User Guide

Novell®
PlateSpin® Portability Suite

8.1

August 3, 2009

www.novell.com

Novell Trademarks

For Novell trademarks, see the Novell Trademark and Service Mark list (http://www.novell.com/company/legal/trademarks/tmlist.html).

Third-Party Materials

All third-party trademarks are the property of their respective owners.

Contents

About This Guide

This guide provides information about using PlateSpin® Portability Suite version 8.1 (PlateSpin® Migrate and PlateSpin® Protect products).

Audience

This guide is intended for IT staff, such as data center administrators and operators, who use Portability Suite in their ongoing workload migration and protection 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).

Documentation Updates

For the most recent version of this text, visit the PlateSpin® Portability Suite version 8.1 Web site (http://www.novell.com/documentation/platespin_portabilitysuite_810/).

Additional Documentation

This guide is part of the PlateSpin® Portability Suite documentation set. Below is a complete list of the set:

Portability Suite 8.1 Release Notes	Provides late-breaking information, as well as information about known issues and suggested workarounds.
Portability Suite 8.1 Installation Guide	Provides information about installing, upgrading, and uninstalling the product.
Portability Suite 8.1 Configuration Guide	Provides information about licensing your product and managing license keys, setting up your network environment, and configuring your product's default settings.
Portability Suite 8.1 User Guide	Provides information about using the product in your workload migration and protection projects.

Portability Suite 8.1 Integrated Help	All product documentation in compiled HTML help (CHM) format, integrated with the product.
Portability Suite 8.1 Online Documentation	All product documentation in WebHelp format for browser-based access over the Internet. This distribution should contain the latest updates at all times.

Additional Resources

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

- PlateSpin User Forum (http://forum.platespin.com): A Web-based community with a variety of discussion topics.
- PlateSpin Knowledge Base (http://support.platespin.com/kb2/): 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 visit the PlateSpin Technical Support Web site (http://www.platespin.com/support/).

Documentation Conventions

In Novell documentation, a greater-than symbol (>) is used to separate actions within a step and items in a cross-reference path.

A trademark symbol ($^{®}$, ™, etc.) denotes a Novell trademark. An asterisk (*) denotes a third-party trademark.

Portability Suite Feature Overview 1

This section provides information about PlateSpin® Portability Suite software, which comprises the PlateSpin® Migrate and PlateSpin® Protect products.

Both products are based on a common workload portability and protection engine (Portability Suite Server) and use the same client (Portability Suite Client).

Portability Suite products enable you to migrate and protect 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.

1.1 About Workload Portability and Protection

Portability Suite automates the migration of workloads among three infrastructures: physical, virtual machine, and volume imaging. On virtual machine and imaging infrastructures, Portability Suite provides protection capabilities for business continuity.

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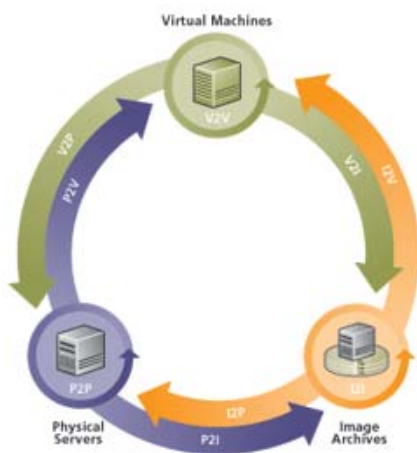
Table 1-1   *Workload Portability and Protection Operations*

Category of Operation	Conversion Infrastructures
Workload Portability: Peer-to-peer	◆ Physical to Virtual (P2V) ◆ Virtual to Virtual (V2V) ◆ Virtual to Physical (V2P) ◆ Physical to Physical (P2P)
Workload Portability: Imaging	◆ Physical to Image (P2I) ◆ Virtual to Image (V2I) ◆ Image to Virtual (I2V) ◆ Image to Physical (I2P)
Workload Protection	◆ Physical to Virtual (P2V) ◆ Virtual to Virtual (V2V) ◆ Physical to Image (P2I) ◆ Virtual to Image (V2I)

Portability Suite supports multiple workload types and virtualization platforms. Protection and imaging are 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 13.

1.1.1 Business Applications for Portability Suite

Portability Suite 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.

- **Disaster Recovery.** Managing planned/unplanned outages through automated replication of physical or virtual production workloads to virtual machines. You can also switch to a virtual machine replica while the physical production workload is unavailable. This reduces downtime and hardware resource utilization, while maintaining business continuity.

- **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 Flexible 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 Conversion Tasks

Portability Suite enables you to define, save, schedule, execute, and monitor the following conversion tasks.

Table 1-2 *Portability Suite Workload Conversion Tasks*

Conversion	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 conversion 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 conversion task when you intend to retire or repurpose the original infrastructure.
Protect Workload	Results in a standby virtual duplicate of a selected physical or virtual workload under a protection contract (a regular, automatic schedule for synchronizing the duplicate with incremental changes on the source).
Server Sync	Synchronizes a virtual or physical workload with another virtual or physical workload without transferring the entire source volume data over the network. You can synchronize two workloads either for a one-time portability task (similar to the Copy or Move tasks) or for ongoing protection (similar to the Protect Workload task).
Capture Image	Creates an image of a physical or virtual workload as a single entity, in Flexible Image format. You can also synchronize the state of an image with the image of its source workload on demand or automatically according to a predefined schedule.
Import Image	Imports third-party volume archives or raw volume data into a Flexible Image.
Deploy Image	Converts a Flexible 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

Portability Suite supports the following operating system families:

- Microsoft Windows
- Linux*
- UNIX*

Not all Portability Suite features and transfer methods are supported for all source operating systems. For version specifics and the most up-to-date list, see Knowledge Base Article Q20002 (http://support.platespin.com/kb2/article.aspx?id=20002).

- "Supported Microsoft Windows Workloads" on page 14
- "Supported Linux Workloads" on page 15
- "Supported UNIX Workloads" on page 15

Supported Microsoft Windows Workloads

For the Windows platform, Portability Suite supports the following portability and protection features:

- Workload portability: Peer-to-peer conversions (P2V, V2V, V2P, P2P).
- Workload protection: Peer-to-peer conversion (P2V and V2V) with incremental synchronization of changes. Covers all supported Windows workloads except Windows NT 4. Supported target infrastructures for protection are VMware* ESX Servers, versions 3.0 and later.
- Peer-to-peer workload synchronization with Server Sync (P2V, V2V, P2P, V2P).
- Workload imaging with Flexible Images (X2I, I2X), with full or partial support for several third-party backup and archiving applications.
- Workload image protection through incremental image synchronization (X2I).

Table 1-3 *Supported Microsoft Windows Workloads*

Version	Bit Spec	Remarks
Windows Vista*, Business/Enterprise/ Ultimate, SP1 and later	32-bit, 64-bit	Not supported on Microsoft Virtual Server (MSVS), VMware Server, and VMware ESX versions earlier than 3.5 Update 2.
Windows Server* 2008, Standard/ Enterprise	32-bit, 64-bit	
Windows Server 2003, Standard/ Enterprise	32-bit, 64-bit	
Windows Server* 2000, Windows Advanced Server 2000	32-bit	
Windows XP Professional, x86 only	32-bit, 64-bit	
Windows NT* 4.0 Server	32-bit	

Portability Suite supports the following localized versions of source Windows workloads:

- English
- French
- German
- Japanese

Supported Linux Workloads

For the Linux platform, Portability Suite provides only peer-to-peer workload portability support (P2P, P2V, V2P, V2V), including workload synchronization with Server Sync.

Supported Linux file systems are EXT2, EXT3, and REISERFS.

Workload protection and imaging are not supported for Linux.

Table 1-4   *Supported Linux Workloads*

Version	Bit Spec
Red Hat* Linux 7.3, 8	32-bit
Red Hat Enterprise Linux AS/ES/WS 2.1	32-bit
Red Hat Enterprise Linux AS/ES/WS 3, 4, 5	32-bit, 64-bit
SUSE® Linux Enterprise Server (SLES) 8, 9, 10, 11	32-bit, 64-bit

NOTE:

- Linux-based source workloads must be running a Secure Shell (SSH) server.
- Customized kernels or kernels rebuilt from source are not supported. For more information about supported Linux kernels, see Knowledge Base Article Q20947 (http://support.platespin.com/kb2/Article.aspx?id=20947).

Supported UNIX Workloads

Portability Suite 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 protection and imaging are 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 version details, as well as the most up-to-date list, see Knowledge Base Article Q20002 (http://support.platespin.com/kb2/article.aspx?id=20002).

Table 1-5 *Supported Target Virtualization Platforms*

Platform	Remarks
VMware ESX 4	
VMware ESX Server 3.5; ESXi 3.5	
VMware Virtual Infrastructure version 3.0 - 3.0.1; 3.5	
VMware ESX Server 3.0.x	
VMware ESX Server 2.5.x	
VMware ESX Server 2.1.1	
VMware Server 1.0	Windows only (x86 and x64)
VMware GSX Server 3.1; 3.2 (Windows only)	Windows NT 4.0 SP6a, 2000, 2003 source workloads only
Microsoft Virtual Server (MSVS) 2005 (build 465; build 465.254 with Service Pack 1)	Windows NT 4.0 SP6a, 2000, 2003 source workloads only (x86 and x64)
Solaris 10 (U5 and U6) zone server	For source workloads that are Solaris 10 (U5 and U6) physical machines or zones
Citrix* XenServer 4.1, 5	
Novell® Xen* Hypervisor on SLES 10 SP2	Supported through semi-automated virtualization using the X2P conversion workflow. See "Semi-Automated Workload Virtualization through X2P" on page 55.
Microsoft Hyper-V	
Virtual Iron*	

1.2.3 Supported Third-Party Image Formats

The following are supported third-party image formats:

 * Acronis* True Image (True Image 9.1, True Image Server 8.0, True Image Enterprise Server 8.0, 9.1).

 * Symantec* LiveState* Recovery (LiveState Recovery Standard Server, LiveState Recovery Advanced Server, LiveState Recovery Advanced Server Suite, versions 3.0 and 6.0).

1.2.4 Hardware Compatibility

A hardware compatibility list for Windows source workloads and Take Control ISO images is available. See Knowledge Base Article Q20314 (http://support.platespin.com/kb2/article.aspx?id=20314).

A list is also available for Linux source workloads. See Knowledge Base Article Q20570 (http://support.platespin.com/kb2/article.aspx?id=20570).

1.3 Supported Transfer Methods

Depending on the selected workload and the conversion type, Portability Suite enables you to select different methods for transferring workload data from the source to the target.

For a list of workload types and conversions arranged by supported transfer method, see Knowledge Base Article Q20002 (http://support.platespin.com/kb2/article.aspx?id=20002).

For information on how to select a transfer method, see "Specifying Transfer Options" on page 79.

For information on how to fine-tune data transfer performance, see "Fine-Tuning Data Transfer Performance" in your *Configuration Guide*.

1.3.1 Offline Transfer with the Take Control Mechanism

Offline Take Control Transfer of Windows and Linux Workloads

This method enables Portability Suite 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 Take Control transfer method depends on the source workload's operating system:

- For Windows workloads, Portability Suite uses the Microsoft Windows Preinstallation Environment (WinPE).
- For Linux workloads, Portability Suite uses a Linux Ramdisk.

To ensure that the source operating system loads the temporary pre-execution environment on reboot, Portability Suite temporarily modifies its boot files and restores them to their original state after the Take Control environment has successfully loaded. For a list of files being changed on the source, see Knowledge Base Article Q20349 (http://support.platespin.com/kb2/article.aspx?id=20349).

Use Take Control transfer to migrate Windows NT 4.0, Windows Server 2000, Windows Server 2003, and Linux workloads, or in other cases where you can afford source workload downtime.

Offline Take Control Transfer of Solaris Workloads

For Solaris workloads, Portability Suite uses a different Take Control 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, Portability Suite 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.

1.3.2 Live Transfer (File-Based)

The File-based Live Transfer method copies data and replicates changes at the file level. During File-based Live Transfer, Portability Suite transfers all files from the source volumes while monitoring them for changes. When the transfer is complete, files that have changed are resent. If present, Microsoft SQL Server* or Microsoft Exchange Server services are stopped and their respective database files are transferred to the target.

You can configure your job to stop these services when you are using the File-based Live Transfer method (see "Handling Services During Live Transfer (Windows Source Workloads)" on page 95). However, if there are other tools present that manage the backup of these databases, consider leaving services running during the transfer. When the transfer completes, verify that the copied database is the up-to-date.

If file system changes are constant, data transfer is stopped and a job progress warning is shown.

When the initial transfer is complete for a workload protection job, the target is powered off and is powered on again during the next scheduled incremental transfer.

File-level Live Transfer is appropriate for moderately active Windows.

1.3.3 Live Transfer (VSS Block-Based)

This Live Transfer method transfers data at the block level and leverages the Microsoft Volume Snapshot Service (VSS), for Windows workloads (Windows 2003 SP1 and later) with applications and services that support VSS. The VSS Block-based Live Transfer method offers an exact point-in-time copy of the source workload.

During VSS Block-Based Live Transfer, Portability Suite takes a VSS snapshot of the volumes on the source machine and transfers any changed data block by block.

The source workload remains online throughout the transfer except for protection jobs, during which the source requires a single reboot for the initial transfer (if the component is not pre-installed).

The VSS Block-Based Live Transfer method is the preferred data transfer method for Windows workloads that support the VSS feature. Database servers, mail servers, and application servers that would otherwise require a temporary service stoppage can be protected by using this Live Transfer method. This method is also recommended for workload protection jobs in networks with high latency.

1.3.4 Live Transfer (Block-Based)

The Block-based Live Transfer method copies data and replicates changes at the block level instead of replicating an entire file.

During data transfer to the target, changes on the source volumes are monitored and continuously retransferred to the target at the block level until full synchronization is achieved.

Because the Block-based Live Transfer method transmits only changed blocks rather than entire files, it transfers significantly less data.

When the initial transfer is complete for a workload protection job, the target is booted into the Take Control pre-execution environment and awaits the next scheduled incremental transfer.

Use the Block-based Live Transfer method to reduce the service downtime when you are converting Windows workloads that do not support the Microsoft Volume Snapshot Service (VSS) feature. Using the Block-based Live Transfer method, you can protect critical database servers, mail servers, and application servers with large databases and with high disk activity. In addition, the Block-based Live Transfer method is recommended for networks with high latency because the size of block-level changes is significantly smaller than the size of file-level changes; when file-level changes are detected during file-level data transfer, the changed files are transferred in their entirety.

If your source workload is running Microsoft Exchange Server 2000 and 2003, and Microsoft SQL Server 2000, Portability Suite automatically detects the Windows services of these applications. You can configure your job to stop these services when you are using the Block-based Live Transfer method. However, if there are other tools present that manage the backup of these databases, consider leaving services running during the transfer. When the transfer completes, verify that the copied database is up-to-date.

1.3.5 Live Transfer (VSS File-Based)

This Live Transfer method transfers data at the file level and uses the Microsoft Volume Snapshot Service (VSS) feature, also known as Shadow Copy, for Windows workloads (Windows 2003 SP1 and later) with applications and services that support VSS. The VSS File-based Live Transfer method offers an exact point-in-time copy of the source workload.

During VSS File-based Live Transfer, Portability Suite takes a VSS snapshot of the volumes on the source machine and transfers the data file-by-file.

When the initial transfer is complete for a workload protection job, the target is powered off; it powers on again during the next scheduled incremental transfer.

Use the VSS File-based Live Transfer method to reduce service downtime when you are converting Windows workloads that support the VSS feature. Database servers, mail servers, and application servers that would otherwise require extended service stoppage can be protected by using this Live Transfer method. This method is also recommended for workload protection jobs in networks with high latency. Because this is a point-in-time solution, data does not need to be retransmitted as it does with other methods.

1.3.6 Installing Live Transfer Components

During the execution of a migration job the appropriate Live Transfer components are automatically installed on source and target workloads.

For block-based transfer components (the Block-based Transfer Component and the VSS Block-based Transfer Component) you have the additional option to manually install or uninstall a component on the appropriate source workload. This enables you to:

◆ Minimize the impact on the operational continuity of your source workloads. See "Continuity of Source Workloads Impacted by Transfer Method Selection" on page 20.

◆ Upgrade a component (for example, one distributed in a hotfix) without reinstalling your Portability Suite product.

The following topics provide more details:

Checking Whether the Block-Based Component is Up to Date

If the Block-based Transfer Component on a source workload is obsolete, selecting the Block-level Live Transfer method for a conversion job causes Portability Suite to display a warning.

A warning message is also shown for existing synchronization schedules in the Synchronization Schedules window, so you can choose to upgrade the block-level transfer component on the source by right-clicking the applicable schedule and selecting *Upgrade Block-based Components*.

To check if the Block-based Transfer Component on a source is current:

1 In the *Servers* view, right-click the source. If the menu includes *Upgrade Block Based Component*, the Block-based Transfer Component is obsolete.

This applies to the Block-based Transfer Component only. The VSS Block-based Transfer Component is new in version 8.1.

Manually Installing and Uninstalling Block-Based Components

You can manually install and uninstall both block-based components included with Portability Suite. Use this option to eliminate or reduce the impact of component installation on the operational continuity of your source workloads. See "Continuity of Source Workloads Impacted by Transfer Method Selection" on page 20.

To manually install or uninstall Block-based Transfer Components:

1 In the Servers view, right-click the required source and select the required action:
   ◆ *Install Block Based Component* or *Install VSS Block Based Component*: To install either component.
   ◆ *Uninstall Block Based Component* or *Uninstall Block Based Component*: To uninstall either component.
2 Provide valid credentials for the source machine, then click *Uninstall*.

1.4 Workload Continuity

1.4.1 Continuity of Source Workloads Impacted by Transfer Method Selection

The operational continuity of your source workload depends on the selected transfer method.

The following table provides information about the different transfer methods along with their impact on the operational continuity of your source workloads.

Table 1-6   *Continuity Impact by Transfer Method*

Transfer Method	Continuity Impact on Source Workload
Take Control	The source is offline during the entire migration.
Block-Based	A single reboot is required at the beginning of the conversion job's execution, when the Block-based Transfer Component is installed on the source. On upgrading the component, existing protection contracts are suspended.
VSS Block-Based	No impact on migration jobs. Impact on protection jobs: A single reboot is required at the beginning of the job's execution, when the VSS Block-based Transfer Component is installed on the source.
File-Based; VSS File-Based	No impact on migration jobs.

1.4.2 Default Post-Conversion End States of Source and Target Workloads

The following table lists the default post-conversion end states of source and target workloads by conversion task. You can also manually specify a post-conversion end state for your source and target. See .

Table 1-7   *Default Post-Conversion End States*

Conversion Task	Transfer Method	Source State	Target State
Move	Take Control	• Shut down • Leave under control, then shut down after transfer, if transfer is OK • Leave under Control, then restart after transfer, if the transfer is not OK	• Power On • Power Off • Suspend (only for target VMs except Solaris zones)
Move	Live Transfer	• Shut down	• Power On • Power Off • Suspend (only for target VMs except Solaris zones)
Copy	Take Control	• Restart • Shut down • Leave under control to delay the restart decision (unlikely, no advantage)	• Power On • Power Off • Suspend (only for target VMs except Solaris zones)

Conversion Task	Transfer Method	Source State	Target State
Copy	Live Transfer	◆ Continue (shutdown not selected) ◆ Shut down	◆ Power On ◆ Power Off ◆ Suspend (only for target VMs except Solaris zones)

1.5 Security and Privacy

Portability Suite provides several features to help you safeguard your data and increase security.

1.5.1 Security of the Workload Data Transfer

To make the transfer of your workload data more secure, Portability Suite allows you to:

- Enable data transfer encryption for an individual conversion job. When encryption is enabled, over-the-network data transfer from the source to the target is encrypted by using AES (Advanced Encryption Standard).
- Configure your Portability Suite 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 Portability Suite 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 conversion job, see "Specifying Transfer Options" on page 79.

1.5.2 Security of Saved Credentials

Portability Suite products cache credentials in the Portability Suite Client user interface and securely store them by using operating system APIs. In addition, credentials are stored on the Portability Suite Server and the Portability Suite database, as well as embedded in conversion jobs.

Data transmission between the Portability Suite Server and the Portability Suite 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 Portability Suite Server host and use HTTPS when specifying the server URL. See "Connecting to a Portability Suite Server" on page 25.

1.5.3 User Authorization and Authentication

Portability Suite provides a role-based user authorization and authentication mechanism. See "Setting Up User Authorization and Authentication" in your *Configuration Guide*.

1.6 Performance and Scalability

1.6.1 Performance

The performance characteristics of your Portability Suite product depend on a number of factors, including:

- Hardware and software profiles of your source and target
- Hardware and software profiles of your Portability Suite Server host
- The specifics of your network bandwidth, configuration, and conditions
- The number of your source workload volumes and their sizes
- The number of files on your source workload volumes (file density)
- Source I/O levels (how busy your workload is)
- 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 conversion job goes through, as applicable to your project:

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

1.6.2 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 conversions to a single VM server: no more than 20
- Multiple conversions against multiple VM servers (ESX): no more than 30

Getting Started with Portability Suite

2

This section introduces the main features of the PlateSpin® Portability Suite products (PlateSpin® Migrate and PlateSpin® Protect) and how you use the Portability Suite Client to interact with either product.

Portability Suite Client connects to a specified Portability Suite Server (the workload portability and protection engine underlying both products of the suite) and provides access to information in the Portability Suite 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.

- Section 2.1, "Connecting to a Portability Suite Server," on page 25
- Section 2.2, "About the Portability Suite Client User Interface," on page 26
- Section 2.3, "Discovering Source Workloads and Migration Targets," on page 32
- Section 2.4, "Analyzing Machines with PlateSpin Analyzer," on page 38
- Section 2.5, "Working with Device Drivers," on page 39
- Section 2.6, "Managing Custom Post-conversion Actions," on page 40
- Section 2.7, "Setting Up, Executing, and Managing Jobs," on page 42

2.1 Connecting to a Portability Suite Server

Every time you start Portability Suite Client, it performs the following actions:

- Obtains authentication from Portability Suite Server for your user account. See "Setting Up User Authorization and Authentication" in your *Configuration Guide*.
- Connects to a specified Portability Suite Server.
- Loads a specified Portability Suite Network, which is a collection of discovered source workloads and targets that you work with at one time.

You specify your connection credentials, the Portability Suite Server instance, and the required Portability Suite Network in the Portability Suite Server settings.

