MICRO FOCUS

PlateSpin® Protect 11.2 SP1
User Guide

December 2017

Legal Notice

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

License Grant

Licenses purchased for PlateSpin Protect 11 and later versions cannot be used for PlateSpin Protect 10.3 or prior versions.

Contents

About This Guide

The *User Guide* provides information about using PlateSpin Protect. The guide provides conceptual information, an overview of the user interface, and step-by-step guidance for common tasks. It also defines terminology and includes troubleshooting information.

Intended Audience

This document is intended for data center administrators and operators who use PlateSpin Protect in their ongoing workload protection and disaster recovery solution.

Additional Documentation

For the most recent version of this guide and other PlateSpin Protect documentation resources for this release, visit the PlateSpin Protect Documentation (https://www.netiq.com/documentation/platespin-protect-11-2-1/) website.

In addition to English, online documentation is available in these national languages: Chinese Simplified, Chinese Traditional, French, German, Japanese, and Spanish.

Contacting Micro Focus

Our goal is to provide documentation that meets your needs. If you have suggestions for documentation improvements, you can use the **comment on this topic** link at the bottom of any HTML page of the online English documentation, or send an email to Documentation-Feedback@microfocus.com. We value your input and look forward to hearing from you.

For specific product issues, contact Micro Focus Customer Care at https://www.microfocus.com/support-and-services/.

Additional technical information or advice is available from several sources:

- Product documentation, Knowledge Base articles, and videos: https://www.microfocus.com/support-and-services/
- The Micro Focus Communities pages for High Availability and Disaster Recovery: https://forums.novell.com/forumdisplay.php/1870-HIGH-AVAILABILITY-DISASTER-RECOVERY

Planning

PlateSpin Protect is business continuity and disaster recovery software that protects physical and virtual workloads (operating systems, middleware, and data) by using virtualization technology. If there is a production server outage or disaster, a virtualized replica of a workload can be rapidly powered on within the target *container* (a VM host), and continue to run as normal until the production environment is restored.

PlateSpin Protect enables you to:

- Quickly recover workloads upon failure
- Simultaneously protect multiple workloads
- Test the failover workload without interfering with your production environment
- Fail back failover workloads to either their original infrastructures or to completely new infrastructures, physical or virtual
- Take advantage of existing external storage solutions, such as SANs
- Chapter 1, "Planning Your PlateSpin Environment," on page 13
- Chapter 2, "Basic Workflow for Workload Protection and Recovery," on page 35

1 Planning Your PlateSpin Environment

Use the information in this section to plan your PlateSpin protection and recovery environment.

1.1 Supported Configurations

PlateSpin Protect supports most major versions of the Microsoft Windows, SUSE Linux Enterprise Server, and Red Hat Enterprise Linux operating systems. It also supports selected versions of Novell Open Enterprise Server, Oracle Enterprise Linux, and CentOS operating systems.

This section describes all of the platform configurations supported by PlateSpin Protect, as well as the software, hardware, and virtualization environments that are required for workload protection and recovery. Some configurations, as noted, require special handling for workload setup and recovery. Ensure that you review the referenced information elsewhere in the online documentation or Knowledgebase Articles before you attempt to set up the workload.

NOTE: Although configurations not mentioned here are not supported, many of the improvements we make to PlateSpin Protect will be in direct response to suggestions from our customers. You can help us ensure our product meets all your needs. If you are interested in a platform configuration not listed, please contact Technical Support. We value your input and look forward to hearing from you.

1.1.1 Supported Windows Workloads

PlateSpin Protect supports workloads for the Microsoft Windows operating system versions listed in Table 1-1.

Both file-level and block-level replications are supported, with certain restrictions. See Section 1.2, "Supported Data Transfer Methods," on page 22.

NOTE: Protection is not supported for desktop (workstation) workloads.

Table 1-1 *Supported Windows Workloads*

Operating System	Notes
Servers	
Windows Server 2016	Protection of Windows Server 2016 servers requires VMware 6.0 or later.
Windows Server 2012 R2 Windows Server 2012	Includes domain controllers (DC) and Small Business Server (SBS) editions. For information about conversion of Active Directory domain controllers, see Knowledgebase Article 7920501 (https://www.netiq.com/support/kb/doc.php?id=7920501).
Windows Server 2008 R2 (64-bit) Windows Server 2008 (64-bit) Windows Server 2008 latest SP (32-bit)	Includes domain controllers (DC) and Small Business Server (SBS) editions. For information about conversion of Active Directory domain controllers, see Knowledgebase Article 7920501 (https://www.netiq.com/support/kb/doc.php?id=7920501).
Windows Server 2003 R2 (64-bit) Windows Server 2003 R2 (32-bit) Windows Server 2003 latest SP (64-bit) Windows Server 2003 latest SP (32-bit)	Windows 2003 requires SP1 or higher for Block-based replication.
Clusters	
Windows Server 2016 server-based Microsoft Failover Cluster	Protection of Windows Server 2016 Cluster requires VMware 6.0 or later.
Windows Server 2012 R2 server-based Microsoft Failover Cluster	Supported models: *Node and Disk Majority Quorum* and *No Majority: Disk Only Quorum*.
Windows Server 2008 R2 server-based Microsoft Failover Cluster	Support includes block-based data transfer with a driver (Fibre Channel SANs only) or without a driver for incremental replications for clusters. File-based replication is not supported. **WARNING:** Do not attempt to use the block-based driver on clusters with shared iSCSI drives. It renders the cluster unusable. See "Preparing for Windows Clusters Protection" on page 111.

Operating System	Notes
Windows Server 2003 R2 server-based Windows Cluster Server	Supported model: *Single-Quorum Device Cluster*. Support includes only driverless block-based data transfer for incremental replications for clusters. File-based replication is not supported. See "Preparing for Windows Clusters Protection" on page 111.
Hyper-V Hosts	
Windows Server 2012 R2 with Hyper-V Role Windows Server 2012 with Hyper-V Role	Protect a Windows server that is functioning as a Hyper-V host and its volumes. Protect the individual VMs separately.

Configuration Requirements for Windows

Windows Updates

Ensure that you apply Windows updates on your source system before you run the first full replication.

Domain Controller and Antivirus Software

If the Windows machine is a Domain Controller, ensure that you also disable antivirus software on the system during the replication.

1.1.2 Supported Linux Workloads

PlateSpin Protect supports workloads for the Linux operating system distributions listed in Table 1-2.

Replication of protected Linux workloads occurs only at the block level. See "Requirement for a blkwatch Driver" on page 17.

Table 1-2 *Supported Linux Workloads*

Operating System	Versions	Notes
Servers		
Red Hat Enterprise Linux (RHEL)	7.0 to 7.3 6.0 to 6.9 5.x 4.x	See "Linux Distributions Supported by Protect" on page 127 for a list of supported Linux kernel versions and architectures for RHEL distributions. PlateSpin Protect does not support the XFS version 5 (v5) file system on RHEL 7.3 and on distributions based on RHEL 7.3. For Red Hat Enterprise Linux 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version `2.6.32-642.13.1.el6.x86_64`) for the RHEL 6.7 distribution. For Red Hat Enterprise Linux 6.8, Oracle Linux 6.8, and CentOS 6.8 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version `2.6.32-696.20.1.el6.x86_64`) for the 6.8 distribution.
SUSE Linux Enterprise Server (SLES)	11 SP1 to 11 SP4 10.x 9.x	See "Linux Distributions Supported by Protect" on page 127 for a list of supported Linux kernel versions and architectures for SLES distributions. Kernel version 3.0.13 of SLES 11 SP3 is not supported. Upgrade to kernel version 3.0.27 or later before you inventory the workload.
Open Enterprise Server (OES)	2015 SP1 11 SP1 to 11 SP3 2 SP3 See SUSE Linux Enterprise Server (SLES).	For OES 2015 SP1, Protect supports NSS32-bit pools up to 8 TB in size; NSS64-bit pools are not supported. See "Linux Distributions Supported by Protect" on page 127 for a list of supported Linux kernel versions and architectures for SLES distributions. The default kernel version 3.0.13 on OES 11 SP2 is not supported. Upgrade to kernel version 3.0.27 or later before you inventory the workload.

Operating System	Versions	Notes
Oracle Linux (OL) (formerly Oracle Enterprise Linux (OEL))	See Red Hat Enterprise Linux (RHEL).	See "Linux Distributions Supported by Protect" on page 127 for a list of supported Linux kernel versions and architectures for RHEL distributions. Blkwatch drivers are available for the standard Red Hat Compatible Kernel (RHCK) and Unbreakable Enterprise Kernel (UEK) in OEL 6 U7 and later, as noted in the "List of Distributions" on page 128. Workloads using the Unbreakable Enterprise Kernel are not supported for PlateSpin Protect 11.2 and earlier. For Oracle Linux 6 U7, blkwatch drivers for kernel version 2.6.32-573 do not support incremental replication for workloads with LVM volumes. Update the kernel, then use RHEL 6 U7 drivers for kernel 2.6.32-642.
CentOS	See Red Hat Enterprise Linux (RHEL).	See "Linux Distributions Supported by Protect" on page 127 for a list of supported Linux kernel versions and architectures for RHEL distributions. CentOS 7.*x* requires VMware 5.5 or higher.

Configuration Requirement for Linux Workloads

Requirement for a `blkwatch` Driver

The block-based transfer of data for a Linux workload in requires a `blkwatch` driver that is compiled for the particular Linux distribution being protected. PlateSpin Protect software includes pre-compiled versions of the `blkwatch` driver for many non-debug Linux distributions (32-bit and 64-bit). You can also create a custom driver. For more information, see "Linux Distributions Supported by Protect" on page 127.

1.1.3 Supported VM Containers

A VM container is a protection infrastructure that acts as the host of a protected workload's regularly updated and bootable virtual replica.

- "Supported VMware Platforms" on page 18
- "Support for VMware DRS Clusters as Containers" on page 18
- "Support for VMware vCenter Site Recovery Manager" on page 19
- "Support for Protect Multitenancy on VMware" on page 19

Supported VMware Platforms

See Table 1-3 for a list of supported VMware platforms. The platforms are supported as protection containers and failback containers.

NOTE: Protection of workloads to a target VM container is subject to the support of the guest operating system on the target host by the host vendor. For information about your target VMware hosts, refer to the *VMware Compatibility Guide* (http://www.vmware.com/resources/compatibility/).

The container infrastructure can be either a VMware ESXi Server or a VMware DRS Cluster. For information about VMware DRS Cluster configuration requirements, see "Support for VMware DRS Clusters as Containers" on page 18.

Table 1-3 *Platforms Supported as VM Containers*

Container	Versions	Notes
VMware vCenter or ESXi	6.5	As a VM Container, the DRS Cluster must consist of ESXi 6.5 servers only, and can be managed by vCenter 6.5 only.
VMware vCenter or ESXi	6.0 (GA2, U2, U3)	As a VM Container, the DRS Cluster must consist of ESXi 6.0 servers only, and can be managed by vCenter 6.0 only.
VMware vCenter or ESXi	5.5 (GA2, U2, U3)	As a VM Container, the DRS Cluster must consist of ESXi 5.5 servers only, and can be managed by vCenter 5.5 only.
VMware vCenter or ESXi	5.1 (GA2, U2, U3)	As a VM Container, the DRS Cluster must consist of ESXi 5.1 servers only, and can be managed by vCenter 5.1 only.
VMware vCenter or ESXi	4.1 (GA2, U3)	As a VM Container, the DRS Cluster must consist of ESXi 4.1 servers only, and can be managed by vCenter 4.1 only.

NOTE: Your VMware ESXi hosts must have a paid license; protection is unsupported with these systems if they are operating with a free license.

Support for VMware DRS Clusters as Containers

To be a valid protection target, your VMware DRS Cluster must be added to the set of containers (inventoried) as a VMware Cluster. You should not attempt to add a DRS Cluster as a set of individual ESX servers. See "Adding Containers (Protection Targets)" on page 90.

In addition, your VMware DRS cluster must meet the following configuration requirements:

- DRS must be enabled and set to either **Partially Automated** or **Fully Automated**. (It must not be set to **Manual**.)
- At least one datastore must be shared among all the VMware hosts in the VMware Cluster.

- At least one vSwitch and virtual port-group, or vNetwork Distributed Switch, must be common to all the VMware hosts in the VMware Cluster.
- The failover workloads (VMs) for each protection contract must be placed exclusively on datastores, vSwitches, and virtual port-groups that are shared among all the VMware hosts in the VMware Cluster.

Support for VMware vCenter Site Recovery Manager

PlateSpin Protect supports copying replicated VMs to a remote recovery site, using VMware vCenter Site Recovery Manager (SRM). See Section 6.7, "Configuring Support for VMware vCenter Site Recovery Manager," on page 70.

Support for Protect Multitenancy on VMware

PlateSpin Protect supports multitenancy in VMware. Multiple Protect servers can share the same VMware cluster backend. See "Setting Up Protect Multitenancy on VMware" on page 54.

1.1.4 Supported Workload Architectures

PlateSpin Protect supports the following x86-based computer architectures:

- "Processor and OS Architecture" on page 19
- "Cores and Sockets for Target VMs" on page 19
- "Virtual CPUs for Target VMs" on page 20
- "UEFI and BIOS Firmware" on page 20

Processor and OS Architecture

PlateSpin Protect supports protection and recovery of x64 and x86 architectures for physical and virtual workloads in your data center:

- 64-bit
- 32-bit

Cores and Sockets for Target VMs

For supported VM containers using VMware 5.1 and higher with a minimum VM hardware Level 8, PlateSpin Protect enables you to specify the number of sockets and the number of cores per socket for the failover workload. It automatically calculates the total cores. This parameter applies on the initial setup of a workload with an initial replication setting of **Full**.

NOTE: The maximum number of cores the workload can use is subject to external factors such as the guest operating system, the VM hardware version, VMware licensing for the ESXi host, and ESXi host compute maximums for vSphere. See *ESXi/ESX Configuration Maximums* (VMware Knowledge Base 1003497) (https://kb.vmware.com/kb/1003497).

Some distributions of a guest OS might not honor the cores and cores per socket configuration. For example, guest OSes using SLES 10 SP4 and OES 2 SP3 retain their original cores and sockets settings as installed, whereas other SLES, RHEL, and OES distributions honor the configuration.

Virtual CPUs for Target VMs

For VM containers using VMware 4.1, PlateSpin Protect enables you to specify the required number of vCPUs (virtual CPUs) to assign to the failover workload. This parameter applies on the initial setup of a workload with an initial replication setting of *Full*. Each vCPU is presented to the guest OS on the VM container as a single core, single socket.

UEFI and BIOS Firmware

PlateSpin Protect supports the UEFI and BIOS firmware interfaces for Windows and Linux workloads.

NOTE: If you are protecting a UEFI-based workload and you want to continue using the same firmware boot mode throughout the protected workload lifecycle, you must target a vSphere 5.0 or newer container.

The following are examples of Protect behavior when protecting and failing back between UEFI and BIOS-based systems:

- When you transfer a UEFI-based workload to a VMware vSphere 4.*x* container (which does not support UEFI), Protect transitions the workload's UEFI firmware at failover time to BIOS firmware. Then, when failback is selected on a UEFI-based physical machine, Protect reverses the firmware transition from BIOS to UEFI.

- If you attempt to failback a protected Windows 2003 workload to a UEFI-based physical machine, Protect analyzes the choice and notifies you that it is not valid. That is, the firmware transition from BIOS to UEFI is not supported because Windows 2003 does not support the UEFI boot mode.

- When you protect a UEFI-based source on a BIOS-based target, Protect migrates the UEFI system's boot disks, which were GPT, to MBR disks. Failing back this BIOS workload to a UEFI-based physical machine converts the boot disks back to GPT.

On Windows workloads, PlateSpin Protect mirrors the Microsoft support of UEFI or BIOS-based Windows workloads. It transfers workloads from source to target while enforcing the supported firmware for the respective source and target operating systems. Both Block-based and File-based transfers are supported. It does the same for the failback to a physical machine. When any transition (failover and failback) between UEFI and BIOS systems are initiated, Protect analyzes the transition and alerts you about its validity.

1.1.5 Supported Storage

PlateSpin Protect supports the following storage configurations for Windows and Linux workloads.

- "Storage Disks" on page 21
- "Partitioning Schemes" on page 21
- "Windows File Systems" on page 21
- "Linux File Systems" on page 21
- "Linux Storage Features" on page 21

Storage Disks

PlateSpin Protect supports several types of source storage disks, including basic disks, Windows dynamic disks, LVM2, RAID, and SAN.

You can specify whether virtual disks on the protected VM replica are thin provisioned or thick provisioned.

NOTE: The following caveats apply for storage disks:

- **Windows Dynamic Disks:** PlateSpin Protect does not support Windows dynamic disks at the target.

  For dynamic disks, the storage does not follow the Same as Source mapping strategy. Both Simple Dynamic Volumes and Spanned Dynamic Volumes will reside on the target workload as Simple Basic Volume disks. The target disk is partitioned as GPT if the total combined size of the dynamic volume's member disks exceeds MBR partition size limits. For more information, see *Microsoft TechNet: Understanding the 2 TB limit in Windows Storage* (https://blogs.technet.microsoft.com/askcore/2010/02/18/understanding-the-2-tb-limit-in-windows-storage/).

- **Linux Software RAID:** PlateSpin Protect does not support Linux workloads with volumes on software RAID.

Partitioning Schemes

PlateSpin Protect supports MBR (Master Boot Record) and GPT (GUID Partition Table) partitioning schemes for Windows and Linux workloads. Workloads and storage for protection must be configured on disks partitioned with the MBR or GPT. Although GPT allows up to 128 partitions per single disk, PlateSpin Protect supports only 57 or fewer GPT partitions per disk.

Windows File Systems

PlateSpin Protect supports only the NTFS file system on any supported Windows system.

Linux File Systems

PlateSpin Protect supports EXT2, EXT3, EXT4, REISERFS, XFS, and NSS (Open Enterprise Server only) file systems, with block-based transfer only.

NOTE: The XFS v5 file system is not supported for Red Hat Enterprise Linux 7.3 and distributions based on that version.

NOTE: Encrypted volumes of workloads on the source are decrypted in the failover VM.

Linux Storage Features

For Linux workloads, PlateSpin Protect provides the following additional storage support:

- Non-volume storage, such as a swap partition that is associated with the source workload, is recreated in the failover workload.

- The layout of volume groups and logical volumes is preserved so that you can re-create it during failback.
- LVM raw disk volumes are supported in Same as Source configurations on Linux workloads.
- (OES 11) Novell Linux Volume Management (NLVM) layout of source workloads are preserved and re-created in the VM container. NSS pools are copied from the source to the recovery VM.
- (OES 2) EVMS layouts of source workloads are preserved and re-created in the VM container. NSS pools are copied from the source to the recovery VM.

1.1.6 Supported International Languages

In addition to English, PlateSpin Protect provides National Language Support (NLS) for installation and use on machines configured for the following international languages:

- Chinese Simplified (zh-cn)
- Chinese Traditional (zn-tw)
- French (fr)
- German (de)
- Japanese (ja)

TIP: Other international versions have limited support. Updating system files could be affected in languages other than those listed above.

Localized online documentation is available in these languages, as well as in Spanish (es).

To use the Web Interface in one of these languages, see "Configuring Language Settings for International Versions" on page 61.

1.1.7 Supported Web Browsers

Most of your interaction with the product takes place through the browser-based Web Interface.

The supported browsers are:

- *Google Chrome*, version 34.0 and later
- *Microsoft Internet Explorer*, version 11.0 and later
- *Mozilla Firefox*, version 29.0 and later

NOTE: JavaScript (Active Scripting) must be enabled in your browser.

To use the PlateSpin Protect Web Interface in one of the supported international languages, see "Configuring Language Settings for International Versions" on page 61.

1.2 Supported Data Transfer Methods

A data transfer method describes the way data is replicated from a source workload to a target workload. PlateSpin Protect provides different data transfer capabilities, which depend on the protected workload's operating system.

1.2.1 Supported Transfer Methods for Windows Workloads

For Windows workloads, PlateSpin Protect provides mechanisms to transfer workload volume data at either block level or file level.

- **Windows File-level Replication:** (Windows only) Data is replicated on a file-by-file basis.
- **Windows Block-Level Replication:** Data is replicated at a volume's block level. For this transfer method, PlateSpin Protect provides two mechanisms that differ by their continuity impact and performance. You can toggle between these mechanisms as required.
  - **Replication using the Block-Based Component:** This option uses a dedicated software component for block-level data transfer. It leverages the Microsoft Volume Snapshot Service (VSS) and the applications and services that support VSS. The installation of the component on your protected workload is automatic.

    NOTE: Installation and uninstallation of the block-based component requires a reboot of your protected workload. No reboot is required when you are protecting Windows clusters with block-level data transfer. When you configure the workload protection details, you can opt to install the component at a later time, deferring the required reboot until the time of the first replication.

  - **Replication without the Block-Based Component:** This option uses an internal 'hashing' mechanism in combination with Microsoft VSS to track changes on the protected volumes. The replication compares each block on the disk and copies only changes.

    This option requires no reboot, but its performance is inferior to that of the block-based component.

1.2.2 Supported Transfer Method for Linux Workloads

For Linux workloads, PlateSpin Protect supports only block-based data transfer with a block-watch (`blkwatch`) driver.

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

The PlateSpin Protect distribution includes precompiled `blkwatch` drivers for workloads running the standard, non-debug kernels of supported Linux distributions. See Section B.2, "Pre-compiled blkwatch Drivers for Linux Distributions," on page 128.

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

1.3 Security and Privacy

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

1.3.1 Encryption of Data in Transmission

Transfer encryption makes the transmission of your workload data more secure during workload replication. When encryption is enabled, over-the-network data transfer from the source to the target is encrypted by using AES (Advanced Encryption Standard).

NOTE: Data encryption has a performance impact and might significantly slow down the data transfer rate by up to 30%.

You can enable or disable encryption individually for each workload by selecting the **Encrypt Data Transfer** option. See "Workload Protection Details" on page 154.

1.3.2 Security of Client/Server Communications

The PlateSpin Server enables SSL on the PlateSpin Server host, providing secure data transmission between your web browser and the PlateSpin Server with HTTPS (Hypertext Transfer Protocol Secure). The installation also adds a self signed certificate if no valid certificates are found.

1.3.3 Security of Credentials

PlateSpin Protect protects credentials by using an SSL connection for communications and the Windows cryptographic library to encrypt passwords.

Credentials that you use to access various systems (such as workloads and failback targets) are stored in the PlateSpin Protect database and are therefore covered by the same security safeguards that you have in place for your PlateSpin Server host.

In addition, credentials are included within diagnostics, which are accessible to accredited users. You should ensure that workload protection projects are handled by authorized staff.

1.3.4 User Authorization and Authentication

PlateSpin Protect provides a comprehensive and secure user authorization and authentication mechanism based on user roles, and controls application access and operations that users can perform. See "Configuring User Authorization and Authentication" on page 51.

1.3.5 Windows Authentication for Microsoft SQL Server Database

PlateSpin Protect provides the ability to use Windows Authentication for access to the Microsoft SQL Server database. See "Requirements for Windows Authentication to the Microsoft SQL Server Database" on page 32.

1.3.6 Port Settings and Firewalls

Table 1-4 lists the default ports used by PlateSpin Protect. If you configure custom ports, you must open those ports instead. For communications between the PlateSpin Server and the source and target machines it manages, ensure that you also open the appropriate ports on any firewalls between them. Traffic for communications is bidirectional (incoming and outgoing). See also "Access and Communication Requirements across Your Protection Network" on page 29.

Table 1-4 *Default Ports Used by PlateSpin Protect*

Port Number	Protocol	Function	Details
80	TCP	HTTP	(Not secure) Used for HTTP communications between the PlateSpin Server host and the source and target machines it manages. Open this port on your PlateSpin Server host, the source and target workloads, and the VMware ESXi hosts.
443	TCP	HTTPS	(Secure) Used for HTTPS communications, if SSL is enabled, between the PlateSpin Server host and the source and target machines. Open this port on your PlateSpin Server host, the source and target workloads, the VMware ESXi hosts, and the vCenter host server.
3725	TCP	Data transfer	Used for data transfer between the source and target machines, including file-based transfer and block-based transfer. Open this port on the source and target machines for all workloads. Any firewall between a source and its target must also allow TCP port 3725. See "Supported Configurations" on page 13.
135 445	TCP	RPC/DCOM	Used for RPC/DCOM communications on Windows machines during the discovery process, and when taking control and rebooting the source machine. Open these ports for communications between the source and target machines for all Windows workloads. See "Supported Windows Workloads" on page 14.
137 138 139	TCP	NetBIOS	Used for NetBIOS communications (name service, datagram service, and session service). Open these ports for communications between the source and target machines for all Windows workloads. See "Supported Windows Workloads" on page 14.
137 138	UDP	SMB	Used for SMB communications for the file transfer of the Take Control folder and files from the PlateSpin Server to the source machine. Open these ports on your PlateSpin Server host and the source workloads.
139 445	TCP	SMB	

Port Number	Protocol	Function	Details
22	TCP		Used for SSH and SCP communications on Linux machines during the discovery process. Open this port on the source and target machines for all Linux workloads. See "Supported Linux Workloads" on page 15.
25	TCP	SMTP	Used for SMTP traffic if email notification is enabled.
25	UDP	SMTP	Open this port on the PlateSpin Server host and the mail relay host.
1433	TCP	SQL	Used for Microsoft SQL Server communications for authentication and data exchange to a remote SQL Server. Open the SQL ports on your PlateSpin Server host and the remote SQL Server host, as well as on any firewalls between them. For more information the SQL Server port requirements, see *Configure the Firewall to Allow Server Access* in the Microsoft Developers Network.
1434	TCP	SQL	Used for the Microsoft SQL Server dedicated administrator connection.
1434	UDP	SQL	Used for the Microsoft SQL Server named instances. This port might be required when you use named instances on a remote SQL Server.
49152 to 65535	TCP	SQL	Used for the Microsoft SQL Server or RPC for LSA, SAM, and Netlogon. If you have configured Microsoft SQL Server to use a specific TCP port, you must open that port on the firewall. See "Requirements for Windows Authentication to the Microsoft SQL Server Database" on page 32.

1.4 Performance

1.4.1 About Product Performance Characteristics

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

- Hardware and software profiles of your source workloads

- Hardware and software profiles of your target containers

- Hardware and software profile of your PlateSpin Server host

- The specifics of your network bandwidth, configuration, and conditions

- The number of protected workloads

- The number of volumes under protection

- The size of volumes under protection

- File density (number of files per unit of capacity) on your source workloads' volumes

- Source I/O levels (how busy your workloads are)

- The number of concurrent replications

- Whether data encryption is enabled or disabled

- Whether data compression is enabled or disabled

For large-scale workload protection plans, you should perform a test protection of a typical workload, run some replications, and use the result as a benchmark, fine-tuning your metrics regularly throughout the project.

1.4.2 RPO, RTO, and TTO Specifications

In your protection environment, you will have different expectations for recovery points and recovery times required for a variety of workloads.

- **Recovery Point Objective (RPO):** The RPO setting describes the tolerable amount of data loss as measured in time in the event of a major IT outage. You define the RPO with a configurable interval between incremental replications of a protected workload.

  The RPO is affected by current utilization levels of PlateSpin Protect, the rate and scope of changes on the workload, your network speed, and the chosen replication schedule.

- **Recovery Time Objective (RTO):** The RTO setting describes a workload's tolerable downtime as measured by the time a failover operation takes to complete. The failover operation brings a failover workload online to temporarily replace a protected production workload.

  The RTO is affected by the time it takes to configure and execute the failover operation (10 to 45 minutes). See "Failover" on page 158.

- **Test Time Objective (TTO):** The TTO setting describes the time required for testing disaster recovery with some confidence of service restoration. It is similar to RTO, but includes the time needed for a user to test the failover workload.

  Use the **Test Failover** feature to run through different scenarios and generate benchmark data. See "Using the Test Failover Feature" on page 159.

Among factors that have an impact on RPO, RTO, and TTO is the number of required concurrent failover operations. A single failed-over workload has more memory and CPU resources available to it than multiple failed-over workloads, which share the resources of their underlying infrastructure.

When you test the failover response, you should note the actual values associated with the configured RPO, RTO, and TTO:

- **Recovery Point Actual (RPA):** The RPA is the actual data loss measured in time and defined by the actual measured interval between incremental replications of a protected workload that occurs during a failover test. RPA is also known as *Actual Recovery Point Objective* (Actual RPO).

- **Recovery Time Actual (RTA):** The RTA is a measure of a workload's actual downtime defined by the time a failover operation takes to complete. RTA is also known as *Actual Recovery Time Objective* (Actual RTO).

- **Test Time Actual (TTA):** The TTA is a measure of the actual time in which a disaster recovery plan can be tested. It is similar to Actual RTO, but includes the time needed for a user to test the failover workload. TTA is also known as *Actual Test Time Objective* (Actual TTO).

You should determine average failover times for workloads in your environment by doing test failovers at various times, then use them as benchmark data in your overall data recovery plans. See "Generating Workload and Workload Protection Reports" on page 178.

1.4.3 Data Compression

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

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

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

Data Compression can be configured individually for each workload or in a Protection Tier. See "Protection Tiers" on page 166.

1.4.4 Bandwidth Throttling

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

Bandwidth throttling can be configured individual for each workload or in a Protection Tier. See "Protection Tiers" on page 166.

1.4.5 Scalability

Scalability encompasses (and depends on) the following major characteristics of your PlateSpin Protect product:

- **Workloads per Server:** The number of workloads per PlateSpin Server might vary between 10 and 50, depending on several factors, including your RPO requirements and the hardware characteristics of the server host.

- **Protections per Container:** The maximum number of protections per container is related to (but is not the same as) the VMware specifications pertaining to the maximum number of VMs supported per ESXi host. Additional factors include recovery statistics (including concurrent replications and fail-overs) and hardware vendor specifications.

You should conduct tests, incrementally adjust your capacity numbers, and use them in determining your scalability ceiling.

1.4.6 Database Server

PlateSpin Protect includes Microsoft SQL Server Express Edition. The capabilities of SQL Server Express are sufficient for a single PlateSpin Server that protects up to 50 workloads (see Section 1.4.5, "Scalability," on page 28).

NOTE: Microsoft SQL Server Express has a database size limit of 10 GB and can use only one CPU core at a time. For more information about requirements and limitations for SQL Server Express, see the Microsoft SQL Server 2014 Express documentation (https://www.microsoft.com/en-us/download/details.aspx?id=42299).

The PlateSpin Server database instance might grow up to 0.5 GB per month per workload, depending on the number of incremental replications that are scheduled. We recommend that you periodically archive or discard the historical reporting data to make room for new reporting data.

In a VMware DRS Cluster, ensure that you balance protection targets across multiple hosts in the cluster for best performance.

We recommend that you configure the PlateSpin Server to use a database instance on your existing Microsoft SQL Server Standard Edition or Enterprise Edition database server in the following environments:

- Deployments of multiple PlateSpin Servers that use the same remote Microsoft SQL Server database server for their database instances
- Deployments where keeping all history of the reporting data is important

While multiple PlateSpin Servers can use the same remote database server, each server requires a separate database instance.

To set up a remote database instance for your PlateSpin Server, see "Configuring Your Remote Microsoft SQL Server Database Server" in the *PlateSpin Protect Installation and Upgrade Guide*.

1.5 Access and Communication Requirements across Your Protection Network

Before you set up workloads for protection and recovery, ensure that you configure your network with the access and communications settings described in this section.

- Section 1.5.1, "Network Requirements for the PlateSpin Server Host Web Interface," on page 30
- Section 1.5.2, "Network Requirements for Containers," on page 30
- Section 1.5.3, "Network Requirements for Workloads," on page 30
- Section 1.5.4, "Requirements for Windows Authentication to the Microsoft SQL Server Database," on page 32
- Section 1.5.5, "Requirements for Protection across Public and Private Networks through NAT," on page 33

1.5.1 Network Requirements for the PlateSpin Server Host Web Interface

Table 1-5 describes the ports that must be open for on the PlateSpin Server host to allow access to the Web Interface.

Table 1-5 *Open Port Requirements for the PlateSpin Server Host*

Port (Default)	Remarks
TCP 80	For HTTP communication
TCP 443	For HTTPS communication (if SSL is enabled)

1.5.2 Network Requirements for Containers

Table 1-6 describes the software, network, and firewall requirements for the supported workload containers.

Table 1-6 *Access and Communication Requirements for Containers*

System	Prerequisites	Required Ports (Defaults)
All containers	Ping (ICMP echo request and response) capability.	
All VMware containers. See "Supported VM Containers" on page 17.	◆ VMware account with an Administrator role ◆ VMware Web services API and file management API	HTTPS (TCP 443)
vCenter Server	The user with access must be assigned the appropriate roles and permissions. Refer to the pertinent release of VMware documentation for more information.	HTTPS (TCP 443)

1.5.3 Network Requirements for Workloads

Table 1-7 describes the software, network, and firewall requirements for workloads that you intend to protect by using PlateSpin Protect.

Table 1-7 *Access and Communication Requirements for Workloads*

Workload Type	Prerequisites	Required Ports (Defaults)
All workloads	Ping (ICMP echo request and response) support	

Workload Type	Prerequisites	Required Ports (Defaults)
All Windows workloads. See "Supported Windows Workloads" on page 14.	◆ Microsoft .NET Framework 3.5 Service Pack 1 ◆ Microsoft .NET Framework 4.0 For discovery, source workloads must be running Microsoft .NET Framework 2 SP2 or later.	
All Windows Server Cluster workloads. See Clusters in "Supported Windows Workloads" on page 14.	Ensure that the PlateSpin Server can resolve DNS forward lookup and reverse lookup for the IP addresses of the Windows Server Cluster and its cluster nodes. You can update the DNS server or update the local `hosts` file (`%systemroot%\system32\drivers\etc\hosts`) on the PlateSpin Server host.	

All Windows workloads. See "Supported Windows Workloads" on page 14.	◆ Built-in `Administrator` or domain administrator account credentials (membership only in the local Administrators group is insufficient).
	
	◆ The Windows Firewall configured to allow **File and Printer Sharing**. Use one of these options:
	
	  ◆ **Option 1, using Windows Firewall:** Use the basic **Windows Firewall** Control Panel item (`firewall.cpl`) and select **File and printer Sharing** in the list of exceptions.
	
	    - OR -
	
	  ◆ **Option 2, using Firewall with Advanced Security:** Use the Windows Firewall with Advanced Security utility (`wf.msc`) with the following **Inbound Rules** enabled and set to `Allow`:
	
	    ◆ **File and Printer Sharing (Echo Request - ICMPv4In)**
	
	    ◆ **File and Printer Sharing (Echo Request - ICMPv6In)**
	
	    ◆ **File and Printer Sharing (NB-Datagram-In)**
	
	    ◆ **File and Printer Sharing (NB-Name-In)**
	
	    ◆ **File and Printer Sharing (NB-Session-In)**
	
	    ◆ **File and Printer Sharing (SMB-In)**
	
	    ◆ **File and Printer Sharing (Spooler Service - RPC)**
	
	    ◆ **File and Printer Sharing (Spooler Service - RPC-EPMAP)**
		TCP 3725
		
		NetBIOS (TCP 137 - 139)
		
		SMB (TCP 139, 445 and UDP 137, 138)
		
		RPC (TCP 135, 445)

Workload Type	Prerequisites	Required Ports (Defaults)
Windows Server 2003 (including SP1 Standard, SP2 Enterprise, and R2 SP2 Enterprise).	**NOTE:** After enabling the required ports, run the following command at the server prompt to enable PlateSpin remote administration: `netsh firewall set service RemoteAdmin enable` For more information about netsh, see the Microsoft TechNet article, *The Netsh Command Line Utility* (http://technet.microsoft.com/en-us/library/cc785383%28v=ws.10%29.aspx).	TCP 3725, 135, 139, 445 UDP 137, 138, 139
All Linux workloads. See "Supported Linux Workloads" on page 15.	Secure Shell (SSH) server	TCP 22, 3725

1.5.4 Requirements for Windows Authentication to the Microsoft SQL Server Database

PlateSpin Protect provides the ability to use Windows Authentication for access to the Microsoft SQL Server database. You must configure Active Directory settings and open up ports in the firewall to allow authentication.

To enable Windows Authentication to the SQL database:

1 Ensure that you configure Microsoft SQL Server to allow both TCP/IP and Named Pipe connections.

2 (Conditional) If you plan to use Windows Authentication to access the Microsoft SQL Server database, you must configure the following in Active Directory:

  ◆ You must add the Microsoft SQL Server database server to the domain.

  ◆ You need two domain user accounts for the PlateSpin Protect installation.

    ◆ **A Domain user with the `sysadmin` role set:** This user with SQL Admin rights is required to create databases, tables, and other schema objects.

    ◆ **PlateSpin Service user:** The service user can be a low-privileged domain user in the domain. However, the service user must be a local administrator on the PlateSpin Protect Server and should be granted that permission prior to the installation.

      If the Windows user's password changes, you must update the password for the PlateSpin Service user and for the IIS App Pool. Consider using a Windows user whose password never expires to avoid the situation.

NOTE: If you use Windows Authentication, you must log in as the domain user with SQL Admin rights when you upgrade or update your PlateSpin Server.

3 Open the following ports on the firewall to support authentication to the SQL Server:

  ◆ **Ports 49152-65535/TCP:** Allow traffic for RPC for LSA, SAM, Netlogon.

  ◆ **Port 1433/TCP:** Allow traffic for Microsoft SQL Server.

  ◆ **Custom ports:** If you configure SQL Server to use a custom TCP port, you must open that port on the firewall.

NOTE: If you do not use dynamic ports, you must specify the dedicated port in the **Database Server** field.

4 (Conditional) If you want to use dedicated ports with PlateSpin Protect, you must open the ports on the firewall:

    4a On the database server, determine which ports need to be opened:

        4a1 In the SQL Server Configuration Manager, select **Protocols for SQLExpress > TCP/IP**, then right-click and select **Properties**.

        4a2 In the dialog, select the **IP Addresses** tab.

        4a3 Under **IPAll** (or under the desired protocol), if **TCP Port** or **TCP Dynamic Ports** is set to any value other than 0, open the specified ports on the firewall. These are the ports you use to connect to the SQL Server.

            For example, if the **TCP Dynamic Ports** field is set to 60664, and the **TCP Port** field is set to 1555, then you must enable Port 60664 and 1555 in the firewall rules on the SQL server.

    4b Open the ports on the firewall.

NOTE: If you have a value set for dynamic ports, you may not see your server in the list of SQL servers when you click **Browse**. In this case, you must specify the server manually in the **Database Server** input field of the PlateSpin Protect installation.

For example, if your server name is MYSQLSERVER, the database instance name SQLEXPRESS, and the dedicated port set for the dynamic port is 60664, you type the following text, and then select the desired authentication type:

```
MYSQLSERVER\SQLEXPRESS,60664
```

You must open the ports on the firewall.

1.5.5 Requirements for Protection across Public and Private Networks through NAT

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

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

- ◆ **PlateSpin Server:** Using your server's PlateSpin Configuration tool, record the additional IP addresses assigned to the PlateSpin Server host. See "Requirements for the PlateSpin Server to Function through NAT" on page 34.
- ◆ **Target Container:** When you attempt to discover a container (such as VMware ESX), specify the public (external) IP address of that host in the discovery parameters.
- ◆ **Workload:** When you attempt to add a workload, specify the public (external) IP address of that workload in the discovery parameters.
- ◆ **Failed-over VM:** During failback, you can specify an alternative IP address for the failed-over workload in Failback Details (Workload to VM) (page 161).

◆ **Failback Target:** During an attempt to register a failback target, when you are prompted to provide the IP address of the PlateSpin Server, provide either the local address of the PlateSpin Server host or one of its public (external) addresses recorded in the server's PlateSpin Configuration database. See "Requirements for the PlateSpin Server to Function through NAT" on page 34.

