User Guide

PlateSpin® Migrate 9.1

November 19, 2012

Novell.

Contents

About This Guide

This guide provides information about using PlateSpin Migrate version 9.1.

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.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation, or submit your comments through the Novell Documentation Feedback site (http://www.novell.com/documentation/feedback.html).

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 9 Online Documentation Web Site (http://www.novell.com/documentation/platespin_migrate_9).

Documentation Updates

The most recent version of this guide can be found at the PlateSpin Migrate 9 Online Documentation Web Site (http://www.novell.com/documentation/platespin_migrate_9).

Additional Resources

We encourage you to use the following additional resources on the Web:

- Novell User Forum (http://forums.novell.com/): A Web-based community with a variety of discussion topics.
- Novell Support Knowledgebase (http://www.novell.com/support/): A collection of in-depth technical articles.

Technical Support

- Telephone (North America): +1-877-528-3774 (1 87 PlateSpin)
- Telephone (global): +1-416-203-4799
- E-mail: support@platespin.com

You can also request support through the online Service Request Web page (http://support.novell.com/contact/getsupport.html).

1 PlateSpin Migrate Feature Overview

PlateSpin Migrate enables you to migrate heterogeneous workloads (operating systems along with their applications and data) across x86 server and desktop infrastructure in the data center, decoupling a workload's hardware from software.

- Section 1.1, "About Workload Portability," on page 9
- Section 1.2, "Supported Source Workloads and Target Platforms," on page 11
- Section 1.3, "Supported Transfer Methods," on page 14
- Section 1.4, "Security and Privacy," on page 16
- Section 1.5, "Performance," on page 17

1.1 About Workload Portability

PlateSpin Migrate automates the migration of workloads among three infrastructures: physical, virtual machine, and volume imaging.

Figure 1-1 *Workload Portability*
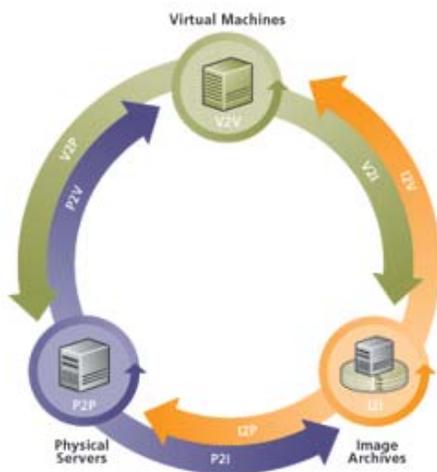
Table 1-1 *Workload Portability Operations*

Category of Operation	Migration Infrastructures
Workload Portability: Peer-to-peer	◆ Physical to Virtual (P2V)
	◆ Virtual to Virtual (V2V)
	◆ Virtual to Physical (V2P)
	◆ Physical to Physical (P2P)

Category of Operation	Migration Infrastructures
Workload Portability: Imaging	◆ Physical to Image (P2I)
	◆ Virtual to Image (V2I)
	◆ Image to Virtual (I2V)
	◆ Image to Physical (I2P)

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 Source Workloads and Target Platforms" on page 11.

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.

Task	Description
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.
Import Image	Imports third-party volume archives or raw volume data into a PlateSpin Image.
Deploy Image	Converts a PlateSpin Image into a booted or bootable workload on a physical or virtual machine.

1.2 Supported Source Workloads and Target Platforms

1.2.1 Supported Source Workloads

PlateSpin Migrate supports the following operating system families:

- Microsoft Windows, including Windows clusters
- Linux
- UNIX

The following topics provide more details:

NOTE: Not all workloads are supported on all target VM platforms. For details, see KB Article 7005871 (http://www.novell.com/support/viewContent.do?externalId=7005871).

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 imaging with PlateSpin Images (X2I, I2X), with full or partial support for several third-party backup and archiving applications.

Table 1-3 *Supported Microsoft Windows Workloads*

Operating System	Remarks
Windows 7	

Operating System	Remarks
Windows Server 2008 R2	Including domain controller (DC) systems and Small Business Server (SBS) editions
Windows Server 2008	
Windows Server 2003	
Windows Vista, Business/Enterprise/Ultimate, SP1 and later	
Windows Server 2000	
Windows Advanced Server 2000	
Windows XP Professional	
Windows NT 4.0 Server	SP 6a (required). Only file-level transfer method supported. 32-bit only
Windows clusters	Supported only to targets on VMware ESX 3.0.2 and later. See "Migrating Windows Clusters" on page 78

NOTE:

Live migration is not supported for the following workloads (use the Offline method):

- Windows NT 4
- Windows Server 2000 systems that do not have Service Pack 4
- Windows workloads with FAT volumes

Offline migration is not supported for workloads running Windows 7 and Windows Server 2008 R2.

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

- English
- French
- German
- Japanese
- Chinese

Supported Linux Workloads

For the Linux platform, PlateSpin Migrate provides Live and Offline peer-to-peer workload portability support (P2P, P2V, V2P, V2V), including workload synchronization with Server Sync.

Supported Linux file systems are EXT2, EXT3, REISERFS, and NSS (OES 2 workloads).

Workload imaging is not supported for Linux.

Table 1-4   *Supported Linux Workloads*

Distribution	Bit Spec
Red Hat Enterprise Linux (RHEL) AS/ES/WS 4, 5	32-bit, 64-bit
SUSE Linux Enterprise Server (SLES) 9, 10, 11	32-bit, 64-bit
Open Enterprise Server (OES) 2, SP2 and SP3	32-bit, 64-bit
Oracle Enterprise Linux (OEL) 5.3, 5.4	32-bit, 64-bit
Red Hat Linux 8	32-bit

NOTE:

Linux-based source workloads must be running a Secure Shell (SSH) server.

Live transfer is supported for these Linux workloads:

- RHEL 4, 5
- SLES 9, 10, 11

For live block-level migration of Linux workloads with customized kernels, you might have to rebuild the PlateSpin block-based module included in your PlateSpin Migrate distribution. See KB Article 7005873 (http://www.novell.com/support/viewContent.do?externalId=7005873).

Supported UNIX Workloads

PlateSpin Migrate provides migration support for Solaris workloads running on SPARC systems.

The supported versions are Solaris 10, Update 5 (U5) and Solaris 10, Update 6 (U6). You can migrate these workloads into native zones in Solaris 10 zone servers that are running the same update version.

Workload imaging is not supported for Solaris.

NOTE: Solaris source workloads must be running a Secure Shell (SSH) server.

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 7005871 (http://www.novell.com/support/viewContent.do?externalId=7005871).

Table 1-5 *Supported Target Virtualization Platforms*

Platform	Notes
VMware vCenter 4.0, 4.1, 4.1 Update 1	
VMware ESX 3.5, 4.0, 4.1, 4.1 Update 1	
VMware ESX**i** 3.5, 4.0, 4.1, 4.1 Update 1	All ESX**i** versions must have a paid license; migration is unsupported with these systems if they are operating with a free license.
Solaris 10 (U5 and U6) zone server	For source workloads that are Solaris 10 (U5 and U6) physical machines or zones
Citrix XenServer 5.5	Supported only for Windows workloads except Windows NT 4 (no support for Linux workloads)
Novell Xen Hypervisor on SLES 10 SP2	Fully virtualized VM targets: supported through Semi-Automated Workload Virtualization Using the X2P Workflow. Paravirtualized VM targets: see Linux Migration to a Paravirtualized VM on Xen-on-SLES.
Microsoft Hyper-V 2008	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 Replication Options" on page 97.

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 117). 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 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, PlateSpin Migrate leverages the LVM snapshot feature, if it is available on the host. If the feature is unavailable, PlateSpin Migrate uses its own block-level data transfer mechanism. For live block-level migration of Linux workloads with customized kernels, you might have to rebuild the PlateSpin `blkwatch` module included in your PlateSpin Migrate distribution. See KB Article 7005873 (http://www.novell.com/support/viewContent.do?externalId=7005873).

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

Offline Transfer of Windows and Linux Workloads

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.

The mechanism of the pre-execution environment underlying the Offline transfer method depends on the source workload's operating system:

 ◆ For Windows workloads, PlateSpin Migrate uses the Microsoft Windows Preinstallation Environment (WinPE).
 ◆ For Linux workloads, PlateSpin Migrate uses a Linux Ramdisk.

To ensure that the source operating system loads the temporary pre-execution environment on reboot, PlateSpin Migrate temporarily modifies its boot files and restores them to their original state after the pre-execution environment has successfully loaded.

Offline Transfer of Solaris Workloads

For Solaris workloads, PlateSpin Migrate uses a different offline mechanism from the one it uses for Windows and Linux workloads; there is no temporary pre-execution environment involved in the preparation of the workload for a migration. Instead, PlateSpin Migrate changes the source workload's run level while the migration is underway.

NOTE: If you have customized services on your Solaris source workload, the system might fail to shut them down during data transfer, and this might result in a failure of the migration job. Ensure that customized services on your source are shut down before setting up the migration job.

1.4 Security and Privacy

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

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 Migrate 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 Migrate Server installation. See "Enabling Support for FIPS-Compliant Data Encryption Algorithms (Optional)" in your *Installation Guide*.

For information on how to enable file encryption for a migration job, see "Specifying Replication Options" on page 97.

1.4.2 Security of Client-Server Communications

Data transmission between the PlateSpin Migrate 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 Migrate Server host and use HTTPS when specifying the server URL. See "Connecting to a PlateSpin Migrate Server" on page 39.

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 Migrate 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 2.2.1, "Setting Up User Authorization and Authentication," on page 24.

1.5 Performance

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 Migrate 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 35).

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 Migrate Server.

Bandwidth throttling is a parameter of a workload migration job's configuration properties. See

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 20
- Multiple migrations against multiple VM servers (ESX): no more than 30

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*.

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.

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 Novell Customer Center Web site (http://www.novell.com/customercenter/). A license activation code will be e-mailed to you.

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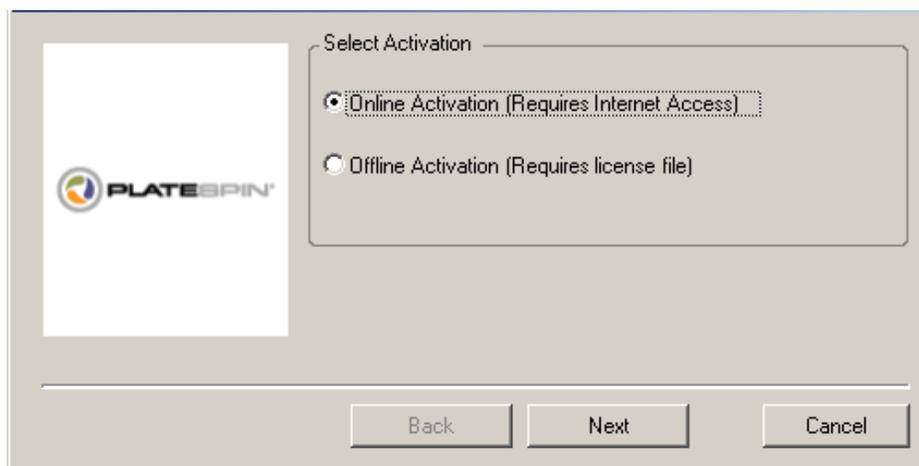
You have two options for activating your product license: online or offline.

- ◆ "Online Activation" on page 20
- ◆ "Offline Activation" on page 20

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).

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	Not applicable
Import Image	Not applicable
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 "Selecting a License Key for a Migration Job" in your *User Guide*.

- ◆ "Managing Available License Keys" on page 21
- ◆ "Managing Workload Designations" on page 22

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*
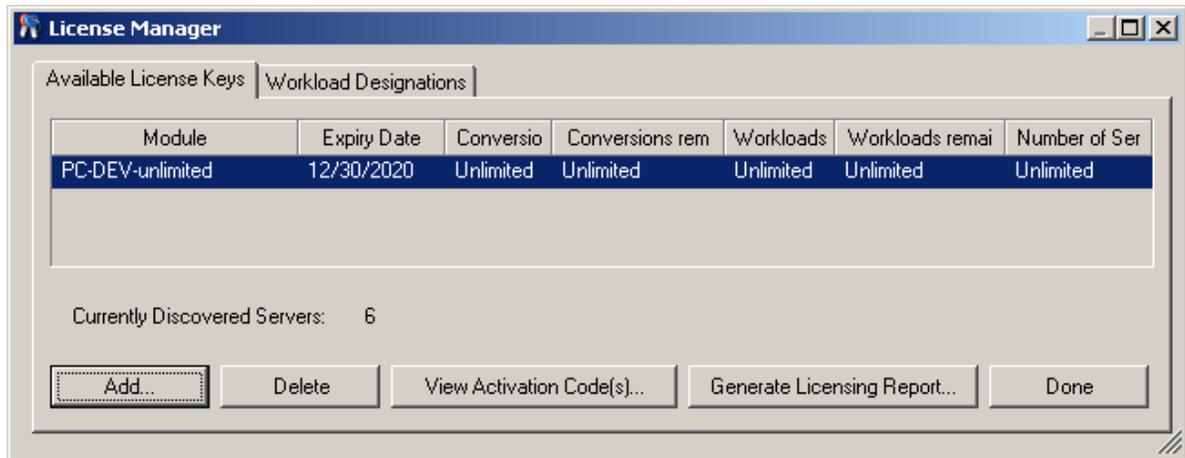
The tab displays the license name (*Module*) along with its expiry date and entitlements. These depend on the license type. The *Number of Servers* column indicates the number of machines you can discover. This is generally the same as the number of machines that you can migrate. Use the buttons at the 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 `*.psl` 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 Migrate 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 (http://www.novell.com/support/viewContent.do?externalId=7920876).

2.2 Configuring PlateSpin Migrate and Your Network Environment

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

2.2.1 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).

PlateSpin Migrate's user auditing functionality is provided through the capability to log user actions (see "Setting Up User Activity Logging" on page 25).

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 Migrate 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:

- **PlateSpin Administrators:** Have unlimited access to all features and functions of the application. A local administrator is implicitly part of this group.
- **PlateSpin 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:** 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 Migrate 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 Migrate Server host, that account is implicitly regarded as a PlateSpin Migrate Administrator.

The following table lists permissions for each role.

Table 2-3 *PlateSpin Migrate Roles and Permission Details*

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

Role Details	Administrators	Power Users	Operators
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 Migrate Server access: View Web services, download client software	yes	yes	yes
PlateSpin Migrate Server settings: Edit settings that control user activity logging and SMTP notifications	yes	no	no
PlateSpin Migrate 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

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 Migrate Server host. For more information, see your Windows documentation.

2.2.2 Setting Up User Activity Logging

By default, PlateSpin Migrate records all user activities in a log file, `PlateSpin.UserActivityLogging.log`, located on your PlateSpin Migrate 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.*

2.2.3 Access and Communication Requirements across your Migration Network

Migration sources and targets (such as production workloads and VM hosts), must meet certain access and network requirements.

- ◆ "Global Requirements" on page 26
- ◆ "Requirements for Discovery" on page 26
- ◆ "Requirements for Migration" on page 28

Global Requirements

All sources and targets must support ping traffic (ICMP echo request and response).

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 "Discovering Source Workloads and Migration Targets" in your *User Guide*.

Table 2-4 *Network Communication Prerequisites for Discovery Operations*

System	Prerequisites
All Windows sources	◆ Microsoft .NEt Framework version 2.0 or 3.5 SP1 ◆ Credentials with local or domain admin privileges

System	Prerequisites

Windows 7;
 	
Windows Server 2008;
 	
Windows Vista
	1. Built-in `Administrator` or a domain account credentials (mere membership in the local Administrators group is insufficient). On Vista, the account must be enabled (it is disabled by default).
	
	2. The Windows Firewall configured to allow *File and Printer Sharing.* Use one of these options:
	
	    ◆ **Option 1, using Windows Firewall:** Use the basic *Windows Firewall* Control Panel item (`firewell.cpl`) 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 (`wf.msc`) with the following *Inbound Rules* enabled and set to `Allow`:
	
	        ◆ *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.

Windows Server 2000; Windows XP; Windows NT 4	◆ Windows Management Instrumentation (WMI) installed ◆ Open ports 135/445 (TCP) for DCOM/RPC Windows NT Server does not include WMI as part of the default installation. Obtain the WMI Core from the Microsoft Web site. If WMI is not installed, discovery of the workload fails. WMI (RPC/DCOM) can use TCP ports 135 and 445 as well as random or dynamically assigned ports above 1024. If problems occur during the discovery process, consider temporarily placing the workload in a DMZ or temporarily opening the firewalled ports for the discovery process only. For additional information, such as guidance in limiting the port range for DCOM and RPC, see the following Microsoft technical articles. ◆ How to configure RPC dynamic port allocation to work with firewalls (http://support.microsoft.com/kb/154596/EN-US) ◆ Configuring RPC dynamic port allocation to work with firewalls (http://support.microsoft.com/default.aspx?scid=kb;en-us;154596) ◆ Configuring DCOM to work over a NAT-based firewall (http://support.microsoft.com/kb/248809)

System	Prerequisites
All Linux sources Citrix Xen Server	◆ Secure Shell (SSH) server ◆ Open port 22 (TCP) ◆ Root-level access. For information on using an account other than root, see KB Article 7920711 (http://www.novell.com/support/viewContent.do?externalId=7920711). ◆ Custom SSH ports are supported; specify the port number during discovery: *<hostname	IP_address>:port_number*.
VMware ESX Server 3.5, ESX 4, 4.1, ESXi; vCenter Server	◆ VMware account with an Administrator role ◆ VMware Web services API and file management API (HTTPS / port 443 TCP)
VMware ESX Server 3.0.*x*	◆ VMware account with an Administrator role ◆ Secure Shell (SSH) server ◆ Open port 22 (TCP) ◆ If you are using the root account, configure the ESX server to enable shell access for the root account. ◆ Custom SSH ports are supported; specify the port number during discovery: *<hostname	IP_address>:port_number*.
All Solaris source workloads and targets	The user must be assigned a primary administrator profile through Solaris RBAC (role-based access control). See KB Article 7920711 (http://www.novell.com/support/viewContent.do?externalId=7920711).

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 2-5 *Network Communication Prerequisites for Workload Portability*

System	Open Port	Remarks
PlateSpin Migrate Server hosts	Either 80 (TCP) or 443 (TCP)	◆ Port 80 (TCP) is required for HTTP communication among the PlateSpin Migrate Server, sources, and targets. ◆ Port 443 (TCP) is required for HTTPS communication (if SSL is used) between the PlateSpin Migrate Server and the source or target machines.
All source workloads except those in image deployment jobs.	3725 (TCP)	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.

System	Open Port	Remarks
All targets	3725 (TCP)	Required for: File-based Server Sync / Image Sync ◆ File-level Server Sync ◆ Image synchronization jobs
All Windows sources and targets	137 - 139 (NetBIOS)	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 VMware ESX Server 2.5 VMware ESX Server 3.0.x Citrix Xen Server	22 (TCP)	Required for communication during offline migration.
PlateSpin Migrate Server hosts; All Windows sources	135/445 (TCP)	For DCOM/RPC communication between PlateSpin Migrate 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.

2.2.4 Optimizing Data Transfer over WAN Connections (Windows)

When doing Windows migrations, you can fine-tune your over-the-network data transfer for optimum performance in your specific environment. For example, you might need to control the number of your TCP connections or impose a packet-level compression threshold.

This functionality is supported for migration jobs that use the following data transfer methods:

◆ File-level
◆ Block-level with Microsoft Volume Shadow Copy Service (VSS)

Fine-tuning is done by modifying the product's productinternal.config configuration file, located on your PlateSpin Migrate Server host in the following directory:

..\PlateSpin Migrate Server\Web

The following table lists the configuration parameters with two sets of values: the defaults and the values recommended for optimum operation in a high-latency WAN environment.

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 2-6 *Parameters for Fine-Tuning File-Level Data Transfer Performance*

Parameter	Default Value	For High-Latency WANs
fileTransferThreadcount	2	4 to 6 (max)
Controls the number of TCP connections opened for file-based data transfer.		
fileTransferMinCompressionLimit	0	max 65536 (64 KB)
Specifies the packet-level compression threshold in bytes.	(disabled)	
fileTransferCompressionThreadsCount	2	n/a
Controls the number of threads used for packet-level data compression. Ignored if compression is disabled.		
Because the compression is CPU-bound, this setting might have a performance impact during Live Transfer.		
fileTransferSendReceiveBufferSize	0 (8192 bytes)	max 5242880 (5 MB)
TCP window size setting for file transfer connections; controls the number of bytes sent without TCP acknowledgement.		
When the value is set to 0, the default TCP window size is used (8 KB). For custom sizes, specify the size in bytes.		
Use the following formula to determine the proper value:		
`((LINK_SPEED(Mbps)/8)*DELAY(sec))*1000*1000`		
For example, for a 100 Mbps link with 10 ms latency, the proper buffer size would be:		
`(100/8)*0.01 * 1000 * 1000 = 125000 bytes`		

2.2.5 Increasing the Size Limit on Post-migration Actions Uploaded to the PlateSpin Migrate Server

By default, PlateSpin Migrate sets a 64 MB upload limit for each individual post-migration action, including its dependencies. See "Managing Custom Actions" in your *User Guide*.

You can increase this limit by modifying the PlateSpin Migrate Server's `web.config` configuration file.

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

1 Using a text editor, open the `web.config` configuration file, located in the following directory:

   `..\PlateSpin Migrate Server\Web`

2 Locate the line defining the value of the `httpRuntime` element's `maxRequestLength` attribute:

   `<httpRuntime maxRequestLength="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`:

   `<httpRuntime maxRequestLength="16384" />`

4 Save the `web.config` file.

5 Restart the PlateSpin Migrate Client.

2.2.6 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 Migrate Server:** In your server's `web.config` configuration file, record the additional IP addresses assigned to that host:

    1. In a text editor, open the `\Program Files\PlateSpin Migrate Server\Web\Web.config` file.

    2. Locate the following line:

        `<add key="AlternateServerAddresses" value="" />`

    3. Add the additional IP addresses, delimited by a semicolon (;), for example:

        `<add key="AlternateServerAddresses" value="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 99.

- ◆ **Target:** When you are attempting to discover a target, such as VMware ESX, specify the public (or external) IP address in the discovery parameters.

2.2.7 Configuring PlateSpin Migrate Default Options

Default options control PlateSpin Migrate's global settings and its default behavior. Configure the default options following installation or to reflect changes in your environment.

- ◆ "General Options" on page 32
- ◆ "Notification Service" on page 33
- ◆ "User Activity Logging" on page 34
- ◆ "Default Job Values" on page 35
- ◆ "Source Service Defaults" on page 36
- ◆ "Target Service Defaults" on page 37

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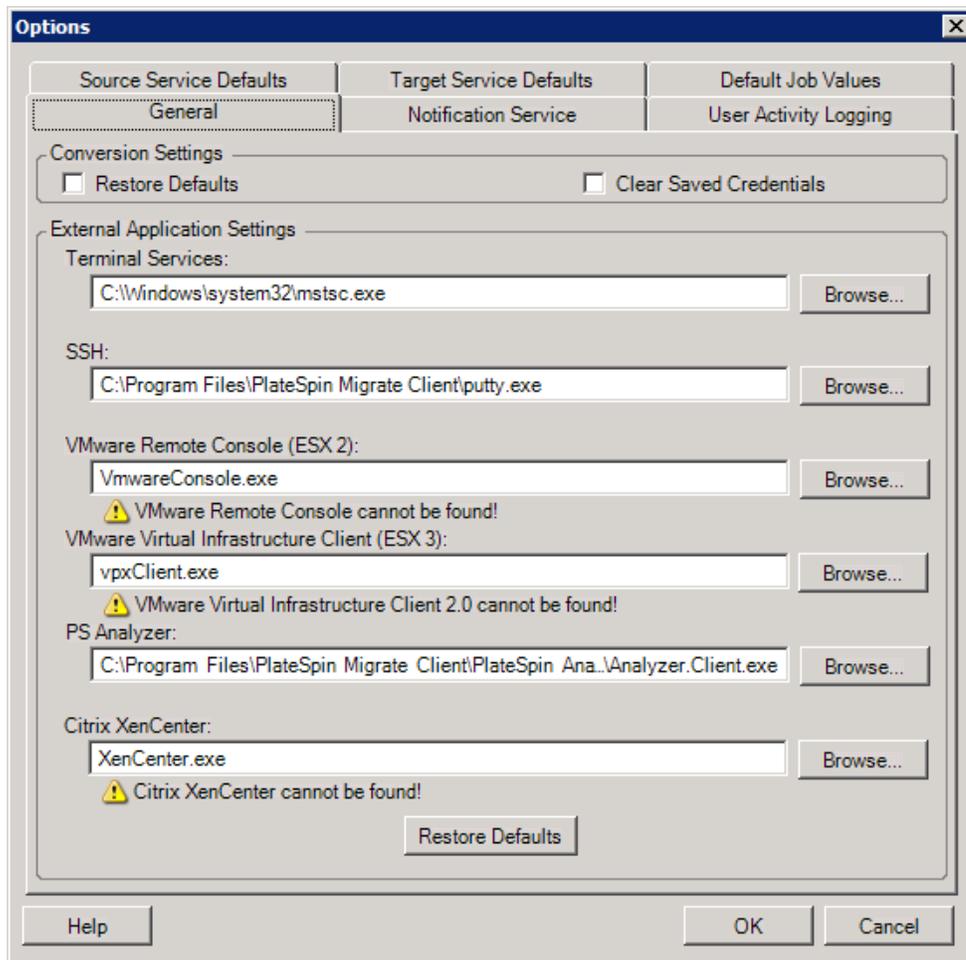
Restore Defaults: When this option is selected, PlateSpin Migrate resets the job configuration method (launches the Actions dialog box after a drag-and-drop) and resumes checking for software updates on the Client startup.

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.

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*.

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 "Configuring Automatic E-Mail Alerts of Job Status and Progress" in your *User Guide*.

User Activity Logging

Use this tab to specify options related to user activity logging. See "Setting Up User Activity Logging" on page 25.

To access this dialog box in the PlateSpin Migrate Client, click *Tools > Options*.
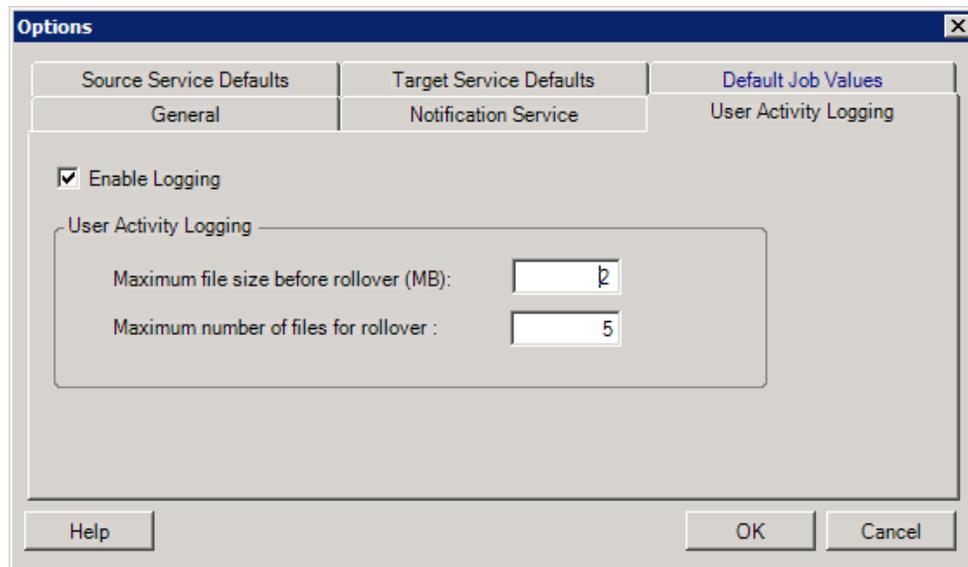
Enable Logging: When this option is selected, PlateSpin Migrate logs all user activities.