1 In the Portability Suite Client, click *Tools > Portability Suite Server Settings*.

or

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

| Server: http://dev3/PortabilitySuite | Network: Sample Environment | User: devdomain\john.smith |
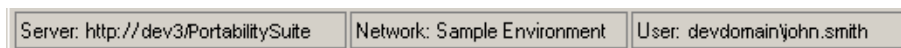
The Portability Suite Server Settings dialog box opens.
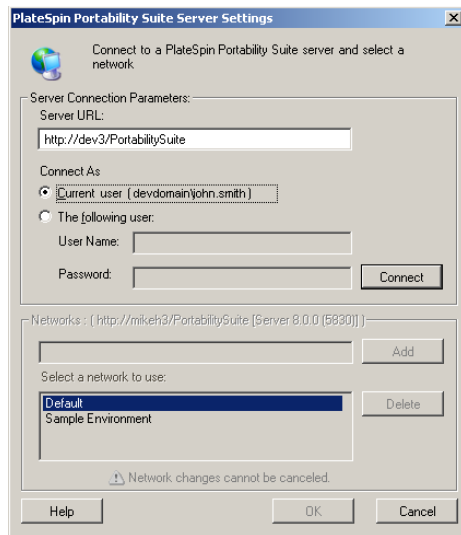
2 Specify the required Portability Suite Server, user, and network parameters as required:

Interface Element	Description
Server URL	Type the Portability Suite Server URL in the following format: `http://`**server_host**`/PortabilitySuite`
Connect As	To connect to a Portability Suite Server, you must have administrative access to the Portability Suite Server host or be a member of one of the Portability Suite roles. See "Setting Up User Authorization and Authentication" in your *Configuration Guide*.
Networks	To familiarize yourself with Portability Suite 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*.

2.2 About the Portability Suite Client User Interface

The Portability Suite 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 conversion 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 Portability Suite Client window, the status bar displays the Portability Suite Server that the client is currently connected to, the Portability Suite 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 2.2.1, "Servers View," on page 27
- Section 2.2.2, "Jobs View," on page 30
- Section 2.2.3, "Tasks Pane," on page 31
- Section 2.2.4, "Status Bar," on page 31

2.2.1 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 2-1 *Portability Suite Client's Servers View*
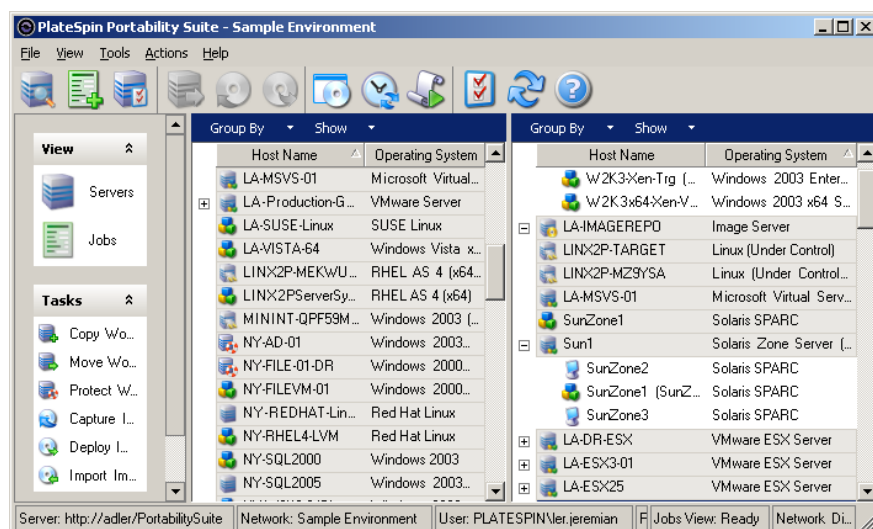
The hierarchical display of items in the Servers view reflects the organization of items on their respective platforms; for example: VMs are shown nested beneath their VM hosts, PlateSpin® Flexible Images are beneath their image servers, and Solaris zones (regarded as VMs) are beneath their respective 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 28.

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 29.

- "Organizing the Servers View" on page 28
- "Viewing the Properties of Source Workloads and Targets" on page 29
- "List of Machine-Specific Icons in the Servers View" on page 30

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 2-2   *Servers View Options for Sorting Items by Type*
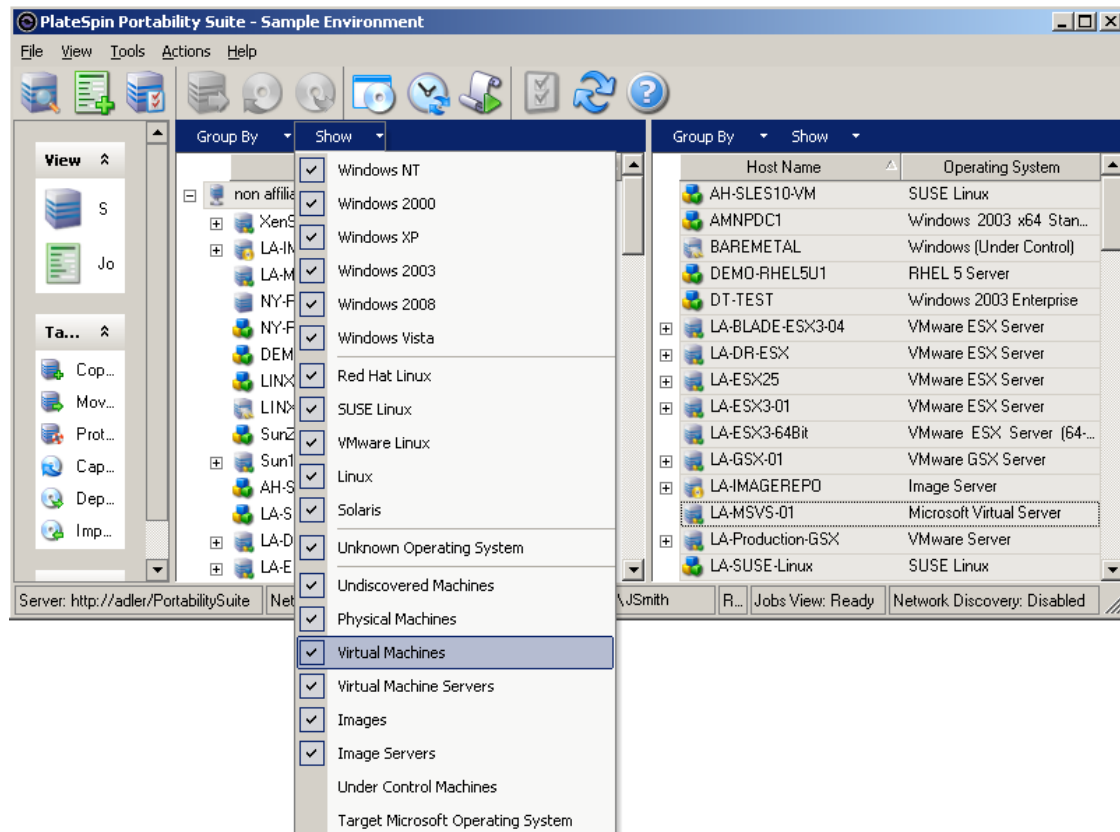
Viewing the Properties of Source Workloads and Targets

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

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

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

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

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

VMware ESX servers that are assigned to a Distributed Resource Scheduler (DRS) cluster provide information about the name of the cluster and the DRS automation level (full, manual, or partially automated).

The properties for VMware ESX servers that are part of VMware vCenter platforms also indicate this.

The following figure shows the properties of a discovered Solaris zone server.

Figure 2-3   Solaris Zone Server-Specific Information in the System's Properties
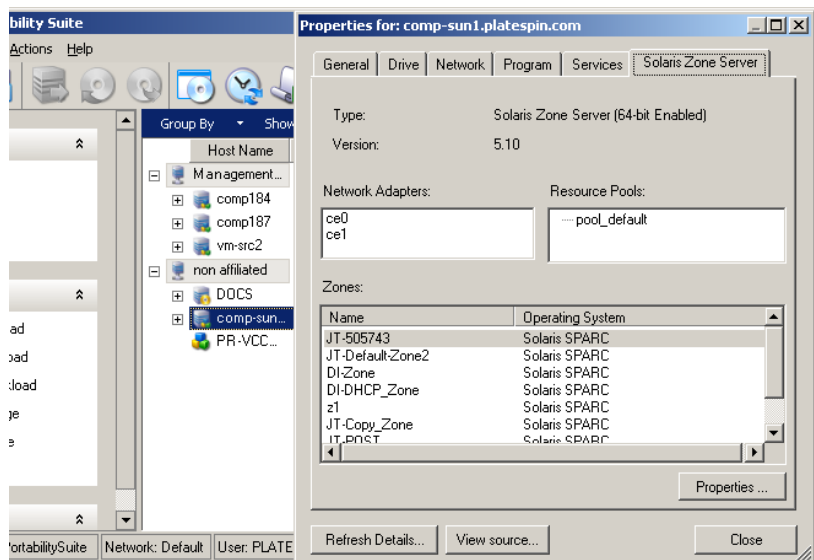
List of Machine-Specific Icons in the Servers View

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

Table 2-1 *Machine-Specific Icons in the Servers View*

Physical machine

Physical machine in Take Control pre-execution environment

Physical machine with workload license

Protected physical machine

Protected 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

Flexible Image Server

Flexible Image

Protected Flexible Image

2.2.2 Jobs View

The Jobs view displays all discovery, conversion, and image capture jobs run by the Portability Suite
Server. Jobs can be filtered based on their current state.

Figure 2-4   *Portability Suite Client's Jobs View*
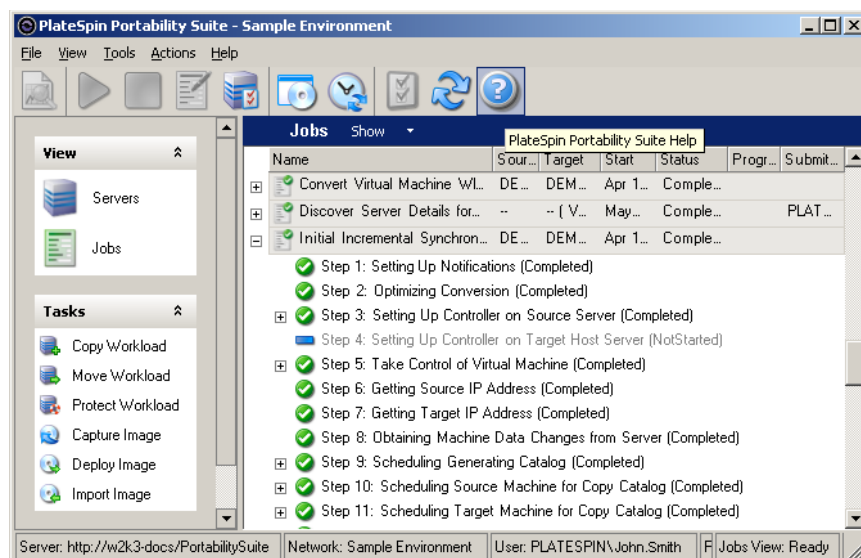
You can use the *Show* drop-down menu to filter the Jobs view to display specific sets of jobs: Recent Jobs Only (past 24 hours), Conversion Jobs Only, Failed, Running, Completed, and/or Saved jobs. When a job is selected from the list, the lower left pane of the screen displays a summary of the job details.

2.2.3 Tasks Pane

The Tasks pane of the Portability Suite Client window contains most essential conversion actions. Clicking a task opens the Action window, which you can use to select the conversion source, target, and setup method. See Section 2.7.1, "Selecting a Conversion Job Setup Method," on page 42.

2.2.4 Status Bar

The status bar of the Portability Suite Client window displays information about:

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

Figure 2-5   *Status Bar of the Portability Suite Client Window*
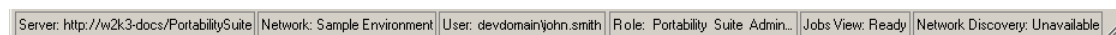
Double-clicking any of the first three status items opens the Portability Suite Server Settings window. See "Connecting to a Portability Suite Server" on page 25.

Double-clicking the Network Discovery status item turns Automatic Windows Network Discovery on or off. See "Discovering and Preparing a Windows VM for a Server Sync Job" on page 38.

2.3 Discovering Source Workloads and Migration Targets

For any conversion job, you must have a discovered source and a discovered target. A discovery operation populates the Portability Suite 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 conversion job.

- Section 2.3.1, "Prepopulating the Servers View with Windows Computers in the Network," on page 32
- Section 2.3.2, "Discovering the Details of Source Workloads and Targets," on page 32
- Section 2.3.3, "Discovering All Machines in a Windows Domain," on page 34
- Section 2.3.4, "Refreshing Machine Details," on page 34
- Section 2.3.5, "Resetting the Access Mechanism Used to Authenticate Access to ESX Hosts," on page 35
- Section 2.3.6, "Discovering and Registering Target Physical Machines," on page 35
- Section 2.3.7, "Discovering and Preparing a Windows VM for a Server Sync Job," on page 38

2.3.1 Prepopulating the Servers View with Windows Computers in the Network

Portability Suite's Network Discovery feature automatically populates Portability Suite 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 and protection 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 34.

Portability Suite uses the standard Windows network browser function for the auto-discovery feature. Because Linux-based workloads and virtual machine servers do not advertise to the Windows network browser, these 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 Portability Suite Client window.

2.3.2 Discovering the Details of Source Workloads and Targets

You can discover an individual physical or virtual machine, multiple machines, an individual virtual machine server, or multiple virtual machine servers. You can also discover all VMware ESX hosts affiliated with a VMware vCenter Server.

Before starting discovery operations, make sure Portability Suite Server can communicate with your source workloads and targets. See "Network Communication Prerequisites" in your *Configuration Guide*.

1 On the Portability Suite toolbar, click *Discover Details*.

or

In the Servers view, right-click in a blank area, then select *Discover Details*.
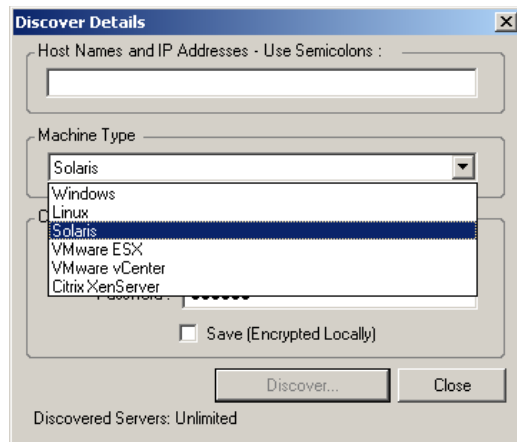
2 In the *Discover Details* dialog box, type the hostname or IP address of the source or target. To discover multiple machines, specify multiple hostnames or IP addresses separated by semicolons.

To discover Linux workloads or VMware ESX Server 2.5 and 3.0.x systems that are using non-default SSH ports, specify the port number using the standard URL syntax (*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.

3 Select the machine type and provide administrator credentials for the machine you are discovering. Use the following guidelines for selecting the machine type and specifying credentials:

To Discover	Machine Type Selection	Credentials	Remarks
All Windows workloads MSVS VM hosts Image Servers	*Windows*	Local or domain admin credentials.	For the username, use this syntax: ◆ For domain member machines: *authority\principal* ◆ For workgroup member machines: *hostname\principal*
All Linux workloads	*Linux*	Root-level username and password	Non-root accounts must be properly configured to use `sudo`. See Knowledge Base Article 20711 (http://support.platespin.com/kb2/Article.aspx?id=20711).
VMware ESX hosts	*VMware ESX*	Root-level username and password	

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)	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 Access Mechanism Used to Authenticate Access to ESX Hosts" on page 35.
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 Knowledge Base Article 20711 (http://support.platespin.com/kb2/Article.aspx?id=20711).

NOTE: Discovering hosts with Microsoft Hyper-V, Virtual Iron, and Xen Hypervisor virtualization software results in these systems being registered as Portability Suite source workloads (as opposed to VM host targets). To use these platforms as workload portability targets, see "Semi-Automated Workload Virtualization through X2P" on page 55.

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.

2.3.3 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 32.

2 In Portability Suite 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.

2.3.4 Refreshing Machine Details

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

To refresh a source or a target machine:

1 In the Servers view, right-click the required item, then select *Refresh Details*.
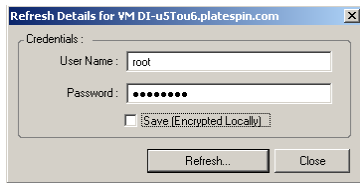
2 Specify the credentials appropriate for the system being refreshed, then click *Refresh*.

Portability Suite starts a discovery job, which you can monitor in the Jobs view.

2.3.5 Resetting the Access 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 .

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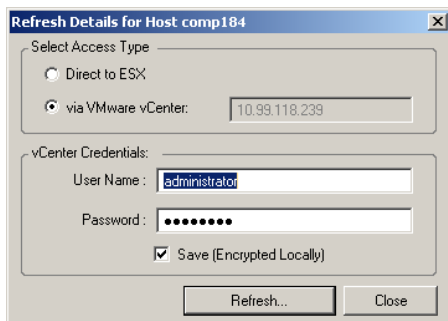
2 Select the required option, then click *Refresh*.

Portability Suite rediscovers the required ESX host by using the specified access mechanism, which it uses for all subsequent communications with that host.

2.3.6 Discovering and Registering Target Physical Machines

To discover a physical target and inventory its hardware components, you must boot the target machine with a Physical Target Take Control 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 through X2P" on page 55.

To use an ISO image, download it from the Novell Customer Center (http://www.novell.com/customercenter/). Use the image appropriate to your target machine:

Table 2-2 *Take Control ISO Images for Target Physical Machines*

Target Type	Filename	
Windows machines with 384 MB RAM or more	`winperamdisk.iso`	
Windows machines with 256 to 384 MB RAM	`winpe.iso`	
Linux machines	`bootofxx2p.iso`	
Windows VMs on Virtual Iron with 384 MB RAM or more or more	`winperamdisk_virtualiron.iso`	Regarded by Portability Suite Server as physical targets. See "Semi-Automated Workload Virtualization through X2P" on page 55.
Windows VMs on Virtual Iron with 256 to 384 MB RAM	`winpe_virtualiron.iso`	
Unisys* ES7000 machines regardless of memory size	`winperamdisk_es7000.iso`	

- "Discovering and Registering Target Physical Machines (Windows)" on page 36
- "Registering Physical Targets with Portability Suite Server (Linux)" on page 37
- "Performing an Unattended Registration of a Target Physical Machine" on page 37

Discovering and Registering Target Physical Machines (Windows)

Complete these steps to register your target with Portability Suite Server when you are using an unmodified Take Control ISO image. For information on how to modify the image for an unattended registration process, see Knowledge Base Article 20832 (http://support.platespin.com/kb2/article.aspx?id=20832).

All discovered and registered physical machines can be also used as Server Sync targets. See "Synchronizing Workloads with Server Sync" on page 53.

1 Burn the appropriate Physical Target Take Control ISO image on a CD or save it to the required media, from which your target can boot. See Table 2-2, "Take Control ISO Images for Target Physical Machines," on page 36.

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 CD, boot the target physical machine and wait for the command prompt window to open.

4 Enter `Y` to transfer data at the block level or enter `N` to transfer data at the file level. The block-level transfer selection applies to the Block-based Live Transfer method (not the VSS Block-based Live Transfer method). See "Supported Transfer Methods" on page 16.

5 When prompted, enter the Portability Suite Server URL, using the following syntax:

`http://`*server_host*`/PortabilitySuite`

Replace *server_host* with the actual Portability Suite Server host's name or IP address.

6 Enter administrator-level credentials for the Portability Suite Server, specifying an authority. Use this syntax:

domain\username or *hostname\username*

7 Enter the password for the provided credentials.

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

8 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.

   8a Enter the IP address

   8b Enter the subnet mask

   8c Enter the gateway

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

10 Enter *Yes* if you have enabled SSL on the Portability Suite Server; otherwise enter *No*.

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

After a few moments, Portability Suite displays the physical target in the Servers view.

Registering Physical Targets with Portability Suite Server (Linux)

Complete these steps to register your target with Portability Suite Server when you are using an unmodified Take Control ISO image. For information on how to modify the image for an unattended registration process, see Knowledge Base Article 20832 (http://support.platespin.com/kb2/article.aspx?id=20832).

1 Burn the appropriate Physical Target Take Control ISO image on a CD or save it to the required media, from which your target can boot. See Table 2-2, "Take Control ISO Images for Target Physical Machines," on page 36.

2 Boot the target machine, using the ISO Ramdisk image.

3 When prompted, enter the Portability Suite Server URL, using the following syntax:

`http://`*server_host*`/PortabilitySuite`

Replace *server_host* with the actual Portability Suite Server host's name or IP address.

4 Enter your credentials for the Portability Suite Server.

5 Specify a static IP address or indicate that the machine should dynamically obtain an IP address from a DHCP server.

6 Enter a name for the target host and specify an existing Portability Suite Network to work with.

After a few moments, Portability Suite displays the physical target in the Servers view.

Performing an Unattended Registration of a Target Physical Machine

Portability Suite provides a mechanism for automating a target physical machine's registration with the Portability Suite Server. This involves updating the Take Control ISO image with specific registration information before booting the target.

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

2.3.7 Discovering and Preparing a Windows VM for a Server Sync Job

If you want to synchronize two Windows 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 53.

1 On your virtual machine server, create a virtual machine with an operating system profile matching the intended source workload.

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:

2.4 Analyzing Machines with PlateSpin Analyzer

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

To open PlateSpin® Analyzer:

1 On the *Tools* menu, click *Analyze Servers*.

The PlateSpin Analyzer window opens.

2 Select the required Portability Suite 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 Portability Suite Server, click *Connect*, then specify the required Portability Suite 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 2-3 *Status Messages in PlateSpin Analyzer Test Results*

Result	Description
Passed	The machine passed the PlateSpin® Analyzer tests.

Result	Description
Warning	One or more tests returned warnings for the machine, indicating potential conversion 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 2-4 *PlateSpin Analyzer Test Results Tab*

Section	Details
System Test	Validates that the machine fulfills Portability Suite's minimum hardware and operating system requirements.
Take Control Hardware Support	Checks for source hardware compatibility.
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 conversion 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 Knowledge Base Article Q20478 (http://support.platespin.com/kb2/article.aspx?id=20478).

2.5 Working with Device Drivers

Portability Suite ships with a library of device drivers, and during conversion 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 38.

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 Portability Suite driver database.

If you need to add drivers, use the Device Driver Manager, as detailed in the following procedure.

For problem-free operation of your conversion 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 systems

---

NOTE: On upload, Portability Suite 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.