1.5.6 Requirements for the PlateSpin Server to Function through NAT

The PlateSpin Server needs additional IP addresses in order to function across environments that are enabled for Network Address Translation. See "Requirements for the PlateSpin Server to Function through NAT" on page 34.

1.5.7 Overriding the Default bash Shell for Executing Commands on Linux Workloads

By default, the PlateSpin Server uses the `/bin/bash` shell when executing commands on a Linux source workload.

If required, you can override the default shell by modifying the corresponding registry key on the PlateSpin Server. See Knowledgebase Article 7010676 *Linux Default Shell Override Procedure* (https://www.netiq.com/support/kb/doc.php?id=7010676).

2 Basic Workflow for Workload Protection and Recovery

PlateSpin Protect defines the following workflow for workload protection and recovery. Most of these steps are represented by workload commands on the Workloads page. See "Workload Protection and Recovery Commands" on page 43.

Table 2-1   *Protection and Recovery Lifecycle*

Task	Action	Remarks
Preparation		
Ensure that your workloads, containers, and environment meet the required criteria.		
	1. Ensure that PlateSpin Protect supports your workload.	See "Supported Configurations" on page 13.
	2. Ensure that your workloads and VM containers meet access and network prerequisites.	See "Access and Communication Requirements across Your Protection Network" on page 29.
Inventory		
Workloads that you want to protect and containers that host failover workloads must be properly inventoried. You can add workloads and containers in any order; however, every protection contract requires a defined workload and container that were inventoried by the PlateSpin Server.		
	3. Add target containers to the PlateSpin Server.	See "Adding Containers (Protection Targets)" on page 90.
	4. Add source workloads to the PlateSpin Server.	See "Adding Workloads (Protection Sources)" on page 94.
	5. For a physical protection target, prepare device drivers.	See Chapter 11, "Preparing Device Drivers for Physical Failback Targets," on page 97.
	6. For a Linux workload, prepare for workload protection:	See Chapter 12, "Preparing Linux Workloads for Protection," on page 107.
	7. For Windows Server Cluster workloads, prepare for cluster workload protection.	See Chapter 13, "Preparing for Windows Clusters Protection," on page 111.
Define Protection Contract		
	8. Define the details and specifications of a protection contract.	See "Configuring Protection Details and Preparing the Replication" on page 153.
	9. Prepare the replication.	
Initiate Protection		
	10. Begin the protection contract according to your requirements.	See "Starting the Workload Protection" on page 157.

Task	Action	Remarks
Protection Lifecycle Tasks (Optional)		

These steps are outside the automated replication schedule but are often useful in different situations or might be dictated by your business continuity strategy.

Task	Action	Remarks
	11. *Manual incremental*. You can run an incremental replication manually, outside the workload protection contract.	Select the workload, then click **Run Incremental**.
	12. *Testing*. You can test failover functionality in a controlled manner and environment.	See Using the Test Failover Feature.
Failover		
	13. This step carries out a failover of your protected workload to its replica running in your VM container.	See "Failover" on page 158.
Failback		
	14. This step corresponds to the business resumption phase after you have addressed any problems with your production workload.	See "Failback" on page 160.
Reprotection		
	15. This step enables you to redefine the original protection contract for your workload.	See "Reprotecting a Workload" on page 164. A **Reprotect** command becomes available after a successful failback.

Managing the PlateSpin Server

This section provides the information you need to activate your PlateSpin Protect license and customize the PlateSpin product for your environment. Familiarize yourself with the PlateSpin tools and configuration options. You can return to this section whenever you need to manage licenses or users, or to customize settings.

3 Using PlateSpin Tools

Most of your interaction with the product takes place through the browser-based Web Interface. You can also configure global parameters for the PlateSpin Server application using the web-based PlateSpin Configuration page.

- Section 3.1, "Launching the Web Interface," on page 39
- Section 3.2, "Dashboard Overview," on page 40
- Section 3.3, "Workloads Overview," on page 42
- Section 3.4, "Workload Protection and Recovery Commands," on page 43
- Section 3.5, "Other PlateSpin Server Management Tools," on page 44

3.1 Launching the Web Interface

1 (Optional) Configure PlateSpin Server and your web browser to use one of the supported international languages instead of English. See "Configuring Language Settings for International Versions" on page 61.

2 Open a supported web browser and go to:

```
https://Your_PlateSpin_Server/Protect
```

Replace `Your_PlateSpin_Server` with the DNS host name or the IP address of your PlateSpin Server host.

If SSL is not enabled, use `http` in the URL.

3 Log in using the local Administrator user credentials for the PlateSpin Server host.

For information about setting up additional users for PlateSpin, see Chapter 5, "Configuring User Authorization and Authentication," on page 51.

3.2 Dashboard Overview

The Dashboard page of the PlateSpin Protect Web Interface contains elements for navigating to different functional areas of the interface and carrying out workload protection and recovery operations.

Figure 3-1   The Default Dashboard Page of the PlateSpin Protect Web Interface
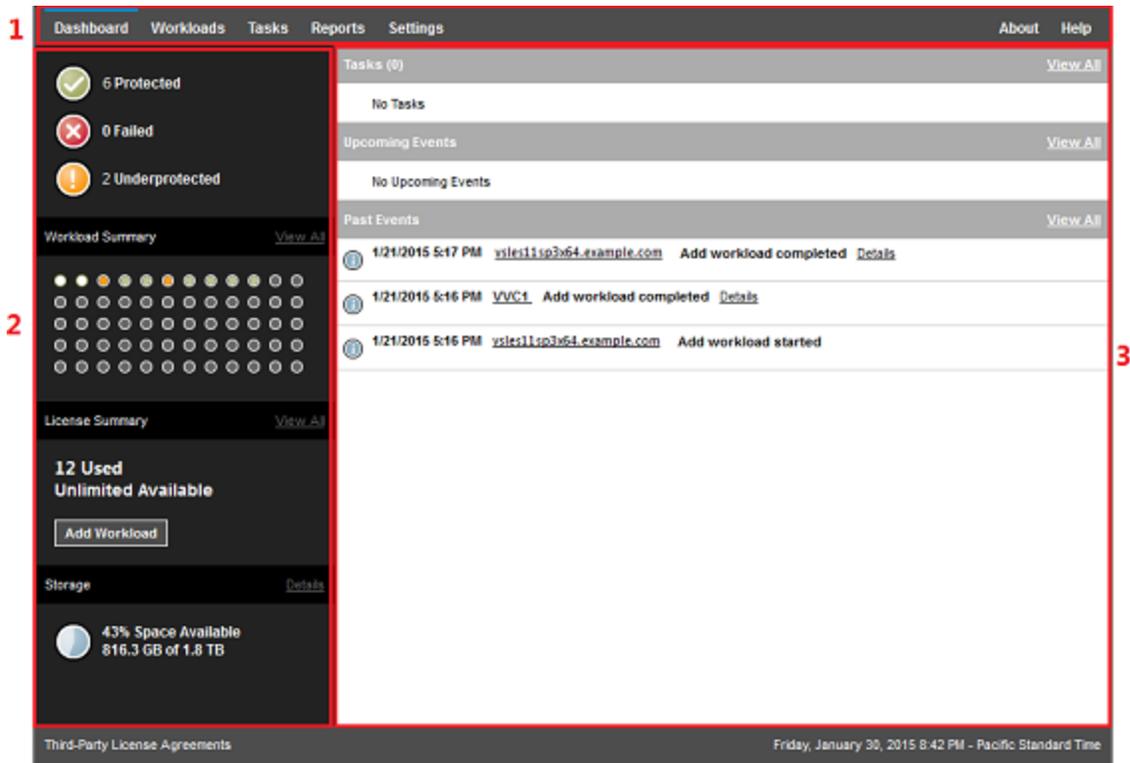
The Dashboard page consists of the following elements:

1. **Navigation bar:** Found on most pages of the PlateSpin Protect Web Interface.
2. **Visual Summary panel:** Provides a high-level view of the overall state of the PlateSpin Protect workload inventory,
3. **Tasks and Events panel:** Provides information about events and tasks requiring user attention.

The following topics provide more details:

- Section 3.2.1, "Navigation Bar," on page 41
- Section 3.2.2, "Visual Summary Panel," on page 41
- Section 3.2.3, "Tasks and Events Panel," on page 42

NOTE: You can alter certain elements of the Web Interface to match your organization branding. For more information, see "Rebranding the PlateSpin Protect Web Interface" on page 81.

3.2.1 Navigation Bar

The Navigation bar provides the following links:

- **Dashboard:** Displays the default Dashboard page.
- **Workloads:** Displays the Workloads page. See "Workloads Overview" on page 42.
- **Tasks:** Displays the Tasks page, which lists items requiring user intervention.
- **Reports:** Displays the Reports page. See "Generating Workload and Workload Protection Reports" on page 178.
- **Settings:** Displays the Settings page, which provides access to the following configuration options:
  - **Protection Tiers:** See "Protection Tiers" on page 166.
  - **Workload Tags:** See "Creating and Managing Workload Tags" on page 73.
  - **Permissions:** See "Configuring User Authorization and Authentication" on page 51.
  - **Containers:** See "Adding Containers (Protection Targets)" on page 90.
  - **Notification Settings:** "Enabling Event Notifications" on page 63.
  - **Replication Reports Settings:** "Enabling Replication Reports" on page 65
  - **SMTP:** See "Configuring SMTP for the Email Notification Service" on page 62.
  - **Licenses:** See "Activating Your Product License" on page 47.

3.2.2 Visual Summary Panel

The Visual Summary panel provides a the high-level protection status of inventoried workloads, the status of each licensed workloads, a license usage summary, and the amount of available storage.

Protection Status

The overall protection status of inventoried workloads are represented by three categories:

- **Protected:** Indicates the number of workloads under active protection.
- **Failed:** Indicates the number of protected workloads that the system has rendered as failed according to that workload's Protection Tier.
- **Underprotected:** Indicates the number of protected workloads that require user attention.

Workload Summary

The Workload Summary presents the health status of each licensed workload listed on the Workloads page. The maximum number of workload status dot icons matches the number of installed workload licenses on the PlateSpin Server. For an unlimited license, the summary displays 96 dot icons. Table 3-1 describes the different workload states represented by the dot icons.

The icons represent workloads in alphabetical order, according to the workload name. Mouse over a dot icon to display the workload name, or click the icon to display the corresponding Workload Details page.

Table 3-1 *Dot Icon Workload Representation*

● Protected	● Unprotected
● Failed	○ Unprotected – Error
● Underprotected	● Expired
	● Unused

License Summary

The License Summary displays the number installed licenses, and the number of licenses currently used by the workloads.

Storage

Storage provides information about the total amount of container storage space available to PlateSpin Protect, and the amount of space that is currently in use.

3.2.3 Tasks and Events Panel

The Tasks and Events panel shows the most recent Tasks, the most recent Past Events, and the next Upcoming Events.

Events are logged whenever something relevant to the system or to the workload occurs. For example, an event could be the addition of a new protected workload, the replication of a workload starting or failing, or the detection of the failure of a protected workload. Some events generate automatic notifications by email if SMTP is configured. See "Configuring Email Notification Services for Events and Replication Reports" on page 62.

Tasks are special commands that are tied to events that require user intervention. For example, upon completion of a Test Failover command, the system generates an event associated with two tasks: `Mark Test as Success` and `Mark Test as Failure`. Clicking either task results in the Test Failover operation being canceled and a corresponding event being written in the history. Another example is the `FullReplicationFailed` event, which is shown coupled with a `StartFull` task. You can view a complete list of current tasks on the **Tasks** tab.

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

3.3 Workloads Overview

The Workloads page displays a table with a row for each inventoried workload. Click a workload name to display a Workload Details page for viewing or editing configurations relevant to the workload and its state. The Workloads list displays information about the workload's availability (online or offline), tag, protection tier, replication status and run times, and last test failover time.

Figure 3-2 *The Workloads Page*
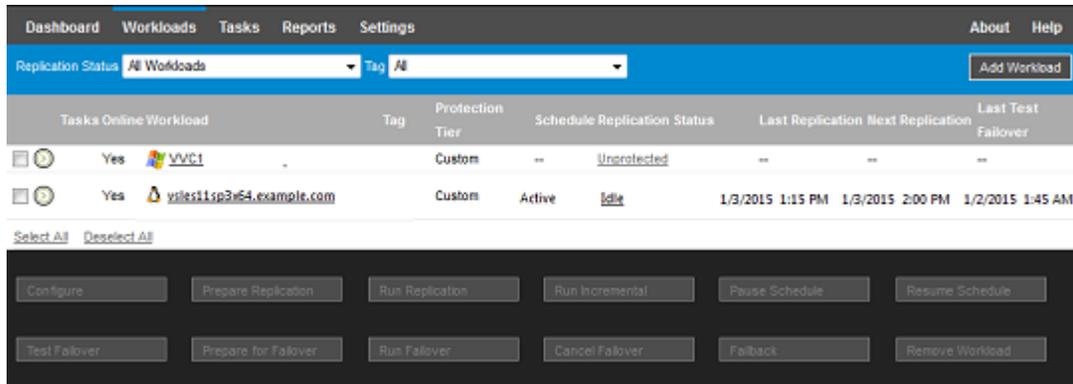
NOTE: All time stamps reflect the time zone of the PlateSpin Server host. This might be different from the time zone of the protected workload or the time zone of the host on which you are running the Web Interface. A display of the server date and time appears at the bottom right of the client window.

3.4   Workload Protection and Recovery Commands

Commands reflect the workflow of workload protection and recovery. To perform a command for a workload, select the corresponding check box at the left. Applicable commands depend on the current state of a workload.

Figure 3-3 *Workload Commands*
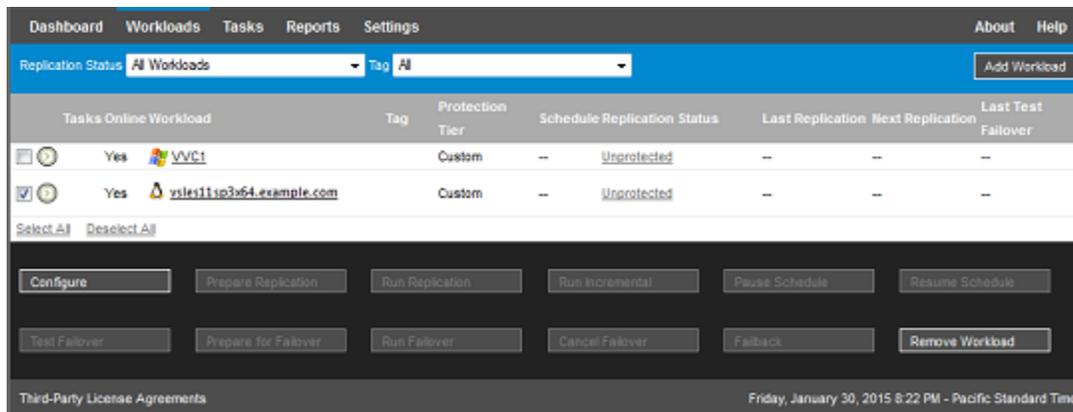
Table 3-2 summarizes workload commands along with their functional descriptions.

Table 3-2 *Workload Protection and Recovery Commands*

Workload Command	Description
Configure	Starts the workload protection configuration with parameters applicable to an inventoried workload.
Prepare Replication	Installs required data transfer software on the source and creates a failover workload (a virtual machine) on the target container in preparation for workload replication.

Workload Command	Description
Run Replication	Starts replicating the workload according to specified parameters (full replication).
Run Incremental	Performs an incremental transfer of changed data from the source to the target outside the workload protection contract.
Pause Schedule	Suspends the protection; all scheduled replications are skipped until the schedule is resumed.
Resume Schedule	Resumes the protection according to saved protection settings.
Test Failover	Boots and configures the failover workload in an isolated environment within the container for testing purposes.
Prepare for Failover	Boots the failover workload in preparation for a failover operation.
Run Failover	Boots and configures the failover workload, which takes over the business services of a failed workload.
Cancel Failover	Aborts the failover process.
Failback	Following a failover operation, fails the failover workload back to its original infrastructure or to a new infrastructure (virtual or physical).
Reprotect	Following a successful Failback operation, the Reprotect option becomes available.
Remove Workload	Removes a workload from the inventory.

3.5 Other PlateSpin Server Management Tools

3.5.1 PlateSpin Configuration

Some aspects of your PlateSpin Server's behavior are controlled by configuration parameters that you set on a configuration web page residing your PlateSpin Server host at:

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

NOTE: Under normal circumstances you should not need to modify these settings unless you are advised to do so by PlateSpin Support.

To change and apply any configuration parameters:

1 From any web browser, open

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

2 Search to locate the required server parameter and change its value.

3 Save your settings and exit the page.

A reboot or restart of PlateSpin services is not required to apply the changes.

The following topics provide information on specific situations when you might need to change product behavior using PlateSpin Configuration parameters:

- "Requirements for the PlateSpin Server to Function through NAT" on page 34
- "Optimizing Data Transfer over WAN Connections" on page 66
- "Optimizing Replication Environment Performance" on page 69
- "Setting Reboot Method for the Configuration Service" on page 69
- "Configuring Support for VMware vCenter Site Recovery Manager" on page 70
- "Rebranding the Web Interface By Using Configuration Parameters" on page 81
- "Configuring Windows Active Node Discovery" on page 116
- "Troubleshooting the Configuration Service" on page 179

3.5.2 Protect Agent Utility

The Protect Agent utility (ProtectAgent.cli.exe) is a command line utility that you can use to install, upgrade, query, or uninstall the block-based transfer drivers. Although a reboot is always required when you install, uninstall, or upgrade drivers, the Protect Agent utility allows you to better control when the action occurs and therefore, when the server reboots. For example, you can use the Protect Agent utility to install the drivers during scheduled down time, instead of during the first replication. See Appendix D, "Protect Agent Utility," on page 145.

3.5.3 VMware Role Tool

The VMware Roles tool (PlateSpin.VMwareRoleTool.exe) is a command line utility that you can use to create unique user roles a VMware data center in support of multitenancy. The roles enable you to allow non-administrative VMware users (or "enabled users") to perform Protect lifecycle operations in the VMware environment. See Section 5.4, "Setting Up Protect Multitenancy on VMware," on page 54.

4 Managing Licenses

After you have activated one license for the product, you can monitor the availability of workload licenses, add new licenses, and remove expired licenses.

4.1 Activating Your Product License

Your PlateSpin Protect product license entitles you to a specific number or unlimited number of workloads for protection through workload licensing.

For PlateSpin Protect product licensing, you must have a license activation code. If you do not have a license activation code, request one through the Customer Center (http://www.netiq.com/customercenter/). A Customer Care representative will contact you and provide the license activation code.

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

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

4.1.1 Online License Activation

For online activation, PlateSpin Protect must have Internet access.

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

To set up online license activation:

1 In the Web Interface, select **Settings > Licenses**, then click **Add License**.
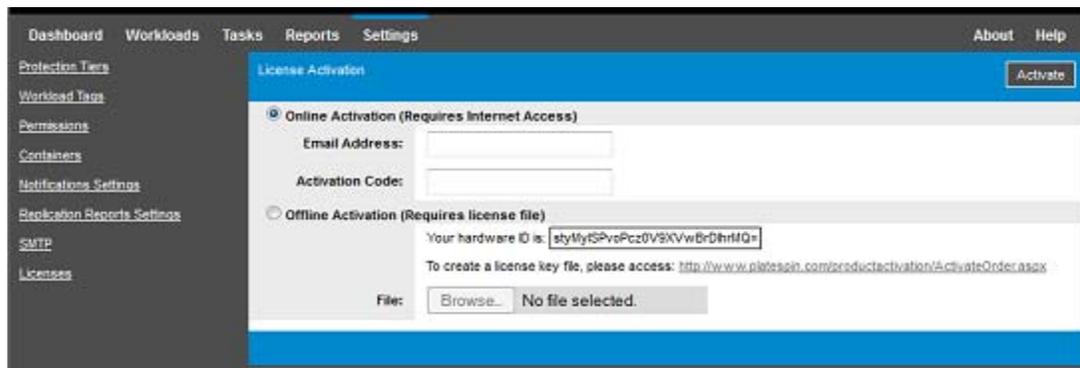
2 Select **Online Activation**.

3 Specify the email address that you provided when you placed your order and the activation code you received, then click **Activate**.

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

4.1.2 Offline License Activation

For offline activation, you obtain a PlateSpin Protect license key by using a computer that has Internet access.

1 In the Web Interface, select **Settings > Licenses**, then click **Add License**.

2 Select **Offline Activation** and copy the hardware ID shown.

3 Use a web browser on a computer that has Internet access to navigate to the PlateSpin Product Activation website (http://www.platespin.com/productactivation/ActivateOrder.aspx). Log in with your Customer Center user name and password.

4 Use the hardware ID to create a license key file. This process requires the following information:

   ◆ the activation code that you received

   ◆ the email address that you provided when you placed your order

   ◆ the hardware ID that you copied in Step 2

5 Save the generated license key file, transfer it to the product host that does not have Internet connectivity, and use it to activate the product.

6 In the Web Interface on the License Activation page, type the path to the file or browse to its location, then click **Activate**.

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

4.2 About Workload License Consumption

Your PlateSpin Protect product license entitles you to a specific or unlimited number of workloads for protection through workload licensing. Every time you add a workload for protection, the system consumes a single workload license from your license pool. You can recover a consumed license, if you remove a workload, up to a maximum of five times.

On the Dashboard page of the PlateSpin Protect Web Interface, the License Summary displays the current count of installed and consumed licenses.

The Licenses page (**Settings** > **Licenses**) lists each installed license with its current counts for consumed workload licenses and the remaining reassignments available for those licenses. The page also shows the total number of remaining unused workload licenses for the PlateSpin Server.

Figure 4-1 *License Count and Remaining Reassignments*

4.3 Viewing License Information

The product Dashboard provides a License Summary that displays the total number of installed licenses and the current number of consumed licenses.

You can view information about the workload licenses installed on a PlateSpin Server on the Licenses page. For each license, you can view the current number of used workload licenses and the current number of remaining reassignments available for used licenses.

To view license information:

1 In the Web Interface, select **Settings > Licenses**.

2 View the license information:
   - Activation Code
   - Expiry Date
   - Workloads
   - Remaining Reassignments
3 View **Workloads remaining** for the number of available unused licenses.

4.4 Adding a License

You use the same process for adding a new license as for activating the first license. See the following for information:

- Section 4.1.1, "Online License Activation," on page 47
- Section 4.1.2, "Offline License Activation," on page 48

4.5 Deleting a License

You can delete an expired license on the Licenses page.

1 In the Web Interface, select **Settings > Licenses**.

2 View the license information.

3 Click **Delete** next to the expired license, then confirm the deletion.

4.6 Generating a Licensing Report for Technical Support

If you have licensing issues, Technical Support might request you to generate a Licensing Report. This diagnostic report contains encoded product information about the licenses you have activated for your PlateSpin Server.

1 In the Web Interface, select **Settings > Licenses**.

2 Below the list of licenses, click **View Licensing Report**.

The `LicenseReport.txt` file opens in a new browser tab or window, depending on your browser settings.

3 Save the `LicenseReport.txt` file as `LicenseReport.psl` on your local computer.

5 Configuring User Authorization and Authentication

The following information is included in this section:

5.1 About PlateSpin Protect Role-Based Access

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

The role-based access mechanism enables you to implement user authorization and authentication in several ways:

- Restricting application access to specific users
- Allowing only specific operations to specific users
- Granting each user access to specific workloads for performing operations defined by the assigned role

Every PlateSpin Protect instance has the following set of operating system-level user groups that define related functional roles:

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

When a user attempts to connect to PlateSpin Protect, the credentials provided through the browser are validated by IIS. If the user is not a member of one of the Workload Protection roles, connection is refused.

Table 5-1 *Workload Protection Roles and Permission Details*

Workload Protection Role Details	Administrators	Power Users	Operators
Add Workload	Allowed	Allowed	Denied
Remove Workload	Allowed	Allowed	Denied

Workload Protection Role Details	Administrators	Power Users	Operators
Configure Protection	Allowed	Allowed	Denied
Prepare Replication	Allowed	Allowed	Denied
Run (Full) Replication	Allowed	Allowed	Allowed
Run Incremental	Allowed	Allowed	Allowed
Pause/Resume Schedule	Allowed	Allowed	Allowed
Test Failover	Allowed	Allowed	Allowed
Failover	Allowed	Allowed	Allowed
Cancel Failover	Allowed	Allowed	Allowed
Abort	Allowed	Allowed	Allowed
Dismiss (Task)	Allowed	Allowed	Allowed
Settings (All)	Allowed	Denied	Denied
Run Reports/Diagnostics	Allowed	Allowed	Allowed
Failback	Allowed	Denied	Denied
Reprotect	Allowed	Allowed	Denied

In addition, PlateSpin Protect software provides a mechanism based on *security groups* that define which users should have access to which workloads in the PlateSpin Protect workload inventory.

To set up a proper role-based access to PlateSpin Protect:

1 Add users to the required user groups detailed in Table 5-1. See your Windows documentation.

2 Create application-level security groups that associate these users with specified workloads. See "Managing PlateSpin Protect Security Groups and Workload Permissions" on page 53.

5.2 Managing PlateSpin Protect Access and Permissions

The following sections provide more information:

◆ Section 5.2.1, "Adding PlateSpin Protect Users," on page 52

◆ Section 5.2.2, "Assigning a Workload Protection Role to a PlateSpin Protect User," on page 53

5.2.1 Adding PlateSpin Protect Users

Use the procedure in this section to add a new PlateSpin Protect user.

If you want to grant specific role permissions to an existing user on the PlateSpin Server host, see "Assigning a Workload Protection Role to a PlateSpin Protect User" on page 53.

1 On your PlateSpin Server host, access the system's Local Users and Groups console (**Start** > **Run** > `lusrmgr.msc` > **Enter**).

2 Right-click the **Users** node, select **New User**.

3 Specify the required details, and click **Create**.

You can now assign a workload protection role to the newly created user. See "Assigning a Workload Protection Role to a PlateSpin Protect User" on page 53.

5.2.2 Assigning a Workload Protection Role to a PlateSpin Protect User

Before assigning a role to a user, determine the collection of permissions that best suits that user. See Table 5-1, "Workload Protection Roles and Permission Details," on page 51.

1 On your PlateSpin Server host, access the system's Local Users and Groups console (**Start** > **Run** > `lusrmgr.msc` > **Enter**).

2 Click the **Users** node, and double-click the required user in the right pane.

3 On the **Member Of** tab, click **Add**.

4 Find the required Workload Protection group, and assign it to the user.

It might take several minutes for the change to take effect. To attempt applying the changes manually, restart your server by using the `RestartPlateSpinServer.exe` executable.

To restart the PlateSpin Server:

1 Before you attempt to restart the PlateSpin Server, pause all of your contracts, or verify that no replications, failovers, or failbacks are in progress. Do not continue until all workloads are idle.

2 On the PlateSpin Server host, navigate to the `..\bin\RestartPlateSpinServer` subdirectory.

3 Double-click the `RestartPlateSpinServer.exe` executable.

   A command prompt window opens, requesting confirmation.

4 Confirm by typing `Y` and pressing `Enter`.

You can now add this user to a PlateSpin Protect security group and associate a specified collection of workloads. See "Managing PlateSpin Protect Security Groups and Workload Permissions" on page 53.

5.3 Managing PlateSpin Protect Security Groups and Workload Permissions

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

1 Assign a PlateSpin Protect user a Workload Protection Role whose permissions best suit that role in your organization. See "Assigning a Workload Protection Role to a PlateSpin Protect User" on page 53.

2 Access PlateSpin Protect as an administrator by using the PlateSpin Protect Web Interface, then click **Settings** > **Permissions**.

   The Security Groups page opens.

3 Click **Create Security Group**.

4 In the **Security Group Name** field, type a name for your security group.

5 Click **Add Users** and select the required users for this security group.

If you want to add a PlateSpin Protect user who was recently added to the PlateSpin Server host, it might not be immediately available in the user interface. In this case, first click **Refresh User Accounts**.
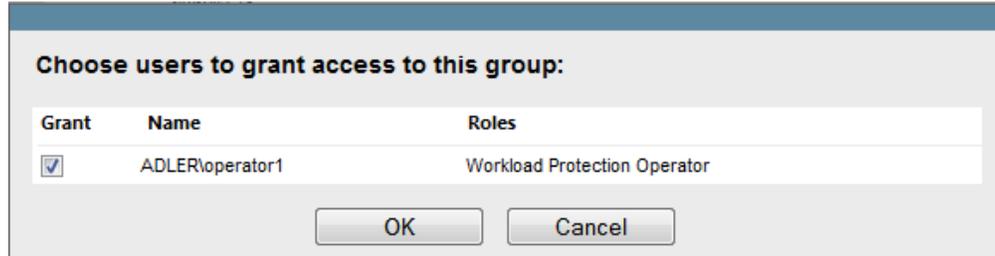
6 Click **Add Workloads** and select the required workloads:
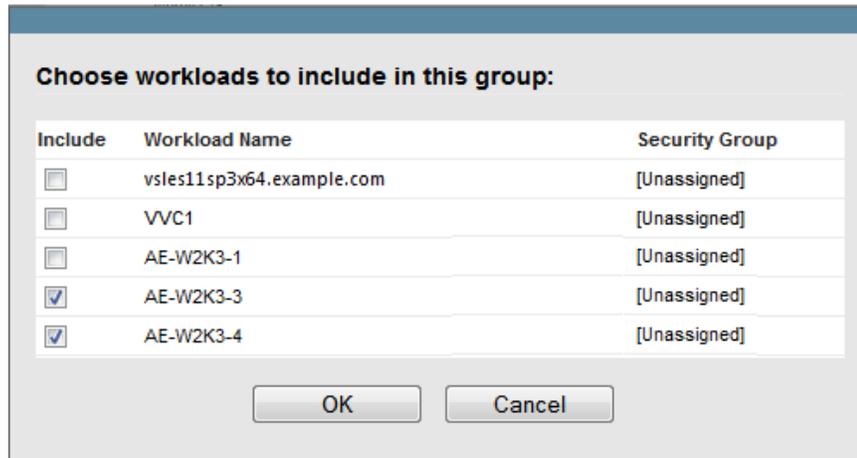
Only users in this security group will have access to the selected workloads.

7 Click **Create**.

The page reloads and displays the your new group in the list of security groups.

To edit a security group, click its name in the list of security groups.

5.4 Setting Up Protect Multitenancy on VMware

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

This section includes the following information:

5.4.1 Defining VMware Roles for Multitenancy

PlateSpin Protect requires certain privileges to access and perform tasks in the VMware Infrastructure (that is, VMware "containers"), making the Protect workflow and functionality possible in that environment. The `PlateSpinRole.xml` file defines the minimum required privileges and aggregates them respectively into three VMware custom roles:

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

This file is included in the PlateSpin Protect Server installation. An accompanying executable, `PlateSpin.VMware.Role.Tool.exe,` accesses the file to enable the creation of these custom PlateSpin roles in a target vCenter environment.

By default, the role definition file (`PlateSpinRole.xml`) and the role definition tool (`PlateSpin.VMwareRoleTool.exe`) are located in the `VMwareRolesTool` folder:

`<install-directory>\PlateSpin Protect Server\bin\VMwareRolesTool`

This section includes the following information:

- "Basic Command Line Syntax" on page 55
- "Additional Command Line Parameters and Flags" on page 56
- "Tool Usage Example" on page 56
- "(Option) Manually Defining the PlateSpin Roles in vCenter" on page 57
- "Using vCenter to View Privileges for PlateSpin Custom Roles" on page 57

Basic Command Line Syntax

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

```
PlateSpin.VMware.Role.Tool.exe /host=[host name or IP address of vCenter or ESX
host] /user=[user name] /role=[PlateSpinRole.xml] /create
```

Where `PlateSpinRole.xml` is the role definition file name.

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

Additional Command Line Parameters and Flags

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

Parameters

`/create`	(Mandatory) Creates the roles defined by the `/role` parameter
`/get_all_prvileges`	Display all server-defined privileges
`/get_compatible_roles`	Display all roles compatible to the one defined by `/role`
`/check_role=[role name]`	Check the given role for compatibility with the one defined by `/role`

Optional Flags

`/interactive`	Run the tool with interactive options that allow you to choose to create individual roles, check role compatibility, or list all compatible roles.
	For information about using the tool in interactive mode, see *VMware Role Tool to Verify Permissions to Roles* (KB 7018547) (https://www.netiq.com/support/kb/doc.php?id=7018547).
`/password=[password]`	Provide the VMware password (bypasses the password prompt)
`/verbose`	Display detailed information

Tool Usage Example

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

Resulting Actions:

1. The role definition tool runs on the `houston_sales` vCenter server, which has an administrator with the user name `pedrom`.

2. In the absence of the `/password` parameter, the tool prompts for the user password, which you enter.

3. The tool accesses the role definition file, `PlateSpinRole.xml`, which is located in the same directory as the tool executable (there was no need to further define its path).

4. The tool locates the definition file and is instructed (`/create`) to create the roles defined in the contents of that file in the vCenter environment.

5. The tool accesses the definition file and creates the new roles (including the appropriate minimum privileges for defined, limited access) inside vCenter.

   The new custom roles are to be assigned to users later in vCenter.

(Option) Manually Defining the PlateSpin Roles in vCenter

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

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

Using vCenter to View Privileges for PlateSpin Custom Roles

You use the vCenter client to view the minimal privileges set for the PlateSpin custom roles.

1 In vCenter, select a custom role:

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

2 Click **Edit** to view the privileges settings in the Edit Role dialog.

   For example, the following figure shows some of the privileges set for the PlateSpin Virtual Machine Manager role.
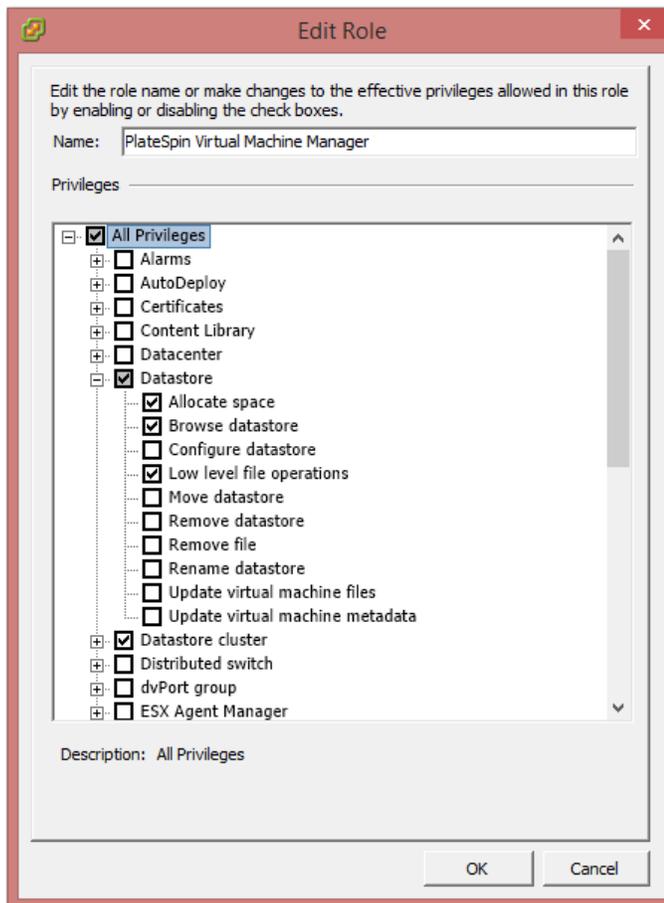
5.4.2 Assigning Roles In vCenter

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

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

vCenter Container for Role Assignment	Role Assignment Specifics	Propagate Instructions	More Information
Root of vCenter inventory tree.	Assign the enabled user the *PlateSpin Infrastructure Manager* (or equivalent) role.	For security reasons, define the permission as non-propagating.	This role is needed to monitor tasks being performed by the Protect software and to end any stale VMware sessions.
All data center objects where the enabled user needs access	Assign the enabled user the *PlateSpin Infrastructure Manager* (or equivalent) role.	For security reasons, define the permission as non-propagating.	This role is needed to allow access to the data center's datastores for file upload/download. Define the permission as non-propagating.
Each cluster to be added to Protect as a container, and each host contained in the cluster	Assign the enabled user the *PlateSpin Infrastructure Manager* (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	To assign to a host, propagate the permission from the cluster object or create an additional permission on each cluster host. If the role is assigned on the cluster object and is propagated, no further changes are necessary when you add a new host to the cluster. However, propagating this permission has security implications.
Each Resource Pool where the enabled user needs access.	Assign the enabled user the *PlateSpin Virtual Machine Manager* (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	Although you can assign access to any number of Resource Pools in any location in the tree, you must assign the enabled user this role on at least one Resource Pool.
Each VM folder where the enabled user needs access	Assign the enabled user the *PlateSpin Virtual Machine Manager* (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	Although you can assign access to any number of VM Folders in any location in the tree, you must assign the enabled user this role on at least one folder.

vCenter Container for Role Assignment	Role Assignment Specifics	Propagate Instructions	More Information
Each Network where the enabled user needs access. Distributed Virtual Networks with a dvSwitch and a dvPortgroup	Assign the enabled user the *PlateSpin Virtual Machine Manager* (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	Although you can assign access to any number of networks in any location in the tree, you must assign the enabled user this role on at least one folder. ◆ To assign the correct role to the dvSwitch, propagate the role on the data center (resulting in an additional object receiving the role) or place the dvSwitch in a folder and assign the role on that folder. ◆ For a standard port group to be listed as an available network in the Protect UI, create a definition for it on every host in the cluster.
Each Datastore and Datastore Cluster where the enabled user needs access	Assign the enabled user the *PlateSpin Virtual Machine Manager* (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	The enabled user must have been assigned this role on at least one Datastore or Datastore Cluster. For Datastore Clusters, the permission must be propagated to the contained datastores. Not providing access to an individual member of the cluster causes both prepare and full replications to fail.

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

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

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

Figure 5-1   *Roles assigned in vCenter*
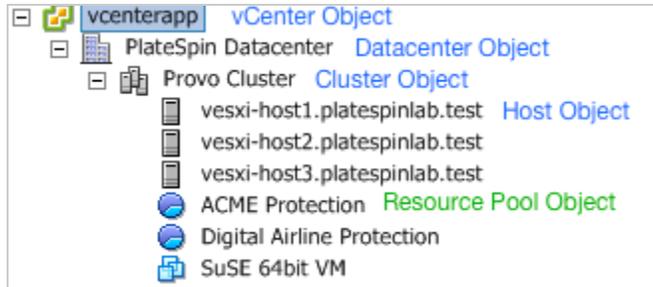
Security Implications of Assigning VMware Roles

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

* With the *PlateSpin Infrastructure Manager* role assigned to the vCenter object, every enabled user can see (but not affect) the tasks performed by every other user.

* Because there is no way to set permissions on datastore folders or subfolders, all enabled users with permissions on a datastore have access to all other enabled users' disks stored on that datastore.

* With the *PlateSpin Infrastructure Manager* role assigned to the cluster object, every enabled user is able to turn off/on HA or DRS on the entire cluster

* With the *PlateSpin User* role assigned at the storage cluster object, every enabled user is able to turn off/on SDRS for the entire cluster

* Setting the *PlateSpin Infrastructure Manager Role* on the DRS Cluster object and propagating this role allows the enabled user to see all VMs placed in the default resource pool and/or default VM folder. Also, propagation requires the administrator to explicitly set the enabled user to have a "no-access" role on every resource pool/VM folder that he or she should not have access to.

* Setting the *PlateSpin Infrastructure Manager Role* on the vCenter object allows the enabled user to end sessions of any other user connected to the vCenter.

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

6 Configuring the PlateSpin Server Application

This section describes configuration requirements and setup for PlateSpin Protect.

6.1 Configuring Language Settings for International Versions

In addition to English, PlateSpin Protect provides National Language Support (NLS) for the following international languages:

- Chinese Simplified
- Chinese Traditional
- French
- German
- Japanese

To manage your PlateSpin Server in one of these supported languages, configure the language code for the operating system on the PlateSpin Server host and in your Web browser.

