
PlateSpin® Protect 11.0

User Guide

February 2017

Legal Notice

THIS DOCUMENT AND THE SOFTWARE DESCRIBED IN THIS DOCUMENT ARE FURNISHED UNDER AND ARE SUBJECT TO THE TERMS OF A LICENSE AGREEMENT OR A NON-DISCLOSURE AGREEMENT. EXCEPT AS EXPRESSLY SET FORTH IN SUCH LICENSE AGREEMENT OR NON-DISCLOSURE AGREEMENT, NETIQ CORPORATION PROVIDES THIS DOCUMENT AND THE SOFTWARE DESCRIBED IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. SOME STATES DO NOT ALLOW DISCLAIMERS OF EXPRESS OR IMPLIED WARRANTIES IN CERTAIN TRANSACTIONS; THEREFORE, THIS STATEMENT MAY NOT APPLY TO YOU.

For purposes of clarity, any module, adapter or other similar material ("Module") is licensed under the terms and conditions of the End User License Agreement for the applicable version of the NetIQ product or software to which it relates or interoperates with, and by accessing, copying or using a Module you agree to be bound by such terms. If you do not agree to the terms of the End User License Agreement you are not authorized to use, access or copy a Module and you must destroy all copies of the Module and contact NetIQ for further instructions.

This document and the software described in this document may not be lent, sold, or given away without the prior written permission of NetIQ Corporation, except as otherwise permitted by law. Except as expressly set forth in such license agreement or non-disclosure agreement, no part of this document or the software described in this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, or otherwise, without the prior written consent of NetIQ Corporation. Some companies, names, and data in this document are used for illustration purposes and may not represent real companies, individuals, or data.

This document could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein. These changes may be incorporated in new editions of this document. NetIQ Corporation may make improvements in or changes to the software described in this document at any time.

U.S. Government Restricted Rights: If the software and documentation are being acquired by or on behalf of the U.S. Government or by a U.S. Government prime contractor or subcontractor (at any tier), in accordance with 48 C.F.R. 227.7202-4 (for Department of Defense (DOD) acquisitions) and 48 C.F.R. 2.101 and 12.212 (for non-DOD acquisitions), the government's rights in the software and documentation, including its rights to use, modify, reproduce, release, perform, display or disclose the software or documentation, will be subject in all respects to the commercial license rights and restrictions provided in the license agreement.

© 2015 NetIQ Corporation. All Rights Reserved.

For information about NetIQ trademarks, see <https://www.netiq.com/company/legal/>.

License Grant

Licenses for PlateSpin Protect 10.4 cannot be used for prior versions of PlateSpin Protect.

Third-Party Software

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

Contents

About NetIQ Corporation	7
About This Guide	9
1 Product Overview	11
1.1 About PlateSpin Protect	11
1.2 Supported Configurations	11
1.2.1 Supported Windows Workloads	11
1.2.2 Supported Linux Workloads	13
1.2.3 Supported VM Containers	14
1.2.4 Supported System Firmware	14
1.3 Security and Privacy	15
1.3.1 Security of Workload Data in Transmission	15
1.3.2 Security of Client/Server Communications	15
1.3.3 Security of Credentials	15
1.3.4 User Authorization and Authentication	15
1.4 Performance	15
1.4.1 About Product Performance Characteristics	16
1.4.2 Data Compression	16
1.4.3 Bandwidth Throttling	16
1.4.4 RPO, RTO, and TTO Specifications	17
1.4.5 Scalability	17
2 PlateSpin Protect Application Configuration	19
2.1 Product Licensing	19
2.1.1 Obtaining a License Activation Code	19
2.1.2 Online License Activation	19
2.1.3 Offline License Activation	20
2.2 Setting Up User Authorization and Authentication	20
2.2.1 About PlateSpin Protect User Authorization and Authentication	21
2.2.2 Managing PlateSpin Protect Access and Permissions	22
2.2.3 Managing PlateSpin Protect Security Groups and Workload Permissions	23
2.3 Access and Communication Requirements across your Protection Network	24
2.3.1 Access and Communication Requirements for Workloads	24
2.3.2 Access and Communication Requirements for Containers	25
2.3.3 Open Port Requirements for PlateSpin Server Hosts	26
2.3.4 Protection Across Public and Private Networks Through NAT	26
2.3.5 Overriding the Default bash Shell for Executing Commands on Linux Workloads	27
2.3.6 Requirements for VMware DRS Clusters as Containers	27
2.4 Configuring PlateSpin Protect Default Options	27
2.4.1 Setting Up Automatic E-Mail Notifications of Events and Reports	28
2.4.2 Language Setup for International Versions of PlateSpin Protect	31
2.4.3 Language Setup for International Versions of PlateSpin Protect	31
2.4.4 Configuring PlateSpin Server Behavior through XML Configuration Parameters	32
2.4.5 Configuring Support for VMware vCenter Site Recovery Manager	34
3 Up and Running	37
3.1 Launching the PlateSpin Protect Web Interface	37
3.2 Elements of the PlateSpin Protect Web Interface	38

3.2.1	Navigation Bar	39
3.2.2	Visual Summary Panel.	39
3.2.3	Tasks and Events Panel.	40
3.3	Workloads and Workload Commands	40
3.3.1	Workload Protection and Recovery Commands	41
3.4	Managing Multiple Instances of PlateSpin Protect and PlateSpin Forge	42
3.4.1	Using the PlateSpin Protect Management Console.	42
3.4.2	About PlateSpin Protect Management Console Cards	42
3.4.3	Adding Instances of PlateSpin Protect and PlateSpin Forge to the Management Console	43
3.4.4	Managing Cards on the Management Console	44
3.5	Generating Workload and Workload Protection Reports	45
4	Workload Protection	47
4.1	Basic Workflow for Workload Protection and Recovery	47
4.2	Adding Containers (Protection Targets)	48
4.3	Adding Workloads for Protection	50
4.4	Configuring Protection Details and Preparing the Replication	51
4.4.1	Workload Protection Details.	52
4.5	Starting the Workload Protection	54
4.6	Aborting Commands	55
4.7	Failover	56
4.7.1	Detecting Offline Workloads	56
4.7.2	Performing a Failover	57
4.7.3	Using the Test Failover Feature.	58
4.8	Failback	58
4.8.1	Automated Failback to a VM Platform	59
4.8.2	Semi-Automated Failback to a Physical Machine	61
4.8.3	Semi-Automated Failback to a Virtual Machine.	62
4.9	Reprotecting a Workload	63
5	Essentials of Workload Protection	65
5.1	Workload License Consumption.	65
5.2	Guidelines for Workload and Container Credentials	66
5.3	Setting Up Protect Multitenancy on VMware	66
5.3.1	Using Tools to Define VMware Roles	67
5.3.2	Assigning Roles In vCenter	68
5.4	Data Transfer	71
5.4.1	Transfer Methods.	72
5.4.2	Data Encryption	73
5.5	Protection Tiers	73
5.6	Recovery Points	74
5.7	Initial Replication Method (Full and Incremental)	74
5.8	Service and Daemon Control.	75
5.9	Using Freeze and Thaw Scripts for Every Replication (Linux)	76
5.10	Volumes	77
5.11	Networking	78
5.12	Failback to Physical Machines.	78
5.12.1	Downloading the PlateSpin Boot ISO Image	79
5.12.2	Injecting Additional Device Drivers into the Boot ISO Image.	79
5.12.3	Registering Physical Machines as Failback Targets with PlateSpin Protect	80
5.13	Advanced Workload Protection Topics.	81
5.13.1	Protecting Windows Clusters.	81

5.13.2	Using Workload Protection Features through the PlateSpin Protect Web Services API	83
6	Auxiliary Tools for Working with Physical Machines	87
6.1	Managing Device Drivers	87
6.1.1	Packaging Device Drivers for Windows Systems	87
6.1.2	Packaging Device Drivers for Linux Systems	88
6.1.3	Uploading Drivers to the PlateSpin Protect Device Driver Database	88
6.1.4	Using the Plug and Play (PnP) ID Translator Feature	90
7	Troubleshooting	97
7.1	Troubleshooting Workload Inventory (Windows)	97
7.1.1	Performing Connectivity Tests	98
7.1.2	Disabling Antivirus Software	99
7.1.3	Enabling File/Share Permissions and Access	100
7.2	Troubleshooting Workload Inventory (Linux)	101
7.3	Troubleshooting Problems during the Prepare Replication Command (Windows)	101
7.3.1	Group Policy and User Rights	101
7.4	Troubleshooting Workload Replication	102
7.5	Troubleshooting Traffic-forwarding Workloads	103
7.6	Troubleshooting Online Help	104
7.7	Generating and Viewing Diagnostic Reports	104
7.8	Removing Workloads	104
7.9	Post-Protection Workload Cleanup	105
7.9.1	Cleaning Up Windows Workloads	105
7.9.2	Cleaning Up Linux Workloads	106
7.10	Shrinking the PlateSpin Protect Databases	107
A	Linux Distributions Supported by Protect	109
A.1	Analyzing Your Linux Workload	109
A.1.1	Determining the Release String	109
A.1.2	Determining the Architecture	110
A.2	Pre-compiled "blkwatch" driver (Linux)	110
B	Synchronizing Serial Numbers on Cluster Node Local Storage	119
C	Documentation Updates	121
C.1	February 23, 2015	121
C.2	February 2, 2015	121
C.3	October 24, 2014	122
C.4	October 13, 2014	122
C.5	September 11, 2014	122
C.6	September 5, 2014	122
	Glossary	123

About NetIQ Corporation

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

Contacting Sales Support

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

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

Contacting Technical Support

For specific product issues, please contact our Technical Support team.

Worldwide: +1-416-203-4799
North America: +1-800-858-4000
Email: support@platespin.com
Web Site: www.netiq.com/support

Contacting Documentation Support

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

Contacting the Online User Community

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

About This Guide

This guide provides information about using PlateSpin Protect.

- ◆ Chapter 1, “Product Overview,” on page 11
- ◆ Chapter 2, “PlateSpin Protect Application Configuration,” on page 19
- ◆ Chapter 3, “Up and Running,” on page 37
- ◆ Chapter 4, “Workload Protection,” on page 47
- ◆ Chapter 5, “Essentials of Workload Protection,” on page 65
- ◆ Chapter 6, “Auxiliary Tools for Working with Physical Machines,” on page 87
- ◆ Chapter 7, “Troubleshooting,” on page 97
- ◆ Appendix A, “Linux Distributions Supported by Protect,” on page 109
- ◆ Appendix B, “Synchronizing Serial Numbers on Cluster Node Local Storage,” on page 119
- ◆ Appendix C, “Documentation Updates,” on page 121
- ◆ “Glossary” on page 123

Audience

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

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the **User Comments** feature at the top and bottom of each page of the online documentation.

Additional Documentation

This guide is part of the PlateSpin Protect documentation set. For a complete list of publications supporting this release, visit the product’s Online Documentation Web Site:

[PlateSpin Protect 11 online documentation \(https://www.netiq.com/documentation/platespin_protect_11/\)](https://www.netiq.com/documentation/platespin_protect_11/)

Documentation Updates

The most recent version of in this guide can be found at the [PlateSpin Protect 11 Online Documentation Web Site \(https://www.netiq.com/documentation/platespin_protect_11/\)](https://www.netiq.com/documentation/platespin_protect_11/):

1 Product Overview

This section includes the following information:

- ♦ [Section 1.1, “About PlateSpin Protect,” on page 11](#)
- ♦ [Section 1.2, “Supported Configurations,” on page 11](#)
- ♦ [Section 1.3, “Security and Privacy,” on page 15](#)
- ♦ [Section 1.4, “Performance,” on page 15](#)

1.1 About PlateSpin Protect

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 or to completely new infrastructures, physical or virtual
- ♦ Take advantage of existing external storage solutions, such as SANs

1.2 Supported Configurations

- ♦ [Section 1.2.1, “Supported Windows Workloads,” on page 11](#)
- ♦ [Section 1.2.2, “Supported Linux Workloads,” on page 13](#)
- ♦ [Section 1.2.3, “Supported VM Containers,” on page 14](#)
- ♦ [Section 1.2.4, “Supported System Firmware,” on page 14](#)

1.2.1 Supported Windows Workloads

PlateSpin Protect supports most Windows-based workloads.

Both file-level and block-level replications are supported, with certain restrictions. See [Section 5.4, “Data Transfer,” on page 71](#).

Table 1-1 Supported Windows Workloads

Operating System	Remarks
Servers	
Windows Server 2012 R2 Windows Server 2012	
Windows Server 2008 R2 (64-bit) Windows Server 2008 (64-bit) Windows Server 2008 latest SP (32-bit)	Including domain controllers (DC) and Small Business Server (SBS) editions
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 2008 R2 Server-based Microsoft Failover Cluster	
Desktops	
Windows 8.1 Windows 8	<p>WARNING: You must select the High Performance power plan on the Windows 8 source so that the workload failover and fallback function correctly.</p> <p>To configure this power plan from the Windows Control Panel:</p> <ol style="list-style-type: none"> 1. Select All Control Panel Items > Power Options. 2. In the Choose or customize power plan dialog box, select Show additional plans > High Performance. 3. Close the control panel.
<p>Windows Cluster Support: See “Protecting Windows Clusters” on page 81 for specific cluster configurations supported.</p> <p>Supported international versions: French, German, Japanese, Chinese Traditional, and Chinese Simplified.</p>	
<p>TIP: Other international versions have limited support: updating system files could be affected in languages other than those listed above.</p>	
<p>Workload Firmware (UEFI and BIOS) Support: PlateSpin Protect mirrors the Microsoft support of UEFI or BIOS-based Windows workloads. It transfers workloads (both Block and File transfers are supported) from source to target while enforcing the supported firmware for the respective source and target operating systems. It does the same for the fallback to a physical machine. When any transition (failover and fallback) between UEFI and BIOS systems are initiated, Protect analyzes the transition and alerts you about its validity.</p>	
<p>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 need to target a vSphere 5.0 container or newer.</p>	

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

- ◆ When transferring 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 – Windows 2003 does not support the UEFI boot mode).
- ◆ When protecting 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.

1.2.2 Supported Linux Workloads

PlateSpin Protect supports a number of Linux distributions.

Table 1-2 Supported Linux Workloads

Operating System	Remarks
Servers	
Red Hat Enterprise Linux (RHEL) 6.2	Block-based transfer only.
RHEL 5.0-5.5, 6.0	Block-based transfer only.
RHEL 5.6-5.8, 6.3	Block-based transfer only. You must compile the PlateSpin <code>blkwatch</code> driver before inventorying these workloads.
RHEL 4 (32-bit)	Block-based transfer only.
SUSE Linux Enterprise Server (SLES) 9, 10, 11 (SP1, SP2, SP3)	NOTE: Kernel version 3.0.13 of SLES 11 SP3 is not supported. Upgrade to kernel version 3.0.27 or later before inventorying the workload.
Novell Open Enterprise Server (OES) 11 SP1 and SP2	NOTE: The default kernel version 3.0.13 of OES 11 SP2 is not supported. Upgrade to kernel version 3.0.27 or later before inventorying the workload. Block-based transfer only.
OES 2 (SP2 and SP3)	Block-based transfer only.
Oracle Enterprise Linux (OEL)	Block-based transfer only. Same level of support as that for workloads running RHEL.

Replication is done at the block level, for which your PlateSpin software requires a `blkwatch` driver compiled for a particular Linux distribution being protected. For a list of distributions that include the `blkwatch` driver, see [Appendix A, "Linux Distributions Supported by Protect," on page 109](#). If your distribution does not include the `blkwatch` driver, you can create a custom driver by following the steps found in the Knowledgebase Article [KB 7005873](#).

Supported Linux File Systems

EXT2, EXT3, EXT4, REISERFS, and NSS (OES 2 and OES 11 workloads) file systems are supported, block-based transfer only.

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

1.2.3 Supported VM Containers

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 ESXi Server or a VMware DRS Cluster.

Table 1-3 Platforms Supported as VM Containers

Container	Notes
VMware ESXi 5.5 (GA, Update 1)	<ul style="list-style-type: none">◆ Supported as a protection and failback container◆ Includes support for VMware DRS cluster. The DRS configuration must be either Partially Automated or Fully Automated (it must not be set to Manual)◆ 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 ESXi 5.1 (GA, Update 1)	<ul style="list-style-type: none">◆ Supported as a protection and failback container◆ Includes support for VMware DRS cluster. The DRS configuration must be either Partially Automated or Fully Automated (it must not be set to Manual)◆ 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 ESXi 4.1 (GA, Update 3)	<ul style="list-style-type: none">◆ Supported as a protection and failback container◆ Includes support for VMware DRS cluster. The DRS configuration must be either Partially Automated or Fully Automated (it must not be set to Manual)◆ 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: ESXi versions must have a paid license; protection is unsupported with these systems if they are operating with a free license.

1.2.4 Supported System Firmware

PlateSpin Protect mirrors the Microsoft support of UEFI. For more information see [Section 1.2.1, "Supported Windows Workloads,"](#) on page 11.

1.3 Security and Privacy

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

- ◆ [Section 1.3.1, “Security of Workload Data in Transmission,” on page 15](#)
- ◆ [Section 1.3.2, “Security of Client/Server Communications,” on page 15](#)
- ◆ [Section 1.3.3, “Security of Credentials,” on page 15](#)
- ◆ [Section 1.3.4, “User Authorization and Authentication,” on page 15](#)

1.3.1 Security of Workload Data in Transmission

To make the transfer of your workload data more secure, you can configure the workload protection to encrypt the data. When encryption is enabled, data replicated over the network is encrypted by using AES (Advanced Encryption Standard).

If necessary, you can configure your PlateSpin Server to use a data encryption algorithm that is compliant with FIPS (Federal Information Processing Standards, Publication 140-2). See [“Enabling Support for FIPS-Compliant Data Encryption Algorithms \(Optional\)”](#) in your *Installation Guide*.

You can enable or disable encryption individually for each workload. See [“Workload Protection Details” on page 52](#).

1.3.2 Security of Client/Server Communications

Because the PlateSpin Server installation enables SSL on the PlateSpin Server host, secure data transmission between your Web browser and the PlateSpin Server is already configured to 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

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 [Section 2.2, “Setting Up User Authorization and Authentication,” on page 20](#).

1.4 Performance

- ◆ [Section 1.4.1, “About Product Performance Characteristics,” on page 16](#)
- ◆ [Section 1.4.2, “Data Compression,” on page 16](#)
- ◆ [Section 1.4.3, “Bandwidth Throttling,” on page 16](#)

- ♦ [Section 1.4.4, “RPO, RTO, and TTO Specifications,” on page 17](#)
- ♦ [Section 1.4.5, “Scalability,” on page 17](#)

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 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 73](#).

1.4.3 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 73](#).

1.4.4 RPO, RTO, and TTO Specifications

- ♦ **Recovery Point Objective (RPO):** Describes the acceptable amount of data loss measured in time. The RPO is determined by the time between incremental replications of a protected workload and 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):** Describes the time required for a failover operation (bringing a failover workload online to temporarily replace a protected production workload).

The RTO for failing a workload over to its virtual replica is affected by the time it takes to configure and execute the failover operation (10 to 45 minutes). See [“Failover” on page 56](#).

- ♦ **Test Time Objective (TTO):** Describes the time required for testing disaster recovery with some confidence of service restoration.

Use the **Test Failover** feature to run through different scenarios and generate benchmark data. See [“Using the Test Failover Feature” on page 58](#).

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.

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 45](#).

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 failovers) and hardware vendor specifications.

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

2 PlateSpin Protect Application Configuration

This section includes information about the following:

- ◆ [Section 2.1, “Product Licensing,” on page 19](#)
- ◆ [Section 2.2, “Setting Up User Authorization and Authentication,” on page 20](#)
- ◆ [Section 2.3, “Access and Communication Requirements across your Protection Network,” on page 24](#)
- ◆ [Section 2.4, “Configuring PlateSpin Protect Default Options,” on page 27](#)

2.1 Product Licensing

This section provides information about activating your PlateSpin Protect software.

- ◆ [Section 2.1.1, “Obtaining a License Activation Code,” on page 19](#)
- ◆ [Section 2.1.2, “Online License Activation,” on page 19](#)
- ◆ [Section 2.1.3, “Offline License Activation,” on page 20](#)

2.1.1 Obtaining a License Activation Code

For product licensing, you must have a license activation code. If you do not have a license activation code, request one through the [Customer Center Web site \(http://www.netiq.com/center/\)](http://www.netiq.com/center/). A license activation code will be e-mailed to you.

The first time you log into PlateSpin Protect, the browser is automatically redirected to the License Activation page. You have two options for activating your product license: [Online License Activation](#) or [Offline License Activation](#).aa

2.1.2 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.

- 1 In the PlateSpin Protect Web Interface, click **Settings > Licenses > Add License**. The License Activation page is displayed.

- 2 Select **Online Activation**, specify the e-mail address that you provided when placing your order and the activation code you received, then click **Activate**.

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

2.1.3 Offline License Activation

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

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

- 1 Click **Settings > License**, then click **Add license**. The License Activation page is displayed.
- 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 Web Site \(http://www.platespin.com/productactivation/ActivateOrder.aspx\)](http://www.platespin.com/productactivation/ActivateOrder.aspx). Log in with your Customer Center username.
- 4 Enter in the appropriate fields:
 - ◆ the activation code that you received
 - ◆ the e-mail address that you provided when placing your order
 - ◆ the hardware ID that you copied in [Step 2](#)
- 5 Click **Activate**.

The system generates a license key file and prompts you to save it.

- 6 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.

2.2 Setting Up User Authorization and Authentication

The following information is included in this section:

- ◆ [Section 2.2.1, "About PlateSpin Protect User Authorization and Authentication,"](#) on page 21
- ◆ [Section 2.2.2, "Managing PlateSpin Protect Access and Permissions,"](#) on page 22
- ◆ [Section 2.2.3, "Managing PlateSpin Protect Security Groups and Workload Permissions,"](#) on page 23

2.2.1 About PlateSpin Protect User Authorization and Authentication

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 2-1 Workload Protection Roles and Permission Details

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

Workload Protection Role Details	Administrators	Power Users	Operators
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.

