
PlateSpin Forge® 11.2

User Guide

March 2018

Legal Notice

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

Copyright © 2016 NetIQ Corporation. All rights reserved.

License Grant

Licenses purchased for PlateSpin Forge 11 or later versions cannot be used for PlateSpin Forge 3.3 or prior versions.

Third-Party Software

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

Contents

About this Book and the Library	7
1 Planning Your PlateSpin Environment	9
1.1 Supported Configurations	9
1.1.1 Supported Windows Workloads	10
1.1.2 Supported Linux Workloads	12
1.1.3 Supported VM Containers	13
1.1.4 Supported System Firmware	13
1.1.5 Supported Storage	13
1.1.6 Supported Browsers for the PlateSpin Forge Web Interface	14
1.2 Security and Privacy	15
1.2.1 Security of Workload Data in Transmission	15
1.2.2 Security of Client/Server Communications	15
1.2.3 Security of Credentials	15
1.2.4 User Authorization and Authentication	15
1.2.5 Network Port Settings	16
1.2.6 Additional Security Enhancements	17
1.3 Performance	17
1.3.1 About Product Performance Characteristics	18
1.3.2 Data Compression	18
1.3.3 Bandwidth Throttling	18
1.3.4 RPO, RTO, and TTO Specifications	19
1.3.5 Scalability	19
2 PlateSpin Forge Application Configuration	21
2.1 Launching the PlateSpin Forge Web Interface	21
2.2 Activating Your Product License	22
2.2.1 Online License Activation	22
2.2.2 Offline License Activation	23
2.2.3 Managing Licenses	23
2.3 Configuring User Authorization and Authentication	24
2.3.1 About PlateSpin Forge Role-Based Access	24
2.3.2 Managing PlateSpin Forge Access and Permissions	26
2.3.3 Managing PlateSpin Forge Security Groups and Workload Permissions	27
2.4 Configuring Access and Communication Settings across your Protection Network	28
2.4.1 Open Port Requirements for the PlateSpin Server HostForge VM Web Interface	28
2.4.2 Access and Communication Requirements for Workloads	28
2.4.3 Access and Communication Requirements for Containers	30
2.4.4 Protection Across Public and Private Networks Through NAT	30
2.4.5 Overriding the Default bash Shell for Executing Commands on Linux Workloads	31
2.4.6 Requirements for VMware DRS Clusters as Containers	31
2.5 Configuring Automatic Email Notifications of Events and Reports	31
2.5.1 SMTP Configuration	31
2.5.2 Setting Up Automatic Event Notifications by Email	32
2.5.3 Setting Up Automatic Replication Reports by Email	34
2.6 Configuring Language Settings for International Versions of PlateSpin Forge	35
2.7 Using Tags to Help Sort Workloads	36
2.8 Configuring PlateSpin Server Behavior through XML Configuration Parameters	37
2.9 Optimizing Data Transfer over WAN Connections	37

2.9.1	Tuning Parameters	38
2.9.2	Tuning FileTransferSendReceiveBufferSize	39
2.10	Configuring Support for VMware vCenter Site Recovery Manager	41
2.10.1	Setting Up Workload Files on the Same Datastore	41
2.10.2	Setting Up VMware Tools for Failover Targets	41
2.10.3	Expediting the Configuration Process	42
2.11	Configuring the Refresh Rates for the Web Interface	43
3	Configuring the Forge Appliance	45
3.1	Setting Up Appliance Networking	45
3.1.1	About Appliance Host Networking	46
3.1.2	Moving or Creating a vSwitch	46
3.1.3	Assigning VLAN Tags to Forge Port Groups	46
3.2	Using External Storage Solutions with PlateSpin Forge	46
3.2.1	Using Forge with SAN Storage	47
3.2.2	Adding a SAN LUN to Forge	47
3.3	Managing the Forge VM with vSphere Client	48
3.3.1	Downloading and Installing the vSphere Client Program	48
3.3.2	Launching the vSphere Client	48
3.3.3	Accessing the Forge VM	48
3.3.4	Starting and Shutting Down the Forge VM	49
3.3.5	Managing Snapshots of the Forge VM on the Appliance Host	49
3.3.6	Manually Importing VMs into the Appliance Host's Datastore	50
3.3.7	Guidelines for Applying Security Updates to the Forge VM	50
3.4	Physically Relocating the Appliance	50
3.4.1	Prerequisites for Relocating Forge	51
3.4.2	Scenario 1: Relocating Forge When the New IP Address Is Known	51
3.4.3	Scenario 2: Relocating Forge When the New IP Address Is Unknown	52
3.5	Recovering the Forge Management VM to Factory Defaults	54
3.6	Resetting the Forge Appliance to Factory Defaults	55
4	Up and Running	59
4.1	Accessing the PlateSpin Forge Web Interface	59
4.2	Elements of the PlateSpin Forge Web Interface	60
4.2.1	Navigation Bar	61
4.2.2	Visual Summary Panel	61
4.2.3	Tasks and Events Panel	62
4.3	Workloads and Workload Commands	62
4.3.1	Workload Protection and Recovery Commands	63
4.4	Managing Multiple Instances of PlateSpin Protect and PlateSpin Forge	64
4.4.1	Using the PlateSpin Forge Management Console	64
4.4.2	About PlateSpin Forge Management Console Cards	65
4.4.3	Adding Instances of PlateSpin Protect and PlateSpin Forge to the Management Console	66
4.4.4	Managing Cards on the Management Console	66
4.5	Generating Workload and Workload Protection Reports	67
5	Workload Protection and Recovery	69
5.1	Basic Workflow for Workload Protection and Recovery	69
5.2	Adding Containers (Protection Targets)	70
5.3	Adding Workloads	72
5.4	Configuring Protection Details and Preparing the Replication	73
5.4.1	Workload Protection Details	73
5.5	Starting the Workload Protection	76

5.6	Aborting Commands	77
5.7	Failover	77
5.7.1	Detecting Offline Workloads	77
5.7.2	Performing a Failover	78
5.7.3	Using the Test Failover Feature	78
5.8	Failback	79
5.8.1	Automated Failback to a VM Platform	79
5.8.2	Semi-Automated Failback to a Physical Machine	82
5.8.3	Semi-Automated Failback to a Virtual Machine	82
5.9	Reprotecting a Workload	83
6	Essentials of Workload Protection	85
6.1	Workload License Consumption	85
6.2	Guidelines for Workload and Container Credentials	86
6.3	Data Transfer	86
6.3.1	Transfer Methods	87
6.3.2	Data Encryption	88
6.3.3	Modifying the Location of the Volume Snapshots Directory for Windows Workloads	88
6.3.4	Excluding or Including Files in Block Transfers for Incremental Replications	88
6.4	Protection Tiers	89
6.5	Recovery Points	90
6.6	Initial Replication Method (Full and Incremental)	91
6.7	Service and Daemon Control	92
6.8	Using Freeze and Thaw Scripts for Every Replication (Linux)	92
6.9	Volumes Storage	93
6.10	Networking	96
6.11	Failback to Physical Machines	96
6.11.1	Downloading the PlateSpin Boot ISO Image	96
6.11.2	Injecting Additional Device Drivers into the Boot ISO Image	96
6.11.3	Registering Physical Machines as Failback Targets with PlateSpin Forge	98
6.12	Protecting Windows Clusters	99
6.12.1	Planning Your Cluster Workload Protection	99
6.12.2	Enabling or Disabling Windows Cluster Discovery	101
6.12.3	Adding Resource Name Search Values	102
6.12.4	Quorum Arbitration Timeout	103
6.12.5	Setting Local Volume Serial Numbers	103
6.12.6	PlateSpin Failover	103
6.12.7	PlateSpin Failback	103
7	Auxiliary Tools for Working with Physical Machines	105
7.1	Managing Device Drivers	105
7.1.1	Packaging Device Drivers for Windows Systems	105
7.1.2	Packaging Device Drivers for Linux Systems	106
7.1.3	Uploading Drivers to the PlateSpin Device Driver Database	106
7.1.4	Using the Plug and Play (PnP) ID Translator Feature	108
8	ProtectAgent Utility	115
9	Troubleshooting	119
9.1	Troubleshooting Workload Inventory (Windows)	119
9.1.1	Performing Connectivity Tests	120
9.1.2	Disabling Antivirus Software	122
9.1.3	Enabling File/Share Permissions and Access	122

9.2	Troubleshooting Workload Inventory (Linux)	123
9.3	Troubleshooting Problems during the Prepare Replication Command (Windows)	123
9.3.1	Group Policy and User Rights	123
9.4	Troubleshooting Workload Replication	124
9.5	Troubleshooting the Configuration Service	125
9.5.1	Understanding What Is Causing the Problem	126
9.5.2	What Can Be Done to Resolve the Problem	126
9.5.3	Additional Troubleshooting Tips	129
9.6	Troubleshooting Traffic-forwarding Workloads	130
9.7	Troubleshooting Online Help	130
9.8	Generating and Viewing Diagnostic Reports	131
9.9	Removing Workloads	131
9.10	Post-Protection Workload Cleanup	131
9.10.1	Cleaning Up Windows Workloads	131
9.10.2	Cleaning Up Linux Workloads	132
9.11	Shrinking the PlateSpin Forge Databases	133
9.12	Active Directory Domain Services Are Not Available After a Failback (Windows)	134
A	Linux Distributions Supported by Forge	135
A.1	Analyzing Your Linux Workload	135
A.1.1	Determining the Release String	135
A.1.2	Determining the Architecture	135
A.2	PlateSpin Forge Pre-compiled “blkwatch” Driver (Linux)	136
A.2.1	List Item Syntax	136
A.2.2	List of Distributions	136
A.2.3	Other Linux Distributions That Use blkwatch Drivers	145
B	Synchronizing Serial Numbers on Cluster Node Local Storage	147
C	Rebranding the PlateSpin Forge Web Interface	149
C.1	Rebranding the Interface By Using Configuration Parameters	149
C.2	Rebranding the Product Name in the Windows Registry	152
D	Using Workload Protection Features through the PlateSpin Protect Server API	155
D.1	API Overview	155
D.2	PlateSpin Protect Server API Documentation	155
D.3	Samples and Other References	156
E	Documentation Updates	159
E.1	March 2018	159
E.2	January 2018	159
E.3	October 2017	160
E.4	June 2017	160
E.5	February 2017	160
E.6	January 2017	160
E.7	November 2016	160
E.8	September 2016	161
E.9	April 2016	161
E.10	February 2016	161

About this Book and the Library

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

Intended Audience

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

Additional Documentation

For the most recent version of this guide and other PlateSpin Forge documentation resources, visit the [PlateSpin Forge Documentation \(https://www.netiq.com/documentation/platespin-forge/\)](https://www.netiq.com/documentation/platespin-forge/) website.

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

Contact Information

We want to hear your comments and suggestions about this book and the other documentation included with this product. You can use the [comment on this topic](#) link at the bottom of any page of the online documentation, or send an email to Documentation-Feedback@microfocus.com.

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

1

Planning Your PlateSpin Environment

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 infrastructures 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.

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

- ♦ [Section 1.1, “Supported Configurations,” on page 9](#)
- ♦ [Section 1.2, “Security and Privacy,” on page 15](#)
- ♦ [Section 1.3, “Performance,” on page 17](#)

1.1 Supported Configurations

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

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

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

- ♦ [Section 1.1.1, “Supported Windows Workloads,” on page 10](#)
- ♦ [Section 1.1.2, “Supported Linux Workloads,” on page 12](#)
- ♦ [Section 1.1.3, “Supported VM Containers,” on page 13](#)

- ♦ [Section 1.1.4, “Supported System Firmware,” on page 13](#)
- ♦ [Section 1.1.5, “Supported Storage,” on page 13](#)
- ♦ [Section 1.1.6, “Supported Browsers for the PlateSpin Forge Web Interface,” on page 14](#)

1.1.1 Supported Windows Workloads

PlateSpin Forge supports workloads for most Microsoft Windows versions. For a list of supported Windows versions, see [Table 1-1](#).

Both file-level and block-level replications are supported, with certain restrictions. See [Section 6.3, “Data Transfer,” on page 86](#).

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

Table 1-1 *Supported Windows Workloads*

Operating System	Notes
Servers	
Windows Server 2012 R2 Windows Server 2012	Includes domain controllers (DC) and Small Business Server (SBS) editions. For information about conversion of Active Directory domain controllers, see Knowledgebase Article 7920501 (https://www.netiq.com/support/kb/doc.php?id=7920501) .
Windows Server 2008 R2 (64-bit) Windows Server 2008 (64-bit) Windows Server 2008 latest SP (32-bit)	Includes domain controllers (DC) and Small Business Server (SBS) editions. For information about conversion of Active Directory domain controllers, see Knowledgebase Article 7920501 (https://www.netiq.com/support/kb/doc.php?id=7920501) .
Windows Server 2003 R2 (64-bit) Windows Server 2003 R2 (32-bit) Windows Server 2003 latest SP (64-bit) Windows Server 2003 latest SP (32-bit)	Windows 2003 requires SP1 or higher for Block-based replication.
Clusters	
Windows Server 2012 R2 server-based Microsoft Failover Cluster	The supported clustering technologies are: <i>Node and Disk Majority Quorum</i> and <i>No Majority: Disk Only Quorum</i> models. Block-based transfer only.
Windows Server 2008 R2 server-based Microsoft Failover Cluster	The supported clustering technologies are: <i>Node and Disk Majority Quorum</i> and <i>No Majority: Disk Only Quorum</i> models. Block-based transfer only.
Windows Server 2003 R2 server-based Windows Cluster Server	The supported clustering technology is: <i>Single-Quorum Device Cluster</i> model. Block-based transfer only.

Operating System	Notes
Hyper-V Hosts	
Windows Server 2012 R2 with Hyper-V Role	Protect a Windows server that is functioning as a Hyper-V host and its volumes. Protect the individual VMs separately.
Windows Server 2012 with Hyper-V Role	

Supported Windows File Systems

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

Supported Windows Clusters

For detailed information about how to protect workloads in a supported cluster, see [“Protecting Windows Clusters” on page 99](#). If you have local storage on cluster nodes, see also [“Synchronizing Serial Numbers on Cluster Node Local Storage” on page 147](#).

Supported International Versions

PlateSpin Forge supports French, German, Japanese, Chinese Traditional, and Chinese Simplified versions of Microsoft Windows. See [“Configuring Language Settings for International Versions of PlateSpin Forge” on page 35](#).

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

Workload Firmware (UEFI and BIOS) Support

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

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

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

- ♦ When you transfer a UEFI-based workload to a VMware vSphere 4.x container (which does not support UEFI), 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 because Windows 2003 does not support the UEFI boot mode.
- ♦ When you protect 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.

Workload Complex Disk Partitioning Support

In addition to the MBR partitioning scheme, PlateSpin Forge supports the GPT partitioning of disks for Windows workloads. Full replication is supported for 57 or fewer partitions or volumes on a single disk.

Windows Update

Ensure that you update Windows (Windows Update) on your source system before you run the first full replication. If the Windows machine is a Domain Controller, ensure that you also disable anti-virus software on the system during the replication.

1.1.2 Supported Linux Workloads

PlateSpin Forge supports a number of Linux distributions. For a list of supported Linux operating systems, see [Table 1-2](#).

Replication of protected Linux workloads occurs only at the block level. See [“Requirement for a blkwatch Driver” on page 13](#).

Table 1-2 Supported Linux Workloads

Operating System	Notes
Servers	
Red Hat Enterprise Linux (RHEL) 7 Red Hat Enterprise Linux 6 Red Hat Enterprise Linux 5 Red Hat Enterprise Linux 4	See “Linux Distributions Supported by Forge” on page 135 for a list of supported Linux kernel versions and architectures for RHEL distributions.
SUSE Linux Enterprise Server (SLES) 11 SUSE Linux Enterprise Server 10 SUSE Linux Enterprise Server 9	See “Linux Distributions Supported by Forge” on page 135 for a list of supported Linux kernel versions and architectures for SLES distributions. NOTE: Kernel version 3.0.13 of SLES 11 SP3 is not supported. Upgrade to kernel version 3.0.27 or later before you inventory the workload.
Novell Open Enterprise Server (OES) 11 Novell Open Enterprise Server 2	PlateSpin Forge supports workloads for an OES 2 or OES 11 version if it is based on a supported SLES distribution, except as noted. See “Linux Distributions Supported by Forge” on page 135 for a list of supported Linux kernel versions and architectures for SLES distributions. NOTE: The default kernel version 3.0.13 of OES 11 SP2 is not supported. Upgrade to kernel version 3.0.27 or later before you inventory the workload.
Oracle Enterprise Linux (OEL)	PlateSpin Forge supports workloads for an OEL version if it is based on a supported RHEL distribution, except as noted. See “Linux Distributions Supported by Forge” on page 135 for a list of supported Linux kernel versions and architectures for RHEL distributions. NOTE: Workloads using the Unbreakable Enterprise Kernel are not supported.