To upload drivers to the Portability Suite device driver database:

1 Save the driver files to a local directory.

2 Click *Tools* > *Manage Device Drivers*.

The Device Driver Manager window opens.
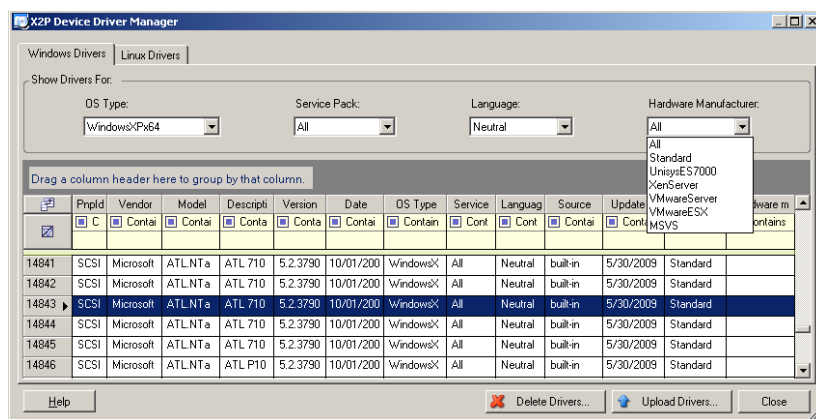
3 Select the tab corresponding to the required operating system, then select the applicable operating system and target hardware manufacturer details.

4 Click *Upload Drivers*.

5 Use the *Device Driver Uploader* dialog box to specify the path to the directory where your saved the required driver files.

6 Click *Upload,* then click *OK*.

2.6 Managing Custom Post-conversion Actions

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

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

◆ One-time Server Sync
◆ Peer-to-peer workload conversion

For the capability to select a post-conversion action to run as part of a conversion job, you must first save the action and its dependencies in a dedicated directory and add it to the Portability Suite Server's library. The maximum size of the directory must not exceed 64 MB. For information about raising this limit, see "Increasing the Size Limit on Post-conversion Actions Uploaded to Portability Suite Server" in your *Configuration Guide*.

Use the following procedure for adding a post-conversion action to the Portability Suite 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 Portability Suite Server can access.

Take special care when developing post-conversion 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 Knowledge Base Article 21121 (http://support.platespin.com/kb2/article.aspx?id=21121).

2 In the Portability Suite Client, click *Tools > Manage Actions*.

3 Click *Add*. The Add Action window opens:

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.

Portability Suite 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*.

Portability Suite packages and uploads the library.

The action is now available for selection in conversion jobs. See "Including a Custom Post-conversion Action in a Conversion Job" on page 111.

2.7 Setting Up, Executing, and Managing Jobs

Most discovery, workload portability, and workload protection actions and operations are defined, executed, and monitored as Portability Suite jobs.

2.7.1 Selecting a Conversion Job Setup Method

You can set up a conversion 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 conversion 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, Portability Suite provides a convenient subset of the Advanced method, interactively captures the essential parameters of a conversion job, itemizes settings specified in previous wizard pages, and provides a summary at the end.

Figure 2-6 *Job Configuration Window in Wizard Mode*
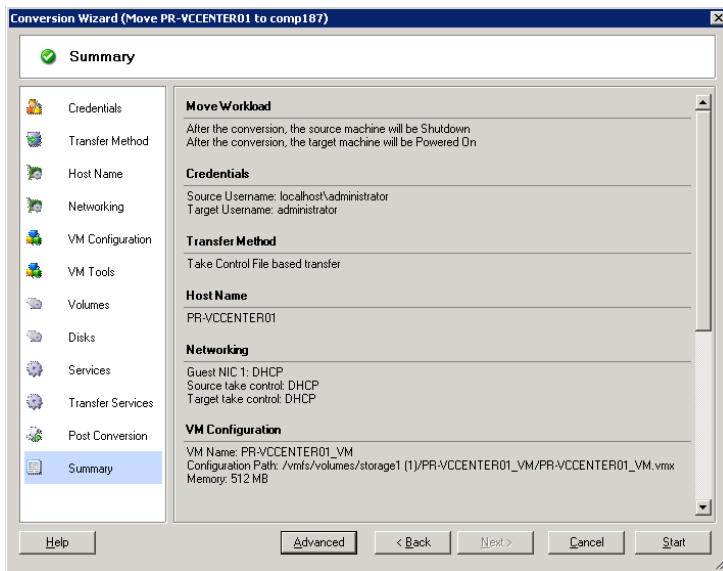
In Advanced mode, Portability Suite provides access to enhanced job configuration options and advanced settings for some aspects of a conversion job. In Advanced mode, you access settings for an aspect of a conversion job by clicking links that are grouped in categories. To hide or show category details, click the category bar.

Figure 2-7 *Job Configuration Window in Advanced Mode*
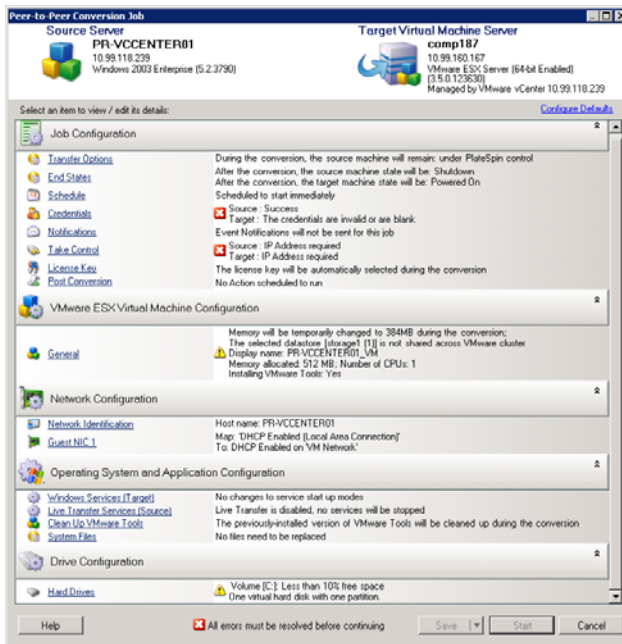
You can access the following enhanced job configuration options in Advanced mode, but cannot access them in Wizard mode:

◆ Specifying a non-default license for a conversion job. See "Selecting a License Key for a Migration or Protection Conversion" on page 77.

* Enabling automatic job progress and status notifications. See "Configuring Automatic E-Mail Notifications of Job Status and Progress" on page 78.

* Job scheduling options, except when protecting a workload. See "Scheduling and Synchronization" on page 83.

* Advanced options for target virtual machine configuration. See "Virtualization Platform-Specific and Advanced VM Configuration Options" on page 90.

* Saving the job for on-demand execution at a later time. See "Saving Conversion Jobs" on page 45.

Starting a Conversion Job

To start setting up a conversion 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:
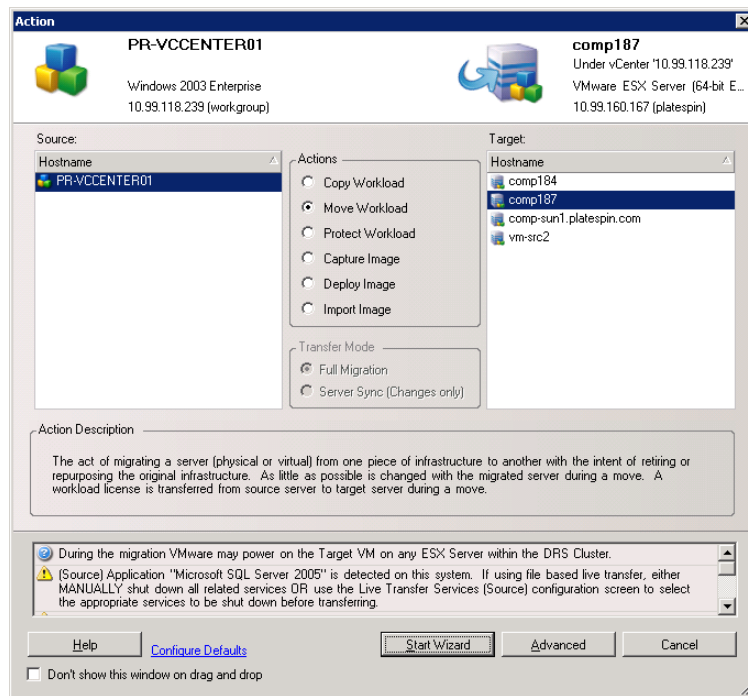
The Source and Target panes display workloads and targets applicable to the selected type of a conversion job.

The *Full Transfer* and *Server Sync* options are enabled only if:

* The system detects an existing operating system on the target

* The operating system profile of the target matches that of the source workload

See .

2 Check validation messages at the bottom of the window.

3 To start configuring your conversion 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 conversion startup behavior, restore application defaults. See "General Options" in your *Configuration Guide*.

2.7.2 Pre-conversion Job Validation

When starting a conversion job, Portability Suite 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 conversion job can start.

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

In a default Portability Suite configuration, validation messages display at the bottom of the Action window. See .

If you have configured Portability Suite to bypass the Action window on drag-and-drop, errors and warnings are displayed in a separate window:

Figure 2-8 *Conversion Validation Window*
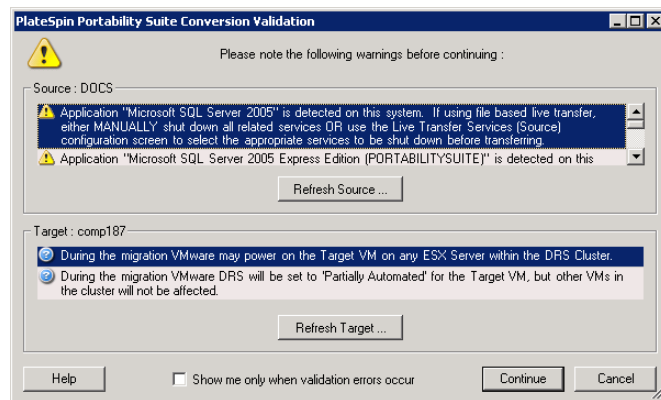
To force this window to open only on errors, select the *Show me only when validation errors occur* option.

2.7.3 Saving Conversion Jobs

When you start a conversion 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:
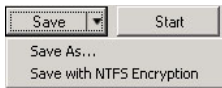
- **Save As:** Select this option to save the conversion job to run at a later time.
- **Save with NTFS Encryption:** Select this option to save the conversion job to run at a later time.

2.7.4 Starting a Saved Conversion Job in the Jobs View

1 In the Portability Suite Client's Jobs view, locate the required saved job.

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

Portability Suite starts the conversion job.

2.7.5 Changing the Start Time of a Saved Conversion Job

1 In the Portability Suite 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.
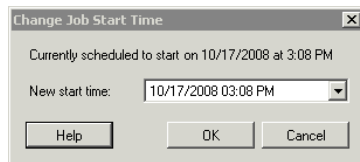
3 Specify the required new start date and time, then click *OK*.

Portability Suite reschedules the job and executes it at the specified time.

2.7.6 Viewing the Parameters of a Completed or In-Progress Conversion Job

1 In the Portability Suite Client's Jobs view, locate the required job.

2 Right-click the job and select *View*.

Portability Suite opens the job configuration window in Advanced mode and displays the job's configuration parameters in read-only mode.

2.7.7 Canceling an In-Progress Job

1 In the Portability Suite Client's Jobs view, locate the required job.

2 Right-click the job and select *Abort*.

2.7.8 Controlling the State of Source Workloads in the Jobs View

You can restart or shut down your source workload in the Jobs view.

1 In the Portability Suite 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-conversion state in your conversion job. See .

2.7.9 Running Job Diagnostics

1 In the Portability Suite Client's Jobs view, locate the required job.

2 Right-click the job, select *Run Diagnostics*, and wait for the process to complete.

After a few moments, Portability Suite launches a Web browser window with the results of comprehensive diagnostics on the job.

2.7.10 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 Portability Suite Client's *Jobs* view, locate the required job.

2 Right-click the job and select *Report*.
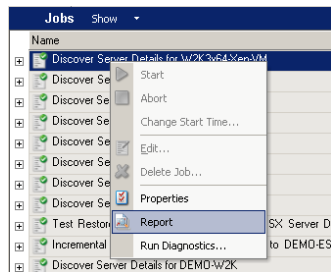
A Web browser window displays the requested report.

2.7.11 Using Windows Scheduled Tasks and a Batch File to Run Jobs

You can schedule and run a job by using Windows Scheduled Tasks and a batch file. At the scheduled time Windows calls the batch file, which launches the Portability Suite in console mode and executes the specified job.

1 Set up and save the required job, and note its location on the hard drive.

2 Use a text editor to create a batch file containing the following line:

```
PortabilitySuite.Console.exe run /job="path_to_job_file"
```

Replace *path_to_job_file* with the actual path to the required job file.

3 Name the batch file and save it to the same directory as the `PortabilitySuite.Console.exe` application:

```
..\PlateSpin Portability Suite Client\CommandLine
```

4 Use Windows Scheduled Tasks to set up a task to run the batch file at the desired date and time. For more information about Windows Scheduled Tasks, see your Windows documentation.

Below is a sample command for a batch file:

```
PortabilitySuite.Console.exe run /job=
"C:\documents and settings\Administrator\Local Settings\
Application Data\PlateSpin ULC\PlateSpin PortabilitySuite\8.0.0.xxxx\
jobs\4f035cda-5c2b-4b73-9923-148ba743e560\32025d86-8850-4c79-8bf3-
bb9c22b0ac1d.job"
```

Portability and Protection with Peer-to-Peer Workload Conversions

3

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

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

- Section 3.1, "Converting a Workload to a Virtual Machine (P2V, V2V)," on page 49
- Section 3.2, "Converting a Workload to a Physical Machine (P2P, V2P)," on page 50
- Section 3.3, "Windows Workload Protection through Virtualization," on page 51
- Section 3.4, "Managing Workload Protection Contracts," on page 52
- Section 3.5, "Synchronizing Workloads with Server Sync," on page 53
- Section 3.6, "Semi-Automated Workload Virtualization through X2P," on page 55

3.1 Converting a Workload to a Virtual Machine (P2V, V2V)

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

1 Discover or refresh your source workload and your target VM host. See "Discovering Source Workloads and Migration Targets" on page 32.

2 Start a conversion job. See "Selecting a Conversion Job Setup Method" on page 42.

3 Configure the required parameters of the job. See "Essentials of Workload Migration and Protection Jobs" on page 77.

4 (Optional) If you are converting a workload to a virtual machine for business continuity purposes, enable incremental synchronizations. See "Scheduling and Synchronization" on page 83.

5 Run the job.

Figure 3-1 *Peer-to-Peer Conversion Job Window (P2V in Advanced Mode)*

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

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

1 Discover your source workload. See "Discovering Source Workloads and Migration Targets" on page 32.

2 Discover and register your target physical machine with Portability Suite. See "Discovering and Registering Target Physical Machines" on page 35.

3 (Recommended) Use PlateSpin® Analyzer to ensure that:

◆ Your source operating system and hardware are supported by Portability Suite.

◆ Portability Suite'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 38.

4 (Conditional) If drivers for the physical target are missing, upload the required drivers to the Portability Suite's X2P device driver library. See "Working with Device Drivers" on page 39.

5 Start a conversion job. See "Selecting a Conversion Job Setup Method" on page 42.

6 Configure the required parameters of the job. See "Essentials of Workload Migration and Protection Jobs" on page 77.

7 Run the job.

Figure 3-2 *Peer-to-Peer Conversion Job Window (V2P in Advanced Mode)*
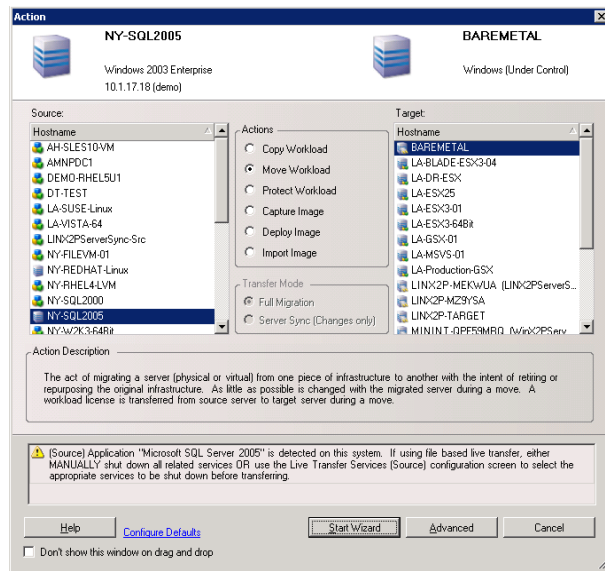
3.2.1 Best Practices (X2P)

◆ 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 94.

◆ When you are using the Take Control 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 Conversion Networking" on page 84.

◆ Ensure that vendor partitions are not being copied from the source. See "Handling a Workload's Storage Media and Volumes" on page 100.

3.3 Windows Workload Protection through Virtualization

Peer-to-peer workload protection is a function of workload portability to a virtual machine infrastructure. You use a Live Transfer method to convert a workload to a virtual machine and simultaneously establish an optional protection contract in your conversion job to incrementally update the target virtual machine on a recurring schedule. See "Scheduling and Synchronization" on page 83.

In terms of scope of data transferred from your source workload to the target, you can protect a workload by using one of the following two methods:

◆ **Straight P2V or V2V Conversion:** To protect a workload with this method, you convert your workload to a virtual machine and enable incremental synchronizations as a parameter of the conversion job. During the conversion process, Portability Suite creates an initial virtual replica

of the source machine, transfers the workload to the VM, and regularly updates it according to the specified synchronization schedule. See "Converting a Workload to a Virtual Machine (P2V, V2V)" on page 49.

  ◆ **Server Sync virtualization with incremental synchronization:** To protect a workload with this method, you enable incremental synchronizations as a parameter of the Server Sync job, which creates the initial virtual replica of the source machine without transferring the entire volume data; it transfers only the differences between your source and an existing *base* VM with a matching OS profile. Subsequent changes are regularly synchronized according to the specified synchronization schedule. See "Synchronizing Workloads with Server Sync" on page 53.

In either case, the virtual machine is left offline as a warm back-up. To support disaster recovery plans and disaster recover-readiness exercises, Portability Suite provides mechanisms for testing and implementing workload failover to the virtual machine. See "Managing Workload Protection Contracts" on page 52.

3.4 Managing Workload Protection Contracts

Portability Suite enables you to create and manage your workload protection contracts. For example, at any time during the period, in which a protection contract is active, you can manually run full and incremental transfers, pause and resume synchronization schedules, test and run failover jobs, edit protection jobs, and view logs.

NOTE: Workload protection through virtualization is supported for Windows workloads only. See "Supported Microsoft Windows Workloads" on page 14.

---

  1 On the Portability Suite Client toolbar, click the *Synchronization Schedules* button. The *Synchronization Schedules* window opens.
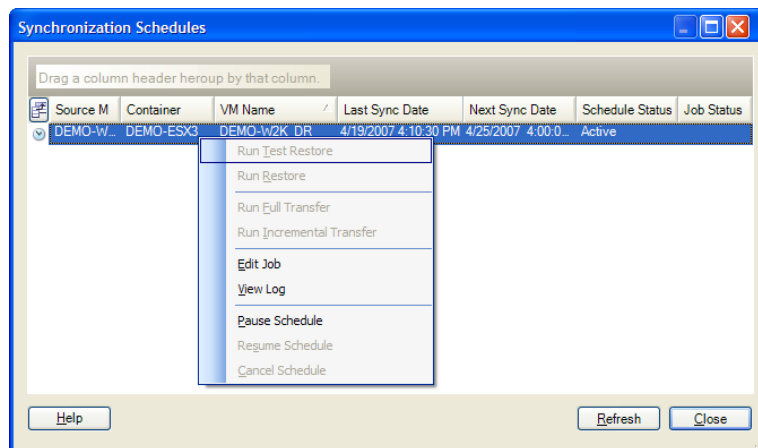
2 Right-click the synchronization schedule that you want to manage and select one of the following commands:

Run Test Restore	Verifies that the replica of the protected source server is functional. This feature performs the following actions:
	1. Creates a conversion job that you can use to specify required temporary network settings, such as placing the VM on an isolated network or assigning it a different IP address.
	2. Creates a snapshot of the VM before powering it on, and removes the snapshot after completion, reverting the VM to its original state. The job also reverses any VM configuration changes, such as NIC connection settings.
	3. Suspends scheduled synchronizations and resumes after completion.
	NOTE: Before executing a Test Restore, change the server name and IP address of the test restore server or connect it to a different physical or logical network. Also, be aware that incremental transfers are paused, and do not resume until the test restore server is powered off.
Run Restore	Brings the target virtual machine online. You would most likely choose this option when the source server has failed. An abbreviated Conversion Job screen opens where you can change and/or confirm the job configuration prior to bringing the server online.
Run Full Transfer	Forces a full transfer from the source server to the target virtual machine.
Run Incremental Transfer	Forces an incremental transfer.
Edit Job	Enables you to edit the incremental transfer schedule.
View Log	Displays status information for the various transfers that have occurred: Initial Copy, Full Transfer, and Incremental Transfer.
Pause Schedule	Puts incremental transfers on hold.
Resume Schedule	Resumes incremental transfers.
Cancel Schedule	Cancels the synchronization schedule.
	NOTE: Before canceling a synchronization schedule, you need to determine what to do with the target virtual machine. To keep the virtual machine, run a recovery job before canceling the schedule. The virtual machine can then be used on any virtual machine server.

3.5 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 or protection 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. Portability Suite 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.

You can use Server Sync for:

- A single migration of a supported Windows or Linux workload (a one-time Server Sync job)
- An ongoing protection of a Windows workload

The capability depends on your product license.

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.

To perform a Server Sync job:

1 Discover your source workload. See "Discovering Source Workloads and Migration Targets" on page 32.

2 (Conditional) If your intended target is a virtual machine:

> **2a** Create a target virtual machine by using one of the following methods:
>
> - Do an initial conversion of your workload to a virtual machine. See "Converting a Workload to a Virtual Machine (P2V, V2V)" on page 49.
> - Using your virtualization platform's native interface, manually install a virtual machine with the same operating system profile as that of your source.
> - Capture your workload to a Flexible Image, and deploy it to a virtual machine on your virtualization platform. See "Capturing a Workload to a Flexible Image" on page 65 and "Deploying a Flexible Image" on page 69.
>
> **2b** Discover and prepare the target VM for synchronization. See "Discovering and Preparing a Windows VM for a Server Sync Job" on page 38.

3 (Conditional) If your intended target is a physical machine, discover your physical target by using the Take Control ISO boot image. See "Discovering and Registering Target Physical Machines" on page 35.

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 (Optional) If you are protecting a workload to a VM, enable incremental synchronizations. See "Scheduling and Synchronization" on page 83.

8 Click *Start*.

Portability Suite starts the job and lists it in the Jobs view.

