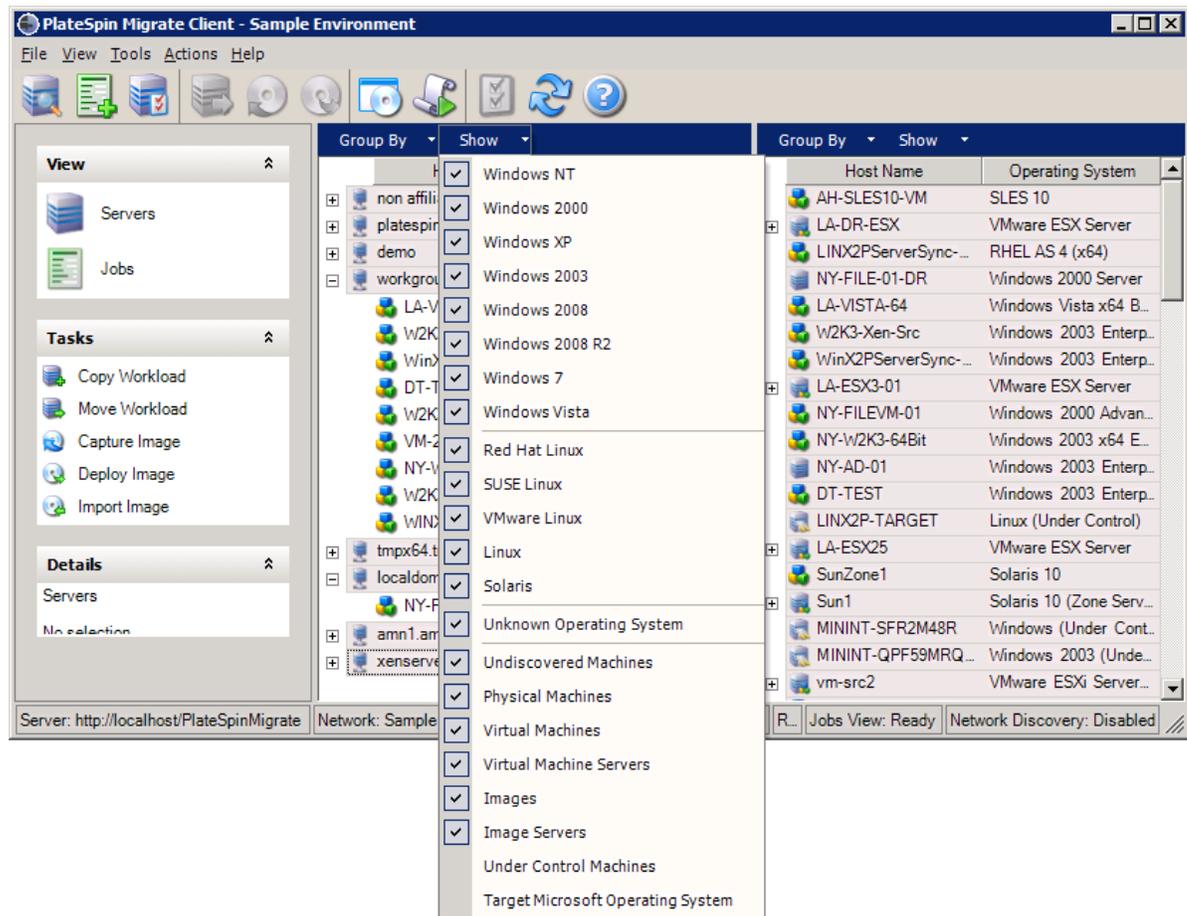
You can filter source workloads and targets based on operating system, domain, name, and type by using the **Group By** and **Show** drop-down menus. You can use the **Group By** drop-down menu to group the items in the Servers view by:

- ◆ Domain affiliation

- ◆ Hostname
- ◆ Affiliation to a VMware vCenter Server

To further control the scope of items shown in either pane of the view, you can also use the **Show** drop-down menu to filter machines by workload type; for example, Windows 2000, Red Hat Linux, and so on, as shown in the figure below:

Figure 3-2 Servers View Options for Sorting Items by Type



Viewing the Properties of Source Workloads and Targets

In the Servers view, you can access the essential properties of your discovered source workloads and targets by right-clicking an item and selecting **Properties**.

For each machine, the system provides information about the selected system's:

- ◆ Hardware, operating system, and network profile
- ◆ Volumes, partitions, and disk usage
- ◆ Programs and services

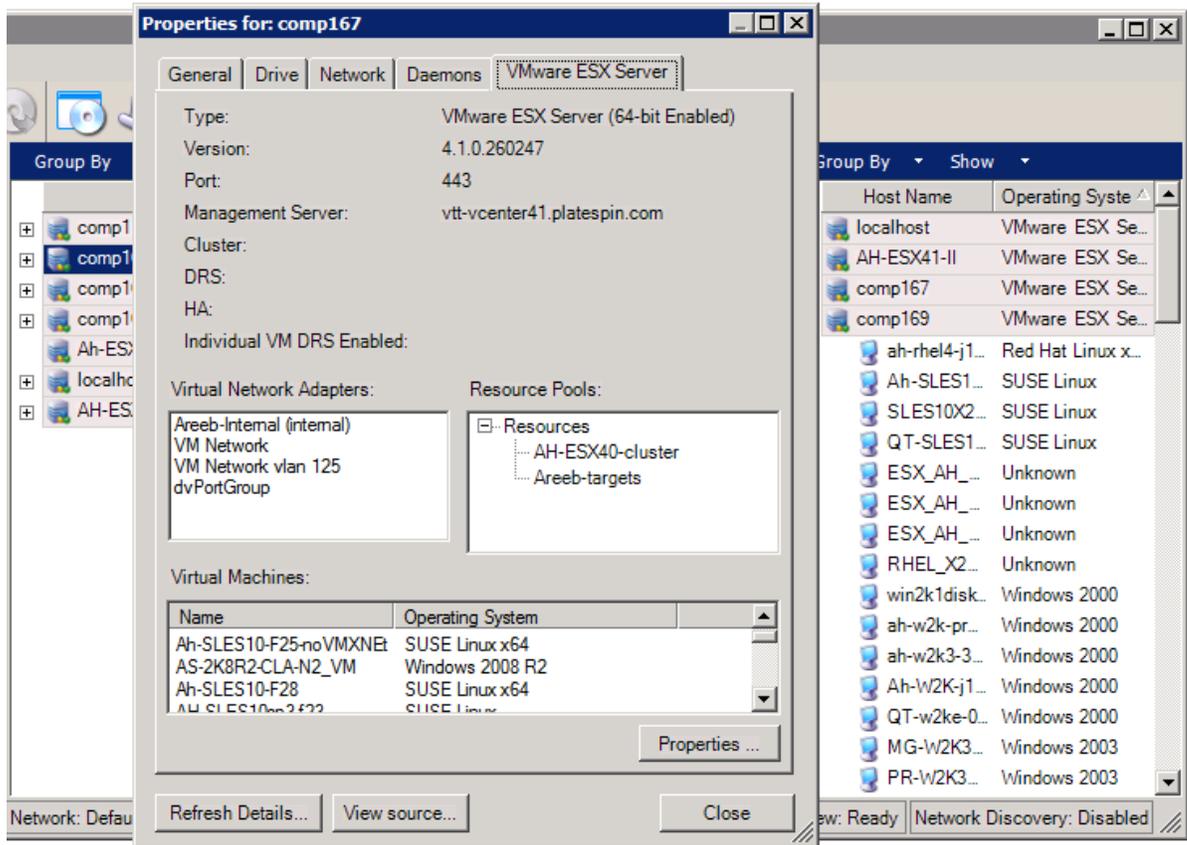
A virtual machine's properties provide information related to the machine's environment on its corresponding virtualization platform, including information about the host, and the amount of allocated memory and processing power.

The properties for virtual machine hosts provide information specific to the selected system. For example, you can view what virtual machines are running on a selected VMware ESX server, what virtual network adapters are in use, and what resource pools are configured on them.

VMware ESX servers that are assigned to a Distributed Resource Scheduler (DRS) cluster provide information about the name of the cluster and the DRS automation level (full, manual, or partially automated). The properties for VMware ESX servers that are part of VMware vCenter platforms also indicate this.

The following figure shows the properties of a discovered VMware ESX Server.

Figure 3-3 VMware ESX Server-Specific Information in the System's Properties



List of Machine-Specific Icons in the Servers View

Discovered source workloads and targets are associated with unique icons to help identify the type of workload or workload host.

Table 3-1 Machine-Specific Icons in the Servers View

	Physical machine
	Physical machine in pre-execution environment for offline migration
	Physical machine with workload license

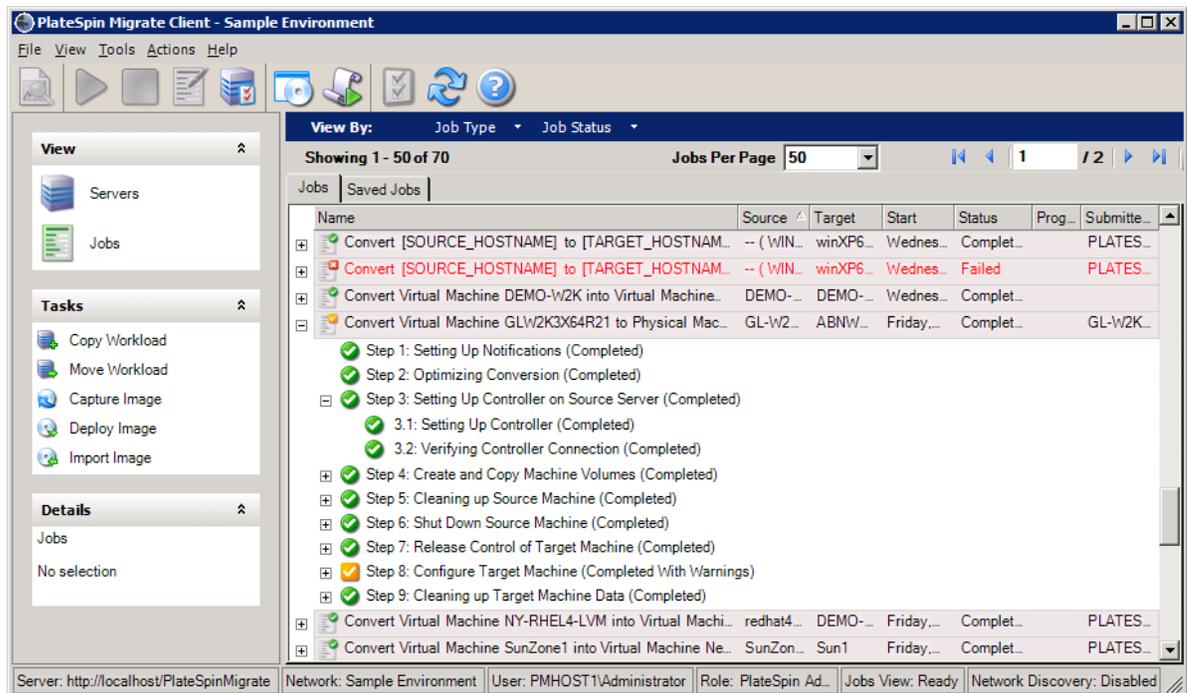
-  Virtual machine server
-  Virtual machine
-  Virtual machine with workload license
-  Undiscovered virtual machine
-  Virtual machine - Server Sync target
-  Virtual machine - Server Sync target with workload license
-  PlateSpin Image Server
-  PlateSpin Image

3.2.2 Jobs View

The Jobs view displays all jobs, such as discovery, migration, and image capture, organized into two tabs:

- ◆ **Jobs:** All jobs submitted for execution.
- ◆ **Saved Jobs:** All saved jobs not yet submitted for execution. See [Section 3.8.4, “Starting a Saved Migration Job in the Jobs View,”](#) on page 85.

Figure 3-4 PlateSpin Migrate Client’s Jobs View



You can limit the scope of jobs displayed in the view. Use the **Job Type** and **Jobs Status** menus to specify filters for the view:

- ♦ **Job Type:** To view discovery, migration, or all other job types.
- ♦ **Job Status:** To view failed, currently running, and completed jobs.

3.2.3 Tasks Pane

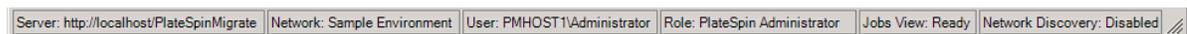
The Tasks pane of the PlateSpin Migrate Client window contains most essential migration actions. Clicking a task opens the Action window, which you can use to select the migration source, target, and setup method.

3.2.4 Status Bar

The status bar of the PlateSpin Migrate Client window displays information about:

- ♦ The PlateSpin Server that you are currently connected to.
- ♦ The PlateSpin Migrate Network that you are currently working with.
- ♦ The User that you are logged in as, and the PlateSpin Migrate role assigned to your user account.
- ♦ The status of the Automatic Network Discovery feature.

Figure 3-5 Status Bar of the PlateSpin Migrate Client Window



Double-clicking any of the first three status items opens the PlateSpin Server Settings window. See [“Connecting to a PlateSpin Server” on page 57](#).

Double-clicking the Network Discovery status item turns Automatic Windows Network Discovery on or off. See [“Discovering, Registering and Preparing a VM for a Server Sync Job” on page 69](#).

3.3 Discovering Source Workloads and Migration Targets

For any migration job, you must have a discovered source and a discovered target. A discovery operation populates the PlateSpin Migrate database with detailed inventory information about a machine. This information provides the data necessary to determine the machine’s use and to properly configure a migration job.

NOTE: All the workloads and the targets that you discover using the PlateSpin Migrate Web Interface are listed in the PlateSpin Migrate Client also. However, if you discover workloads and targets through the PlateSpin Migrate Client, then only the workloads and targets that you discover in the default network displays on the PlateSpin Migrate Web Interface.

- ♦ [Section 3.3.1, “Pre-populating the Servers View with Windows Computers in the Network,” on page 65](#)
- ♦ [Section 3.3.2, “Discovering the Details of Source Workloads and Targets,” on page 65](#)
- ♦ [Section 3.3.3, “Discovering All Machines in a Windows Domain,” on page 66](#)
- ♦ [Section 3.3.4, “Refreshing Machine Details,” on page 66](#)

- ◆ [Section 3.3.5, “Resetting the Mechanism Used to Authenticate Access to ESX Hosts,” on page 67](#)
- ◆ [Section 3.3.6, “Discovering Target Physical Machines,” on page 67](#)
- ◆ [Section 3.3.7, “Discovering, Registering and Preparing a VM for a Server Sync Job,” on page 69](#)
- ◆ [Section 3.3.8, “Discovery Guidelines for Machine Types and Credentials,” on page 70](#)

3.3.1 Pre-populating the Servers View with Windows Computers in the Network

PlateSpin Migrate’s Network Discovery feature automatically populates PlateSpin Migrate Client with a list of Windows physical and virtual machines that are online. Unlike a full discovery, Network discovery creates a list of machines but does not retrieve detailed inventory information about each one, which is required for workload portability jobs.

The Network Discovery feature is necessary to discover all machines of a Windows domain in a single job. See [“Discovering All Machines in a Windows Domain” on page 66](#).

PlateSpin Migrate uses the standard Windows network browser function for the auto-discovery feature. Because Linux workloads and virtual machine servers do not advertise to the Windows network browser, they are not auto-discovered.

The Network Discovery feature is enabled by default. To disable the feature or toggle between its enabled/disabled modes, double-click **Network Discovery** in the bottom right corner of the PlateSpin Migrate Client window.

3.3.2 Discovering the Details of Source Workloads and Targets

You can discover:

- ◆ An individual physical machine
- ◆ An individual virtual machine
- ◆ An individual virtual machine server
- ◆ Multiple virtual machine servers
- ◆ All VMware ESX hosts affiliated with a VMware vCenter Server
- ◆ Hyper-V hosts
- ◆ Multiple machines

Before starting discovery operations, make sure PlateSpin Server can communicate with your source workloads and targets. See [Section 2.3.1, “Requirements for Discovery,” on page 40](#).

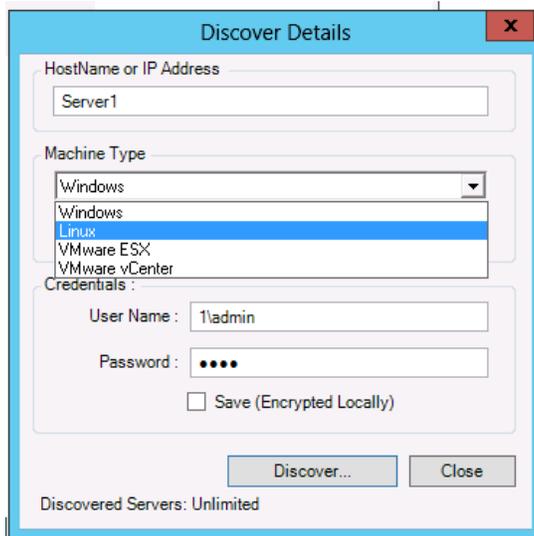
1 On the PlateSpin Migrate toolbar, click **Discover Details**.

or

In the Servers view, right-click in a blank area, then select **Discover Details**.

2 In the **Discover Details** dialog box, type the hostname or IP address of the source or target. To discover multiple machines, specify multiple hostnames or IP addresses separated by semicolons. If the target is behind a NAT device, specify its public (external) IP address.

See [“Migrations Across Public and Private Networks through NAT” on page 42](#).



- 3 Select the machine type and provide administrator credentials for the machine you are discovering.

See [“Discovery Guidelines for Machine Types and Credentials”](#) on page 70.

Discovering hosts with Xen Hypervisor systems results in these systems being registered as PlateSpin Migrate source workloads (as opposed to VM host targets). For information about using these platforms as workload portability targets, see [“Semi-Automated Workload Virtualization Using the X2P Workflow”](#) on page 122.

- 4 (Optional) If you want to store these credentials for use during future jobs, enable the **Save (Encrypted Locally)** option.
- 5 Click **Discover** and wait for the process to complete.
- 6 (Optional) If you want to monitor the progress of the job, switch to the Jobs view.

3.3.3 Discovering All Machines in a Windows Domain

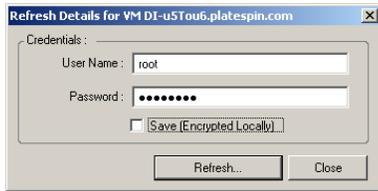
- 1 Enable the automatic Network Discovery feature.
See [“Pre-populating the Servers View with Windows Computers in the Network”](#) on page 65.
- 2 In PlateSpin Migrate Client, expand the domain listing containing the machines to be discovered.
- 3 Right-click the domain name, then select **Discover All Servers**.
- 4 Specify domain-level administrator credentials.
- 5 Click **Discover** and wait for the process to complete.
- 6 (Optional) If you want to monitor the progress of the job, switch to the Jobs view.

3.3.4 Refreshing Machine Details

You should routinely refresh your source workloads and targets before setting up a migration job.

To refresh a source or a target machine:

- 1 In the Servers view, right-click the required item, then select **Refresh Details**.



- 2 Specify the credentials appropriate for the system being refreshed, then click **Refresh**.
PlateSpin Migrate starts a discovery job, which you can monitor in the Jobs view.

3.3.5 Resetting the Mechanism Used to Authenticate Access to ESX Hosts

If you discover ESX hosts through a vCenter Server, all subsequent communications with these VM hosts take place through the vCenter Server.

You can reset this access mechanism to **Direct to ESX** or back to **Via vCenter**:

- 1 Refresh the required ESX host.
See [“Refreshing Machine Details”](#) on page 66.

The system prompts you to select an access type:



- 2 Select the required option, then click **Refresh**.
PlateSpin Migrate rediscovers the required ESX host by using the specified access mechanism, which it uses for all subsequent communications with that host.

3.3.6 Discovering Target Physical Machines

To discover a physical target and inventory its hardware components, you must boot the target machine with the PlateSpin boot ISO image on a CD or other media, from which your target can be booted. A similar process is used to migrate workloads to virtualization platforms that use the semi-automated virtualization process, whereby a target virtual machine on a virtual machine host is regarded as a physical target (see [“Semi-Automated Workload Virtualization Using the X2P Workflow”](#) on page 122).

- ♦ [“Downloading the PlateSpin Boot ISO Images”](#) on page 68
- ♦ [“Injecting Additional Device Drivers into the Boot ISO Images”](#) on page 68

- ◆ “Registering Physical Targets Using the PlateSpin Boot ISO Image” on page 69
- ◆ “Performing an Unattended Registration of a Target Physical Machine” on page 69

Downloading the PlateSpin Boot ISO Images

You can download the PlateSpin boot ISO image (`bootofx.x2p.iso`) from the PlateSpin Migrate area of [NetIQ Downloads](https://dl.netiq.com/index.jsp) (<https://dl.netiq.com/index.jsp>) by doing a search with the following parameters:

- ◆ **Product or Technology:** PlateSpin Migrate
- ◆ **Select Version:** 12.0
- ◆ **Date Range:** All Dates

IMPORTANT: The `.iso` file is compressed in `PhysicalTarget.zip` at the download site.

Injecting Additional Device Drivers into the Boot ISO Images

The PlateSpin boot ISO image contains a large library of device drivers sufficient to boot most common targets. However, occasionally you might want to use your own, such as lesser-known, vendor-specific or custom-developed drivers.

To inject drivers into the PlateSpin boot ISO image:

- 1 Download and extract the PlateSpin boot ISO images from [NetIQ Downloads](#). See [Downloading the PlateSpin Boot ISO Images](#).
- 2 Obtain or compile the required `*.ko` driver files.

IMPORTANT: The `rebuildiso.sh` script that helps you rebuild the ISO file is currently supported only on 64-bit version of the OS. Ensure that the drivers are valid for the kernel (`3.0.93-0.8-default`) included with the ISO file.

- 3 Mount the ISO image in any Linux machine (`root` credentials required). Use the following command syntax:

```
mount -o loop <path-to-ISO> <mount_point>
```

- 4 Copy the `rebuildiso.sh` script, located in the `/tools` subdirectory of the mounted ISO file, into a temporary working directory.
- 5 Create another working directory for the required driver files and save them in that directory.
- 6 In the directory where you saved the `rebuildiso.sh` script, run the following command as root:

```
./rebuildiso.sh -i <ISO_file> -d <driver_dir> -m32|-m64
```

On completion, the ISO file is updated with the additional drivers.

NOTE: To rebuild Migrate 12.0 LRD ISO, a minimum of `genisoimage 1.1.11` is required. By default, operating systems such as RHEL 7, CentOS 7, and Ubuntu 14.04.1 have the required `genisoimage` version.

- 7 Unmount the ISO file (execute the command `umount <mount_point>`).

Registering Physical Targets Using the PlateSpin Boot ISO Image

Complete these steps to register your target with PlateSpin Server when you are using an unmodified boot ISO image. If you want to inject additional device drivers into the image (Linux workloads only), first complete the steps in [Injecting Additional Device Drivers into the Boot ISO Images \(page 68\)](#).

To boot the target physical machine and register it with the PlateSpin Migrate Server:

- 1 Burn the PlateSpin boot ISO image on a CD or save it to the required media, from which your target can boot.
- 2 Boot the target machine from the ISO image.
- 3 (Conditional) For 64-bit systems, at the initial boot prompt, type the following:

```
ps64
```


Press Enter.
- 4 When prompted, enter the PlateSpin Server URL, using the following format:

```
http://<hostname | IP_address>/platespinmigrate
```


Replace `<hostname | IP_address>` with the hostname or the IP address of your PlateSpin Server host.
- 5 Enter your credentials for the PlateSpin Server.
- 6 Specify a static IP address or indicate that the machine should dynamically obtain an IP address from a DHCP server.
- 7 Enter a name for the target host and specify an existing PlateSpin Migrate Network to work with.

After a few moments, PlateSpin Migrate displays the physical target in the Servers view.

Performing an Unattended Registration of a Target Physical Machine

PlateSpin Migrate provides a mechanism for automating a target physical machine's registration with the PlateSpin Server. This involves updating the boot ISO image with specific registration information before booting the target.

For details, see [KB Article 7013485 \(https://www.netiq.com/support/kb/doc.php?id=7013485\)](https://www.netiq.com/support/kb/doc.php?id=7013485).

3.3.7 Discovering, Registering and Preparing a VM for a Server Sync Job

If you want to synchronize two workloads, and if your synchronization target is a virtual machine, you must discover and register an appropriate virtual machine first. For information about the Server Sync feature, see [“Synchronizing Workloads with Server Sync” on page 119](#).

- 1 On your virtual machine host, create a virtual machine with the desired specifications and install the operating system that matches the intended source workload, including the exact service pack.
- 2 Discover the virtual machine host or refresh its details.
- 3 In the Servers view, right-click the newly created virtual machine underneath the virtual machine server, then select **Prepare for synchronization**.
- 4 Specify administrator credentials for the virtual machine server.
- 5 (Optional) If you want to store these credentials for use during future jobs, enable the **Save (Encrypted Locally)** option.

- 6 (Optional) To configure the temporary (Take Control) network settings, such as choosing which virtual network to use from those available on the virtual machine server and configuring TCP/IP settings, click **Configure**, then configure the network settings as required.
- 7 Click **Prepare** and wait for the job to complete.

On completion, the Servers view lists a new Server Sync target underneath the VM host: 

3.3.8 Discovery Guidelines for Machine Types and Credentials

The following are guidelines for machine type selection, credential format, and syntax for discovery parameters.

Table 3-2 Guidelines for Discovery Parameters

To Discover	Machine Type Selection	Credentials	Remarks
All Windows workloads PlateSpin Image Servers	Windows	Local or domain admin credentials.	For the username, use this format: <ul style="list-style-type: none"> ◆ For domain member machines: <i>authority\principal</i> ◆ For workgroup member machines: <i>hostname\principal</i>
All Linux workloads	Linux	Root-level username and password	Non-root accounts must be properly configured to use <code>sudo</code> . See KB Article 7920711 (https://www.netiq.com/support/kb/doc.php?id=7920711).
VMware ESX hosts	VMware ESX	ESX account with admin role OR Windows domain credentials (versions 4 and 4.1 only)	
VMware ESX hosts affiliated with a VMware vCenter Server	VMware vCenter	VMware vCenter Web service credentials (username and password) OR Windows domain credentials (versions 4 and 4.1 only)	All subsequent communications with ESX hosts take place through the vCenter Server. To reset the access mechanism to Direct to ESX or back to Via vCenter , refresh the details of the required ESX host and select the required option. See “Resetting the Mechanism Used to Authenticate Access to ESX Hosts” on page 67.
Hyper-V hosts	Windows	Local or domain admin credentials.	For the username, use this format: <ul style="list-style-type: none"> ◆ For domain member machines: <i>authority\principal</i> ◆ For workgroup member machines: <i>hostname\principal</i>

3.4 Undiscovering Source Workloads and Migration Targets

You can undiscover objects (workloads or targets) through the PlateSpin Migrate Client or the PlateSpin Migrate Web Interface:

- ♦ To undiscover a object through the PlateSpin Migrate Client, right-click the object and select **Undiscover Server**.
- ♦ To undiscover or remove a workload through the PlateSpin Web Interface, on the Workloads page, select the workload and click **Remove Workload**.
- ♦ To undiscover or remove a target through the PlateSpin Web Interface, on the Targets page, click **Remove** next to the target you want to remove.

IMPORTANT

- ♦ If an object is listed both in the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface, then you must use the PlateSpin Migrate Web Interface to remove the object.
 - ♦ Before you delete an object that has jobs configured, you must ensure that all the configured jobs are completed.
-

3.5 Analyzing Machines with PlateSpin Analyzer

Before you begin any large-scale migration projects, you should identify potential migration problems and correct them beforehand. Use the PlateSpin Analyzer utility to determine whether discovered machines are suitable for migration jobs.

NOTE: PlateSpin Analyzer currently supports only Windows workloads.

To open PlateSpin Analyzer:

- 1 On the **Tools** menu, click **Analyze Servers**.
The PlateSpin Analyzer window opens.
- 2 Select the required PlateSpin Migrate Network and the required machines to analyze.
- 3 (Optional) To reduce the analysis time, limit the scope of machines to a specific language.
- 4 (Optional) To analyze machines in the inventory of a different PlateSpin Server, click **Connect**, then specify the required PlateSpin Server URL and valid credentials.
- 5 Click **Analyze**.

Depending on the number of discovered machines you select, the analysis might take a few seconds to several minutes.

Analyzed servers are listed in the left pane. Select a server to view test results in the right pane. Test results can be any combination of the following:

Table 3-3 Status Messages in PlateSpin Analyzer Test Results

Result	Description
Passed	The machine passed the PlateSpin Analyzer tests.
Warning	One or more tests returned warnings for the machine, indicating potential migration issues. Click the hostname to see the details.
Failed	One or more tests failed for this machine. Click the hostname to see the details and obtain more information.

The **Summary** tab provides a listing of the number of machines analyzed and not checked, as well as those that passed the test, failed the test, or were assigned a warning status.

The **Test Results** tab provides the following information:

Table 3-4 PlateSpin Analyzer Test Results Tab

Section	Details
System Test	Validates that the machine fulfills PlateSpin Migrate's minimum hardware and operating system requirements.
Take Control Hardware Support	Checks for source hardware compatibility for offline migration.
Target Hardware Support	Checks hardware compatibility for use as a target physical machine.
Software Test	Checks for applications that must be shut down for Live Transfer, and databases that should be shut down during Live Transfer to guarantee transactional integrity.
Incompatible Application Test	Verifies that applications known to interfere with the migration process are not installed on the system. These applications are stored in the Incompatible Application Database. To add, delete, or edit entries in this database, select Incompatible Application from the Tools menu.

The **Properties** tab provides detailed information about a selected machine.

For more information about using PlateSpin Analyzer and understanding the results, see [KB Article 7920478](https://www.netiq.com/support/kb/doc.php?id=7920478) (<https://www.netiq.com/support/kb/doc.php?id=7920478>).

3.6 Working with Device Drivers

PlateSpin Migrate ships with a library of device drivers, and during migration jobs it automatically installs the appropriate ones for the target. To determine if the required drivers are available, use PlateSpin Analyzer. See “[Analyzing Machines with PlateSpin Analyzer](#)” on page 71.

If PlateSpin Analyzer encounters missing or incompatible drivers, or if you require specific drivers for your target infrastructure, you might need to add (upload) drivers to the PlateSpin Migrate driver database.

In addition,

- ♦ [Section 3.6.1, “Packaging Device Drivers for Windows Systems,” on page 73](#)
- ♦ [Section 3.6.2, “Packaging Device Drivers for Linux Systems,” on page 73](#)
- ♦ [Section 3.6.3, “Uploading Drivers to the PlateSpin Migrate Device Driver Database,” on page 74](#)
- ♦ [Section 3.6.4, “Using the Plug and Play \(PnP\) ID Translator Feature,” on page 75](#)

3.6.1 Packaging Device Drivers for Windows Systems

To package your Windows device drivers for uploading to the PlateSpin Migrate driver database:

- 1 Prepare all interdependent driver files (*.sys, *.inf, *.dll, etc.) for your target infrastructure and device. If you have obtained manufacturer-specific drivers as a .zip archive or an executable, extract them first.
- 2 Save the driver files in separate folders, with a discrete folder per device.

The drivers are now ready for upload. See [Section 3.6.3, “Uploading Drivers to the PlateSpin Migrate Device Driver Database,” on page 74](#).

NOTE: For problem-free operation of your migration job and the target workload, upload *only digitally signed* drivers for:

- ♦ All 64-bit Windows systems
 - ♦ 32-bit versions of Windows Vista and Windows Server 2008, and Windows 7 systems
-

3.6.2 Packaging Device Drivers for Linux Systems

To package your Linux device drivers for uploading to the PlateSpin Migrate driver database, you can use a custom utility included in your Linux ISO boot image.

- 1 Find a Linux workstation that has the same kernel version as the source machine. Source machine itself is one of the best choices. On the Linux workstation, create a directory for your device driver files. All the drivers in the directory must be for the same kernel and architecture.
- 2 Download the boot image and mount it.

For example, assuming that the ISO has been copied under the /root directory, issue these commands:

```
# mkdir /mnt/ps
# mount -o loop /root/bootofx.x2p.iso /mnt/ps
```

- 3 From the /tools subdirectory of the mounted ISO image, copy the packageModules.tar.gz archive into a another working directory and extract it.

For example, with the .gz file is inside your current working directory, issue this command:

```
tar -xvzf packageModules.tar.gz
```

- 4 Enter the working directory and execute the following command:

```
./PackageModules.sh -d <path_to_driver_dir> -o <package name>
```

Replace <path_to_driver_dir> with the actual path to the directory where you saved you driver files, and <package name> with the actual package name, using the following format:

Drivername-driverversion-dist-kernelversion-arch.pkg

For example, bnx2x-1.48.107-RHEL4-2.6.9-11.EL-i686.pkg

The package is now ready for upload. See [Section 3.6.3, “Uploading Drivers to the PlateSpin Migrate Device Driver Database,”](#) on page 74

3.6.3 Uploading Drivers to the PlateSpin Migrate Device Driver Database

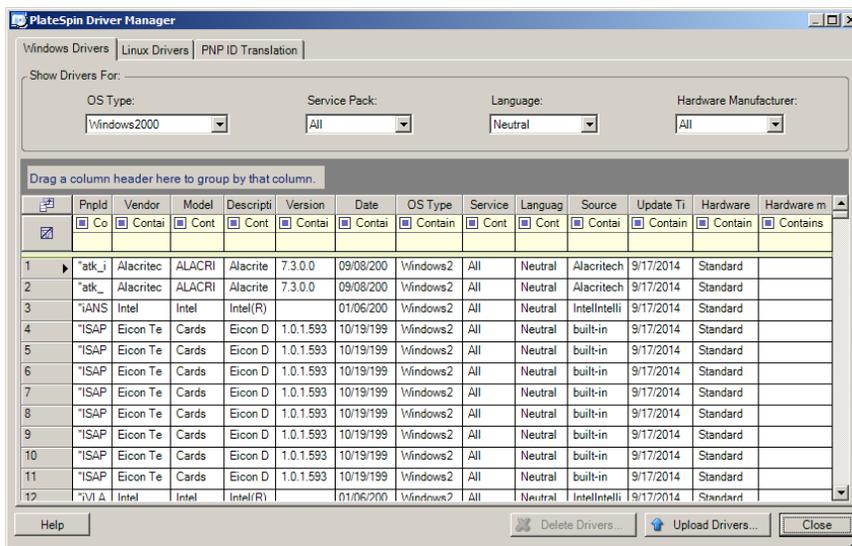
Use the PlateSpin Driver Manager to upload device drivers to the driver database.

NOTE: On upload, PlateSpin Migrate does not validate drivers against selected operating system types or their bit specifications; make sure that you only upload drivers that are appropriate for your target infrastructure.

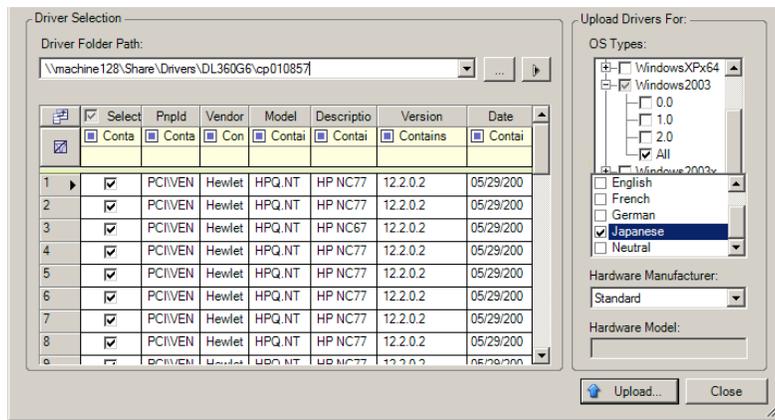
- ◆ “[Device Driver Upload Procedure \(Windows\)](#)” on page 74
- ◆ “[Device Driver Upload Procedure \(Linux\)](#)” on page 75

Device Driver Upload Procedure (Windows)

- 1 Obtain and prepare the required device drivers.
See [Packaging Device Drivers for Windows Systems](#).
- 2 Click **Tools > Manage Device Drivers** and select the **Windows Drivers** tab:



- 3 Click **Upload Drivers**, browse to the folder that contains the required driver files, and select applicable OS type, language, and hardware manufacturer options.



For most X2P migrations, select **Standard** as the **Hardware Manufacturer** option, unless your drivers are designed specifically for any of the target environments listed.

- 4 Click **Upload** and confirm your selections when prompted.

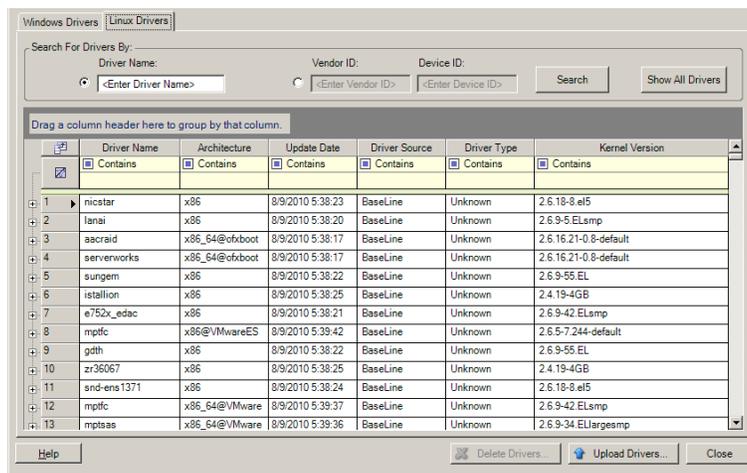
The system uploads the selected drivers to the driver database.

Device Driver Upload Procedure (Linux)

- 1 Obtain and prepare the required device drivers.

See [Packaging Device Drivers for Linux Systems](#).

- 2 Click **Tools > Manage Device Drivers** and select the **Linux Drivers** tab:



- 3 Click **Upload Drivers**, browse to the folder that contains the required driver package (*.pkg), and click **Upload All Drivers**.

The system uploads the selected drivers to the driver database.

3.6.4 Using the Plug and Play (PnP) ID Translator Feature

“Plug and Play” (PnP) refers to Windows operating system functionality that supports connectivity, configuration, and management with native plug and play devices. In Windows, the feature facilitates discovery of PnP compliant hardware devices attached to a PnP compliant bus. PnP compliant

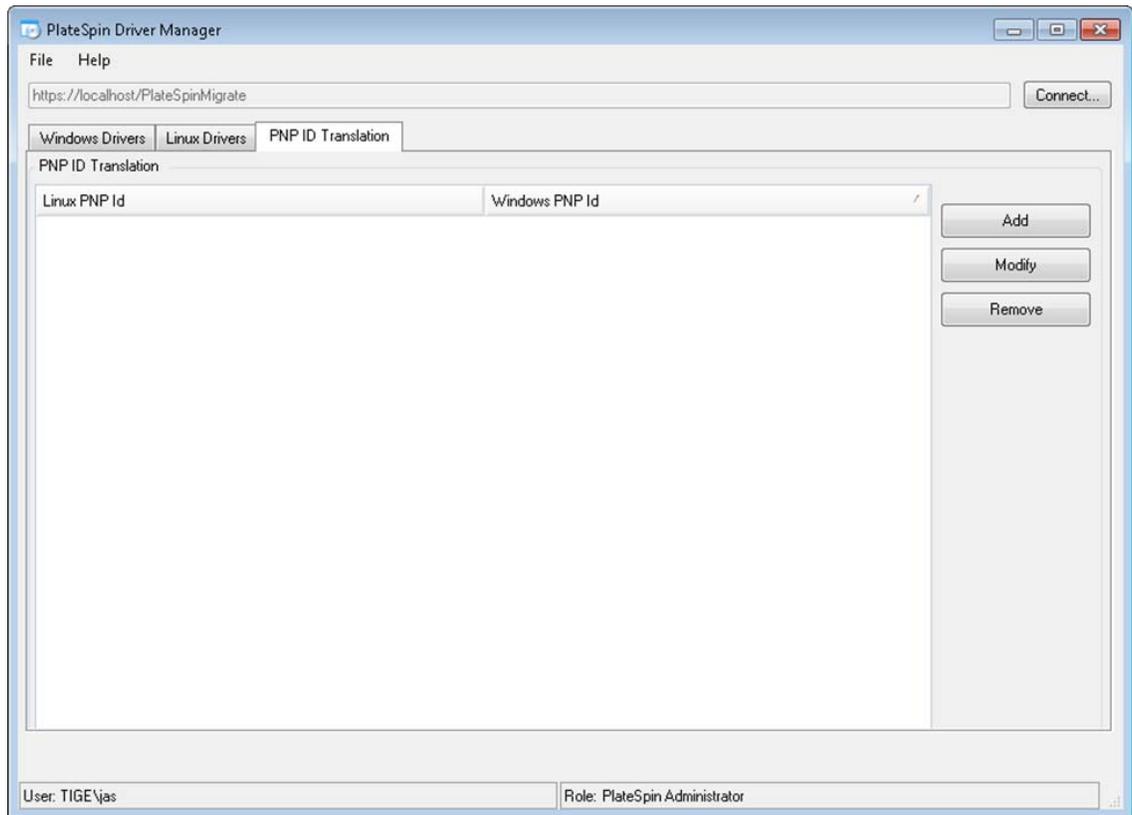
devices are assigned a set of Device Identification Strings by their manufacturer. These strings are programmed into the device when it is built. These strings are fundamental to how PnP works: they are part of the Windows' information source used to match the device with a suitable driver.

When the PlateSpin Server discovers workloads and their available hardware, the discovery includes these PnP IDs and the storage of that data as part of the workload's details. PlateSpin uses the IDs to determine which, if any, drivers need to be injected during a conversion operation. The PlateSpin Server maintains a database of PnP IDs for the associated drivers of each of the supported operating systems. Because Windows and Linux use different formats for PnP IDs, a Windows workload discovered by the Migrate Linux RAM disk contains Linux-style PnP IDs.

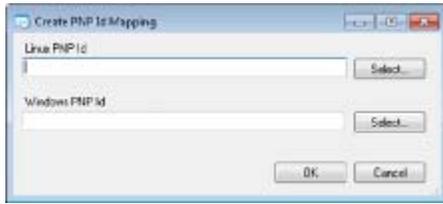
These IDs are formatted consistently, so PlateSpin can apply a standard transformation to each of them to determine its corresponding Windows PnP ID. The translation occurs automatically within the PlateSpin product. The feature lets you or a support technician add, edit or remove custom PnP mappings.

Follow these steps to use the PnP ID Translation feature:

- 1 Launch the PlateSpin Driver Manager tool and connect to the PlateSpin Server.
- 2 In the Driver Manager tool, select the PNP ID Translation tab to open the **PNP ID Translation** list, which includes the currently known custom PnP ID mappings.



- 3 On the list page, click **Add** to display the Create PNP ID Mapping dialog box.



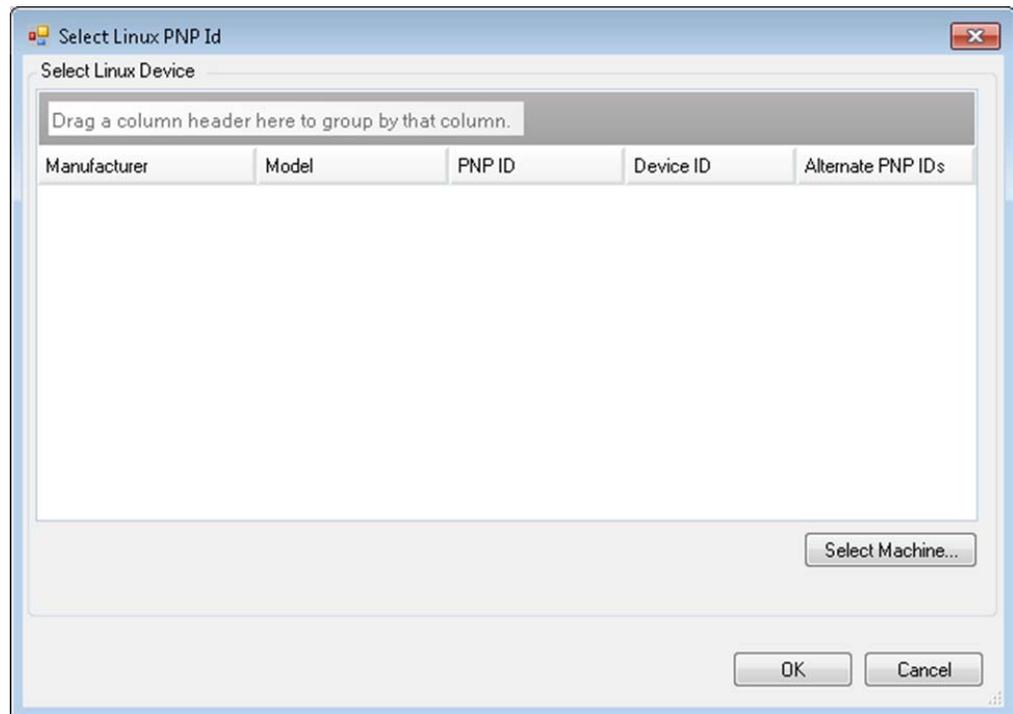
4 In the **Linux PnP ID** field, add a Linux PnP ID.

4a (Conditional) If you know it, type the Linux PnP ID you want to use.

or

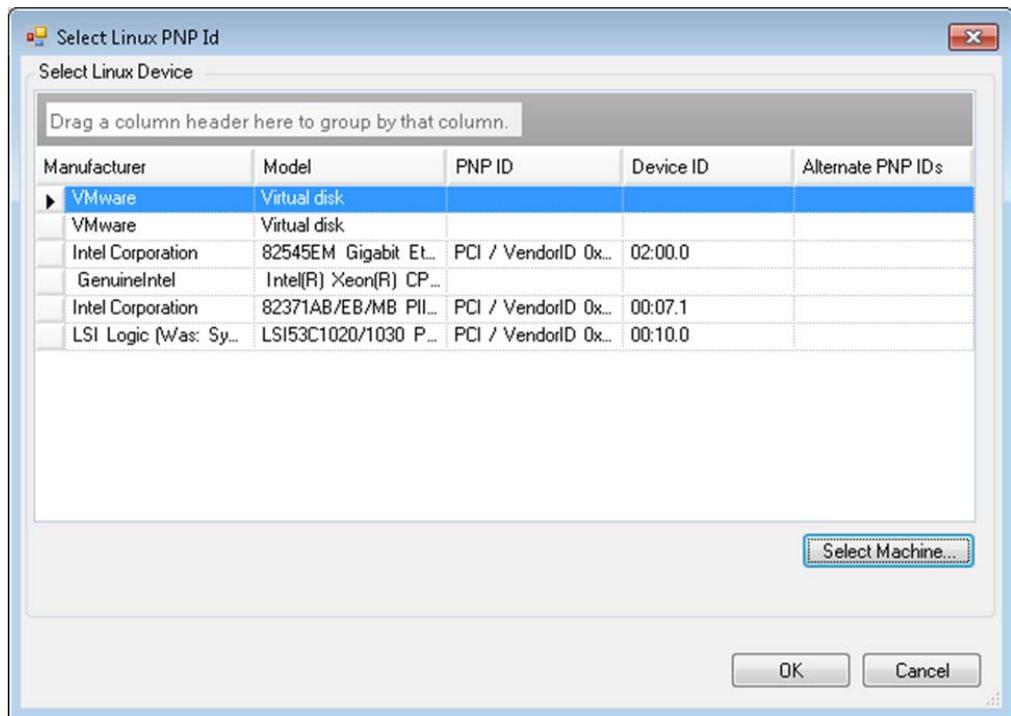
4b (Conditional) Select an ID from a previously discovered workload:

4b1 Adjacent to the **Linux PnP ID** field, click **Select** to open the Select Linux PnP ID dialog box.