Operating System	Notes
CentOS 7	PlateSpin Forge supports workloads for a CentOS version if it is based on a supported RHEL distribution. See “Linux Distributions Supported by Forge” on page 135 for a list of supported Linux kernel versions and architectures for RHEL distributions.
CentOS 6	
CentOS 5	
CentOS 4	

Supported Linux File Systems

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

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

Workload Firmware (UEFI and BIOS) Support

PlateSpin Forge provides support for the UEFI and BIOS firmware interfaces.

Workload Complex Disk Partitioning Support

In addition to the MBR partitioning scheme, PlateSpin Forge supports the GPT partitioning of disks for Linux workloads. Full replication is supported for 57 or fewer partitions or volumes on a single disk.

Requirement for a blkwatch Driver

The block-based transfer of data for a Linux workload in PlateSpin Forge requires a `blkwatch` driver that is compiled for the particular Linux distribution being protected. PlateSpin Forge software includes pre-compiled versions of the `blkwatch` driver for many non-debug Linux distributions (32-bit and 64-bit). You can also create a custom driver. For more information, see [“Linux Distributions Supported by Forge” on page 135](#).

1.1.3 Supported VM Containers

PlateSpin Forge ships with VMware ESXi 5.5 GA2 Update 2 as the protection VM container.

1.1.4 Supported System Firmware

PlateSpin Forge supports the UEFI and BIOS firmware interfaces.

On Windows systems, PlateSpin Forge mirrors the Microsoft support of UEFI. For more information see [Workload Firmware \(UEFI and BIOS\) Support](#) in [“Supported Windows Workloads” on page 10](#).

1.1.5 Supported Storage

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

PlateSpin Forge supports several types of storage, including basic disks, 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 11 workloads) Novell Linux Volume Management (NLVM) layout of source workloads are preserved and re-created in the appliance host. NSS pools are copied from the source to the recovery VM.
- ♦ (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.

1.1.6 Supported Browsers for the PlateSpin Forge Web Interface

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

The supported browsers are:

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

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

To enable JavaScript:

- ♦ **Chrome:**
 1. From the Chrome menu, select **Settings**, then scroll to and click **Show advanced settings**.
 2. Under **Privacy**, click **Content Settings**.
 3. Scroll to **JavaScript**, then select **Allow all sites to run JavaScript**.
 4. Click **Done**.
- ♦ **Firefox:**
 1. In the Location bar, type `about:config` and press Enter.
 2. Click **I'll be careful, I promise!**
 3. In the **Search** bar, type `javascript.enabled`, then press Enter.
 4. In the search results, view the value for the `javascript.enabled` parameter. If its value is `false`, right-click `javascript.enabled` and select **Toggle** to set its value to `true`.
- ♦ **Internet Explorer:**
 1. From the Tools menu, select **Internet Options**.
 2. Select **Security**, then click **Custom level**.
 3. Scroll to **Scripting > Active scripting**, then select **Enable**.
 4. Click **Yes** at the warning dialog box, then click **OK**.
 5. Click **Apply > OK**.

To use the PlateSpin Forge Web Interface and integrated help in one of the supported languages, see [“Configuring Language Settings for International Versions of PlateSpin Forge” on page 35](#).

1.2 Security and Privacy

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

- ♦ [Section 1.2.1, “Security of Workload Data in Transmission,” on page 15](#)
- ♦ [Section 1.2.2, “Security of Client/Server Communications,” on page 15](#)
- ♦ [Section 1.2.3, “Security of Credentials,” on page 15](#)
- ♦ [Section 1.2.4, “User Authorization and Authentication,” on page 15](#)
- ♦ [Section 1.2.5, “Network Port Settings,” on page 16](#)
- ♦ [Section 1.2.6, “Additional Security Enhancements,” on page 17](#)

1.2.1 Security of Workload Data in Transmission

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

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

You can enable or disable encryption individually for each workload by selecting the **Encrypt Data Transfer** option. See [“Workload Protection Details” on page 73](#).

1.2.2 Security of Client/Server Communications

Because the PlateSpin Server enables SSL on the Forge VM, secure data transmission between your web browser and the PlateSpin Server is already configured to HTTPS (Hypertext Transfer Protocol Secure).

1.2.3 Security of Credentials

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

Credentials that you use to access various systems (such as workloads and failback targets) are stored in the PlateSpin 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.

1.2.4 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 [“Configuring User Authorization and Authentication” on page 24](#).

1.2.5 Network Port Settings

[Table 1-3](#) lists the default ports used by PlateSpin Forge. If you configure custom ports, you must open those ports instead. For communications between the PlateSpin Forge Server and the source and target machines it manages, ensure that you also open the appropriate ports on any firewalls between them. Traffic for communications is bidirectional (incoming and outgoing). For more information about network access configuration for your PlateSpin Server environment, see [“Configuring Access and Communication Settings across your Protection Network” on page 28](#).

Table 1-3 Default Ports Used by PlateSpin Forge

Port Number	Protocol	Function	Details
80	TCP	HTTP	(Not secure) Used for HTTP communications between the Forge VM and the source and target machines it manages. Open this port on your Forge VM, the source and target workloads, and the VMware ESXi hosts.
443	TCP	HTTPS	(Secure) Used for HTTPS communications, if SSL is enabled, between the Forge VM and the source and target machines. Open this port on your Forge VM, the source and target workloads, the VMware ESXi hosts, and the vCenter host server.
3725	TCP	Data transfer	Used for data transfer between the source and target machines, including file-based transfer and block-based transfer. Open this port on the source and target machines for all workloads. Any firewall between a source and its target must also allow TCP port 3725. See “Supported Configurations” on page 9 .
135 445	TCP	RPC/DCOM	Used for RPC/DCOM communications on Windows machines during the discovery process, and when taking control and rebooting the source machine. Open these ports for communications between the source and target machines for all Windows workloads. See “Supported Windows Workloads” on page 10 .
137 138 139	TCP	NetBIOS	Used for NetBIOS communications (name service, datagram service, and session service). Open these ports for communications between the source and target machines for all Windows workloads. See “Supported Windows Workloads” on page 10 .
137 138	UDP	SMB	Used for SMB communications for the file transfer of the Take Control folder and files from the PlateSpin Server to the source machine.
139 445	TCP	SMB	Open these ports on your Forge VM and the source workloads.

Port Number	Protocol	Function	Details
22	TCP		Used for SSH and SCP communications on Linux machines during the discovery process. Open this port on the source and target machines for all Linux workloads. See “Supported Linux Workloads” on page 12 .
25	TCP	SMTP	Used for SMTP traffic if email notification is enabled.
25	UDP	SMTP	Open this port on the Forge VM and the mail relay host.
1433	TCP	SQL	Used for Microsoft SQL Server communications for authentication and data exchange to a remote SQL Server. Open the SQL ports on your Forge VM and the remote SQL Server host, as well as on any firewalls between them. For more information the SQL Server port requirements, see Configure the Firewall to Allow Server Access in the Microsoft Developers Network.
1434	TCP	SQL	Used for the Microsoft SQL Server dedicated admin connection.
1434	UDP	SQL	Used for the Microsoft SQL Server named instances. This port might be required when you use named instances on a remote SQL Server.
49152 to 65535	TCP	SQL	Used for the Microsoft SQL Server or RPC for LSA, SAM, and Netlogon. If you have configured Microsoft SQL Server to use a specific TCP port, you must open that port on the firewall.

1.2.6 Additional Security Enhancements

PlateSpin Forge provides information in [Knowledgebase Article 7015818 \(https://www.netiq.com/support/kb/doc.php?id=7015818\)](https://www.netiq.com/support/kb/doc.php?id=7015818) about how to remove the vulnerability to potential POODLE (Padding Oracle On Downgraded Legacy Encryption) attacks from your PlateSpin servers.

1.3 Performance

- ♦ [Section 1.3.1, “About Product Performance Characteristics,” on page 18](#)
- ♦ [Section 1.3.2, “Data Compression,” on page 18](#)
- ♦ [Section 1.3.3, “Bandwidth Throttling,” on page 18](#)
- ♦ [Section 1.3.4, “RPO, RTO, and TTO Specifications,” on page 19](#)
- ♦ [Section 1.3.5, “Scalability,” on page 19](#)

1.3.1 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
- ♦ Hardware and software profiles of your target containers
- ♦ 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.

1.3.2 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 [“Protection Tiers” on page 89](#).

1.3.3 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 [“Protection Tiers” on page 89](#).

1.3.4 RPO, RTO, and TTO Specifications

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

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

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

The RTO is affected by the time it takes to configure and execute the failover operation (10 to 45 minutes). See [“Failover” on page 77](#).

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

Use the **Test Failover** feature to run through different scenarios and generate benchmark data. See [“Using the Test Failover Feature” on page 78](#).

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

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

- ♦ **Recovery Point Actual (RPA):** The RPA is the actual data loss measured in time and defined by the actual measured interval between incremental replications of a protected workload that occurs during a failover test. RPA is also known as *Actual Recovery Point Objective (Actual RPO)*.
- ♦ **Recovery Time Actual (RTA):** The RTA is a measure of a workload's actual downtime defined by the time a failover operation takes to complete. RTA is also known as *Actual Recovery Time Objective (Actual RTO)*.
- ♦ **Test Time Actual (TTA):** The TTA is a measure of the actual time in which a disaster recovery plan can be tested. It is similar to Actual RTO, but includes the time needed for a user to test the failover workload. TTA is also known as *Actual Test Time Objective (Actual TTO)*.

You should determine average failover times for workloads in your environment by doing test failovers at various times, then use them as benchmark data in your overall data recovery plans. See [“Generating Workload and Workload Protection Reports” on page 67](#).

1.3.5 Scalability

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

- ♦ **Workloads per Server:** The number of workloads per PlateSpin Server might vary between 10 and 50, depending on several factors, including your RPO requirements and the hardware characteristics of the server host.
- ♦ **Protections per Container:** The maximum number of protections per container is related to (but is not the same as) the VMware specifications pertaining to the maximum number of VMs supported per ESXi host. Additional factors include recovery statistics (including concurrent replications and failovers) and hardware vendor specifications.

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

2 PlateSpin Forge Application Configuration

This section describes configuration requirements and setup for PlateSpin Forge.

- [Section 2.1, “Launching the PlateSpin Forge Web Interface,” on page 21](#)
- [Section 2.2, “Activating Your Product License,” on page 22](#)
- [Section 2.3, “Configuring User Authorization and Authentication,” on page 24](#)
- [Section 2.4, “Configuring Access and Communication Settings across your Protection Network,” on page 28](#)
- [Section 2.5, “Configuring Automatic Email Notifications of Events and Reports,” on page 31](#)
- [Section 2.6, “Configuring Language Settings for International Versions of PlateSpin Forge,” on page 35](#)
- [Section 2.7, “Using Tags to Help Sort Workloads,” on page 36](#)
- [Section 2.8, “Configuring PlateSpin Server Behavior through XML Configuration Parameters,” on page 37](#)
- [Section 2.9, “Optimizing Data Transfer over WAN Connections,” on page 37](#)
- [Section 2.10, “Configuring Support for VMware vCenter Site Recovery Manager,” on page 41](#)
- [Section 2.11, “Configuring the Refresh Rates for the Web Interface,” on page 43](#)

2.1 Launching the PlateSpin Forge Web Interface

To use the PlateSpin Forge Web Interface and integrated help in one of the supported languages, see [“Configuring Language Settings for International Versions of PlateSpin Forge” on page 35](#).

To launching the PlateSpin Forge Web Interface:

- 1 Open a [supported web browser](#) and go to:

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

Replace `<hostname | IP_address>` with the DNS hostname or the IP address of your Forge VM.

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

The first time you log into PlateSpin Forge, the browser redirects to the License Activation page.

- 2 Log in using the local Administrator user credentials for the Forge VM.

The default credentials for the Forge Management VM are user name `Administrator` and password `Password1`. To modify the Administrator user password, you can log in remotely to the VM's Windows desktop, and use Windows management tools to set a new password.

For information about setting up additional users for PlateSpin, see [Section 2.3, “Configuring User Authorization and Authentication,” on page 24](#).

2.2 Activating Your Product License

Your PlateSpin Forge product license entitles you to a specific number or unlimited number of workloads for protection through workload licensing. For more information, see [Section 6.1, “Workload License Consumption,” on page 85](#).

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

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

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

- ♦ [Section 2.2.1, “Online License Activation,” on page 22](#)
- ♦ [Section 2.2.2, “Offline License Activation,” on page 23](#)
- ♦ [Section 2.2.3, “Managing Licenses,” on page 23](#)

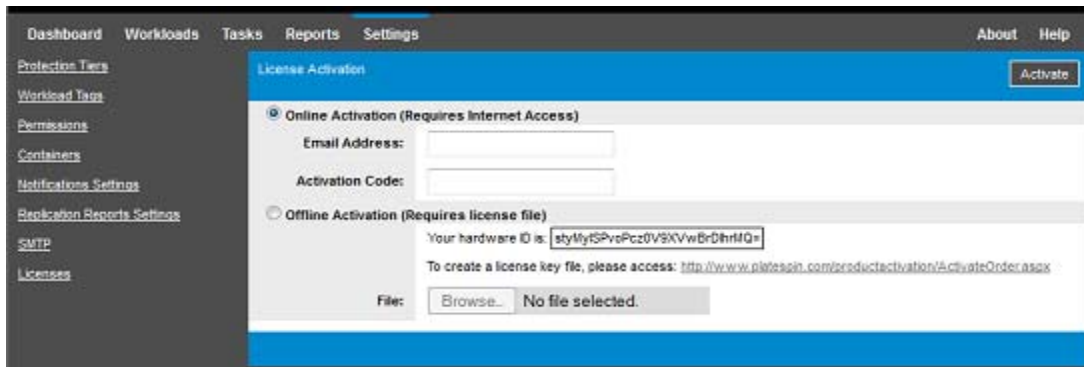
2.2.1 Online License Activation

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 set up online license activation:

- 1 In the PlateSpin Forge Web Interface, click **Add PlateSpin Forge License > Add License**.



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

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

2.2.2 Offline License Activation

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

- 1 In the PlateSpin Forge Web Interface, click **Add PlateSpin Forge License > Add License**.
- 2 Select **Offline Activation** and copy the hardware ID shown.
- 3 Use a web browser on a computer that has Internet access to navigate to the [PlateSpin Product Activation website \(http://www.platespin.com/productactivation/ActivateOrder.aspx\)](http://www.platespin.com/productactivation/ActivateOrder.aspx). Log in with your Customer Center user name and password.
- 4 Use the hardware ID to create a license key file. This process requires the following information:
 - ♦ the activation code that you received
 - ♦ the email address that you provided when you placed your order
 - ♦ the hardware ID that you copied in [Step 2](#)
- 5 Save the generated license key file, transfer it to the product host that does not have Internet connectivity, and use it to activate the product.
- 6 In the PlateSpin Forge Web Interface on the License Activation page, type the path to the file or browse to its location, then click **Activate**.

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

2.2.3 Managing Licenses

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

- ♦ [“Viewing License Information” on page 23](#)
- ♦ [“Adding a License” on page 24](#)
- ♦ [“Deleting a License” on page 24](#)
- ♦ [“Generating a Licensing Report for Technical Support” on page 24](#)

Viewing License Information

The product Dashboard provides a License Summary that displays the number installed licenses and the current number of consumed licenses. Licenses are consumed as you add a workload for protection. See [Section 6.1, “Workload License Consumption,” on page 85](#).

You can view information about product licenses on the Licenses page.

- 1 In the PlateSpin Forge Web Interface, select **Settings > Licenses**.
- 2 View the license information:
 - ♦ **Activation Code**
 - ♦ **Expiry Date**
 - ♦ **Workloads**
- 3 View **Workloads remaining** for the number of unused licenses.

Adding a License

You use the same process for adding a new license as for activating the first license:

- ♦ [Section 2.2.1, “Online License Activation,” on page 22](#)
- ♦ [Section 2.2.2, “Offline License Activation,” on page 23](#)

Deleting a License

You can delete an expired license on the Licenses page.