6.1.1 Setting the Language on the Operating System

The language of a small portion of system messages generated by the PlateSpin Server depends on the operating system interface language selected in your PlateSpin Server host.

To change the operating system language:

1 Access your PlateSpin Server host.

2 Start the Regional and Language Options applet (click **Start > Run**, type `intl.cpl`, and press Enter), then click the **Languages** (Windows Server 2003) or **Keyboards and Languages** (Windows Server 2008) tab, as applicable.

3 If it is not already installed, install the required language pack. You might need access to your OS installation media.

4 Select the required language as the interface language of the operating system. When you are prompted, log out or restart the system.

6.1.2 Setting the Language in Your Web Browser

To use the PlateSpin Protect Web Interface in one of these languages, the corresponding language must be added in your web browser and moved to the top of the order of preference:

1 Access the Languages setting in your web browser:

- **Chrome:**
    1. From the Chrome menu, click **Settings**, then scroll to and click **Show advanced settings**.
    2. Scroll to **Languages**, then click **Language and input settings**.
- **Firefox:**
    1. From the **Tools** menu, select **Options**, then select the **Content** tab.
    2. Under **Languages**, click **Choose**.
- **Internet Explorer:**
    1. From the **Tools** menu, select **Internet Options**, then select the **General** tab.
    2. Under **Appearance**, click **Languages**.

2 Add the required language and move it up the top of the list.

3 Save the settings, then start the client application by connecting to your PlateSpin Server. See "Launching the Web Interface" on page 39.

NOTE: (For users of Chinese Traditional and Chinese Simplified languages) Attempting to connect to the PlateSpin Server with a browser that does not have a specific version of Chinese added might result in web server errors. For correct operation, use your browser's configuration settings to add a specific Chinese language (for example, `Chinese [zh-cn]` or `Chinese [zh-tw]`). Do not use the culture-neutral `Chinese [zh]` language.

6.2 Configuring Email Notification Services for Events and Replication Reports

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

- Section 6.2.1, "Configuring SMTP for the Email Notification Service," on page 62
- Section 6.2.2, "Enabling Event Notifications," on page 63
- Section 6.2.3, "Enabling Replication Reports," on page 65

6.2.1 Configuring SMTP for the Email Notification Service

Use the PlateSpin Protect Web Interface to configure SMTP (Simple Mail Transfer Protocol) settings for the server used to deliver email notifications of events and replication reports.

Figure 6-1 *Simple Mail Transfer Protocol Settings*

To configure SMTP settings:

1 In your PlateSpin Protect Web Interface, click **Settings > SMTP**.

2 Specify the SMTP server settings for receiving email event and progress notifications:

  ◆ **Address**

  ◆ **Port** (the default is 25)

  ◆ **Reply Address**

3 Type a user name and password, then confirm the password.

4 Click **Save**.

6.2.2 Enabling Event Notifications

Events are always added to the System Application Event Log, according to the log entry types of Warning, Error, and Information. You can also enable Notifications to automatically send event notifications to appropriate recipients.

1 Set up an SMTP server for PlateSpin Protect to use.

See "Configuring SMTP for the Email Notification Service" on page 62.

2 In your PlateSpin Protect Web Interface, click **Settings > Notification Settings**.

3 Select the **Enable Notifications** option.

4 Click **Edit Recipients**, type the required email addresses separated by commas, then click **OK**.

5 Click **Save**.

To delete listed email addresses, click **Remove** next to the address.

The event types shown in Table 6-1 can trigger email notifications if Email Notification is enabled.

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

Table 6-1 *Events Types Organized by Log Entry Types*

Event Types	Remarks
Log Entry Type: Warning	
FullReplicationMissed	Similar to the `Incremental Replication Missed` event.
IncrementalReplicationMissed	Generated when any of the following applies: ◆ A replication is manually paused while a scheduled incremental replication is due. ◆ The system attempts to carry out a scheduled incremental replication while a manually-triggered replication is underway. ◆ The system determines that the target has insufficient free disk space.
WorkloadOfflineDetected	Generated when the system detects that a previously online workload is now offline. Applies to workloads whose protection contract's state is not **Paused**.
Log Entry Type: Error	
FailoverFailed	
FullReplicationFailed	
IncrementalReplicationFailed	
PrepareFailoverFailed	
Log Entry Type: Information	
FailoverCompleted	
FullReplicationCompleted	
IncrementalReplicationCompleted	
PrepareFailoverCompleted	
TestFailoverCompleted	Generated upon manually marking a Test Failover operation a success or a failure.
WorkloadOnlineDetected	Generated when the system detects that a previously offline workload is now online. Applies to workloads whose protection contract's state is not **Paused**.

6.2.3 Enabling Replication Reports

You can enable Replication Reports to automatically send reports to appropriate recipients.

1 Set up an SMTP server for PlateSpin Protect to use.

See "Configuring SMTP for the Email Notification Service" on page 62.

2 In your PlateSpin Protect Web Interface, click **Settings > Replication Reports Settings**.

3 Select the **Enable Replication Reports** option.

4 In the **Report Recurrence** section, click **Edit**, then specify the appropriate recurrence pattern for the reports. You can click **Close** to collapse the section.

5 In the **Recipients** section, click **Edit Recipients**, type the appropriate email addresses separated by commas, then click **OK**. You can click **Remove** next to an email address to delete the recipient from the list.
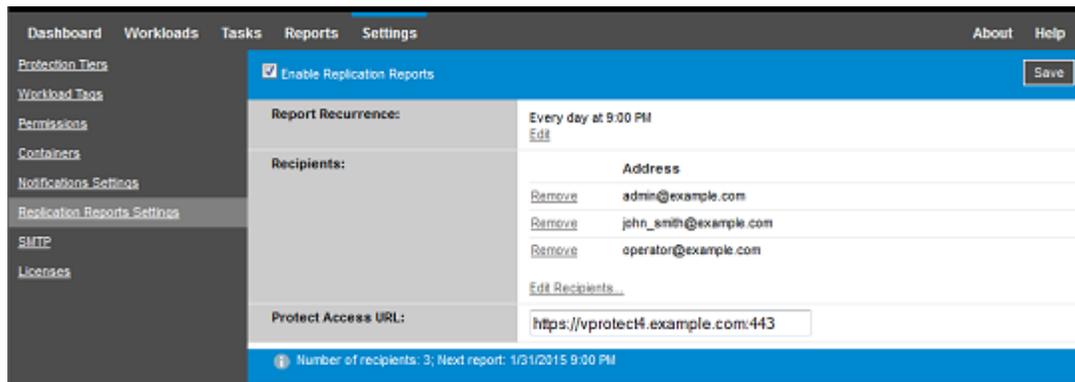
6 (Optional) In the Protect **Access URL** section, specify a non-default URL for your PlateSpin Server (for example, when your PlateSpin Server host has more than one NIC or if it is located behind a NAT server). This URL affects the title of the report and the functionality of accessing relevant content on the server through hyperlinks within emailed reports.

7 Click **Save**.

For information on other types of reports that you can generate and view on demand, see "Generating Workload and Workload Protection Reports" on page 178.

6.3 Configuring Alternate IP Addresses for PlateSpin Server

You can add alternate IP addresses to the PlateSpin Configuration `AlternateServerAddresses` parameter in order to enable the PlateSpin Server to function across NAT-enabled environments.

To add alternate IP addresses for PlateSpin Server:

1 From any web browser, open

`https://Your_PlateSpin_Server/platespinconfiguration/`

2 Search to locate the `AlternateServerAddresses` parameter and add IP addresses for the PlateSpin Server.

3 Save your settings and exit the page.

A reboot or restart of PlateSpin services is not required to apply the changes.

6.4 Optimizing Data Transfer over WAN Connections

You can optimize data transfer performance and fine tune it for WAN connections. You do this by modifying configuration parameters that the system reads from settings you make in a configuration tool residing on your PlateSpin Server host. For more information, see Section 3.5.1, "PlateSpin Configuration," on page 44.

- ◆ Section 6.4.1, "Tuning Parameters," on page 66
- ◆ Section 6.4.2, "Tuning FileTransferSendReceiveBufferSize," on page 68

6.4.1 Tuning Parameters

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

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

Table 6-2 lists the configuration parameters on the PlateSpin Configuration page (`https://Your_PlateSpin_Server/platespinconfiguration/`) that control file transfer speeds with the defaults and maximum values. You can modify these values through trial-and-error testing in order to optimize operation in a high-latency WAN environment.

Table 6-2   Default and Optimized File Transfer Configuration Parameters

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

Parameter	Default Value	Maximum Value
`FileTransferMaxBufferSizeInBytes` Specifies internal buffer size for holding packet data.	31457280	
`FileTransferMaxPacketSizeInButes` Determines the largest packets that will be sent.	1048576	
`FileTransferMinCompressionLimit` Specifies the packet-level compression threshold in bytes.	0 (disabled)	max 65536 (64 KB)
`FileTransferPort`	3725	

`FileTransferSendReceiveBufferSize`
 	
Defines the maximum size (in bytes) of the send and receive buffers for TCP connections in the replication network. The buffer size affects the TCP Receive Window (RWIN) size, which sets the number of bytes that can be sent without TCP acknowledgment. This setting is relevant for both file-based and block-based transfers. Tuning the buffer size based on your network bandwidth and latency improves throughput and reduces CPU processing.
 	
When the value is set to zero (off), the default TCP window size is used (8 KB). For custom sizes, specify the size in bytes.
 	
Use the following formula to determine the proper value:
 	
((LINK_SPEED in Mbps / 8) * DELAY in sec)) * 1000 * 1024
 	
For example, for a 100 Mbps link with 10 ms latency, the proper buffer size would be:
 	
`(100/8)* 0.01 * 1000 * 1024 = 128000 bytes`
 	
For tuning information, see Section 6.4.2, "Tuning FileTransferSendReceiveBufferSize," on page 68.
	0 (8192 bytes)	max 5242880 (5 MB)

`FileTransferSendReceiveBufferSizeLinux` Specifies the TCP/IP Receive Window (RWIN) Size setting for file transfer connections for Linux. It controls the number of bytes sent without TCP acknowledgment, in bytes. When the value is set to zero (off), the TCP/IP window size value for Linux is automatically calculated from the `FileTransferSendReceiveBufferSize` setting. If both parameters are set to zero (off), the default value is 248 KB. For custom sizes, specify the size in bytes. **NOTE:** In previous release versions, you were required to set this parameter to 1/2 the desired value, but this is no longer required.	0 (253952 bytes)	
`FileTransferShutDownTimeOutInMinutes`	1090	
`FileTransferTCPTimeOutMilliSec` Sets both the TCP Send and TCP Receive Timeout values.	30000	
`PostFileTransferActionsRequiredTimeInMinutes`	60	

6.4.2 Tuning FileTransferSendReceiveBufferSize

The FileTransferSendReceiveBufferSize parameter defines the maximum size (in bytes) of the send and receive buffers for TCP connections in the replication network. The buffer size affects the TCP Receive Window (RWIN) size, which sets the number of bytes that can be sent without TCP acknowledgment. This setting is relevant for both file-based and block-based transfers. Tuning the buffer size based on your network bandwidth and latency improves throughput and reduces CPU processing.

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

To calculate the optimum buffer size:

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

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

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

   1b Enter the following at a command prompt:

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

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

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

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

   1d Note the latency in seconds.

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

2 Calculate a byte value for your initial buffer size:

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

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

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

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

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

   3a Determine the MSS:

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

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

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

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

   3b Calculate the optimal buffer size:

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

To continue the example:

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

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

```
Non-optimal Buffer Size = 32080 * 1430 = 45874400
```

6.5 Optimizing Replication Environment Performance

Use the Take Control and Snapshot configuration parameters settings to optimize replication performance. These settings are global and affect all replications.

Table 6-3 lists the configuration parameters on the PlateSpin Configuration page (`https://Your_PlateSpin_Server/platespinconfiguration/`) that control the replication environment with the default values.

Table 6-3 *Default Configuration Parameters for the Replication Environment*

Parameter	Default Value
`TakeControlMemorySizeInMB`	768
The memory size (in MB) to set when taking control for replication.	
`TakeControlCoresPerSocket`	1
The number of virtual cores per socket to use for taking control, when the target is booted to LRD or `bootofx.iso`.	
`TakeControlSockets`	1
The number of virtual sockets to use for taking control, when the target is booted into LRD or `bootofx.iso`.	
`MaximumConcurrentReplications`	25
The number of concurrent replications that can be running at the same time.	
`VssSnapshotCreationDelay`	120
The number of seconds to delay between retry attempts when creating a VSS snapshot during replication.	
`VssSnapshotCreationRetryCount`	5
The maximum number of attempts to create a VSS snapshot during replication before the replication attempt fails.	

6.6 Setting Reboot Method for the Configuration Service

During a failover action, the Configuration Service optimizes reboots by minimizing the number of reboots and controlling when they occur. If you experience a Configuration Service hang during a failover action for a Windows workload with an error `Configuration Service Not Started`, you

might need to allow reboots to occur as they are requested during the configuration. You can configure the single affected workload to skip reboot optimization, or configure a global `SkipRebootOptimization` setting on the PlateSpin Server to skip reboot optimization for all Windows workloads.

To skip reboot optimization for a single Windows workload:

1 Log on as an Administrator user on the source workload.

2 Add a file at the root of the system drive (usually `C:`) called `PlateSpin.ConfigService.LegacyReboot` with no file extension. From a command prompt, enter

```
echo $null >> %SYSTEMDRIVE%\PlateSpin.ConfigService.LegacyReboot
```

3 Run the failed Test Failover or Failover action again.

To skip reboot optimization for all Windows workloads:

1 Log in to the PlateSpin Server, then open the PlateSpin Server Configuration page at:

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

2 Search for the configuration parameter **ConfigurationServiceValues**, then click **Edit** for the parameter.

3 Change the setting **SkipRebootOptimization** from `false` to `true`.

4 Click **Save**.

5 Run an incremental or full replication.

The replication also propagates the modified configuration settings to the target VM.

6 Run the Test Failover or Failover again for affected Windows workloads.

6.7 Configuring Support for VMware vCenter Site Recovery Manager

You might use PlateSpin Protect to protect your workloads locally and then use some additional method to replicate those workloads to a remote location, such as a SAN. For example, you might choose to use VMware vCenter Site Recovery Manager (SRM) to replicate an entire datastore of replicated target VMs to a remote site. In this case, specific configuration steps are needed to ensure that the target VMs can be replicated and behave correctly when powered on at the remote site.

Workloads replicated by PlateSpin Protect and managed on VMware vCenter SRM can behave seamlessly if you configure PlateSpin Protect to support SRM by making the following adjustments:

◆ Configure a setting to keep the PlateSpin Protect ISO and floppies on the same datastore as the VMware `.vmx` and `.vmdk` files.

◆ Prepare the PlateSpin Protect environment to copy VMware Tools to the failover target. This involves some manual file creation and copying in addition to making some configuration settings that expedite the VMware Tools installation process.

6.7.1 Setting Up Workload Files on the Same Datastore

To ensure that the workload files are kept on the same datastore:

1 From any web browser, open `https://Your_PlateSpin_Server/platespinconfiguration/` to display the configuration web page.

2 On the configuration web page, locate the `CreatePSFilesInVmDatastore` server parameter and change its value to `true`.

> **NOTE:** The person configuring the replication contract is responsible to ensure that the same datastore is specified for all target VM disk files.

3 Save your settings and exit the page.

6.7.2 Setting Up VMware Tools for Failover Targets

VMware Tools setup packages can be copied to the failover target during replication so that they can be installed by the configuration service when the VM is booted. This happens automatically when the failover target is able to contact the PlateSpin Server. In cases where this cannot happen, you need to prepare your environment prior to replication.

To prepare your environment:

1 Retrieve the VMware Tools packages from an ESX host:

  1a Secure copy (`scp`) the `windows.iso` image from the `/usr/lib/vmware/isoimages` directory on an accessible VMware host to a local temporary folder.

  1b Open the ISO and extract its setup packages, saving them to an accessible location:

   ◆ **VMware 5.*x* and later:** The setup packages are `setup.exe` and `setup64.exe`.

   ◆ **VMware 4.*x*:** The setup packages are `VMware Tools.msi` and `VMware Tools64.msi`.

2 Create OFX packages from the setup packages you extracted:

  2a Zip the package you want, making sure that the setup installer file is at the root of the `.zip` archive.

  2b Rename the `.zip` archive to `1.package` so that it can be used as an OFX package.

> **NOTE:** If you want to create an OFX package for more than one of the setup packages, remember that each setup package must have its own unique `.zip` archive.
>
> Because each package must have the same name (`1.package`), if you want to save multiple `.zip` archives as OFX packages, you need to save each in its own unique subdirectory.

3 Copy the appropriate OFX package (`1.package`) to the `%ProgramFiles(x86)%\PlateSpin\Packages\%GUID%` directory on the PlateSpin Server.

  The value of `%GUID%` depends on the version of your VMware Server and its VMware Tools architecture, as shown in Table 6-4. Use the appropriate GUID value to copy the package to the correct directory.

Table 6-4 *GUIDs for the VMware Tools Directory Names*

VMware Server Version	VMware Tools Architecture	GUID
6.5	x86	`D61C0FCA-058B-42C3-9F02-898F568A3071`
6.5	x64	`5D3947B7-BE73-4A00-A549-B15E84B98803`
6.0	x86	`311E672E-05BA-4CAF-A948-B26DF0C6C5A6`
6.0	x64	`D7F55AED-DA64-423F-BBBE-F1215529AD03`
5.5	x86	`660C345A-7A91-458b-BC47-6A3914723EF7`
5.5	x64	`8546D4EF-8CA5-4a51-A3A3-6240171BE278`
5.1	x86	`34DD2CBE-183E-492f-9B36-7A8326080755`
5.1	x64	`AD4FDE1D-DE86-4d05-B147-071F4E1D0326`
5.0	x86	`AD4FDE1D-DE86-4d05-B147-071F4E1D0326`
5.0	x64	`F7C9BC91-7733-4790-B7AF-62E074B73882`
4.1	x86	`F2957064-65D7-4bda-A52B-3F5859624602`
4.1	x64	`80B1C53C-6B43-4843-9D63-E9911E9A15D5`
4.0	x86	`D052CBAC-0A98-4880-8BCC-FE0608F0930F`
4.0	x64	`80B50267-B30C-4001-ABDF-EA288D1FD09C`

6.7.3 Expediting the Configuration Process

After the failover target boots, the configuration service launches to prepare the VM for use, but sits inactive for several minutes, waiting for data from the PlateSpin Server or looking for VMware Tools on the CD ROM.

To shorten this wait time:

1 On the configuration web page, locate the `ConfigurationServiceValues` configuration setting, and then change the value of its `WaitForFloppyTimeoutInSecs` subsetting to zero (0).

2 On the configuration web page, locate the `ForceInstallVMToolsCustomPackage` and change the value to `true`.

With these settings in place, the configuration process takes less than 15 minutes: the target machine reboots (up to two times), the VMware tools are installed, and SRM accesses the tools to help it configure networking at the remote site.

7 Configuring PlateSpin Web Interface

PlateSpin Web Interface enables you to configure tags to use to track logical associates among workloads. In addition, you can control screen refresh rates for several pages. Use the information in this section to configure your Web Interface.

- Section 7.1, "Creating and Managing Workload Tags," on page 73
- Section 7.2, "Configuring Refresh Rates for the Web Interface," on page 75
- Section 7.3, "Customizing the UI for the Web Interface," on page 76

7.1 Creating and Managing Workload Tags

When you have a large number of workloads to manage, it can be time-consuming to browse the list and select similar workloads for concurrent operation actions. Sorting on name or feature can help. Another alternative is to use a tag to set up a custom association among workloads that you want to manage as a group. You can easily sort the workloads by the Tag column, select the appropriate tagged workloads, and run available operations on them at the same time.

A tag can represent any logical or physical association for a workload that is meaningful to you. You affiliate a unique color and name for each tag. You can create as many unique tags as you like, although the choice of unique colors is limited. Each workload can have a single tag associated with it. When you export a workload to a new server, its tag setting persists.

- Section 7.1.1, "Creating a Workload Tag," on page 73
- Section 7.1.2, "Editing a Workload Tag," on page 74
- Section 7.1.3, "Adding a Tag to a Workload," on page 74
- Section 7.1.4, "Removing a Tag from a Workload," on page 74
- Section 7.1.5, "Deleting a Workload Tag," on page 75

7.1.1 Creating a Workload Tag

1 In the PlateSpin Protect Web Interface, click **Settings** > **Workload Tags** > **Create Workload Tag**.
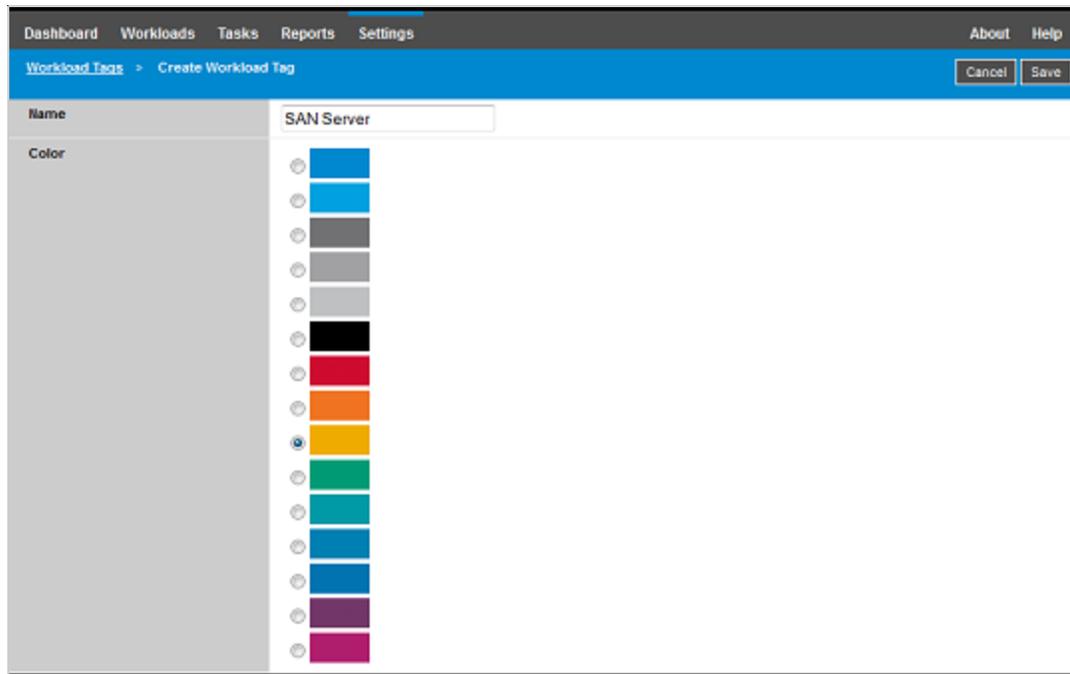
2 Specify a unique tag name (25-character limit) and associate a color with that description.

3 Click **Save** to add this new tag to the list of available workload tags in the Workload Tags view of Settings page.

7.1.2 Editing a Workload Tag

1 In the PlateSpin Protect Web Interface, click **Settings** > **Workload Tags**.

2 Edit any of the available tags. Click the tag name, modify its name or affiliated color, then click **Save**.

7.1.3 Adding a Tag to a Workload

1 In the workload list, select the active workload you want to tag, then click **Configure** to open its configuration page.

2 Expand the **Tag** section to view the **Tag** drop-down box.

3 Select the name of the tag you want to associate with the workload, then click **Save**.

7.1.4 Removing a Tag from a Workload

1 In the workload list, select the workload, then click **Configure** to open its configuration page.

2 Expand the **Tag** section to view the **Tag** drop-down box.

3 Select the "empty" line in the list of available tag names, then click **Save**.

7.1.5 Deleting a Workload Tag

You can delete any tags that you no longer use. You cannot delete a tag if it is associated with any workload.

1 In the PlateSpin Protect Web Interface, click **Settings** > **Workload Tags**.

2 Disassociate the tag of interest from workloads.

3 Click **Delete** next to the tag, then click **OK** to confirm.

7.2 Configuring Refresh Rates for the Web Interface

Several pages in the Web Interface have configurable refresh intervals, as shown in Table 7-1. You can modify the interval setting to meet the needs of your PlateSpin environment.

Table 7-1 *Web Interface Default Refresh Intervals*

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

1 Open the following file in a text editor:

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

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

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

3 Save the file.

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

7.3   Customizing the UI for the Web Interface

You can modify the appearance of PlateSpin Web Interface to match the look and feel of your corporate identity. You can modify colors, logo, and product name. For more information, see Appendix A, "Rebranding the PlateSpin Protect Web Interface," on page 81.

8 Managing Multiple PlateSpin Servers in the Management Console

PlateSpin Protect includes a Web-based client application, the Management Console, that provides centralized access to multiple instances of PlateSpin Protect and PlateSpin Forge.

In a data center with more than one instance of PlateSpin Protect and PlateSpin Forge, you can designate one of the instances as the manager and run the management console from there. Other instances are added under the Manager, providing a single point of control and interaction.

8.1 Using the PlateSpin Protect Management Console

To begin using the Management Console:

1 Open a web browser on a machine that has access to your PlateSpin Protect instances and navigate to:

`https://`*Your_PlateSpin_Server*`/console`

Replace *Your_PlateSpin_Server* with either the IP address or the DNS host name of the PlateSpin Server host that is designated as the Manager.

2 Log in with your user name and password.

3 (Initial Logon) On the Welcome page, click **Add PlateSpin Server**, then set up a PlateSpin Server instance as described in Section 8.3, "Adding Instances of PlateSpin Protect and PlateSpin Forge to the Management Console," on page 79.

4 (Subsequent Logons) View the dashboard.
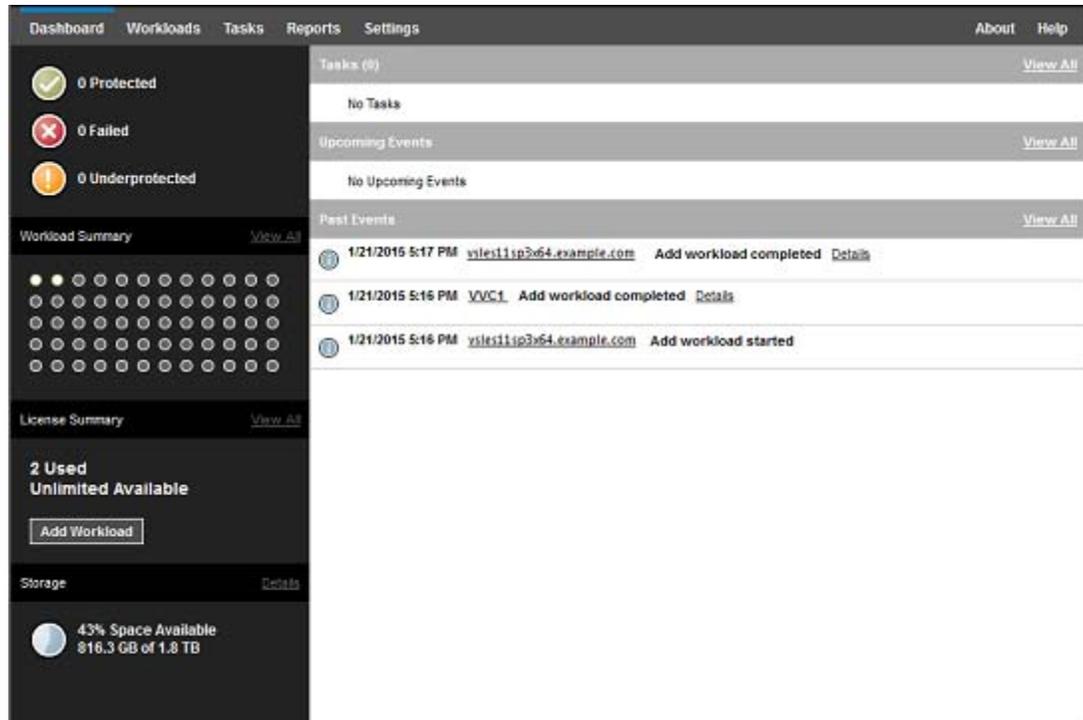
8.2 About PlateSpin Protect Management Console Cards

Individual instances of PlateSpin Protect and PlateSpin Forge, when added to the Management Console, are represented by cards.

Figure 8-1 *PlateSpin Protect Instance Card*
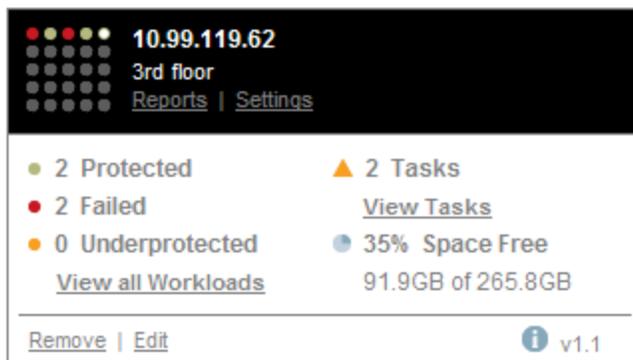
A card displays basic information about the specific instance of PlateSpin Protect and PlateSpin Forge, such as:

- ◆ IP address/hostname
- ◆ Location
- ◆ Version number

- Workload count
- Workload status
- Storage capacity
- Remaining free space

Hyperlinks on each card allow you to navigate to that particular instance's Workloads, Reports, Settings, and Tasks pages. There are also hyperlinks that allow you to edit a card's configuration or remove a card from the display.

8.3 Adding Instances of PlateSpin Protect and PlateSpin Forge to the Management Console

Adding a PlateSpin Protect or PlateSpin Forge instance to the Management Console results in a new card on the Management Console's dashboard.

NOTE: When you log in to the Management Console running on an instance of PlateSpin Protect or PlateSpin Forge, that instance is not automatically added to the console. It must be manually added.

To add a PlateSpin Protect or PlateSpin Forge instance to the console:

1 On the console's main dashboard, click **Add PlateSpin Server**.

2 Specify the URL of the PlateSpin Server host or Forge VM. Use HTTPS if SSL is enabled.

3 (Optional) Enable the **Use Management Console Credentials** check box to use the same credentials as those used by the console. When it is selected, the console automatically populates the **Domain\Username** field.

4 In the **Domain\Username** field, type a domain name and a user name valid for the instance of PlateSpin Protect or PlateSpin Forge that you are adding. In the **Password** field, type the corresponding password.

5 (Optional) Specify a unique descriptive **Display Name** (up to 15 characters) for the PlateSpin Server, its **Location** (up to 20 characters), and any **Notes** you might require (up to 400 characters).

6 Click **Add**.

A new card is added to the dashboard.

8.4 Editing Cards on the Management Console

To modify the details of a card on the Management Console:

1 In the Management Console, locate the card instance for the PlateSpin Protect server or PlateSpin Forge server that you want to modify.

2 Click the **Edit** hyperlink on the card.

  The console's **Add/Edit** page is displayed.

3 Make any desired changes, then click **Add/Save**.

  The updated console dashboard is displayed.

8.5 Removing Cards on the Management Console

To remove a card from the Management Console:

1 In the Management Console, locate the card instance for the PlateSpin Protect server or PlateSpin Forge server that you want to remove.

2 Click the **Remove** hyperlink on the card.

  A confirmation prompt is displayed.

3 Click **OK** to confirm.

  The card instance is removed from the dashboard.

A Rebranding the PlateSpin Protect Web Interface

You can modify the appearance of the Web Interface to match the look and feel of your corporate identity, including colors, logo, and product name. You can even eliminate the links to **About** tab and **Help** tab in the product interface.

This section includes information to help you change the branding of the product:

- Section A.1, "Rebranding the Web Interface By Using Configuration Parameters," on page 81
- Section A.2, "Rebranding the Product Name in the Windows Registry," on page 84

A.1 Rebranding the Web Interface By Using Configuration Parameters

You can change the look and feel of the Web Interface to match the proprietary look of your organization websites. To customize the branding of the Web Interface, modify the configuration parameters of your PlateSpin Server host.

To modify Web Interface branding parameters:

1 From any web browser, open `https://Your_PlateSpin_Server/platespinconfiguration/`, then log in as Administrator.

2 Locate the required server parameter, click **Edit**, then change its value.

For more information, see Figure A-1 to view the configurable elements in the UI. See Table A-1 to view the setting name, description, and default value information for each configurable element.

3 Save and your settings and exit the page.

Although no reboot or restart of services is required after the change is made in the configuration tool, it might take up to 30 seconds for the change to take effect in the interface.

A.1.1 Web Interface Configurable Elements

The look and feel of the Web Interface is consistent throughout. The illustration of the PlateSpin Protect Dashboard in Figure A-1 identifies the elements you can modify with numbered callouts. The inset shows the configurable elements in the Settings panel.

Figure A-1 *Protect Web Interface with Configurable Elements Labeled*
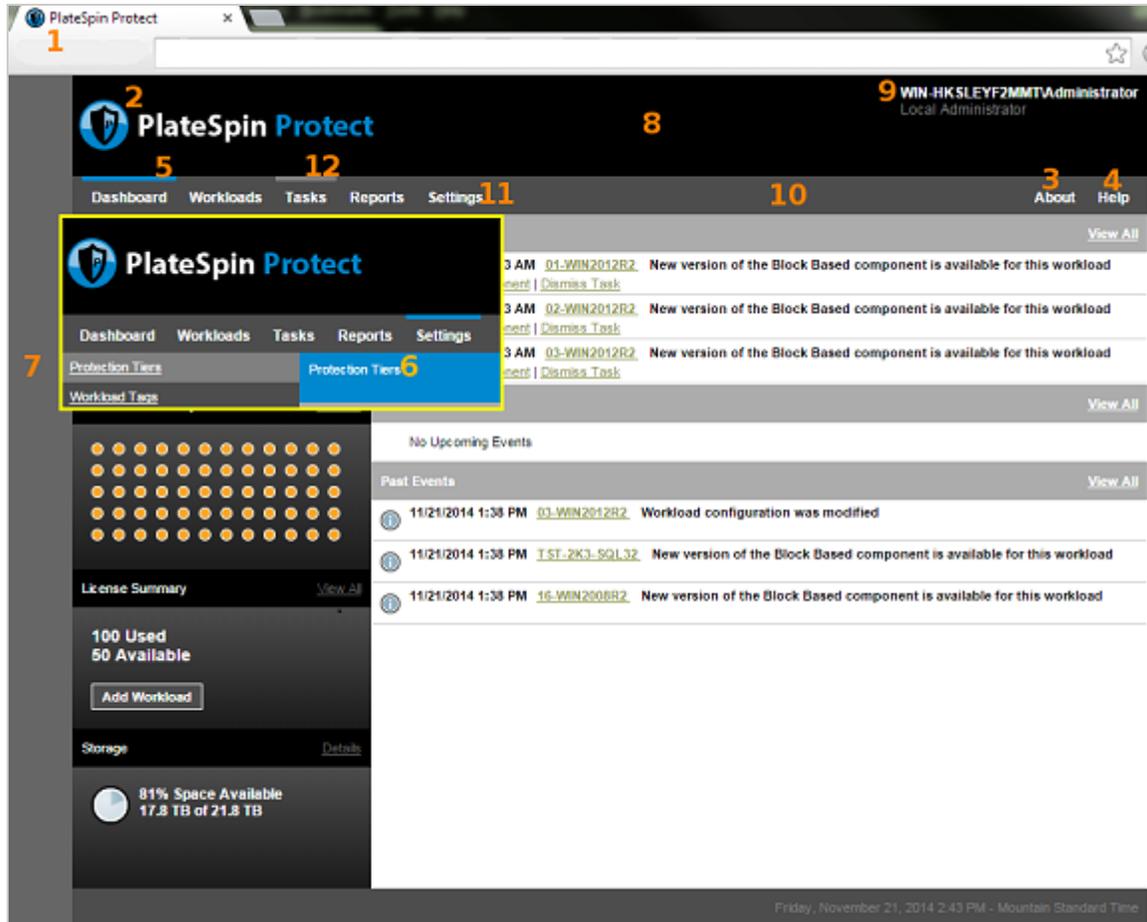
A.1.2 Web Interface Configurable Parameters

The table below lists the identified interface element (or "ID") in the screen shot above, the setting name, description, and default value. Use the PlateSpin Server Configuration Settings page to change these values (that is, on the settings page, click **Edit** on a configuration value), according to the new "look and feel" you want.

Table A-1 *Web Interface Configuration Parameters and Default Values*

ID	Setting Name and Description	Default Value
1	**WebUIFaviconUrl** Location of a valid `.ico` graphic file. Specify one of the following: ◆ A valid URL to the appropriate `.ico` file on a different machine. For example: https://myserver.example.com/dir1/dir2/icons/mycompany_favicon.ico ◆ A relative path below the root of the local web server where you have uploaded the appropriate `.ico` file. For example, if you create a path called `mycompany\images\icons` at the root of the web server to store your custom icon graphics: `~/mycompany/images/icons/ mycompany_favicon.ico` In this example, the actual file system path that contains the file is `C:\Program Files (x86)\PlateSpin Protect Server\PlateSpin Forge\web\mycompany\images\icons\mycompany_f avicon.ico`.	`~/doc/en/favicon.ico` [1]
2	**WebUILogoUrl** Location of product logo graphic file. Specify one of the following: ◆ A valid URL to the appropriate graphics file on a different machine. For example: https://myserver.example.com/dir1/dir2/logos/mycompany_logo.png ◆ A relative path below the root of the local web server where you have uploaded the appropriate graphics file. For example, if you create a path called `mycompany\images\logos` at the root of the web server to store your custom logo images: `~/mycompany/images/logos/mycompany_logo.png` In this example, the actual file system path that contains the file is `C:\Program Files (x86)\PlateSpin Protect Server\PlateSpin Forge\web\mycompany\images\logos\mycompany_l ogo.png`.	`~/Resources/protectLogo.png` [2]
3	**WebUIShowAboutTab** Toggle the visibility of the **About** tab on (**True**) or off (**False**).	True
4	**WebUIShowHelpTab** Toggle the visibility of the **Help** tab on (**True**) or off (**False**).	True

ID	Setting Name and Description	Default Value
5	`WebUISiteAccentColor` Accent color (RGB hex value)	#0088CE
6	`WebUISiteAccentFontColor` Font color to display with accent color in Web UI (RGB hex value)	#FFFFFF
7	`WebUISiteBackgroundColor` Site background color (RGB hex value)	#666666
8	`WebUISiteHeaderBackgroundColor` Site header background color (RGB hex value)	#000000
9	`WebUISiteHeaderFontColor` Site header font color in Web UI (RGB hex value)	#FFFFFF
10	`WebUISiteNavigationBackgroundColor` Color of site navigation background in Web UI (RGB hex value)	#4D4D4D
11	`WebUISiteNavigationFontColor` Color of site navigation link font color in Web UI (RGB hex value)	#FFFFFF
12	`WebUISiteNavigationLinkHoverBackgroundColor` Color of site navigation link background in hover state (RGB hex value)	#808080

[1] Actual file path is `C:\Program Files (x86)\PlateSpin Protect Server\PlateSpin Forge\web\doc\en\favicon.ico`.

[2] Actual file path is `C:\Program Files (x86)\PlateSpin Protect Server\PlateSpin Forge\web\Resources\protectLogo.png`.

A.2   Rebranding the Product Name in the Windows Registry

The masthead at the top of the product interface provides space for both a corporate logo, and the name of the product itself. You can change the logo, which commonly includes the product name, using a configuration parameter. To change or eliminate the product name in a browser tab, you need to make a change in the Windows Registry.

To change the product name:

1 At the PlateSpin Server, run `regedit`.

2 In the Windows Registry Editor, navigate to the following registry key:

`HKEY_LOCAL_MACHINE\SOFTWARE\PlateSpin\ProtectServer\ProductName`

NOTE: In some cases, the registry key can be found in this location:

`HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\PlateSpin\Protect`

3 Double-click the `ProductName` key and change the **Value data** for the key as you prefer, then click **OK**.

4 Restart the IIS Server for the interface change to take effect.
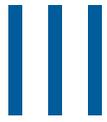
III Preparing Protection Targets and Sources

Before you can configure protection contracts, you must identify your planned target containers and source workloads. You get details about targets and workloads through an inventory process.