3.6 Semi-Automated Workload Virtualization through X2P

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

◆ Microsoft Hyper-V

◆ Xen Hypervisor on SUSE® Linux Enterprise Server (SLES) 10

◆ Virtual Iron

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 Novell® PlateSpin® Take Control ISO image.

This discovers and registers the target virtual machine as a Portability Suite physical machine target. This is similar to the process described in "Discovering and Registering Target Physical Machines" on page 35. 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 Portability Suite Client to create and execute an X2P conversion job.

4 Upon completion of the conversion job, install virtualization enhancement software specific to the target virtualization platform.

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

◆ Section 3.6.1, "Migrating a Workload to Microsoft Hyper-V," on page 55

◆ Section 3.6.2, "Migrating a Workload to Xen Hypervisor on SLES," on page 58

◆ Section 3.6.3, "Migrating a Workload to Virtual Iron," on page 60

3.6.1 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 Portability Suite and Hyper-V. See "Supported Target Virtualization Platforms" on page 15.

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 Take Control Boot ISO Image (Hyper-V)" on page 56
- "Creating and Configuring the Target Virtual Machine (Hyper-V)" on page 56
- "Registering the Virtual Machine with Portability Suite Server (Hyper-V)" on page 57
- "MIgrating Your Source Workload to the Target Virtual Machine (Hyper-V)" on page 57
- "Post-Conversion Steps (Hyper-V)" on page 58

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

1 Download the appropriate Take Control ISO image for your workload. See Table 2-2 on page 36.

2 (Optional) If you want to perform an unattended registration of your target, update the ISO image with the required registration parameters. See Knowledge Base Article 20832 (http://support.platespin.com/kb2/article.aspx?id=20832).

3 Save the ISO image in a location that Hyper-V server can access. For example: `c:\temp`.

This ensures that the Take Control 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 Take Control 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 Portability Suite.

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 Portability Suite. If you are using the unattended registration process, the required parameters are read from an answer file.

Registering the Virtual Machine with Portability Suite Server (Hyper-V)

Complete these steps to register your target with Portability Suite Server when you are using an unmodified Take Control ISO image. For information on how to modify the image for an unattended registration process, see Knowledge Base Article 20832 (http://support.platespin.com/kb2/article.aspx?id=20832).

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

- ◆ **Block transfer:** Enter Y to transfer data at the block level or enter N to transfer data at the file level. The block-level transfer selection applies to the Block-based Live Transfer method (not the VSS Block-based Live Transfer method). See "Supported Transfer Methods" on page 16.

- ◆ **Portability Suite Server:** Use the following syntax:

  http://*server_host*/PortabilitySuite

  Replace *server_host* with the actual Portability Suite Server host's name or IP address.

- ◆ **Credentials (User Name/Password):** Enter the name of an admin-level user on the Portability Suite 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 Portability Suite Client to use to list the newly registered VM. The workload's target hostname you select in the conversion job overwrites this name.

- ◆ **SSL encryption:** If your Portability Suite is installed on a host with SSL encryption enabled, enter Yes. If not, enter No.

- ◆ **Portability Suite Network:** Unless you have defined your own Portability Suite Network in Portability Suite Client, press the Enter key. If you are working with a non-default Portability Suite Network, type its name, then press the Enter key.

  A controller on your target virtual machine communicates with Portability Suite Server and registers the virtual machine as a physical target for a conversion job.

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

1 Use Portability Suite Client to start an X2P conversion job with your source workload being the job's conversion source and the target being the new VM on Hyper-V. See "Converting a Workload to a Physical Machine (P2P, V2P)" on page 50.

2 Monitor the conversion job in Portability Suite 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 Take Control 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 conversion job resumes, reboots the target, and completes the workload configuration.

Post-Conversion Steps (Hyper-V)

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

3.6.2 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 (not paravirtualized) virtual machine.

◆ Your source workload is supported by Portability Suite and the Xen hypervisor. See "Supported Target Virtualization Platforms" on page 15.

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 Take Control Boot ISO Image (Xen on SLES)" on page 58
◆ "Creating and Configuring the Target Virtual Machine (Xen on SLES)" on page 58
◆ "Registering the Virtual Machine with Portability Suite Server (Xen on SLES)" on page 59
◆ "Migrating Your Source Workload to the Target Virtual Machine (Xen on SLES)" on page 60
◆ "Post-Conversion Steps (Xen on SLES)" on page 60

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

1 Download the appropriate Take Control ISO image for your workload. See Table 2-2 on page 36.

2 (Optional) If you want to perform an unattended registration of your target, update the ISO image with the required registration parameters. See Knowledge Base Article 20832 (http://support.platespin.com/kb2/article.aspx?id=20832).

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 conversion and improves transfer speed. If the virtual machine requires less memory after the conversion, reduce the assigned memory through after the conversion 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 Take Control 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 Portability Suite. If you are using the unattended registration process, the required parameters are read from an answer file.

Registering the Virtual Machine with Portability Suite Server (Xen on SLES)

Complete these steps to register your target with Portability Suite Server when you are using an unmodified Take Control ISO image. For information on how to modify the image for an unattended registration process, see Knowledge Base Article 20832 (http://support.platespin.com/kb2/article.aspx?id=20832)

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

- **Block transfer:** Enter `Y` to transfer data at the block level or enter `N` to transfer data at the file level. The block-level transfer selection applies to the Block-based Live Transfer method (not the VSS Block-based Live Transfer method). See "Supported Transfer Methods" on page 16.

- **Portability Suite Server:** Use the following syntax:

    `http://`*server_host*`/PortabilitySuite`

    Replace *server_host* with the actual Portability Suite Server host's name or IP address.

- **Credentials (User Name/Password):** Enter the name of an admin-level user on the Portability Suite 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 Portability Suite Client to use to list the newly registered VM. The workload's target hostname you select in the conversion job overwrites this name.

- **SSL encryption:** If your Portability Suite is installed on a host with SSL encryption enabled, enter `Yes`. If not, enter `No`.

- **Portability Suite Network:** Unless you have defined your own Portability Suite Network in Portability Suite Client, press they Enter key. If you are working with a non-default Portability Suite Network, type its name, then press the Enter key.

  A controller on your target virtual machine communicates with Portability Suite Server and registers the virtual machine as a physical target for a conversion job.

Migrating Your Source Workload to the Target Virtual Machine (Xen on SLES)

1 Use Portability Suite Client to start an X2P conversion job with your source workload being the job's conversion source and the target being the new VM on the Xen hypervisor. See "Converting a Workload to a Physical Machine (P2P, V2P)" on page 50.

2 Monitor the conversion job in Portability Suite 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 Take Control 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 conversion job resumes, reboots the target, and completes the workload configuration.

Post-Conversion 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).

3.6.3 Migrating a Workload to Virtual Iron

You can use Virtual Iron as the target virtualization platform in a semi-automated workload virtualization.

Prerequisites:

- Your source workload is supported by Portability Suite and Virtual Iron. See "Supported Target Virtualization Platforms" on page 15.

Additional information:

- *Virtual Iron Administrator's Guide*

This section includes the following topics:

- "Downloading and Saving the Take Control Boot ISO Image (Virtual Iron)" on page 61
- "Creating and Configuring the Target Virtual Machine (Virtual Iron)" on page 61
- "Registering the Virtual Machine with Portability Suite Server (Virtual Iron)" on page 61
- "Migrating Your Source Workload to the Target Virtual Machine (Virtual Iron)" on page 62
- "Post-Conversion Steps (Virtual Iron)" on page 62

Downloading and Saving the Take Control Boot ISO Image (Virtual Iron)

1 Download the appropriate Take Control ISO image for your workload. See Table 2-2 on page 36.

2 (Optional) If you want to perform an unattended registration of your target, update the ISO image with the required registration parameters. See Knowledge Base Article 20832 (http://support.platespin.com/kb2/article.aspx?id=20832).

3 Save the ISO image in the appropriate directory:

- For Linux Virtualization Manager: `/opt/VirtualIron/nbd`

- For Windows Virtualization Manager: `..\Program Files\nbd`

This ensures that the Take Control ISO image is available as a boot option for your target virtual machine.

Creating and Configuring the Target Virtual Machine (Virtual Iron)

1 Use the Virtual Iron Virtualization Manager to create a new *Virtual Server* with the following settings:

- **Storage:** 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.

- **Virtual Memory:** Assign at least 384 MB of RAM to the VM. This ensures that the VM has sufficient resources during the conversion and improves transfer speed. If the virtual machine requires less memory after the conversion, reduce the assigned memory through the Virtual Iron Virtualization Manager.

- **Operating System:** Must match the operating system profile of your source workload.

- **VS Tools Installed:** Make sure this options is deselected. You need to install virtualization enhancement software after completing the conversion.

- **Boot Options:** Select the *Network (Image) Boot* option, then select the `*.iso` image file you downloaded saved.

2 When the configuration is complete, commit the new virtual machine.

3 Launch the selected virtual machine's 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 Portability Suite. If you are using the unattended registration process, the required parameters are read from an answer file.

Registering the Virtual Machine with Portability Suite Server (Virtual Iron)

Complete these steps to register your target with Portability Suite Server when you are using an unmodified Take Control ISO image. For information on how to modify the image for an unattended registration process, see Knowledge Base Article 20832 (http://support.platespin.com/kb2/article.aspx?id=20832).

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

- **Block transfer:** Enter `Y` to transfer data at the block level or enter `N` to transfer data at the file level. The block-level transfer selection applies to the Block-based Live Transfer method (not the VSS Block-based Live Transfer method). See "Supported Transfer Methods" on page 16.

- **Portability Suite Server:** Use the following syntax:

```
http://server_host/PortabilitySuite
```

Replace *server_host* with the actual Portability Suite Server host's name or IP address.

- ◆ **Credentials (User Name/Password):** Enter the name of an admin-level user on the Portability Suite 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 Portability Suite Client to use to list the newly registered VM. The workload's target hostname you select in the conversion job overwrites this name.

- ◆ **Number of processors:** Enter the number of processors supported for your source workload as a guest on Virtual Iron.

- ◆ **SSL encryption:** If your Portability Suite is installed on a host with SSL encryption enabled, enter `Yes`. If not, enter `No`.

- ◆ **Portability Suite Network:** Unless you have defined your own Portability Suite Network in Portability Suite Client, press the Enter key. If you are working with a non-default Portability Suite Network, type its name, then press the Enter key.

   A controller on your target virtual machine communicates with Portability Suite Server and registers the virtual machine as a physical target for a conversion job.

Migrating Your Source Workload to the Target Virtual Machine (Virtual Iron)

1 Use Portability Suite Client to start an X2P conversion job with your source workload being the job's conversion source and the target being the new VM on Virtual Iron. See "Converting a Workload to a Physical Machine (P2P, V2P)" on page 50.

2 Monitor the conversion job in Portability Suite Client's Jobs view.

When the job reaches the step *Configure Target Machine* step, the virtual machine's console returns to the boot prompt of the Take Control 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 conversion job resumes, reboots the target, and completes the workload configuration.

Post-Conversion Steps (Virtual Iron)

Install VSTools (virtualization enhancement software). You must install VSTools to the disk from which the guest operating system boots. For more information, see your *Virtual Iron Administrator's Guide*.

Portability and Protection with a Flexible Image

4

This section provides information about using the PlateSpin® Flexible Image volume archiving feature for workload protection and portability.

4.1 About Flexible Images

One of Portability Suite's three fundamental workload infrastructures, a PlateSpin® Flexible 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.

You can use Flexible Images for provisioning, data recovery, basic disaster recovery, and as an intermediary format in peer-to-peer workload protection jobs, such as Server Sync. See "Synchronizing Workloads with Server Sync" on page 53.

Image configurations are maintained in an XML (`config.xml`) file with each image having one or more sets of associated volume data.

Flexible Images and the image server's `config.xml` configuration file are stored on the designated Flexible Image Server host in the following directory:

```
..\Program Files\PlateSpin Image Server
```

In addition to volume data directly captured during an X2I conversion, Portability Suite supports:

* Third-party volume archives, such as Acronis True Image or Symantec LiveState
* Raw volume data

Like peer-to-peer conversions, 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.

4.2 Designating a Flexible Image Server

To work with Flexible Images, you must first designate a machine as an image server by installing on it the Flexible Image Server software. You can install a Flexible Image Server instance either on a dedicated host or on your Portability Suite Server host.

Dedicated Flexible Image Server hosts must meet the following requirements:

Table 4-1 *Flexible Image Server Host Requirements*

Requirement	Details
Operating System	Any of the following:
	◆ Microsoft Windows Server 2008
	◆ Microsoft Windows Vista
	◆ Microsoft Windows Server 2003
	◆ Microsoft Windows XP
	◆ Microsoft Windows 2000
	◆ Microsoft Windows NT 4.0
Disk Space	Minimum 100 MB for basic controller software.
	Additional space might be necessary, depending on the number and size of workload images that you intend to store on a given image server.
Software	◆ Microsoft .NET Framework 2.0

To designate a machine as a Flexible Image Server:

1 Discover the system you plan to designate as a Flexible Image Server.

2 In the Servers view, right-click the discovered server and select *Install Image Server*.

The Install Image Server dialog box opens.
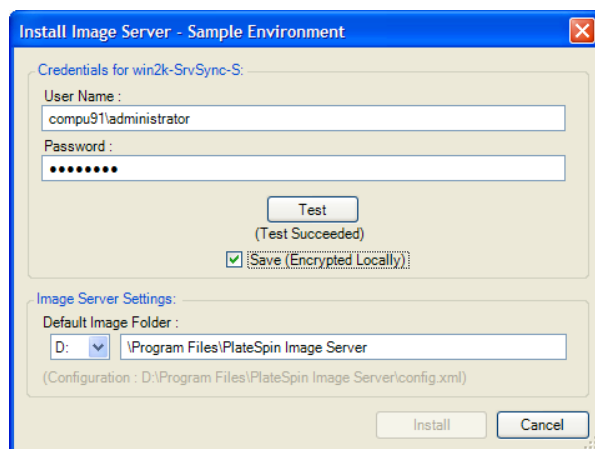
3 Provide administrator credentials for the selected host and specify the desired directory for image files.

4 Click *Install*.

Portability Suite installs a controller on the selected host and configures it to run as a Flexible Image Server. On completion, the Servers view lists a new Flexible Image Server item:

4.3 Capturing a Workload to a Flexible Image

Use this procedure to capture a physical or virtual workload as a PlateSpin® Flexible Image.

1 Discover, or refresh the details of, your source workload and your Flexible 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 Portability Suite 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.
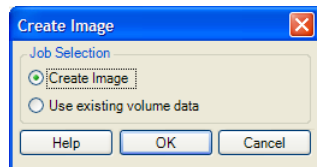
3 Select *Create Image*, then click *OK*. For information about the *Use existing volume data* option, see "Creating Images by Using Raw or Existing Volume Data" on page 66.

4 The *Create Image Job* window opens.
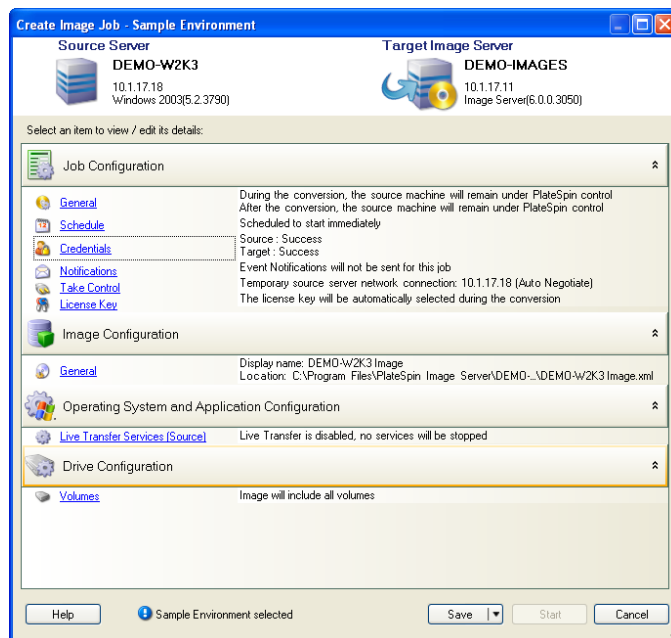
5 Specify the required settings for the conversion job by clicking the links in each category:

- ◆ **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*).
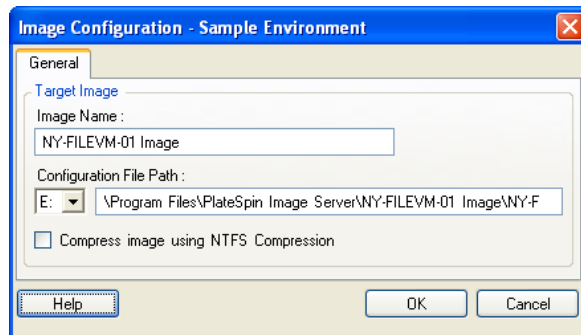
- ◆ **Operating System and Application Configuration:** If you selected the Live Transfer method, specify how you want Portability Suite to handle operating system and application services on your source (*Live Transfer Services*).

- ◆ **Drive Configuration:** Select the volumes that you want Portability Suite to include in the image and specify the path for the package file (under Drive Configuration, click *Volumes*).
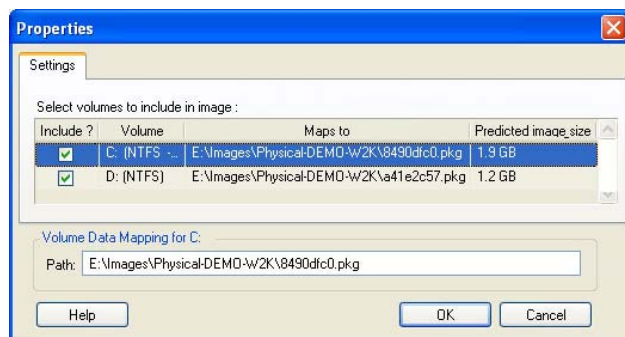
4.4 Creating Images by Using Raw or Existing Volume Data

Portability Suite can use existing system images created with Acronis True Image or Symantec LiveState as volume data for a PlateSpin® Flexible Image. To use this functionality, Acronis True Image or Symantec LiveState must be installed on the Flexible Image Server host.

For image import, Portability Suite supports the following third-party volume backup formats:

◆ Acronis (True Image Server 8.0, True Image 9.1, True Image Enterprise Server 8.0, True Image Enterprise Server 8.0).

◆ Symantec (LiveState Recovery Advanced Server Suite 3.0 and 6.0, LiveState Recovery Advanced Server 3.0 and 6.0, LiveState Recovery Standard Server 3.0 and 6.0).

Portability Suite can also use raw volume data when deploying an image.

◆

◆

4.4.1 Importing Raw Volume Data or Volume Data Extracted from Acronis True Image and Symantec LiveState

1 Discover the source machine with the raw volume data you want to capture, or that corresponds to the existing Acronis True Image or Symantec LiveState image.

2 Drag and drop the source server onto the target Flexible 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*. Alternatively, to use the Wizard, click *Start Wizard*.

◆ If you configured Portability Suite 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*.
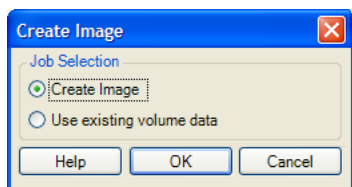

3 In the *Create Image Job* window, click *Volumes* under the *Drive Configuration* category. A Properties dialog box opens.
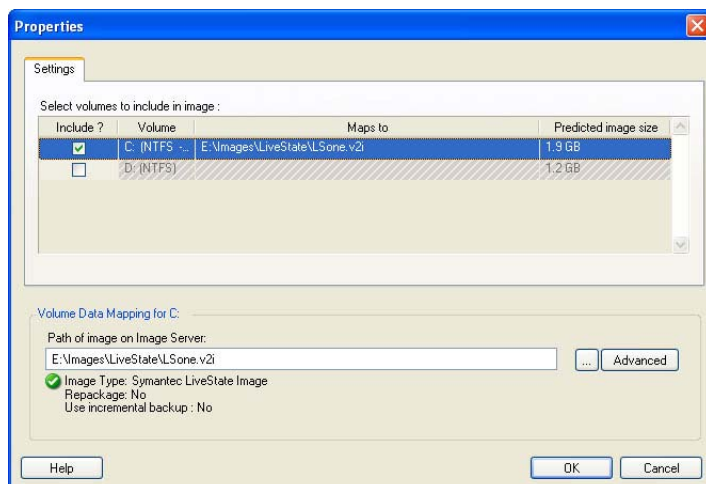
4 Select each volume. In the text box at the bottom of the dialog box, type the full path to the directory containing the raw volume data, or to the existing image file for each volume.

If you are using an incremental LiveState image, the `*.iv2i` and `*.v2i` files must exist in the same directory. In that case, enter the path of the `*.iv2i` file in the *Path of image on Image Server* text box.

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

   ◆ Raw volume data

Repackage: Select this option to create a Flexible Image from the provided image file.
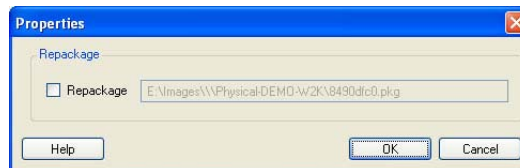
   ◆ Symantec LiveState

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

Security: If the provided image is password protected, provide the required password.

Use incremental backup: Select this option if the provided image is part of an incremental backup set.

   ◆ Acronis True Image

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

Security: If the provided image is password protected, provide the required password.

Use incremental backup: Select this option if the provided image is part of an incremental backup set.

Index: Enter the index number for the volume contained in the provided image. The following tables show the index numbers for each volume.
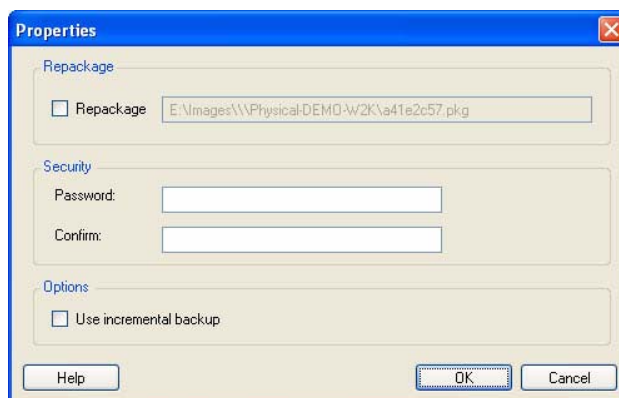
Volume:	A	B	C	D	E	F	G	H	I	J	K	L	M
Index:	0	1	2	3	4	5	6	7	8	9	10	11	12

Volume:	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
Index:	13	14	15	16	17	18	19	20	21	22	23	24	25

If the image contains data from the E: volume on the source server, the index should be set to 4.

