



PlateSpin[®] Forge 11.0

User Guide

August 15, 2014

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.

© 2014 NetIQ Corporation. All Rights Reserved.

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

Contents

| | |
|---|-----------|
| About NetIQ Corporation | 7 |
| About This Guide | 9 |
| | |
| 1 Product Overview | 11 |
| 1.1 About PlateSpin Forge | 11 |
| 1.2 Supported Configurations | 11 |
| 1.2.1 Supported Windows Workloads | 11 |
| 1.2.2 Supported Linux Workloads | 12 |
| 1.2.3 Supported VM Containers | 13 |
| 1.3 Security and Privacy | 14 |
| 1.3.1 Security of Workload Data in Transmission | 14 |
| 1.3.2 Security of Credentials | 14 |
| 1.3.3 User Authorization and Authentication | 14 |
| 1.4 Performance | 14 |
| 1.4.1 About Product Performance Characteristics | 14 |
| 1.4.2 Data Compression | 15 |
| 1.4.3 Bandwidth Throttling | 15 |
| 1.4.4 RPO, RTO, and TTO Specifications | 15 |
| | |
| 2 PlateSpin Forge Application Configuration | 17 |
| 2.1 Product Licensing | 17 |
| 2.1.1 Obtaining a License Activation Code | 17 |
| 2.1.2 Online License Activation | 17 |
| 2.1.3 Offline License Activation | 18 |
| 2.2 Setting Up User Authorization and Authentication | 18 |
| 2.2.1 About PlateSpin Forge User Authorization and Authentication | 18 |
| 2.2.2 Managing PlateSpin Forge Security Groups and Workload Permissions | 20 |
| 2.3 Access and Communication Requirements across your Protection Network | 21 |
| 2.3.1 Access and Communication Requirements for Workloads | 21 |
| 2.3.2 Protection Across Public and Private Networks Through NAT | 22 |
| 2.3.3 Overriding the Default bash Shell for Executing Commands on Linux Workloads | 23 |
| 2.4 Configuring PlateSpin Forge Default Options | 23 |
| 2.4.1 Setting Up Automatic Email Notifications of Events and Reports | 23 |
| 2.4.2 Language Setup for International Versions of PlateSpin Forge | 26 |
| 2.4.3 Configuring PlateSpin Server Behavior through XML Configuration Parameters | 26 |
| 2.4.4 Configuring Support for VMware vCenter Site Recovery Manager | 28 |
| | |
| 3 Appliance Setup and Maintenance | 31 |
| 3.1 Setting up Appliance Networking | 31 |
| 3.1.1 Setting up Appliance Host Networking | 31 |
| 3.2 Physically Relocating the Appliance | 32 |
| 3.2.1 Scenario 1 - Relocating Forge (New IP Address Known) | 32 |
| 3.2.2 Scenario 2 - Relocating Forge (New IP Address Unknown) | 33 |
| 3.3 Using External Storage Solutions with PlateSpin Forge | 34 |
| 3.3.1 Using Forge with SAN Storage | 34 |
| 3.3.2 Adding a SAN LUN to Forge | 35 |
| 3.4 Accessing and Working with the Forge Management VM in the Appliance Host | 36 |
| 3.4.1 Downloading the vSphere Client Program | 36 |

| | | |
|-------|---|----|
| 3.4.2 | Launching the vSphere Client and Accessing the Forge Management VM | 36 |
| 3.4.3 | Starting and Shutting Down the Forge Management VM | 37 |
| 3.4.4 | Managing Forge Snapshots on the Appliance Host | 37 |
| 3.4.5 | Manually Importing VMs into the Appliance Host's Datastore | 37 |
| 3.4.6 | Guidelines for Applying Security Updates to the PlateSpin Forge Management VM | 38 |
| 3.5 | Resetting Forge to Factory Defaults | 38 |

4 Up and Running 43

| | | |
|-------|---|----|
| 4.1 | Launching the PlateSpin Forge Web Interface | 43 |
| 4.2 | Elements of the PlateSpin Forge Web Interface | 44 |
| 4.2.1 | Navigation Bar | 45 |
| 4.2.2 | Visual Summary Panel | 45 |
| 4.2.3 | Tasks and Events Panel | 46 |
| 4.3 | Workloads and Workload Commands | 46 |
| 4.3.1 | Workload Protection and Recovery Commands | 47 |
| 4.4 | Managing Multiple Instances of PlateSpin Protect and PlateSpin Forge | 48 |
| 4.4.1 | Using the PlateSpin Forge Management Console | 48 |
| 4.4.2 | About PlateSpin Forge Management Console Cards | 48 |
| 4.4.3 | Adding Instances of PlateSpin Protect and PlateSpin Forge to the Management Console | 49 |
| 4.4.4 | Managing Cards on the Management Console | 50 |
| 4.5 | Generating Workload and Workload Protection Reports | 50 |

5 Workload Protection 53

| | | |
|-------|--|----|
| 5.1 | Basic Workflow for Workload Protection and Recovery | 53 |
| 5.2 | Adding Workloads for Protection | 54 |
| 5.3 | Configuring Protection Details and Preparing the Replication | 55 |
| 5.3.1 | Workload Protection Details | 56 |
| 5.4 | Starting the Workload Protection | 58 |
| 5.5 | Aborting Commands | 58 |
| 5.6 | Failover | 59 |
| 5.6.1 | Detecting Offline Workloads | 59 |
| 5.6.2 | Performing a Failover | 60 |
| 5.6.3 | Using the Test Failover Feature | 60 |
| 5.7 | Failback | 61 |
| 5.7.1 | Automated Failback to a VM Platform | 61 |
| 5.7.2 | Semi-Automated Failback to a Physical Machine | 64 |
| 5.8 | Reprotecting a Workload | 65 |

6 Essentials of Workload Protection 67

| | | |
|-------|---|----|
| 6.1 | Workload License Consumption | 67 |
| 6.2 | Guidelines for Workload Credentials | 68 |
| 6.3 | Data Transfer | 68 |
| 6.3.1 | Transfer Methods | 68 |
| 6.3.2 | Data Encryption | 69 |
| 6.4 | Protection Tiers | 69 |
| 6.5 | Recovery Points | 71 |
| 6.6 | Initial Replication Method (Full and Incremental) | 71 |
| 6.7 | Service and Daemon Control | 72 |
| 6.8 | Using Freeze and Thaw Scripts for Every Replication (Linux) | 72 |
| 6.9 | Volumes | 73 |
| 6.10 | Networking | 75 |
| 6.11 | Failback to Physical Machines | 75 |

| | | |
|----------|---|------------|
| 6.11.1 | Downloading the PlateSpin Boot ISO Images | 75 |
| 6.11.2 | Injecting Additional Device Drivers into the Boot ISO Image | 75 |
| 6.11.3 | Registering Physical Machines as Failback Targets with PlateSpin Forge | 76 |
| 6.12 | Advanced Workload Protection Topics | 77 |
| 6.12.1 | Protecting Windows Clusters | 77 |
| 6.12.2 | Using Workload Protection Features through the PlateSpin Forge Web Services API | 79 |
| 7 | Auxiliary Tools for Working with Physical Machines | 81 |
| 7.1 | Managing Device Drivers | 81 |
| 7.1.1 | Packaging Device Drivers for Windows Systems | 81 |
| 7.1.2 | Packaging Device Drivers for Linux Systems | 82 |
| 7.1.3 | Uploading Drivers to the PlateSpin Forge Device Driver Database | 82 |
| 7.1.4 | Using the Plug and Play (PnP) ID Translator Feature | 84 |
| 8 | Troubleshooting | 87 |
| 8.1 | Troubleshooting Workload Inventory (Windows) | 87 |
| 8.1.1 | Performing Connectivity Tests | 88 |
| 8.1.2 | Disabling AntiVirus Software | 90 |
| 8.1.3 | Enabling File/Share Permissions and Access | 90 |
| 8.2 | Troubleshooting Workload Inventory (Linux) | 91 |
| 8.3 | Troubleshooting Problems during the Prepare Replication Command (Windows) | 91 |
| 8.3.1 | Group Policy and User Rights | 91 |
| 8.4 | Troubleshooting Workload Replication | 92 |
| 8.5 | Generating and Viewing Diagnostic Reports | 93 |
| 8.6 | Removing Workloads | 94 |
| 8.7 | Post-Protection Workload Cleanup | 94 |
| 8.7.1 | Cleaning Up Windows Workloads | 94 |
| 8.7.2 | Cleaning Up Linux Workloads | 95 |
| 8.8 | Shrinking the PlateSpin Forge Databases | 96 |
| A | Linux Distributions Supported by Forge | 97 |
| A.1 | Analyzing Your Linux Workload | 97 |
| A.1.1 | Determining the Release String | 97 |
| A.1.2 | Determining the Architecture | 97 |
| A.2 | Pre-compiled "blkwatch" driver (Linux) | 98 |
| B | Synchronizing Cluster Node Local Storage | 109 |
| | Glossary | 111 |

About NetIQ Corporation

We are a global, enterprise software company, with a focus on the three persistent challenges in your environment: Change, complexity and risk—and how we can help you control them.

Our Viewpoint

Adapting to change and managing complexity and risk are nothing new

In fact, of all the challenges you face, these are perhaps the most prominent variables that deny you the control you need to securely measure, monitor, and manage your physical, virtual, and cloud computing environments.

Enabling critical business services, better and faster

We believe that providing as much control as possible to IT organizations is the only way to enable timelier and cost effective delivery of services. Persistent pressures like change and complexity will only continue to increase as organizations continue to change and the technologies needed to manage them become inherently more complex.

Our Philosophy

Selling intelligent solutions, not just software

In order to provide reliable control, we first make sure we understand the real-world scenarios in which IT organizations like yours operate — day in and day out. That's the only way we can develop practical, intelligent IT solutions that successfully yield proven, measurable results. And that's so much more rewarding than simply selling software.

Driving your success is our passion

We place your success at the heart of how we do business. From product inception to deployment, we understand that you need IT solutions that work well and integrate seamlessly with your existing investments; you need ongoing support and training post-deployment; and you need someone that is truly easy to work with — for a change. Ultimately, when you succeed, we all succeed.

Our Solutions

- ◆ Identity & Access Governance
- ◆ Access Management
- ◆ Security Management
- ◆ Systems & Application Management
- ◆ Workload Management
- ◆ Service Management

Contacting Sales Support

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

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

Contacting Technical Support

For specific product issues, contact our Technical Support team.

| | |
|---|---|
| Worldwide: | www.netiq.com/support/contactinfo.asp |
| North and South America: | 1-713-418-5555 |
| Europe, Middle East, and Africa: | +353 (0) 91-782 677 |
| Email: | support@platespin.com |
| Web Site: | www.netiq.com/support |
| Technical Support Guide: | https://www.netiq.com/Support/process.asp#_Maintenance_Programs_and |
| Product Specific Information: | https://www.netiq.com/support/kb/product.php?id=SG_XPLATESPINFORGE_1_2 |

Contacting Documentation Support

Our goal is to provide documentation that meets your needs. The documentation for this product is available on the NetIQ Web site in HTML and PDF formats on a page that does not require you to log in. If you have suggestions for documentation improvements, click **Add Comment** at the bottom of any page in the HTML version 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

NetIQ Communities, 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, NetIQ Communities helps ensure you are mastering the knowledge you need to realize the full potential of IT investments upon which you rely. For more information, visit <http://community.netiq.com>.

About This Guide

This *User Guide* provides information about using PlateSpin Forge. provides conceptual information about the PlateSpin Forge product. It also defines terminology and includes troubleshooting information.

Intended Audience

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

Other Information in the Library

The library provides the following information resources:

Getting Started Guide

Provides information about the basic steps required for setting up your PlateSpin Forge appliance.

Upgrade Guide

Provides conceptual information about upgrading the PlateSpin Forge Appliance from version 3.1, 3.3 or 3.4 to version 11.0.

Rebuild Guide

Provides information about rebuilding and reconfiguring the PlateSpin Forge 11 Hardware Appliance by using the *Forge 11.0.0 Field Rebuild Kit*.

Help

Provides embedded guidance to assist users with common tasks as they access the user interface.

Documentation Updates

The most recent version of this guide can be found at [PlateSpin Forge 11 Online Documentation Web Site \(https://www.netiq.com/documentation/platespin_forge_11/\)](https://www.netiq.com/documentation/platespin_forge_11/).

1 Product Overview

This section includes the following information:

- ♦ [Section 1.1, “About PlateSpin Forge,” on page 11](#)
- ♦ [Section 1.2, “Supported Configurations,” on page 11](#)
- ♦ [Section 1.3, “Security and Privacy,” on page 14](#)
- ♦ [Section 1.4, “Performance,” on page 14](#)

1.1 About PlateSpin Forge

PlateSpin Forge is a consolidated recovery hardware appliance that protects physical and virtual workloads (operating systems, middleware, and data) by using embedded virtualization technology. If there is a production server outage or disaster, workloads can be rapidly powered on within the PlateSpin Forge recovery environment and continue to run as normal until the production environment is restored.

PlateSpin Forge enables you to:

- ♦ Quickly recover workloads upon failure
- ♦ Simultaneously protect multiple workloads (10 to 50, depending on the model)
- ♦ 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

With internal, prepackaged storage, Forge has a total storage capacity of up to 20 terabytes, although the capacity is almost unlimited when external storage configurations are used by adding iSCSI or Fibre Channel cards.

1.2 Supported Configurations

- ♦ [Section 1.2.1, “Supported Windows Workloads,” on page 11](#)
- ♦ [Section 1.2.2, “Supported Linux Workloads,” on page 12](#)
- ♦ [Section 1.2.3, “Supported VM Containers,” on page 13](#)

1.2.1 Supported Windows Workloads

PlateSpin Forge supports most Windows-based workloads.

Both file-level and block-level replication are supported, with certain restrictions. See [Section 6.3, “Data Transfer,” on page 68](#).