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.

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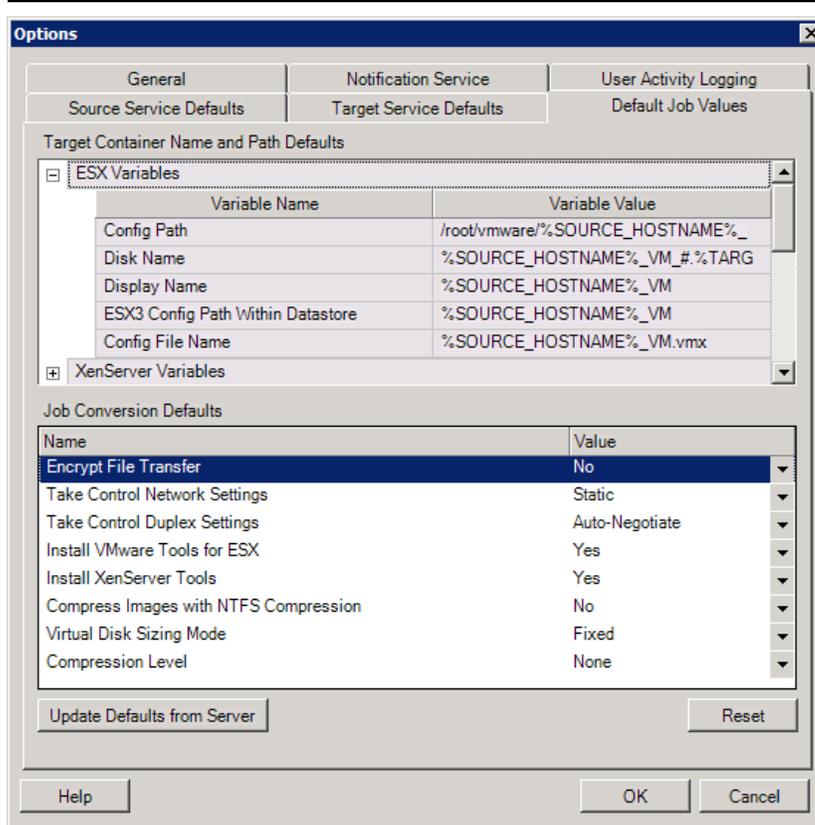
Target Container Name and Path Details: These variables control naming conventions for target paths and containers. Expand a variable set and click a variable to edit its value.

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" (*User Guide*).

- ◆ *Take Control Network* and Duplex Settings: See "Temporary (Take Control) Network Settings" (*User Guide*).

- ◆ *Install VMware Tools for ESX*/XenServer Tools: See "Handling Virtualization Enhancement Software" (*User Guide*).

- ◆ *Compress Images with NTFS Compression*: See "Capturing a Workload to a PlateSpin Image" (*User Guide*). Unrelated to data compression for over-the-network transfer.

- ◆ *Virtual Disk Sizing Mode*: (ESX only) *Fixed*—space is preallocated for the virtual disk; *Dynamic*—the virtual disk is assigned a minimum amount of space, which grows when needed.

- ◆ *Compression Level*: See "Data Compression" (*User Guide*).

Reset: Restores default job values.

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

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)" in your *User Guide*.

To access this dialog box in the PlateSpin Migrate Client, click *Tools > Options*.
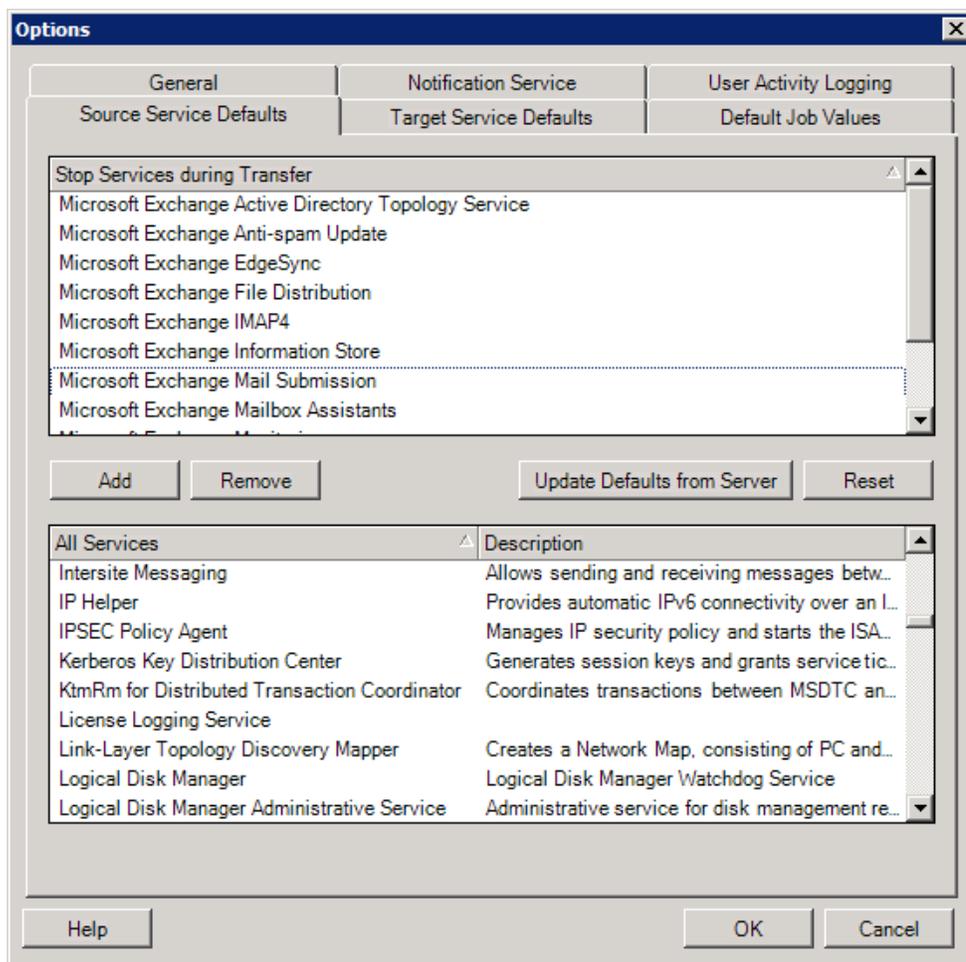
Stop Services during Transfer section: Lists services that are stopped by default. To stop a service during data transfer that uses a specific transfer method by default, select the corresponding check box. A deselected check box means the service remains active during Live Transfer.

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 Migrate Server.

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)" in your *User Guide*.

To access this dialog box in the PlateSpin Migrate Client, click *Tools > Options*.
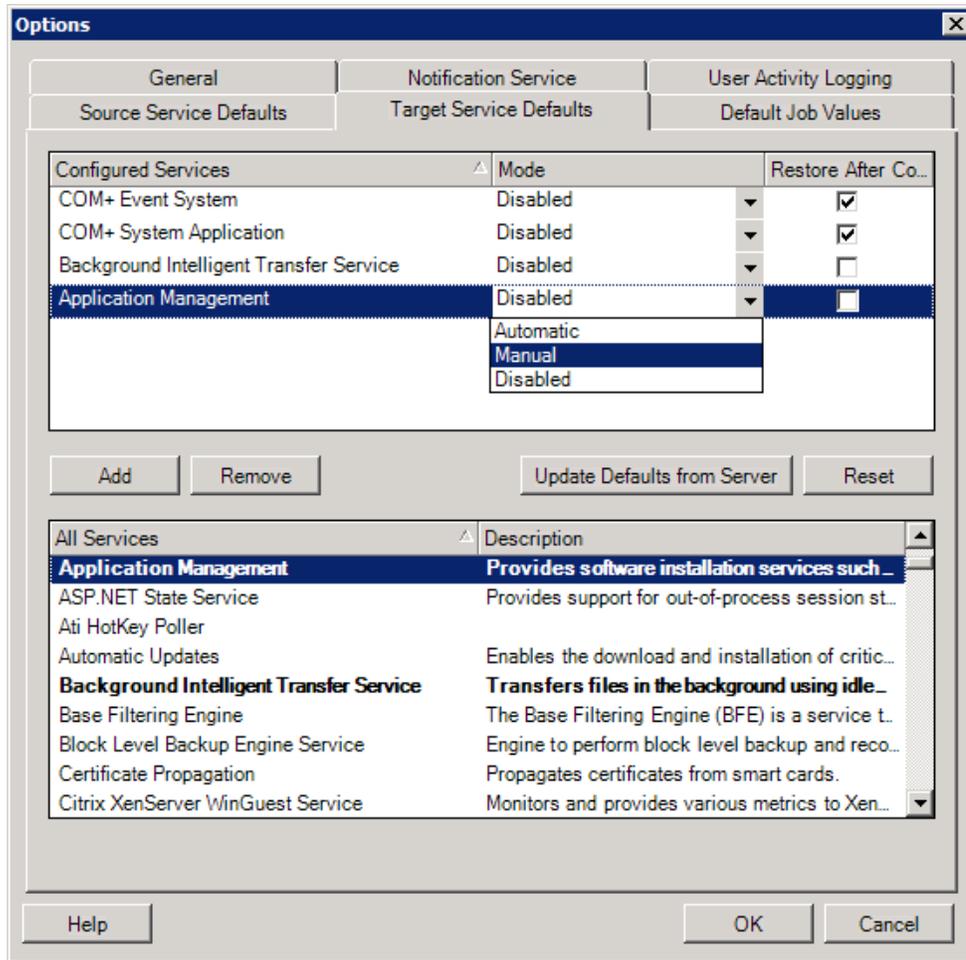
Configure Services section: Lists services and their target startup modes. Select the *Restore After Conversion* check box to use the selected mode during the migration. The service is then restored to match the source after the migration is complete and the target machine is ready to run.

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.

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 Migrate Server and provides access to information in the PlateSpin Migrate 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.

3.1 Connecting to a PlateSpin Migrate Server

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

- Performs authentication of the specified user account with the PlateSpin Migrate Server.

  See Section 2.2.1, "Setting Up User Authorization and Authentication," on page 24.
- Connects to a specified PlateSpin Migrate 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 Migrate Server instance, and the required PlateSpin Migrate Network in the PlateSpin Migrate Server settings.

1 In the PlateSpin Migrate Client, click *Tools* > *PlateSpin Migrate 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*.

Server: http://dev3/PlateSpinMigrate	Network: Sample Environment	User: devdomain\john.smith
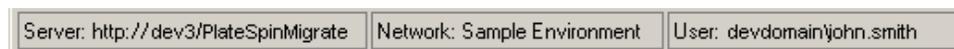
The PlateSpin Migrate Server Settings dialog box opens.
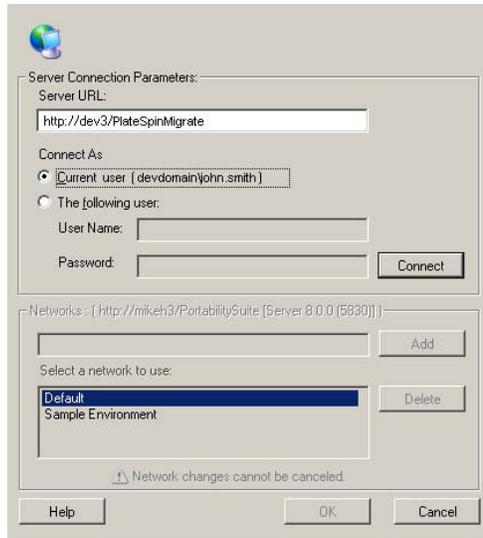
2 Specify the required PlateSpin Migrate Server URL, user, and network parameters as required:

Interface Element	Description
Server URL	Type the PlateSpin Migrate Server URL in the following format: `http://`***server_host***`/platespinmigrate` (If SSL is enabled on the PlateSpin Migrate Server host, replace `http` in the URL with `https`).
Connect As	To connect to a PlateSpin Migrate Server, you must have administrative access to the PlateSpin Migrate Server host or be a member of one of the PlateSpin Migrate roles. See Section 2.2.1, "Setting Up User Authorization and Authentication," on page 24.
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 Migrate 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 41
- Section 3.2.2, "Jobs View," on page 45
- Section 3.2.3, "Tasks Pane," on page 46
- Section 3.2.4, "Status Bar," on page 46

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*
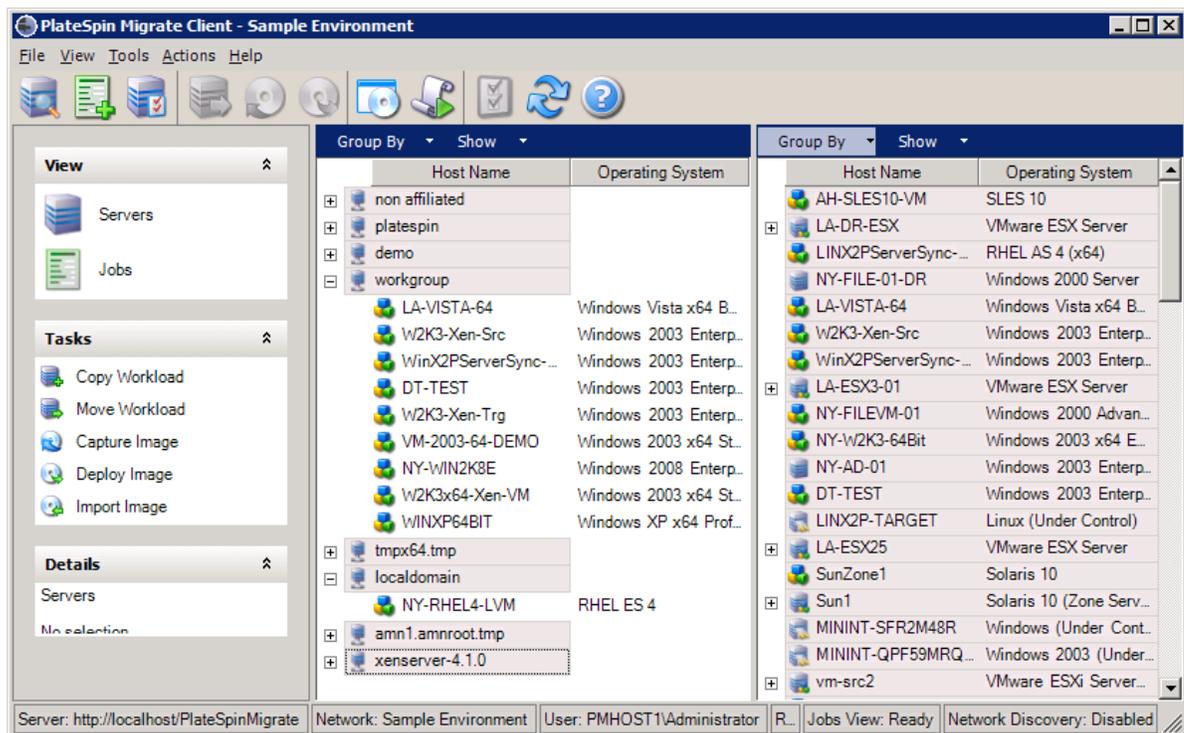
The hierarchical display of items in the Servers view reflects the organization of items on their respective platforms; for example: VMs are shown nested beneath their VM hosts, PlateSpin Images are beneath their image servers, and Solaris zones (regarded as VMs) are beneath their respective Solaris zone 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 42.

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 43.

- "Organizing the Servers View" on page 42
- "Viewing the Properties of Source Workloads and Targets" on page 43
- "List of Machine-Specific Icons in the Servers View" on page 44

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 NT 4, 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*
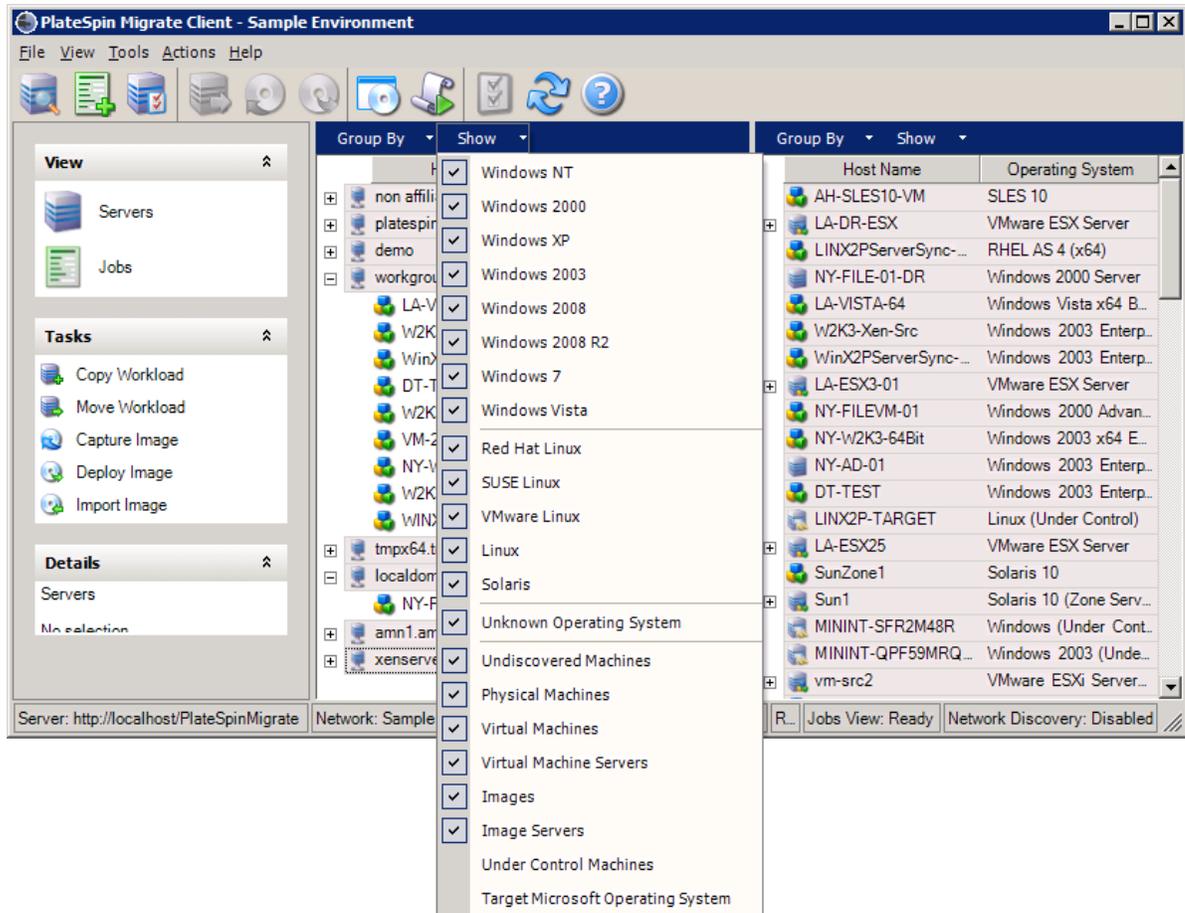
Viewing the Properties of Source Workloads and Targets

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 properties of Windows Cluster systems provide information about the basic cluster parameters of the inventoried cluster, such as cluster name, IP address, quorum, shared disks, and node type.

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*
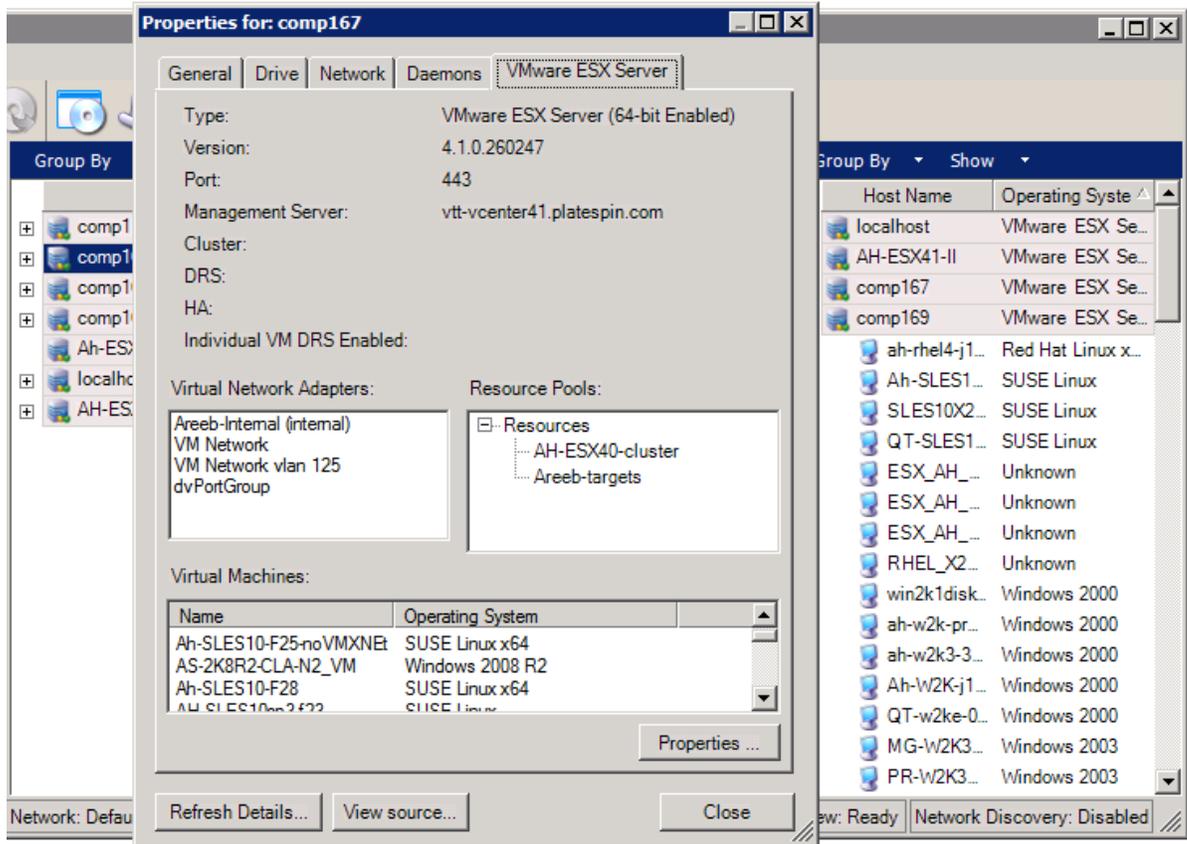
List of Machine-Specific Icons in the Servers View

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 .

Figure 3-4 *PlateSpin Migrate Client's Jobs View*
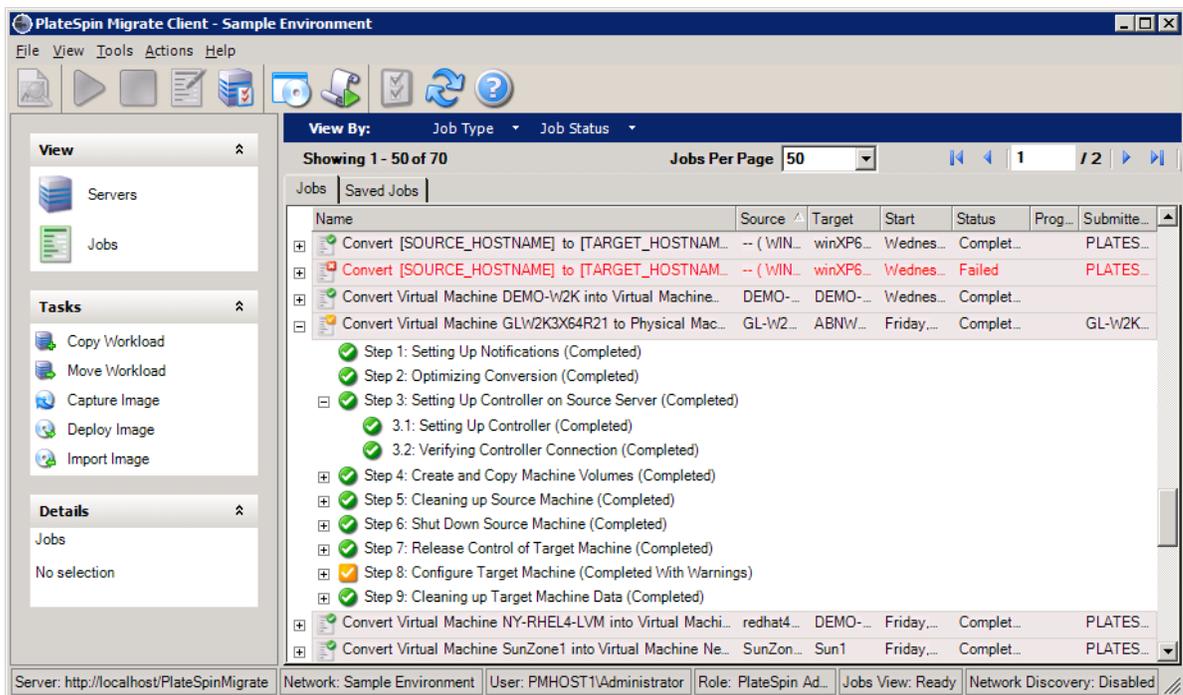
You can limit the scope of jobs displayed in the view. Use the *Job Type* and *Jobs Status* menus to specify filters for the view:

- **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. See Section 3.7.1, "Selecting a Migration Job Setup Method," on page 62.

3.2.4 Status Bar

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

- The PlateSpin Migrate 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*

| Server: http://localhost/PlateSpinMigrate | Network: Sample Environment | User: PMHOST1\Administrator | Role: PlateSpin Administrator | Jobs View: Ready | Network Discovery: Disabled |
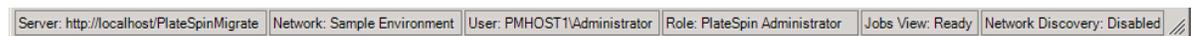
Double-clicking any of the first three status items opens the PlateSpin Migrate Server Settings window. See "Connecting to a PlateSpin Migrate Server" on page 39.

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 53.

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.

- Section 3.3.1, "Prepopulating the Servers View with Windows Computers in the Network," on page 47
- Section 3.3.2, "Discovering the Details of Source Workloads and Targets," on page 47
- Section 3.3.3, "Discovering All Machines in a Windows Domain," on page 48
- Section 3.3.4, "Refreshing Machine Details," on page 48
- Section 3.3.5, "Resetting the Mechanism Used to Authenticate Access to ESX Hosts," on page 49
- Section 3.3.6, "Discovering and Registering Target Physical Machines," on page 49
- Section 3.3.7, "Discovering, Registering and Preparing a VM for a Server Sync Job," on page 53
- Section 3.3.8, "Discovery Guidelines for Machine Types and Credentials," on page 54

3.3.1 Prepopulating 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 48.

PlateSpin Migrate uses the standard Windows network browser function for the auto-discovery feature. Because Linux and Solaris 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
- The active or passive node of a Windows cluster
- An individual virtual machine server
- Multiple virtual machine servers
- All VMware ESX hosts affiliated with a VMware vCenter Server
- Multiple machines

Before starting discovery operations, make sure PlateSpin Migrate Server can communicate with your source workloads and targets. See Section 2.2.3, "Access and Communication Requirements across your Migration Network," on page 26.

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 31.
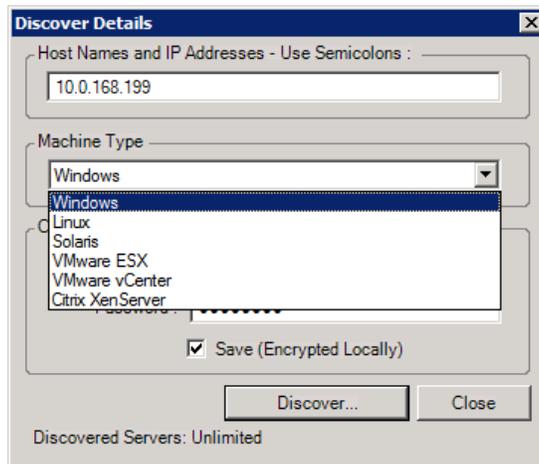
3 Select the machine type and provide administrator credentials for the machine you are discovering.