4b2 On the dialog box, click **Select Machine** to display a list of the machines previously discovered by the PlateSpin Linux RAM disk.

4b3 Highlight one of the devices in the list, then click **Select** to populate the list in the Select Linux PnP ID dialog box.



4b4 Select a device on the list, then click **OK** to apply the standard transformation to the PnP ID and display it in the Create PnP ID Mapping dialog box.

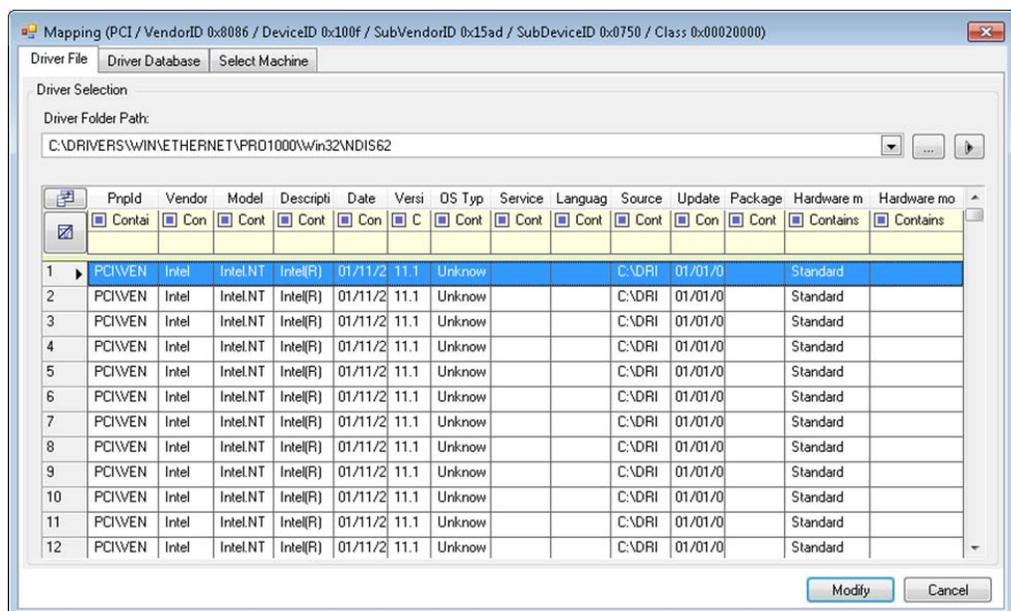
5 In the **Windows PNP ID** field, add a Windows PnP ID:

5a (Conditional) If you know it, type the Windows PnP ID you want to use.

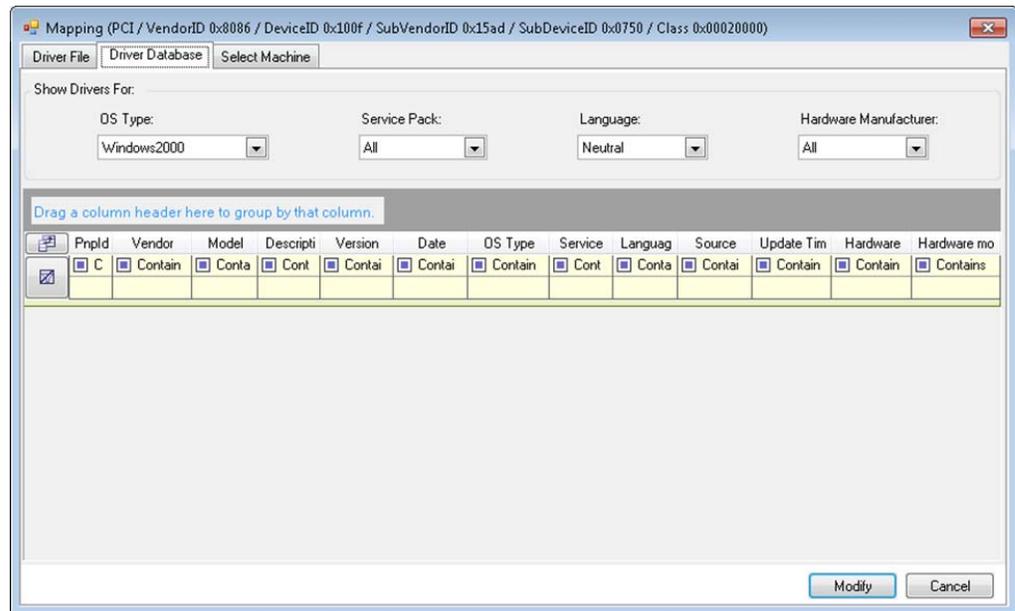
or

5b (Conditional) Adjacent to the **Windows PNP ID** field, click **Select** to open a mapping tool that presents three methods for helping you map a the Windows PnP ID:

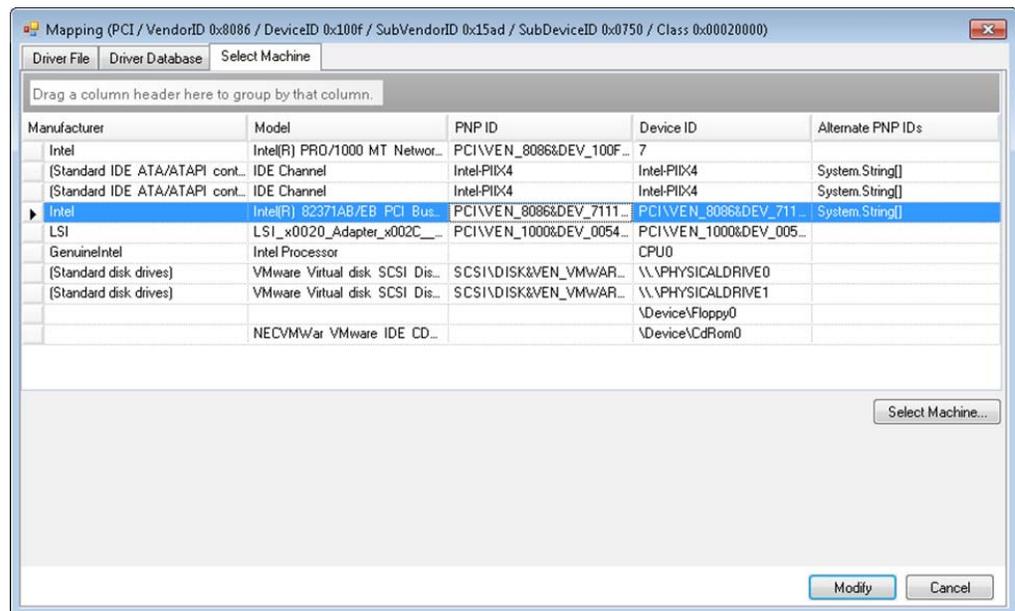
- Under the **Driver File** tab, browse to and select a Windows driver file (that is, a file with the *.inf extension), select the desired PnP ID, then click **Modify**.



- Under the **Driver Database** tab, browse to and select the existing driver database, select the correct PnP ID, then select **Modify**.

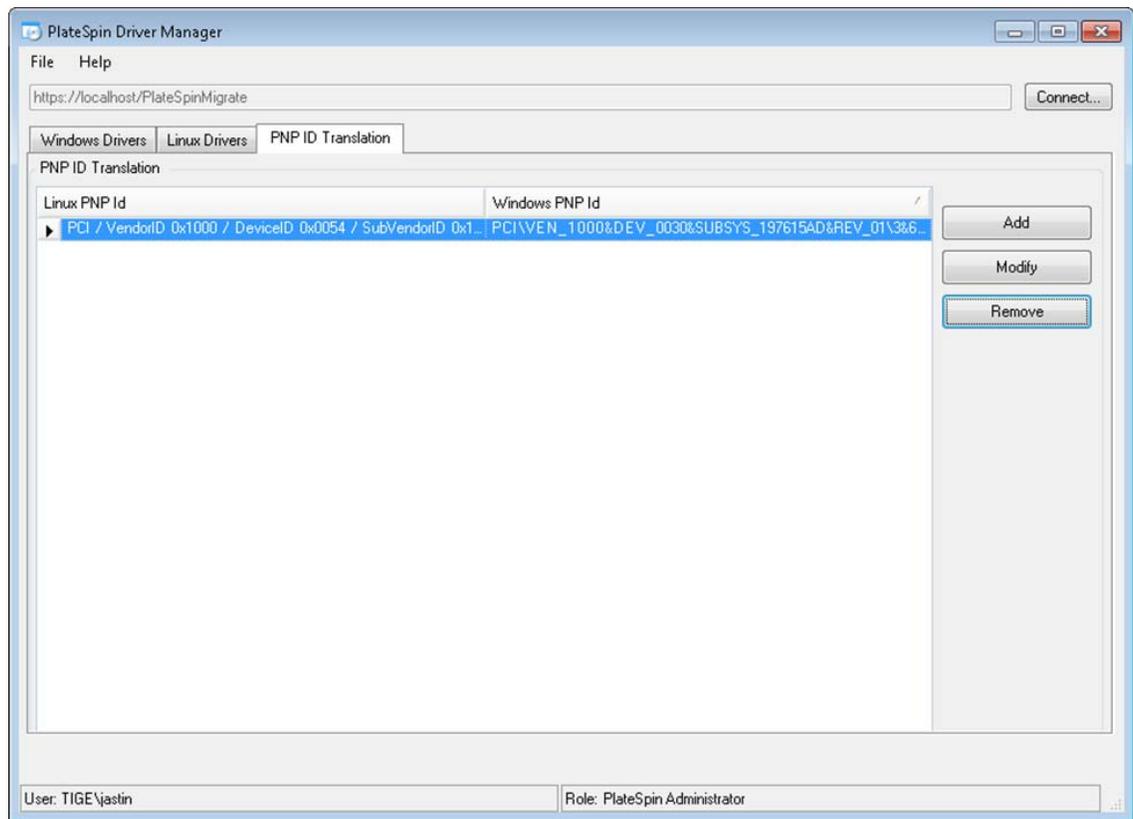


- Under the **Select Machine** tab, click **Select Machine**, then, from the list of Windows machines discovered using live discovery, select a machine, click **OK** to display its devices, select the desired PnP ID, then click **Modify**.



IMPORTANT: Selecting a Windows PnP ID that does not have an associated driver package installed might result in a failure at conversion time.

- In the Create PnP Id Mapping dialog box, confirm that the correct Linux PnP ID and the correct Windows PnP are selected, then click **OK** to display the PNP ID Translation page of the PlateSpin Driver Manager.



- 7 (Optional) To modify or remove the mapping in the PNP ID Translation list, select the mapping pattern, then click **Remove** or **Modify**, depending on the operation you want to perform.

Remove simply deletes the mapping (after displaying a confirmation dialog box).

To modify,

- 7a Click **Modify** to open the Create PNP id Mapping dialog box.
- 7b Repeat [Step 5 on page 78](#) to modify the Windows PnP ID.

NOTE: You cannot select or modify the Linux PnP ID.

3.7 Managing Custom Actions

PlateSpin Migrate provides you with the capability to automatically execute custom actions, such as batch files and scripts.

- ♦ [Section 3.7.1, “Managing Post-migration Actions \(Windows and Linux\),” on page 81](#)
- ♦ [Section 3.7.2, “Freeze and Thaw Scripting Capabilities \(Linux Block-Level Migrations\),” on page 82](#)

3.7.1 Managing Post-migration Actions (Windows and Linux)

To automate specific post-migration tasks on your target, you can include a custom action, such as a batch file, a shell script, or a program executable, in your migration job. At the end of the migration process, PlateSpin Migrate uploads the specified action, along with its dependencies, to the target and executes it.

Custom post-migration actions are supported for the following job types:

- One-time Server Sync
- Peer-to-peer workload migration

For the capability to select a post-migration action to run as part of a migration job, you must first save the action and its dependencies in a dedicated directory and add it to the PlateSpin Server's library. The maximum size of the directory must not exceed 64 MB. For information about raising this limit, see [Section 2.5.2, "Increasing the Size Limit on Post-Migration Actions Uploaded to the PlateSpin Server,"](#) on page 51.

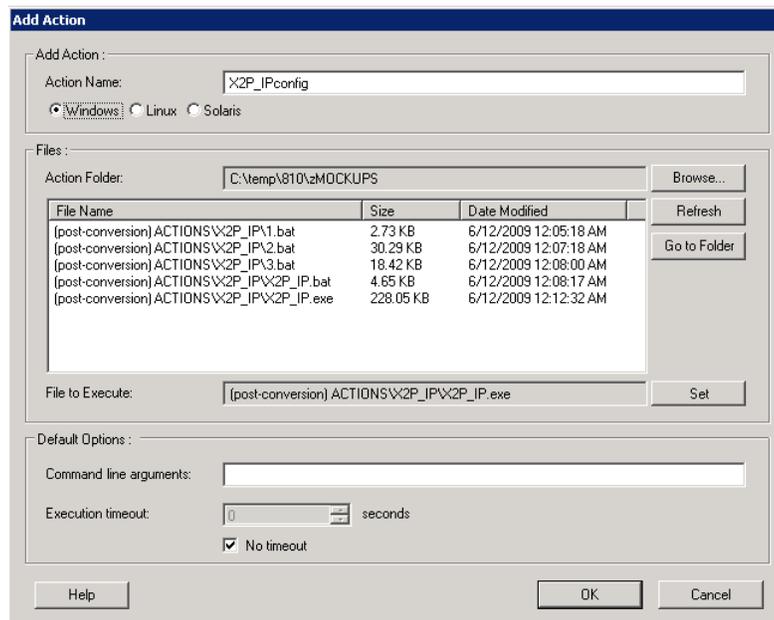
Use the following procedure for adding a post-migration action to the PlateSpin Server's library of custom actions.

- 1 Create the action, test it on a sample workload, and save it together with its dependencies in a directory that the PlateSpin Server can access.

Take special care when developing post-migration actions for Linux workloads, which allow different characters in file names and support different ACL (Access Control List) permissions. For Linux operating systems, amalgamate the action's directory structure into a single file.

See [KB Article 7970214 \(https://www.netiq.com/support/kb/doc.php?id=7970214\)](https://www.netiq.com/support/kb/doc.php?id=7970214).

- 2 In the PlateSpin Migrate Client, click **Tools > Manage Actions**.
- 3 Click **Add**.



- 4 In the Add Action window, type a name for your custom action, select the target operating system type, then browse to and select the directory that contains the required action with its dependencies.

PlateSpin Migrate populates the list with the contents of the selected folder.

- 5 In the **File Name** column, select the required executable, then click **Set**.
- 6 In the **Default Options** section, specify any required command line arguments and an execution timeout, then click **OK**.

PlateSpin Migrate packages and uploads the library.

The action is now available for selection in migration jobs. See [“Including a Custom Post-migration Action in a Migration Job” on page 178](#).

3.7.2 Freeze and Thaw Scripting Capabilities (Linux Block-Level Migrations)

PlateSpin Migrate provides an additional means of control over your Linux block-level migration process — the `freeze` and `thaw` shell scripts.

These scripts are executed during Linux workload migrations, at the beginning and end of block-level data transfer sessions. Specifically, they interject in the migration process in the following fashion:

1. First pass of all volumes without snapshots:
 - ◆ Regular (non-LVM) volumes
 - ◆ LVM without enough space to take a snapshot
2. *Freeze script*
3. Take snapshots
4. Second pass of all non-snapshot volumes
5. *Thaw script*
6. Transfer volume snapshots

You can use this capability to complement the automated daemon control feature provided through the user interface (see [“Handling Source Workload Services or Daemons During Live Transfer \(Windows and Linux\)” on page 166](#)).

For example, you might want to use this feature to cause an application to flush its data to disk so that the workload remains in a more consistent state during a Live Transfer migration.

To use the feature, do the following before setting up your migration job:

- 1 Create the following files:
 - ◆ `platespin.freeze.sh`— to contain the `freeze` shell script logic
 - ◆ `platespin.thaw.sh`— to contain the `thaw` shell script logic
 - ◆ `platespin.conf`. A text file defining any required arguments, along with a timeout value.

The required format for the contents of the `platespin.conf` file is:

```
[ServiceControl]
(optional) FreezeArguments=<arguments>
(optional) ThawArguments=<arguments>
(optional) TimeOut=<timeout>
```

Replace `<arguments>` with the required command arguments, separated by a space, and `<timeout>` with a timeout value in seconds. If unspecified, the default timeout is used (60 seconds).

- 2 Save the scripts, along with the `.conf` file, on your Linux source workload, in the following directory:

```
/etc/platespin/
```

3.8 Setting Up, Executing, and Managing Jobs

Most discovery and portability operations are defined, executed, and monitored as PlateSpin Migrate jobs.

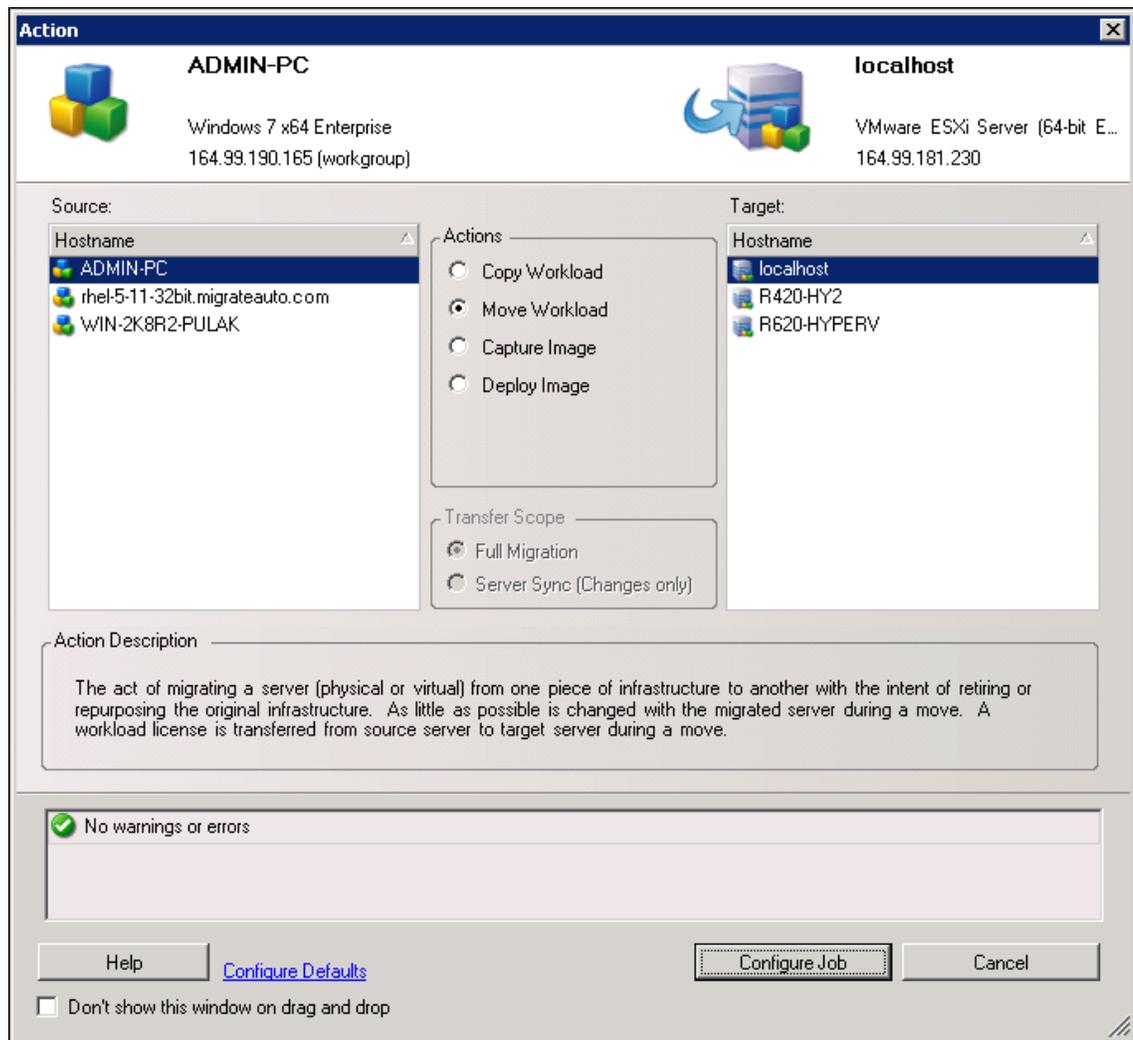
- ◆ [Section 3.8.1, “Setting a Migration Job,” on page 83](#)
- ◆ [Section 3.8.2, “Pre-Migration Job Validation,” on page 85](#)
- ◆ [Section 3.8.3, “Saving Migration Jobs,” on page 85](#)
- ◆ [Section 3.8.4, “Starting a Saved Migration Job in the Jobs View,” on page 85](#)
- ◆ [Section 3.8.5, “Changing the Start Time of a Scheduled Migration Job,” on page 86](#)
- ◆ [Section 3.8.6, “Viewing the Parameters of a Completed or In-Progress Migration Job,” on page 86](#)
- ◆ [Section 3.8.7, “Canceling an In-Progress Job,” on page 86](#)
- ◆ [Section 3.8.8, “Controlling the State of Source Workloads in the Jobs View,” on page 86](#)
- ◆ [Section 3.8.9, “Generating Job Reports,” on page 86](#)
- ◆ [Section 3.8.10, “Job Diagnostics,” on page 87](#)

3.8.1 Setting a Migration Job

For any migration job, you must have a discovered source and a discovered target. For information about discovering source workloads and migration targets, see [Section 3.3, “Discovering Source Workloads and Migration Targets,” on page 64](#).

To start a migration job:

- 1 Do one of the following to open the Action window:
 - ◆ Drag a discovered source and drop it on a discovered target.
 - ◆ Click a task in the Tasks pane.
 - ◆ Click the New Job toolbar.
 - ◆ In the Jobs view, right-click a source and select a command from the context menu. Available commands depend on the type of source.



The Source and Target panes display workloads and targets applicable to the selected type of a migration job.

The **Full Transfer** and **Server Sync** options are enabled in the following circumstances:

- ◆ The system detects an existing operating system on the target
- ◆ The operating system profile of the target matches that of the source workload

See [“Synchronizing Workloads with Server Sync” on page 119](#).

- 2 Check the validation messages at the bottom of the window.
- 3 To start configuring your migration job, click **Configure Job**.
- 4 (Optional) For convenience, to avoid displaying the **Action** window on drag-and-drop, select **Don't show this window on drag and drop** before you proceed. Subsequent drag-and-drop actions bypass the Action window and directly open a Conversion Job window.

To restore the job migration startup behavior, restore the application defaults. See [“General Options” on page 44](#).

3.8.2 Pre-Migration Job Validation

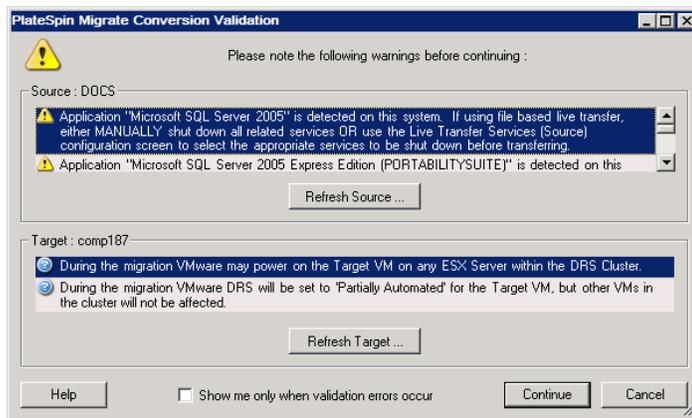
When starting a migration job, PlateSpin Migrate validates the job type, the source, the target, and the selected parameters, and might generate errors and warnings.

 Error markers show configurations that you need to change before the migration job can start.

 Warning markers alert you to settings that should be confirmed prior to starting the migration.

In a default PlateSpin Migrate configuration, validation messages display at the bottom of the Action window. However, if you have configured PlateSpin Migrate to bypass the Action window on drag-and-drop, errors and warnings are displayed in a separate window:

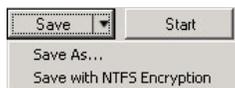
Figure 3-6 Migration Validation Window



To force this window to open only on errors, select the **Show me only when validation errors occur** option.

3.8.3 Saving Migration Jobs

You can save a migration job for executing it at a later time. After you start a job and configure all the options, click the arrow at the right side of the **Save** button, then select a command from the drop-down list:



3.8.4 Starting a Saved Migration Job in the Jobs View

- 1 In the PlateSpin Migrate Client's Jobs view, click the **Saved Jobs** tab and locate the required saved job.
- 2 Right-click the job and select **Start**.
PlateSpin Migrate starts the saved migration job.

3.8.5 Changing the Start Time of a Scheduled Migration Job

- 1 In the PlateSpin Migrate Client's Jobs view, locate the required scheduled job.
- 2 Right-click the job and select **Change Start Time**.

The Change Job Start Time dialog box opens.



- 3 Specify the required start date and time and click **OK**.
PlateSpin Migrate reschedules the job and executes it at the specified time.

3.8.6 Viewing the Parameters of a Completed or In-Progress Migration Job

- 1 In the PlateSpin Migrate Client's Jobs view, locate the required job.
- 2 Right-click the job and select **View**.

PlateSpin Migrate opens the job configuration window and displays the job's configuration parameters in read-only mode.

3.8.7 Canceling an In-Progress Job

- 1 In the PlateSpin Migrate Client's Jobs view, locate the required job.
- 2 Right-click the job and select **Abort**.

3.8.8 Controlling the State of Source Workloads in the Jobs View

You can restart or shut down your source workload in the Jobs view, unless it is in an active migration job.

- 1 In the PlateSpin Migrate Client's Jobs view, locate the required job.
- 2 Right-click the job and select **Restart Source** or **Shutdown Source** as applicable.

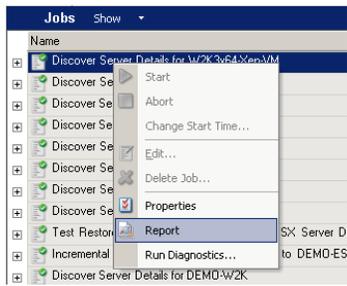
To automate the startup state of source and target workloads, specify the required post-migration state in your migration job. See [“Specifying Conversion Options” on page 147](#).

3.8.9 Generating Job Reports

You can generate detailed reports of running and completed jobs. A report records the tasks performed during the job.

To generate a job report:

- 1 In the PlateSpin Migrate Client's **Jobs** view, locate the required job.
- 2 Right-click the job and select **Report**.



A Web browser window displays the requested report.

3.8.10 Job Diagnostics

PlateSpin Migrate provides a tool that can produce a diagnostics report for any running or completed job.

To view a diagnostics report:

- 1 In the PlateSpin Migrate Client Jobs view, right-click the required job and select **Run Diagnostics**.

The diagnostics report is displayed in a browser window. This process might take a few moments.

The diagnostics report lists several statistics:

- ♦ All the operations involved in the job. Click any operation to view its XML representation.
- ♦ The status of each operation.
- ♦ The controller that ran the operation. Click the controller to view its XML representation, or click **Logs** to view its event log.

In addition, the report contains links to:

- ♦ The XML representations of the source machine, original target machine, and the target VM host.
- ♦ The root operation for the job, as well as a variety of logs and reports.

Diagnostics reports can be sent directly to technical support by following the instructions provided on the report.

4 Working with the PlateSpin Migrate Web Interface

PlateSpin Migrate includes a Web Interface that supports large scale migration of physical and virtual source workloads to target VMware containers or Microsoft Azure containers. Use this interface when you want to simultaneously migrate a large number of workloads.

Before you use the PlateSpin Migrate Web Interface to migrate workloads to Azure, you must plan the migration. See [Section 4.1, “Planning For Migrating Workloads to Azure,” on page 90](#).

IMPORTANT

- ◆ Do not use the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface interchangeably to perform the migration tasks during the migration process of a workload. NetIQ Corporation strongly recommends that you either use the PlateSpin Migrate Client or the PlateSpin Migrate Web Interface throughout the migration cycle of a workload. For information about the migration tasks you can perform using the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface, see [Section 1.1.3, “Migration Tasks Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface,” on page 15](#).
- ◆ You cannot use the PlateSpin Migrate Web Interface to migrate the following:
 - ◆ PlateSpin Images
 - ◆ Workloads to the target non-VMware hypervisor containers
 - ◆ Workloads to target physical machines
 - ◆ Azure workloads to target Azure containers.
- ◆ The default filename for a VM's virtual disk uses the workload's FQDN, such as `<fqdn>_VM_N.vmdk`, where N is the disk number. The total length limit of the virtual disk filename is 63 characters, including the fixed naming conventions and file extension.

You can alter certain elements of the Web Interface to match your organization branding. For more information, see [“Rebranding the PlateSpin Migrate Web Interface” on page 221](#).

Review the following sections:

- ◆ [Section 4.1, “Planning For Migrating Workloads to Azure,” on page 90](#)
- ◆ [Section 4.2, “Enabling an Azure Subscription to Deploy the Replication Environment VM,” on page 91](#)
- ◆ [Section 4.3, “Azure Networking Guidelines,” on page 92](#)
- ◆ [Section 4.4, “Accessing the PlateSpin Migrate Web Interface,” on page 93](#)
- ◆ [Section 4.5, “Navigating the Web Interface,” on page 94](#)
- ◆ [Section 4.6, “Dashboard,” on page 95](#)
- ◆ [Section 4.7, “Discovering Source Workloads and Targets,” on page 95](#)
- ◆ [Section 4.8, “Workloads,” on page 96](#)
- ◆ [Section 4.9, “Targets,” on page 108](#)
- ◆ [Section 4.10, “Tasks,” on page 108](#)

- ♦ [Section 4.11, “Generating Reports Using the PlateSpin Migrate Web Interface,” on page 109](#)
- ♦ [Section 4.12, “Configuring the Settings For the Web Interface,” on page 110](#)
- ♦ [Section 4.13, “Configuring the Refresh Rates for the Web Interface,” on page 115](#)

4.1 Planning For Migrating Workloads to Azure

PlateSpin Migrate 12.1 enhances the Web Interface to let you migrate the following Windows and Linux workloads to Microsoft Azure:

Windows:

- ♦ Microsoft Windows Server 2012 R2
- ♦ Microsoft Windows Server 2012
- ♦ Microsoft Windows Server 2008 R2

Linux:

- ♦ Red Hat Enterprise Linux (RHEL) 7.1
- ♦ Red Hat Enterprise Linux (RHEL) 6.7
- ♦ SUSE Linux Enterprise Server (SLES) 11 SP4
- ♦ SUSE Linux Enterprise Server (SLES) 11 SP3

NOTE

- ♦ Migration of Microsoft Azure source workloads to Microsoft Azure target containers is not supported.
 - ♦ Migration of Linux workloads to an Azure target is restricted to workloads with only a single NIC.
-

Consider the following points before you use the PlateSpin Migrate Web Interface to migrate workloads to Azure:

- ♦ Migration of Windows cluster workloads is not supported because Microsoft Azure does not support Windows clusters.
- ♦ Migration of UEFI workloads is not supported.
- ♦ The PlateSpin Migrate Client does not support migration of workloads to Microsoft Azure. You can use only the PlateSpin Migrate Web Interface to migrate the workloads to Microsoft Azure.
- ♦ Test Cutover of workloads is not supported. You can perform only Run Cutover of workloads.
- ♦ PlateSpin Migrate supports Azure VM sizes with up to 64 data disks. For the maximum instance size in a selected Azure Region, Migrate will use one data disk for the OS disk replication in the PlateSpin Replication Environment. After migration, this disk becomes the OS disk, and you can add a data disk.

Each data disk must have a maximum size of 1TB (1024 GB).

- ♦ Migrate recommends an Azure VM instance size that meets or exceeds the source workload's settings for cores, memory, data disks, and NICs. However, you can choose a smaller or larger instance size based on your requirements for the target workload, as limited by the maximum instance sizes available in the selected Azure Region.
- ♦ The size of the disk created on the Azure VM is the size of the source disk partition plus about 1 GB because of the granularity of available disk sizes on Azure.

- ◆ You need an OS license for the migrated target workload. For Azure target workloads, you must provide Azure with the license information or Microsoft will charge you for the OS license.
- ◆ For each target Azure subscription, you must enable programmatic deployment for the PlateSpin Migrate Replication Environment VM. See [Section 4.2, “Enabling an Azure Subscription to Deploy the Replication Environment VM,”](#) on page 91.
- ◆ Currently, when the time on the PlateSpin Server goes out of sync, the cutover will fail with a 403 forbidden error. Can we either detect that the root cause is the time problem, and state that as the error message or add that to it - or otherwise build in some sort of very visible warning that the time is out of sync
- ◆ Ensure that the PlateSpin Server host displays the correct time for the time zone it is in. If the time on the PlateSpin Server host is incorrect, the cutover process fails with a 403 forbidden error.

4.1.1 Prerequisites for Migrating Workloads to Azure

PlateSpin Migrate uses the Microsoft Azure Resource Manager for migrating workloads to Microsoft Azure. Before you use PlateSpin Migrate to migrate workloads to Microsoft Azure, ensure that the following cloud access prerequisites are correctly configured and available:

- ◆ A Microsoft Azure Account and a subscription for your account
- ◆ A Client ID required to access Microsoft Azure
- ◆ An Azure Active Directory user created as a contributor for the subscription
- ◆ A Virtual Network with a Subnet different from the default Gateway Subnet
- ◆ The minimum network-related prerequisites for a successful migration are:
 - ◆ The source and the target workload must be able to communicate with the PlateSpin Migrate server on port 443. The target workload is the replica of the source workload that will reside in Microsoft Azure.
 - ◆ The PlateSpin Migrate server must be able to communicate with the Microsoft Azure API endpoint on port 443.
 - ◆ The PlateSpin Migrate server must be able to communicate with the source workloads on the ports that are used for DCOM, WMI, and RPC.
 - ◆ The target workload must be able to reach the source workload on port 3725 (default). The direction of this connection can be reversed (source to target), and the port number is configurable. For information about changing the default settings, see [Chapter 2, “Post-Installation Configuration,”](#) on page 27.

For information about setting up your Azure cloud account to work with PlateSpin Migrate, see “PlateSpin Migrate Best Practices for Migrating Windows Workloads to Microsoft Azure” on the [PlateSpin Migrate Resources Webpage](#).

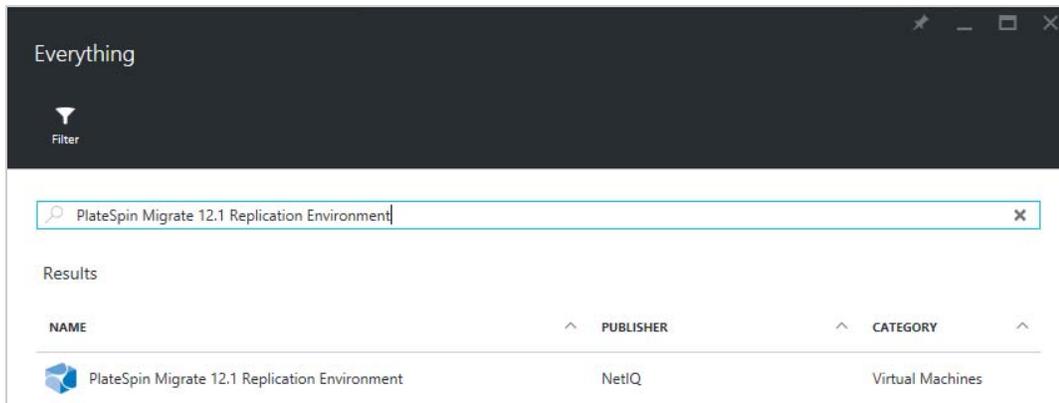
4.2 Enabling an Azure Subscription to Deploy the Replication Environment VM

PlateSpin Migrate must be able to programmatically deploy a PlateSpin Migrate Replication Environment VM during the replication of workloads to Azure. The required VM image is available in the Azure Marketplace. You must enable programmatic deployment of the image for each subscription that will perform migrations with PlateSpin Migrate.

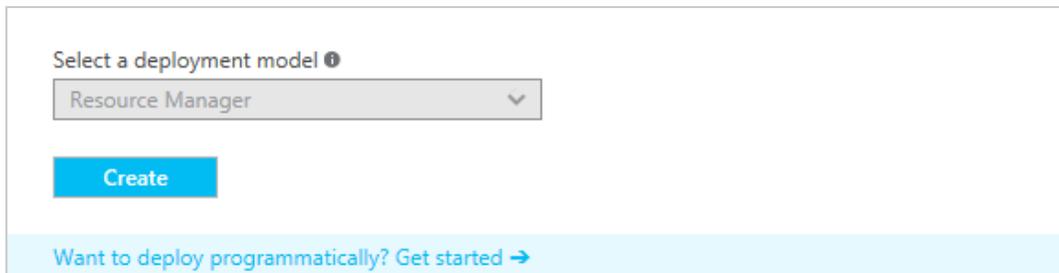
NOTE: All migrations for the target subscription will fail until you enable the subscription to deploy the Replication Environment VM.

To enable programmatic deployment of the Replication Environment VM:

- 1 Go to the [Microsoft Azure Portal \(https://portal.azure.com\)](https://portal.azure.com) and log in to your Azure account.
- 2 In the portal menu, click **New**, then search for the **PlateSpin Migrate 12.1 Replication Environment** image in the Azure Marketplace.
- 3 In the Results panel, select the **PlateSpin Migrate 12.1 Replication Environment**.



- 4 At the bottom of the PlateSpin Migrate 12.1 Replication Environment page under **Select a deployment model**, click **Want to deploy programmatically? Get Started**.



- 5 On the Configure Programmatic Deployment page, read the Terms of Use.
- 6 For each subscription that will perform migrations with PlateSpin Migrate, change its status from **Disable** to **Enable**.
- 7 Click **Save**.

4.3 Azure Networking Guidelines

Consider the following guidelines when you are migrating workloads to Microsoft Azure.

- ♦ [Section 4.3.1, "Windows Workloads with Multiple NICs," on page 93](#)
- ♦ [Section 4.3.2, "Private and Public IP Addresses for Workloads Connected on an Azure VPN," on page 93](#)

4.3.1 Windows Workloads with Multiple NICs

For Windows workloads with multiple NICs, Azure configures the VM with a default gateway that is associated with the primary network interface. Azure removes the gateway information for all secondary NICs, which limits their communications to the same subnet as the primary interface.

To enable a secondary NIC to communicate outside its own subnet, you can use the Windows `route add` command to add a different gateway entry for the secondary NIC in the routing table. See “Configure Windows VMs” in *Create a VM with Multiple NICs* (<https://azure.microsoft.com/en-us/documentation/articles/virtual-networks-multiple-nics/>) on the Microsoft Azure website (<https://azure.microsoft.com/>).

4.3.2 Private and Public IP Addresses for Workloads Connected on an Azure VPN

Each Azure VM has both a public IP address and a private IP address for communications from machines outside the Microsoft Azure cloud environment. Azure automatically associates these IP addresses with the primary network interface for the VM.

You can use the Microsoft Remote Desktop client to connect remotely to the Azure VM. Specify the IP address as follows:

- ♦ **Private IP address:** Use the VM’s private IP address if your machine is part of the address space for the Azure VPN.
- ♦ **Public IP address:** Use the VM’s public IP address if your machine is not part of the address space for the Azure VPN.

You can alternatively use the **Connect** option in the Microsoft Azure portal (<https://azure.microsoft.com/en-us/features/azure-portal/>) from a machine with an address space that is not part of the Azure VPN. This option automatically launches the Microsoft Remote Desktop client configured to connect to the VM’s public IP address for the primary NIC.

NOTE: This portal operation fails if your machine is in the address space of the Azure VPN.

4.4 Accessing the PlateSpin Migrate Web Interface

To access the PlateSpin Migrate Web Interface, use one of the following browsers:

- ♦ **Google Chrome:** Version 34.0 and later
- ♦ **Microsoft Internet Explorer:** Version 11.0 and later
- ♦ **Mozilla Firefox:** Version 29.0 and later

You must ensure that JavaScript (Active Scripting) is enabled in the browser. To enable JavaScript in the browser, do the following:

- ♦ **Google Chrome browser:**
 1. From the menu, click **Settings** and scroll to click **Show advanced settings**.
 2. Under **Privacy**, click **Content settings**.
 3. Scroll to **JavaScript** and select **Allow all sites to run JavaScript**.
 4. Click **Done**.

- ♦ **Firefox browser:**
 1. In the Location bar, type `about:config` and press Enter.
 2. Click **I'll be careful, I promise!**
 3. In the **Search** bar, type `javascript.enabled` and press Enter.
 4. In the search results, view the value for the `javascript.enabled` parameter. If its value is `false`, right-click `javascript.enabled` and select **Toggle** to set its value to `true`.
- ♦ **Internet Explorer:**
 1. From the Tools menu, select **Internet Options**.
 2. Click **Security** and then click **Custom level**.
 3. Scroll to **Scripting > Active scripting** and select **Enable**.
 4. Click **Yes** on the warning dialog box, then click **OK**.
 5. Click **Apply > OK**.

To launch PlateSpin Migrate Web Interface:

- 1 Open a web browser and go to:


```
https://hostname or IP_address/Migrate
```

 Replace *hostname* or *IP_address* with the DNS hostname or IP address of your PlateSpin Migrate Server.
 If SSL is not enabled, use `http` in the URL.
- 2 Log in using the local Administrator user credentials for the PlateSpin Server host or as an authorized user.
 For information about setting up additional users for PlateSpin, see [Section 2.2, "Setting Up User Authorization and Authentication,"](#) on page 31.

4.5 Navigating the Web Interface

The Web Interface displays a navigation bar with the following links:

- ♦ **Dashboard:** Displays the default Dashboard page that provides information about the Migrate licenses, latest tasks, running events, upcoming events, and past events. See [Section 4.6, "Dashboard,"](#) on page 95.
- ♦ **Workloads:** Displays the Workloads page that lists all the discovered workloads. To add or discover a workload, click **Add Workload** option on the Dashboard page or the Workloads page. For more information about adding or discovering a workload, see [Section 4.8.1, "Adding or Discovering a Source Workload,"](#) on page 97.
 You can perform various other tasks such as configuring a workload, preparing a workload for migration, and migrating a workload. See [Section 4.8, "Workloads,"](#) on page 96.
- ♦ **Targets:** Displays the Targets page that lists the already added target containers and lets you add new targets. See [Section 4.9, "Targets,"](#) on page 108.
- ♦ **Tasks:** Displays the Tasks page that lists items requiring user intervention. See [Section 4.10, "Tasks,"](#) on page 108.
- ♦ **Reports:** Displays the Reports page. See [Section 4.11, "Generating Reports Using the PlateSpin Migrate Web Interface,"](#) on page 109.
- ♦ **Settings:** Displays the Settings page that lets you configure the following:
 - ♦ **Workload Tags:** See [Section 4.12.1, "Using Tags to Help Sort Workloads,"](#) on page 110.

- ♦ **Permissions:** See [Section 4.12.2, “Managing PlateSpin Migrate Security Groups and Workload Permissions,”](#) on page 111.
- ♦ **General Notification Settings:** See [“Setting Up Automatic Event Notifications by Email”](#) on page 112.
- ♦ **Report Notification Settings:** See [“Setting Up Automatic Replication Reports by Email”](#) on page 113.
- ♦ **SMTP:** See [“SMTP Configuration”](#) on page 112.
- ♦ **Licenses:** See [Section 4.12.4, “Configuring Product Licenses,”](#) on page 114.

4.6 Dashboard

The Dashboard page provides information about the Migrate licenses, tasks, running events, upcoming events, and past events.

The left pane of the Dashboard page provides a high-level view of the overall state of the PlateSpin Migrate workload inventory, summary of the license information and also lets you add or discover a new workload. For more information about adding or discovering a workload, see [Section 4.8.1, “Adding or Discovering a Source Workload,”](#) on page 97.

The right pane of the Dashboard page provides information about events and tasks that requires user attention.

4.7 Discovering Source Workloads and Targets

Before you migrate a workload through the Web Interface, you must first add or discover the workload and a target. You must also ensure that PlateSpin Server can communicate with your source workloads and targets. For more information, see [section Section 2.3.1, “Requirements for Discovery,”](#) on page 40.

For more information about workloads and how to add workloads, see [Section 4.8, “Workloads,”](#) on page 96.

For more information about targets and how to add targets, see [Section 4.9, “Targets,”](#) on page 108.

NOTE

- ♦ All the workloads and the targets that you discover using the PlateSpin Migrate Web Interface are listed in the PlateSpin Migrate Client also. However, if you discover workloads and targets through the Platespin Migrate Client, then only the workloads and targets that you discover in the default network displays on the PlateSpin Migrate Web Interface.
- ♦ When you use the Web Interface to discover workloads and targets, a default heartbeat delay of 15 seconds (15000 ms) is set on the controller.

To enable a heartbeat delay of shorter or longer duration, do the following:

1. On the Migrate server computer, open the registry editor.
 2. Go to `HKLM\SOFTWARE\PlateSpin\OperationsFramework\Controller`.
 3. Add a key named `HeartbeatStartupDelayInMS` of type `REG_SZ` and set its value to the desired value in milliseconds. The default setting should be 15000.
 4. Restart the server computer.
-

4.8 Workloads

The Workloads page displays the following information for each discovered workload:

Item	Description
Tasks	Displays a warning icon for a task that might require user attention. For example: if a workload goes offline, then a warning icon displays. Click the icon to see more details.
Online	Displays one of the following: <ul style="list-style-type: none">◆ Yes: If the workload is online.◆ No: If the workload is offline.
Workload	Displays the workload name. Click the workload name to configure the workload for migration.
Tag	Displays the tag associated with the workload. For more information about the tags, see Section 4.12.1, “Using Tags to Help Sort Workloads,” on page 110 .
Schedule	Displays the status of the schedule if you have configured a schedule for the workload migration. For example: if the schedule is configured, it displays Active after you have prepared the workload for migration until the end of the migration cycle, unless you pause the schedule. If you click Pause Schedule , then Paused displays. If you click Resume Schedule , then Active displays again.
Status	Displays the current status of the workload. For example: <ul style="list-style-type: none">◆ Adding Workload: The process of adding or discovering a workload is in progress.◆ Not Configured: The workload has been discovered but is not yet configured.◆ Migration Configured: The workload has been configured for migration.◆ Preparing Migration: The workload is being prepared for migration.◆ Running First Replication: The workload is being replicated for the first time.
Last Replication	Displays the date when the workload was last replicated.
Next Replication	Displays the date when the workload is scheduled for the next replication.
Last Test Cutover	Displays the date when the target workload was last tested.

NOTE: All time stamps reflect the time zone of the PlateSpin Server host. This might be different from the time zone of the source workload or the time zone of the host on which you are running the PlateSpin Migrate Web Interface. A display of the server date and time appears at the bottom right of the client window.