- 1 In the PlateSpin Forge Web Interface, select **Settings > Licenses**.
- 2 View the license information.
- 3 Click **Delete** next to the expired license, then confirm the deletion.

Generating a Licensing Report for Technical Support

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

- 1 In the PlateSpin Forge Web Interface, select **Settings > Licenses**.
- 2 Below the list of licenses, click **View Licensing Report**.
The `LicenseReport.txt` file opens in a new browser tab or window, depending on your browser settings.
- 3 Save the `LicenseReport.txt` file as `LicenseReport.psl` on your local computer.

2.3 Configuring User Authorization and Authentication

The following information is included in this section:

- ♦ [Section 2.3.1, “About PlateSpin Forge Role-Based Access,” on page 24](#)
- ♦ [Section 2.3.2, “Managing PlateSpin Forge Access and Permissions,” on page 26](#)
- ♦ [Section 2.3.3, “Managing PlateSpin Forge Security Groups and Workload Permissions,” on page 27](#)

2.3.1 About PlateSpin Forge Role-Based Access

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.

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

- 1 Add users to the required user groups detailed in [Table 2-1](#). See your Windows documentation.
- 2 Create application-level security groups that associate these users with specified workloads. See [“Managing PlateSpin Forge Security Groups and Workload Permissions” on page 27](#).

2.3.2 Managing PlateSpin Forge Access and Permissions

The following sections provide more information:

- ♦ [“Modifying the Administrator User Password for the Forge Management VM” on page 26](#)
- ♦ [“Adding PlateSpin Forge Users” on page 26](#)
- ♦ [“Assigning a Workload Protection Role to a PlateSpin Forge User” on page 26](#)

Modifying the Administrator User Password for the Forge Management VM

The default credentials for the Forge Management VM are user name `Administrator` and password `Password1`.

To modify the Administrator user password:

- 1 Launch a Remote Desktop connection to the Forge Management VM, using the IP address that you configured for the VM.
- 2 Log in as the Administrator user, using the current credentials.
- 3 Use Windows management tools to set a new password for the Administrator user.
- 4 Log out and close the Remote Desktop connection.

Adding PlateSpin Forge Users

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

If you want to grant specific role permissions to an existing user on the Forge VM, see [“Assigning a Workload Protection Role to a PlateSpin Forge User” on page 26](#).

You can now assign a workload protection role to the newly created user. See [“Assigning a Workload Protection Role to a PlateSpin Forge User” on page 26](#).

Assigning a Workload Protection Role to a PlateSpin Forge User

Before assigning a role to a user, determine the collection of permissions that best suits that user. See [Table 2-1, “Workload Protection Roles and Permission Details,” on page 25](#).

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

To restart the PlateSpin Server:

- 1 Before you attempt to restart the PlateSpin Server, pause all of your contracts, or verify that no replications, failovers, or failbacks are in progress. Do not continue until all workloads are idle.
- 2 Go to the PlateSpin Server’s `bin\RestartPlateSpinServer` subdirectory.
- 3 Double-click the `RestartPlateSpinServer.exe` executable.
A command prompt window opens, requesting confirmation.
- 4 Confirm by typing `Y` and pressing `Enter`.

You can now add this user to a PlateSpin Forge security group and associate a specified collection of workloads. See [“Managing PlateSpin Forge Security Groups and Workload Permissions” on page 27](#).

2.3.3 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*.

- 1 Assign a PlateSpin Forge user a Workload Protection Role whose permissions best suit that role in your organization. See [“Assigning a Workload Protection Role to a PlateSpin Forge User” on page 26](#).

- 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 who 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**.

The dialog box titled "Choose users to grant access to this group:" contains a table with three columns: "Grant", "Name", and "Roles".

Grant	Name	Roles
<input checked="" type="checkbox"/>	ADLER\operator1	Workload Protection Operator

At the bottom of the dialog are "OK" and "Cancel" buttons.

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

The dialog box titled "Choose workloads to include in this group:" contains a table with three columns: "Include", "Workload Name", and "Security Group".

Include	Workload Name	Security Group
<input type="checkbox"/>	vsles11sp3x64.example.com	[Unassigned]
<input type="checkbox"/>	VVC1	[Unassigned]
<input type="checkbox"/>	AE-W2K3-1	[Unassigned]
<input checked="" type="checkbox"/>	AE-W2K3-3	[Unassigned]
<input checked="" type="checkbox"/>	AE-W2K3-4	[Unassigned]

At the bottom of the dialog are "OK" and "Cancel" buttons.

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.

2.4 Configuring Access and Communication Settings across your Protection Network

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

- [Section 2.4.1, “Open Port Requirements for the PlateSpin Server HostForge VM Web Interface,” on page 28](#)
- [Section 2.4.2, “Access and Communication Requirements for Workloads,” on page 28](#)
- [Section 2.4.3, “Access and Communication Requirements for Containers,” on page 30](#)
- [Section 2.4.4, “Protection Across Public and Private Networks Through NAT,” on page 30](#)
- [Section 2.4.5, “Overriding the Default bash Shell for Executing Commands on Linux Workloads,” on page 31](#)
- [Section 2.4.6, “Requirements for VMware DRS Clusters as Containers,” on page 31](#)

2.4.1 Open Port Requirements for the PlateSpin Server HostForge VM Web Interface

[Table 2-2](#) describes the ports that must be open for on the Forge VM to allow access to the PlateSpin Forge Web Interface.

Table 2-2 Open Port Requirements for the PlateSpin Server HostForge VM

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

2.4.2 Access and Communication Requirements for Workloads

[Table 2-3](#) describes the software, network, and firewall requirements for workloads that you intend to protect by using PlateSpin Forge.

Table 2-3 Access and Communication Requirements for Workloads

Workload Type	Prerequisites	Required Ports (Defaults)
All workloads	Ping (ICMP echo request and response) support	
All Windows workloads. See “Supported Windows Workloads” on page 10 .	<ul style="list-style-type: none">• Microsoft .NET Framework 3.5 Service Pack 1• Microsoft .NET Framework 4.0 <p>For discovery, source workloads must be running Microsoft .NET Framework 2 SP2 or later.</p>	

Workload Type	Prerequisites	Required Ports (Defaults)
All Windows workloads. See “Supported Windows Workloads” on page 10.	<ul style="list-style-type: none"> ♦ Built-in Administrator or domain administrator account credentials (membership only in the local Administrators group is insufficient). ♦ The Windows Firewall configured to allow File and Printer Sharing. Use one of these options: <ul style="list-style-type: none"> ♦ Option 1, using Windows Firewall: Use the basic Windows Firewall Control Panel item (<code>firewall.cpl</code>) and select File and printer Sharing in the list of exceptions. - OR - ♦ Option 2, using Firewall with Advanced Security: Use the Windows Firewall with Advanced Security utility (<code>wf.msc</code>) with the following Inbound Rules enabled and set to Allow: <ul style="list-style-type: none"> ♦ File and Printer Sharing (Echo Request - ICMPv4In) ♦ File and Printer Sharing (Echo Request - ICMPv6In) ♦ File and Printer Sharing (NB-Datagram-In) ♦ File and Printer Sharing (NB-Name-In) ♦ File and Printer Sharing (NB-Session-In) ♦ File and Printer Sharing (SMB-In) ♦ File and Printer Sharing (Spooler Service - RPC) ♦ File and Printer Sharing (Spooler Service - RPC-EPMAP) 	TCP 3725 NetBIOS (TCP 137 - 139) SMB (TCP 139, 445 and UDP 137, 138) RPC (TCP 135, 445)
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, <i>The Netsh Command Line Utility</i> (http://technet.microsoft.com/en-us/library/cc785383%28v=ws.10%29.aspx).</p>	TCP 3725, 135, 139, 445 UDP 137, 138, 139
All Linux workloads. See “Supported Linux Workloads” on page 12.	Secure Shell (SSH) server	TCP 22, 3725

2.4.3 Access and Communication Requirements for Containers

Table 2-4 describes the software, network, and firewall requirements for the supported workload containers.

Table 2-4 Access and Communication Requirements for Containers

System	Prerequisites	Required Ports (Defaults)
All containers	Ping (ICMP echo request and response) capability.	
All VMware containers. See “Supported VM Containers” on page 13.	<ul style="list-style-type: none">◆ VMware account with an Administrator role◆ VMware Web services API and file management API	HTTPS (TCP 443)
vCenter Server	The user with access must be assigned the appropriate roles and permissions. Refer to the pertinent release of VMware documentation for more information.	HTTPS (TCP 443)

2.4.4 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 30.
- ◆ **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 80).
- ◆ **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 Forge VM or one of its public (external) addresses recorded in the server's *PlateSpin Server Configuration* tool (see *PlateSpin Server* above).

Configuring the Application to Function through NAT

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

For information on the update procedure, see [“Configuring PlateSpin Server Behavior through XML Configuration Parameters”](#) on page 37.

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

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

If required, you can override the default shell by modifying the corresponding registry key on the PlateSpin Server.

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

2.4.6 Requirements for VMware DRS Clusters as Containers

To be a valid protection target, your VMware DRS cluster must be added to the set of containers (inventoried) as a VMware Cluster. You should not attempt to add a DRS Cluster as a set of individual ESX servers. See [“Adding Containers \(Protection Targets\)” on page 70](#).

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

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

2.5 Configuring 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.

- [Section 2.5.1, “SMTP Configuration,” on page 31](#)
- [Section 2.5.2, “Setting Up Automatic Event Notifications by Email,” on page 32](#)
- [Section 2.5.3, “Setting Up Automatic Replication Reports by Email,” on page 34](#)

2.5.1 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 the 'SMTP Settings' page in the PlateSpin Forge Web Interface. The left sidebar contains a navigation menu with options: Protection Tiers, Workload Tags, Permissions, Containers, Notifications Settings, Replication Reports Settings, SMTP (selected), and Licenses. The main content area is titled 'SMTP Settings' and includes a 'Save' button in the top right corner. The settings are organized into a table with the following fields:

Field	Value
SMTP Server Address:	
Port:	25
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 user name and password, then confirm the password.
- 4 Click **Save**.

2.5.2 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 31.
- 2 In your PlateSpin Forge Web Interface, click **Settings > Notification Settings**.
- 3 Select the **Enable Notifications** option.
- 4 Click **Edit Recipients**, type the required email addresses separated by commas, then click **OK**.

The screenshot shows the 'Notification Settings' page in the PlateSpin Forge Web Interface. The left sidebar contains a navigation menu with options: Protection Tiers, Workload Tags, Permissions, Containers, Notifications Settings (selected), Replication Reports Settings, SMTP, and Licenses. The main content area is titled 'Notification Settings' and includes a 'Save' button in the top right corner. The settings are organized into a table with the following fields:

Field	Value
Enable Notifications	<input checked="" type="checkbox"/>
Recipients:	
Address	
Remove	dradmin@example.com
Remove	john_smith@example.com
Remove	sysadmin@example.com
Remove	webadmin@example.com
Edit Recipients...	

At the bottom of the page, a status bar indicates: "Email notifications are ON; Number of recipients: 4".

- 5 Click **Save**.

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

The event types shown in [Table 2-5](#) can trigger email notifications if notification is configured. The events are always added to the System Application Event Log, according to the log entry types of Warning, Error, and Information.

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

Table 2-5 Events Types Organized by Log Entry Types

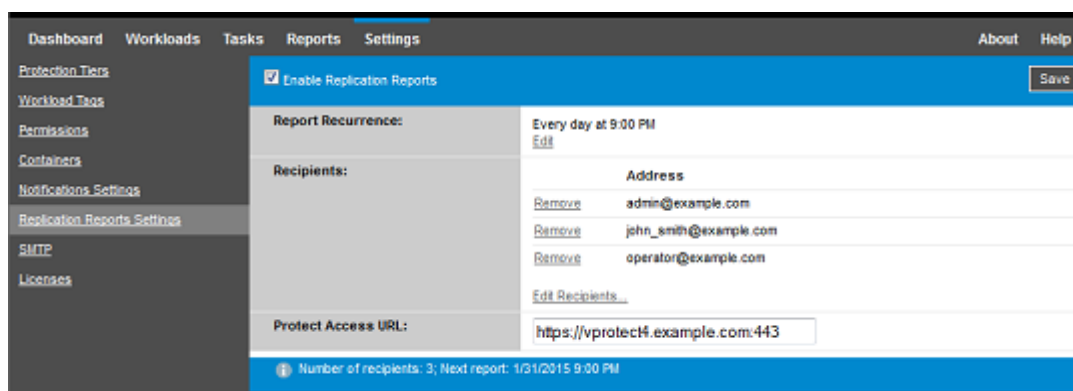
Event Types	Remarks
Log Entry Type: Warning	
FullReplicationMissed	Similar to the Incremental Replication Missed event.
IncrementalReplicationMissed	Generated when any of the following applies: <ul style="list-style-type: none"> ♦ A replication is manually paused while a scheduled incremental replication 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.
WorkloadOfflineDetected	Generated when the system detects that a previously online workload is now offline. Applies to workloads whose protection contract's state is not Paused .
Log Entry Type: Error	
FailoverFailed	
FullReplicationFailed	
IncrementalReplicationFailed	
PrepareFailoverFailed	
Log Entry Type: Information	
FailoverCompleted	
FullReplicationCompleted	
IncrementalReplicationCompleted	
PrepareFailoverCompleted	
TestFailoverCompleted	Generated upon manually marking a Test Failover operation a success or a failure.

Event Types	Remarks
WorkloadOnlineDetected	Generated when the system detects that a previously offline workload is now online.
	Applies to workloads whose protection contract's state is not Paused .

2.5.3 Setting Up Automatic Replication Reports by Email

To set up PlateSpin Forge to automatically send out replication reports by email, follow these steps:

- 1 Set up an SMTP server for PlateSpin Forge to use. See [SMTP Configuration \(page 31\)](#).
- 2 In your PlateSpin Forge Web Interface, click **Settings > Replication Reports Settings**.
- 3 Select the **Enable Replication Reports** option.
- 4 In the **Report Recurrence** section, click **Edit**, then specify the appropriate recurrence pattern for the reports. You can click **Close** to collapse the section.
- 5 In the **Recipients** section, click **Edit Recipients**, type the appropriate email addresses separated by commas, then click **OK**. You can click **Remove** next to an email address to delete the recipient from the list.



- 6 (Optional) In the Forge **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 [“Generating Workload and Workload Protection Reports” on page 67](#).

2.6 Configuring Language Settings 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:

- 1 Access the Languages setting in your web browser:
 - ♦ **Chrome:**
 1. From the Chrome menu, click **Settings**, then scroll to and click **Show advanced settings**.
 2. Scroll to **Languages**, then click **Language and input settings**.
 - ♦ **Firefox:**
 1. From the **Tools** menu, select **Options**, then select the **Content** tab.
 2. Under **Languages**, click **Choose**.
 - ♦ **Internet Explorer:**
 1. From the **Tools** menu, select **Internet Options**, then select the **General** tab.
 2. Under **Appearance**, click **Languages**.
- 2 Add the required language and move it up the top of the list.
- 3 Save the settings, then start the client application by connecting to your PlateSpin Forge Server. See [“Launching the PlateSpin Forge Web Interface” on page 21](#).

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 operating system language:

- 1 Access your Forge VM.
See [“Accessing the Forge VM” on page 48](#).
- 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.