Setting up a proper role-based access to PlateSpin Protect involves two tasks:

1. Adding users to the required user groups detailed in [Table 2-1](#) (see your Windows documentation).
2. Creating application-level security groups that associate these users with specified workloads (see [“Managing PlateSpin Protect Security Groups and Workload Permissions” on page 23](#)).

2.2.2 Managing PlateSpin Protect Access and Permissions

The following sections provide more information:

- ♦ [“Adding PlateSpin Protect Users” on page 22](#)
- ♦ [“Assigning a Workload Protection Role to a PlateSpin Protect User” on page 22](#)

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 22](#).

- 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**, 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 22](#).

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 2-1, “Workload Protection Roles and Permission Details,” on page 21](#).

- 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**, 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 following these steps:

- 1 Go to the PlateSpin Server's `bin\RestartPlateSpinServer` subdirectory.
- 2 Double-click the `RestartPlateSpinServer.exe` executable.
A command prompt window opens, requesting confirmation.
- 3 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 23](#).

2.2.3 Managing PlateSpin Protect Security Groups and Workload Permissions

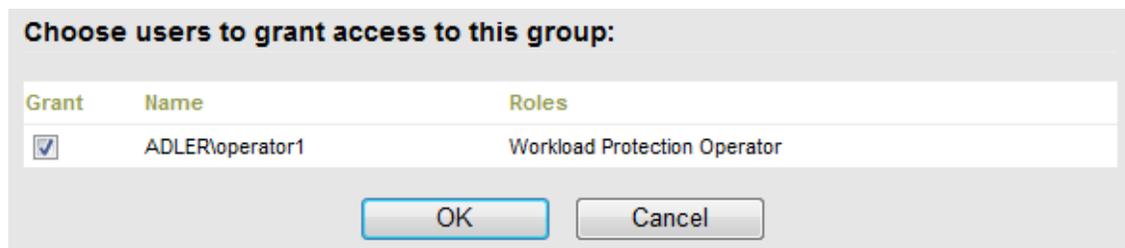
PlateSpin Protect provides a granular application-level access mechanism that allows specific users to carry out specific workload protection tasks on specified workloads. This is accomplished by setting up *security groups*.

- 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 22](#).
- 2 Access PlateSpin Protect as an administrator by using the PlateSpin Protect Web Interface, then click **Settings > Permissions**.

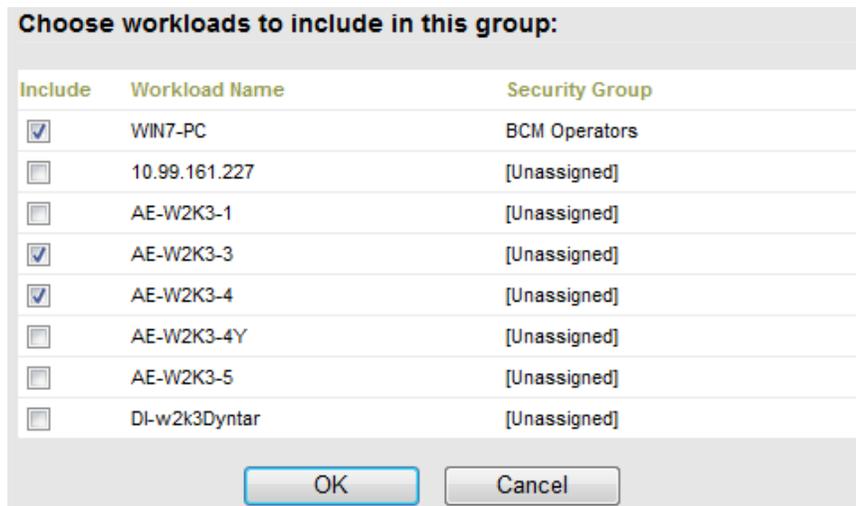
The Security Groups page opens:

- 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 that 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.

2.3 Access and Communication Requirements across your Protection Network

This section includes the following information:

- [Section 2.3.1, “Access and Communication Requirements for Workloads,” on page 24](#)
- [Section 2.3.2, “Access and Communication Requirements for Containers,” on page 25](#)
- [Section 2.3.3, “Open Port Requirements for PlateSpin Server Hosts,” on page 26](#)
- [Section 2.3.4, “Protection Across Public and Private Networks Through NAT,” on page 26](#)
- [Section 2.3.5, “Overriding the Default bash Shell for Executing Commands on Linux Workloads,” on page 27](#)
- [Section 2.3.6, “Requirements for VMware DRS Clusters as Containers,” on page 27](#)

2.3.1 Access and Communication Requirements for Workloads

The following software, network, and firewall requirements are for workloads that you intend to protect by using PlateSpin Protect.

Table 2-2 Access and Communication Requirements for Workloads

Workload Type	Prerequisites	Required Ports (Defaults)
All workloads	Ping (ICMP echo request and response) support	
All Windows workloads	Microsoft .NET Framework version 2.0 or 3.5 SP1	

Workload Type	Prerequisites	Required Ports (Defaults)
Windows Server 2008	<ul style="list-style-type: none"> ◆ 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: <ul style="list-style-type: none"> ◆ Option 1, using Windows Firewall: Use the basic Windows Firewall Control Panel item (<code>firewall.cpl</code>) and select File and printer Sharing in the list of exceptions. - OR - ◆ Option 2, using Firewall with Advanced Security: Use the Windows Firewall with Advanced Security utility (<code>wf.msc</code>) with the following Inbound Rules enabled and set to Allow: <ul style="list-style-type: none"> ◆ File and Printer Sharing (Echo Request - ICMPv4In) ◆ File and Printer Sharing (Echo Request - ICMPv6In) ◆ File and Printer Sharing (NB-Datagram-In) ◆ File and Printer Sharing (NB-Name-In) ◆ File and Printer Sharing (NB-Session-In) ◆ File and Printer Sharing (SMB-In) ◆ File and Printer Sharing (Spooler Service - RPC) ◆ File and Printer Sharing (Spooler Service - RPC-EPMAP) 	<p>TCP 3725</p> <p>NetBIOS 137 - 139</p> <p>SMB (TCP 139, 445 and UDP 137, 138)</p> <p>TCP 135/445</p>
Windows Server 2003 (including SP1 Standard, SP2 Enterprise, and R2 SP2 Enterprise)	<p>NOTE: After enabling the required ports, run the following command at the server prompt to enable PlateSpin remote administration:</p> <pre>netsh firewall set service RemoteAdmin enable</pre> <p>For more information about netsh, see the Microsoft TechNet article, http://technet.microsoft.com/en-us/library/cc785383%28v=ws.10%29.aspx. (http://technet.microsoft.com/en-us/library/cc785383%28v=ws.10%29.aspx).</p>	<ul style="list-style-type: none"> ◆ TCP: 3725, 135, 139, 445 ◆ UDP: 137, 138, 139
All Linux workloads	Secure Shell (SSH) server	TCP 22, 3725

2.3.2 Access and Communication Requirements for Containers

The following software, network, and firewall requirements are for the supported workload containers.

Table 2-3 Access and Communication Requirements for Containers

System	Prerequisites	Required Ports (Defaults)
All containers	Ping (ICMP echo request and response) capability.	

System	Prerequisites	Required Ports (Defaults)
VMware ESX/ESXi 4.1	<ul style="list-style-type: none"> VMware account with an Administrator role 	HTTPS (TCP 443)
VMware ESXi 5.0	<ul style="list-style-type: none"> VMware Web services API and file management API 	
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)

2.3.3 Open Port Requirements for PlateSpin Server Hosts

The following open port requirements are for PlateSpin Server hosts.

Table 2-4 Open Port Requirements for PlateSpin Server Hosts

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

2.3.4 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:** In your server's *PlateSpin Server Configuration* tool, record the additional IP addresses assigned to that host. See [“Configuring the Application to Function through NAT” on page 27](#).
- Target Container:** When you are attempting to discover a container (such as VMware ESX), specify the public (or external) IP address of that host in the discovery parameters.
- Workload:** When you are attempting 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 \(Workload to VM\) \(page 60\)](#).
- Failback Target:** During an attempt to register a failback target, when prompted to provide the IP address of the PlateSpin Server, provide either the local address of the Protect Server host or one of its public (external) addresses recorded in the server's *PlateSpin Server Configuration* tool (see *PlateSpin Server* above).

Configuring the Application to Function through NAT

To enable the PlateSpin Server to function across NAT-enabled environments, you must record additional IP addresses of your PlateSpin Server in the *PlateSpin Server Configuration* tool's database that the server reads upon startup.

For information on the update procedure, see [“Configuring PlateSpin Server Behavior through XML Configuration Parameters” on page 32.](#)

2.3.5 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 [KB Article 7010676 \(https://www.netiq.com/support/kb/doc.php?id=7010676\)](https://www.netiq.com/support/kb/doc.php?id=7010676).

2.3.6 Requirements 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 48.](#)

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

- ◆ DRS is enabled and set to either `Partially Automated` or `Fully Automated`.
- ◆ At least one datastore is shared among all the ESX servers in the VMware Cluster.
- ◆ At least one vSwitch and virtual port-group, or vNetwork Distributed Switch, is common to all the ESX servers in the VMware Cluster.
- ◆ The failover workloads (VMs) for each Protection contract is placed exclusively on datastores, vSwitches and virtual port-groups that are shared among all the ESX servers in the VMware Cluster.

2.4 Configuring PlateSpin Protect Default Options

This section includes the following information:

- ◆ [Section 2.4.1, “Setting Up Automatic E-Mail Notifications of Events and Reports,” on page 28](#)
- ◆ [Section 2.4.2, “Language Setup for International Versions of PlateSpin Protect,” on page 31](#)
- ◆ [Section 2.4.3, “Language Setup for International Versions of PlateSpin Protect,” on page 31](#)
- ◆ [Section 2.4.4, “Configuring PlateSpin Server Behavior through XML Configuration Parameters,” on page 32](#)
- ◆ [Section 2.4.5, “Configuring Support for VMware vCenter Site Recovery Manager,” on page 34](#)

2.4.1 Setting Up Automatic E-Mail Notifications of Events and Reports

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

- ♦ [“SMTP Configuration” on page 28](#)
- ♦ [“Setting Up Automatic Event Notifications by E-Mail” on page 28](#)
- ♦ [“Setting Up Automatic Replication Reports by E-Mail” on page 30](#)

SMTP Configuration

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

Figure 2-1 Simple Mail Transfer Protocol Settings

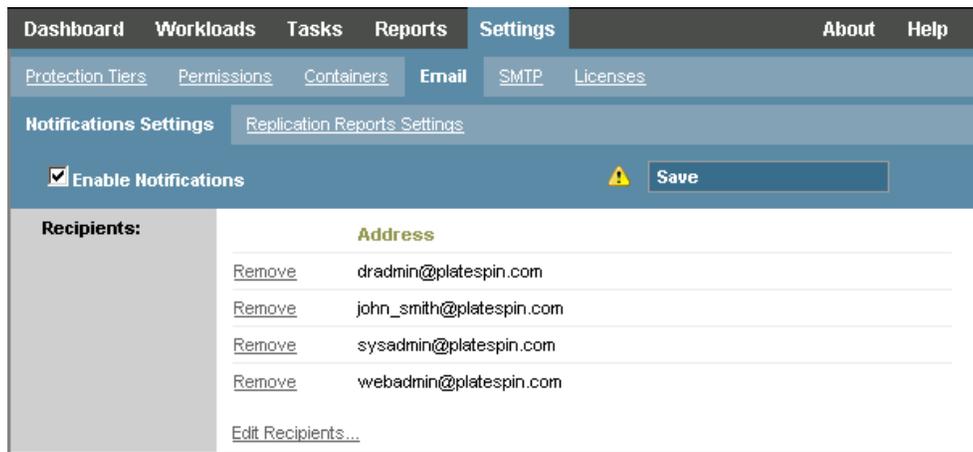
SMTP Settings		Save
SMTP Server Address:	<input type="text"/>	
Port:	<input type="text" value="25"/>	
Reply Address:	<input type="text"/>	
Username:	<input type="text"/>	
Password:	<input type="text"/>	
Confirm:	<input type="text"/>	

To configure SMTP settings:

- 1 In your PlateSpin Protect Web Interface, click **Settings > SMTP**.
- 2 Specify an SMTP server **Address**, a **Port** (the default is 25), and a **Reply Address** for receiving e-mail event and progress notifications.
- 3 Type a **Username** and **Password**, then confirm the password.
- 4 Click **Save**.

Setting Up Automatic Event Notifications by E-Mail

- 1 Set up an SMTP server for PlateSpin Protect to use. See [“SMTP Configuration” on page 28](#).
- 2 In your PlateSpin Protect Web Interface, click **Settings > Email > Notification Settings**.
- 3 Select the **Enable Notifications** option.
- 4 Click **Edit Recipients**, type the required e-mail addresses separated by commas, then click **OK**.



5 Click **Save**.

To delete listed e-mail addresses, click **Delete** next to the address that you want to remove.

The following events trigger e-mail notifications:

Event	Remarks
Workload Online Detected	Generated when the system detects that a previously offline workload is now online. Applies to workloads whose protection contract's state is not Paused .
Workload Offline Detected	Generated when the system detects that a previously online workload is now offline. Applies to workloads whose protection contract's state is not Paused .
Full Replication Successfully Completed	
Full Replication Failed	
Full Replication Missed	Similar to the Incremental Replication Missed event.
Incremental Replication Successfully Completed	
Incremental Replication Failed	
Incremental Replication Missed	Generated when any of the following applies: <ul style="list-style-type: none"> ◆ 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.

Event	Remarks
Test Failover Completed	Generated upon manually marking a Test Failover operation a success or a failure.
Prepare Failover Completed	
Prepare Failover Failed	
Failover Completed	
Failover Failed	

Setting Up Automatic Replication Reports by E-Mail

To set up PlateSpin Protect to automatically send out replication reports by e-mail, follow these steps:

- 1 Set up an SMTP server for PlateSpin Protect to use. See [SMTP Configuration \(page 28\)](#).
- 2 In your PlateSpin Protect Web Interface, click **Settings > Email > Replication Reports Settings**.
- 3 Select the **Enable Replication Reports** option.
- 4 In the **Report Recurrence** section, click **Configure** and specify the required recurrence pattern for the reports.
- 5 In the **Recipients** section, click **Edit Recipients**, type the required e-mail addresses separated by commas, then click **OK**.

- 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 e-mailed 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 45](#).

2.4.2 Language Setup for International Versions of PlateSpin Protect

PlateSpin Protect provides National Language Support (NLS) for Chinese Simplified, Chinese Traditional, French, German, and Japanese.

To use the PlateSpin Protect Web Interface and integrated help 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:
 - ◆ **Internet Explorer:** Click **Tools > Internet Options > General tab > Languages**.
 - ◆ **Firefox:** Click **Tools > Options > Content tab > 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 PlateSpin Protect Web Interface” on page 37](#).

NOTE: (For users of Chinese Traditional and Chinese Simplified versions) 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.

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:

2.4.3 Language Setup for International Versions of PlateSpin Protect

PlateSpin Protect provides National Language Support (NLS) for Chinese Simplified, Chinese Traditional, French, German, and Japanese.

To use the PlateSpin Protect Web Interface and integrated help 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 Lanaguages setting in your Web browser:
 - ◆ **Internet Explorer:** Click **Tools > Internet Options > General tab > Languages**.
 - ◆ **Firefox:** Click **Tools > Options > Content tab > 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 PlateSpin Protect Web Interface” on page 37](#).

NOTE: (For users of Chinese Traditional and Chinese Simplified versions) 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.

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:

- 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.

2.4.4 Configuring PlateSpin Server Behavior through XML Configuration Parameters

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

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

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

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

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

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

- ♦ [“Optimizing Data Transfer over WAN Connections” on page 32](#)
- ♦ [“Setting up Support for SRM” on page 33](#)

Optimizing Data Transfer over WAN Connections

You can optimize data transfer performance and fine tune it for WAN connections. You do this by modifying configuration parameters that the system reads from settings you make in a configuration tool residing on your PlateSpin Server host. For the generic procedure, see [“Configuring PlateSpin Server Behavior through XML Configuration Parameters” on page 32](#).

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

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

[Table 2-5](#) lists the configuration parameters with the defaults and with the values recommended for optimum operation in a high-latency WAN environment.

Table 2-5 Default and Optimized Configuration Parameters in https://Your_PlateSpin_Server/platespinconfiguration/

Parameter	Default Value	Optimized Value
<code>fileTransferMinCompressionLimit</code>	0 (disabled)	max 65536 (64 KB)
Specifies the packet-level compression threshold in bytes.		
<code>fileTransferCompressionThreadsCount</code>	2	N/A
Controls the number of threads used for packet-level data compression. This is ignored if compression is disabled. Because the compression is CPU-bound, this setting might have a performance impact.		
<code>fileTransferSendReceiveBufferSize</code>	0 (8192 bytes)	max 5242880 (5 MB)
TCP/IP window size setting for file transfer connections. It controls the number of bytes sent without TCP acknowledgement, in bytes.		
When the value is set to 0, the default TCP window size is used (8 KB). For custom sizes, specify the size in bytes. Use the following formula to determine the proper value:		
$((\text{LINK_SPEED}(\text{Mbps})/8) * \text{DELAY}(\text{sec})) * 1000 * 1000$		
For example, for a 100 Mbps link with 10 ms latency, the proper buffer size would be:		
$(100/8) * 0.01 * 1000 * 1000 = 125000 \text{ bytes}$		

Setting up Support for SRM

Workloads replicated by PlateSpin Protect and managed on VMware vCenter Site Recovery Manager (SRM) can behave seamlessly if you configure the product to support SRM. Some of the configuration requires a change to the XML configuration parameters of the PlateSpin Server. For information about these configuration changes, see the complete section: [Section 2.4.5, “Configuring Support for VMware vCenter Site Recovery Manager,” on page 34.](#)

2.4.5 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.

Protect configuration for SRM support consists of the following adjustments:

- ◆ Configure a setting to keep the PlateSpin Protect ISO and floppies on the same datastore as the VMware `.vmtx` 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.

Use the following steps to make sure 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.

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 Protect Server. In cases where this cannot happen, you need to prepare your environment prior to replication by following these steps:

- 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 5.0 and 5.1:** The setup packages are `setup.exe` and `setup64.exe`.
 - ◆ **VMware 4.0 and 4.1:** The setup packages are `VMware Tools.msi` and `VMware Tools64.msi`.
- 2 Create OFX packages from the setup packages you extracted from the VMware Server:
 - 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

%ProgramFiles(x86)%\PlateSpin\Packages\%GUID% on the PlateSpin Server. The value of %GUID% depends on the version of your VMware Server and its VMware Tools architecture.

The following table lists the server versions, VMware Tools architecture and the GUID identifier you need to copy the package to the correct directory:

VMware Server Version	VMware Tools Architecture	GUID
4.0	x86	D052CBAC-0A98-4880-8BCC-FE0608F0930F
4.0	x64	80B50267-B30C-4001-ABDF-EA288D1FD09C
4.1	x86	F2957064-65D7-4bda-A52B-3F5859624602
4.1	x64	80B1C53C-6B43-4843-9D63-E9911E9A15D5
5.0	x86	AD4FDE1D-DE86-4d05-B147-071F4E1D0326
5.0	x64	F7C9BC91-7733-4790-B7AF-62E074B73882
5.1	x86	34DD2CBE-183E-492f-9B36-7A8326080755
5.1	x64	AD4FDE1D-DE86-4d05-B147-071F4E1D0326

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.

3 Up and Running

This section provides information about the essential features of PlateSpin Protect and its interface.

- ♦ [Section 3.1, “Launching the PlateSpin Protect Web Interface,” on page 37](#)
- ♦ [Section 3.2, “Elements of the PlateSpin Protect Web Interface,” on page 38](#)
- ♦ [Section 3.3, “Workloads and Workload Commands,” on page 40](#)
- ♦ [Section 3.4, “Managing Multiple Instances of PlateSpin Protect and PlateSpin Forge,” on page 42](#)
- ♦ [Section 3.5, “Generating Workload and Workload Protection Reports,” on page 45](#)

3.1 Launching the PlateSpin Protect Web Interface

Most of your interaction with the appliance takes place through the browser-based PlateSpin Protect Web Client.

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:

- ♦ **Chrome:** From the Chrome menu, select **Settings**, scroll to and select **Show advanced settings...**, select **Content Settings > Allow all sites to run JavaScript**.
- ♦ **IE:** From the Tools menu, select **Internet Options > Security**, click **Custom level. . .**, scroll to and select **Active scripting**, select **Enable**, select **Yes** at the warning dialog box, click **OK**, click **Apply > OK**.
- ♦ **Firefox:** Click **Tools > Options > Content**, then select the **Enable JavaScript** option.

To use the PlateSpin Protect Web Interface and integrated help in one of the supported languages, see [Section 2.4.3, “Language Setup for International Versions of PlateSpin Protect,” on page 31](#).