Table 1-1 Supported Windows Workloads

| Operating System | Remarks |
|--|---|
| Server Class workloads | |
| Windows Server 2012 R2 Windows Server 2012 | Including domain controllers (DC) and Small Business Server (SBS) editions |
| Windows Server 2008 R2 Windows Server 2008 latest SP (64-bit) Windows Server 2008 latest SP (32-bit) Windows Server 2008 (64-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) | Including domain controllers (DC) and Small Business Server (SBS) editions. Windows 2003 requires SP1 or higher for block-based replication. |
| Windows Server 2000 SP4 (32-bit) | |
| Windows 2008 R2 Server-based Microsoft Failover Cluster | |
| Workstation Class workloads | |
| Windows 8.1 | |
| Windows 8 | |
| Windows 7 | Professional, Enterprise, and Ultimate editions only |
| Windows Vista | Business, Enterprise, and Ultimate editions; SP1 and later |
| Windows XP | |

The following are examples of Forge 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), Forge transitions the workload's UEFI firmware at failover time to BIOS firmware. Then, when failback is selected on a UEFI-based physical machine, Forge reverses the firmware transition from BIOS to UEFI.
- ◆ If you attempt to failback a protected Windows 2003 workload to a UEFI-based physical machine, Forge 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, Forge 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 Forge supports a number of Linux distributions.

Replication of protected Linux workloads occurs at the block level. PlateSpin Forge software includes pre-compiled versions of the `blkwatch` module/driver. These are only for non-debug kernels of the following Linux distributions (both 32 and 64 bit):

Table 1-2 Linux Distributions that have a corresponding, pre-compiled `blkwatch` module/driver

| Operating System | Remarks |
|--|--|
| Red Hat Enterprise Linux 4 | See Appendix A, “Linux Distributions Supported by Forge,” on page 97 for a list of supported kernel versions. |
| Red Hat Enterprise Linux 5 | See Appendix A, “Linux Distributions Supported by Forge,” on page 97 for a list of supported kernel versions. |
| RedHat Enterprise Linux 6 | See Appendix A, “Linux Distributions Supported by Forge,” on page 97 for a list of supported kernel versions. |
| SUSE Linux Enterprise Server 9 | See Appendix A, “Linux Distributions Supported by Forge,” on page 97 for a list of supported kernel versions. |
| SUSE Linux Enterprise Server 10 | See Appendix A, “Linux Distributions Supported by Forge,” on page 97 for a list of supported kernel versions. |
| SUSE Linux Enterprise Server 11 | See Appendix A, “Linux Distributions Supported by Forge,” on page 97 for a list of supported kernel versions. |
| <ul style="list-style-type: none"> ◆ Novell Open Enterprise Server (OES) 11 SP1 and SP2 ◆ OES 2 (SP2, SP3) | See Appendix A, “Linux Distributions Supported by Forge,” on page 97 for a list of supported SLES kernel versions. |
| Oracle Enterprise Linux (OEL) | <ul style="list-style-type: none"> ◆ See Appendix A, “Linux Distributions Supported by Forge,” on page 97 for a list of supported RedHat kernel versions. ◆ Workloads using the Unbreakable Enterprise Kernel are not supported. |

Supported Linux file systems: EXT2, EXT3, EXT4, REISERFS, and NSS (OES 2 workloads).

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

For a list of Linux distributions for which Forge software has pre-compiled versions of the `blkwatch` module please, see [Appendix A, “Linux Distributions Supported by Forge,” on page 97](#).

If a pre-compiled `blkwatch` driver doesn't exist for your distribution, you might able to create a **custom** `blkwatch` driver by following the steps found in [KB Article 7005873 \(https://www.netiq.com/support/kb/doc.php?id=7005873\)](https://www.netiq.com/support/kb/doc.php?id=7005873).

1.2.3 Supported VM Containers

PlateSpin Forge ships with a VMware ESXi 5.5.0 Update 1, as the protection VM container.

The following table lists the supported VM containers.

Table 1-3 Supported VM Containers

| Operating System | Remarks |
|-----------------------------------|-------------------------------|
| VMware DRS Cluster in vSphere 5.5 | Supported failback container. |
| VMware DRS Cluster in vSphere 5.1 | Supported failback container. |
| VMware DRS Cluster in vSphere 4.1 | Supported failback container. |

1.3 Security and Privacy

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

- ◆ [Section 1.3.1, “Security of Workload Data in Transmission,” on page 14](#)
- ◆ [Section 1.3.2, “Security of Credentials,” on page 14](#)
- ◆ [Section 1.3.3, “User Authorization and Authentication,” on page 14](#)

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).

You can enable or disable encryption individually for each workload. See [Section 5.3.1, “Workload Protection Details,” on page 56](#).

1.3.2 Security of Credentials

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

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.

1.3.3 User Authorization and Authentication

PlateSpin Forge 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 18](#).

1.4 Performance

- ◆ [Section 1.4.1, “About Product Performance Characteristics,” on page 14](#)
- ◆ [Section 1.4.2, “Data Compression,” on page 15](#)
- ◆ [Section 1.4.3, “Bandwidth Throttling,” on page 15](#)
- ◆ [Section 1.4.4, “RPO, RTO, and TTO Specifications,” on page 15](#)

1.4.1 About Product Performance Characteristics

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

- ◆ Hardware and software profiles of your source workloads
- ◆ 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 Forge 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 [Section 6.4, "Protection Tiers," on page 69](#).

1.4.3 Bandwidth Throttling

PlateSpin Forge 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 [Section 6.4, "Protection Tiers," on page 69](#).

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 Forge, 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 [Section 5.6, "Failover," on page 59](#).
- ◆ **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 60](#).

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 [Section 4.5, “Generating Workload and Workload Protection Reports,” on page 50](#).

2 PlateSpin Forge Application Configuration

This section includes the following information:

- ♦ [Section 2.1, “Product Licensing,” on page 17](#)
- ♦ [Section 2.2, “Setting Up User Authorization and Authentication,” on page 18](#)
- ♦ [Section 2.3, “Access and Communication Requirements across your Protection Network,” on page 21](#)
- ♦ [Section 2.4, “Configuring PlateSpin Forge Default Options,” on page 23](#)

2.1 Product Licensing

This section provides information about activating your PlateSpin Forge software.

- ♦ [Section 2.1.1, “Obtaining a License Activation Code,” on page 17](#)
- ♦ [Section 2.1.2, “Online License Activation,” on page 17](#)
- ♦ [Section 2.1.3, “Offline License Activation,” on page 18](#)

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 [Novell Customer Center Web site \(http://www.novell.com/customercenter/\)](http://www.novell.com/customercenter/). A license activation code will be emailed to you.

The first time you log into PlateSpin Forge, 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](#).

2.1.2 Online License Activation

For online activation, PlateSpin Forge must have Internet access.

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

To activate an online license:

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

- 2 Select **Online Activation**, specify the email 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 Novell account. If you are an existing PlateSpin customer and you don't have a Novell account, you must first create one. Use your existing PlateSpin username (a valid email address registered with PlateSpin) as input for your Novell account username.

To activate an offline license:

- 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 Novell username.
- 4 Enter in the appropriate fields:
 - ♦ the activation code that you received
 - ♦ the email 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

- ♦ [Section 2.2.1, "About PlateSpin Forge User Authorization and Authentication," on page 18](#)
- ♦ [Section 2.2.2, "Managing PlateSpin Forge Security Groups and Workload Permissions," on page 20](#)

2.2.1 About PlateSpin Forge User Authorization and Authentication

The user authorization and authentication mechanism of PlateSpin Forge 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 Forge 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 Forge, 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 |
| Run Reports/Diagnostics | Allowed | Allowed | Allowed |
| Failback | Allowed | Denied | Denied |
| Reprotect | Allowed | Allowed | Denied |

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

Setting up a proper role-based access to PlateSpin Forge 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 [Section 2.2.2, “Managing PlateSpin Forge Security Groups and Workload Permissions,”](#) on [page 20](#)).

2.2.2 Managing PlateSpin Forge Security Groups and Workload Permissions

PlateSpin Forge 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*.

To set up a security group:

- 1 Assign a PlateSpin Forge user a Workload Protection Role whose permissions best suit that role in your organization.
- 2 Access PlateSpin Forge as an administrator by using the PlateSpin Forge 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 Forge user that was recently added to the Forge VM, it might not be immediately available in the user interface. In this case, first click **Refresh User Accounts**.

| Grant | Name | Roles |
|-------------------------------------|----------------|------------------------------|
| <input checked="" type="checkbox"/> | ADLERoperator1 | Workload Protection Operator |

- 6 Click **Add Workloads** and select the required workloads:

| Include | Workload Name | Security Group |
|-------------------------------------|---------------|----------------|
| <input checked="" type="checkbox"/> | WIN7-PC | BCM Operators |
| <input type="checkbox"/> | 10.99.161.227 | [Unassigned] |
| <input type="checkbox"/> | AE-W2K3-1 | [Unassigned] |
| <input checked="" type="checkbox"/> | AE-W2K3-3 | [Unassigned] |
| <input checked="" type="checkbox"/> | AE-W2K3-4 | [Unassigned] |
| <input type="checkbox"/> | AE-W2K3-4Y | [Unassigned] |
| <input type="checkbox"/> | AE-W2K3-5 | [Unassigned] |
| <input type="checkbox"/> | DI-w2k3Dyntar | [Unassigned] |

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

- ◆ [Section 2.3.1, “Access and Communication Requirements for Workloads,” on page 21](#)
- ◆ [Section 2.3.2, “Protection Across Public and Private Networks Through NAT,” on page 22](#)
- ◆ [Section 2.3.3, “Overriding the Default bash Shell for Executing Commands on Linux Workloads,” on page 23](#)

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 Forge.

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, 3.5 SP1, 4.0, or 4.5 | |
| Windows Vista and later | <ul style="list-style-type: none"> ◆ Built-in Administrator or domain administrator account credentials (membership only in the local Administrators group is insufficient). On Vista, the account must be enabled (it is disabled by default). ◆ 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> |

| Workload Type | Prerequisites | Required Ports (Defaults) |
|---|--|--|
| 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 |
| Windows Server 2000; Windows XP | <ul style="list-style-type: none"> ◆ Windows Management Instrumentation (WMI) installed <p>WMI (RPC/DCOM) can use TCP ports 135 and 445 as well as random or dynamically assigned ports above 1024. If problems occur when adding the workload, consider temporarily placing the workload in a DMZ or temporarily opening the firewalled ports while adding the workload to PlateSpin Forge.</p> <p>For additional information, such as guidance in limiting the port range for DCOM and RPC, see the following Microsoft technical articles.</p> <ul style="list-style-type: none"> ◆ Using DCOM with Firewalls (http://msdn.microsoft.com/en-us/library/ms809327.aspx) ◆ Configuring RPC dynamic port allocation to work with firewalls (http://support.microsoft.com/default.aspx?scid=kb;en-us;154596) ◆ Configuring DCOM to work over a NAT-based firewall (http://support.microsoft.com/kb/248809) | <p>TCP 3725</p> <p>NetBIOS 137 - 139</p> <p>SMB (TCP 139, 445 and UDP 137, 138)</p> <p>RPC (TCP 135)</p> |
| All Linux workloads | Secure Shell (SSH) server | TCP 22, 3725 |

2.3.2 Protection Across Public and Private Networks Through NAT

In some cases, a source, a target, or PlateSpin Forge 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 Forge 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 23](#).
- ◆ **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 63](#)).
- ◆ **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 [Section 2.4.3, "Configuring PlateSpin Server Behavior through XML Configuration Parameters,"](#) on page 26.

2.3.3 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).

2.4 Configuring PlateSpin Forge Default Options

- ◆ [Section 2.4.1, "Setting Up Automatic Email Notifications of Events and Reports,"](#) on page 23
- ◆ [Section 2.4.2, "Language Setup for International Versions of PlateSpin Forge,"](#) on page 26
- ◆ [Section 2.4.3, "Configuring PlateSpin Server Behavior through XML Configuration Parameters,"](#) on page 26
- ◆ [Section 2.4.4, "Configuring Support for VMware vCenter Site Recovery Manager,"](#) on page 28

2.4.1 Setting Up Automatic Email Notifications of Events and Reports

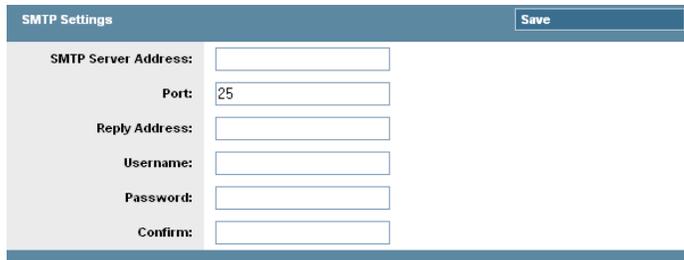
You can configure PlateSpin Forge to automatically send notifications of events and replication reports to specified email addresses. This functionality requires that you first specify a valid SMTP server for PlateSpin Forge to use.

- ◆ ["SMTP Configuration" on page 23](#)
- ◆ ["Setting Up Automatic Event Notifications by Email" on page 24](#)
- ◆ ["Setting Up Automatic Replication Reports by Email" on page 25](#)

SMTP Configuration

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

Figure 2-1 Simple Mail Transfer Protocol Settings



The screenshot shows a web form titled "SMTP Settings" with a "Save" button in the top right. The form contains the following fields:

- SMTP Server Address:
- Port:
- Reply Address:
- Username:
- Password:
- Confirm:

To configure SMTP settings:

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

Setting Up Automatic Event Notifications by Email

To set up automatic event notifications:

- 1 Set up an SMTP server for PlateSpin Forge to use. See [“SMTP Configuration” on page 23](#).
- 2 In your PlateSpin Forge Web Interface, click **Settings > Email > Notification Settings**.
- 3 Select the **Enable Notifications** option.
- 4 Click **Edit Recipients**, type the required email addresses separated by commas, then click **OK**.
- 5 Click **Save**.

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

The following events trigger email 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 | |

| Event | Remarks |
|--------------------------------|--|
| 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. |
| 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 Email

To set up PlateSpin Forge to automatically send out replication reports by email:

- 1 Set up an SMTP server for PlateSpin Forge to use. See [“SMTP Configuration” on page 23](#).
- 2 In your PlateSpin Forge 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 email 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 Forge VM has more than one NIC or if it is located behind a NAT server). This URL affects the title of the report and the functionality of accessing relevant content on the server through hyperlinks within emailed reports.
- 7 Click **Save**.

For information on other types of reports that you can generate and view on demand, see [Section 4.5, “Generating Workload and Workload Protection Reports,” on page 50](#).