See "Discovery Guidelines for Machine Types and Credentials" on page 54.

Discovering hosts with Microsoft Hyper-V, Virtual Iron, and 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 73.

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 "Prepopulating the Servers View with Windows Computers in the Network" on page 47.

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 48.

The system prompts you to select an access type:
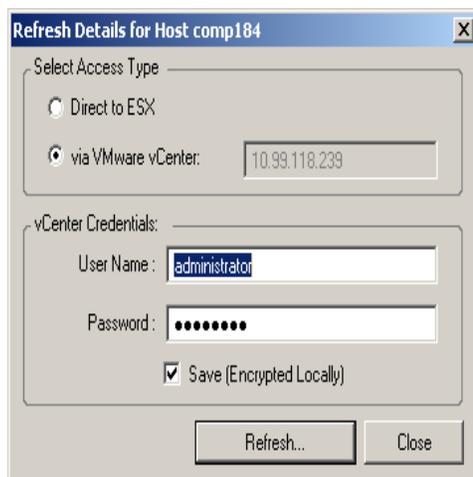
2 Select the required option, then click *Refresh*.

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 and Registering 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 73.

To use a boot ISO image, download it from the PlateSpin Migrate 9 area of Novell Downloads (http://download.novell.com/Download?buildid=DMYwAdg1ZWs). Use the image appropriate to your target machine:

Table 3-2 *Boot ISO Images for Target Physical Machines*

Filename	Remarks
`WindowsPhysicalTarget.zip` (contains `WindowsPhysicalTarget.iso`)	Windows systems
`LinuxFailback.zip` (contains `LinuxPhysicalTarget.iso`)	Linux systems
`WindowsPhysicalTarget-Cisco.zip` (contains `WindowsPhysicalTarget-Cisco.iso`)	Windows systems on Cisco hardware
`WindowsPhysicalTarget-Dell.zip` (contains `WindowsPhysicalTarget-Dell.iso`)	Windows systems on Dell hardware
`WindowsPhysicalTarget-Fujitsu.zip` (contains `WindowsPhysicalTarget-Fujitsu.iso`)	Windows systems on Fujitsu hardware

After downloading the required file, unzip and save the extracted ISO file.

- "Discovering and Registering Target Physical Machines (Windows)" on page 50
- "Discovering and Registering Target Physical Machines (Linux)" on page 52
- "Performing an Unattended Registration of a Target Physical Machine" on page 53

Discovering and Registering Target Physical Machines (Windows)

Complete these steps to register your target with PlateSpin Migrate 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 (http://www.novell.com/support/viewContent.do?externalId=7920832).

NOTE: All discovered and registered physical machines can be also used as Server Sync targets. See "Synchronizing Workloads with Server Sync" on page 72.

1 Burn the appropriate boot ISO image on a CD or save it to the required media, from which your target can boot.

See Table 3-2, "Boot ISO Images for Target Physical Machines," on page 50.

2 Ensure that the network switch port connected to the target is set to Auto Full Duplex.

Because the Windows version of the boot CD image supports Auto Negotiate Full Duplex only, this ensures that there are no conflicts in the duplex settings.

3 Using the boot image, boot the target physical machine and wait for the *REGISTERMACHINE* and *Recovery Console* command boxes to open.
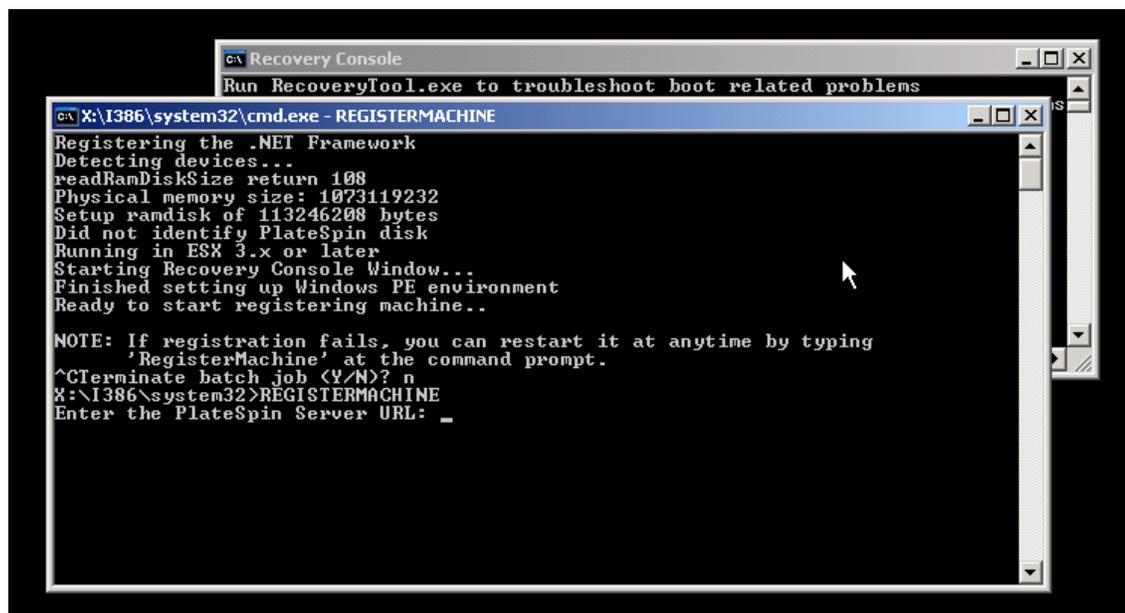
Use the REGISTERMACHINE command line utility. For information on the Recovery Console utility, see Using the Recovery Tool Command Line Utility (Windows) (page 52).

4 When prompted, enter the following URL:

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

Replace *<hostname | IP_address>* with the hostname or the IP address of your PlateSpin Migrate Server host.

5 Enter administrator-level credentials for the PlateSpin Migrate Server, specifying an authority. Use this format:

domain\username or *hostname\username*

6 Enter the password for the provided credentials.

Available network cards are detected and displayed by their MAC addresses.

7 If DHCP is available on the NIC to be used, press the Enter key to continue. If DHCP is not available, select the desired NIC to configure with a static IP address.

   7a Enter the IP address

   7b Enter the subnet mask

   7c Enter the gateway

8 Enter a hostname for the physical machine or press the Enter key to accept the default values.

9 Enter *Yes* if you have enabled SSL on the PlateSpin Migrate Server; otherwise enter *No*.

10 Enter the name of the PlateSpin Migrate Network that you want the discovered machine to be part of. If you are unsure, just press the Enter key. PlateSpin Migrate Server uses the Default Network.

After a few moments you should see the new physical target in the Servers view of the PlateSpin Migrate Client. Should the target machine throw driver-related general protection errors, consider using the Recovery Console utility to inject alternative drivers.

Using the Recovery Tool Command Line Utility (Windows)

The Recovery Console command line utility enables you to dynamically inject Windows device drivers into the target physical machine without having to restart the entire physical target registration process.

The utility is loaded in a secondary command box upon the initial attempt to boot from the Windows boot image (see).

To use the Recovery Tool, type its command name, RECOVERYTOOL, followed by an applicable parameter, in the Recovery Console window.

You can use:

- /L - to list any driver services installed on the target OS
- /J - to inject drivers into the target OS

You can specify whether the drivers are to be downloaded from the PlateSpin Migrate Server or from a local path. If you intend to use a local path, you should group multiple drivers for the same device together. If you want to download drivers from the PlateSpin Migrate Server, the utility prompts you to specify which driver you want to use (if there is more than one).

Discovering and Registering Target Physical Machines (Linux)

Complete these steps to register your target with PlateSpin Migrate 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 (http://www.novell.com/support/viewContent.do?externalId=7920832).

You can use a custom utility to package and inject additional Linux device drivers into the ISO boot image. See Table 3-2 on page 50. If you want to inject drivers into the ISO image, follow these steps first:

1 Obtain or compile the required *.ko driver files.

IMPORTANT: Make sure the drivers are valid for the kernel included with the ISO file (2.6.32.12-0.7-default) and are appropriate for the target architecture.

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

mount –o loop *<path-to-ISO>* *<mount_point>*

3 Copy the rebuildiso.sh script, located in the /tools subdirectory of the mounted ISO file, into a temporary working directory.

4 Create another working directory for the required driver files and save them in that directory.

5 In the directory where you saved the rebuildiso.sh script, run the following command as root:

./rebuildiso.sh –i *<ISO_file>* -d *<driver_dir>* -m i586|x86_64

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

6 Unmount the ISO file (execute the command `unmount <mount_point>`).

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

1 Burn the appropriate 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` (for systems with up to 512 MB RAM)

   ◆ `ps64_512m` (for systems with more than 512 MB RAM)

   Press Enter.

4 When prompted, enter the PlateSpin Migrate 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 Migrate Server host.

5 Enter your credentials for the PlateSpin Migrate 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 Migrate Server. This involves updating the boot ISO image with specific registration information before booting the target.

For details, see KB Article 20832 (http://support.platespin.com/kb2/article.aspx?id=20832).

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 72.

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 3-3 *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: ◆ For domain member machines: *authority\principal* ◆ For workgroup member machines: *hostname\principal*
Windows Clusters	Windows	Cluster admin credentials	In discovery parameters, use the cluster's virtual IP address. If you use the IP address of an individual Windows cluster node, that node is discovered as a regular (cluster-unaware) Windows workload.
All Linux workloads	*Linux*	Root-level username and password	Non-root accounts must be properly configured to use sudo. See KB Article 7920711 (http://www.novell.com/support/viewContent.do?externalId=7920711).
VMware ESX hosts	*VMware ESX*	ESX account with admin role OR Windows domain credentials (versions 4 and 4.1 only)	
Linux, VMware ESX Server 2.5 and 3.0.x systems with non-default SSH ports	*Linux* or *VMware ESX*, as applicable	ESX account with admin role	Use the standard URL format including the port number: *hostname:portnumber* or *IP_address:portnumber* as applicable. Port numbers are not required for VMware ESX Server 3.5, ESXi 3.5, and ESX 4 systems.

To Discover	Machine Type Selection	Credentials	Remarks
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 49.
Citrix XenServer VM hosts	*Citrix XenServer*	username `root` along with the associated password	XenServer hosts currently do not support multiple root-level user accounts
Solaris workloads and zone servers	*Solaris*	username and password	The user must be assigned a primary administrator profile through Solaris RBAC (role-based access control). See KB Article 7920711 (http://www.novell.com/support/viewContent.do?externalId=7920711).

3.4 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 Migrate Server, click *Connect*, then specify the required PlateSpin Migrate 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 3-4 *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 3-5 *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 (http://www.novell.com/support/viewContent.do?externalId=7920478).

3.5 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 55.

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 3.5.1, "Packaging Device Drivers for Windows Systems," on page 57
- Section 3.5.2, "Packaging Device Drivers for Linux Systems," on page 57
- Section 3.5.3, "Uploading Drivers to the PlateSpin Migrate Device Driver Database," on page 58

3.5.1 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 3.5.3, "Uploading Drivers to the PlateSpin Migrate Device Driver Database," on page 58.

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

---

3.5.2 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. See Table 3-2 on page 50.

1 On a 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/linuxphysicaltarget.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 .

3.5.3 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 58
- "Device Driver Upload Procedure (Linux)" on page 59

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:
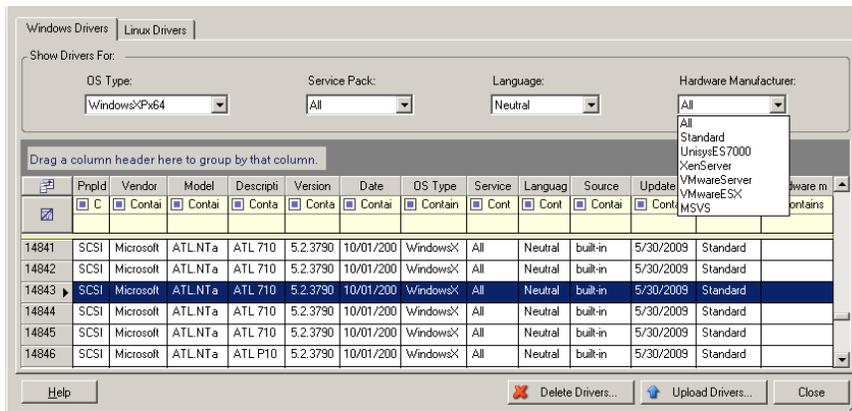
3 Click *Upload Drivers*, browse to the folder that contains the required driver files, and select applicable OS type, language, and hardware manufacturer options.
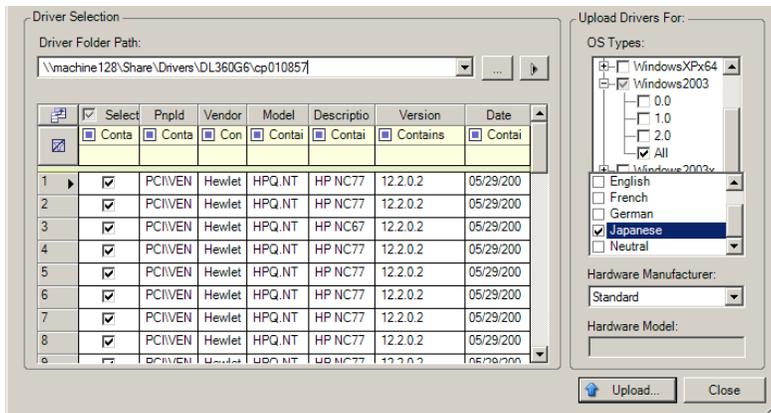
For most X2P migrations, select *Standard* as the *Hardware Manufacturer* option, unless your drivers are designed specifically for any of the target environments listed.

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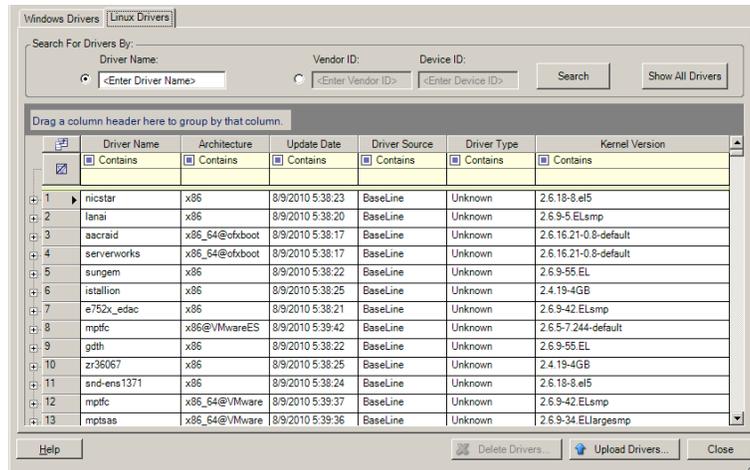
3 Click *Upload Drivers*, browse to the folder that contains the required driver package (`*.pkg`), and click *Upload All Drivers*.

The system uploads the selected drivers to the driver database.

3.6 Managing Custom Actions

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

- Section 3.6.1, "Managing Post-migration Actions (Windows and Linux)," on page 60
- Section 3.6.2, "Freeze and Thaw Scripting Capabilities (Linux Block-Level Migrations)," on page 61

3.6.1 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 Migrate Server's library. The maximum size of the directory must not exceed 64 MB. For information about raising this limit, see Section 2.2.5, "Increasing the Size Limit on Post-migration Actions Uploaded to the PlateSpin Migrate Server," on page 30.

Use the following procedure for adding a post-migration action to the PlateSpin Migrate 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 Migrate Server can access.

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

See KB Article 7970214 (http://www.novell.com/support/viewContent.do?externalId=7970214).

2 In the PlateSpin Migrate Client, click *Tools > Manage Actions*.

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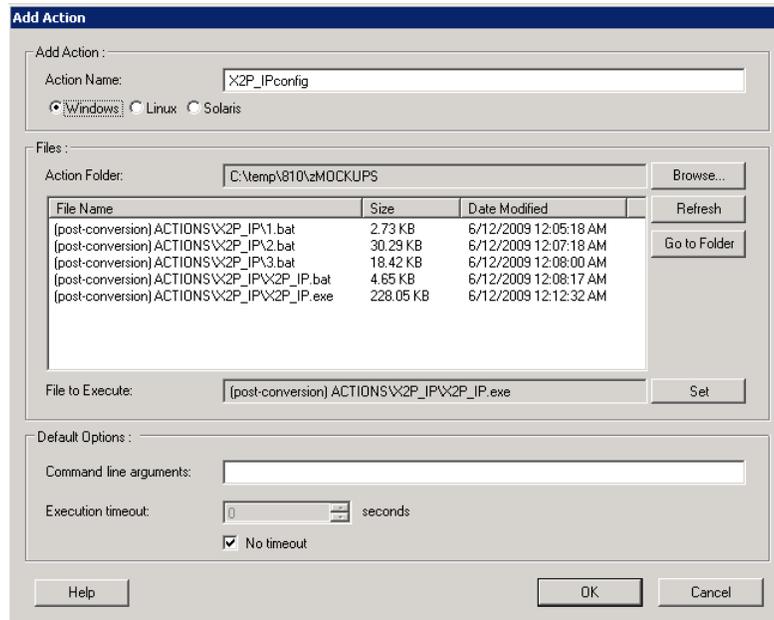
4 In the Add Action window, type a name for your custom action, select the target operating system type, then browse to and select the directory that contains the required action with its dependencies.

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 138.

3.6.2 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 117.

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/`

3.7 Setting Up, Executing, and Managing Jobs

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

3.7.1 Selecting a Migration Job Setup Method

You can set up a migration job and specify the job configuration parameters in one of the following two modes:

- **Advanced mode:** Uses the Conversion Job window, which provides access to all job configuration parameters.
- **Wizard mode:** Provides a convenient subset of the Advanced method and interactively captures only the essential parameters of a migration job.

The following sections provide more information:

Differences Between the Advanced and Wizard Job Configuration Modes

The Advanced and Wizard job configuration modes differ in the extent of available configuration options.

In Wizard mode, PlateSpin Migrate provides a convenient subset of the Advanced method, interactively captures the essential parameters of a migration job, itemizes settings specified in previous wizard pages, and provides a summary at the end.

Figure 3-6   *Job Configuration Window in Wizard Mode*
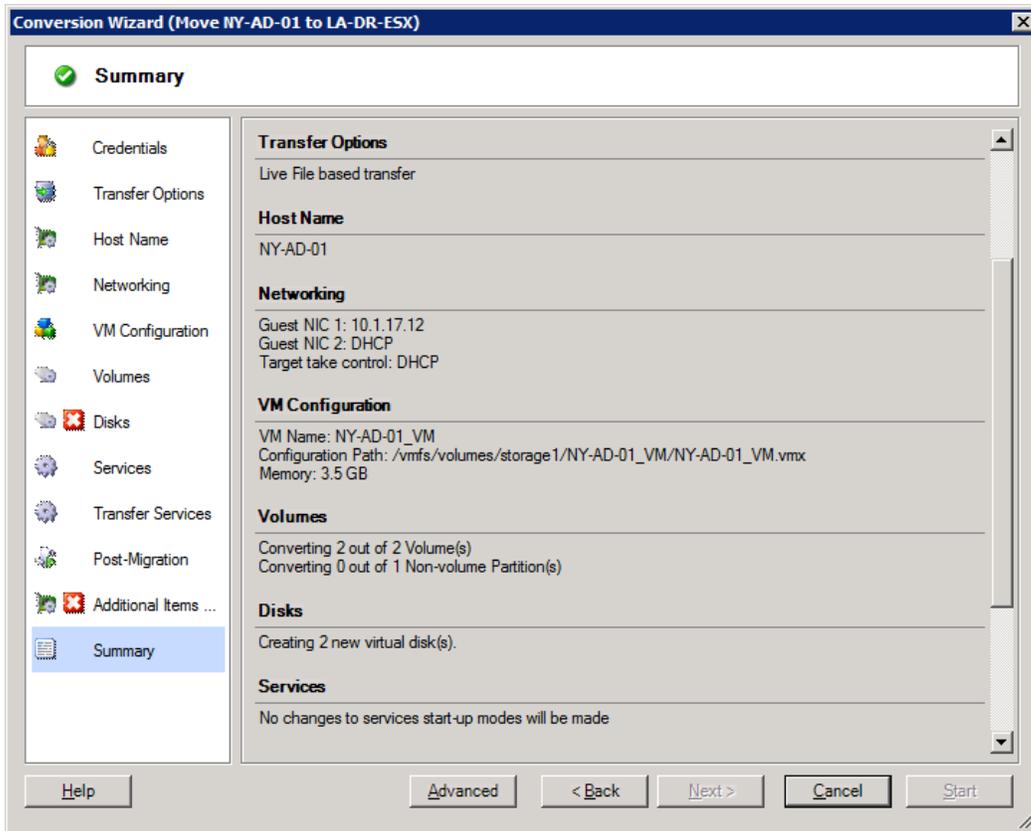
In Advanced mode, PlateSpin Migrate provides access to enhanced job configuration options and advanced settings for some aspects of a migration job. In Advanced mode, you access settings for an aspect of a migration job by clicking links that are grouped in categories. To hide or show category details, click the category bar.

Figure 3-7 *Job Configuration Window in Advanced Mode*
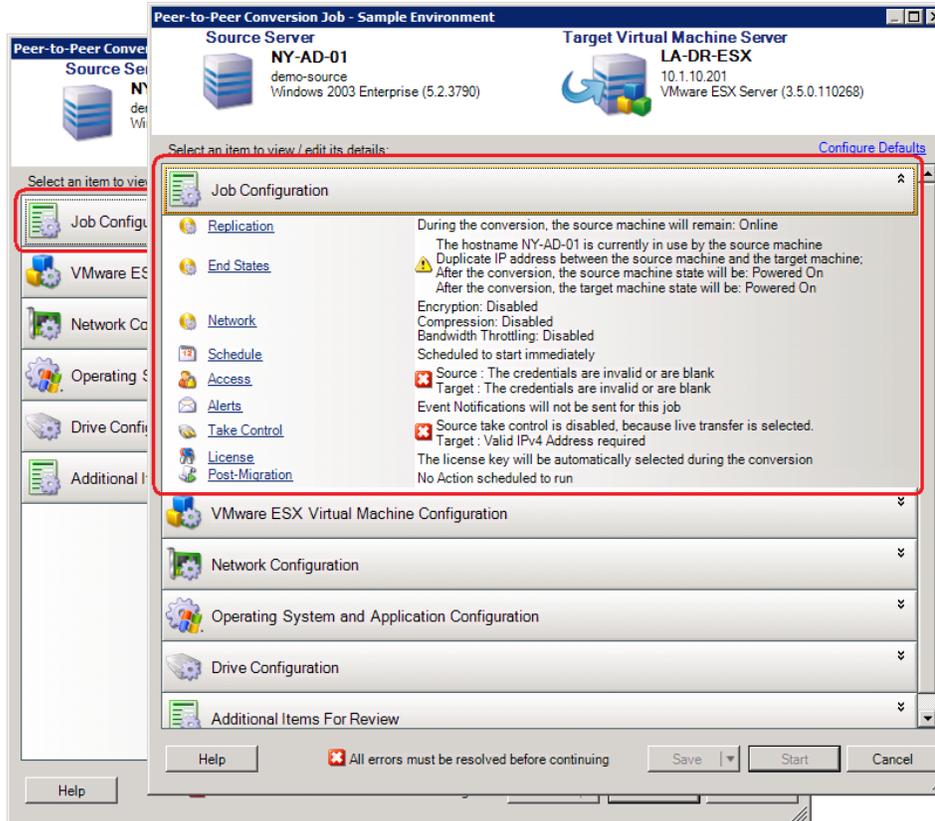
You can access the following enhanced job configuration options in Advanced mode, but cannot access them in Wizard mode:

- Specifying a non-default license for a migration job.

  See "Selecting a License Key for a Migration Job" on page 95.

- Enabling automatic job progress and status notifications.

  See "Configuring Automatic E-Mail Alerts of Job Status and Progress" on page 96.

- Job scheduling options.

  See "Job Scheduling" on page 103.

- Advanced options for target virtual machine configuration.

  See "Virtualization Platform-Specific and Advanced VM Configuration Options" on page 112.

- Saving the job for on-demand execution at a later time.

  See "Saving Migration Jobs" on page 66.

Starting a Migration Job

To start setting up a migration job and select either the Advanced or the Wizard method:

1 Open the Action window by using one of the following methods:

- 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 Action window opens:
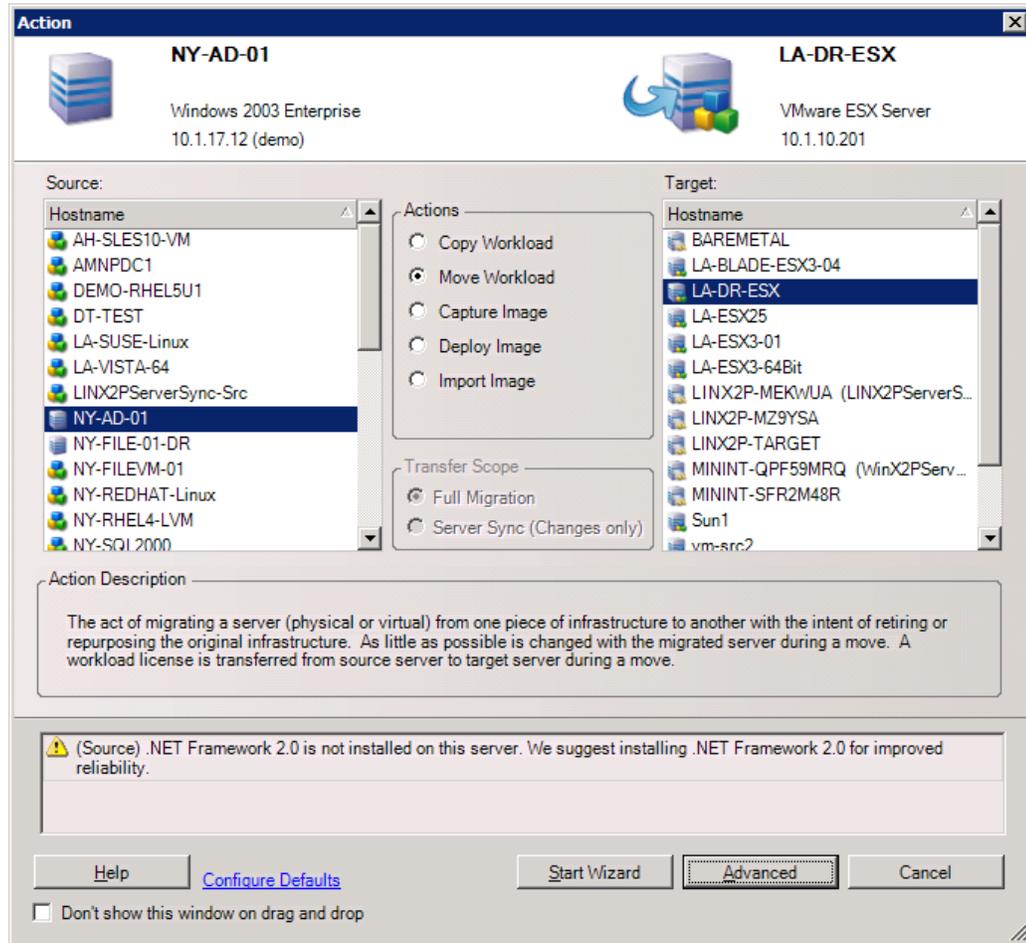
The Source and Target panes display workloads and targets applicable to the selected type of a migration job.

The *Full Transfer* and *Server Sync* options are enabled under 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 72.

2 Check validation messages at the bottom of the window.

3 To start configuring your migration job in either Wizard or Advanced mode, click *Start Wizard* or click *Advanced*.

4 (Optional) For convenience, to avoid displaying the *Action* window on drag-and-drop, select *Don't show on drag and drop* before proceeding. Subsequent drag-and-drops bypass the Action window and directly open a Conversion Job window in Advanced mode.