To launch the PlateSpin Protect Web Interface:

- 1 Open a Web browser and go to:

`https://<hostname | IP_address>/Protect`

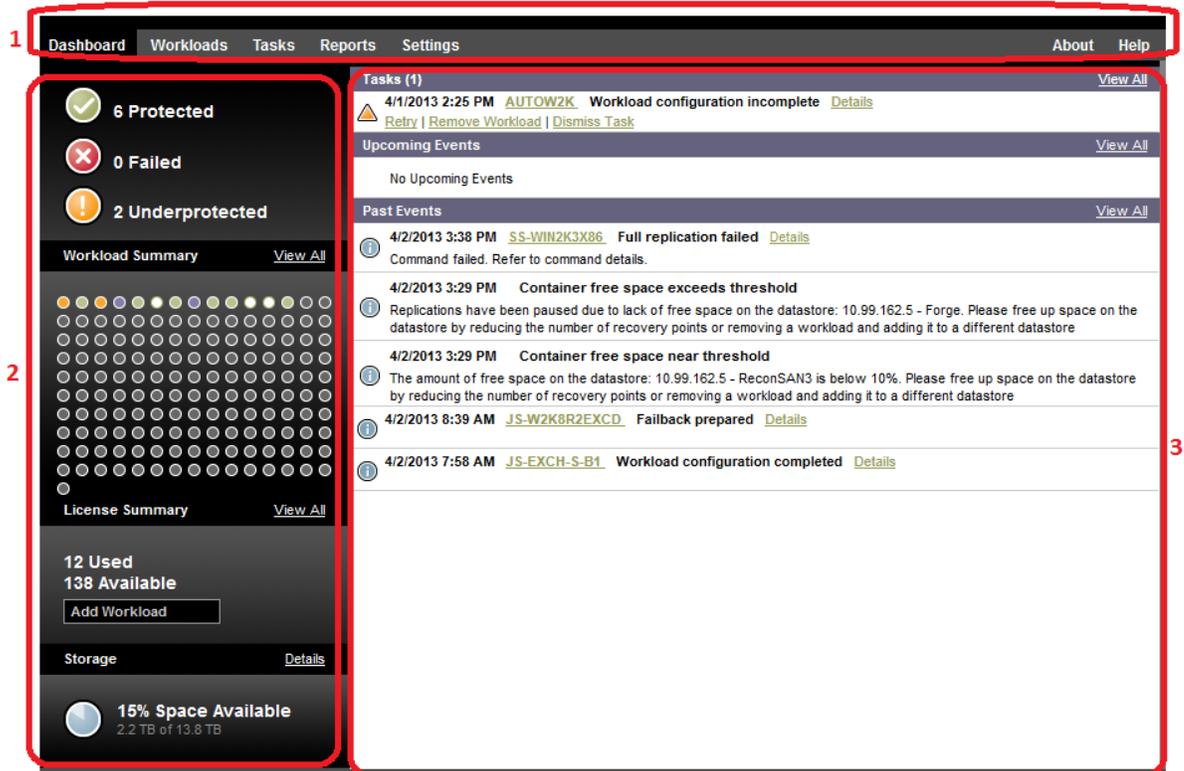
Replace `<hostname | IP_address>` with the hostname or the IP address of your PlateSpin Server host.

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

3.2 Elements of the PlateSpin Protect Web Interface

The default interface of the PlateSpin Protect Web Interface is the Dashboard page, which 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 39
- ♦ Section 3.2.2, “Visual Summary Panel,” on page 39
- ♦ Section 3.2.3, “Tasks and Events Panel,” on page 40

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 and Workload Commands” on page 40.](#)
- ◆ **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 45.](#)
- ◆ **Settings:** Displays the Settings page, which provides access to the following configuration options:
 - ◆ **Protection Tiers:** See [“Protection Tiers” on page 73.](#)
 - ◆ **Permissions:** See [“Setting Up User Authorization and Authentication” on page 20.](#)
 - ◆ **Containers:** See [“Adding Containers \(Protection Targets\)” on page 48.](#)
 - ◆ **Email/SMTP:** See [“Setting Up Automatic E-Mail Notifications of Events and Reports” on page 28.](#)
 - ◆ **Licenses/License Designations:** See [“Product Licensing” on page 19.](#)

3.2.2 Visual Summary Panel

The Visual Summary panel provides a high-level view of all licensed workloads and the amount of available storage.

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.

The area in the center of the left panel represents a graphical summary of the Workloads page. It uses the following dot icons to represent workloads in different states:

Table 3-1 Dot Icon Workload Representation

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

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

Storage provides information about container storage space available to PlateSpin Protect.

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 e-mail if SMTP is configured. See “[Setting Up Automatic E-Mail Notifications of Events and Reports](#)” on page 28.

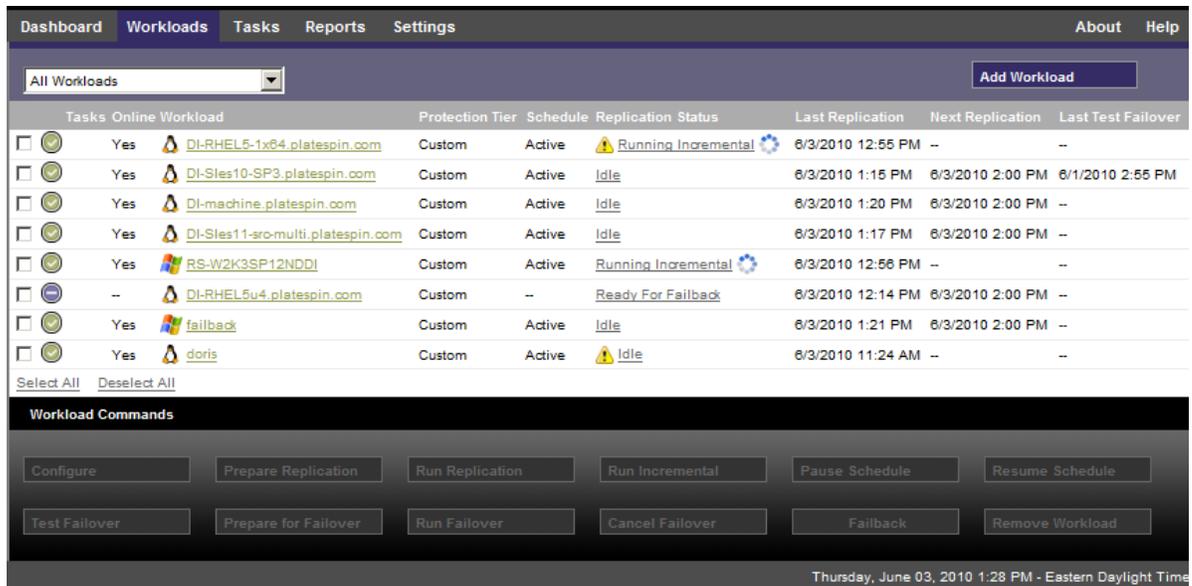
Tasks are special commands that are tied to events that require user intervention. For example, upon completion of a Test Failover command, the system generates an event associated with two tasks: Mark Test as Success and Mark Test as Failure. Clicking either task results in the Test Failover operation 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 and Workload Commands

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.

Figure 3-2 The Workloads Page

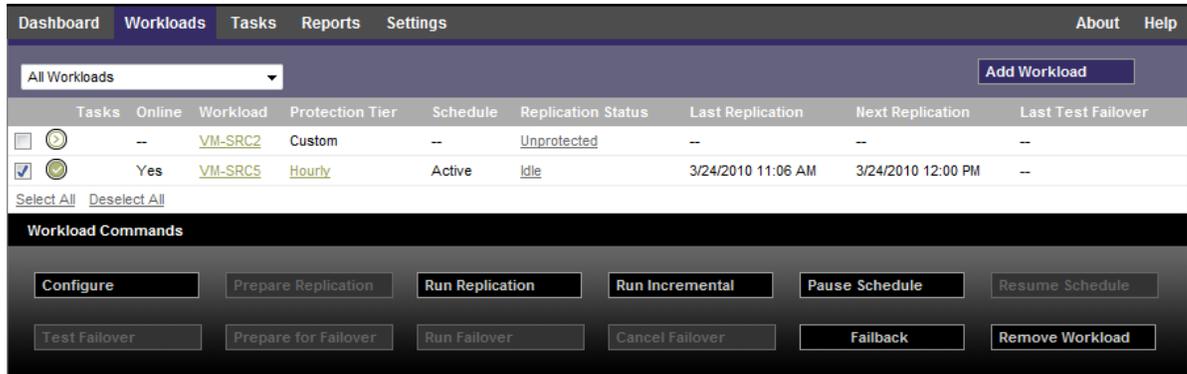


NOTE: All time stamps reflect the time zone of the PlateSpin Server host. This might be different from the time zone of the protected workload or the time zone of the host on which you are running the PlateSpin Protect Web Interface. A display of the server date and time appears at the bottom right of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

3.3.1 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



The following table 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.
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).
Remove Workload	Removes a workload from the inventory.

3.4 Managing Multiple Instances of PlateSpin Protect and PlateSpin Forge

PlateSpin Protect includes a Web-based client application, the PlateSpin Protect 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, 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.

- [Section 3.4.1, “Using the PlateSpin Protect Management Console,” on page 42](#)
- [Section 3.4.2, “About PlateSpin Protect Management Console Cards,” on page 42](#)
- [Section 3.4.3, “Adding Instances of PlateSpin Protect and PlateSpin Forge to the Management Console,” on page 43](#)
- [Section 3.4.4, “Managing Cards on the Management Console,” on page 44](#)

3.4.1 Using the PlateSpin Protect Management Console

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

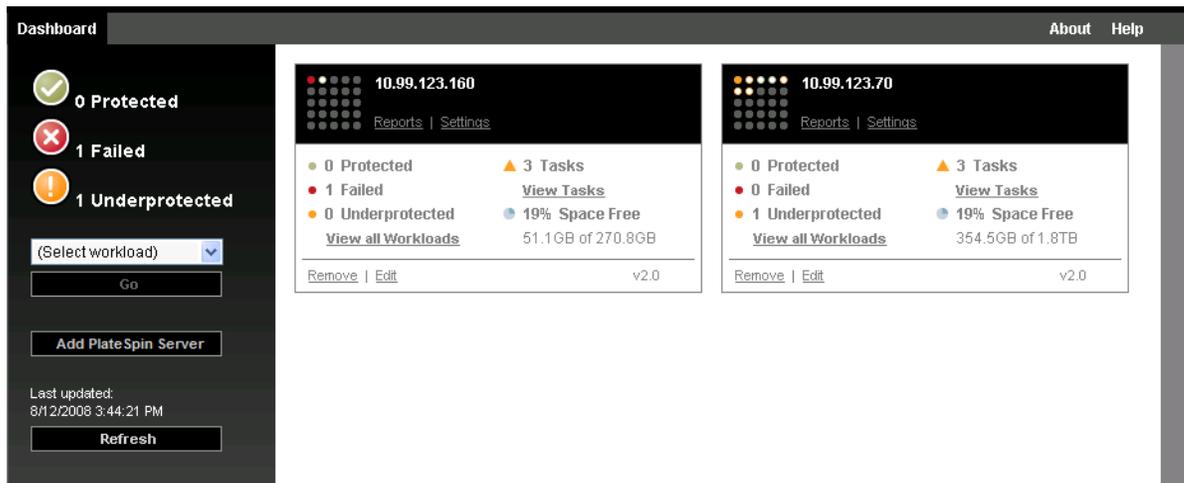
```
https://<IP_address | hostname>/console
```

Replace *<IP_address | hostname>* with either the IP address or the hostname of the PlateSpin Server host that is designated as the Manager.

- 2 Log in with your username and password.

The console’s default Dashboard page is displayed.

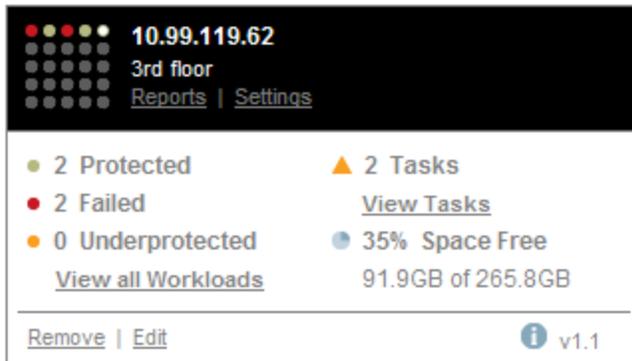
Figure 3-4 The Management Console’s Default Dashboard Page



3.4.2 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 3-5 PlateSpin Protect Instance Card



A card displays basic information about the specific instance of PlateSpin Protect or 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.

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

Adding a PlateSpin Protect or 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 Forge instance to the console:

- 1 On the console's main dashboard, click **Add PlateSpin Server**.
The **Add/Edit** page is displayed.
- 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 username 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 descriptive or identifying **Display Name** (15 characters max), a **Location** (20 characters max), and any **Notes** you might require (400 characters max).
- 6 Click **Add/Save**.
A new card is added to the dashboard.

3.4.4 Managing Cards on the Management Console

You can modify the details of a card on the Management Console.

- 1 Click the **Edit** hyperlink on the card that you want to edit.
The console's **Add/Edit** page is displayed.
- 2 Make any desired changes, then click **Add/Save**.
The updated console dashboard is displayed.

To remove a card from the Management Console:

- 1 Click the **Remove** hyperlink on the card you want to remove.
A confirmation prompt is displayed.
- 2 Click **OK**.
The individual card is removed from the dashboard.

3.5 Generating Workload and Workload Protection Reports

PlateSpin Protect enables you to generate reports that provide analytical insight into your workload protection contracts over time.

The following report types are supported:

- ♦ **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 3-6 Options for a Replication History Report

The screenshot shows the 'Reports' section of the PlateSpin Protect web interface. The 'Replication History' report is selected. The interface includes a navigation bar with 'Dashboard', 'Workloads', 'Tasks', 'Reports', and 'Settings'. Below the navigation bar, there is a search bar with the text 'What are the replication events relevant to my workload?'. The report configuration area includes a dropdown menu set to 'Custom', two date pickers for the time range (4/4/2011 12:00:00 AM to 4/18/2011 4:15:41 PM), a 'Workload' dropdown set to 'SES-2K8-1', and a checkbox for 'All Replication Events' with a link to 'Diagnostics View'. Below this is a table with the following data:

Date	Replication Event	Total Time	Transfer Time	Transfer Size	Transfer Speed
4/17/2011 4:01 AM	Incremental replication did not run as scheduled because the workload was busy	--	--	.0 MB	0.00 Mbps
4/17/2011 4:00 AM	Incremental replication did not run as scheduled because the workload was busy	--	--	.0 MB	0.00 Mbps
4/10/2011 4:01 AM	Incremental replication did not run as scheduled because the workload was busy	--	--	.0 MB	0.00 Mbps
4/10/2011 4:00 AM	Incremental replication did not run as scheduled because the workload was busy	--	--	.0 MB	0.00 Mbps

At the bottom of the report area, there are links for 'Printable View' and 'Export To Xml'. The footer of the interface shows the date and time: 'Monday, April 18, 2011 4:15 PM - Eastern Daylight Time'.

To generate a report:

- 1 In your PlateSpin Protect Web Interface, click **Reports**.
A list of the report types is displayed.
- 2 Click the name of the required report type.

4 Workload Protection

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 in the VM container of PlateSpin Protect that takes over the business function of your production workload in case of a disruption at the production site.

- ◆ [Section 4.1, “Basic Workflow for Workload Protection and Recovery,” on page 47](#)
- ◆ [Section 4.2, “Adding Containers \(Protection Targets\),” on page 48](#)
- ◆ [Section 4.3, “Adding Workloads for Protection,” on page 50](#)
- ◆ [Section 4.4, “Configuring Protection Details and Preparing the Replication,” on page 51](#)
- ◆ [Section 4.5, “Starting the Workload Protection,” on page 54](#)
- ◆ [Section 4.6, “Aborting Commands,” on page 55](#)
- ◆ [Section 4.7, “Failover,” on page 56](#)
- ◆ [Section 4.8, “Failback,” on page 58](#)
- ◆ [Section 4.9, “Reprotecting a Workload,” on page 63](#)

4.1 Basic Workflow for Workload Protection and Recovery

PlateSpin Protect defines the following workflow for workload protection and recovery:

- 1 Preparation:** This step involves preparatory steps to ensure that your workloads, containers, and environment meet the required criteria.
 - 1a** Make sure that PlateSpin Protect supports your workload.
See [“Supported Configurations” on page 11](#).
 - 1b** Make sure that your workloads and containers meet access and network prerequisites.
See [“Access and Communication Requirements across your Protection Network” on page 24](#).
 - 1c** (Linux only)
 - ◆ (Conditional) 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 [KB Article 7005873 \(https://www.netiq.com/support/kb/doc.php?id=7005873\)](https://www.netiq.com/support/kb/doc.php?id=7005873).
 - ◆ (Recommended) Prepare LVM snapshots for block-level data transfer. Ensure that each volume group has sufficient free space for LVM snapshots (at least 10 % of the sum of all partitions).
See [KB Article 7005872 \(https://www.netiq.com/support/kb/doc.php?id=7005872\)](https://www.netiq.com/support/kb/doc.php?id=7005872).
 - ◆ (Optional) Prepare `freeze` and `thaw` scripts to execute on your source workload upon each replication.
See [“Using Freeze and Thaw Scripts for Every Replication \(Linux\)” on page 76](#).

2 Inventory: This step involves adding workloads and containers to the PlateSpin Server database.

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. See [“Adding Containers \(Protection Targets\)” on page 48](#) and [“Adding Workloads for Protection” on page 50](#).

3 Definition of the protection contract: In this step, you define the details and specifications of a protection contract and prepare the replication.

See [“Configuring Protection Details and Preparing the Replication” on page 51](#).

4 Initiating the Protection: This step commences the protection contract according to your requirements.

See [“Starting the Workload Protection” on page 54](#).

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

- ◆ *Manual incremental.* You can run an incremental replication manually, outside the workload protection contract, by clicking **Run Incremental**.
- ◆ *Testing.* You can test failover functionality in a controlled manner and environment. See [Using the Test Failover Feature](#).

6 Failover: This step carries out a failover of your protected workload to its replica running in your VM container. See [“Failover” on page 56](#).

7 Failback: This step corresponds to the business resumption phase after you have addressed any problems with your production workload. See [“Failback” on page 58](#).

8 ReProtection: This step enables you to redefine the original protection contract for your workload. See [“Reprotecting a Workload” on page 63](#)

Most of these steps are represented by workload commands on the Workloads page. See [“Workloads and Workload Commands” on page 40](#).

A **Reprotect** command becomes available following a successful Failback operation.

4.2 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.

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 PlateSpin Protect Web Interface, click **Settings > Containers > Add Container**.

Containers						
Name	Description	Purpose	CPU	Memory	Free Space	Last Refresh
comp129	VMware ESX Server 4.0.0.261974	Protection	8 x Intel(R) Xeon(R) CPU X5355 @ 2.66GHz	15.6 GB	--	60 Day(s) ago Remove
Comp164	VMware ESX Server 3.5.0.207095	Failback/Deployment	32 x Intel(R) Xeon(R) CPU X735Q @ 2.93GHz	31.2 GB	--	60 Day(s) ago Remove
Add Container						

Monday, April 18, 2011 4:22 PM - Eastern Daylight Time

2 Specify the following parameters:

- ◆ **Type:** Select the type of the container (**VMware ESX Server** or **VMware DRS Cluster**). Make sure the container is supported.

For more information, see [“Supported VM Containers” on page 14](#).

- ◆ **Hostname or IP:** Type the container’s hostname or IP address.
- ◆ **vCenter Hostname or IP:** (DRS clusters only) Type the vCenter server’s hostname or IP address.
- ◆ **Cluster Name:** (DRS clusters only) Type the name of the required DRS cluster.

When you attempt to add or refresh a DRS cluster, the underlying discovery operation might fail if:

- ◆ A cluster contains no ESX hosts.
- ◆ A cluster name is not unique across the vCenter server (even if it has a unique inventory path).
- ◆ None of the cluster members are accessible (for example, because the vCenter server is in maintenance mode).
- ◆ **Username/Password:** Provide administrator-level credentials for accessing the required host. See [“Guidelines for Workload and Container Credentials” on page 66](#).
- ◆ **Purpose:** (VM containers only) Select the required item (**Protection**, **Failback/Deployment**, or both). Selecting both (**Protection** and **Failback/Deployment**) results in that container being available for selection as a target in both protection and failback/deployment operations.

3 Click **Add**.

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 .

To refresh a container, click the **Refresh** icon next to the container you want to refresh. This performs a re-inventory of the container.

To remove a container, click **Remove** adjacent to the container that you want to remove.

Select the container name to edit its vCenter hostname, IP address, username, or password.

4.3 Adding Workloads for Protection

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 [Step 1](#) in “[Basic Workflow for Workload Protection and Recovery](#)” on page 47.
- 2 On the Dashboard or Workloads page, click **Add Workload**.

The PlateSpin Protect Web Interface displays the Add Workload page.

The screenshot shows the 'Add Workload' page in the PlateSpin Protect Web Interface. The page has a navigation bar with 'Dashboard', 'Workloads', 'Tasks', 'Reports', 'Settings', 'About', and 'Help'. The main content area is titled 'Add Workload' and features a progress bar with four steps: 'ADD WORKLOAD' (highlighted in green), 'CONFIGURE PROTECTION', 'PREPARE REPLICATION', and 'RUN REPLICATION'. Below the progress bar is the 'Workload Settings' section, which includes the following fields:

- Hostname or IP:** 10.99.123.170
- Workload Type:** Radio buttons for 'Windows' and 'Linux' (selected).
- Credentials:** Fields for 'User Name' (root) and 'Password' (masked with dots). A 'Test Credentials' link is located below the password field.

At the bottom of the page, there is a 'Workload Commands' section with two buttons: 'Add Workload' and 'Add and New'.

- 3 Specify the required workload details:
 - ♦ **Workload Settings:** Specify your workload’s hostname or IP address, the operating system, and administrator-level credentials.

Use the required credential format. See “[Guidelines for Workload and Container Credentials](#)” on page 66.

To make sure 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.

If you haven’t added a container yet, add one to prepare for protecting the workload, otherwise, skip to “[Configuring Protection Details and Preparing the Replication](#)” on page 51

4.4 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 (see [“Basic Workflow for Workload Protection and Recovery” on page 47](#)), relevant settings are read from the protection details.

To configure your workload’s protection details:

- 1 Add a workload. See [“Adding Workloads for Protection” on page 50](#).
- 2 Add a container. See [“Adding Containers \(Protection Targets\)” on page 48](#).
- 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 doesn’t 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 74](#).
- 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 74](#).