9 Preparing Containers (Protection Targets)

A container is a protection infrastructure that acts as the host of a protected workload's regularly-updated replica. Adding a target container populates the PlateSpin Protect database with detailed inventory information about the container and its resources. The inventory provides the data necessary to determine the container's use and to properly configure one or more workload protection contracts for the target container.

- ◆ Section 9.1, "About Containers (Protection Targets)," on page 89
- ◆ Section 9.2, "Adding Containers (Protection Targets)," on page 90
- ◆ Section 9.3, "Refreshing Container Details," on page 91
- ◆ Section 9.4, "Removing Containers (Protection Targets)," on page 92

9.1 About Containers (Protection Targets)

PlateSpin Web Interface provides automated inventory of supported target container platforms.

- ◆ Section 9.1.1, "Supported Containers," on page 89
- ◆ Section 9.1.2, "Network Access Requirements for Containers," on page 89
- ◆ Section 9.1.3, "Parameter Guidelines for Containers," on page 89

9.1.1 Supported Containers

Before you add a container to the PlateSpin Server, ensure that the VM container version is supported. See "Supported VM Containers" on page 17.

9.1.2 Network Access Requirements for Containers

Before you begin inventory operations, ensure that PlateSpin Server can communicate with your source workloads and targets. See Section 1.5.2, "Network Requirements for Containers," on page 30.

9.1.3 Parameter Guidelines for Containers

Table 9-1 provides guidelines for machine type selection, credential format, and syntax for inventory parameters for target hosts using the Web Interface.

Table 9-1 *Guidelines for Web Interface Discovery Parameters for Target Containers*

To Discover	Target Type	Credentials
VMware vCenter Cluster	**VMware DRS Cluster**	VMware vCenter Web service credentials (user name and password)

To Discover	Target Type	Credentials
VMware ESXi Server	**VMware ESX Server**	ESX account with administrator role OR Windows domain credentials (versions 4 and 4.1 only)

9.2 Adding Containers (Protection Targets)

A container is a protection infrastructure that acts as the host of a protected workload's regularly-updated replica. That infrastructure can be either a VMware ESX Server or a VMware DRS Cluster. PlateSpin Protect allows you to use containers for both protection and Failback.

To be able to protect a workload, you must have a a workload and a container inventoried by (or *added to*) the PlateSpin Server.

To add a container:

1 In your Web Interface, select **Settings > Containers > Add Container**.
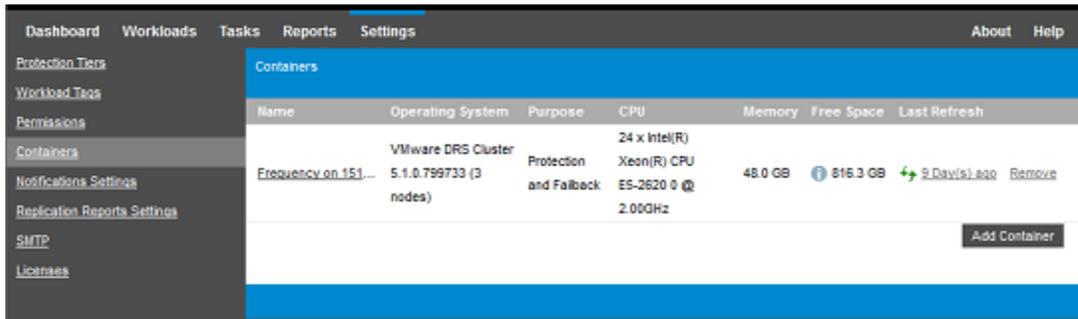
2 Specify the type of container:
- ◆ **VMware ESX Server**
- ◆ **VMware DRS Cluster**

3 Depending on the type of targets you selected in the previous step, specify the appropriate access information.

Table 9-2 *Options for VMware DRS Cluster Target*

Option	Description
vCenter Hostname or IP	Specify the host name or IP address of the vCenter server.

Option	Description
vCenter Hostname or IP	Specify the host name or IP address of the vCenter server.

- ◆ **VMware DRS Cluster:** See Table 9-3.
- ◆ **VMware ESX Server:** See Table 9-4.

Table 9-3 *Options for VMware DRS Cluster Target*

Option	Description
vCenter Hostname or IP	Specify the host name or IP address of the vCenter server.

Option	Description
vCenter Hostname or IP	Specify the host name or IP address of the vCenter server.

Table 9-4 *Options for VMware ESX Server Target*

Option	Description
Hostname or IP	Specify the host name or IP address of the VMware ESX server.
Username and Password	Specify administrator-level credentials for accessing the target container. See "Guidelines for Workload and Container Credentials" on page 165.

4 Click **Test Credentials** to validate the credential values you specified.

5 Select the purpose for the VM container:
   - **Protection**
   - **Failback**
   - **Protection and Failback**

   Selecting both **Protection** and **Failback** results in that container being available for selection as a target in both protection and Failback operations.

6 Click **Add** to add and discover details about the container and list it on the Containers page.

   PlateSpin Protect reloads the Containers page and displays a process indicator for the container being added 🔄. On completion, the process indicator icon turns into a **Refresh** icon ↺.

9.3   Refreshing Container Details

You should routinely refresh details about your target containers before setting up or executing a protection contract. PlateSpin Web Interface enables you to refresh the discovered resources for virtual target containers.

When you refresh the target, its associated resources are automatically rediscovered and updated. You can refresh one container at a time.

To refresh details for a target container:

1 In the PlateSpin Web Interface, select **Settings > Containers**.

2 Click the **Refresh** icon ↺ next to the container you want to refresh.

   This performs a re-inventory of the container.

3 Expand the panels on the Container Details page for information about the inventory changes.

9.4 Removing Containers (Protection Targets)

If you remove all protection contracts for a target container, you can remove (undiscover) the target container. You might also remove a container that will not be used.

IMPORTANT: Before you delete an target container that is in use for configured workload protection contract, you must ensure that all the affected contracts removed or reconfigured for a different target container.

To remove a target through the Web Interface:

1 In the PlateSpin Web Interface, select **Settings > Containers**.

2 On the Containers page, click **Remove** next to the container you want to remove from Protect.

10 Preparing Workloads (Protection Sources)

For any protection contract, you must have a source workload and a target container. Adding a workload to the PlateSpin Protect Server populates the PlateSpin database with detailed inventory information about the machine. This information provides the data necessary to determine the machine's use and to properly configure a protection contract.

- Section 10.1, "About Workloads (Protection Sources)," on page 93
- Section 10.2, "Adding Workloads (Protection Sources)," on page 94
- Section 10.3, "Tagging Workloads," on page 95
- Section 10.4, "Refreshing Workload Details," on page 96
- Section 10.5, "Removing Workloads," on page 96

10.1 About Workloads (Protection Sources)

PlateSpin Web Interface provides automated inventory of supported source workload configurations

- Section 10.1.1, "Supported Workloads," on page 93
- Section 10.1.2, "Network Access Requirements for Source Workloads," on page 93
- Section 10.1.3, "Parameter Guidelines for Source Workloads," on page 94

10.1.1 Supported Workloads

Before you add a workload to the PlateSpin Server, ensure that the workload operating system version and hardware is supported. See the following sections in Section 1.1, "Supported Configurations," on page 13:

- "Supported Windows Workloads" on page 14
- "Supported Linux Workloads" on page 15
- "Supported Workload Architectures" on page 19
- "Supported Storage" on page 20

10.1.2 Network Access Requirements for Source Workloads

For information about network access requirements for inventory for Windows and Linux workloads, see Section 1.5.3, "Network Requirements for Workloads," on page 30.

10.1.3 Parameter Guidelines for Source Workloads

Table 10-1 provides guidelines for machine type selection, credential format, and syntax for inventory parameters for workloads.

Table 10-1 *Guidelines for Discovery Parameters for Workloads*

To Discover	Machine Type	Credentials	Remarks
All Windows workloads	Windows	Local or Domain Admin credentials.	For the username, use this format: ◆ 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 KB Article 7920711 (https://www.netiq.com/support/kb/doc.php?id=7920711).

10.2 Adding Workloads (Protection Sources)

A workload, the basic object of protection in a data store, is an operating system, along with its middleware and data, decoupled from the underlying physical or virtual infrastructure.

To protect a workload, you must have a workload and a container inventoried by (or *added to*) the PlateSpin Server.

To add a workload:

1 Follow the required preparatory steps.

   See Preparation in "Basic Workflow for Workload Protection and Recovery" on page 35.

2 On the Dashboard or Workloads page, click **Add Workload**.

   The Web Interface displays the Add Workload page.

3 Specify the required workload details:

 ◆ **Workload Settings:** Specify your workload's host name or IP address, the operating system, and administrator-level credentials.

 Use the required credential format. See "Guidelines for Workload and Container Credentials" on page 165.

 To ensure that PlateSpin Protect can access the workload, click **Test Credentials**.

4 Click **Add Workload**.

 PlateSpin Protect reloads the Workloads page and displays a process indicator for the workload being added . Wait for the process to complete. Upon completion, a **Workload Added** event is shown on the Dashboard, and the new workload becomes available on the Workloads page.

5 (Conditional) If you haven't added a container yet for use with this workload, add one to prepare for protecting the workload. See "Preparing Containers (Protection Targets)" on page 89.

6 Continue with "Configuring Protection Details and Preparing the Replication" on page 153.

10.3 Tagging Workloads

In the PlateSpin Web Interface, the Workloads page might display a long list of workloads. Searching through these workloads to manage operations for similar workloads can be time-consuming. To overcome this issue, you can create tags for various workload categories, departments, or other logical associations appropriate to your environment.

For information about creating, modifying, or deleting workload tags, see Section 7.1, "Creating and Managing Workload Tags," on page 73.

After you create tags, they are available at the bottom of the Edit Target Details page where you can assign a tag to the appropriate workloads. The Workloads page includes a **Tag** column where the single tag you associate with a workload is displayed. You can sort on this column to group similar workloads together. This enables you to easily locate and run operations on the tagged workloads at the same time.

NOTE: When you export a workload with a tag setting to a new server, the tag settings persist.

To associate a tag with a workload during Configure Protection:

1 In the Protect Web Interface, click **Workloads.**

2 In the workload list, select the workload you want to tag and click **Configure Protection**.

3 Configure the workload.

4 In the Tag section at the bottom of the Edit Target Details page, select the tag name you want to associate with the workload

5 Click **Save**.

To add or modify a tag associated with configured workload:

1 In the Protect Web Interface, click **Workloads.**

2 In the workload list, click the workload you want to tag to open the Target Details page.

3 Click **Edit**.

4 In the Tag section at the bottom of the Edit Target Details page, select the tag name you want to associate with the workload.

5 Click **Save**.

To disassociate a tag from a workload:

1 In the Protect Web Interface, click **Workloads.**

2 In the workload list, select the workload for which you want to remove the tag and click **Configure Protection**.

3 In the Tag section of the configuration page, select the empty string and click **Save**.

10.4 Refreshing Workload Details

PlateSpin Web Interface does not support refreshing details for the discovered workloads. To update details about a discovered workload, you must remove the workload, and then add and discover its details again. Configuration details are lost if the workload is in a configured state when you remove it. If a protection license is in use, it is removed from the workload and returned to the license pool. See Section 10.5, "Removing Workloads," on page 96.

10.5 Removing Workloads

In some circumstances you might need to remove a workload from the Protect inventory and re-add it later.

1 On the Workloads page, select the workload that you want to remove, then click **Remove Workload**.

2 (Conditional, Windows) For Windows workloads previously protected through block-level replication, the Web Interface prompts you to indicate whether you also want to remove the Block-Based Components. You can make the following selections:

   ◆ **Do not remove components:** The components will not be removed.

   ◆ **Remove components but do not restart workload:** The components will be removed. However, a reboot of the workload will be required to complete the uninstallation process.

   ◆ **Remove components and restart workload:** The components will be removed, and the workload will be automatically rebooted. Ensure that you carry out this operation during scheduled downtime.

3 On the Command Confirmation page, click **Confirm** to execute the command.

   Wait for the process to complete.

4 (Conditional, Linux) For Linux workloads, manually uninstall the block-based driver from the source workload. See Block-level data transfer software in Cleaning Up Linux Workloads.

11 Preparing Device Drivers for Physical Failback Targets

PlateSpin Protect provides a library of device drivers and PnP (Plug and Play) IDs that are needed if you have physical machines as failback targets. You can add custom device drivers and PnP ID mappings by using the PlateSpin Device Driver tool (`DeviceDriver.exe`).

11.1 Managing Device Drivers

PlateSpin Protect ships with a library of device drivers. It automatically installs the appropriate device drivers on target workloads. If some drivers are missing or incompatible on the physical failback target machine, or if you require specific drivers for a target infrastructure, you might need to add (upload) drivers to the PlateSpin Protect driver database.

11.1.1 Packaging Device Drivers for Windows Workloads

You must package your Windows device drivers to prepare them for upload to the PlateSpin Protect driver database.

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

- All 64-bit Windows systems
- 32-bit versions of Windows Server 2008 systems

To package Windows device drivers:

1 Prepare all interdependent driver files (`*.sys`, `*.inf`, `*.dll`, and so on) for your target infrastructure and device.

   If you have obtained manufacturer-specific drivers as a `.zip` archive or an executable, extract them first.

2 Save the driver files in separate folders, with one folder per device.

The package is now ready for upload. See "Uploading Driver Packages to the PlateSpin Device Driver Database" on page 98.

11.1.2   Packaging Device Drivers for Linux Workloads

You must package your Linux device drivers to prepare them for upload to the PlateSpin Protect driver database. A custom utility for this purpose is included in your PlateSpin boot ISO image (`bootofx.x2p.iso`).

1 On a Linux workstation, create a directory for your device driver files. All the drivers in the directory must be for the same kernel and architecture.

2 Download the boot image and mount it.

   For example, assuming that the ISO has been copied under the `/root` directory, issue this command for BIOS firmware-based targets and for UEFI firmware-based targets:

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

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

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

   `tar -xvzf packageModules.tar.gz`

4 Enter the working directory and execute the following command:

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

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

   Drivername-driverversion-dist-kernelversion-arch`.pkg`

   For example,

   `bnx2x-1.48.107-RHEL4-2.6.9-11.EL-i686.pkg`

The package is now ready for uploading. See "Uploading Driver Packages to the PlateSpin Device Driver Database" on page 98.

11.1.3   Uploading Driver Packages to the PlateSpin Device Driver Database

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

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

- "Device Driver Upload Procedure (Windows)" on page 98
- "Device Driver Upload Procedure (Linux)" on page 100

Device Driver Upload Procedure (Windows)

1 Obtain and prepare the required device drivers. See "Packaging Device Drivers for Windows Workloads".

2 Log in as an Administrator user to the PlateSpin Server host.

3 Launch the PlateSpin Driver Manager tool. Navigate to `C:\Program Files\PlateSpin Protect Server\DriverManager`, then start the `DriverManager.exe` program.

4 Select **Tools** > **Manage Device Drivers**, then select the **Windows Drivers** tab.
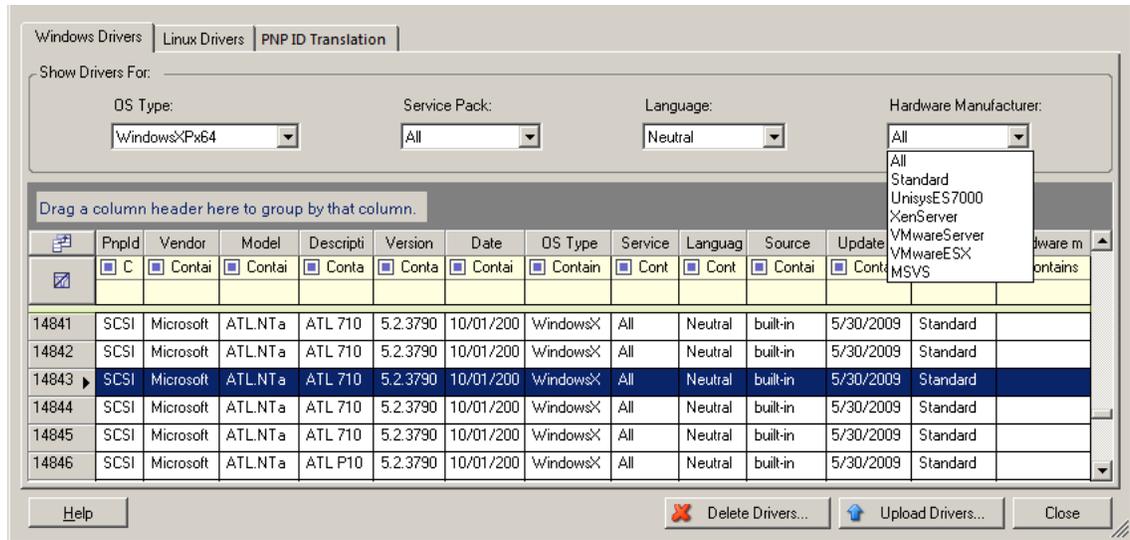
5 At the bottom of the dialog, click **Upload Drivers**.

6 On the Driver Selection dialog, browse to the folder that contains the required driver files, and select applicable OS type, language, and hardware manufacturer options.

Select **Standard** as the **Hardware Manufacturer** option, unless your drivers are designed specifically for any of the target environments listed.
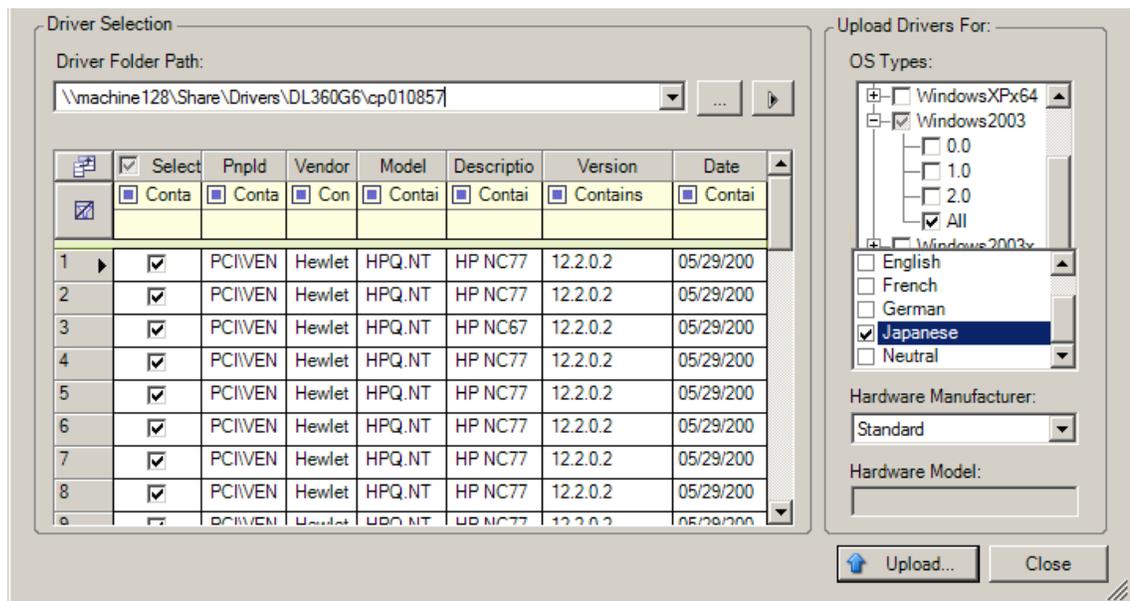
7 Click **Upload** and confirm your selections when you are prompted.

The system uploads the selected drivers to the driver database.

Device Driver Upload Procedure (Linux)

1 Obtain and prepare the required device drivers. See "Packaging Device Drivers for Linux Workloads".

2 Log in as an Administrator user to the PlateSpin Server host.

3 Launch the PlateSpin Driver Manager tool. Navigate to `C:\Program Files\PlateSpin Protect Server\DriverManager`, then start the `DriverManager.exe` program.

4 Select **Tools** > **Manage Device Drivers**, then select the **Linux Drivers** tab.

	Driver Name	Architecture	Update Date	Driver Source	Driver Type	Kernel Version
1	nicstar	x86	8/9/2010 5:38:23	BaseLine	Unknown	2.6.18-8.el5
2	lanai	x86	8/9/2010 5:38:20	BaseLine	Unknown	2.6.9-5.ELsmp
3	aacraid	x86_64@ofxboot	8/9/2010 5:38:17	BaseLine	Unknown	2.6.16.21-0.8-default
4	serverworks	x86_64@ofxboot	8/9/2010 5:38:17	BaseLine	Unknown	2.6.16.21-0.8-default
5	sungem	x86	8/9/2010 5:38:22	BaseLine	Unknown	2.6.9-55.EL
6	istallion	x86	8/9/2010 5:38:25	BaseLine	Unknown	2.4.19-4GB
7	e752x_edac	x86	8/9/2010 5:38:21	BaseLine	Unknown	2.6.9-42.ELsmp
8	mptfc	x86@VMwareES	8/9/2010 5:39:42	BaseLine	Unknown	2.6.5-7.244-default
9	gdth	x86	8/9/2010 5:38:22	BaseLine	Unknown	2.6.9-55.EL
10	zr36067	x86	8/9/2010 5:38:25	BaseLine	Unknown	2.4.19-4GB
11	snd-ens1371	x86	8/9/2010 5:38:24	BaseLine	Unknown	2.6.18-8.el5
12	mptfc	x86_64@VMware	8/9/2010 5:39:37	BaseLine	Unknown	2.6.9-42.ELsmp
13	mptsas	x86_64@VMware	8/9/2010 5:39:36	BaseLine	Unknown	2.6.9-34.ELlargesmp

5 At the bottom of the dialog, click **Upload Drivers**.

6 Browse to the folder that contains the required driver package (`*.pkg`), and click **Upload All Drivers**.

The system uploads the selected drivers to the driver database.

11.2 Managing the PlateSpin PnP ID Mappings

"Plug and Play" (PnP) refers to Windows operating system functionality that supports connectivity, configuration, and management with native plug and play devices. In Windows, the feature facilitates discovery of PnP compliant hardware devices attached to a PnP compliant bus. PnP compliant devices are assigned a set of Device Identification Strings by their manufacturer. These strings are programmed into the device when it is built. These strings are fundamental to how PnP works: they are part of the Windows' information source used to match the device with a suitable driver.

When the PlateSpin Server discovers workloads and their available hardware, the discovery includes these PnP IDs and the storage of that data as part of the workload's details. PlateSpin uses the IDs to determine which, if any, drivers need to be injected during a failover/failback operation. The PlateSpin

Server maintains a database of PnP IDs for the associated drivers of each of the supported operating systems. Because Windows and Linux use different formats for PnP IDs, a Windows workload discovered by the Protect Linux RAM disk (LRD) contains Linux-style PnP IDs.

These IDs are formatted consistently, so PlateSpin can apply a standard transformation to each of them to determine its corresponding Windows PnP ID. The translation occurs automatically within the PlateSpin product.

You (or a support technician) can use the PNP ID Translation option in the PlateSpin Device Driver tool to add, edit, or remove custom PnP ID mappings.

To add custom PnP ID mappings:

1 Log in as an Administrator user to the PlateSpin Server host.

2 Launch the PlateSpin Driver Manager tool. Navigate to `C:\Program Files\PlateSpin Protect Server\DriverManager`, then start the `DriverManager.exe` program.

3 Connect to the PlateSpin Server.

   `https://localhost/Protect`

4 In the Driver Manager tool, select the **PNP ID Translation** tab to open the **PNP ID Translation** list, which includes the currently known custom PnP ID mappings.
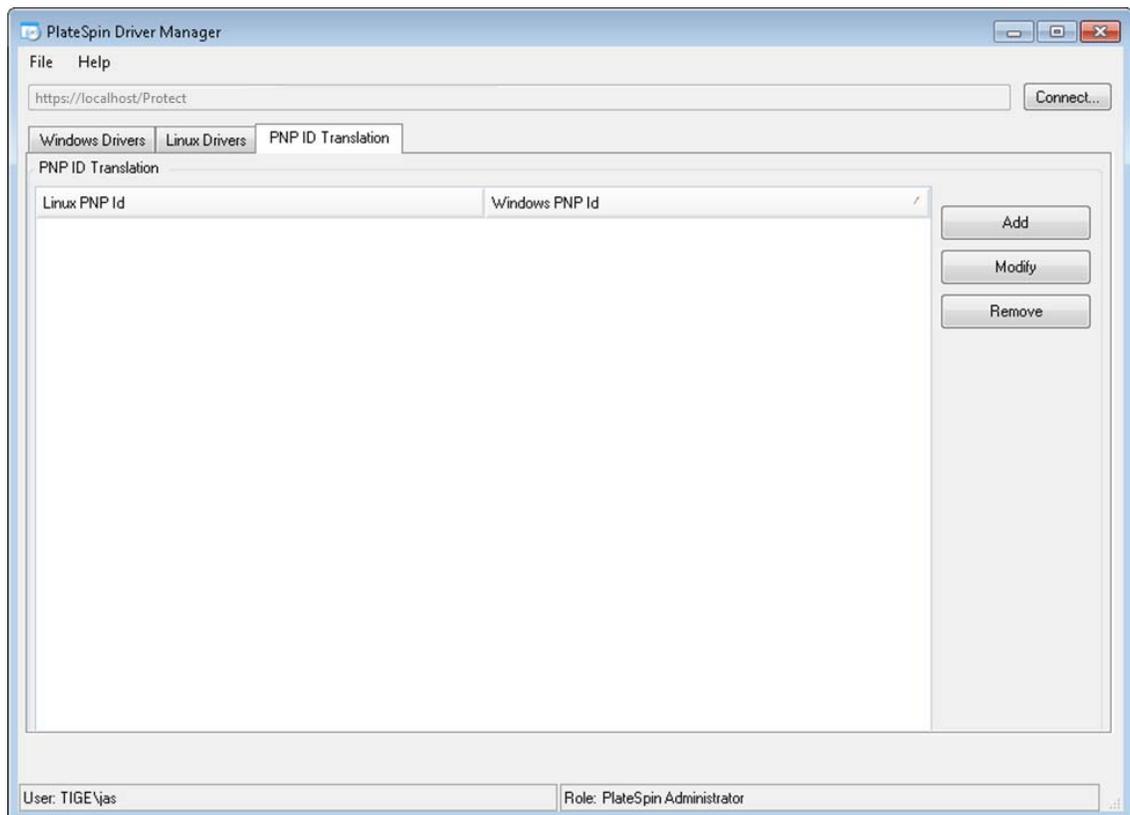
5 On the list page, click **Add** to display the Create PNP ID Mapping dialog.
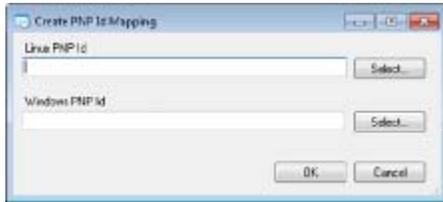
6 In the **Linux PNP ID** field, add a Linux PnP ID.

   6a (Conditional) If you know it, type the Linux PnP ID you want to use.

   or

   6b (Conditional) Select an ID from a previously discovered workload:

       6b1 Adjacent to the **Linux PnP ID** field, click **Select** to open the Select Linux PNP ID dialog.
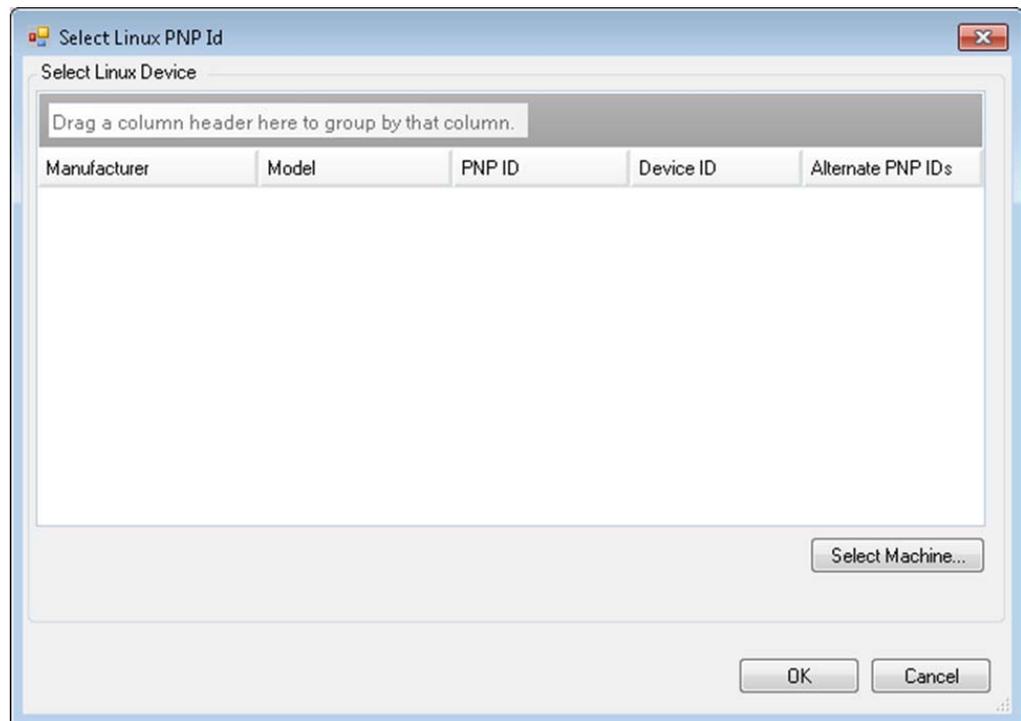
       6b2 On the dialog, click **Select Machine** to display a list of the machines previously discovered by the PlateSpin Linux RAM disk.

       6b3 Highlight one of the devices in the list, then click **Select** to populate the list in the Select Linux PnP ID dialog.
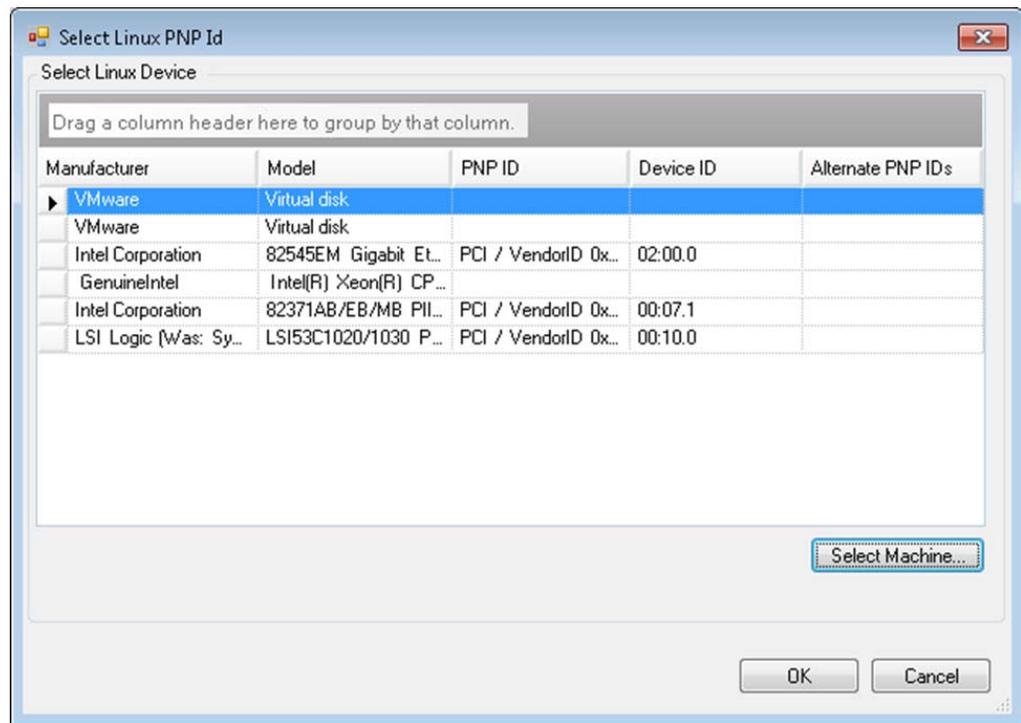
6b4 Select a device on the list, then click **OK** to apply the standard transformation to the PnP ID and display it in the Create PNP ID Mapping dialog.

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

7a (Conditional) If you know it, type the Windows PnP ID you want to use.

or

7b (Conditional) Adjacent to the **Windows PNP ID** field, click **Select** to open a mapping tool that presents three methods for helping you map a the Windows PnP ID:

◆ Under the **Driver File** tab, browse to and select a Windows driver file (that is, a file with the *.inf extension), select the desired PnP ID, then click **Modify**.
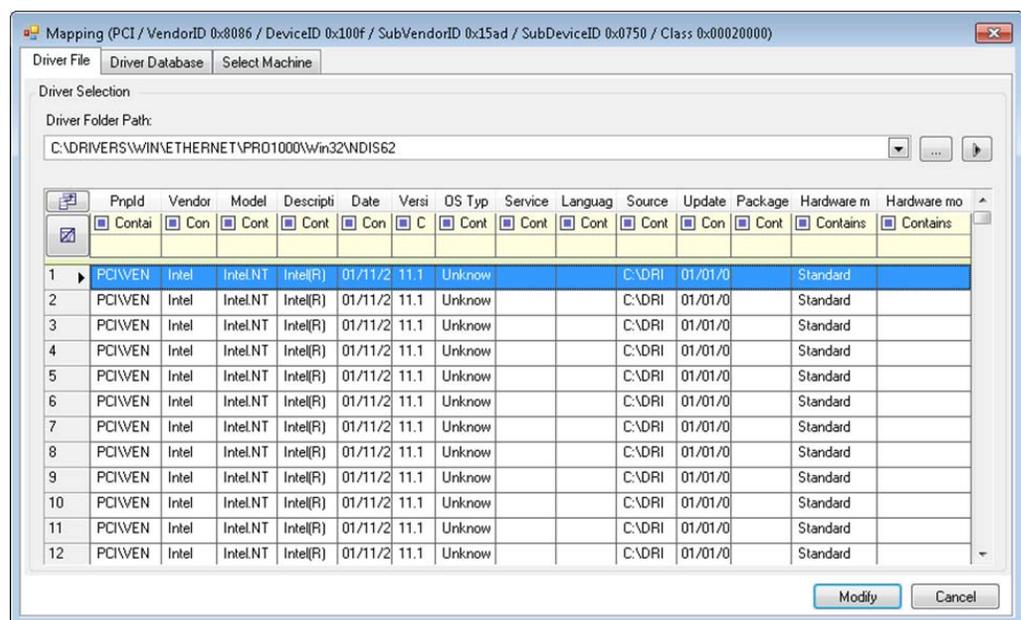
- Under the **Driver Database** tab, browse to and select the existing driver database, select the correct PnP ID, then select **Modify**.
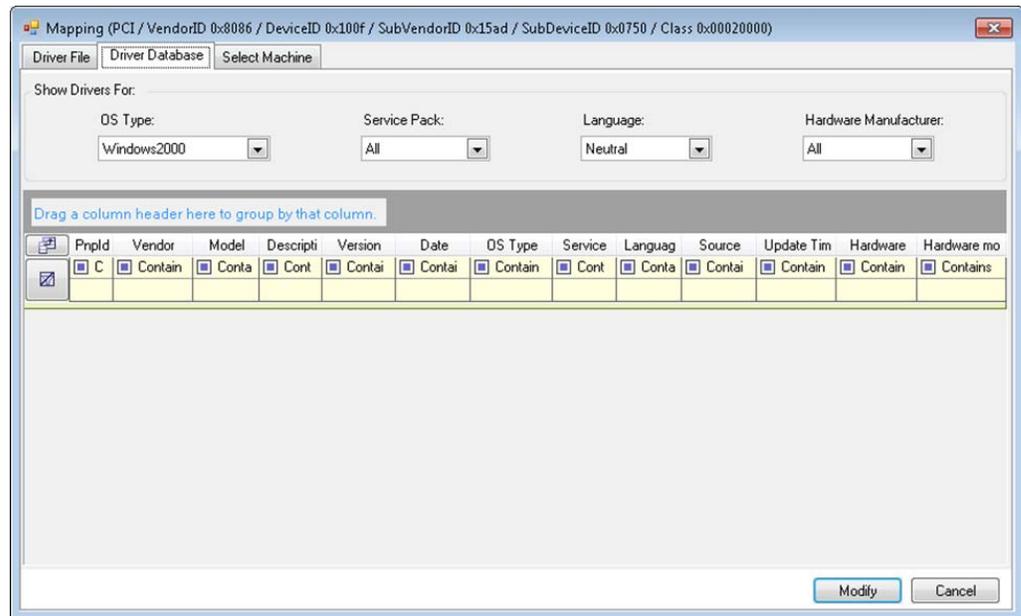
- Under the **Select Machine** tab, click **Select Machine**, then, from the list of Windows machines discovered using live discovery, select a machine, click **OK** to display its devices, select the desired PnP ID, then click Modify.
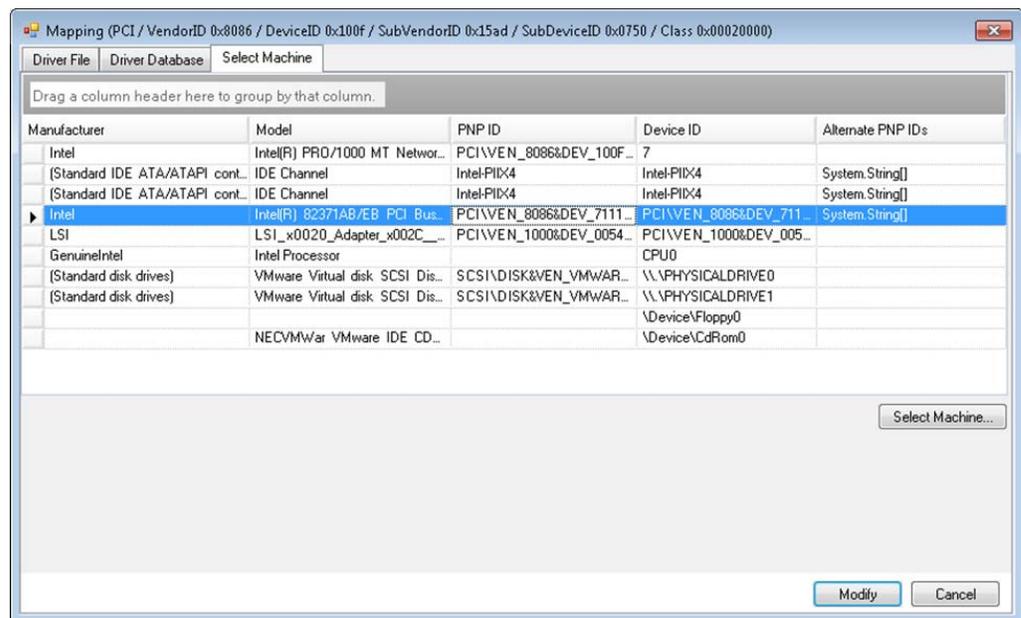
IMPORTANT: Selecting a Windows PnP ID that does not have an associated driver package installed might result in a failure at failover/failback time.