To restore the job migration startup behavior, restore application defaults. See "General Options" on page 32.

3.7.2 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. See Step 1.

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 3-8 *Migration Validation Window*
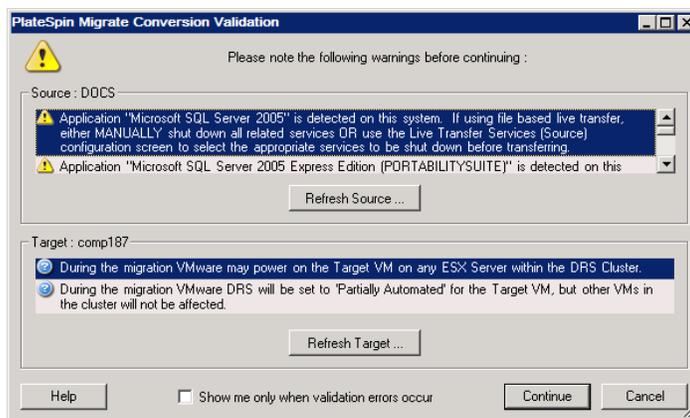
To force this window to open only on errors, select the *Show me only when validation errors occur* option.

3.7.3 Saving Migration Jobs

When you start a migration job in Advanced mode, you have additional options for saving the job and executing it at a later time.

After all options have been configured for the job, click the arrow at the right side of the *Save* button, then select a command from the drop-down list:
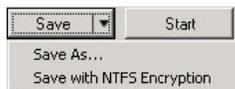
3.7.4 Starting a Saved Migration Job in the Jobs View

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 migration job.

3.7.5 Changing the Start Time of a Saved Migration Job

1 In the PlateSpin Migrate Client's Jobs view, locate the required saved job.

2 Right-click the job and select *Change Start Time*.

The Change Job Start Time dialog box opens.
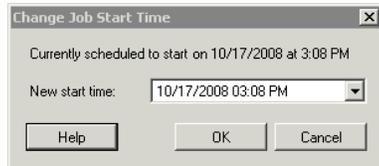
3 Specify the required new start date and time, then click *OK*.

PlateSpin Migrate reschedules the job and executes it at the specified time.

3.7.6 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 in Advanced mode and displays the job's configuration parameters in read-only mode.

3.7.7 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*.

3.7.8 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 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 Replication Options" on page 97.

3.7.9 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*.
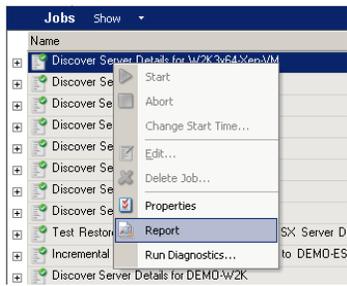
A Web browser window displays the requested report.

3.7.10 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 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 4.1, "Migrating a Workload to a VM Host (P2V, V2V)," on page 69
- Section 4.2, "Converting a Workload to a Physical Machine (P2P, V2P)," on page 70
- Section 4.3, "Synchronizing Workloads with Server Sync," on page 72
- Section 4.4, "Semi-Automated Workload Virtualization Using the X2P Workflow," on page 73
- Section 4.5, "Advanced Workload Migration Notes," on page 78

4.1 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 46.

2 Start a migration job.

See "Selecting a Migration Job Setup Method" on page 62.

3 Configure the required parameters of the job.

See "Essentials of Workload Migration Jobs" on page 95.

4 Run the job.

Figure 4-1   *Peer-to-Peer Migration Job Window (P2V in Advanced Mode)*
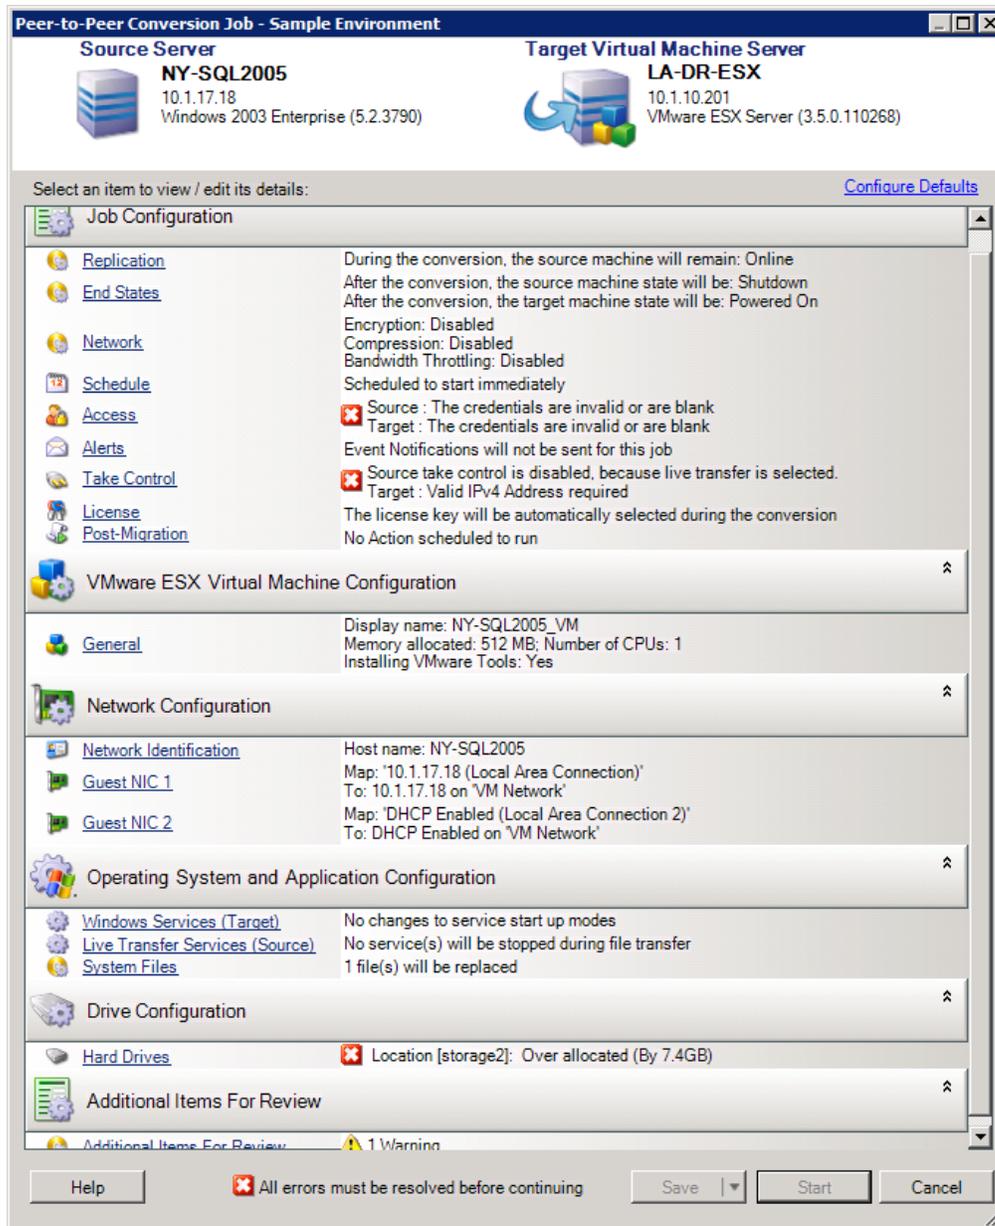
4.2 Converting a Workload to a Physical Machine (P2P, V2P)

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 46.

2 Discover and register your target physical machine with PlateSpin Migrate.

See "Discovering and Registering Target Physical Machines" on page 49.

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 55.

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 56.

5 Start a migration job.

See "Selecting a Migration Job Setup Method" on page 62.

6 Configure the required parameters of the job.

See "Essentials of Workload Migration Jobs" on page 95.

7 Run the job.

Figure 4-2 *Peer-to-Peer Conversion Job Window (V2P in Advanced Mode)*
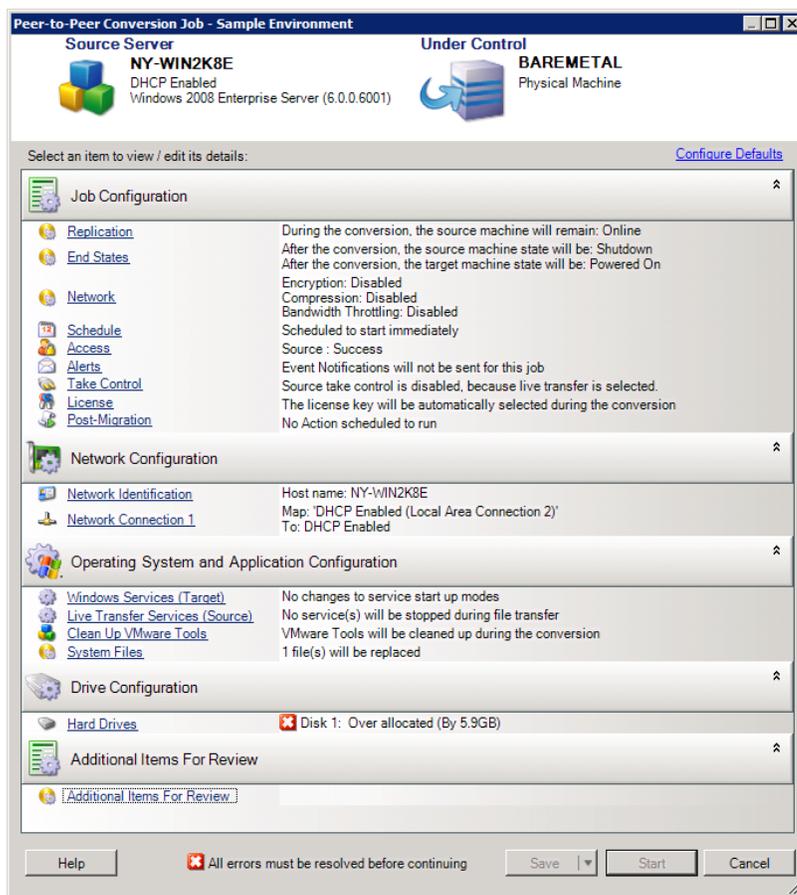
4.2.1 Best Practices (X2P)

- ◆ 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 116.

◆ 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 104.

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

See "Handling a Workload's Storage Media and Volumes" on page 122.

4.3 Synchronizing Workloads with Server Sync

The Server Sync feature enables you to control the scope of data that is transferred from your source to your target. 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 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 workload virtualization over the network.

◆ Section 4.3.1, "Server Sync to a Virtual Target," on page 72
◆ Section 4.3.2, "Server Sync to a Physical Target," on page 73

4.3.1 Server Sync to a Virtual Target

1 Discover your source workload. See "Discovering Source Workloads and Migration Targets" on page 46.

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

   2a Do an initial migration of your workload to a virtual machine.

      See "Migrating a Workload to a VM Host (P2V, V2V)" on page 69.

   2b Using your virtualization platform's native interface, manually install a virtual machine with the same operating system profile as that of your source.

   2c (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 85 and "Deploying a PlateSpin Image" on page 89.

3 Discover and prepare the target VM for synchronization.

   See "Discovering, Registering and Preparing a VM for a Server Sync Job" on page 53.

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).

   If an operating system is detected on the target, and if the detected operating system matches that of the source, the system prompts you to select the scope of data to transfer (complete source volume data or only files that are different between the source and the target).

5 Select the *Server Sync* option, then click Start.

   A Server Sync job starts in Advanced mode.

6 Configure the parameters of the job as dictated by the purpose of the operation. Make sure that you map the required volumes on the source to those on the target.

7 Click *Start*.

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

4.3.2 Server Sync to a Physical Target

1 Discover your source workload.

See "Discovering Source Workloads and Migration Targets" on page 46.

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

See "Discovering and Registering Target Physical Machines" on page 49.

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).

If an operating system is detected on the target, and if the detected operating system matches that of the source, the system prompts you to select the scope of data to transfer (complete source volume data or only files that are different between the source and the target).

4 Select the *Server Sync* option, then click Start.

A Server Sync job starts in Advanced mode.

5 Configure the parameters of the job as dictated by the purpose of the operation. Make sure that you map the required volumes on the source to those on the target.

6 Click *Start*.

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

4.4 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:

◆ Microsoft Hyper-V

◆ Xen Hypervisor on SUSE Linux Enterprise Server (SLES) 10 (applicable to workloads migrating to fully virtualized targets; for information on migrations to paravirtualized targets, see "Linux Migration to a Paravirtualized VM on Xen-on-SLES" on page 79).

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 Boot the newly created virtual machine by using the appropriate boot ISO image.

This discovers and registers the target virtual machine as a PlateSpin Migrate physical machine target. This is similar to the process described in "Discovering and Registering Target Physical Machines" on page 49. 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.

4.4.1 Migrating a Workload to Microsoft Hyper-V

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

Prerequisites:

- Your source workload is supported by PlateSpin Migrate and Hyper-V.

  See "Supported Target Virtualization Platforms" on page 13.

Additional information:

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

This section includes the following topics:

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

1 Download the appropriate boot ISO image for your workload.

   See Table 3-2 on page 50.

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 20832 (http://www.novell.com/support/viewContent.do?externalId=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.

4 Start the VM and use the Virtual Machine Connection Console to 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 Migrate Server (Hyper-V)

Complete these steps to register your target with PlateSpin Migrate 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 20832 (http://www.novell.com/support/viewContent.do?externalId=7920832).

1 At the command line in the Hyper-V Virtual Machine Connection Console, provide the required information at each individual prompt:

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

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

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

- **Credentials (User Name/Password):** Enter the name of an admin-level user on the PlateSpin Migrate 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 Migrate Server and registers the virtual machine as a physical target for a migration job.

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 70.

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*.

4.4.2 Migrating a Workload to Xen Hypervisor on SLES

You can use the Xen Hypervisor on SLES as the target virtualization platform in a semi-automated workload virtualization.

Prerequisites:

- Your target is a fully virtualized VM. (For information on migrations to paravirtualized targets, see "Linux Migration to a Paravirtualized VM on Xen-on-SLES" on page 79).
- Your source workload is supported by PlateSpin Migrate and the Xen hypervisor.

  See "Supported Target Virtualization Platforms" on page 13.

Additional information:

- Virtualization with Xen (http://www.novell.com/documentation/sles10/xen_admin/index.html?page=/documentation/sles10/xen_admin/data/bookinfo.html)

This section includes the following topics:

- "Downloading and Saving the Boot ISO Image (Xen on SLES)" on page 76
- "Creating and Configuring the Target Virtual Machine (Xen on SLES)" on page 77
- "Registering the Virtual Machine with PlateSpin Migrate Server (Xen on SLES)" on page 77
- "Migrating Your Source Workload to the Target Virtual Machine (Xen on SLES)" on page 78
- "Post-Migration Steps (Xen on SLES)" on page 78

Downloading and Saving the Boot ISO Image (Xen on SLES)

1 Download the appropriate boot ISO image for your workload.

See Table 3-2 on page 50.

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 20832 (http://www.novell.com/support/viewContent.do?externalId=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 10, use the Virtual Machine Manager Wizard or the Create Virtual Machines program shortcut to create a new virtual machine and install an operating system that matches the profile of your source.

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 "Configuring a Virtual Machine by Modifying its xend Settings" (http://www.novell.com/documentation/sles10/xen_admin/index.html?page=/documentation/sles10/xen_admin/data/sec_xen_virtualization_xend_vmachine.html) in the SLES 10 documentation.

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 Migrate Server (Xen on SLES)

Complete these steps to register your target with PlateSpin Migrate 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 20832 (http://www.novell.com/support/viewContent.do?externalId=7920832).

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

- ◆ **PlateSpin Migrate Server:** Use the following format:

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

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

- ◆ **Credentials (User Name/Password):** Enter the name of an admin-level user on the PlateSpin Migrate 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 they 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 Migrate 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 70.

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:

Virtualization with Xen (http://www.novell.com/documentation/sles10/xen_admin/index.html?page=/documentation/sles10/xen_admin/data/bookinfo.html).

4.5 Advanced Workload Migration Notes

4.5.1 Migrating Windows Clusters

You can migrate a Microsoft Windows cluster's business services (to ESX 3.0.2 and later). PlateSpin Migrate supports single-active-node versions of the following clustering technologies:

- ◆ Windows 2003 Server-based Windows Cluster Server (*Single-Quorum Device Cluster* model)
- ◆ Windows 2008 Server-based Microsoft Failover Cluster (*Node and Disk Majority* and *No Majority: Disk Only* models)

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 scope of support for cluster migrations in the current release is subject to the following conditions:

- ◆ All shared disks belong to the active node.

- The source workload of the migration must be the active node—the node that currently owns the quorum resource of the cluster. To discover a cluster, specify the IP address of one of the cluster's resource groups.

- A cluster's quorum resource must be collocated with the cluster's resource group (service) being protected.

- For a successful X2P or Server Sync operation, target disks must have discrete SCSI controllers to separate the cluster's shared disks from disks that host system volumes of individual nodes.

- To be able to function, 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 or simultaneously migrating it as well.

The workflow of migrating a Windows cluster is similar to that of migrating a standalone server:

1 Discover the active node, specifying the cluster IP address and cluster admin credentials.

2 In the Servers view, use drag-and-drop to start a migration job, then configure the job's parameters.

NOTE:

Selecting *Shut down* for the source's post-migration end state results in all source nodes of the cluster being shut down.

If a cluster failover occurs prior to the completion of file transfer, the migration job is aborted. If this happens, refresh the source and retry the migration job.

4.5.2 Linux Migration to a Paravirtualized VM on Xen-on-SLES

You can do a migration to a paravirtualized VM on Xen-on-SLES (version 10 only). This is done indirectly, through a two-stage process. The paravirtualized VM needs to be transformed into a fully virtualized VM first and later transformed back. A utility (xmps), included in your PlateSpin boot ISO image, is used to transform the VM.

The procedure varies slightly, depending on whether the target is a new or an existing paravirtualized VM.

- "Linux Migration to a New Paravirtualized VM" on page 79
- "Linux Migration to an Existing Paravirtualized VM" on page 81

Linux Migration to a New Paravirtualized VM

1 Copy the PlateSpin Linux boot ISO to the target Xen/SLES server.

See Table 3-2, "Boot ISO Images for Target Physical Machines," on page 50.

2 Start the Virtual Machine manager and create a fully virtualized VM:

2a Select the *I need to install an operating system* option.

2b Choose a suitable size for the disk image (the disk size should be equal to or bigger than that of the source machine).

2c Select the boot ISO as the installation source.

The VM boots into the PlateSpin OS environment, used in X2P settings.

3 Complete the migration.

See "Migrating a Workload to Xen Hypervisor on SLES" on page 76.

Upon completion, the VM should be fully functional as a fully virtualized machine.

4 Reboot the VM, making sure that it still boots into the PlateSpin OS environment.

```
Available boot options (type the name to boot into):

  ps         - PlateSpin Linux for Taking Control (press ENTER to boot into)
  ps64       - PlateSpin Linux(x86_64) for Taking Control
  ps64_512m  - PlateSpin Linux(x86_64) for Taking Control a Virtual Machine
               which has more than 512M memory
  next       - Boot from Next Boot Device Set in BIOS (timeout)
  debug      - PlateSpin Linux for Trouble Shooting
  switch     - PlateSpin Linux for switching kernel to Xen PV

  When no key is pressed for 20 seconds, it will boot from the next boot device.


boot: switch_
```

5 At the `boot:` prompt, type `switch` and press Enter.

This reconfigures the operating system to be bootable as a paravirtualized machine. Upon completion, the output should look similar to the one shown below:

```
about to find other volumes in native off-line OS
kjournald starting.  Commit interval 5 seconds
EXT3-fs: mounted filesystem with ordered data mode.
found volume /boot in off-line OS
found other 1 volume(s)
mount all the system volumes
kjournald starting.  Commit interval 5 seconds
EXT3 FS on hda1, internal journal
EXT3-fs: mounted filesystem with ordered data mode.
volume /boot has been mounted.
all the system volumes are mounted
Switching to Xen kernel for Para-virt machine....
unmount all the system volumes for clean up.
volume /boot has been unmounted
volume / has been unmounted

###########################################################################
Please apply the following data as bootloader_args for
switching Xen fully-virt machine to Para-virt machine:

'--entry=xvda1:/vmlinuz-2.6.16.60-0.54.5-xen,/initrd-2.6.16.60-0.54.5-xen'

###########################################################################
[DB]$ _
```

Note the bootloader arguments in the final segment of the output:

```
Please apply the following data as bootloader_args for
switching Xen fully-virt machine to Para-virt machine:

'-entry=xvda1:/vmlinuz-2.6.16.60-0.54.5-xen, /initrd-2.6.16.60-0.54.5-xen'
```

These are used by the `xmps` utility to set up the location of the kernel and the `initrd` image, from which the paravirtualized machine boots from.

6 Power off the virtual machine:

```
[DB]$ poweroff
```

7 Login to the XEN/SLES server as `root` and mount the PlateSpin Linux boot ISO (the command example assumes that the ISO has been copied under the `/root` directory):

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

8 Run the `xmps` utility to create a paravirtualized VM based on the configuration of the fully virtualized VM:

```
# /mnt/ps/tools/xmps --pv --vm_name=SLES10-FV --new_vm_name=SLES10-PV --
bootloader_args="--entry=xvda1:/vmlinuz-2.6.16.60-0.54.5-xen, /initrd-
2.6.16.60-0.54.5-xen"
```

The utility takes as input:

- the name of the fully virtualized VM on which the configuration of the paravirtualized machine will be based (`SLES10-FV`)
- the name of the virtual machine to create (`SLES10-PV`)
- the paravirtualized machine's bootloader arguments `"--bootloader_args"` (shown at Step 5)

NOTE: If a VM with the same name as the one passed as `new_vm_name` already exists, the `xmps` utility fails.

The newly created paravirtualized VM (`SLES10-PV`) should now be available in the Virtual Machine Manager, ready to be turned on. The corresponding fully virtualized machine is retired and will fail to boot. This VM can be deleted safely (only the VM configuration will be removed).

9 Unmount the PlateSpin Linux boot ISO:

```
# umount /mnt/ps
```

Linux Migration to an Existing Paravirtualized VM

1 Copy the PlateSpin Linux boot ISO to the target Xen/SLES server.

See Table 3-2, "Boot ISO Images for Target Physical Machines," on page 50.

2 Login to the XEN/SLES server as `root` and mount the PlateSpin Linux boot ISO:

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

3 Run the `xmps` utility to create a fully virtualized VM based on the configuration of the paravirtualized VM (the intended failback target):

```
# /mnt/ps/tools/xmps --fv --vm_name=SLES10-PV  --new_vm_name=SLES10-FV --
bootiso=/root/linuxphysicaltarget.iso
```

The utility takes as input:

- the name of the existing paravirtualized machine (`SLES10-PV`), which is the intended failback target
- the name of the temporary fully virtualized machine (`SLES10-FV`) to be created for the two-stage failback operation
- the full path of the boot ISO (assuming that the ISO file is located uner `/root: /root/linuxphysicaltarget.iso`)

NOTE: If a VM with the same name as the one passed as `new_vm_name` already exists, the `xmps` utility fails.

The newly created fully virtualized machine (`SLES10-FV`) should now be available in the Virtual Machine Manager.

4 Turn on the newly created fully virtualized machine (`SLES10-FV`).

The VM boots into the PlateSpin OS environment, used in X2P settings.

5 Complete the migration.

See "Migrating a Workload to Xen Hypervisor on SLES" on page 76.

6 Reboot the VM, run `switch`, and reconfigure the workload as described in "Linux Migration to a New Paravirtualized VM" on page 79 (from Step 4 to Step 9 only).

5 Windows Workload Portability with a PlateSpin Image

This section provides information about using the PlateSpin Image volume archiving feature (Windows only).

5.1 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.

5.2 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 Migrate Server host. For information about storing PlateSpin PlateSpin Images on a NAS (Network Attached Storage) device or a remote share, see KB Article 7921021 (http://www.novell.com/support/viewContent.do?externalId=7921021).

NOTE: Although collocation of the PlateSpin Migrate 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 5-1   *PlateSpin Image Server Host Requirements*

Requirement	Details
Operating System	Any of the following:
	◆ Microsoft Windows Server 2008 R2
	◆ Microsoft Windows Server 2008
	◆ Microsoft Windows Vista
	◆ Microsoft Windows Server 2003
	◆ Microsoft Windows 2000
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 2.0 or later
	◆ (Windows Server 2008 and Vista systems only) Remote Registry service enabled

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*.
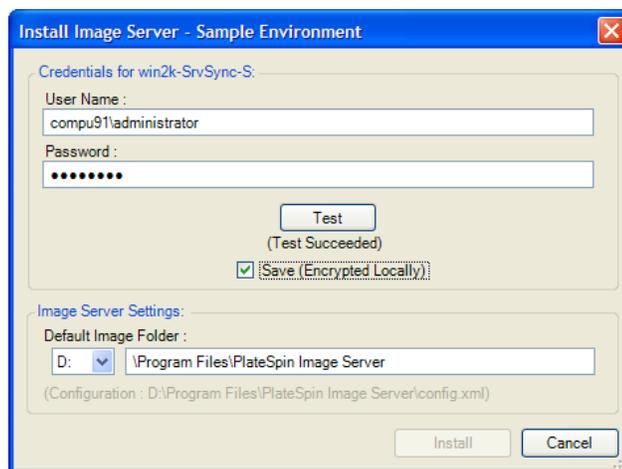
3 Provide administrator credentials for the selected host and specify the desired directory for image files.