NOTE: If your inventory has only one container, your workload is automatically assigned to it.

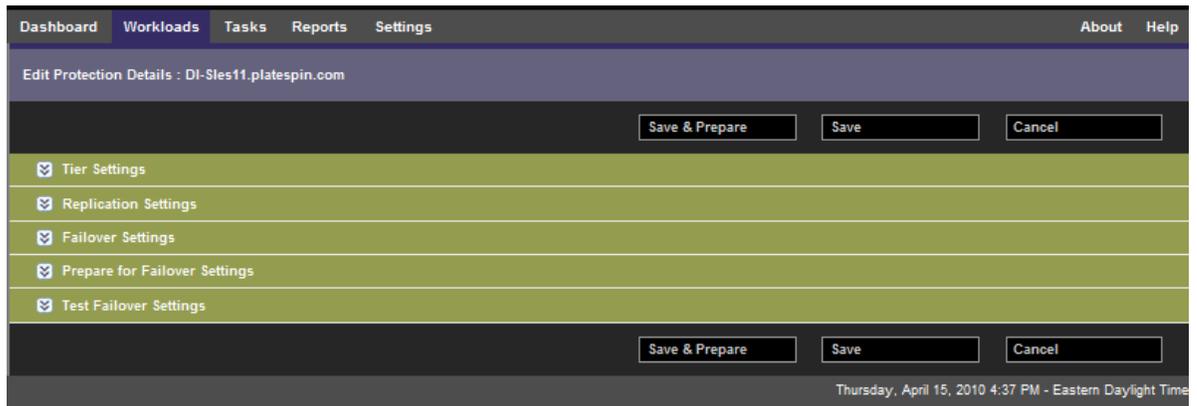
- 6 Configure the protection details in each set of settings as dictated by your business continuity needs. See [“Workload Protection Details” on page 52](#).
- 7 Correct any validation errors, if displayed by the PlateSpin Protect Web Interface.
- 8 Click **Save**.

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.

4.4.1 Workload Protection Details

Workload protection details are represented by five sets of parameters:



You can expand or collapse each parameter set by clicking the  icon at the left.

The following are the details of the five parameter sets:

Table 4-1 Workload Protection Details

Parameter Set	Details (Settings)
---------------	--------------------

Tier	Indicates the Protection Tier that the current protection uses. See “Protection Tiers” on page 73 .
------	---

**Parameter Set Details
(Settings)**

Replication

Transfer Method: (Windows) Select a data transfer mechanism and security through encryption. See [“Data Transfer” on page 71](#).

Transfer Encryption: To enable encryption, select the **Encrypt Data Transfer** option. See [“Supported Configurations” on page 11](#).

Source Credentials: Required for accessing the workload. See [“Guidelines for Workload and Container Credentials” on page 66](#).

Number of CPUs: Specify the required number of vCPUs assigned to the failover workload (applicable only when the selected method of initial replication is **Full**).

Replication Network: Lets you separate replication traffic based on virtual networks defined on your VM container. See [“Networking” on page 78](#). 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. If you do not specify an MTU value, one is supplied for you.

Allowed Networks: Specifies one or more network interfaces (NIC or IP address) on the source to use for replication traffic.

Resource Pool for Target VM: (Setting available only if the container is a DRS Cluster.) The Resource Pool location where the failover VM is to be created in the Protect container.

VM Folder for Target VM: (Setting available only if the container is a DRS Cluster.) The VM folder location where the failover VM is to be created in the Protect container.

Configuration File Datastore: Lets you select a datastore associated with your VM container for storing VM configuration files. See [“Recovery Points” on page 74](#).

Protected Volumes: Use these options to select volumes for protection and to assign their replicas to specific datastores on your VM container.

- ♦ **Thin Disk option:** Enables 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: Specifies the LVM logical volumes to be protected for a Linux workload or the NSS Pools on an OES workload.

Non-volume Storage: Specifies 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: In Linux, this setting specifies the LVM volume groups to be protected with the LVM logical volumes listed in the [Protected Logical Volumes](#) section of the settings.

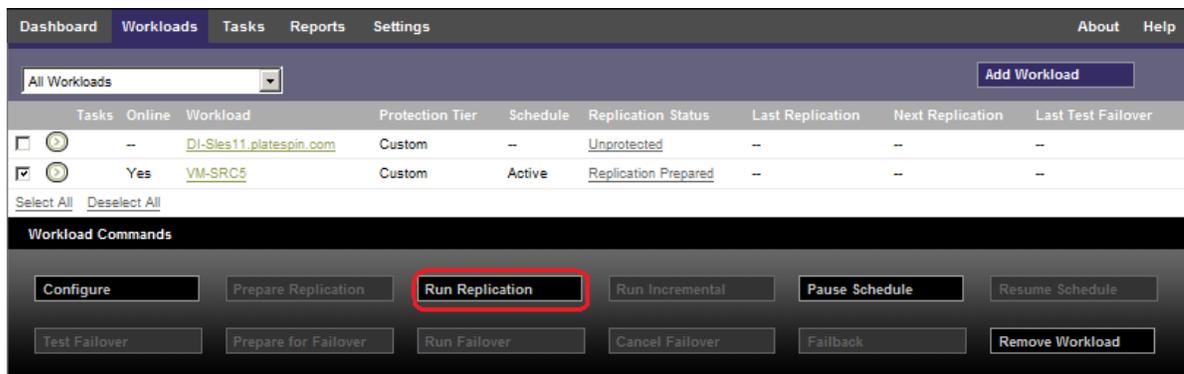
Services/Daemons to Stop During Replication: Lets you select Windows services or Linux Daemons that are automatically stopped during the replication. See [Section 5.8, “Service and Daemon Control,” on page 75](#).

Parameter Set Details (Settings)

Failover	<p>VM Memory: Specifies the amount of memory allocated to the failover workload.</p> <p>Hostname and Domain/Workgroup affiliation: Use these options to control the identity and domain/workgroup affiliation of the failover workload when it is live. For domain affiliation, domain administrator credentials are required.</p> <p>Network Connections: Use these options to control the LAN settings of the failover workload. See “Networking” on page 78.</p> <p>Service/Daemon States to Change: Enables you to control the startup state of specific application services (Windows) or daemons (Linux). See Section 5.8, “Service and Daemon Control,” on page 75.</p>
Prepare for Failover	<p>Enables you to control the temporary network settings of the failover workload during the optional Prepare for Failover operation. See “Networking” on page 78.</p>
Test Failover	<p>VM Memory: Enables you to assign the required RAM to the temporary workload.</p> <p>Hostname: Enables you to assign a hostname to the temporary workload.</p> <p>Domain/Workgroup: Enables you to affiliate the temporary workload with a domain or a workgroup. For domain affiliation, domain administrator credentials are required.</p> <p>Network Connections: Controls the LAN settings of the temporary workload. See “Networking” on page 78.</p> <p>Service/Daemon States to Change: Enables you to control the startup state of specific application services (Windows) or daemons (Linux). See “Service and Daemon Control” on page 75.</p>

4.5 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.
-

4.6 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.

<input type="checkbox"/>		No	CL-2K8R2-VM1	Custom	Active	Idle	3/5/2012 12:23 AM	4/11/2012 12:00 AM	--
<input type="checkbox"/>		Yes	DI-Sles11x64-Src	every 4 hours	Active	Failover Prepared	3/29/2012 8:13 AM	4/9/2012 12:00 PM	3/23/2012 3:32 PM
<input type="checkbox"/>		--	ma-cl-slessp2.site	every 4 hours	--	Live	3/15/2012 2:49 PM	--	3/9/2012 2:44 PM
<input type="checkbox"/>		Yes	VISTACLIENT	Custom	Active	Running Incremental	3/26/2012 10:21 AM	4/9/2012 12:00 PM	3/23/2012 5:14 PM
<input type="checkbox"/>		--	CL-VISTASP1-SRC	every 4 hours	--	Live	2/22/2012 2:55 PM	--	--
<input type="checkbox"/>		Yes	CL-XPX64-SRC	Custom	Active	Idle	4/9/2012 10:17 AM	4/9/2012 12:00 PM	3/23/2012 5:15 PM

The PlateSpin Protect Web Interface displays the appropriate Command Details page:

The screenshot shows the VISTACLIENT interface for a 'Running Incremental' command. The status is 'Running' with a warning icon. The duration is 3d 21h 31m 37s. The step is 'Copy data (2%)' with a progress bar. The 'Workload Commands' section at the bottom has three buttons: 'Abort', 'Configure', and 'Pause Schedule'. The 'Abort' button is highlighted with a red rectangle.

Event	Details	User	Date
Incremental replication started		DEV-MORTA.ZAA\PlateSpin	4/5/2012 2:00 PM

Step	Status	Start Time	End Time	Duration	Diagnostics
Revert to snapshot	Completed	4/5/2012 2:00 PM	4/5/2012 2:01 PM	1m 7s	--
Copy data	Running (2%)	4/5/2012 2:01 PM	--	3d 21h 30m 30s	--

3 Click **Abort**.

4.7 Failover

A *Failover* results in the business function of a failed workload being taken over by a failover workload within a PlateSpin Protect VM container.

- ◆ Section 4.7.1, “Detecting Offline Workloads,” on page 56
- ◆ Section 4.7.2, “Performing a Failover,” on page 57
- ◆ Section 4.7.3, “Using the Test Failover Feature,” on page 58

4.7.1 Detecting Offline Workloads

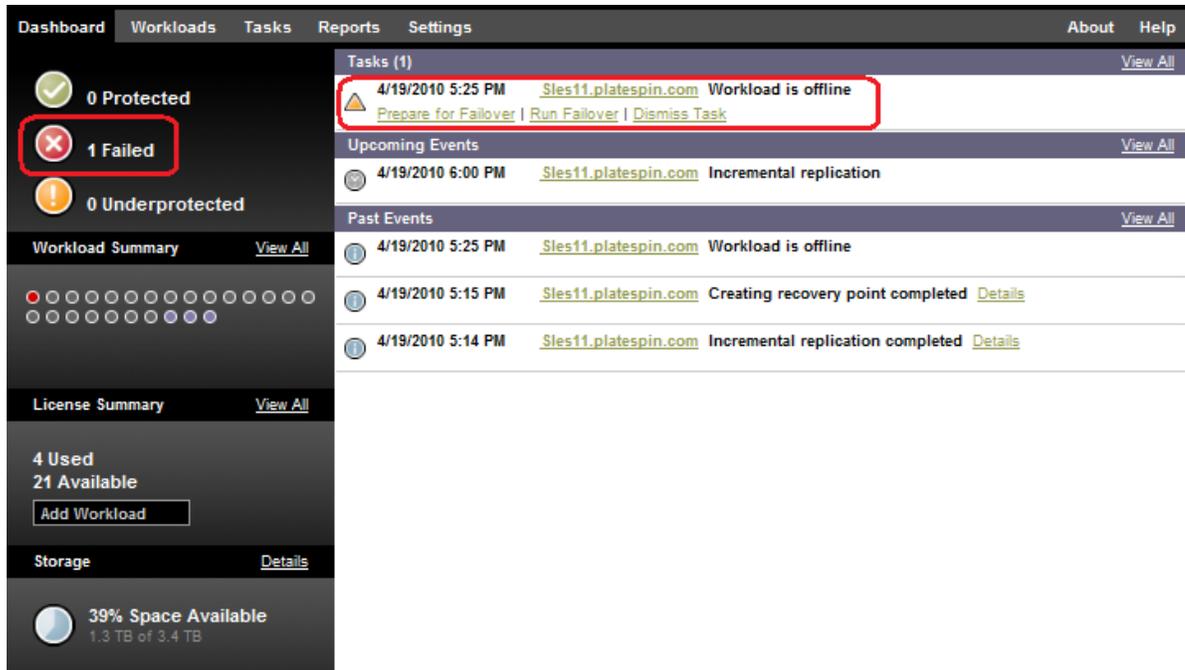
PlateSpin Protect constantly monitors your protected workloads. If an attempt to monitor a workload fails for a predefined number of times, PlateSpin Protect generates a **Workload is offline** event. Criteria that determine and log a workload failure are part of a workload protection’s Tier settings (see the [Tier](#) row in “[Workload Protection Details](#)” on page 52).

If notifications are configured along with SMTP settings, PlateSpin Protect simultaneously sends a notification e-mail to the specified recipients. See “[Setting Up Automatic E-Mail Notifications of Events and Reports](#)” on page 28.

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 55), and then proceed to the **Run Failover** command. See “[Performing a Failover](#)” on page 57.

The following figure shows the PlateSpin Protect Web Interface’s Dashboard page upon detecting a workload failure. Note the applicable tasks in the Tasks and Events pane:

Figure 4-1 The Dashboard Page upon Workload Failure Detection ('Workload Offline')



4.7.2 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 the [Failover](#) row in "Workload Protection Details" on page 52.

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 4-1](#).
- ◆ 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 74). 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 58.

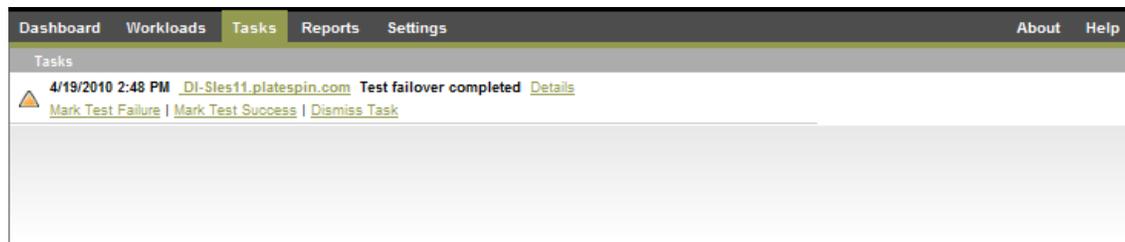
4.7.3 Using the Test Failover Feature

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 a restricted network environment for testing.

When you execute the command, PlateSpin Protect applies the Test Failover Settings, as saved in the workload protection details, to the failover workload (see the **Test Failover** row in “[Workload Protection Details](#)” on page 52).

- 1 Define an appropriate time window for testing and make sure 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 74), 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.

4.8 Failback

A Failback operation 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 depend 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 Xen on SLES and Microsoft Hyper-V platforms.

The following topics provide more information:

- ♦ [Section 4.8.1, “Automated Failback to a VM Platform,” on page 59](#)
- ♦ [Section 4.8.2, “Semi-Automated Failback to a Physical Machine,” on page 61](#)
- ♦ [Section 4.8.3, “Semi-Automated Failback to a Virtual Machine,” on page 62](#)

4.8.1 Automated Failback to a VM Platform

PlateSpin Protect supports automated failback for failback/deployment containers on a supported VMware ESXi Server or a VMware DRS Cluster. See [Section 1.2.3, “Supported VM Containers,” on page 14](#).

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 hostname or IP address and provide administrator-level credentials. Use the required credential format (see [“Guidelines for Workload and Container Credentials” on page 66](#)).
 - ♦ **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 74](#).
 - ♦ **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 60](#).
- 5 Click **Save and Failback** and monitor the progress on the Command Details page. See [Figure 4-2](#).

PlateSpin Protect executes the command. If you selected **Reprotect after Failback** in the Post-Failback parameter set, a **Reprotect** command is shown in the PlateSpin Protect Web Interface.

Figure 4-2 Failback Command Details

The screenshot displays the 'Running Failback' command details for VM-SRC18. The status is 'Running' with a progress bar for 'Copy data (91%)' and 'Uninstalling VMware Tools (1%)'. The command summary shows the start time as 4/20/2010 7:55 PM and a duration of 42m 22s. The replication transfer summary shows an average transfer speed of 40.88 Mbps, total data transferred of 5.5 GB, total files transferred of 27138, total folders transferred of 4063, and a duration of 19m 25s.

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.

Table 4-2 Failback Details (VM)

Parameter Set (Settings)	Details
--------------------------	---------

Failback	<p>Transfer Method: Enables you to select a data transfer mechanism and security through encryption. See “Data Transfer” on page 71.</p> <p>Failback Network: Enables you to direct failback traffic over a dedicated network based on virtual networks defined on your VM container. See “Networking” on page 78.</p> <p>VM Datastore: Enables you to select a datastore associated with your failback container for the target workload.</p> <p>Volume Mapping: When the initial replication method is specified as “incremental”, enables you to select source volumes and map to volumes on the failback target for synchronization.</p> <p>Services/Daemons to stop: Enables you to select Windows services or Linux daemons that are automatically stopped during the failback. See “Service and Daemon Control” on page 75.</p> <p>Alternative Address for Source: Accepts input of an additional IP address for the failed-over VM if applicable. See “Protection Across Public and Private Networks Through NAT” on page 26.</p>
----------	---

Parameter Set (Settings)	Details
Workload	<p>Number of CPUs: Enables you to specify the required number of vCPUs assigned to the target workload.</p> <p>VM Memory: Enables you to assign the required RAM to the target workload .</p> <p>Hostname, Domain/Workgroup: Use these options to control the identity and domain/workgroup affiliation of the target workload. For domain affiliation, domain administrator credentials are required.</p> <p>Network Connections: Use these options to specify the network mapping of the target workload based on the virtual networks of the underlying VM container.</p> <p>Service States to Change: Enables you to control the startup state of specific application services (Windows) or daemons (Linux). See “Service and Daemon Control” on page 75.</p>
Post-Failback	<p>Reprotect Workload: Use this option if you plan to re-create the protection contract for the target workload after deployment. This maintains a continuous event history for the workload and auto-assigns/designates a workload license.</p> <ul style="list-style-type: none"> ◆ 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 PlateSpin Protect Web Interface for the failed-back workload. ◆ No reprotect: Select this option if you don't 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.

4.8.2 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 78](#).
- 2 If there are missing or incompatible drivers, upload the required drivers to the PlateSpin Protect device driver database. See [“Managing Device Drivers” on page 87](#).
- 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 hostname or IP address and provide administrator-level credentials. Use the required credential format (see [“Guidelines for Workload and Container Credentials” on page 66](#)).
 - ◆ **Failback Target Settings:** Specify the following parameters:
 - ◆ **Replication Method:** Select the scope of data replication. See [“Initial Replication Method \(Full and Incremental\)” on page 74](#).
 - ◆ **Target Type:** Select the **Physical Target** option and then select the physical machine you registered in [Step 1](#).

The screenshot shows the 'CONFIGURE FAILBACK' step of a three-stage process (PREPARE FAILBACK, CONFIGURE FAILBACK, RUN FAILBACK). The interface is divided into two main sections: 'Workload Settings' and 'Failback Target Settings'.

Workload Settings:

- Hostname or IP:** MA--Rhel5u3
- Credentials:**
 - User Name: root
 - Password: [masked]
 - [Test Credentials](#)

Failback Target Settings:

- Replication Method:**
 - Full Replication
 - Incremental Replication
- Target type:**
 - Virtual Targets
 - Physical Targets
- Failback Target:** [Selection required below] ❌

A message below the target selection states: "No physical targets available." A note at the bottom of this section reads: "Note: To add a physical target, boot up and register the physical server with PlateSpin Failback ISO Image. To download, visit the [PlateSpin Resource Centre](#)."

Workload Commands:

- Save and Prepare** ▶

- 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.