4.8.1 Adding or Discovering a Source Workload

To migrate a workload through the Web Interface, you must first add or discover the workload. Before you discover a workload, ensure that PlateSpin Server can communicate with your source workloads and targets. For more information, see section [Section 2.3.1, “Requirements for Discovery,”](#) on page 40.

- 1 In the PlateSpin Migrate Web Interface, click **Workloads > Add Workload**.
Alternatively, you can click the **Add Workload** option on the Dashboard page.
- 2 Specify the hostname or the IP address of the workload you want to add.
- 3 Select the type of workload.
- 4 Specify the credentials to connect to the workload.
- 5 Click **Add Workload** to discover the workload and list it on the Workloads page.

4.8.2 Modifying the Listing of the Workloads

On the Workloads page, you can limit the display of the discovered workloads. For example:

- ♦ To display all the workloads that are not yet configured, select the **Workload Status** option as **Not Configured** and the **Tag** option as **All**.
- ♦ To display all the failed Windows workloads, select the **Workload Status** option as **Failed Workloads** and the **Tag** option as **Windows**.

For information about how to create tags and associate them with workloads, see [Section 4.12.1, “Using Tags to Help Sort Workloads,”](#) on page 110.

To modify the listing of the workloads:

- 1 In the **Workload Status** option, select one of the following:
 - ♦ **All Workloads:** Displays all the discovered workloads.
 - ♦ **Ready For Cutover:** Displays the workloads that are replicated at least once.
 - ♦ **Pending:** Displays the workloads on which some action is still in progress.
 - ♦ **Failed Workloads:** Displays the workloads that have failed an operation such as configuring migration, preparing migration, or running migration.
 - ♦ **Not Configured:** Displays all the workloads that are not yet configured.
- 2 In the **Tag** option, select the tag associated with the workloads you want to list.

For information about how to create tags and associate them with workloads, see [Section 4.12.1, “Using Tags to Help Sort Workloads,”](#) on page 110.

4.8.3 Understanding the Migration Tasks

You can perform the following tasks on a workload that you want to migrate depending on the current state of the workload. For example: After you add or discover a workload, you must configure the migration details for the discovered workload. You cannot directly prepare the discovered workload for migration, unless you have configured the migration details for the workload.

Task	Description
Configure Migration	Allows you to configure the settings required for migrating the workload. See Section 4.8.4, “Configuring the Workload for Migration,” on page 98.

Task	Description
Prepare Migration	Installs the required data transfer software on the source and creates a target workload (a virtual machine) on the target container in preparation for the workload replication. See Section 4.8.5, “Preparing the Migration,” on page 107 .
Run Migration	Starts replicating the workload according to specified parameters. See Section 4.8.6, “Running the Migration,” on page 107 .
Test Cutover	Boots and configures the target workload within the container for testing purposes.
Run Cutover	Boots and configures the target workload to complete the migration cycle.
Remove Workload	Removes a workload from the inventory.
Pause Schedule	Suspends all scheduled replications until the schedule is resumed.
Resume Schedule	Resumes the workload replication according to the saved replication settings.

4.8.4 Configuring the Workload for Migration

When you add or discover a workload, the workload is listed on the Workloads page and the status is set as **Not Configured**. Before you migrate the workload, you must configure the workload for migration:

- 1 Launch the PlateSpin Migrate Web Interface.
- 2 On the Workloads page, select the workload you want to configure.
- 3 Click **Configure Migration**.
- 4 Select one of the following based on the scope of data you want to transfer from the source to the target:
 - ♦ **Full Replication:** A full volume of data transfer takes place from the source to the target.
 - ♦ **Incremental Replication:** Only differences are transferred from the source to the target, provided they have similar operating system and volume profiles.

NOTE

Incremental Replication of data to Microsoft Azure is not supported.

- 5 Select the target to which you want to migrate the source data.
- 6 Click **Configure Migration**.
- 7 Configure the following settings:

Setting Name	Description
Schedule Settings	

Setting Name	Description
Incremental Recurrence	<p>Specify the following:</p> <ul style="list-style-type: none"> ◆ Start of Recurrence: The date when you want to start the replication. You can specify the date or click the calendar icon to select the date. By default, the time is 12.00 a.m. ◆ Recurrence Pattern: The pattern to follow for the recurrence of the replication. For example: <ul style="list-style-type: none"> ◆ To use incremental recurrence everyday, select Daily. ◆ To never use incremental recurrence, select None. <p>NOTE: When you schedule incremental recurrence, the replication takes place for a maximum period of 60 days from the starting time of replication. For example:</p> <ul style="list-style-type: none"> ◆ If you select Daily, then the replication takes place for 60 days from the time the replication starts. ◆ If you select Weekly, then the replication takes place for 8 weeks from the time the replication starts. ◆ If you select Monthly, then the replication takes place for 2 months from the time the replication starts.
Full Replication	<p>Do one of the following:</p> <ul style="list-style-type: none"> ◆ To specify a schedule for the replication, click Start and specify the date when you want to start the full replication. ◆ To start full replication manually without setting a schedule, click None. <p>NOTE: You must prepare the workload prior to the scheduled time. If the workload is not prepared by the scheduled time, then full replication is skipped.</p>
Blackout Window	<p>Use these settings to force a replication blackout. The replication blackout suspends scheduled replications during peak utilization hours or prevents conflicts between VSS-aware software and the PlateSpin VSS block-level data transfer component.</p> <p>To specify a blackout window, click Edit and do the following:</p> <ul style="list-style-type: none"> ◆ Specify the start and end time for the blackout period. ◆ Select one of the blackout recurrence pattern such as daily, weekly, or monthly. If you do not want to force a replication blackout, select None. <p>NOTE: The blackout start and end times are based on the system clock on the PlateSpin Server.</p>
Compression Level	<p>These settings control how workload data is compressed before transmission. See Section 1.5.2, “Data Compression,” on page 25. Select one of the following options:</p> <ul style="list-style-type: none"> ◆ Fast: Consumes the least CPU resources on the source, but yields a lower compression ratio. ◆ Optimal: Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option. ◆ Maximum: Consumes the most CPU resources on the source, but yields a higher compression ratio.

Setting Name	Description
Bandwidth Throttling	<p>These settings control the bandwidth throttling. PlateSpin Migrate enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. You can specify a throughput rate for each migration job. This provides a way to prevent migration traffic from congesting your production network and reduces the overall load of your PlateSpin Server.</p> <p>To throttle replications to a specified rate, specify the required throughput value in Mbps and the time pattern.</p>
Migration Settings	
Transfer Method	<p>(Windows) Select a data transfer mechanism and security through encryption. See Section 1.3, "Supported Transfer Methods," on page 21.</p> <p>To enable encryption, select the Encrypt Data Transfer option. See Section 1.4, "Security and Privacy," on page 23</p> <p>NOTE: The Offline Transfer with Temporary Boot Environment transfer method is not applicable for the Web interface.</p>
Transfer Encryption	<p>(Linux) To enable encryption, select the Encrypt Data Transfer option. See Section 1.4, "Security and Privacy," on page 23.</p>
Source Credentials	<p>Specify the credentials required for accessing the workload. See Section 3.3.8, "Discovery Guidelines for Machine Types and Credentials," on page 70.</p>
Cloud Instance Size	<p>(For migration to Azure) Select the cloud instance size appropriate for your workload. By default, the cloud instance size that most closely matches your workload is selected. The default selection is based on Azure's definition of instance for the following component types:</p> <ul style="list-style-type: none"> ◆ Number of cores ◆ Amount of RAM ◆ Number of data disks ◆ Number of NICs <p>The default instance either meets or exceed the settings for these components on the source workload. However, you can choose an instance based on your requirements. For example:</p> <ul style="list-style-type: none"> ◆ To reduce the number of cores or amount of RAM on the target workload, select a smaller cloud instance size in the list. ◆ To migrate only some of the data disks, select a cloud instance size with fewer data disks that best fits your need. Deselect the data disks that should not be migrated until the number of disks for migration fits the selected instance

Setting Name	Description
CPU	<p>(VM containers using VMware 5.1, 5.5, and 6.0 with a minimum VM hardware Level 8) Specify the number of sockets and the number of cores per socket for the target workload. It automatically calculates the total cores. This parameter applies on the initial setup of a workload with an initial replication setting of Full Replication.</p> <p>NOTE: The maximum number of cores the workload can use is subject to external factors such as the guest operating system, the VM hardware version, VMware licensing for the ESXi host, and ESXi host compute maximums for vSphere (see <i>vSphere 5.1 Configuration Maximums</i> (http://www.vmware.com/pdf/vsphere5/r51/vsphere-51-configuration-maximums.pdf)).</p> <p>Some distributions of a guest OS might not honor the cores and cores per socket configuration. For example, guest OSES using SLES 10 SP4 and OES 2 SP3 retain their original cores and sockets settings as installed, whereas other SLES, RHEL, and OES distributions honor the configuration.</p>
Number of CPUs	<p>(VM containers using VMware 4.1) Specify the required number of vCPUs (virtual CPUs) to assign to the target workload. This parameter applies on the initial setup of a workload with an initial replication setting of Full Replication. Each vCPU is presented to the guest OS on the VM container as a single core, single socket.</p>
Replication Network	<p>Separate replication traffic based on virtual networks defined on your VM container.</p> <p>For this setting, you can also specify an MTU value that the PlateSpin Migrate Linux RAM Disk (LRD) replication network can use. Setting a low value helps to avoid jabber over networks. For example: a VPN.</p> <p>The default value is an empty string. When networking is configured in the LRD, it allows the network device to set its own default, which is usually 1500. However, if you specify a value, PlateSpin Migrate adjusts the MTU when it configures the network interface.</p>
Networks Allowed for Replication	<p>Specify one or more network interfaces (NIC or IP address) on the source to use for replication traffic.</p>
Resource Pool for Target VM	<p>(VM container is part of a DRS Cluster.) Specify the Resource Pool location where the migrated VM is to be created.</p>
VM Folder for Target VM	<p>(VM container is part of a DRS Cluster.) Specify the VM folder location where the migrated VM is to be created.</p>
Virtual Machine Name	<p>Specify a display name for the new virtual machine.</p>
Configuration File Datastore	<p>Select a datastore associated with your VM for storing VM configuration files.</p>
Virtual Machine Configuration Path	<p>Specify the path to the configuration file on the target virtual machine.</p>
Disks	<p>Specify the path to the hard disk on the target virtual machine.</p> <p>Select Thin Disk to enable the thin-provisioned virtual disk feature. The virtual disk then appears to the VM to have a set size, but only consumes the amount of disk space that is actually required by data on that disk.</p>
Volumes	<p>Select volumes to be included in the target for migration.</p>
Converted Logical Volumes	<p>(Linux) Specify one or more LVM logical volumes to be migrated for a Linux workload.</p>

Setting Name	Description
Non-volume Storage	(Linux) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups	(Linux) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.
Volume Groups	(Linux) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the Converted Logical Volumes section of the settings.
Services to Stop Temporarily on Source Before Any Replication	(Windows) We recommend that all the non-VSS compliant services or anti-virus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored back as soon as the VSS snapshot creation completes.
Services to Stop Permanently on Source For Cutover With Any Replication	(Windows) Select the Windows services that should be permanently stopped on the source workload throughout the cutover process to ensure application data consistency. The services stopped on the source machine during the cutover process are not restored after the cutover process is complete even if you do not choose to shut down the source machine during the cutover process.
Daemons to Stop During Replication	(Linux) Select Linux daemons to be automatically stopped during the replication.
Target Workload Settings	
(These settings are applied during the Run Cutover.)	
VM Memory	Specify the amount of memory allocated to the target workload.
VM Tools	To install the VM tools, select the Install VM Tools option . This option is selected by default.
Hostname	Do one of the following: <ul style="list-style-type: none"> ◆ To retain the same hostname, select No Change. ◆ To change the hostname, select Set To and specify the new name.

Setting Name	Description
<p>System Identifier (SID) - (This Setting is applicable only for Windows Server 2008, Windows Server 2003, Windows Server 2000, Windows XP, and Windows Vista)</p>	<p>Before you generate a new SID for the Windows Server 2003, Windows Server 2000, or Windows XP target workload computer, you must do the following:</p> <ul style="list-style-type: none"> ◆ Enable the SID generation: <ol style="list-style-type: none"> 1. Open a web browser and go to: <code>https://hostname or IP_address/platespinconfiguration</code> Replace <i>hostname</i> or <i>IP_address</i> with the DNS hostname or IP address of your PlateSpin Migrate Server. If SSL is not enabled, use <code>http</code> in the URL. 2. On the PlateSpin Server Configuration page, set alwaysGenerateNewSid to True. ◆ Ensure that the hostname of the source and target workloads are different. <p>To generate a new system identifier for the target workload, select Generate New System Identifier (SID) in the Target Workload Test Settings section of the Web Interface. For Windows Server 2008 and Windows Vista, you must specify the local Administrator account credentials. If this account has been locally renamed on the source, provide the new name. If this account is disabled on Windows Vista (by default), you must first enable it.</p>

Setting Name	Description
Domain / Workgroup	<p>(Windows) Depending on whether the source workload belongs to workgroup or domain, one of the following displays:</p> <ul style="list-style-type: none"> ◆ Workgroup: <i>Workgroup_name</i> where <i>Workgroup_name</i> is the workgroup name to which the source belongs. ◆ Domain: <i>Domain_name</i> where <i>Domain_name</i> is the domain name to which the source belongs. <p>Do one of the following depending on where you want the target workload to join:</p> <ul style="list-style-type: none"> ◆ When the source workload belongs to a workgroup: Assume that the source workload belongs to a workgroup named WorkGroup1. <ul style="list-style-type: none"> ◆ For the target workload to join the same workgroup (WorkGroup1), retain the following existing selection: Workgroup: Workgroup1 ◆ For the target workload to join a different workgroup (say WorkGroup2), select Join Workgroup and specify the name as WorkGroup2. ◆ For the target workload to join a domain, select Join Domain and specify the domain name you want the target to join. ◆ When the source workload belongs to a domain: Assume that the source workload belongs to a domain named Domain1. <ul style="list-style-type: none"> ◆ For the target workload to join a workgroup, click Join Workgroup and specify the name of the workgroup you want the target to join. ◆ For the target workload to join the same domain (Domain1) with the domain registration settings preserved, retain the following existing selection: Domain: Domain1 ◆ For the target workload to join the same domain (Domain1) without preserving the domain registration settings, select Join Domain and specify the domain name as Domain1. ◆ For the target workload to join a different domain, select Join Domain and specify the domain name you want the target to join.
Domain Credentials	<p>(Windows) If you select Join Domain, specify the domain administrator credentials.</p>
Network Connections	<p>Select the local area connection and then select one of the following:</p> <ul style="list-style-type: none"> ◆ DHCP: Obtain an IP address automatically assigned by a DHCP server. ◆ Static: Specify a static IP address. <p>For Windows workloads that have more than one NIC, select the connection for each NIC.</p> <p>(For migration to Azure) For Windows workloads that have more than one NIC, the Primary option displays (by default) next to the first network name in the list. To make any other network as primary, click the Make Primary option next to the corresponding network name. For more information, see Section 4.3, "Azure Networking Guidelines," on page 92.</p>

Setting Name	Description
DNS Servers	<p>Specify the DNS Servers for the target workloads. This is applicable only if you select Static in the Network Connections option:</p> <ul style="list-style-type: none"> ◆ Primary DNS server: Specify the primary DNS server address. ◆ Alternative DNS server: Specify an alternate DNS server address. ◆ Additional DNS server: To specify additional DNS server addresses: <ol style="list-style-type: none"> 1. Click Advanced. 2. Specify the DNS server address. 3. Click Add to add the server in the DNS Server Addresses list. 4. Click OK.
Services States on Target VM	(Windows) Select Windows services that must be automatically stopped on the target VM.
Daemons States to Change	(Linux) Select Linux daemons that must be automatically stopped on the target VM.
Target Workload Test Settings	
(These settings are applied during the Test Cutover)	
VM Memory	Specify the amount of memory allocated to the target workload.
VM Tools	To install the VM tools, select the Install VM Tools option . This option is selected by default.
Hostname	<p>Do one of the following:</p> <ul style="list-style-type: none"> ◆ To retain the same hostname, select No Change. ◆ To change the hostname, select Set To and specify the new name.
System Identifier (SID) - (This Setting is applicable only for Windows Server 2008, Windows Server 2003, Windows Server 2000, Windows XP, and Windows Vista)	<p>Before you generate a new SID for the Windows Server 2003, Windows Server 2000, or Windows XP target workload computer, you must do the following:</p> <ul style="list-style-type: none"> ◆ Enable the SID generation: <ol style="list-style-type: none"> 1. Open a web browser and go to: <code>https://hostname or IP_address/platespinconfiguration</code> Replace <i>hostname</i> or <i>IP_address</i> with the DNS hostname or IP address of your PlateSpin Migrate Server. If SSL is not enabled, use <code>http</code> in the URL. 2. On the PlateSpin Server Configuration page, set alwaysGenerateNewSid to True. ◆ Ensure that the hostname of the source and target workloads are different. <p>To generate a new system identifier for the target workload, select Generate New System Identifier (SID) in the Target Workload Test Settings section of the Web Interface. For Windows Server 2008 and Windows Vista, you must specify the local Administrator account credentials. If this account has been locally renamed on the source, provide the new name. If this account is disabled on Windows Vista (by default), you must first enable it.</p>

Setting Name	Description
Domain / Workgroup	<p>(Windows) Depending on whether the source workload belongs to workgroup or domain, one of the following displays:</p> <ul style="list-style-type: none"> ◆ Workgroup: <i>Workgroup_name</i> where <i>Workgroup_name</i> is the workgroup name to which the source belongs. ◆ Domain: <i>Domain_name</i> where <i>Domain_name</i> is the domain name to which the source belongs. <p>Do one of the following depending on where you want the target workload to join:</p> <ul style="list-style-type: none"> ◆ When the source workload belongs to a workgroup: Assume that the source workload belongs to a workgroup named WorkGroup1. <ul style="list-style-type: none"> ◆ For the target workload to join the same workgroup (WorkGroup1), retain the following existing selection: Workgroup: Workgroup1 ◆ For the target workload to join a different workgroup (say WorkGroup2), select Join Workgroup and specify the name as WorkGroup2. ◆ For the target workload to join a domain, select Join Domain and specify the domain name you want the target to join. ◆ When the source workload belongs to a domain: Assume that the source workload belongs to a domain named Domain1. <ul style="list-style-type: none"> ◆ For the target workload to join a workgroup, click Join Workgroup and specify the name of the workgroup you want the target to join. ◆ For the target workload to join the same domain (Domain1) with the domain registration settings preserved, retain the following existing selection: Domain: Domain1 ◆ For the target workload to join the same domain (Domain1) without preserving the domain registration settings, select Join Domain and specify the domain name as Domain1. ◆ For the target workload to join a different domain, select Join Domain and specify the domain name you want the target to join.
Domain Credentials	<p>(Windows) If you select Join Domain, specify the domain administrator credentials.</p>
Network Connections	<p>Select the network connection and then select one of the following:</p> <ul style="list-style-type: none"> ◆ DHCP: Obtain an IP address automatically assigned by a DHCP server. ◆ Static: Specify a static IP address.

Setting Name	Description
DNS Servers	Specify the DNS Servers for the target workloads. This is applicable only if you select Static in the Network Connections option: <ul style="list-style-type: none"> ◆ Primary DNS server: Specify the primary DNS server address. ◆ Alternative DNS server: Specify an alternate DNS server address. ◆ Additional DNS server: To specify additional DNS server addresses: <ol style="list-style-type: none"> 1. Click Advanced. 2. Specify the DNS server address. 3. Click Add to add the server in the DNS Server Addresses list. 4. Click OK.
Services States on Target VM	(Windows) Select Windows services that must be automatically stopped on the target VM.
Daemons States to Change	(Linux) Select Linux daemons that must be automatically stopped on the target VM.
Tag	
Tag	Select a tag to assign to the workload. For more information about tags, see Section 4.12.1, "Using Tags to Help Sort Workloads," on page 110 .

8 (Optional) To change the target, click **Change Target**.

NOTE: If you change the target, all the settings you specified will be cleared.

9 Do one of the following:

- ◆ Click **Save** to save the settings.
- ◆ Click **Save and Prepare** to save the settings and start preparing the workload migration.
- ◆ Click **Cancel** to exit.

4.8.5 Preparing the Migration

After you configure the workload, you must prepare the workload for migration:

- 1 On the Workloads page, select the workload you want to migrate.
- 2 Click **Prepare Migration**, and then click **Execute**.

4.8.6 Running the Migration

- 1 On the Workloads page, select the workload you want to migrate.
- 2 Click **Run Migration**.
- 3 (Optional) To boot and configure the target workload immediately after full replication, select **Run cutover after successful replication**, and then select the following if required:
 - ◆ (Optional) Select **Shutdown source after cutover** to shutdown the source after cutover.
 - ◆ (Optional) Select **Shutdown target after cutover** to shutdown the target after cutover.
- 4 Click **Execute**.

4.9 Targets

For any migration job, you must have a source workload that you want to migrate and a target to which you can migrate the source workload. The Targets page displays all the available targets and lets you add a new target.

To add target:

- 1 In the PlateSpin Migrate Web Interface, click **Targets > Add Target**.
- 2 Select one of the following VMware targets type:
 - ◆ Microsoft Azure Location
 - ◆ VMware DRS Cluster
 - ◆ VMware ESX Server
- 3 Depending on the type of VMware targets you selected in the previous step, specify the following. For information about the fields, see “PlateSpin Migrate Best Practices for Migrating Windows Workloads to Microsoft Azure” on the [PlateSpin Migrate Resources Webpage](#).

Option	Description
Subscription Id	Specify the subscription id for your Microsoft Azure account. This is applicable only for Microsoft Azure Location target type.
Client Id	Specify your Azure Client id required to allow the PlateSpin software access your Azure subscription. This is applicable only for Microsoft Azure Location target type.
Hostname or IP	Specify the hostname or IP address of the target. This is applicable for all target types other than Microsoft Azure Location.
Cluster Name	Specify the name of the DRS cluster. This is applicable only for VMware DRS Cluster.
Username and Password	Specify administrator-level credentials for accessing the target host
Location Name	Select the location for the Microsoft Azure target. This is applicable only for Microsoft Azure Location target type.

- 4 Click **Add**.

4.10 Tasks

The Tasks panel displays the most recent tasks, the most recent events, and the upcoming events.

Events are logged whenever some action related to the system or the workload occurs. For example, an event could be the addition of a new workload, the replication of a workload starting or failing, or the detection of the failure of a migrated workload. Some events also email automatic notifications if SMTP is configured. For more information, see [Section 4.12.3, “Setting Up Automatic Email Notifications of Events and Reports,” on page 112](#).

Tasks are special operations tied to events that require user intervention. For example, upon completion of a Test Cutover operation, the system generates an event associated with two tasks: **Mark Test as Success** and **Mark Test as Failure**. When you click either of the tasks, the Test Cutover operation is canceled and a corresponding event is logged.

The Tasks and Events panel on the dashboard displays a maximum of three entries. To see all tasks or to see past and upcoming events, click **View All** in the appropriate section.

4.11 Generating Reports Using the PlateSpin Migrate Web Interface

You can use the Web Interface to generate various reports:

- ♦ [Section 4.11.1, “Generating Workload and Workload Migration Reports,” on page 109](#)
- ♦ [Section 4.11.2, “Generating Licensing Report,” on page 109](#)

4.11.1 Generating Workload and Workload Migration Reports

You can generate reports that provide analytical insight into your workload migration contracts over time.

The following report types are supported:

- ♦ **Workload Migration:** Reports replication events for all workloads over a selectable time window.
- ♦ **Migration History:** Reports replication size, time, and transfer speed per selectable workload over a selectable time window.
- ♦ **Replication Statistics:** Reports the dynamics of full and incremental replications that can be summarized by **Average**, **Most Recent**, **Sum**, and **Peak** perspectives.
- ♦ **Current Migration Status:** Displays the migration status such last test cutover, last replication date, and the test age statistics.
- ♦ **Events:** Reports system events for all workloads over a selectable time window.
- ♦ **Scheduled Events:** Reports only upcoming workload migration events.
- ♦ **Resource Usage:** Displays the resources configured for the target workload.

To generate a report:

- 1 In your PlateSpin Migrate Interface, click **Reports**.
A list of the report types is displayed.
- 2 Click the name of the required report type.
- 3 Select one or more workloads for which you want to create the report.
- 4 Configure the time period for which you want to view the report.
- 5 Do one of the following:
 - ♦ Click **Printable View** to print the report.
 - ♦ Click **Export to XML** to export the report to an XML format.

4.11.2 Generating Licensing Report

Use the Web Interface to generate a licensing report that is used by Technical Support to troubleshoot licensing issues.

To generate a licensing report:

- 1 In your PlateSpin Migrate Web Interface, do one of the following:
 - ♦ Click **Settings > Licenses**.
 - ♦ In the License Summary panel on the dashboard, click **View All**
 - 2 Click **Generate Licensing Report**.
 - 3 Right-click **View Licensing Report** and save the report to a location on your computer.
- This report is used by Technical Support to troubleshoot licensing issues.

4.12 Configuring the Settings For the Web Interface

The Settings page lets you perform various configurations such as creating licenses, creating workload tags, creating notification settings, and setting permissions. For more information about the various configuration options, review the following sections:

- ♦ [Section 4.12.1, “Using Tags to Help Sort Workloads,” on page 110](#)
- ♦ [Section 4.12.2, “Managing PlateSpin Migrate Security Groups and Workload Permissions,” on page 111](#)
- ♦ [Section 4.12.3, “Setting Up Automatic Email Notifications of Events and Reports,” on page 112](#)
- ♦ [Section 4.12.4, “Configuring Product Licenses,” on page 114](#)

4.12.1 Using Tags to Help Sort Workloads

In the PlateSpin Migrate Web Interface, the Workloads page might display a long list of workloads. Sorting through these workloads to manage operations for similar workloads can be time-consuming. To overcome this issue, you can create a tag and attach the tag to all similar workloads. Tags enable you to group similar workloads so that you can easily manage operations on these workloads.

The Workload Tags page (**Settings > Workload Tags**) displays all the available tags. You can create new tags and edit or delete any existing tags.

To create workload tags:

- 1 In the PlateSpin Migrate Web Interface, click **Settings > Workload Tags > Create Workload Tag**.
- 2 On the Workload Tag Creation page, specify a tag name (25-character limit) and select a color to associate with the tag.
- 3 Click **Save** to list the tag on the Workload Tags page.

The Workloads page includes a **Tag** column where the single tag you associate with a workload is displayed. You can sort on this column to group similar workloads together. This enables you to run operations on the tagged workloads at the same time.

To associate a tag with a workload:

- 1 In the PlateSpin Migrate Web Interface, click **Settings > Workloads**.
- 2 In the workload list, select the workload you want to tag and click **Configure Migration**.
- 3 In the Tag section of the configuration page, select the tag name you want to associate with the workload and click **Save**.

To disassociate a tag from a workload:

- 1 In the PlateSpin Migrate Web Interface, click **Settings > Workloads**.
- 2 In the workload list, select the workload for which you want to remove the tag and click **Configure Migration**.
- 3 In the Tag section of the configuration page, select the empty string and click **Save**.

NOTE

- ♦ When you export a workload with a tag setting to a new server, the tag settings persist.
 - ♦ You cannot delete a tag if it is associated with any workload in the list.
-

4.12.2 Managing PlateSpin Migrate Security Groups and Workload Permissions

PlateSpin Migrate provides a granular application-level access mechanism that allows specific users to carry out specific workload migration tasks on specified workloads. This is accomplished by setting up *security groups*.

- 1 Assign a PlateSpin Migrate user a Workload Migration Role whose permissions best suit that role in your organization. For more information about the workload migration roles, [Section 2.2.1, “PlateSpin Migrate Roles,” on page 32](#).
- 2 In the PlateSpin Migrate Web Interface, click **Settings > Permissions**.
- 3 On the Security Groups page, click **Create Security Group**.
- 4 In the **Security Group Name** field, specify a name for the security group.
- 5 (Optional) Click **Add Users** to select the users you want to grant access to this security group and click **OK**.

A PlateSpin Migrate user you recently added to the PlateSpin Server host might not immediately list in the user interface. To list such newly added users, click **Refresh User Accounts**.

- 6 (Optional) Click **Assign Workloads**, select the workloads you want to include in this group and click **OK**.

Only the users of this security group have access to these workloads.

- 7 Click **Create** to add the new group to the security groups list on the Security Groups page.

NOTE: The default users created during Migrate installation are added to every security group you create, by default.

- 8 (Optional) To edit a security group, click the security group name on the Security Groups page and edit the group information as required.

To delete a security group, click **Delete** next to the security group name you want to delete.

4.12.3 Setting Up Automatic Email Notifications of Events and Reports

You can configure PlateSpin Migrate to automatically send notifications of events and replication reports to specified email addresses. This functionality requires that you first specify a valid SMTP server for PlateSpin Migrate to use.

- ◆ [“SMTP Configuration” on page 112](#)
- ◆ [“Setting Up Automatic Event Notifications by Email” on page 112](#)
- ◆ [“Setting Up Automatic Replication Reports by Email” on page 113](#)

SMTP Configuration

- 1 In the PlateSpin Migrate Web Interface, click **Settings > SMTP**.
- 2 Specify the following:
 - ◆ **SMTP Server Address:** The address of the SMTP server.
 - ◆ **Port:** The port at which the SMTP server is listening. By default, it is 25.
 - ◆ **Reply Address:** The address from which you want to send email event and progress notifications.
 - ◆ **Username and Password:** Provide valid credentials if your SMTP server requires authentication.
- 3 Click **Save**.

Setting Up Automatic Event Notifications by Email

To set up automatic event notifications:

- 1 Configure an SMTP server for PlateSpin Migrate to use. See [“SMTP Configuration” on page 112](#).
- 2 In the PlateSpin Migrate Web Interface, click **Settings > General Notification Settings**.
- 3 Select the **Enable Notifications** option.
- 4 Click **Edit Recipients**, specify the required email addresses separated by commas and click **OK**.
- 5 Click **Save**.

To delete an email address, click **Remove** next to the address that you want to delete.

The following event types triggers email notifications if notification is configured. The events are always added to the System Application Event Log according to the log entry types such as Warning, Error, and Information.

Event Types	Remarks
Log Entry Type: Warning	
IncrementalReplicationMissed	Generates when any of the following applies: <ul style="list-style-type: none"> ◆ A replication is manually paused when a scheduled incremental replication is due. ◆ The system attempts to carry out a scheduled incremental replication when a manually-triggered replication is in progress. ◆ The system determines that the target has insufficient free disk space.
FullReplicationMissed	Similar to IncrementalReplicationMissed event.
WorkloadOfflineDetected	Generated when the system detects that a previously online workload is now offline. Applies to workloads whose migration state is not Paused .
Log Entry Type: Error	
FailoverFailed	
FullReplicationFailed	
IncrementalReplicationFailed	
PrepareFailoverFailed	
Log Entry Type: Information	
FailoverCompleted	
FullReplicationCompleted	
IncrementalReplicationCompleted	
PrepareFailoverCompleted	
TestFailoverCompleted	Generates upon manually marking a Test Cutover operation a success or a failure.
WorkloadOnlineDetected	Generates when the system detects that a previously offline workload is now online. Applies to workloads whose migration state is not Paused .

NOTE: Although event log entries have unique IDs, the IDs are not guaranteed to remain the same in future releases.

Setting Up Automatic Replication Reports by Email

- 1 Set up an SMTP server for PlateSpin Migrate to use. See [“SMTP Configuration” on page 112](#).
- 2 In the PlateSpin Migrate Web Interface, click **Settings > Email > Report Notification Settings**.

- 3 Select the **Enable Report Notifications Reports** option.
- 4 In the **Report Recurrence** section, click **Edit** and specify the required recurrence pattern for the reports.
- 5 In the **Recipients** section, click **Edit Recipients** to specify the required email addresses separated by commas and click **OK**.
- 6 (Optional) In the **Migrate Access URL** section, specify a non-default URL for your PlateSpin Server. For example, if your PlateSpin Server host has more than one NIC or is located behind a NAT server. This URL affects the title of the report and the functionality of accessing relevant content on the server through hyperlinks within emailed reports.
- 7 Click **Save**.

For information on other types of reports that you can generate and view on demand, see [“Generating Reports Using the PlateSpin Migrate Web Interface” on page 109](#).

4.12.4 Configuring Product Licenses

To configure product licenses, click **Settings > Licenses**. The Licenses page displays information about the existing licenses, allows you to add new licenses and delete existing licenses, and generate the licensing report.

For PlateSpin Migrate product licensing, you must have a license activation code. If you do not have a license activation code, request one through the [Customer Center](http://www.netiq.com/customercenter/) (<http://www.netiq.com/customercenter/>). A license activation code will be emailed to you.

NOTE: If you are an existing PlateSpin customer and you do not have a Customer Center account, you must first create an account using the same email address as specified in your purchase order. See [Create Account](https://www.netiq.com/selfreg/jsp/createAccount.jsp) (<https://www.netiq.com/selfreg/jsp/createAccount.jsp>).

To activate your product license, you have the following two options:

- ♦ [“Online License Activation” on page 114](#)
- ♦ [“Offline License Activation” on page 115](#)

Online License Activation

Online activation requires that your PlateSpin Migrate Web Interface has Internet access.

NOTE: HTTP proxies might cause failures during online activation. Offline activation is recommended for users in environments that use HTTP proxy.

To set up online license activation:

- 1 In the PlateSpin Migrate Web Interface, click **Settings > Licenses**, then click **Add license**.
- 2 Click **Online Activation**.
- 3 Specify the email address that you provided when you placed your order and the activation code you received, then click **Activate**.

The system obtains the required license over the Internet and activates the product.

Offline License Activation

For offline activation, you must first use a computer that has Internet access to obtain a PlateSpin Migrate license key.

- 1 In the PlateSpin Migrate Web Interface, click **Settings > Licenses**, then click **Add license**.
- 2 Click **Offline Activation** and copy the hardware ID displayed in the interface.
- 3 Use a web browser on a computer that has Internet access to navigate to the [PlateSpin Product Activation website \(http://www.platespin.com/productactivation/ActivateOrder.aspx\)](http://www.platespin.com/productactivation/ActivateOrder.aspx). Log in with your Customer Center user name and password.
- 4 Open the [PlateSpin Activate Order page](#) to generate a license key file. You need the following information:
 - ◆ activation code that you received
 - ◆ email address that you provided when you placed your order
 - ◆ hardware ID that you copied in [Step 2](#)
- 5 Save the generated license key file, transfer it to the product host that does not have Internet connectivity, and use it to activate the product.
- 6 In the PlateSpin Migrate Web Interface on the License Activation page, browse to the location of the license key file, then click **Activate**.

The license key file is saved and the product is activated based on this file.

4.13 Configuring the Refresh Rates for the Web Interface

Several pages in the Web Interface have configurable refresh intervals, as shown in [Table 4-1](#). You can modify the interval setting to meet the needs of your PlateSpin environment.

Table 4-1 Web Interface Default Refresh Intervals

Web Interface Parameter	Default Refresh Interval (in Seconds)
DashboardUpdateIntervalSeconds	60
WorkloadsUpdateIntervalSeconds	60
WorkloadTargetsUpdateIntervalSeconds	30
WorkloadDetailsUpdateIntervalSeconds	15
TasksUpdateIntervalSeconds	15

- 1 Open the following file in a text editor:

```
\Program Files\PlateSpin Migrate Server\Platespin Forge\web\web.config
```

- 2 Modify the value for any of the following interval settings as appropriate for your PlateSpin environment:

```
<add key="DashboardUpdateIntervalSeconds" value="60" />
<add key="WorkloadsUpdateIntervalSeconds" value="60" />
<add key="WorkloadTargetsUpdateIntervalSeconds" value="30" />
<add key="WorkloadDetailsUpdateIntervalSeconds" value="15" />
<add key="TasksUpdateIntervalSeconds" value="15" />
```

3 Save the file.

The new settings apply in your next Web Interface session. It is not necessary to restart the PlateSpin Server service or server.

5 Workload Portability with Peer-to-Peer Workload Migrations

Peer-to-peer refers to workload portability operations that result in a booted or bootable workload on physical hardware or a virtual machine. P2P, P2V, V2V, and V2V migrations are all peer-to-peer.

This section provides information about peer-to-peer workload portability jobs and best practices.

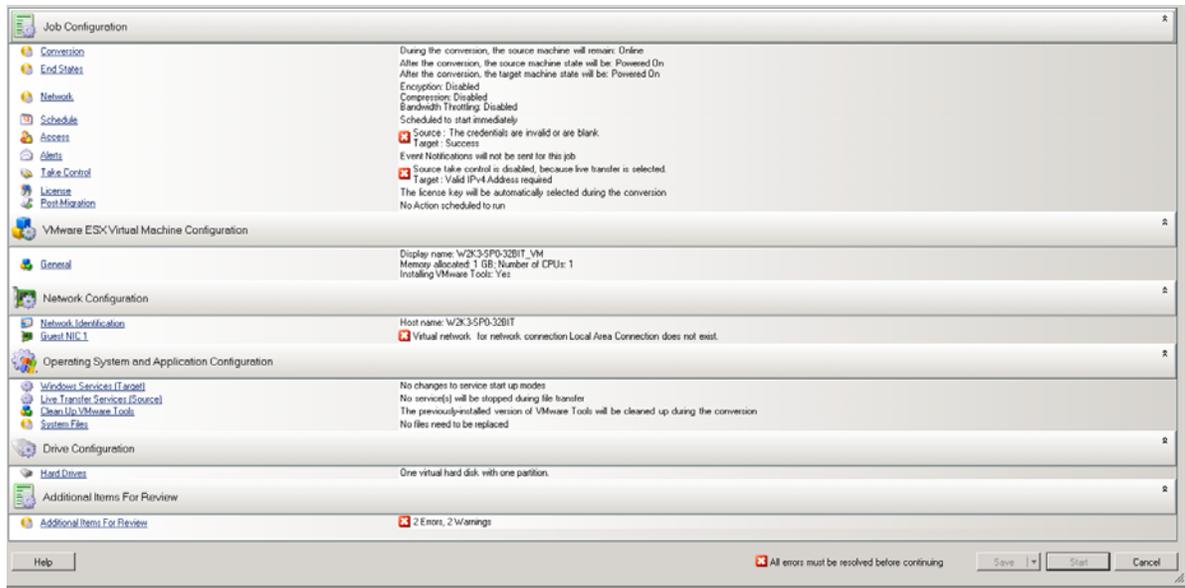
- ◆ [Section 5.1, “Migrating a Workload to a VM Host \(P2V, V2V\),” on page 117](#)
- ◆ [Section 5.2, “Converting a Workload to a Physical Machine \(P2P, V2P\),” on page 118](#)
- ◆ [Section 5.3, “Synchronizing Workloads with Server Sync,” on page 119](#)
- ◆ [Section 5.4, “Semi-Automated Workload Virtualization Using the X2P Workflow,” on page 122](#)
- ◆ [Section 5.5, “Migrating Windows Clusters,” on page 131](#)

5.1 Migrating a Workload to a VM Host (P2V, V2V)

To initiate a peer-to-peer workload virtualization migration job:

- 1 Discover or refresh your source workload and your target VM host.
See [“Discovering Source Workloads and Migration Targets” on page 64](#).
- 2 Start a migration job.
See [“Setting a Migration Job” on page 83](#).
- 3 Configure the required parameters of the job.
See [“Essentials of Workload Migration Jobs” on page 145](#).
- 4 Run the job.

Figure 5-1 Peer-to-Peer Migration Job Window



5.2 Converting a Workload to a Physical Machine (P2P, V2P)

To initiate a peer-to-peer workload migration to a physical machine:

- 1 Discover your source workload.
See [“Discovering Source Workloads and Migration Targets”](#) on page 64.
- 2 Discover and register your target physical machine with PlateSpin Migrate.
See [“Discovering Target Physical Machines”](#) on page 67.
- 3 (Recommended) Use PlateSpin Analyzer to ensure that:
 - ◆ Your source operating system and hardware are supported by PlateSpin Migrate.
 - ◆ PlateSpin Migrate’s X2P device driver library contains device drivers that your target requires for the operating system being ported.See [“Analyzing Machines with PlateSpin Analyzer”](#) on page 71.
- 4 (Conditional) If drivers for the physical target are missing, upload the required drivers to the PlateSpin Migrate’s X2P device driver library.
See [“Working with Device Drivers”](#) on page 72.
- 5 Start a migration job.
See [“Setting a Migration Job”](#) on page 83.
- 6 Configure the required parameters of the job.
See [“Essentials of Workload Migration Jobs”](#) on page 145.
- 7 Run the job.

5.2.1 Best Practices (X2P)

- ◆ When you are migrating a workload from one vendor to a target hardware infrastructure from another vendor (for example, from HP to Dell), or if your source is a virtual machine, make sure that you disable vendor-specific or VM-specific services during the transfer. For example, the HP Insight service and the VMware Tools service.

See [“Handling Operating System Services and Applications” on page 164](#).

- ◆ When you are using the Offline transfer method for P2P and V2P migrations, make sure that you select the appropriate Full Duplex speed that matches your network Full Duplex mode.

See [“Setting Up Migration Networking” on page 151](#).

- ◆ Ensure that vendor partitions are not being copied from the source.

See [“Handling a Workload’s Storage Media and Volumes” on page 169](#).

5.3 Synchronizing Workloads with Server Sync

The Server Sync feature enables you to reduce the scope of data that is transferred from your source to your target to just data that is different between a source and a target, effectively synchronizing their volume contents.

For example, when setting up a job for a workload portability operation, you can choose to update an existing physical or virtual machine to match the state of your source workload without transferring volume data in its entirety. PlateSpin Migrate compares the target physical or virtual workload with the selected source and transfers only data that is different between the two, overwriting files on the target with those on the source workload.

Server Sync is useful in situations where the size of volume data or network conditions are prohibitive for a direct source-to-target virtualization over the network.

- ◆ [Section 5.3.1, “Server Sync to a Virtual Target,” on page 119](#)
- ◆ [Section 5.3.2, “Server Sync to a Physical Target,” on page 122](#)

5.3.1 Server Sync to a Virtual Target

- 1 Discover your source workload.

See [“Discovering Source Workloads and Migration Targets” on page 64](#).

- 2 Create a target virtual machine by using one of the following methods:

- ◆ Do an initial migration of your workload to a virtual machine. See [“Migrating a Workload to a VM Host \(P2V, V2V\)” on page 117](#).

- OR -

- ◆ Using your virtualization platform’s native interface, manually install a virtual machine with the same operating system profile as that of your source.

NOTE: When you are creating a virtual target for Server Sync, you should also manually install the appropriate virtualization enhancement tools, such as VMware Tools or XenServer Tools.

- OR -

- ♦ (Windows only) Capture your workload to a PlateSpin Image, and deploy it to a virtual machine on your virtualization platform. See “[Capturing a Workload to a PlateSpin Image](#)” on page 139 and “[Deploying a PlateSpin Image](#)” on page 140.
- 3 (Conditional; Hyper-V changes to Hyper-V 2008 R2 only) Because the Server Sync option is disabled for a Hyper-V VM, it is necessary to use the following steps, as documented in [KB 7010748](#) (<https://www.netiq.com/support/kb/doc.php?id=7010748>):

NOTE: Hyper-V 2012 and Hyper-V 2012 R2 automatic server sync is available.

3a After booting the target VM with the LRD ISO ((`bootofx.x2p.iso`)) wait for the Migrate Server URL prompt, Then press Alt+F7 to launch the debug console.

3b From the debug console, run the following command to determine which devices are /, /boot and swap:

```
fdisk -l
```

3c Using the information obtained from the debug console, mount the appropriate devices as under:

```
mount /dev/%root device% /
mount /dev/%boot device% /boot
```

3d Press Alt+F1 to switch to the server command line.

3e At the command line, provide the required information at each individual prompt:

- ♦ **PlateSpin Server:** Use the following format:

```
http://server_host/platespinmigrate
```

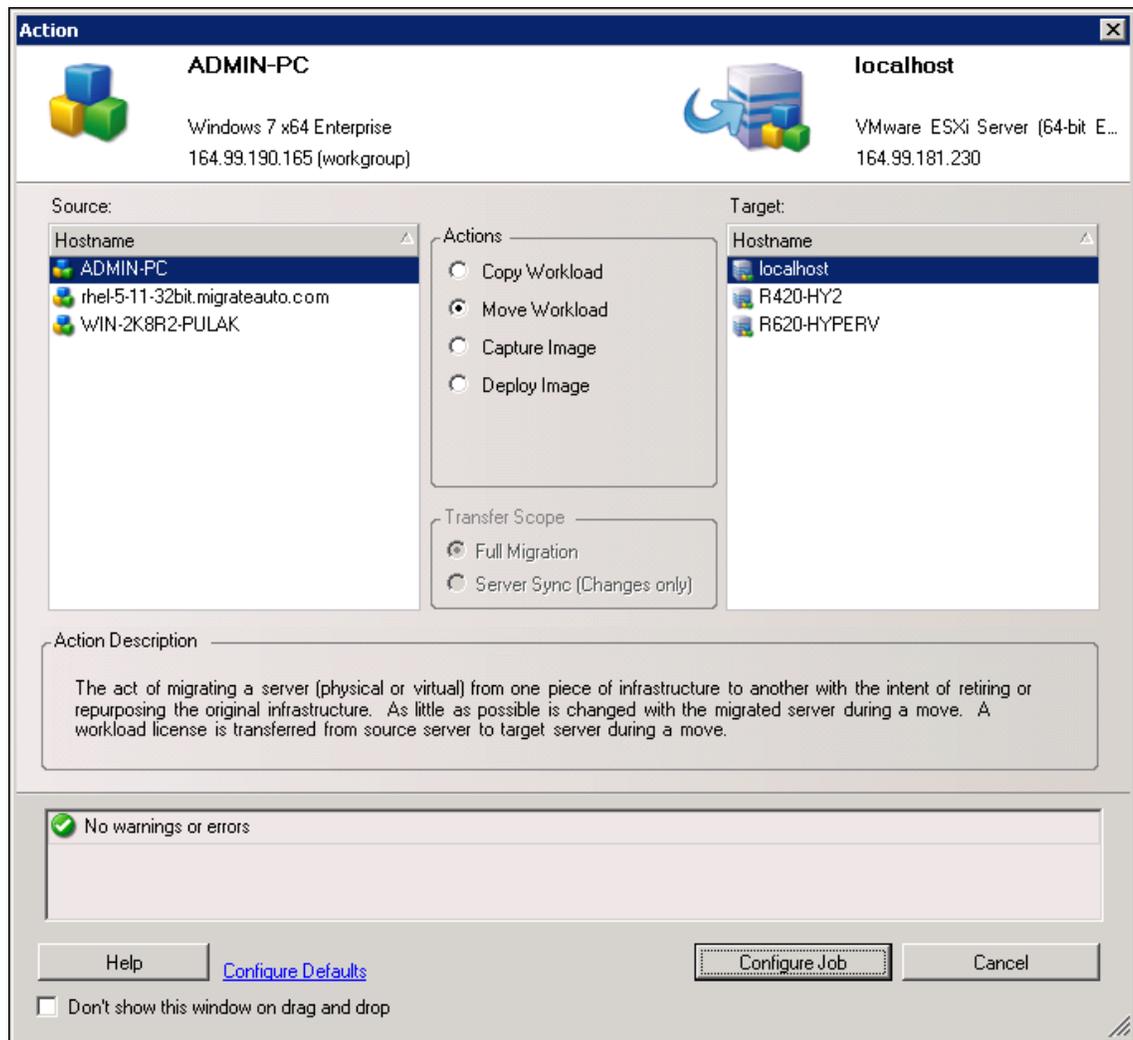
Replace *server_host* with the actual PlateSpin Server host’s name or IP address.