2.4.2 Language Setup for International Versions of PlateSpin Forge

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

To use the PlateSpin Forge 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.

To add a language to your Web browser:

- 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 Forge Server. See [Section 4.1, “Launching the PlateSpin Forge Web Interface,”](#) on page 43.

NOTE: (For users of Chinese Traditional and Chinese Simplified versions) Attempting to connect to the PlateSpin Forge 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 Forge Server depends on the operating system interface language selected in your Forge VM.

To change the OS language:

- 1 Access your Forge VM.
See [Section 3.4, “Accessing and Working with the Forge Management VM in the Appliance Host,”](#) on page 36.
- 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.

2.4.3 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 on your Forge VM (https://Your_Forge_VM/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.

To change and apply any configuration parameters:

- 1 On your Forge VM, go to the indicated directory.
- 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, in which you might need to change product behavior using an XML configuration value.

- ♦ [“Optimizing Data Transfer over WAN Connections” on page 27](#)
- ♦ [“Setting up Support for SRM” on page 28](#)

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 Forge VM. For the generic procedure, see [Section 2.4.3, “Configuring PlateSpin Server Behavior through XML Configuration Parameters,” on page 26](#).

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 2-3](#) lists the configuration parameters with the defaults and with the values recommended for optimum operation in a high-latency WAN environment.

Table 2-3 *Default and Optimized Configuration Parameters in https://Your_Forge_VM/platespinconfiguration/*

| Parameter | Default Value | Optimized Value |
|--|---------------|-------------------|
| fileTransferMinCompressionLimit | 0 (disabled) | max 65536 (64 KB) |
| Specifies the packet-level compression threshold in bytes. | | |
| fileTransferCompressionThreadsCount | 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. | | |

| Parameter | Default Value | Optimized Value |
|--|----------------|--------------------|
| <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 Forge 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 2.4.4, “Configuring Support for VMware vCenter Site Recovery Manager,” on page 28](#)

2.4.4 Configuring Support for VMware vCenter Site Recovery Manager

You might use PlateSpin Forge 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.

Configuration for Forge SRM support consists of the following adjustments:

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

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 PlateSpin Forge Server. In cases where this cannot happen, you need to prepare your environment prior to replication.

To prepare your environment:

- 1 Retrieve the VMware Tools packages from an ESXi 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 5.5:** The setup packages are `setup.exe` and `setup64.exe`.
 - ♦ **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 |
| 5.5 | x86 | 660C345A-7A91-458b-BC47-6A3914723EF7 |
| 5.5 | x64 | 8546D4EF-8CA5-4a51-A3A3-6240171BE278 |

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 Appliance Setup and Maintenance

This section provides information about appliance setup and maintenance tasks that you might need to complete on a regular basis.

- ♦ [Section 3.1, “Setting up Appliance Networking,” on page 31](#)
- ♦ [Section 3.2, “Physically Relocating the Appliance,” on page 32](#)
- ♦ [Section 3.3, “Using External Storage Solutions with PlateSpin Forge,” on page 34](#)
- ♦ [Section 3.4, “Accessing and Working with the Forge Management VM in the Appliance Host,” on page 36](#)
- ♦ [Section 3.5, “Resetting Forge to Factory Defaults,” on page 38](#)

3.1 Setting up Appliance Networking

This section provides information about customizing the networking settings of your appliance host.

- ♦ [Section 3.1.1, “Setting up Appliance Host Networking,” on page 31](#)

3.1.1 Setting up Appliance Host Networking

Your PlateSpin Forge appliance has six physical network interfaces configured for external access:

- ♦ **External Test Network:** To isolate network traffic when testing a failover workload with the Test Failover feature.
- ♦ **Internal Test Network:** For testing a failover workload in complete isolation from the production network.
- ♦ **Replication Network:** To provide the system with networking designated for ongoing traffic between your production workload and its replica in the Management VM.
- ♦ **Production Network:** For real-life business continuity networking when performing a failover or a failback.
- ♦ **Forge VM Management Network:** The management network that hosts the Forge management VM.
- ♦ **Management Network:** Hypervisor management network. This network is unavailable for selection in the PlateSpin Forge Web Client.

By default, PlateSpin Forge ships with all 6 physical network interfaces mapped to one vSwitch in the hypervisor. You can customize the mapping to better suit your environment. For example, you can protect a workload that has two NICs, one of which is used for production connectivity, and the other strictly for replications. For additional information, see [KB Article 7921062 \(https://www.netiq.com/support/kb/doc.php?id=7921062\)](https://www.netiq.com/support/kb/doc.php?id=7921062).

In addition, to further fine-tune the control of your network traffic, consider assigning a different VLAN ID to each of these individual port groups. This ensures that your production network is not interfered with by traffic from workload protection and recovery operations. See [KB Article 21057 \(https://www.netiq.com/support/kb/doc.php?id=7921057\)](https://www.netiq.com/support/kb/doc.php?id=7921057).

3.2 Physically Relocating the Appliance

Relocating your PlateSpin Forge appliance (version 3) involves changing the IP addresses of its components to reflect the new environment. These are the IP addresses you specified during the initial setup of the appliance (see your *Forge Getting Started Guide*).

Before starting the relocation procedure:

- 1 Pause all replication schedules, ensuring that at least one incremental has run for each workload:
 - 1a In your PlateSpin Forge Web Client, select all workloads, click **Pause**, then click **Execute**.
 - 1b Ensure that the status **Paused** is displayed for all the workloads.

The specifics of the relocation process vary depending on whether the new IP address of the appliance at the target site is known (scenario 1) or unknown (scenario 2).

- ♦ [Section 3.2.1, “Scenario 1 - Relocating Forge \(New IP Address Known\),” on page 32](#)
- ♦ [Section 3.2.2, “Scenario 2 - Relocating Forge \(New IP Address Unknown\),” on page 33](#)

3.2.1 Scenario 1 - Relocating Forge (New IP Address Known)

To relocate the Forge Appliance Hardware when you know its new IP address:

- 1 Pause all replications. See [Step 1a](#) and [Step 1b](#) above.
- 2 Launch the Forge Appliance Configuration Console (Forge ACC): open a browser and go to `http://<Forge_IP_address>:10000`.
- 3 Log in using the `forgeuser` account and click **Configure Host**.
- 4 Enter the new network parameters and click **Apply**.
- 5 Wait for the configuration process to complete and for the browser window to display the Configuration Successful popup window.

NOTE: The link in the popup window for the new Forge ACC address will not work until you physically disconnect your appliance and connect it to the new subnet.

- 6 Shut down the appliance:
 - 6a Shut down the Forge Management VM. See [Section 3.4.3, “Starting and Shutting Down the Forge Management VM,” on page 37](#).
 - 6b Shut down the Appliance Host:
 - 6b1 At the Forge Console, switch to the ESX Server console by pressing Alt-F2.
 - 6b2 Log in as the superuser (user `root` with the associated password).
 - 6b3 Type the following command and press Enter:

```
shutdown -h now
```
 - 6c Power down the appliance.
- 7 Disconnect your appliance, move it to the new site, attach it to the new subnet, and power it on. The new IP address should now be valid.

- 8 Launch the Forge ACC and log in using the `forgeuser` account, click **Configure Forge VM**, specify the required parameters, then click **Apply**.
- 9 Verify that the setting are correct, click **Continue**, and wait for the process to complete.

NOTE: If you configured the Forge VM to use DHCP, do the following after the relocation:

1. Determine the Forge VM's new IP address (use the vSphere Client program to access the Forge VM and look it up in the VM's Windows interface. See [Section 3.4.2, "Launching the vSphere Client and Accessing the Forge Management VM,"](#) on page 36).
2. Use the new IP address to launch the PlateSpin Forge Web Interface and refresh the container (click **Settings** > **Containers** > then click ↻).

-
- 10 Resume the paused replications.

3.2.2 Scenario 2 - Relocating Forge (New IP Address Unknown)

To relocate the Forge Appliance Hardware when its new IP address is unknown:

- 1 Pause all replications. See [Step 1a](#) and [Step 1b](#) on page 32.
- 2 Shut down the appliance:
 - 2a Shut down the Forge Management VM. See [Section 3.4.3, "Starting and Shutting Down the Forge Management VM,"](#) on page 37.
 - 2b Shut down the Appliance Host:
 - 2b1 At the Forge Console, switch to the ESX Server console by pressing Alt-F2.
 - 2b2 Log in as the superuser (user `root` with the associated password).
 - 2b3 Type the following command and press Enter:

```
shutdown -h now
```
 - 2c Power the appliance off.
- 3 Disconnect your appliance, move it, attach to the new network, then power it on.
- 4 Set up a computer (notebook computer recommended) so that it is able to communicate with Forge at its current IP address (the IP address at the old site), then connect it to the appliance. See ["Appliance Configuration Procedure"](#) in the *PlateSpin Forge 11.0 Getting Started Guide*.
- 5 Launch the Forge ACC: open a browser and go to `http://<Forge_IP_address>:10000`.
- 6 Log in using the `forgeuser` account and click **Configure Host**.
- 7 Enter the new network parameters and click **Apply**.
- 8 Wait for the configuration process to complete and for the browser window to display the Configuration Successful popup window.

NOTE: The link in the popup window for the new Forge ACC address will not work until you now physically disconnect your appliance and connect it to the new subnet.

-
- 9 Disconnect the computer from the appliance and connect the appliance to the new subnet. The new IP address should now be valid.
 - 10 Launch the Forge ACC and log in using the `forgeuser` account, click **Configure Forge VM**, specify the required parameters, then click **Apply**.
 - 11 Verify that the setting are correct, click **Continue**, and wait for the process to complete.

NOTE: If you configured the Forge VM to use DHCP, do the following after the relocation:

1. Determine the Forge VM's new IP address (use the vSphere Client program to access the Forge VM and look it up in the VM's Windows interface. See [“Launching the vSphere Client and Accessing the Forge Management VM” on page 36](#)).
 2. Use the new IP address to launch the PlateSpin Forge Web Interface and refresh the container (click **Settings** > **Containers** > then click ↻).
-

12 Resume the paused replications.

3.3 Using External Storage Solutions with PlateSpin Forge

The following sections contain information to help you with the setup and configuration of external storage for PlateSpin Forge.

- ♦ [Section 3.3.1, “Using Forge with SAN Storage,” on page 34](#)
- ♦ [Section 3.3.2, “Adding a SAN LUN to Forge,” on page 35](#)

3.3.1 Using Forge with SAN Storage

PlateSpin Forge supports existing external storage solutions, such as Storage Area Network (SAN) implementations. Both Fibre Channel and iSCSI solutions are supported. SAN support for Fibre Channel and iSCSI HBAs allows a Forge appliance to be connected to a SAN array. You can then use SAN-array LUNs (Logical Units) to store workload data. Using Forge with a SAN improves flexibility, efficiency, and reliability.

Each SAN product has its own nuances and differences that do not migrate from one hardware manufacturer to the next. This is especially true when considering how these products connect and interact with the Forge Management VM. As such, specific configuration steps for each possible environment and context are beyond the scope of this guide.

The best place to find this type of information is from your hardware vendor or your SAN product sales representative. Many hardware vendors have support guides available describing these tasks in detail. You can find a wealth of information at the following sites:

The [VMware Documentation Web site](http://www.vmware.com/support/pubs/) (<http://www.vmware.com/support/pubs/>).

- ♦ The *Fibre Channel SAN Configuration Guide* discusses the use of ESX Server with Fibre Channel storage area networks.
- ♦ The *iSCSI SAN Configuration Guide* discusses the use of ESX Server with iSCSI storage area networks.
- ♦ The *VMware I/O Compatibility Guide* lists the currently approved HBAs, HBA drivers, and driver versions.
- ♦ The *VMware Storage/SAN Compatibility Guide* lists currently approved storage arrays.
- ♦ The *VMware Release Notes* give information about known issues and workarounds.
- ♦ The *VMware Knowledge Bases* have information on common issues and workarounds.

The following vendors provide storage products that have all been tested by VMware:

- ♦ [3PAR](http://www.3par.com) (<http://www.3par.com>)
- ♦ [Bull](http://www.bull.com) (<http://www.bull.com>) (FC only)
- ♦ [Compellent](http://www.compellent.com) (<http://www.compellent.com>)
- ♦ [Dell](http://www.dell.com) (<http://www.dell.com>)

- ♦ [EMC \(http://www.emc.com\)](http://www.emc.com)
- ♦ [EqualLogic \(http://www.equallogic.com\)](http://www.equallogic.com) (iSCSI only)
- ♦ [Fujitsu \(http://www.fujitsu.com\)](http://www.fujitsu.com)
- ♦ [HP \(http://www.hp.com\)](http://www.hp.com)
- ♦ [Hitachi \(http://www.hitachi.com\)](http://www.hitachi.com) and [Hitachi Data Systems \(http://www.hds.com\)](http://www.hds.com) (FC only)
- ♦ [IBM \(http://www.ibm.com\)](http://www.ibm.com)
- ♦ [NEC \(http://www.nec.com\)](http://www.nec.com) (FC only)
- ♦ [Network Appliance \(NetApp\) \(http://www.netapp.com\)](http://www.netapp.com)
- ♦ [Nihon Unisys \(http://www.unisys.com\)](http://www.unisys.com) (FC only)
- ♦ [Pillar Data \(http://www.pillardata.com\)](http://www.pillardata.com) (FC only)
- ♦ [Sun Microsystems \(http://www.sun.com\)](http://www.sun.com)
- ♦ [Xiotech \(http://www.xiotech.com\)](http://www.xiotech.com) (FC only)

You can also learn more about iSCSI by visiting the Storage Networking Industry Association Web site at http://www.snia.org/tech_activities/ip_storage/iscsi/.

3.3.2 Adding a SAN LUN to Forge

PlateSpin Forge supports the use of Storage Area Network (SAN) storage, but before Forge can access an existing SAN, a SAN Logical Unit (LUN) needs to be added to Forge's ESX.