4.8.3 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
```

4.9 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 PlateSpin Protect Web Interface 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 58](#).

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.

5 Essentials of Workload Protection

This section provides information about the different functional areas of a workload protection contract.

- ◆ [Section 5.1, “Workload License Consumption,” on page 65](#)
- ◆ [Section 5.2, “Guidelines for Workload and Container Credentials,” on page 66](#)
- ◆ [Section 5.3, “Setting Up Protect Multitenancy on VMware,” on page 66](#)
- ◆ [Section 5.4, “Data Transfer,” on page 71](#)
- ◆ [Section 5.5, “Protection Tiers,” on page 73](#)
- ◆ [Section 5.6, “Recovery Points,” on page 74](#)
- ◆ [Section 5.7, “Initial Replication Method \(Full and Incremental\),” on page 74](#)
- ◆ [Section 5.8, “Service and Daemon Control,” on page 75](#)
- ◆ [Section 5.9, “Using Freeze and Thaw Scripts for Every Replication \(Linux\),” on page 76](#)
- ◆ [Section 5.10, “Volumes,” on page 77](#)
- ◆ [Section 5.11, “Networking,” on page 78](#)
- ◆ [Section 5.12, “Failback to Physical Machines,” on page 78](#)
- ◆ [Section 5.13, “Advanced Workload Protection Topics,” on page 81](#)

5.1 Workload License Consumption

Your PlateSpin Protect product license entitles you to a specific 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.

For information about product licensing and license activation, see [“Product Licensing” on page 19](#).

5.2 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 5-1 Workload and Container Credentials

To Discover	Credentials	Remarks
All Windows workloads	Local or domain administrator credentials.	For the username, use this format: <ul style="list-style-type: none">◆ For domain member machines: <i>authority\principal</i>◆ For workgroup member machines: <i>hostname\principal</i>
Windows Clusters	Domain administrator credentials	
All Linux workloads	Root-level username and password	Non-root accounts must be properly configured to use <code>sudo</code> . See KB Article 7920711 .
VMware ESX/ESXi 4.1; ESXi 5.0, ESXi 5.1, ESXi 5.5	VMware account with an appropriate role configuration. See Section 5.3.1, “Using Tools to Define VMware Roles,” on page 67.	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. See Section 5.3.1, “Using Tools to Define VMware Roles,” on page 67.	

5.3 Setting Up Protect Multitenancy on VMware

PlateSpin Protect includes unique user roles (and a tool for creating them in a VMware datacenter) that make it possible non-administrative VMware users (or “enabled users”) to perform Protect lifecycle operations in the VMware environment. These roles makes it possible for you, as a service provider, to segment your VMware cluster to allow multitenancy: where multiple Protect containers are instantiated in your datacenter 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 datacenter.

This section includes the following information:

- ◆ [Section 5.3.1, “Using Tools to Define VMware Roles,”](#) on page 67
- ◆ [Section 5.3.2, “Assigning Roles In vCenter,”](#) on page 68

5.3.1 Using Tools to Define VMware Roles

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. Because there are many of these required privileges, NetIQ has created a file that defines the minimum required privileges and aggregates them respectively into three VMware custom roles:

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

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

This section includes the following information:

- ◆ [“Basic Command Line Syntax” on page 67](#)
- ◆ [“Additional Command Line Parameters and Flags” on page 67](#)
- ◆ [“Tool Usage Example” on page 68](#)
- ◆ [“\(Option\) Manually Defining the PlateSpin Roles in vCenter” on page 68](#)

Basic Command Line Syntax

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

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

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

Additional Command Line Parameters and Flags

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

<code>/create</code>	(mandatory) Creates the roles defined by the <code>/role</code> parameter
<code>/get_all_privileges</code>	Display all server-defined privileges

Optional Flags

<code>/interactive</code>	Run the tool with interactive options that allow you to choose to create individual roles, check role compatibility, or list all compatible roles.
<code>/password=[password]</code>	Provide the VMware password (bypasses the password prompt)
<code>/verbose</code>	Display detailed information

Tool Usage Example

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

Resulting Actions:

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

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

(Option) Manually Defining the PlateSpin Roles in vCenter

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

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

5.3.2 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 <i>PlateSpin Infrastructure Manager</i> (or equivalent) role.	For security reasons, define the permission as non-propagating.	This role is needed to monitor tasks being performed by the Protect software and to end any stale VMware sessions.

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

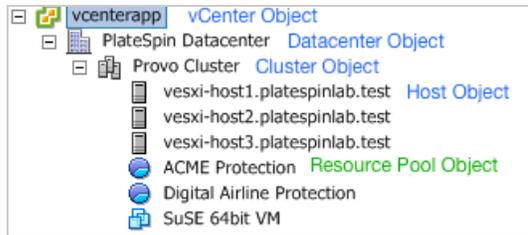
vCenter Container for Role Assignment	Role Assignment Specifics	Propagate Instructions	More Information
Each Datastore and Datastore Cluster where the enabled user needs access	Assign the enabled user the <i>PlateSpin Virtual Machine Manager</i> (or equivalent) role.	Propagation is at the discretion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

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

vCenter Container for Role Assignment	Role Assignment Specifics	Propagate Instructions	More Information
Each resource pool(s) and folder(s) where the customer's VMs will be created.	Assign the tenant user the <i>PlateSpin User</i> (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	This tenant is a member of the PlateSpin Administrators group on the PlateSpin 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 you want to allow the customer to change the network where their VM is attached, this user should be assigned the Read-only role (or better) on all of the networks being made accessible to the customer.

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

Figure 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/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.

5.4 Data Transfer

The following topics provide information about the mechanisms and options of data transfer from your workloads to their replicas.

- [Section 5.4.1, "Transfer Methods," on page 72](#)
- [Section 5.4.2, "Data Encryption," on page 73](#)

5.4.1 Transfer Methods

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

- ♦ “Transfer Methods Supported for Windows Workloads” on page 72
- ♦ “Transfer Methods Supported for Linux Workloads” on page 72

Transfer Methods Supported for Windows Workloads

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

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 and leverages the Microsoft Volume Snapshot Service (VSS) with 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 are configuring 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.

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

Windows File-level Replication: Data is replicated on a file-by-file basis (Windows only).

Transfer Methods Supported for Linux Workloads

For Linux workloads, PlateSpin Protect provides a mechanism to transfer workload volume data at block level only. Data transfer is powered by a block-level data transfer component that leverages LVM snapshots if available (this is the default and recommended option). See [KB Article 7005872 \(https://www.netiq.com/support/kb/doc.php?id=7005872\)](https://www.netiq.com/support/kb/doc.php?id=7005872).

The Linux block-based component included in your PlateSpin Protect distribution is precompiled for the standard, non-debug kernels of the supported Linux distributions. If you have a non-standard, customized, or newer kernel, you can rebuild the block-based component for your specific kernel. See [KB Article 7005873 \(https://www.netiq.com/support/kb/doc.php?id=7005873\)](https://www.netiq.com/support/kb/doc.php?id=7005873).

Deployment or removal of the component is transparent, has no continuity impact, and requires no intervention and no reboot.

5.4.2 Data Encryption

To make the transfer of workload data more secure, PlateSpin Protect enables you to encrypt data replication. When encryption is enabled, over-the-network data transfer from the source to the target is encrypted by using AES (Advanced Encryption Standard) or 3DES if FIPS-compliant encryption is enabled (see “[Enabling Support for FIPS-Compliant Data Encryption Algorithms \(Optional\)](#)” in your *Installation and Upgrade Guide*).

NOTE: Data encryption has a performance impact and might significantly slow down the data transfer.

5.5 Protection Tiers

A Protection Tier is a customizable 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.

You can also create custom Protection Tiers in advance:

- 1 In your PlateSpin Protect Web Interface, click **Settings > Protection Tiers > Create Protection Tier**.
- 2 Specify the parameters for the new Protection Tier:

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.
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	<p>These settings control whether and how workload data is compressed before transmission. See “Data Compression” on page 16.</p> <p>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.</p>
Bandwidth Throttling	<p>These settings control bandwidth throttling. See “Bandwidth Throttling” on page 16.</p> <p>To throttle replications to a specified rate, specify the required throughput value in Mbps and indicate the time pattern.</p>
Recovery Points to Keep	<p>Specify the number of recovery points to keep for workloads that use this Protection Tier. See “Recovery Points” on page 74.</p>
Workload Failure	<p>Specify the number of workload detection attempts before it is considered failed.</p>
Workload Detection	<p>Specify the time interval (in seconds) between workload detection attempts.</p>

5.6 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.

5.7 Initial Replication Method (Full and Incremental)

In workload protection and failback operations, the Initial Replication parameter determines the scope of data transferred from a source to a target.

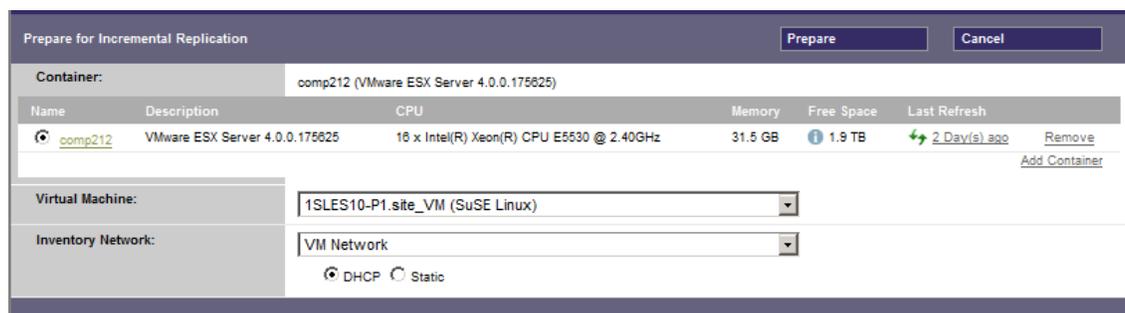
- ♦ **Full:** A full volume transfer takes place from a production workload to its replica (the failover workload), or from a failover workload to its original virtual or physical infrastructure.
- ♦ **Incremental:** Only differences are transferred from a source to its target, provided that they have 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 into the VM container, such as a workload VM physically moved on portable media from the production site to a remote recovery site. For details, see your VMware documentation.

- ◆ 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 61).

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.

- 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 PlateSpin Protect Web Interface displays the Prepare for Incremental Replication page.



- 4 Select the required container, the virtual machine, and the inventory 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 system to assign the workload to.
- 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.

5.8 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 76.

- ◆ **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.

5.9 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 [“Source service/daemon control:”](#) on page 75). 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
```

5.10 Volumes

Upon adding a workload for protection, PlateSpin Protect inventories your source workload's storage media and automatically sets up options in the PlateSpin Protect Web Interface for you to specify the volumes you require for protection.

PlateSpin Protect supports several types of storage, including Windows dynamic disks, LVM (version 2 only), RAID, and SAN.

For Linux workloads, PlateSpin Protect provides the following additional features:

- ◆ 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.
- ◆ (OES 2 workloads) 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.

The following figures show the Replication Settings parameter set for a Linux workload with multiple volumes and two logical volumes in a volume group.

Figure 5-2 Volumes, Logical Volumes, and Volume Groups of a Protected Linux Workload

Tier Settings				
Replication Settings				
Encrypt Data Transfer:	No			
Source Credentials:	root			
Number of CPUs:	1			
Replication Network:	DHCP - VM Network			
Recovery Point Datastore:	Storage2 (689.7 GB free)			
Protected Volumes:	Include	Name	Total Size	Datastore
	<input checked="" type="checkbox"/>	/usr	2.9 GB	Storage2
	<input checked="" type="checkbox"/>	/boot	2.0 GB	Storage2
	<input checked="" type="checkbox"/>	/new2 (EXT3)	151.9 MB	Storage2
Protected Logical Volumes:	Include	Name	Total Size	Volume Group
	<input checked="" type="checkbox"/>	/LogicalVolume1 (EXT3)	484.2 MB	group
	<input checked="" type="checkbox"/>	/LogicalVolume2 (EXT3)	193.7 MB	group
Volume Groups:	Include	Name	Total Size	Datastore
	<input checked="" type="checkbox"/>	group	1016.0 MB	Storage2
Non-volume Storage:	--			
Daemons to Stop During Replication:	--			
Failover Settings				
Prepare for Failover Settings				
Test Failover Settings				
Recovery Points				
Workload Details				

The following figure shows volume protection options of an OES 2 workload with options indicating that the EVMS layout should be preserved and re-created for the failover workload:

Figure 5-3 Replication Settings, Volume-Related Options (OES 2 Workload)

Protected Logical Volumes:	<input checked="" type="checkbox"/>	/ (REISERFS)	2.2 GB	2.2 GB	system
	<input checked="" type="checkbox"/>	/boot (EXT2)	13.0 MB	55.3 MB	/dev/evms/sda1
	<input checked="" type="checkbox"/>	/opt/novell/nss/mnt/pools/NEWPOOL (NSSFS)	23.3 MB	999.6 MB	NEWPOOL
Non-volume Storage:	<input checked="" type="checkbox"/>	/dev/system/swap	Yes	1.48 GB	system
Volume Groups:	<input checked="" type="checkbox"/>	system	5.9 GB	dev-comp124:storage	<input type="checkbox"/>
EVMS Volumes:	<input checked="" type="checkbox"/>	/dev/evms/sda1	70.6 MB	dev-comp124:storage	<input type="checkbox"/>
	<input checked="" type="checkbox"/>	NEWPOOL	1023.0 MB	dev-comp124:storage	<input type="checkbox"/>
Daemons to Stop During Replication:	Add Daemons				

5.11 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](#) parameter set) For separating regular replication traffic from your production traffic.
- ♦ **Failover:** ([Failover](#) parameter set) For the failover workload to become part of your production network when it goes live.
- ♦ **Prepare for Failover:** ([Prepare for Failover](#) network parameter) For network settings during the optional Prepare for Failover stage.
- ♦ **Test Failover:** ([Test Failover](#) parameter set) For network settings to apply to the failover workload during a Test Failover stage.

5.12 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 ISO image.

- ♦ [Section 5.12.1, "Downloading the PlateSpin Boot ISO Image," on page 79](#)
- ♦ [Section 5.12.2, "Injecting Additional Device Drivers into the Boot ISO Image," on page 79](#)
- ♦ [Section 5.12.3, "Registering Physical Machines as Failback Targets with PlateSpin Protect," on page 80](#)

5.12.1 Downloading the PlateSpin Boot ISO Image

You can download the PlateSpin boot ISO images (`bootofx.x2p.iso` for BIOS firmware-based targets and `bootofx.x2p.uefi.iso` for UEFI firmware-based targets) from the PlateSpin Protect area of [NetIQ Downloads \(https://dl.netiq.com\)](https://dl.netiq.com) by doing a search with the following parameters:

- ◆ **Product or Technology:** PlateSpin Protect
- ◆ **Select Version:** PlateSpin Protect 11.0.1
- ◆ **Date Range:** <All Dates>

5.12.2 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:

- 1 Obtain or compile *.ko driver files appropriate for the target hardware manufacturer.

IMPORTANT: Make sure 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 [KB Article 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>	<p><path> specifies the directory that contains the drivers (that is, *.ko files) that you want to inject.</p> <p>On completion of the command, the ISO file is updated with the added drivers.</p>
-c <path>	<path> specifies where a ConfigureTakeControl.xml file resides.
-l [<type>]	<p><type> specifies a subset of drivers you want to list. The default is "all" types.</p> <p>Listed driver types beginning with a forward slash (/) are assumed to be located in <kernel_module_directory>/kernel/</p> <p>Listed driver types without a leading forward slash (/) are assumed to be located in <kernel_module_directory>/kernel/drivers/</p> <p>Driver Subset Examples:</p> <pre>-l scsi -l 'net video' -l '/net net'</pre> <p>Special Usage of this Argument:</p> <p>If you want to list the available subdirectories of each of the subsets, use the argument like this: -l INDEX</p>

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
```

5.12.3 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` (for systems with up to 512 MB RAM)
 - ◆ `ps64_512m` (for systems with more than 512 MB RAM)
- 5 Press Enter.
- 6 When you are prompted, enter the hostname 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 hostname 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` if you have enabled SSL, and `N` if you have not.

After a few minutes, the physical machine should be available in the failback settings of the PlateSpin Protect Web Interface.

5.13 Advanced Workload Protection Topics

- ◆ [Section 5.13.1, “Protecting Windows Clusters,” on page 81](#)
- ◆ [Section 5.13.2, “Using Workload Protection Features through the PlateSpin Protect Web Services API,” on page 83](#)

5.13.1 Protecting Windows Clusters

PlateSpin Protect supports the protection of a Microsoft Windows cluster’s business services. The supported clustering technologies are:

- ◆ Windows 2008 R2 Server-based Microsoft Failover Cluster

This section includes the following information:

- ◆ [“Workload Protection” on page 82](#)
- ◆ [“Protection Failover” on page 83](#)
- ◆ [“Protection Failback” on page 83](#)

NOTE: For information about rebuilding the Windows 2008/2008R2 Failover Cluster environment after it has been protected by PlateSpin Protect during failover/failback, see the [KB Article](#) describing the procedure.

Workload Protection

Protection of a 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.

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

- ◆ When you perform an **Add Workload** operation, you must identify the active node—the node that currently owns the quorum resource of the cluster—identified by the cluster’s IP address (*virtual IP address*). Specifying the IP address of an individual node results in that node being inventoried as a regular, cluster-unaware Windows workload.
- ◆ A cluster’s quorum resource must be co-located with the cluster’s resource group (service) being protected.

When using block-based transfer, the block-based driver components are not installed on the cluster nodes. The block-based transfer occurs using a driverless sync with an MD5 based replication. Because the block-based driver is not installed, no reboot is required on the source cluster nodes.

NOTE: File-based transfer is not supported for protecting Microsoft Windows clusters.

If the new cluster node becomes active between incremental replications of a protected cluster and if the new active node’s profile is similar to the failed active node, the protection contract continues as scheduled; otherwise, the command fails. The profiles of cluster nodes are considered similar if:

- ◆ They have the same number of volumes.
- ◆ Each volume is exactly the same size on each node.
- ◆ They have an identical number of network connections.
- ◆ Serial numbers for local volumes (System volume and System Reserved volume) must be the same on each cluster node.

If the local drives on each node of the cluster have different serial numbers, you cannot run an incremental replication after the active node “flips” in the event of a node failure. For example, the active node is Node 1, and it then “flips” to Node 2).

There are two supported options to support clusters in this scenario for Protect 11.0.1:

- ◆ (Recommended) Use the customized *Volume Manager* utility to change the local volume serial numbers to match each node of the cluster. For more information, see [Appendix B, “Synchronizing Serial Numbers on Cluster Node Local Storage,” on page 119](#).
- ◆ (Conditional and Optional) If you see this error:

```
Volume mappings does not contain source serial number: xxxx-xxxx,
```

it might have been caused by a change in the active node prior to running the incremental replication. In this case, you can run a full replication to ensure the cluster is once again protected. Incremental replications should function again after the full replication.

If you choose not to match the volume serial numbers on each node in the cluster, then a full replication is required before each incremental replication whenever the active node performs a “flip” to a new node in the cluster.

If a new node becomes active prior to the completion of the copy process during a full or an incremental replication, the command aborts and a message displays indicating that the replication needs to be re-run.

To protect a Windows cluster, follow the normal workload protection workflow (see [“Basic Workflow for Workload Protection and Recovery”](#) on page 47).

Protection Failover

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

To failover (or to test the 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.)

Protection Failback

Only a failback using full replication for Windows Cluster workloads is supported for this release.

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

- ◆ Map all disks on the Failover machine to a single local disk on the failback target.
- ◆ Add another disk (`Disk 2`) to the physical failback machine. You can then configure the failback operation to restore the failover's system volume to `Disk 1` and the failover'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 failback is complete, you can rejoin other nodes to the newly restored cluster.

5.13.2 Using Workload Protection Features through the PlateSpin Protect Web Services API

You can use workload protection functionality programmatically through the `protectionservices` API from within your applications. You can use any programming or scripting language that supports an HTTP client and JSON serialization framework.

```
https://<hostname | IP_address>/protectionservices
```

Replace `<hostname | IP_address>` with the hostname or the IP address of your PlateSpin Server host. If SSL is not enabled, use `http` in the URI.

Figure 5-4 The Front Page of the Protect Server API

PlateSpin Protect Server API

Version 4.0.0.1161

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](#)
- [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.

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.

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

Protect administrators can leverage a Jscript sample (<https://localhost/protectionservices/Documentation/Samples/protect.js>) from the command line to access the product through the API. 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

The PlateSpin Protect REST API home page (<https://localhost/protectionservices/> or <https://<server page>/protectionservices/>) includes links to the content that can be useful for developers and administrators.

This technology preview will be fully developed with more features in subsequent releases.

6 Auxiliary Tools for Working with Physical Machines

Your PlateSpin Protect distribution includes tools for use when working with physical machines as failback targets.

- ◆ [Section 6.1, “Managing Device Drivers,” on page 87](#)

6.1 Managing Device Drivers

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

The following sections provide more details:

- ◆ [Section 6.1.1, “Packaging Device Drivers for Windows Systems,” on page 87](#)
- ◆ [Section 6.1.2, “Packaging Device Drivers for Linux Systems,” on page 88](#)
- ◆ [Section 6.1.3, “Uploading Drivers to the PlateSpin Protect Device Driver Database,” on page 88](#)
- ◆ [Section 6.1.4, “Using the Plug and Play \(PnP\) ID Translator Feature,” on page 90](#)

6.1.1 Packaging Device Drivers for Windows Systems

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

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

The drivers are now ready for upload. See [“Uploading Drivers to the PlateSpin Protect Device Driver Database” on page 88](#).

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

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

6.1.2 Packaging Device Drivers for Linux Systems

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

- 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:

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

or this command for UEFI firmware-based targets:

```
# mkdir /mnt/ps # mount -o loop /root/bootofx.x2p.uefi.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 Drivers to the PlateSpin Protect Device Driver Database” on page 88](#).

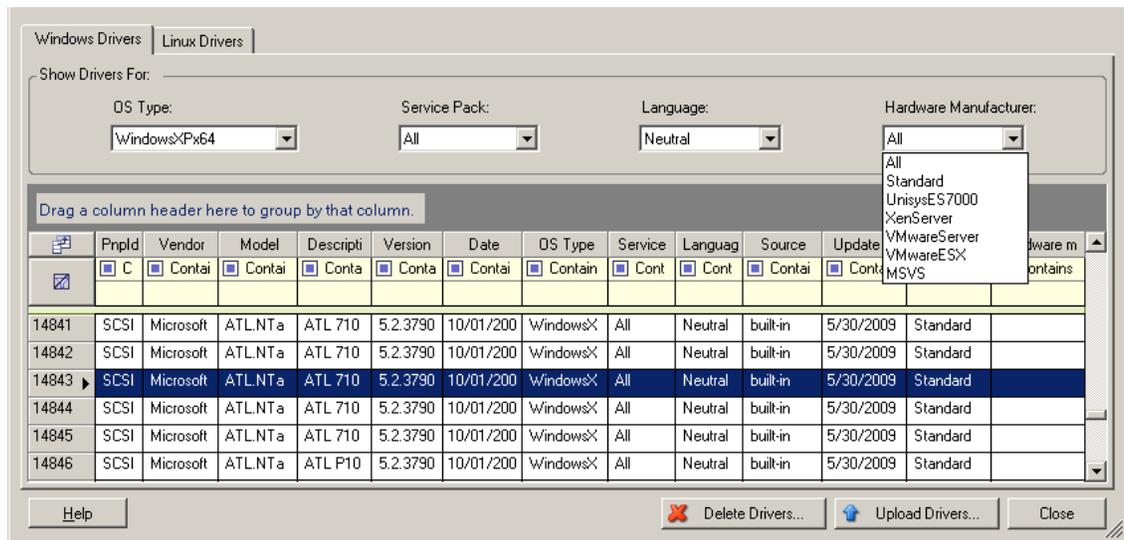
6.1.3 Uploading Drivers to the PlateSpin Protect Device Driver Database

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

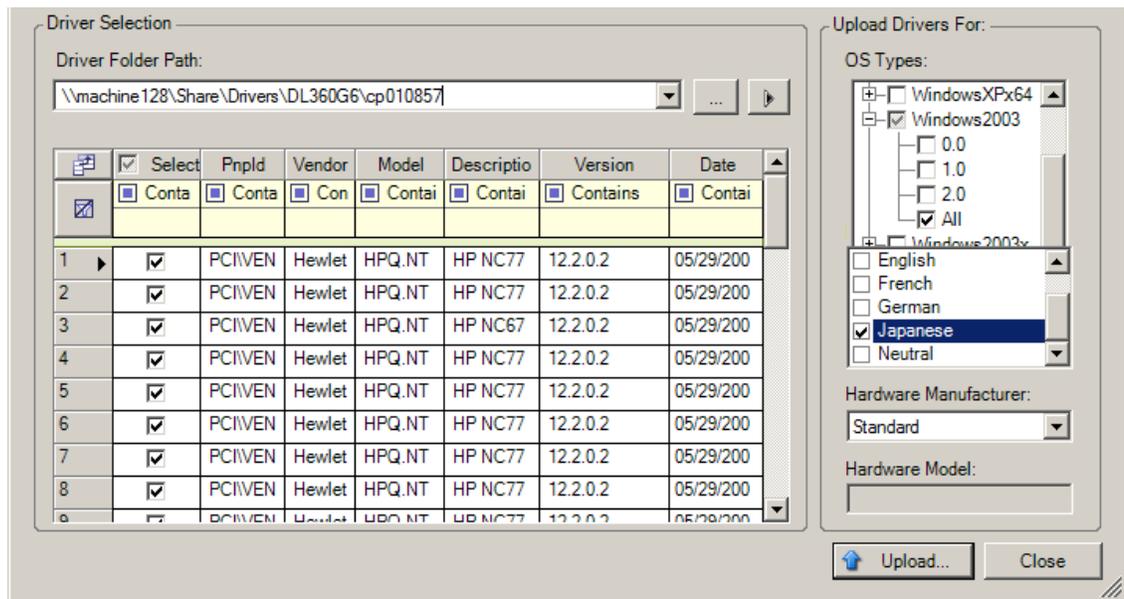
NOTE: On upload, PlateSpin Protect does not validate drivers against selected operating system types or their bit specifications; make sure that you only upload drivers that are appropriate for your target infrastructure.