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.

4.4.2 Additional Information About Creating Flexible Images

Table 4-2 *Knowledge Base Articles Related to Flexible Images*

ID	Description
Q20768 (http://support.platespin.com/kb2/article.aspx?id=20768)	INFO: Portability Suite and Double-Take* Application Note Download [PDF].
Q20392 (http://support.platespin.com/kb2/article.aspx?id=20392)	HOWTO: How to create a Flexible Image from NT backup restore.
Q20417 (http://support.platespin.com/kb2/article.aspx?id=20417)	HOWTO: How to create a Flexible Image from Computer Associates' BrightStor* restore.
Q20400 (http://support.platespin.com/kb2/article.aspx?id=20400)	HOWTO: How to create a Flexible Image from Symantec's Backup Exec* system restore.
Q20461 (http://support.platespin.com/kb2/article.aspx?id=20461)	HOWTO: How to create a Flexible Image from Ghost* Image extraction.
Q20723 (http://support.platespin.com/kb2/article.aspx?id=20723)	FAQ: How to create a Flexible Image from IBM* Tivoli* Storage Manager (TSM).

4.5 Deploying a Flexible Image

Use this procedure to deploy a Flexible Image on a supported physical machine or virtualization platform.

1 Drag and drop the required Flexible Image to a discovered target physical machine or VM host.
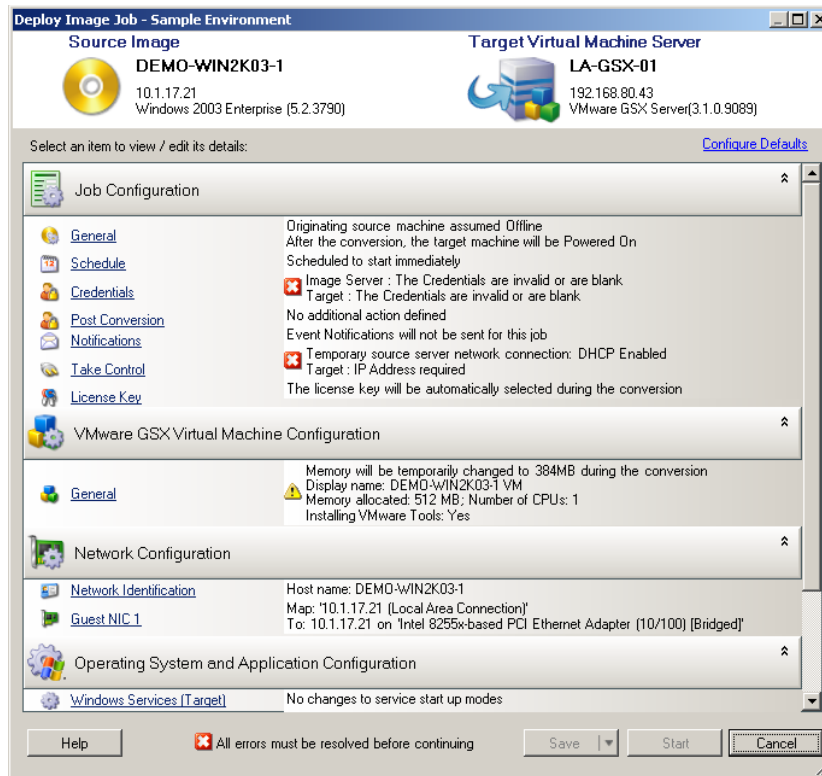
2 Specify the required settings for the conversion job by clicking the links in each category.

Conversion jobs are auto-configured to create the target machine with the same settings as the source server. Depending on the objectives of the conversion, 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 conversion.

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.

4.6 Incremental Imaging

Incremental imaging enables you to schedule recurring incremental synchronizations that capture only files that change over time.

Incremental imaging introduces the following image-related entities:

- **Base (or full) image:** A standalone entity that can be restored without additional files or data; contains complete workload information available at the time of capture.

◆ **Increment:** A point-in-time delta representing only changes occurring since the capture of the last full image or increment. You can restore a specific point-in-time state of an image by selecting the required increment during deployment. Applying an increment to the base image requires that all intervening increments be available.

◆ **Differential:** An aggregate of multiple increments. You can create an image differential by consolidating multiple consecutive increments into one.

4.6.1 Automatically Capturing Image Increments

Capturing image increments involves setting up a Capture Image job with incremental synchronizations enabled.

1 Discover the workload source and the target image server. If they are already discovered, refresh their details (right-click > *Refresh Server Details*) to ensure that Portability Suite has the most up-to-date information about them.

2 Using the *Capture Image* task in the *Tasks* pane, start a new image capture job with the required source and target, then click *Advanced* to start the *Create Image Job* window.

3 In the *Create Image Job* window, click *General*, then in the Job Configuration dialog box, click the *Schedule* tab. See "Scheduling and Synchronization" on page 83.

4 Select the *Enable Synchronizations option.*

5 Click *Edit Schedule* to set up the required recurrence for full and incremental synchronizations.

6 Set the number of required Rolling Increments.

This determines the number of incremental images that are stored at a given time. When this number is exceeded, the oldest increment is merged with the base image to form a new base.

7 Click *OK* and wait for the conversion to complete.

On completion, the Servers view lists the new image under the corresponding Flexible Image Server:

4.6.2 Viewing Increments

To view all increments associated with an image:

1 In the Servers view, right-click the required image and select *Properties*.

2 Click the *Image Increments* tab.

A list of all increments associated with the image is displayed. This includes each increment's name, creation date and time, and the path to the file containing its meta information.
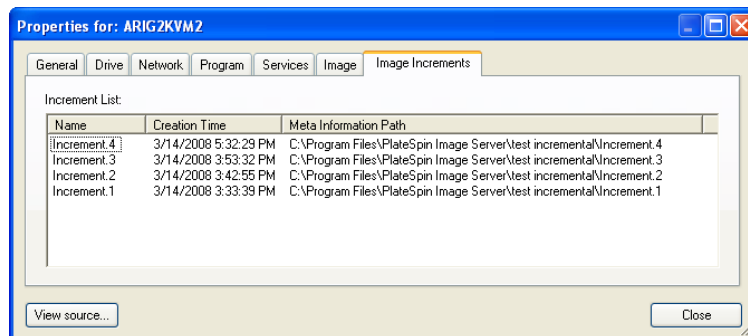
4.6.3 Selecting an Increment to Deploy with an Image

To deploy an image with a specific increment:

1 In the Tasks pane, click *Deploy Image*. The Action window opens.

2 Select the image and the target image server, then click *Advanced*.

3 Configure the conversion as job as described in "Deploying a Flexible Image" on page 69.

4 Under *Drive Configuration* in the *Deploy Image Job* window, click *Source Increment Selection*.

5 In the *Drive Configuration* dialog box, deselect *Use Latest Increment* and select the increment to be deployed.
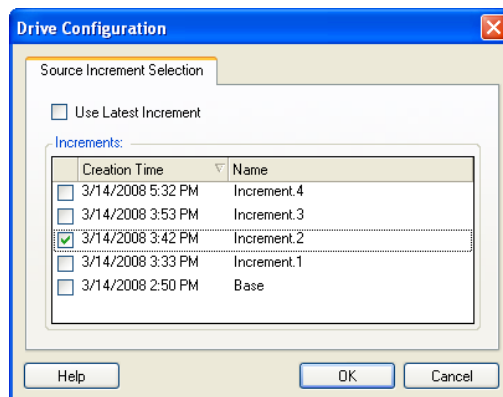
6 Click *OK*.

7 Click *Start* to start the conversion.

Portability Suite deploys the image, using the base image and all increments up to and including the one selected.

4.6.4 Consolidating Increments

Increment consolidation replaces two or more increments with a single, merged image increment. You can consolidate only consecutive increments.

1 In the Servers view, right-click the required image and select *Consolidate Increments*.
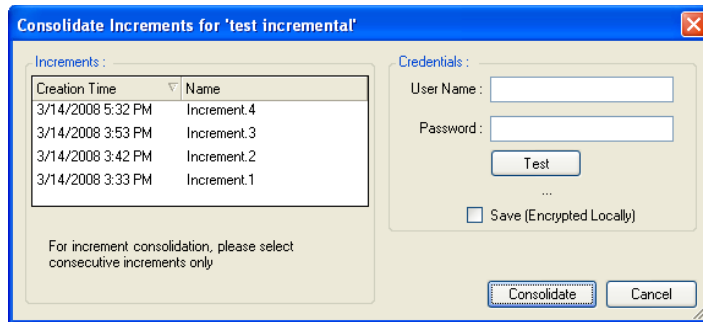
2 Select any number of consecutive increments to consolidate by holding down the Ctrl key while selecting with the mouse or the arrow keys.

3 Enter valid credentials for the image and test them.

4 Click *Consolidate*.

The increments are replaced with a single differential image. This image appears as an increment in the list of increments, but its name includes the IDs of all the consolidated increments. For example, `Increment.2-4`.

4.6.5 Applying Increments to the Base Image

Use this feature to update a base image to include one or more subsequent increments. Applying an increment merges all preceding increments with the base image.

1 In the Servers view, right-click the required image and select *Apply Increment to Base*.

2 Select the increment to apply.

All increments created before the selected increment are also applied.

3 Type valid credentials for the image and test them.

4 Click *Apply*.

Portability Suite applies the selected increments to the base image, replacing it with a new, updated image.

4.6.6 Deleting Increments

You can delete captured image increments if required. Deleting an increment results in the removal of all increments created after the selected increment.

1 In the Servers view, right-click the required image and select *Delete Increments*.

2 Select the increment you want to delete.

3 Enter valid credentials for the image and test them.

4 Click *Delete*.

Portability Suite removes the selected increments.

4.7 Managing Flexible Images

◆

4.7.1 Moving Images from One Flexible Image Server to Another

1. Copy the image directory from the old Flexible Image Server host's file system to a location on the new Flexible Image Server host.
2. Update the new Flexible Image Server's `config.xml` file to identify the path to and the name of the image that was moved from the old Flexible Image Server.
3. Refresh the new image server's details in Portability Suite Client's Servers view.

For more information, see Knowledge Base Article Q20189 (http://support.platespin.com/kb2/article.aspx?id=20189).

4.7.2 Automating Image Operations

You can use the ImageOperations command line utility, included with Portability Suite, to automate several tasks related to images, such as regularly moving multiple base images, along with related increments, between Flexible 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 Flexible Image and its volumes into a specified subdirectory.

To use the `ImageOperations command line utility`:

1. On your Flexible 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.

4.7.3 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® Flexible Images.

To do this, you can use the PlateSpin® Image Browser utility, which enables you to browse, search, sort, and extract files from:

* An image file
* A specific image increment file

You can work with both base images and image increments by loading:

- 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 75
- "Sorting and Searching Items in the Image Browser Interface" on page 75
- "Extracting Items" on page 76
- "Browsing and Extracting Image Files at the Command Line" on page 76

Starting the Image Browser and Loading Image Files

1 Start the ImageBrowser program (`ImageBrowser.exe`), located in one of the following directories:

- On your Portability Suite Server host:

    `..\PlateSpin Portability Suite Server\bin\ImageOperations`

- On your Flexible 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, then 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.
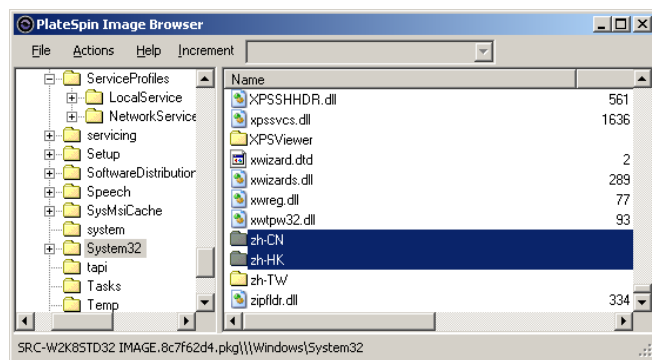
Depending on the size of the image, it might take a few seconds to several minutes for the utility to load the required file.

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 75.

2 In the left pane, select a volume or a subdirectory.

3 On the *Actions* menu, click *Search*.

Alternatively, 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 75.

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*.

Alternatively, 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 Flexible 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.

Essentials of Workload Migration and Protection Jobs

<div style="text-align: right">5</div>

This section provides information about configuring the key aspects of a workload conversion job, using the Portability Suite Client with either PlateSpin® Migrate or PlateSpin® Protect product.

Where applicable, differences between the Advanced and Wizard modes are noted. See "Setting Up, Executing, and Managing Jobs" on page 42.

- Section 5.1, "Selecting a License Key for a Migration or Protection Conversion," on page 77
- Section 5.2, "Configuring Automatic E-Mail Notifications of Job Status and Progress," on page 78
- Section 5.3, "Specifying Transfer Options," on page 79
- Section 5.4, "Specifying End States for Source and Target Workloads," on page 80
- Section 5.5, "Specifying Source and Target Workload Credentials," on page 80
- Section 5.6, "Managing a Workload's Network Identity," on page 81
- Section 5.7, "Scheduling and Synchronization," on page 83
- Section 5.8, "Setting Up Conversion Networking," on page 84
- Section 5.9, "Configuring Target Virtual Machines," on page 89
- Section 5.10, "Handling Operating System Services and Applications," on page 94
- Section 5.11, "Handling a Workload's Storage Media and Volumes," on page 100
- Section 5.12, "Including a Custom Post-conversion Action in a Conversion Job," on page 111

5.1 Selecting a License Key for a Migration or Protection Conversion

By default, Portability Suite automatically selects the best license key for a particular conversion job. If you have multiple license keys, you can select a specific license key to apply to a particular conversion job. For information about product licensing and license key management, see "Portability Suite Product Licensing" in your *Configuration Guide*.

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

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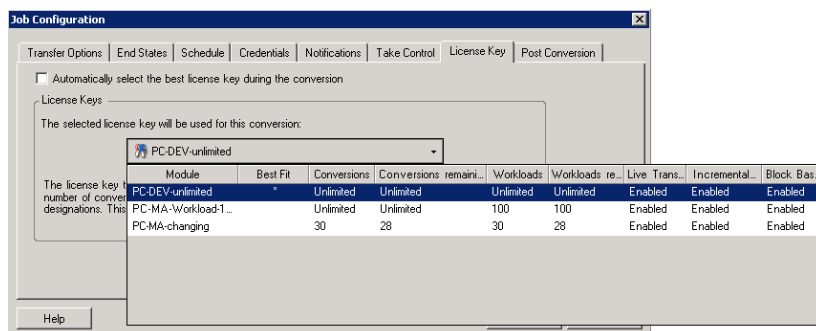
* **In Wizard mode:** Not available.

To manually choose a different key, deselect the *Automatically select the best license key during the conversion* 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 conversion. Licenses can be invalid for reasons such as:

* There are no remaining conversions for the license.
* The license does not allow X2V conversions and the current conversion is a P2V.
* The license does not support live transfer conversions and the current conversion is marked for live transfer.

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

5.2 Configuring Automatic E-Mail Notifications of Job Status and Progress

You can set up a conversion job to automatically send e-mail notifications to a specified address about:

* **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 conversion job or globally. See "SMTP Notifications" in your *Configuration Guide*.

To configure automatic e-mail notifications:

- **In Advanced mode:** In the Conversion Job window, under the Job Configuration section, click *Notifications*.
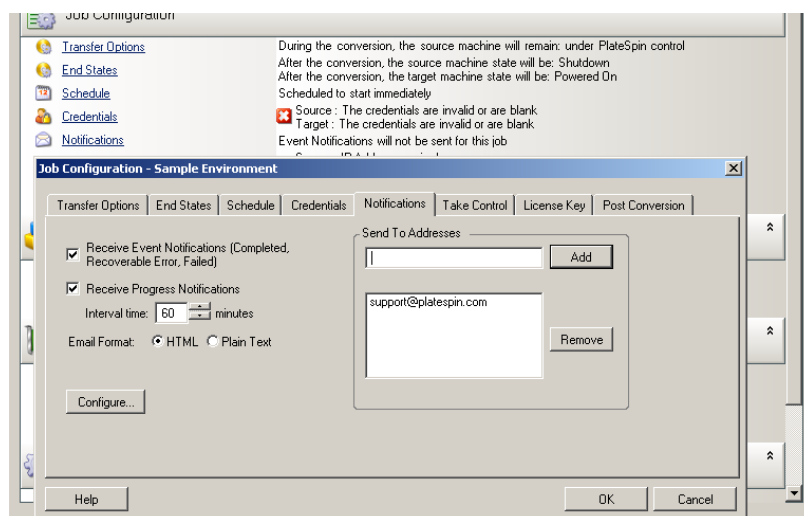
- **In Wizard mode:** Not available.

5.3 Specifying Transfer Options

Transfer options enable you to specify:

- How data is transferred from source to target. Portability Suite supports multiple transfer methods, and their availability depends on your workload and conversion job type. See "Supported Transfer Methods" on page 16.

- 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 53.

- Whether to encrypt the data transferred from source to target. See "Security and Privacy" on page 22.

To specify the transfer options for a conversion job:

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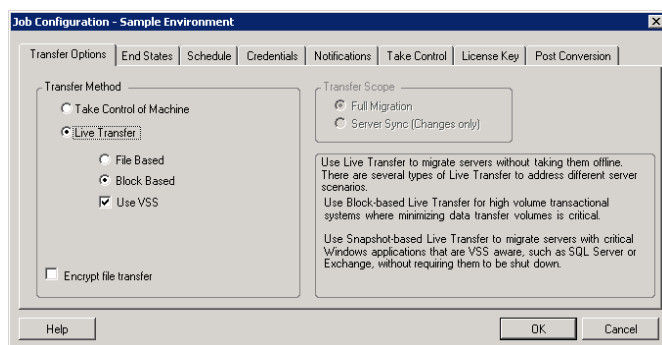
- **In Wizard mode:** In the wizard's navigation pane, click *Transfer Method*.

5.4 Specifying End States for Source and Target Workloads

You can specify a non-default post-conversion end state for your source and target. After completing the conversion, Portability Suite shuts down or boots the selected workloads as specified. For information about the default end states, see "Default Post-Conversion End States of Source and Target Workloads" on page 21.

To specify the required post-conversion state of a source and a target:

- **In Advanced mode:** In the Conversion Job window, under the Job Configuration section, click *End States*.

- **In Wizard mode:** Not available; source and target end states are preconcerted to their defaults.

5.5 Specifying Source and Target Workload Credentials

For a conversion job to execute properly, you must provide valid credentials for your source and target.

Use the following guidelines:

Table 5-1 *Source and Target Credentials*

Credentials	Windows	Linux
User Name	Username of account with local or domain-level administrative privileges. Use this syntax: ◆ 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 conversion job, you can validate the provided credentials and save them for future conversion jobs that use the same source and target.

To specify source and target credentials:

- **In Advanced mode:** In the Conversion Job window, under the Job Configuration section, click *Credentials*.
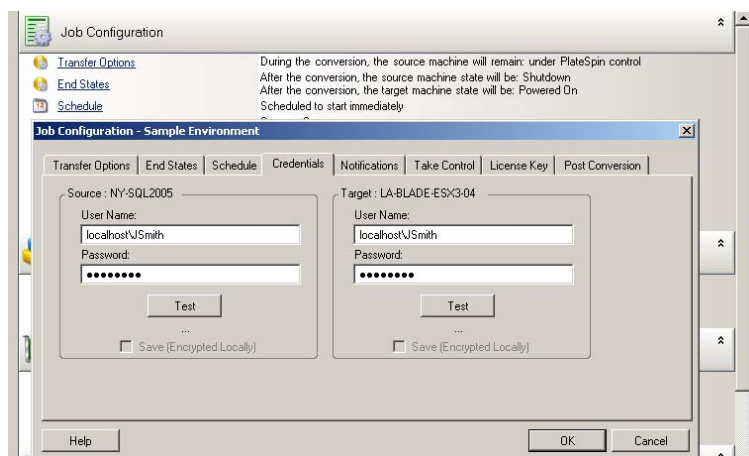
- **In Wizard mode:** In the wizard's navigation pane, click *Credentials*.
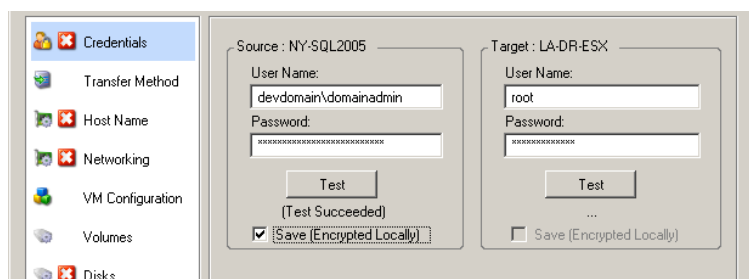
5.6 Managing a Workload's Network Identity

Portability Suite enables you to manage the network identity and domain registration of your conversion target workload and specify related preferences as part of a conversion 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 conversion job.

Proper configuration of conversion 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 Conversion 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.

- Section 5.6.1, "Managing the Identity of Windows Workloads," on page 82
- Section 5.6.2, "Managing the Network Identity of Linux and Solaris Workloads," on page 82

5.6.1 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 conversion. 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.
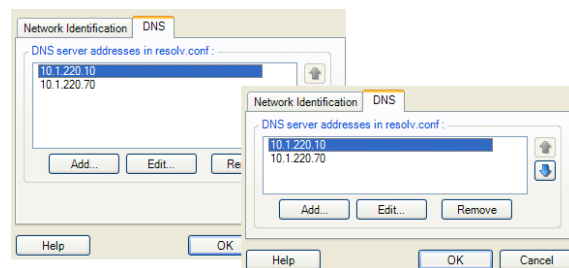
5.6.2 Managing the Network Identity of Linux and Solaris Workloads

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).
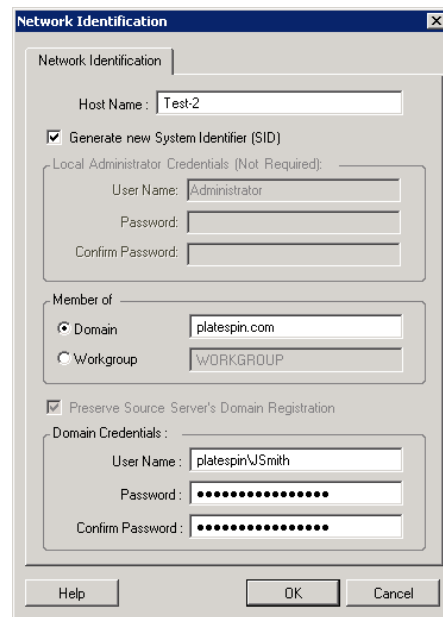
Network Identification tab: Specify the desired hostname for the target server.