To add a SAN LUN to Forge:

- 1 Set up and configure your SAN system.
- 2 Access the appliance host (see "[Downloading the vSphere Client Program](#)" on page 36).
- 3 In the vSphere Client interface, click the root (top-level) node in the Inventory panel, then click the **Configuration** tab.
- 4 Click the **Add Storage** hyperlink in the upper right.
- 5 In the Add Storage Wizard, click **Next** until you are prompted to specify datastore information.
- 6 Specify a datastore name and click **Next** in the subsequent wizard pages. When the wizard finishes, click **Finish**.
- 7 Click **Storage** under **Hardware** to see the Forge datastores. The newly added SAN LUN should appear in the window.
- 8 Quit the vSphere Client program.

In the PlateSpin Forge Web Client, the new datastore does not appear until the next replication runs and the Application Host is refreshed. You can force a refresh by selecting **Settings > Containers** and clicking  near the appliance hostname.

3.4 Accessing and Working with the Forge Management VM in the Appliance Host

Occasionally you might need to access the Forge Management VM and perform maintenance tasks as described here or when you are advised to do so by PlateSpin Support.

Use the vSphere Client software to access the Forge Management VM, including its OS interface and VM settings.

- ♦ [Section 3.4.1, “Downloading the vSphere Client Program,” on page 36](#)
- ♦ [Section 3.4.2, “Launching the vSphere Client and Accessing the Forge Management VM,” on page 36](#)
- ♦ [Section 3.4.3, “Starting and Shutting Down the Forge Management VM,” on page 37](#)
- ♦ [Section 3.4.4, “Managing Forge Snapshots on the Appliance Host,” on page 37](#)
- ♦ [Section 3.4.5, “Manually Importing VMs into the Appliance Host’s Datastore,” on page 37](#)
- ♦ [Section 3.4.6, “Guidelines for Applying Security Updates to the PlateSpin Forge Management VM,” on page 38](#)

3.4.1 Downloading the vSphere Client Program

Download the client software from the appliance host and install it on a Windows workstation external to PlateSpin Forge.

To download the vSphere Client:

- 1 Download the client software:
 - ♦ For Forge appliance version 3 with VMware ESXi 5.5 Update 1, download the [VMware vSphere Client 5.5 Update 1 program](#).
- 2 Launch the downloaded installation program and follow the instructions to install the software.

3.4.2 Launching the vSphere Client and Accessing the Forge Management VM

To launch the vSphere Client:

- 1 Clicking **Start > Programs > VMWare > VMware vSphere | Virtual InfrastructureClient**.
The vSphere Client login window is displayed.
- 2 Specify your administrator-level credentials and log in, ignoring any certificate warnings.
The vSphere Client program opens.
- 3 In the inventory panel at the left, locate and select the **PlateSpin Forge VM** item. At the top of the right panel, click the **Console** tab.
The Client’s console area displays the Forge Management VM’s Windows interface.

Use the console to work with the Management VM the same way as you would work with Windows on a physical machine.

To unlock the Management VM, click inside the console and press **Ctrl+Alt+Insert**.

To release the cursor for working outside the vSphere Client program, press **Ctrl+Alt**.

3.4.3 Starting and Shutting Down the Forge Management VM

Occasionally you might need to shut down and then restart the Forge Management VM, such as when you relocate the appliance.

To shut down and then restart the VM:

- 1 Use the vSphere Client to access the Forge Management VM host. See [“Downloading the vSphere Client Program” on page 36](#).
- 2 Use the standard Windows procedure to shut down the VM (**Start > Shut Down**).

To restart the Management VM:

- 1 In the inventory panel at the left, right-click the **PlateSpin Forge Management VM** item and select **Power on**.

3.4.4 Managing Forge Snapshots on the Appliance Host

Occasionally you might need to take a point-in-time snapshot of your management VM, such as when you upgrade Forge software or when carry out troubleshooting tasks. You might also need to remove snapshots (recovery points) to free storage space.

To manage snapshots of the Forge Management VM:

- 1 Use the vSphere Client to access the appliance host. See [“Downloading the vSphere Client Program” on page 36](#).
- 2 In the inventory panel at the left, right-click the **PlateSpin Forge Management VM** item and select **Snapshot > Take Snapshot**.
- 3 Type a name and a description for the snapshot, then click **OK**.

To revert the management VM to a previous state:

- 1 In the inventory panel at the left, right-click the **PlateSpin Forge Management VM** item and select **Snapshot > Snapshot Manager**.
- 2 In the tree representation of the VM states, select a snapshot, then click **Go to**.

To remove snapshots that represent recovery points:

- 1 In the inventory panel at the left, right-click the **PlateSpin Forge Management VM** item and select **Snapshot > Snapshot Manager**.
- 2 In the tree representation of the VM states, select a snapshot, then click **Remove**.

3.4.5 Manually Importing VMs into the Appliance Host’s Datastore

To manually import a VM into the appliance host’s datastore:

- 1 At the production site, create a VM (ESX 3.5 and later) from your production workload (for example, by using PlateSpin Migrate) and copy the VM files from the ESX host’s datastore to portable media, such as a portable hard drive or a USB flash drive. Use the Datastore Browser of the client software to browse and locate the files.
- 2 At the disaster recovery site, attach the media to a workstation that has network access to Forge and has the vSphere Client program installed. See [“Downloading the vSphere Client Program” on page 36](#).

- 3 Use the vSphere Client's Datastore Browser to access the Forge datastore (**Storage1**) and upload the VM files from the temporary media. Use the uploaded VM to register it with the appliance host (right-click > **Add to Inventory**).
- 4 Refresh the PlateSpin Forge inventory (in the PlateSpin Forge Web Client, click **Settings** > **Containers**, then click ↗ adjacent to the appliance host).

TIP: You might want to consider this option when you want your failover workload to be created differentially (see [“Initial Replication Method \(Full and Incremental\)”](#) on page 71).

3.4.6 Guidelines for Applying Security Updates to the PlateSpin Forge Management VM

To apply security patches to the Forge Management VM:

- 1 During a maintenance window, access the Forge Management VM by using the VMware vSphere Client program. See [“Downloading the vSphere Client Program”](#) on page 36.
- 2 From within the Forge Management VM's Windows interface, check for security updates from Microsoft.
- 3 Use the PlateSpin Forge Web Client to put PlateSpin Forge into maintenance mode by pausing all replication schedules and ensuring that any incomplete replications are complete.
- 4 Take a snapshot of the Forge Management VM. See [“Managing Forge Snapshots on the Appliance Host”](#) on page 37.
- 5 Download and install the required security patches. When the installation finishes, reboot the Forge Management VM.
- 6 Use the PlateSpin Forge Web Client to resume replications paused in [Step 3](#) and verify that replications are working properly.
- 7 Remove the snapshot of the Forge Management VM that you took in [Step 4](#). See [“Managing Forge Snapshots on the Appliance Host”](#) on page 37.

3.5 Resetting Forge to Factory Defaults

TIP: Depending on your Forge model, this process might take up to 45 minutes or longer.

To reset your Forge 11 Appliance (Version 3) unit to its factory default state:

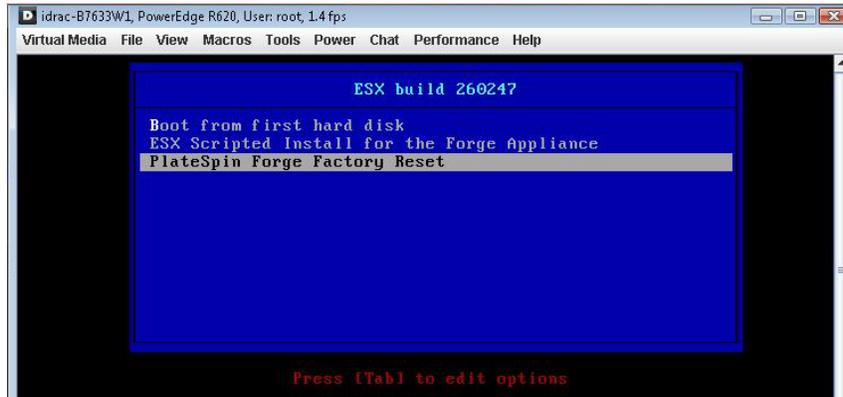
- 1 Disconnect all external/remote/shared storage systems from Forge (iSCSI, Fibre Channel, NFS).
- 2 Disconnect all network cables from Forge.

WARNING: If you are performing a factory reset on multiple Forge appliances connected to the same physical switch, skipping this step might cause IP address conflicts and result in failure.

- 3 Reboot the appliance host:
 - 3a Log in to the hypervisor (VMware ESXi) either directly or by using iDRAC.
 - 3b Press F2 to open the ESXi console.

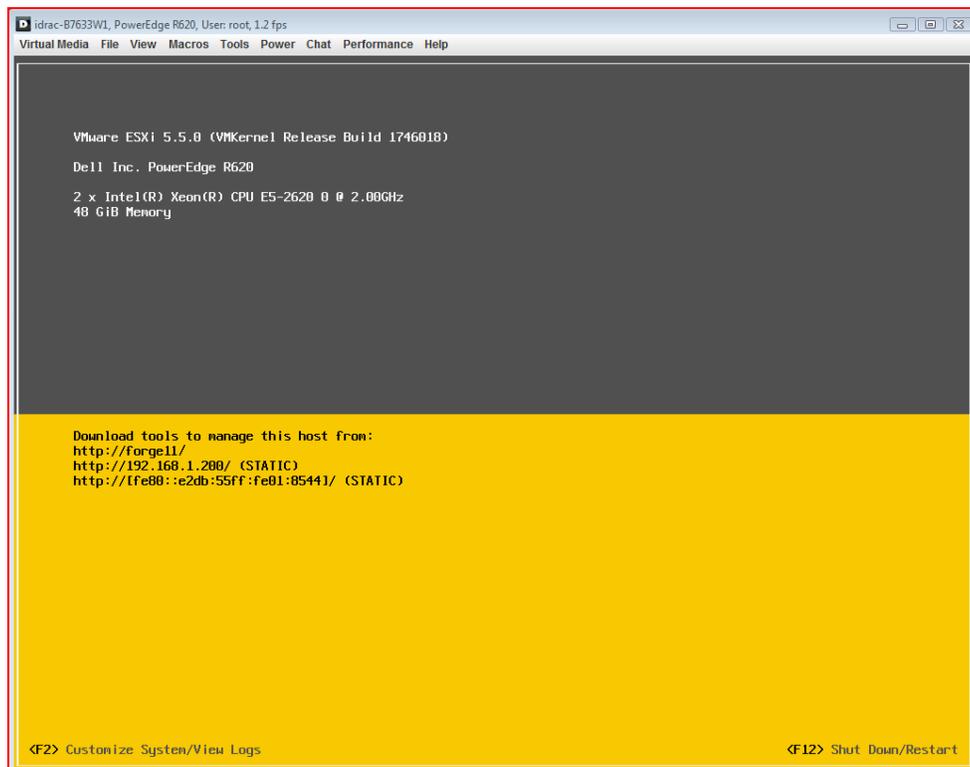
IMPORTANT: You need to remember the factory reset IP address of the appliance shown on this page. You need this address to log in to the Forge ACC and “relocate” the container to a known, valid IP address. Use the procedure documented in [Section 3.2, “Physically Relocating the Appliance,”](#) on page 32 to reset the IP properly.

- 3c Press F12 to shut down the ESXi console.
- 3d Log in with your administrator-level credentials.
- 3e Press F2 to shut down ESXi, or reboot the appliance.
- 3f Boot from the Forge CD media (or connect to the ISO via iDRAC) and wait for the SYSLINUX menu to display.

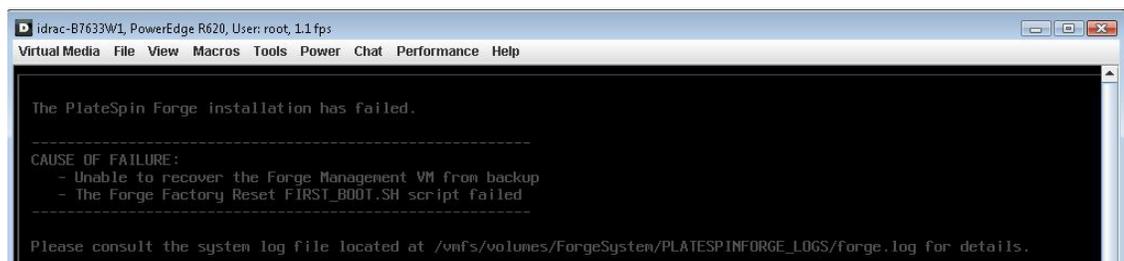


- 4 Select the **PlateSpin Forge Factory Reset** option and press Enter. Make sure that you do this before the default configuration is automatically applied. (about 10 seconds).
- 5 Wait for the Factory Reset process to complete.

If the reset process is successful, the command prompt window should look similar to the one below:



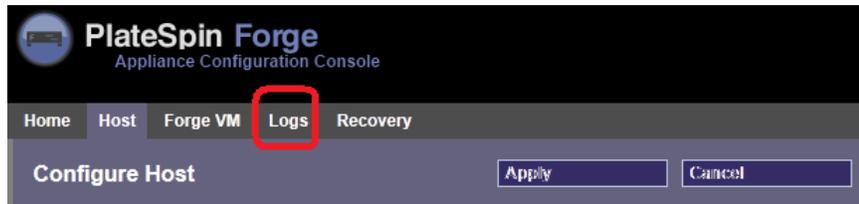
If the reset process is unsuccessful, the screen might look like the following:



In case of failure:

- ◆ Call PlateSpin Support and be prepared to provide the log files. Log files required for troubleshooting the reset process are:
 - ◆ /var/log/forge/forge-recovery.log
 - ◆ /var/log/forge/INSTALL_LOG.log
 - ◆ /var/log/weasel.log
 - ◆ /vmfs/volumes/forgeSystem/PLATESPINFORGE_LOGS/forge.log

The contents of these log files should also be available through the Forge ACC interface.



- ◆ Consider rebuilding Forge using a [Field Rebuild Kit](#) that you can obtain from PlateSpin Support.