Device Driver Upload Procedure (Windows)

- 1 Obtain and prepare the required device drivers. See [Packaging Device Drivers for Windows Systems](#).
- 2 On your PlateSpin Server host, under `\Program Files\PlateSpin Protect Server\DriverManager`, start the `DriverManager.exe` program and select the **Windows Drivers** tab.



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

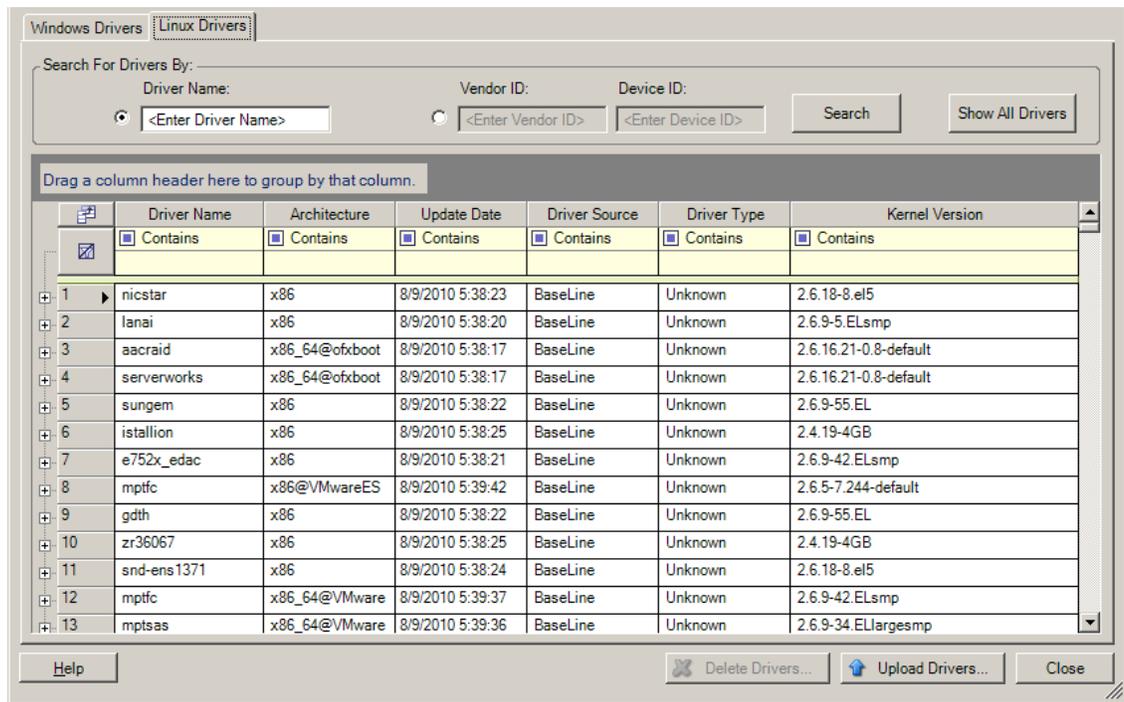


Select **Standard** as the **Hardware Manufacturer** option, unless your drivers are designed specifically for any of the target environments listed.

- 4 Click **Upload** and confirm your selections when prompted.
The system uploads the selected drivers to the driver database.

Device Driver Upload Procedure (Linux)

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



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

The system uploads the selected drivers to the driver database.

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

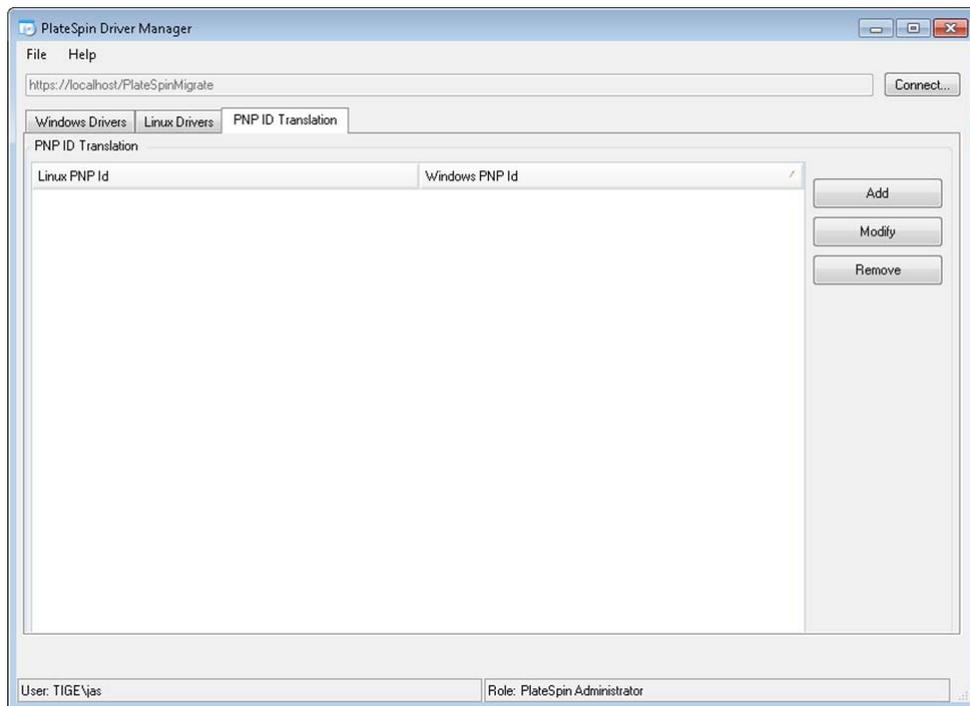
“Plug and Play” (PnP) refers to Windows operating system functionality that supports connectivity, configuration, and management with native plug and play devices. In Windows, the feature facilitates discovery of PnP compliant hardware devices attached to a PnP compliant bus. PnP compliant devices are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

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

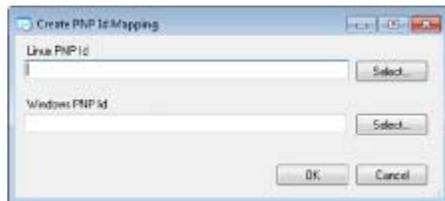
These IDs are formatted consistently, so PlateSpin can apply a standard transformation to each of them to determine its corresponding Windows PnP ID. The translation occurs automatically within the PlateSpin product. The feature lets you or a support technician add, edit or remove custom PnP mappings.

Follow these steps to use the PnP ID Translation feature:

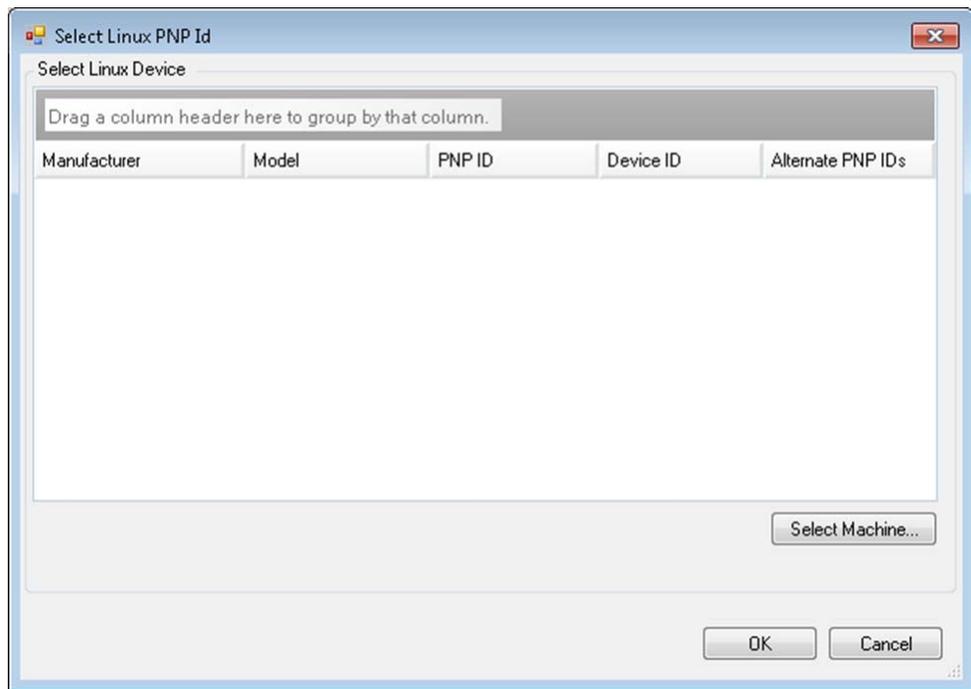
- 1 Launch the PlateSpin Driver Manager tool and connect to the PlateSpin Server.
- 2 In the Driver Manager tool, select the PNP ID Translation tab to open the **PNP ID Translation** list, which includes the currently known custom PnP ID mappings.



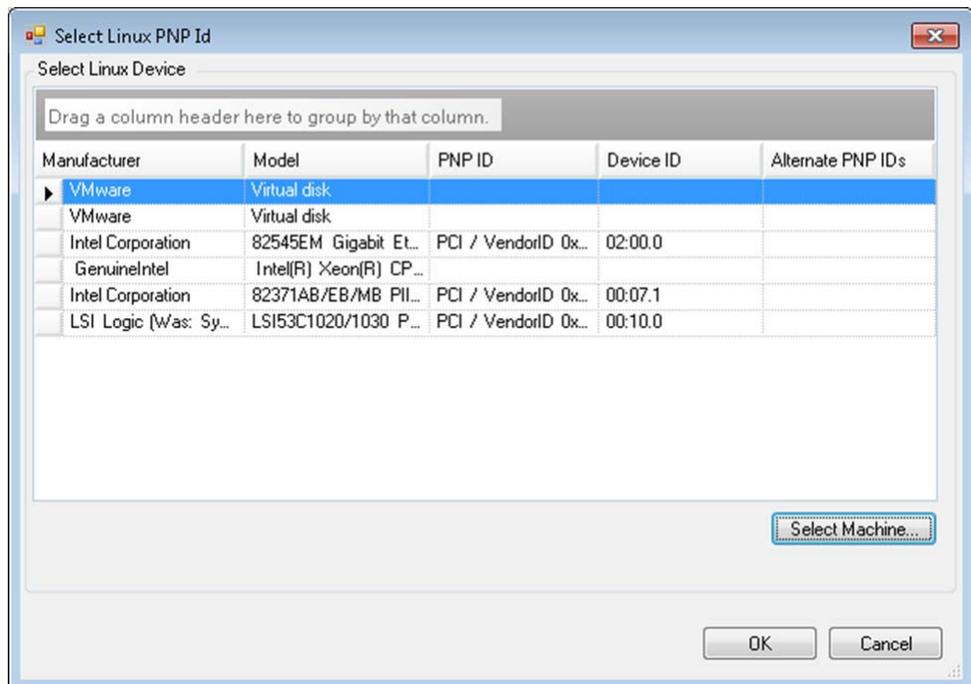
- 3 On the list page, click **Add** to display the Create PNP ID Mapping dialog box.



- 4 In the **Linux PNP ID** field, add a Linux PnP ID.
 - 4a (Conditional) If you know it, type the Linux PnP ID you want to use.
or
 - 4b (Conditional) Select an ID from a previously discovered workload:
 - 4b1 Adjacent to the **Linux PnP ID** field, click **Select** to open the Select Linux PnP ID dialog box.



- 4b2 On the dialog box, click **Select Machine** to display a list of the machines previously discovered by the PlateSpin Linux RAM disk.
- 4b3 Highlight one of the devices in the list, then click **Select** to populate the list in the Select Linux PnP ID dialog box.



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

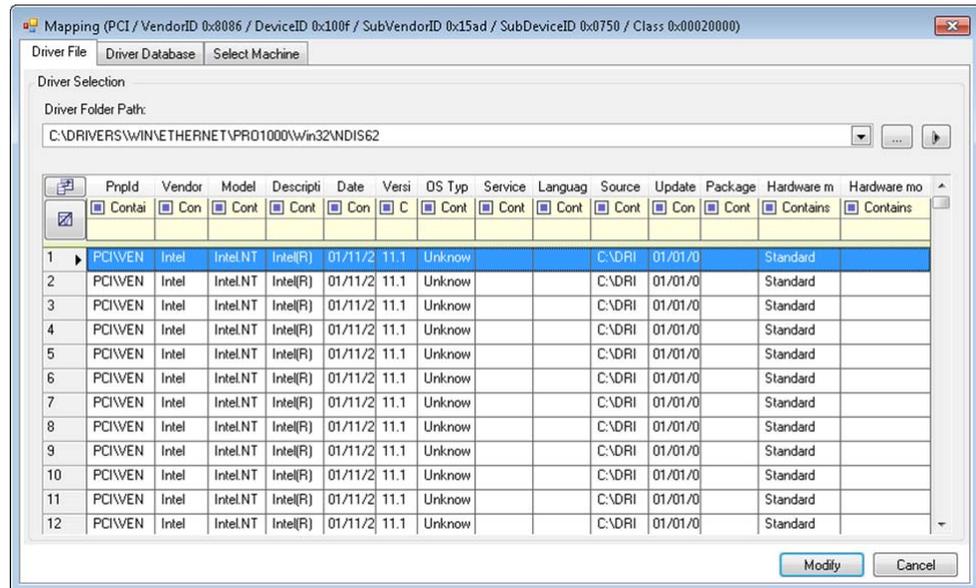
5 In the **Windows PNP ID** field, add a Windows PnP ID:

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

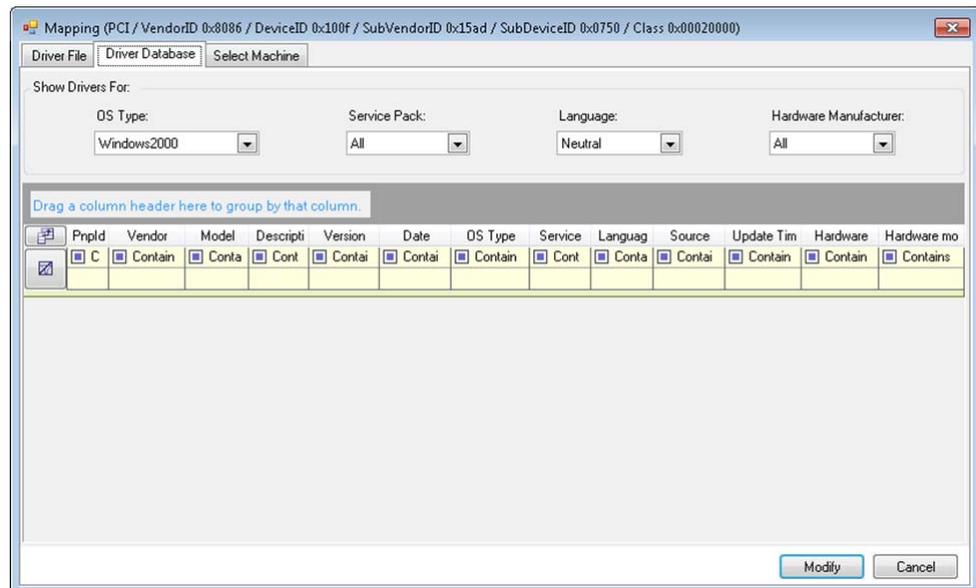
or

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

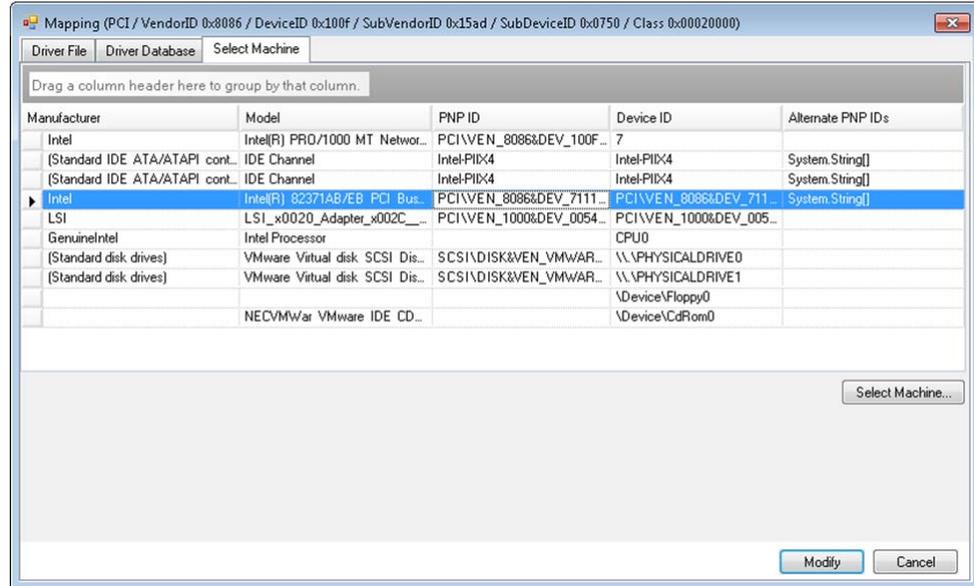
- ◆ Under the **Driver File** tab, browse to and select a Windows driver file (that is, a file with the *.inf extension), select the desired PnP ID, then click **Modify**.



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

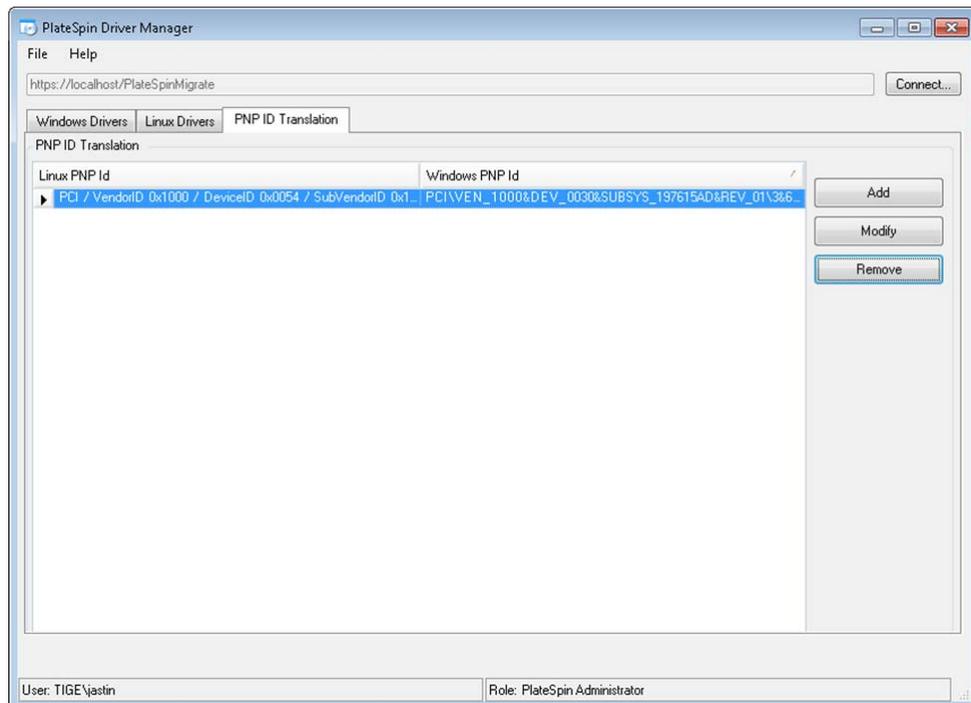


- Under the **Select Machine** tab, click **Select Machine**, then, from the list of Windows machines discovered using live discovery, select a machine, click **OK** to display its devices, select the desired PnP ID, then click **Modify**.



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

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



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

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

To modify,

7a Click **Modify** to open the Create PNP id Mapping dialog box.

7b Repeat [Step 5 on page 93](#) to modify the Windows PnP ID.