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

5.7 Scheduling and Synchronization

An important element of a conversion job is scheduling. Scheduling options enable you to specify:

◆ Whether the conversion job is to run immediately or on a specific date and a specific time. For information on other options that Portability Suite provides for saving and executing jobs, see "Setting Up, Executing, and Managing Jobs" on page 42.

◆ Whether to enable incremental synchronization of the target with the source. This establishes a protection contract for workload or workload image protection jobs with two synchronization schedules: full and incremental.

To access scheduling options of a conversion job:

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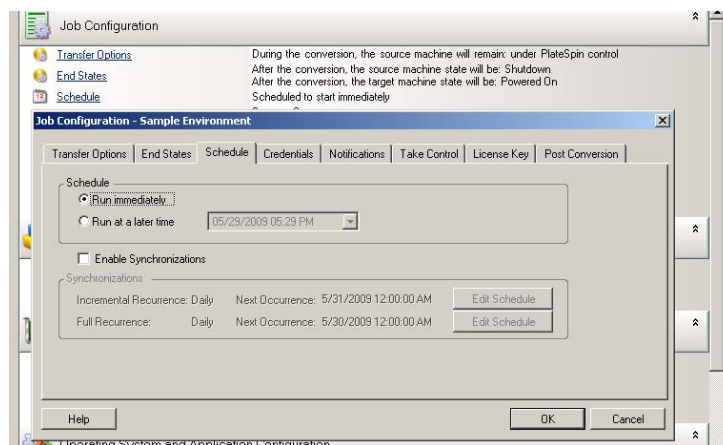
◆ **In Wizard mode:** In the wizard's navigation pane, click *Schedule*.

To edit either schedule, select the *Enable Synchronization* option, then click *Edit Schedule*.

NOTE: Synchronizations require the Live Transfer method. If Live Transfer is not supported for a particular conversion, this option is disabled.

If the schedules of an incremental and a full synchronization coincide, the full synchronization takes precedence. For example, if a full synchronization is scheduled for every Sunday at 1:00:00 a.m., and an incremental synchronization is scheduled to run every day at 1:00:00 a.m., on Sundays at 1:00:00 a.m. only a full synchronization runs.

5.7.1 Working with the Recurrence Editor

The Recurrence Editor enables you to specify scheduling options for synchronizing your target VM or Flexible Image with changes on your source workload. Depending on your conversion requirements, you can specify a schedule for a full synchronization or incremental synchronization.

To edit either schedule, click *Edit Schedule* in your job's schedule. Options displayed in the Recurrence Editor dialog box depend on prior job configuration selections.

These settings control the recurrence pattern of incremental and full transfers.

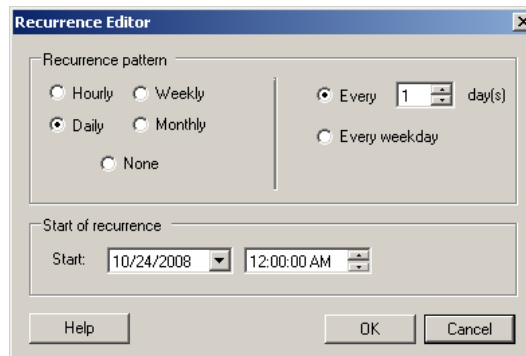*Hourly:* Use this option to set the hourly interval for your transfers.

Daily: Use this option to set the number of days between transfers.

Weekly: Use this option to schedule the number of weeks between transfers and on what days of the week to do the transfer.

Monthly: Use this option to define a monthly interval for incremental transfers.

None: Select this option to disable incremental or full transfers; for example, when you are planning to perform transfers at a later time or manually.

Start of Recurrence: Specify the start date and time for the transfers.

5.8 Setting Up Conversion Networking

For each workload portability and protection job, you must properly configure workload networking so that:

- Source workloads and targets can communicate with each other and Portability Suite Server during the conversion process.
- The network configuration of a target workload is in line with its end state.

A job configuration interface provides configuration settings for:

Take Control Networking: Also called *Temporary Network Settings*; they apply to source and target workloads booted into a temporary pre-execution environment. See "Offline Transfer with the Take Control Mechanism" on page 17.

Target Post-conversion Networking: In peer-to-peer conversion and image deployment jobs, they apply to target virtual or physical NICs.

- Section 5.8.1, "Take Control (Temporary) Network Settings," on page 84
- Section 5.8.2, "Target Post-conversion Networking," on page 87
- Section 5.8.3, "TCP/IP and Advanced Network Settings," on page 88

5.8.1 Take Control (Temporary) Network Settings

Take Control (Temporary) Network Settings control how source workloads, targets, and the Portability Suite Server communicate among each other during the conversion. 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 conversion.

During Windows and Linux workload conversions, the Take Control (Temporary) Network Settings control the Portability Suite Server's communication with the source and target workloads that are booted into a temporary pre-execution environment. See "Offline Take Control Transfer of Windows and Linux Workloads" on page 17.

During Solaris workload conversions, 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 .

To configure Take Control (Temporary) network settings:

- **In Advanced mode:** In the Conversion 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.
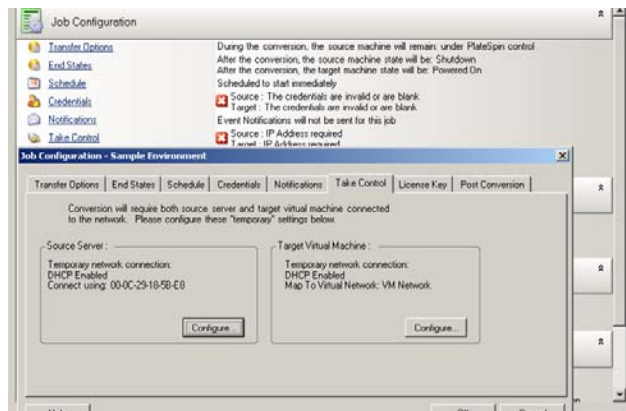
- **In Wizard mode:** In the wizard's navigation pane, click *Networking*. In the *Source Take Control* or *Target Take Control* row, in either *Temp IP* or *Map To* drop-down menu, select *Configure*. To quickly select DHCP without opening configuration options, select *DHCP*.
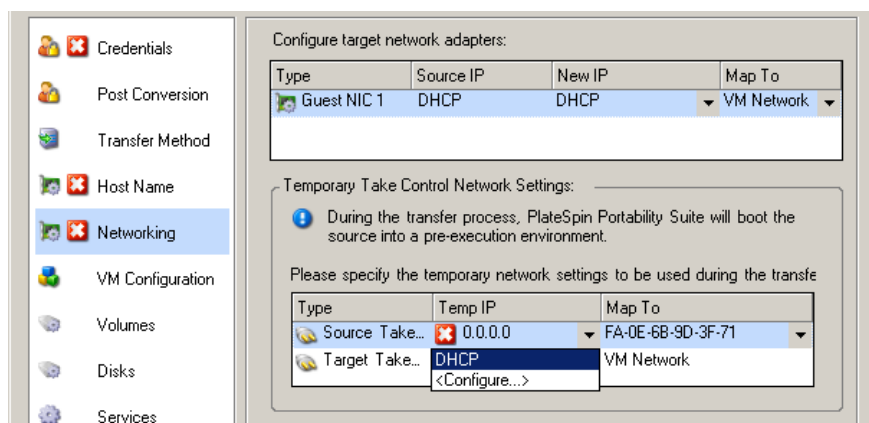
Configuration options for the Take Control (Temporary) networking vary and depend on whether the network interface is virtual or physical, and whether it is connecting a Windows or a Linux workload.

Target Take Control network settings are only used during a Take Control conversion process. On completion, target network settings are read from settings you specify for Target Post-conversion Networking. See "Target Post-conversion Networking:" on page 84.

- ◆ "Take Control (Temporary) Network Settings: Physical Network Interfaces" on page 86
- ◆ "Take Control (Temporary) Network Settings: Virtual Network Interfaces" on page 86

Take Control (Temporary) Network Settings: Physical Network Interfaces

These settings apply only to source physical machines. For target physical machines, Take Control (Temporary) Network settings are configured during the boot process that uses the Take Control ISO image. See "Discovering and Registering Target Physical Machines" on page 35.

Connect using: If multiple network adapters are present, select the adapter that can communicate with both the Portability Suite 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 to access TCP/IP and advanced network settings. See "TCP/IP and Advanced Network Settings" on page 88.
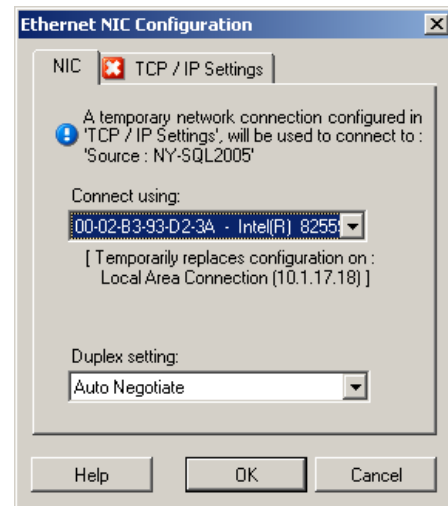
Take Control (Temporary) Network Settings: Virtual Network Interfaces

These settings apply to both source and target Take Control (Temporary) network settings.

Map to Virtual Network: From the drop-down list, select the virtual switch or network to use for communication during a Take Control conversion. If multiple virtual network adapters are present, select the adapter that can communicate with both the Portability Suite Server and the source machine. This network can differ from the network on which the target virtual machine will run after the conversion.

TCP/IP Settings tab: Click to access TCP/IP and advanced network settings. See "TCP/IP and Advanced Network Settings" on page 88.

5.8.2 Target Post-conversion Networking

Target post-conversion network settings defined in a conversion job control the network configuration of a target after the conversion is complete. This applies to both physical and virtual network interfaces.

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

To configure target post-conversion network settings:

- **In Advanced mode:** In the Conversion 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-conversion 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-conversion Networking for Physical Network Interfaces (Windows and Linux)" on page 87
- "Post-conversion Networking for Virtual Network Interfaces (Windows and Linux)" on page 88

Post-conversion Networking for Physical Network Interfaces (Windows and Linux)

Use these settings to configure the post-conversion 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 Portability Suite Server.

TCP/IP Settings tab: Click to access TCP/IP and advanced network settings. See "TCP/IP and Advanced Network Settings" on page 88.
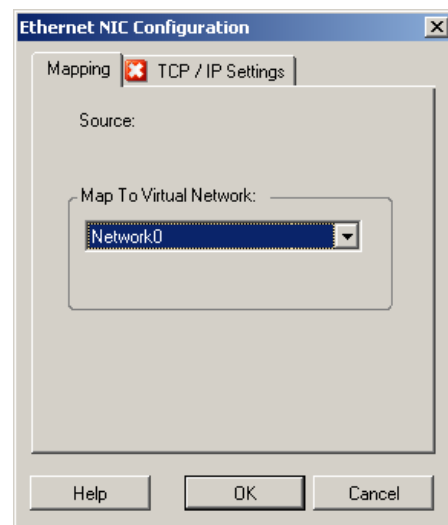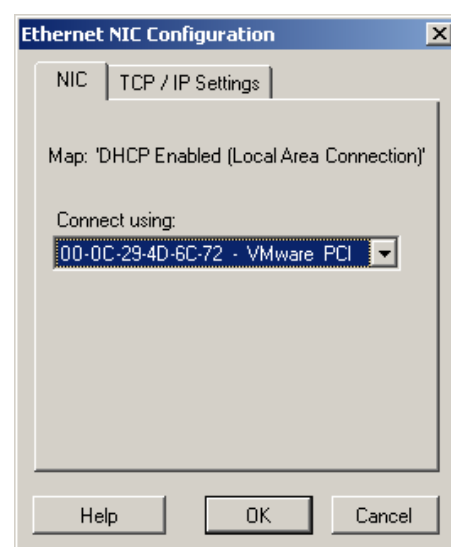
Post-conversion Networking for Virtual Network Interfaces (Windows and Linux)

By default, Portability Suite configures a conversion job to create a virtual NIC for each NIC found on the source. For post-conversion 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, Portability Suite creates a virtual NIC for a source NIC.

Map to Virtual Adapter: Select the virtual adapter that will be used on the target machine after the conversion is complete.

Start connected: Enable this option to connect the virtual network interface when starting the target machine.

TCP/IP Settings tab: Click to access TCP/IP and advanced network settings. See "TCP/IP and Advanced Network Settings" on page 88.

5.8.3 TCP/IP and Advanced Network Settings

Portability Suite provides a standard network configuration interface to both source and target network settings, and for both Take Control (Temporary) and target post-conversion networking. Configuration settings vary slightly, depending on the operating system.
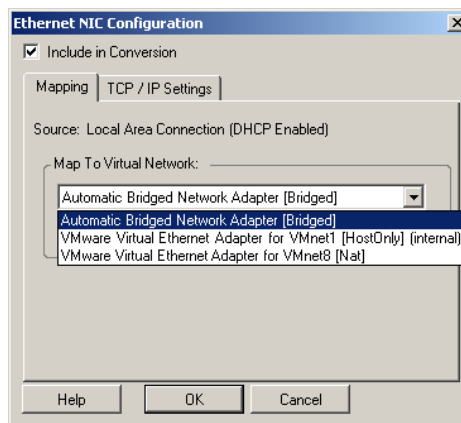
- ◆ "TCP/IP and Advanced Network Settings (Windows)" on page 88
- ◆ "TCP/IP and Advanced Network Settings (Linux and Solaris)" on page 89

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 conversion 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 option to access advanced TCP/IP configuration settings, then specify or edit default gateway, DNS server, and WINS server information as required.

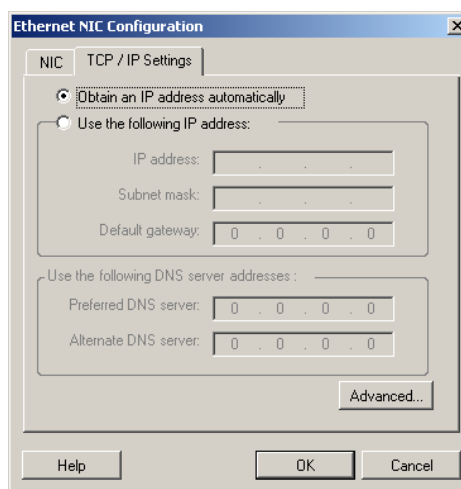**TCP/IP and Advanced Network Settings (Linux and Solaris)**

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 conversion 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.
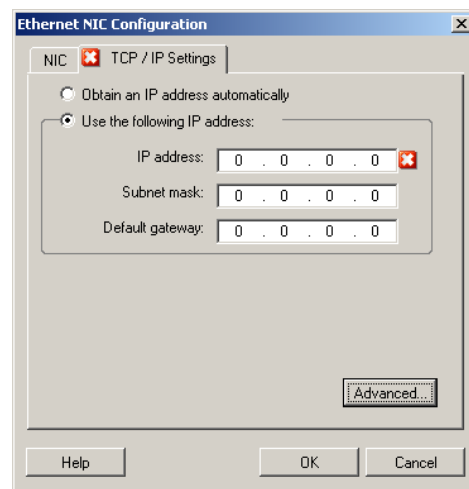
5.9 Configuring Target Virtual Machines

For jobs that involve workload virtualization, Portability Suite 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 it 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 conversion. 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 Conversion 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, Portability Suite 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.
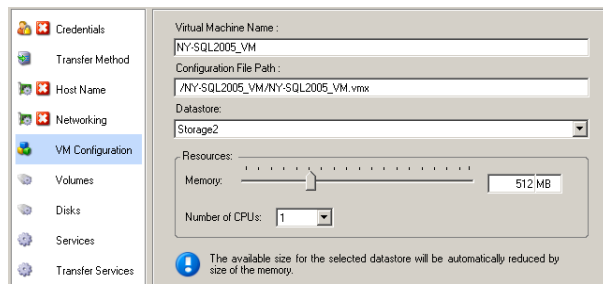
5.9.1 Virtualization Platform-Specific and Advanced VM Configuration Options

In advanced mode, Portability Suite displays target virtual machine configuration options specific to the selected target and provides access to more advanced configuration options.

Virtual Machine Configuration: MSVS

The following are configuration options specific to Microsoft Virtual Server.

Virtual Machine Name: Specify the display name for the new virtual machine.

Configuration Folder: Select the directory where the configuration files will be stored for the virtual machine.

Virtual Machine Memory Allocation: Specify a value for the amount of virtual RAM to be assigned to the virtual machine.

Install Virtual Machine Additions: Enable this option to install VM Additions during the conversion process.
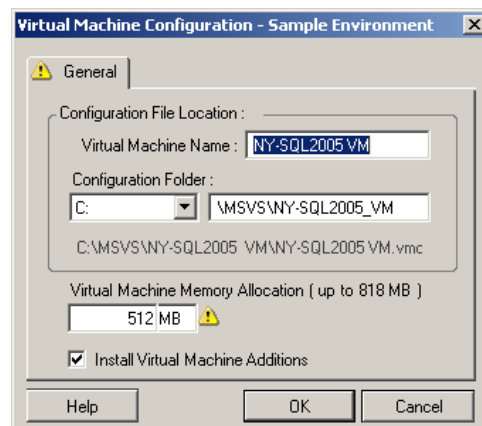
Virtual Machine Configuration: VMware ESX 2.x

The following are configuration options specific to VMware ESX 2. To access settings for the ESX 2 memory allocation, the number of CPUs, and CPU scheduling affinity, click *Advanced*.

Virtual Machine Name: Specify the display name for the new virtual machine.

Configuration File Path: Select the path to the directory where the VM's configuration files will be stored.

Virtual Machine Memory Allocation: Specify the amount of virtual RAM to assign to the VM.

Install VMware Tools: Enable this option to install VMware Tools during the conversion process (recommended).

Network Adaptors: Select either *vlance* or *vmxnet*. *Vmxnet* is the recommended selection for optimum performance.

SCSI Drives: Select either *BusLogic* (most conversions) or *LSIlogic* (recommended for Windows 2003 workloads).

Advanced: Click this button to view or modify advanced VM configuration settings.
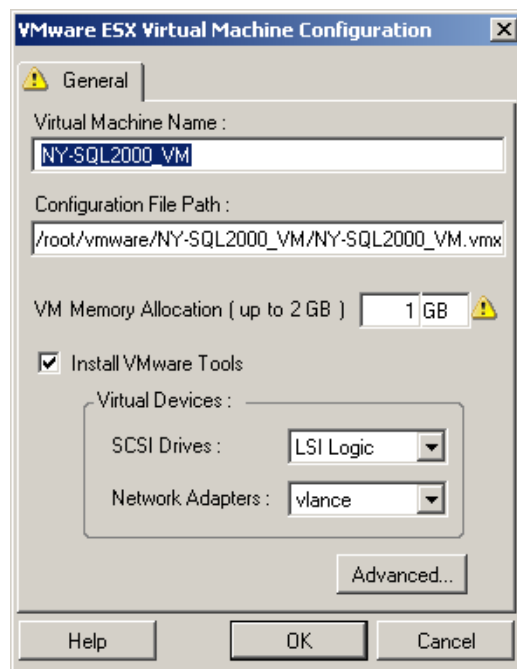
Memory Allocation: Minimum and Maximum memory allocation controls how memory is distributed among different VMs on the VM host. *Shares* define the memory allocation priority of a given VM relative to other VMs. In most cases you shouldn't need to change these settings.

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.

CPU Allocation: The minimum percentage represents an absolute, fixed lower limit. The maximum percentage represents an absolute, fixed upper limit. A VM uses at least as much CPU time as specified by the minimum percentage and never uses more CPU time than the specified maximum percentage.

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.
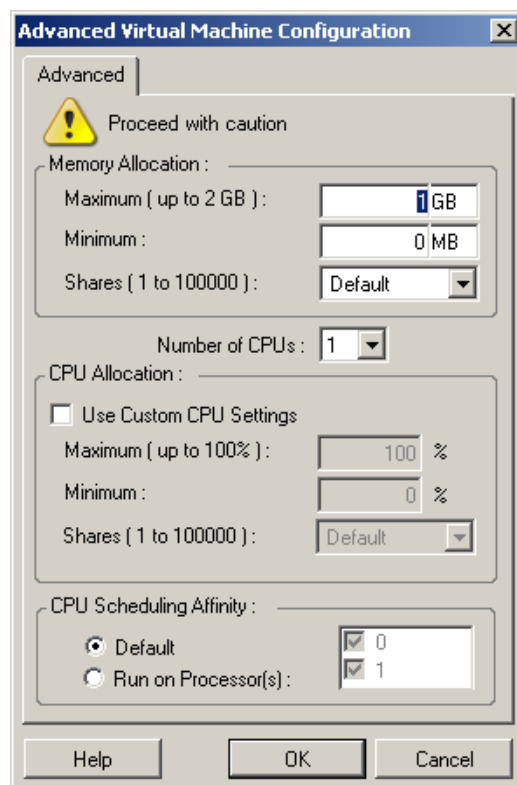
Virtual Machine Configuration: VMware ESX 3/3.5

The following are configuration options specific to VMware ESX 3 and 3.5. To access settings that control resource pools, the number of CPUs, and CPU scheduling affinity, click *Advanced.*

Virtual Machine Name: Specify the display name for the new virtual machine.

Datastore: Select the datastore where you want to create the *.vmx file.

Configuration File Path: Specify a name and the directory path for the virtual machine's *.vmx configuration file.

Virtual Machine Memory Allocation: Specify a value for the amount of virtual RAM to be assigned to the virtual machine.

Install VMware Tools: Enable this option to install VMware tools during the conversion process (recommended).

SCSI Drives: Select either *BusLogic* (most conversions) or *LSIlogic* (recommended for Windows 2003 workloads).

*Advanced*: Click this button to view or modify advanced VM configuration settings.
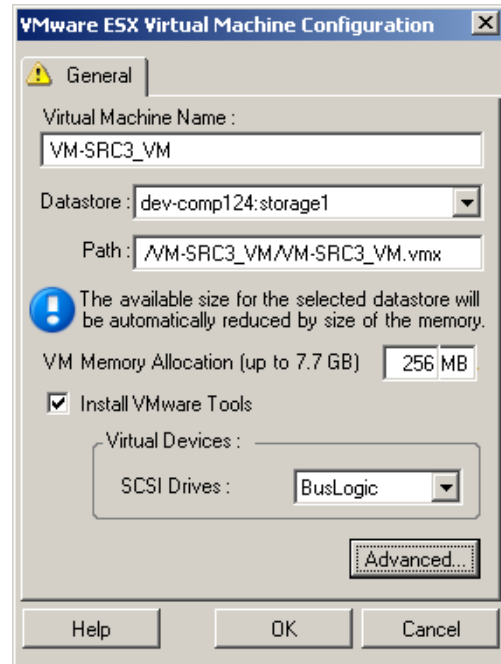
Resource Pool: If required, assign your target VM to a resource pool. When no resource pool is specified, the VM is assigned to the root resource pool.