4 Up and Running

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

- ♦ Section 4.1, “Launching the PlateSpin Forge Web Interface,” on page 43
- ♦ Section 4.2, “Elements of the PlateSpin Forge Web Interface,” on page 44
- ♦ Section 4.3, “Workloads and Workload Commands,” on page 46
- ♦ Section 4.4, “Managing Multiple Instances of PlateSpin Protect and PlateSpin Forge,” on page 48
- ♦ Section 4.5, “Generating Workload and Workload Protection Reports,” on page 50

4.1 Launching the PlateSpin Forge Web Interface

Most of your interaction with the appliance takes place through the browser-based PlateSpin Forge 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 launch the PlateSpin Forge Web Client:

- 1 Open a Web browser and go to:

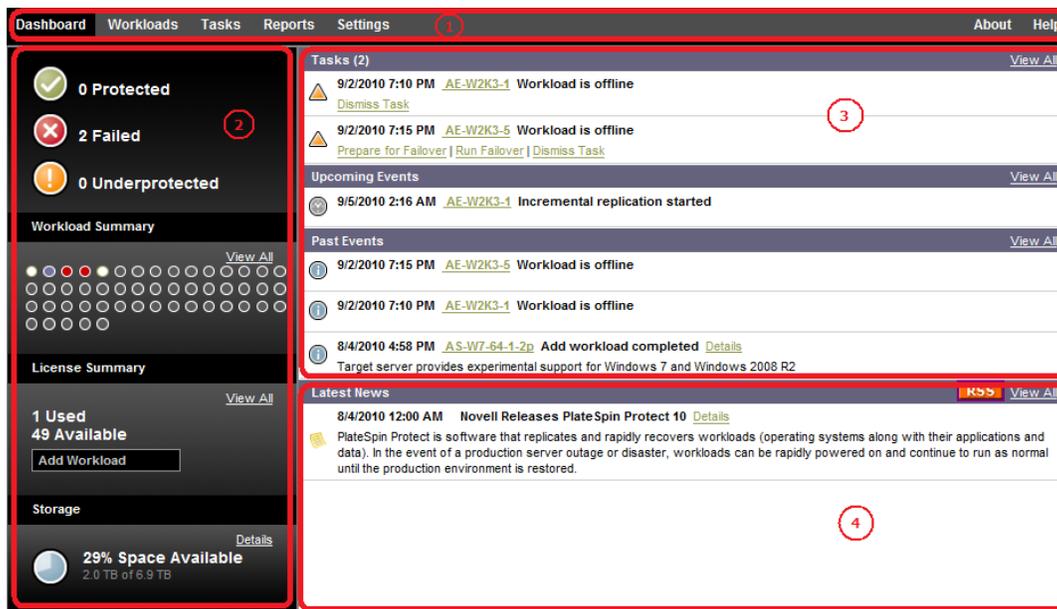
`http://<hostname | IP_address>/Forge`

NOTE: Replace `<hostname | IP_address>` with the hostname or the IP address of your Forge VM. If SSL is enabled, use `https` in the URL.

4.2 Elements of the PlateSpin Forge Web Interface

The default interface of the PlateSpin Forge 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 4-1 The Default Dashboard Page of the PlateSpin Forge Web Interface



The Dashboard page consists of the following elements:

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

The following topics provide more details:

- ◆ Section 4.2.1, “Navigation Bar,” on page 45
- ◆ Section 4.2.2, “Visual Summary Panel,” on page 45
- ◆ Section 4.2.3, “Tasks and Events Panel,” on page 46

4.2.1 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 46](#).
- ◆ **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 50](#).
- ◆ **Settings:** Displays the Settings page, which provides access to the following configuration options:
 - ◆ **Protection Tiers:** See [“Protection Tiers” on page 69](#).
 - ◆ **Permissions:** See [“Setting Up User Authorization and Authentication” on page 18](#).
 - ◆ **Email/SMTP:** See [“Setting Up Automatic Email Notifications of Events and Reports” on page 23](#).
 - ◆ **Licenses/License Designations:** See [“Product Licensing” on page 17](#).

4.2.2 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 4-1 *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 Forge.

4.2.3 Tasks and Events Panel

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

Events are logged whenever something relevant to the system or to the workload occurs. For example, an event could be the addition of a new protected workload, the replication of a workload starting or failing, or the detection of the failure of a protected workload. Some events generate automatic notifications by email if SMTP is configured. See [“Setting Up Automatic Email Notifications of Events and Reports”](#) on page 23.

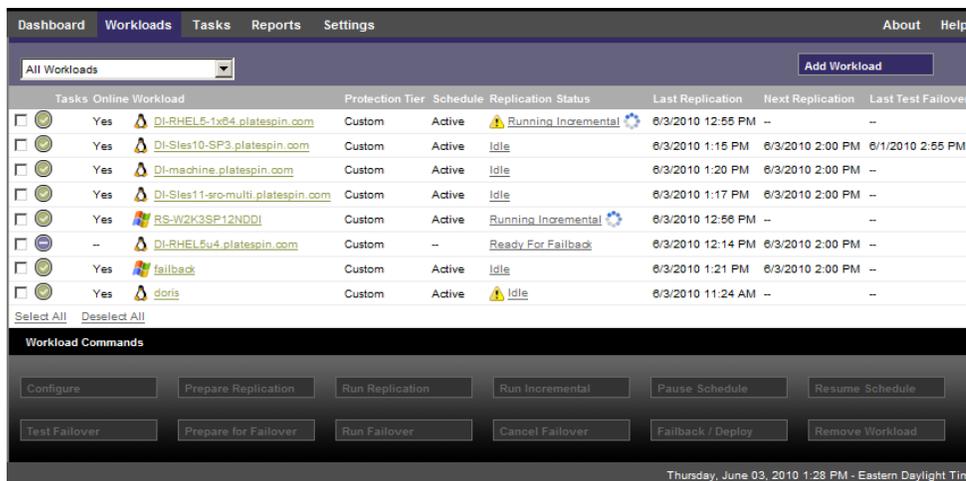
Tasks are special commands that are tied to events that require user intervention. For example, upon completion of a Test Failover command, the system generates an event associated with two tasks: Mark Test as Success and Mark Test as Failure. Clicking either task results in the Test Failover operation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

4.3 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 4-2 *The Workloads Page*

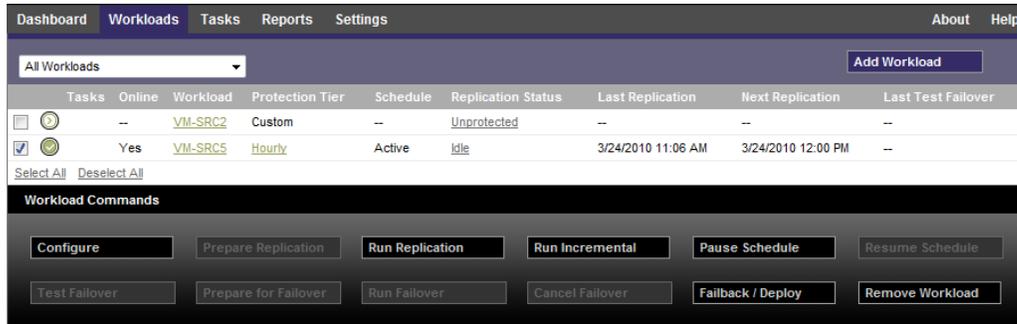


NOTE: All time stamps reflect the time zone of the Forge VM. This might be different from the time zone of the protected workload or the time zone of the host on which you are running the PlateSpin Forge Web Interface. A display of the server date and time appears at the bottom right of the client window.

4.3.1 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 4-3 Workload Commands



The following table summarizes workload commands along with their functional descriptions.

Table 4-2 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. |

4.4 Managing Multiple Instances of PlateSpin Protect and PlateSpin Forge

PlateSpin Forge includes a Web-based client application, the PlateSpin Forge 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 Forge, you can designate one of the instances as the manager and run the management console from there. Other instances are added under the Manager, providing a single point of control and interaction.

- [Section 4.4.1, “Using the PlateSpin Forge Management Console,” on page 48](#)
- [Section 4.4.2, “About PlateSpin Forge Management Console Cards,” on page 48](#)
- [Section 4.4.3, “Adding Instances of PlateSpin Protect and PlateSpin Forge to the Management Console,” on page 49](#)
- [Section 4.4.4, “Managing Cards on the Management Console,” on page 50](#)

4.4.1 Using the PlateSpin Forge Management Console

To begin using the Management Console:

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

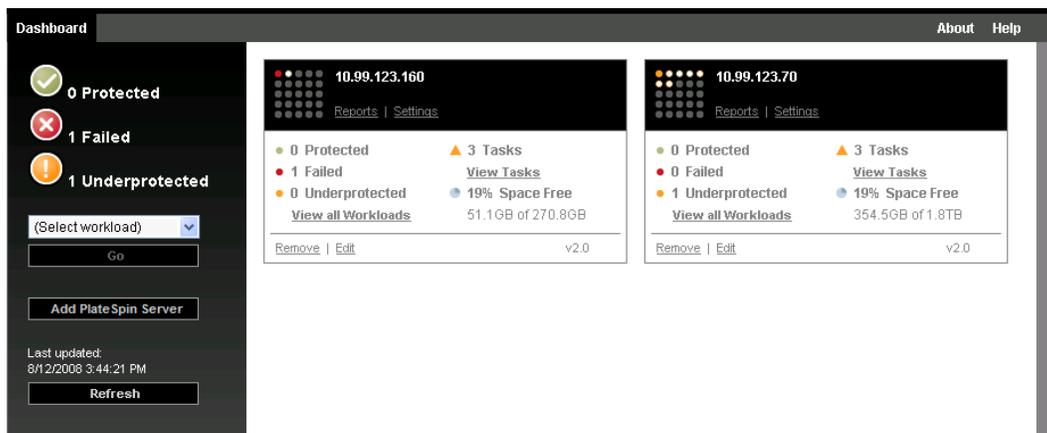
```
https://<IP_address | hostname>/console
```

Replace <IP_address | hostname> with either the IP address or the hostname of the Forge VM that is designated as the Manager.

- 2 Log in with your username and password.

The console’s default Dashboard page is displayed.

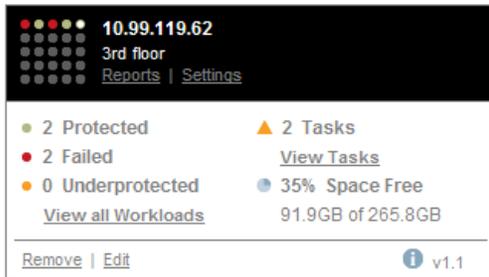
Figure 4-4 The Management Console’s Default Dashboard Page



4.4.2 About PlateSpin Forge Management Console Cards

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

Figure 4-5 PlateSpin Forge 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.

4.4.3 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.

4.4.4 Managing Cards on the Management Console

To 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.

4.5 Generating Workload and Workload Protection Reports

PlateSpin Forge 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 4-6 Options for a Replication History Report

| 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 |

To generate a report:

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

5 Workload Protection

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

- ♦ [Section 5.1, “Basic Workflow for Workload Protection and Recovery,” on page 53](#)
- ♦ [Section 5.2, “Adding Workloads for Protection,” on page 54](#)
- ♦ [Section 5.3, “Configuring Protection Details and Preparing the Replication,” on page 55](#)
- ♦ [Section 5.4, “Starting the Workload Protection,” on page 58](#)
- ♦ [Section 5.5, “Aborting Commands,” on page 58](#)
- ♦ [Section 5.6, “Failover,” on page 59](#)
- ♦ [Section 5.7, “Failback,” on page 61](#)
- ♦ [Section 5.8, “Reprotecting a Workload,” on page 65](#)

5.1 Basic Workflow for Workload Protection and Recovery

PlateSpin Forge 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 Forge supports your workload.
See [Section 1.2, “Supported Configurations,” on page 11](#).
 - 1b** Make sure that your workloads meet access and network prerequisites.
See [Section 2.3, “Access and Communication Requirements across your Protection Network,” on page 21](#).
 - 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 your freeze and thaw scripts to execute on your source workload upon each replication.
See [Section 6.8, “Using Freeze and Thaw Scripts for Every Replication \(Linux\),”](#) on page 72.
- 2 Inventory:** This step involves adding workloads to the PlateSpin Server database.
See [Section 5.2, “Adding Workloads for Protection,”](#) on page 54.
 - 3 Definition of the protection contract:** In this step, you define the details and specifications of a protection contract and prepare the replication.
See [Section 5.3, “Configuring Protection Details and Preparing the Replication,”](#) on page 55.
 - 4 Initiating the Protection:** This step commences the protection contract according to your requirements.
See [Section 5.4, “Starting the Workload Protection,”](#) on page 58.
 - 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 [Section 5.6.3, “Using the Test Failover Feature,”](#) on page 60.
 - 6 Failover:** This step carries out a failover of your protected workload to its replica running in your appliance host. See [Section 5.6, “Failover,”](#) on page 59.
 - 7 Failback:** This step corresponds to the business resumption phase after you have addressed any problems with your production workload. See [Section 5.7, “Failback,”](#) on page 61.
 - 8 Reprotection:** This step enables you to redefine the original protection contract for your workload. See [Section 5.8, “Reprotecting a Workload,”](#) on page 65

Most of these steps are represented by workload commands on the Workloads page. See [Section 4.3, “Workloads and Workload Commands,”](#) on page 46.

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

5.2 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 53.
- 2** On the Dashboard or Workloads page, click **Add Workload**.
The PlateSpin Forge Web Interface displays the Add Workload page.

| Name | Description | CPU | Memory | Free Space | Last Refresh |
|---------|---------------------------------|---|---------|------------|----------------------|
| comp129 | VMware ESX Server 4.0.0.261974 | 8 x Intel(R) Xeon(R) CPU X5355 @ 2.66GHz | 15.6 GB | -- | 48 Day(s) ago Remove |
| comp213 | VMware ESXi Server 4.1.0.260247 | 16 x Intel(R) Xeon(R) CPU E5530 @ 2.40GHz | 32.0 GB | 1.9 TB | 0 Hour(s) ago Remove |

3 Specify the required workload details:

- ◆ **Workload Settings:** Specify your workload’s hostname or IP address, the operating system, administrator-level credentials.

Use the required credential format. See [Section 6.2, “Guidelines for Workload Credentials,” on page 68.](#)

To make sure that PlateSpin Forge can access the workload, click **Test Credentials**.

4 Click **Add Workload**.

PlateSpin Forge 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 [Section 5.3, “Configuring Protection Details and Preparing the Replication,” on page 55](#)

5.3 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 [Section 5.1, “Basic Workflow for Workload Protection and Recovery,” on page 53](#)), relevant settings are read from the protection details.