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:

5.3 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*. For information about the *Use existing volume data* option, see "Creating Images by Using Raw or Existing Volume Data" on page 87.
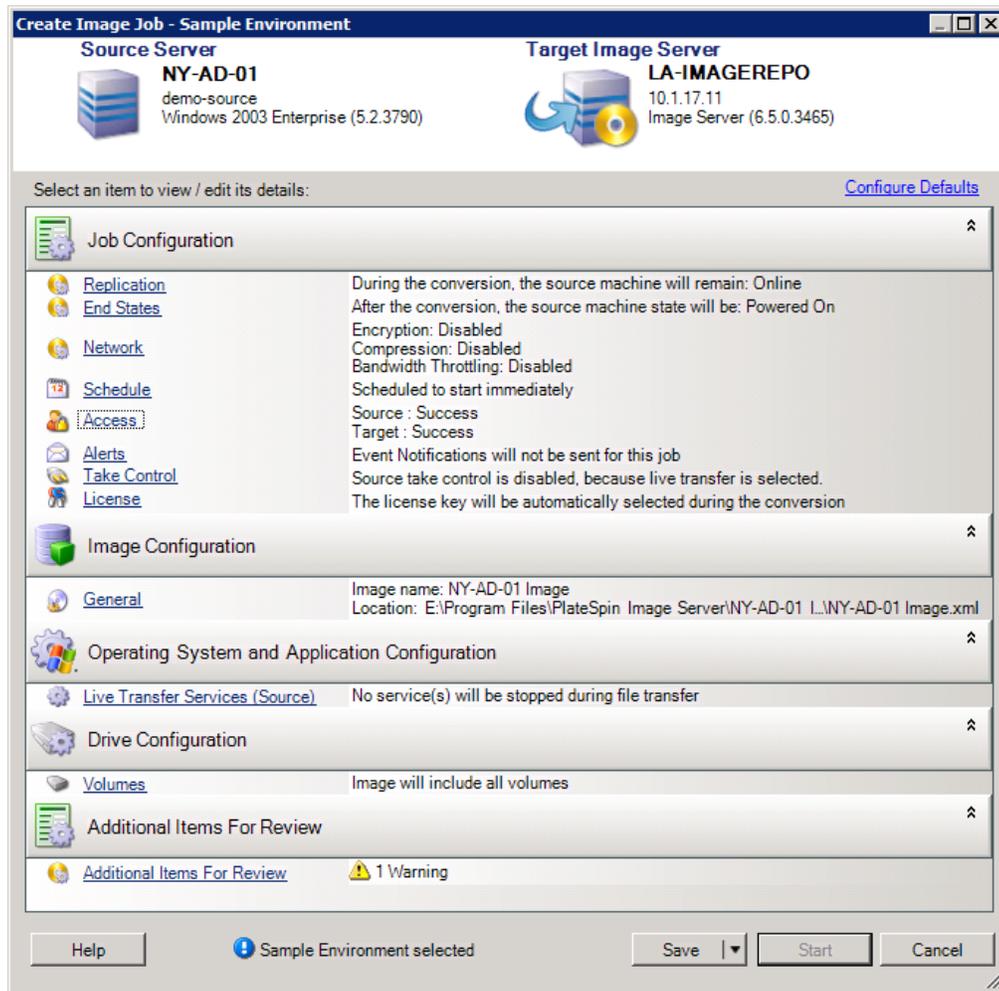
4 Specify the required settings for the migration job by clicking the links in each category:

- ◆ **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*).
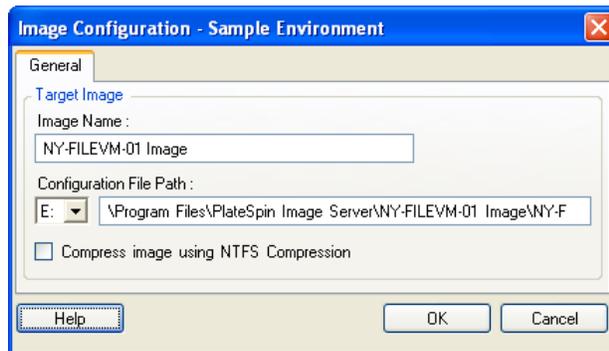
- **Operating System and Application Configuration:** If you selected the Live Transfer method, specify how you want PlateSpin Migrate to handle operating system and application services on your source (*Live Transfer Services*).
- **Drive Configuration:** Select the volumes that you want PlateSpin Migrate to include in the image and specify the path for the package file (under Drive Configuration, click *Volumes*).
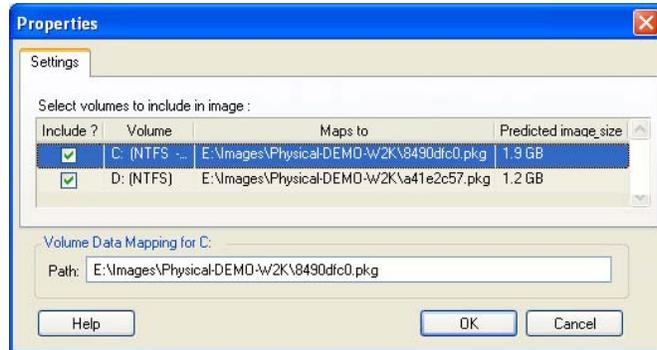
5.4 Creating Images by Using Raw or Existing Volume Data

PlateSpin Migrate can use existing system images to create a PlateSpin Image.

5.4.1 Importing Existing or Raw Volume Data

1 Discover the source machine with the raw volume data you want to capture.

2 Drag and drop the source server onto the target PlateSpin Image Server.

One of the following happens:

- The *Action* dialog box opens. Ensure that the *Capture Image* option is selected in the *Actions* area. Click *Advanced*. Alternately, to use the Wizard, click *Start Wizard*.
- 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. Select *Use existing volume data* and click *OK*.
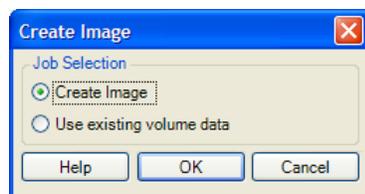
3 In the *Create Image Job* window, click *Volumes* under the *Drive Configuration* category. A Properties dialog box opens.
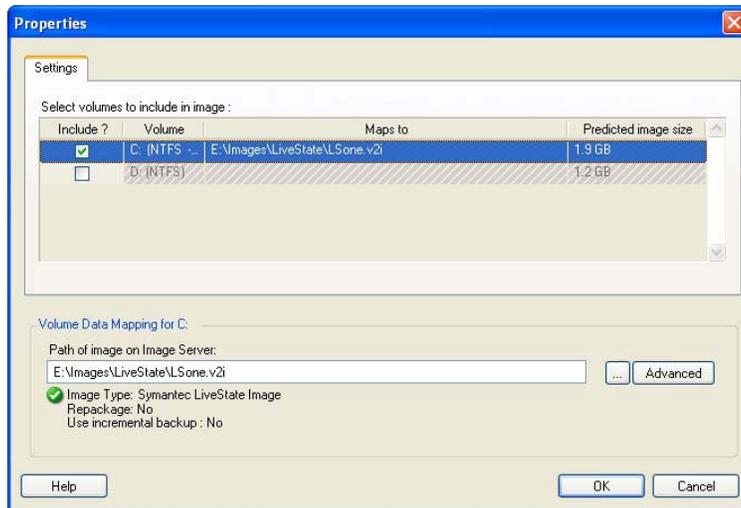
4 Select each volume. In the text box at the bottom of the dialog box, type the full path to the directory containing the raw volume data, or to the existing image file for each volume.

5 For each image, further properties can be configured by clicking the *Advanced* button:

Repackage: Select this option to create a PlateSpin Image from the provided image file.

---
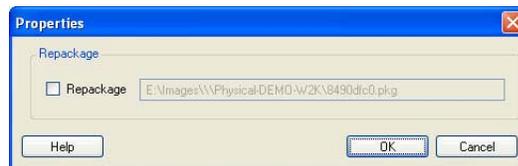
6 After selecting suitable images or raw volume data for each volume to be included, confirm the settings for the *Job Configuration* and *Image Configuration* sections.

7 Click *Start* to begin the image creation process.

5.4.2 Additional Information About Creating PlateSpin Images

Table 5-2   *KB Articles Related to PlateSpin Images*

ID	Description
7920392 (http://www.novell.com/support/viewContent.do?externalId=7920392)	HOWTO: How to create a PlateSpin Image from NT backup restore.
2790417 (http://www.novell.com/support/viewContent.do?externalId=7920417)	HOWTO: How to create a PlateSpin Image from Computer Associates' BrightStor restore.
7920400 (http://www.novell.com/support/viewContent.do?externalId=7920400)	HOWTO: How to create a PlateSpin Image from Symantec's Backup Exec system restore.
7920461 (http://www.novell.com/support/viewContent.do?externalId=7920461)	HOWTO: How to create a PlateSpin Image from Ghost Image extraction.
7920723 (http://www.novell.com/support/viewContent.do?externalId=7920723)	FAQ: How to create a PlateSpin Image from IBM Tivoli Storage Manager (TSM).

5.5 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.

5.6 Managing PlateSpin Images

5.6.1 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 20189 (http://www.novell.com/support/viewContent.do?externalId=7920189).

5.6.2 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.

5.6.3 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 91
- ◆ "Sorting and Searching Items in the Image Browser Interface" on page 92
- ◆ "Extracting Items" on page 92
- ◆ "Browsing and Extracting Image Files at the Command Line" on page 93

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 Migrate 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.
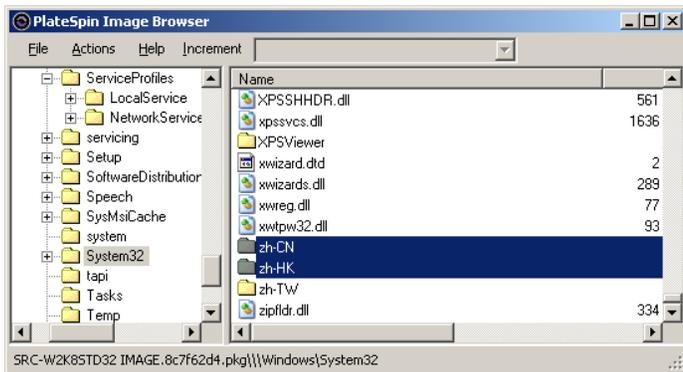
Depending on the size of the image, it might take a few seconds to several minutes for the utility to load the required file.

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).

To search for an item:

1 Load the required image or image increment. See "Starting the Image Browser and Loading Image Files" on page 91.

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 91.

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.

6 Essentials of Workload Migration Jobs

This section provides information about configuring the key aspects of a workload migration job using the PlateSpin Migrate Client.

Where applicable, differences between the Advanced and Wizard modes are noted. See "Setting Up, Executing, and Managing Jobs" on page 62.

## 6.1 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 19.

To view or modify the license key selected for a migration job:

- **In Advanced mode:** In the Migration Job window, under the Job Configuration section, click *License Key*.
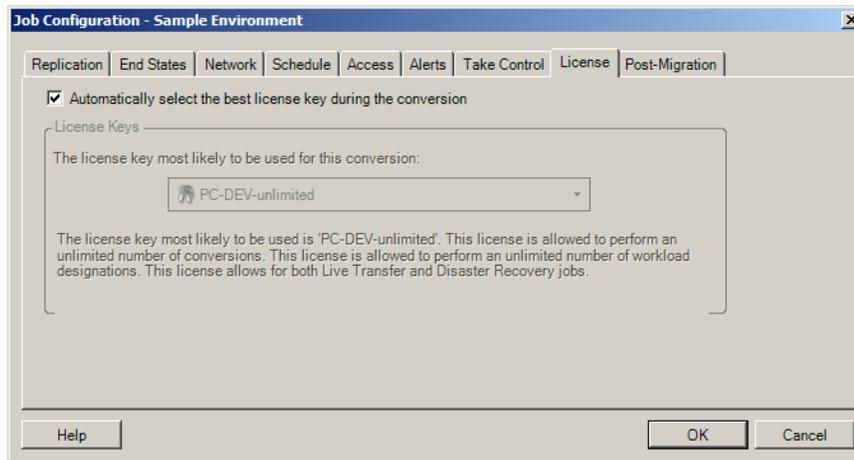
- **In Wizard mode:** Not available.

To manually choose a different key, deselect the *Automatically select the best license key during the migration* check box and choose the required license key from the drop-down list.

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.

The selected license key is displayed on the *License Key* tab and the description is updated accordingly.

6.2 Configuring Automatic E-Mail Alerts of Job Status and Progress

You can set up a migration job to automatically send e-mail notifications about status and progress to a specified address:

- **Job events:** Job status messages `Completed`, `Recoverable Error`, and `Failed`.
- **Job progress:** Detailed job progress messages at configurable intervals.

You can specify SMTP server and e-mail account details either during the migration job or globally. See "Notification Service" on page 33.

To configure automatic e-mail notifications:

- **In Advanced mode:** In the Migration Job window, under the Job Configuration section, click *Alerts*.
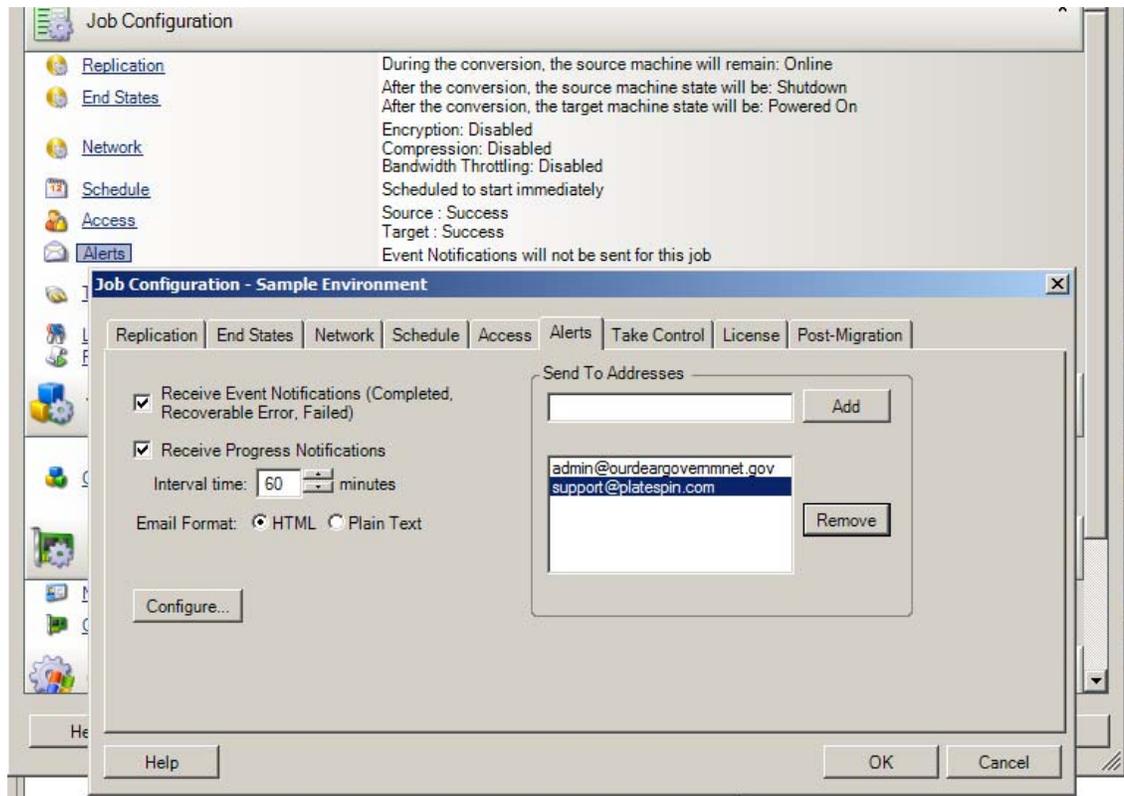
◆ **In Wizard mode:** Not available.

6.3   Specifying Replication Options

Replication 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 14.

◆ 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 72.

To specify the transfer options for a migration job:

◆ **In Advanced mode:** In the Migration Job window, under the Job Configuration section, click *Replication*.
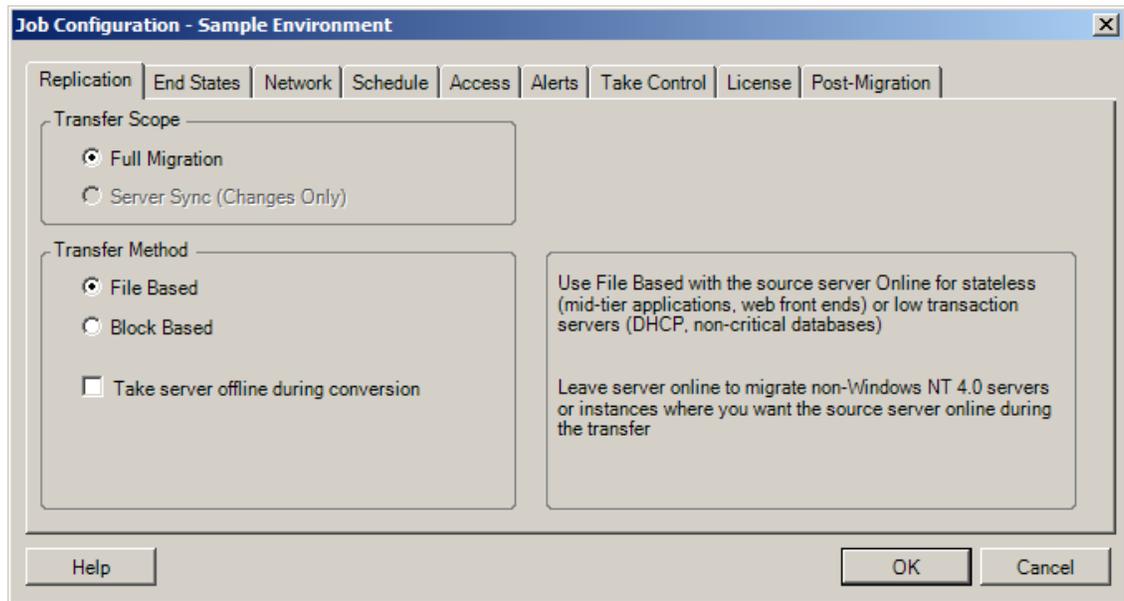
- **In Wizard mode:** In the wizard's navigation pane, click *Transfer Method*.

6.4 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.

You can specify non-default post-migration end states for your source and target:

- **In Advanced mode:** In the Migration Job window, under the Job Configuration section, click *End States*.
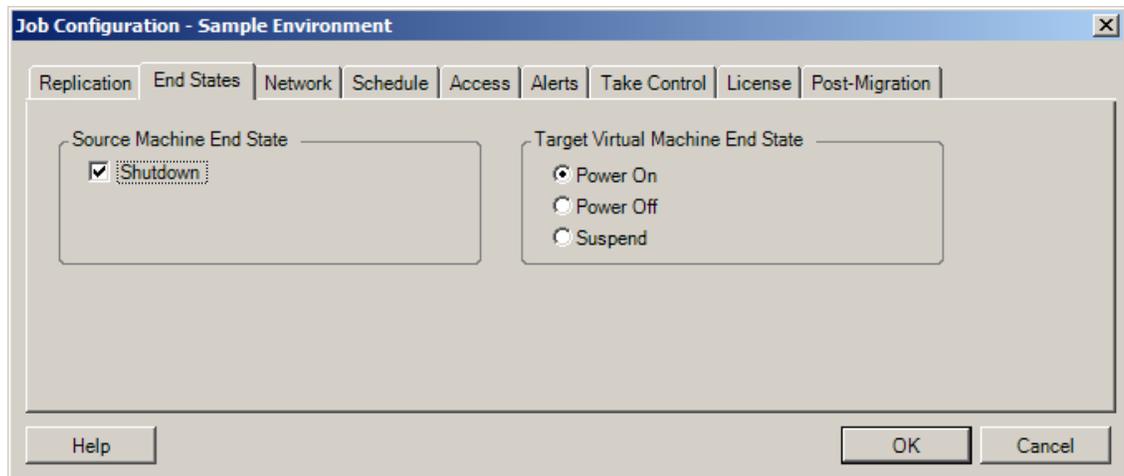
- **In Wizard mode:** Not available; source and target end states are preset to their defaults.

6.5 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 17.

  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 16.

- Whether you want to apply bandwidth throttling for the current migration job.

  See "Bandwidth Throttling" on page 18.

  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 Migrate Server, see "Migrations Across Public and Private Networks Through NAT" on page 31.

To specify network options:

- **In Advanced mode:** In the Migration Job window, under the Job Configuration section, click *Network*.

- **In Wizard mode:** Not available.

6.6   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 6-1   *Source and Target Credentials*

Credentials	Windows	Linux
User Name	Account username with local or domain-level administrative privileges. Use this format: • For domain member machines: *authority\principal* • For workgroup member machines: *hostname\principal*	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:

• **In Advanced mode:** In the Migration Job window, under the Job Configuration section, click *Access*.
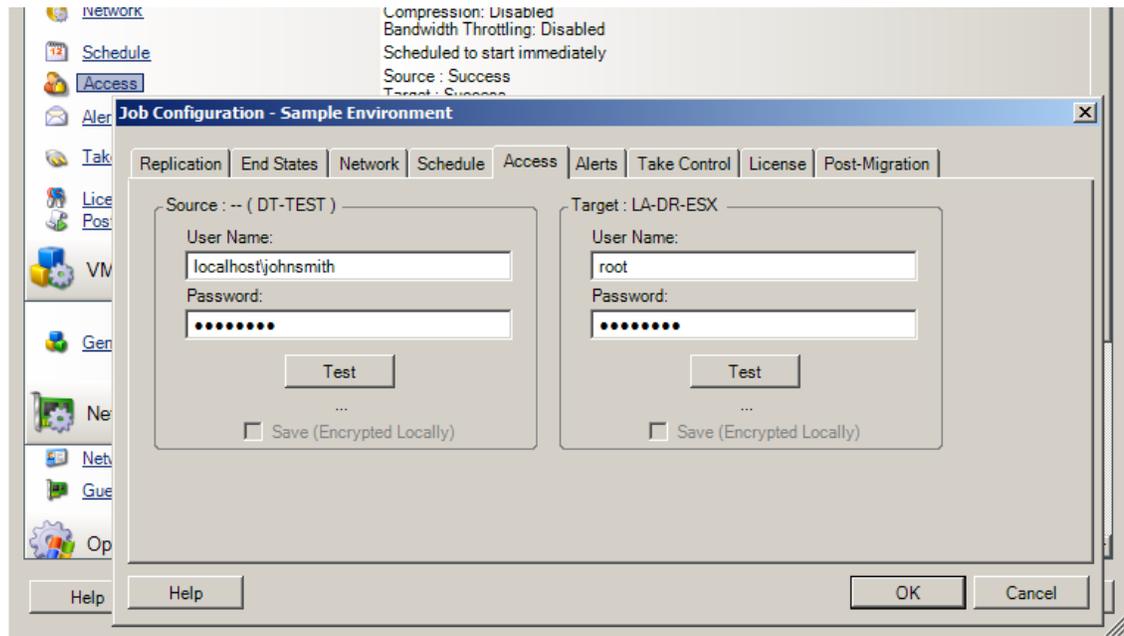
◆ **In Wizard mode:** In the wizard's navigation pane, click *Credentials*.
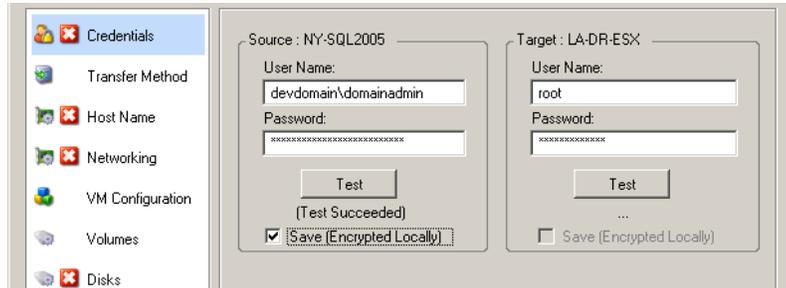
6.7 Managing a Workload's Network Identity

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:

◆ **In Advanced mode:** In the Migration Job window, under the Network Configuration section, click *Network Identification*.
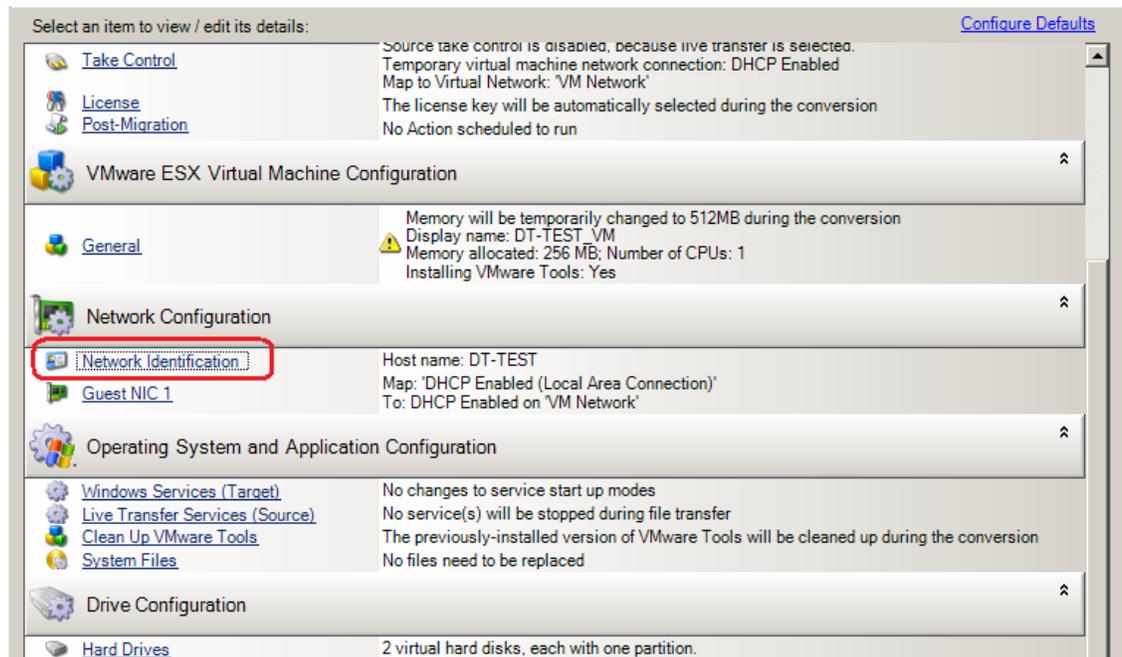
◆ **In Wizard mode:** In the wizard's navigation pane, click *Host Name.*

Configuration options vary depending on whether the target machine is Windows or Linux.

6.7.1 Managing the Identity of Windows Workloads

Use these settings to configure the network identity of your target Windows workload (in either Wizard or Advanced mode).

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 that have permission to add servers to the domain.

6.7.2 Managing the Network Identity of Linux and Solaris Workloads

Use these settings to configure the network identity of your target Linux workload (in either Wizard or Advanced mode) and DNS server addresses as required (Advanced mode).
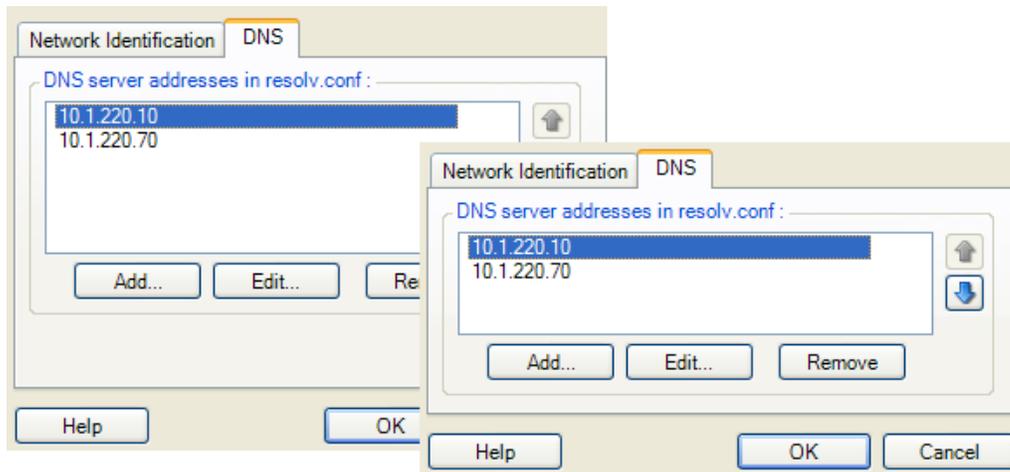
Network Identification tab: Specify the desired hostname for the target server.

DNS tab: Use the *Add*, *Edit*, and *Remove* buttons to manage DNS server entries for the new virtual machine.

6.8 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 62

To access scheduling options of a migration job:

 - **In Advanced mode:** In the Migration Job window, under the Job Configuration section, click *Schedule*.
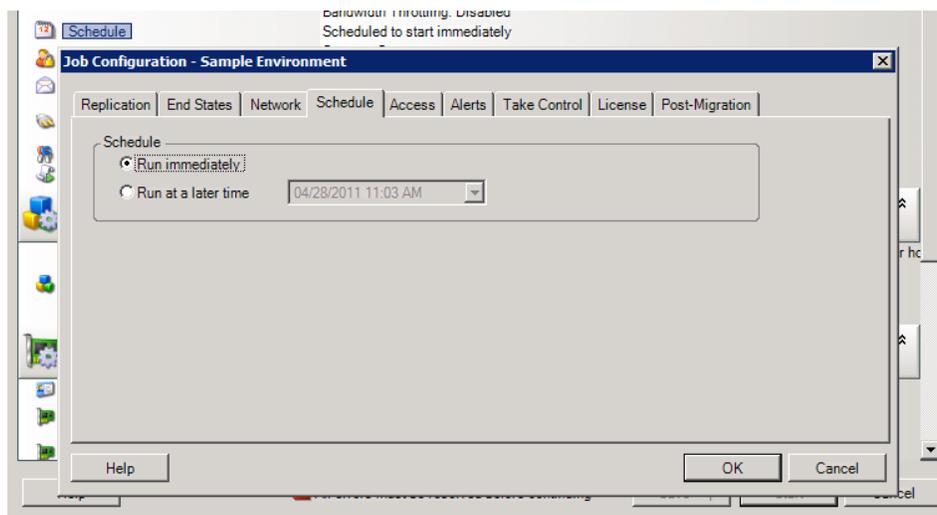
 - **In Wizard mode:** In the wizard's navigation pane, click *Schedule*.