- ♦ **Credentials (User Name/Password):** Enter the name of an admin-level user on the PlateSpin Server host, including the domain or machine name. For example: *domain\username*, or *localhostAdministrator*. Provide a valid password for the specified user.
- ♦ **Network Card:** Select the network card that is active, then either enter a temporary static IP address for this card or press the Enter key to use a DHCP server.
- ♦ **Temporary hostname:** Provide a temporary VM name for PlateSpin Migrate Client to use to list the newly registered VM. The workload’s target hostname you select in the migration job overwrites this name.
- ♦ **SSL encryption:** If your PlateSpin Migrate is installed on a host with SSL encryption enabled, enter *Yes*. If not, enter *No*.
- ♦ **PlateSpin Migrate Network:** Unless you have defined your own PlateSpin Migrate Network in PlateSpin Migrate Client, press the Enter key. If you are working with a non-default PlateSpin Migrate Network, type its name, then press the Enter key.

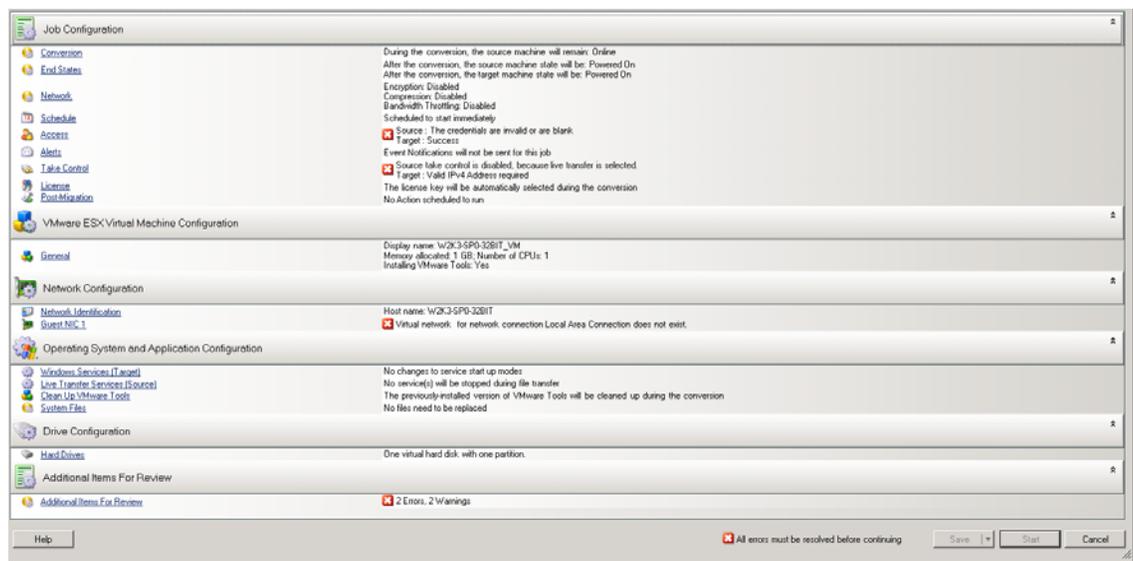
A controller on your target virtual machine communicates with PlateSpin Server and registers the virtual machine as a physical target for a migration job.

- 4 In the Servers view, drag your source workload and drop it on the required target (Server Sync target or discovered physical machine under control).

The system validates the selected source and target and, if it detects matching operating systems on them, provides you with two **Transfer Scope** options, **Full Migration** and **Server Sync**:



5 Select the **Server Sync** option, then click **Configure Job**.



- 6 In the job configuration window, specify the parameters of the job as dictated by the purpose of the operation, address any warnings and errors, and make sure that you map the required volumes on the source to those on the target (see [“Volume Configuration in Server Sync” on page 177](#)).

When you have finished, click **Start**.

PlateSpin Migrate starts the job and lists it in the Jobs view.

5.3.2 Server Sync to a Physical Target

- 1 Discover your source workload.

See [“Discovering Source Workloads and Migration Targets” on page 64](#).

- 2 Discover your physical target by using the appropriate boot ISO boot image.

See [“Discovering Target Physical Machines” on page 67](#).

- 3 In the Servers view, drag your source workload and drop it on the required target (Server Sync target or discovered physical machine under control).

The system validates the selected source and target and, if it detects matching operating systems on them, it provides you with two **Transfer Scope** options, **Full Migration** and **Server Sync**, similar to the [Server Sync to a Virtual Target](#) scenario (see [Step 4](#)).

- 4 Select the **Server Sync** option, then click **Configure Job**.

- 5 In the job configuration window, specify the parameters of the job as dictated by the purpose of the operation, address any warnings and errors, and make sure that you map the required volumes on the source to those on the target.

- 6 When you have finished, click **Start**.

PlateSpin Migrate starts the job and lists it in the Jobs view.

5.4 Semi-Automated Workload Virtualization Using the X2P Workflow

Semi-automated refers to the process of workload virtualization in which a workload's target infrastructure is a virtual machine that PlateSpin Migrate regards as a physical machine. This applies to the following virtualization platforms:

- ♦ Windows Server 2008 R2 Hyper-V
- ♦ Citrix XenServer 6.1
- ♦ RHEL KVM 7.0
- ♦ SUSE Linux Enterprise Server (SUSE) 11 SP3 Xen (fully virtualized guest)

To migrate a workload to one of these platforms:

- 1 Use the native interface of the required virtualization platform to create a virtual machine whose operating system profile matches that of your source workload.
- 2 Begin booting the newly created virtual machine by using the appropriate PlateSpin boot ISO image, load the appropriate driver, if needed, then continue the boot process.

This discovers and registers the target virtual machine as a PlateSpin Migrate physical machine target. This is similar to the process described in [“Discovering Target Physical Machines” on page 67](#). However, instead of burning the ISO image to a CD to boot physical hardware, you save the image on the VM host and use it to boot the target VM.

- 3 Use the PlateSpin Migrate Client to create and execute an X2P migration job.
- 4 Upon completion of the migration job, install virtualization enhancement software specific to the target virtualization platform.

The following topics provide additional notes about migrating workloads to the specific virtualization platform involved.

- ♦ [Section 5.4.1, “Migrating a Workload to Microsoft Hyper-V,” on page 123](#)
- ♦ [Section 5.4.2, “Migrating a Workload to Citrix XenServer,” on page 125](#)
- ♦ [Section 5.4.3, “Migrating a Workload to RHEL KVM,” on page 127](#)
- ♦ [Section 5.4.4, “Migrating a Workload to the SLES 11 SP3 Xen Hypervisor,” on page 129](#)

5.4.1 Migrating a Workload to Microsoft Hyper-V

You can use Microsoft Hyper-V as the target virtualization platform in a semi-automated workload virtualization.

NOTE: Hyper-V 2012 and Hyper-V 2012 R2 automatic workload virtualization is also available.

Prerequisites:

- ♦ Your source workload is supported by PlateSpin Migrate and Hyper-V.
See [“Supported Target Virtualization Platforms” on page 20](#).

Additional information:

- ♦ [Microsoft Hyper-V Getting Started Guide \(http://technet.microsoft.com/en-us/library/cc732470.aspx\)](http://technet.microsoft.com/en-us/library/cc732470.aspx)
- ♦ [Virtualization with Hyper-V \(http://www.microsoft.com/windowsserver2008/en/us/hyperv.aspx\)](http://www.microsoft.com/windowsserver2008/en/us/hyperv.aspx)

This section includes the following topics:

- ♦ [“Downloading and Saving the Boot ISO Image \(Hyper-V\)” on page 123](#)
- ♦ [“Creating and Configuring the Target Virtual Machine \(Hyper-V\)” on page 124](#)
- ♦ [“Registering the Virtual Machine with PlateSpin Server \(Hyper-V\)” on page 124](#)
- ♦ [“Migrating Your Source Workload to the Target Virtual Machine \(Hyper-V\)” on page 125](#)
- ♦ [“Post-Migration Steps \(Hyper-V\)” on page 125](#)

Downloading and Saving the Boot ISO Image (Hyper-V)

- 1 Download the appropriate boot ISO image for your workload.
- 2 (Optional) If you want to perform an unattended registration of your target, update the ISO image with the required registration parameters.
See [KB Article 7920832 \(https://www.netiq.com/support/kb/doc.php?id=7920832\)](https://www.netiq.com/support/kb/doc.php?id=7920832).
- 3 Save the ISO image in a location that Hyper-V server can access. For example: `c:\temp`.
This ensures that the boot ISO image is available to the target VM as a bootable CD-ROM image.

Creating and Configuring the Target Virtual Machine (Hyper-V)

- 1 Use the Hyper-V Manager's New Virtual Machine Wizard to create a new virtual machine with the following settings:
 - ♦ **Name and Location:** Specify a name for your new target and accept the default location.
 - ♦ **Assign Memory:** Assign at least 384 MB of RAM to the VM.
 - ♦ **Connect Virtual Hard Disk:** Ensure that the virtual disk size is equal to or greater than the system disk of your source workload.
 - ♦ **Installation Options:** Configure the VM to boot from an ISO image file, and point the wizard to the downloaded boot ISO image.
 - ♦ **Summary:** Configure the VM to not start upon creation (deselect the **Start the virtual machine after it is created** option).
- 2 After creating the VM, remove the default NIC and replace it with a generic one, called *Legacy Network Adapter*.

This is required because the New Virtual Machine Wizard creates a NIC of a custom Microsoft type, which is currently unsupported by PlateSpin Migrate.
- 3 Connect the newly added NIC (*Legacy Network Adapter*) to the external virtual network.

Registering the Virtual Machine with PlateSpin Server (Hyper-V)

Complete these steps to register your target with PlateSpin Server when you are using an unmodified boot ISO image.

- 1 At the command line in the Hyper-V Virtual Machine Connection Console, provide the required information at each individual prompt:
 - ♦ **PlateSpin Server:** Use the following format:

```
http://server_host/platespinmigrate
```

Replace *server_host* with the actual PlateSpin Server host's name or IP address.
 - ♦ **Credentials (User Name/Password):** Enter the name of an admin-level user on the PlateSpin Server host, including the domain or machine name. For example: *domain\username*, or *localhost\Administrator*. Provide a valid password for the specified user.
 - ♦ **Network Card:** Select the network card that is active, then either enter a temporary static IP address for this card or press the Enter key to use a DHCP server.
 - ♦ **Temporary hostname:** Provide a temporary VM name for PlateSpin Migrate Client to use to list the newly registered VM. The workload's target hostname you select in the migration job overwrites this name.
 - ♦ **SSL encryption:** If your PlateSpin Migrate is installed on a host with SSL encryption enabled, enter *Yes*. If not, enter *No*.
 - ♦ **PlateSpin Migrate Network:** Unless you have defined your own PlateSpin Migrate Network in PlateSpin Migrate Client, press the Enter key. If you are working with a non-default PlateSpin Migrate Network, type its name, then press the Enter key.

A controller on your target virtual machine communicates with PlateSpin Server and registers the virtual machine as a physical target for a migration job.

NOTE: If registration fails with an authorization error, you might need to synchronize the clocks of the source and the target, modify the LAN Manager Authentication Level on the target, or both. See [Table F-1, “Common Issues and Solutions Related to Discovery Operations,” on page 231.](#)

Migrating Your Source Workload to the Target Virtual Machine (Hyper-V)

- 1 Use PlateSpin Migrate Client to start an X2P migration job with your source workload being the job's migration source and the target being the new VM on Hyper-V.
See [“Converting a Workload to a Physical Machine \(P2P, V2P\)” on page 118.](#)
- 2 Monitor the migration job in PlateSpin Migrate Client's Jobs view.
When the job reaches the **Configure Target Machine** step, the virtual machine's console returns to the boot prompt of the boot ISO image.
- 3 Shut down the virtual machine and reconfigure it to boot from disk rather than from the boot image.
- 4 Power on the virtual machine.
The migration job resumes, reboots the target, and completes the workload configuration.

Post-Migration Steps (Hyper-V)

Install Hyper-V Integration Services (virtualization enhancement software). For more information, see your [Microsoft Hyper-V Getting Started Guide](#).

5.4.2 Migrating a Workload to Citrix XenServer

You can use Citrix XenServer as the target virtualization platform in a semi-automated workload virtualization.

Prerequisites:

- ◆ Your target is a fully virtualized (not paravirtualized) VM.
- ◆ Your source workload is supported by PlateSpin Migrate and Citrix XenServer.
See [“Supported Target Virtualization Platforms” on page 20.](#)

Additional information:

- ◆ [Citrix XenServer® 6.1.0 Administrator's Guide \(http://docs.vmd.citrix.com/XenServer/6.1.0/1.0/en_gb/reference.html\)](http://docs.vmd.citrix.com/XenServer/6.1.0/1.0/en_gb/reference.html)

This section includes the following topics:

- ◆ [“Downloading and Saving the Boot ISO Image \(Citrix XenServer\)” on page 126](#)
- ◆ [“Creating and Configuring the Target Virtual Machine \(Citrix XenServer\)” on page 126](#)
- ◆ [“Registering the Virtual Machine with PlateSpin Server \(Citrix XenServer\)” on page 126](#)
- ◆ [“Migrating Your Source Workload to the Target Virtual Machine \(Citrix XenServer\)” on page 127](#)

Downloading and Saving the Boot ISO Image (Citrix XenServer)

- 1 Download the appropriate PlateSpin boot ISO image.
- 2 (Optional) If you want to perform an unattended registration of your target, update the ISO image with the required registration parameters.

See [KB Article 7013485 \(https://www.netiq.com/support/kb/doc.php?id=7013485\)](https://www.netiq.com/support/kb/doc.php?id=7013485).

- 3 Save the downloaded image file in the following directory:

```
/var/lib/xen/images
```

Creating and Configuring the Target Virtual Machine (Citrix XenServer)

- 1 On Citrix XenServer, use the Virtual Machine Manager Wizard or the Create Virtual Machines program shortcut to create a new virtual machine.

Ensure that the new virtual machine is created with the following settings:

- ♦ **Virtualization method:** Fully virtualized.
 - ♦ **Memory:** Assign at least 384 MB of RAM to the VM. This ensures that the VM has sufficient resources during the migration and improves transfer speed. If the virtual machine requires less memory after the migration, reduce the assigned memory after the migration completes.
 - ♦ **Disks:** Assign one or more disks, depending on the requirements of your source workload. The storage can be either a raw SAN LUN or a virtual disk. Also, create a Virtual CD-ROM assigned to the downloaded boot ISO image.
 - ♦ **Operating System:** Must match the operating system profile of your source workload.
- 2 Ensure that the VM is configured to restart on reboot by exporting the VM's settings from the xend database to a text file and making sure that the `on_reboot` parameter is set to `restart`. If not, shut down the VM, update the settings, and reimport them into the xend database.

For detailed instructions, see the [XenServer 6.1.0 Virtual Machine User's Guide \(http://support.citrix.com/article/CTX134587\)](http://support.citrix.com/article/CTX134587).

- 3 From the Virtual Machine Manager, launch the virtual machine console and monitor the boot process.

When the virtual machine completes the boot process, it prompts you for parameters that control the registration of the machine and its profile with PlateSpin Migrate. If you are using the unattended registration process, the required parameters are read from an answer file.

Registering the Virtual Machine with PlateSpin Server (Citrix XenServer)

Complete these steps to register your target with PlateSpin Server when you are using an unmodified boot ISO image.

- 1 At the command line, provide the required information at each individual prompt:

- ♦ **PlateSpin Server:** Use the following format:

```
http://server_host/platespinmigrate
```

Replace `server_host` with the actual PlateSpin Server host's name or IP address.

- ♦ **Credentials (User Name/Password):** Enter the name of an admin-level user on the PlateSpin Server host, including the domain or machine name. For example: *domain\username*, or *localhost\Administrator*. Provide a valid password for the specified user.
 - ♦ **Network Card:** Select the network card that is active, then either enter a temporary static IP address for this card or press the Enter key to use a DHCP server.
 - ♦ **Temporary hostname:** Provide a temporary VM name for PlateSpin Migrate Client to use to list the newly registered VM. The workload's target hostname you select in the migration job overwrites this name.
 - ♦ **SSL encryption:** If your PlateSpin Migrate is installed on a host with SSL encryption enabled, enter *Yes*. If not, enter *No*.
 - ♦ **PlateSpin Migrate Network:** Unless you have defined your own PlateSpin Migrate Network in PlateSpin Migrate Client, press the Enter key. If you are working with a non-default PlateSpin Migrate Network, type its name, then press the Enter key.
- A controller on your target virtual machine communicates with PlateSpin Server and registers the virtual machine as a physical target for a migration job.

Migrating Your Source Workload to the Target Virtual Machine (Citrix XenServer)

- 1 Use PlateSpin Migrate Client to start an X2P migration job with your source workload being the job's migration source and the target being the new VM on the Citrix XenServer hypervisor.
See [“Converting a Workload to a Physical Machine \(P2P, V2P\)” on page 118](#).
- 2 Monitor the migration job in the PlateSpin Migrate Client's Jobs view.
When the job reaches the **Configure Target Machine** step, the virtual machine's console returns to the boot prompt of the boot ISO image.
- 3 Shut down the virtual machine, reconfigure it to boot from disk rather than from the boot image, and deselect the **VS Tools Installed** option.
- 4 Power on the virtual machine.
The migration job resumes, reboots the target, and completes the workload configuration.

5.4.3 Migrating a Workload to RHEL KVM

You can use RHEL KVM as the target virtualization platform in a semi-automated workload virtualization.

Prerequisites:

- ♦ Your target is a fully virtualized (not paravirtualized) VM.
- ♦ Your source workload is supported by PlateSpin Migrate and RHEL KVM.
See [“Supported Target Virtualization Platforms” on page 20](#).

This section includes the following topics:

- ♦ [“Downloading and Saving the Boot ISO Image \(RHEL KVM\)” on page 128](#)
- ♦ [“Creating and Configuring the Target Virtual Machine \(RHEL KVM\)” on page 128](#)
- ♦ [“Registering the Virtual Machine with PlateSpin Server \(RHEL KVM\)” on page 128](#)
- ♦ [“Migrating Your Source Workload to the Target Virtual Machine \(RHEL KVM\)” on page 129](#)

Downloading and Saving the Boot ISO Image (RHEL KVM)

- 1 Download the appropriate PlateSpin boot ISO image to a location on your RHEL KVM machine.
- 2 (Optional) If you want to perform an unattended registration of your target, update the ISO image with the required registration parameters.

See [KB Article 7013485](https://www.netiq.com/support/kb/doc.php?id=7013485) (<https://www.netiq.com/support/kb/doc.php?id=7013485>).

Creating and Configuring the Target Virtual Machine (RHEL KVM)

- 1 On RHEL KVM, use the Virtual Machine Manager Wizard or the Create Virtual Machines program shortcut to create a new virtual machine.

Ensure that the new virtual machine is created with the following settings:

- ◆ **Virtualization method:** Fully virtualized.
 - ◆ **Memory:** Assign at least 384 MB of RAM to the VM. This ensures that the VM has sufficient resources during the migration and improves transfer speed. If the virtual machine requires less memory after the migration, reduce the assigned memory after the migration completes.
 - ◆ **Disks:** Assign one or more disks, depending on the requirements of your source workload. The storage can be either a raw SAN LUN or a virtual disk. Also, create a Virtual CD-ROM assigned to the downloaded boot ISO image.
 - ◆ **Operating System:** Must match the operating system profile of your source workload.
- 2 Ensure that the VM is configured to restart on reboot.
 - 3 From the Virtual Machine Manager, launch the virtual machine console and monitor the boot process.

When the virtual machine completes the boot process, it prompts you for parameters that control the registration of the machine and its profile with PlateSpin Migrate. If you are using the unattended registration process, the required parameters are read from an answer file.

Registering the Virtual Machine with PlateSpin Server (RHEL KVM)

Complete these steps to register your target with PlateSpin Server when you are using an unmodified boot ISO image.

- 1 At the command line, provide the required information at each individual prompt:
 - ◆ **PlateSpin Server:** Use the following format:

```
http://server_host/platespinmigrate
```

Replace *server_host* with the actual PlateSpin Server host's name or IP address.
 - ◆ **Credentials (User Name/Password):** Enter the name of an admin-level user on the PlateSpin Server host, including the domain or machine name. For example: *domain\username*, or *localhostAdministrator*. Provide a valid password for the specified user.
 - ◆ **Network Card:** Select the network card that is active, then either enter a temporary static IP address for this card or press the Enter key to use a DHCP server.
 - ◆ **Temporary hostname:** Provide a temporary VM name for PlateSpin Migrate Client to use to list the newly registered VM. The workload's target hostname you select in the migration job overwrites this name.
 - ◆ **SSL encryption:** If your PlateSpin Migrate is installed on a host with SSL encryption enabled, enter `Yes`. If not, enter `No`.

- ♦ **PlateSpin Migrate Network:** Unless you have defined your own PlateSpin Migrate Network in PlateSpin Migrate Client, press the Enter key. If you are working with a non-default PlateSpin Migrate Network, type its name, then press the Enter key.

A controller on your target virtual machine communicates with PlateSpin Server and registers the virtual machine as a physical target for a migration job.

Migrating Your Source Workload to the Target Virtual Machine (RHEL KVM)

- 1 Use PlateSpin Migrate Client to start an X2P migration job with your source workload being the job's migration source and the target being the new VM on the RHEL KVM hypervisor.
See [“Converting a Workload to a Physical Machine \(P2P, V2P\)” on page 118](#).
- 2 Monitor the migration job in the PlateSpin Migrate Client's Jobs view.
When the job reaches the **Configure Target Machine** step, the virtual machine's console returns to the boot prompt of the boot ISO image.
- 3 Shut down the virtual machine, reconfigure it to boot from disk rather than from the boot image.
- 4 Power on the virtual machine.
The migration job resumes, reboots the target, and completes the workload configuration.

5.4.4 Migrating a Workload to the SLES 11 SP3 Xen Hypervisor

You can use the Xen Hypervisor on SLES as the target virtualization platform in a semi-automated workload virtualization.

Prerequisites:

- ♦ Your source workload is supported by PlateSpin Migrate and the Xen hypervisor.
See [“Supported Target Virtualization Platforms” on page 20](#).

Additional information:

- ♦ [SUSE Linux Enterprise Server 11 SP3 Virtualization with Xen \(https://www.suse.com/documentation/sles11/singlehtml/book_xen/book_xen.html\)](https://www.suse.com/documentation/sles11/singlehtml/book_xen/book_xen.html)

This section includes the following topics:

- ♦ [“Downloading and Saving the Boot ISO Image \(Xen on SLES\)” on page 129](#)
- ♦ [“Creating and Configuring the Target Virtual Machine \(Xen on SLES\)” on page 130](#)
- ♦ [“Registering the Virtual Machine with PlateSpin Server \(Xen on SLES\)” on page 130](#)
- ♦ [“Migrating Your Source Workload to the Target Virtual Machine \(Xen on SLES\)” on page 131](#)
- ♦ [“Post-Migration Steps \(Xen on SLES\)” on page 131](#)

Downloading and Saving the Boot ISO Image (Xen on SLES)

- 1 Download the appropriate boot ISO image for your workload.
See [“Downloading the PlateSpin Boot ISO Images” on page 68](#).
- 2 (Optional) If you want to perform an unattended registration of your target, update the ISO image with the required registration parameters.
See [KB Article 7920832 \(https://www.netiq.com/support/kb/doc.php?id=7920832\)](https://www.netiq.com/support/kb/doc.php?id=7920832).

- 3 Save the downloaded image file in the following directory:

```
/var/lib/xen/images
```

Creating and Configuring the Target Virtual Machine (Xen on SLES)

- 1 On SLES 11, use the Virtual Machine Manager Wizard or the Create Virtual Machines program shortcut to create a new virtual machine.

Ensure that the new virtual machine is created with the following settings:

- ♦ **Virtualization method:** Fully virtualized.
 - ♦ **Memory:** Assign at least 384 MB of RAM to the VM. This ensures that the VM has sufficient resources during the migration and improves transfer speed. If the virtual machine requires less memory after the migration, reduce the assigned memory after the migration completes.
 - ♦ **Disks:** Assign one or more disks, depending on the requirements of your source workload. The storage can be either a raw SAN LUN or a virtual disk. Also, create a Virtual CD-ROM assigned to the downloaded boot ISO image.
 - ♦ **Operating System:** Must match the operating system profile of your source workload.
- 2 Ensure that the VM is configured to restart on reboot by exporting the VM's settings from the xend database to a text file and making sure that the `on_reboot` parameter is set to `restart`. If not, shut down the VM, update the settings, and reimport them into the xend database.

For detailed instructions, see your [SLES 11 documentation \(https://www.suse.com/documentation/sles11/\)](https://www.suse.com/documentation/sles11/).

- 3 From the Virtual Machine Manager, launch the virtual machine console and monitor the boot process.

When the virtual machine completes the boot process, it prompts you for parameters that control the registration of the machine and its profile with PlateSpin Migrate. If you are using the unattended registration process, the required parameters are read from an answer file.

Registering the Virtual Machine with PlateSpin Server (Xen on SLES)

Complete these steps to register your target with PlateSpin Server when you are using an unmodified boot ISO image. For information on how to modify the image for an unattended registration process, see [KB Article 7920832 \(https://www.netiq.com/support/kb/doc.php?id=7920832\)](https://www.netiq.com/support/kb/doc.php?id=7920832).

- 1 At the command line, provide the required information at each individual prompt:
 - ♦ **PlateSpin Server:** Use the following format:

```
http://server_host/platespinmigrate
```

Replace `server_host` with the actual PlateSpin Server host's name or IP address.
 - ♦ **Credentials (User Name/Password):** Enter the name of an admin-level user on the PlateSpin Server host, including the domain or machine name. For example: `domain\username`, or `localhostAdministrator`. Provide a valid password for the specified user.
 - ♦ **Network Card:** Select the network card that is active, then either enter a temporary static IP address for this card or press the Enter key to use a DHCP server.
 - ♦ **Temporary hostname:** Provide a temporary VM name for PlateSpin Migrate Client to use to list the newly registered VM. The workload's target hostname you select in the migration job overwrites this name.

- ♦ **SSL encryption:** If your PlateSpin Migrate is installed on a host with SSL encryption enabled, enter `Yes`. If not, enter `No`.
 - ♦ **PlateSpin Migrate Network:** Unless you have defined your own PlateSpin Migrate Network in PlateSpin Migrate Client, press the `Enter` key. If you are working with a non-default PlateSpin Migrate Network, type its name, then press the `Enter` key.
- A controller on your target virtual machine communicates with PlateSpin Server and registers the virtual machine as a physical target for a migration job.

Migrating Your Source Workload to the Target Virtual Machine (Xen on SLES)

- 1 Use PlateSpin Migrate Client to start an X2P migration job with your source workload being the job's migration source and the target being the new VM on the Xen hypervisor.
See [“Converting a Workload to a Physical Machine \(P2P, V2P\)”](#) on page 118.
- 2 Monitor the migration job in the PlateSpin Migrate Client's Jobs view.
When the job reaches the **Configure Target Machine** step, the virtual machine's console returns to the boot prompt of the boot ISO image.
- 3 Shut down the virtual machine, reconfigure it to boot from disk rather than from the boot image, and deselect the **VS Tools Installed** option.
- 4 Power on the virtual machine.
The migration job resumes, reboots the target, and completes the workload configuration.

Post-Migration Steps (Xen on SLES)

Install SUSE Drivers for Xen (virtualization enhancement software). For more information, see the following online document:

[SUSE Linux Enterprise Server 11 SP3 Virtualization with Xen \(https://www.suse.com/documentation/sles11/singlehtml/book_xen/book_xen.html\)](https://www.suse.com/documentation/sles11/singlehtml/book_xen/book_xen.html)

5.5 Migrating Windows Clusters

You can migrate a Microsoft Windows cluster's business services. The supported clustering technologies are:

- ♦ **Windows Server 2012 R2:** Server-based Microsoft Failover Cluster (*Node and Disk Majority Quorum* and *No Majority: Disk Only Quorum* models)
- ♦ **Windows Server 2008 R2:** Server-based Microsoft Failover Cluster (*Node and Disk Majority Quorum* and *No Majority: Disk Only Quorum* models)
- ♦ **Windows Server 2003 R2:** Server-based Windows Cluster Server (*Single-Quorum Device Cluster* model)

You can enable or disable Windows cluster discovery for your PlateSpin Migrate environment. See [Section 5.5.2, “Enabling or Disabling Windows Cluster Discovery,”](#) on page 134.

NOTE: The Windows cluster management software provides the failover and failback control for the resources running on its cluster nodes. This document refers to this action as a *cluster node failover* or a *cluster node failback*.

- ◆ [Section 5.5.1, “Planning Your Cluster Migration,” on page 132](#)
- ◆ [Section 5.5.2, “Enabling or Disabling Windows Cluster Discovery,” on page 134](#)
- ◆ [Section 5.5.3, “Adding Resource Name Search Values,” on page 134](#)
- ◆ [Section 5.5.4, “Setting the Quorum Arbitration Timeout,” on page 135](#)
- ◆ [Section 5.5.5, “Setting Local Volume Serial Numbers,” on page 135](#)
- ◆ [Section 5.5.6, “PlateSpin CutOver,” on page 135](#)
- ◆ [Section 5.5.7, “PlateSpin Cluster Migration,” on page 136](#)
- ◆ [Section 5.5.8, “Migrating Windows Clusters with the Web Interface,” on page 136](#)
- ◆ [Section 5.5.9, “Migrating Windows Clusters Using the Migrate Client,” on page 136](#)

5.5.1 Planning Your Cluster Migration

Migration of a cluster is achieved through incremental replications of changes on the active node streamed to a virtual one node cluster. Before you configure Windows clusters for migration, ensure that your environment meets the prerequisites and that you understand the conditions for migrating cluster workloads.

- ◆ [“Prerequisites” on page 132](#)
- ◆ [“Block-Based Transfer” on page 133](#)
- ◆ [“Cluster Node Failover during the First Full Replication” on page 133](#)
- ◆ [“Cluster Node Failover during Replication” on page 133](#)
- ◆ [“Cluster Node Failover between Replications” on page 133](#)

Prerequisites

The scope of support for cluster migration is subject to the following conditions:

- ◆ **Active node hostname or IP address:** You must specify the hostname or IP address of the cluster’s active node when you perform an Add Workload operation. Because of security changes made by Microsoft, Windows clusters can no longer be discovered by using the virtual cluster name (that is, the shared cluster IP address).
- ◆ **Active node discovery:** Ensure that the PlateSpin global configuration setting `DiscoverActiveNodeAsWindowsCluster` is set to `True` on the PlateSpin Server Configuration page. This is the default setting. See [Section 5.5.2, “Enabling or Disabling Windows Cluster Discovery,” on page 134](#).
- ◆ **Resource name search values:** You must specify search values to use that can help differentiate the name of the shared Cluster IP Address resource from the name of other IP address resources on the cluster. See [Section 5.5.3, “Adding Resource Name Search Values,” on page 134](#)
- ◆ **Resolvable hostname:** The PlateSpin Server must be able to resolve the hostname of each of the nodes in the cluster.

NOTE: The hostname must be resolvable by the IP address. That is, both hostname lookup and reverse lookup are required.

- ♦ **Quorum resource:** A cluster's quorum resource must be co-located on the node with the cluster's resource group (service) being migrated.
- ♦ **PowerShell 2.0:** Windows PowerShell 2.0 Engine must be installed on each node of the cluster.
- ♦ **Shared Disks:** All shared disks belong to the active node.
- ♦ **Domain Controller:** The migrated single-cluster virtual machine requires access to a domain controller with the same parameters as the original domain controller. To address this requirement, consider either leaving the original domain controller online, migrate it concurrently with the cluster.

Block-Based Transfer

When you use block-based transfer for cluster workloads, the block-based driver components are not installed on the cluster nodes. The block-based transfer occurs using a driverless synchronization with an MD5-based replication. Because the block-based driver is not installed, no reboot is required on the source cluster nodes.

NOTE: File based transfer and block-based transfer with a BBT driver are not supported for migrating Microsoft Windows clusters.

Cluster Node Failover during the First Full Replication

A cluster workload requires that the first full replication completes successfully without a cluster node failover. If a cluster node failover occurs prior to the completion of the first full replication, you must remove the existing workload, re-add the cluster using the active node, and try again.

Cluster Node Failover during Replication

If a cluster node failover occurs prior to the completion of the copy process during a full replication or an incremental replication, the command aborts and a message displays indicating that the replication needs to be re-run.

Cluster Node Failover between Replications

The nodes must have similar profiles to prevent interruptions in the replication process. If a cluster node failover occurs between the incremental replications of a cluster workload and if the new active node's profile is similar to the failed active node, the migration contract continues as scheduled for the next incremental replication. Otherwise, the next incremental replication command fails.

The profiles of cluster nodes are considered similar if all of the following conditions are met:

- ♦ Serial numbers for the nodes' local volumes (System volume and System Reserved volume) must be the same on each cluster node.

NOTE: Use the customized *Volume Manager* utility to change the local volume serial numbers to match each node of the cluster. See [“Synchronizing Serial Numbers on Cluster Node Local Storage” on page 227](#).

If the local volumes on each node of the cluster have different serial numbers, you cannot run a replication after a cluster node failover occurs. For example, during a cluster node failover, the active node Node 1 fails, and the cluster software makes Node 2 the active node. If the local drives on the two nodes have different serial numbers, the next replication command for the workload fails.

- ◆ The nodes must have the same number of volumes.
- ◆ Each volume must be exactly the same size on each node.
- ◆ The nodes must have an identical number of network connections.

5.5.2 Enabling or Disabling Windows Cluster Discovery

PlateSpin Migrate can discover and inventory Windows Server failover clusters in your PlateSpin environment based on the active node in each cluster. Alternatively, it can treat all active and non-active cluster nodes as standalone machines.

To enable cluster discovery for all Windows clusters, ensure that the parameter `DiscoverActiveNodeAsWindowsCluster` is set to `True`. This is the default setting. Cluster discovery, inventory, and workload migration use the hostname or IP address of a cluster's active node, instead of using its cluster name and an administration share. You do not configure separate workloads for the cluster's non-active nodes. For other cluster workload migration requirements, see ["Prerequisites" on page 132](#).

To disable cluster discovery for all Windows clusters, set the parameter `DiscoverActiveNodeAsWindowsCluster` to `False`. This setting allows the PlateSpin Server to discover all nodes in a Windows failover cluster as standalone machines. That is, it inventories a cluster's active node and non-active nodes as a regular, cluster-unaware Windows workloads.

To enable or disable cluster discovery:

- 1 Go to the PlateSpin Server configuration page at <https://<platespin-server-ip-address>/PlateSpinConfiguration>
- 2 Search for `DiscoverActiveNodeAsWindowsCluster`, then click **Edit**.
- 3 In the **Value** field, select **True** to enable cluster discovery, or select **False** to disable cluster discovery.
- 4 Click **Save**.

5.5.3 Adding Resource Name Search Values

To help identify the active node in a Windows failover cluster, PlateSpin Migrate must differentiate the name of the shared Cluster IP Address resource from the names of other IP address resources on the cluster. The shared Cluster IP Address resource resides on the cluster's active node.

The global parameter `MicrosoftClusterIPAddressNames` on the PlateSpin Server Configuration page contains a list of search values to use in discovery for a Windows cluster workload. When you add a Windows cluster workload, you must specify the IP address of the cluster's currently active node. PlateSpin Migrate searches the names of the cluster's IP address resources on that node to find one that *starts with* the specified characters of any value in the list. Thus, each search value must contain enough characters to differentiate the shared Cluster IP Address resource on a specific cluster, but it can be short enough to apply to discovery in other Windows clusters.

For example, a search value of `Clust IP Address` or `Clust IP` matches the resource names `Clust IP Address` for 10.10.10.201 and `Clust IP Address` for 10.10.10.101.

The default name for the shared Cluster IP Address resource is `Cluster IP Address` in English, or the equivalent if the cluster node is configured in another language. The default search values in the `MicrosoftClusterIPAddressNames` list include the resource name `Cluster IP Address` in English and each of the supported languages.

Because the resource name of the shared Cluster IP Address resource is user-configurable, you must add other search values to the list, as needed. If you change the resource name, you must add a related search value to the `MicrosoftClusterIPAddressNames` list. For example, if you specify a resource name of `Win2012-CLUS10-IP-ADDRESS`, you should add that value to the list. If you have multiple clusters using the same naming convention, an entry of `Win2012-CLUS` matches any resource name that starts with that sequence of characters.

To add search values in the `MicrosoftClusterIPAddressNames` list:

- 1 Go to the PlateSpin Server configuration page at `https://<platespin-server-ip-address>/PlateSpinConfiguration`
- 2 Search for `MicrosoftClusterIPAddressNames`, then click **Edit**.
- 3 In the **Value** field, add one or more search values to the list.
- 4 Click **Save**.

5.5.4 Setting the Quorum Arbitration Timeout

You can set the `QuorumArbitrationTimeMax` registry key for Windows Server failover clusters in your PlateSpin environment by using the global parameter `FailoverQuorumArbitrationTimeout` on the PlateSpin Server Configuration page. The default timeout is 60 seconds, in keeping with the Microsoft default value for this setting. See [QuorumArbitrationTimeMax \(https://msdn.microsoft.com/en-us/library/aa369123%28v=vs.85%29.aspx?f=255&MSPPErr=-2147217396\)](https://msdn.microsoft.com/en-us/library/aa369123%28v=vs.85%29.aspx?f=255&MSPPErr=-2147217396) on the Microsoft Developer Network website. The specified timeout interval is honored for quorum arbitration at failover and failback.

To set the quorum arbitration timeout for all Windows failover clusters:

- 1 Go to the PlateSpin Server configuration page at `https://<platespin-server-ip-address>/PlateSpinConfiguration`
- 2 Search for `FailoverQuorumArbitrationTimeout`, then click **Edit**.
- 3 In the **Value** field, specify the maximum number of seconds to allow for quorum arbitration.
- 4 Click **Save**.

5.5.5 Setting Local Volume Serial Numbers

You can use the *Volume Manager* utility to change the local volume serial numbers to match in each node of the cluster. See [“Synchronizing Serial Numbers on Cluster Node Local Storage” on page 227](#).

5.5.6 PlateSpin CutOver

When the PlateSpin Cutover operation is complete and the cluster comes online, you see a multi-node cluster with one active node (all other nodes are unavailable).

To perform a PlateSpin cutover (or to test the PlateSpin cutover) on a Windows cluster, the cluster must be able to connect to a domain controller. To leverage the test cutover functionality, you need to migrate the domain controller along with the cluster. During the test, bring up the domain controller, followed by the Windows cluster workload (on an isolated network).

5.5.7 PlateSpin Cluster Migration

A PlateSpin cluster migration operation requires a full replication for Windows Cluster workloads.

After PlateSpin cluster migration is complete, you must reattach the shared storage and rebuild the cluster environment before you can rejoin additional nodes to the newly restored cluster.

For information about rebuilding the cluster environment after a PlateSpin migration, see [Knowledgebase Article 7016770](#).

5.5.8 Migrating Windows Clusters with the Web Interface

Use the PlateSpin Migrate Web Interface to add the active node by specifying the IP address of the active node, configure migration for the active node, and then run the migration.

5.5.9 Migrating Windows Clusters Using the Migrate Client

You can use a **Move** job to migrate the essential services of a cluster that results in a functional single-node cluster in a virtual machine.

The workflow of migrating a Windows cluster is similar to that of migrating a standalone server:

- 1 Discover the active node by specifying the IP address of the active node.
- 2 In the Servers view, use drag-and-drop to start a migration job, then configure the job's parameters.
- 3 (Conditional: successful migration) If the migration job completes successfully, perform a [Server Sync operation](#) on the active node.

NOTE: If the active node in the cluster fails over before you can perform a Server Sync operation, perform a full migration using the *new* active node, and then perform a Server Sync on this new node.

- 4 (Conditional: failover prior to migration) If a cluster failover occurs prior to the completion of file transfer, the migration job aborts. If this happens, refresh the source and retry the migration job.

NOTE: If you select **Shut down** for the source's post-migration end state, a shutdown of all source nodes of the cluster results.

6 Windows Workload Portability with a PlateSpin Image

This section provides information about using the PlateSpin Image volume archiving feature (Windows only).

- ♦ [Section 6.1, “About PlateSpin Images,” on page 137](#)
- ♦ [Section 6.2, “Designating a PlateSpin Image Server,” on page 137](#)
- ♦ [Section 6.3, “Capturing a Workload to a PlateSpin Image,” on page 139](#)
- ♦ [Section 6.4, “Deploying a PlateSpin Image,” on page 140](#)
- ♦ [Section 6.5, “Managing PlateSpin Images,” on page 141](#)

6.1 About PlateSpin Images

One of PlateSpin Migrate’s three fundamental workload infrastructures, a PlateSpin Image is an image of a supported Windows workload consisting of volume data along with configuration specifics of the source server’s hardware, operating system, and network identity.

Image configurations are maintained in an XML (`config.xml`) file with each image having one or more sets of associated volume data.

PlateSpin Images and the image server’s `config.xml` configuration file are stored on the designated PlateSpin Image Server host in the following directory:

```
..\Program Files\PlateSpin Image Server
```

In addition to volume data directly captured during an X2I migration, PlateSpin Migrate supports existing or raw volume data.

Like peer-to-peer migrations, image deployment allows for key workload configuration options, such as those for managing the workload’s disk layout, volume sizes, network identity, and domain or workgroup affiliation.

6.2 Designating a PlateSpin Image Server

To work with PlateSpin Images, you must first designate a machine as an image server by installing the PlateSpin Image Server software on it. You can install a PlateSpin Image Server instance either on a dedicated host or on your PlateSpin Server host. For information about storing PlateSpin PlateSpin Images on a NAS (Network Attached Storage) device or a remote share, see [KB Article 7921021 \(https://www.netiq.com/support/kb/doc.php?id=7921021\)](https://www.netiq.com/support/kb/doc.php?id=7921021).

NOTE: Although collocation of the PlateSpin Server with a PlateSpin Image Server instance on the same host is supported, the recommended setup is to install a PlateSpin Image Server on a dedicated host, which simplifies troubleshooting related to imaging functionality.

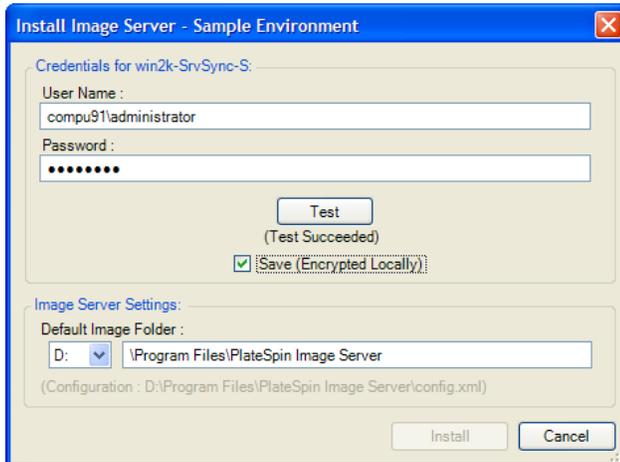
Dedicated PlateSpin Image Server hosts must meet the following requirements:

Table 6-1 PlateSpin Image Server Host Requirements

Requirement	Details
Operating System	Any of the following, running on dedicated hardware or in a virtual machine: <ul style="list-style-type: none"> ◆ Microsoft Windows Server 2012 R2 ◆ Microsoft Windows Server 2012 ◆ Microsoft Windows Server 2008 R2
Disk Space	Minimum 100 MB for basic controller software. Additional space requirements depend on the number and size of workload images that you intend to store on a given image server.
Software	◆ Microsoft .NET Framework 3.5 SP1

To designate a machine as a PlateSpin Image Server:

- 1 Discover the system you plan to designate as a PlateSpin Image Server.
- 2 In the Servers view, right-click the discovered server and select **Install Image Server**.



- 3 Provide administrator credentials for the selected host and specify the desired directory for image files.
- 4 Click **Install**.

PlateSpin Migrate installs a controller on the selected host and configures it to run as a PlateSpin Image Server. On completion, the Servers view lists a new PlateSpin Image Server item: 

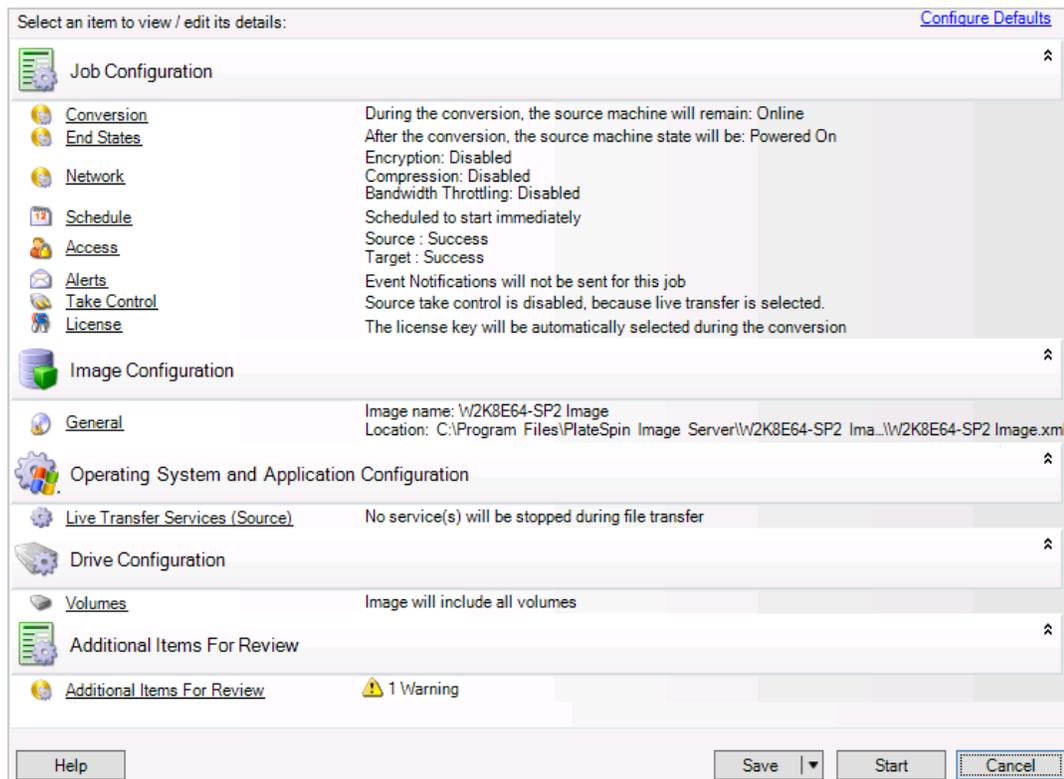
6.3 Capturing a Workload to a PlateSpin Image

Use this procedure to capture a physical or virtual workload as a PlateSpin Image.