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.
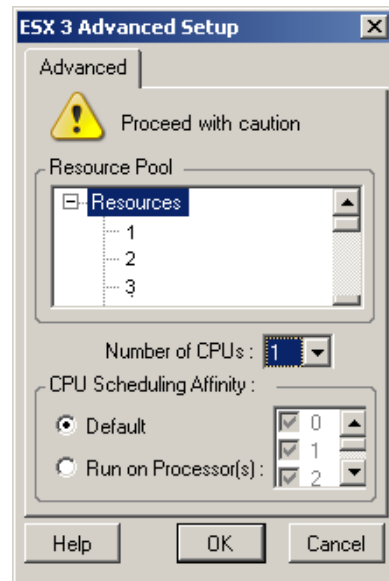
Virtual Machine Configuration: Citrix XenServer
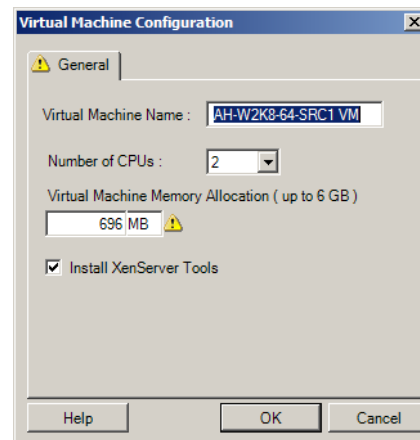
The following are configuration options specific to Citrix XenServer.

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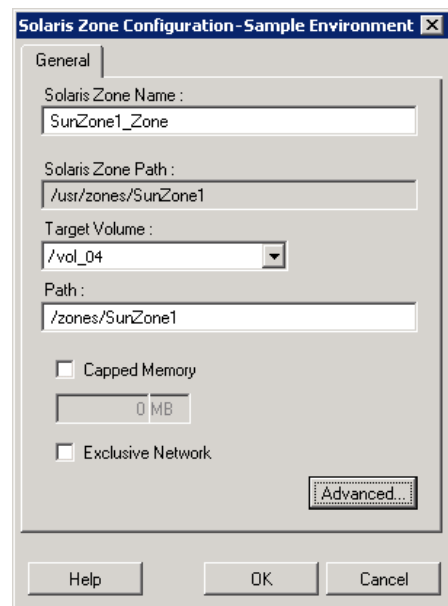
Target Volume: This drop-down list contains discovered volumes on the target zone host. Select the required volume for the target zone. Do not use system hierarchies, such as / usr, /sbin, /lib, /platform, and /var.

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 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, and specify the required number of CPU shares in the corresponding field.

5.10 Handling Operating System Services and Applications

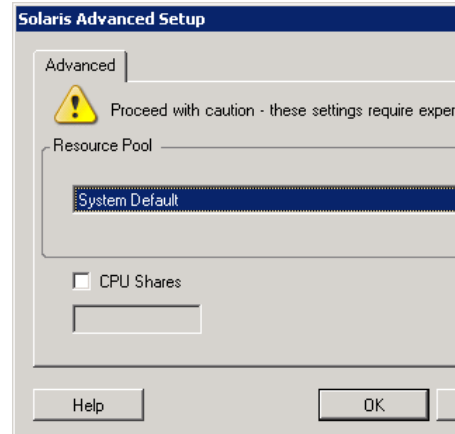Portability Suite provides a mechanism for conversion jobs to handle Windows services (including special features for those related to Microsoft SQL Server and Microsoft Exchange Server software), Linux daemons, Solaris services, and virtualization enhancements (such as VMware Tools). In addition, when Windows HAL or kernel file replacements are required, you can view which ones were selected for update during a conversion job.

- Section 5.10.1, "Handling the Startup Mode of Services (Windows Targets)," on page 94
- Section 5.10.2, "Handling Services During Live Transfer (Windows Source Workloads)," on page 95
- Section 5.10.3, "Viewing Windows System Files Selected for Replacement During a Conversion," on page 96
- Section 5.10.4, "Handling the Run Level of Daemons (Linux Targets)," on page 97
- Section 5.10.5, "Handling Services on Solaris Targets," on page 98
- Section 5.10.6, "Handling Virtualization Enhancement Software," on page 99

5.10.1 Handling the Startup Mode of Services (Windows Targets)

You can configure a job to modify the startup mode of selected Windows services after the conversion 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 conversion is complete. For example, you might require a virus scanner to be disabled during the conversion, but its startup type to be restored after the conversion completes.

TIP: You can globally configure your startup mode preferences for selected Windows services in Portability Suite Server default options. See "Target Service Defaults" in your *Configuration Guide*.

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

- ◆ **In Advanced mode:** In the Conversion Job window, under the Operating System and Application Configuration section, click *Windows Services (Target)*, then click an item in the *Start Mode* column.
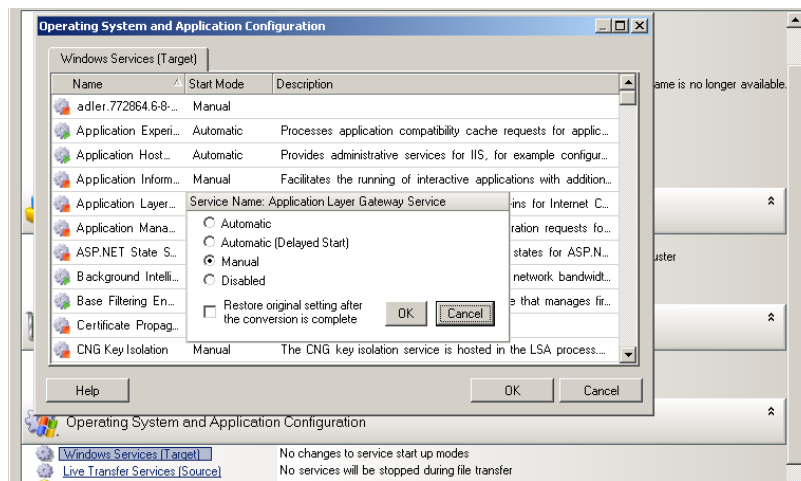
- ◆ **In Wizard mode:** In the wizard's navigation pane, click *Services*, then click an item in the *Target Mode* column.
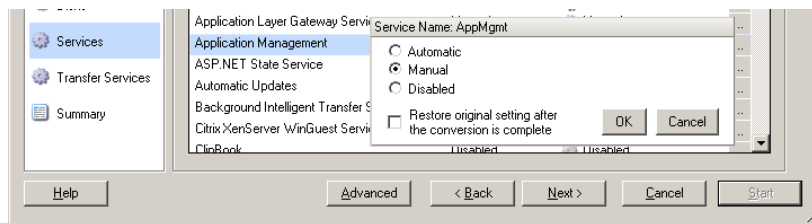
5.10.2 Handling Services During Live Transfer (Windows Source Workloads)

For jobs that are using the VSS File-based and VSS Block-based Live Transfer methods, Portability Suite provides a mechanism to stop selected services during the conversion. 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 conversion job to automatically copy the database files of these servers. If you do not require the conversion 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 conversion.

After the completion of the conversion, services that you select to stop during a Live Transfer conversion restart on the source automatically, unless you explicitly configure your conversion job to power off the source on completion.

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" in your *Configuration Guide*.

To specify which Windows services to stop during a VSS-aware Live Transfer:

- ◆ **In Advanced mode:** In the Conversion Job window, under the Operating System and Application Configuration section, click *Live Transfer Services (Source)*. To indicate that you want SQL Server and Exchange Server database files copied during the conversion, click *Advanced*.
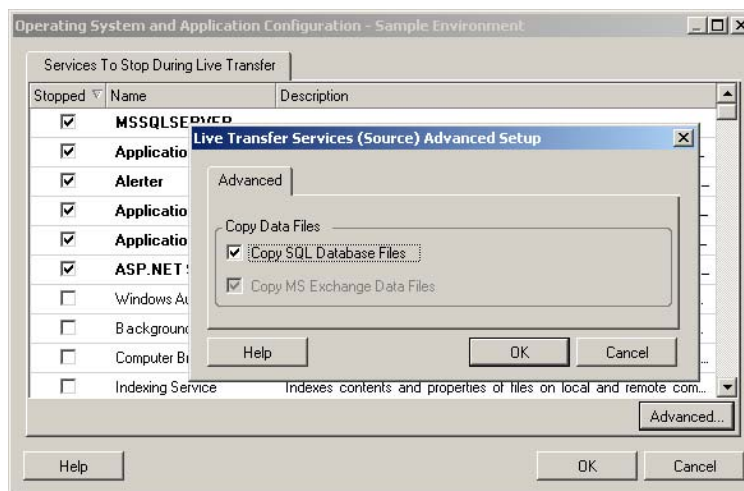
- ◆ **In Wizard mode:** In the wizard's navigation pane, click *Transfer Services*.
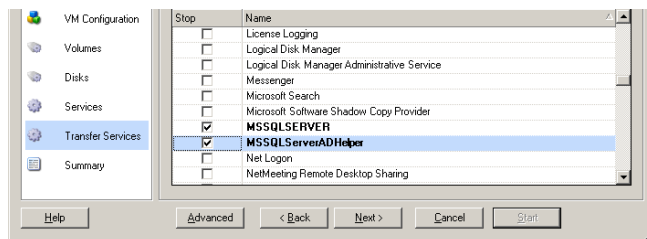
5.10.3 Viewing Windows System Files Selected for Replacement During a Conversion

When converting Windows workloads with system files (such as a HAL or kernel files) that are incompatible with the target infrastructure, Portability Suite 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 Portability Suite identifies as those requiring replacement:

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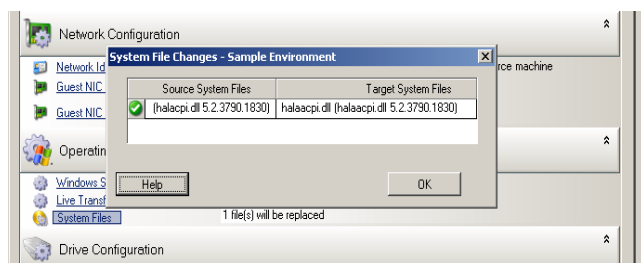
◆ **In Wizard mode:** Not available.

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	Portability Suite has a partial match with its matrix but the driver cache contains a higher version of one or more system files than what Portability Suite will use.
File <filename> will be replaced with lower version	Portability Suite has not found a match for the system files in its matrix. It will replace the system files with a version that is lower than what was discovered as the source machine's original system files.
File <filename> will be replaced with higher version	Portability Suite has not found a match for the system files in its matrix. It will replace the system files with a version that is higher than what was 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 Knowledge Base Articles:

◆ Q20815 (http://support.platespin.com/kb2/article.aspx?id=20815): FAQ: Understanding the System Files Information screen

◆ Q20817 (http://support.platespin.com/kb2/article.aspx?id=20817): INFO: How to Restore the Driver Cache on a source machine

5.10.4 Handling the Run Level of Daemons (Linux Targets)

You can configure a job to modify the run level of selected Linux daemons after the conversion 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-conversion run level of Linux daemons:

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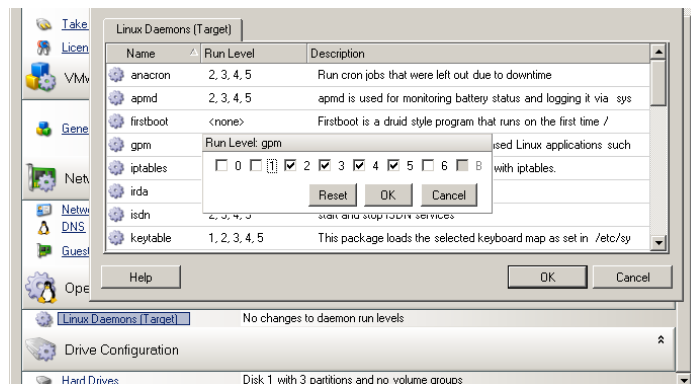
◆ In Wizard mode: In the wizard's navigation pane, click *Daemons*, then click an item in the *Target Run Level* column.
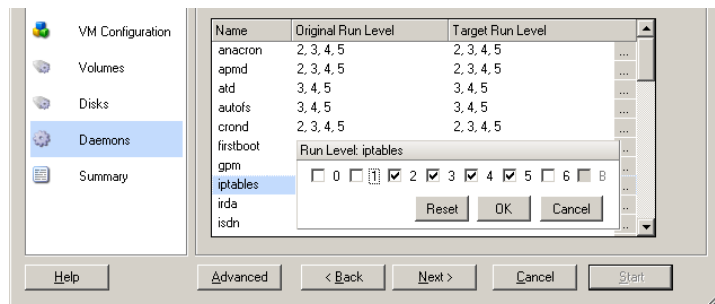
5.10.5 Handling Services on Solaris Targets

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

To enable or disable the required Solaris services:

◆ **In Advanced mode:** In the Conversion 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.
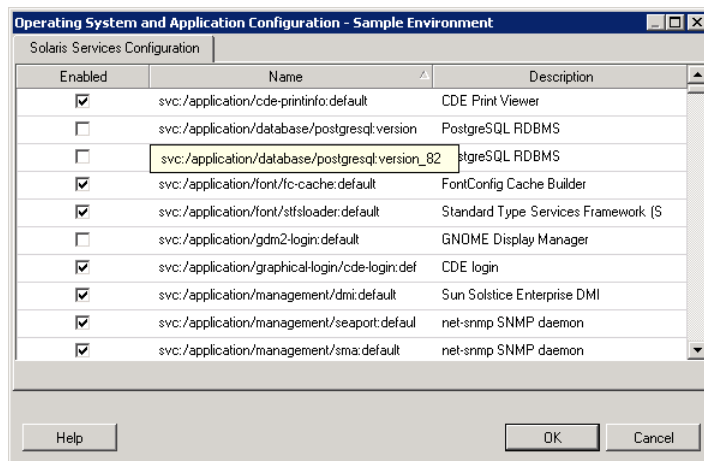
◆ In Wizard mode: Not available.

5.10.6 Handling Virtualization Enhancement Software

For V2X conversions, Portability Suite 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, Portability Suite identifies the presence of obsolete software and adds a dedicated VMware Tools Cleanup step in the conversion 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, Portability Suite 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 conversion:

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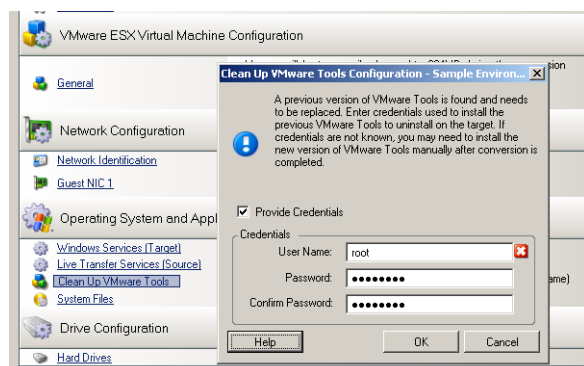
- **In Wizard mode:** In the wizard's navigation pane, click *VM Tools*.
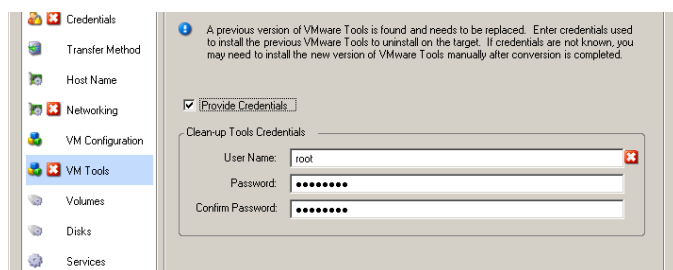
Depending on the target, Portability Suite identifies existing instances of VMware Tools and prompts to either replace or remove them, as applicable:

- **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 conversion.

- **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 conversion completes.

5.11 Handling a Workload's Storage Media and Volumes

Portability Suite provides mechanisms for configuring your conversion 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

Portability Suite 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, Portability Suite successfully completes the migration. Portability Suite does not support software implementations of RAID.

SAN Storage

Portability Suite supports SAN storage. As long as the driver for the associated host bus adapter (HBA) is present, Portability Suite successfully completes the migration. Portability Suite treats SAN LUNs like any other disk with logical volumes.

NAS

Portability Suite supports Network Attached Storage (NAS) systems. Portability Suite treats NAS like any other disk with logical volumes.

Windows Dynamic Disks

Portability Suite supports Windows dynamic disks, including mirrored, spanned, and RAID 5 configurations.

Portability Suite 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 conversion, you can upgrade the required disks on your targets from basic to dynamic disks.

Linux Logical Volumes

Portability Suite 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: Portability Suite does not support:

- LVM snapshots. You can create LVM snapshots on your target after the migration completes.
- LVM mirrors. You can re-create mirrored logical volumes on the target after the conversion completes.

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

- Section 5.11.1, "Storage Layout and Volume Configuration (Wizard Mode)," on page 102
- Section 5.11.2, "Drive Configuration (Advanced Mode)," on page 103
- Section 5.11.3, "Volume Configuration in Server Sync," on page 109
- Section 5.11.4, "Handling Volumes and Image Files in X2I (Imaging) Conversions," on page 110

5.11.1 Storage Layout and Volume Configuration (Wizard Mode)

- "Storage Layout Configuration: X2V (Wizard Mode)" on page 102
- "Volume Configuration: X2V, X2P (Wizard Mode)" on page 103
- "Linux Volume Groups Configuration (Wizard Mode)" on page 103

Storage Layout Configuration: X2V (Wizard Mode)

Use this Conversion 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 conversion. 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.
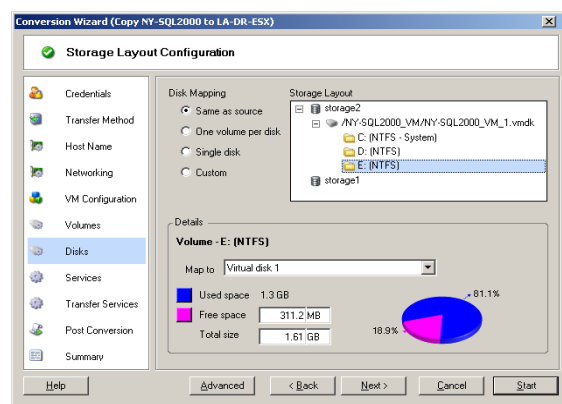
Volume Configuration: X2V, X2P (Wizard Mode)

Use this Conversion Wizard page to select volumes to include in the conversion 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.
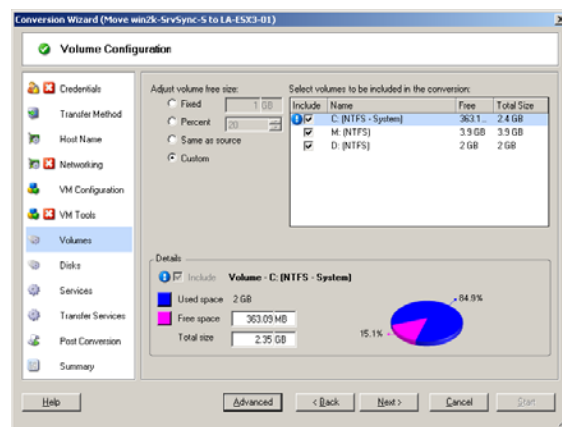
Linux Volume Groups Configuration (Wizard Mode)

Use this Conversion 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 conversion, 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 conversion.
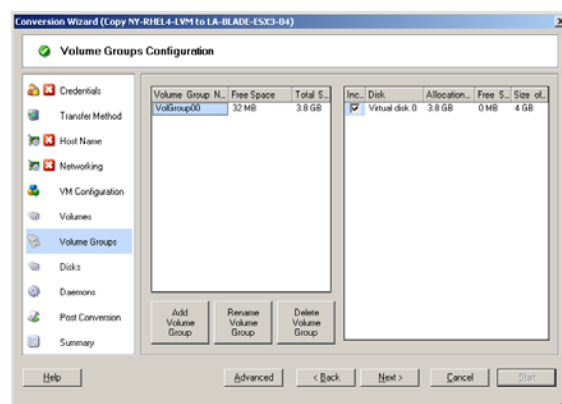
5.11.2 Drive Configuration (Advanced Mode)

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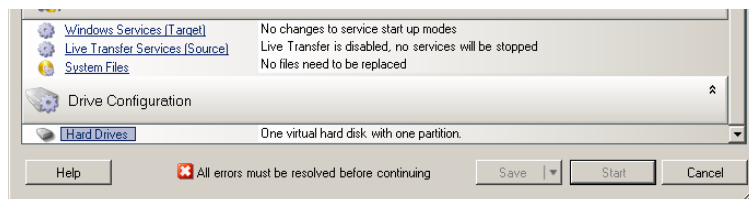
Settings vary depending on the target system.

◆ "Windows X2P Drive Configuration (Advanced Mode)" on page 104
◆ "Linux X2P Drive and LVM Volume Configuration (Advanced Mode)" on page 104
◆ "Target VM-Specific P2V/V2V Drive Configuration (Advanced Mode)" on page 106
◆ "Solaris Zone Volume Configuration" on page 109

Windows X2P Drive Configuration (Advanced Mode)

Use these settings to select the volumes to copy during the conversion:

Copy: Select the volumes to be copied during the conversion.

New Free Space: To resize the volume during the conversion, specify the desired amount of free space. Portability Suite automatically adjusts *New Size*.

New Size: To resize the volume during the conversion, specify the desired size. Portability Suite 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 conversion. If the partitions are not selected, Portability Suite permanently removes the partitions from the server.
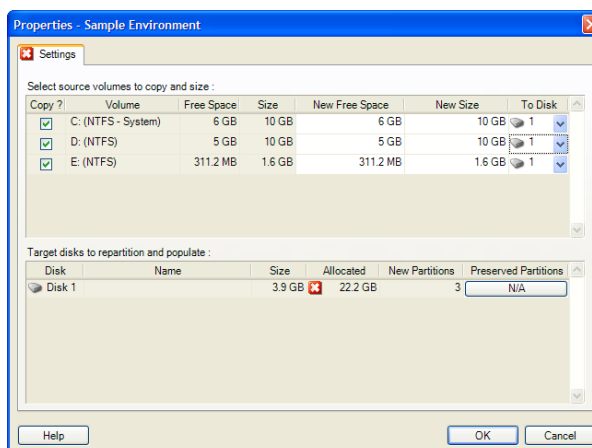
Linux X2P Drive and LVM Volume Configuration (Advanced Mode)

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

◆ "Linux X2P Drive and LVM Volume Configuration (Settings Tab)" on page 105
◆ "Linux Drive and LVM Volume Configuration (Volume Groups Tab)" on page 105

Linux X2P Drive and LVM Volume Configuration (Settings Tab)

Use these settings to select source volumes to copy, non-volume source spaces to recreated 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 conversion.