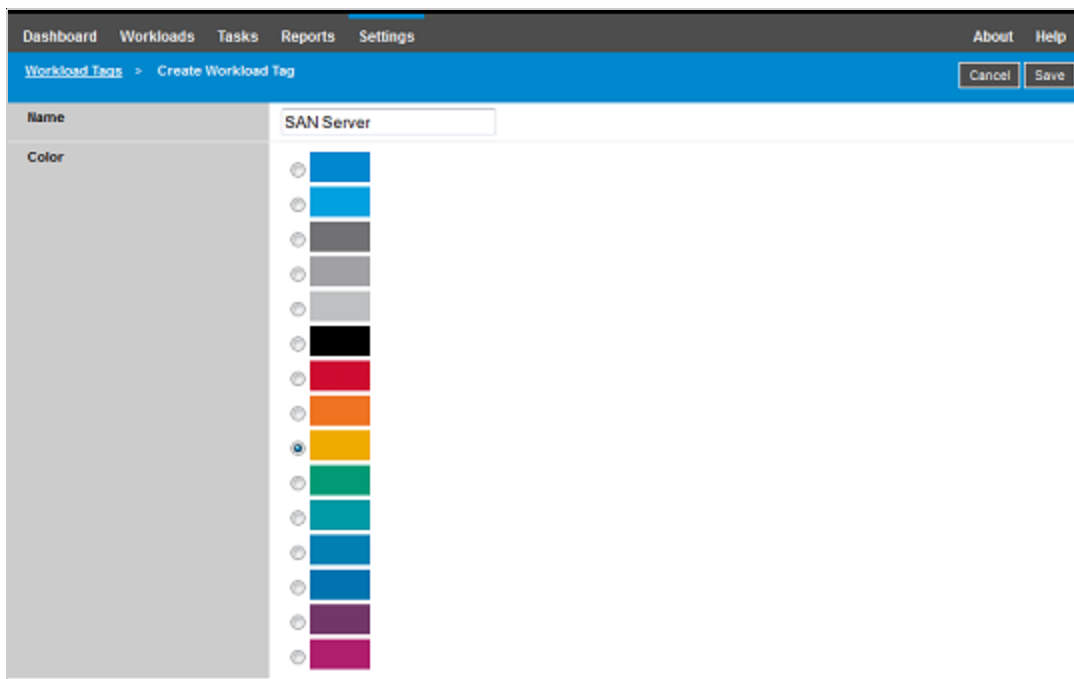
2.7 Using Tags to Help Sort Workloads

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

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

To set up a workload tag:

- 1 In the PlateSpin Forge Web Interface, click **Settings** > **Workload Tags** > **Create Workload Tag**.



- 2 Specify a unique tag name (25-character limit) and associate a color with that description.
- 3 Click **Save** to add this new tag to the list of available workload tags in the Workload Tags view of Settings page.

To edit or delete an available workload tag:

- 1 In the PlateSpin Forge Web Interface, click **Settings** > **Workload Tags**.
- 2 Edit any of the available tags. Click the tag name, modify its name or affiliated color, then click **Save**.
- 3 Delete any tags that you no longer use. Click **Delete** next to the tag, then click **OK** to confirm. You cannot delete a tag if it is associated with any workload.

To add, or associate, a single tag with a workload:

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

- 2 Expand the **Tag** section to view the **Tag** drop-down box.
- 3 Select the name of the tag you want to associate with the workload, then click **Save**.



To remove, or *dissociate*, a single tag from a workload:

- 1 In the workload list, select the workload, then click **Configure** to open its configuration page.
- 2 Expand the **Tag** section to view the **Tag** drop-down box.
- 3 Select the “empty” line in the list of available tag names, then click **Save**.



2.8 Configuring PlateSpin Server Behavior through XML Configuration Parameters

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

https://Your_PlateSpin_Server/platespinconfiguration/

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

To change and apply any configuration parameters:

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

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

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

- ♦ [“Optimizing Data Transfer over WAN Connections” on page 37](#)
- ♦ [“Rebranding the PlateSpin Forge Web Interface” on page 149](#)

2.9 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 [“Configuring PlateSpin Server Behavior through XML Configuration Parameters” on page 37](#).

- ♦ [Section 2.9.1, “Tuning Parameters,” on page 38](#)
- ♦ [Section 2.9.2, “Tuning FileTransferSendReceiveBufferSize,” on page 39](#)

2.9.1 Tuning Parameters

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

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

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

Table 2-6 Default and Optimized File Transfer Configuration Parameters

Parameter	Default Value	Maximum Value
AlwaysUseNonVSSFileTransferForWindows2003	False	
FileTransferCompressionThreadsCount	2	N/A
Controls the number of threads used for packet-level data compression. This setting is ignored if compression is disabled. Because the compression is CPU-bound, this setting might have a performance impact.		
FileTransferBufferThresholdPercentage	10	
Determines the minimum amount of data that must be buffered before creating and sending new network packets.		
FileTransferKeepAliveTimeOutMilliSec	120000	
Specifies how long to wait to start sending keep alive messages if TCP times out.		
FileTransferLongerThan24HoursSupport	True	
FileTransferLowMemoryThresholdInBytes	536870912	
Determines when the server considers itself to be in a low memory state, which causes augmentation of some networking behavior.		
FileTransferMaxBufferSizeForLowMemoryInBytes	5242880	
Specifies the internal buffer size used in a low memory state.		
FileTransferMaxBufferSizeInBytes	31457280	
Specifies internal buffer size for holding packet data.		
FileTransferMaxPacketSizeInButes	1048576	
Determines the largest packets that will be sent.		
FileTransferMinCompressionLimit	0 (disabled)	max 65536 (64 KB)
Specifies the packet-level compression threshold in bytes.		
FileTransferPort	3725	

Parameter	Default Value	Maximum Value
FileTransferSendReceiveBufferSize	0 (8192 bytes)	max 5242880 (5 MB)
<p>Defines the maximum size (in bytes) of the send and receive buffers for TCP connections in the replication network. The buffer size affects the TCP Receive Window (RWIN) size, which sets the number of bytes that can be sent without TCP acknowledgement. This setting is relevant for both file-based and block-based transfers. Tuning the buffer size based on your network bandwidth and latency improves throughput and reduces CPU processing.</p> <p>When the value is set to zero (off), the default TCP window size is used (8 KB). For custom sizes, specify the size in bytes.</p> <p>Use the following formula to determine the proper value:</p> $((\text{LINK_SPEED in Mbps} / 8) * \text{DELAY in sec}) * 1000 * 1024$ <p>For example, for a 100 Mbps link with 10 ms latency, the proper buffer size would be:</p> $(100/8) * 0.01 * 1000 * 1024 = 128000 \text{ bytes}$ <p>For tuning information, see Section 2.9.2, "Tuning FileTransferSendReceiveBufferSize," on page 39.</p>		
FileTransferSendReceiveBufferSizeLinux	0 (253952 bytes)	
<p>Specifies the TCP/IP Receive Window (RWIN) Size setting for file transfer connections for Linux. It controls the number of bytes sent without TCP acknowledgement, in bytes.</p> <p>When the value is set to zero (off), the TCP/IP window size value for Linux is automatically calculated from the FileTransferSendReceiveBufferSize setting. If both parameters are set to zero (off), the default value is 248 KB. For custom sizes, specify the size in bytes.</p> <p>NOTE: In previous release versions, you were required to set this parameter to 1/2 the desired value, but this is no longer required.</p>		
FileTransferShutDownTimeOutInMinutes	1090	
FileTransferTCPTimeOutMilliSec	30000	
Sets both the TCP Send and TCP Receive Timeout values.		
PostFileTransferActionsRequiredTimeInMinutes	60	

2.9.2 Tuning FileTransferSendReceiveBufferSize

The FileTransferSendReceiveBufferSize parameter defines the maximum size (in bytes) of the send and receive buffers for TCP connections in the replication network. The buffer size affects the TCP Receive Window (RWIN) size, which sets the number of bytes that can be sent without TCP

acknowledgement. This setting is relevant for both file-based and block-based transfers. Tuning the buffer size based on your network bandwidth and latency improves throughput and reduces CPU processing.

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

To calculate the optimum buffer size:

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

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

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

1b Enter the following at a command prompt:

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

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

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

The packet needs to be fragmented.

1d Note the latency in seconds.

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

- 2 Calculate a byte value for your initial buffer size:

$$\text{Buffer Size} = (\text{Bandwidth in Mbps} / 8) * \text{Latency in seconds} * 1000 * 1024$$

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

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

$$\text{Buffer Size} = (10240 / 8) * 0.035 * 1000 * 1024 = 45875200 \text{ bytes}$$

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

3a Determine the MSS:

$$\text{MSS} = \text{MTU Size in bytes} - (\text{IP Header Size} + \text{TCP Header Size})$$

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

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

$$\text{MSS} = 1470 \text{ bytes} - (20 \text{ bytes} + 20 \text{ bytes}) = 1430 \text{ bytes}$$

3b Calculate the optimal buffer size:

$$\text{Optimal Buffer Size} = (\text{roundup}(\text{Buffer Size} / \text{MSS})) * \text{MSS}$$

To continue the example:

$$\begin{aligned} \text{Optimal Buffer Size} &= (\text{roundup}(45875200 / 1430)) * 1430 \\ &= 32081 * 1430 \\ &= 45875830 \end{aligned}$$

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

Non-optimal Buffer Size = 32080 * 1430 = 45874400

2.10 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.

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

- Configure a setting to keep the PlateSpin Forge ISO and floppies on the same datastore as the VMware `.vmtx` 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.
- [Section 2.10.1, “Setting Up Workload Files on the Same Datastore,” on page 41](#)
- [Section 2.10.2, “Setting Up VMware Tools for Failover Targets,” on page 41](#)
- [Section 2.10.3, “Expediting the Configuration Process,” on page 42](#)

2.10.1 Setting Up Workload Files on the Same Datastore

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

- 1 From any web browser, open `https://Your_PlateSpin_Server/platespinconfiguration/` to display the configuration web page.
- 2 On the configuration web page, locate the `CreatePSFilesInVmDatastore` server parameter and change its value to `true`.

NOTE: The person configuring the [replication contract](#) is responsible to ensure that the same datastore is specified for all target VM disk files.

- 3 Save your settings and exit the page.

2.10.2 Setting Up VMware Tools for Failover Targets

VMware Tools setup packages can be copied to the failover target during replication so that they can be installed by the configuration service when the VM is booted. This happens automatically when the failover target is able to contact the PlateSpin 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 ESX host:
 - 1a Secure copy (scp) the windows.iso image from the /usr/lib/vmware/isoimages directory on an accessible VMware host to a local temporary folder.
 - 1b Open the ISO and extract its setup packages, saving them to an accessible location:
 - ♦ **VMware 5.x:** The setup packages are setup.exe and setup64.exe.
 - ♦ **VMware 4.x:** The setup packages are VMware Tools.msi and VMware Tools64.msi.
- 2 Create OFX packages from the setup packages you extracted 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

2.10.3 Expediting the Configuration Process

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

To shorten this wait time:

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

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

2.11 Configuring the Refresh Rates for the Web Interface

Several pages in the Web Interface have configurable refresh intervals, as shown in [Table 2-7](#). You can modify the interval setting to meet the needs of your PlateSpin environment.

Table 2-7 Web Interface Default Refresh Intervals

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

- 1 Open the following file in a text editor:

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

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

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

- 3 Save the file.

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

3 Configuring the Forge Appliance

The PlateSpin Forge Appliance consists of the host hardware and a virtual machine running a Microsoft Windows operating system with the PlateSpin Forge software installed. This section provides information about appliance setup and maintenance tasks that you might need to perform regularly on the Forge VM.

- ♦ [Section 3.1, “Setting Up Appliance Networking,” on page 45](#)
- ♦ [Section 3.2, “Using External Storage Solutions with PlateSpin Forge,” on page 46](#)
- ♦ [Section 3.3, “Managing the Forge VM with vSphere Client,” on page 48](#)
- ♦ [Section 3.4, “Physically Relocating the Appliance,” on page 50](#)
- ♦ [Section 3.5, “Recovering the Forge Management VM to Factory Defaults,” on page 54](#)
- ♦ [Section 3.6, “Resetting the Forge Appliance to Factory Defaults,” on page 55](#)

3.1 Setting Up Appliance Networking

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

- ♦ [Section 3.1.1, “About Appliance Host Networking,” on page 46](#)
- ♦ [Section 3.1.2, “Moving or Creating a vSwitch,” on page 46](#)
- ♦ [Section 3.1.3, “Assigning VLAN Tags to Forge Port Groups,” on page 46](#)

3.1.1 About 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 fallback.
- ♦ **Management Network:** The Forge VM network.
- ♦ **Appliance Host Network:** Hypervisor management network. This network is unavailable for selection in the PlateSpin Forge Web Interface.

3.1.2 Moving or Creating a vSwitch

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. See [Move / Create a vSwitch on your Forge Appliance, Knowledgebase Article 7921062](https://www.netiq.com/support/kb/doc.php?id=7921062) (<https://www.netiq.com/support/kb/doc.php?id=7921062>).

3.1.3 Assigning VLAN Tags to Forge Port Groups

To further fine-tune the control of your network traffic, consider assigning a different VLAN ID to each of the individual port groups. Separate VLANs help to ensure that your production network is not interfered with by traffic from workload protection and recovery operations. See [Assigning VLAN Tags to Forge Port Groups, Knowledgebase Article 21057](https://www.netiq.com/support/kb/doc.php?id=7921057) (<https://www.netiq.com/support/kb/doc.php?id=7921057>).

3.2 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 3.2.1, “Using Forge with SAN Storage,” on page 47](#)
- ♦ [Section 3.2.2, “Adding a SAN LUN to Forge,” on page 47](#)

3.2.1 Using Forge with SAN Storage

PlateSpin Forge supports existing external storage solutions, such as Storage Area Network (SAN) implementations. Both Fibre Channel (FC) 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 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 also find information at the following sites:

- ♦ See the documentation for your VMware ESX version on the [VMware Documentation website](http://www.vmware.com/support/pubs/) (<http://www.vmware.com/support/pubs/>):
 - ♦ The *vSphere Storage Guide* discusses the use of ESX Server with storage area networks using Fibre Channel, iSCSI, and Fibre Channel over Ethernet.
 - ♦ The interactive *VMware Compatibility Guide for IO Devices* (<http://www.vmware.com/resources/compatibility/search.php?deviceCategory=io>) lists the currently supported HBAs, HBA drivers, and driver versions.
 - ♦ The interactive *VMware Compatibility Guide for Storage/SAN* (<http://www.vmware.com/resources/compatibility/search.php?deviceCategory=san>) lists the 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.
- ♦ You can learn more about iSCSI by visiting the [Storage Networking Industry Association website](http://www.snia.org/education/storage_networking_primer/ipstorage/) (http://www.snia.org/education/storage_networking_primer/ipstorage/).

3.2.2 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.

- 1 Set up and configure your SAN system.
- 2 Access the appliance host (see “[Downloading and Installing the vSphere Client Program](#)” on [page 48](#)).
- 3 In the VMware 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 VMware client program.

In the PlateSpin Forge Web Interface, 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.

3.3 Managing the Forge VM with vSphere Client

Occasionally you might need to work directly with the Forge 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 VM, including its OS interface and VM settings.

- ♦ [Section 3.3.1, “Downloading and Installing the vSphere Client Program,” on page 48](#)
- ♦ [Section 3.3.2, “Launching the vSphere Client,” on page 48](#)
- ♦ [Section 3.3.3, “Accessing the Forge VM,” on page 48](#)
- ♦ [Section 3.3.4, “Starting and Shutting Down the Forge VM,” on page 49](#)
- ♦ [Section 3.3.5, “Managing Snapshots of the Forge VM on the Appliance Host,” on page 49](#)
- ♦ [Section 3.3.6, “Manually Importing VMs into the Appliance Host’s Datastore,” on page 50](#)
- ♦ [Section 3.3.7, “Guidelines for Applying Security Updates to the Forge VM,” on page 50](#)

3.3.1 Downloading and Installing the vSphere Client Program

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

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

3.3.2 Launching the vSphere Client

- 1 Clicking **Start > Programs > VMware > VMware vSphere | Virtual Infrastructure Client**.
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.

3.3.3 Accessing the Forge VM

- 1 Launch the vSphere Client, then log in with your administrator-level credentials.
- 2 In the inventory panel at the left, locate and select the **PlateSpin Forge VM** item.
- 3 At the top of the right panel, click the **Console** tab.

The Client’s console area displays the Forge 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.

TIP

- ♦ 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.
-

3.3.4 Starting and Shutting Down the Forge VM

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

To shut down the Forge VM:

- 1 Use the vSphere Client to access the Forge VM host.
- 2 Open the VM console for the Forge VM.
- 3 In the VM Console, use the standard Windows procedure to shut down the VM. (Select **Start > Shut Down**.)
Wait for the VM to shut down gracefully.

To start or restart the Forge VM:

- 1 Use the vSphere Client to access the Forge VM host.
- 2 In the inventory panel at the left, right-click the **PlateSpin Forge VM** item and select **Power on**.

3.3.5 Managing Snapshots of the Forge VM 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 you carry out troubleshooting tasks. You might also need to remove snapshots (recovery points) to free storage space.

To manage snapshots of the Forge VM:

- 1 Launch the vSphere Client, then log in to the appliance host with administrator-level credentials.
- 2 In the inventory panel at the left, right-click the **PlateSpin Forge VM** item and select **Snapshot > Take Snapshot**.
- 3 Type a name and a description for the snapshot, then click **OK**.