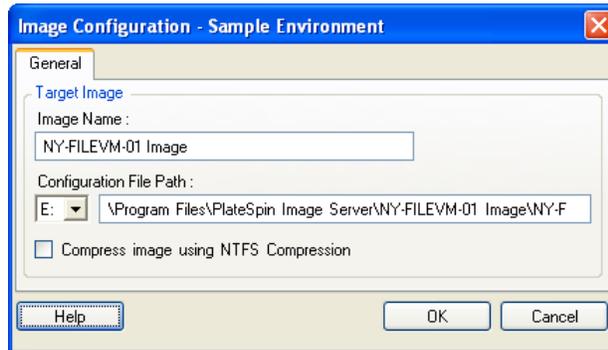
- 1 Discover, or refresh the details of, your source workload and your PlateSpin Image Server.
- 2 Start a new Capture Image job by using one of the following methods:
 - ♦ In the Servers view, right-click the source workload, then select **Capture Image**. In the Action window, select the source workload and the target image server.
 - ♦ In the Tasks pane, click **Capture Image**. In the Action window, select the source workload and the target image server.
 - ♦ In the Servers view, drag the source workload and drop it on the image server. If you configured PlateSpin Migrate to bypass the Action window on drag-and-drop, the Create Image dialog box prompts you to specify whether you want to create a new image or use existing volume data.



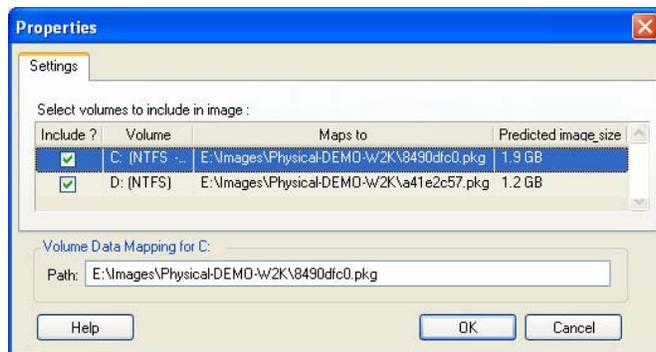
- 3 Select **Create Image**, then click **OK**.



- 4 Specify the required settings for the migration job by clicking the links in each category:
 - ♦ **Job Configuration:** Specify the required transfer method and operational continuity settings for your source and target (**General**), scheduling options (**Schedule**), source and target credentials (**Credentials**), job status and progress notification options, temporary network settings (**Take Control**), and the required license key to use (**License Key**).
 - ♦ **Image Configuration:** Specify the image name, the path to the location where the you want the image to be stored, and whether or not to use NTFS compression (under Image Configuration, click **General**).



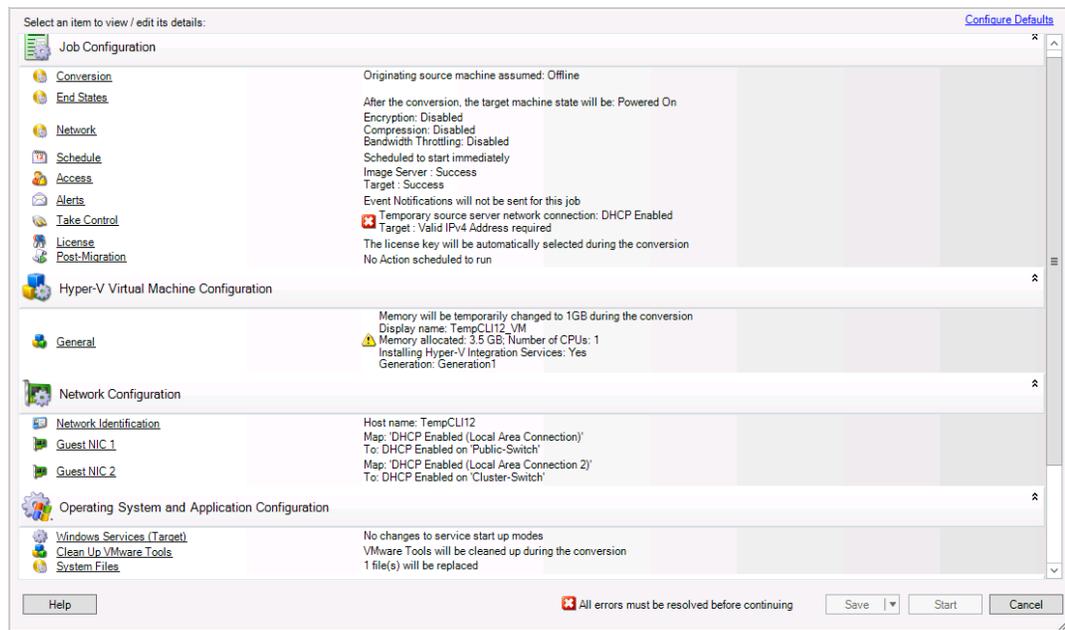
- ♦ **Operating System and Application Configuration:** If you selected the Live Transfer method, specify how you want PlateSpin Migrate to handle operating system and application services on your source (**Live Transfer Services**).
- ♦ **Drive Configuration:** Select the volumes that you want PlateSpin Migrate to include in the image and specify the path for the package file (under Drive Configuration, click **Volumes**).



6.4 Deploying a PlateSpin Image

Use this procedure to deploy a PlateSpin Image on a supported physical machine or virtualization platform.

- 1 Drag and drop the required PlateSpin Image to a discovered target physical machine or VM host.



2 Specify the required settings for the migration job by clicking the links in each category.

Migration jobs are auto-configured to create the target machine with the same settings as the source server. Depending on the objectives of the migration, you can:

- ◆ Modify the **Network Identification** settings to configure the hostname and domain/workgroup registration of the target machine.
- ◆ Modify the **Guest NIC** settings to configure the TCP/IP properties for the network adapters on the target machine.
- ◆ Modify the **Drive Configuration** settings to select the volumes to copy during the migration.

3 If the intended target is a virtual machine, specify the required virtual machine parameters and select the options you require, such as memory allocation, or automatic installation of VMware Tools or VMAdditions.

4 Review and address errors and warnings.

5 Click **Start** to deploy the image.

6.5 Managing PlateSpin Images

- ◆ [Section 6.5.1, “Moving Images from One PlateSpin Image Server to Another,” on page 142](#)
- ◆ [Section 6.5.2, “Automating Image Operations,” on page 142](#)
- ◆ [Section 6.5.3, “Browsing and Extracting Image Files,” on page 142](#)

6.5.1 Moving Images from One PlateSpin Image Server to Another

- 1 Copy the image directory from the old PlateSpin Image Server host's file system to a location on the new PlateSpin Image Server host.
- 2 Update the new PlateSpin Image Server's `config.xml` file to identify the path to and the name of the image that was moved from the old PlateSpin Image Server.
- 3 Refresh the new image server's details in the PlateSpin Migrate Client's Servers view.

For more information, see [KB Article 7920189 \(https://www.netiq.com/support/kb/doc.php?id=7920189\)](https://www.netiq.com/support/kb/doc.php?id=7920189).

6.5.2 Automating Image Operations

You can use the `ImageOperations` command line utility, included with PlateSpin Migrate, to automate several tasks related to images, such as regularly moving multiple base images, along with related increments, between PlateSpin Image Servers.

The utility provides the capability to automate the following operations:

- ♦ **Register:** Associate an image or image increments with a specified image server.
- ♦ **Unregister:** Disassociate a registered image from a specified image server.
- ♦ **Gather:** Assemble a package of a PlateSpin Image and its volumes into a specified subdirectory.

To use the `ImageOperations` command line utility:

- 1 On your PlateSpin Image Server host, open a command interpreter (`cmd.exe`) and change the current directory to `..\Program Files\PlateSpin Image Server\ImageOperations`.
- 2 Type `ImageOperations` followed by the required command and parameters, then press Enter. For command syntax and usage details, type `ImageOperations`, then press Enter.
- 3 When you have finished, refresh the image server's details in the Servers view.

6.5.3 Browsing and Extracting Image Files

During a disaster recovery effort or a business continuity exercise you can selectively restore files in your production server's file system, using backup versions of those files that are stored in PlateSpin Images.

To do this, you can use the PlateSpin Image Browser utility, which enables you to browse, search, sort, and extract files from different sources:

- ♦ An image file
- ♦ A specific image increment file

You can work with both base images and image increments by loading different files:

- ♦ A base image's corresponding binary file (`volume-x.pkg`) or text configuration file (`image_name.xml`).
- ♦ An image increment's binary (`image_increment.pkg`) file. You cannot use an increment's text configuration file (`image_increment_name.xml`).

The utility enables you to work with image files in a Windows Explorer-like environment. A command line version enables you to extract files at the command line.

- ◆ “Starting the Image Browser and Loading Image Files” on page 143
- ◆ “Sorting and Searching Items in the Image Browser Interface” on page 143
- ◆ “Extracting Items” on page 144
- ◆ “Browsing and Extracting Image Files at the Command Line” on page 144

Starting the Image Browser and Loading Image Files

- 1 Start the ImageBrowser program (ImageBrowser.exe), located in one of the following directories:

- ◆ On your PlateSpin Server host:

..\PlateSpin Migrate Server\bin\ImageOperations

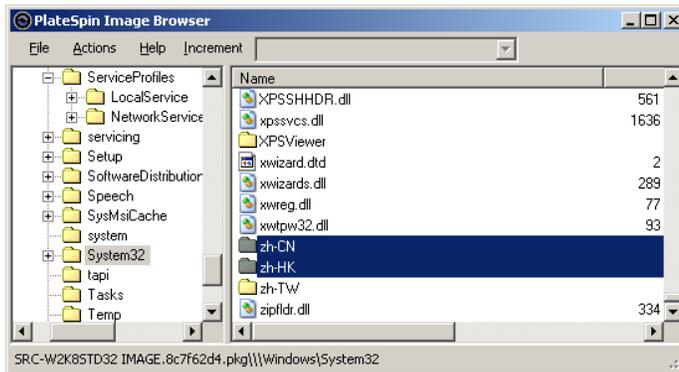
- ◆ On your PlateSpin Image Server host:

..\Program Files\PlateSpin Image Server\ImageOperations

The utility starts and displays the Open dialog box. At any time after the program's initial startup, you can load an image file by clicking **File > Open**.

- 2 In the Open dialog box, select the file type, navigate to and select the required image or image increment file, then click **OK**.

The utility loads the required file and displays its contents in a two-pane interface.



Depending on the size of the image, it might take a few seconds to several minutes for the utility to load the required file.

Sorting and Searching Items in the Image Browser Interface

You can sort the contents of a selected directory by name, size, type, date last modified, and by file attribute. To sort items in a selected view, click the corresponding bar at the top of the right pane.

You can search for a specific directory name or file name. You can use alphanumeric text, wildcards, and regular expressions. Regular expression search patterns that you specify must adhere to the Microsoft .NET Framework regular expression syntax requirements. See the [Microsoft .NET Framework Regular Expressions page on MSDN \(http://msdn.microsoft.com/en-us/library/hs600312.aspx\)](http://msdn.microsoft.com/en-us/library/hs600312.aspx).

To search for an item:

- 1 Load the required image or image increment. See [“Starting the Image Browser and Loading Image Files” on page 143](#).
- 2 In the left pane, select a volume or a subdirectory.
- 3 On the **Actions** menu, click **Search**.

Alternately, you can right-click the required volume or subdirectory in the left pane and click **Search** in the context menu.

The Image Browser Search window opens.

- 4 Specify the name of the file you are searching. If you are using a regular expression, select the corresponding option.
- 5 Click **Search**.

The results are shown in the right pane.

Extracting Items

- 1 Load the required image or image increment. See [“Starting the Image Browser and Loading Image Files” on page 143](#).
- 2 Locate and select the required file or directory. You can select multiple files and directories in the right pane.
- 3 On the **Actions** menu, click **Extract**.

Alternately, you can right-click the required item and click **Extract** in the context menu.

The Browse for Folder dialog box opens.

- 4 Browse to the required destination, then click **OK**.

The selected items are extracted to the specified destination.

NOTE: Files that you choose to overwrite are deleted if you interrupt the extraction process.

Browsing and Extracting Image Files at the Command Line

To browse and extract files from images and image increments at the command line, you can use the `ImageBrowser.Console` utility.

To start the utility:

- 1 On your PlateSpin Image Server host, open a command interpreter (`cmd.exe`) and change the current directory to `..\Program Files\PlateSpin Image Server\ImageOperations`.
- 2 At the command prompt, type `ImageBrowser.Console`, then press Enter.

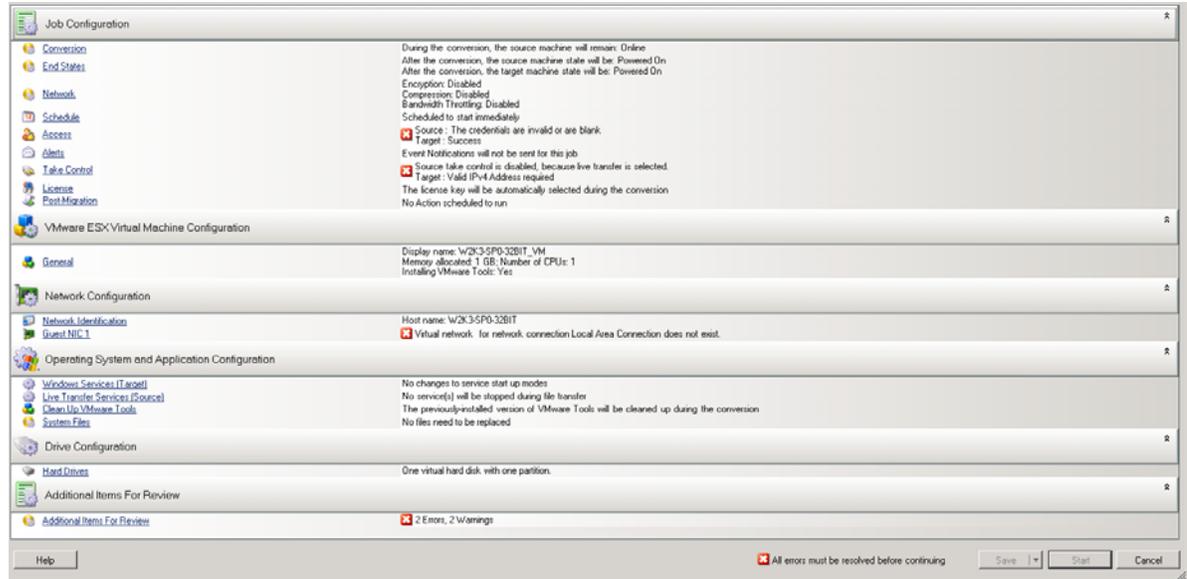
For command syntax and usage details, type `ImageBrowser.Console /help`, then press Enter.

7 Essentials of Workload Migration Jobs

This section provides information about configuring the key aspects of a workload migration job using the PlateSpin Migrate Client.

- ◆ Section 7.1, “Selecting a License Key for a Migration Job,” on page 146
- ◆ Section 7.2, “Configuring Automatic Email Alerts of Job Status and Progress,” on page 146
- ◆ Section 7.3, “Specifying Conversion Options,” on page 147
- ◆ Section 7.4, “Specifying End States for Source and Target Workloads,” on page 147
- ◆ Section 7.5, “Specifying Network Options,” on page 147
- ◆ Section 7.6, “Specifying Credentials to Access Sources and Targets,” on page 148
- ◆ Section 7.7, “Managing a Workload’s Network Identity,” on page 149
- ◆ Section 7.8, “Job Scheduling,” on page 150
- ◆ Section 7.9, “Setting Up Migration Networking,” on page 151
- ◆ Section 7.10, “Configuring Target Virtual Machines,” on page 157
- ◆ Section 7.11, “Handling Operating System Services and Applications,” on page 164
- ◆ Section 7.12, “Handling a Workload’s Storage Media and Volumes,” on page 169
- ◆ Section 7.13, “Including a Custom Post-migration Action in a Migration Job,” on page 178

Figure 7-1 Sample of a Migration Job Window



7.1 Selecting a License Key for a Migration Job

By default, PlateSpin Migrate automatically selects the best license key for a particular migration job. If you have multiple license keys, you can select a specific license key to apply to a particular migration job. For information about product licensing and license key management, see [Section 2.1, “PlateSpin Migrate Product Licensing,” on page 27](#).

To view or modify the license key selected for a migration job:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,” on page 83](#).
- 2 In the Job Configuration section of the Migration Job window, click **License**.
- 3 To manually choose a different key, deselect the **Automatically select the best license key during the conversion** check box and choose the required license key from the drop-down list.
- 4 Click **OK**.

The selected license key is displayed on the **License** tab and the description is updated accordingly.

Certain licenses cannot be selected if they are invalid for the current migration. Licenses can be invalid for reasons such as:

- ♦ There are no remaining migrations for the license.
- ♦ The license does not allow X2V migrations and the current migration is a P2V.
- ♦ The license does not support live transfer migrations and the current migration is marked for live transfer.

7.2 Configuring Automatic Email Alerts of Job Status and Progress

You can set up a migration job to automatically send email notifications about status and progress to a specified address:

- ♦ **Job events:** Job status messages such as `Completed`, `Recoverable Error`, and `Failed`.
- ♦ **Job progress:** Detailed job progress messages at configurable intervals.

You can specify SMTP server and email account details either during the migration job or globally. See [“Notification Service” on page 45](#).

To configure automatic email notifications:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,” on page 83](#).
- 2 In the Job Configuration section of the Migration Job window, click **Alerts** and configure the required options.
- 3 Click **OK**.

7.3 Specifying Conversion Options

Conversion options enable you to specify:

- ♦ How data is transferred from source to target. PlateSpin Migrate supports multiple transfer methods, and their availability depends on your workload and migration job type.

See [“Supported Transfer Methods” on page 21](#).

- ♦ The scope of workload data to transfer from the source to the target (**Full Migration and Changes only**). Applicable only to Server Sync jobs.

See [“Synchronizing Workloads with Server Sync” on page 119](#).

To specify the transfer options for a migration job:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,” on page 83](#).
- 2 In the Job Configuration section of the Migration Job window, click **Conversion**.
- 3 Select the scope and method of data transfer.
- 4 Click **OK**.

7.4 Specifying End States for Source and Target Workloads

After completing a migration, PlateSpin Migrate shuts down or boots up the source and target workloads depending on the nature of the migration. For example, a Copy workload job results in the source workload being left to run, and a Move workload job results in the source workload being shut down.

To specify non-default post-migration end states for your source and target:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,” on page 83](#).
- 2 In the Job Configuration section of the Migration Job window, click **End States**.
- 3 Select the end state for the source and the target machine.
- 4 Click **OK**.

7.5 Specifying Network Options

Network options are settings for security, performance, and connectivity, and enable you to specify:

- ♦ Whether you want the system to compress workload data that is being transferred over the network.

See [“Data Compression” on page 25](#).

Fast consumes the least CPU resources on the source but yields a lower compression ratio, **Maximum** consumes the most, but yields a higher compression ratio. **Optimal**, the middle ground, is the recommended option.

- ♦ Whether to encrypt the data transferred from source to target.

See [“Security and Privacy” on page 23](#).

- ◆ Whether you want to apply bandwidth throttling for the current migration job.

See [“Bandwidth Throttling” on page 25](#).

To throttle this job’s bandwidth to a specified rate, select the **Enable Bandwidth Throttling** option, specify the required maximum value in Mbps, and optionally a time period during which to enforce the throttling. If specified, the **from** and **to** time values are based on the source workload’s system time.

If no time interval is defined, bandwidth is throttled to the specified rate at all times by default. If time interval is defined and the migration job executes outside this interval, data is transferred at full speed.

- ◆ Additional IP addresses for source workloads to enable communication in environments that use network address translation (NAT).

For information on how to specify additional IP addresses for your PlateSpin Server, see [“Migrations Across Public and Private Networks through NAT” on page 42](#).

To specify network options:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,” on page 83](#).
- 2 In the Job Configuration section of the Migration Job window, click **Network**.
- 3 Select the required network options.
- 4 Click **OK**.

7.6 Specifying Credentials to Access Sources and Targets

For a migration job to execute properly, you must provide valid credentials for your source and target.

Table 7-1 Source and Target Credentials

Credentials	Windows	Linux
User Name	Account username with local or domain-level administrative privileges. Use this format: <ul style="list-style-type: none"> ◆ For domain member machines: <i>authority\principal</i> ◆ For workgroup member machines: <i>hostname\principal</i> 	Root or root-level username.
Password	Valid password for the specified username.	Valid password for the specified username.

When setting up a migration job, you can validate the provided credentials and save them for future migration jobs that use the same source and target.

To specify source and target credentials:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,” on page 83](#).
- 2 In the Job Configuration section of the Migration Job window, click **Access**.

- 3 Specify the credentials.
- 4 Click **OK**.

7.7 Managing a Workload's Network Identity

PlateSpin Migrate enables you to manage the network identity and domain registration of your migration target workload and specify related preferences as part of a migration job. By default, a job is configured to preserve a source workload's network identity and domain registration. You can modify the default configuration to suit the objectives of your migration job.

Proper configuration of migration target's network identity is especially important when you are migrating a workload to a different domain, planning to take it off a domain, or if you intend to change the hostname of a workload while it is in the domain.

To configure a target workload's network identity options:

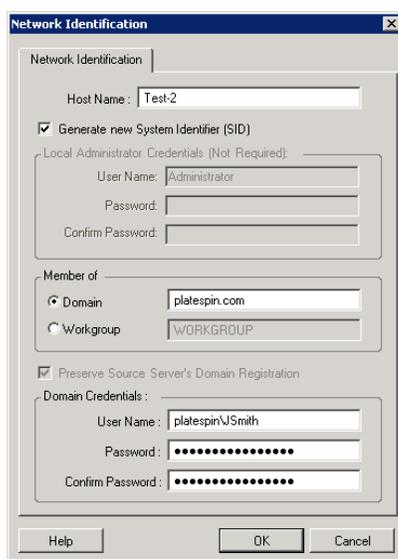
- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, "Setting a Migration Job,"](#) on page 83.
- 2 In the Network Configuration section of the Migration Job window, click **Network Identification**.
- 3 Specify the options and then click **OK**.

Configuration options vary depending on whether the target machine is Windows or Linux. For information about the configuration options, see the following sections:

- ♦ [Section 7.7.1, "Managing the Identity of Windows Workloads,"](#) on page 149
- ♦ [Section 7.7.2, "Managing the Network Identity of Linux Workloads,"](#) on page 150

7.7.1 Managing the Identity of Windows Workloads

Use these settings to configure the network identity of your target Windows workload.



Host Name: Specify the desired hostname for the target machine.

Generate New SID: When this option is selected, the target workload is assigned a new System Identifier (SID). Credentials are required only for Windows 2008 and Vista systems, and must be the credentials for the local (embedded) `Administrator` account. If this account has been locally renamed on the source, provide the new name. If this account is disabled on Vista (default), enable it first.

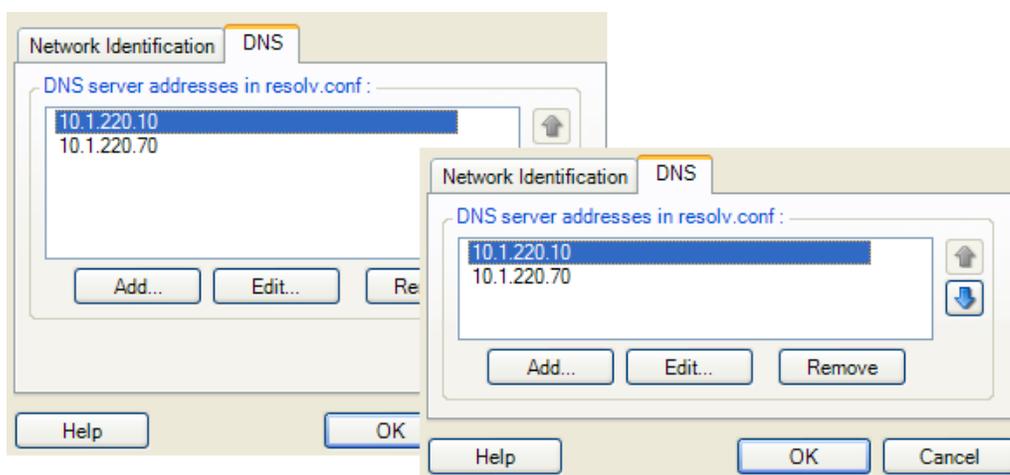
Member of (Domain / Workgroup): Select the required option and type the name of the domain or workgroup that you want the target machine to join.

Preserve Source Server's Domain Registration: Preserves domain registration and ensures that the source server domain registration remains intact during migration. If you disable this option, the source machine's domain account is transferred to the target machine. The source server still appears to be on the domain, but does not have a valid connection.

Domain Credentials: If the target machine is to be part of a domain, specify valid credentials for a user account with permission to add servers to the domain, such as a member of the Domain Admins group or Enterprise Admins group.

7.7.2 Managing the Network Identity of Linux Workloads

Use these settings to configure the network identity of your target Linux workload and DNS server addresses as required.



Network Identification tab: Specify the desired hostname for the target server.

DNS tab: Use the **Add**, **Edit**, and **Remove** buttons to manage DNS server entries for the new virtual machine.

7.8 Job Scheduling

Scheduling options enable you to specify whether the migration job is to run immediately or on a specific date and a specific time. For information on other options that PlateSpin Migrate provides for saving and executing jobs, see [“Setting Up, Executing, and Managing Jobs” on page 83](#)

To access scheduling options of a migration job:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,” on page 83](#).
- 2 In the Job Configuration section of the Migration Job window, click **Schedule**.

- 3 Specify the schedule.
- 4 Click **OK**.

7.9 Setting Up Migration Networking

For each workload portability job, you must properly configure workload networking so that source workloads and targets can communicate with each other and the PlateSpin Server during the migration process, and that the network configuration of a target workload is in line with its end state.

- ♦ [Section 7.9.1, “Temporary \(Take Control\) Network Settings,” on page 151](#)
- ♦ [Section 7.9.2, “Target Post-Migration Networking,” on page 153](#)
- ♦ [Section 7.9.3, “TCP/IP and Advanced Network Settings,” on page 155](#)

Temporary Networking: Also called *Take Control Network Settings*; they apply to source and target workloads booted into a temporary pre-execution environment. See [“Offline Transfer with Temporary Boot Environment” on page 22](#).

7.9.1 Temporary (Take Control) Network Settings

Temporary (Take Control) Network Settings control how source workloads, targets, and the PlateSpin Server communicate among each other during the migration. If required, you can manually specify a temporary network address to your source and target, or configure them to use a DHCP-assigned IP address during the migration.

During Windows and Linux workload migrations, the Temporary Network Settings control the PlateSpin Server’s communication with the source and target workloads that are booted into a temporary pre-execution environment. See [“Offline Transfer with Temporary Boot Environment” on page 22](#).

To configure Temporary (Take Control) network settings:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,” on page 83](#).
- 2 In the Job Configuration section of the Migration Job window, click **Take Control**.
- 3 To access network interface mapping and TCP/IP settings, click **Configure** in the source and target areas as applicable.
- 4 Click **OK**.

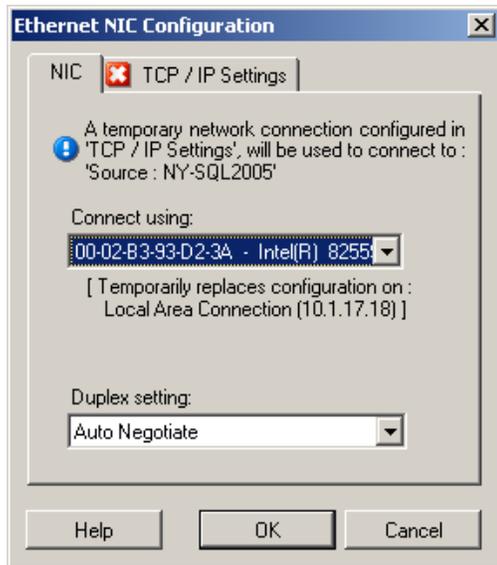
Configuration options for the Temporary networking vary and depend on whether the network interface is virtual or physical, and whether it is connecting a Windows or a Linux workload.

- ♦ [“Temporary \(Take Control\) Network Settings: Physical Network Interfaces” on page 152](#)
- ♦ [“Temporary \(Take Control\) Network Settings: Virtual Network Interfaces” on page 153](#)

Target Take Control network settings are only used during an Offline migration process. On completion, target network settings are read from settings you specify for Target Post-Migration Networking. See [“Target Post-Migration Networking” on page 153](#).

Temporary (Take Control) Network Settings: Physical Network Interfaces

These settings apply only to source physical machines. For target physical machines, Temporary (Take Control) network settings are configured during the boot process that uses the boot ISO image. See [“Discovering Target Physical Machines” on page 67](#).



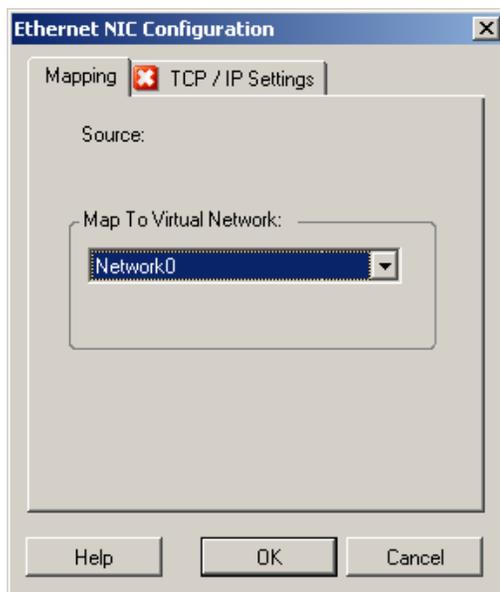
Connect using: If multiple network adapters are present, select the adapter that can communicate with both the PlateSpin Server and the target.

Duplex setting: Use the drop-down list to select network card duplexing. It must match the duplex setting for the switch to which the network interface is connected. When the source is connected to switch ports that are set to 100 Mbit full duplex and cannot be changed to auto negotiation, select **Force NIC to Full Duplex**.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See [“TCP/IP and Advanced Network Settings” on page 155](#).

Temporary (Take Control) Network Settings: Virtual Network Interfaces

These settings apply to both source and target Take Control network settings.



Map to Virtual Network: From the drop-down list, select the virtual switch or network to use for communication during an Offline migration. If multiple virtual network adapters are present, select the adapter that can communicate with both the PlateSpin Server and the source machine. This network can differ from the network on which the target virtual machine will run after the migration.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See [“TCP/IP and Advanced Network Settings” on page 155](#).

7.9.2 Target Post-Migration Networking

Target post-migration network settings defined in a migration job control the network configuration of a target after the migration is complete. This applies to both physical and virtual network interfaces.

During workload migration, the target workload’s post-migration network settings are configured while the workload is booted into a pre-execution environment.

To configure target post-migration network settings:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,” on page 83](#).
- 2 In the Network Configuration section of the Migration Job window, do one of the following:
 - ◆ **For target virtual machines:** click **Guest NIC**.
 - ◆ **For target physical machines:** click **Network Connection**.

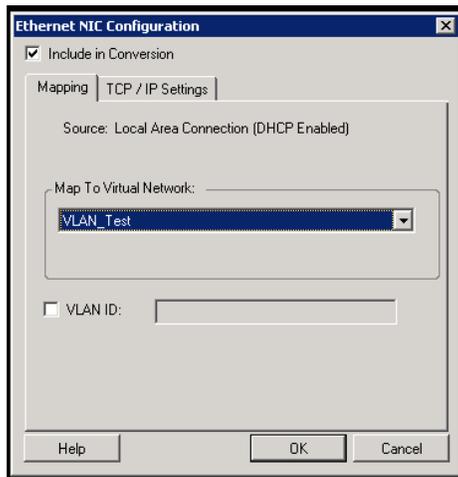
- 3 Configure the options as required and click **OK**.

The Configuration options for the target post-migration network settings vary and depend on whether the network interface is virtual or physical, and whether it is connecting a Windows or a Linux workload. For more information about the options, review the following sections:

- ♦ [“Post-Migration Networking for Physical Network Interfaces \(Windows and Linux\)”](#) on page 154
- ♦ [“Post-Migration Networking for Virtual Network Interfaces \(Windows and Linux\)”](#) on page 155

Post-Migration Networking for Physical Network Interfaces (Windows and Linux)

Use these settings to configure the post-migration network settings of a workload being migrated to physical hardware.

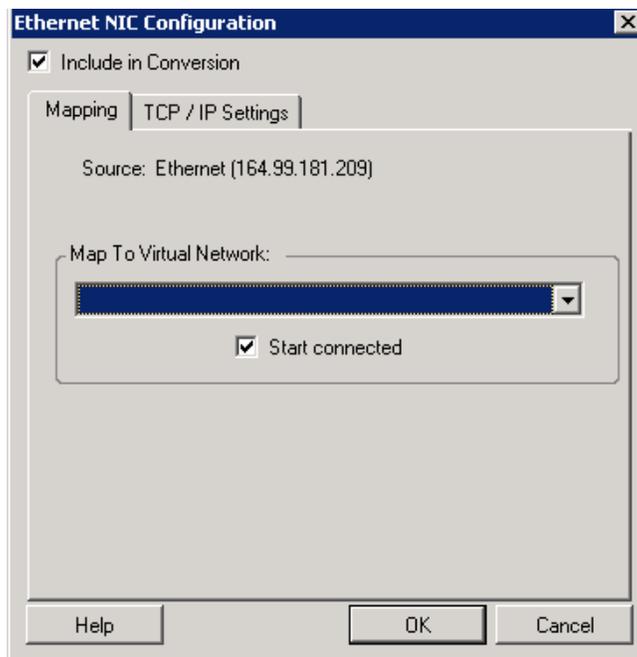


Connect using: If multiple network adapters are present, select the adapter that can communicate with the PlateSpin Server.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See [“TCP/IP and Advanced Network Settings”](#) on page 155.

Post-Migration Networking for Virtual Network Interfaces (Windows and Linux)

By default, PlateSpin Migrate configures a migration job to create a virtual NIC for each NIC found on the source. For post-migration connectivity, ensure that the target virtual NIC is mapped to the appropriate virtual network on the target virtualization platform.



Include in Conversion: When this option is selected, PlateSpin Migrate creates a virtual NIC for a source NIC.

Map to Virtual Network: Select the virtual network that will be used on the target VM. Choose a virtual network that allows the target VM to communicate with the server.

Start connected: Enable this option to connect the virtual network interface when starting the ESX target machine.

VLAN ID: (Applicable for target machine on a Hyper-V server only) Enable this option to specify the virtual network ID to be used on the target machine. If you do not specify this ID, then the virtual network ID of the source machine is used by default.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See [“TCP/IP and Advanced Network Settings” on page 155](#).

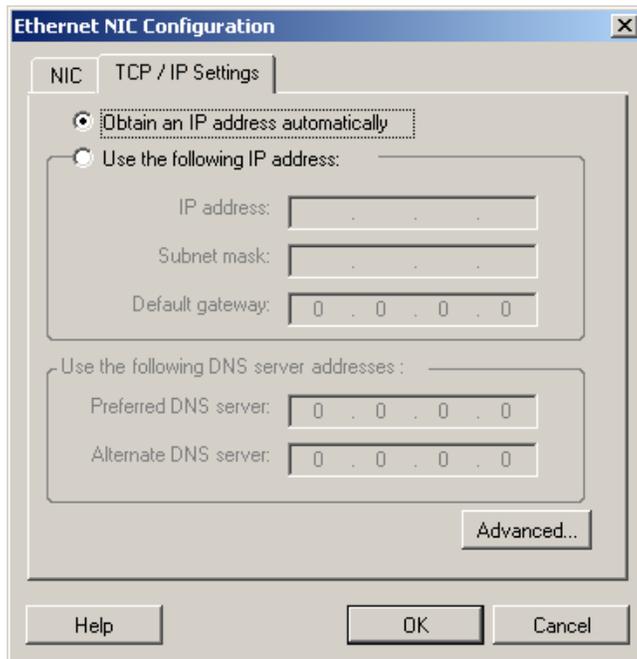
7.9.3 TCP/IP and Advanced Network Settings

PlateSpin Migrate provides a standard network configuration interface to both source and target network settings, and for both Temporary and target post-migration networking. Configuration settings vary slightly, depending on the operating system.

- ◆ [“TCP/IP and Advanced Network Settings \(Windows\)” on page 156](#)
- ◆ [“TCP/IP and Advanced Network Settings \(Linux\)” on page 156](#)

TCP/IP and Advanced Network Settings (Windows)

The following are standard TCP/IP and advanced network settings for Windows workloads:



Obtain an IP address automatically: When this option is selected, the workload uses an IP address automatically assigned by a DHCP server during the migration process.

Use the following IP address: Select this option to specify a static IP address.

Use the following DNS server addresses: If required, specify preferred and alternative DNS server addresses.

Advanced: Click this button to access advanced TCP/IP configuration settings, then specify or edit default gateway, DNS server, and WINS server information as required.

TCP/IP and Advanced Network Settings (Linux)

The following are standard TCP/IP and advanced network settings for Linux workloads:



Obtain an IP address automatically: When this option is selected, the workload uses an IP address automatically assigned by a DHCP server during the migration process.

Use the following IP address: Select this option to specify a static IP address.

Advanced: Click this button to access DNS configuration settings, then specify preferred and alternate DNS server addresses as required. You can also indicate whether you want DNS addresses copied to the `resolv.conf` file located in your target's `/etc` directory.

7.10 Configuring Target Virtual Machines

For jobs that involve workload virtualization, PlateSpin Migrate provides a mechanism for specifying target VM configuration options, such as providing a target VM name and a configuration file path, selecting a datastore to use, and allocating virtual memory, in accordance with the features and capabilities of the selected virtualization platform.

If you have resource pools configured on your target virtualization platform, you can select a resource pool for your VM to be assigned to.

NOTE: If your target VMware ESX server is part of a fully automated Distributed Resource Scheduler (DRS) cluster (a cluster with its VM migration automation level set to **Fully Automated**), the newly created target VM's automation level is changed to **Partially Automated** for the duration of the migration. This means that your target VM might power up on a different ESX server from the one initially selected, but migration is prevented from automatic execution.

To specify target VM configuration options:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, "Setting a Migration Job," on page 83](#).
- 2 In the Virtual Machine Configuration section of the Migration Job window, click **General**.

- 3 Specify the values for the configuration options and click **OK**.

For more information about the virtual machine configuration options, see [Section 7.10.1, “Virtualization Platform-Specific Virtual Machine Configuration Options,”](#) on page 158

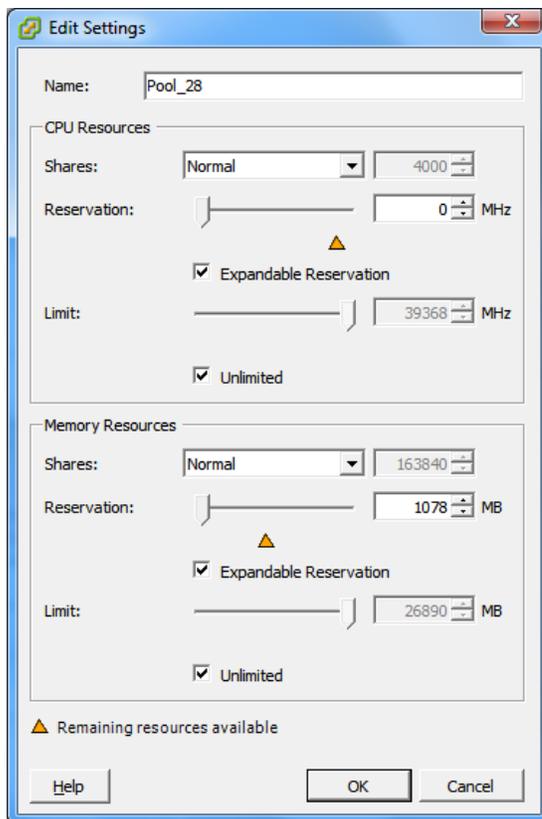
7.10.1 Virtualization Platform-Specific Virtual Machine Configuration Options

PlateSpin Migrate displays target virtual machine configuration options specific to the selected target and also provides access to advanced configuration options.

- ◆ [“Virtual Machine Configuration: VMware ESXi 5”](#) on page 158
- ◆ [“Virtual Machine Configuration: VMware ESX 4.1”](#) on page 160
- ◆ [“Virtual Machine Configuration: Microsoft Hyper-V”](#) on page 162
- ◆ [“Virtual Machine Configuration: Citrix XenServer”](#) on page 164

Virtual Machine Configuration: VMware ESXi 5

The following are configuration options specific to VMware vSphere 5 (applicable to all VMs under the containing resource pool).



Name: Specify the display name for the new virtual machine.

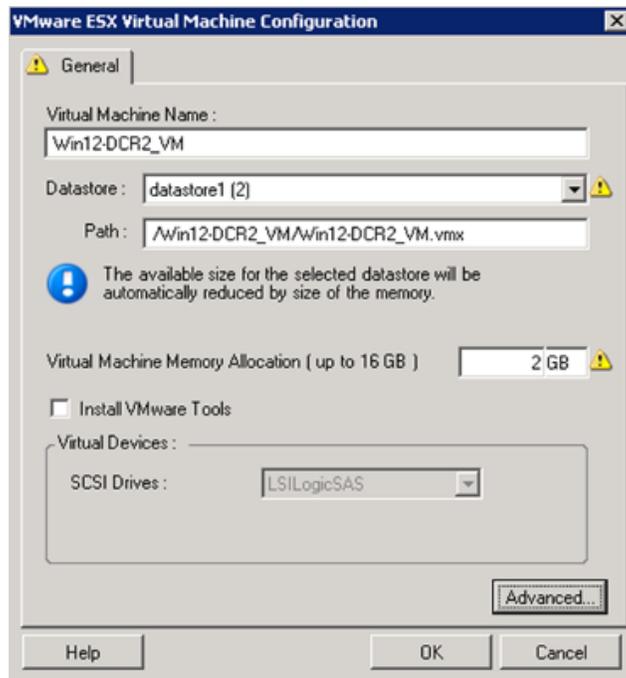
CPU Resources

- ◆ **Shares:** CPU shares for this virtual machine with respect to the parent's total. Peer VMs share resources according to their relative share values bounded by the **Reservation** and **Limit**. Select **Low**, **Normal**, or **High**, which specify share values respectively in a 1:2:4 ratio. Select **Custom** to give each virtual machine a specific number of shares, which express a proportional weight.
- ◆ **Reservation:** Guaranteed CPU allocation for this VM.
Expandable Reservation: Select this option to specify that more than the specified reservation is allocated if resources are available in a parent.
- ◆ **Limit:** Upper limit for this virtual machine's CPU allocation.
Unlimited: Select this option to specify no upper limit.

Memory Resources: (these are similar to CPU resource settings, but apply to memory resources)

Virtual Machine Configuration: VMware ESX 4.1

The following are configuration options specific to VMware ESX systems prior to vSphere 5. To access settings that control resource pools, the number of CPUs, and CPU scheduling affinity, click [Advanced](#).



Virtual Machine Name: Specify the display name for the new virtual machine.

Datastore: Select the datastore where you want to create the *.vmx file.

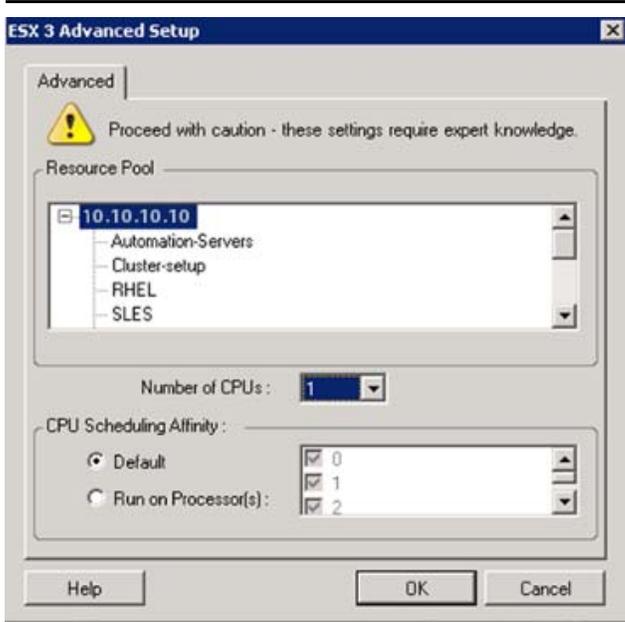
Configuration File Path: Specify a name and the directory path for the virtual machine's *.vmx configuration file.

Virtual Machine Memory Allocation: Specify a value for the amount of virtual RAM to be assigned to the virtual machine.

Install VMware Tools: Enable this option to install VMware tools during the migration process (recommended).

SCSI Drives: Select either **BusLogic** or **LSILogic** (the recommended option).

Advanced: Click this button to view or modify advanced VM configuration settings.



Resource Pool: If required, assign your target VM to a resource pool. When no resource pool is specified, the VM is assigned to the root resource pool.

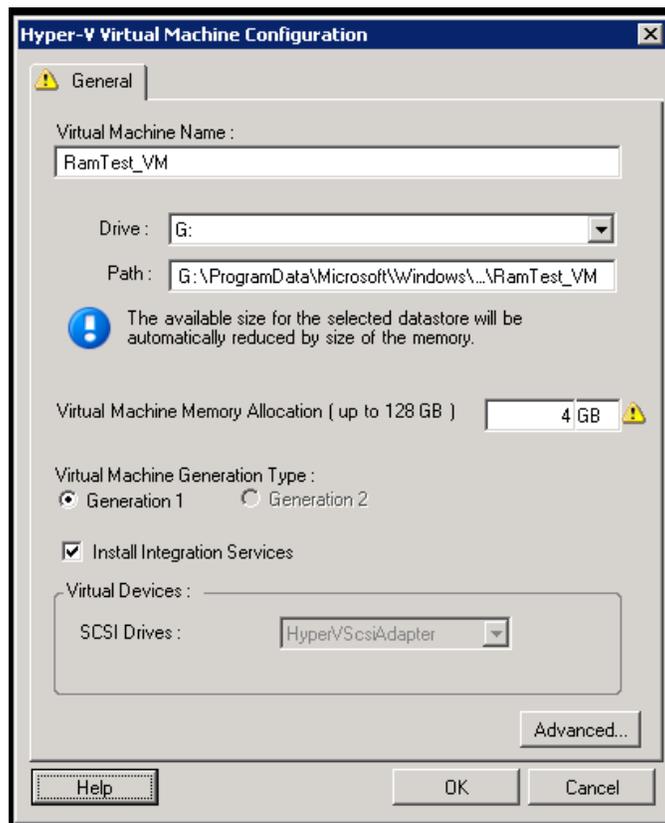
Number of CPUs: Select the required number of CPUs to assign to the target VM. For example, you can convert a single-processor workload to a multi-processor VM, or a multi-processor workload to a single-processor VM.

CPU Scheduling Affinity: Represents which ESX Server processors the virtual machine can run on (if your ESX Server is a multiprocessor system). Specify the required processor or select **Default** (recommended).

For details, see your VMware documentation.

Virtual Machine Configuration: Microsoft Hyper-V

The following are configuration options specific to Hyper-V 2012 systems.



Virtual Machine Name: Specify the display name for the new virtual machine.

Datastore: Select the datastore where you want to create the *.vmx file.

Configuration File Path: Specify a name and the directory path for the virtual machine's *.vmx configuration file.

Virtual Machine Memory Allocation: Specify a value for the amount of virtual RAM to be assigned to the virtual machine.