8 In the Create PnP Id Mapping dialog, confirm that the correct Linux PnP ID and the correct Windows PnP are selected, then click **OK** to display the PNP ID Translation page of the PlateSpin Driver Manager.
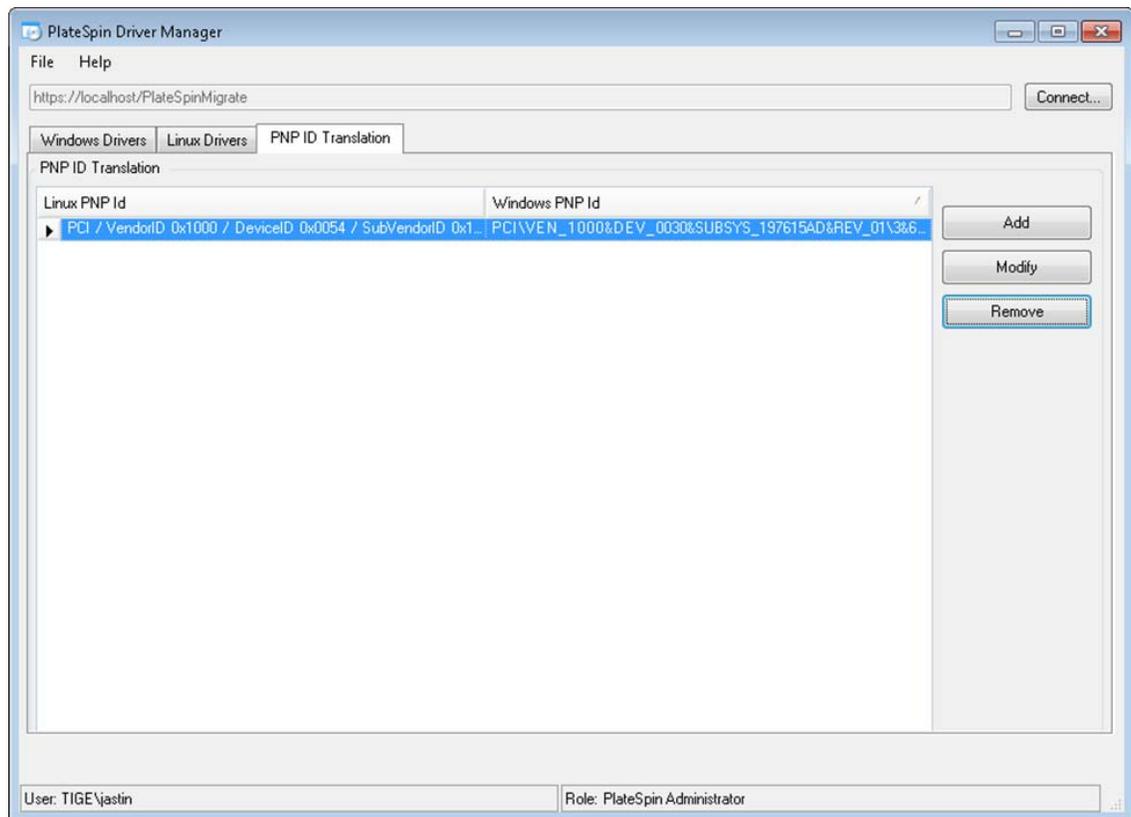
9 (Optional) To modify or remove the mapping in the PNP ID Translation list, select the mapping pattern, then click **Remove** or **Modify**, depending on the operation you want to perform.

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

To modify,

9a Click **Modify** to open the Create PNP ID Mapping dialog.

9b Repeat Step 7 on page 103 to modify the Windows PnP ID.

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

12 Preparing Linux Workloads for Protection

Perform the tasks in this section to prepare your Linux workloads for protection in PlateSpin Protect.

12.1 Verifying Block-Based Drivers for Linux

Verify that a blkwatch module is available for the workload's Linux distribution. For a list of preconfigured drivers, see "Linux Distributions Supported by Protect" on page 127.

If you plan to protect a supported Linux workload that has a non-standard, customized, or newer kernel, rebuild the PlateSpin `blkwatch` module, which is required for block-level data replication.

See Knowledgebase Article 7005873 (https://www.netiq.com/support/kb/doc.php?id=7005873).

12.2 Preparing Snapshots for Block-Level Transfer (Linux)

We recommend that you prepare snapshots for block-level data transfer. Ensure that each volume group has sufficient free space for snapshots (at least 10% of the sum of all partitions). If snapshots are not available, Protect locks and releases each block in turn on the source workload for data transfer.

12.2.1 Configuring LVM Snapshots for Linux Volume Replication

The `blkwatch` driver leverages LVM snapshots if they are available. Copying blocks from the snapshot helps avoid potential open file conflicts.

For LVM storage, see Knowledgebase Article 7005872 (https://www.netiq.com/support/kb/doc.php?id=7005872).

12.2.2 Configuring NSS Snapshots for NSS Pool Replication

For Linux workloads running Open Enterprise Server, the LVM snapshot solution is not available for NSS pools. During replication for NSS pools, Protect locks and releases each block in turn for data transfer. To avoid potential open file conflicts and to improve replication performance, you can leverage NSS pool snapshots for replication.

You can add a single unformatted disk to use for all NSS pool snapshots, or you can add a separate unformatted disk for each NSS pool. The best performance occurs when you add a separate disk for each pool. Add the disk before you set up the workload protection. You will prepare the disk to use, and PlateSpin will configure the NSS snapshots for the pool during replication.

NOTE: By default, PlateSpin uses the NLVM managed disk that has the largest amount of free space (unpartitioned space) for NSS pool snapshots. If you see the NSS pool snapshots for replication being located on the same disk as your root file system or on another disk that will be under constant disk IO, you should use the /etc/platespin/platespin.conf file to direct the NSS snapshots to an appropriate disk.

For information about how NSS snapshots work on Open Enterprise Server, see "Guidelines for Using and Managing Pool Snapshots" (http://www.novell.com/documentation/oes2015/stor_nss_lx/data/br18up4.html) in the *NSS File System Administration Guide for Linux*.

To set up one or multiple disks to use for snapshots of NSS pools:

1 On the OES source workload, add an unformatted Linux disk to use for snapshots of all NSS pools. You can alternatively create a separate disk for each NSS pool.

   The size of the disk should be about 20% of the amount of used data on the NSS pool. Adjust the size according to the amount of data change or growth that might occur during the interval of time for a replication.

2 For each disk that you created in Step 1, initialize the disk to be managed by NLVM.

   You can use NSSMU or NLVM commands to initialize the disk. The device format can be either GPT or DOS.

   ◆ To use NSSMU:

      1. Launch NSSMU, then select **Devices**.

      2. Select the new disk, then press F3 to initialize it.

   ◆ To use NLVM commands:

      1. On the command line, enter

         NLVM init <*device_name*> [*format*]

3 You might need to specify which disk to use for each NSS pool's snapshots. Create a platespin.conf file on the OES source workload, and associate the NSS pools with the new disks:

   3a In a text editor, create a file at /etc/platespin/platespin.conf.

   3b For each NSS pool, add device and size information under the Customlocation parameter using the following syntax:

```
[Customlocation]
/dev/pool/<yourPoolName>=<device>:<maxUnpartitionSize-in-MB>
```

For example, specify the following entry for a pool named `NSSPOOL` to add snapshots on device `sdc` with a maximum size of 12228 MB.

```
[Customlocation]
/dev/pool/NSSPOOL=sdc:12288
```

4 Save the file.

5 Continue with setting up workload protection for the source OES workload.

12.3 Using Freeze and Thaw Scripts for Every Replication (Linux)

For Linux systems, PlateSpin Protect provides you with the capability to automatically execute custom scripts, `freeze` and `thaw`, that complement the automatic daemon control feature.

The `freeze` script is executed at the beginning of a replication, and `thaw` is executed at the end of a replication.

Consider using this capability to complement the automated daemon control feature provided through the user interface (see . For example, you might want to use this feature to temporarily freeze certain daemons instead of shutting them down during replications.

To implement the feature, use the following procedure before setting up your Linux workload protection:

1 Create the following files:

- ◆ `platespin.freeze.sh`: A shell script to execute at the beginning of the replication
- ◆ `platespin.thaw.sh`: A shell script to execute at the end of the replication
- ◆ `platespin.conf`: A text file defining any required arguments, along with a timeout value.

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

   ```
   [ServiceControl]

   FreezeArguments=<arguments>

   ThawArguments=<arguments>

   TimeOut=<timeout>
   ```

   Replace *<arguments>* with the required command arguments, separated by a space, and *<timeout>* with a timeout value in seconds. If a value is not specified, the default timeout is used (60 seconds).

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

`/etc/platespin`

13 Preparing for Windows Clusters Protection

PlateSpin Protect supports the protection of a Microsoft Windows cluster's business services. The supported Microsoft Windows cluster operating systems are:

- Windows Server 2016
- Windows Server 2012 R2
- Windows Server 2008 R2
- Windows Server 2003 R2

For more information, see "Clusters" in Section 1.1.1, "Supported Windows Workloads," on page 14.

NOTE: The Windows cluster management software provides the failover and failback control for the resources running on its cluster nodes. This document refers to this action as a *cluster node failover* or a *cluster node failback*.

The PlateSpin Server provides the failover and failback control for the protected workload that represents the cluster. This document refers to this action as a *Platespin failover* or a *PlateSpin failback*.

- Section 13.1, "Planning Your Cluster Workload Protection," on page 111
- Section 13.2, "Configuring Windows Active Node Discovery," on page 116
- Section 13.3, "Configuring the Block-Based Transfer Method for Clusters," on page 117
- Section 13.4, "Adding Resource Name Search Values," on page 117
- Section 13.5, "Quorum Arbitration Timeout," on page 118
- Section 13.6, "Setting Local Volume Serial Numbers," on page 118
- Section 13.7, "PlateSpin Failover," on page 118
- Section 13.8, "PlateSpin Failback," on page 119

13.1 Planning Your Cluster Workload Protection

When active node discovery is enabled (the default) for the PlateSpin environment, protection of a Windows cluster is achieved through incremental replications of changes on the active node streamed to a virtual one node cluster, which you can use while troubleshooting the source infrastructure. If you disable active node discovery, each node of a Windows cluster can be discovered and protected as a standalone node.

Before you configure Windows clusters for protection, ensure that your environment meets the prerequisites and that you understand the conditions for protecting cluster workloads.

- Section 13.1.1, "Requirements for Cluster Protection," on page 112
- Section 13.1.2, "Block-Based Transfer for Clusters," on page 113
- Section 13.1.3, "Impact of Cluster Node Failover on Replication," on page 114

13.1.1 Requirements for Cluster Protection

The scope of support for cluster protection is subject to the conditions described in Table 13-1. Consider these requirements when you configure protection for clusters in your PlateSpin environment.

Table 13-1   Cluster Protection Requirements

Requirement	Description
Discover the active node as a Windows Cluster	The PlateSpin global configuration setting `DiscoverActiveNodeAsWindowsCluster` determines whether Windows clusters are protected as clusters or as separate standalone machines: ◆ **True (Default):** The active node is discovered as a Windows cluster. ◆ **False:** Individual nodes can be discovered as standalone machines. See Section 13.2, "Configuring Windows Active Node Discovery," on page 116.
Resource name search values	The PlateSpin global configuration setting `MicrosoftClusterIPAddressNames` determines the cluster resource names that can be discovered in your PlateSpin environment. You must configure search values that help to differentiate the name of the shared Cluster IP Address resource from the name of other IP address resources on the cluster. See Section 13.4, "Adding Resource Name Search Values," on page 117.
Windows Cluster Mode	The PlateSpin global configuration setting `WindowsClusterMode` determines the method of block-based data transfer for incremental replications: ◆ **Default:** Driverless synchronization. ◆ **SingleNodeBBT:** Driver-based block-based transfer. See the following: ◆ "Block-Based Transfer for Clusters" on page 113 ◆ "Configuring the Block-Based Transfer Method for Clusters" on page 117
Active node hostname or IP address	You must specify the host name or IP address of the cluster's active node when you perform an **Add Workload** operation. Because of security changes made by Microsoft, Windows clusters can no longer be discovered by using the virtual cluster name (that is, the shared cluster IP address).
Resolvable host name	The PlateSpin Server must be able to resolve the host name of each of the nodes in the cluster by their IP addresses. **NOTE:** DNS forward lookup and reverse lookup are required to resolve the host name by its IP address.
Quorum resource	A cluster's quorum resource must be co-located on the node with the cluster's resource group (service) being protected.

Requirement	Description
Similarity of cluster nodes	In the default Windows Cluster Mode, driverless sync can continue from any node that becomes active if the nodes are similar. If they do not match, replications can occur only on the originally discovered active node. See "Cluster Node Similarity" on page 116.
PowerShell 2.0	Windows PowerShell 2.0 must be installed on each node of the cluster.

13.1.2 Block-Based Transfer for Clusters

Block-based transfer for clusters works differently than for standalone servers. The initial replication either makes a complete copy (full) or uses a driverless synchronization method performed on the active node of the cluster. Subsequent incremental replications can use a driverless method or driver-based method for block-based data transfer.

NOTE: Protect does not support file-based transfer for clusters.

The PlateSpin global configuration setting `WindowsClusterMode` determines the method of block-based data transfer for incremental replications:

- **Default:** Driverless synchronization.
- **SingleNodeBBT:** Driver-based block-based transfer. Use only with Fibre Channel SANs.

    WARNING: Do not attempt to use SingleNodeBBT on clusters with shared iSCSI drives. It renders the cluster unusable.

Table 13-2 describes and compares the two methods.

Table 13-2 *Comparison Block-Based Data Transfer Methods for Incremental Replication*

Consideration	Default BBT	Single-Node BBT
Data transfer method	Uses driverless synchronization with an MD5-based replication on the currently active node.	Uses a BBT driver installed on the originally discovered active node.
Performance	Potentially slow incremental replications.	Significantly improves performance for incremental replications.

Consideration	Default BBT	Single-Node BBT
Drivers	◆ No BBT driver to install. ◆ No reboot is required on the source cluster nodes.	◆ Use the Protect Agent utility to install a BBT driver on the originally discovered active node of the cluster. ◆ Reboot the node to apply the driver. This initiates a failover to another node in the cluster. After the reboot, make the originally discovered node the active node again. ◆ The same node must remain active for replications to occur and to use single-node block-based transfer. ◆ After you install the BBT driver, either a full replication or a driverless incremental replication must occur before the driver-based incremental replications can begin.
Supported Windows Clusters	Works with any supported Windows Server clusters.	Works with Windows Server 2008 R2 and later clusters. Other supported Windows clusters use the driverless synchronization method for replication.
First incremental replication	Uses driverless sync on the active node.	Uses driver-based block-based transfer on the originally discovered active node if a full replication was completed after the BBT driver was installed. Otherwise, it uses driverless sync on the originally discovered active node.
Subsequent incremental replication	Uses driverless sync on the active node.	Uses driver-based block-based transfer on the originally discovered active node. If a cluster switches nodes, the driverless sync method is used for the first incremental replication after the originally active node becomes active again. See "Impact of Cluster Node Failover on Replication" on page 114.

13.1.3 Impact of Cluster Node Failover on Replication

Table 13-3 describes the impact of cluster node failover on replication and the required actions for the Protect administrator.

Table 13-3 *Impact of Cluster Node Failover on Replication*

Cluster Node Failover or Failback	Default BBT	Single-Node BBT
Cluster node failover occurs during the first full replication	Replication fails. The first full replication must complete successfully without a cluster node failover. 1. Remove the cluster from Protect. 2. (Optional) Make the originally discovered active node the active node again. 3. Re-add the cluster using the active node. 4. Re-run the first full replication.	
Cluster node failover occurs during a subsequent full replication or a subsequent incremental replication	The replication command aborts and a message displays indicating that the replication needs to be re-run. If the new active node's profile is similar to the failed active node, the protection contract remains valid. 1. Re-run the replication on the now-active node. If the new active node's profile is not similar to the failed active node, the protection contract is valid only on the originally active node. 1. Make the originally discovered active node the active node again. 2. Re-run the replication on the active node.	The replication command aborts and a message displays indicating that the replication needs to be re-run. The protection contract is valid only on the originally discovered active node. 1. Make the originally discovered active node the active node again. 2. Re-run the replication on the active node. This first incremental replication after a cluster failover/failback event automatically uses driverless sync. Subsequent incremental replications will use the block-based driver as specified by single-node BBT.
Cluster node failover occurs between replications	If the new active node's profile is similar to the failed active node, the protection contract continues as scheduled for the next incremental replication. Otherwise, the next incremental replication command fails. If a scheduled incremental replication fails: 1. Make the originally discovered active node the active node again. 2. Run an incremental replication.	Incremental replication fails if the active node switches between replications. 1. Ensure that the originally discovered active node is again the active node. 2. Run an incremental replication. This first incremental replication after a cluster failover/failback event automatically uses driverless sync. Subsequent incremental replications will use the block-based driver as specified by single-node BBT.

13.1.4 Cluster Node Similarity

In the default Windows Cluster Mode, the cluster nodes must have similar profiles to prevent interruptions in the replication process. The profiles of cluster nodes are considered similar if all of the following conditions are met:

- Serial numbers for the nodes' local volumes (System volume and System Reserved volume) must be the same on each cluster node.

  NOTE: Use the customized *Volume Manager* utility to change the local volume serial numbers to match each node of the cluster. See "Synchronizing Serial Numbers on Cluster Node Local Storage" on page 143.

  If the local volumes on each node of the cluster have different serial numbers, you cannot run a replication after a cluster node failover occurs. For example, during a cluster node failover, the active node Node 1 fails, and the cluster software makes Node 2 the active node. If the local drives on the two nodes have different serial numbers, the next replication command for the workload fails.

- The nodes must have the same number of volumes.
- Each volume must be exactly the same size on each node.
- The nodes must have an identical number of network connections.

13.1.5 Protection Setup

To configure protection for a Windows cluster, follow the normal workload protection workflow. Ensure that you provide the host name or IP address of the cluster's active node. See "Basic Workflow for Workload Protection and Recovery" on page 35.

13.2 Configuring Windows Active Node Discovery

You can discover Windows Server clusters as clusters or as individual standalone machines, depending on the PlateSpin global configuration setting `DiscoverActiveNodeAsWindowsCluster`.

To discover Windows clusters as clusters, set the `DiscoverActiveNodeAsWindowsCluster` parameter to `True`. This is the default setting. Cluster discovery, inventory, and workload protection use the host name or IP address of a cluster's active node, instead of using its cluster name and an administration share. You do not configure separate workloads for the cluster's non-active nodes. For other cluster workload protection requirements, see "Requirements for Cluster Protection" on page 112.

To discover all Windows clusters as individual standalone machines, set the `DiscoverActiveNodeAsWindowsCluster` parameter to `False`. This setting allows the PlateSpin Server to discover all nodes in a Windows failover cluster as standalone machines. That is, it inventories a cluster's active node and non-active nodes as a regular, cluster-unaware Windows workloads.

To enable or disable cluster discovery:

1 Go to the PlateSpin Server configuration page at

   https://<p*latespin-server-ip-address*>/PlateSpinConfiguration

2 Search for `DiscoverActiveNodeAsWindowsCluster`, then click **Edit**.

3 In the **Value** field, select **True** to enable cluster discovery, or select **False** to disable cluster discovery.

4 Click **Save**.

13.3 Configuring the Block-Based Transfer Method for Clusters

Incremental replications for Windows clusters can use a driverless method (Default) or driver-based method (SingleNodeBBT) for block-based data transfer, depending on the PlateSpin global configuration setting `WindowsClusterMode`. For more information, see "Block-Based Transfer for Clusters" on page 113.

To configure WindowsClusterMode:

1 Go to the PlateSpin Server configuration page at

https://<p*latespin-server-ip-address*>/PlateSpinConfiguration

2 Search for `WindowsClusterMode`, then click **Edit**.

3 In the **Value** field, select **Default** to use driverless synchronization for incremental replication, or select **SingleNodeBBT** to use block-based drivers for incremental replication.

4 Click **Save**.

13.4 Adding Resource Name Search Values

To help identify the active node in a Windows failover cluster, PlateSpin Protect must differentiate the name of the shared Cluster IP Address resource from the names of other IP address resources on the cluster. The shared Cluster IP Address resource resides on the cluster's active node.

The global parameter `MicrosoftClusterIPAddressNames` on the PlateSpin Server Configuration page contains a list of search values to use in discovery for a Windows cluster workload. When you add a Windows cluster workload, you must specify the IP address of the cluster's currently active node. PlateSpin Protect searches the names of the cluster's IP address resources on that node to find one that *starts with* the specified characters of any value in the list. Thus, each search value must contain enough characters to differentiate the shared Cluster IP Address resource on a specific cluster, but it can be short enough to apply to discovery in other Windows clusters.

For example, a search value of `Clust IP Address` or `Clust IP` matches the resource names *Clust IP Address* for 10.10.10.201 and *Clust IP Address* for 10.10.10.101.

The default name for the shared Cluster IP Address resource is `Cluster IP Address` in English, or the equivalent if the cluster node is configured in another language. The default search values in the `MicrosoftClusterIPAddressNames` list include the resource name `Cluster IP Address` in English and each of the supported languages.

Because the resource name of the shared Cluster IP Address resource is user-configurable, you must add other search values to the list, as needed. If you change the resource name, you must add a related search value to the `MicrosoftClusterIPAddressNames` list. For example, if you specify a resource name of `Win2012-CLUS10-IP-ADDRESS`, you should add that value to the list. If you have multiple clusters using the same naming convention, an entry of `Win2012-CLUS` matches any resource name that starts with that sequence of characters.

To add search values in the `MicrosoftClusterIPAddressNames` list:

1 Go to the PlateSpin Server configuration page at

https://<p*latespin-server-ip-address*>/PlateSpinConfiguration

  2 Search for `MicrosoftClusterIPAddressNames`, then click **Edit**.

  3 In the **Value** field, add one or more search values to the list.

  4 Click **Save**.

13.5 Quorum Arbitration Timeout

You can set the QuorumArbitrationTimeMax registry key for Windows Server failover clusters in your PlateSpin environment by using the global parameter `FailoverQuorumArbitrationTimeout` on the PlateSpin Server Configuration page. The default timeout is 60 seconds, in keeping with the Microsoft default value for this setting. See *QuorumArbitrationTimeMax* (https://msdn.microsoft.com/en-us/library/aa369123%28v=vs.85%29.aspx?f=255&MSPPError=-2147217396) on the Microsoft Developer Network website. The specified timeout interval is honored for quorum arbitration at failover and failback.

To set the quorum arbitration timeout for all Windows failover clusters:

  1 Go to the PlateSpin Server configuration page at

     https://<p*latespin-server-ip-address*>/PlatespinConfiguration

  2 Search for `FailoverQuorumArbitrationTimeout`, then click **Edit**.

  3 In the **Value** field, specify the maximum number of seconds to allow for quorum arbitration.

  4 Click **Save**.

13.6 Setting Local Volume Serial Numbers

You can use the *Volume Manager* utility to change the local volume serial numbers to match in each node of the cluster. See "Synchronizing Serial Numbers on Cluster Node Local Storage" on page 143.

13.7 PlateSpin Failover

When the PlateSpin failover operation is complete and the virtual one-node cluster comes online, you see a multi-node cluster with one active node (all other nodes are unavailable).

To perform a PlateSpin failover (or to test the PlateSpin failover on) a Windows cluster, the cluster must be able to connect to a domain controller. To leverage the test failover functionality, you need to protect the domain controller along with the cluster. During the test, bring up the domain controller, followed by the Windows cluster workload (on an isolated network).

13.8 PlateSpin Failback

A PlateSpin failback operation requires a full replication for Windows Cluster workloads.

If you configure the PlateSpin failback as a full replication to a physical target, you can use one of these methods:

- Map all disks on the PlateSpin virtual one-node cluster to a single local disk on the failback target.
- Add another disk (`Disk 2`) to the physical failback machine. You can then configure the PlateSpin failback operation to restore the failover machine's system volume to `Disk 1` and the failover machine's additional disks (previous shared disks) to `Disk 2`. This allows the system disk to be restored to the same size storage disk as the original source.

After a PlateSpin failback is complete, you must reattach the shared storage and rebuild the cluster environment before you can rejoin additional nodes to the newly restored cluster.

NOTE: When the cluster is at the stage of **Ready To Reprotect**, ensure that you first rebuild and restore the failback target so that it gets discovered as a cluster. You must manually uninstall the PlateSpin Cluster Driver as part of the rebuild process.

For information about rebuilding the cluster environment after a PlateSpin failover and failback occurs, see the following resources:

- **Windows Server 2012 R2 Failover Cluster (failback to physical or virtual rebuild):** See Knowledgebase Article 7016770 (http://www.netiq.com/support/kb/doc.php?id=7016770).
- **Windows Server 2008 R2 Failover Cluster (failback to physical or virtual rebuild):** See Knowledgebase Article 7015576 (http://www.netiq.com/support/kb/doc.php?id=7015576).

---

14 Troubleshooting Workload Discovery and Inventory

This section can help you troubleshoot common problems during the workload discovery and inventory.

-
-
-

14.1 Troubleshooting Discovery for Windows Workloads

Use the information in this section to troubleshoot and resolve issues during the workload inventory and discovery of Windows workloads:

-
-
-
-
-

14.1.1 Common Problems and Solutions

Problems or Messages	Solutions
`The domain in the credentials is invalid or blank`	This error occurs when the Credential Format is incorrect. Try the discovery by using a local administrator account with the credential format `hostname\LocalAdmin`. Or, try the discovery by using a domain administrator account with the credential format `domain\DomainAdmin`.
`Unable to connect to Windows server...Access is denied`	A non-administrator account was used when trying to add a workload. Use an administrator account or add the user to the Administrators group and try again. This message might also indicate WMI connectivity failure. For each of the following possible resolutions, attempt the solution and then perform the "WMI Connectivity Test" on page 123 again. If the test succeeds, try adding the workload again. • "Troubleshooting DCOM Connectivity" on page 123 • "Troubleshooting RPC Service Connectivity" on page 124

Problems or Messages	Solutions
`Unable to connect to Windows server...The network path was not found`	Network connectivity failure. Perform the tests in "Performing Connectivity Tests" on page 123. If a test fails, ensure that PlateSpin Protect and the workload are on the same network. Reconfigure the network and try again.
`Discover Server Details {hostname}" Failed Progress: 0% Status: NotStarted`	This error can occur for several reasons and each has a unique solution: ◆ For environments using a local proxy with authentication, bypass the proxy or add the proper permissions. See Knowledgebase Article 7920339 (https://www.netiq.com/support/kb/doc.php?id=7920339) for more details. ◆ If local or domain policies restrict required permissions, follow the steps outlined in Knowledgebase Article 7920862 (https://www.netiq.com/support/kb/doc.php?id=7920862).

Workload Discovery fails with error message
 	
`Could not find file output.xml`
 	
or
 	
`Network path not found`
 	
or (upon attempting to discover a Windows cluster)
 	
`Inventory failed to discover. Inventory result returned nothing.`
	There are several possible reasons for the `Could not find file output.xml` error:
	
	◆ Antivirus software on the source could be interfering with the discovery. Disable the antivirus software to determine whether or not it is the cause of the problem. See "Disabling Antivirus Software" on page 124.
	
	◆ File and Printer Sharing for Microsoft Networks might not be enabled. Enable it under the Network Interface Card properties.
	
	◆ The Admin$ shares on the source might not be accessible. Ensure that Protect can access those shares. See "Enabling File/Share Permissions and Access" on page 125.
	
	◆ The Server or the Workstation service might not be running. If this is the case, enable them and set the startup mode to `automatic`.
	
	◆ The Windows remote registry service is disabled. Start the service and set the startup type to automatic.

14.1.2 Modifying the OFX Controller Heartbeat Startup Delay

To avoid discovery failures caused by timing issues, a default heartbeat startup delay of 15 seconds (15000 ms) is set on the OFX Controller. The setting is configurable by adding the `HeartbeatStartupDelayInMS` registry key on the source workload. This registry key is not configured by default.

To enable a heartbeat delay of shorter or longer duration:

1 On the source workload, open the Windows Registry Editor.

2 Go to the following location in the Registry Editor, depending on the operating system architecture on the source workload:

Path for a 64-bit source workload:

`HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\PlateSpin\OperationsFramework\Controller`

Path for a 32-bit source workload:

`HKEY_LOCAL_MACHINE\SOFTWARE\PlateSpin\OperationsFramework\Controller`

3 Add a key named `HeartbeatStartupDelayInMS` of type `REG_SZ` and set its value to the desired value in milliseconds. The default setting should be 15000.

```
REG_SZ: HeartbeatStartupDelayInMS
Value: "15000"
```

4 Restart the source workload.

14.1.3 Performing Connectivity Tests

Network Connectivity Test

Perform this basic network connectivity test to determine whether Protect can communicate with the workload that you are trying to protect.

1 Go to your PlateSpin Server host.

2 Open a command prompt and ping your workload:

`ping` *workload_ip*

WMI Connectivity Test

1 Go to your PlateSpin Server host.

2 Click **Start > Run**, type `Wbemtest` and press `Enter`.

3 Click **Connect**.

4 In the **Namespace**, type the name of the workload you are trying to discover with `\root\cimv2` appended to it. For example, if the host name is `win2k`, type:

`\\win2k\root\cimv2`

5 Enter the appropriate credentials, using either the `hostname\LocalAdmin` or `domain\DomainAdmin` format.

6 Click **Connect** to test the WMI connection.

If an error message is returned, a WMI connection cannot be established between Protect and your workload.

Troubleshooting DCOM Connectivity

1 Log into the workload that you want to protect.

2 Click **Start > Run**.

3 Type `dcomcnfg` and press Enter.

4 Check connectivity:

- ◆ For Windows systems (XP/Vista/2003/2008/7), the Component Services window is displayed. In the **Computers** folder of the console tree of the Component Services administrative tool, right-click the computer that you want to check for DCOM connectivity, then click **Properties**. Click the **Default Properties** tab and ensure that **Enable Distributed COM on this computer** is selected.

- ◆ On a Windows 2000 Server machine, the DCOM Configuration dialog is displayed. Click the **Default Properties** tab and ensure that **Enable Distributed COM on this computer** is selected.

5 If DCOM was not enabled, enable it and either reboot the server or restart the Windows Management Instrumentation Service. Then try adding the workload again.

Troubleshooting RPC Service Connectivity

There are three potential blockages for the RPC service:

- ◆ The Windows Service
- ◆ A Windows firewall
- ◆ A network firewall

For the Windows Service, ensure that the RPC service is running on the workload. To access the services panel, run `services.msc` from a command prompt. For a Windows firewall, add an RPC exception. For hardware firewalls, you can try the following strategies:

- ◆ Putting Protect and the workload on the same side of the firewall
- ◆ Opening up specific ports between Protect and the workload (See "Access and Communication Requirements across Your Protection Network" on page 29.

14.1.4   Disabling Antivirus Software

Antivirus software might occasionally block some of the Protect functionality related to WMI and Remote Registry. In order to ensure that workload inventory is successful, it might be necessary to first disable the antivirus service on a workload.

In addition, antivirus software might occasionally lock access to certain files, allowing access only to certain processes or executables. This lock might occasionally obstruct file-based data replication. In this case, when you configure the workload protection, you can select services to disable, such as services installed and used by antivirus software. These services are disabled only for the duration of the file transfer, and are restarted when the process completes. Disabling services is not necessary during block-level data replication.

14.1.5 Enabling File/Share Permissions and Access

To successfully protect a workload, PlateSpin Protect needs to successfully deploy and install software within the workload. Upon deployment of these components to a workload, as well as during the Add Workload process, Protect uses the workload's administrative shares. Protect needs administrative access to the shares, using either a local administrator account or a domain administrator account for this to work.

To ensure that the Administrative shares are enabled:

1 Right-click **My Computer** on the desktop and select **Manage**.

2 Expand **System Tools > Shared Folders > Shares**

3 In the `Shared Folders` directory, you should see Admin$, among other shares.

After confirming that the shares are enabled, ensure that they are accessible from within the PlateSpin Server host:

1 Go to your PlateSpin Server host.

2 Click **Start > Run**, type `\\<server_host>\Admin$`, then click **OK**.

3 If you are prompted, use the same credentials as those you will use to add the workload to the Protect workload inventory.

   The directory is opened and you should be able to browse and modify its contents.

4 Repeat the process for all shares with the exception of the IPC$ share.

   Windows uses the IPC$ share for credential validation and authentication purposes. It is not mapped to a folder or file on the workload, so the test always fails; however, the share should still be visible.

PlateSpin Protect does not modify the existing content of the volume; however, it creates its own directory, to which it requires access and permissions.

14.2 Troubleshooting Discovery for Linux Workloads

Problems or Messages	Solutions
`Unable to connect neither to the SSH server running on <IP_address> nor to VMware Virtual Infrastructure web-services at <ip_address>/sdk`	This message has a number of possible causes: ◆ The workload is unreachable. ◆ The workload does not have SSH running. ◆ The firewall is on and the required ports have not been opened. ◆ The workload's specific operating system is not supported. For network and access requirements for a workload, see "Access and Communication Requirements across Your Protection Network" on page 29.
`Access denied`	This authentication problem indicates either an invalid user name or password. For information on proper workload access credentials, see "Guidelines for Workload and Container Credentials" on page 165.

14.3   Troubleshooting Discovery for Target Hosts

Problems or Messages	Solutions
For ESXi 4.1, Direct host discovery results in missing VM port groups if dvSwitch port groups share the same name.	Ensure that port group names are unique on the target VMware host.

B B Linux Distributions Supported by Protect

PlateSpin Protect software includes pre-compiled versions of the `blkwatch` driver for many non-debug Linux distributions (32-bit and 64-bit).

- Section B.1, "Analyzing Your Linux Workload," on page 127
- Section B.2, "Pre-compiled blkwatch Drivers for Linux Distributions," on page 128

B.1 Analyzing Your Linux Workload

Prior to determining whether PlateSpin Protect has a `blkwatch` driver for your Linux distribution, you need to learn more about the kernel of your Linux workload so that you can use it as a search term against the list of supported distributions.

- Section B.1.1, "Determining the Release String," on page 127
- Section B.1.2, "Determining the Architecture," on page 127

B.1.1 Determining the Release String

You can determine the release string of the kernel of your Linux workload by running the following command at the workload's Linux terminal:

```
uname -r
```

For example, if you run `uname -r`, you might see the following output:

```
3.0.76-0.11-default
```

If you search the list of distributions, you see there are two entries that match this string:

- `SLES11SP3-GA-3.0.76-0.11-default-x86`
- `SLES11SP3-GA-3.0.76-0.11-default-x86_64`

The search results indicate that the product has drivers for both 32-bit (x86) and 64-bit (x86_64) architectures.

B.1.2 Determining the Architecture

You can determine the architecture of your Linux workload by running the following command at the workload's Linux terminal:

```
uname -m
```

For example, if you run `uname -m`, you might see the following output:

```
x86_64
```

With this information, you can determine that the workload has 64-bit architecture.

B.2 Pre-compiled blkwatch Drivers for Linux Distributions

PlateSpin Protect provides precompiled blkwatch drivers for many non-debug Linux distributions. You can search the List of Distributions to determine if the release string and architecture of your Linux workload kernel matches a supported distribution in the list. If you find your release string and architecture, PlateSpin Protect has a pre-compiled version of the `blkwatch` driver.

If your search is unsuccessful, you can create a custom `blkwatch` driver by following the steps found in the Knowledgebase Article 7005873 (https://www.netiq.com/support/kb/doc.php?id=7005873). Self-compiled drivers are supported only for the Linux major and minor kernel versions that appear in the List of Distributions, or a patched version thereof. If the major and minor kernel version in the release string of your Linux workload kernel matches a major and minor kernel version in the list, your self-compiled driver will be supported.

- Section B.2.1, "List Item Syntax," on page 128
- Section B.2.2, "List of Distributions," on page 128
- Section B.2.3, "Other Linux Distributions That Use blkwatch Drivers," on page 140

B.2.1 List Item Syntax

Each item in the list is formatted using the following syntax:

<Distro>-<Patch>-<Kernel_Release_String>-<Kernel_Architecture>

So, for a SLES 9 SP1 distribution with a kernel release string of `2.6.5-7.139-bigsmp` for 32-bit (x86) architecture, the item is listed in a format like this:

`SLES9-SP1-2.6.5-7.139-bigsmp-x86`

B.2.2 List of Distributions

The following distributions have a pre-compiled blkwatch driver. See also Section B.2.3, "Other Linux Distributions That Use blkwatch Drivers," on page 140.

Oracle Linux 6 U7

   NOTE: Blkwatch drivers for kernel version 2.6.32-573 do not support incremental replication for workloads with LVM volumes. Update the kernel, then use RHEL 6 U7 drivers for kernel 2.6.32-642.