6.9 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 Migrate Server during the migration process, and that the network configuration of a target workload is in line with its end state.

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 15.

6.9.1 Temporary (Take Control) Network Settings

Temporary (Take Control) Network Settings control how source workloads, targets, and the PlateSpin Migrate 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 Migrate Server's communication with the source and target workloads that are booted into a temporary pre-execution environment. See "Offline Transfer of Windows and Linux Workloads" on page 15.

During Solaris workload migrations, temporary network settings for the target are not required, and the corresponding option is disabled. For source workloads, the capability to assign temporary network settings is subject to the following conditions:

- Temporary network settings can be assigned to a Solaris source workload if it is a physical machine with no non-global zones.
- Temporary network settings cannot be assigned to a source workload that is a Solaris zone with its network interface in exclusive mode.

See "Offline Transfer of Solaris Workloads" on page 15.

To configure Temporary (Take Control) network settings:

- **In Advanced mode:** In the Migration Job window, under the Job Configuration section, click *Take Control*. To access network interface mapping and TCP/IP settings, click *Configure* in the source and target areas as applicable.
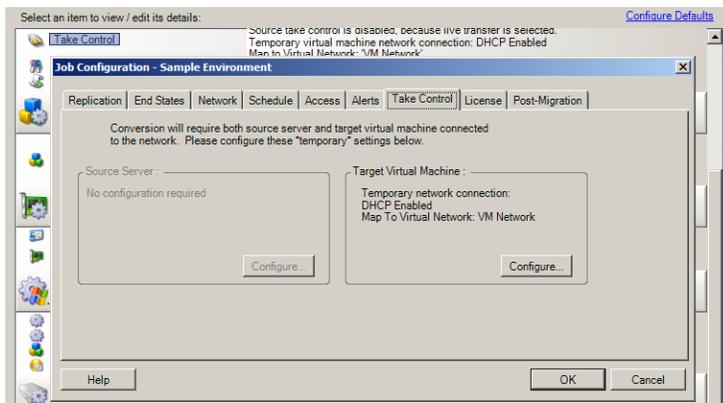
In Wizard mode: In the wizard's navigation pane, click *Networking.* In the *Source Take Control* or *Target Take Control* row, in either *Temp IP* or *Map To* drop-down menu, select *Configure*. To quickly select DHCP without opening configuration options, select *DHCP*.
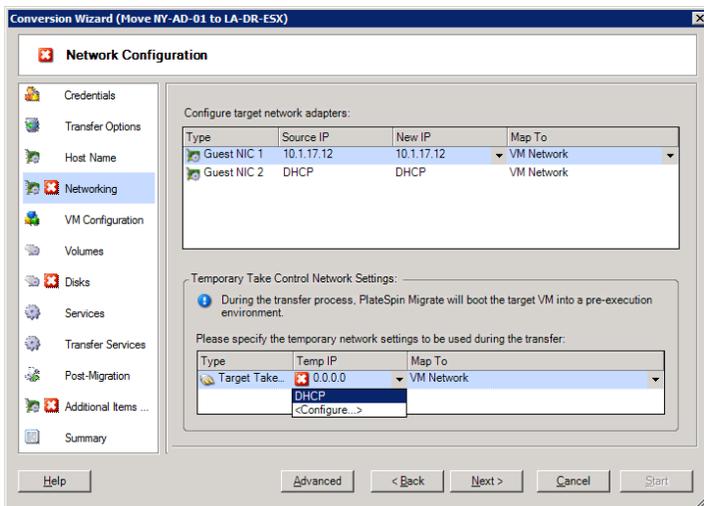
Configuration options for the Temporary networking vary and depend on whether the network interface is virtual or physical, and whether it is connecting a Windows or a Linux workload.

- "Temporary (Take Control) Network Settings: Physical Network Interfaces" on page 106
- "Temporary (Take Control) Network Settings: Virtual Network Interfaces" on page 107

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 107.

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 and Registering Target Physical Machines" on page 49.
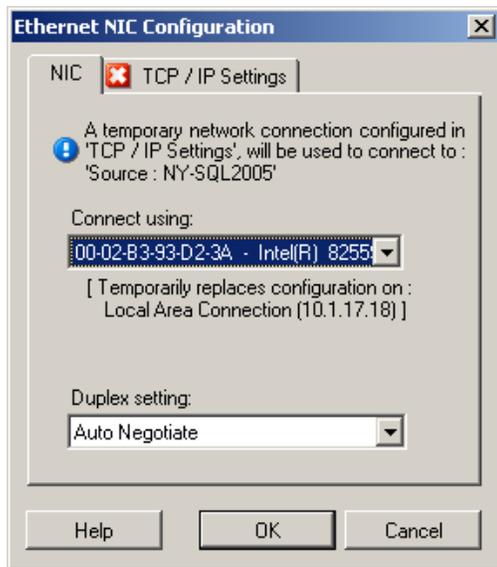
Connect using: If multiple network adapters are present, select the adapter that can communicate with both the PlateSpin Migrate Server and the target.

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 109.

Temporary (Take Control) Network Settings: Virtual Network Interfaces

These settings apply to both source and target Take Control network settings.
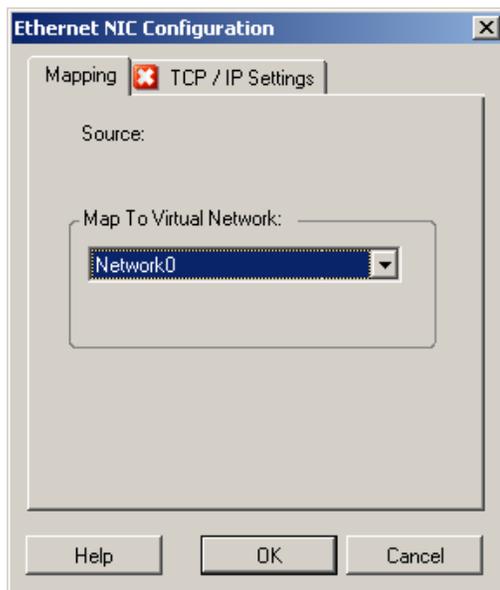
Map to Virtual Network: From the drop-down list, select the virtual switch or network to use for communication during an Offline migration. If multiple virtual network adapters are present, select the adapter that can communicate with both the PlateSpin Migrate Server and the source machine. This network 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 109.

6.9.2 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 the migration of Windows and Linux workloads, the target workload's post-migration network settings are configured while the workload is booted into a pre-execution environment. During the migration of Solaris workloads, the target's post-migration network settings are configured through the target zone's host.

To configure target post-migration network settings:

- **In Advanced mode:** In the Migration Job window, under the Network Configuration section, click *Guest NIC* (for target virtual machines) or *Network Connection* (for target physical machines).

- **In Wizard mode:** In the wizard's navigation pane, click *Networking*. In the Configure Target Network Adapters section, in either *New IP* or *Map To* drop-down menu, select *Configure*. To quickly select DHCP without opening configuration options, select *DHCP*.

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.

- ◆ "Post-Migration Networking for Physical Network Interfaces (Windows and Linux)" on page 108
- ◆ "Post-Migration Networking for Virtual Network Interfaces (Windows and Linux)" on page 109

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.
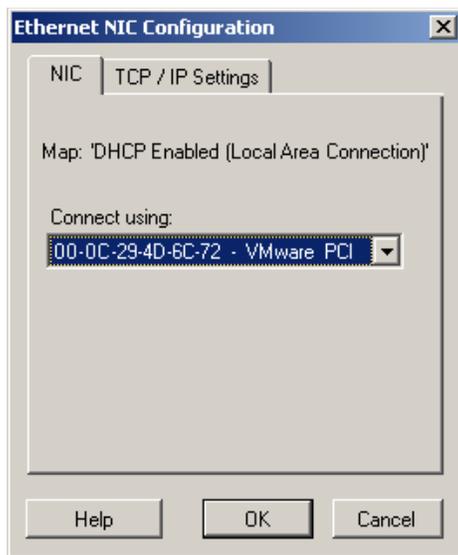
Connect using: If multiple network adapters are present, select the adapter that can communicate with the PlateSpin Migrate Server.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See "TCP/IP and Advanced Network Settings" on page 109.

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.
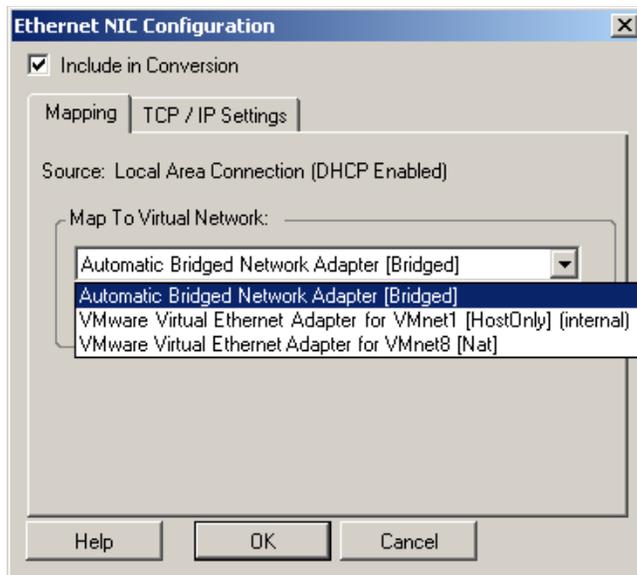
Include in Conversion: When this options is selected, PlateSpin Migrate creates a virtual NIC for a source NIC.

Map to Virtual Adapter: 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 target machine.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See "TCP/IP and Advanced Network Settings" on page 109.

6.9.3 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 110
- "TCP/IP and Advanced Network Settings (Linux and Solaris)" on page 110

TCP/IP and Advanced Network Settings (Windows)

The following are standard TCP/IP and advanced network settings for Windows workloads:
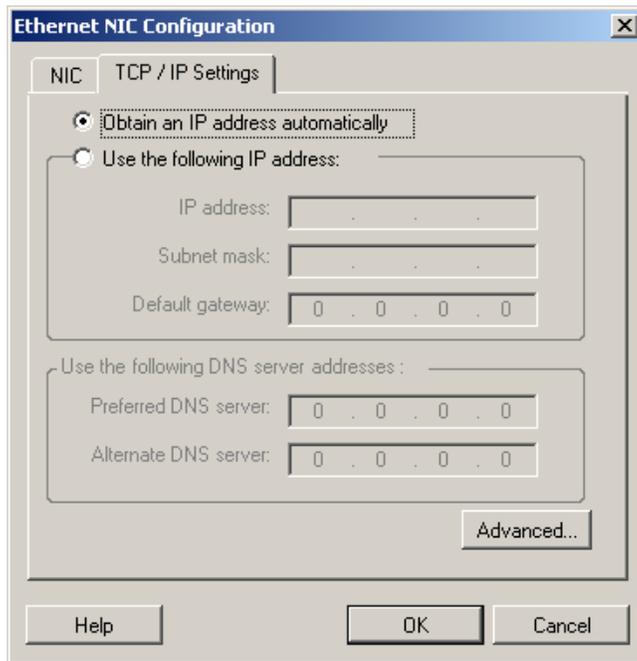
Obtain an IP address automatically: When this option is selected, the workload uses an IP address automatically assigned by a DHCP server during the migration process.

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 and Solaris)

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.

6.10 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:

- ◆ **In Advanced mode:** In the Migration Job window, in the Virtual Machine Configuration section, click *General*.
- ◆ **In Wizard mode:** In the wizard's navigation pane, click *Host Name.*

In Wizard mode, PlateSpin Migrate prompts only for basic virtual machine configuration options. For configuring advanced virtual machine options, such as those related to resource pools and CPU scheduling, switch to the Advanced mode.
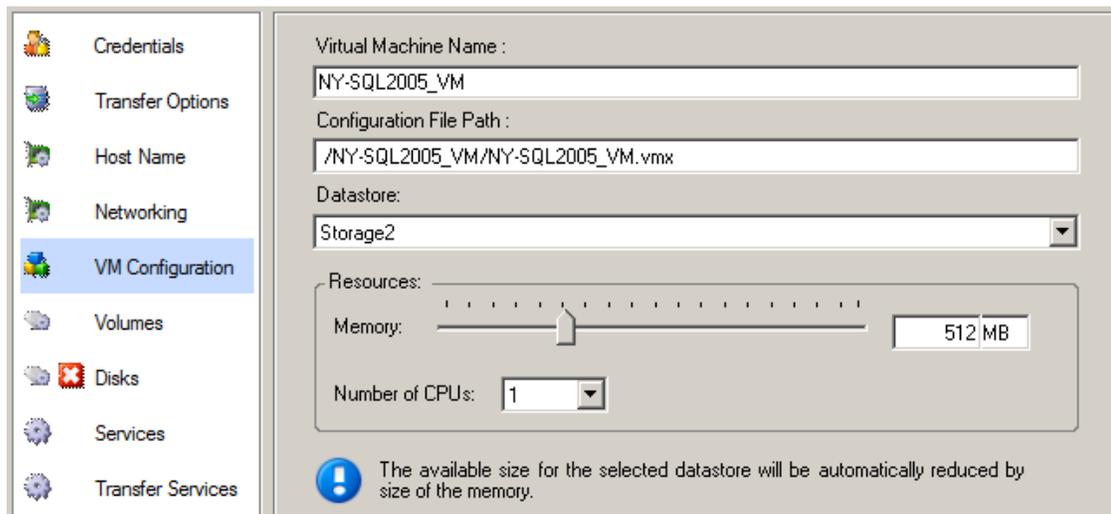
VM Name: Specify a display name for the new virtual machine.

Configuration File Path: Specify the path to the target virtual machine's configuration file.

Datastore: Select the required virtual machine's datastore.

Resources: Specify the amount of virtual RAM and the number of CPUs to be assigned to the virtual machine.

6.10.1 Virtualization Platform-Specific and Advanced VM Configuration Options

In Advanced mode, PlateSpin Migrate displays target virtual machine configuration options specific to the selected target and provides access to more advanced configuration options.

- ◆ "Virtual Machine Configuration: VMware ESX 4.1 and Earlier" on page 113
- ◆ "Virtual Machine Configuration: Citrix XenServer" on page 114
- ◆ "Solaris Zone Configuration" on page 115

Virtual Machine Configuration: VMware ESX 4.1 and Earlier

The following are configuration options specific to VMware ESX systems. To access settings that control resource pools, the number of CPUs, and CPU scheduling affinity, click *Advanced.*
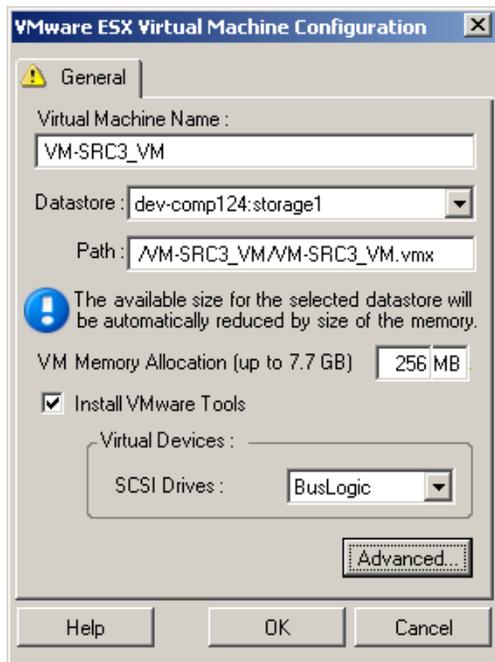
Virtual Machine Name: Specify the display name for the new virtual machine.

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.
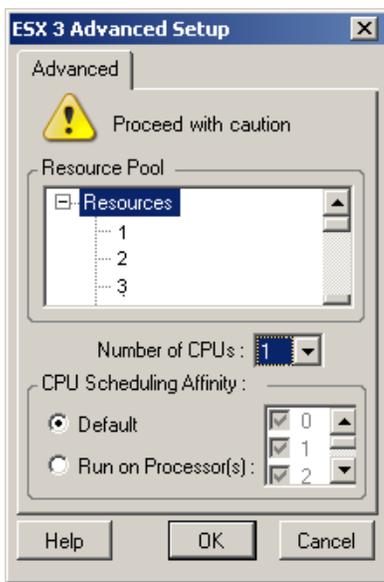
Resource Pool: If required, assign your target VM to a resource pool. When no resource pool is specified, the VM is assigned to the root resource pool.

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: Citrix XenServer

The following are configuration options specific to Citrix XenServer.
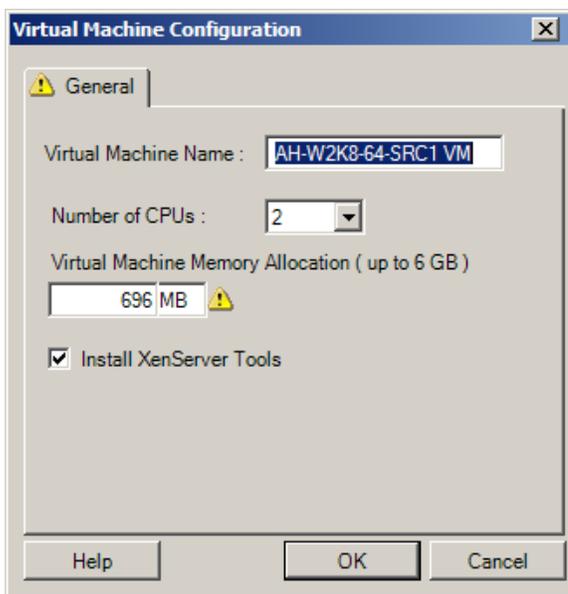
Virtual Machine Name: Specify the display name for the new virtual machine.

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).

Solaris Zone Configuration

The following are configuration options specific to Solaris. To access settings that control resource pools, CPU shares, click *Advanced.*
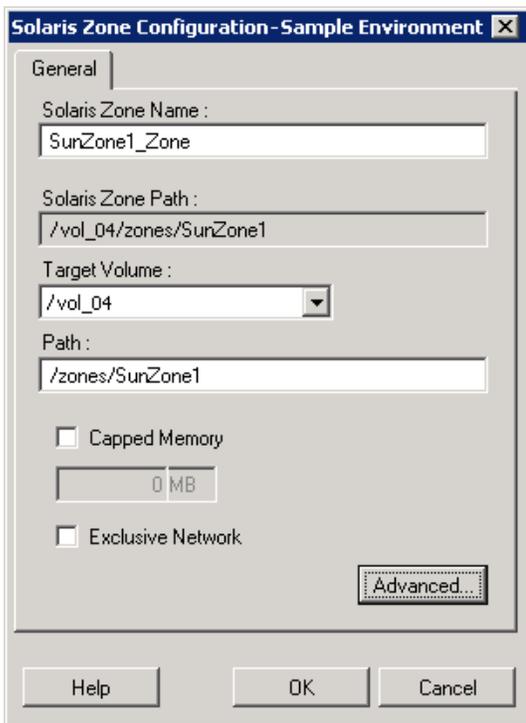
Solaris Zone Name: Type a name for your target Solaris zone.

Solaris Zone Path: The path is determined based on the two subsequent values.

Target Volume: This drop-down list contains discovered volumes on the target zone host. Select the required volume for the target zone. Do not use system hierarchies, such as /usr, /sbin, /lib, /var, and /platform.

Path: Type a pathname for the target zone.

Capped Memory: Specify the number or maximum memory to be allocated to the guest zone.

Exclusive Network: If this option is selected, the network interface is used exclusively by the guest zone. If this options is deselected, the NIC is shared by the guest zone, other zones if available, and the host.

Advanced: Click this button to access options that control resource pool assignment and CPU shares allocated to the guest zone.
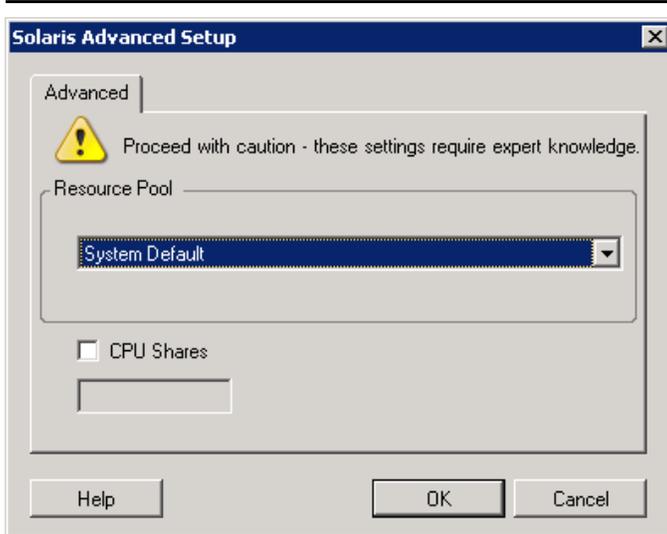
Resource Pool: This drop-down list contains discovered resource pools on the target zone server. Select the required resource pool for the guest zone.

CPU Shares: Select this option to allocate CPU shares to the guest zone, then specify the required number of CPU shares in the corresponding field.

6.11 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, Solaris services, 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 6.11.1, "Handling the Startup Mode of Services (Windows Targets)," on page 116
- Section 6.11.2, "Handling Source Workload Services or Daemons During Live Transfer (Windows and Linux)," on page 117
- Section 6.11.3, "Viewing Windows System Files Selected for Replacement During a Migration," on page 119
- Section 6.11.4, "Handling the Run Level of Daemons (Linux Targets)," on page 119
- Section 6.11.5, "Handling Services on Solaris Targets," on page 120
- Section 6.11.6, "Handling Virtualization Enhancement Software," on page 121

6.11.1 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 Migrate Server default options. See "Target Service Defaults" on page 37.

To configure the post-migration startup mode of Windows services:

 ◆ **In Advanced mode:** In the Migration Job window, under the Operating System and Application Configuration section, click *Windows Services (Target)*, then click an item in the *Start Mode* column.
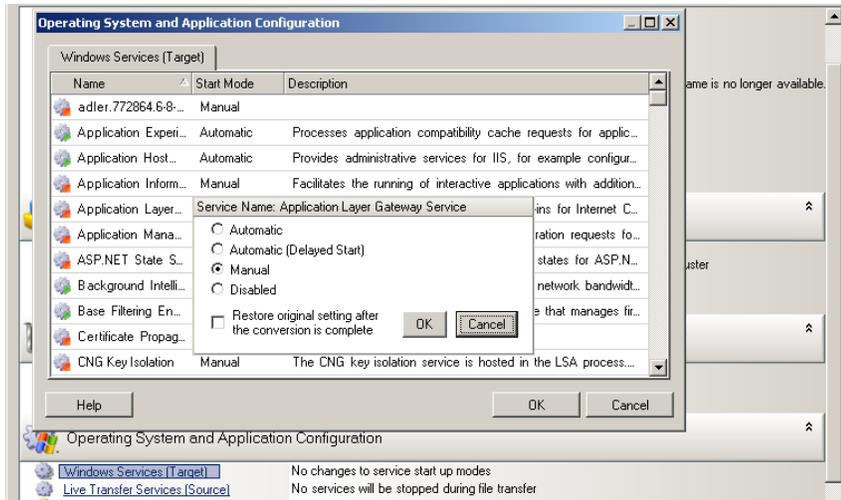
 ◆ **In Wizard mode:** In the wizard's navigation pane, click *Services*, then click an item in the *Target Mode* column.

6.11.2 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, consder using the custom `freeze` and `thaw` scripting capability. See "Freeze and Thaw Scripting Capabilities (Linux Block-Level Migrations)" on page 61.

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 36.

To specify which services or daemons you want the system to stop during Live Transfer:

◆ **In Advanced mode:** In the Migration Job window, under the Operating System and Application Configuration section, click *Live Transfer Services/Daemons (Source)*. To indicate that you want SQL Server and Exchange Server database files copied during the migration, click *Advanced* (applicable to Windows systems only).
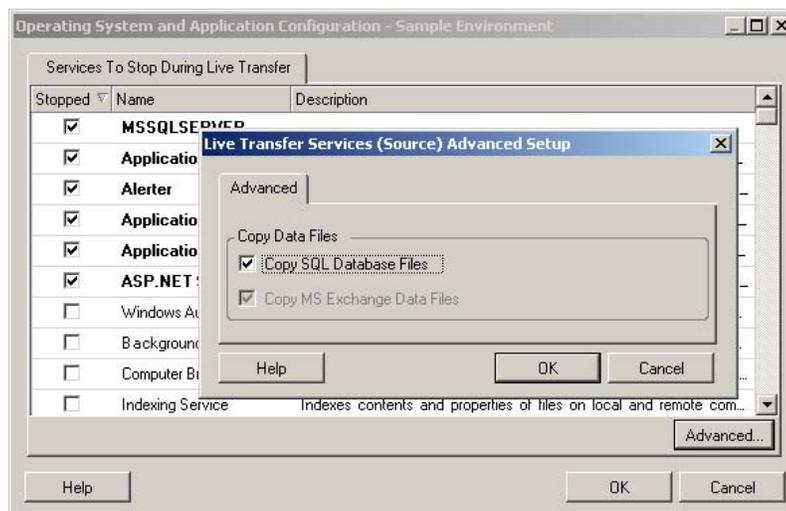
◆ **In Wizard mode:** In the wizard's navigation pane, click *Transfer Services.*
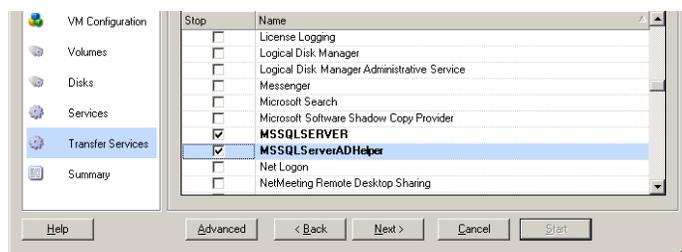
6.11.3 Viewing Windows System Files Selected for Replacement During a Migration

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:

◆ **In Advanced mode:** In the Migration Job window, under the Operating System and Application Configuration section, click *System Files*.
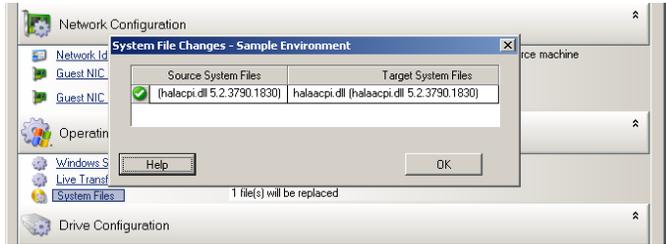
◆ **In Wizard mode:** Not available.

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 (http://www.novell.com/support/viewContent.do?externalId=7920815) (FAQ: Understanding the System Files Information screen).