To revert the Forge VM to a previous state:


- 1 Launch the vSphere Client, then log in to the appliance host with administrator-level credentials.
- 2 In the inventory panel at the left, right-click the **PlateSpin Forge VM** item and select **Snapshot > Snapshot Manager**.
- 3 In the tree representation of the VM states, select a snapshot, then click **Go to**.

To remove snapshots that represent recovery points:

- 1 Launch the vSphere Client, then log in to the appliance host with administrator-level credentials.
- 2 In the inventory panel at the left, right-click the **PlateSpin Forge VM** item and select **Snapshot > Snapshot Manager**.
- 3 In the tree representation of the VM states, select a snapshot, then click **Remove**.

3.3.6 Manually Importing VMs into the Appliance Host's Datastore

You might want to manually import VMs into the Appliance host's datastore when you want your failover workload to be created differentially. See [“Initial Replication Method \(Full and Incremental\)” on page 91](#).

- 1 At the production site, create a VM 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 and Installing the vSphere Client Program” on page 48](#).
- 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 Interface, click **Settings > Containers**, then click the Refresh icon  next to the appliance host.

3.3.7 Guidelines for Applying Security Updates to the Forge VM

- 1 During a maintenance window, access the Forge VM by using the VMware vSphere Client program.
- 2 From within the Forge VM's Windows interface, check for security updates from Microsoft.
- 3 Use the PlateSpin Forge Web Interface 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 VM. See [“Managing Snapshots of the Forge VM on the Appliance Host” on page 49](#).
- 5 Download and install the required security patches. When the installation finishes, reboot the Forge VM.
- 6 Use the PlateSpin Forge Web Interface to resume replications paused in [Step 3](#) and verify that replications are working properly.
- 7 Remove the snapshot of the Forge VM that you took in [Step 4](#). See [“Managing Snapshots of the Forge VM on the Appliance Host” on page 49](#).

3.4 Physically Relocating the Appliance

If you physically relocate your PlateSpin Forge appliance you must modify 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 the *PlateSpin Forge Getting Started Guide*.

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 3.4.1, “Prerequisites for Relocating Forge,” on page 51](#)
- [Section 3.4.2, “Scenario 1: Relocating Forge When the New IP Address Is Known,” on page 51](#)
- [Section 3.4.3, “Scenario 2: Relocating Forge When the New IP Address Is Unknown,” on page 52](#)

3.4.1 Prerequisites for Relocating Forge

Before you start the relocation procedure:

- 1 Pause all replication schedules, ensuring that at least one incremental has run for each workload:
 - 1a Launch your PlateSpin Forge Web Interface.
 - 1b In the Workloads list, verify that at least one incremental has run for each workload.
 - 1c Select all workloads, click **Pause**, then click **Execute**.
 - 1d Ensure that the status **Paused** is displayed for all the workloads.

3.4.2 Scenario 1: Relocating Forge When the New IP Address Is Known

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

- 1 Ensure that replication schedules for all workloads are in a Paused state.
See [“Prerequisites for Relocating Forge” on page 51](#).
- 2 Launch the Forge Appliance Configuration Console (Forge ACC). Launch a web 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**, then click **Continue** to confirm the change.
If you are relocating the Forge Appliance to a new IP address that is not currently accessible (until you physically relocate the Appliance), you must deselect **Verify network settings before saving** on the Confirmation dialog before you continue, or the configuration will not be able to set the new IP address.
- 5 Wait for the configuration process to complete and for the browser window to display the Configuration Successful popup window.
If the configuration process returns a timeout error, ensure you have the Forge VM correctly licensed with a single container license, then try again.

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 In the vSphere Client, shut down the appliance:
 - 6a Shut down the Forge VM. (In the Forge VM Console, select **Start > Shut Down**.)
See [“Starting and Shutting Down the Forge VM” on page 49](#). Wait for the VM to shut down gracefully.
 - 6b Shut down the Appliance Host:
 - 6b1 At the Forge VM 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
```


Wait for the Appliance Host to shut down gracefully.
 - 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, then look up the IP address in the VM's Windows interface.

See ["Accessing the Forge VM" on page 48](#).

2. Use the new IP address to launch the PlateSpin Forge Web Interface, then refresh the container. Click **Settings** > **Containers** > then click the Refresh icon .

- 10 Resume the paused replications.


3.4.3 Scenario 2: Relocating Forge When the New IP Address Is Unknown

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

- 1 Ensure that replication schedules for all workloads are in a Paused state.
See ["Prerequisites for Relocating Forge" on page 51](#).
- 2 In the vSphere Client, shut down the appliance:
 - 2a Shut down the Forge VM. (In the Forge VM Console, select **Start** > **Shut Down**.)
See ["Starting and Shutting Down the Forge VM" on page 49](#). Wait for the VM to shut down gracefully.
 - 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 down the appliance.
- 3 Disconnect your appliance, move it, attach it to the new network, then power it on.

- 4 Set up an administrative computer (a laptop computer is recommended) so that it is able to communicate with Forge at its currently configured IP address (the IP address at the old site), then connect it to the Appliance using either of the following methods:
 - ♦ Connect both machines to the same Ethernet switch or hub.



-OR-

- ♦ Connect the two machines using an Ethernet crossover cable.



- 5 Launch the Forge ACC: open a browser and go to `http://<Forge_IP_address>:10000`.
- 6 Log in using the `forgeuser` account, then 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.
If the configuration process returns a timeout error, ensure you have the Forge VM correctly licensed with a single container license, then try again.

NOTE: The link in the popup window for the new Forge ACC address will not work until after you connect the appliance 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 settings 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 [“Accessing the Forge VM” on page 48](#)

2. Use the new IP address to launch the PlateSpin Forge Web Interface, then refresh the container. Click **Settings > Containers** > then click the Refresh icon ↻.

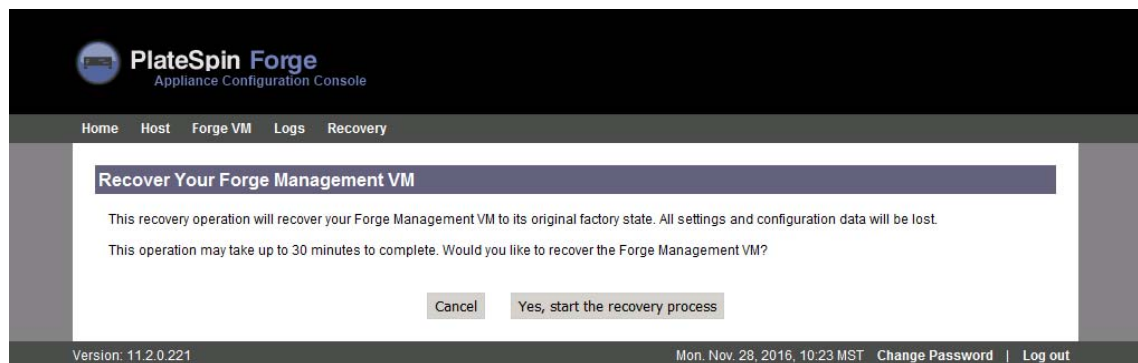
-
- 12 Resume the paused replications.

3.5 Recovering the Forge Management VM to Factory Defaults

To recover the Forge Management VM to factory defaults:

- 1 Pause all replication schedules and ensure that no replications are running:
 - 1a In your PlateSpin Forge Web Interface, select all workloads, click **Pause Schedule**, then click **Execute**.
 - 1b Ensure that the status **Paused** is displayed for all the workloads.
 - 1c Ensure that the Replication Status of each of the protected workloads is **Idle**. Wait for any replications that are underway to complete, or abort the running operations.
- 2 (Optional) Export the workload contracts data if you want to use the same contracts after the VM recovery to factory defaults:
 - 2a Log in as an Administrator user to the Forge Management VM.
 - 2b Open a Command Prompt and navigate to `D:\Program Files\PlateSpin Forge Server\PlateSpin Forge\bin\ImportExport`.
 - 2c Enter

```
ImportExportAll.bat /export "C:\temp"
```
 - 2d Copy the exported XML file to a location where it can be available for import after the VM recovery.
- 3 Shut down the Forge Management VM. See [“Starting and Shutting Down the Forge VM” on page 49](#).
- 4 Launch the Forge Appliance Configuration Console (Forge ACC): open a web browser and go to `http://<Forge_IP_address>:10000`.
- 5 Log in using the `forgeuser` account, then click **Recovery**.
- 6 Review the **Recover Your Forge Management VM** advisory, then confirm by clicking **Yes, start the recovery process**.



- 7 Wait for the configuration process to complete and for the browser window to display the Recovery Successful message.
- 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 recovery:

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 ["Accessing the Forge VM" on page 48](#)).
 2. Use the new IP address to launch the PlateSpin Forge Web Interface and refresh the container (click **Settings** > **Containers** > then click ↻).
-

- 10 If you exported contracts in [Step 2](#), import them to the recovered VM:

10a Log in as an Administrator user to the Forge Management VM.

10b Copy the exported XML file to `C:\temp`.

10c Open a Command Prompt and navigate to `D:\Program Files\PlateSpin Forge Server\PlateSpin Forge\bin\ImportExport`.

10d Enter

```
ImportExportAll.bat /import "C:\temp"
```


10e After the import is complete, connect to the Forge Web Interface to confirm that the data is intact.

- 11 Resume the paused replications.

In case of failure, call Micro Focus Customer Care and be prepared to provide the log files. Log files required for troubleshooting the recovery 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.



3.6 Resetting the Forge Appliance to Factory Defaults

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

To reset your Forge Appliance 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 [“Physically Relocating the Appliance” on page 50](#) 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 the VMware host, then power down the appliance.

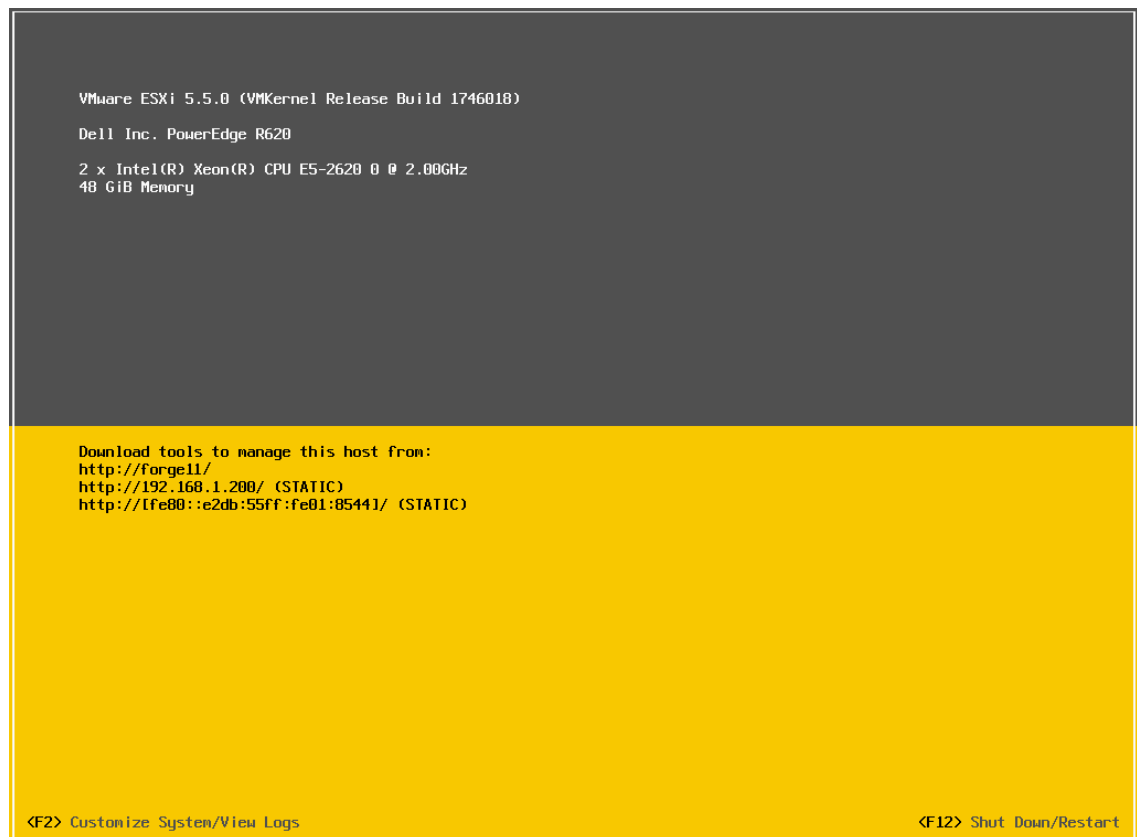
3f If the Forge model does not have an optical media player, attach a CD/DVD player to the Forge Appliance.

3g Boot from the Forge CD media and wait for the SYSLINUX menu to display.

4 Select the **PlateSpin Forge Factory Reset** option and press Enter. Ensure 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 VMware window should look similar to the following:



If the reset process fails, the console message should look similar to the following:

```
The PlateSpin Forge Installation has failed.


-----
CAUSE OF FAILURE:
- Unable to recover the Forge Management VM from backup
- The Forge Factory Reset FIRST_BOOT.SH script failed
-----