   `OEL6-U7-2.6.32-573.el6.i686-x86`

   `OEL6-U7-2.6.32-573.el6.x86_64-x86_64`

   `OEL6-U7_UEK-2.6.39-400.250.7.el6uek.i686-x86`

   `OEL6-U7_UEK-3.8.13-68.3.4.el6uek.x86_64-x86_64`

Oracle Linux 6 U8

   `OEL6-U8-2.6.32-642.el6.i686-x86`

   `OEL6-U8-2.6.32-642.el6.x86_64-x86_64`

   `OEL6-U8_UEK-2.6.39-400.278.2.el6uek.i686-x86`

   `OEL6-U8_UEK-4.1.12-37.4.1.el6uek.x86_64-x86_64`

Oracle Linux 6 U9

   `OEL6-U9_UEK-2.6.39-400.294.3.el6uek.i686-x86`

   `OEL6-U9_UEK-4.1.12-61.1.28.el6uek.x86_64-x86_64`

Oracle Linux 7 GA

```
OEL7-GA-3.10.0-123.el7.x86_64-x86_64
OEL7-GA_UEK-3.8.13-35.3.1.el7uek.x86_64-x86_64
```

Oracle Linux 7 U1

```
OEL7-U1-3.10.0-229.el7.x86_64-x86_64
OEL7-U1_UEK-3.8.13-55.1.6.el7uek.x86_64-x86_64
```

Oracle Linux 7 U2

```
OEL7-U2-3.10.0-327.el7.x86_64-x86_64
OEL7-U2_UEK-3.8.13-98.7.1.el7uek.x86_64-x86_64
```

Oracle Linux 7 U3

```
OEL7-U3_UEK-4.1.12-61.1.18.el7uek.x86_64-x86_64
```

Red Hat Enterprise Linux 4 GA

```
RHEL4-GA-2.6.9-5.EL-x86
RHEL4-GA-2.6.9-5.EL-x86_64
RHEL4-GA-2.6.9-5.ELhugemem-x86
RHEL4-GA-2.6.9-5.ELsmp-x86
RHEL4-GA-2.6.9-5.ELsmp-x86_64
```

Red Hat Enterprise Linux 4 U1

```
RHEL4-U1-2.6.9-11.EL-x86
RHEL4-U1-2.6.9-11.EL-x86_64
RHEL4-U1-2.6.9-11.ELhugemem-x86
RHEL4-U1-2.6.9-11.ELsmp-x86
RHEL4-U1-2.6.9-11.ELsmp-x86_64
```

Red Hat Enterprise Linux 4 U2

```
RHEL4-U2-2.6.9-22.EL-x86
RHEL4-U2-2.6.9-22.EL-x86_64
RHEL4-U2-2.6.9-22.ELhugemem-x86
RHEL4-U2-2.6.9-22.ELsmp-x86
RHEL4-U2-2.6.9-22.ELsmp-x86_64
```

Red Hat Enterprise Linux 4 U3

```
RHEL4-U3-2.6.9-34.EL-x86
RHEL4-U3-2.6.9-34.EL-x86_64
RHEL4-U3-2.6.9-34.ELhugemem-x86
RHEL4-U3-2.6.9-34.ELlargesmp-x86_64
RHEL4-U3-2.6.9-34.ELsmp-x86
RHEL4-U3-2.6.9-34.ELsmp-x86_64
```

Red Hat Enterprise Linux 4 U4

```
RHEL4-U4-2.6.9-42.EL-x86
RHEL4-U4-2.6.9-42.EL-x86_64
RHEL4-U4-2.6.9-42.ELhugemem-x86
RHEL4-U4-2.6.9-42.ELlargesmp-x86_64
RHEL4-U4-2.6.9-42.ELsmp-x86
RHEL4-U4-2.6.9-42.ELsmp-x86_64
```

Red Hat Enterprise Linux 4 U5

```
RHEL4-U5-2.6.9-55.EL-x86
RHEL4-U5-2.6.9-55.EL-x86_64
```

```
RHEL4-U5-2.6.9-55.ELhugemem-x86
RHEL4-U5-2.6.9-55.ELlargesmp-x86_64
RHEL4-U5-2.6.9-55.ELsmp-x86
RHEL4-U5-2.6.9-55.ELsmp-x86_64
```

Red Hat Enterprise Linux 4 U6

```
RHEL4-U6-2.6.9-67.EL-x86
RHEL4-U6-2.6.9-67.EL-x86_64
RHEL4-U6-2.6.9-67.ELhugemem-x86
RHEL4-U6-2.6.9-67.ELlargesmp-x86_64
RHEL4-U6-2.6.9-67.ELsmp-x86
RHEL4-U6-2.6.9-67.ELsmp-x86_64
```

Red Hat Enterprise Linux 4 U7

```
RHEL4-U7-2.6.9-78.EL-x86
RHEL4-U7-2.6.9-78.EL-x86_64
RHEL4-U7-2.6.9-78.ELhugemem-x86
RHEL4-U7-2.6.9-78.ELlargesmp-x86_64
RHEL4-U7-2.6.9-78.ELsmp-x86
RHEL4-U7-2.6.9-78.ELsmp-x86_64
```

Red Hat Enterprise Linux 4 U8

```
RHEL4-U8-2.6.9-89.EL-x86
RHEL4-U8-2.6.9-89.EL-x86_64
RHEL4-U8-2.6.9-89.ELhugemem-x86
RHEL4-U8-2.6.9-89.ELlargesmp-x86_64
RHEL4-U8-2.6.9-89.ELsmp-x86
RHEL4-U8-2.6.9-89.ELsmp-x86_64
```

Red Hat Enterprise Linux 4 U9

```
RHEL4-U9-2.6.9-100.EL-x86
RHEL4-U9-2.6.9-100.EL-x86_64
RHEL4-U9-2.6.9-100.ELhugemem-x86
RHEL4-U9-2.6.9-100.ELlargesmp-x86_64
RHEL4-U9-2.6.9-100.ELsmp-x86
RHEL4-U9-2.6.9-100.ELsmp-x86_64
```

Red Hat Enterprise Linux 5 GA

```
RHEL5-GA-2.6.18-8.el5-x86
RHEL5-GA-2.6.18-8.el5-x86_64
RHEL5-GA-2.6.18-8.el5PAE-x86
```

Red Hat Enterprise Linux 5 U1

```
RHEL5-U1-2.6.18-53.el5-x86
RHEL5-U1-2.6.18-53.el5-x86_64
RHEL5-U1-2.6.18-53.el5PAE-x86
```

Red Hat Enterprise Linux 5 U2

```
RHEL5-U2-2.6.18-92.el5-x86
RHEL5-U2-2.6.18-92.el5-x86_64
RHEL5-U2-2.6.18-92.el5PAE-x86
```

Red Hat Enterprise Linux 5 U3

```
RHEL5-U3-2.6.18-128.el5-x86
RHEL5-U3-2.6.18-128.el5-x86_64
```

```
RHEL5-U3-2.6.18-128.el5PAE-x86
```
Red Hat Enterprise Linux 5 U4
```
RHEL5-U4-2.6.18-164.el5-x86
RHEL5-U4-2.6.18-164.el5-x86_64
RHEL5-U4-2.6.18-164.el5PAE-x86
```
Red Hat Enterprise Linux 5 U5
```
RHEL5-U5-2.6.18-194.el5-x86
RHEL5-U5-2.6.18-194.el5-x86_64
RHEL5-U5-2.6.18-194.el5PAE-x86
```
Red Hat Enterprise Linux 5 U6
```
RHEL5-U6-2.6.18-238.el5-x86
RHEL5-U6-2.6.18-238.el5-x86_64
RHEL5-U6-2.6.18-238.el5PAE-x86
```
Red Hat Enterprise Linux 5 U7
```
RHEL5-U7-2.6.18-274.el5-x86
RHEL5-U7-2.6.18-274.el5-x86_64
RHEL5-U7-2.6.18-274.el5PAE-x86
```
Red Hat Enterprise Linux 5 U8
```
RHEL5-U8-2.6.18-308.el5-x86
RHEL5-U8-2.6.18-308.el5-x86_64
RHEL5-U8-2.6.18-308.el5PAE-x86
```
Red Hat Enterprise Linux 5 U9
```
RHEL5-U9-2.6.18-348.el5-x86
RHEL5-U9-2.6.18-348.el5-x86_64
RHEL5-U9-2.6.18-348.el5PAE-x86
```
Red Hat Enterprise Linux 5 U10
```
RHEL5-U10-2.6.18-371.el5-x86
RHEL5-U10-2.6.18-371.el5-x86_64
RHEL5-U10-2.6.18-371.el5PAE-x86
```
Red Hat Enterprise Linux 5 U11
```
RHEL5-U11-2.6.18-398.el5-x86
RHEL5-U11-2.6.18-398.el5-x86_64
RHEL5-U11-2.6.18-398.el5PAE-x86
```
Red Hat Enterprise Linux 6 GA
```
RHEL6-GA-2.6.32-71.el6.i686-x86
RHEL6-GA-2.6.32-71.el6.x86_64-x86_64
```
Red Hat Enterprise Linux 6 U1
```
RHEL6-U1-2.6.32-131.0.15.el6.i686-x86
RHEL6-U1-2.6.32-131.0.15.el6.x86_64-x86_64
```
Red Hat Enterprise Linux 6 U2
```
RHEL6-U2-2.6.32-220.el6.i686-x86
RHEL6-U2-2.6.32-220.el6.x86_64-x86_64
```
Red Hat Enterprise Linux 6 U3
```
RHEL6-U3-2.6.32-279.el6.i686-x86
RHEL6-U3-2.6.32-279.el6.x86_64-x86_64
```
Red Hat Enterprise Linux 6 U4
```
RHEL6-U4-2.6.32-358.el6.i686-x86
```

```
RHEL6-U4-2.6.32-358.el6.x86_64-x86_64
```
Red Hat Enterprise Linux 6 U5
```
RHEL6-U5-2.6.32-431.el6.i686-x86
RHEL6-U5-2.6.32-431.el6.x86_64-x86_64
```
Red Hat Enterprise Linux 6 U6
```
RHEL6-U6-2.6.32-504.el6-x86
RHEL6-U6-2.6.32-504.el6-x86_64
```
Red Hat Enterprise Linux 6 U7

NOTE: Blkwatch drivers for kernel version 2.6.32-573 do not support incremental replication for workloads with LVM volumes. Update the kernel, then use drivers for kernel 2.6.32-642.
```
RHEL6-U7-2.6.32-573.el6.i686-x86
RHEL6-U7-2.6.32-573.el6.x86_64-x86_64
RHEL6-RHSA201700361-2.6.32-642.13.1.el6.i686-x86
RHEL6-RHSA201700361-2.6.32-642.13.1.el6.x86_64-x86_64
```
Red Hat Enterprise Linux 6 U8
```
RHEL6-U8-2.6.32-642.el6.i686-x86
RHEL6-U8-2.6.32-642.el6.x86_64-x86_64
```
Red Hat Enterprise Linux 6 U9
```
RHEL6-U9-2.6.32-696.el6.i686-x86
RHEL6-U9-2.6.32-696.el6.x86_64-x86_64
```
Red Hat Enterprise Linux 7 GA
```
RHEL7-GA-3.10.0-123.el7.x86_64-x86_64
```
Red Hat Enterprise Linux 7 U1
```
RHEL7-U1-3.10.0-229.el7.x86_64-x86_64
```
Red Hat Enterprise Linux 7 U2
```
RHEL7-U2-3.10.0-327.el7.x86_64-x86_64
```
Red Hat Enterprise Linux 7 U3
```
RHEL7-U3-3.10.0-514.el7.x86_64-x86_64
```

SUSE Linux Enterprise Server 9 GA
```
SLES9-GA-2.6.5-7.97-bigsmp-x86
SLES9-GA-2.6.5-7.97-default-x86
SLES9-GA-2.6.5-7.97-default-x86_64
SLES9-GA-2.6.5-7.97-smp-x86
SLES9-GA-2.6.5-7.97-smp-x86_64
```
SUSE Linux Enterprise Server 9 SP 1
```
SLES9-SP1-2.6.5-7.139-bigsmp-x86
SLES9-SP1-2.6.5-7.139-default-x86
SLES9-SP1-2.6.5-7.139-default-x86_64
SLES9-SP1-2.6.5-7.139-smp-x86
SLES9-SP1-2.6.5-7.139-smp-x86_64
```
SUSE Linux Enterprise Server 9 SP 2
```
SLES9-SP2-2.6.5-7.191-bigsmp-x86
SLES9-SP2-2.6.5-7.191-default-x86
SLES9-SP2-2.6.5-7.191-default-x86_64
SLES9-SP2-2.6.5-7.191-smp-x86
SLES9-SP2-2.6.5-7.191-smp-x86_64
```

SUSE Linux Enterprise Server 9 SP 3

```
SLES9-SP3-2.6.5-7.244-bigsmp-x86
SLES9-SP3-2.6.5-7.244-default-x86
SLES9-SP3-2.6.5-7.244-default-x86_64
SLES9-SP3-2.6.5-7.244-smp-x86
SLES9-SP3-2.6.5-7.244-smp-x86_64
```

SUSE Linux Enterprise Server 9 SP 4

```
SLES9-SP4-2.6.5-7.308-bigsmp-x86
SLES9-SP4-2.6.5-7.308-default-x86
SLES9-SP4-2.6.5-7.308-default-x86_64
SLES9-SP4-2.6.5-7.308-smp-x86
SLES9-SP4-2.6.5-7.308-smp-x86_64
```

SUSE Linux Enterprise Server 10 GA

```
SLES10-GA-2.6.16.21-0.8-bigsmp-x86
SLES10-GA-2.6.16.21-0.8-default-x86
SLES10-GA-2.6.16.21-0.8-default-x86_64
SLES10-GA-2.6.16.21-0.8-smp-x86
SLES10-GA-2.6.16.21-0.8-smp-x86_64
SLES10-GA-2.6.16.21-0.8-xen-x86
SLES10-GA-2.6.16.21-0.8-xen-x86_64
SLES10-GA-2.6.16.21-0.8-xenpae-x86
```

SUSE Linux Enterprise Server 10 SP 1

```
SLES10-SP1-2.6.16.46-0.12-bigsmp-x86
SLES10-SP1-2.6.16.46-0.12-default-x86
SLES10-SP1-2.6.16.46-0.12-default-x86_64
SLES10-SP1-2.6.16.46-0.12-smp-x86
SLES10-SP1-2.6.16.46-0.12-smp-x86_64
SLES10-SP1-2.6.16.46-0.12-xen-x86
SLES10-SP1-2.6.16.46-0.12-xen-x86_64
SLES10-SP1-2.6.16.46-0.12-xenpae-x86
```

SUSE Linux Enterprise Server 10 SP 2

```
SLES10-SP2-2.6.16.60-0.21-bigsmp-x86
SLES10-SP2-2.6.16.60-0.21-default-x86
SLES10-SP2-2.6.16.60-0.21-default-x86_64
SLES10-SP2-2.6.16.60-0.21-smp-x86
SLES10-SP2-2.6.16.60-0.21-smp-x86_64
SLES10-SP2-2.6.16.60-0.21-xen-x86
SLES10-SP2-2.6.16.60-0.21-xen-x86_64
SLES10-SP2-2.6.16.60-0.21-xenpae-x86
```

SUSE Linux Enterprise Server 10 SP 2 LTSS U2

```
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-bigsmp-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-default-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-default-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-smp-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-smp-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xen-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xen-x86_64
```

```
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xenpae-x86
```
SUSE Linux Enterprise Server 10 SP 3
```
SLES10-SP3-2.6.16.60-0.54.5-bigsmp-x86
SLES10-SP3-2.6.16.60-0.54.5-default-x86
SLES10-SP3-2.6.16.60-0.54.5-default-x86_64
SLES10-SP3-2.6.16.60-0.54.5-smp-x86
SLES10-SP3-2.6.16.60-0.54.5-smp-x86_64
SLES10-SP3-2.6.16.60-0.54.5-xen-x86
SLES10-SP3-2.6.16.60-0.54.5-xen-x86_64
SLES10-SP3-2.6.16.60-0.54.5-xenpae-x86
```
SUSE Linux Enterprise Server 10 SP 3 LTSS U1
```
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-bigsmp-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-default-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-default-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-smp-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-smp-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xen-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xen-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xenpae-x86
```
SUSE Linux Enterprise Server 10 SP 3 LTSS U2
```
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-bigsmp-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-default-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-default-x86_64
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-smp-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-smp-x86_64
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-xen-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-xen-x86_64
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-xenpae-x86
```
SUSE Linux Enterprise Server 10 SP 4
```
SLES10-SP4-2.6.16.60-0.85.1-bigsmp-x86
SLES10-SP4-2.6.16.60-0.85.1-default-x86
SLES10-SP4-2.6.16.60-0.85.1-default-x86_64
SLES10-SP4-2.6.16.60-0.85.1-smp-x86
SLES10-SP4-2.6.16.60-0.85.1-smp-x86_64
SLES10-SP4-2.6.16.60-0.85.1-xen-x86
SLES10-SP4-2.6.16.60-0.85.1-xen-x86_64
SLES10-SP4-2.6.16.60-0.85.1-xenpae-x86
```
SUSE Linux Enterprise Server 10 SP 4 U4
```
SLES10-SP4_U4-2.6.16.60-0.93.1-bigsmp-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-default-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-default-x86_64
SLES10-SP4_U4-2.6.16.60-0.93.1-smp-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-smp-x86_64
SLES10-SP4_U4-2.6.16.60-0.93.1-xen-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-xen-x86_64
SLES10-SP4_U4-2.6.16.60-0.93.1-xenpae-x86
```
SUSE Linux Enterprise Server 10 SP 4 U5

```
SLES10-SP4_U5-2.6.16.60-0.97.1-bigsmp-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-default-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-default-x86_64
SLES10-SP4_U5-2.6.16.60-0.97.1-smp-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-smp-x86_64
SLES10-SP4_U5-2.6.16.60-0.97.1-xen-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-xen-x86_64
SLES10-SP4_U5-2.6.16.60-0.97.1-xenpae-x86
```

SUSE Linux Enterprise Server 10 SP 4 U6

```
SLES10-SP4_U6-2.6.16.60-0.99.1-bigsmp-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-default-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-default-x86_64
SLES10-SP4_U6-2.6.16.60-0.99.1-smp-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-smp-x86_64
SLES10-SP4_U6-2.6.16.60-0.99.1-xen-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-xen-x86_64
SLES10-SP4_U6-2.6.16.60-0.99.1-xenpae-x86
```

SUSE Linux Enterprise Server 10 SP 4 U7

```
SLES10-SP4_U7-2.6.16.60-0.101.1-bigsmp-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-default-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-default-x86_64
SLES10-SP4_U7-2.6.16.60-0.101.1-smp-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-smp-x86_64
SLES10-SP4_U7-2.6.16.60-0.101.1-xen-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-xen-x86_64
SLES10-SP4_U7-2.6.16.60-0.101.1-xenpae-x86
```

SUSE Linux Enterprise Server 10 SP 4 U8

```
SLES10-SP4_U8-2.6.16.60-0.103.1-bigsmp-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-default-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-default-x86_64
SLES10-SP4_U8-2.6.16.60-0.103.1-smp-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-smp-x86_64
SLES10-SP4_U8-2.6.16.60-0.103.1-xen-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-xen-x86_64
SLES10-SP4_U8-2.6.16.60-0.103.1-xenpae-x86
```

SUSE Linux Enterprise Server 10 SP 4 LTSS U1

```
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-bigsmp-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-default-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-default-x86_64
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-smp-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-smp-x86_64
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-xen-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-xen-x86_64
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-xenpae-x86
```

SUSE Linux Enterprise Server 10 SP 4 LTSS U2

```
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-bigsmp-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-default-x86
```

```
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-default-x86_64
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-smp-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-smp-x86_64
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-xen-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-xen-x86_64
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-xenpae-x86
```

SUSE Linux Enterprise Server 11 GA

```
SLES11-GA-2.6.27.19-5-default-x86
SLES11-GA-2.6.27.19-5-default-x86_64
SLES11-GA-2.6.27.19-5-pae-x86
```

SUSE Linux Enterprise Server 11 SP 1

```
SLES11-SP1-2.6.32.12-0.6-default-x86
SLES11-SP1-2.6.32.12-0.6-default-x86_64
SLES11-SP1-2.6.32.12-0.6-pae-x86
```

SUSE Linux Enterprise Server 11 SP 1 U14

```
SLES11-SP1_U14-2.6.32.54-0.3-default-x86
SLES11-SP1_U14-2.6.32.54-0.3-default-x86_64
SLES11-SP1_U14-2.6.32.54-0.3-pae-x86
```

SUSE Linux Enterprise Server 11 SP 1 U15

```
SLES11-SP1_U15-2.6.32.59-0.3-default-x86
SLES11-SP1_U15-2.6.32.59-0.3-default-x86_64
SLES11-SP1_U15-2.6.32.59-0.3-pae-x86
```

SUSE Linux Enterprise Server 11 SP 1 U16

```
SLES11-SP1_U16-2.6.32.59-0.7-default-x86
SLES11-SP1_U16-2.6.32.59-0.7-default-x86_64
SLES11-SP1_U16-2.6.32.59-0.7-pae-x86
```

SUSE Linux Enterprise Server 11 SP 1 LTSS U1

```
SLES11-SP1_LTSS_U1-2.6.32.59-0.9-default-x86
SLES11-SP1_LTSS_U1-2.6.32.59-0.9-default-x86_64
SLES11-SP1_LTSS_U1-2.6.32.59-0.9-pae-x86
```

SUSE Linux Enterprise Server 11 SP 1 LTSS U2

```
SLES11-SP1_LTSS_U2-2.6.32.59-0.13-default-x86
SLES11-SP1_LTSS_U2-2.6.32.59-0.13-default-x86_64
SLES11-SP1_LTSS_U2-2.6.32.59-0.13-pae-x86
```

SUSE Linux Enterprise Server 11 SP 2 GA

```
SLES11SP2-GA-3.0.13-0.27-default-x86
SLES11SP2-GA-3.0.13-0.27-default-x86_64
SLES11SP2-GA-3.0.13-0.27-pae-x86
SLES11SP2-GA-3.0.13-0.27-xen-x86
SLES11SP2-GA-3.0.13-0.27-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U1

```
SLES11SP2-U1-3.0.26-0.7-default-x86
SLES11SP2-U1-3.0.26-0.7-default-x86_64
SLES11SP2-U1-3.0.26-0.7-pae-x86
SLES11SP2-U1-3.0.26-0.7-xen-x86
SLES11SP2-U1-3.0.26-0.7-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U2

```
SLES11SP2-U2-3.0.31-0.9-default-x86
SLES11SP2-U2-3.0.31-0.9-default-x86_64
SLES11SP2-U2-3.0.31-0.9-pae-x86
SLES11SP2-U2-3.0.31-0.9-xen-x86
SLES11SP2-U2-3.0.31-0.9-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U3

```
SLES11SP2-U3-3.0.34-0.7-default-x86
SLES11SP2-U3-3.0.34-0.7-default-x86_64
SLES11SP2-U3-3.0.34-0.7-pae-x86
SLES11SP2-U3-3.0.34-0.7-xen-x86
SLES11SP2-U3-3.0.34-0.7-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U4

```
SLES11SP2-U4-3.0.38-0.5-default-x86
SLES11SP2-U4-3.0.38-0.5-default-x86_64
SLES11SP2-U4-3.0.38-0.5-pae-x86
SLES11SP2-U4-3.0.38-0.5-xen-x86
SLES11SP2-U4-3.0.38-0.5-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U5

```
SLES11SP2-U5-3.0.42-0.7-default-x86
SLES11SP2-U5-3.0.42-0.7-default-x86_64
SLES11SP2-U5-3.0.42-0.7-pae-x86
SLES11SP2-U5-3.0.42-0.7-xen-x86
SLES11SP2-U5-3.0.42-0.7-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U6

```
SLES11SP2-U6-3.0.51-0.7.9-default-x86
SLES11SP2-U6-3.0.51-0.7.9-default-x86_64
SLES11SP2-U6-3.0.51-0.7.9-pae-x86
SLES11SP2-U6-3.0.51-0.7.9-xen-x86
SLES11SP2-U6-3.0.51-0.7.9-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U7

```
SLES11SP2-U7-3.0.58-0.6.2-default-x86
SLES11SP2-U7-3.0.58-0.6.2-default-x86_64
SLES11SP2-U7-3.0.58-0.6.2-pae-x86
SLES11SP2-U7-3.0.58-0.6.2-xen-x86
SLES11SP2-U7-3.0.58-0.6.2-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U8

```
SLES11SP2-U8-3.0.58-0.6.6-default-x86
SLES11SP2-U8-3.0.58-0.6.6-default-x86_64
SLES11SP2-U8-3.0.58-0.6.6-pae-x86
SLES11SP2-U8-3.0.58-0.6.6-xen-x86
SLES11SP2-U8-3.0.58-0.6.6-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U9

```
SLES11SP2-U9-3.0.74-0.6.6-default-x86
SLES11SP2-U9-3.0.74-0.6.6-default-x86_64
SLES11SP2-U9-3.0.74-0.6.6-pae-x86
SLES11SP2-U9-3.0.74-0.6.6-xen-x86
SLES11SP2-U9-3.0.74-0.6.6-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U10

```
SLES11SP2-U10-3.0.74-0.6.8-default-x86
SLES11SP2-U10-3.0.74-0.6.8-default-x86_64
SLES11SP2-U10-3.0.74-0.6.8-pae-x86
SLES11SP2-U10-3.0.74-0.6.8-xen-x86
SLES11SP2-U10-3.0.74-0.6.8-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U11

```
SLES11SP2-U11-3.0.74-0.6.10-default-x86
SLES11SP2-U11-3.0.74-0.6.10-default-x86_64
SLES11SP2-U11-3.0.74-0.6.10-pae-x86
SLES11SP2-U11-3.0.74-0.6.10-xen-x86
SLES11SP2-U11-3.0.74-0.6.10-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U12

```
SLES11SP2-U12-3.0.80-0.5-default-x86
SLES11SP2-U12-3.0.80-0.5-default-x86_64
SLES11SP2-U12-3.0.80-0.5-pae-x86
SLES11SP2-U12-3.0.80-0.5-xen-x86
SLES11SP2-U12-3.0.80-0.5-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U13

```
SLES11SP2-U13-3.0.80-0.7-default-x86
SLES11SP2-U13-3.0.80-0.7-default-x86_64
SLES11SP2-U13-3.0.80-0.7-pae-x86
SLES11SP2-U13-3.0.80-0.7-xen-x86
SLES11SP2-U13-3.0.80-0.7-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U14

```
SLES11SP2-U14-3.0.93-0.5-default-x86
SLES11SP2-U14-3.0.93-0.5-default-x86_64
SLES11SP2-U14-3.0.93-0.5-pae-x86
SLES11SP2-U14-3.0.93-0.5-xen-x86
SLES11SP2-U14-3.0.93-0.5-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U15

```
SLES11SP2-U15-3.0.101-0.5-default-x86
SLES11SP2-U15-3.0.101-0.5-default-x86_64
SLES11SP2-U15-3.0.101-0.5-pae-x86
SLES11SP2-U15-3.0.101-0.5-xen-x86
SLES11SP2-U15-3.0.101-0.5-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U16

```
SLES11SP2-U16-3.0.101-0.7.15-default-x86
SLES11SP2-U16-3.0.101-0.7.15-default-x86_64
SLES11SP2-U16-3.0.101-0.7.15-pae-x86
SLES11SP2-U16-3.0.101-0.7.15-xen-x86
SLES11SP2-U16-3.0.101-0.7.15-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 U17

```
SLES11SP2-U17-3.0.101-0.7.17-default-x86
SLES11SP2-U17-3.0.101-0.7.17-default-x86_64
SLES11SP2-U17-3.0.101-0.7.17-pae-x86
SLES11SP2-U17-3.0.101-0.7.17-xen-x86
```

```
SLES11SP2-U17-3.0.101-0.7.17-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 LTSS U1

```
SLES11SP2-LTSS_U1-3.0.101-0.7.19-default-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-default-x86_64
SLES11SP2-LTSS_U1-3.0.101-0.7.19-pae-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-xen-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 2 LTSS U2

```
SLES11SP2-LTSS_U2-3.0.101-0.7.21-default-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-default-x86_64
SLES11SP2-LTSS_U2-3.0.101-0.7.21-pae-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-xen-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 3 GA

```
SLES11SP3-GA-3.0.76-0.11-default-x86
SLES11SP3-GA-3.0.76-0.11-default-x86_64
SLES11SP3-GA-3.0.76-0.11-pae-x86
SLES11SP3-GA-3.0.76-0.11-xen-x86
SLES11SP3-GA-3.0.76-0.11-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 3 U1

```
SLES11SP3-U1-3.0.82-0.7-default-x86
SLES11SP3-U1-3.0.82-0.7-default-x86_64
SLES11SP3-U1-3.0.82-0.7-pae-x86
SLES11SP3-U1-3.0.82-0.7-xen-x86
SLES11SP3-U1-3.0.82-0.7-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 3 U2

```
SLES11SP3-U2-3.0.93-0.8-default-x86
SLES11SP3-U2-3.0.93-0.8-default-x86_64
SLES11SP3-U2-3.0.93-0.8-pae-x86
SLES11SP3-U2-3.0.93-0.8-xen-x86
SLES11SP3-U2-3.0.93-0.8-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 3 U3

```
SLES11SP3-U3-3.0.101-0.8-default-x86
SLES11SP3-U3-3.0.101-0.8-default-x86_64
SLES11SP3-U3-3.0.101-0.8-pae-x86
SLES11SP3-U3-3.0.101-0.8-xen-x86
SLES11SP3-U3-3.0.101-0.8-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 3 U4

```
SLES11SP3-U4-3.0.101-0.15-default-x86
SLES11SP3-U4-3.0.101-0.15-default-x86_64
SLES11SP3-U4-3.0.101-0.15-pae-x86
SLES11SP3-U4-3.0.101-0.15-xen-x86
SLES11SP3-U4-3.0.101-0.15-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 3 U5

```
SLES11SP3-U5-3.0.101-0.21-default-x86
SLES11SP3-U5-3.0.101-0.21-default-x86_64
SLES11SP3-U5-3.0.101-0.21-pae-x86
```

```
SLES11SP3-U5-3.0.101-0.21-xen-x86
SLES11SP3-U5-3.0.101-0.21-xen-x86_64
```
SUSE Linux Enterprise Server 11 SP 3 U6
```
SLES11SP3-U6-3.0.101-0.29-default-x86
SLES11SP3-U6-3.0.101-0.29-default-x86_64
SLES11SP3-U6-3.0.101-0.29-pae-x86
SLES11SP3-U6-3.0.101-0.29-xen-x86
SLES11SP3-U6-3.0.101-0.29-xen-x86_64
```
SUSE Linux Enterprise Server 11 SP 3 U7
```
SLES11SP3-U7-3.0.101-0.31-default-x86
SLES11SP3-U7-3.0.101-0.31-default-x86_64
SLES11SP3-U7-3.0.101-0.31-pae-x86
SLES11SP3-U7-3.0.101-0.31-xen-x86
SLES11SP3-U7-3.0.101-0.31-xen-x86_64
```
SUSE Linux Enterprise Server 11 SP 3 U8
```
SLES11SP3-U8-3.0.101-0.35-default-x86
SLES11SP3-U8-3.0.101-0.35-default-x86_64
SLES11SP3-U8-3.0.101-0.35-pae-x86
SLES11SP3-U8-3.0.101-0.35-xen-x86
SLES11SP3-U8-3.0.101-0.35-xen-x86_64
```
SUSE Linux Enterprise Server 11 SP 4 GA
```
SLES11SP4-GA-3.0.101-63-default-x86
SLES11SP4-GA-3.0.101-63-default-x86_64
SLES11SP4-GA-3.0.101-63-pae-x86
SLES11SP4-GA-3.0.101-63-xen-x86
SLES11SP4-GA-3.0.101-63-xen-x86_64
```
SUSE Linux Enterprise Server 11 SP 4 U1
```
SLES11SP4-U1-3.0.101-65-default-x86
SLES11SP4-U1-3.0.101-65-default-x86_64
SLES11SP4-U1-3.0.101-65-pae-x86
SLES11SP4-U1-3.0.101-65-xen-x86
SLES11SP4-U1-3.0.101-65-xen-x86_64
```
SUSE Linux Enterprise Server 11 SP 4 U2
```
SLES11SP4-U2-3.0.101-68-default-x86
SLES11SP4-U2-3.0.101-68-default-x86_64
SLES11SP4-U2-3.0.101-68-pae-x86
SLES11SP4-U2-3.0.101-68-xen-x86
SLES11SP4-U2-3.0.101-68-xen-x86_64
```

B.2.3   Other Linux Distributions That Use blkwatch Drivers

PlateSpin Protect supports other Linux distributions listed in Table B-1 if the distribution is based on a supported release version of Red Hat Enterprise Linux or SUSE Linux Enterprise Server. You can use the pre-compiled blkwatch driver for the supported Linux Distribution.

Table B-1 *Blkwatch Driver Support for Other Linux Distributions*

Other Linux Distribution	Based on a Supported Release Version for RHEL or SLES	Notes
CentOS	Red Hat Enterprise Linux	
Open Enterprise Server (OES)	SUSE Linux Enterprise Server 11 SP 1 or later	The default kernel version 3.0.13 of OES 11 SP2 is not supported. Upgrade to kernel version 3.0.27 or higher version before you inventory the workload.
Oracle Linux (OL) (formerly Oracle Enterprise Linux (OEL))	Red Hat Enterprise Linux	Blkwatch drivers are available for the standard kernel and the Unbreakable Enterprise Kernel (UEK) as noted in the Section B.2.2, "List of Distributions," on page 128. For other Oracle Linux distributions, precompiled drivers are available only for the corresponding Red Hat Compatible Kernel (RHCK). Workloads using the Oracle Linux Unbreakable Enterprise Kernel are not supported in PlateSpin Protect 11.2 and lower versions.

For a list of supported kernel distributions, see "List of Distributions" on page 128.

C Synchronizing Serial Numbers on Cluster Node Local Storage

This section details the procedure you can use to change local volume serial numbers to match each node of the Windows cluster that you want to protect. The information includes the use of the Volume Manager utility (`VolumeManager.exe`) to synchronize serial numbers on cluster node local storage.

To download and run the utility:

1 Download the `VolumeManager.exe` file from the PlateSpin Protect download page:

   1a Go to Micro Focus Downloads (https://www.microfocus.com/support-and-services/download/).

   1b Select PlateSpin Protect from the **Browse by Product** list, or type the product name in the **Browse by Product** field to find the product and then select it.

   1c If a list of releases is available, select PlateSpin Protect 11.2.1.

   1d On the Download overview page, click **proceed to download**, then log in with your customer account credentials.

   1e Click **accept** to acknowledge and agree to the U.S. Export Laws and Regulations.

   1f On the Download page, click **download** next to the `VolumeManager.exe` file, then save the file.

2 Copy the downloaded file to an accessible location on each cluster node.

3 On the active node of the cluster, open an administrative command prompt, navigate to the location of the downloaded utility, and run the following command:

```
VolumeManager.exe -l
```

A listing of the local volumes and their respective serial numbers is displayed. For example:

```
Volume Listing:

--------------------

DriveLetter (*:) VolumeId="System Reserved" SerialNumber: AABB-CCDD

DriveLetter (C:) VolumeId=C:\ SerialNumber: 1122-3344
```

Make note of these serial numbers or keep them displayed for later comparison.

4 Verify that all local storage serial numbers of the active node match the local storage serial numbers on each of the other nodes in the cluster.

   4a On each cluster node, run the `VolumeManager.exe -l` command to obtain its volume serial numbers.

   4b Compare the local storage serial numbers of the active node (Step 3) against the local storage serial numbers of the node (Step 4a).

   4c (Conditional) If there are any differences in the serial numbers between the active node and this node, take note of the serial number you want to propagate on this node and run the following command to set, and then to verify the serial number:

   ```
   VolumeManager -s <VolumeId> <serial-number>
   ```

Following are two examples of how this command could be used:

- `VolumeManager -s "System Reserved" AAAA-AAAA`
- `VolumeManager -s C:\ 1111-1111`

4d When you have successfully changed all of the volume serial numbers on a node of the cluster, you need to restart that node.

4e Repeat Step 4a through Step 4d for each node of the cluster.

5 (Conditional) If the cluster has already been protected in a PlateSpin environment, we recommend running a full replication on the active node to ensure that any changes are propagated to the database.

D Protect Agent Utility

Protect Agent is a command line utility that you can use to install, upgrade, query, or uninstall the block-based transfer drivers.

Although a reboot is always required when you install, uninstall, or upgrade drivers, the Protect Agent allows you to better control when the action occurs and therefore, when the server reboots. For example, you can use the Protect Agent to install the drivers during scheduled down time, instead of during the first replication.

- Section D.1, "Using the Protect Agent Utility for Windows," on page 145
- Section D.2, "Using Protect Agent with Block-Based Transfer Drivers," on page 146

D.1 Using the Protect Agent Utility for Windows

To download the Protect Agent utility for Windows to the source workload:

1 Log in to the source Windows computer as the Administrator user.

2 In a web browser, launch the Web Interface and log in.

3 Click the **Downloads** tab.

4 Click the Protect Agent application link for the Windows target platform, then save the compressed `ProtectAgent.cli.exe` file.

5 Extract the contents of the file to access the executable file.

6 (Optional) View the Protect Agent Help by entering

   `Protect.Agent.cli.exe -h`

The utility is available on the PlateSpin Server host in a compressed file. Extract the contents of the file to access the executable file.

`C:\Program Files\PlateSpin Protect Server\bin\ProtectAgent`

The syntax for running the Protect Agent utility for Windows is:

`ProtectAgent.cli.exe {command} [command_option] [/psserver=%IP%]`

Table D-1 describes the commands, command option, and switch available for the `ProtectAgent.cli.exe` command.

Table D-1 *Protect Agent Utility for Windows Commands, Command Option, and Switch*

Usage	Description
Commands	
`h |?	help`	Displays usage and options for the command.
`logs	view-logs`	Opens the application log directory.

Usage	Description

`status`
 	
`/status [/psserver=%IP%]`
	Shows installation status for the PlateSpin controller and drivers on this workload.
	
	If you specify the PlateSpin Server, it checks for driver upgrades from the server.

`din | driver-install`
 	
`/din [/psserver=%IP%]`
	Installs the PlateSpin drivers.
	
	If you specify the PlateSpin Server, it checks for driver upgrades from the server.

`dup | driver-upgrade`
 	
`/dup [/psserver=%IP%]`
	Upgrades the PlateSpin drivers.
	
	If you specify the PlateSpin Server, it checks for driver upgrades from the server.

`dun | driver-uninstall`
 	
`[/dun /psserver=%IP%]`
	Uninstalls the PlateSpin drivers.

`con | config`
 	
`/con /setting=<`*`setting_name`*`>:<`*`value`*`>`
 	
Example:
 	
`ProtectAgent.cli.exe /config / setting=psserver:10.10.10.202`
	Specifies the name of the setting and its value to change in the configuration file on this workload.
	
	The `psserver` option stops the OFX Controller (`ofxcontroller`) service, modifies the `OfxController.exe.config` file with the new IP address, and restarts the service. If you modify the public IP address of the PlateSpin Server, you must run this command on each of the source workloads that are configured for the server.

Switch	
`/psserver=%IP%`	Downloads the block-based transfer drivers from the specified server when you invoke the `status`, `driver-install`, or `driver-upgrade` options.
Command Option	

`setting`
 	
`/setting=<`*`setting_name`*`>:<`*`value`*`>`
	Specifies the setting name and value of the configuration setting to modify.
	
	Supported setting names are:
	
	`psserver`
	`altAddress`
	`heartbeat`

D.2 Using Protect Agent with Block-Based Transfer Drivers

A copy of the block-based transfer drivers is bundled with the Protect Agent utility. You can alternatively specify the `/psserver=` command line switch in order to download the drivers from the PlateSpin Server when you invoke the `status`, `driver-install`, or `driver-upgrade` options. This is useful when the server is patched with a new driver package, but the Protect Agent command line utility is not patched.

NOTE: To avoid confusion, the recommended method of using the Protect Agent is to install, uninstall, or upgrade the drivers and then reboot prior to doing a replication.

You should reboot the source workload each time that you install, upgrade, or uninstall the drivers. The reboot forces the running driver to stop and the new driver to be applied on system restart. If you do not reboot the system prior to replication, the source continues to act as if the operation has not been completed. For example, if you install drivers without rebooting the system, the source acts as if no driver is installed during replication. Similarly, if you upgrade the drivers without rebooting, the source continues to use the running driver during replication until you reboot the system.

If the version of the installed driver is different than the version of the running driver, the `status` option will remind the user to reboot. For example:

```
C:\ProtectAgent\ProtectAgent.cli.exe status
Step 1 of 2: Querying the PlateSpin controller service
    Done
Step 2 of 2: Querying the installed PlateSpin driver version
    Done

The task completed successfully
PlateSpin Controller Service Status
    Status: Running
    Version: 9.9.9.9
    Last Successful Contact: 1/5/2015 12:14:25 PM