6.11.4 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:

* **In Advanced mode:** In the Migration Job window, under the Operating System and Application Configuration section, click *Linux Daemons (Target)*, then click an item in the *Run Level* column.
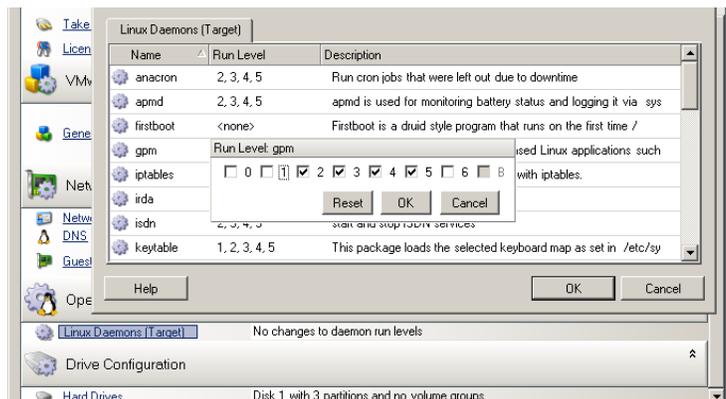
* In Wizard mode: In the wizard's navigation pane, click *Daemons*, then click an item in the *Target Run Level* column.
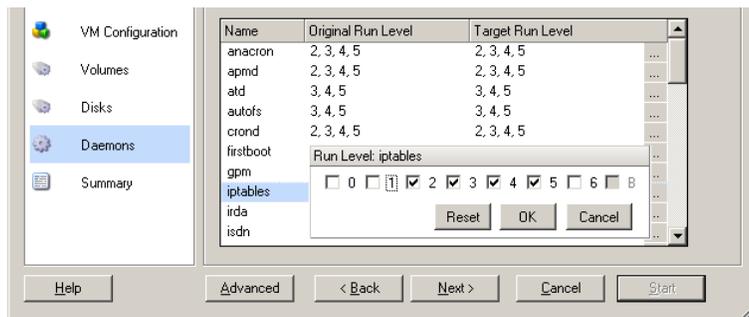
6.11.5 Handling Services on Solaris Targets

You can select Solaris services for the system to enable or disable after the migration is complete.

* **In Advanced mode:** In the Migration Job window, under the Operating System and Application Configuration section, click *Solaris Services Configuration*, then select or deselect the required service in the *Enabled* column.
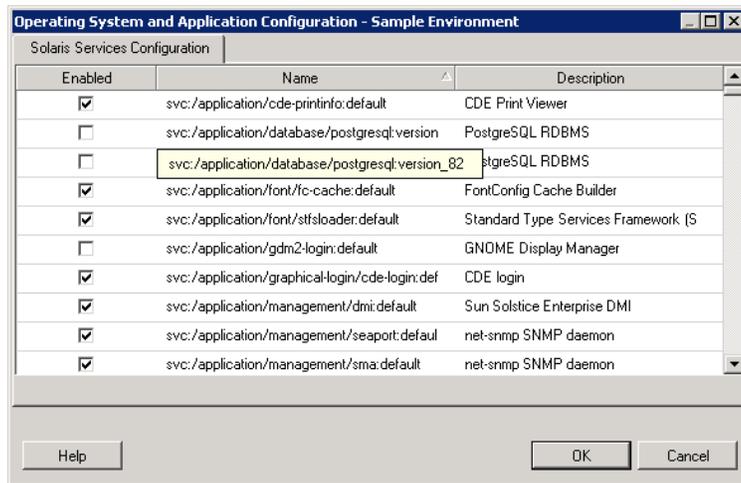
◆ In Wizard mode: Not available.

6.11.6 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:

◆ **In Advanced mode:** In the Migration Job window, under the Operating System and Application Configuration section, click *Clean up VMware Tools*.
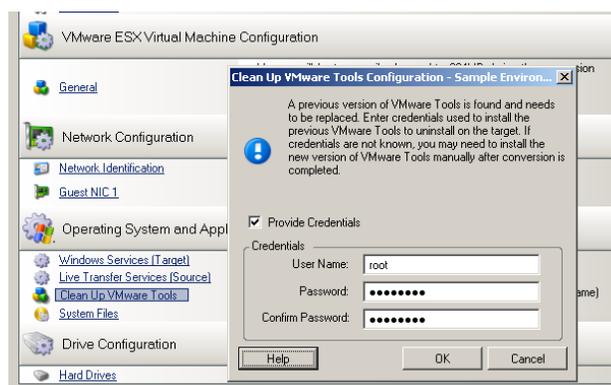
- **In Wizard mode:** In the wizard's navigation pane, click *VM Tools.*
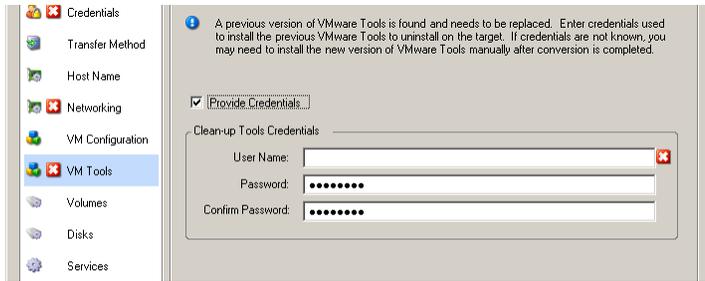
Depending on the target, PlateSpin Migrate identifies existing instances of VMware Tools and prompts to either replace or remove them, as applicable:

- **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.

6.12   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.

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, 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.

Storage layout and volume configuration settings depend on the job configuration mode (Advanced or Wizard), migration type, target virtualization platform, and source operating system.

OES 2: NSS File Systems and EVMS

PlateSpin Migrate supports NSS file systems of OES 2 workloads along with EVMS volume layout features. You can:

- Preserve source EVMS layouts on your target workload
- Create, delete EVMS volumes, and manage EVMS volume layouts
- Copy NSS pools from your source to the target
- Make use of NSS snapshots for consistent copying of your source volumes.

The following topics provide additional information:

- Section 6.12.1, "Storage Layout and Volume Configuration (Wizard Mode)," on page 124
- Section 6.12.2, "Drive Configuration (Advanced Mode)," on page 127
- Section 6.12.3, "Volume Configuration in Server Sync," on page 134
- Section 6.12.4, "Handling Volumes and Image Files in X2I (Imaging) Migrations," on page 136

6.12.1 Storage Layout and Volume Configuration (Wizard Mode)

Storage Layout Configuration: X2V (Wizard Mode)

Use this Migration Wizard page to specify a disk mapping scheme during workload virtualization operations.

In the wizard's navigation pane, click *Disks*.
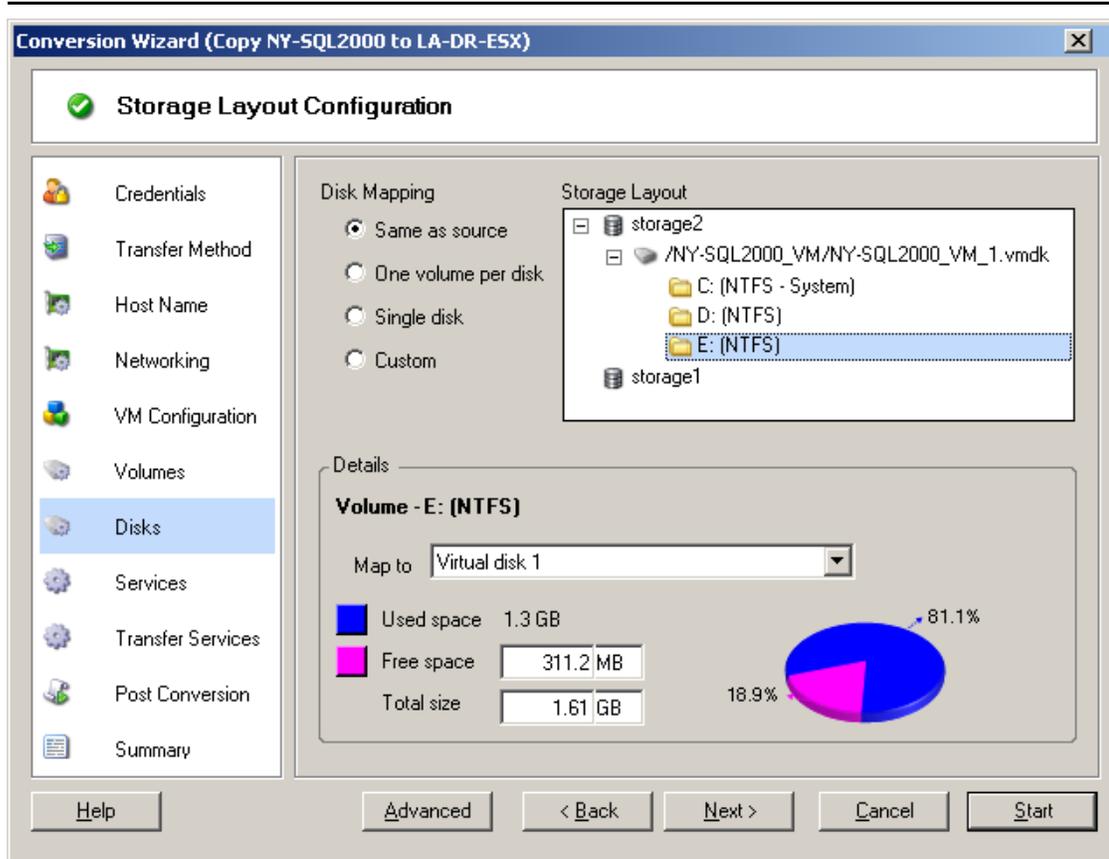
Storage Layout section: Displays information in a tree format according to the options selected.

Disk Mapping options: These options control how the source disk arrangement is propagated on the target. Select the option that best suits your migration. To add a virtual disk, select *Custom*, then click *Add Virtual Disk* in the Details section.

Details section: Displays information applicable to the selected item in the Storage Layout tree.

You can rearrange disks by dragging and dropping.

To add a virtual disk, select a storage item in the upper section, and in the Details section click *Add Virtual Disk*.

To view the path and name of a virtual disk and datastore it is assigned to, select the virtual disk.

Volume Configuration: X2V, X2P (Wizard Mode)

Use this Migration Wizard page to select volumes to include in the migration and to adjust the volume size on the target.

In the wizard's navigation pane, click *Volumes*.
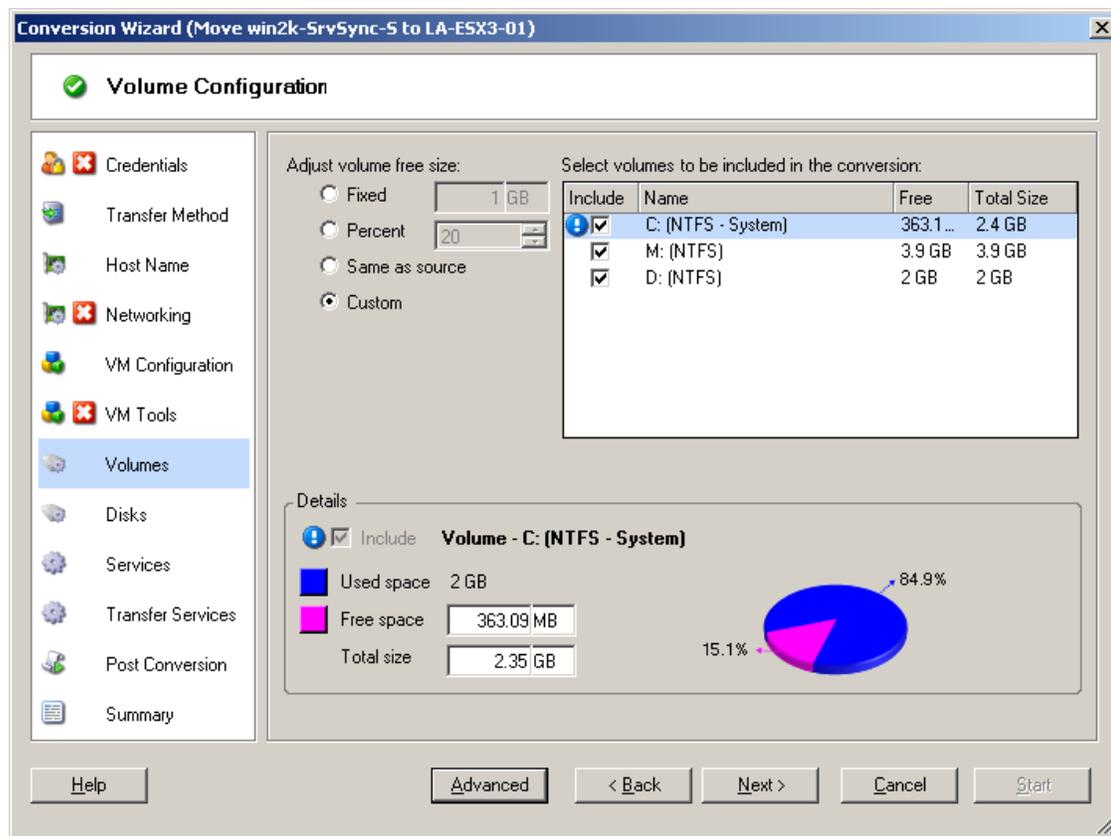
Adjust Volume Free Size: Use one of the options to specify the volume free size. *Custom* lets you enter a free space value in the Details area.

Select volumes to include in the conversion: Select the volumes to be converted. System or boot volumes must be selected.

Details: View information about the selected volume. You can modify free space and total size values.

Linux Volume Groups Configuration (Wizard Mode)

Use this wizard page to manage LVM (Logical Volume Manager) volume groups. This page is displayed only if the source has LVM installed.

In the wizard's navigation pane, click *Volume Groups*.
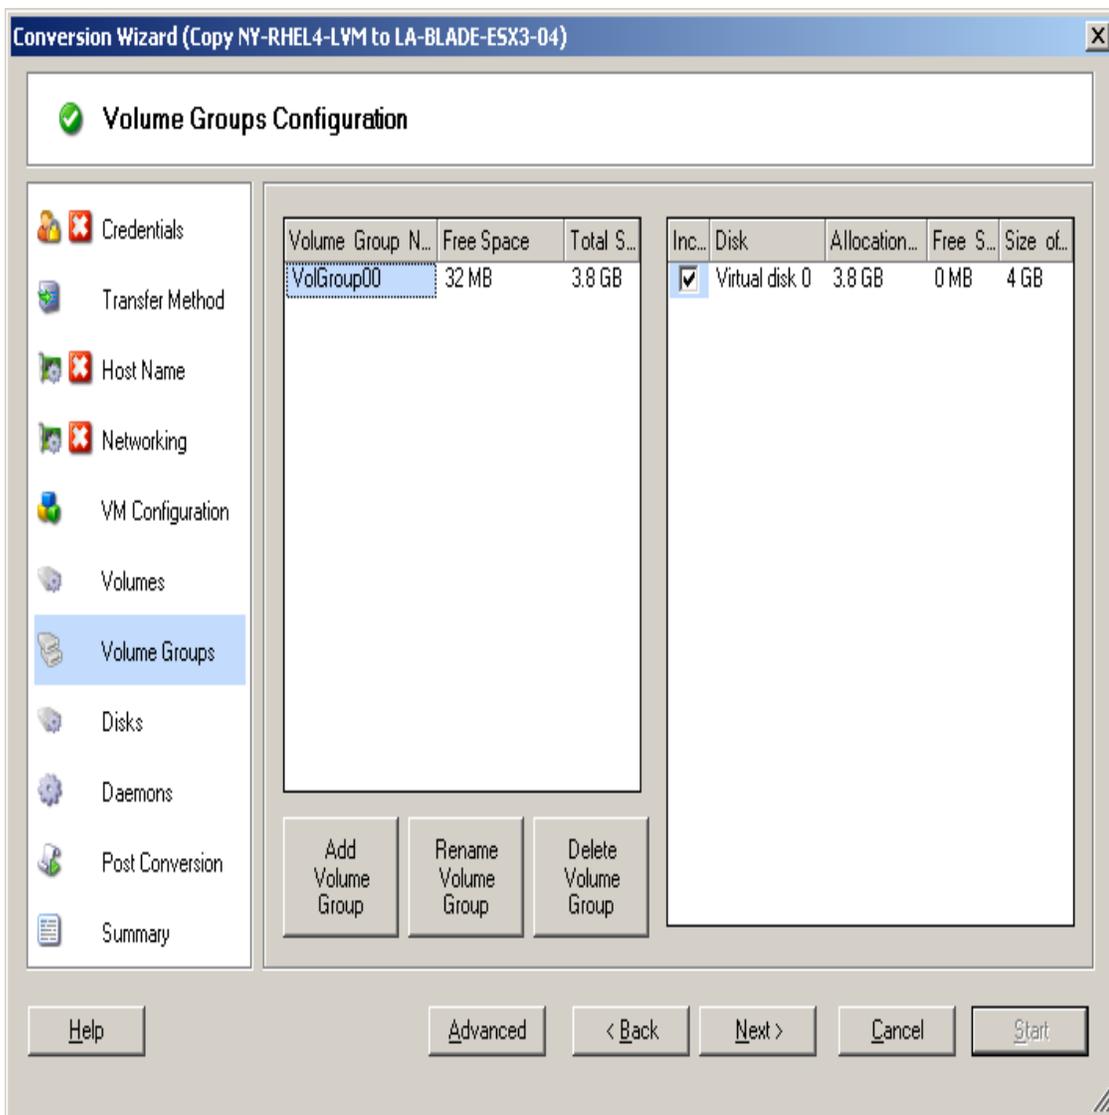
To add, rename or delete a volume group, click the corresponding button.

To include a volume group in the migration, select the corresponding check box in the right pane.

To allocate storage to a volume group, click in the *Allocation for Volume Group* row.

NOTE: Unassigned volume groups are removed prior to migration.

6.12.2 Drive Configuration (Advanced Mode)

When you are working in Advanced mode, the Peer-to-Peer Conversion job window provides access to a single configuration interface that combines the wizard's Volumes and Disks screens.

To access drive configuration options in Advanced mode:

 * In the Migration Job window, under the Drive Configuration section, click *Hard Drives*.
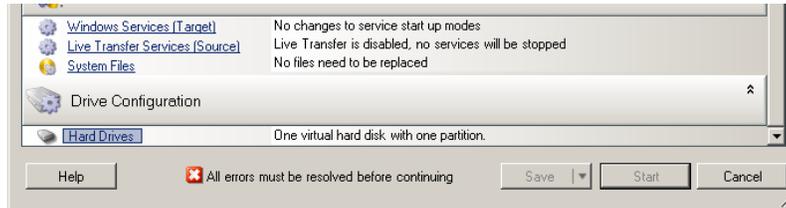
Settings vary depending on the target system.

Windows X2P Drive Configuration (Advanced Mode)

Use these settings to select the volumes to copy during the migration:
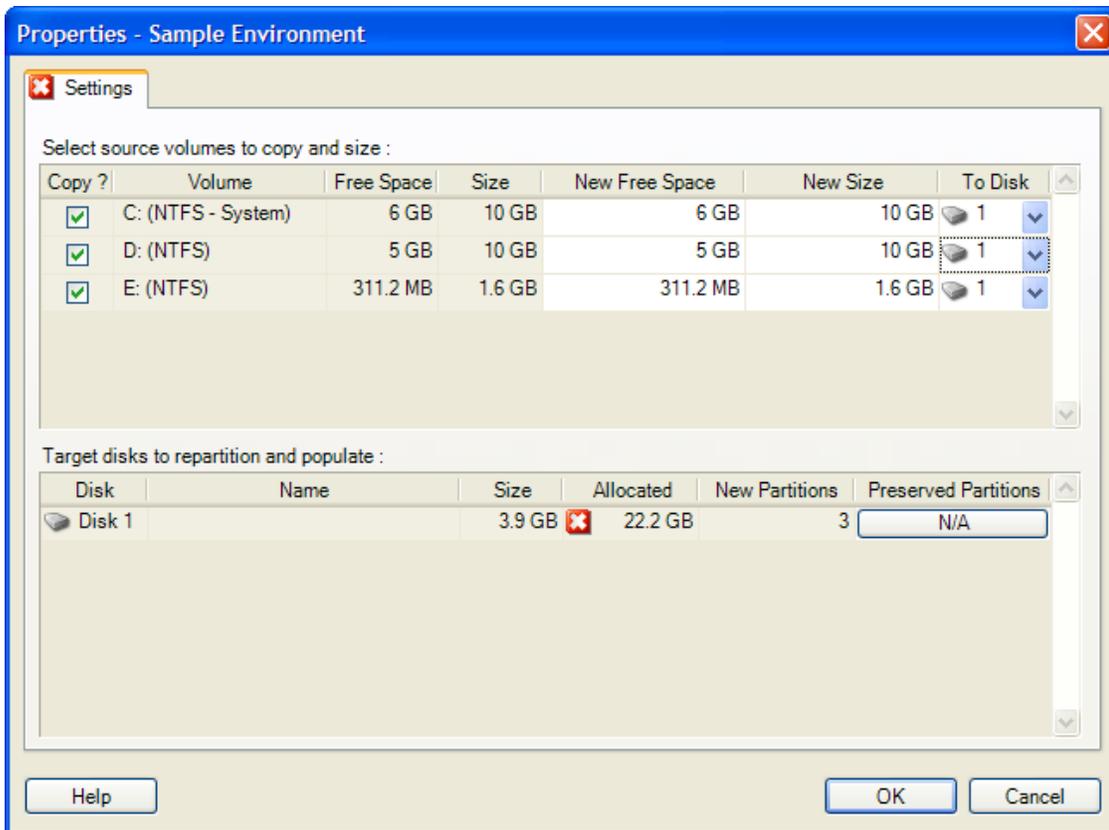
Copy: Select the volumes to be copied during the migration.

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 (Advanced Mode)

Use these settings to select the volumes and non-volume source spaces to copy and size during the migration. The *Volume Group* tab is available only if LVM is installed on the source.

- ◆ "Linux Drive and LVM Volume Configuration (Settings Tab)" on page 129
- ◆ "Linux Drive and LVM Volume Configuration (Volume Groups Tab)" on page 130
- ◆ "Linux (OES 2) Drive and Volume Configuration (EVMS Volumes tab)" on page 131

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.

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.
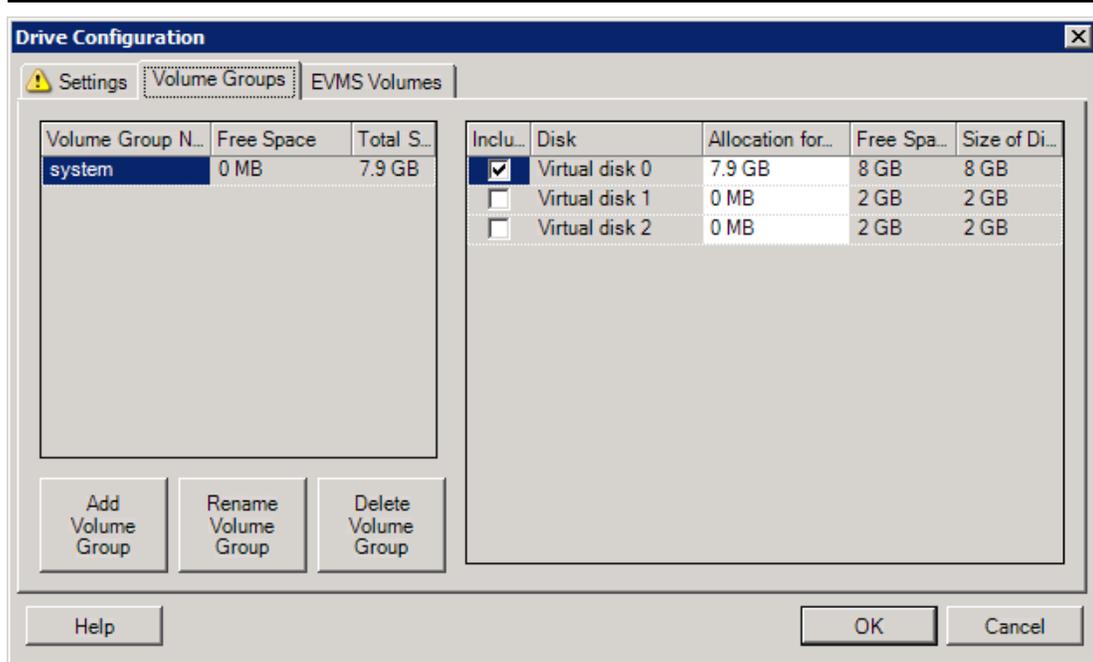
Add Volume Group: Creates a volume group on the target machine that is not present on the source machine.

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.

Linux (OES 2) Drive and Volume Configuration (EVMS Volumes tab)

Use these settings to manage EVMS volumes (OES 2 workloads with NSS file systems).
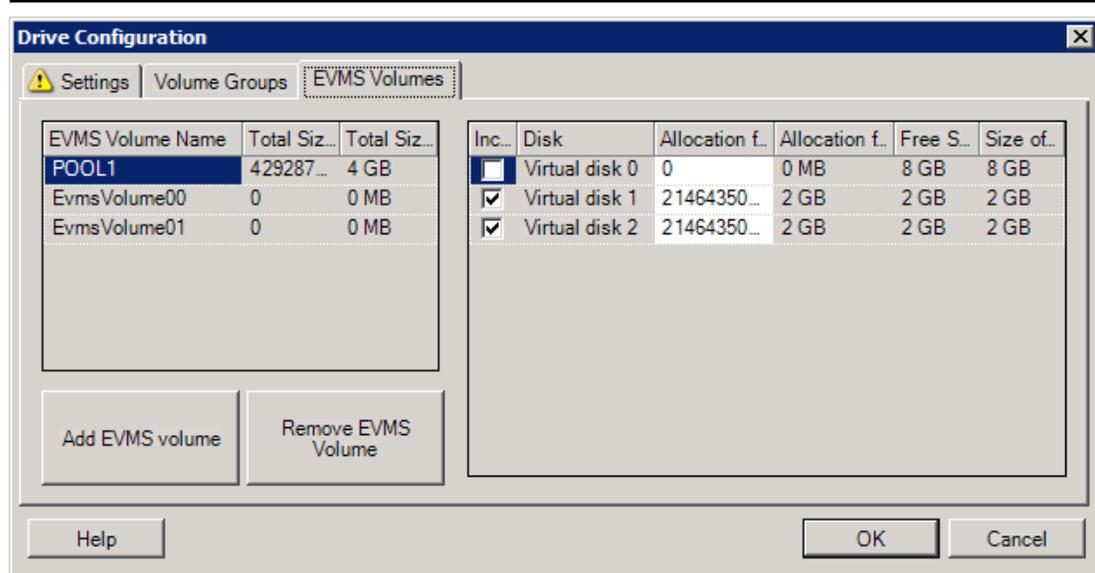
(**Left Section**)

EVMS Volume Name: Displays a list of EVMS volumes found on the source.

Add EVMS volume and *Remove EVMS Volume* buttons: Click to create or delete EVMS volumes.

(**Right Section**)

Include: Select this option to include the indicated volume in the migration.

Allocation for EVMS Volume: To allocate space, select a volume, then select the disks to include in it. Specify the amount of space to be allocated to it on each included disk.

NOTE:

* Selecting multiple disks for a single EVMS volume is possible only for volumes with NSS pools.
* The source and target EVMS volumes must be equal in size.

Target VM-Specific P2V/V2V Drive Configuration (Advanced Mode)

When you configure a peer-to-peer virtualization job in Advanced mode, the job configuration window provides access to settings specific to the target virtualization platform.