Virtual Machine Generation Type: Select one of the following generations type for the new virtual machine:

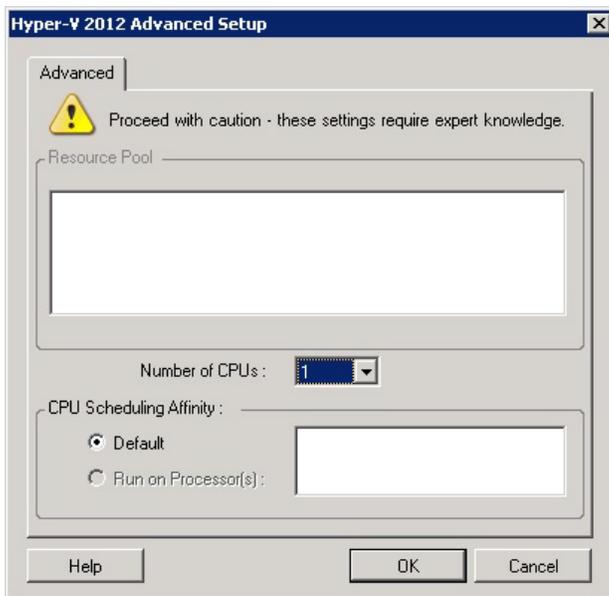
- ♦ **Generation 1:** Select this option to deploy the target virtual machine with Hyper-V BIOS architecture.
- ♦ **Generation 2:** Select this option to deploy the target virtual machine with Hyper-V UEFI architecture

This option is displayed only for Hyper-V machines.

Install Integration Services: Enable this option to install integration services during the migration process (recommended).

SCSI Drives: Select either **BusLogic** or **LSILogic** (the recommended option).

Advanced: Click this button to view or modify advanced VM configuration settings.



Number of CPUs: Select the required number of CPUs to assign to the target VM. For example, you can convert a single-processor workload to a multi-processor VM, or a multi-processor workload to a single-processor VM.

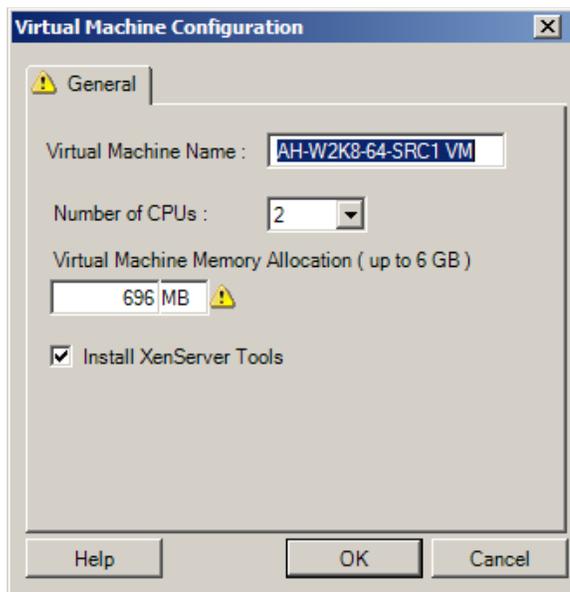
NOTE: For Generation 1, you can create four legacy network cards and eight synthetic network cards (if integration service is enabled). For Generation 2, you can create eight synthetic network cards.

CPU Scheduling Affinity: Represents which Hyper-V Server processors the virtual machine can run on (if your Hyper-V Server is a multiprocessor system). Specify the required processor or select **Default** (recommended).

For details, see your Hyper-V documentation.

Virtual Machine Configuration: Citrix XenServer

The following are configuration options specific to Citrix XenServer.



Virtual Machine Name: Specify the display name for the new virtual machine.

Number of CPUs: Select the number of CPUs to assign to the target VM. For example, you can convert a single-processor workload to a multi-processor VM, or a multi-processor workload to a single-processor VM.

Virtual Machine Memory Allocation: Specify a value for the amount of virtual RAM to be assigned to the virtual machine.

Install XenServer Tools: Enable this option to install XenServer Tools during the migration process (recommended).

7.11 Handling Operating System Services and Applications

PlateSpin Migrate provides a mechanism for migration jobs to handle Windows services (including special features for those services related to Microsoft SQL Server and Microsoft Exchange Server software), Linux daemons, and virtualization enhancements (such as VMware Tools). In addition, when Windows HAL or kernel file replacements are required, you can view which ones were selected for update during a migration job.

- ◆ [Section 7.11.1, “Handling the Startup Mode of Services \(Windows Targets\),” on page 165](#)
- ◆ [Section 7.11.2, “Handling Source Workload Services or Daemons During Live Transfer \(Windows and Linux\),” on page 166](#)
- ◆ [Section 7.11.3, “Viewing Windows System Files Selected for Replacement During a Migration,” on page 167](#)
- ◆ [Section 7.11.4, “Handling the Run Level of Daemons \(Linux Targets\),” on page 168](#)
- ◆ [Section 7.11.5, “Handling Virtualization Enhancement Software,” on page 168](#)

7.11.1 Handling the Startup Mode of Services (Windows Targets)

You can configure a job to modify the startup mode of selected Windows services after the migration is complete. For example, if you don't need a certain Windows service to continue running on a virtualized workload, you can configure your job to set that service's target startup type to Disabled.

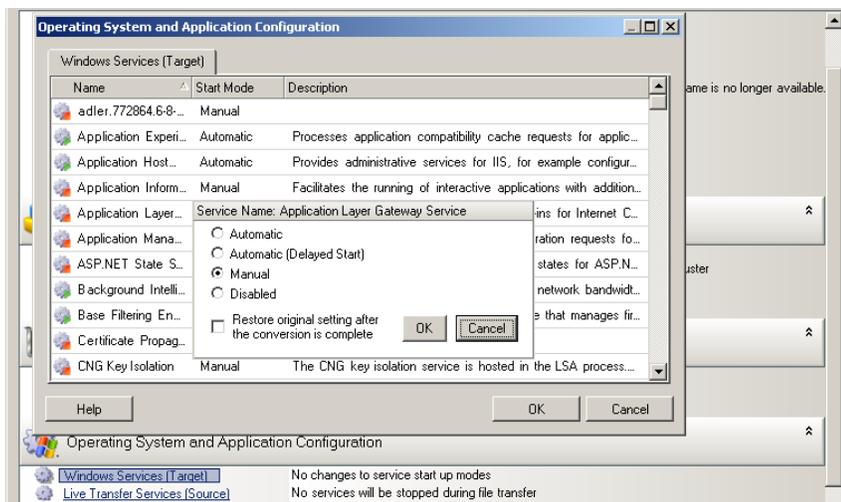
If you require a service to start upon another service's request, set the startup type of the required service to Manual.

You can also configure the job to restore a service's original startup type after the migration is complete. For example, you might require a virus scanner to be disabled during the migration, but its startup type to be restored after the migration completes.

TIP: You can globally configure your startup mode preferences for selected Windows services in PlateSpin Server default options. See ["Target Service Defaults"](#) on page 49.

To configure post-migration startup mode of Windows services:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, "Setting a Migration Job,"](#) on page 83.
- 2 In the Operating System and Application Configuration section of the Migration Job window, click **Windows Services (Target)** and then click an item in the **Start Mode** column.



- 3 Select the desired startup mode.
- 4 To restore the original setting after conversion is complete, select the checkbox.
- 5 Click **OK**.

7.11.2 Handling Source Workload Services or Daemons During Live Transfer (Windows and Linux)

For Live Transfer jobs, PlateSpin Migrate provides a mechanism to stop selected services or daemons during the migration. This ensures that data on your source is captured in a consistent state.

If your source workload is running Microsoft SQL Server or Microsoft Exchange Server software, you can configure your migration job to automatically copy the database files of these servers. If you do not require the migration to include the volume containing the databases, consider not stopping these services.

If your source workload includes I/O-intensive application services that might inhibit the ability of the file transfer process to keep up with the changes, consider stopping them during a Live Transfer migration.

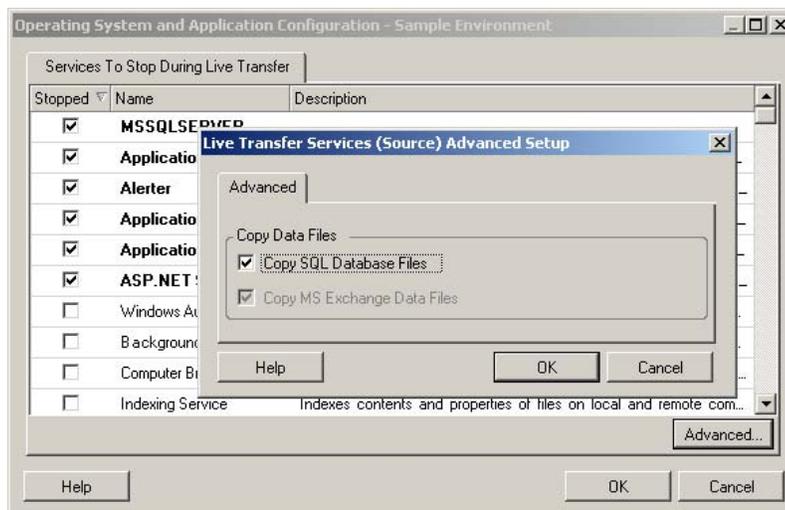
After the completion of the migration, services that you select to stop during a Live Transfer migration are automatically restarted on the source, unless you explicitly configure your migration job to power off the source on completion.

For Linux systems, consider using the custom `freeze` and `thaw` scripting capability. See [“Freeze and Thaw Scripting Capabilities \(Linux Block-Level Migrations\)”](#) on page 82.

TIP: You can globally configure your preferences for stopping selected Windows services during VSS File-based or VSS Block-based Live Transfer. See [“Source Service Defaults”](#) on page 48.

To specify which services or daemons you want the system to stop during Live Transfer:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,”](#) on page 83.
- 2 In the Operating System and Application Configuration section of the Migration Job window, click [Live Transfer Services/Daemons \(Source\)](#).
- 3 To indicate that you want SQL Server and Exchange Server database files copied during the migration, click [Advanced](#) (applicable to Windows systems only)



- 4 Click **OK**.

7.11.3 Viewing Windows System Files Selected for Replacement During a Migration

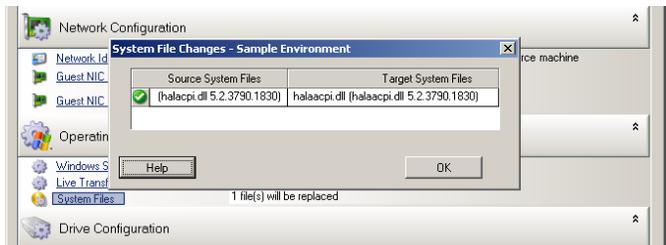
When converting Windows workloads with system files (such as a HAL or kernel files) that are incompatible with the target infrastructure, PlateSpin Migrate uses an appropriate file from its library and saves a backup copy of the source file (*.bak) on the target, in the same system directory.

You can view the HAL or kernel files that PlateSpin Migrate identifies as those requiring replacement:

To view the files selected for replacement during migration:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,”](#) on page 83.
- 2 In the Operating System and Application Configuration section of the Migration Job window, click **System Files**.

Files selected for replacement during migration are listed.



- 3 Click **OK**.

The following warnings might display at the bottom of the dialog box:

Driver Cache is empty	Indicates that you might need to place the necessary files into the local driver cache on the source Windows server (..\Windows\Driver Cache).
The driver cache contains a higher version	PlateSpin Migrate has a partial match with its matrix but the driver cache contains a later version of one or more system files than the one that PlateSpin Migrate will use.
File <filename> will be replaced with lower version	PlateSpin Migrate has not found a match for the system files in its matrix. It will replace the system files with a version that is earlier than the ones that were discovered as the source machine's original system files.
File <filename> will be replaced with higher version	PlateSpin Migrate has not found a match for the system files in its matrix. It will replace the system files with a version that is later than the ones that were discovered as the source machine's original system files.

If warnings appear on the screen, click **More Help** (only available if warnings exist) to learn more.

See also the following [KB Article 7920815](https://www.netiq.com/support/kb/doc.php?id=7920815) (https://www.netiq.com/support/kb/doc.php?id=7920815) (FAQ: Understanding the System Files Information screen).

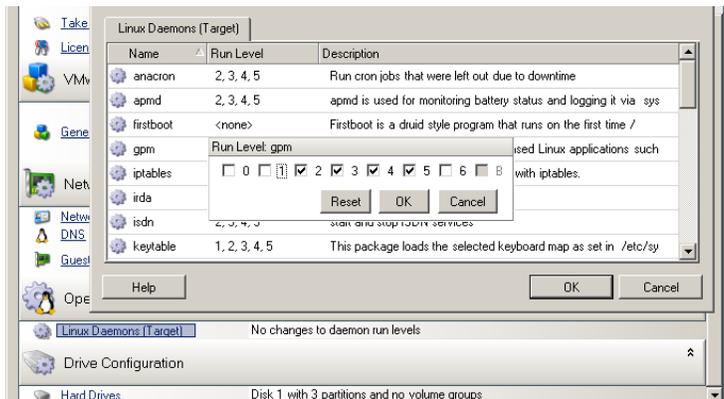
7.11.4 Handling the Run Level of Daemons (Linux Targets)

You can configure a job to modify the run level of selected Linux daemons after the migration is complete. Use the following numeric values:

0	Shutdown
1	Single-user mode
2	Unused (user-defined)
3	Full multi user-mode (no GUI)
4	Unused (user-defined)
5	Full multi-user mode with display manager (GUI)
6	Reboot

To configure the post-migration run level of Linux daemons:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,”](#) on page 83.
- 2 In the Operating System and Application Configuration section of the Migration Job window, click **Linux Daemons (Target)** and then click an item in the **Run Level** column



- 3 Select the desired run level. Click **OK**.

7.11.5 Handling Virtualization Enhancement Software

For V2X migrations, PlateSpin Migrate provides a mechanism to automatically uninstall virtualization enhancement software, such as VMware Tools.

When converting a workload on a VMware platform that has an earlier version of VMware Tools installed, PlateSpin Migrate identifies the presence of obsolete software and adds a VMware Tools Cleanup step in the migration job.

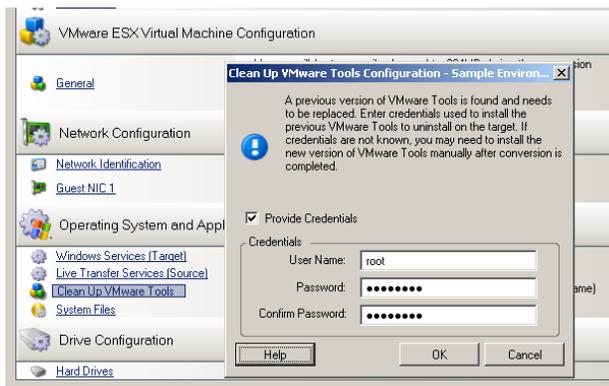
You must provide administrator credentials to uninstall VMware Tools. The credentials provided must match the admin-level user account that was logged in during the installation of VMware Tools.

When the earlier version is uninstalled, PlateSpin Migrate proceeds with the installation of the new version of VMware Tools.

NOTE: If you are downgrading a virtual machine that has VMware Tools installed, or if you are converting a virtual machine to another VMware target that has an older version of VMware Tools, the installation of VMware Tools during the configuration of the target will fail.

To configure a job to remove or replace VMware Tools during the migration:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,”](#) on page 83.
- 2 In the Operating System and Application Configuration section of the Migration Job window, click **Clean up VMware Tools**.



- 3 Depending on the target, PlateSpin Migrate identifies existing instances of VMware Tools and prompts to either replace or remove them, as applicable:
 - ♦ **For non-VMware targets:** The job configuration interface prompts you to uninstall VMware Tools. Provide the same admin-level credentials used to install the software. If the credentials are unknown, VMware Tools remains on the target machine after migration.
 - ♦ **For VMware targets:** The job configuration interface prompts you to replace VMware Tools. Provide the same admin-level credentials used to install the obsolete version of VMware Tools. If the credentials are unknown, install the new version of VMware Tools manually after the migration completes.
- 4 Click **OK**.

7.12 Handling a Workload’s Storage Media and Volumes

PlateSpin Migrate provides mechanisms for configuring your migration job to handle your workload volumes and their physical or virtual layout in the target infrastructure.

NOTE: PlateSpin Migrate does not support NSS file system and EVMS volumes.

Workload Virtualization (X2V)

When virtualizing a workload, you can select which volumes you want included on the target and manage their free space sizes. You can also control how physical disk arrangement on your source is propagated on the peer virtual machine according to the target virtualization platform's storage media configuration and virtual disk handling features and capabilities.

Workload Deployment to Physical Hardware

When you are migrating workloads to physical hardware, you can select which source volumes to include and size, and which target disks to repartition and populate.

RAID Storage

PlateSpin Migrate supports RAID (Redundant Array of Independent Disks) storage, which it treats like any other storage hardware. As long as the associated storage controller driver is present, PlateSpin Migrate successfully completes the migration. PlateSpin Migrate does not support software implementations of RAID.

SAN Storage

PlateSpin Migrate supports SAN storage. As long as the driver for the associated host bus adapter (HBA) is present, PlateSpin Migrate successfully completes the migration. PlateSpin Migrate treats SAN LUNs like any other disk with logical volumes.

NAS

PlateSpin Migrate supports Network Attached Storage (NAS) systems. PlateSpin Migrate treats NAS like any other disk with logical volumes.

Windows Dynamic Disks

PlateSpin Migrate supports Windows dynamic disks, including mirrored, striped, spanned, and RAID 5 configurations.

PlateSpin Migrate treats dynamic disks like it treats any other logical volume. When you are converting workloads that have dynamic disks, the disks on the target workload are created as basic disks, which you can use to remove unnecessary or obsolete dynamic disk configurations. After the migration, you can upgrade the required disks on your targets from basic to dynamic disks.

Linux Logical Volumes

PlateSpin Migrate supports logical volumes of Linux workloads. If Logical Volume Manager (LVM) is installed on your Linux source, you can use several LVM1 and LVM2 features to better manage your target workload's volume layout and organization.

You can set up your workload migration job to:

- ◆ Re-create logical volumes of the source on the target, or create logical volumes on the target even if the source is not using LVM.
- ◆ Create new volume groups on the target that are not on the source, or omit volume groups from the target that are present on the source.
- ◆ Rename volume groups on the target.
- ◆ Distribute volumes to different volume groups and disks.

NOTE: PlateSpin Migrate does not support:

- ♦ Copying LVM snapshots and LVM mirrors. You can create LVM snapshots and mirrored logical volumes on the target after the migration completes.
 - ♦ Migration of Linux workloads with encrypted volumes.
-

Storage layout and volume configuration settings depend on the job configuration mode (Advanced or Wizard), migration type, target virtualization platform, and source operating system.

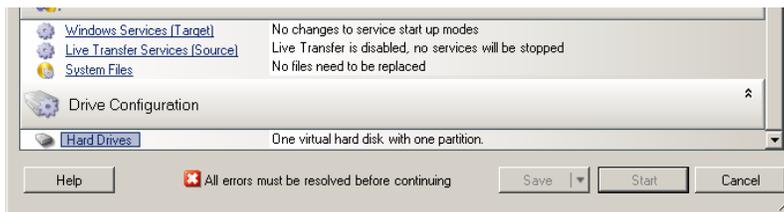
The following topics provide additional information:

- ♦ [Section 7.12.1, “Storage Layout and Volume Configuration,” on page 171](#)
- ♦ [Section 7.12.2, “Volume Configuration in Server Sync,” on page 177](#)

7.12.1 Storage Layout and Volume Configuration

To access drive configuration options:

- ♦ In the **Drive Configuration** of the Migration Job windows, click **Hard Drives**.

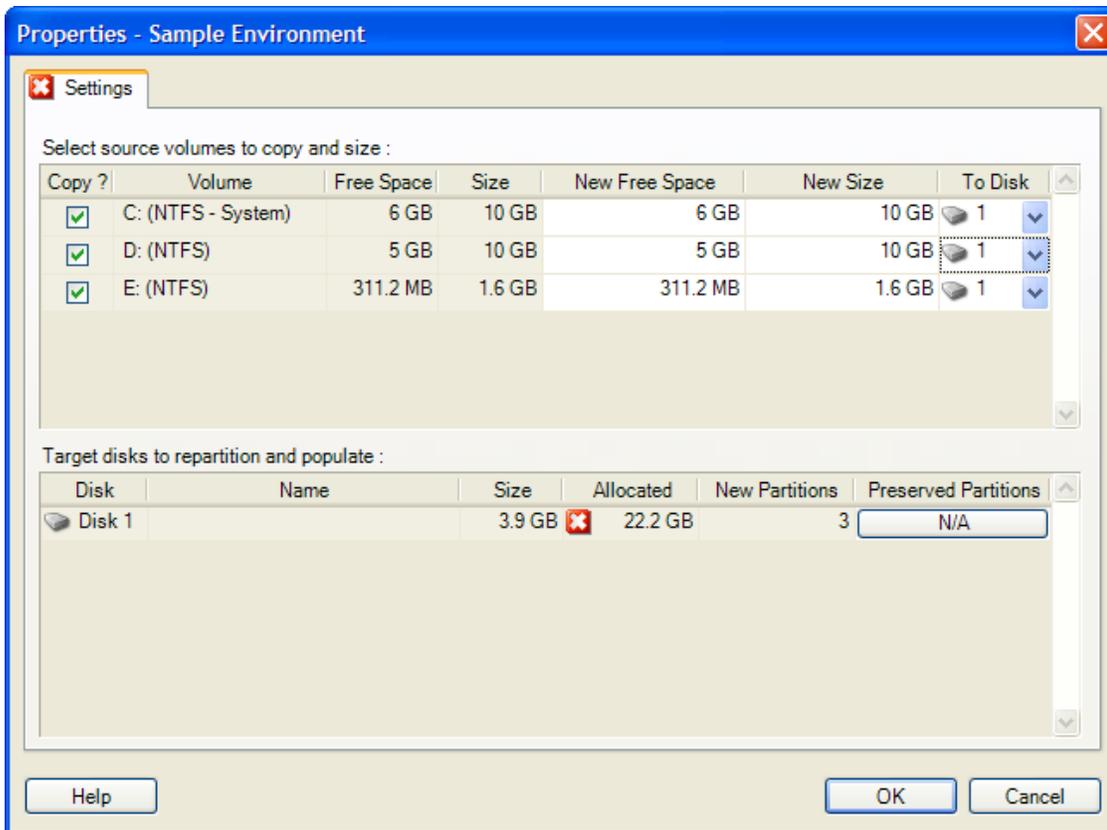


Settings vary depending on the target system.

- ♦ [“Windows X2P Drive Configuration” on page 171](#)
- ♦ [“Linux Drive and LVM Volume Configuration” on page 172](#)
- ♦ [“Target VM-Specific P2V/V2V Drive Configuration” on page 175](#)

Windows X2P Drive Configuration

Use these settings to select the volumes to copy during the migration:



Copy: Select the volumes to be copied during the migration.

New Free Space: To resize the volume during the migration, specify the desired amount of free space. PlateSpin Migrate automatically adjusts **New Size**.

New Size: To resize the volume during the migration, specify the desired size. PlateSpin Migrate automatically adjusts **New Free Space**.

To Disk: Select which hard drive the volume will be copied to on the physical target machine.

Preserve Partitions: Click this column to determine if an existing vendor partition should remain intact during the migration. If the partitions are not selected, PlateSpin Migrate permanently removes the partitions from the server.

Linux Drive and LVM Volume Configuration

Use these settings to select the volumes and non-volume source spaces to copy and size during the migration. If LVM is installed on the source, a **Volume Group** tab provides you with corresponding options.

- ◆ [“Handling Linux Disks and Volume Groups” on page 173](#)
- ◆ [“Linux Drive and LVM Volume Configuration \(Settings Tab\)” on page 173](#)
- ◆ [“Linux Drive and LVM Volume Configuration \(Volume Groups Tab\)” on page 174](#)

Handling Linux Disks and Volume Groups

The PlateSpin Migrate Client provides you with Linux-specific user interface elements that provide you with options to properly handle your Linux storage.

Note the following sequence of steps that you must take for properly configuring and mapping newly-added disks and volume groups.

- 1 After adding a new disk, go to the **Volume Groups** tab and map the required volume group name by selecting the **Include** option.

See [Linux Drive and LVM Volume Configuration \(Volume Groups Tab\)](#).

- 2 Specify Size in Allocation for Volume Group Box

- 3 For each added disk, specify the required size in the corresponding **Allocation for Volume Group** field.

After the system focus shifts away from the field, the size of the newly-added disk is updated dynamically.

Linux Drive and LVM Volume Configuration (Settings Tab)

Use these settings to select source volumes to copy, non-volume source spaces to re-create and size, and target disks to repartition and populate.

The screenshot shows the 'Drive Configuration' dialog box with the 'Volume Groups' tab selected. The dialog is divided into three main sections:

- Virtual disks to create:** A table with columns: Disk, Datastore, Size, Thin, File Name. One entry is visible: 'Virtual disk 0' with Datastore 'newLUN-VC...', Size '10 GB', Thin 'No', and File Name '/linux_VM/linux_VM_1.vmdk'. Buttons 'Add' and 'Remove Unused Disks' are present.
- Select volumes to copy and size:** A table with columns: Include, Volume, Free Space, Size, New Free Space, New Si..., Disk/Volume Group/EVMS Volu... One entry is visible: '/' with Free Space '6.6 GB', Size '9 GB', New Free Space '6.6 GB', New Si... '9 GB', and Disk/Volume Group/EVMS Volu... 'Disk 0'. The 'Include' checkbox is checked.
- Select non-volume storage to recreate and size:** A table with columns: Include, Type, Partition, Size, Is Swap, New Size, Disk/Volume Group/EVMS Volume. One entry is visible: with Partition '/dev/sda1', Size '1 GB', Is Swap 'Yes', New Size '1 GB', and Disk/Volume Group/EVMS Volume 'Disk 0'. The 'Include' checkbox is checked.

Buttons 'Help', 'OK', and 'Cancel' are located at the bottom of the dialog.

Include: Select the volumes or non-volume source spaces to be copied or re-created and sized during the migration.

New Free Space: To resize the volume during the migration, enter the desired amount of free space. PlateSpin Migrate automatically adjusts **New Size**.

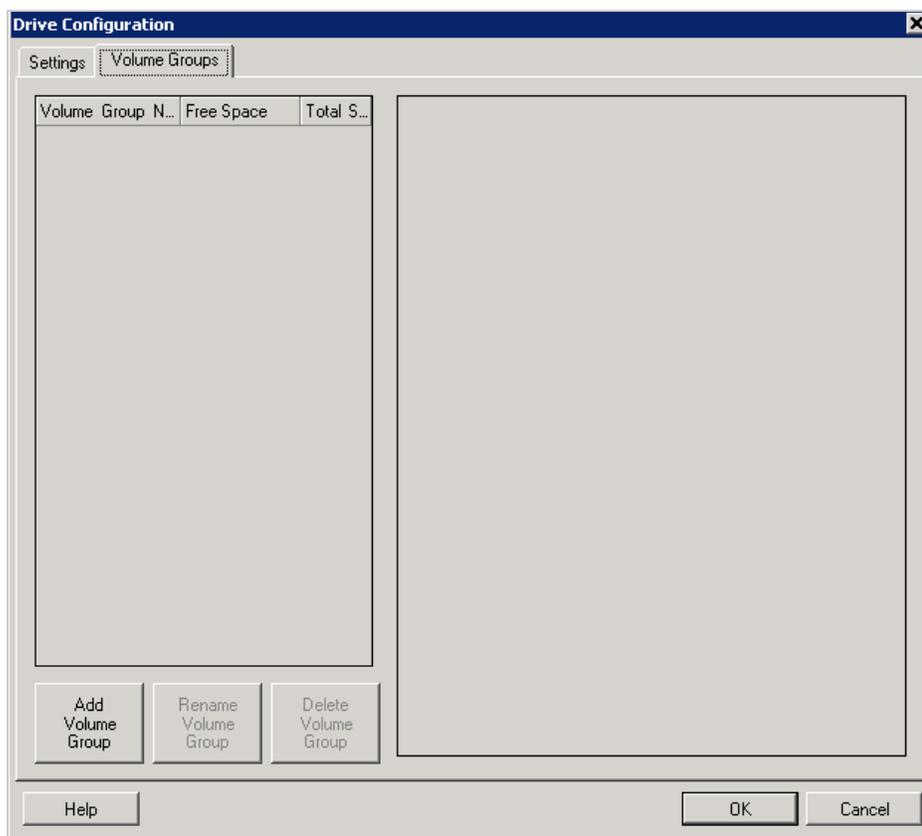
New Size: To resize the volume during the migration, enter the desired size. PlateSpin Migrate automatically adjusts **New Free Space**.

Disk/Volume Group: Select which hard drive or volume group the volume will be copied to on the physical target machine.

Preserve Partitions: For each disk, click the corresponding cell in this column to select existing vendor partitions to preserve during the migration. If the partitions are not selected, PlateSpin Migrate permanently removes them from the server.

Linux Drive and LVM Volume Configuration (Volume Groups Tab)

Use these settings to manage volume groups.



Add Volume Group: Creates a volume group on the target machine that is not present on the source machine.

Rename Volume Group: Renames a volume group that is being copied from the source to the target.

Delete Volume Group: Deletes a volume group so that it is not created on the target machine. The volumes assigned to the volume group can be reassigned to other locations by using the **Settings** tab (by default, they are assigned to disk).

Allocation for Volume Group: To allocate space on disks to a volume group, select the volume group, then select the disks to include in it. Specify the amount of space to be allocated to it on each included disk.

Target VM-Specific P2V/V2V Drive Configuration

When you configure a peer-to-peer virtualization job, the job configuration window provides access to settings specific to the target virtualization platform.

This section includes the following information:

- “Drive Configuration: VMware ESX” on page 175
- “Drive Configuration: Hyper-V” on page 176

Drive Configuration: VMware ESX

The following are drive configuration settings specific to VMware ESX:

Virtual disks to create:

Disk	Datastore	Size	File Name
Virtual disk 0	storage1	4 GB	/NY-RHEL4-LVM_VM/NY-RHEL4-LVM_VM_1.vmdk

Select volumes to copy and size:

Include	Volume	Free Space	Size	New Free Space	New Size	Disk/Volume Group
<input checked="" type="checkbox"/>	/	2.4 GB	3.2 GB	2.4 GB	3.2 GB	VolGroup00
<input checked="" type="checkbox"/>	/boot	81.2 MB	98.7 MB	81.2 MB	98.7 MB	Disk 0
<input checked="" type="checkbox"/>	/home	88.1 MB	98.7 MB	88.1 MB	98.7 MB	Disk 0

Select non-volume storage to recreate and size:

Include	Type	Partition	Size	Is Swap	Disk/Volu...	New Size
<input checked="" type="checkbox"/>		/dev/VolGroup00/Lo...	512 MB	<input checked="" type="checkbox"/>	Vol...	512 MB

Datastore: Select the datastore volume on the ESX server where you want to place the vmdk files.

Copy: Select the volumes to be copied during the migration.

New Free Space: To resize the volume during the migration, specify the desired amount of free space. PlateSpin Migrate automatically adjusts New Size.

New Size: To resize the volume during the migration, specify the desired size. PlateSpin Migrate automatically adjusts New Free Space.

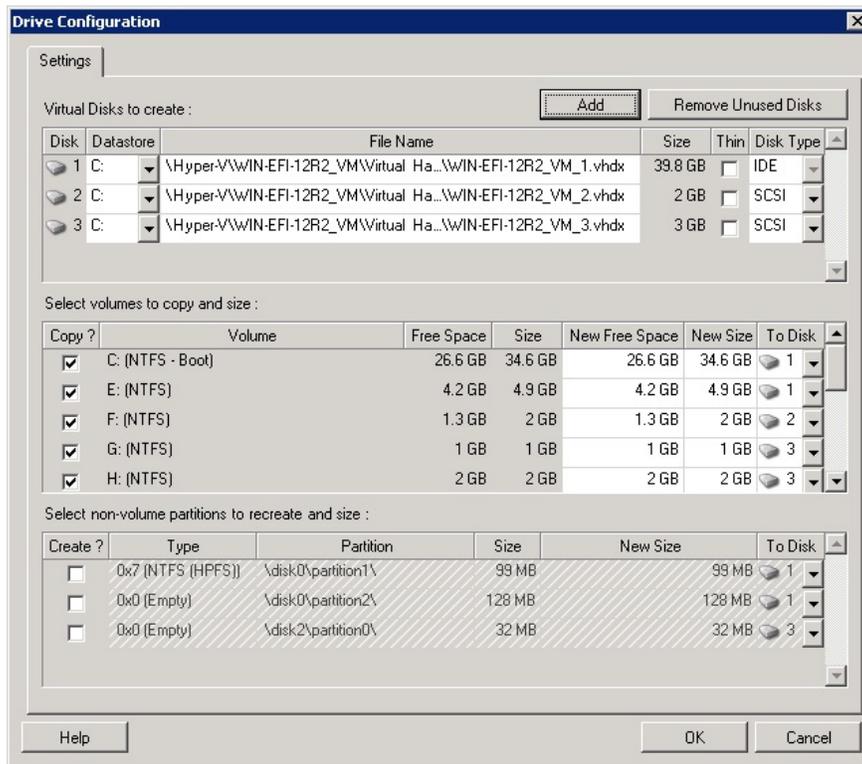
Disk/Volume Group: Assign the volume to a disk or, if LVM is enabled, to a volume group. The volume will be copied to this disk or volume group on the target machine.

Create: Select any non-volume disk partitions that should be created on the target machine (for example, a Linux swap partition).

New Size: To resize the non-volume partition during the migration, specify the desired size.

Drive Configuration: Hyper-V

The following are drive configuration settings specific to Hyper-V:



Datastore: Select the datastore volume on the Hyper-V server where you want to place the .vhd and .vhdx files.

Disk Type: A Generation 1 disk containing the System/Boot volume should be on an IDE disk. (You can create a maximum of three IDE disks.)

NOTE: For a Generation 1 disk, the values of second and third disk are chained. For example, if you select the third disk (from the top of the **Disk Type** list) as **IDE**, the second disk autoselects as **IDE**. If you select the second disk as a **SCSI** then the third disk autoselects to **SCSI**.

Copy?: Select the volumes to be copied during the migration.

New Free Space: To resize the volume during the migration, specify the desired amount of free space. PlateSpin Migrate automatically adjusts New Size.

New Size: To resize the volume during the migration, specify the desired size. PlateSpin Migrate automatically adjusts New Free Space.

To Disk: Assign the volume to a disk or, if LVM is enabled, to a volume group. The volume is copied to this disk or volume group on the target machine.

Create?: Select any non-volume disk partitions that should be created on the target machine (for example, a Linux swap partition).

New Size: To resize the non-volume partition during the migration, specify the desired size.

7.12.2 Volume Configuration in Server Sync

When you are using Server Sync to synchronize two Windows or Linux workloads, PlateSpin Migrate provides you with the capability to specify the required mapping between source volumes and existing volumes on the target. See [“Synchronizing Workloads with Server Sync” on page 119](#).

To access volume configuration options in a Server Sync job:

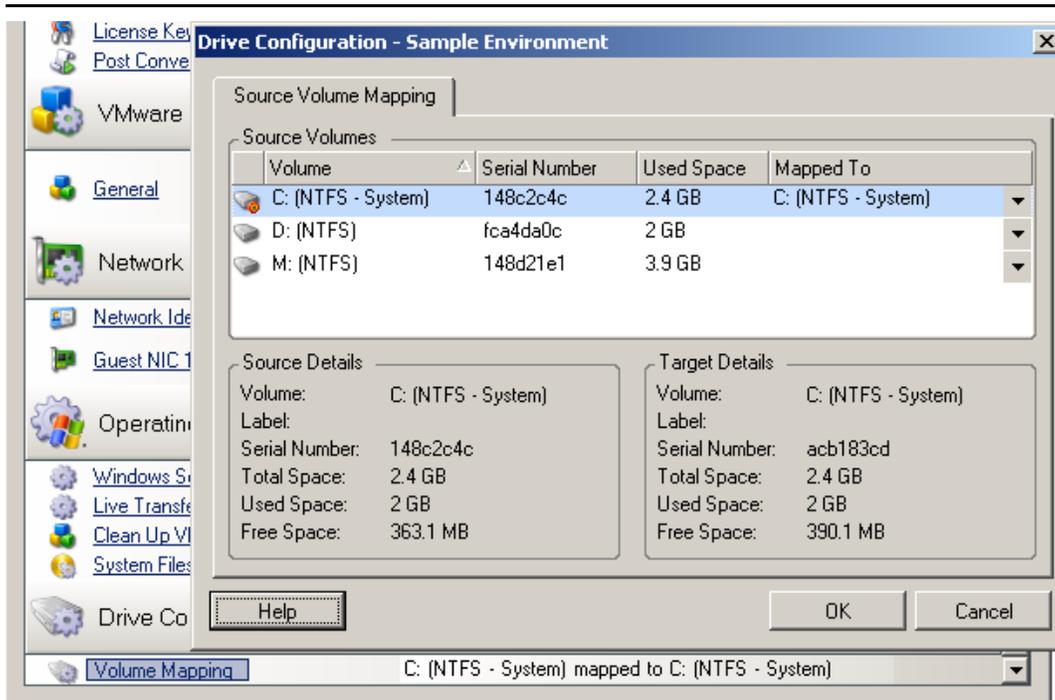
- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,” on page 83](#).
- 2 In the **Drive Configuration** section of the Migration Job window, do one of the following:
 - ♦ **On Windows:** Click **Volume Mapping**.
 - ♦ **On Linux:** Click **Drives and Volumes**.
- 3 Configure the Server Sync volume configuration options.

The following topics provide information about Server Sync volume configuration options specific to Windows and Linux workloads.

- ♦ [“Server Sync Volume Configuration \(Windows\)” on page 177](#)
- ♦ [“Server Sync Volume Configuration \(Linux\)” on page 178](#)

Server Sync Volume Configuration (Windows)

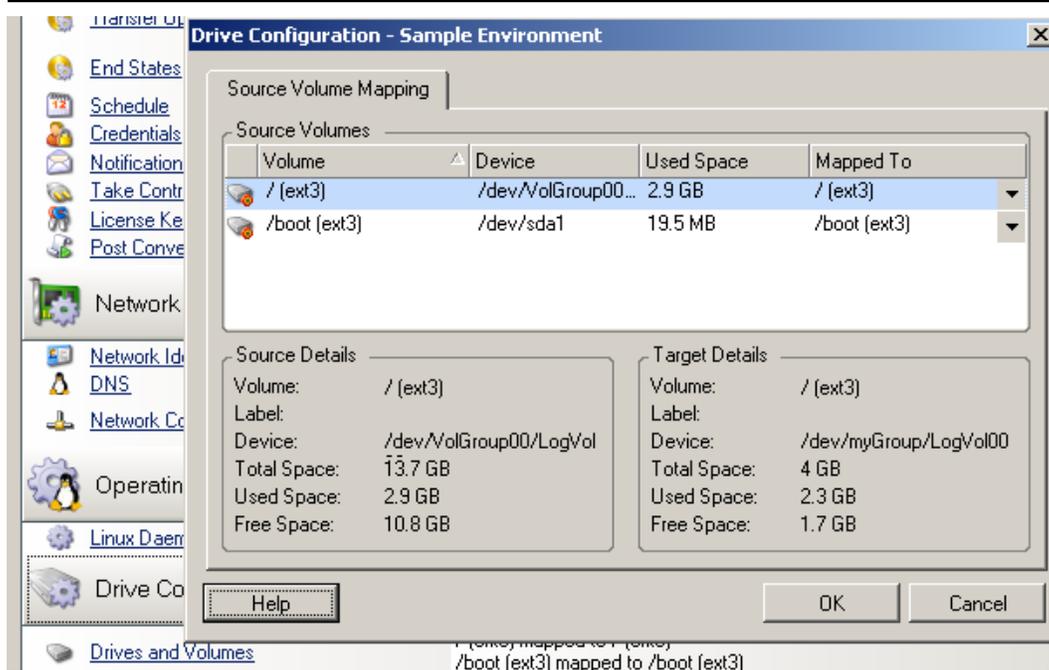
A Server Sync job for Windows workloads provides detailed drive and volume information for both the source and the target, and enables you to specify the required mapping.



Mapped To: Map each volume on the source to an existing volume on the target.

Server Sync Volume Configuration (Linux)

A Server Sync job for Linux workloads provides detailed mount point and volume information for both the source and the target, and enables you to specify the required mapping.



Mapped To: Map each volume on the source to an existing volume on the target.

7.13 Including a Custom Post-migration Action in a Migration Job

You can set up your migration job to execute a custom action on your target. You must define and save your custom actions and their dependencies in advance. See [“Managing Custom Actions” on page 80](#).

NOTE: Post-migration actions are supported for peer-to-peer and one-time Server Sync migrations only.

When you are setting up a migration job, select the required action, any required command line parameters, and a timeout as required. You must also provide valid credentials for the target workload. If the target workload credentials are unknown, you can use the credentials of the source workload.

To specify a custom post-migration action for your migration job:

- 1 Start the migration job. For information about starting a migration job, see [Section 3.8.1, “Setting a Migration Job,” on page 83](#).
- 2 In the Virtual Machine Configuration section of the Migration Job window, click **Post Conversion**.

Select Action: IPconfig_X2V

Execution Parameters

Command Line:

Execution Timeout: 5 seconds

No Timeout

Credentials

Username:

Password:

Use Source Credentials

3 Specify the following options:

- ♦ **Select Action:** From the drop-down list, select a custom action previously saved in your library of post-migration actions.
- ♦ **Execution Parameters:** Specify any required command line parameters for the action. If required, specify a timeout.
- ♦ **Credentials:** Provide administrator credentials for the target. If they are the same as those for the source, and if they have been saved, select **Use Source Credentials**.

8 MigrateAgent Utility

The MigrateAgent (`MigrateAgent.cli.exe`) is a command line utility that you can use to install, upgrade, query, or uninstall the block-based transfer drivers. Although a reboot is always required when you install, uninstall, or upgrade drivers, the MigrateAgent utility allows you to better control when the action occurs and therefore, when the server reboots. For example, you can use the MigrateAgent utility to install the drivers during scheduled down time, instead of during the first replication.

The syntax of the MigrateAgent utility is:

```
MigrateAgent.cli.exe [Option] [/psserver=%IP%]
```

Table 8-1 describes the options and switch available for the `MigrateAgent.cli.exe` command.

Table 8-1 *MigrateAgentCommand Options and Switch*

Usage	Description
Options	
<code>h ? help</code>	Displays usage and options for the command.
<code>logs view-logs</code>	Opens the application log directory.
<code>status</code>	Shows installation status for the PlateSpin controller and drivers.
<code>din driver-install</code>	Installs the PlateSpin drivers.
<code>dup driver-upgrade</code>	Upgrades the PlateSpin drivers.
<code>dun driver-uninstall</code>	Uninstalls the PlateSpin drivers.
Switch	
<code>/psserver=%IP%</code>	Downloads the block-based transfer drivers from the specified server when you invoke the <code>status</code> , <code>driver-install</code> , or <code>driver-upgrade</code> options.

A copy of the block-based transfer drivers is bundled with the MigrateAgent utility. You can alternatively specify the `/psserver=` command line switch in order to download the drivers from the PlateSpin Server when you invoke the `status`, `driver-install`, or `driver-upgrade` options. This is useful when the server is patched with a new driver package, but the MigrateAgent command line utility is not patched.

NOTE: To avoid confusion, the recommended method of using the MigrateAgent is to install, uninstall, or upgrade the drivers and then reboot prior to doing a replication.

You should reboot the system each time you install, upgrade, or uninstall the drivers. The reboot forces the running driver to stop and the new driver to be applied on system restart. If you do not reboot the system prior to replication, the source continues to act as if the operation has not been completed. For example, if you install drivers without rebooting the system, the source acts as if no driver is installed during replication. Similarly, if you upgrade the drivers without rebooting, the source continues to use the already running driver during replication until you reboot the system.

If the version of the installed driver is different than the version of the running driver, the status option will remind the user to reboot. For example:

```
C:\MigrateAgent\MigrateAgent.cli.exe status
Step 1 of 2: Querying the PlateSpin controller service
  Done
Step 2 of 2: Querying the installed PlateSpin driver version
  Done

The task completed successfully
PlateSpin Controller Service Status
  The PlateSpin Controller service is not installed

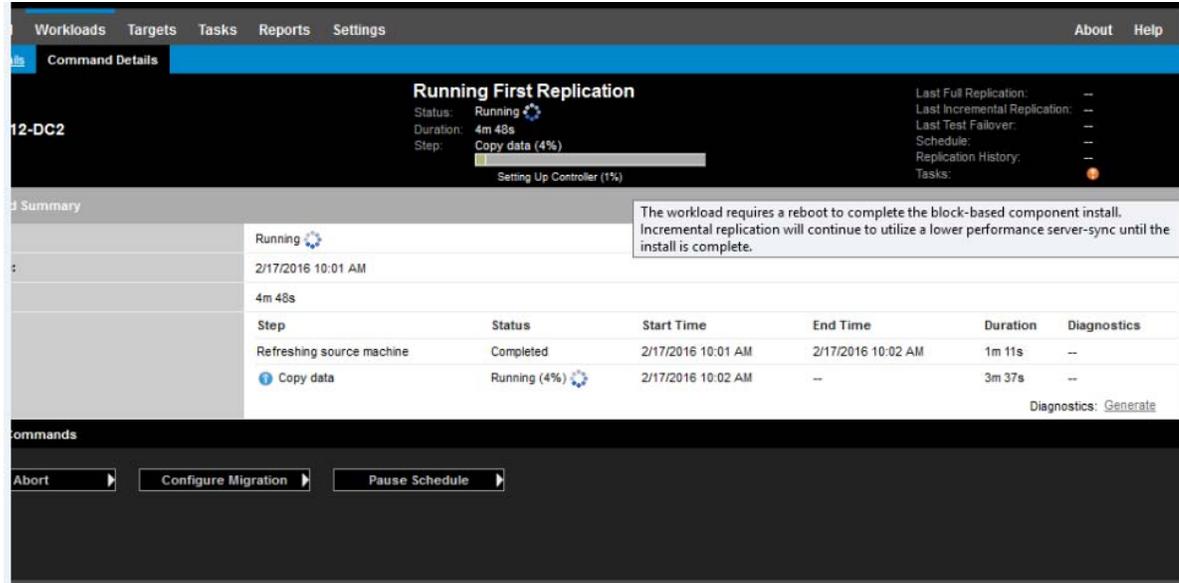
PlateSpin Driver Status
  Installed Driver Version: 8.0.0.11
  Running Driver Version: Not running. Reboot to load the driver.
  Upgrade Available: No
```

PlateSpin creates a task to warn the user that a reboot is necessary in order to complete the driver installation or upgrade. The notification appears in the Tasks list (Figure 8-1). During replication, the notification appears on the Command Details page (Figure 8-2).

Figure 8-1 Reboot Notification Task



Figure 8-2 Reboot Notification During Replication



Rebooting the source machine applies and starts the installed or upgraded drivers. If the driver was recently installed, after the reboot, one full replication or a server-sync replication is required in order to ensure that all of a source's changes are captured. This server-sync replication will be represented to the user in the Status field as a warning (Figure 8-3). Subsequent incremental replications will complete without warning.