New Free Space: To resize the volume during the conversion, enter the desired amount of free space. Portability Suite automatically adjusts *New Size*.

New Size: To resize the volume during the conversion, enter the desired size. Portability Suite 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 conversion. If the partitions are not selected, Portability Suite permanently removes them from the server.

Linux Drive and LVM Volume Configuration (Volume Groups Tab)

Use these settings to manage volume groups.

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

Rename Volume Group: Renames a volume group that is being copied from the source to the target.

Delete Volume Group: Deletes a volume group so that it is not created on the target machine. The volumes assigned to the volume group can be reassigned to other locations by using the *Settings* tab (by default, they are assigned to disk).

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

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

When configuring 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: Microsoft Virtual Server, VMware Server/GSX" on page 106
- "Drive Configuration: VMware ESX 3" on page 107
- "Drive Configuration: VMware ESX 2" on page 108

Drive Configuration: Microsoft Virtual Server, VMware Server/GSX

The following are drive configuration settings specific to Microsoft Virtual Server and VMware Server (formerly GSX):

Copy: Select the volumes to be copied during the conversion.

New Free Space: To resize the volume during the conversion, specify the desired amount of free space. Portability Suite automatically adjusts *New Size*.

New Size: To resize the volume during the conversion, specify the desired size. Portability Suite automatically adjusts *New Free Space*.
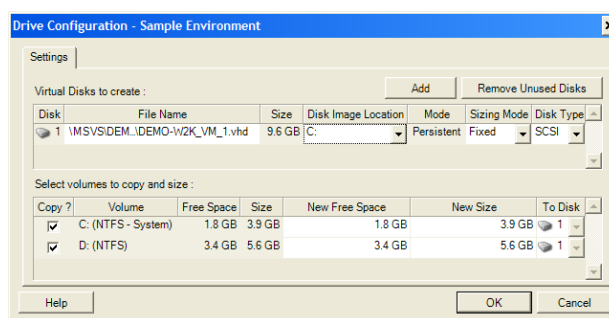
NOTE: Operating systems have partition size limitations. For instance, Windows NT 4 SP4 and SP6 targets must not exceed 7.8 GB.

To Disk: Select which hard drive the volume will be copied to on the physical target server.

Disk Type: IDE is permitted when the source machine uses IDE and the target is MSVS.

Drive Configuration: VMware ESX 3

The following are VMware ESX 3-specific drive configuration settings:

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 conversion.

New Free Space: To resize the volume during the conversion, specify the desired amount of free space. Portability Suite automatically adjusts *New Size*.

New Size: To resize the volume during the conversion, specify the desired size. Portability Suite 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 conversion, specify the desired size.

Drive Configuration: VMware ESX 2

The following are VMware ESX 2-specific drive configuration settings:

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. Portability Suite automatically adjusts the *New Size*.

New Size: To resize the volume during the conversion, specify the desired size. Portability Suite 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 conversion, specify the desired size.
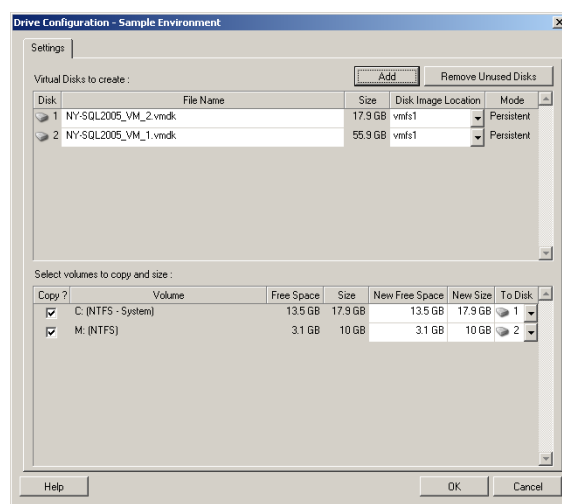
Solaris Zone Volume Configuration

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 (ZFS) on the target. For the *Zone FS 0* file system, you cannot modify the *Datastore* and *Location* values, which are computed 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 93).

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 conversion); it is always selected and assigned to the *Zone FS 0* file system.
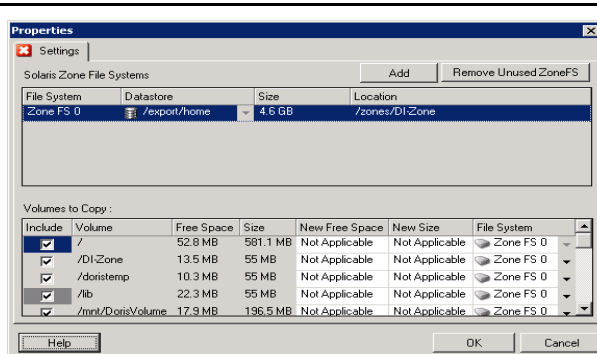
5.11.3 Volume Configuration in Server Sync

When you are using Server Sync to synchronize two Windows or Linux workloads, Portability Suite 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 53.

To access volume configuration options in a Server Sync job:

◆ **In Advanced mode:** In the Conversion 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)" on page 110
◆ "Server Sync Volume Configuration (Linux)" on page 110

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.
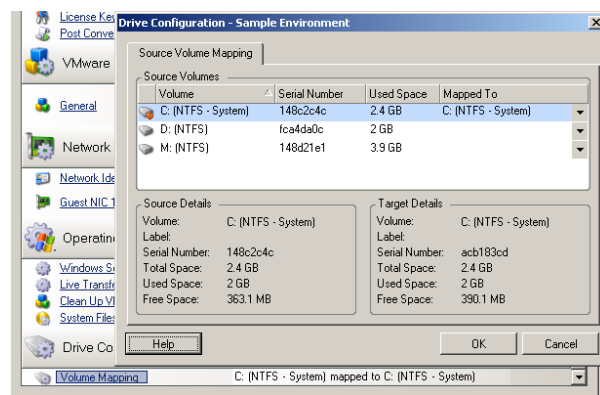
Server Sync Volume Configuration (Linux)

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.
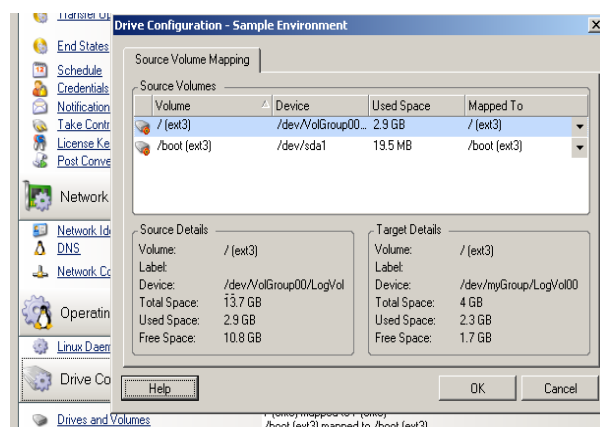
5.11.4 Handling Volumes and Image Files in X2I (Imaging) Conversions

When capturing a Flexible Image or importing volumes into a Flexible Image, Portability Suite 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 111
- "Image Configuration: X2I (Wizard Mode)" on page 111

Target Volume Configuration: X2I (Wizard Mode)

Use this page, when setting up a Capture Image or Import Image job in wizard mode, 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.
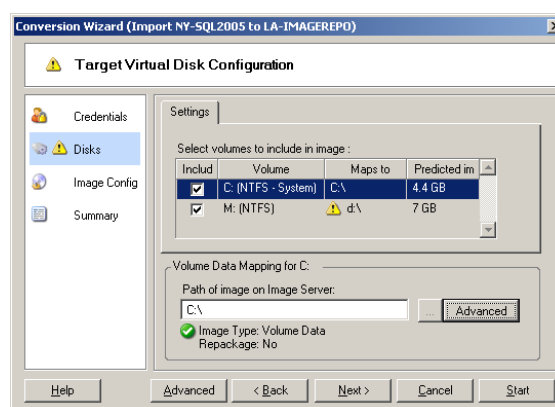
Image Configuration: X2I (Wizard Mode)

Use this page, when setting up a Capture Image or Import Image job in wizard mode, 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.
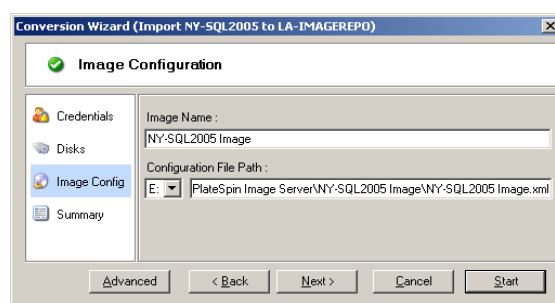
5.12 Including a Custom Post-conversion Action in a Conversion Job

You can set up your conversion 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 Post-conversion Actions" on page 40.

NOTE: Post-conversion actions are supported for peer-to-peer and one-time Server Sync conversions only.

When setting up a conversion 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.
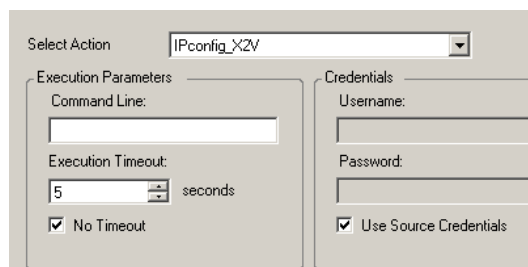
To specify a custom post-conversion action for your conversion job:

- ◆ **In Advanced mode:** In the Conversion 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-conversion 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*.

Select Action   IPconfig_X2V   ▼

Execution Parameters

   Command Line:

   Execution Timeout:

   5   ↕   seconds

   ☑ No Timeout

Credentials

   Username:

   Password:

   ☑ Use Source Credentials

Frequently Asked Questions

A

This section provides answers to frequently asked questions.

What are the performance and scalability characteristics of my Portability Suite product?

Your Portability Suite product's overall performance, including data transfer speeds and scalability, depend on a variety of factors in your specific environment. See "Performance and Scalability" on page 23.

How secure is my Portability Suite product?

Portability Suite provides several features to help you safeguard your data and increase security. See "Security and Privacy" on page 22.

Does Portability Suite support my workload's data storage technology?

Portability Suite 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 32.

Can multiple conversions run simultaneously?

Yes. See "Performance and Scalability" on page 23.

Troubleshooting Portability Suite

B

This section provides a series of topics about troubleshooting PlateSpin® Portability Suite.

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 Portability Suite 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 Portability Suite Server.
My physical server has completed the registration process, but is not seen in Portability Suite 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 Portability Suite Client.
Problems discovering source and target servers	Knowledge Base Article Q20291 (http://support.platespin.com/kb2/article.aspx?id=20291) contains troubleshooting checklists for discovering: ◆ 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$\{randomID}.xml"	This error might occur when you are attempting to discover a Windows NT 4.0 Server with an older version of WMI Core installed. It can also occur with a Windows Server 2000 or 2003. For Windows NT 4.0 servers, verify the version of WMI installed on the server by doing the following:
	1. Locate the `wbemcore.dll` file in the `\winnt\system32\wbem` directory.
	2. Right-click `wbemcore.dll` and click *Properties*.
	3. Check the version number of the file. If the version number is not 1.50.xxxx, upgrade the WMI core to V1.5. See the following Microsoft Web site (http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=C174CFB1-EF67-471D-9277-4C2B1014A31E)).
	After WMI v1.5 has been installed, rediscover the server.
	In some cases, reinstalling WMI 1.5 might resolve the issue. If this discovery error still occurs after installing WMI 1.5 on a Windows NT 4.0 Server or if it is still occurring when you attempt to discover a Windows 2000/2003 Server, do the following:
	1. Ensure that the `Admin$` 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 Portability Suite Server\Web` directory.
	3. Using a text editor, 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.
Rebooting the target machine as soon as the job indicates Recoverable Error status	During the configuration service step, there might be a delay when the Portability Suite Server receives notification from the machine that is being configured. The job then indicates a Recoverable Error status.
	It is not necessary to reboot the target machine. Wait at least 30 minutes before attempting to reboot or perform other interventions.

Related Knowledge Base Articles:

ID	Description
Q20339 (http://support.platespin.com/kb2/article.aspx?id=20339)	ERRMSG: Discovery fails with "The request failed with HTTP status 407" message
Q20894 (http://support.platespin.com/kb2/article.aspx?id=20894)	ERRMSG: Recoverable Error: ControllerConnectionBroken during discovery

ID	Description
Q20291 (http://support.platespin.com/kb2/article.aspx?id=20291)	ERRMSG: Server details discovery problems
Q20853 (http://support.platespin.com/kb2/article.aspx?id=20853)	FAQ: MSVS is discovered as a regular server in Portability Suite
Q20732 (http://support.platespin.com/kb2/article.aspx?id=20732)	FAQ: Network path not found error during discovery
Q20868 (http://support.platespin.com/kb2/article.aspx?id=20868)	INFO: Files and directories are not cleaned up after discovery of a Windows VM

B.2 Peer-to-Peer Conversions

Table B-2 *Common Issues and Solutions Related to Peer-to-Peer Conversions*

Problems or Messages	Solutions
One of the following errors displays during Take Control of source or target virtual/physical machine: ◆ 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 Portability Suite Server. ◆ The source/target machine was not able to fully boot into the Take Control 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. For a troubleshooting checklist and list of information required if technical support is necessary, see Knowledge Base Article Q20405 (http://support.platespin.com/kb2/article.aspx?id=20405).
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 conversion. 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 conversion. If this option is set incorrectly, a duplex mismatch occurs on the network.

Problems or Messages	Solutions
Unable to determine suitable boot partition	When converting existing source servers, the boot volume must pass the following checks:

<div style="margin-left:2em">

- 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 conversion fails while attempting to take control of the source server.

</div>

Problems or Messages	Solutions
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 Novell® PlateSpin® Operations Framework Controller on the Portability Suite Server. Use the Windows services plug-in to confirm that the Controller is running. See Knowledge Base Article Q20323 (http://support.platespin.com/kb2/article.aspx?id=20323) for other troubleshooting instructions.
Troubleshooting failures at the *Configuring Operating System* stage (also applicable to *Configure Target Machine* or *Configuring Virtual Machine* conversion 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 conversion job appears to have failed, the overall migration is probably successful and the configuration service running on the target will likely continue its operations. Knowledge Base Article Q20327 (http://support.platespin.com/kb2/article.aspx?id=20327) 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 "Live Transfer (File-Based)" on page 18.

Related Knowledge Base Articles:

ID	Description
Q20323 (http://support.platespin.com/kb2/article.aspx?id=20323)	ERRMSG: Portability Suite Job remains at a "Scheduled" or "Recoverable Error" state
Q20810 (http://support.platespin.com/kb2/article.aspx?id=20810)	INFO: Restore job stalls - "The configuration service in the target machine..."
Q20327 (http://support.platespin.com/kb2/article.aspx?id=20327)	INFO: Troubleshooting failures at "Configuring Operating System"
Q20405 (http://support.platespin.com/kb2/article.aspx?id=20405)	INFO: Waiting for Controller to Start - Controller Connection Not Established/Broken

ID	Description
Q20341 (http://support.platespin.com/kb2/article.aspx?id=20341)	INFO: What ports does Portability Suite use during discovery, conversion and file transfer?

B.3 Using Images

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

Problems or Messages	Solutions
Cannot see Flexible Images on Flexible 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 Flexible Image Server on Windows 2003. To resolve the error, use the Windows services plug-in on the Flexible Image Server. Modify the logon properties for the Portability Suite Operations Management Controller service to use an account with local administrative privileges. Restart the service after making this change.
Security descriptors are not intact on deployed server when you are using volume data from a Symantec Ghost image	When you are creating a Flexible 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, without having any knowledge of the original source server's permissions.

Related Knowledge Base Articles:

ID	Description
Q20879 (http://support.platespin.com/kb2/article.aspx?id=20879)	ERRMSG: The file cannot be accessed by the system
Q20865 (http://support.platespin.com/kb2/article.aspx?id=20865)	FAQ: How many images can an Image Server store?
Q20872 (http://support.platespin.com/kb2/article.aspx?id=20872)	HOWTO: Deploying physical Take Control CD using RIS (Remote Installation Services)

B.4 Diagnostics

Portability Suite provides a tool that can produce a diagnostics report for any running or completed job.

To view a diagnostics report:

1 In the Portability Suite 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 Novell® PlateSpin® technical support by following the instructions provided on the report.

Reference of Tables

C

This section provides a list of links to important tables used throughout the documentation accompanying PlateSpin® Portability Suite version 8.1 (the PlateSpin® Migrate and PlateSpin® Protect products).

Table	In Publication
"Portability Suite Roles and Permission Details"	*Portability Suite Configuration Guide*
"Network Communication Prerequisites for Discovery"	*Portability Suite Configuration Guide*
"Network Communication Prerequisites for Workload Migration and Protection"	*Portability Suite Configuration Guide*
Table 1-2, "Portability Suite Workload Conversion Tasks," on page 13	*Portability Suite User Guide*
Table 1-3, "Supported Microsoft Windows Workloads," on page 14	*Portability Suite User Guide*
Table 1-4, "Supported Linux Workloads," on page 15	*Portability Suite User Guide*
Table 1-5, "Supported Target Virtualization Platforms," on page 16	*Portability Suite User Guide*
Table 1-7, "Default Post-Conversion End States," on page 21	*Portability Suite User Guide*
Table 2-1, "Machine-Specific Icons in the Servers View," on page 30	*Portability Suite User Guide*
Table 2-2, "Take Control ISO Images for Target Physical Machines," on page 36	*Portability Suite User Guide*

Glossary

auto-discovery

See Network Discovery.

Block-based Transfer Component

A Portability Suite 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 Portability Suite job that captures an image of a workload in Flexible 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 Portability Suite 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 Portability Suite job that converts a Flexible Image to a bootable workload on physical hardware or on a virtual machine. *See also* Capture Image.

discover

See Discover Details.

Discover Details

A Portability Suite job that inventories the details of a supported workload or target machine in preparation for a migration or protection operation. *Compare with* Network Discovery.

File-based Transfer Component

A Portability Suite 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.

Flexible Image

One of Portability Suite's three fundamental workload infrastructures along with physical machines and virtualization platforms. 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 Flexible Image bootable, you convert it to a workload on physical hardware or a virtualization platform. *See also* Capture Image, Deploy Image, and Import Image.

Flexible Image Server

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

I2P

Image-to-physical. A workload portability operation in which the source is a workload's Flexible 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 Flexible Images 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 Flexible Image and the target is a bootable workload either on physical hardware or on a virtualization platform.

image

See Flexible Image.

image server

See Flexible Image Server.

Import Image

A Portability Suite job that creates a Flexible Image by using either raw volume data or third-party volume archives.

increment

In a workload protection job, any change on the monitored workload's inventoried volumes that Portability Suite captures for updating the corresponding workload or image replica. See protection.

job

A collection of the essential parameters of a workload portability, protection, discovery, removal, or related operation that you define, save, schedule, execute, and monitor in the Portability Suite Client.

Migrate

See PlateSpin Migrate.

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 Portability Suite peer-to-peer workload migration task that moves a workload to a physical or virtual machine. *Compare with* Copy.

Network Discovery

A Portability Suite feature, based on standard Windows network browsing functionality, that automatically populates the Portability Suite 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 Portability Suite installs on source and target machines to enable them to communicate with Portability Suite Server and with each other.

P2I

Physical-to-image. A workload portability operation in which the source is a physical machine and the target is the workload's Flexible 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. P2V or V2V portability is an integral part and a prerequisite for workload protection.

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 protection), or a Flexible Image (P2I conversion).

peer-to-peer

A workload portability or protection 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 Flexible Image). P2P, P2V, V2V, and V2P are all peer-to-peer portability operations; I2X and X2I are not.

PlateSpin Migrate

One of the two Portability Suite products, packaged for workload migration operations. *See also* PlateSpin Protect.

PlateSpin Protect

One of the two Portability Suite products, packaged for workload protection operations. *See also* PlateSpin Migrate.

Portability Suite Client

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

Portability Suite Network

A named collection of discovered workloads and target machines you work with at any given time. A parameter of all discovery jobs.

Portability Suite Server

The remote runtime execution engine underlying both Portability Suite products (PlateSpin® Migrate and PlateSpin® Protect).

Prepare for Synchronization

A step in Windows workload synchronization operations where the target is a virtual machine in a Server Sync job. Automates the configuration and preliminary booting of the target virtual machine, and registers it as a valid Server Sync target in the Portability Suite Client's user interface. *See* Server Sync.

Protect

See PlateSpin Protect.

protection

A workload portability operation that creates a VM replica of a source workload and regularly synchronizes the state of the replica with the state of the source workload.

raw volume data

A directory of files that constitute a volume suitable for a Flexible Image.

role

One of the three kinds of Windows local user groups that define Portability Suite users with specific access permissions: Portability Suite Administrator, Portability Suite Power User, and Portability Suite Operator.

Run Restore

A Portability Suite operation that brings a protected workload's target virtual machine online.

semi-automated virtualization

A workload virtualization operation in which a workload's target infrastructure is a virtual machine that is treated by Portability Suite as a physical machine. Applies to target virtualization platforms for which Portability Suite provides limited workload portability automation. *Compare with* X2P.

Server

See Portability Suite 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. In protection jobs (Windows only), can serve as an intermediary stage in situations where the initial replication of the source workload by transferring the entire volume data is difficult, impractical, or impossible.

source

The origin, or the infrastructure of the origin, of a workload portability and protection operation (such as a physical machine, a virtual machine, or a Flexible Image). *Compare with* target.

Take Control

The temporary pre-execution environment into which a Windows or Linux workload is booted to facilitate communication and data transfer during a workload portability operation. For Windows workloads, Portability Suite uses Microsoft WinPE; for Linux, it uses a RAM disk.

target

The outcome, or the infrastructure of the outcome, of a workload portability or protection operation (such as a physical machine, a virtual machine, or a Flexible Image). *Compare with* source.

Test Restore

A Portability Suite job that creates, powers up, and configures a snapshot of a protected workload's virtual machine on the target VM host for testing purposes. *See also* Run Restore.

Undiscover

A Portability Suite 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 Flexible 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 Flexible Image.

VSS Block-based Transfer Component

A Portability Suite software component, dynamically installed on Windows source machines 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 Flexible 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 Flexible 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 Flexible Image.