Please consult the system log file located at /vmfs/volumes/ForgeSystem/PLATESPINFORGE_LOGS/
forge.log for details.
```

In case of failure:

- ♦ Call Micro Focus Customer Care 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, “Accessing the PlateSpin Forge Web Interface,” on page 59](#)
- ♦ [Section 4.2, “Elements of the PlateSpin Forge Web Interface,” on page 60](#)
- ♦ [Section 4.3, “Workloads and Workload Commands,” on page 62](#)
- ♦ [Section 4.4, “Managing Multiple Instances of PlateSpin Protect and PlateSpin Forge,” on page 64](#)
- ♦ [Section 4.5, “Generating Workload and Workload Protection Reports,” on page 67](#)

4.1 Accessing the PlateSpin Forge Web Interface

To launch the PlateSpin Forge Web Interface:

- 1 Open a web browser and go to:

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

Replace `<hostname | IP_address>` with the DNS hostname or the IP address of your Forge VM.

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


- 2 Log in using the local Administrator user credentials for the Forge VM, or log in as an authorized user.

For information about setting up additional users for PlateSpin, see [“Configuring User Authorization and Authentication” on page 24](#).

4.2 Elements of the PlateSpin Forge Web Interface

The Dashboard page of the PlateSpin Forge Web Interface 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 61](#)
- ♦ [Section 4.2.2, “Visual Summary Panel,” on page 61](#)
- ♦ [Section 4.2.3, “Tasks and Events Panel,” on page 62](#)

NOTE: You can alter certain elements of the Web Interface to match your organization branding. For more information, see [“Rebranding the PlateSpin Forge Web Interface” on page 149](#).

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 62](#).
- ♦ **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 67](#).
- ♦ **Settings:** Displays the Settings page, which provides access to the following configuration options:
 - ♦ **Protection Tiers:** See [“Protection Tiers” on page 89](#).
 - ♦ **Workload Tags:** See [“Using Tags to Help Sort Workloads” on page 36](#).
 - ♦ **Permissions:** See [“Configuring User Authorization and Authentication” on page 24](#).
 - ♦ **Notification Settings:** [“Setting Up Automatic Event Notifications by Email” on page 32](#).
 - ♦ **Replication Reports Settings:** [“Setting Up Automatic Replication Reports by Email” on page 34](#)
 - ♦ **SMTP:** See [“SMTP Configuration” on page 31](#).
 - ♦ **Licenses:** See [“Activating Your Product License” on page 22](#).

4.2.2 Visual Summary Panel

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

Protection Status

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








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

Workload Summary

The Workload Summary presents the health status of each licensed workload listed on the Workloads page. The maximum number of workload status dot icons matches the number of installed workload licenses on the PlateSpin Server. For an unlimited license, the summary displays 96 dot icons. [Table 4-1](#) describes the different workload states represented by the dot icons.

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

Table 4-1 Dot Icon Workload Representation

 Protected	 Unprotected
 Failed	 Unprotected – Error
 Underprotected	 Expired
	 Unused

License Summary

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

Storage

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

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 [“Configuring Automatic Email Notifications of Events and Reports” on page 31](#).


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




Table 4-2 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.

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

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

- [Section 4.4.1, “Using the PlateSpin Forge Management Console,” on page 64](#)
- [Section 4.4.2, “About PlateSpin Forge Management Console Cards,” on page 65](#)
- [Section 4.4.3, “Adding Instances of PlateSpin Protect and PlateSpin Forge to the Management Console,” on page 66](#)
- [Section 4.4.4, “Managing Cards on the Management Console,” on page 66](#)

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 DNS hostname of the Forge VM that is designated as the Manager.

- 2 Log in with your user name 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 and PlateSpin Forge, such as:

- ♦ IP address/hostname
- ♦ Location
- ♦ Version number
- ♦ Workload count
- ♦ Workload status

- ♦ Storage capacity
- ♦ Remaining free space

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

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

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

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

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

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

- 2 Specify the URL of the PlateSpin Server host or Forge VM. Use HTTPS if SSL is enabled.
- 3 (Optional) Enable the **Use Management Console Credentials** check box to use the same credentials as those used by the console. When it is selected, the console automatically populates the **Domain\Username** field.
- 4 In the **Domain\Username** field, type a domain name and a user name valid for the instance of PlateSpin Protect or PlateSpin Forge that you are adding. In the **Password** field, type the corresponding password.
- 5 (Optional) Specify a unique descriptive **Display Name** (up to 15 characters) for the PlateSpin Server, its **Location** (up to 20 characters), and any **Notes** you might require (up to 400 characters).
- 6 Click **Add**.

A new card is added to the dashboard.

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 the following reports that provide analytical insight into your workload protection contracts over time:

- ♦ **Workload Protection:** Reports replication events for all workloads over a selectable time window.
- ♦ **Replication History:** Reports replication type, size, time, and transfer speed per selectable workload over a selectable time window.
- ♦ **Replication Window:** Reports the dynamics of full and incremental replications that can be summarized by **Average**, **Most Recent**, **Sum**, and **Peak** perspectives.
- ♦ **Current Protection Status:** Reports **Target RPO**, **Actual RPO**, **Actual TTO**, **Actual RTO**, **Last Test Failover**, **Last Replication**, and **Test Age** statistics.
- ♦ **Events:** Reports system events for all workloads over a selectable time window.
- ♦ **Scheduled Events:** Reports only upcoming workload protection events.

Figure 4-6 Options for a Replication History Report

Date	Replication Event	Total Time	Transfer Time	Transfer Size	Transfer Speed
1/17/2015 4:01 AM	Incremental replication did not run as scheduled because the workload was busy	—	—	0 MB	0.00 Mbps
1/17/2015 4:00 AM	Incremental replication did not run as scheduled because the workload was busy	—	—	0 MB	0.00 Mbps
1/16/2015 4:01 AM	Incremental replication did not run as scheduled because the workload was busy	—	—	0 MB	0.00 Mbps
1/16/2015 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 and Recovery

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 managed by 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 69](#)
- ♦ [Section 5.2, “Adding Containers \(Protection Targets\),” on page 70](#)
- ♦ [Section 5.3, “Adding Workloads,” on page 72](#)
- ♦ [Section 5.4, “Configuring Protection Details and Preparing the Replication,” on page 73](#)
- ♦ [Section 5.5, “Starting the Workload Protection,” on page 76](#)
- ♦ [Section 5.6, “Aborting Commands,” on page 77](#)
- ♦ [Section 5.7, “Failover,” on page 77](#)
- ♦ [Section 5.8, “Failback,” on page 79](#)
- ♦ [Section 5.9, “Reprotecting a Workload,” on page 83](#)

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** Ensure that PlateSpin Forge supports your workload.
See [“Supported Configurations” on page 9](#).
 - 1b** Ensure that your workloads and VM containers meet access and network prerequisites.
See [“Configuring Access and Communication Settings across your Protection Network” on page 28](#).
 - 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 [Knowledgebase 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 [Knowledgebase 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 [“Using Freeze and Thaw Scripts for Every Replication \(Linux\)” on page 92](#).

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

Workloads that you want to protect and containers that host failover workloads must be properly inventoried. You can add workloads and containers in any order; however, every protection contract requires a defined workload and container that were inventoried by the PlateSpin Server. See [“Adding Containers \(Protection Targets\)” on page 70](#) and [“Adding Workloads” on page 72](#).

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

See [“Configuring Protection Details and Preparing the Replication” on page 73](#).

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

See [“Starting the Workload Protection” on page 76](#).

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

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

6 Failover: This step carries out a failover of your protected workload to its replica running in your appliance host. See [“Failover” on page 77](#).

7 Failback: This step corresponds to the business resumption phase after you have addressed any problems with your production workload. See [“Failback” on page 79](#).

8 Reprotection: This step enables you to redefine the original protection contract for your workload. See [“Reprotecting a Workload” on page 83](#)

Most of these steps are represented by workload commands on the Workloads page. See [“Workloads and Workload Commands” on page 62](#).

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

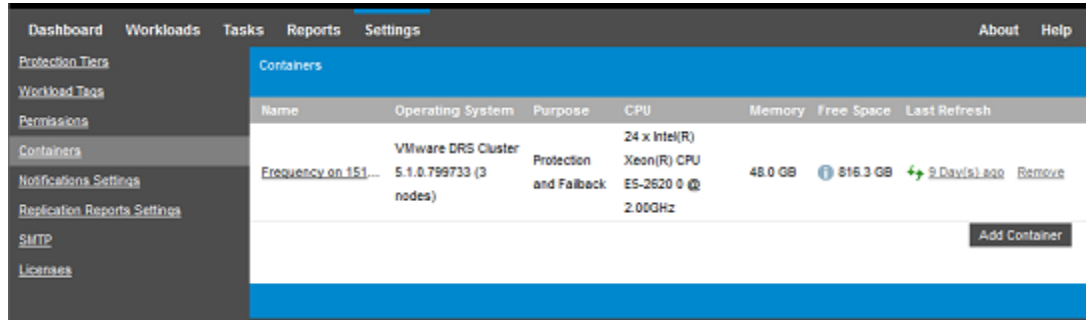
5.2 Adding Containers (Protection Targets)

A container is a protection infrastructure that acts as the host of a protected workload’s regularly-updated replica. That infrastructure can be either a VMware ESX Server or a VMware DRS Cluster. PlateSpin Forge resides in a protection container on the appliance. You can define only Failback that reside on a target VM infrastructures.

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

To add a container:

- 1 In your PlateSpin Forge Web Interface, click **Settings > Containers > Add Container**.



- 2 Specify the following parameters:

- ♦ **Type:** Select the type of the container:
 - ♦ **VMware ESX Server**
 - ♦ **VMware DRS Cluster**

Ensure that the VM container is supported. See [“Supported VM Containers” on page 13](#).



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

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


- ♦ A cluster contains no ESX hosts.
- ♦ A cluster name is not unique across the vCenter server (even if it has a unique inventory path).
- ♦ None of the cluster members are accessible (for example, because the vCenter server is in maintenance mode).
- ♦ **Username/Password:** Provide administrator-level credentials for accessing the target host. See [“Guidelines for Workload and Container Credentials” on page 86](#).
- ♦ **Purpose:** (VM containers only) Select the required purpose for the VM container:
 - ♦ **Failback**

PlateSpin Forge allows you to add containers only for Failback operations.

- 3 Click **Add**.

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

To refresh a container:

- 1 Click the **Refresh** icon  next to the container you want to refresh.
This performs a re-inventory of the container.

To remove a container:

- 1 Click **Remove** next the container that you want to remove.

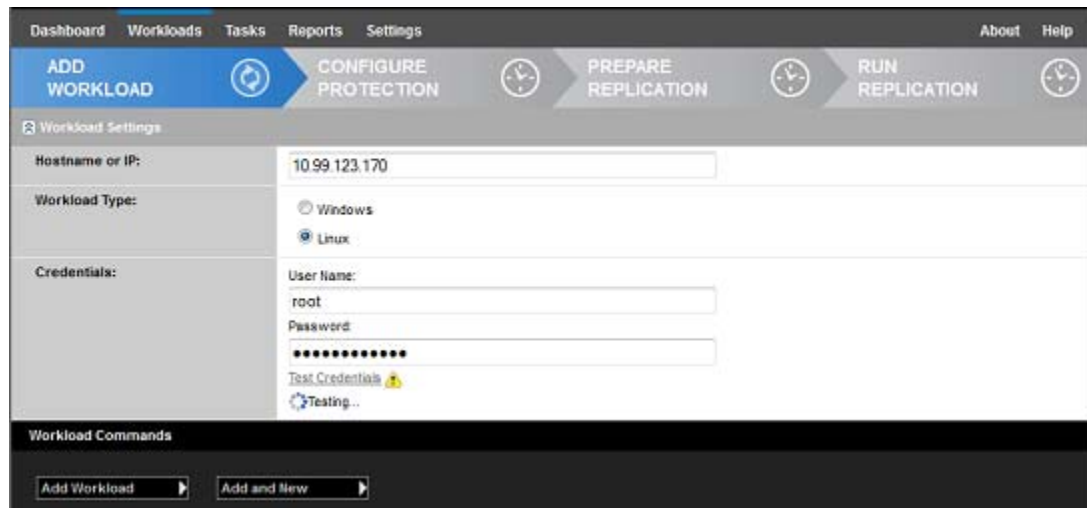
5.3 Adding Workloads


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 69.
- 2 On the Dashboard or Workloads page, click **Add Workload**.
The PlateSpin Forge Web Interface displays the Add Workload page.



- 3 Specify the required workload details:
 - ♦ **Workload Settings:** Specify your workload’s hostname or IP address, the operating system, and administrator-level credentials.
Use the required credential format. See “[Guidelines for Workload and Container Credentials](#)” on page 86.
To ensure 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.
- 5 (Conditional) If you haven’t added a container yet for use with this workload, add one to prepare for protecting the workload. See “[Adding Containers \(Protection Targets\)](#)” on page 70.
- 6 Continue with “[Configuring Protection Details and Preparing the Replication](#)” on page 73.

5.4 Configuring Protection Details and Preparing the Replication

Protection details control the workload protection and recovery settings and behavior over the entire life cycle of a workload under protection. At each phase of the protection and recovery workflow (Add-inventory, initial and ongoing Replications, Failover, Failback, and Reprotect), relevant settings are read from the protection details. See [“Basic Workflow for Workload Protection and Recovery” on page 69](#). This collection of currently-active settings pertaining to the complete lifecycle of a workload’s protection is referred to as the workload’s *protection contract*.

To configure your workload’s protection details:

- 1 Add a workload. See [“Adding Workloads” on page 72](#).
- 2 Add a container. See [“Adding Containers \(Protection Targets\)” on page 70](#).
- 3 On the Workloads page, select the required workload and click **Configure**.

Alternatively, you can click the name of the workload.

NOTE: If the PlateSpin Forge inventory does not have a container yet, the system prompts you to add one; do so by clicking **Add Container** at the bottom.

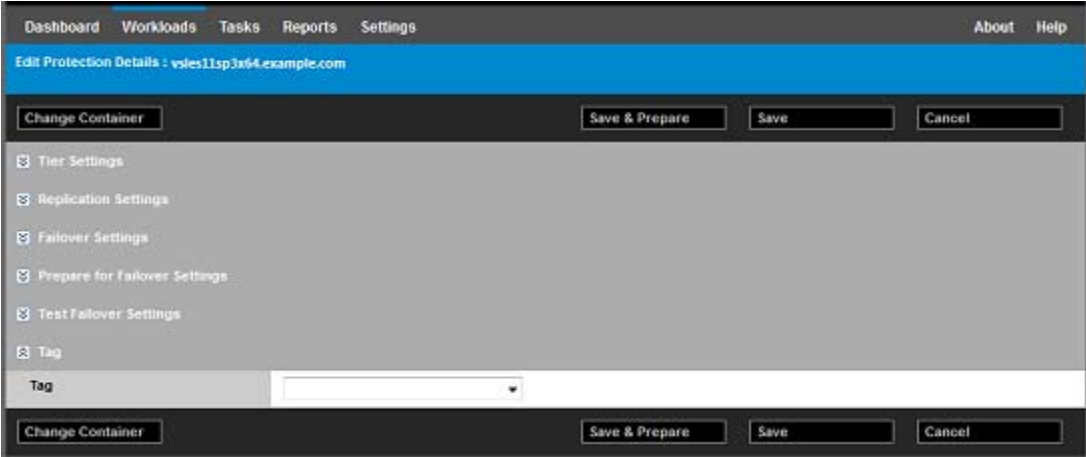
- 4 Select an **Initial Replication Method**. This indicates whether you want volume data transferred entirely from your workload to the failover VM or synchronized with volumes on an existing VM. See [“Initial Replication Method \(Full and Incremental\)” on page 91](#).
- 5 Configure the protection details in each set of settings as dictated by your business continuity needs. See [“Workload Protection Details” on page 73](#).
- 6 Correct any validation errors, if displayed by the PlateSpin Forge Web Interface.
- 7 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.

5.4.1 Workload Protection Details

Workload protection details are represented by five sets of parameters, as described in [Table 5-1](#):




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

Table 5-1 Workload Protection Details

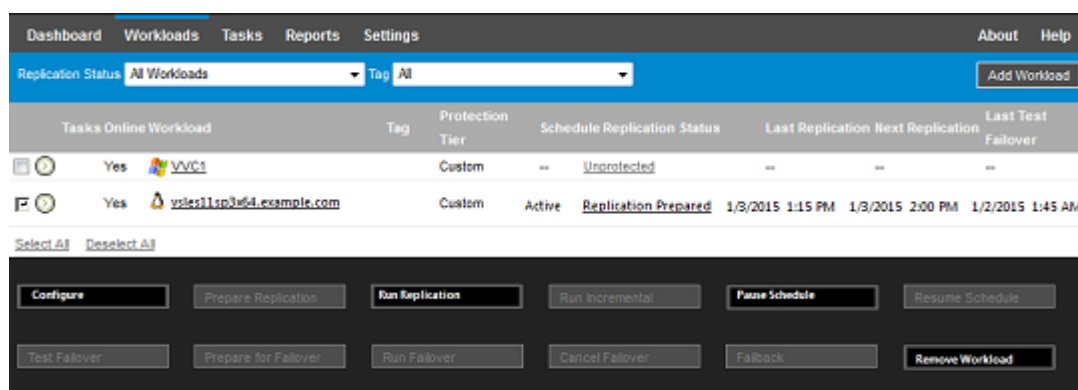
Parameter Settings	Details
Tier Settings	
Protection Tier	Specify the Protection Tier that the current protection uses. See “Protection Tiers” on page 89 .
Replication Settings	
Transfer Method	<p>(Windows) Select a file-based or block-based data transfer mechanism. For information about block-level replication with or without block-based components, see “Data Transfer” on page 86.</p> <p>To enable encryption, select the Encrypt Data Transfer option. See “Data Encryption” on page 88.</p>
Transfer Encryption	(Linux) To enable encryption, select the Encrypt Data Transfer option. See “Data Encryption” on page 88 .
Source Credentials	Specify the credentials required for accessing the workload. See “Guidelines for Workload and Container Credentials” on page 86 .
CPU	<p>(VM containers using VMware 5.1, 5.5, and 6.0 with a minimum VM hardware Level 8) Specify the number of sockets and the number of cores per socket for the failover workload. It automatically calculates the total cores. This parameter applies on the initial setup of a workload with an initial replication setting of Full.</p> <p>NOTE: The maximum number of cores the workload can use is subject to external factors such as the guest operating system, the VM hardware version, VMware licensing for the ESXi host, and ESXi host compute maximums for vSphere (see vSphere 5.1 Configuration Maximums (http://www.vmware.com/pdf/vsphere5/r51/vsphere-51-configuration-maximums.pdf)).</p> <p>Some distributions of a guest OS might not honor the cores and cores per socket configuration. For example, guest OSes using SLES 10 SP4 and OES 2 SP3 retain their original cores and sockets settings as installed, whereas other SLES, RHEL, and OES distributions honor the configuration.</p>