Figure 8-3 Server-Sync Required Notification

Dashboard Workloads Targets Tasks Reports Settings About Help

Migration Details Command Details

Win12-DC2

Running Incremental
 Status: ● Running
 Duration: 8m 29s
 Step: Copy data (28%)
 Copying Volume Data from Source to Target (39%)

Last Full Replication: 2/17/2016 10:25 AM
 Last Incremental Replication: --
 Last Test Failover: --
 Schedule: --
 Replication History: [View](#)
 Tasks: --

Command Summary

Events:	Event	Details	User	Date
	Incremental replication of workload started		MIGRATEAUTO/build	2/18/2016 1:25 AM

Status: ● Running
 As the block-based driver was just installed, this first replication is performed via a server-sync and may take more time. All subsequent replications will be performed using the block-based driver.

Start Time: 2/18/2016 1:25 AM
 Duration: 8m 29s

Steps:	Step	Status	Start Time	End Time	Duration	Diagnostics
	Refreshing source machine	Completed	2/18/2016 1:25 AM	2/18/2016 1:27 AM	1m 22s	--
	Revert to snapshot	Completed	2/18/2016 1:27 AM	2/18/2016 1:28 AM	40s	--
	Copy data	● Running (28%)	2/18/2016 1:28 AM	--	6m 27s	--

Diagnostics: [Generate](#)

Replication Transfer Summary

Duration:	1m 31s
Total Data Transferred:	27.9 MB
Total Files Transferred:	463

Workload Commands

A Linux Distributions Supported by Migrate

PlateSpin Migrate software includes pre-compiled versions of the `blkwatch` driver for many non-debug Linux distributions (32-bit and 64-bit). This section includes the following information:

- ♦ [Section A.1, “Analyzing Your Linux Workload,” on page 185](#)
- ♦ [Section A.2, “Pre-compiled “blkwatch” Driver Support \(Linux Distros\),” on page 186](#)

A.1 Analyzing Your Linux Workload

Prior to determining whether PlateSpin Migrate has a `blkwatch` driver for your distribution, you need to learn more about the kernel of your Linux workload so that you can use it as a search term against the list of supported distributions. This section includes the following information:

- ♦ [Section A.1.1, “Determining the Release String,” on page 185](#)
- ♦ [Section A.1.2, “Determining the Architecture,” on page 185](#)

A.1.1 Determining the Release String

You can determine the release string of the kernel of your Linux workload by running the following command at the workload’s Linux terminal:

```
uname -r
```

For example, if you run `uname -r`, you might see the following output:

```
3.0.76-0.11-default
```

If you search the list of distributions, you see there are two entries that match this string:

- ♦ `SLES11SP3-GA-3.0.76-0.11-default-x86`
- ♦ `SLES11SP3-GA-3.0.76-0.11-default-x86_64`

The search results indicate that the product has drivers for both 32-bit (x86) and 64-bit (x86_64) architectures.

A.1.2 Determining the Architecture

You can determine the architecture of your Linux workload by running the following command at the workload’s Linux terminal:

```
uname -m
```

For example, if you run `uname -m`, you might see the following output:

```
x86_64
```

With this information, you can determine that the workload has 64-bit architecture.

A.2 Pre-compiled “blkwatch” Driver Support (Linux Distros)

Following is a list of non-debug Linux distributions for which Migrate has a `blkwatch` driver. You can search the list to determine if the release string and architecture of your Linux workload kernel matches a supported distribution in the list. If you find your release string and architecture, PlateSpin Migrate has a pre-compiled version the `blkwatch` driver.

If your search is unsuccessful, you can create a custom `blkwatch` driver by following the steps found in the Knowledgebase Article [KB 7005873](#).

List Item Syntax

Each item in the list is formatted using the following syntax:

```
<Distro>-<Patch>-<Kernel_Release_String>-<Kernel_Architecture>
```

So, for a SLES 9 SP1 distribution with a kernel release string of `2.6.5-7.139-bigsm` for 32-bit (x86) architecture, the item is listed in a format like this:

```
SLES9-SP1-2.6.5-7.139-bigsm-x86
```

List of Distributions

```
RHEL4-GA-2.6.9-5.EL-x86
RHEL4-GA-2.6.9-5.EL-x86_64
RHEL4-GA-2.6.9-5.ELhugemem-x86
RHEL4-GA-2.6.9-5.ELsmp-x86
RHEL4-GA-2.6.9-5.ELsmp-x86_64
RHEL4-U1-2.6.9-11.EL-x86
RHEL4-U1-2.6.9-11.EL-x86_64
RHEL4-U1-2.6.9-11.ELhugemem-x86
RHEL4-U1-2.6.9-11.ELsmp-x86
RHEL4-U1-2.6.9-11.ELsmp-x86_64
RHEL4-U2-2.6.9-22.EL-x86
RHEL4-U2-2.6.9-22.EL-x86_64
RHEL4-U2-2.6.9-22.ELhugemem-x86
RHEL4-U2-2.6.9-22.ELsmp-x86
RHEL4-U2-2.6.9-22.ELsmp-x86_64
RHEL4-U3-2.6.9-34.EL-x86
RHEL4-U3-2.6.9-34.EL-x86_64
RHEL4-U3-2.6.9-34.ELhugemem-x86
RHEL4-U3-2.6.9-34.ELlargesmp-x86_64
RHEL4-U3-2.6.9-34.ELsmp-x86
RHEL4-U3-2.6.9-34.ELsmp-x86_64
RHEL4-U4-2.6.9-42.EL-x86
RHEL4-U4-2.6.9-42.EL-x86_64
RHEL4-U4-2.6.9-42.ELhugemem-x86
RHEL4-U4-2.6.9-42.ELlargesmp-x86_64
RHEL4-U4-2.6.9-42.ELsmp-x86
RHEL4-U4-2.6.9-42.ELsmp-x86_64
RHEL4-U5-2.6.9-55.EL-x86
RHEL4-U5-2.6.9-55.EL-x86_64
RHEL4-U5-2.6.9-55.ELhugemem-x86
RHEL4-U5-2.6.9-55.ELlargesmp-x86_64
RHEL4-U5-2.6.9-55.ELsmp-x86
RHEL4-U5-2.6.9-55.ELsmp-x86_64
RHEL4-U6-2.6.9-67.EL-x86
```

RHEL4-U6-2.6.9-67.EL-x86_64
RHEL4-U6-2.6.9-67.ELhugemem-x86
RHEL4-U6-2.6.9-67.ELlargesmp-x86_64
RHEL4-U6-2.6.9-67.ELsmp-x86
RHEL4-U6-2.6.9-67.ELsmp-x86_64
RHEL4-U7-2.6.9-78.EL-x86
RHEL4-U7-2.6.9-78.EL-x86_64
RHEL4-U7-2.6.9-78.ELhugemem-x86
RHEL4-U7-2.6.9-78.ELlargesmp-x86_64
RHEL4-U7-2.6.9-78.ELsmp-x86
RHEL4-U7-2.6.9-78.ELsmp-x86_64
RHEL4-U8-2.6.9-89.EL-x86
RHEL4-U8-2.6.9-89.EL-x86_64
RHEL4-U8-2.6.9-89.ELhugemem-x86
RHEL4-U8-2.6.9-89.ELlargesmp-x86_64
RHEL4-U8-2.6.9-89.ELsmp-x86
RHEL4-U8-2.6.9-89.ELsmp-x86_64
RHEL4-U9-2.6.9-100.EL-x86
RHEL4-U9-2.6.9-100.EL-x86_64
RHEL4-U9-2.6.9-100.ELhugemem-x86
RHEL4-U9-2.6.9-100.ELlargesmp-x86_64
RHEL4-U9-2.6.9-100.ELsmp-x86
RHEL4-U9-2.6.9-100.ELsmp-x86_64
RHEL5-GA-2.6.18-8.el5-x86
RHEL5-GA-2.6.18-8.el5-x86_64
RHEL5-GA-2.6.18-8.el5PAE-x86
RHEL5-U1-2.6.18-53.el5-x86
RHEL5-U1-2.6.18-53.el5-x86_64
RHEL5-U1-2.6.18-53.el5PAE-x86
RHEL5-U10-2.6.18-371.el5-x86
RHEL5-U10-2.6.18-371.el5-x86_64
RHEL5-U10-2.6.18-371.el5PAE-x86
RHEL5-U2-2.6.18-92.el5-x86
RHEL5-U2-2.6.18-92.el5-x86_64
RHEL5-U2-2.6.18-92.el5PAE-x86
RHEL5-U3-2.6.18-128.el5-x86
RHEL5-U3-2.6.18-128.el5-x86_64
RHEL5-U3-2.6.18-128.el5PAE-x86
RHEL5-U4-2.6.18-164.el5-x86
RHEL5-U4-2.6.18-164.el5-x86_64
RHEL5-U4-2.6.18-164.el5PAE-x86
RHEL5-U5-2.6.18-194.el5-x86
RHEL5-U5-2.6.18-194.el5-x86_64
RHEL5-U5-2.6.18-194.el5PAE-x86
RHEL5-U6-2.6.18-238.el5-x86
RHEL5-U6-2.6.18-238.el5-x86_64
RHEL5-U6-2.6.18-238.el5PAE-x86
RHEL5-U7-2.6.18-274.el5-x86
RHEL5-U7-2.6.18-274.el5-x86_64
RHEL5-U7-2.6.18-274.el5PAE-x86
RHEL5-U8-2.6.18-308.el5-x86
RHEL5-U8-2.6.18-308.el5-x86_64
RHEL5-U8-2.6.18-308.el5PAE-x86
RHEL5-U9-2.6.18-348.el5-x86
RHEL5-U9-2.6.18-348.el5-x86_64
RHEL5-U9-2.6.18-348.el5PAE-x86
RHEL5-U11-2.6.18-398.el5-x86
RHEL5-U11-2.6.18-398.el5-x86_64
RHEL5-U11-2.6.18-398.el5PAE-x86
RHEL6-GA-2.6.32-71.el6.i686-x86

RHEL6-GA-2.6.32-71.el6.x86_64-x86_64
RHEL6-U1-2.6.32-131.0.15.el6.i686-x86
RHEL6-U1-2.6.32-131.0.15.el6.x86_64-x86_64
RHEL6-U2-2.6.32-220.el6.i686-x86
RHEL6-U2-2.6.32-220.el6.x86_64-x86_64
RHEL6-U3-2.6.32-279.el6.i686-x86
RHEL6-U3-2.6.32-279.el6.x86_64-x86_64
RHEL6-U4-2.6.32-358.el6.i686-x86
RHEL6-U4-2.6.32-358.el6.x86_64-x86_64
RHEL6-U5-2.6.32-431.el6.i686-x86
RHEL6-U5-2.6.32-431.el6.x86_64-x86_64
RHEL6-U6-2.6.32-504.el6.i686-x86
RHEL6-U6-2.6.32-504.el6.x86_64-x86_64
RHEL6-U6-2.6.32-504.el6PAE-x86
RHEL6-U7-2.6.32-573.el6.i686-x86
RHEL6-U7-2.6.32-573.el6.x86_64-x86_64
RHEL7-GA-3.10.0-123.el7.x86_64-x86_64
RHEL7-U1-3.10.0-229.el7.x86_64-x86_64
RHEL7-U2-3.10.0-327.el7.x86_64-x86_64
SLES10-GA-2.6.16.21-0.8-bigsmp-x86
SLES10-GA-2.6.16.21-0.8-default-x86
SLES10-GA-2.6.16.21-0.8-default-x86_64
SLES10-GA-2.6.16.21-0.8-smp-x86
SLES10-GA-2.6.16.21-0.8-smp-x86_64
SLES10-GA-2.6.16.21-0.8-xen-x86
SLES10-GA-2.6.16.21-0.8-xen-x86_64
SLES10-GA-2.6.16.21-0.8-xenpae-x86
SLES10-SP1-2.6.16.46-0.12-bigsmp-x86
SLES10-SP1-2.6.16.46-0.12-default-x86
SLES10-SP1-2.6.16.46-0.12-default-x86_64
SLES10-SP1-2.6.16.46-0.12-smp-x86
SLES10-SP1-2.6.16.46-0.12-smp-x86_64
SLES10-SP1-2.6.16.46-0.12-xen-x86
SLES10-SP1-2.6.16.46-0.12-xen-x86_64
SLES10-SP1-2.6.16.46-0.12-xenpae-x86
SLES10-SP2-2.6.16.60-0.21-bigsmp-x86
SLES10-SP2-2.6.16.60-0.21-default-x86
SLES10-SP2-2.6.16.60-0.21-default-x86_64
SLES10-SP2-2.6.16.60-0.21-smp-x86
SLES10-SP2-2.6.16.60-0.21-smp-x86_64
SLES10-SP2-2.6.16.60-0.21-xen-x86
SLES10-SP2-2.6.16.60-0.21-xen-x86_64
SLES10-SP2-2.6.16.60-0.21-xenpae-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-bigsmp-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-default-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-default-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-smp-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-smp-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xen-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xen-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xenpae-x86
SLES10-SP3-2.6.16.60-0.54.5-bigsmp-x86
SLES10-SP3-2.6.16.60-0.54.5-default-x86
SLES10-SP3-2.6.16.60-0.54.5-default-x86_64
SLES10-SP3-2.6.16.60-0.54.5-smp-x86
SLES10-SP3-2.6.16.60-0.54.5-smp-x86_64
SLES10-SP3-2.6.16.60-0.54.5-xen-x86
SLES10-SP3-2.6.16.60-0.54.5-xen-x86_64
SLES10-SP3-2.6.16.60-0.54.5-xenpae-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-bigsmp-x86

SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-default-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-default-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-smp-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-smp-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xen-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xen-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xenpae-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-bigsmp-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-default-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-default-x86_64
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-smp-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-smp-x86_64
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-xen-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-xen-x86_64
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-xenpae-x86
SLES10-SP4-2.6.16.60-0.85.1-bigsmp-x86
SLES10-SP4-2.6.16.60-0.85.1-default-x86
SLES10-SP4-2.6.16.60-0.85.1-default-x86_64
SLES10-SP4-2.6.16.60-0.85.1-smp-x86
SLES10-SP4-2.6.16.60-0.85.1-smp-x86_64
SLES10-SP4-2.6.16.60-0.85.1-xen-x86
SLES10-SP4-2.6.16.60-0.85.1-xen-x86_64
SLES10-SP4-2.6.16.60-0.85.1-xenpae-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-bigsmp-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-default-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-default-x86_64
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-smp-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-smp-x86_64
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-xen-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-xen-x86_64
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-xenpae-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-bigsmp-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-default-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-default-x86_64
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-smp-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-smp-x86_64
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-xen-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-xen-x86_64
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-xenpae-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-bigsmp-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-default-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-default-x86_64
SLES10-SP4_U4-2.6.16.60-0.93.1-smp-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-smp-x86_64
SLES10-SP4_U4-2.6.16.60-0.93.1-xen-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-xen-x86_64
SLES10-SP4_U4-2.6.16.60-0.93.1-xenpae-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-bigsmp-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-default-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-default-x86_64
SLES10-SP4_U5-2.6.16.60-0.97.1-smp-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-smp-x86_64
SLES10-SP4_U5-2.6.16.60-0.97.1-xen-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-xen-x86_64
SLES10-SP4_U5-2.6.16.60-0.97.1-xenpae-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-bigsmp-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-default-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-default-x86_64
SLES10-SP4_U6-2.6.16.60-0.99.1-smp-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-smp-x86_64

SLES10-SP4_U6-2.6.16.60-0.99.1-xen-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-xen-x86_64
SLES10-SP4_U6-2.6.16.60-0.99.1-xenpae-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-bigsmp-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-default-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-default-x86_64
SLES10-SP4_U7-2.6.16.60-0.101.1-smp-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-smp-x86_64
SLES10-SP4_U7-2.6.16.60-0.101.1-xen-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-xen-x86_64
SLES10-SP4_U7-2.6.16.60-0.101.1-xenpae-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-bigsmp-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-default-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-default-x86_64
SLES10-SP4_U8-2.6.16.60-0.103.1-smp-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-smp-x86_64
SLES10-SP4_U8-2.6.16.60-0.103.1-xen-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-xen-x86_64
SLES10-SP4_U8-2.6.16.60-0.103.1-xenpae-x86
SLES11-GA-2.6.27.19-5-default-x86
SLES11-GA-2.6.27.19-5-default-x86_64
SLES11-GA-2.6.27.19-5-pae-x86
SLES11-SP1-2.6.32.12-0.6-default-x86
SLES11-SP1-2.6.32.12-0.6-default-x86_64
SLES11-SP1-2.6.32.12-0.6-pae-x86
SLES11-SP1_LTSS_U1-2.6.32.59-0.9-default-x86
SLES11-SP1_LTSS_U1-2.6.32.59-0.9-default-x86_64
SLES11-SP1_LTSS_U1-2.6.32.59-0.9-pae-x86
SLES11-SP1_LTSS_U2-2.6.32.59-0.13-default-x86
SLES11-SP1_LTSS_U2-2.6.32.59-0.13-default-x86_64
SLES11-SP1_LTSS_U2-2.6.32.59-0.13-pae-x86
SLES11-SP1_U14-2.6.32.54-0.3-default-x86
SLES11-SP1_U14-2.6.32.54-0.3-default-x86_64
SLES11-SP1_U14-2.6.32.54-0.3-pae-x86
SLES11-SP1_U15-2.6.32.59-0.3-default-x86
SLES11-SP1_U15-2.6.32.59-0.3-default-x86_64
SLES11-SP1_U15-2.6.32.59-0.3-pae-x86
SLES11-SP1_U16-2.6.32.59-0.7-default-x86
SLES11-SP1_U16-2.6.32.59-0.7-default-x86_64
SLES11-SP1_U16-2.6.32.59-0.7-pae-x86
SLES11SP2-GA-3.0.13-0.27-default-x86
SLES11SP2-GA-3.0.13-0.27-default-x86_64
SLES11SP2-GA-3.0.13-0.27-pae-x86
SLES11SP2-GA-3.0.13-0.27-xen-x86
SLES11SP2-GA-3.0.13-0.27-xen-x86_64
SLES11SP2-LTSS_U1-3.0.101-0.7.19-default-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-default-x86_64
SLES11SP2-LTSS_U1-3.0.101-0.7.19-pae-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-xen-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-xen-x86_64
SLES11SP2-LTSS_U2-3.0.101-0.7.21-default-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-default-x86_64
SLES11SP2-LTSS_U2-3.0.101-0.7.21-pae-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-xen-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-xen-x86_64
SLES11SP2-U1-3.0.26-0.7-default-x86
SLES11SP2-U1-3.0.26-0.7-default-x86_64
SLES11SP2-U1-3.0.26-0.7-pae-x86
SLES11SP2-U1-3.0.26-0.7-xen-x86
SLES11SP2-U1-3.0.26-0.7-xen-x86_64

SLES11SP2-U10-3.0.74-0.6.8-default-x86
SLES11SP2-U10-3.0.74-0.6.8-default-x86_64
SLES11SP2-U10-3.0.74-0.6.8-pae-x86
SLES11SP2-U10-3.0.74-0.6.8-xen-x86
SLES11SP2-U10-3.0.74-0.6.8-xen-x86_64
SLES11SP2-U11-3.0.74-0.6.10-default-x86
SLES11SP2-U11-3.0.74-0.6.10-default-x86_64
SLES11SP2-U11-3.0.74-0.6.10-pae-x86
SLES11SP2-U11-3.0.74-0.6.10-xen-x86
SLES11SP2-U11-3.0.74-0.6.10-xen-x86_64
SLES11SP2-U12-3.0.80-0.5-default-x86
SLES11SP2-U12-3.0.80-0.5-default-x86_64
SLES11SP2-U12-3.0.80-0.5-pae-x86
SLES11SP2-U12-3.0.80-0.5-xen-x86
SLES11SP2-U12-3.0.80-0.5-xen-x86_64
SLES11SP2-U13-3.0.80-0.7-default-x86
SLES11SP2-U13-3.0.80-0.7-default-x86_64
SLES11SP2-U13-3.0.80-0.7-pae-x86
SLES11SP2-U13-3.0.80-0.7-xen-x86
SLES11SP2-U13-3.0.80-0.7-xen-x86_64
SLES11SP2-U14-3.0.93-0.5-default-x86
SLES11SP2-U14-3.0.93-0.5-default-x86_64
SLES11SP2-U14-3.0.93-0.5-pae-x86
SLES11SP2-U14-3.0.93-0.5-xen-x86
SLES11SP2-U14-3.0.93-0.5-xen-x86_64
SLES11SP2-U15-3.0.101-0.5-default-x86
SLES11SP2-U15-3.0.101-0.5-default-x86_64
SLES11SP2-U15-3.0.101-0.5-pae-x86
SLES11SP2-U15-3.0.101-0.5-xen-x86
SLES11SP2-U15-3.0.101-0.5-xen-x86_64
SLES11SP2-U16-3.0.101-0.7.15-default-x86
SLES11SP2-U16-3.0.101-0.7.15-default-x86_64
SLES11SP2-U16-3.0.101-0.7.15-pae-x86
SLES11SP2-U16-3.0.101-0.7.15-xen-x86
SLES11SP2-U16-3.0.101-0.7.15-xen-x86_64
SLES11SP2-U17-3.0.101-0.7.17-default-x86
SLES11SP2-U17-3.0.101-0.7.17-default-x86_64
SLES11SP2-U17-3.0.101-0.7.17-pae-x86
SLES11SP2-U17-3.0.101-0.7.17-xen-x86
SLES11SP2-U17-3.0.101-0.7.17-xen-x86_64
SLES11SP2-U2-3.0.31-0.9-default-x86
SLES11SP2-U2-3.0.31-0.9-default-x86_64
SLES11SP2-U2-3.0.31-0.9-pae-x86
SLES11SP2-U2-3.0.31-0.9-xen-x86
SLES11SP2-U2-3.0.31-0.9-xen-x86_64
SLES11SP2-U3-3.0.34-0.7-default-x86
SLES11SP2-U3-3.0.34-0.7-default-x86_64
SLES11SP2-U3-3.0.34-0.7-pae-x86
SLES11SP2-U3-3.0.34-0.7-xen-x86
SLES11SP2-U3-3.0.34-0.7-xen-x86_64
SLES11SP2-U4-3.0.38-0.5-default-x86
SLES11SP2-U4-3.0.38-0.5-default-x86_64
SLES11SP2-U4-3.0.38-0.5-pae-x86
SLES11SP2-U4-3.0.38-0.5-xen-x86
SLES11SP2-U4-3.0.38-0.5-xen-x86_64
SLES11SP2-U5-3.0.42-0.7-default-x86
SLES11SP2-U5-3.0.42-0.7-default-x86_64
SLES11SP2-U5-3.0.42-0.7-pae-x86
SLES11SP2-U5-3.0.42-0.7-xen-x86
SLES11SP2-U5-3.0.42-0.7-xen-x86_64

SLES11SP2-U6-3.0.51-0.7.9-default-x86
SLES11SP2-U6-3.0.51-0.7.9-default-x86_64
SLES11SP2-U6-3.0.51-0.7.9-pae-x86
SLES11SP2-U6-3.0.51-0.7.9-xen-x86
SLES11SP2-U6-3.0.51-0.7.9-xen-x86_64
SLES11SP2-U7-3.0.58-0.6.2-default-x86
SLES11SP2-U7-3.0.58-0.6.2-default-x86_64
SLES11SP2-U7-3.0.58-0.6.2-pae-x86
SLES11SP2-U7-3.0.58-0.6.2-xen-x86
SLES11SP2-U7-3.0.58-0.6.2-xen-x86_64
SLES11SP2-U8-3.0.58-0.6.6-default-x86
SLES11SP2-U8-3.0.58-0.6.6-default-x86_64
SLES11SP2-U8-3.0.58-0.6.6-pae-x86
SLES11SP2-U8-3.0.58-0.6.6-xen-x86
SLES11SP2-U8-3.0.58-0.6.6-xen-x86_64
SLES11SP2-U9-3.0.74-0.6.6-default-x86
SLES11SP2-U9-3.0.74-0.6.6-default-x86_64
SLES11SP2-U9-3.0.74-0.6.6-pae-x86
SLES11SP2-U9-3.0.74-0.6.6-xen-x86
SLES11SP2-U9-3.0.74-0.6.6-xen-x86_64
SLES11SP3-GA-3.0.76-0.11-default-x86
SLES11SP3-GA-3.0.76-0.11-default-x86_64
SLES11SP3-GA-3.0.76-0.11-pae-x86
SLES11SP3-GA-3.0.76-0.11-xen-x86
SLES11SP3-GA-3.0.76-0.11-xen-x86_64
SLES11SP3-U1-3.0.82-0.7-default-x86
SLES11SP3-U1-3.0.82-0.7-default-x86_64
SLES11SP3-U1-3.0.82-0.7-pae-x86
SLES11SP3-U1-3.0.82-0.7-xen-x86
SLES11SP3-U1-3.0.82-0.7-xen-x86_64
SLES11SP3-U2-3.0.93-0.8-default-x86
SLES11SP3-U2-3.0.93-0.8-default-x86_64
SLES11SP3-U2-3.0.93-0.8-pae-x86
SLES11SP3-U2-3.0.93-0.8-xen-x86
SLES11SP3-U2-3.0.93-0.8-xen-x86_64
SLES11SP3-U3-3.0.101-0.8-default-x86
SLES11SP3-U3-3.0.101-0.8-default-x86_64
SLES11SP3-U3-3.0.101-0.8-pae-x86
SLES11SP3-U3-3.0.101-0.8-xen-x86
SLES11SP3-U3-3.0.101-0.8-xen-x86_64
SLES11SP3-U4-3.0.101-0.15-default-x86
SLES11SP3-U4-3.0.101-0.15-default-x86_64
SLES11SP3-U4-3.0.101-0.15-pae-x86
SLES11SP3-U4-3.0.101-0.15-xen-x86
SLES11SP3-U4-3.0.101-0.15-xen-x86_64
SLES11SP3-U5-3.0.101-0.21-default-x86
SLES11SP3-U5-3.0.101-0.21-default-x86_64
SLES11SP3-U5-3.0.101-0.21-pae-x86
SLES11SP3-U5-3.0.101-0.21-xen-x86
SLES11SP3-U5-3.0.101-0.21-xen-x86_64
SLES11SP3-U6-3.0.101-0.29-default-x86
SLES11SP3-U6-3.0.101-0.29-default-x86_64
SLES11SP3-U6-3.0.101-0.29-pae-x86
SLES11SP3-U6-3.0.101-0.29-xen-x86
SLES11SP3-U6-3.0.101-0.29-xen-x86_64
SLES11SP3-U7-3.0.101-0.31-default-x86
SLES11SP3-U7-3.0.101-0.31-default-x86_64
SLES11SP3-U7-3.0.101-0.31-pae-x86
SLES11SP3-U7-3.0.101-0.31-xen-x86
SLES11SP3-U7-3.0.101-0.31-xen-x86_64

SLES11SP3-U8-3.0.101-0.35-default-x86
SLES11SP3-U8-3.0.101-0.35-default-x86_64
SLES11SP3-U8-3.0.101-0.35-pae-x86
SLES11SP3-U8-3.0.101-0.35-xen-x86
SLES11SP3-U8-3.0.101-0.35-xen-x86_64
SLES11SP4-GA-3.0.101-63-default-x86
SLES11SP4-GA-3.0.101-63-default-x86_64
SLES11SP4-GA-3.0.101-63-pae-x86
SLES11SP4-GA-3.0.101-63-xen-x86
SLES11SP4-GA-3.0.101-63-xen-x86_64
SLES11SP4-U1-3.0.101-65-default-x86
SLES11SP4-U1-3.0.101-65-default-x86_64
SLES11SP4-U1-3.0.101-65-pae-x86
SLES11SP4-U1-3.0.101-65-xen-x86
SLES11SP4-U1-3.0.101-65-xen-x86_64
SLES11SP4-U2-3.0.101-68-default-x86
SLES11SP4-U2-3.0.101-68-default-x86_64
SLES11SP4-U2-3.0.101-68-pae-x86
SLES11SP4-U2-3.0.101-68-xen-x86
SLES11SP4-U2-3.0.101-68-xen-x86_64
SLES9-GA-2.6.5-7.97-bigsmmp-x86
SLES9-GA-2.6.5-7.97-default-x86
SLES9-GA-2.6.5-7.97-default-x86_64
SLES9-GA-2.6.5-7.97-smp-x86
SLES9-GA-2.6.5-7.97-smp-x86_64
SLES9-SP1-2.6.5-7.139-bigsmmp-x86
SLES9-SP1-2.6.5-7.139-default-x86
SLES9-SP1-2.6.5-7.139-default-x86_64
SLES9-SP1-2.6.5-7.139-smp-x86
SLES9-SP1-2.6.5-7.139-smp-x86_64
SLES9-SP2-2.6.5-7.191-bigsmmp-x86
SLES9-SP2-2.6.5-7.191-default-x86
SLES9-SP2-2.6.5-7.191-default-x86_64
SLES9-SP2-2.6.5-7.191-smp-x86
SLES9-SP2-2.6.5-7.191-smp-x86_64
SLES9-SP3-2.6.5-7.244-bigsmmp-x86
SLES9-SP3-2.6.5-7.244-default-x86
SLES9-SP3-2.6.5-7.244-default-x86_64
SLES9-SP3-2.6.5-7.244-smp-x86
SLES9-SP3-2.6.5-7.244-smp-x86_64
SLES9-SP4-2.6.5-7.308-bigsmmp-x86
SLES9-SP4-2.6.5-7.308-default-x86
SLES9-SP4-2.6.5-7.308-default-x86_64
SLES9-SP4-2.6.5-7.308-smp-x86
SLES9-SP4-2.6.5-7.308-smp-x86_64

B Using the PlateSpin Migrate Command Line Interface

The PlateSpin Migrate Client installation includes a command line interface (CLI) tool to help you perform common migrations tasks. Conversion jobs using .ini files is supported onto VMware and Hyper-V targets only. Using this tool, you can

- ◆ Discover and subsequently refresh a host or target server to populate the Migrate Server with server information.
- ◆ Migrate (also known as "convert") heterogeneous workloads across x86 server and desktop infrastructure in the data center.
- ◆ Prepare the target host for its new workload and then, after a conversion, synchronize the host and the target.
- ◆ Install an image server, capture an image, deploy an image, or incrementally migrate an image.
- ◆ Check the status of a job as it is running, and if necessary, abort it.

This section includes information that can help you use the CLI effectively. The content includes:

- ◆ [Section B.1, "Where Is the Tool Located?," on page 195](#)
- ◆ [Section B.2, "Before You Use the Tool," on page 195](#)
- ◆ [Section B.3, "Configurable .ini Files \(Jobs\) You Can Use with the Tool," on page 198](#)

B.1 Where Is the Tool Located?

The CLI tool, `PlateSpin.Migrate.Console.exe`, is installed with the PlateSpin Migrate Client at the following location:

- ◆ **32-bit host:** `C:\Program Files\PlateSpin Migrate Client\CommandLine\PlateSpin.Migrate.Console.exe`
- ◆ **64-bit host:** `C:\Program Files(x86)\PlateSpin Migrate Client\CommandLine\PlateSpin.Migrate.Console.exe`

B.2 Before You Use the Tool

This section includes the following information:

- ◆ [Section B.2.1, "Pre-configuring the Migrate Server Values for CLI," on page 196](#)
- ◆ [Section B.2.2, "Becoming Familiar with the Commands," on page 196](#)

B.2.1 Pre-configuring the Migrate Server Values for CLI

Before you begin using the command line utility, you need to make sure that the Migrate Server is properly configured. You can check the configuration in the `PlateSpin.Migrate.Console.config` file, located in the same path as the command line utility. After the Migrate installation, the following config file should already be populated with values.

```
<?xml version="1.0" encoding="utf-8 ?>"
<configuration>
  <appSettings>
    <add key="MigrateServerURL" value="https://localhost/PlateSpinMigrate/" />
    <add key="ServerDomain" value="" />
    <add key="psuser" value="administrator" />
    <add key="pspassword" value="encoded_password" />
    <add key="encoded" value="yes" />
  </appSettings>
</configuration>
```

The tool uses these values as it executes commands. You need to reconcile the values in the file with the settings for the Migrate Server with which you want to connect.

The value for the `pspassword` key is blank by default and you must specify an encoded password as the value. To encode the password, use the `encode` command. For more information about commands, see [Section B.2.2, "Becoming Familiar with the Commands," on page 196](#).

If you choose to provide encoded passwords for source workload and target container, set the value for the `encoded` key in the following line of the `PlateSpin.Migrate.Console.config` file to `yes`, otherwise set value to `no`.

```
<add key="encoded" value="no" />
```

B.2.2 Becoming Familiar with the Commands

You can display the commands supported in the tool by running it with the `Help` option or with the `?` option from the command prompt, like this:

```
C:\Program Files\PlateSpin Migrate
Client\CommandLine>PlateSpin.Migrate.Console.exe Help
```

The tool displays a matrix that includes information similar to what is included in the following table:

Table B-1 Commands available from the Migrate CLI tool

Command	Description
<code>run</code>	Runs a configured <code>.ini</code> file as a scheduled job. When the you add the <code>/wait=no</code> parameter and the job starts to run, its Job ID is displayed in the interface.
<code>query</code>	Runs a query on the job (when you specify a Job ID) to display its current status.
<code>discover</code>	Runs an operation that inventories the details of a supported workload or target computer in preparation for a migration or "conversion" job.
<code>refresh</code>	Refreshes a discovered server.
<code>unDiscover</code>	Undiscovers a server.
<code>imageserver</code>	Performs imaging operations on a workload (that is, <i>install server</i> , <i>uninstall server</i> , <i>update tools</i>) on a server.

Command	Description
abort	Aborts a scheduled job.
licenseInfo	Displays the license information of the migrate server.
serversync	Prepares the server for the Server Sync operation and then runs a serversync job using the configuration file.
encode	Encodes the text input or the data in the text file.
massdiscover	Performs mass discovery of source workloads and targets. The discovered workloads and targets are displayed both in the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface To mass discover workloads and targets, you must first list the workloads and targets that you want to discover in a CSV file. To create this CSV file, refer to the sample CSV file located at \PlateSpin Migrate Client\CommandLine\Sample INI\MassDiscovery.csv.

When you run any of these commands, you must include its required parameter(s) in the command line. You can also include some optional parameters, when needed. For example, `savejob=` parameter saves the job in default location.

To display a list of these parameters at the command prompt, run the command without any parameter. For example, if you run the `discover` command without parameters, like this:

```
C:\Program Files\PlateSpin Migrate
Client\CommandLine>PlateSpin.Migrate.Console.exe discover
```

the command line interface displays these following:

```
[discover]                discovers a server
Required Parameters:
  /machineAddress=        machine address to discover
  /userName=              the username to use
  /password=              the password to use
  /type=                  type like windows, linux,vmware_esx,vmware_vcenter,
Optional Parameters:
  /network=               network name to connect to
  /address=               server address to connect to
  /psuser=                Username used for accessing PlateSpin Migrate server as
user different from the one logged on this computer
  /pspassword=            Password used for accessing Platespin Migrate server for
the user different from the one logged on this computer
  /wait=                  wait for completion of job [yes,no]
  /clusterName=           clustername to be discovered
  /verbose=               verbose mode for output [on,off]
  /output=                the output file
  /format=                the ouput format to display in [text,html,xml]
  /sslcertificatewarnings= Whether to Ignore or Enforce SSL Certificate
Warnings [Ignore| Enforce]
```

NOTE: You should become familiar with the different CLI commands and their respective required and optional parameters.

Command Line Syntax

If you were to run the discover command (which is also a job), you would use a syntax similar to this example, at the command prompt:

```
C:\Program Files\PlateSpin Migrate
Client\CommandLine>PlateSpin.Migrate.Console.exe discover /
machineaddress=10.10.8.100 /username=administrator /password=password /
type=windows /wait=no
```

Note that all required parameters and one optional parameter are included in this example.

When the discover command (job) starts, the CLI tool displays its job ID, similar to this example:

```
8be8d306-7665-4869-9795-a9dbb3ce1471
```

You can leverage this ID to learn the status of the job, just by using the query command, like this:

```
C:\Program Files\PlateSpin Migrate
Client\CommandLine>PlateSpin.Migrate.Console.exe query /id=8be8d306-7665-4869-
9795-a9dbb3ce1471
```

The query command yields a status report that includes all of the details of the job. This is the same kind of information you might see from the Migrate Client Jobs view.

B.3 Configurable .ini Files (Jobs) You Can Use with the Tool

When you install the PlateSpin Migrate Client, the installation creates a separate directory for a number of preconfigured jobs (actually, .ini files) that can do the following:

- ♦ Workload conversion (that is, a migration operation)
- ♦ Server Sync
- ♦ Imaging capture and deployment of image target

You execute a job by using the run command at the command line. The values in the files are the optional parameters that run along with the job. Each of these functions has a “default” .ini file version that runs with basic settings, and one or more “platform-specific” .ini file(s) that run with custom settings:

- ♦ Conversion-Default.ini
- ♦ Conversion-Windows.ini (customized)
- ♦ Conversion-Linux.ini (customized)
- ♦ ServerSync-Default.ini
- ♦ ServerSync-Windows.ini (customized)
- ♦ ServerSync-Linux.ini (customized)
- ♦ CaptureImage-Default.ini
- ♦ CaptureImage.ini (customized)
- ♦ DeployImage-Default.ini
- ♦ DeployImage.ini (customized)
- ♦ IncrementalImaging-Default.ini
- ♦ IncrementalImaging.ini (customized)

This section includes more details about these jobs in the following subsections:

- ♦ [Section B.3.1, “Conversion Jobs,” on page 199](#)
- ♦ [Section B.3.2, “ServerSync Jobs,” on page 206](#)
- ♦ [Section B.3.3, “Imaging Jobs,” on page 210](#)

B.3.1 Conversion Jobs

The CLI tool supports converting Windows and Linux workloads (source) to Hyper-V, vCenter, or ESX servers (target). There are two types of `.ini` files, one for a basic job configuration, and one for custom configurations. While the job is running you can abort the job or check its status.

Before you start a conversion job, make sure you run the `discover` command on the source computer and then on the target container. The following is example syntax for running the `discover` command:

```
discover /machineaddress=10.10.10.10 /username=administrator /  
password=anything@123 /type=vmware_vcenter
```

The tables in this section are named by the respective conversion jobs `.ini` files they represent. The table contents include the file section names within the `.ini` and the available settings you can configure according to your conversion needs:

IMPORTANT: For conversions to Hyper-V Generation 1 or BIOS machines, you must ensure that the boot volume is always mapped to Disk1 irrespective of the number of disks on the target machine. So, in the `.ini` file, you must ensure that the `MapTo=` setting in the `[Volume]` section that has `VolumeToCopy` mapped to boot volume is set to `Disk1`.

Sample of the settings in `Conversion-Windows.ini` file:

```
[Volume1]  
VolumeToCopy=boot_volume  
  
FreeSpace=  
  
MapTo=Disk1  
  
[Volume2]  
VolumeToCopy=non_boot_volume  
  
FreeSpace=  
  
MapTo=Disk2
```

Conversion-Default.ini

Table B-2 *Details of Conversion-Default.ini*

File Sections and Default Settings	Comment
[Type]	
Conversion=X2V	{required} This value must be used for every conversion.

File Sections and Default Settings	Comment
[JobConfig]	
Default=true	
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
TakeControl=static/dhcp	{conditional} Use this value only if the source is WINXP, Win2k3, Windows 2000. Otherwise, it is not required.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
DNS=	
[TargetContainer]	
Address=	<p>{required} Specify the IP address for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ◆ If ESX is discovered, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery. ◆ If Hyper-V is discovered, specify the Hyper-V IP Address.
UserName=	<p>{required} Specify the username for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ◆ If ESX is discovered via VCenter, specify the vCenter username. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root username. ◆ If Hyper-V is discovered, specify the Hyper-V username.
Password=	<p>{required} Specify the password for the target container depending on how it is discovered. For example:</p> <ul style="list-style-type: none"> ◆ If ESX is discovered via VCenter, specify the vCenter password. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root password. ◆ If Hyper-V is discovered, specify the Hyper-V password.
[NewMachine]	
DisplayName=	{required} Specify the name you want to display in the target container console.
HostName=	{required} Hostname of the target machine.

Conversion-Windows.ini

You can skip system volume.

Table B-3 Details of *Conversion-Windows.ini*

File Sections and Default Settings	Comment
[Type]	
Conversion=X2V	{required} This value must be used for every conversion.
[JobConfig]	
Default=false	
[Transfer]	
TransferType=VSSFileBased/VSSblockBased/FileBased	Possible settings shown. If the Windows source machine support VSS snapshotting, use the <i>vss...</i> settings, if it doesn't support VSS, use the <i>Filebased</i> setting.
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the <i>TransferType</i> setting. If that setting is <i>Filebased</i> and you want to perform an offline conversion, this setting must be set to <i>false</i> .
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
EndState=ShutDown/Donothing/Reboot	Possible settings shown.
TakeControl=static/dhcp	{conditional} Use this value only if the source is WINXP, Win2k3, Windows 2000. Otherwise, it is not required.
TakeControlAddress=	
VirtualNetwork=	For offline conversions, specify the MAC address of the source workload.
SubnetMask=	
DefaultGateway=	
DNS=	
[TargetContainer]	

File Sections and Default Settings	Comment
Address=	<p>{required} Specify the IP address for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ◆ If ESX is discovered, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery. ◆ If Hyper-V is discovered, specify the Hyper-V IP Address.
UserName=	<p>{required} Specify the username for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ◆ If ESX is discovered via VCenter, specify the vCenter username. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root username. ◆ If Hyper-V is discovered, specify the Hyper-V username.
Password=	<p>{required} Specify the password for the target container depending on how it is discovered. For example:</p> <ul style="list-style-type: none"> ◆ If ESX is discovered via VCenter, specify the vCenter password. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root password. ◆ If Hyper-V is discovered, specify the Hyper-V password.
VirtualNetwork=	Specify the target container virtual network name you want to use.
TakeControl=static/dhcp	Specify <code>static</code> or <code>dhcp</code> , depending on your networking configuration.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
DNS=	
[NewMachine]	
DisplayName=	{required} Specify the name you want to display in the target container console.

File Sections and Default Settings	Comment
DataStore=	Specify the name of datastore you want to use for configuration files. <ul style="list-style-type: none"> ◆ On ESX: datastore1 ◆ On Hyper-V: E:
ConfigPath=	<ul style="list-style-type: none"> ◆ On ESX: Specify the complete path where you want to create the .vmx file. For example: <i>/folder_name/vmx_file_name</i> The .vmx file is created in the specified folder within the datasource. ◆ On Hyper-V: Specify the path to the folder where you want to create the configuration file. For example: <i>Drive:\folder_name\config_file_name</i>
Memory=	Specify the amount of RAM you want for the target computer. The setting can be in MB or GB and must be specified with integers (no decimal values).
InstallTools=true/false	Possible settings shown. Default is true.
NumberOfCPU=	Specify the number of CPUs you want for the target computer.
HostName=	{required} Specify the target hostname.
WorkGroup=	{optional} Specify the workgroup name you want to join.
Domain=	
DomainUserName=	
DomainUserPassword=	
EndState=VMPowerOFF/VMPowerON/VMSuspend	Possible settings shown.
ScsiType=	(On VMware) Specify the Scsi Adapter type. If you do not specify a type or specify an unsupported adapter type, the default adapter type is used.
ResourcePool=	(On VMware) Specify the ResourcePool name in the vCenter. If the resource pool is nested, then use \ to separate names. For example, windows\local.
UseThinDisks=	To use thin disks, specify true . Else, specify false .
BootMode=	(On Hyper-V for Windows workload) Specify the boot mode supported on the target machine. For example: <ul style="list-style-type: none"> ◆ If the target machine is Windows Server 2012, specify either BIOS or UEFI. ◆ If the target machine is Windows Server 2008, specify BIOS.

File Sections and Default Settings	Comment
[EthernetNic1]	You can repeat this section of the .ini file for every NIC at the target container. For example, the second NIC section would be named [EthernetNic2] . Configuration settings would be specified for each NIC section in the file.
DHCPEnabled=true/false	Specify true for DHCP and false for static IP.
VirtualNetwork=	Specify the target container virtual network name you want to use.
Address=	Specify the IP address for the target machine.
SubnetMask=	
DefaultGateway=	
DNS=	Specify one or more DNS names separated by commas.
[DriveGeneral]	If you have multiple disks at the source, you can specify them here. You can specify as many disks as there are at the source.
DataStore1=	Specify the datastore on the target container. For example: <ul style="list-style-type: none"> ◆ On ESX: datastore1 ◆ On Hyper-V: E:
Disk1=	Specify the path to the configuration file on the target container. For example: <ul style="list-style-type: none"> ◆ On ESX: /win2k8r2/win2k8r2.vmdk ◆ On Hyper-V: \win2k8r2\win2k8r2.vhdx
DataStore2=	
Disk2=	
[Volume1]	You can repeat this section of the .ini file for every volume at the target container. For example, the second volume section would be named [Volume2] . Configuration settings would be specified for each volume section in the file.
VolumeToCopy=	Specify the volume to copy to the target.
MapTo=	Specify the disk to map.
FreeSpace=	Specify the amount of free space, in MB or GB, available on the target for File-Based conversion.

Conversion-Linux.ini

The sections in the `Conversion-Windows.ini` and in the `Conversion-Linux.ini` file are identical, except for the settings in **[Transfer]** section, along with the settings for workgroup and domain configuration. The differences for the Linux source job are shown in the following table.

Table B-4 *Conversion-Linux.ini: Differences in Setting Details of the [Transfer] section*

File Sections and Default Settings (differences only)	Comment
[Transfer]	
TransferType=BlockBased/FileBased	Possible settings shown. Linux does not support VSS.
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the TransferType setting. If that setting is Filebased and you want to perform an offline conversion, this setting must be set to false.
[Source]	
VirtualNetwork=	For offline conversions, specify the MAC address of the source workload.
[NewMachine]	
ScsiType=	(On VMware) Specify the Scsi Adapter type. If you do not specify a type or specify an unsupported adapter type, the default adapter type is used.
ResourcePool=	(On VMware) Specify the ResourcePool name in the vCenter. If the resource pool is nested, then use \ to separate names. For example, windows\local.
UseThinDisks=	To use thin disks, specify true . Else, specify false .
[EthernetNic1]	
DNS=	Specify one or more DNS names separated by commas.
[LVMGroup]	
Group1= Add entries depending on the number of groups you want. If you have two groups, then add the following: Group1= Group2=	Name of the LVM group in the source.
[Volume1]	
FreeSpace=	Specify the amount of free space, in MB or GB, available on the target for File-Based conversion.

B.3.2 ServerSync Jobs

Use `serversync` command to perform the Server Sync operation. There are two types of `.ini` files, one for a basic job configuration, and one for custom configurations. While the job is running you can abort the job or check its status. If you specify the required settings, it will start the job. Then, when it runs, the job populates the other values with default settings.