Drive Configuration: VMware ESX 3

The following are drive configuration settings specific to VMware ESX 3:
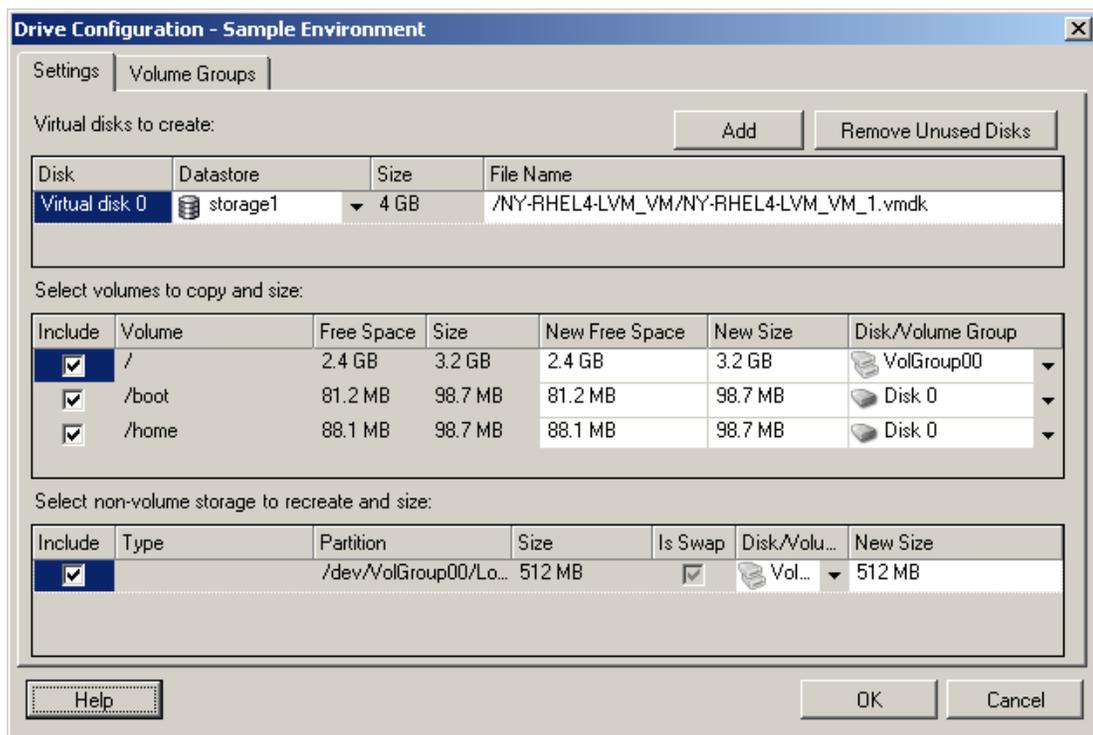
Datastore: Select the datastore volume on the ESX 3.0 server where you want to place the vmdk files.

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: VMware ESX 2

The following are drive configuration settings specific to VMware ESX 2:
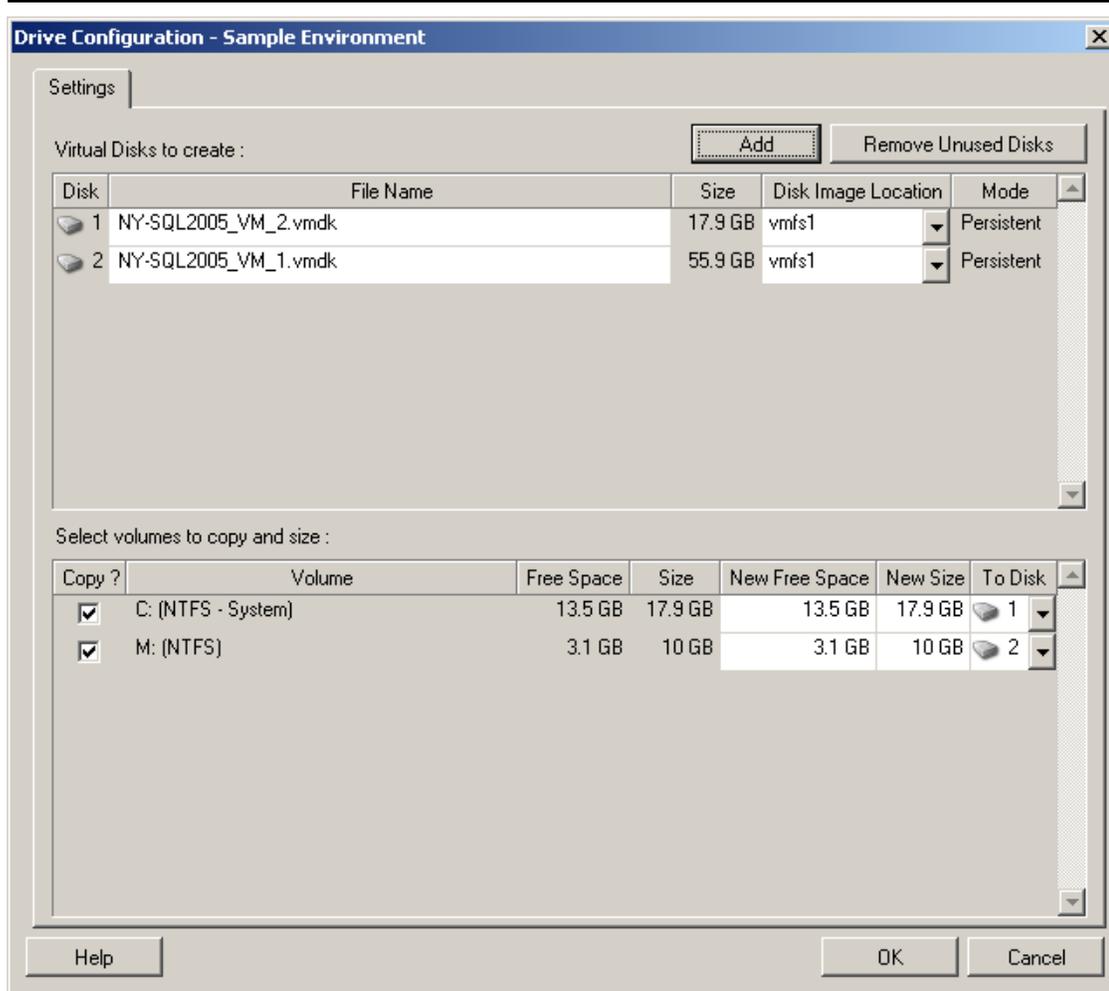
Disk Image Location: Select the location on the ESX 2.0 server where you want to place the `vmdk` files.

Copy: Select the required source volumes to copy during the migration.

New Free Space: To resize the volume during the migration, specify the desired amount of free space on the target. PlateSpin Migrate automatically adjusts the *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 to a volume group, if LVM is enabled. 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.

Solaris Zone Volume Configuration

The following are volume configuration options for Solaris zones.
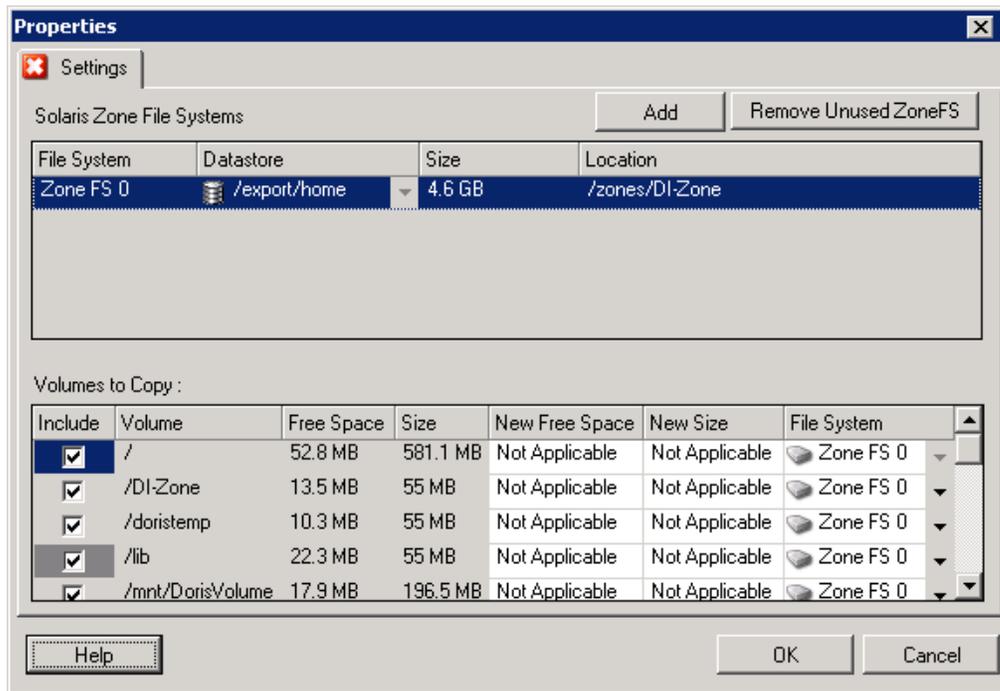
Solaris Zone File Systems: Options in this area provide information about and a means to manage zone file systems on the target. For the *Zone FS 0* file system, you cannot modify the *Datastore* and *Location* values, which are based on the zone path you specify for the zone. To change the *Datastore* and *Location* values, change the *Target Volume* and *Path* values in the target zone configuration settings (see "Solaris Zone Configuration" on page 115).

To add a new file system, click *Add* and specify the required target directory in the Datastore column. To remove file systems that you do not need on the target, click *Remove Unused FS*.

Volumes to Copy: Options in this area provide information about volumes that were discovered on the source and enable you to select which ones to reproduce on the target zone. For volumes that you select to reproduce, you can specify a file system from the list in the upper area. The root volume cannot be deselected (excluded from migration); it is always selected and assigned to the *Zone FS 0* file system.

6.12.3 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 72.

To access volume configuration options in a Server Sync job:

- ◆ **In Advanced mode:** In the Migration Job window, under the Drive Configuration section, click *Volume Mapping* (for Windows machines) or *Drives and Volumes* (Linux machines).
- ◆ **In Wizard mode:** Not available.

The following topics provide information about Server Sync volume configuration options specific to Windows and Linux workloads.

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 enebles you to specify the required mapping.
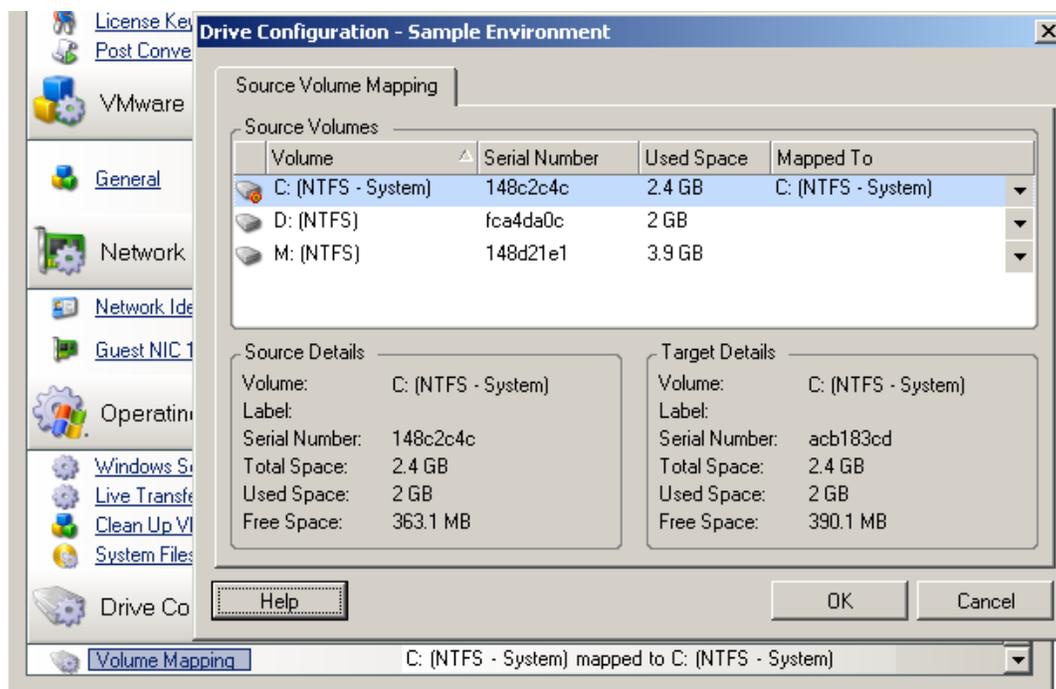
Mapped To: Map each volume on the source to an existing volume on the target.

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 enebles you to specify the required mapping.
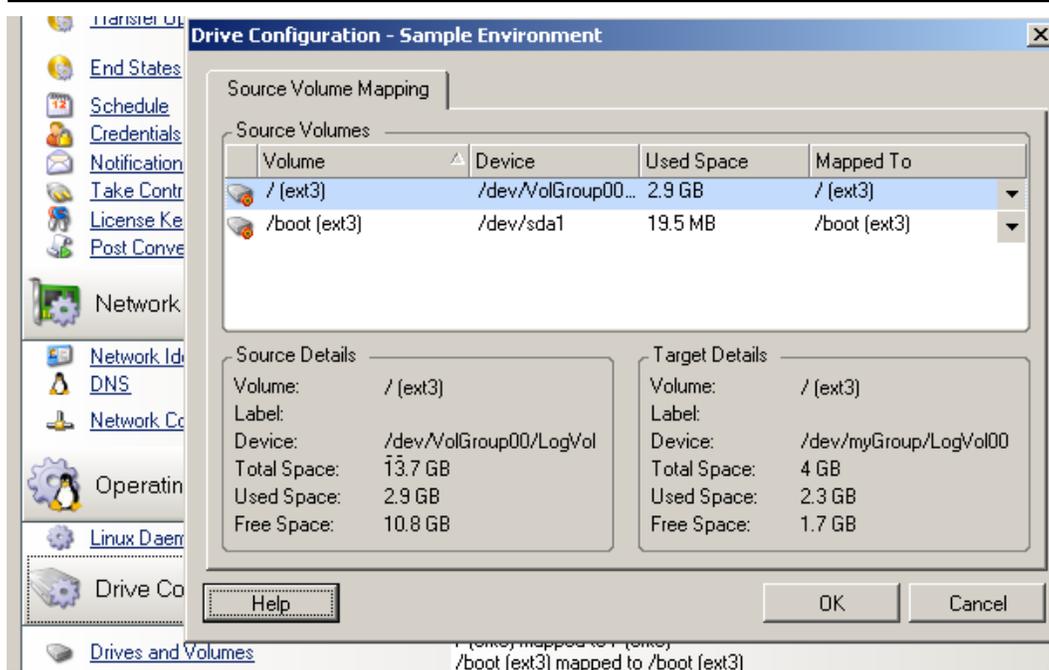
Mapped To: Map each volume on the source to an existing volume on the target.

6.12.4   Handling Volumes and Image Files in X2I (Imaging) Migrations

When capturing a PlateSpin Image or importing volumes into a PlateSpin Image, PlateSpin Migrate provides a mechanism for including required volumes in an image, and specifying volume data mapping and image configuration options.

 * "Target Volume Configuration: X2I (Wizard Mode)" on page 137
 * "Image Configuration: X2I (Wizard Mode)" on page 138

Target Volume Configuration: X2I (Wizard Mode)

When you are setting up a Capture Image or Import Image job in wizard mode, use this page to select volumes to include in the image and to specify paths to existing volume data.
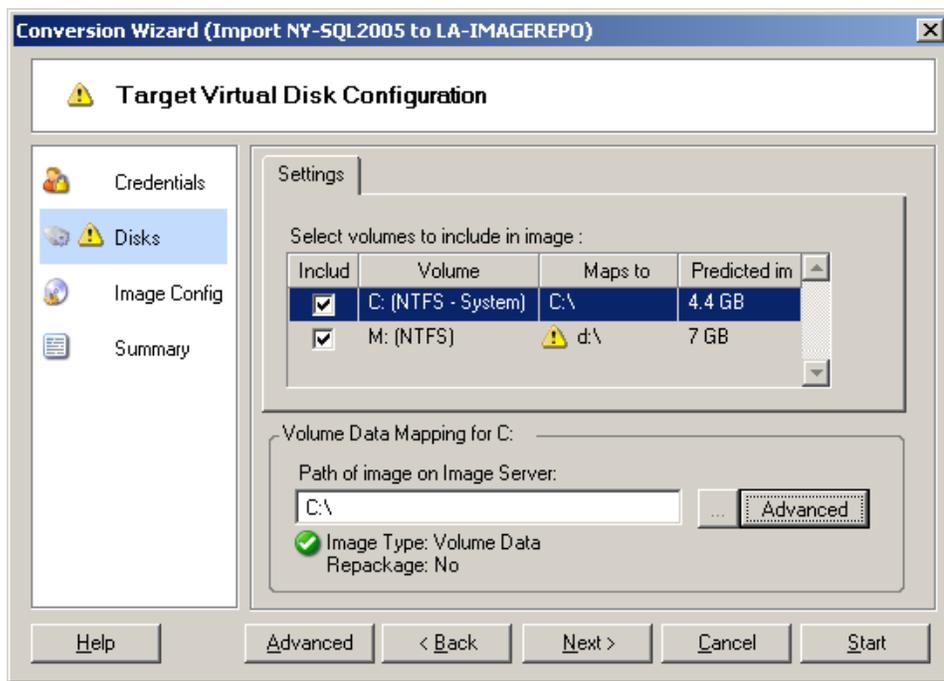
For each volume selected for the image, specify the path to the corresponding image data.

For a Capture Image job, specify the path to the image package file (`*.pkg`).

For an Import Image job, specify the path to the directory containing the raw volume data.

To repackage the image, click *Advanced* in the *Volume Data Mapping* area and select *Repackage* in the Properties dialog box. Specify a path for the new image file, or use the default path.

Image Configuration: X2I (Wizard Mode)

When you are setting up a Capture Image or Import Image job in wizard mode, use this page to specify the image name and the path where it is stored.
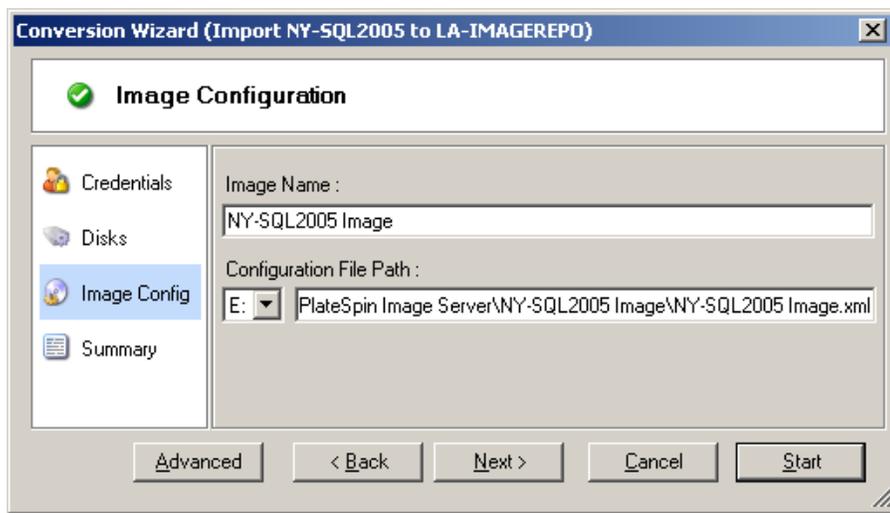
Image Name: Type a name for the workload image or accept the default.

Configuration file path: Type the full path to the image's XML configuration file or accept the default.

6.13 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 59.

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:

- **In Advanced mode:** In the Migration Job window, under the Virtual Machine Configuration section, click *Post Conversion*.
- **In Wizard mode:** In the wizard's navigation pane, click *Post Conversion*.

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*.

A 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 17.

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 16.

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/ESX Servers?

Yes. See "Discovering the Details of Source Workloads and Targets" on page 47.

Can multiple migrations run simultaneously?

Yes. See "Performance" on page 17.

B Troubleshooting PlateSpin Migrate

This section provides a series of topics about troubleshooting PlateSpin Migrate.

B.1 Discovery

Table B-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 Migrate Server. A common cause is incorrect information entered during the registration process. To restart the registration process, enter `RegisterMachine.bat`. Ping to confirm basic connectivity with the PlateSpin Migrate 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	KB Article 7920291 (http://www.novell.com/support/viewContent.do?externalId=7920291) contains troubleshooting checklists for discovering the following: ◆ Linux servers and VMware ESX Servers ◆ Windows-based source and target servers The article also has instructions for troubleshooting WMI connections and checking if DCOM is enabled.
"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$\{randomID}.xml"	This error might occur when you are attempting to discover a Windows NT 4.0 Server with an older version of WMI Core installed. It can also occur with Windows Server 2000 or 2003. For Windows NT 4.0 servers, verify the version of WMI installed on the server by doing the following: 1. Locate the `wbemcore.dll` file in the `\winnt\system32\wbem` directory. 2. Right-click `wbemcore.dll` and click *Properties*. 3. Check the version number of the file. If the version number is not 1.50.xxxx, upgrade the WMI core to V1.5. See the following Microsoft Web site (http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=C174CFB1-EF67-471D-9277-4C2B1014A31E)). After WMI v1.5 has been installed, rediscover the server. In some cases, reinstalling WMI 1.5 might resolve the issue. If this discovery error still occurs after installing WMI 1.5 on a Windows NT 4.0 Server or if it is still occurring when you attempt to discover a Windows 2000/2003 Server, do the following: 1. Ensure that the `Admin$` share on the remote machine is accessible and then proceed to the next step. If the share is not accessible, enable it and try the discovery again. 2. Navigate to the `..\PlateSpin Migrate Server\Web` directory. 3. Use a text editor to open the `web.config` file. Locate the `<add key="MachineDiscoveryUsingService" value="false" />` entry near the end of the file and change the value to `"true"`. 4. Save the `web.config` file and retry the discovery.

Related KB Articles:

ID	Description
7920339 (http://www.novell.com/support/viewContent.do?externalId=7920339)	ERRMSG: Discovery fails with "The request failed with HTTP status 407" message
7920862 (http://www.novell.com/support/viewContent.do?externalId=7920862)	ERRMSG: Recoverable Error: ControllerConnectionBroken during discovery
7920291 (http://www.novell.com/support/viewContent.do?externalId=7920291)	ERRMSG: Server details discovery problems

B.2 Peer-to-Peer Migrations (Windows)

Table B-2 *Common Issues and Solutions Related to Peer-to-Peer Migrations (Windows)*

Problems or Messages	Solutions
One of the following errors displays during offline migration: ◆ Waiting for Controller to start (Failed) ◆ Controller Connection Not Established ◆ Controller Connection Broken ◆ Unable to start the Heartbeat Service	This indicates one of the following problems: ◆ 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 Migrate Server. ◆ The source/target machine was not able to fully boot into the pre-execution environment. To diagnose the exact cause of failure, check the state of the system where the controller failed to start. Commands such as `ipconfig` and `ping` are available to verify basic network connectivity.
File transfer hangs at 1% or progresses at a slow pace	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.
Unable to determine suitable boot partition	When converting existing source servers, the boot volume must pass the following checks: ◆ 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)	There is a problem with the Operations Framework Controller on the PlateSpin Migrate Server. Use the Windows services plug-in to confirm that the Controller is running. See KB Article 7920862 (http://www.novell.com/support/viewContent.do?externalId=7920862) for other troubleshooting instructions.

Problems or Messages	Solutions
Troubleshooting failures at the *Configuring Operating System* stage (also applicable to *Configure Target Machine* or *Configuring Virtual Machine* migration steps)	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.
	KB Article 7920327 (http://www.novell.com/support/viewContent.do?externalId=7920327) contains a detailed troubleshooting checklist and lists information required if technical support is necessary.
Live Transfer is unavailable	Either an unsupported file system or operating system exists on the server. Only NTFS file systems are supported. If a FAT partition is selected in the Drive Configuration dialog box. Clearing the option should enable the Live Transfer option in Job Configuration. For more information, see "File-Level Transfer (Live)" on page 14.

Related KB Articles:

ID	Description
7920862 (http://www.novell.com/support/viewContent.do?externalId=7920862)	ERRMSG: PlateSpin Migrate Job remains at a "Scheduled" or "Recoverable Error" state
7920810 (http://www.novell.com/support/viewContent.do?externalId=7920810)	INFO: Restore job stalls - "The configuration service in the target machine..."
2790341 (http://www.novell.com/support/viewContent.do?externalId=7920341)	INFO: What ports does PlateSpin Migrate use during discovery, migration and file transfer?

B.3 Using Images

Table B-3   *Common Issues and Solutions Related to PlateSpin Images*

Problems or Messages	Solutions
Cannot see PlateSpin Images on PlateSpin Image Server	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.
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.

Problems or Messages	Solutions
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 (http://www.novell.com/support/viewContent.do?externalId=7920879)	ERRMSG: The file cannot be accessed by the system

B.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 B.4.1, "Cleaning Up Windows Workloads," on page 147
- Section B.4.2, "Cleaning Up Linux Workloads," on page 148

B.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
PlateSpin Block-Based Transfer Component	All migrations	See KB Article 7005616 (http://www.novell.com/support/viewContent.do?externalId=7005616).
SteelEye Block-Based Transfer Component (discontinued)	One-time Migration (*Move*)	Reboot source workload to remove driver.

	Migration (*Copy*)	1. Use the Windows Add/Remove Programs applet (run `appwiz.cpl`) and remove the component. Depending on the source, you might have either of the following versions:
		
		    - SteelEye Data Replication for Windows v6 Update2
		    - SteelEye DataKeeper For Windows v7
		
		2. Reboot the machine.

File-based Transfer Component	All Migrations	At root level for each volume migrated, remove all files named `PlateSpinCatalog*.dat`

Component	Use Case	Removal Instructions

Workload discovery software	All migrations	1. In the Servers view, undiscover the source (right-click, then select *Undiscover*).
		
		2. In the source workload's `Windows` directory:
		
		    ◆ Remove all files named `machinediscovery*`.
		
		    ◆ Remove the subdirectory named `platespin`.

Controller software	All migrations	1. In the Servers view, undiscover the source (right-click, then select *Undiscover*).
		
		2. Open a command prompt and change the current directory to:
		
		    ◆ `\Program Files\platespin*` (32-bit systems)
		
		    ◆ `\Program Files (x86)\platespin` (64-bit systems)
		
		3. Run the following command:
		
		    `ofxcontroller.exe /uninstall`
		
		4. Remove the `platespin*` directory

B.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	◆ Kill these processes:
		
		    ◆ `pkill -9 ofxcontrollerd`
		
		    ◆ `pkill -9 ofxjobexec`
		
		◆ remove the OFX controller rpm package:
		
		    `rpm -e ofxcontrollerd`
		
		◆ 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	1. Check if the driver is active:
		
		`lsmod | grep blkwatch`
		
		If the driver is still loaded in memory, the result should contain a line, similar to the following:
		
		`blkwatch_7616  70924  0`
		
		2. (Conditional) If the driver is still loaded, remove it from memory:
		
		`rmmod blkwatch_7616`
		
		3. Remove the driver from the boot sequence:
		
		`blkconfig -u`
		
		4. Remove the driver files by deleting the following directory with its contents:
		
		`/lib/modules/`*`[Kernel_Version]`*`/Platespin`
		
		5. Delete the following file:
		
		`/etc/blkwatch.conf`

LVM snapshots	Block-level migrations using LVM snapshots	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: `lvremove` *`snapshot_name`*

C Reference of Tables

This section provides a list of links to important tables used throughout the documentation accompanying PlateSpin Migrate version 9.1.

Table
Table 2-3Table 2-3, "PlateSpin Migrate Roles and Permission Details," on page 24
Table 2-4, "Network Communication Prerequisites for Discovery Operations," on page 26
Table 2-5, "Network Communication Prerequisites for Workload Portability," on page 28
Table 1-2, "PlateSpin Migrate Workload Migration Tasks," on page 10
Table 1-3, "Supported Microsoft Windows Workloads," on page 11
Table 1-4, "Supported Linux Workloads," on page 13
Table 1-5, "Supported Target Virtualization Platforms," on page 14
Table 3-1, "Machine-Specific Icons in the Servers View," on page 44
Table 3-2, "Boot ISO Images for Target Physical Machines," on page 50

Glossary

Auto-Discovery. *See* Network Discovery.

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.

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.

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.

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 Image Server. A machine on which you have installed the PlateSpin Image Server software. *See* PlateSpin Image.

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.

OFX Controller. Operations Framework Controller. Software component that PlateSpin Migrate installs on source and target machines to enable them to communicate with PlateSpin Migrate Server and with each other.

Offline Migration. A type of migration, during which the source is shut down and booted into a temporary pre-execution environment. For Windows workloads, PlateSpin Migrate uses Microsoft WinPE; for Linux, it uses a RAM disk.

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.

PlateSpin Migrate Client. The client application you use to interact with the PlateSpin Migrate Server, discover source workloads and targets; set up, execute, and monitor jobs; manage license keys; and configure the default behavior of the server.

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 Migrate Server. The remote runtime execution engine underlying the PlateSpin Migrate product.

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 Migrate 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.

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.

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.

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.