To configure your workload’s protection details:

- 1 Add a workload. See [Section 5.2, “Adding Workloads for Protection,” on page 54.](#)
- 2 On the Workloads page, select the required workload and click **Configure**.

Alternatively, you can click the name of the workload.

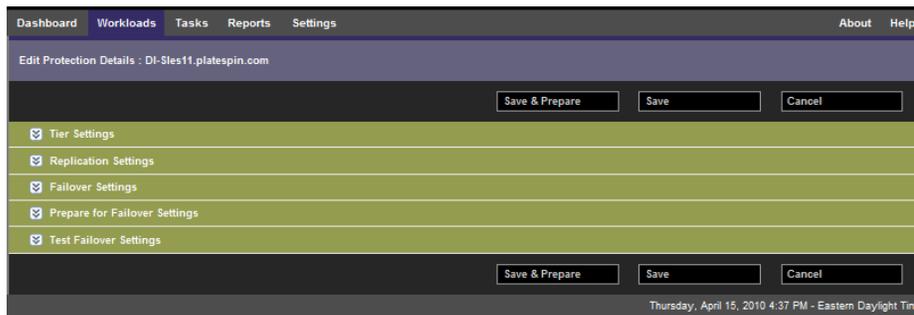
- 3 Select an **Initial Replication Method**. This indicates whether you want volume data transferred entirely from your workload to the failover VM or synchronized with volumes on an existing VM. See [Section 6.6, “Initial Replication Method \(Full and Incremental\),” on page 71](#).
- 4 Configure the protection details in each set of settings as dictated by your business continuity needs. See [“Workload Protection Details” on page 56](#).
- 5 Correct any validation errors, if displayed by the PlateSpin Forge Web Interface.
- 6 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.

5.3.1 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:

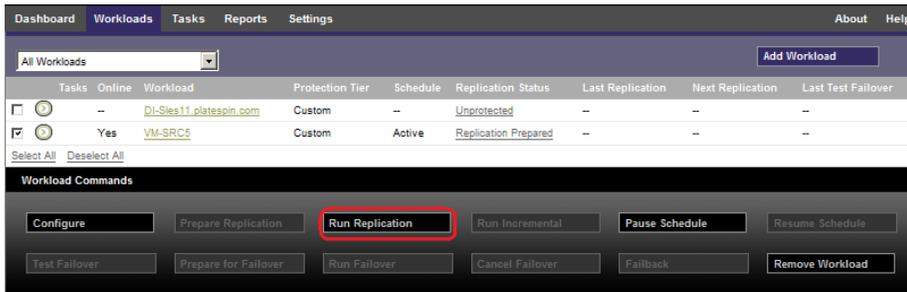
Table 5-1 Workload Protection Details

| Parameter Set (Settings) | Details |
|--------------------------|---|
| Tier | Indicates the Protection Tier that the current protection uses. See Section 6.4, “Protection Tiers,” on page 69 . |

| Parameter Set (Settings) | Details |
|--------------------------|---|
| Replication | <p>Transfer Method: (Windows) Enables you to select a data transfer mechanism and security through encryption. See Section 6.3, “Data Transfer,” on page 68.</p> <p>Transfer Encryption: To enable encryption, select the Encrypt Data Transfer option. See Section 1.3, “Security and Privacy,” on page 14.</p> <p>Source Credentials: Required for accessing the workload. See Section 6.2, “Guidelines for Workload Credentials,” on page 68.</p> <p>Number of CPUs: Enables you to specify the required number of vCPUs assigned to the failover workload (applicable only when the selected method of initial replication is Full).</p> <p>Replication Network: Enables you to separate replication traffic based on virtual networks defined on your appliance host. See Section 6.10, “Networking,” on page 75.</p> <p>Configuration File Datastore: Enables you to select a datastore associated with your appliance host for storing VM configuration files. See Section 6.5, “Recovery Points,” on page 71.</p> <p>Protected Volumes: Use these options to select volumes for protection and to assign their replicas to specific datastores on your appliance host.</p> <p>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.</p> <p>Services/Daemons to Stop During Replication: Enables you to select Windows services or Linux Daemons that are automatically stopped during the replication. See Section 6.7, “Service and Daemon Control,” on page 72.</p> |
| Failover | <p>VM Memory: Enables you to specify 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 Section 6.10, “Networking,” on page 75.</p> <p>Service/Daemon States to Change: Enables you to control the startup state of specific application services (Windows) or daemons (Linux). See Section 6.7, “Service and Daemon Control,” on page 72.</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 75.</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 Section 6.10, “Networking,” on page 75.</p> <p>Service/Daemon States to Change: Enables you to control the startup state of specific application services (Windows) or daemons (Linux). See Section 6.7, “Service and Daemon Control,” on page 72.</p> |

5.4 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 Forge 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.
-

5.5 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-2K6R2-VM1 | Custom | Active | | Idle | 3/5/2012 12:23 AM | 4/11/2012 12:00 AM | -- |
| <input type="checkbox"/> | | Yes | | DL-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-slesp2_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 | 3/28/2012 10:21 AM | 4/9/2012 12:00 PM | 3/23/2012 5:14 PM |
| <input type="checkbox"/> | | -- | | CL-VISTASPI-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 Forge Web Interface displays the appropriate Command Details page:

The screenshot shows the 'Command Details' page for a 'Running Incremental' command. The page includes a summary section with the following information:

- Command Summary:**
 - Events:** Incremental replication started (DEV-MORTAZAA\PlateSpin, 4/5/2012 2:00 PM)
 - Status:** Running (Warning: Controller installation has not finished in a timely fashion. A controller has already been installed on 10.99.123.164.)
 - Start Time:** 4/5/2012 2:00 PM
 - Duration:** 3d 21h 31m 37s
 - Steps:**
 - 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)

At the bottom of the page, under 'Workload Commands', there are three buttons: 'Abort', 'Configure', and 'Pause Schedule'. The 'Abort' button is highlighted with a red box.

3 Click **Abort**.

5.6 Failover

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

- ◆ [Section 5.6.1, "Detecting Offline Workloads," on page 59](#)
- ◆ [Section 5.6.2, "Performing a Failover," on page 60](#)
- ◆ [Section 5.6.3, "Using the Test Failover Feature," on page 60](#)

5.6.1 Detecting Offline Workloads

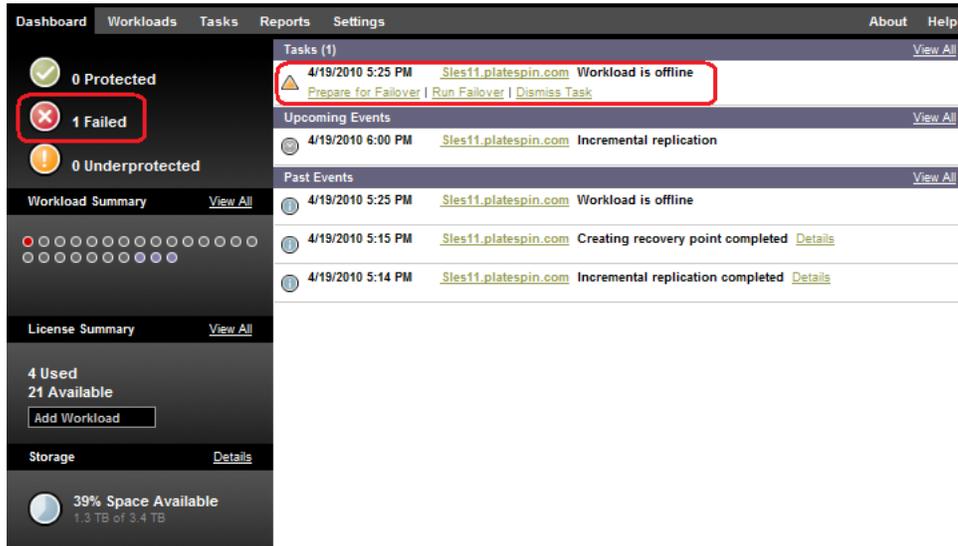
PlateSpin Forge constantly monitors your protected workloads. If an attempt to monitor a workload fails for a predefined number of times, PlateSpin Forge generates a **Workload is offline** event. Criteria that determine and log a workload failure are part of a workload protection's Tier settings (see the **Tier** row in [Section 5.3.1, "Workload Protection Details," on page 56](#)).

If notifications are configured along with SMTP settings, PlateSpin Forge simultaneously sends a notification email to the specified recipients. See [Section 2.4.1, "Setting Up Automatic Email Notifications of Events and Reports," on page 23](#).

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 [Section 5.5, "Aborting Commands," on page 58](#)), and then proceed to the **Run Failover** command. See [Section 5.6.2, "Performing a Failover," on page 60](#).

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

Figure 5-1 The Dashboard Page upon Workload Failure Detection ('Workload Offline')



5.6.2 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 [Section 5.3.1, "Workload Protection Details,"](#) on page 56.

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 5-1](#).
- ◆ 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 [Section 6.5, "Recovery Points,"](#) on page 71). 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 [Section 5.6.3, "Using the Test Failover Feature,"](#) on page 60.

5.6.3 Using the Test Failover Feature

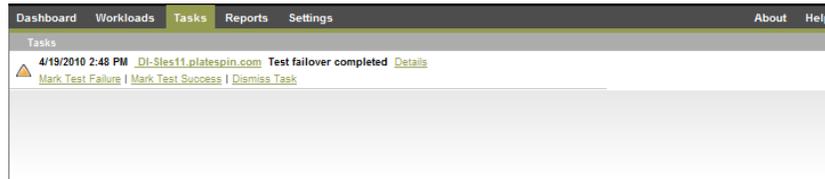
PlateSpin Forge provides you with the capability to test the failover functionality and the integrity of the failover workload. This is done by using the **Test Failover** command, which boots the failover workload in a restricted network environment for testing.

When you execute the command, PlateSpin Forge applies the Test Failover Settings, as saved in the workload protection details, to the failover workload (see the [Test Failover](#) row in [Section 5.3.1, "Workload Protection Details,"](#) on page 56).

To use the Test Failover feature:

- 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 [Section 6.5, “Recovery Points,”](#) on page 71), and then click **Execute**.

Upon completion, PlateSpin Forge 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 appliance host.
See [Section 3.4.1, “Downloading the vSphere Client Program,”](#) on page 36.
- 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 Forge discards temporary settings that were applied to the failover workload, and the protection returns to its pre-test state.

5.7 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.

The following topics provide more information:

- ♦ [Section 5.7.1, “Automated Failback to a VM Platform,”](#) on page 61
- ♦ [Section 5.7.2, “Semi-Automated Failback to a Physical Machine,”](#) on page 64

5.7.1 Automated Failback to a VM Platform

The following containers are supported as automated failback targets:

| Target | Notes |
|-----------------------------------|--|
| VMware DRS Cluster in vSphere 5.5 | <ul style="list-style-type: none"> ◆ 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 DRS Cluster in vSphere 5.1 | <ul style="list-style-type: none"> ◆ 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 DRS Cluster in vSphere 5.0 | <ul style="list-style-type: none"> ◆ 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 5.0 servers only, and can be managed by vCenter 5.0 only. |
| VMware DRS Cluster in vSphere 4.1 | <ul style="list-style-type: none"> ◆ The DRS configuration must be either Partially Automated or Fully Automated (it must not be set to Manual) ◆ As a VM Container, the Cluster, as a container, can use a combination of ESX 4.1 and ESXi 4.1 servers, and can be managed by vCenter 4.1 only |
| VMware ESXi 4.1, 5.0, 5.1 | ESXi versions must have a paid license; protection is unsupported with these systems if they are operating with a free license. |
| VMware ESX 4.1 | |

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 [Section 6.2, "Guidelines for Workload Credentials,"](#) on page 68).
- ◆ **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 [Section 6.6, "Initial Replication Method \(Full and Incremental\),"](#) on page 71.
 - ◆ **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.

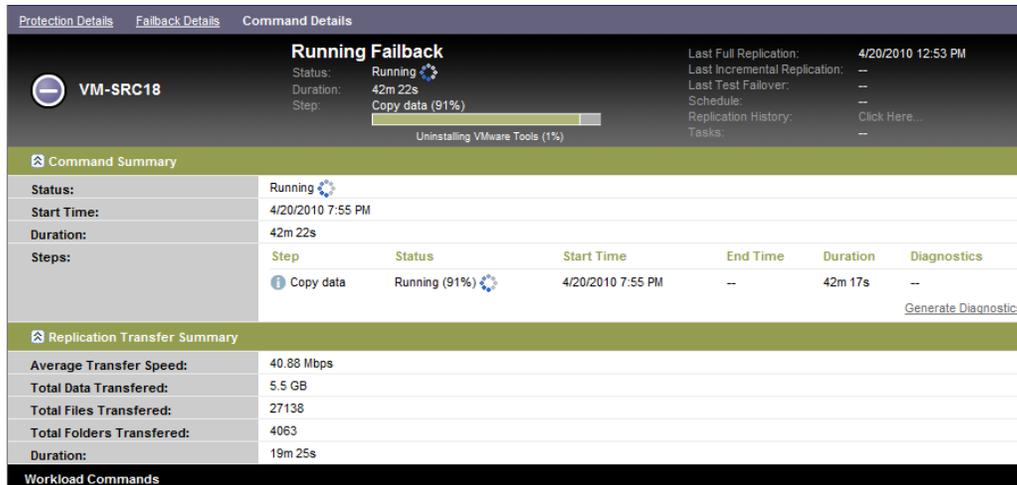
Upon successful completion, PlateSpin Forge loads the Ready for Failback screen, prompting you to specify the details of the failback operation.

- 4 Configure the failback details. See ["Failback Details \(Workload to VM\)"](#) on page 63.

- 5 Click **Save and Failback** and monitor the progress on the Command Details page. See [Figure 5-2](#).

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

Figure 5-2 Failback Command Details



Failback Details (Workload to VM)

Failback details are represented by three sets of parameters that you configure when you are performing a workload failback operation to a virtual machine.