The tables in this section are named by the respective `serversync` jobs `.ini` files they represent. The table contents include the file section names within the `.ini` and the available settings you can configure according to your conversion needs:

ServerSync-Default.ini

Table B-5 Details of ServerSync-Default.ini

File Sections and Default Settings	Comment
[Type]	
Conversion=Sync2V	{required} This value must be used for every Server Sync operation.
[JobConfig]	
Default=true	
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
TakeControl=static/dhcp	{conditional} Use this value only if the source is WINXP, Win2k3, Windows 2000, or for offline conversion.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
DNS=	
[TargetContainer]	

File Sections and Default Settings	Comment
Address=	{required} Specify the IP address for the target container depending on how it is discovered. For example: <ul style="list-style-type: none"> ◆ If ESX is discovered, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery. ◆ If Hyper-V is discovered, specify the Hyper-V IP Address.
UserName=	{required} Specify the username for the target container depending on how it is discovered. For example: <ul style="list-style-type: none"> ◆ If ESX is discovered via VCenter, specify the vCenter username. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root username. ◆ If Hyper-V is discovered, specify the Hyper-V username.
Password=	{required} Specify the password for the target container depending on how it is discovered. For example: <ul style="list-style-type: none"> ◆ If ESX is discovered via VCenter, specify the vCenter password. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root password. ◆ If Hyper-V is discovered, specify the Hyper-V password.
[ExistingTargetMachine]	
DisplayName=	{required} Specify the display name of the target machine where you want to sync.
HostName=	{required}

ServerSync-Windows.ini

For prepare for Sync, the ServerSync command uses target container and network details from TargetContainer and machine name from ExistingTargetMachine file sections.

Table B-6 Details of ServerSync-Windows.ini

File Sections and Default Settings	Comment
[Type]	
Conversion=Sync2V	{required} This value must be used for every Server Sync operation.

File Sections and Default Settings	Comment
[JobConfig]	
Default=false	
[Transfer]	
TransferType=VSSFileBased/VSSblockBased/FileBased	Possible settings shown. If the Windows source machine support VSS snapshotting, use the VSS... settings, if it doesn't support VSS, use the Filebased setting.
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the TransferType setting. If that setting is Filebased and you want to perform an offline conversion, this setting must be set to false.
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
EndState=ShutDown/Donothing/Reboot	Possible settings shown.
TakeControl=static/dhcp	{conditional} Use this value only if the source is WINXP, Win2k3, Windows 2000. Otherwise, it is not required.
TakeControlAddress=	
VirtualNetwork=	For offline conversions, specify the MAC address of the source workload.
SubnetMask=	
DefaultGateway=	
DNS=	
[TargetContainer]	
Address=	<p>{required} Specify the IP address for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ◆ If ESX is discovered, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery. ◆ If Hyper-V is discovered, specify the Hyper-V IP address.

File Sections and Default Settings	Comment
UserName=	{required} Specify the username for the target container depending on how it is discovered. For example: <ul style="list-style-type: none"> ◆ If ESX is discovered via VCenter, specify the vCenter username. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root username. ◆ If Hyper-V is discovered, specify the Hyper-V username.
Password=	{required} Specify the password for the target container depending on how it is discovered. For example: <ul style="list-style-type: none"> ◆ If ESX is discovered via VCenter, specify the vCenter password. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root password. ◆ If Hyper-V is discovered, specify the Hyper-V password.
VirtualNetwork=	Specify the target container virtual network name you want to use.
TakeControl=static/dhcp	Specify <i>static</i> or <i>dhcp</i> depending on your networking configuration.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
DNS=	
[ExistingTargetMachine]	
DisplayName=	{required} Specify the display name of the target machine where you want to sync.
HostName=	.
InstallTools=true/false	.
WorkGroup=	Specify the workgroup name if you want to join workgroup.
Domain=	.
DomainUserName=	.
DomainUserPassword=	.
EndState=VMPowerOFF/VMPowerON/VMSuspend	Possible settings shown.

File Sections and Default Settings	Comment
[EthernetNic1]	You can repeat this section of the .ini file for every NIC at the target container. For example, the second NIC section would be named [EthernetNic2]. Configuration settings would be specified for each NIC section in the file.
DHCPEnabled=true/false	Specify true for DHCP and false for static IP.
VirtualNetwork=	Specify the target container virtual network name you want to use.
Address=	Specify the IP address for the target machine.
SubnetMask=	
DefaultGateway=	
DNS=	

ServerSync-Linux.ini

The sections in the `ServerSync-Windows.ini` and in the `ServerSync-Linux.ini` file are identical, except for the settings in [Transfer] section, along with the settings for the workgroup and domain configuration. For prepare for Sync, the `ServerSync` command uses target container and network details from `TargetContainer` and machine name from `ExistingTargetMachine` file sections.

The differences for the Linux source job are shown in the following table.

Table B-7 *ServerSync-Linux.ini: Differences in Setting Details of the [Transfer] section*

File Sections and Default Settings (differences only)	Comment
[Transfer]	
TransferType=BlockBased/FileBased	Possible settings shown. Linux does not support VSS.
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the <code>TransferType</code> setting. If that setting is <code>Filebased</code> and you want to perform an offline conversion, this setting must be set to <code>false</code> .
[Source]	
VirtualNetwork=	For offline conversions, specify the MAC address of the source workload.

B.3.3 Imaging Jobs

The CLI tool supports several imaging operations (for example, install, uninstall, and update tools) through its `imageserver` command. Before you start an `imageserver` job, make sure you run the `discover` command on the source computer and then on the target container.

In addition to the `imageserver` job, the CLI tool supports imaging Windows workloads (source) to the target. There are two types of imaging .ini files, one for a basic job configuration, and one for custom configurations. While the job is running you can abort the job or check its status.

The tables in this section are named by the respective imaging jobs `.ini` files they represent. The table contents include the file section names within the `.ini` and the available settings you can configure according to your conversion needs:

CaptureImage-Default.ini

Table B-8 Details of `CaptureImage-Default.ini`

File Sections and Default Settings	Comment
[Type]	
Conversion=X2I	{required} This value must be used for every image capture.
[JobConfig]	
Default=true	
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
TakeControl=static/dhcp	{conditional} Use this value only if the source is WINXP, Win2k3, Windows 2000. Otherwise, it is not required.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
[ImageConfiguration]	
ImageDisplayName=	{required} Specify the display name of the image in the image server.
[TargetContainer]	
Address=	{required} Specify IP address of image server.
UserName=	{required} Specify the username of the image server.
Password=	{required} Specify the password of the image server.

CaptureImage.ini

You can skip system volume for capture image.

Table B-9 Details of `CaptureImage.ini`

File Sections and Default Settings	Comment
[Type]	
Conversion=X2I	{required} This value must be used for every image capture.
[JobConfig]	

File Sections and Default Settings	Comment
Default=false	
[Transfer]	
TransferType=VSSFileBased/FileBased	Possible settings are shown. If the Windows source machine support VSS snapshotting, use the <code>VSSFileBased</code> setting, if it doesn't support VSS, use the <code>Filebased</code> setting.
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the <code>TransferType</code> setting. If that setting is <code>Filebased</code> and you want to perform an offline conversion, this setting must be set to <code>false</code> .
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
EndState=ShutDown/Donothing/Reboot	Possible settings are shown.
TakeControl=static/dhcp	{conditional} Use this value only if the source is WINXP, Win2k3, Windows 2000. Otherwise, it is not required.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
DNS=	
[ImageConfiguration]	
ImageDisplayName=	Specify a name that you want to appear in the Image Server.
CompresionEnabled=true/false	Specify whether or not to use NTFS file compression. Default is <code>false</code> .
ConfigurationPath=	Specify the absolute path of the configuration file.
[TargetContainer]	
Address=	{required} Specify IP address of image server.
UserName=	{required} Specify the username of the image server.
Password=	{required} Specify the password of the image server.
[Volume1]	

File Sections and Default Settings	Comment
VolumeToCopy=	Specify the name of the volume you want to capture.
MapTo=	Specify the path where you want to create the package file for the volume.

DeployImage-Default.ini

Table B-10 Details of *DeployImage-Default.ini*

File Sections and Default Settings	Comment
[Type]	
Conversion=I2V	{required} This value must be used for every image deployment.
[JobConfig]	
Default=true	
[Source]	
Address=	{required} Specify an IP address for the image container.
UserName=	{required} Specify a username credential for the image container.
Password=	{required} Specify a password credential for the image container.
ImageDisplayName=	Specify the name of the image that you want to deploy.
[TargetContainer]	
Address=	<p>{required} Specify the IP address for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ◆ If ESX is discovered, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery. ◆ If Hyper-V is discovered, specify the Hyper-V IP Address.
UserName=	<p>{required} Specify the username for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ◆ If ESX is discovered via VCenter, specify the vCenter username. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root username. ◆ If Hyper-V is discovered, specify the Hyper-V username.

File Sections and Default Settings	Comment
Password=	{required} Specify the password for the target container depending on how it is discovered. For example: <ul style="list-style-type: none"> ◆ If ESX is discovered via VCenter, specify the vCenter password. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root password. ◆ If Hyper-V is discovered, specify the Hyper-V password.
[NewMachine]	
DisplayName=	{required} Specify the name you want to display in the target container.
HostName=	Specify the target hostname.

DeployImage.ini

You can skip system volume.

IMPORTANT: For deploying image to Hyper-V Generation 1 or BIOS machines, you must ensure that the boot volume is always mapped to Disk1 irrespective of the number of disks on the target machine. So, in the .ini file, you must ensure that the MapTo= setting in the [Volume] section that has VolumeToCopy mapped to boot volume is set to Disk1.

Sample of the settings in Conversion-Windows.ini file:

```
[Volume1]
VolumeToCopy=boot_volume
FreeSpace=
MapTo=Disk1
[Volume2]
VolumeToCopy=non_boot_volume
FreeSpace=
MapTo=Disk2
```

Table B-11 Details of DeployImage.ini

File Sections and Default Settings	Comment
[Type]	
Conversion=I2V	{required} This value must be used for every image deployment.
[JobConfig]	
Default=false	

File Sections and Default Settings	Comment
[Source]	
Address=	{required} Specify an IP address for the image container.
UserName=	{required} Specify a username credential for the image container.
Password=	{required} Specify a password credential for the image container.
ImageDisplayName=	Specify a name for the image that you want to deploy.
[TargetContainer]	
Address=	<p>{required} Specify the IP address for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ◆ If ESX is discovered, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery. ◆ If Hyper-V is discovered, specify the Hyper-V IP Address.
UserName=	<p>{required} Specify the username for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ◆ If ESX is discovered via VCenter, specify the vCenter username. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root username. ◆ If Hyper-V is discovered, specify the Hyper-V username.
Password=	<p>{required} Specify the password for the target container depending on how it is discovered. For example:</p> <ul style="list-style-type: none"> ◆ If ESX is discovered via VCenter, specify the vCenter password. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root password. ◆ If Hyper-V is discovered, specify the Hyper-V password.
TakeControl=static/dhcp	Specify <i>static</i> or <i>dhcp</i> , depending on your networking configuration.
VirtualNetwork=	
TakeControlAddress=	
SubnetMask=	

File Sections and Default Settings	Comment
DefaultGateway=	
[NewMachine]	
DisplayName=	{required} Specify the name you want to display in target container.
DataStore=	Specify the name of datastore you want to use. <ul style="list-style-type: none"> ◆ On ESX: datastore1 ◆ On Hyper-V: E:
ConfigPath=	<ul style="list-style-type: none"> ◆ On ESX: Specify the complete path where you want to create the .vmx file. For example: <i>/folder_name/vmx_file_name</i> The .vmx file is created in the specified folder within the datasource. ◆ On Hyper-V: Specify the path to the folder where you want to create the configuration file. For example: <i>Drive:\folder_name\config_file_name</i>
Memory=	Specify the amount of RAM you want for the target computer in MB or GB.
WorkGroup=	Specify the name of the workgroup you want to join.
InstallTools=true/false	Possible settings shown. Default is true.
NumberOfCPU=	Specify the number of CPUs you want for the target computer.
Memory=	.Specify the amount of RAM you want for the target computer.
Domain=	
DomainUserName=	
DomainUserPassword=	
HostName=	
EndState=VMPowerOFF/VMPowerON/VMSuspend	Possible settings shown.
ScsiType=	(On VMware) Specify the Scsi Adapter type. If you do not specify a type or specify an unsupported adapter type, the default adapter type is used.
ResourcePool=	(On VMware) Specify the ResourcePool name in the vCenter. If the resource pool is nested, then use \ to separate names. For example, windows\local.
UseThinDisks=	To use thin disks, specify true . Else, specify false .

File Sections and Default Settings	Comment
BootMode=	(On Hyper-V for Windows workload) Specify the boot mode supported on the target machine. For example: <ul style="list-style-type: none"> ◆ If the target machine is Windows Server 2012, specify either BIOS or UEFI. ◆ If the target machine is Windows Server 2008, specify BIOS.
[EthernetNic1]	If you have two (or more) NICs at the target, you can specify them and their configurations
DHCPEnabled=true/false	Specify true for DHCP and false for static IP.
VirtualNetwork=	Specify the target container virtual network name you want to use
Address=	
SubnetMask=	
DefaultGateway=	
DNS=	Specify one or more DNS names separated by commas.
[DriveGeneral]	If you have multiple disks at the source, you specify them here (create more if needed). The <code>.ini</code> file shows examples of two disks being specified. You can specify as many disks as there are at the source.
DataStore1=	Specify the datastore on the target container. For example: <ul style="list-style-type: none"> ◆ On ESX: <code>datastore1</code> ◆ On Hyper-V: <code>E:</code>
Disk1=	Specify the path to the configuration file on the target container. For example: <ul style="list-style-type: none"> ◆ On ESX: <code>/win2k8r2/win2k8r2.vmdk</code> ◆ On Hyper-V: <code>\win2k8r2\win2k8r2.vhdx</code>
DataStore2=	
Disk2=	
[Volume1]	
VolumeToCopy=	Specify the volume to copy to the target.
MapTo=	Specify the disk to map.
FreeSpace=	Specify the amount of free space, in MB or GB, available on the target for File-Based conversion.

IncrementalImaging-Default.ini

Table B-12 Details of *IncrementalImaging-Default.ini*

File Sections and Default Settings	Comment
[Type]	
Conversion=Sync2I	{“Existing Image”: required} Every incremental image capture uses this setting.
[JobConfig]	
Default=true	
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
TakeControl=static/dhcp	{conditional} Use this value only if the source is WINXP, Win2k3, Windows 2000. Otherwise, it is not required.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
ImageDisplayName=	Specify an image name that already exists in the image server.
[TargetContainer]	
Address=	{required} Specify IP address of image server.
UserName=	{required} Specify the username of the image server.
Password=	{required} Specify the password of the image server.

IncrementalImaging.ini

Table B-13 Details of *IncrementalImaging.ini*

File Sections and Default Settings	Comment
[Type]	
Conversion=Sync2I	{“Existing Image: required} Every incremental image capture uses this setting.
[JobConfig]	
Default=false	
[Transfer]	

File Sections and Default Settings	Comment
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the <code>TransferType</code> setting. If that setting is <code>Filebased</code> and you want to perform an offline conversion, this setting must be set to <code>false</code> .
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
TakeControl=static/dhcp	{conditional} Use this value only if the source is WINXP, Win2k3, Windows 2000. Otherwise, it is not required.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
EndState=ShutDown/Donothing/Reboot	Possible settings shown.
ImageDisplayName=	Specify an image name that already exists in image server.
[TargetContainer]	
Address=	{required} Specify IP address of image server.
UserName=	{required} Specify the username of the image server.
Password=	{required} Specify the password of the image server.

C Rebranding the PlateSpin Migrate Web Interface

You can modify the appearance of the PlateSpin Migrate Web Interface to match the look and feel of your corporate identity. You can modify colors, logo, and product name. You can even eliminate the links to **About** tab and **Help** tab in the product interface.

This section includes information to help you change the branding of the product:

- ◆ [Section C.1, “Rebranding the Interface Using Configuration Parameters,” on page 221](#)
- ◆ [Section C.2, “Rebranding the Product Name in the Windows Registry,” on page 224](#)

C.1 Rebranding the Interface Using Configuration Parameters

You can change the look and feel of the Web Interface to match the proprietary look of your organization websites. To customize the branding of the Web Interface, modify the configuration parameters of your PlateSpin Server host:

- 1 In the address bar of a web browser, type the following URL to launch the PlateSpin Server Configuration Settings page:

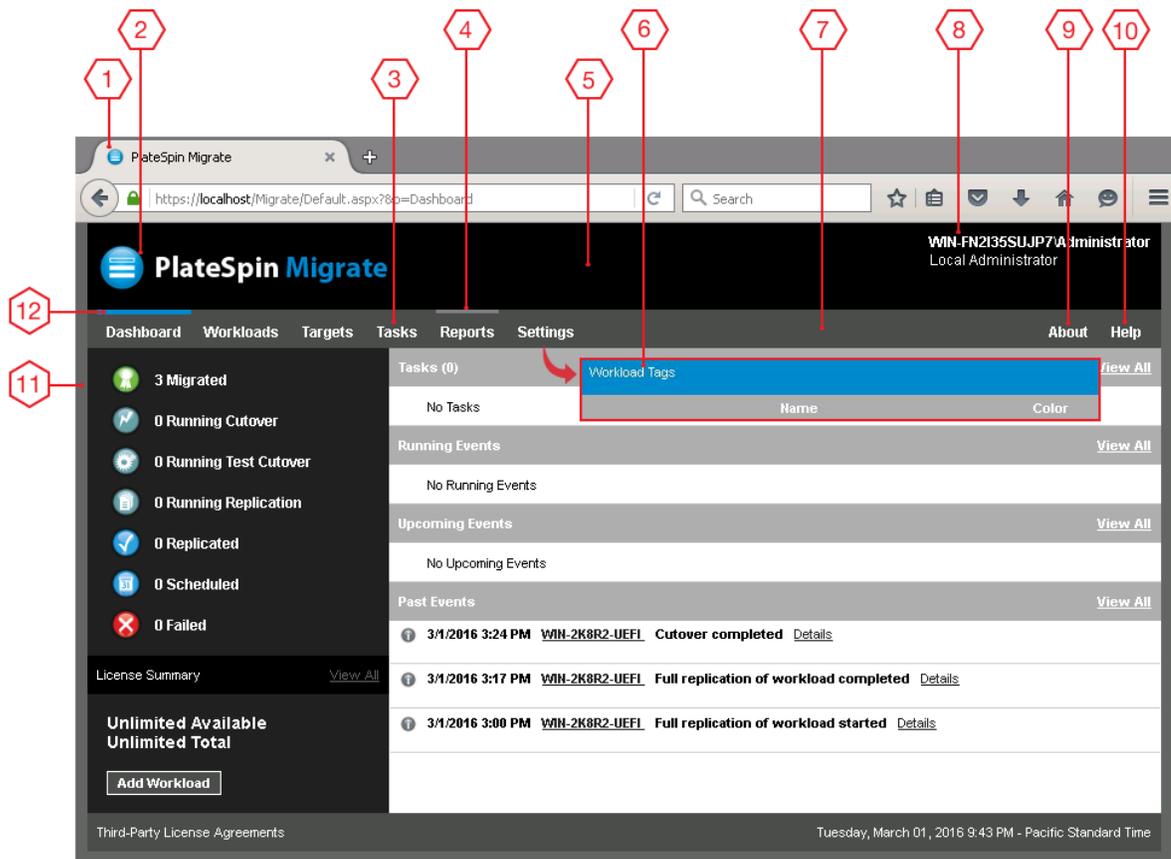
```
https://Your_PlateSpin_Server/platespinconfiguration/
```

- 2 Log in as Administrator.
- 3 Locate the required server parameter and click **Edit** to change its value.
- 4 Click Save.

After you modify the settings in the configuration tool, it might take up to 30 seconds for the change to take reflect on the interface. You need not reboot or restart the services.

The look and feel of the Web Interface is consistent across the various pages. In [Section C.1.1, “PlateSpin Migrate Web Interface with Configurable Elements Labeled,” on page 222](#), the illustration of the PlateSpin Migrate Dashboard identifies the elements that you can modify with numbered callouts.

C.1.1 PlateSpin Migrate Web Interface with Configurable Elements Labeled



C.1.2 PlateSpin Server Configuration Settings

The following table provides information about the setting you must use to modify the corresponding interface element. The ID column in the table lists the ID of the interface element identified in the Web Interface figure provided in Section C.1.2, “PlateSpin Server Configuration Settings,” on page 222.

ID	Setting Name and Description	Default Value
1	<p>WebUIFaviconUrl</p> <p>Location of a valid <code>.ico</code> graphic file. Specify one of the following:</p> <ul style="list-style-type: none"> ◆ A valid URL to the appropriate <code>.ico</code> file on a different machine. For example: <code>https://myserver.example.com/dir1/dir2/icons/mycompany_favicon.ico</code> ◆ A relative path below the root of the local web server where you have uploaded the appropriate <code>.ico</code> file. For example, if you create a path called <code>mycompany\images\icons</code> at the root of the web server to store your custom icon graphics: <code>~/mycompany/images/icons/mycompany_favicon.ico</code> In this example, the actual file system path that contains the file is <code>C:\Program Files (x86)\PlateSpin Migrate Server\PlateSpin Forge\web\mycompany\images\icons\mycompany_favicon.ico</code>. 	<code>~/doc/en/favicon.ico</code> ¹
2	<p>WebUILogoUrl</p> <p>Location of product logo graphic file. Specify one of the following:</p> <ul style="list-style-type: none"> ◆ A valid URL to the appropriate graphics file on a different machine. For example: <code>https://myserver.example.com/dir1/dir2/logos/mycompany_logo.png</code> ◆ A relative path below the root of the local web server where you have uploaded the appropriate graphics file. For example, if you create a path called <code>mycompany\images\logos</code> at the root of the web server to store your custom logo images: <code>~/mycompany/images/logos/mycompany_logo.png</code> In this example, the actual file system path that contains the file is <code>C:\Program Files (x86)\PlateSpin Migrate Server\PlateSpin Forge\web\mycompany\images\logos\mycompany_logo.png</code>. 	<code>~/Resources/protectLogo.png</code> ²
3	<p>WebUISiteNavigationFontColor</p> <p>Color of site navigation link font color in Web UI (RGB hex value)</p>	<code>#FFFFFF</code>
4	<p>WebUISiteNavigationLinkHoverBackgroundColor</p> <p>Color of site navigation link background in hover state (RGB hex value)</p>	<code>#808080</code>

ID	Setting Name and Description	Default Value
5	WebUISiteHeaderBackgroundColor Site header background color (RGB hex value)	#000000
6	WebUISiteAccentFontColor Font color to display with accent color in Web UI (RGB hex value)	#FFFFFF
7	WebUISiteNavigationBackgroundColor Color of site navigation background in Web UI (RGB hex value)	#4D4D4D
8	WebUISiteHeaderFontColor Site header font color in Web UI (RGB hex value)	#FFFFFF
9	WebUIShowAboutTab Toggles the visibility of the About tab: <ul style="list-style-type: none"> ◆ True: The About tab is visible on the interface. ◆ False: The About tab is not visible on the interface. 	True
10	WebUIShowHelpTab Toggle the visibility of the Help tab: <ul style="list-style-type: none"> ◆ True: The Help tab is visible on the interface. ◆ False: The Help tab is not visible on the interface. 	True
11	WebUISiteBackgroundColor Site background color (RGB hex value)	#666666
12	WebUISiteAccentColor Accent color (RGB hex value)	#0088CE

¹ Actual file path is C:\Program Files (x86)\PlateSpin Migrate Server\PlateSpin Forge\web\doc\en\favicon.ico.

² Actual file path is C:\Program Files (x86)\PlateSpin Migrate Server\PlateSpin Forge\web\Resources\protectLogo.png.

C.2 Rebranding the Product Name in the Windows Registry

The masthead at the top of the product interface provides space for the corporate logo and the product name. To change the logo, which commonly includes the product name, see [Section C.1, “Rebranding the Interface Using Configuration Parameters,”](#) on page 221.

To edit or eliminate the product name in a browser tab, do the following:

- 1 On the PlateSpin Migrate server, run `regedit`.
- 2 In the Windows Registry Editor, navigate to the following registry key:

HKEY_LOCAL_MACHINE\SOFTWARE\PlateSpin\MigrateServer\ProductName

NOTE: In some cases, the registry key can be found in this location:

HKEY_LOCAL_MACHINE\SOFTWARE Wow6432Node\PlateSpin\MigrateServer

- 3 Double-click the `ProductName` key and change the **Value data** for the key as required and then click **OK**.
- 4 Restart the IIS Server.

D Synchronizing Serial Numbers on Cluster Node Local Storage

This section details the procedure you can use to change local volume serial numbers to match each node of the Windows cluster that you want to migrate. The information includes the use of the Volume Manager utility (`VolumeManager.exe`) to synchronize serial numbers on cluster node local storage.

To download and run the utility:

- 1 From the [NetIQ Downloads site](#), search for the PlateSpin Migrate product, then click **Submit Query**.
- 2 On the Products tab, select PlateSpin Migrate 12.1 to go to the product-specific download page, then click **proceed to download**.
- 3 On the download page, click **download** on the `VolumeManager.exe` line or select the comparable download manager link.
- 4 Download the utility, then copy it to an accessible location on each cluster node.
- 5 On the active node of the cluster, open an administrative command prompt, navigate to the location of the downloaded utility, and run the following command:

```
VolumeManager.exe -l
```

A listing of the local volumes and their respective serial numbers is displayed. For example:

```
Volume Listing:
```

```
-----
```

```
DriveLetter (*) VolumeId="System Reserved" SerialNumber: AABB-CCDD
```

```
DriveLetter (C:) VolumeId=C:\ SerialNumber: 1122-3344
```

Make note of these serial numbers or keep them displayed for later comparison.

- 6 Verify that all local storage serial numbers of the active node match the local storage serial numbers on each of the other nodes in the cluster.
 - 6a On each cluster node, run the `VolumeManager.exe -l` command to obtain its volume serial numbers.
 - 6b Compare the local storage serial numbers of the active node ([Step 5](#)) against the local storage serial numbers of the node ([Step 6a](#)).
 - 6c (Conditional) If there are any differences in the serial numbers between the active node and this node, take note of the serial number you want to propagate on this node and run the following command to set, and then to verify the serial number:

```
VolumeManager -s <VolumeId> <serial-number>
```

Following are two examples of how this command could be used:

- ◆ `VolumeManager -s "System Reserved" AAAA-AAAA`
- ◆ `VolumeManager -s C:\ 1111-1111`

- 6d** When you have successfully changed all of the volume serial numbers on a node of the cluster, you need to restart that node.
- 6e** Repeat [Step 6a](#) through [Step 6d](#) for each node of the cluster.
- 7** (Conditional) If the cluster has already been migrated in a PlateSpin environment, we recommend running a full replication on the active node to ensure that any changes are propagated to the database.

E Frequently Asked Questions

This section provides answers to frequently asked questions.

What are the performance and scalability characteristics of my PlateSpin Migrate product?

Your PlateSpin Migrate product's overall performance, including data transfer speeds and scalability, depend on a variety of factors in your specific environment. See ["Performance" on page 24](#).

How secure is my PlateSpin Migrate product?

PlateSpin Migrate provides several features to help you safeguard your data and increase security. See ["Security and Privacy" on page 23](#).

Does PlateSpin Migrate support my workload's data storage technology?

PlateSpin Migrate products support a number of data storage and management technologies, including Windows dynamic disks, Linux logical volumes, RAID (Redundant Array of Independent Disks) systems, and SAN (Storage Area Network) systems.

Can I use custom SSH ports to communicate with my workloads?

Yes. See ["Discovering the Details of Source Workloads and Targets" on page 65](#).

Can multiple migrations run simultaneously?

Yes. See ["Performance" on page 24](#).

F Troubleshooting PlateSpin Migrate

This section provides a series of topics about troubleshooting PlateSpin Migrate.

- ◆ [Section F.1, “Discovery,” on page 231](#)
- ◆ [Section F.2, “Peer-to-Peer Migrations \(Windows\),” on page 232](#)
- ◆ [Section F.3, “Using Images,” on page 233](#)
- ◆ [Section F.4, “Post-migration Cleanup of Source Workloads,” on page 234](#)
- ◆ [Section F.5, “Shrinking the PlateSpin Migrate Databases,” on page 236](#)

F.1 Discovery

Table F-1 Common Issues and Solutions Related to Discovery Operations

Problems or Messages	Solutions
"Application has generated an error" during registration of physical server	This error occurs if the physical server is unable to contact the PlateSpin Server. A common cause is incorrect information entered during the registration process. To restart the registration process, enter <code>RegisterMachine.bat</code> . Ping to confirm basic connectivity with the PlateSpin Server.
My physical server has completed the registration process, but is not seen in PlateSpin Migrate Client	The full registration process can take some time to complete. After the second command prompt window has closed on the physical server, wait a few minutes before clicking the Refresh button in PlateSpin Migrate Client.
Problems discovering source and target servers	<p>KB Article 7920291 (https://www.netiq.com/support/kb/doc.php?id=7920291) contains troubleshooting checklists for discovering the following:</p> <ul style="list-style-type: none">◆ Linux servers and VMware ESX Servers◆ Windows-based source and target servers <p>The article also has instructions for troubleshooting WMI connections and checking if DCOM is enabled.</p>
"Package <...> Not Found" during discovery of existing Windows servers	Check for IIS configuration and network settings.

Problems or Messages	Solutions
Could not find file "\\{servername}\admin\$\{random ID}.xml"	<p>This error might occur on Windows Server 2000 or 2003 hosts.</p> <p>In some cases, either of these troubleshooting steps addresses the issue:</p> <ul style="list-style-type: none"> ◆ Ensure that the Admin\$ share on the PlateSpin Server host is accessible. If not, enable it and try the discovery again. <p>- OR -</p> <ul style="list-style-type: none"> ◆ 1. From any Web browser, open <code>https://Your_PlateSpin_Server/platespinconfiguration</code>. 2. Locate and edit the <code>ForceMachineDiscoveryUsingService</code> entry and change it to <code>true</code>. 3. Save the value and retry the discovery.

Related KB Articles:

ID	Description
7920339 (https://www.netiq.com/support/kb/doc.php?id=7920339)	ERRMSG: Discovery fails with "The request failed with HTTP status 407" message
7920862 (https://www.netiq.com/support/kb/doc.php?id=7920862)	ERRMSG: Recoverable Error: ControllerConnectionBroken during discovery
7920291 (https://www.netiq.com/support/kb/doc.php?id=7920291)	ERRMSG: Server details discovery problems

F.2 Peer-to-Peer Migrations (Windows)

Table F-2 Common Issues and Solutions Related to Peer-to-Peer Migrations (Windows)

Problems or Messages	Solutions
<p>One of the following errors displays during offline migration:</p> <ul style="list-style-type: none"> ◆ Waiting for Controller to start (Failed) ◆ Controller Connection Not Established ◆ Controller Connection Broken ◆ Unable to start the Heartbeat Service 	<p>This indicates one of the following problems:</p> <ul style="list-style-type: none"> ◆ The network settings for the temporary IP addresses under Job Configuration > Advanced might not be configured properly. ◆ There was a possible network outage that prevented the source/target machine from communicating with the PlateSpin Server. ◆ The source/target machine was not able to fully boot into the pre-execution environment. <p>To diagnose the exact cause of failure, check the state of the system where the controller failed to start. Commands such as <code>ipconfig</code> and <code>ping</code> are available to verify basic network connectivity.</p>
File transfer hangs at 1% or progresses at a slow pace	<p>By default, a link type of AUTO is used on the source server during a migration. If the source server is connected to a switch port that is forced to 100/FULL, the Force Full Duplex option must be enabled when configuring the migration. If this option is set incorrectly, a duplex mismatch occurs on the network.</p>

Problems or Messages	Solutions
Unable to determine suitable boot partition	<p>When converting existing source servers, the boot volume must pass the following checks:</p> <ul style="list-style-type: none"> ◆ It must be on a basic disk ◆ It must have 175 MB of free space ◆ It must be a primary partition ◆ If any of these are not true for the system volume, the migration fails while attempting to take control of the source server.
Job remains in a Scheduled state for a long period and then changes to Recoverable error (all sub-steps display NotStarted status)	<p>There is a problem with the Operations Framework Controller on the PlateSpin Server. Use the Windows services plug-in to confirm that the Controller is running. See KB Article 7920862 for other troubleshooting instructions.</p>
Troubleshooting failures at the Configuring Operating System stage (also applicable to Configure Target Machine or Configuring Virtual Machine migration steps)	<p>Generally, failures during the configuration step indicate that a time-out occurred when attempting to configure the target physical or virtual machine. Although the migration job appears to have failed, the overall migration is probably successful and the configuration service running on the target will likely continue its operations.</p> <p>KB Article 7920327 contains a detailed troubleshooting checklist and lists information required if technical support is necessary.</p>
Live Transfer is unavailable	<p>Either an unsupported file system or operating system exists on the server.</p>

Related KB Articles:

ID	Description
7920862	ERRMSG: PlateSpin Migrate Job remains at a "Scheduled" or "Recoverable Error" state
7920810	INFO: Restore job stalls - "The configuration service in the target machine..."
2790341	INFO: What ports does PlateSpin Migrate use during discovery, migration and file transfer?

F.3 Using Images

Table F-3 Common Issues and Solutions Related to PlateSpin Images

Problems or Messages	Solutions
Cannot see PlateSpin Images on PlateSpin Image Server	<p>If the Servers view is configured to group servers by machine, discovered image servers cannot be expanded. To display the images, reconfigure the Servers View so the servers are grouped by domain instead of machine.</p>

Problems or Messages	Solutions
Failed to mount image. The volume does not contain a recognized file system	This error message might appear when you are importing or deploying volume data while installing a PlateSpin Image Server on Windows 2003. To resolve the error, use the Windows services plug-in on the PlateSpin Image Server. Modify the logon properties for the PlateSpin Migrate Operations Management Controller service to use an account with local administrative privileges. Restart the service after making this change.
Security descriptors are not intact on deployed server when you are using volume data from a Symantec Ghost image	When you are creating a PlateSpin Image using raw volume data that was extracted from a Ghost Image, the security descriptors are not preserved on the VM. This is because the extracted files inherit permissions of their parent folder.

Related KB Articles:

ID	Description
7920879 (https://www.netiq.com/support/kb/doc.php?id=7920879)	ERRMSG: The file cannot be accessed by the system

F.4 Post-migration Cleanup of Source Workloads

Source workloads might occasionally need a post-migration cleanup. For example, you might need to clean up a source workload of all PlateSpin software components after an unsuccessful migration.

- ◆ [Section F.4.1, “Cleaning Up Windows Workloads,” on page 234](#)
- ◆ [Section F.4.2, “Cleaning Up Linux Workloads,” on page 235](#)

F.4.1 Cleaning Up Windows Workloads

The following are instructions for cleaning up Windows workloads by component and use case.

Component	Use Case	Removal Instructions
File-based Transfer Component	All Migrations	At root level for each volume migrated, remove all files named <code>PlateSpinCatalog*.dat</code>
Workload discovery software	All migrations	<ol style="list-style-type: none"> 1. In the Servers view, undiscover the source (right-click, then select Undiscover). 2. In the source workload's windows directory: <ul style="list-style-type: none"> ◆ Remove all files named <code>machinediscovery*</code>. ◆ Remove the subdirectory named <code>platespin</code>.

Component	Use Case	Removal Instructions
Controller software	All migrations	<ol style="list-style-type: none"> In the Servers view, undiscover the source (right-click, then select Undiscover). Open a command prompt and change the current directory to: <ul style="list-style-type: none"> \Program Files\platespin* (32-bit systems) \Program Files (x86)\platespin (64-bit systems) Run the following command: <code>ofxcontroller.exe /uninstall</code> Remove the platespin* directory

F.4.2 Cleaning Up Linux Workloads

The following are instructions for cleaning up Linux workloads by component and use case.

Component	Use Case	Removal Instructions
Controller software	Offline migrations	In the source workload's file system, under /boot, remove the ofx directory with its contents.
	All live migrations	<ul style="list-style-type: none"> Kill these processes: <ul style="list-style-type: none"> <code>kill -9 ofxcontrollerd</code> <code>kill -9 ofxjobexec</code> remove the OFX controller rpm package: <code>rpm -e ofxcontrollerd</code> In the source workload's file system, remove the /usr/lib/ofx directory with its contents.

Component	Use Case	Removal Instructions
Block-level data transfer software	All block-level migrations	<ol style="list-style-type: none"> 1. Check if the driver is active: <pre>lsmod grep blkwatch</pre> <p>If the driver is still loaded in memory, the result should contain a line, similar to the following:</p> <pre>blkwatch_7616 70924 0</pre> 2. (Conditional) If the driver is still loaded, remove it from memory: <pre>rmmmod blkwatch_7616</pre> 3. Remove the driver from the boot sequence: <pre>blkconfig -u</pre> 4. Remove the driver files by deleting the following directory with its contents: <pre>/lib/modules/[Kernel_Version]/Platespin</pre> 5. Delete the following file: <pre>/etc/blkwatch.conf</pre>
LVM snapshots	Block-level migrations using LVM snapshots	<ol style="list-style-type: none"> 1. In the Jobs view, generate a Job Report for the failed job, then note the name of the snapshot. 2. Remove the snapshot device by using the following command: <pre>lvremove <i>snapshot_name</i></pre>

F.5 Shrinking the PlateSpin Migrate Databases

When the PlateSpin Migrate databases (OFX and PortabilitySuite) reach a predetermined capacity, cleanup on those databases occurs at regular intervals. If there is a need to further regulate the size or content of those databases, Migrate provides a utility (`PlateSpin.DBCleanup.exe`) to further clean up and shrink those databases. [KB Article 7006458 \(https://www.netiq.com/support/kb/doc.php?id=7006458\)](https://www.netiq.com/support/kb/doc.php?id=7006458) explains the location of the tool and the options available for it, should you decide to use it for offline database operations.

G Reference of Tables

This section provides a list of links to important tables used throughout the documentation accompanying PlateSpin Migrate.

Table

[Table 1-1, "Workload Portability Operations," on page 14](#)

[Table 1-2, "PlateSpin Migrate Workload Migration Tasks," on page 15](#)

["Supported Microsoft Windows Workloads" on page 17](#)

[Table 1-4, "Supported Linux Workloads," on page 19](#)

[Table 1-5, "Supported Target Virtualization Platforms," on page 20](#)

[Table 2-1, "PlateSpin Migrate Workload License Assignment by Migration Type," on page 29](#)

[Table 2-2, "License Manager Command Buttons," on page 30](#)

[Table 2-3, "PlateSpin Migrate Roles and Permission Details For PlateSpin Migrate Client Users," on page 32](#)

[Table 2-5, "Network Communication Prerequisites for Discovery Operations," on page 40](#)

[Table 2-6, "Network Communication Prerequisites for Workload Portability," on page 42](#)

[Table 2-7, "Default and Optimized File Transfer Configuration Parameters," on page 52](#)

[Table 2-8, "List of Common Use Cases for changing settings in the Web configuration tool," on page 55](#)

[Table 3-1, "Machine-Specific Icons in the Servers View," on page 62](#)

[Table 3-2, "Guidelines for Discovery Parameters," on page 70](#)

[Table 3-3, "Status Messages in PlateSpin Analyzer Test Results," on page 72](#)

[Table 3-4, "PlateSpin Analyzer Test Results Tab," on page 72](#)

[Table 6-1, "PlateSpin Image Server Host Requirements," on page 138](#)

[Table 7-1, "Source and Target Credentials," on page 148](#)

[Table F-1, "Common Issues and Solutions Related to Discovery Operations," on page 231](#)

[Table F-2, "Common Issues and Solutions Related to Peer-to-Peer Migrations \(Windows\)," on page 232](#)

[Table F-3, "Common Issues and Solutions Related to PlateSpin Images," on page 233](#)

Glossary

PlateSpin Migrate Client. The client application you use to interact with the PlateSpin Server, discover source workloads and targets; set up, execute, and monitor jobs; manage license keys; and configure the default behavior of the server.

PlateSpin Image Server. A machine on which you have installed the PlateSpin Image Server software. See [PlateSpin Image](#).

PlateSpin Image. (Formerly *PlateSpin Flexible Image*) One of PlateSpin Migrate's three fundamental workload infrastructures along with physical machines and virtualization platforms. A PlateSpin Image is a static stored copy of the state of a physical or virtual machine (including volume data and the configuration details of the workload's hardware profile, operating system, and network identity), captured at a specific point in time. To make a PlateSpin Image bootable, you convert it to a workload on physical hardware or a virtualization platform. See also [Capture Image](#), [Deploy Image](#), and [Import Image](#).

PlateSpin Migrate Network. A named collection of discovered workloads and target machines you work with at any given time. A parameter of all discovery jobs.

PlateSpin Server. The remote runtime execution engine underlying the PlateSpin Migrate product.

VSS Block-based Transfer Component. A PlateSpin Migrate software component, dynamically installed on Windows source workloads to facilitate volume data transfer at the block level by using the Microsoft Volume Snapshot Service (VSS). See also [Block-Based Transfer Component](#), [File-based Transfer Component](#).

Block-Based Transfer Component. A PlateSpin Migrate software component, dynamically installed on source workloads to facilitate volume data transfer at the block level. *Compare with* [VSS Block-based Transfer Component](#); see also [File-based Transfer Component](#).

File-based Transfer Component. A PlateSpin Migrate software component, dynamically installed on a source workload to facilitate volume data transfer at the file level. See also [Block-Based Transfer Component](#), [VSS Block-based Transfer Component](#).

Auto-Discovery. See [Network Discovery](#).

Capture Image. A PlateSpin Migrate job that captures an image of a workload in [PlateSpin Image](#) format with an optional synchronization schedule for updating incremental changes. See also [Deploy Image](#).

Controller. See [OFX Controller](#).

Conversion. Any workload portability operation irrespective of target infrastructure or business purpose (workload-hardware decoupling, virtualization, or volume synchronization). Used interchangeably with the term [Migration](#) throughout product collateral, the user interface, error messages, and logs.

Copy. A PlateSpin Migrate peer-to-peer migration job that creates a duplicate of a workload, with a new network identity, on a physical or virtual machine. *Compare with* [Move](#).

Deploy Image. A PlateSpin Migrate job that converts a PlateSpin Image to a bootable workload on physical hardware or on a virtual machine. See also [Capture Image](#).

Discover. See [Discover Details](#).

Discover Details. A PlateSpin Migrate job that inventories the details of a supported workload or target machine in preparation for a migration operation. *Compare with* [Network Discovery](#).

I2P. Image-to-physical. A workload portability operation in which the source is a workload's PlateSpin Image and the target is a bootable workload on physical hardware.

I2V. Image-to-virtual. A workload portability operation in which the source is a workload's PlateSpin Image and the target is a bootable workload on a virtualization platform.

I2X. Image-to-anything. A workload portability operation in which the source is a workload's PlateSpin Image and the target is a bootable workload either on physical hardware or on a virtualization platform.

Image. See [PlateSpin Image](#).

Image server. See [PlateSpin Image Server](#).

Import Image. A PlateSpin Migrate job that creates a PlateSpin Image by using either raw volume data or third-party volume archives.

Job. A collection of the essential parameters of a workload portability, discovery, removal (undiscovery), or related operation that you define, save, schedule, execute, and monitor in the PlateSpin Migrate Client.

Migration. Any peer-to-peer workload portability operation in which a physical or virtual workload is copied or moved into another physical or virtual infrastructure.

Move. A PlateSpin Migrate peer-to-peer workload migration task that moves a workload to a physical or virtual machine. *Compare with* [Copy](#).

Network Discovery. A PlateSpin Migrate feature, based on standard Windows network browsing functionality, that automatically populates the PlateSpin Migrate Client's Servers view with the names of Windows machines in the neighborhood. *Compare with* [Discover Details](#).

Offline Migration. A type of migration, during which the source is shut down and booted into a temporary pre-execution environment to enable communication among the source, the target, and the PlateSpin Server.

OFX Controller. Operations Framework Controller. Software component that PlateSpin Migrate installs on source and target machines to enable them to communicate with PlateSpin Server and with each other.

P2C. Physical-to-cloud. A workload portability operation in which the target is Microsoft Azure and the source is a physical machine.

P2I. Physical-to-image. A workload portability operation in which the source is a physical machine and the target is the workload's PlateSpin Image.

P2P. Physical-to-physical. A peer-to-peer workload portability operation in which the source is a physical machine and the target is another physical machine on different hardware.

P2V. Physical-to-virtual. A peer-to-peer workload portability operation in which the source is a physical machine and the target is virtual machine.

P2X. Physical-to-anything. A workload portability operation in which the source is a physical machine and the target is another physical machine (P2P migration), a virtual machine (P2V migration), or a PlateSpin Image (P2I migration).

Peer-to-Peer. A workload portability operation in which both the source and the target are bootable workloads, either physical or virtual (as opposed to operations in which either the source or the target is a static PlateSpin Image). P2P, P2V, V2V, and V2P are all peer-to-peer portability operations; I2X and X2I are not.

Prepare for Synchronization. A step in workload synchronization operations where the target is a virtual machine in a Server Sync job. It automates the configuration and preliminary booting of the target virtual machine, and registers it as a valid Server Sync target in the PlateSpin Migrate Client's user interface. See [Server Sync](#).

Raw Volume Data. A directory of files that constitute a volume suitable for a PlateSpin Image.

Role. One of the three kinds of Windows local user groups that define PlateSpin Migrate users with specific access permissions: PlateSpin Migrate Administrator, PlateSpin Migrate Power User, and PlateSpin Migrate Operator.

Semi-Automated Virtualization. A workload virtualization operation in which a workload's target infrastructure is a virtual machine that is treated by PlateSpin Migrate as a physical machine. Applies to target virtualization platforms for which PlateSpin Migrate provides limited workload portability automation. *Compare with* [X2P](#).

Server. See [PlateSpin Server](#).

Server Sync. A workload portability operation that synchronizes the state of a physical or virtual workload (Windows or Linux) with the state of another physical or virtual workload of the same OS profile.

Source. The origin, or the infrastructure of the origin, of a workload portability operation (such as a physical machine, a virtual machine, or a PlateSpin Image). *Compare with* [Target](#).

Take Control. 1. See [Offline Migration](#)

2. The temporary pre-execution environment into which a Windows or Linux workload is booted during offline migration.

Target. The outcome, or the infrastructure of the outcome, of a workload portability operation (such as a physical machine, a virtual machine, or a PlateSpin Image). *Compare with* [Source](#).

Undiscover. A PlateSpin Migrate job that discards information about a workload from its inventory and removes any controllers installed. See *also* [Discover Details](#).

V2C. Virtual-to-cloud. A workload portability operation in which the target is Microsoft Azure and the source is a virtual machine.

V2I. Virtual-to-image. A workload portability operation in which the source is a virtual machine and the target is a PlateSpin Image.

V2P. Virtual-to-physical. A peer-to-peer workload portability operation in which the source is a virtual machine and the target is a physical machine.

V2V. Virtual-to-virtual. A workload portability operation in which both the source and the target are virtual machines on similar or different virtualization platforms.

V2X. Virtual-to-anything. A workload portability operation in which the source is a virtual machine and the target is another virtual machine, a physical machine, or a workload's PlateSpin Image.

Workload. Any physical or virtual instance of a supported operating system, along with its applications and data.

X2I. Anything-to-image. A workload portability operation in which the target is a workload's PlateSpin Image and the source is a physical machine or a virtual machine.

X2P. Anything-to-physical. A workload portability operation in which the target is a physical machine and the source is another physical machine, a virtual machine, or a workload's PlateSpin Image.

X2V. Anything-to-virtual. A workload portability operation in which the target is a virtual machine and the source is another virtual machine, a physical machine, or a workload's PlateSpin Image.