PlateSpin Driver Status
    Installed Driver Version: 8.0.0.11
    Running Driver Version: Not running. Reboot to load the driver.
    Upgrade Available: No
```

PlateSpin creates a task to warn the user that a reboot is necessary in order to complete the driver installation or upgrade. The notification appears in the Tasks list (Figure D-1).

Figure D-1 *Reboot Notification Task*
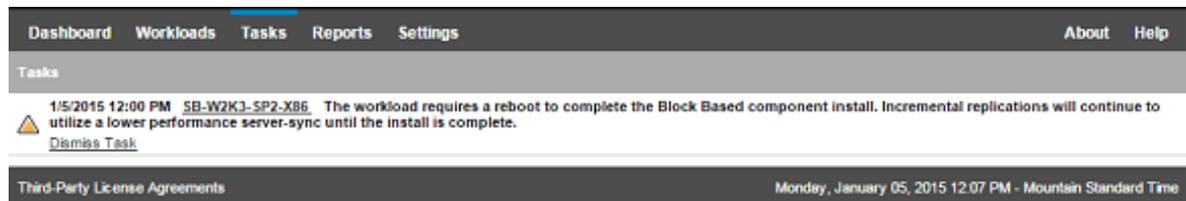
During replication, the notification appears on the Command Details page (Figure D-2).

Figure D-2 *Reboot Notification During Replication*

Rebooting the source machine applies and starts the installed or upgraded drivers. If the driver was recently installed, after the reboot, a full replication or a server-sync replication is required in order to ensure that all of a source's changes are captured. This server-sync replication requirement will be represented to the user in the Status field as a warning, as shown see Figure D-3. Subsequent incremental replications will complete as scheduled without a warning.

Figure D-3 *Server-Sync Required Notification*

IV Protecting Workloads

After you discover targets and workloads, you are ready to prepare for protection by configuring protection contracts for your workloads.

15 Workload Protection and Recovery

PlateSpin Protect creates a replica of your production workload and regularly updates that replica based on a schedule that you define.

The replica, or the *failover workload*, is a virtual machine managed by PlateSpin Protect that takes over the business function of your production workload in case of a disruption at the production site.

15.1 Prerequisites for Workload Protection

Prepare your containers and workloads for protection. See Part III, "Preparing Protection Targets and Sources," on page 87.

In an Active Directory domain, follow these best practices before you run the first full replication:

- Ensure that you update Windows (run Windows Update) on the source workload before you run the first full replication.
- Ensure that you set up your antivirus software with the recommended file and folder exclusions described in the *Microsoft KB 822158 article: Virus scanning recommendations for Enterprise computers that are running currently supported versions of Windows* (https://support.microsoft.com/en-us/kb/822158).
- If the Windows machine is a Domain Controller, ensure that you disable antivirus software on the system during the replication.

15.2 Configuring Protection Details and Preparing the Replication

Protection details control the workload protection and recovery settings and behavior over the entire life cycle of a workload under protection. At each phase of the protection and recovery workflow (Add-inventory, initial and ongoing Replications, Failover, Failback, and Reprotect), relevant settings are read from the protection details. See "Basic Workflow for Workload Protection and Recovery" on page 35. This collection of currently-active settings pertaining to the complete lifecycle of a workload's protection is referred to as the workload's *protection contract*.

To configure your workload's protection details:

1 Add a container. See "Adding Containers (Protection Targets)" on page 90.

2 Add a workload. See "Adding Workloads (Protection Sources)" on page 94.

3 On the Workloads page, select the required workload and click **Configure**.

Alternatively, you can click the name of the workload.

NOTE: If the PlateSpin Protect inventory does not have a container yet, the system prompts you to add one; do so by clicking **Add Container** at the bottom.

---

4 Select an **Initial Replication Method**. This indicates whether you want volume data transferred entirely from your workload to the failover VM or synchronized with volumes on an existing VM. See "Initial Replication Method (Full and Incremental)" on page 167.

5 Assign a protection target. This can be either a container or, if you have selected **Incremental Replication** as the initial replication method, a *prepared* workload. See "Initial Replication Method (Full and Incremental)" on page 167.

NOTE: If your inventory has only one container, your workload is automatically assigned to it.

---

6 Configure the protection details in each set of settings as dictated by your business continuity needs. See "Workload Protection Details" on page 154.

7 Correct any validation errors, if displayed by the PlateSpin Protect Web Interface.

8 Click **Save**.

Alternately, click **Save & Prepare**. This saves the settings and simultaneously executes the **Prepare Replication** command (installing data transfer drivers on the source workload if necessary and creating the initial VM replica of your workload).

Wait for the process to complete. Upon completion, a **Workload configuration completed** event is shown on the Dashboard.

15.2.1 Workload Protection Details

Workload protection details are represented by five sets of parameters, as described in Table 15-1:
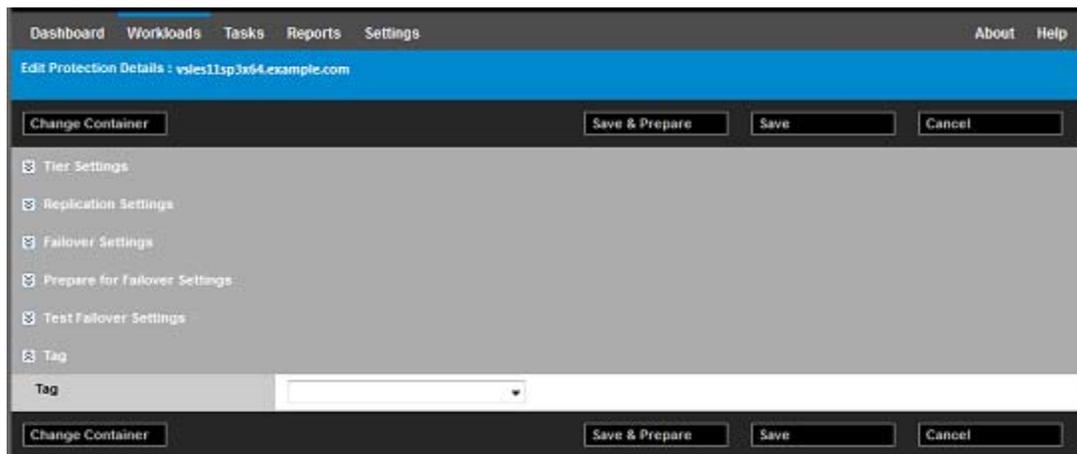
You can expand or collapse each parameter set by clicking the ⊠ icon at the left.

Table 15-1 *Workload Protection Details*

Parameter Settings	Details
Tier Settings	

| Parameter Settings	Details
Protection Tier	Specify the Protection Tier that the current protection uses. See "Protection Tiers" on page 166.
Replication Settings	
Transfer Method	(Windows) Select a file-based or block-based data transfer mechanism. For information about block-level replication with or without block-based components, see "Supported Data Transfer Methods" on page 22. To enable encryption, select the **Encrypt Data Transfer** option. See "Encryption of Data in Transmission" on page 24.
Transfer Encryption	(Linux) To enable encryption, select the **Encrypt Data Transfer** option. See "Encryption of Data in Transmission" on page 24.
Source Credentials	Specify the credentials required for accessing the workload. See "Guidelines for Workload and Container Credentials" on page 165.
CPU	(VM containers using VMware 5.1, 5.5, and 6.0 with a minimum VM hardware Level 8) Specify the number of sockets and the number of cores per socket for the failover workload. It automatically calculates the total cores. This parameter applies on the initial setup of a workload with an initial replication setting of **Full**. **NOTE:** The maximum number of cores the workload can use is subject to external factors such as the guest operating system, the VM hardware version, VMware licensing for the ESXi host, and ESXi host compute maximums for vSphere (see *vSphere 5.1 Configuration Maximums* (http://www.vmware.com/pdf/vsphere5/r51/vsphere-51-configuration-maximums.pdf)). Some distributions of a guest OS might not honor the cores and cores per socket configuration. For example, guest OSes using SLES 10 SP4 and OES 2 SP3 retain their original cores and sockets settings as installed, whereas other SLES, RHEL, and OES distributions honor the configuration.
Number of CPUs	(VM containers using VMware 4.1) Specify the required number of vCPUs (virtual CPUs) to assign to the failover workload. This parameter applies on the initial setup of a workload with an initial replication setting of **Full**. Each vCPU is presented to the guest OS on the VM container as a single core, single socket.
Replication Network	Separate replication traffic based on virtual networks defined on your VM container. See "Networking" on page 172. For this setting, you can also specify an MTU value to be used by the PlateSpin Protect Linux RAM Disk (LRD) replication network. Setting the value can help avoid jabber over networks (for example, a VPN) that have a smaller MTU value. The default value is empty string (nothing listed in the text box). When networking is configured in the LRD, this allows the network device to set its own default (which is usually 1500). If you enter a value, PlateSpin Protect adjusts the MTU while configuring the network interface.
Allowed Networks	Specify one or more network interfaces (NIC or IP address) on the source to use for replication traffic.
Resource Pool for Target VM	(VM container is part of a DRS Cluster) Specify the Resource Pool location where the failover VM is to be created.

Parameter Settings	Details
VM Folder for Target VM	(VM container is part of a DRS Cluster) Specify the VM folder location where the failover VM is to be created.
Configuration File Datastore	Select a datastore associated with your VM container for storing VM configuration files. See "Recovery Points" on page 167.
Protected Volumes	Select volumes for protection and to assign their replicas to specific datastores on your VM container.
Thin Disk	Select to enable the thin-provisioned virtual disk feature, whereby a virtual disk appears to the VM to have a set size, but only consumes the amount of disk space that is actually required by data on that disk.
Protected Logical Volumes	(Linux) Specify one or more LVM logical volumes to be protected for a Linux workload or the NSS Pools on an Open Enterprise Server workload.
Non-volume Storage	(Linux) Specify a storage area (such as a swap partition) that is associated with the source workload. This storage is re-created in the failover workload.
Volume Groups	(Linux) Specify the LVM volume groups to be protected with the LVM logical volumes listed in the Protected Logical Volumes section of the settings.
Services/Daemons to Stop During Replication	Select Windows services or Linux daemons that are automatically stopped during the replication. See "Service and Daemon Control" on page 169.
Failover Settings	
VM Memory	Specify the amount of memory allocated to the failover workload.
Hostname and Domain/Workgroup affiliation	Specify the identity and domain/workgroup affiliation of the failover workload when it is live. For domain affiliation, domain administrator credentials are required.
Network Connections	Specify the LAN settings of the failover workload. See "Networking" on page 172.
DNS Servers	Specify the IP address of the primary DNS server and an alternative DNS (optional).
Services/Daemon States to Change	Specify the startup state of specific application services (Windows) or daemons (Linux) See "Service and Daemon Control" on page 169.
Prepare for Failover Settings	
Temporary Failover Network	Specify the temporary LAN settings of the failover workload during the optional Prepare for Failover operation. See "Networking" on page 172.
Test Failover Settings	
VM Memory	Assign the required RAM to the temporary workload.
Hostname	Assign a host name to the temporary workload.
Domain/Workgroup	Affiliate the temporary workload with a domain or a workgroup. For domain affiliation, domain administrator credentials are required.
Network Connections	Specify the LAN settings of the temporary workload. See "Networking" on page 172.
DNS Servers	Specify the IP address of the primary DNS server and an alternative DNS (optional).

Parameter Settings	Details
Service/Daemon States to Change	Specify the startup state of specific application services (Windows) or daemons (Linux). See "Service and Daemon Control" on page 169.
Tags	
Tag	(Optional) Assign a tag to this workload. See "Tagging Workloads" on page 95.

15.3 Starting the Workload Protection

Workload protection is started by the **Run Replication** command:
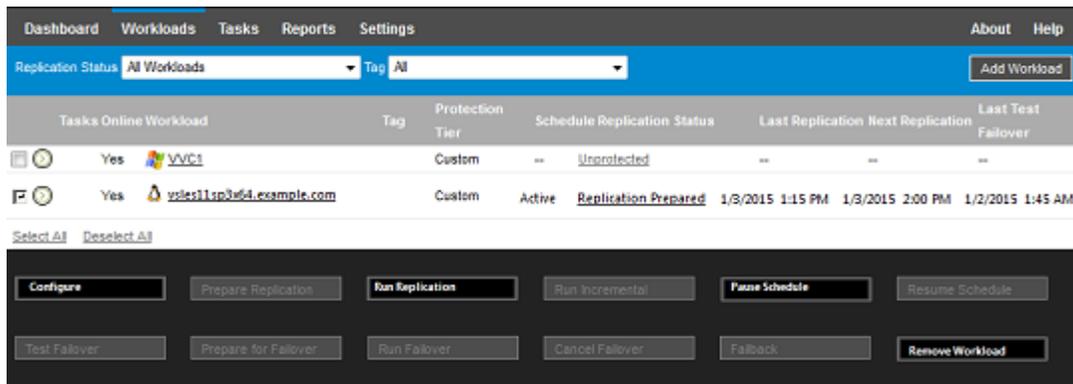
You can execute the Run Replication command after:

- Adding a workload.
- Configuring the workload's protection details.
- Preparing the initial replication.

When you are ready to proceed:

1 On the Workloads page, select the required workload, then click **Run Replication**.

2 Click **Execute**.

PlateSpin Protect starts the execution and displays a process indicator for the **Copy data** step
.

NOTE: After a workload has been protected:

- Changing the size of a volume that is under block-level protection invalidates the protection. The appropriate procedure is to
    1. Remove the workload from protection.
    2. Resize the volumes as required.
    3. Re-establish the protection by re-adding the workload, configuring its protection details, and starting replications.
- Any significant modification of the protected workload requires that the protection be re-established. Examples include adding volumes or network cards to the workload under protection.

15.4 Aborting Commands

You can abort a command after executing it and while it is underway, on the Command Details page of that particular command.

To access the Command Details page of any command that is underway:

1 Go to the Workloads page.

2 Locate the required workload and click the link representing the command currently executing on that workload, such as **Running Incremental**.

The Web Interface displays the appropriate Command Details page:

3 Click **Abort**.

15.5 Failover

In a *failover* operation, the failover workload within a PlateSpin Protect VM container takes over the business function of a failed production workload.

- ◆ Section 15.5.1, "Detecting Offline Workloads," on page 158
- ◆ Section 15.5.2, "Performing a Failover," on page 159
- ◆ Section 15.5.3, "Using the Test Failover Feature," on page 159

15.5.1 Detecting Offline Workloads

PlateSpin Protect constantly monitors your protected workloads. If an attempt to monitor a workload fails for a predefined number of times, PlateSpin Protect generates a **Workload is offline** event. Criteria that determine and log a workload failure are part of a workload protection's Tier settings. See "Tier Settings" row in "Workload Protection Details" on page 154.

If notifications are configured along with SMTP settings, PlateSpin Protect simultaneously sends a notification email to the specified recipients. See "Configuring Email Notification Services for Events and Replication Reports" on page 62.

If a workload failure is detected while the status of the replication is **Idle**, you can proceed to the **Run Failover** command. If a workload fails while an incremental is underway, the job stalls. In this case, abort the command (see "Aborting Commands" on page 158), and then proceed to the **Run Failover** command. See "Performing a Failover" on page 159.

Figure 15-1 shows the Web Interface's Dashboard page upon detecting a workload failure. Note the applicable tasks in the Tasks and Events pane:

15.5.2   Performing a Failover

Failover settings, including the failover workload's network identity and LAN settings, are saved together with the workload's protection details at configuration time. See "Failover Settings" in "Workload Protection Details" on page 154.

You can use the following methods to perform a failover:

- Select the required workload on the Workloads page and click **Run Failover**.
- Click the corresponding command hyperlink of the **Workload is offline** event in the Tasks and Events pane. See Figure 15-1.
- Run a **Prepare for Failover** command to boot the failover VM ahead of time You still have the option to cancel the failover (useful in staged failovers).

Use one of these methods to start the failover process and select a recovery point to apply to the failover workload (see "Recovery Points" on page 167). Click **Execute** and monitor the progress. Upon completion, the replication status of the workload should indicate **Live**.

For testing the failover workload or testing the failover process as part of a planned disaster recovery exercise, see "Using the Test Failover Feature" on page 159.

15.5.3   Using the Test Failover Feature

PlateSpin Protect provides you with the capability to test the failover functionality and the integrity of the failover workload. This is done by using the **Test Failover** command, which boots the failover workload in an isolated networking environment for testing the functionality of the failover and verifying the integrity of the failover workload.

When you execute the command, PlateSpin Protect applies the Test Failover Settings, as saved in the workload protection details, to the failover workload. See "Test Failover Settings" in "Workload Protection Details" on page 154.

To use the Test Failover feature:

1 Define an appropriate time window for testing and ensure that there are no replications underway. The replication status of the workload must be **Idle**.

2 On the Workloads page, select the required workload, click **Test Failover**, select a recovery point (see "Recovery Points" on page 167), and the click **Execute**.

Upon completion, PlateSpin Protect generates a corresponding event and a task with a set of applicable commands:

3 Verify the integrity and business functionality of the failover workload. Use the VMware vSphere Client to access the failover workload in the VM container

4 Mark the test as a **failure** or a **success**. Use the corresponding commands in the task (**Mark Test Failure**, **Mark Test Success**). The selected action is saved in the history of events associated with the workload and is retrievable by reports. **Dismiss Task** discards the task and the event.

Upon completion of the **Mark Test Failure** or **Mark Test Success** tasks, PlateSpin Protect discards temporary settings that were applied to the failover workload, and the protection returns to its pre-test state.

15.6 Failback

A *failback* operation restores the business function of a failed production workload in its original environment when the business function of a temporary failover workload is no longer required. Failback is the next logical step after a failover; it transfers the failover workload to its original infrastructure or, if necessary, a new one.

Supported failback methods depends on the target infrastructure type and the degree of automation of the failback process:

- **Automated Failback to a Virtual Machine:** Supported for VMware ESX platforms and VMware DRS Clusters.
- **Semi-Automated Failback to a Physical Machine:** Supported for all physical machines.
- **Semi-Automated Failback to a Virtual Machine:** Supported for Microsoft Hyper-V platforms.

The following topics provide more information:

15.6.1 Automated Failback to a VM Platform

PlateSpin Protect supports automated failback for Failback containers on a supported VMware ESXi Server or a VMware DRS Cluster. See "Supported VM Containers" on page 17.

To execute an automated failback of a failover workload to a target VMware container:

1 Following a failover, select the workload on the Workloads page and click **Failback**.

The system prompts you to make the following selections

2 Specify the following sets of parameters:

- **Workload Settings:** Specify the failover workload's host name or IP address and provide administrator-level credentials. Use the required credential format. See "Guidelines for Workload and Container Credentials" on page 165.
- **Failback Target Settings:** Specify the following parameters:
  - **Replication Method:** Select the scope of data replication. If you select **Incremental**, you must **Prepare** a target. See "Initial Replication Method (Full and Incremental)" on page 167.
  - **Target Type:** Select **Virtual Target**. If you don't yet have a failback container, click **Add Container** and inventory a supported container.

3 Click **Save and Prepare** and monitor the progress on the Command Details screen.

Upon successful completion, PlateSpin Protect loads the Ready for Failback screen, prompting you to specify the details of the failback operation.

4 Configure the failback details. See "Failback Details (Workload to VM)" on page 161.

5 Click **Save and Failback** and monitor the progress on the Command Details page. See Figure 15-2.

PlateSpin Protect executes the command. If you selected **Reprotect after Failback** in the Post-Failback parameter set, a **Reprotect** command is shown in the Web Interface.

Figure 15-2 *Failback Command Details*
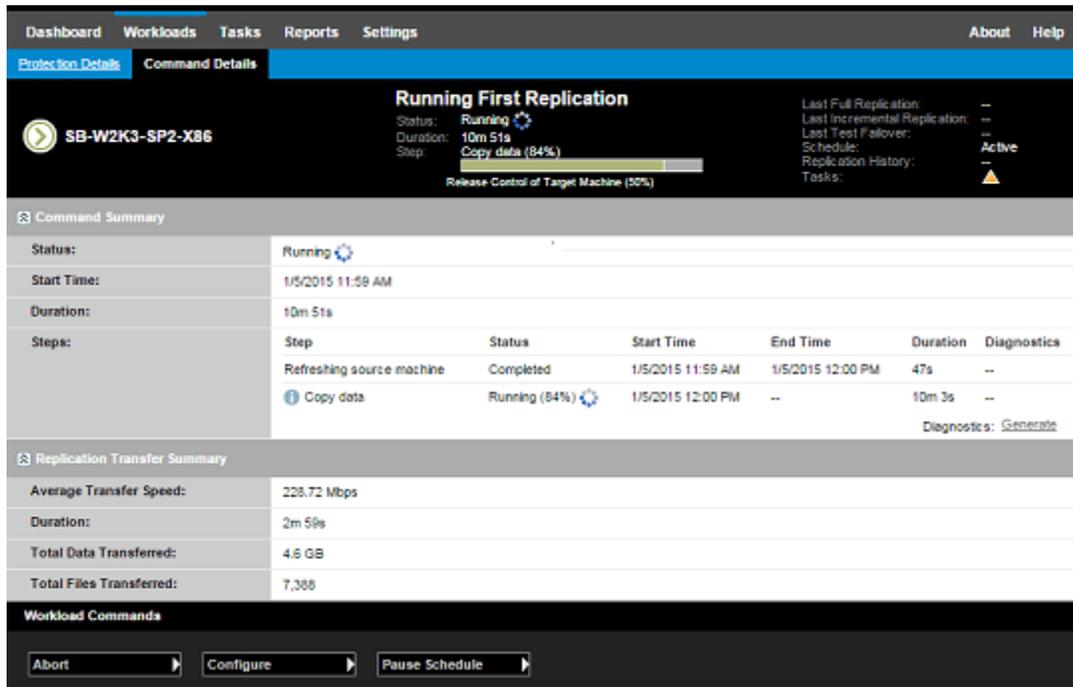
Failback Details (Workload to VM)

Failback details are represented by three sets of parameters that you configure when you are performing a workload failback operation to a virtual machine. See Table 15-2 for information about parameter settings.

Table 15-2 *Failback Details (Workload to VM)*

Parameter Settings	Details
Failback Settings	
Transfer Method	Select a data transfer mechanism and security through encryption. See "Encryption of Data in Transmission" on page 24.
Failback Network	Specify the network to use for failback traffic. This is a dedicated network based on virtual networks defined on your VM container. See "Networking" on page 172.
VM Datastore	Select a datastore associated with your failback container for the target workload.

Parameter Settings	Details
Volume Mapping	If the initial replication method is specified as "incremental", select source volumes and map to volumes on the failback target for synchronization.
Services/Daemons to stop	Specify the application services (Windows) or daemons (Linux) that are automatically stopped during the failback. See "Service and Daemon Control" on page 169.
Alternative Address for Source	Specify an additional IP address for the failed-over VM if applicable. See "Requirements for Protection across Public and Private Networks through NAT" on page 33.
Workload Settings	
CPU	(VM containers using VMware 5.1, 5.5, and 6.0 with a minimum VM hardware Level 8) Specify the number of sockets and the number of cores per socket for the failback to virtual workload. It automatically calculates the total cores. This parameter applies on the initial setup of a workload with an initial replication setting of **Full**.
	NOTE: The maximum number of cores the workload can use is subject to external factors such as the guest operating system, the VM hardware version, VMware licensing for the ESXi host, and ESXi host compute maximums for vSphere (see *vSphere 5.1 Configuration Maximums* (http://www.vmware.com/pdf/vsphere5/r51/vsphere-51-configuration-maximums.pdf)).
	Some distributions of a guest OS might not honor the cores and cores per socket configuration. For example, guest OSes using SLES 10 SP4 and OES 2 SP3 retain their original cores and sockets settings as installed, whereas other SLES, RHEL, and OES distributions honor the configuration.
Number of CPUs	(VM containers using VMware 4.1) Specify the required number of vCPUs (virtual CPUs) to assign to the failback to virtual workload. This parameter applies on the initial setup of a workload with an initial replication setting of **Full**. Each vCPU is presented to the guest OS on the VM container as a single core, single socket.
VM Memory	Assign the required RAM to the target workload.
Hostname, Domain/Workgroup	Specify the identity and domain/workgroup affiliation of the target workload. For domain affiliation, domain administrator credentials are required.
Network Connections	Specify the network mapping of the target workload based on the virtual networks of the underlying VM container.
Service States to Change	Specify the startup state of specific application services (Windows) or daemons (Linux). See "Service and Daemon Control" on page 169.
Post-Failback Settings	
Reprotect Workload	Select this option if you plan to re-create the protection contract for the target workload after deployment. This option maintains a continuous event history for the workload and auto-assigns/designates a workload license.
Reprotect after Failback	Select this option if you intend to re-create a protection contract for the target workload. When the failback is complete, a **Reprotect** command will be available in the Web Interface for the failed-back workload.

Parameter Settings	Details
No reprotect	Select this option if you do not intend to re-create a protection contract for the target workload. To protect the failed-back workload upon completion, you will have to re-inventory that workload and reconfigure its protection details.

15.6.2 Semi-Automated Failback to a Physical Machine

Use these steps to fail a workload back to a physical machine after a failover. The physical machine might be either the original infrastructure or a new one.

1 Register the required physical machine with your PlateSpin Server. See "Failback to Physical Machines" on page 172.

2 If there are missing or incompatible drivers, upload the required drivers to the PlateSpin Protect device driver database. See "Preparing Device Drivers for Physical Failback Targets" on page 97.

3 Following a failover, select the workload on the Workloads page and click **Failback**.

4 Specify the following sets of parameters:

 ◆ **Workload Settings:** Specify the failover workload's host name or IP address and provide administrator-level credentials. Use the required credential format (see "Guidelines for Workload and Container Credentials" on page 165.

 ◆ **Failback Target Settings:** Specify the following parameters:

  ◆ **Replication Method:** Select the scope of data replication.

   See "Initial Replication Method (Full and Incremental)" on page 167.

  ◆ **Target Type:** Select the **Physical Target** option and then select the physical machine you registered in Step 1.

5 Click **Save and Prepare** and monitor the progress on the Command Details screen.

 Upon successful completion, PlateSpin Protect loads the Ready for Failback screen, prompting you to specify the details of the failback operation.

6 Configure the failback details, then click **Save and Failback**.

 Monitor the progress on the Command Details screen.

15.6.3 Semi-Automated Failback to a Virtual Machine

This failback type follows a process similar to the Semi-Automated Failback to a Physical Machine for a VM target other than a natively-supported VMware container. During this process, you direct the system to regard a VM target as a physical machine.

You can do a semi-automated failback to a container, for which there is fully-automated failback support (VMware ESX and DRS Cluster targets).

You can also do a semi-automated failback for target VM platforms on Microsoft Hyper-V Server 2012 hosts.

To start the Hyper-V VMs on failover:

1 In a text editor, modify each Hyper-V host's `/etc/vmware/config` file by adding the following line:

```
vhv.allow = "TRUE"
```

2 In the vSphere Web Client, modify the failover VM Settings for the CPU:

    2a Under the **Virtual Hardware** tab, select **CPU**.

    2b In **Hardware virtualization**, select **Expose hardware assisted virtualization to guest OS**.

3 In the vSphere Web Client, modify the failover VM Settings for the CPU ID:

    3a Under the **VM Options** tab, expand **Advanced**, then select **Edit configuration parameters**.

    3b Verify the following setting:

```
hypervisor.cpuid.v0 = FALSE
```

15.7 Reprotecting a Workload

A **Reprotect** operation, the next logical step after a **Failback**, completes the workload protection lifecycle and starts it anew. Following a successful Failback operation, a **Reprotect** command becomes available in the Web Interface, and the system applies the same protection details as those indicated during the initial configuration of the protection contract.

NOTE: The **Reprotect** command becomes available only if you selected the **Reprotect** option in the Failback details. See "Failback" on page 160.

The rest of the workflow covering the protection lifecycle is the same as that in normal workload protection operations; you can repeat it as many times as required.

16 Essentials of Workload Protection

This section provides information about the different functional areas of a workload protection contract.

16.1 Guidelines for Workload and Container Credentials

PlateSpin Protect must have administrator-level access to workloads and appropriate role configuration for containers. Throughout the workload protection and recovery workflow, PlateSpin Protect prompts you to specify credentials that must be provided in a specific format.

Table 16-1 *Workload and Container Credentials*

To Discover	Credentials	Remarks
All Windows workloads	Local or domain administrator credentials.	For the user name, use this format: ◆ For domain member machines: *authority\principal* ◆ For workgroup member machines: *hostname\principal*
Windows Clusters	Domain administrator credentials	For domain member machines: *authority\principal*
All Linux workloads	Root-level user name and password	Non-root accounts must be properly configured to use `sudo`. See Knowledgebase Article 7920711 (https://www.netiq.com/support/kb/doc.php?id=7920711).

To Discover	Credentials	Remarks
VMware ESX or ESXi host	VMware account with an appropriate role configuration. To set up roles for Protect multitenancy, see "Defining VMware Roles for Multitenancy" on page 55.	If ESX is configured for Windows domain authentication, you can also use your Windows domain credentials.
VMware vCenter Server	VMware account with an appropriate role configuration. To set up roles for Protect multitenancy, see "Defining VMware Roles for Multitenancy" on page 55.	

16.2 Protection Tiers

A Protection Tier is a custom collection of workload protection parameters that define the following:

◆ The frequency and recurrence pattern of replications

◆ Whether to encrypt data transmission

◆ Whether and how to apply data compression

◆ Whether to throttle available bandwidth to a specified throughput rate during data transfer

◆ Criteria for the system to consider a workload as offline (failed)

A Protection Tier is an integral part of every workload protection contract. During the configuration stage of a workload protection contract, you can select one of several built-in Protection Tiers and customize its attributes as required by that specific protection contract.

To create custom Protection Tiers in advance:

1 In your Web Interface, click **Settings > Protection Tiers > Create Protection Tier**.

2 Specify the parameters for the new Protection Tier:

Parameter	Action
Name	Type the name you want to use for the tier.
Incremental Recurrence	Specify the frequency of incremental replications and the incremental recurrence pattern. You can type directly in the **Start of recurrence** field, or click the calendar icon to select a date. Select **None** as the Recurrence Pattern to never use incremental replication.
Full Recurrence	Specify the frequency of full replications and the full recurrence pattern.

Parameter	Action
Blackout Window	Use these settings to force a replication blackout (for suspending scheduled replications during peak utilization hours or to prevent conflicts between VSS-aware software and the PlateSpin VSS block-level data transfer component). To specify a blackout window, click **Edit**, then select a blackout recurrence pattern (daily, weekly, etc.), and the blackout period's start and end times. **NOTE:** The blackout start and end times are based on the system clock on your PlateSpin Server.
Compression Level	These settings control whether and how workload data is compressed before transmission. See "Data Compression" on page 28. Select one of the available options. **Fast** consumes the least CPU resources on the source but yields a lower compression ratio, **Maximum** consumes the most, but yields a higher compression ratio. **Optimal**, the middle ground, is the recommended option.
Bandwidth Throttling	These settings control bandwidth throttling. See "Bandwidth Throttling" on page 28. To throttle replications to a specified rate, specify the required throughput value in Mbps and indicate the time pattern.
Recovery Points to Keep	Specify the number of recovery points to keep for workloads that use this Protection Tier. See "Recovery Points" on page 167.
Workload Failure	Specify the number of workload detection attempts before it is considered failed.
Workload Detection	Specify the time interval (in seconds) between workload detection attempts.

16.3 Recovery Points

A recovery point is a point-in-time snapshot of a workload. It allows a replicated workload to be restored to a specific state.

Each protected workload has at least one recovery point and may have a maximum of 32 recovery points.

WARNING: Recovery points that accumulate over time might cause your PlateSpin Protect storage to run out of space.

16.4 Initial Replication Method (Full and Incremental)

The *initial replication* is the creation of an initial base copy of a production workload to the failover workload (virtual replica) in a protection operation, or from a failover workload to its original virtual or physical infrastructure in preparation for a failback operation for the production workload.

In workload protection and failback operations, the Initial Replication parameter determines the scope of data transferred from a source to a target.

- ◆ **Full:** A full workload transfer takes place based on all of its data.

◆ **Incremental:** Only differences are transferred from a source to its target, provided that they have similar operating system and volume profiles.

  - ◆ **During protection:** The production workload is compared with an existing VM in the VM container. The existing VM might be one of the following:

    - ◆ A previously-protected workload's recovery VM (when a **Remove Workload** command's **Delete VM** option is deselected).

    - ◆ A VM that is manually imported in the VM container, such as a workload VM physically moved on portable media from the production site to a remote recovery site.

  - ◆ **During failback to a virtual machine:** The failover workload is compared with an existing VM in a failback container.

  - ◆ **During failback to a physical machine:** The failover workload is compared with a workload on the target physical machine, if the physical machine is registered with PlateSpin Protect (see "Semi-Automated Failback to a Physical Machine" on page 163).

During workload protection and failback to a VM host, selecting **Incremental** as the initial replication method requires that you browse, locate, and prepare the target VM for synchronization with the selected operation's source.

To set up the initial replication method:

1 Proceed with the required workload command, such as **Configure (Protection Details)** or **Failback.**

2 For the **Initial Replication Method** option, select **Incremental Replication**.

3 Click **Prepare Workload**.

The Web Interface displays the Prepare for Incremental Replication page.
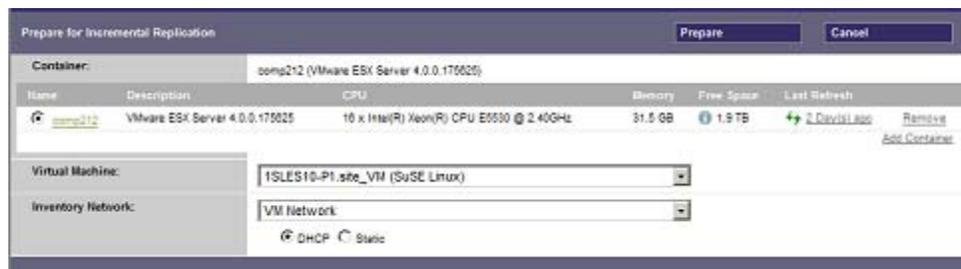
4 Select the required container, the virtual machine, and the network to use for communicating with the VM. If the specified target container is a VMware DRS Cluster, you can also specify a target Resource Pool for the workload.

5 Click **Prepare**.

Wait for the process to complete and for the user interface to return to the original command, then select the prepared workload.

NOTE: (Block-level data replications only) An initial incremental replication takes significantly longer than subsequent replications. This is because the system must compare the volumes on the source and the target block by block. Subsequent replications rely on changes detected by the block-based component while it is monitoring a running workload.

16.5 Service and Daemon Control

PlateSpin Protect enables you to control services and daemons:

◆ **Source service/daemon control:** During data transfer, you can automatically stop Windows services or Linux daemons that are running on your source workload. This ensures that the workload is replicated in a more consistent state than if you leave them running.

For example, for Windows workloads, consider stopping antivirus software services or services of third-party VSS-aware backup software.

For additional control of Linux sources during replication, consider the capability to run custom scripts on your Linux workloads during each replication. See "Using Freeze and Thaw Scripts for Every Replication (Linux)" on page 109.

◆ **Target startup state/run level control:** You can select the startup state (Windows) or the run level (Linux) of services/daemons on the failover VM. When you perform a Failover or Test Failover operation, you can specify which services or daemons you want to be running or stopped when the failover workload has gone live.

Common services that you might want to assign a `disabled` startup state are vendor-specific services that are tied to their underlying physical infrastructure and are not required in a virtual machine.

16.6 Volumes Storage

Upon adding a workload for protection, PlateSpin Protect inventories your source workload's storage media and automatically sets up options in the Web Interface that you use to specify the volumes you require for protection. For more information, see Section 1.1.5, "Supported Storage," on page 20.

Figure 16-1 shows the Replication Settings parameter set for a Linux workload with multiple volumes and two logical volumes in a volume group.

Figure 16-1   Volumes, Logical Volumes, and Volume Groups of a Protected Linux Workload

Figure 16-2 shows volume protection options of an OES 11 workload with options indicating that the LVM2 volume and NSS pool layout should be preserved and re-created for the failover workload:

Figure 16-2   *Replication Settings, Volume-Related Options (OES 11 Workload)*

Protected Volumes:	Include	Name		Total Size	Datastore	Thin Disk
	☑	/ (EXT3 - System)		13.8 GB	BBCSLESSAN	☐
Protected Logical Volumes:	Include	Name		Total Size	Volume Group	
	☑	/lvmtest1 (EXT3)		1007.9 MB	VolGroup1	
	☑	/lvmtest2 (EXT3)		2.0 GB	VolGroup1	
	☑	/opt/novell/nss/mnt/.pools /POOL1 (NSSFS)		12.0 GB	POOL1	
Non-volume Storage:	Include	Partition	Is Swap	Total Size	Datastore	Thin Disk
	☑	/dev/sda1	Yes	2.0 GB	BBCSLESSAN	☐
Volume Groups:	Include	Name		Total Size	Datastore	Thin Disk
	☑	VolGroup1		8.0 GB	BBCSLESSAN	☐
OES Volumes:	Include	Name		Total Size	Datastore	Thin Disk
	☑	POOL1		12.0 GB	BBCSLESSAN	☐
Daemons to Stop During Replication:	--					

Figure 16-3 shows volume protection options of an OES 2 workload with options indicating that the EVMS and NSS pool layout should be preserved and re-created for the failover workload:

Figure 16-3   *Replication Settings, Volume-Related Options (OES 2 Workload)*

Protected Logical Volumes:	Include	Name		Used Space	Free Space	Volume Group / EVMS Volume	
	☑	/ (REISERFS)		2.2 GB	2.2 GB	system	
	☑	/boot (EXT2)		13.0 MB	55.3 MB	/dev/evms/sda1	
	☑	/opt/novell/nss/mnt/.pools/NEWPOOL (NSSFS)		23.3 MB	999.6 MB	NEWPOOL	
Non-volume Storage:	Include	Partition	Is Swap	Total Size		Datastore / Volume Group	
	☑	/dev/system/swap	Yes	1.48 GB ▾		system	
Volume Groups:	Include	Name		Total Size		Datastore	Thin Disk
	☑	system		5.9 GB		dev-comp124:storage ▾	☐
EVMS Volumes:	Include	Name	Datastore	Total Size		Datastore	Thin Disk
	☑	/dev/evms/sda1		70.6 MB		dev-comp124:storage ▾	☐
	☑	NEWPOOL		1023.0 MB		dev-comp124:storage ▾	☐
Daemons to Stop During Replication:	Add Daemons						

16.7 Networking

PlateSpin Protect enables you to control your failover workload's network identity and LAN settings to prevent replication traffic from interfering with your main LAN or WAN traffic.

You can specify distinct networking settings in your workload protection details for use at different stages of the workload protection and recovery workflow:

- **Replication:** (Replication Settings parameter set) For separating regular replication traffic from your production traffic.
- **Failover:** (Failover Settings parameter set) For the failover workload to become part of your production network when it goes live.
- **Prepare for Failover:** (Prepare for Failover Settings network parameter) For network settings during the optional Prepare for Failover stage.
- **Test Failover:** (Test Failover Settings parameter set) For network settings to apply to the failover workload during a Test Failover stage.

16.8 Failback to Physical Machines

If the required target infrastructure for a failback operation is a physical machine, you must register it with PlateSpin Protect.

The registration of a physical machine is carried out by booting the target physical machine with the PlateSpin boot OFX ISO image.

- Section 16.8.1, "Downloading the PlateSpin Boot OFX ISO Image," on page 172
- Section 16.8.2, "Injecting Additional Device Drivers into the Boot ISO Image," on page 172
- Section 16.8.3, "Registering Physical Machines as Failback Targets with PlateSpin Protect," on page 174

16.8.1 Downloading the PlateSpin Boot OFX ISO Image

You can download the PlateSpin boot OFX ISO images (`bootofx.x2p.iso` for BIOS firmware-based targets and for UEFI firmware-based targets) from the PlateSpin Protect software download page.

1 Go to Micro Focus Downloads (https://www.microfocus.com/support-and-services/download/).
2 Select PlateSpin Protect from the **Browse by Product** list, or type the product name in the **Browse by Product** field to find the product and then select it.
3 On the Download overview page, click **proceed to download**, then log in with your customer account credentials.
4 Click **accept** to acknowledge and agree to the U.S. Export Laws and Regulations.
5 On the Download page, click **download** next to the `bootofx.x2p.iso` file, then save the file.

16.8.2 Injecting Additional Device Drivers into the Boot ISO Image

You can use a custom utility to package and inject additional Linux device drivers into the PlateSpin boot image before burning it on a CD.

To use this utility:

1 Obtain or compile `*.ko` driver files appropriate for the target hardware manufacturer.

IMPORTANT: Ensure that the drivers are valid for the kernel included with the ISO file (for x86 systems: `3.0.93-0.8-pae`, for x64 systems: `3.0.93-0.8-default`) and are appropriate for the target architecture. See also Knowledgebase Article 7005990 (https://www.netiq.com/support/kb/doc.php?id=7005990).

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

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

3 Copy the `rebuildiso.sh` script, located in the `/tools` subdirectory of the mounted ISO file, into a temporary working directory. When you have finished, unmount the ISO file (execute the command `umount <mount_point>`).

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

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

```
./rebuildiso.sh <ARGS> [-v] -m32|-m64 -i <ISO_file>
```

The following table lists the possible command line options for this command:

Option	Description
`-i <ISO_file>`	`<ISO_file>` is the ISO to modify, list, etc.
`-v`	If used together with the `-l` argument, the option causes the use of modinfo to obtain verbose driver information.
`-o`	If used together with the `-c` argument or the `-d` argument, the old copy of the ISO file is not overwritten.
`-m32`	Specifies 32-bit initrd injection.
`-m64`	Specifies 64-bit initrd injection.

The next table lists the possible arguments for use with this command. At least one of these arguments must be used in the command:

Argument	Description
`-d <path>`	`<path>` specifies the directory that contains the drivers (that is,. `*.ko` files) that you want to inject. On completion of the command, the ISO file is updated with the added drivers.
`-c <path>`	`<path>` specifies where a `ConfigureTakeControl.xml` file resides.

Argument	Description
`-l [<type>]`	*`<type>`* specifies a subset of drivers you want to list. The default is "all" types.
	Listed driver types beginning with a forward slash (/) are assumed to be located in *`<kernel_module_directory>`*`/kernel/`
	Listed driver types without a leading forward slash (/) are assumed to be located in *`<kernel_module_directory`*`/kernel/drivers/`
	Driver Subset Examples:

	`-l scsi`
	`-l 'net|video'`
	`-l '/net|net'`

	Special Usage of this Argument:
	If you want to list the available subdirectories of each of the subsets, use the argument like this: `-l INDEX`

Syntax Examples

- To list an index of 32-bit drivers:

  ```
  # ./rebuildiso.sh -i bootofx.x2p.iso -m32 -l INDEX
  ```

- To list drivers found in the /misc folder:

  ```
  # ./rebuildiso.sh -i bootofx.x2p.iso -m32 -l misc
  ```

- To inject 32-bit drivers from the /oem-drivers folder:

  ```
  # ./rebuildiso.sh -i bootofx.x2p.iso -m32 -d oem-drivers
  ```

- To inject 64-bit drivers from an /oem-drivers folder and also inject a customized ConfigureTakeControl.xml file:

  ```
  # ./rebuildiso.sh -i bootofx.x2p.iso -m64 -c ConfigureTakeControl.xml -d
  oem-drivers
  ```

16.8.3 Registering Physical Machines as Failback Targets with PlateSpin Protect

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

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

3 Use the boot CD to boot the target physical machine, then wait for the command prompt window to open.

4 (Linux only) For 64-bit systems, at the initial boot prompt, type the following:

`ps64`

5 Press Enter.

6 When you are prompted, enter the host name or the IP address of your PlateSpin Server host.

7 Provide your administrator-level credentials for the PlateSpin Server host, specifying an authority. For the user account, use this format:

domain\username or *hostname\username*

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

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

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

10 When prompted to indicate whether to use HTTPS, enter `Y` (yes) if you have enabled SSL, and `N` (no) if you have not.

After a few minutes, the physical machine should be available in the failback settings of the PlateSpin Protect Web Interface.

16.9 Protecting Windows Clusters

PlateSpin Protect supports the protection of a Microsoft Windows Server Cluster business services. For information about requirements and options for protecting nodes in a Windows Server Cluster, see Chapter 13, "Preparing for Windows Clusters Protection," on page 111.

- Section 16.9.1, "PlateSpin Failover," on page 175
- Section 16.9.2, "PlateSpin Failback," on page 175

16.9.1 PlateSpin Failover

When the PlateSpin failover operation is complete and the virtual one-node cluster comes online, you see a multi-node cluster with one active node (all other nodes are unavailable).

To perform a PlateSpin failover (or to test the PlateSpin failover on) a Windows cluster, the cluster must be able to connect to a domain controller. To leverage the test failover functionality, you need to protect the domain controller along with the cluster. During the test, bring up the domain controller, followed by the Windows cluster workload (on an isolated network).

16.9.2 PlateSpin Failback

A PlateSpin failback operation requires a full replication for Windows Cluster workloads.

If you configure the PlateSpin failback as a full replication to a physical target, you can use one of these methods:

- Map all disks on the PlateSpin virtual one-node cluster to a single local disk on the failback target.
- Add another disk (`Disk 2`) to the physical failback machine. You can then configure the PlateSpin failback operation to restore the failover machine's system volume to `Disk 1` and the failover machine's additional disks (previous shared disks) to `Disk 2`. This allows the system disk to be restored to the same size storage disk as the original source.

After a PlateSpin failback is complete, you must reattach the shared storage and rebuild the cluster environment before you can rejoin additional nodes to the newly restored cluster.

NOTE: When the cluster is at the stage of **Ready To Reprotect**, ensure that you first rebuild and restore the failback target so that it gets discovered as a cluster. You must manually uninstall the PlateSpin Cluster Driver as part of the rebuild process.

For information about rebuilding the cluster environment after a PlateSpin failover and failback occurs, see the following resources:

- **Windows Server 2012 R2 Failover Cluster (failback to physical or virtual rebuild):** See Knowledgebase Article 7016770 (http://www.netiq.com/support/kb/doc.php?id=7016770).

- **Windows Server 2008 R2 Failover Cluster (failback to physical or virtual rebuild):** See Knowledgebase Article 7015576 (http://www.netiq.com/support/kb/doc.php?id=7015576).

17 Generating Reports

You can generate reports about discovered workloads and the workload protection contracts by using PlateSpin Web Interface. For information about generating a Licensing Report, see Section 4.6, "Generating a Licensing Report for Technical Support," on page 50.

17.1 About Protect Reports

PlateSpin Protect enables you to generate the following reports that provide analytical insight into your workload protection contracts over time:

- **Workload Protection:** Reports replication events for all workloads over a selectable time window.
- **Replication History:** Reports replication type, size, time, and transfer speed per selectable workload over a selectable time window.
- **Replication Window:** Reports the dynamics of full and incremental replications that can be summarized by **Average**, **Most Recent**, **Sum**, and **Peak** perspectives.
- **Current Protection Status:** Reports **Target RPO**, **Actual RPO**, **Actual TTO**, **Actual RTO**, **Last Test Failover**, **Last Replication**, and **Test Age** statistics.
- **Events:** Reports system events for all workloads over a selectable time window.
- **Scheduled Events:** Reports only upcoming workload protection events.

Figure 17-1   Options for a Replication History Report

17.2 Generating Workload and Workload Protection Reports

To generate a report:

1 In your Web Interface, click **Reports**.

A list of the report types is displayed.

2 Click the name of the required report type.

3 Select one or more workloads for which you want to create the report.

4 Configure the time period for which you want to view the report.

5 Specify the appropriate parameters for the report.

6 Do one of the following:

- Click **Printable View** to view the report in your web browser.
- Click **Export to XML**, then save the XML file to your computer.

17.3 Generating Diagnostic Reports

In the PlateSpin Protect Web Interface, after you have executed a command, you can generate detailed diagnostic reports about the command's details.

1 Click **Command Details**, then click the **Generate** link in the lower right of the panel.

After a few moments, the page refreshes and displays a **Download** link above the **Generate** link.

2 Click **Download**.

A `.zip` file contains the comprehensive diagnostic information about the current command.

3 Save the file, then extract the diagnostics to view them.

4 Have the `.zip` file ready if you need to contact Technical Support.

18 Troubleshooting Workload Protection and Recovery

This section can help you troubleshoot common problems during the workload protection and recovery.

For issues for discovery and inventory for source workloads and target hosts, see Chapter 14, "Troubleshooting Workload Discovery and Inventory," on page 121.

18.1 Optimizing Throughput for a Connection

If you experience slow throughput, you can test the connection to see if there are any connection or bandwidth issues, and resolve them. See Appendix F, "Using the iPerf Network Test Tool to Optimize Network Throughput for PlateSpin Products," on page 197.

18.2 Troubleshooting Traffic-Forwarding Workloads

In some scenarios, the replica of a workload that is forwarding network traffic (for example, if the workload's purpose is to serve as a network bridge for NAT, VPN, or a firewall) might show significant network performance degradation. This is related to a problem with VMXNET 2 and VMXNET 3 adapters that have LRO (large receive offload) enabled.

To work around this issue, you need to disable LRO on the virtual network adapter. For more information, see Knowledgebase Article 7005495 (https://www.netiq.com/support/kb/doc.php?id=7005495).

18.3 Troubleshooting the Configuration Service

After Test Failover or Failover, an error occurs on the target VM because of non-specific Configuration Service issues. The common error message is:

```
Configuration service in the target machine does not seem to have started
```

Troubleshooting tips in this section explain common Configuration Service issues and some alternative ways to resolve them.

18.3.1 Understanding What Is Causing the Problem

The Configuration Service error indicates that the PlateSpin Server is unable to communicate with the Configuration Service on the Target VM. Analyze your system to determine the possible root cause of the problem.

Target VM Fails to Boot

The operating system must be loaded in the target VM in order for the Configuration Service to start up normally. A failure to boot indicates that there could be a driver conflict, a boot loader error, or possible disk corruption.

We recommend that you open a service ticket with Micro Focus Customer Care if the operating system fails to boot on the target VM.

Network Is Not Set Up Correctly

The network must be set up correctly in order for the Configuration Service on the target workload to communicate with the PlateSpin Server.

Ensure that you have configured your network in a way that the target workload can communicate with the PlateSpin Server. See Section 1.5, "Access and Communication Requirements across Your Protection Network," on page 29.

Unable to Read or Write Status Messages to Floppy Devices

The Configuration Service must be able to communicate with the floppy devices for VMware VMs in order to read and write status messages for the PlateSpin Server.

On the target VM, verify that the machine is able to communicate with the floppy devices:

1 On the VM, open the log file (`C:\windows\platespin\configuration\data\log.txt`).

2 Any of the following messages might be an indication that the floppy is inaccessible:

```
Failed (5) to write to file \\?\Volume{<guid-number>}\log.zip

CopyFile \\?\Volume{<guid-number>}\windows\platespin\configuration\data\result.txt
to \\?\Volume{<guid-number>}\result.txt failed

The output floppy was not accessible after the timeout period
```

18.3.2 What Can Be Done to Resolve the Problem

To resolve a Configuration Service error, you can try any of the solutions in this section.

- "Skip the Target VM Reboot Optimizations" on page 181
- "Reduce the Read/Write Traffic to Floppy Devices" on page 181
- "Change the Startup Type to Increase the Delay" on page 182
- "Configure Conflicting Services to Not Run Automatically at Startup" on page 183

Skip the Target VM Reboot Optimizations

Protect tries to minimize the number of reboots that occur on the target VM by default in order to speed up the Failover process. It is possible that allowing the additional reboots will improve the target VM's ability to communicate with the PlateSpin Server.

To skip reboot optimizations:

1 Log in to the PlateSpin Server, then open the PlateSpin Server Configuration page at:

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

2 Search for the parameter **ConfigurationServiceValues**.

3 Edit the **ConfigurationServiceValues** parameter and set the **SkipRebootOptimization** option to `true`.

4 Click **Save**.

5 Run an incremental or full replication.

   The replication also propagates the modified configuration settings to the target VM.

6 Run the Test Failover or Failover again for affected workloads.

Reduce the Read/Write Traffic to Floppy Devices

You can decrease the number of times the PlateSpin Server attempts to read from and write to the VMware input or output floppy devices if the diagnostic log shows the following error:

`Information:1:Attempting floppy download`

followed by

`Verbose:1:Failed to copy file from remote URL`

-or-

`Exception: The remote server returned an error: (500) Internal Server Error`

This error is caused by VMware locking the resource. It indicates that the PlateSpin Server is detaching and reattaching the floppy each time it checks the status. Locking can cause the target VM to fail to read and write to the floppy device. See *Using the VMware vCenter Server 4.x,5.x and 6.0 Datastore Browser to Download or Copy a Powered-On Virtual Machine's .vmx and .nvram Files Fails (1019286)* (https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1019286).

If you experience floppy device locking issues, you can increase values for the Configuration Service polling settings on the PlateSpin Server:

vmwareConfigServicePollStartDelay

> This parameter determines how long to wait before the PlateSpin Server starts polling for target workload status. The default value is 120 seconds (2 minutes).

vmwareConfigServicePollIntervalInMilliseconds

> This parameter determines how frequently the PlateSpin Server attempts to communicate with the target workload and to read or write to the VMware floppy devices. The poll interval default is 30000 ms (30 seconds).

vmwareConfigServicePollStartTimeout

> This parameter determines how long the PlateSpin Server waits after it starts the target VM before it displays an error in the Web Interface. The default value is 420 seconds (7 minutes).

vmwareConfigServicePollUpdateTimeout

> This parameter determines how long the PlateSpin Server waits after each polling interval before displaying an error in the Web Interface. The default value is 300 seconds (5 minutes).

Higher values for these parameters reduce the frequency that the PlateSpin Server attempts to read from and write to the VMware floppy devices on target VMs.

To reduce read and write traffic for VMware floppy devices:

1 Log in to the PlateSpin Server, then open the PlateSpin Server Configuration page at:

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

2 Search for the Configuration Service polling parameters, modify their settings as appropriate, then click **Save**.

For example:

```
vmwareConfigServicePollStartDelay = 180 (3 minutes)
vmwareConfigServicePollIntervalInMilliseconds = 300000 (5 minutes)
vmwareConfigServicePollStartTimeout = 1200 (20 minutes)
vmwareConfigServicePollUpdateTimeout = 900 (15 minutes)
```

or

```
vmwareConfigServicePollStartDelay = 300 (5 minutes)
vmwareConfigServicePollIntervalInMilliseconds = 480000 (8 minutes)
vmwareConfigServicePollStartTimeout = 1200 (20 minutes)
vmwareConfigServicePollUpdateTimeout = 900 (15 minutes)
```

3 Run an incremental or full replication.

The replication also propagates the modified configuration settings to the target VM.

4 Run the Test Failover or Failover again for affected workloads.

Change the Startup Type to Increase the Delay

The Configuration Service might be coming up before resources are accessible. You can change the Configuration Service startup type to have increase the delay.

To change the startup type:

1 Log in to the PlateSpin Server, then open the PlateSpin Server Configuration page at:

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

2 Search for the parameter **windowsConfigServiceStartType**.

3 Change the **windowsConfigServiceStartType** value to **AutoDelay**.

Options for **windowsConfigServiceStartType** are:

- **GroupDelay** is the default value and adds the Configuration Service to the end of the **ServiceGroupOrder** in the registry.

- **AutoDelay** will maximize the amount of time the service waits before starting (2 minutes after boot). Also modify the **ServicesPipeTimeoutForWindowsConfigService** parameter value in Step 4.

- **NoDelay** is the most efficient option and starts the service as soon as Windows can. However, it is not recommended because of the potential issues connecting to resources.

4 (AutoDelay) Change the **ServicesPipeTimeoutForWindowsConfigService** parameter setting to 180 seconds to account for the 120 seconds that the service will take to start up after boot when AutoDelay is set for **windowsConfigServiceStartType** in Step 3.

5 Click **Save**.

6 Run an incremental or full replication.

The replication also propagates the modified configuration settings to the target VM.

7 Run the Test Failover or Failover again for affected workloads.

Configure Conflicting Services to Not Run Automatically at Startup

During a failover action, a Windows service interferes with the mounting of floppy drivers.

Determine which Windows Services are configured to start up at reboot. Some services are known to interfere with the Configuration Service writing to a floppy, such as Wireless Configuration and some antivirus software. You should configure these services to not run automatically on Test Failover or Failover, then run the Test Failover or Failover again.

You can also try to disable all non-essential services for Test Failover and Failover on the Configuration page, then run the Test Failover or Failover again.

18.3.3 Additional Troubleshooting Tips

If the Configuration Service cannot contact the PlateSpin Server, diagnostics will tell only part of the picture. You must also get logs from the target VM:

- **Windows workloads:** The Configuration Service logs are found in the `C:\windows\platespin\configuration\data` folder.

  - The `log.txt` file contains all of the logging information, but the `Config.ini` file is useful in understanding what is to be configured.

  - The `result.txt` file contains the status of the Configuration Service run.

  - If the target VM cannot read from the input floppy device, it will not have the merged `Config.ini` file, which might include custom network configuration information for the test failover network environment.

  - If the `Config.ini` file has no network related information (such as a `[NIC0]`, the target VM network adapter might have special characters in the name.

    It is a known issue that the `Config.ini` file might not be accurate until it is merged with the one from the floppy device.

  - The target VM tries a reboot if it cannot connect to either the output floppy or input floppy (one time only). You will see a `config.ini.floppyreboot` file if this is the case.

- **Linux workloads:** The Configuration Service logs are found in the `/tmp` folder.
  - The main log files are named `file*.platespin.fileLogger`.

    We recommend examining any configuration folders in `/tmp`. Tar the configuration folders along with the `file*.platespin.fileLogger` files to send to Micro Focus Customer Care.
  - Other config files to check for include the following:

    ```
    /tmp/Ofx.RunCommand.Output*
    /tmp/*DiskHelper*
    /tmp/*VmTools*
    ```
  - The configuration file is `/usr/lib/psconfigservice/data/config.conf`.
  - The end result log file is `/usr/lib/psconfigservice/data/result.txt`.

18.4 Troubleshooting Workload Prepare Replication (Windows)

Problems or Messages	Solutions
Authentication error when verifying the controller connection while setting up the controller on the source.	The account used to add a workload needs to be allowed by this policy. See "Group Policy and User Rights" on page 184.
Failure to determine whether .NET Framework is installed (with exception `The trust relationship between this workstation and the primary domain failed`).	Check whether the Remote Registry service on the source is enabled and started. See also "Troubleshooting Discovery for Windows Workloads" on page 121.

18.4.1 Group Policy and User Rights

Because of the way that PlateSpin Protect interacts with the source workload's operating system, it requires the administrator account that is used to add a workload to have certain user rights on the source machine. In most instances, these settings are defaults of group policy; however, if the environment has been locked down, the following user rights assignments might have been removed:

- Bypass Traverse Checking
- Replace Process Level Token
- Act as part of the Operating System

In order to verify that these Group Policy settings have been set, you can run `gpresult /v` from the command line on the source machine, or alternately `RSOP.msc`. If the policy has not been set, or has been disabled, it can be enabled through either the Local Security Policy of the machine or through any of the Domain Group Policies being applied to the machine.

You can refresh the policy immediately by using `gpupdate /force`.

18.4.2 Two or More Volumes Have the Same Volume Serial Number

Issue: When you attempt to set up a protection for a Windows server, the following error is displayed:

```
[Source] Two or more volumes have the same serial number. Change the serial numbers
so that they are unique and rediscover the machine.
```

Workaround: This problem can occur if the Volume Serial Numbers for two or more volumes are the same. PlateSpin Protect requires the serial numbers to be unique.

To resolve this issue, modify the serial numbers for the data volumes as appropriate, and then rediscover the machine. For information about how to use Windows native tools to modify the serial numbers, see KB Article 7921101.

18.5 Troubleshooting Workload Replication

Problems or Messages	Solutions
Recoverable error during replication either during **Scheduling Taking Snapshot of Virtual Machine** or **Scheduling Reverting Virtual Machine to Snapshot before Starting**.	This problem occurs when the server is under load and the process is taking longer than expected. Wait until the replication is complete.
Incremental file-based replication does not complete with Encryption enabled	After you enable Encryption for a Windows workload that is configured for file-based data transfer, the Windows receiver might hang at the end of the transfer for incremental replications. The hang occurs if the last byte read of the transfer is incorrectly set by the encryption process to a non-zero value, indicating that more files are being transferred and to continue reading from the stream. You can use block-based data transfer for Windows workloads if you want to enable Encryption for replication data transfers.
Workload issue requires user intervention	Several types of issues might cause this message. In most cases the message should contain further specifics about the nature of the problem and the problem area (such as connectivity, credentials,. After troubleshooting, wait for a few minutes. If the message persists, contact PlateSpin Support.
All workloads go into recoverable errors because you are out of disk space.	Verify the free space. If more space is required, remove a workload.
Protection over a WAN takes a long time If the VM container has a large number of datastores	Under some circumstances the process of locating the appropriate ISO image required for booting the target might take longer than expected. This might happen when your PlateSpin Server is connected to the VM container over a WAN and your VM container has a large number of datastores.
Slow network speeds under 1 MB.	Confirm that the source machine's network interface card's duplex setting is on and the switch it is connected to has a matching setting. That is, if the switch is set to auto, the source can't be set to 100 MB.
Slow network speeds over 1 MB.	Measure the latency by running the following command from the source workload: `ping ip-t` (replace *ip* with the IP address of your PlateSpin Server host). Allow it to run for 50 iterations and the average indicates the latency. Also see "Optimizing Data Transfer over WAN Connections" on page 66.

Problems or Messages	Solutions

`The file transfer cannot begin - port 3725 is already in use`
 	
or
 	
`3725 unable to connect`
	Ensure that the port is open and listening:
	
	Run `netstat -ano` on the workload.
	
	Check the firewall.
	
	Retry the replication.

`Controller connection not established` Replication fails at the **Take Control of Virtual Machine** step.	This error occurs when the replication networking information is invalid. Either the DHCP server is not available or the replication virtual network is not routable to the PlateSpin Server host. Change the replication IP to a static IP or enable the DHCP server. Ensure that the virtual network selected for replication is routable to the PlateSpin Server host.
Replication job does not start (stuck at 0%)	This error can occur for different reasons and each has a unique solution: ◆ For environments using a local proxy with authentication, bypass the proxy or add proper permissions to resolve this problem. See Knowledgebase Article 7920339 (https://www.netiq.com/support/kb/doc.php?id=7920339). ◆ If local or domain policies restrict required permissions, follow the steps outlined in Knowledgebase Article 7920862 (https://www.netiq.com/support/kb/doc.php?id=7920862). This is a common issue when PlateSpin Server host is affiliated with a domain and the domain policies are applied with restrictions. See "Group Policy and User Rights" on page 184.
After a Windows Update, some files in the `C:\Windows\SoftwareDistribution` folder are not transferred to the target machine during incremental file-based replication.	This is a Microsoft Windows common practice: For optimization purposes, some files are marked for deletion in the `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\BackupRestore\FilesNotToSnapshot` registry key to prevent them from being included in VSS snapshots. See the Microsoft Developer Network article, *Excluding Files from Shadow Copies* (http://msdn.microsoft.com/en-us/library/aa819132.aspx) for more information. Generally, these files are used to install Windows updates before they are deleted and are no longer necessary after the update. If you choose to restore these files, run Windows Update on the target machine after failover to repopulate the `SoftwareDistribution` folder.

18.6 Troubleshooting Workload Failover or Failback

Problems or Messages	Solutions
Active Directory Domain Services are not available after a failback (Windows)	Active Directory Domain Services might not come up after a Failover, if `chkdsk` errors occur. Two avoidable causes of `chkdsk` errors are: ◆ Log files related to Microsoft Updates if the source machine is not up-to-date with all Microsoft recommended patches or updates when you perform the first full replication. ◆ System files and folders that should be excluded from your antivirus software. To avoid these issues, follow these best practices described in Section 15.1, "Prerequisites for Workload Protection," on page 153 before you run the first full replication.
During failback, the wrong NICs are mapped and failback hangs	Use one of the following workarounds to allow the failback to complete successfully: ◆ Switch the IP configuration to the expected mappings so that the target is successfully configured. ◆ Reboot the 'takecontrol' hardware to the LRD, then repeat the steps to use it as the failback target. There is a good chance that Protect will map to the correct Ethernet interfaces the next time. ◆ In the Web Interface, if the failback seems to hang near the end of completion, the failback target likely cannot communicate with the PlateSpin Server that the failback is complete. Switch the network cables in the back of the failback target to place the correct NIC on the intended networks. This enables the failback target to communicate with the PlateSpin Server, and the failback completes.
X2P Failback of Linux Workloads Causes Failure of the X-Server Graphical Interface	The issue is caused by a reconfiguration of the failed-over VM when VMware tools are installed. To correct this, use the following command to find the files with the string `BeforeVMwareToolsInstall` in the filename: `find / -iname '*BeforeVMwareToolsInstall'` After you identify all such files, move them back to their original locations, then reboot the workload to fix the workload's X Server interface.

18.7 Shrinking the PlateSpin Protect Databases

When the PlateSpin Protect databases (`OFX`, `PortabilitySuite,` and `Protection`) reach a predetermined capacity, cleanup on those databases occurs at regular intervals. If there is a need to further regulate the size or content of those databases, Protect provides a utility (`PlateSpin.DBCleanup.exe`) to further clean up and shrink those databases. Knowledgebase Article 7006458 (https://www.netiq.com/support/kb/doc.php?id=7006458) explains the location of the tool and the options available for it, should you decide to use it for offline database operations.

18.8 Post-Protection Workload Cleanup

Use these steps to clean up your source workload from all PlateSpin software components when required, such as following an unsuccessful or problematic protection.

18.8.1 Cleaning Up Windows Workloads

Component	Removal Instructions
PlateSpin Block-Based Transfer Component	See Knowledgebase Article 7005616 (https://www.netiq.com/support/kb/doc.php?id=7005616).

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

File-based Transfer Component	At root level for each volume under protection, remove all files named `PlateSpinCatalog*.dat`.
Workload Inventory software	In the workload's `Windows` directory: ◆ Remove all files named `machinediscovery*`. ◆ Remove the subdirectory named `platespin`.

Controller software	1. Open a command prompt on the source workload and change the current directory to:
	
	    ◆ `\Program Files\platespin*` (32-bit systems)
	
	    ◆ `\Program Files (x86)\platespin*` (64-bit systems)
	
	2. Run the following command:
	
	    `ofxcontroller.exe /uninstall`
	
	3. Remove the `platespin*` directory.

18.8.2 Cleaning Up Linux Workloads

Component	Removal Instructions
Controller software	- Kill these processes: `pkill -9 ofxcontrollerd` `pkill -9 ofxjobexec` - remove the OFX controller rpm package: `rpm -e ofxcontrollerd` - In the workload's file system, remove the `/usr/lib/ofx` directory with its contents.

Block-level data transfer software	1. Check if the driver is active:
	`lsmod | grep blkwatch`
	If the driver is still loaded in memory, the result should contain a line, similar to the following:
	`blkwatch_7616 70924 0`
	2. (Conditional) If the driver is still loaded, remove it from memory:
	`rmmod blkwatch_7616`
	3. Remove the driver from the boot sequence:
	`blkconfig -u`
	4. Remove the driver files by deleting the following directory with its contents:
	`/lib/modules/[Kernel_Version]/Platespin`
	5. Delete the following file:
	`/etc/blkwatch.conf`

LVM snapshots	LVP snapshots used by ongoing replications are named according to a `volume_name-PS-snapshot` convention. For example, a snapshot of a `LogVol01` volume will be named `LogVol01-PS-snapshot`. To remove these LVM snapshots: 1. Generate a list of snapshot on the required workload by using one of the following ways: - Use the Web Interface to generate a Job Report for the failed job. The report should contain information about LVM snapshots and their names. - OR - - On the required Linux workload, run the following command to display a list of all volumes and snapshots: `# lvdisplay -a` 2. Note the names and locations of the snapshots you want to remove. 3. Remove the snapshots by using the following command: `lvremove snapshot_name`

Component	Removal Instructions

NSS snapshot	NSS snapshot created and used by PlateSpin for ongoing replications. The snapshot name ends with the suffix `PSSNP`.
	
	To remove these NSS snapshots:
	
	1. Generate a list of snapshots on the required workload by using one of the following methods:
	
	   ◆ Use the Web Interface to generate a Job Report for the failed job. The report should contain information about NSS snapshots and their names.
	
	   - OR -
	
	   ◆ On the required Open Enterprise Server workload, enter the following command to display a list of all NSS snapshots:
	
	   `# NLVM list snaps`
	
	   - OR -
	
	   ◆ On the required Open Enterprise Server workload, launch NSSMU and select **Snapshot** to view a list of snapshots.
	
	2. Note the names and locations of the snapshots you want to remove.
	
	3. On the Open Enterprise Server workload, remove the appropriate snapshots by using one of the following methods:
	
	   ◆ Enter the following command:
	
	   `NLVM delete snap` *<snapshot_name>*
	
	   - OR -
	
	   ◆ Launch NSSMU, then select **Snapshot**. For each snapshot you want to delete, highlight the snapshot then press Delete.

Bitmap files	For each volume under protection, at the root of the volume, remove the corresponding `.blocks_bitmap` file.
Tools	On the source workload, under `/sbin`, remove the following files: ◆ `bmaputil` ◆ `blkconfig`

V PlateSpin Tools

PlateSpin Protect provides additional tools to enhance your protection environment.

- Appendix E, "Using Workload Protection Features through the PlateSpin Protect Server API," on page 193
- Appendix F, "Using the iPerf Network Test Tool to Optimize Network Throughput for PlateSpin Products," on page 197

E Using Workload Protection Features through the PlateSpin Protect Server API

You can use workload protection features of PlateSpin Protect programmatically through the PlateSpin Protect Server API (`protectionservices`) from within your applications. You can use any programming or scripting language that supports an HTTP client and JSON serialization framework.

NOTE: The Protect Server API is experimental. Information in this section is provided as a technology preview.

E.1 API Overview

PlateSpin Protect exposes a REST-based API technology preview that developers can use as they build their own applications to work with the product. The API includes information about the following operations:

* discover containers
* discover workloads
* configure protection
* run replications, failover operations and failback
* query for status of workload and container status
* query for status of running operations
* query security groups and their protection ties

E.2 PlateSpin Protect Server API Documentation

The PlateSpin Protect Server API home page for `protectionservices` provides documentation and samples that can be useful for developers and administrators. For information, go to the following location on your PlateSpin Server host:

`https://`*Your_PlateSpin_Server*`/protectionservices`

Replace *Your_PlateSpin_Server* with the host name or the IP address of your PlateSpin Server host. If SSL is not enabled, use `http` in the URI.

Figure E-1   *The Home Page of the Protect Server API*

PlateSpin Protect Server API

Version 11.2.0.81

Documentation

Getting started

- Getting started with API
- Security and authentication
- Developer Guidelines
- Troubleshooting
- FAQ

How to

- Steps to protect workload
- Working with workload
- Working with container
- Working with security groups
- Working with protection tiers
- Adding multiple workloads and containers
- Limitations of the API
- Samples
- Glossary

REST Resources (auto-generated)

- Containers
- Workloads
- Configuration
- Operations
- Protection Tiers
- Security Groups

Resource representations

This section specifies the representations of the resources which this API operates on. The representations are made up of fields, each with a name and value, encoded using a JSON dictionary. The values may be numeric or string literals, lists, or dictionaries, each of which are represented in the obvious way in JSON. These representations typically nest. For example, the representation of a Containers will include representations of the Container which inhabit it, which in turn include representations of the Virtual Machine. Many of the models specify that the representation includes a uri field whose value is the URI of the resource being represented. This is present to support URI discovery in nested representations.

E.3   Samples and Other References

Protect administrators can leverage a JScript sample from the command line to access the product through the API. On the PlateSpin Server host, see the sample at

https://localhost/protectionservices/Documentation/Samples/protect.js

The sample can help you write scripts to help you work with the product. Using the command line utility, you can perform the following operations:

- Add a single workload
- Add a single container
- Run the replication, failover, and failback operations
- Add multiple workloads and containers at one time

> **NOTE:** For more information about this operation, see the API documentation at
>
> https://localhost/protectionservices/Documentation/
> AddWorkloadsAndContainersFromCsvFile.htm

- remove all workloads at one time
- remove all container at one time

To script common workload protection operations, use the referenced samples written in Python as guidance. A Microsoft Silverlight application, along with its source code, is also provided for reference purposes.

F Using the iPerf Network Test Tool to Optimize Network Throughput for PlateSpin Products

Before you execute replication, ensure that you test the connection to see if there are any connection or bandwidth issues, and resolve them. This section describes how to use the open source iPerf Network Test tool to optimize throughput on the connection.

F.1 Introduction

In the interest of helping PlateSpin administrators in their efforts to achieve better network throughput when using PlateSpin products, the iPerf Network Test tool is provided on the PlateSpin LRD (Linux RAM Disk) take-control environment. As stated in the iPerf documentation: "The primary goal of iPerf is to help in tuning TCP connections over a particular path. The most fundamental tuning issue for TCP is the TCP window size, which controls how much data can be in the network at any one point."

The purpose of this README is to describe a basic method for network tuning and testing as it relates to using PlateSpin products. First, you calculate a theoretical optimum TCP window size. Then you use the iPerf tool to validate and fine-tune this calculated size and measure the resulting throughput. Using this method is also useful in determining the real achievable throughput for a given network.

Both the iPerf tool and PlateSpin products are actually using the *TCP send/receive buffer size* in order to affect the eventual internal choice of *TCP window size.* Going forward, these terms will be used interchangeably.

NOTE: There are many factors that affect network throughput. A wealth of information is available on the Internet that can aid in understanding. One such resource is the *Network Throughput Calculator* (http://wintelguy.com/wanperf.pl), which can help in calculating the expected maximum TCP throughput given applicable customer network characteristics. We strongly recommend that this online calculator be used in order to correctly set expectations regarding throughput.

F.2 Calculations

Tuning of the TCP window size is based on a number of factors, including network link speed and network latency. For our purposes relating to PlateSpin products, the initial choice of TCP window size for tuning is based on standard calculations (widely available on the Internet and elsewhere) as follows:

```
WinSizeInBytes=((LINK_SPEED(Mbps)/8)*DELAY(sec))*1000*1024
```

For example, for a 54 Mbps link with 150 ms latency, the proper initial window size would be:

```
(54/8)*0.15*1000*1024 = 1,036,800 bytes
```

For a 1000 Mbps link with 10 ms latency, the proper initial window size would be:

```
(1000/8)*.01*1000*1024 = 1,280,000 bytes
```

In order to get a latency value for the network, use `ping` from the command prompt (Windows) or the terminal (Linux). Although the `ping` round-trip time (RTT) is arguably different than the actual latency, the value obtained is sufficiently close for use in this method.

The following is a sample output from a Windows `ping` command, where the latency is observed to be 164 ms on average:

ping 10.10.10.232 -n 5

```
Pinging 10.10.10.232 with 32 bytes of data:
Reply from 10.10.10.232: bytes=32 time=154ms TTL=61
Reply from 10.10.10.232: bytes=32 time=157ms TTL=61
Reply from 10.10.10.232: bytes=32 time=204ms TTL=61
Reply from 10.10.10.232: bytes=32 time=153ms TTL=61
Reply from 10.10.10.232: bytes=32 time=153ms TTL=61

Ping statistics for 10.10.10.232:
    Packets: Sent = 5, Received = 5, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 153ms, Maximum = 204ms, Average = 164ms
```

The following is a sample output from a Linux `ping` command, where the latency is observed to be 319 ms on average:

ping 10.10.10.232 -c 5

```
PING 10.10.10.232 (10.10.10.232) 56(84) bytes of data.
64 bytes from 10.10.10.232: icmp_seq=1 ttl=62 time=0.328 ms
64 bytes from 10.10.10.232: icmp_seq=2 ttl=62 time=0.280 ms
64 bytes from 10.10.10.232: icmp_seq=3 ttl=62 time=0.322 ms
64 bytes from 10.10.10.232: icmp_seq=4 ttl=62 time=0.349 ms
64 bytes from 10.10.10.232: icmp_seq=5 ttl=62 time=0.316 ms

--- 10.10.10.232 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 3998ms
rtt min/avg/max/mdev = 0.280/0.319/0.349/0.022 ms
```

In practice, you should use the `-n` or `-c` option to specify a larger number of ping packets in order to more closely measure the latency value.

F.3 Setup

The iPerf tool runs in either server mode or client mode.

The basic usage syntax for `iperf` server mode is:

```
iperf -s -w <win_size>
```

The basic usage syntax for `iperf` client mode is:

```
iperf -c <server_ip> -w <win_size>
```

Our intent is to measure and tune the network between a source and target workload. In many cases, these can be the actual source and targets in use. It is possible to complete the testing using a different workload for either source or target, provided that the substitute has the same network characteristics as the original, such as NIC, network connection, and so on.

NOTE: Ensure that you are not testing the throughput from the PlateSpin server to either the source or the target, as this traffic is minimal, and does not represent the traffic that occurs during a migration or replication.

While it is possible to use a live workload (either Windows or Linux) as the target/iperf server, the following steps provide the environment most similar to what happens at migration/replication time, and is strongly recommended.

To set up and run `iperf` on the target:

  1 Boot the target using the LRD.

  2 In the LRD console, use the helper terminal (accessible via Alt-F2) to do the following:

   2a Set up networking using option 5.

   2b Mount the CD media using option 6.

  3 In the LRD console, switch to the debug terminal (accessible via Alt-F7) to go to the location of the iPerf tool:

```
cd /mnt/cdrom/LRDTools/iperf_2.0.X/linux
```

  4 Run the iPerf tool in server mode. Enter

```
./iperf -s -w <win_size>
```

To set up and run `iperf` on the source:

  1 Mount the LRD ISO by using software or physical media.

  2 Open a command prompt (Windows) or terminal (Linux) and go to the location of the iPerf tool:

```
cd <media>/LRDTools/iperf_2.0.X/
```

  3 As determined by the source operating system, go to the `windows` or `linux` subdirectory:

```
cd windows
```

```
-OR-
```

```
cd linux
```

  4 Run the iPerf tool in client mode. Enter

```
iperf -c <target_ip> -w <win_size>
```

NOTE: You can download and use `iperf3` for the calculations, which is helpful in certain scenarios where `iperf2` is unable to generate useful throughput numbers. Although the command syntax and output from `iperf3` differs slightly, it should be fairly straightforward to adapt and interpret the newer output, if necessary.

F.4 Methodology

Starting with the initial `win_size` calculated in the Calculations section, record the output from several iterations of the iPerf tool using the calculated value as well as slightly larger and smaller values. We recommend that you increase and decrease the `win_size` by increments of about 10 percent of the original value.

For example, for the example of 1,280,000 bytes above, you might increase or decrease `win_size` in increments of about 100,000 bytes.

NOTE: The `-w` option of `iperf` allows specifying units such as K (kilobytes) or M (megabytes).

Using the same example, you could use `-w` values of `1.28M`, `1.38M`, `1.18M`, and so on as the `win_size` in Step 4. Of course, it is assumed that only the run step is repeated for each iteration of the iPerf tool.

Sample output from an iperf client iteration looks similar to the following:

```
iperf.exe -c 10.10.10.232 -w 1.1M


------------------------------------------------------------
Client connecting to 10.10.10.232, TCP port 5001
TCP window size: 1.10 MByte
------------------------------------------------------------
[296] local 10.10.10.224 port 64667 connected with 10.10.10.232 port 5001
[ ID] Interval        Transfer      Bandwidth
[296]  0.0-10.2 sec  11.3 MBytes  9.29 Mbits/sec
```

Sample output from the referenced target server looks similar to the following:

```
./iperf -s -w .6M


------------------------------------------------------------
Server listening on TCP port 5001
TCP window size: 1.20 MByte (WARNING: requested 614 Kbyte)
------------------------------------------------------------
[  4] local 10.10.10.232 port 5001 connected with 10.10.10.224 port 64667
[  4] 0.0-10.2 sec  11.3 MBytes  9.29 Mbits/sec
```

NOTE:

  ◆ The client disconnects from the server after a single iteration, while the server continues to listen until stopped by using Ctrl-C.

  ◆ The window size specified for a Linux server is 1/2 the desired value, because Linux doubles the requested TCP buffer size as a matter of course.

Use several iterations to determine the optimal value for the TCP window size. Remember to use only 1/2 the desired value when specifying the `-w` option for `iperf` on Linux.

Increased throughput indicates that you are getting closer to an optimal TCP window size. Finally, as you get closer to an optimal value, use longer iterations in order to more closely simulate real running conditions. To achieve a longer iteration, use the `-t <time_in_seconds>` option to `iperf`. This option needs to be specified only on the client side.

For example:

```
iperf.exe -c 10.10.10.232 -w 1.25M -t 60
```

After an optimal value has been determined, configure this value in the `FileTransferSendReceiveBufferSize` parameter for the appropriate PlateSpin server at:

https://<*my_ps_server*>/PlatespinConfiguration/

This global value applies to all workloads on the PlateSpin server, so care should be taken to group workloads and their respective networks in a sensible manner across available PlateSpin servers.

F.5   Expectations

Modifying the TCP window size indirectly with the TCP send/receive buffer size can be a very effective method for increasing network throughput in some scenarios. Two to three or even more times the original throughput can sometimes be achieved. However, it is important to remember that the network characteristics can (and often do) change over time because of changes in usage patterns, hardware, software, or other infrastructure.

We strongly recommend that you use this method to calculate the optimum value at the same time of day and under the same network usage patterns that are intended to be in use during the planned live migration or replication tasks. We also recommend that you recalculate the setting periodically in order to account for changing network conditions.

VI Documentation Updates

This guide has been updated since the General Availability of PlateSpin Protect 11.2.1.

- ◆ Appendix G, "Documentation Updates," on page 205

G Documentation Updates

This section contains information on documentation content changes that were made in the English version of the *PlateSpin Protect User Guide* since the General Availability of PlateSpin Protect 11.2.1.

- Section G.1, "February 2018," on page 205

G.1 February 2018

Location	Update
Table 1-2, "Supported Linux Workloads," on page 16	For Red Hat Enterprise Linux 6.8, Oracle Linux 6.8, and CentOS 6.8 workloads with LVM volumes, PlateSpin Migrate supports incremental replication only for the latest available kernel (version `2.6.32-696.20.1.el6.x86_64`) for the 6.8 distribution.