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

7 Troubleshooting

This section includes the following information:

- ◆ [Section 7.1, “Troubleshooting Workload Inventory \(Windows\),” on page 97](#)
- ◆ [Section 7.2, “Troubleshooting Workload Inventory \(Linux\),” on page 101](#)
- ◆ [Section 7.3, “Troubleshooting Problems during the Prepare Replication Command \(Windows\),” on page 101](#)
- ◆ [Section 7.4, “Troubleshooting Workload Replication,” on page 102](#)
- ◆ [Section 7.5, “Troubleshooting Traffic-forwarding Workloads,” on page 103](#)
- ◆ [Section 7.6, “Troubleshooting Online Help,” on page 104](#)
- ◆ [Section 7.7, “Generating and Viewing Diagnostic Reports,” on page 104](#)
- ◆ [Section 7.8, “Removing Workloads,” on page 104](#)
- ◆ [Section 7.9, “Post-Protection Workload Cleanup,” on page 105](#)
- ◆ [Section 7.10, “Shrinking the PlateSpin Protect Databases,” on page 107](#)

7.1 Troubleshooting Workload Inventory (Windows)

You might need to troubleshoot the following common problems during the workload inventory.

Problems or Messages	Solutions
The domain in the credentials is invalid or blank	<p>This error occurs when the Credential Format is incorrect.</p> <p>Try the discovery by using a local administrator account with the credential format <code>hostname\LocalAdmin</code></p> <p>Or, try the discovery by using a domain administrator account with the credential format <code>domain\DomainAdmin</code></p>
Unable to connect to Windows server...Access is denied	<p>An invalid account was used when trying to add a workload. Use an administrator account or add the user to the administrators group and try again.</p> <p>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 99 again. If the test succeeds, try adding the workload again.</p> <ul style="list-style-type: none">◆ “Troubleshooting DCOM Connectivity” on page 99◆ “Troubleshooting RPC Service Connectivity” on page 99
Unable to connect to Windows server...The network path was not found	<p>Network connectivity failure. Perform the tests in “Performing Connectivity Tests” on page 98. If a test fails, ensure that PlateSpin Protect and the workload are on the same network. Reconfigure the network and try again.</p>

Problems or Messages	Solutions
"Discover Server Details {hostname}" Failed Progress: 0% Status: NotStarted	This error can occur for several reasons and each has a unique solution: <ul style="list-style-type: none"> ◆ For environments using a local proxy with authentication, bypass the proxy or add the proper permissions. See KB 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 KB 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: <ul style="list-style-type: none"> ◆ 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 99. ◆ 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 PlateSpin Protect can access those shares. See "Enabling File/Share Permissions and Access" on page 100. ◆ 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.

The following sections provide more troubleshooting information on Windows workloads:

- ◆ [Section 7.1.1, "Performing Connectivity Tests," on page 98](#)
- ◆ [Section 7.1.2, "Disabling Antivirus Software," on page 99](#)
- ◆ [Section 7.1.3, "Enabling File/Share Permissions and Access," on page 100](#)

7.1.1 Performing Connectivity Tests

- ◆ ["Network Connectivity Test" on page 98](#)
- ◆ ["WMI Connectivity Test" on page 99](#)
- ◆ ["Troubleshooting DCOM Connectivity" on page 99](#)
- ◆ ["Troubleshooting RPC Service Connectivity" on page 99](#)

Network Connectivity Test

Perform this basic network connectivity test to determine whether PlateSpin 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 hostname 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 PlateSpin 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 (2003/2008/2012/Win 8.x), 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.
- 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 PlateSpin Protect and the workload on the same side of the firewall
- ♦ Opening up specific ports between PlateSpin Protect and the workload (See [“Access and Communication Requirements across your Protection Network”](#) on page 24).

7.1.2 Disabling Antivirus Software

Antivirus software might occasionally block some of the PlateSpin 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 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 only disabled for the duration of the file transfer, and are restarted when the process completes. This is not necessary during block-level data replication.

7.1.3 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, PlateSpin Protect uses the workload's administrative shares. PlateSpin 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 PlateSpin 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.

7.2 Troubleshooting Workload Inventory (Linux)

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	<p>This message has a number of possible causes:</p> <ul style="list-style-type: none">◆ 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. <p>For network and access requirements for a workload, see "Access and Communication Requirements across your Protection Network" on page 24.</p>
Access denied	<p>This authentication problem indicates either an invalid username or password. For information on proper workload access credentials, see "Guidelines for Workload and Container Credentials" on page 66.</p>

7.3 Troubleshooting Problems during the Prepare Replication Command (Windows)

Problems or Messages	Solutions
Authentication error when verifying the controller connection while setting up the controller on the source.	<p>The account used to add a workload needs to be allowed by this policy. See "Group Policy and User Rights" on page 101.</p>
Failure to determine whether .NET Framework is installed (with exception The trust relationship between this workstation and the primarydomain failed).	<p>Check whether the Remote Registry service on the source is enabled and started. See also "Troubleshooting Workload Inventory (Windows)" on page 97.</p>

7.3.1 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` (for Windows 2003/XP) or `secedit /refreshpolicy machine_policy /enforce` (for Windows 2000).

7.4 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.
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.
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: <code>ping ip-t</code> (replace <i>ip</i> 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 32.
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 <code>netstat -ano</code> on the workload. Check the firewall. Retry the replication.

Problems or Messages	Solutions
<p>Controller connection not established</p> <p>Replication fails at the Take Control of Virtual Machine step.</p>	<p>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.</p> <p>Change the replication IP to a static IP or enable the DHCP server.</p> <p>Ensure that the virtual network selected for replication is routable to the PlateSpin Server host.</p>
<p>Replication job does not start (stuck at 0%)</p>	<p>This error can occur for different reasons and each has a unique solution:</p> <ul style="list-style-type: none"> ◆ For environments using a local proxy with authentication, bypass the proxy or add proper permissions to resolve this problem. See KB Article 20339 (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 KB Article 7920862 (https://www.netiq.com/support/kb/doc.php?id=7920862). <p>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 101.</p>
<p>After a Windows Update, some files in the C:\Windows\SoftwareDistribution folder are not transferred to the target machine during incremental file-based replication.</p>	<p>This is a Microsoft Windows common practice: For optimization purposes, some files are marked for deletion in the <code>HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\BackupRestore\FilesNotToSnapshot</code> registry key to prevent them from being included in VSS snapshots. See the Microsoft Developer Network article, Excluding Files from Shadow Copies for more information.</p> <p>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 <code>SoftwareDistribution</code> folder.</p>

7.5 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 [KB Article 7005495](#).

7.6 Troubleshooting Online Help

On some systems with enhanced browser security settings (such as Internet Explorer 8 on Windows Server 2008), the Expand and Collapse icons (+ and -) in the Table of Contents might fail to work. To fix the issue, enable JavaScript in your browser:

- ♦ **Internet Explorer:** Click **Tools > Internet Options > Security** tab > **Internet** zone > **Custom level**, then select the **Enable** option for the **Active Scripting** feature.
- ♦ **Firefox:** Click **Tools > Options > Content** tab, then select the **Enable JavaScript** option.

7.7 Generating and Viewing 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 Diagnostics** link.

The screenshot shows the 'Command Details' page for a replication command. The command is titled 'Running First Replication' and is currently in a 'Running' state. The duration is 14h 49m 6s. The progress bar indicates that 80% of the data has been copied. The table below the command summary lists the steps of the replication process. The 'Copy data' step is currently running at 80% completion. A 'Generate Diagnostics' link is visible in the bottom right corner of the table, highlighted with a red box.

Step	Status	Start Time	End Time	Duration	Diagnostics
Copy data	Running (80%)	3/31/2010 8:24 PM	--	14h 48m 53s	--

After a few moments, the page refreshes and displays a **View** link above the **Generated Diagnostics** link.

- 2 Click **View**.
A new page opens with comprehensive diagnostic information about the current command.
- 3 Save the diagnostics page and have it ready if you need to contact technical support.

7.8 Removing Workloads

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

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

(Conditional) For Windows workloads previously protected through block-level replication, the PlateSpin Protect 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. Make sure you carry out this operation during scheduled downtime.

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

Wait for the process to complete.

7.9 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.

The following sections include more information:

- ◆ [Section 7.9.1, “Cleaning Up Windows Workloads,” on page 105](#)
- ◆ [Section 7.9.2, “Cleaning Up Linux Workloads,” on page 106](#)

7.9.1 Cleaning Up Windows Workloads

Component	Removal Instructions
PlateSpin Block-Based Transfer Component	See KB Article 7005616 (https://www.netiq.com/support/kb/doc.php?id=7005616) .
Third-party Block-based Transfer Component (discontinued)	<ol style="list-style-type: none">1. Use the Windows Add/Remove Programs applet (run <code>appwiz.cpl</code>) and remove the component. Depending on the source, you might have either of the following versions:<ul style="list-style-type: none">◆ SteelEye Data Replication for Windows v6 Update2◆ SteelEye DataKeeper For Windows v72. Reboot the machine.
File-based Transfer Component	At root level for each volume under protection, remove all files named <code>PlateSpinCatalog*.dat</code>
Workload Inventory software	In the workload's <code>Windows</code> directory: <ul style="list-style-type: none">◆ Remove all files named <code>machinediscovery*</code>.◆ Remove the subdirectory named <code>platespin</code>.

Component	Removal Instructions
Controller software	<ol style="list-style-type: none"> Open a command prompt and change the current directory to: <ul style="list-style-type: none"> \Program Files\platespin* (32-bit systems) \Program Files (x86)\platespin* (64-bit systems) Run the following command: <pre>ofxcontroller.exe /uninstall</pre> Remove the platespin* directory

7.9.2 Cleaning Up Linux Workloads

Component	Removal Instructions
Controller software	<ul style="list-style-type: none"> Kill these processes: <ul style="list-style-type: none"> <code>pkill -9 ofxcontrollerd</code> <code>pkill -9 ofxjobexec</code> remove the OFX controller rpm package: <pre>rpm -e ofxcontrollerd</pre> In the workload's file system, remove the <code>/usr/lib/ofx</code> directory with its contents.
Block-level data transfer software	<ol style="list-style-type: none"> Check if the driver is active: <pre>lsmod grep blkwatch</pre> <p>If the driver is still loaded in memory, the result should contain a line, similar to the following:</p> <pre>blkwatch_7616 70924 0</pre> (Conditional) If the driver is still loaded, remove it from memory: <pre>rmmod blkwatch_7616</pre> Remove the driver from the boot sequence: <pre>blkconfig -u</pre> Remove the driver files by deleting the following directory with its contents: <pre>/lib/modules/[Kernel_Version]/Platespin</pre> Delete the following file: <pre>/etc/blkwatch.conf</pre>

Component	Removal Instructions
LVM snapshots	<p>LVP snapshots used by ongoing replications are named according to a <code>volume_name-PS-snapshot</code> convention. For example, a snapshot of a LogVol01 volume will be named <code>LogVol01-PS-snapshot</code>.</p> <p>To remove these LVM snapshots:</p> <ol style="list-style-type: none"> 1. Generate a list of snapshot on the required workload by using one of the following ways: <ul style="list-style-type: none"> ◆ Use the PlateSpin Protect 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: <pre># lvsdisplay -a</pre> 2. Note the names and locations of the snapshots you want to remove. 3. Remove the snapshots by using the following command: <pre>lvremove <i>snapshot_name</i></pre>
Bitmap files	<p>For each volume under protection, at the root of the volume, remove the corresponding <code>.blocks_bitmap</code> file.</p>
Tools	<p>On the source workload, under <code>/sbin</code>, remove the following files:</p> <ul style="list-style-type: none"> ◆ <code>bmaputil</code> ◆ <code>blkconfig</code>

7.10 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. [KB Article 7006458 \(https://www.netiq.com/support/kb/doc.php?id=7006458\)](https://www.netiq.com/support/kb/doc.php?id=7006458) explains the location of the tool and the options available for it, should you decide to use it for offline database operations.

A 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). This section includes the following information:

- ♦ [Section A.1, “Analyzing Your Linux Workload,” on page 109](#)
- ♦ [Section A.2, “Pre-compiled “blkwatch” driver \(Linux\),” on page 110](#)

NOTE: Currently, when you add a Linux workload for protection, the product interface might display a message informing you that the system is

```
Unable to find a compiled version of the block-based kernel driver to exactly match the kernel on this Linux workload....
```

The message advises further to build a block based driver upon the next replication.

This validator message is likely inaccurate, due to the addition of hundreds of newly-supported Linux distributions for which Protect has a pre-compiled version of the `blkwatch` driver. If your workload (that is, the Linux distribution) matches an entry on [the list in this section](#), you can ignore the message and continue.

A.1 Analyzing Your Linux Workload

Prior to determining whether PlateSpin Protect has a `blkwatch` driver for your distribution, you need to learn more about the kernel of your Linux workload so that you can use it as a search term against the list of supported distributions. This section includes the following information:

- ♦ [Section A.1.1, “Determining the Release String,” on page 109](#)
- ♦ [Section A.1.2, “Determining the Architecture,” on page 110](#)

A.1.1 Determining the Release String

You can determine the release string of the kernel of your Linux workload by running the following command at the workload’s Linux terminal:

```
uname -r
```

For example, if you run `uname -r`, you might see the following output:

```
3.0.76-0.11-default
```

If you search the list of distributions, you see there are two entries that match this string:

- ♦ `SLES11SP3-GA-3.0.76-0.11-default-x86`
- ♦ `SLES11SP3-GA-3.0.76-0.11-default-x86_64`

The search results indicate that the product has drivers for both 32-bit (x86) and 64-bit (x86_64) architectures.

A.1.2 Determining the Architecture

You can determine the architecture of your Linux workload by running the following command at the workload's Linux terminal:

```
uname -m
```

For example, if you run `uname -m`, you might see the following output:

```
x86_64
```

With this information, you can determine that the workload has 64-bit architecture.

A.2 Pre-compiled “blkwatch” driver (Linux)

Following is a list of non-debug Linux distributions for which Protect has a `blkwatch` driver. You can search the list to determine if the release string and architecture of your Linux workload kernel matches a supported distribution in the list. If you find your release string and architecture, PlateSpin Protect has a pre-compiled version the `blkwatch` driver.

If your search is unsuccessful, you can create a custom `blkwatch` driver by following the steps found in the Knowledgebase Article [KB 7005873](#).

List Item Syntax

Each item in the list is formatted using the following syntax:

```
<Distro>-<Patch>-<Kernel_Release_String>-<Kernel_Architecture>
```

So, for a SLES 9 SP1 distribution with a kernel release string of `2.6.5-7.139-bigsm` for 32-bit (x86) architecture, the item is listed in a format like this:

```
SLES9-SP1-2.6.5-7.139-bigsm-x86
```

List of Distributions

```
RHEL4-GA-2.6.9-5.EL-x86
RHEL4-GA-2.6.9-5.EL-x86_64
RHEL4-GA-2.6.9-5.ELhugemem-x86
RHEL4-GA-2.6.9-5.ELsmp-x86
RHEL4-GA-2.6.9-5.ELsmp-x86_64
RHEL4-U1-2.6.9-11.EL-x86
RHEL4-U1-2.6.9-11.EL-x86_64
RHEL4-U1-2.6.9-11.ELhugemem-x86
RHEL4-U1-2.6.9-11.ELsmp-x86
RHEL4-U1-2.6.9-11.ELsmp-x86_64
RHEL4-U2-2.6.9-22.EL-x86
RHEL4-U2-2.6.9-22.EL-x86_64
RHEL4-U2-2.6.9-22.ELhugemem-x86
RHEL4-U2-2.6.9-22.ELsmp-x86
RHEL4-U2-2.6.9-22.ELsmp-x86_64
RHEL4-U3-2.6.9-34.EL-x86
RHEL4-U3-2.6.9-34.EL-x86_64
RHEL4-U3-2.6.9-34.ELhugemem-x86
RHEL4-U3-2.6.9-34.ELlargesmp-x86_64
RHEL4-U3-2.6.9-34.ELsmp-x86
RHEL4-U3-2.6.9-34.ELsmp-x86_64
RHEL4-U4-2.6.9-42.EL-x86
```

RHEL4-U4-2.6.9-42.EL-x86_64
RHEL4-U4-2.6.9-42.ELhugemem-x86
RHEL4-U4-2.6.9-42.ELlargesmp-x86_64
RHEL4-U4-2.6.9-42.ELsmp-x86
RHEL4-U4-2.6.9-42.ELsmp-x86_64
RHEL4-U5-2.6.9-55.EL-x86
RHEL4-U5-2.6.9-55.EL-x86_64
RHEL4-U5-2.6.9-55.ELhugemem-x86
RHEL4-U5-2.6.9-55.ELlargesmp-x86_64
RHEL4-U5-2.6.9-55.ELsmp-x86
RHEL4-U5-2.6.9-55.ELsmp-x86_64
RHEL4-U6-2.6.9-67.EL-x86
RHEL4-U6-2.6.9-67.EL-x86_64
RHEL4-U6-2.6.9-67.ELhugemem-x86
RHEL4-U6-2.6.9-67.ELlargesmp-x86_64
RHEL4-U6-2.6.9-67.ELsmp-x86
RHEL4-U6-2.6.9-67.ELsmp-x86_64
RHEL4-U7-2.6.9-78.EL-x86
RHEL4-U7-2.6.9-78.EL-x86_64
RHEL4-U7-2.6.9-78.ELhugemem-x86
RHEL4-U7-2.6.9-78.ELlargesmp-x86_64
RHEL4-U7-2.6.9-78.ELsmp-x86
RHEL4-U7-2.6.9-78.ELsmp-x86_64
RHEL4-U8-2.6.9-89.EL-x86
RHEL4-U8-2.6.9-89.EL-x86_64
RHEL4-U8-2.6.9-89.ELhugemem-x86
RHEL4-U8-2.6.9-89.ELlargesmp-x86_64
RHEL4-U8-2.6.9-89.ELsmp-x86
RHEL4-U8-2.6.9-89.ELsmp-x86_64
RHEL4-U9-2.6.9-100.EL-x86
RHEL4-U9-2.6.9-100.EL-x86_64
RHEL4-U9-2.6.9-100.ELhugemem-x86
RHEL4-U9-2.6.9-100.ELlargesmp-x86_64
RHEL4-U9-2.6.9-100.ELsmp-x86
RHEL4-U9-2.6.9-100.ELsmp-x86_64
RHEL5-GA-2.6.18-8.el5-x86
RHEL5-GA-2.6.18-8.el5-x86_64
RHEL5-GA-2.6.18-8.el5PAE-x86
RHEL5-U1-2.6.18-53.el5-x86
RHEL5-U1-2.6.18-53.el5-x86_64
RHEL5-U1-2.6.18-53.el5PAE-x86
RHEL5-U10-2.6.18-371.el5-x86
RHEL5-U10-2.6.18-371.el5-x86_64
RHEL5-U10-2.6.18-371.el5PAE-x86
RHEL5-U2-2.6.18-92.el5-x86
RHEL5-U2-2.6.18-92.el5-x86_64
RHEL5-U2-2.6.18-92.el5PAE-x86
RHEL5-U3-2.6.18-128.el5-x86
RHEL5-U3-2.6.18-128.el5-x86_64
RHEL5-U3-2.6.18-128.el5PAE-x86
RHEL5-U4-2.6.18-164.el5-x86
RHEL5-U4-2.6.18-164.el5-x86_64
RHEL5-U4-2.6.18-164.el5PAE-x86
RHEL5-U5-2.6.18-194.el5-x86
RHEL5-U5-2.6.18-194.el5-x86_64
RHEL5-U5-2.6.18-194.el5PAE-x86
RHEL5-U6-2.6.18-238.el5-x86
RHEL5-U6-2.6.18-238.el5-x86_64
RHEL5-U6-2.6.18-238.el5PAE-x86
RHEL5-U7-2.6.18-274.el5-x86

RHEL5-U7-2.6.18-274.el5-x86_64
RHEL5-U7-2.6.18-274.el5PAE-x86
RHEL5-U8-2.6.18-308.el5-x86
RHEL5-U8-2.6.18-308.el5-x86_64
RHEL5-U8-2.6.18-308.el5PAE-x86
RHEL5-U9-2.6.18-348.el5-x86
RHEL5-U9-2.6.18-348.el5-x86_64
RHEL5-U9-2.6.18-348.el5PAE-x86
RHEL6-GA-2.6.32-71.el6.i686-x86
RHEL6-GA-2.6.32-71.el6.x86_64-x86_64
RHEL6-U1-2.6.32-131.0.15.el6.i686-x86
RHEL6-U1-2.6.32-131.0.15.el6.x86_64-x86_64
RHEL6-U2-2.6.32-220.el6.i686-x86
RHEL6-U2-2.6.32-220.el6.x86_64-x86_64
RHEL6-U3-2.6.32-279.el6.i686-x86
RHEL6-U3-2.6.32-279.el6.x86_64-x86_64
RHEL6-U4-2.6.32-358.el6.i686-x86
RHEL6-U4-2.6.32-358.el6.x86_64-x86_64
RHEL6-U5-2.6.32-431.el6.i686-x86
RHEL6-U5-2.6.32-431.el6.x86_64-x86_64
SLES10-GA-2.6.16.21-0.8-bigsmp-x86
SLES10-GA-2.6.16.21-0.8-default-x86
SLES10-GA-2.6.16.21-0.8-default-x86_64
SLES10-GA-2.6.16.21-0.8-smp-x86
SLES10-GA-2.6.16.21-0.8-smp-x86_64
SLES10-GA-2.6.16.21-0.8-xen-x86
SLES10-GA-2.6.16.21-0.8-xen-x86_64
SLES10-GA-2.6.16.21-0.8-xenpae-x86
SLES10-SP1-2.6.16.46-0.12-bigsmp-x86
SLES10-SP1-2.6.16.46-0.12-default-x86
SLES10-SP1-2.6.16.46-0.12-default-x86_64
SLES10-SP1-2.6.16.46-0.12-smp-x86
SLES10-SP1-2.6.16.46-0.12-smp-x86_64
SLES10-SP1-2.6.16.46-0.12-xen-x86
SLES10-SP1-2.6.16.46-0.12-xen-x86_64
SLES10-SP1-2.6.16.46-0.12-xenpae-x86
SLES10-SP2-2.6.16.60-0.21-bigsmp-x86
SLES10-SP2-2.6.16.60-0.21-default-x86
SLES10-SP2-2.6.16.60-0.21-default-x86_64
SLES10-SP2-2.6.16.60-0.21-smp-x86
SLES10-SP2-2.6.16.60-0.21-smp-x86_64
SLES10-SP2-2.6.16.60-0.21-xen-x86
SLES10-SP2-2.6.16.60-0.21-xen-x86_64
SLES10-SP2-2.6.16.60-0.21-xenpae-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-bigsmp-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-default-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-default-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-smp-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-smp-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xen-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xen-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xenpae-x86
SLES10-SP3-2.6.16.60-0.54.5-bigsmp-x86
SLES10-SP3-2.6.16.60-0.54.5-default-x86
SLES10-SP3-2.6.16.60-0.54.5-default-x86_64
SLES10-SP3-2.6.16.60-0.54.5-smp-x86
SLES10-SP3-2.6.16.60-0.54.5-smp-x86_64
SLES10-SP3-2.6.16.60-0.54.5-xen-x86
SLES10-SP3-2.6.16.60-0.54.5-xen-x86_64
SLES10-SP3-2.6.16.60-0.54.5-xenpae-x86

SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-bigsmp-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-default-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-default-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-smp-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-smp-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xen-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xen-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xenpae-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-bigsmp-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-default-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-default-x86_64
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-smp-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-smp-x86_64
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-xen-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-xen-x86_64
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-xenpae-x86
SLES10-SP4-2.6.16.60-0.85.1-bigsmp-x86
SLES10-SP4-2.6.16.60-0.85.1-default-x86
SLES10-SP4-2.6.16.60-0.85.1-default-x86_64
SLES10-SP4-2.6.16.60-0.85.1-smp-x86
SLES10-SP4-2.6.16.60-0.85.1-smp-x86_64
SLES10-SP4-2.6.16.60-0.85.1-xen-x86
SLES10-SP4-2.6.16.60-0.85.1-xen-x86_64
SLES10-SP4-2.6.16.60-0.85.1-xenpae-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-bigsmp-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-default-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-default-x86_64
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-smp-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-smp-x86_64
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-xen-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-xen-x86_64
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-xenpae-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-bigsmp-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-default-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-default-x86_64
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-smp-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-smp-x86_64
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-xen-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-xen-x86_64
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-xenpae-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-bigsmp-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-default-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-default-x86_64
SLES10-SP4_U4-2.6.16.60-0.93.1-smp-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-smp-x86_64
SLES10-SP4_U4-2.6.16.60-0.93.1-xen-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-xen-x86_64
SLES10-SP4_U4-2.6.16.60-0.93.1-xenpae-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-bigsmp-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-default-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-default-x86_64
SLES10-SP4_U5-2.6.16.60-0.97.1-smp-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-smp-x86_64
SLES10-SP4_U5-2.6.16.60-0.97.1-xen-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-xen-x86_64
SLES10-SP4_U5-2.6.16.60-0.97.1-xenpae-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-bigsmp-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-default-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-default-x86_64
SLES10-SP4_U6-2.6.16.60-0.99.1-smp-x86

SLES10-SP4_U6-2.6.16.60-0.99.1-smp-x86_64
SLES10-SP4_U6-2.6.16.60-0.99.1-xen-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-xen-x86_64
SLES10-SP4_U6-2.6.16.60-0.99.1-xenpae-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-bigsmp-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-default-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-default-x86_64
SLES10-SP4_U7-2.6.16.60-0.101.1-smp-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-smp-x86_64
SLES10-SP4_U7-2.6.16.60-0.101.1-xen-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-xen-x86_64
SLES10-SP4_U7-2.6.16.60-0.101.1-xenpae-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-bigsmp-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-default-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-default-x86_64
SLES10-SP4_U8-2.6.16.60-0.103.1-smp-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-smp-x86_64
SLES10-SP4_U8-2.6.16.60-0.103.1-xen-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-xen-x86_64
SLES10-SP4_U8-2.6.16.60-0.103.1-xenpae-x86
SLES11-GA-2.6.27.19-5-default-x86
SLES11-GA-2.6.27.19-5-default-x86_64
SLES11-GA-2.6.27.19-5-pae-x86
SLES11-SP1-2.6.32.12-0.6-default-x86
SLES11-SP1-2.6.32.12-0.6-default-x86_64
SLES11-SP1-2.6.32.12-0.6-pae-x86
SLES11-SP1_LTSS_U1-2.6.32.59-0.9-default-x86
SLES11-SP1_LTSS_U1-2.6.32.59-0.9-default-x86_64
SLES11-SP1_LTSS_U1-2.6.32.59-0.9-pae-x86
SLES11-SP1_LTSS_U2-2.6.32.59-0.13-default-x86
SLES11-SP1_LTSS_U2-2.6.32.59-0.13-default-x86_64
SLES11-SP1_LTSS_U2-2.6.32.59-0.13-pae-x86
SLES11-SP1_U14-2.6.32.54-0.3-default-x86
SLES11-SP1_U14-2.6.32.54-0.3-default-x86_64
SLES11-SP1_U14-2.6.32.54-0.3-pae-x86
SLES11-SP1_U15-2.6.32.59-0.3-default-x86
SLES11-SP1_U15-2.6.32.59-0.3-default-x86_64
SLES11-SP1_U15-2.6.32.59-0.3-pae-x86
SLES11-SP1_U16-2.6.32.59-0.7-default-x86
SLES11-SP1_U16-2.6.32.59-0.7-default-x86_64
SLES11-SP1_U16-2.6.32.59-0.7-pae-x86
SLES11SP2-GA-3.0.13-0.27-default-x86
SLES11SP2-GA-3.0.13-0.27-default-x86_64
SLES11SP2-GA-3.0.13-0.27-pae-x86
SLES11SP2-GA-3.0.13-0.27-xen-x86
SLES11SP2-GA-3.0.13-0.27-xen-x86_64
SLES11SP2-LTSS_U1-3.0.101-0.7.19-default-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-default-x86_64
SLES11SP2-LTSS_U1-3.0.101-0.7.19-pae-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-xen-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-xen-x86_64
SLES11SP2-LTSS_U2-3.0.101-0.7.21-default-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-default-x86_64
SLES11SP2-LTSS_U2-3.0.101-0.7.21-pae-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-xen-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-xen-x86_64
SLES11SP2-U1-3.0.26-0.7-default-x86
SLES11SP2-U1-3.0.26-0.7-default-x86_64
SLES11SP2-U1-3.0.26-0.7-pae-x86
SLES11SP2-U1-3.0.26-0.7-xen-x86

SLES11SP2-U1-3.0.26-0.7-xen-x86_64
SLES11SP2-U10-3.0.74-0.6.8-default-x86
SLES11SP2-U10-3.0.74-0.6.8-default-x86_64
SLES11SP2-U10-3.0.74-0.6.8-pae-x86
SLES11SP2-U10-3.0.74-0.6.8-xen-x86
SLES11SP2-U10-3.0.74-0.6.8-xen-x86_64
SLES11SP2-U11-3.0.74-0.6.10-default-x86
SLES11SP2-U11-3.0.74-0.6.10-default-x86_64
SLES11SP2-U11-3.0.74-0.6.10-pae-x86
SLES11SP2-U11-3.0.74-0.6.10-xen-x86
SLES11SP2-U11-3.0.74-0.6.10-xen-x86_64
SLES11SP2-U12-3.0.80-0.5-default-x86
SLES11SP2-U12-3.0.80-0.5-default-x86_64
SLES11SP2-U12-3.0.80-0.5-pae-x86
SLES11SP2-U12-3.0.80-0.5-xen-x86
SLES11SP2-U12-3.0.80-0.5-xen-x86_64
SLES11SP2-U13-3.0.80-0.7-default-x86
SLES11SP2-U13-3.0.80-0.7-default-x86_64
SLES11SP2-U13-3.0.80-0.7-pae-x86
SLES11SP2-U13-3.0.80-0.7-xen-x86
SLES11SP2-U13-3.0.80-0.7-xen-x86_64
SLES11SP2-U14-3.0.93-0.5-default-x86
SLES11SP2-U14-3.0.93-0.5-default-x86_64
SLES11SP2-U14-3.0.93-0.5-pae-x86
SLES11SP2-U14-3.0.93-0.5-xen-x86
SLES11SP2-U14-3.0.93-0.5-xen-x86_64
SLES11SP2-U15-3.0.101-0.5-default-x86
SLES11SP2-U15-3.0.101-0.5-default-x86_64
SLES11SP2-U15-3.0.101-0.5-pae-x86
SLES11SP2-U15-3.0.101-0.5-xen-x86
SLES11SP2-U15-3.0.101-0.5-xen-x86_64
SLES11SP2-U16-3.0.101-0.7.15-default-x86
SLES11SP2-U16-3.0.101-0.7.15-default-x86_64
SLES11SP2-U16-3.0.101-0.7.15-pae-x86
SLES11SP2-U16-3.0.101-0.7.15-xen-x86
SLES11SP2-U16-3.0.101-0.7.15-xen-x86_64
SLES11SP2-U17-3.0.101-0.7.17-default-x86
SLES11SP2-U17-3.0.101-0.7.17-default-x86_64
SLES11SP2-U17-3.0.101-0.7.17-pae-x86
SLES11SP2-U17-3.0.101-0.7.17-xen-x86
SLES11SP2-U17-3.0.101-0.7.17-xen-x86_64
SLES11SP2-U2-3.0.31-0.9-default-x86
SLES11SP2-U2-3.0.31-0.9-default-x86_64
SLES11SP2-U2-3.0.31-0.9-pae-x86
SLES11SP2-U2-3.0.31-0.9-xen-x86
SLES11SP2-U2-3.0.31-0.9-xen-x86_64
SLES11SP2-U3-3.0.34-0.7-default-x86
SLES11SP2-U3-3.0.34-0.7-default-x86_64
SLES11SP2-U3-3.0.34-0.7-pae-x86
SLES11SP2-U3-3.0.34-0.7-xen-x86
SLES11SP2-U3-3.0.34-0.7-xen-x86_64
SLES11SP2-U4-3.0.38-0.5-default-x86
SLES11SP2-U4-3.0.38-0.5-default-x86_64
SLES11SP2-U4-3.0.38-0.5-pae-x86
SLES11SP2-U4-3.0.38-0.5-xen-x86
SLES11SP2-U4-3.0.38-0.5-xen-x86_64
SLES11SP2-U5-3.0.42-0.7-default-x86
SLES11SP2-U5-3.0.42-0.7-default-x86_64
SLES11SP2-U5-3.0.42-0.7-pae-x86
SLES11SP2-U5-3.0.42-0.7-xen-x86

SLES11SP2-U5-3.0.42-0.7-xen-x86_64
SLES11SP2-U6-3.0.51-0.7.9-default-x86
SLES11SP2-U6-3.0.51-0.7.9-default-x86_64
SLES11SP2-U6-3.0.51-0.7.9-pae-x86
SLES11SP2-U6-3.0.51-0.7.9-xen-x86
SLES11SP2-U6-3.0.51-0.7.9-xen-x86_64
SLES11SP2-U7-3.0.58-0.6.2-default-x86
SLES11SP2-U7-3.0.58-0.6.2-default-x86_64
SLES11SP2-U7-3.0.58-0.6.2-pae-x86
SLES11SP2-U7-3.0.58-0.6.2-xen-x86
SLES11SP2-U7-3.0.58-0.6.2-xen-x86_64
SLES11SP2-U8-3.0.58-0.6.6-default-x86
SLES11SP2-U8-3.0.58-0.6.6-default-x86_64
SLES11SP2-U8-3.0.58-0.6.6-pae-x86
SLES11SP2-U8-3.0.58-0.6.6-xen-x86
SLES11SP2-U8-3.0.58-0.6.6-xen-x86_64
SLES11SP2-U9-3.0.74-0.6.6-default-x86
SLES11SP2-U9-3.0.74-0.6.6-default-x86_64
SLES11SP2-U9-3.0.74-0.6.6-pae-x86
SLES11SP2-U9-3.0.74-0.6.6-xen-x86
SLES11SP2-U9-3.0.74-0.6.6-xen-x86_64
SLES11SP3-GA-3.0.76-0.11-default-x86
SLES11SP3-GA-3.0.76-0.11-default-x86_64
SLES11SP3-GA-3.0.76-0.11-pae-x86
SLES11SP3-GA-3.0.76-0.11-xen-x86
SLES11SP3-GA-3.0.76-0.11-xen-x86_64
SLES11SP3-U1-3.0.82-0.7-default-x86
SLES11SP3-U1-3.0.82-0.7-default-x86_64
SLES11SP3-U1-3.0.82-0.7-pae-x86
SLES11SP3-U1-3.0.82-0.7-xen-x86
SLES11SP3-U1-3.0.82-0.7-xen-x86_64
SLES11SP3-U2-3.0.93-0.8-default-x86
SLES11SP3-U2-3.0.93-0.8-default-x86_64
SLES11SP3-U2-3.0.93-0.8-pae-x86
SLES11SP3-U2-3.0.93-0.8-xen-x86
SLES11SP3-U2-3.0.93-0.8-xen-x86_64
SLES11SP3-U3-3.0.101-0.8-default-x86
SLES11SP3-U3-3.0.101-0.8-default-x86_64
SLES11SP3-U3-3.0.101-0.8-pae-x86
SLES11SP3-U3-3.0.101-0.8-xen-x86
SLES11SP3-U3-3.0.101-0.8-xen-x86_64
SLES11SP3-U4-3.0.101-0.15-default-x86
SLES11SP3-U4-3.0.101-0.15-default-x86_64
SLES11SP3-U4-3.0.101-0.15-pae-x86
SLES11SP3-U4-3.0.101-0.15-xen-x86
SLES11SP3-U4-3.0.101-0.15-xen-x86_64
SLES11SP3-U5-3.0.101-0.21-default-x86
SLES11SP3-U5-3.0.101-0.21-default-x86_64
SLES11SP3-U5-3.0.101-0.21-pae-x86
SLES11SP3-U5-3.0.101-0.21-xen-x86
SLES11SP3-U5-3.0.101-0.21-xen-x86_64
SLES11SP3-U6-3.0.101-0.29-default-x86
SLES11SP3-U6-3.0.101-0.29-default-x86_64
SLES11SP3-U6-3.0.101-0.29-pae-x86
SLES11SP3-U6-3.0.101-0.29-xen-x86
SLES11SP3-U6-3.0.101-0.29-xen-x86_64
SLES11SP3-U7-3.0.101-0.31-default-x86
SLES11SP3-U7-3.0.101-0.31-default-x86_64
SLES11SP3-U7-3.0.101-0.31-pae-x86
SLES11SP3-U7-3.0.101-0.31-xen-x86

SLES11SP3-U7-3.0.101-0.31-xen-x86_64
SLES11SP3-U8-3.0.101-0.35-default-x86
SLES11SP3-U8-3.0.101-0.35-default-x86_64
SLES11SP3-U8-3.0.101-0.35-pae-x86
SLES11SP3-U8-3.0.101-0.35-xen-x86
SLES11SP3-U8-3.0.101-0.35-xen-x86_64
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
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
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
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
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

B 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 From the [NetIQ download site](#), search for the Protect 11 product, then click **Submit Query**.
- 2 On the products tab, select **PlateSpin Protect 11.0**, then click **proceed to download**.
- 3 On the download page, click **download** on the `VolumeManager.exe` line or select the comparable download manager link.
- 4 Download the utility, then copy it to an accessible location on each cluster node.
- 5 On the active node of the cluster, open an administrative command prompt, navigate to the location of the downloaded utility, and run the following command:

```
VolumeManager.exe -l
```

A listing of the local volumes and their respective serial numbers is displayed. For example:

```
Volume Listing:
```

```
-----
```

```
DriveLetter (*:) VolumeId="System Reserved" SerialNumber: AABB-CCDD
```

```
DriveLetter (C:) VolumeId=C:\ SerialNumber: 1122-3344
```

Make note of these serial numbers or keep them displayed for later comparison.

- 6 Verify that all local storage serial numbers of the active node match the local storage serial numbers on each of the other nodes in the cluster.
 - 6a On each cluster node, run the `VolumeManager.exe -l` command to obtain its volume serial numbers.
 - 6b Compare the local storage serial numbers of the active node ([Step 5](#)) against the local storage serial numbers of the node ([Step 6a](#)).
 - 6c (Conditional) If there are any differences in the serial numbers between the active node and this node, take note of the serial number you want to propagate on this node and run the following command to set, and then to verify the serial number:

```
VolumeManager -s <VolumeId> <serial-number>
```

Following are two examples of how this command could be used:

- ♦ `VolumeManager -s "System Reserved" AAAA-AAAA`
- ♦ `VolumeManager -s C:\ 1111-1111`

- 6d** When you have successfully changed all of the volume serial numbers on a node of the cluster, you need to restart that node.
- 6e** Repeat [Step 6a](#) through [Step 6d](#) for each node of the cluster.
- 7** (Conditional) If the cluster has already been 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.

C Documentation Updates

This section contains information on documentation content changes that were made in this *User Guide* after the initial release of NetIQ PlateSpin Protect 11.0.1. The changes are listed according to the date they were published.

NOTE: This updated information does not appear in the help content accessible from the product's user interface nor in localized versions of the *User Guide*.

The documentation for this product is provided on the Web in two formats: HTML and PDF. The HTML and PDF documentation are both kept up-to-date with the changes listed in this section.

If you need to know whether a copy of the PDF documentation that you are using is the most recent, the PDF document includes a publication date on the title page.

The documentation was updated on the following dates:

- ◆ [Section C.1, "February 23, 2015," on page 121](#)
- ◆ [Section C.2, "February 2, 2015," on page 121](#)
- ◆ [Section C.3, "October 24, 2014," on page 122](#)
- ◆ [Section C.4, "October 13, 2014," on page 122](#)
- ◆ [Section C.5, "September 11, 2014," on page 122](#)
- ◆ [Section C.6, "September 5, 2014," on page 122](#)

C.1 February 23, 2015

Updates were made to the following sections:

Location	Updates
Section 1.2.3, "Supported VM Containers," on page 14	A paid license for ESXi is required.
Section 4.8, "Failback," on page 58	Added a link to information about supported VM containers.

C.2 February 2, 2015

Updates were made to the following sections:

Location	Updates
Section 4.8.3, "Semi-Automated Failback to a Virtual Machine," on page 62	Added information about using semi-automated failback for target VM platforms on Microsoft Hyper-V Server 2012 hosts. Removed information about Xen and Hyper-V Server 2008.

C.3 October 24, 2014

Updates were made to the following sections:

Location	Update
Section 1.2.1, "Supported Windows Workloads," on page 11.	Revised the table to include Windows Server 2008, latest SP (32-bit) and Windows 2008 R2 Server-based Microsoft Failover Cluster as supported workload types.

C.4 October 13, 2014

Updates were made to the following sections:

Location	Update
Section 1.2.3, "Supported VM Containers," on page 14.	Revised the table to indicate the ESXi Server and DRS Cluster versions supported as VM containers.

C.5 September 11, 2014

Updates were made to the following sections:

Location	Update
Appendix A, "Linux Distributions Supported by Protect," on page 109.	Added a note indicating what to do if a false validator message is displayed when protecting a Linux workload.

C.6 September 5, 2014

Updates were made to the following sections:

Location	Update
Section 1.2.2, "Supported Linux Workloads," on page 13.	Added other Linux platforms as supported workloads to the matrix. Also added information about and a cross reference to the Linux distributions supported by PlateSpin Protect.

Glossary

container. PlateSpin Protect's workload protection infrastructure, such as a VM host.

event. A PlateSpin Server message that contains information about important steps throughout the workload protection lifecycle.

failback. Restoration of the business function of a failed workload in its original environment when the business function of a temporary failover workload within PlateSpin Protect is no longer required.

failover. Taking over the business function of a failed workload by a failover workload within a PlateSpin Protect VM container.

failover workload. A protected workload's bootable virtual replica.

incremental. 1. (noun) An individual scheduled transfer or manual transfer of differences between a protected workload and its replica (the failover workload).

2. (adjective) Describes the scope of *replication (1)*, in which the initial replica of a workload is created differentially, based on differences between the workload and its prepared counterpart.

prepare for failover. A PlateSpin Protect operation that boots the failover workload in preparation of a full Failover operation.

protection tier. A customizable collection of workload protection parameters that define the frequency of replications and criteria for the system to consider a workload as failed.

protection contract. A collection of currently-active settings pertaining to the complete lifecycle of a workload's protection (*Add-inventory*, initial and ongoing *Replications*, *Failover*, *Failback*, and *Reprotect*).

recovery point. A point-in-time snapshot, allowing a replicated workload to be restored to a previous state.

Recovery Point Objective (RPO). Tolerable data loss measured in time and defined by a configurable interval between incremental replications of a protected workload.

Recovery Time Objective (RTO). A measure of a workload's tolerable downtime defined by the time a failover operation takes to complete.

replication. 1. *Initial Replication*, the creation of an initial base copy of a workload. Can be carried out as a *Full Replication* (all workload data is transferred to a 'blank' failover VM), or as an *Incremental Replication* (see [incremental \(2\)](#)).

2. Any transfer of changed data from a protected workload to its replica in the container.

replication schedule. The schedule that is set up to control the frequency and scope of replications.

reprotect. A PlateSpin Protect command that reestablishes a protection contract for a workload following the failover and failback operations.

source. A workload or its infrastructure that is the starting point of a PlateSpin Protect operation. For example, upon initial protection of a workload, the source is your production workload. In a failback operation, it is the failover workload in the container.

See also [target](#).

target. A workload or its infrastructure that is the outcome of a PlateSpin Protect command. For example, upon initial protection of a workload, the target is the failover workload in the container. In a failback operation, it is either your production workload's original infrastructure or any supported container that has been inventoried by PlateSpin Protect.

See *also* [source](#).

test failover. A PlateSpin Protect operation that boots a failover workload in an isolated networking environment for testing the functionality of the failover and verifying the integrity of the failover workload.

Test Time Objective (TTO). A measure of the ease with which a disaster recovery plan can be tested. It is similar to RTO, but includes the time needed for a user to test the failover workload.

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