Table 5-2 Failback Details (VM)

| Parameter Set (Settings) | Details |
|--------------------------|--|
| Failback | <p>Transfer Method: Enables you to select a data transfer mechanism and security through encryption. See Section 6.3, “Data Transfer,” on page 68.</p> <p>Failback Network: Enables you to direct failback traffic over a dedicated network based on virtual networks defined on your appliance host. See Section 6.10, “Networking,” on page 75.</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 Section 6.7, “Service and Daemon Control,” on page 72.</p> <p>Alternative Address for Source: Accepts input of an additional IP address for the failed-over VM if applicable. See Section 2.3.2, “Protection Across Public and Private Networks Through NAT,” on page 22.</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 Section 6.7, “Service and Daemon Control,” on page 72.</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 Forge 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. |

5.7.2 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.

To fail a workload back to a physical machine:

- 1 Register the required physical machine with your PlateSpin Server. See [Section 6.11, “Failback to Physical Machines,”](#) on page 75.
- 2 If the drivers are missing or incompatible, upload the required drivers to the PlateSpin Forge device driver database. See [Section 7.1, “Managing Device Drivers,”](#) on page 81.
- 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 [Section 6.2, “Guidelines for Workload Credentials,”](#) on page 68).
 - ◆ **Failback Target Settings:** Specify the following parameters:
 - ◆ **Replication Method:** Select the scope of data replication. See [Section 6.6, “Initial Replication Method \(Full and Incremental\),”](#) on page 71.
 - ◆ **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 in the PlateSpin Forge interface. It is divided into three main sections:

- Workload Settings:**
 - Hostname or IP: MA-Rhel5u3
 - Credentials:
 - User Name: root
 - Password: [masked]
 - Test Credentials: [link]
- Failback Target Settings:**
 - Replication Method:
 - Full Replication
 - Incremental Replication
 - Target type:
 - Virtual Targets
 - Physical Targets
 - Failback Target: [Selection required below] x
 - No physical targets available.
 - 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 [button]

- 5 Click **Save and Prepare** and monitor the progress on the Command Details screen. Upon successful completion, PlateSpin Forge 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.

5.8 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 Forge 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 [Section 5.7, “Failback,” on page 61](#).

The remainder 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.

6 Essentials of Workload Protection

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

- ◆ [Section 6.1, “Workload License Consumption,” on page 67](#)
- ◆ [Section 6.2, “Guidelines for Workload Credentials,” on page 68](#)
- ◆ [Section 6.3, “Data Transfer,” on page 68](#)
- ◆ [Section 6.4, “Protection Tiers,” on page 69](#)
- ◆ [Section 6.5, “Recovery Points,” on page 71](#)
- ◆ [Section 6.6, “Initial Replication Method \(Full and Incremental\),” on page 71](#)
- ◆ [Section 6.7, “Service and Daemon Control,” on page 72](#)
- ◆ [Section 6.8, “Using Freeze and Thaw Scripts for Every Replication \(Linux\),” on page 72](#)
- ◆ [Section 6.9, “Volumes,” on page 73](#)
- ◆ [Section 6.10, “Networking,” on page 75](#)
- ◆ [Section 6.11, “Failback to Physical Machines,” on page 75](#)
- ◆ [Section 6.12, “Advanced Workload Protection Topics,” on page 77](#)

6.1 Workload License Consumption

Your PlateSpin Forge 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 [Section 2.1, “Product Licensing,” on page 17](#).

6.2 Guidelines for Workload Credentials

PlateSpin Forge must have administrator-level access to workloads. Throughout the workload protection and recovery workflow, PlateSpin Forge prompts you to specify credentials that must be provided in a specific format.

Table 6-1 Workload 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 . |

6.3 Data Transfer

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

- ◆ [Section 6.3.1, “Transfer Methods,” on page 68](#)
- ◆ [Section 6.3.2, “Data Encryption,” on page 69](#)

6.3.1 Transfer Methods

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

- ◆ [“Transfer Methods Supported for Windows Workloads” on page 68](#)
- ◆ [“Transfer Methods Supported for Linux Workloads” on page 69](#)

Transfer Methods Supported for Windows Workloads

For Windows workloads, PlateSpin Forge 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 Forge provides two mechanisms that differ by their continuity impact and performance. You can toggle between these mechanisms as required.

No reboot is required when you are protecting Windows clusters with block-level data transfer.

- ♦ **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. 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 Forge 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 Forge 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.

6.3.2 Data Encryption

To make the transfer of workload data more secure, PlateSpin Forge 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.

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

6.4 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.

To create custom Protection Tiers in advance:

- 1 In your PlateSpin Forge 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 | These settings control whether and how workload data is compressed before transmission. See Section 1.4.2, "Data Compression," on page 15 . Select one of the available options. Fast consumes the least CPU resources on the source but yields a lower compression ratio, Maximum consumes the most, but yields a higher compression ratio. Optimal , the middle ground, is the recommended option. |
| Bandwidth Throttling | These settings control bandwidth throttling. See Section 1.4.3, "Bandwidth Throttling," on page 15 . To throttle replications to a specified rate, specify the required throughput value in Mbps and indicate the time pattern. |
| Recovery Points to Keep | Specify the number of recovery points to keep for workloads that use this Protection Tier. See Section 6.5, "Recovery Points," on page 71 . |
| Workload Failure | Specify the number of workload detection attempts before it is considered failed. |
| Workload Detection | Specify the time interval (in seconds) between workload detection attempts. |

6.5 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 Forge storage to run out of space.

To remove recovery points from your appliance, see [Section 3.4.4, “Managing Forge Snapshots on the Appliance Host,”](#) on page 37.

6.6 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.

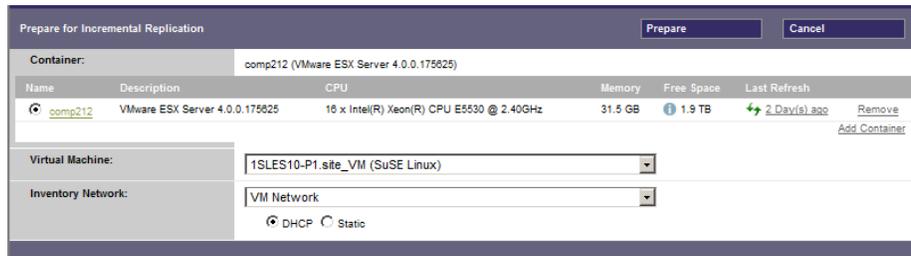
- ♦ **Full:** A full volume transfer takes place from a production workload to its replica (the failover workload), or from a failover workload to its original virtual or physical infrastructure.
- ♦ **Incremental:** Only the differences are transferred from a source to its target, provided that they have similar operating system and volume profiles.
 - ♦ During protection: The production workload is compared with an existing VM in the appliance host. 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 appliance host, such as a workload VM physically moved on portable media from the production site to a remote recovery site. See [Section 3.4.5, “Manually Importing VMs into the Appliance Host’s Datastore,”](#) on page 37.
 - ♦ 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 Forge (see [Section 5.7.2, “Semi-Automated Failback to a Physical Machine,”](#) on page 64).

During workload protection and failback to a VM host, selecting **Incremental** as the initial replication method requires that you browse, locate, and prepare the target VM for synchronization with the selected operation’s source.

To set up the initial replication method:

- 1 Proceed with the required workload command, such as **Configure (Protection Details)** or **Failback**.
- 2 For the **Initial Replication Method** option, select **Incremental Replication**.
- 3 Click **Prepare Workload**.

The PlateSpin Forge 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.

6.7 Service and Daemon Control

PlateSpin Forge 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 [Section 6.8, “Using Freeze and Thaw Scripts for Every Replication \(Linux\),”](#) on page 72.

- ♦ **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.

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

For Linux systems, PlateSpin Forge 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 72. 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 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
```

6.9 Volumes

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

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

For Linux workloads, PlateSpin Forge 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 appliance host. 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 6-1 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 (869.7 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 6-2 Replication Settings, Volume-Related Options (OES 2 Workload)

| | | | | | | |
|-------------------------------------|-------------------------------------|---|---------------------|---------------------|----------------------------|--------------------------|
| Protected Logical Volumes: | Include | Name | Used Space | Free Space | Volume Group / EVMS Volume | |
| | <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/hss/mnt/pools/NEWPOOL (NSSFS) | 23.3 MB | 999.6 MB | NEWPOOL | |
| Non-volume Storage: | Include | Partition | Is Swap | Total Size | Datastore / Volume Group | |
| | <input checked="" type="checkbox"/> | /dev/system/swap | Yes | 1.48 GB | system | |
| Volume Groups: | Include | Name | Total Size | Datastore | Thin Disk | |
| | <input checked="" type="checkbox"/> | system | 5.9 GB | dev-comp124:storage | <input type="checkbox"/> | |
| EVMS Volumes: | Include | Name | Datastore | Total Size | Datastore | Thin Disk |
| | <input checked="" type="checkbox"/> | /dev/evms/sda1 | dev-comp124:storage | 70.6 MB | dev-comp124:storage | <input type="checkbox"/> |
| | <input checked="" type="checkbox"/> | NEWPOOL | dev-comp124:storage | 1023.0 MB | dev-comp124:storage | <input type="checkbox"/> |
| Daemons to Stop During Replication: | Add Daemons | | | | | |

6.10 Networking

PlateSpin Forge 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.

6.11 Failback to Physical Machines

If the required target infrastructure for a failback operation is a physical machine, you must register it with PlateSpin Forge.

The registration of a physical machine is carried out by booting the target physical machine with the PlateSpin boot ISO image.

- ♦ [Section 6.11.1, "Downloading the PlateSpin Boot ISO Images," on page 75](#)
- ♦ [Section 6.11.2, "Injecting Additional Device Drivers into the Boot ISO Image," on page 75](#)
- ♦ [Section 6.11.3, "Registering Physical Machines as Failback Targets with PlateSpin Forge," on page 76](#)

6.11.1 Downloading the PlateSpin Boot ISO Images

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 Forge area of [NetIQ Downloads \(https://dl.netiq.com\)](https://dl.netiq.com) by doing a search with the following parameters:

- ♦ **Product or Technology:** PlateSpin Forge
- ♦ **Select Version:** PlateSpin Forge 11
- ♦ **Date Range:** All Dates

6.11.2 Injecting Additional Device Drivers into the Boot ISO Image

You can use a custom utility to package and inject additional Linux device drivers into the PlateSpin boot image before burning it on a CD.

To use this utility:

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

IMPORTANT: 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 \(https://www.netiq.com/support/kb/doc.php?id=7005990\)](https://www.netiq.com/support/kb/doc.php?id=7005990).

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

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

- 3 Copy the `rebuildiso.sh` script, located in the `/tools` subdirectory of the mounted ISO file, into a temporary working directory. When you have finished, unmount the ISO file (execute the command `umount <mount_point>`).
- 4 Create another working directory for the required driver files and save them in that directory.
- 5 In the directory where you saved the `rebuildiso.sh` script, run the following command as root to copy SOURCE files into the ISO file:

```
./rebuildiso.sh <SOURCE> <-m32|-m64> <-i ISO file>
```

NOTE: SOURCE must be one or more of the following parameters:

- d path to the directory containing drivers (that is, *.ko files) to inject
 - c Path to the `ConfigureTakeControl.xml` file
-

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

6.11.3 Registering Physical Machines as Failback Targets with PlateSpin Forge

To register physical machines as failback targets:

- 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 Forge VM.
- 7 Provide your administrator-level credentials for the Forge VM, 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 Forge Web Interface.

6.12 Advanced Workload Protection Topics

- ◆ [Section 6.12.1, “Protecting Windows Clusters,” on page 77](#)
- ◆ [Section 6.12.2, “Using Workload Protection Features through the PlateSpin Forge Web Services API,” on page 79](#)

6.12.1 Protecting Windows Clusters

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

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

This section includes the following information:

- ◆ [“Workload Protection” on page 77](#)
- ◆ [“Protection Failover” on page 78](#)
- ◆ [“Protection Failback” on page 78](#)

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 collocated 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 a node failover occurs 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 Forge 11:

- ◆ (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 Cluster Node Local Storage,” on page 109](#).
- ◆ (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 failover to a new node in the cluster.

If a node failover occurs 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 53](#)).

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.

6.12.2 Using Workload Protection Features through the PlateSpin Forge Web Services API

PlateSpin Forge 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

Forge 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 Forge 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.

7 Auxiliary Tools for Working with Physical Machines

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

- ♦ [Section 7.1, “Managing Device Drivers,” on page 81](#)

7.1 Managing Device Drivers

PlateSpin Forge 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.

- ♦ [Section 7.1.1, “Packaging Device Drivers for Windows Systems,” on page 81](#)
- ♦ [Section 7.1.2, “Packaging Device Drivers for Linux Systems,” on page 82](#)
- ♦ [Section 7.1.3, “Uploading Drivers to the PlateSpin Forge Device Driver Database,” on page 82](#)
- ♦ [Section 7.1.4, “Using the Plug and Play \(PnP\) ID Translator Feature,” on page 84](#)

7.1.1 Packaging Device Drivers for Windows Systems

To package your Windows device drivers for uploading to the PlateSpin Forge 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 [Section 7.1.3, “Uploading Drivers to the PlateSpin Forge Device Driver Database,” on page 82](#).

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 Vista and Windows Server 2008, and Windows 7 systems
-

7.1.2 Packaging Device Drivers for Linux Systems

To package your Linux device drivers for uploading to the PlateSpin Forge driver database, you can use a custom utility included in one of the PlateSpin boot ISO images:

- 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 appropriate 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 [Section 7.1.3, “Uploading Drivers to the PlateSpin Forge Device Driver Database,”](#) on page 82.

7.1.3 Uploading Drivers to the PlateSpin Forge Device Driver Database

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

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

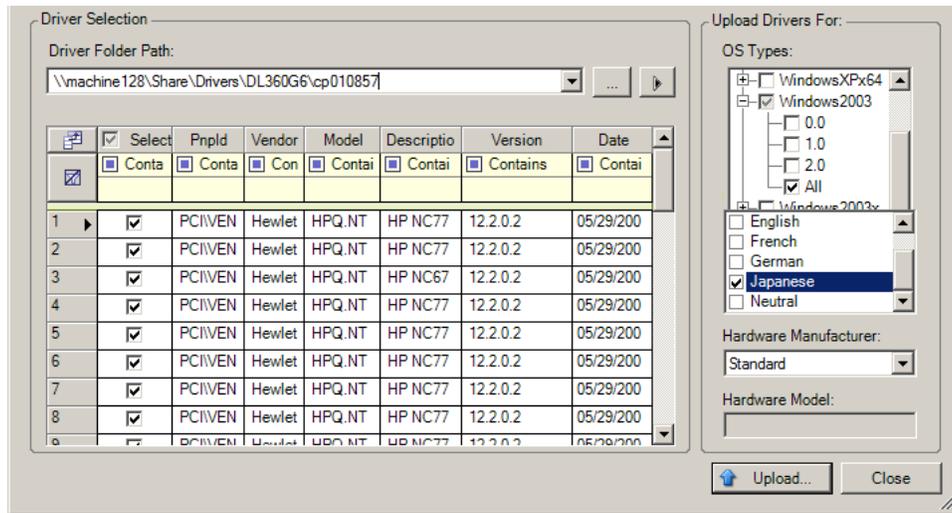
- ♦ [“Device Driver Upload Procedure \(Windows\)”](#) on page 82
- ♦ [“Device Driver Upload Procedure \(Linux\)”](#) on page 83

Device Driver Upload Procedure (Windows)

To upload a Windows device driver:

- 1 Obtain and prepare the required device drivers. See [Section 7.1.1, “Packaging Device Drivers for Windows Systems,”](#) on page 81.
- 2 On your Forge VM, under Program Files\PlateSpin Forge Server\DriverManager, start the DriverManager.exe program and select the **Windows Drivers** tab.

- 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.

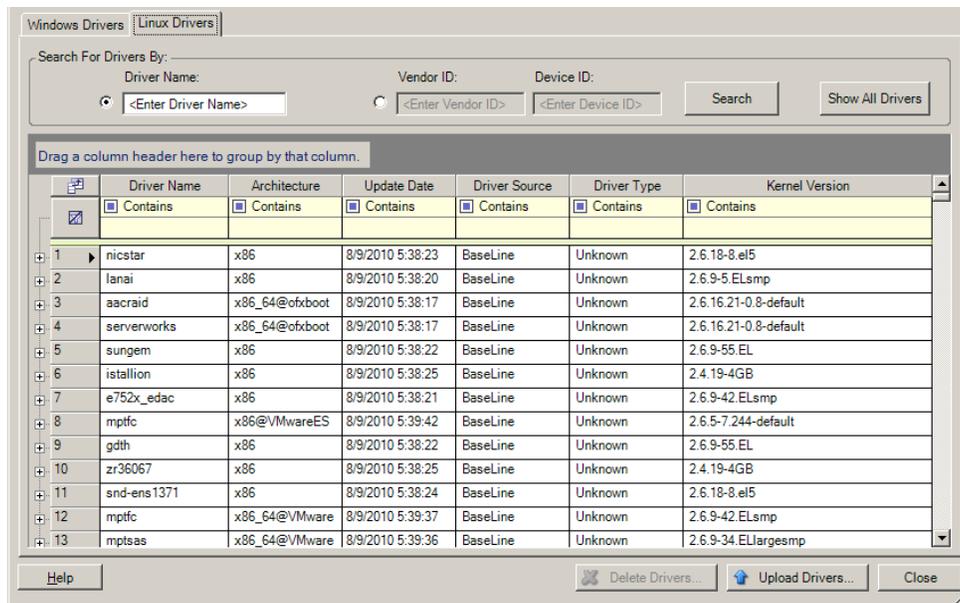
- Click **Upload** and confirm your selections when prompted.

The system uploads the selected drivers to the driver database.

Device Driver Upload Procedure (Linux)

To upload a Linux device driver:

- Obtain and prepare the required device drivers. See [Section 7.1.2, “Packaging Device Drivers for Linux Systems,”](#) on page 82.
- 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.

7.1.4 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.

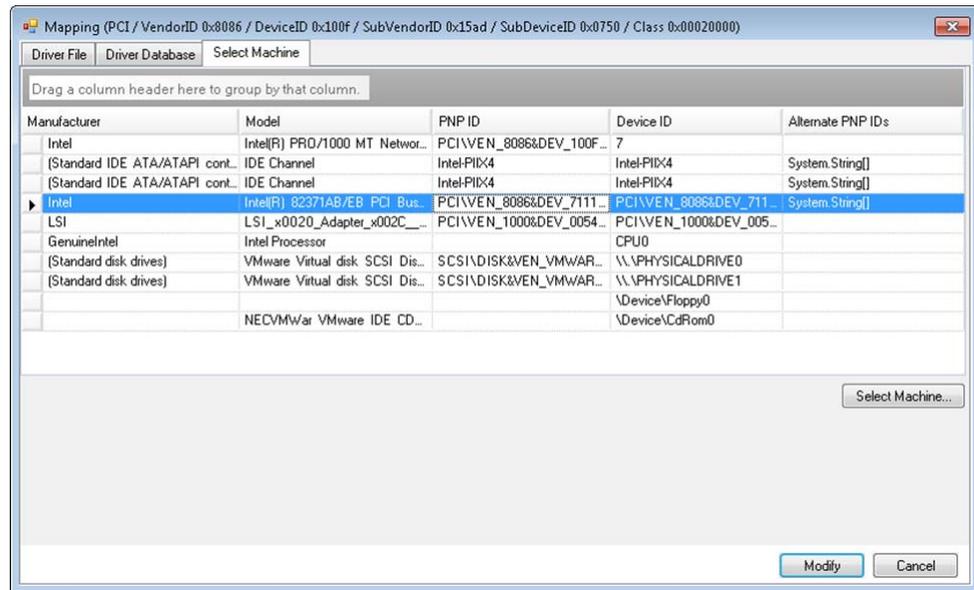
These IDs are formatted consistently, so PlateSpin can apply a standard transformation to each of them to determine its corresponding Windows PnP ID. The translation occurs automatically within the PlateSpin product. The feature lets you or a support technician add, edit or remove custom PnP mappings.

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.

6 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 84](#) to modify the Windows PnP ID.

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

8 Troubleshooting

- ◆ [Section 8.1, “Troubleshooting Workload Inventory \(Windows\),” on page 87](#)
- ◆ [Section 8.2, “Troubleshooting Workload Inventory \(Linux\),” on page 91](#)
- ◆ [Section 8.3, “Troubleshooting Problems during the Prepare Replication Command \(Windows\),” on page 91](#)
- ◆ [Section 8.4, “Troubleshooting Workload Replication,” on page 92](#)
- ◆ [Section 8.5, “Generating and Viewing Diagnostic Reports,” on page 93](#)
- ◆ [Section 8.6, “Removing Workloads,” on page 94](#)
- ◆ [Section 8.7, “Post-Protection Workload Cleanup,” on page 94](#)
- ◆ [Section 8.8, “Shrinking the PlateSpin Forge Databases,” on page 96](#)

8.1 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>A non-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 89 again. If the test succeeds, try adding the workload again.</p> <ul style="list-style-type: none">◆ “Troubleshooting DCOM Connectivity” on page 89◆ “Troubleshooting RPC Service Connectivity” on page 89 |
| 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 88. If a test fails, ensure that PlateSpin Forge and the workload are on the same network. Reconfigure the network and try again.</p> |

| Problems or Messages | Solutions |
|---|---|
| D"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 90. ◆ 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 Forge can access those shares. See "Enabling File/Share Permissions and Access" on page 90. ◆ 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. |

This section also contains the following information:

- ◆ [Section 8.1.1, "Performing Connectivity Tests," on page 88](#)
- ◆ [Section 8.1.2, "Disabling AntiVirus Software," on page 90](#)
- ◆ [Section 8.1.3, "Enabling File/Share Permissions and Access," on page 90](#)

8.1.1 Performing Connectivity Tests

- ◆ ["Network Connectivity Test" on page 88](#)
- ◆ ["WMI Connectivity Test" on page 89](#)
- ◆ ["Troubleshooting DCOM Connectivity" on page 89](#)
- ◆ ["Troubleshooting RPC Service Connectivity" on page 89](#)

Network Connectivity Test

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

- 1 Go to your Forge VM.
See ["Downloading the vSphere Client Program" on page 36](#).

- 2 Open a command prompt and ping your workload:

```
ping workload_ip
```

WMI Connectivity Test

- 1 Go to your Forge VM.
See [Section 3.4.1, “Downloading the vSphere Client Program,”](#) on page 36.
- 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 Forge and your workload.

Troubleshooting DCOM Connectivity

- 1 Log into the workload that you want to protect.
- 2 Click **Start > Run**.
- 3 Type `dcomcnfg` and press Enter.
- 4 Check connectivity:
 - ♦ For Windows systems (XP/Vista/2003/2008/7), the Component Services window is displayed. In the **Computers** folder of the console tree of the Component Services administrative tool, right-click the computer that you want to check for DCOM connectivity, then click **Properties**. Click the **Default Properties** tab and ensure that **Enable Distributed COM on this computer** is selected.
 - ♦ On a Windows 2000 Server machine, the DCOM Configuration dialog box is displayed. Click the **Default Properties** tab and ensure that **Enable Distributed COM on this computer** is selected.
- 5 If DCOM was not enabled, enable it and either reboot the server or restart the Windows Management Instrumentation Service. Then try adding the workload again.

Troubleshooting RPC Service Connectivity

There are three potential blockages for the RPC service:

- ♦ The Windows Service
- ♦ A Windows firewall
- ♦ A network firewall

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

- ♦ Putting PlateSpin Forge and the workload on the same side of the firewall
- ♦ Opening up specific ports between PlateSpin Forge and the workload (See [Section 2.3, “Access and Communication Requirements across your Protection Network,”](#) on page 21.

8.1.2 Disabling AntiVirus Software

Antivirus software might occasionally block some of the PlateSpin Forge 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.

8.1.3 Enabling File/Share Permissions and Access

To successfully protect a workload, PlateSpin Forge 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 Forge uses the workload’s administrative shares. PlateSpin Forge 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 Forge VM:

- 1 Go to your Forge VM.
See [Section 3.4.1, “Downloading the vSphere Client Program,”](#) on page 36.
- 2 Click **Start > Run**, type `\\<server_host>\Admin$`, then click **OK**.
- 3 If you are prompted, use the same credentials as those you will use to add the workload to the PlateSpin Forge workload inventory.
The directory is opened and you should be able to browse and modify its contents.
- 4 Repeat the process for all shares with the exception of the `IPC$` 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 Forge does not modify the existing content of the volume; however, it creates its own directory, to which it requires access and permissions.

8.2 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 Section 2.3, "Access and Communication Requirements across your Protection Network," on page 21.</p> |
| Access denied | <p>This authentication problem indicates either an invalid username or password. For information on proper workload access credentials, see Section 6.2, "Guidelines for Workload Credentials," on page 68.</p> |

8.3 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 Section 8.3.1, "Group Policy and User Rights," on page 91.</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 Section 8.1, "Troubleshooting Workload Inventory (Windows)," on page 87.</p> |

8.3.1 Group Policy and User Rights

Because of the way that PlateSpin Forge 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).

8.4 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 Forge VM). Allow it to run for 50 iterations and the average indicates the latency. Also see "Optimizing Data Transfer over WAN Connections" on page 27. |
| 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. |
| Controller connection not established Replication fails at the Take Control of Virtual Machine step. | This error occurs when the replication networking information is invalid. Either the DHCP server is not available or the replication virtual network is not routable to the Forge VM. Change the replication IP to a static IP or enable the DHCP server. Ensure that the virtual network selected for replication is routable to the Forge VM. |

| Problems or Messages | Solutions |
|--|---|
| Replication job does not start (stuck at 0%) | <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 Forge VM is affiliated with a domain and the domain policies are applied with restrictions. See Section 8.3.1, "Group Policy and User Rights," on page 91.</p> |

8.5 Generating and Viewing Diagnostic Reports

In the PlateSpin Forge 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 job. The job title is 'Running First Replication'. The status is 'Running' with a refresh icon. The duration is '14h 49m 6s'. The step is 'Copy data (80%)' with a progress bar. The 'Generate Diagnostics' link is 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.

8.6 Removing Workloads

In some circumstances you might need to remove a workload from the PlateSpin Forge 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 Forge 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.

8.7 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.

8.7.1 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"> 1. 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) 2. Run the following command: <pre>ofxcontroller.exe /uninstall</pre> 3. Remove the platespin* directory |

8.7.2 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"> 1. Check if the driver is active: <pre>lsmod grep blkwatch</pre> <p>If the driver is still loaded in memory, the result should contain a line, similar to the following:</p> <pre>blkwatch_7616 70924 0</pre> 2. (Conditional) If the driver is still loaded, remove it from memory: <pre>rmmod blkwatch_7616</pre> 3. Remove the driver from the boot sequence: <pre>blkconfig -u</pre> 4. Remove the driver files by deleting the following directory with its contents: <pre>/lib/modules/[Kernel_Version]/Platespin</pre> 5. Delete the following file: <pre>/etc/blkwatch.conf</pre> |

| Component | Removal Instructions |
|---------------|--|
| LVM snapshots | <p>LVP snapshots used by ongoing replications are named according to a <i>volume_name-PS-snapshot</i> 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 Forge 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># lvdisplay -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 snapshot_name</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> |

8.8 Shrinking the PlateSpin Forge Databases

When the PlateSpin Forge 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, Forge 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 Forge

PlateSpin Forge 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 97](#)
- ♦ [Section A.2, “Pre-compiled “blkwatch” driver \(Linux\),” on page 98](#)

A.1 Analyzing Your Linux Workload

Prior to determining whether PlateSpin Forge 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 97](#)
- ♦ [Section A.1.2, “Determining the Architecture,” on page 97](#)

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 Forge 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 Forge 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-bigsmpt for 32-bit (x86) architecture, the item is listed in a format like this:

```
SLES9-SP1-2.6.5-7.139-bigsmpt-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-bigsmg-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-bigsmg-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-bigsmg-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 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 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.

Glossary

appliance Host. See [container](#).

container. The VM host that contains the failover workload (a protected workload's bootable virtual replica).

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 Forge is no longer required.

failover. Taking over the business function of a failed workload by a failover workload within a PlateSpin Forge 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.

Management VM. The management virtual machine containing the PlateSpin Forge software.

prepare for failover. A PlateSpin Forge 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 Forge 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 Forge 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 Forge 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 Forge.

See also [source](#).

test failover. A PlateSpin Forge 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.