Number of CPUs	<p>(VM containers using VMware 4.1) Specify the required number of vCPUs (virtual CPUs) to assign to the failover workload. This parameter applies on the initial setup of a workload with an initial replication setting of Full. Each vCPU is presented to the guest OS on the VM container as a single core, single socket.</p>
Replication Network	<p>Separate replication traffic based on virtual networks defined on your appliance host. See “Networking” on page 96.</p> <p>For this setting, you can also specify an MTU value to be used by the PlateSpin Forge Linux RAM Disk (LRD) replication network. Setting the value can help avoid jabber over networks (for example, a VPN) that have a smaller MTU value. The default value is empty string (nothing listed in the text box). When networking is configured in the LRD, this allows the network device to set its own default (which is usually 1500). If you enter a value, PlateSpin Forge adjusts the MTU while configuring the network interface.</p>

Parameter Settings	Details
Allowed Networks	Specify one or more network interfaces (NIC or IP address) on the source to use for replication traffic.
Resource Pool for Target VM	(VM container is part of a DRS Cluster) Specify the Resource Pool location where the failover VM is to be created.
VM Folder for Target VM	(VM container is part of a DRS Cluster) Specify the VM folder location where the failover VM is to be created.
Configuration File Datastore	Select a datastore associated with your appliance host for storing VM configuration files. See “Recovery Points” on page 90 .
Protected Volumes	Select volumes for protection and to assign their replicas to specific datastores on your appliance host.
Thin Disk	Select to enable the thin-provisioned virtual disk feature, whereby a virtual disk appears to the VM to have a set size, but only consumes the amount of disk space that is actually required by data on that disk.
Protected Logical Volumes	(Linux) Specify one or more LVM logical volumes to be protected for a Linux workload or the NSS Pools on an Open Enterprise Server workload.
Non-volume Storage	(Linux) Specify a storage area (such as a swap partition) that is associated with the source workload. This storage is re-created in the failover workload.
Volume Groups	(Linux) Specify the LVM volume groups to be protected with the LVM logical volumes listed in the Protected Logical Volumes section of the settings.
Services/Daemons to Stop During Replication:	Select Windows services or Linux daemons that are automatically stopped during the replication. See “Service and Daemon Control” on page 92 .
Failover Settings	
VM Memory	Specify the amount of memory allocated to the failover workload.
Hostname and Domain/Workgroup affiliation	Specify the identity and domain/workgroup affiliation of the failover workload when it is live. For domain affiliation, domain administrator credentials are required.
Network Connections	Specify the LAN settings of the failover workload. See “Networking” on page 96 .
DNS Servers	Specify the IP address of the primary DNS server and an alternative DNS (optional).
Services/Daemon States to Change	Specify the startup state of specific application services (Windows) or daemons (Linux) See “Service and Daemon Control” on page 92 .
Prepare for Failover Settings	
Temporary Failover Network	Specify the temporary LAN settings of the failover workload during the optional Prepare for Failover operation. See “Networking” on page 96 .
Test Failover Settings	
VM Memory	Assign the required RAM to the temporary workload.
Hostname	Assign a hostname to the temporary workload.
Domain/Workgroup	Affiliate the temporary workload with a domain or a workgroup. For domain affiliation, domain administrator credentials are required.

Parameter Settings	Details
Network Connections	Specify the LAN settings of the temporary workload. See “Networking” on page 96 .
DNS Servers	Specify the IP address of the primary DNS server and an alternative DNS (optional).
Service/Daemon States to Change	Specify the startup state of specific application services (Windows) or daemons (Linux). See “Service and Daemon Control” on page 92 .
Tags	
Tag	(Optional) Assign a tag to this workload. See “Using Tags to Help Sort Workloads” on page 36 .

5.5 Starting the Workload Protection

Workload protection is started by the **Run Replication** command:



You can execute the Run Replication command after:

- ♦ Adding a workload.
- ♦ Configuring the workload's protection details.
- ♦ Preparing the initial replication.

When you are ready to proceed:

- 1 On the Workloads page, select the required workload, then click **Run Replication**.
- 2 Click **Execute**.

PlateSpin 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.
-


5.6 Aborting Commands

You can abort a command after executing it and while it is underway, on the Command Details page of that particular command.

To access the Command Details page of any command that is underway:

- 1 Go to the Workloads page.
- 2 Locate the required workload and click the link representing the command currently executing on that workload, such as **Running Incremental**.

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



- 3 Click **Abort**.

5.7 Failover

In a *failover* operation, the failover workload within a PlateSpin Forge VM container takes over the business function of a failed production workload.

- ♦ [Section 5.7.1, “Detecting Offline Workloads,” on page 77](#)
- ♦ [Section 5.7.2, “Performing a Failover,” on page 78](#)
- ♦ [Section 5.7.3, “Using the Test Failover Feature,” on page 78](#)

5.7.1 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 “[Tier Settings](#)” row in “[Workload Protection Details](#)” on page 73.

If notifications are configured along with SMTP settings, PlateSpin Forge simultaneously sends a notification email to the specified recipients. See “[Configuring Automatic Email Notifications of Events and Reports](#)” on page 31.

If a workload failure is detected while the status of the replication is **Idle**, you can proceed to the **Run Failover** command. If a workload fails while an incremental is underway, the job stalls. In this case, abort the command (see “[Aborting Commands](#)” on page 77), and then proceed to the **Run Failover** command. See “[Performing a Failover](#)” on page 78.

Figure 5-1 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')



5.7.2 Performing a Failover

Failover settings, including the failover workload's network identity and LAN settings, are saved together with the workload's protection details at configuration time. See [“Failover Settings”](#) in [“Workload Protection Details”](#) on page 73.

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 [“Recovery Points”](#) on page 90). Click **Execute** and monitor the progress. Upon completion, the replication status of the workload should indicate **Live**.

For testing the failover workload or testing the failover process as part of a planned disaster recovery exercise, see [“Using the Test Failover Feature”](#) on page 78.

5.7.3 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 an isolated networking environment for testing the functionality of the failover and verifying the integrity of the failover workload.

When you execute the command, PlateSpin Forge applies the Test Failover Settings, as saved in the workload protection details, to the failover workload. See [“Test Failover Settings”](#) in [“Workload Protection Details”](#) on page 73.

To use the Test Failover feature:

- 1 Define an appropriate time window for testing and ensure that there are no replications underway. The replication status of the workload must be **Idle**.
- 2 On the Workloads page, select the required workload, click **Test Failover**, select a recovery point (see [“Recovery Points”](#) on page 90), and 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 [“Downloading and Installing the vSphere Client Program” on page 48](#).

- 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.

5.8 Failback

A *failback* operation restores the business function of a failed production workload in its original environment when the business function of a temporary failover workload is no longer required. Failback is the next logical step after a failover; it transfers the failover workload to its original infrastructure or, if necessary, a new one.

Supported failback methods depends on the target infrastructure type and the degree of automation of the failback process:

- ♦ **Automated Failback to a Virtual Machine:** Supported for VMware ESX platforms and VMware DRS Clusters.
- ♦ **Semi-Automated Failback to a Physical Machine:** Supported for all physical machines.
- ♦ **Semi-Automated Failback to a Virtual Machine:** Supported for Microsoft Hyper-V platforms.

The following topics provide more information:

- ♦ [Section 5.8.1, “Automated Failback to a VM Platform,” on page 79](#)
- ♦ [Section 5.8.2, “Semi-Automated Failback to a Physical Machine,” on page 82](#)
- ♦ [Section 5.8.3, “Semi-Automated Failback to a Virtual Machine,” on page 82](#)

5.8.1 Automated Failback to a VM Platform

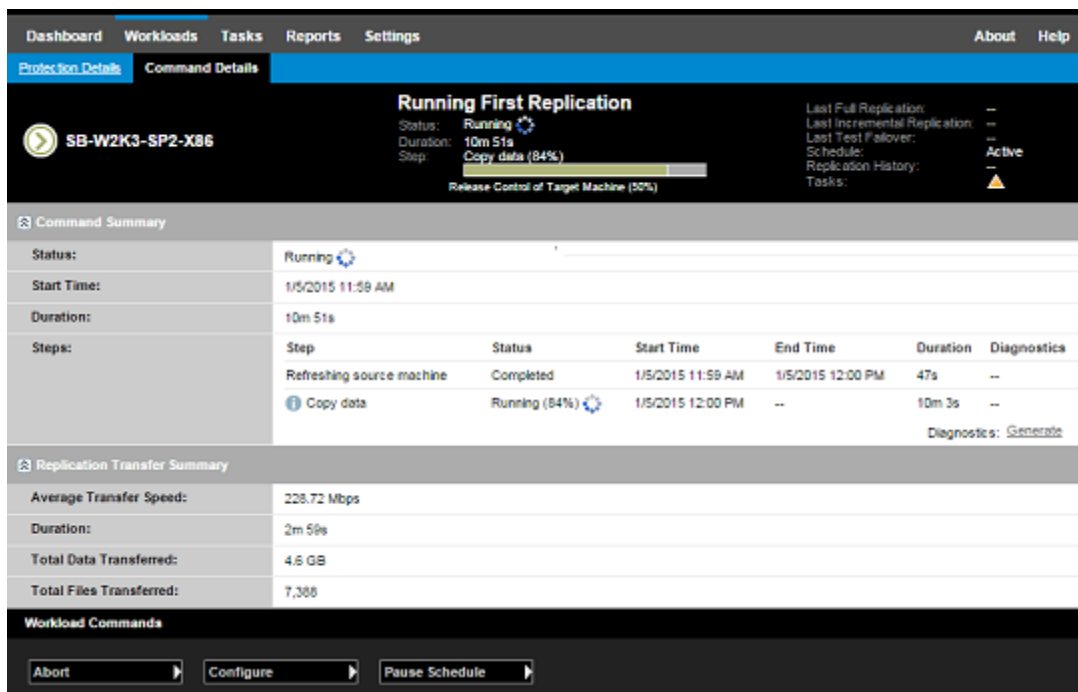
PlateSpin Forge supports automated failback for Failback containers on a supported VMware ESXi Server or a VMware DRS Cluster. See [“Supported VM Containers” on page 13](#).

To execute an automated failback of a failover workload to a target VMware container:

- 1 Following a failover, select the workload on the Workloads page and click **Failback**.
The system prompts you to make the following selections
- 2 Specify the following sets of parameters:
 - ♦ **Workload Settings:** Specify the failover workload’s hostname or IP address and provide administrator-level credentials. Use the required credential format. See [“Guidelines for Workload and Container Credentials” on page 86](#).

- ♦ **Failback Target Settings:** Specify the following parameters:
 - ♦ **Replication Method:** Select the scope of data replication. If you select **Incremental**, you must **Prepare** a target. See “[Initial Replication Method \(Full and Incremental\)](#)” on page 91.
 - ♦ **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 80.
- 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. See [Table 5-2](#) for information about parameter settings.

Table 5-2 Failback Details (Workload to VM)

Parameter Settings	Details
Failback Settings	

Parameter Settings	Details
Transfer Method	Select a data transfer mechanism and security through encryption. See “Data Transfer” on page 86 .
Failback Network	Specify the network to use for failback traffic. This is a dedicated network based on virtual networks defined on your appliance host. See “Networking” on page 96 .
VM Datastore	Select a datastore associated with your failback container for the target workload.
Volume Mapping	If the initial replication method is specified as “incremental”, select source volumes and map to volumes on the failback target for synchronization.
Services/Daemons to stop	Specify the application services (Windows) or daemons (Linux) that are automatically stopped during the failback. See “Service and Daemon Control” on page 92 .
Alternative Address for Source	Specify an additional IP address for the failed-over VM if applicable. See “Protection Across Public and Private Networks Through NAT” on page 30 .
Workload Settings	
CPU	<p>(VM containers using VMware 5.1, 5.5, and 6.0 with a minimum VM hardware Level 8) Specify the number of sockets and the number of cores per socket for the failback to virtual workload. It automatically calculates the total cores. This parameter applies on the initial setup of a workload with an initial replication setting of Full.</p> <p>NOTE: The maximum number of cores the workload can use is subject to external factors such as the guest operating system, the VM hardware version, VMware licensing for the ESXi host, and ESXi host compute maximums for vSphere (see vSphere 5.1 Configuration Maximums (http://www.vmware.com/pdf/vsphere5/r51/vsphere-51-configuration-maximums.pdf)).</p> <p>Some distributions of a guest OS might not honor the cores and cores per socket configuration. For example, guest OSes using SLES 10 SP4 and OES 2 SP3 retain their original cores and sockets settings as installed, whereas other SLES, RHEL, and OES distributions honor the configuration.</p>
Number of CPUs	(VM containers using VMware 4.1) Specify the required number of vCPUs (virtual CPUs) to assign to the failback to virtual workload. This parameter applies on the initial setup of a workload with an initial replication setting of Full . Each vCPU is presented to the guest OS on the VM container as a single core, single socket.
VM Memory	Assign the required RAM to the target workload.
Hostname, Domain/Workgroup	Specify the identity and domain/workgroup affiliation of the target workload. For domain affiliation, domain administrator credentials are required.
Network Connections	Specify the network mapping of the target workload based on the virtual networks of the underlying VM container.
Service States to Change	Specify the startup state of specific application services (Windows) or daemons (Linux). See “Service and Daemon Control” on page 92 .
Post-Failback Settings	

Parameter Settings	Details
Reprotect Workload	Select this option if you plan to re-create the protection contract for the target workload after deployment. This option maintains a continuous event history for the workload and auto-assigns/designates a workload license.
Reprotect after Failback	Select this option if you intend to re-create a protection contract for the target workload. When the failback is complete, a Reprotect command will be available in the PlateSpin Forge Web Interface for the failed-back workload.
No reprotect	Select this option if you do not intend to re-create a protection contract for the target workload. To protect the failed-back workload upon completion, you will have to re-inventory that workload and reconfigure its protection details.

5.8.2 Semi-Automated Failback to a Physical Machine

Use these steps to fail a workload back to a physical machine after a failover. The physical machine might be either the original infrastructure or a new one.

- 1 Register the required physical machine with your PlateSpin Server. See [“Failback to Physical Machines” on page 96](#).
- 2 If there are missing or incompatible drivers, upload the required drivers to the PlateSpin Forge device driver database. See [“Managing Device Drivers” on page 105](#).
- 3 Following a failover, select the workload on the Workloads page and click **Failback**.
- 4 Specify the following sets of parameters:
 - ♦ **Workload Settings:** Specify the failover workload’s hostname or IP address and provide administrator-level credentials. Use the required credential format (see [“Guidelines for Workload and Container Credentials” on page 86](#)).
 - ♦ **Failback Target Settings:** Specify the following parameters:
 - ♦ **Replication Method:** Select the scope of data replication. See [“Initial Replication Method \(Full and Incremental\)” on page 91](#).
 - ♦ **Target Type:** Select the **Physical Target** option and then select the physical machine you registered in [Step 1](#).
- 5 Click **Save and Prepare** and monitor the progress on the Command Details screen.
Upon successful completion, PlateSpin 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.

5.8.3 Semi-Automated Failback to a Virtual Machine

This failback type follows a process similar to the [Semi-Automated Failback to a Physical Machine](#) for a VM target other than a natively-supported VMware container. During this process, you direct the system to regard a VM target as a physical machine.

You can do a semi-automated failback to a container, for which there is fully-automated failback support (VMware ESX and DRS Cluster targets).

You can also do a semi-automated failback for target VM platforms on Microsoft Hyper-V Server 2012 hosts.

To start the Hyper-V VMs on failover:

- 1 In a text editor, modify each Hyper-V host's `/etc/vmware/config` file by adding the following line:

```
vhv.allow = "TRUE"
```

- 2 In the vSphere Web Client, modify the failover VM Settings for the CPU:
 - 2a Under the **Virtual Hardware** tab, select **CPU**.
 - 2b In **Hardware virtualization**, select **Expose hardware assisted virtualization to guest OS**.
- 3 In the vSphere Web Client, modify the failover VM Settings for the CPU ID:
 - 3a Under the **VM Options** tab, expand **Advanced**, then select **Edit configuration parameters**.
 - 3b Verify the following setting:

```
hypervisor.cpuid.v0 = FALSE
```

5.9 Reprotecting a Workload

A **Reprotect** operation, the next logical step after a **Failback**, completes the workload protection lifecycle and starts it anew. Following a successful Failback operation, a **Reprotect** command becomes available in the PlateSpin Forge Web Interface, and the system applies the same protection details as those indicated during the initial configuration of the protection contract.

NOTE: The **Reprotect** command becomes available only if you selected the **Reprotect** option in the Failback details. See [“Failback” on page 79](#).

The rest of the workflow covering the protection lifecycle is the same as that in normal workload protection operations; you can repeat it as many times as required.

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 85](#)
- ♦ [Section 6.2, “Guidelines for Workload and Container Credentials,” on page 86](#)
- ♦ [Section 6.3, “Data Transfer,” on page 86](#)
- ♦ [Section 6.4, “Protection Tiers,” on page 89](#)
- ♦ [Section 6.5, “Recovery Points,” on page 90](#)
- ♦ [Section 6.6, “Initial Replication Method \(Full and Incremental\),” on page 91](#)
- ♦ [Section 6.7, “Service and Daemon Control,” on page 92](#)
- ♦ [Section 6.8, “Using Freeze and Thaw Scripts for Every Replication \(Linux\),” on page 92](#)
- ♦ [Section 6.9, “Volumes Storage,” on page 93](#)
- ♦ [Section 6.10, “Networking,” on page 96](#)
- ♦ [Section 6.11, “Failback to Physical Machines,” on page 96](#)
- ♦ [Section 6.12, “Protecting Windows Clusters,” on page 99](#)

6.1 Workload License Consumption

Your PlateSpin Forge product license entitles you to a specific or unlimited number of workloads for protection through workload licensing. Every time you add a workload for protection, the system consumes a single workload license from your license pool. On the Dashboard page of the PlateSpin Forge Web Interface, the License Summary displays the number installed licenses, and the current number of consumed licenses. 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 [“Activating Your Product License” on page 22](#).

6.2 Guidelines for Workload and Container Credentials

PlateSpin Forge must have administrator-level access to workloads and appropriate role configuration for containers. 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 user name, 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 user name and password	Non-root accounts must be properly configured to use <code>sudo</code> . See Knowledgebase Article 7920711 (https://www.netiq.com/support/kb/doc.php?id=7920711).
VMware ESX/ESXi 4.1; ESXi 5.0, ESXi 5.1, ESXi 5.5	VMware account with an appropriate role configuration.	If ESX is configured for Windows domain authentication, you can also use your Windows domain credentials.
VMware vCenter Server	VMware account with an appropriate role configuration.	

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 87](#)
- ♦ [Section 6.3.2, “Data Encryption,” on page 88](#)
- ♦ [Section 6.3.3, “Modifying the Location of the Volume Snapshots Directory for Windows Workloads,” on page 88](#)
- ♦ [Section 6.3.4, “Excluding or Including Files in Block Transfers for Incremental Replications,” on page 88](#)

6.3.1 Transfer Methods

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

- ♦ [“Transfer Methods Supported for Windows Workloads” on page 87](#)
- ♦ [“Transfer Methods Supported for Linux Workloads” on page 87](#)

Transfer Methods Supported for Windows Workloads

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

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

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

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

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

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 [Knowledgebase Article 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 [Knowledgebase Article 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

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

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

You can enable or disable encryption individually for each workload by selecting the **Encrypt Data Transfer** option. See “[Workload Protection Details](#)” on page 73.

6.3.3 Modifying the Location of the Volume Snapshots Directory for Windows Workloads

The PlateSpin Server saves volume snapshots in the following directory by default:

```
\ProgramData\PlateSpin\Volume Snapshots
```

You might need to modify the path:

- ♦ If the current drive for the path does not have sufficient space available for the Windows workload snapshots
- ♦ If you want to move the location in order to more easily exclude the path from your backup list

You can use the PlateSpin Server global parameter `VssSnapshotMountPath` on the PlateSpin Server Configuration page to specify a custom path on the server where you want to store the snapshots. If the parameter's value is empty, the path will remain at the default.

To specify a custom path for the volume snapshot directory on Windows:

- 1 Go to the PlateSpin Server Configuration page at <https://<platespin-server-ip-address>/PlatespinConfiguration>
- 2 Search for `VssSnapshotMountPath`, then click **Edit**.
- 3 In the **Value** field, specify the full path of the directory on the PlateSpin Server where you want to store volume snapshots for Windows workloads. For example:

```
G:\PlateSpin\Volume Snapshots
```
- 4 Click **Save**.

6.3.4 Excluding or Including Files in Block Transfers for Incremental Replications

PlateSpin Forge excludes and includes some files by default in the data transfer for a block-based incremental replication. The Block Based Volume Server Exclude and Include lists can accept new files in the lists, in addition to the default files. You can add a new list by using the following global parameters on the Platespin Server Configuration page:

```
BlockBasedTransferExcludeFileList  
BlockBasedTransferIncludeFileList
```


To specify a file that must always be transferred during incremental, block-based replications:

- 1 Go to the PlateSpin Server Configuration page at `https://<platespin-server-ip-address>/PlatespinConfiguration`
- 2 Search for `BlockBasedTransferIncludeFileList`, then click **Edit**.
- 3 In the **Value** field, add the file name to the list.
- 4 Click **Save**.

To specify a file that must always be excluded during incremental, block-based replications:

- 1 Go to the PlateSpin Server Configuration page at `https://<platespin-server-ip-address>/PlatespinConfiguration`
- 2 Search for `BlockBasedTransferExcludeFileList`, then click **Edit**.
- 3 In the **Value** field, add the file name to the list.
- 4 Click **Save**.

6.4 Protection Tiers

A Protection Tier is a custom collection of workload protection parameters that define the following:

- ♦ The frequency and recurrence pattern of replications
- ♦ Whether to encrypt data transmission
- ♦ Whether and how to apply data compression
- ♦ Whether to throttle available bandwidth to a specified throughput rate during data transfer
- ♦ Criteria for the system to consider a workload as offline (failed)

A Protection Tier is an integral part of every workload protection contract. During the configuration stage of a workload protection contract, you can select one of several built-in Protection Tiers and customize its attributes as required by that specific protection contract.

To create custom Protection Tiers in advance:

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

Parameter	Action
Name	Type the name you want to use for the tier.
Incremental Recurrence	Specify the frequency of incremental replications and the incremental recurrence pattern. You can type directly in the Start of recurrence field, or click the calendar icon to select a date. Select None as the Recurrence Pattern to never use incremental replication.
Full Recurrence	Specify the frequency of full replications and the full recurrence pattern.

Parameter	Action
Blackout Window	<p>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).</p> <p>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.</p> <p>NOTE: The blackout start and end times are based on the system clock on your PlateSpin Server.</p>
Compression Level	<p>These settings control whether and how workload data is compressed before transmission. See “Data Compression” on page 18.</p> <p>Select one of the available options. Fast consumes the least CPU resources on the source but yields a lower compression ratio, Maximum consumes the most, but yields a higher compression ratio. Optimal, the middle ground, is the recommended option.</p>
Bandwidth Throttling	<p>These settings control bandwidth throttling. See “Bandwidth Throttling” on page 18.</p> <p>To throttle replications to a specified rate, specify the required throughput value in Mbps and indicate the time pattern.</p>
Recovery Points to Keep	Specify the number of recovery points to keep for workloads that use this Protection Tier. See “Recovery Points” on page 90 .
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 [“Managing Snapshots of the Forge VM on the Appliance Host” on page 49](#).

6.6 Initial Replication Method (Full and Incremental)

The *initial replication* is the creation of an initial base copy of a production workload to the failover workload (virtual replica) in a protection operation, or from a failover workload to its original virtual or physical infrastructure in preparation for a failback operation for the production workload.

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

- ♦ **Full:** A full workload transfer takes place based on all of its data.
- ♦ **Incremental:** Only differences are transferred from a source to its target, provided that they have similar operating system and volume profiles.
 - ♦ **During protection:** The production workload is compared with an existing VM in the 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 in the appliance host, such as a workload VM physically moved on portable media from the production site to a remote recovery site.See [“Manually Importing VMs into the Appliance Host's Datastore” on page 50](#).
 - ♦ **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 [“Semi-Automated Failback to a Physical Machine” on page 82](#)).

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.

Name	Description	CPU	Memory	Free Space	Last Refresh	Remove
comp212	VMware ESX Server 4.0.0.175025	10 x Intel(R) Xeon(R) CPU E5530 @ 2.40GHz	31.5 GB	1.9 TB	2 Days Ago	Remove

Container: comp212 (VMware ESX Server 4.0.0.175025)

Virtual Machine: 1SLES10-P1.site_VM (SuSE Linux)

Inventory Network: VM Network

☒ DHCP ☐ Static

- 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 [“Using Freeze and Thaw Scripts for Every Replication \(Linux\)” on page 92](#).

- ♦ **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 92](#)). For example, you might want to use this feature to temporarily freeze certain daemons instead of shutting them down during replications.

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

- 1 Create the following files:
 - ♦ `platespin.freeze.sh`: A shell script to execute at the beginning of the replication
 - ♦ `platespin.thaw.sh`: A shell script to execute at the end of the replication

- ♦ `platespin.conf`: A text file defining any required arguments, along with a timeout value.

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

```
[ServiceControl]

FreezeArguments=<arguments>

ThawArguments=<arguments>

TimeOut=<timeout>
```

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

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

`/etc/platespin`

6.9 Volumes Storage

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 that you use to specify the volumes you require for protection. For more information, see [Section 1.1.5, "Supported Storage," on page 13](#).

[Figure 6-1](#) shows 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

DashboardWorkloadsTasksReportsSettingsAboutHelp

Edit Protection Details : vses11sp3x64.example.com

Change Container

Save & Prepare

Save

Cancel

Tier Settings

Replication Settings

Transfer Encryption:

Source Credentials:

CPU:

Replication Network:

Allowed Networks:

Resource Pool for Target VM:

VM Folder for Target VM:

Configuration File Datastore:

Protected Volumes:

Protected Logical Volumes:

Non-volume Storage:

Volume Groups:

Daemons to Stop During Replication:

☐ Encrypt Data Transfer

User Name:
root
Password:

[Test Credentials](#)

Sockets
3
Cores Per Socket
3
Total Cores
9

VM Network - 10.10.18x
☒ DHCP ☐ Static MTU:

Allow	Name	Address	Uses DHCP
<input checked="" type="checkbox"/>	eth0	10.10.187.153	False

cluster60
[Edit](#)

dc60
[Edit](#)

VOL1-HPSAN-STORAGE (366.5 GB free)

Include	Name	Used Space	Free Space	Datastore	Thin Disk
<input checked="" type="checkbox"/>	/ (EXT3 - System)	5.0 GB	8.73 GB	VOL1-HPSAN-STOI	<input type="checkbox"/>
<input type="checkbox"/>	/opt/hovel/nas/lnmt/pools/POOL1 (NSSFS)	66.9 MB	11.93 GB	VOL1-HPSAN-STOI	<input type="checkbox"/>

Include	Name	Used Space	Free Space	Volume Group / OES Volume
<input checked="" type="checkbox"/>	/var/mtst1 (EXT3)	84.5 MB	923.4 MB	VolGroup1
<input checked="" type="checkbox"/>	/var/mtst2 (EXT3)	169.5 MB	1.8 GB	VolGroup1

Include	Partition	Is Swap	Total Size	Datastore	Thin Disk
<input checked="" type="checkbox"/>	/dev/sda1	Yes	2.01 GB	BBCSLESSAN (3.8	<input type="checkbox"/>

Include	Name	Total Size	Datastore	Thin Disk
<input checked="" type="checkbox"/>	VolGroup1	8.0 GB	BBCSLESSAN (3.8	<input type="checkbox"/>

[Add Daemons](#)

Failover Settings

Prepare for Failover Settings

Test Failover Settings

Tag

Figure 6-2 shows volume protection options of an OES 11 workload with options indicating that the LVM2 volume and NSS pool layout should be preserved and re-created for the failover workload:

Figure 6-2 Replication Settings, Volume-Related Options (OES 11 Workload)

Protected Volumes:	Include	Name	Total Size	Datastore	Thin Disk	
	<input checked="" type="checkbox"/>	/ (EXT3 - System)	13.8 GB	BBCSLESSAN	<input type="checkbox"/>	
Protected Logical Volumes:	Include	Name	Total Size	Volume Group		
	<input checked="" type="checkbox"/>	/vmtest1 (EXT3)	1007.9 MB	VolGroup1		
	<input checked="" type="checkbox"/>	/vmtest2 (EXT3)	2.0 GB	VolGroup1		
	<input checked="" type="checkbox"/>	/opt/novell/nss/mnt/pools /POOL1 (NSSFS)	12.0 GB	POOL1		
Non-volume Storage:	Include	Partition	Is Swap	Total Size	Datastore	Thin Disk
	<input checked="" type="checkbox"/>	/dev/sda1	Yes	2.0 GB	BBCSLESSAN	<input type="checkbox"/>
Volume Groups:	Include	Name	Total Size	Datastore	Thin Disk	
	<input checked="" type="checkbox"/>	VolGroup1	8.0 GB	BBCSLESSAN	<input type="checkbox"/>	
OES Volumes:	Include	Name	Total Size	Datastore	Thin Disk	
	<input checked="" type="checkbox"/>	POOL1	12.0 GB	BBCSLESSAN	<input type="checkbox"/>	
Daemons to Stop During Replication:	--					

Figure 6-3 shows volume protection options of an OES 2 workload with options indicating that the EVMS and NSS pool layout should be preserved and re-created for the failover workload:

Figure 6-3 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"/>	/ (RISDRFS)	2.2 GB	2.2 GB	system	
	<input checked="" type="checkbox"/>	/boot (EXT2)	13.0 MB	55.3 MB	/dev/evms/sda1	
	<input checked="" type="checkbox"/>	/opt/novell/nss/mnt/pools/NEWPOOL (NSSFS)	23.3 MB	999.6 MB	NEWPOOL	
Non-volume Storage:	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		70.6 MB	dev-comp124:storage	<input type="checkbox"/>
	<input checked="" type="checkbox"/>	NEWPOOL		1023.0 MB	dev-comp124:storage	<input type="checkbox"/>
Daemons to Stop During Replication:		Add Daemons				

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

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 Image," on page 96](#)
- ♦ [Section 6.11.2, "Injecting Additional Device Drivers into the Boot ISO Image," on page 96](#)
- ♦ [Section 6.11.3, "Registering Physical Machines as Failback Targets with PlateSpin Forge," on page 98](#)

6.11.1 Downloading the PlateSpin Boot ISO Image

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

- ♦ **Product or Technology:** PlateSpin Forge
- ♦ **Select Version:** PlateSpin Forge 11.2
- ♦ **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: Ensure that the drivers are valid for the kernel included with the ISO file (for x86 systems: 3.0.93-0.8-pae, for x64 systems: 3.0.93-0.8-default) and are appropriate for the target architecture. See also [Knowledgebase Article 7005990 \(https://www.netiq.com/support/kb/doc.php?id=7005990\)](https://www.netiq.com/support/kb/doc.php?id=7005990).

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

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

- 3 Copy the `rebuildiso.sh` script, located in the `/tools` subdirectory of the mounted ISO file, into a temporary working directory. When you have finished, unmount the ISO file (execute the command `umount <mount_point>`).

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

- 5 In the directory where you saved the `rebuildiso.sh` script, run the `rebuildiso.sh` script as root, using the following syntax:

```
./rebuildiso.sh <ARGS> [-v] -m32|-m64 -i <ISO_file>
```

The following table lists the possible command line options for this command:

Option	Description
-i <ISO_file>	<ISO_file> is the ISO to modify, list, etc.
-v	If used together with the -l argument, the option causes the use of modinfo to obtain verbose driver information.
-o	If used together with the -c argument or the -d argument, the old copy of the ISO file is not overwritten.
-m32	Specifies 32-bit initrd injection.
-m64	Specifies 64-bit initrd injection.

The next table lists the possible arguments for use with this command. At least one of these arguments must be used in the command:

Argument	Description
-d <path>	<path> specifies the directory that contains the drivers (that is, *.ko files) that you want to inject. On completion of the command, the ISO file is updated with the added drivers.
-c <path>	<path> specifies where a <code>ConfigureTakeControl.xml</code> file resides.

Argument	Description
-l [<i><type></i>]	<p><i><type></i> specifies a subset of drivers you want to list. The default is “all” types.</p> <p>Listed driver types beginning with a forward slash (/) are assumed to be located in <i><kernel_module_directory>/kernel/</i></p> <p>Listed driver types without a leading forward slash (/) are assumed to be located in <i><kernel_module_directory>/kernel/drivers/</i></p> <p>Driver Subset Examples:</p> <pre>-l scsi -l 'net video' -l '/net net'</pre> <p>Special Usage of this Argument:</p> <p>If you want to list the available subdirectories of each of the subsets, use the argument like this: -l INDEX</p>

Syntax Examples

- ♦ To list an index of 32-bit drivers:

```
# ./rebuildiso.sh -i bootofx.x2p.iso -m32 -l INDEX
```
- ♦ To list drivers found in the /misc folder:

```
# ./rebuildiso.sh -i bootofx.x2p.iso -m32 -l misc
```
- ♦ To inject 32-bit drivers from the /oem-drivers folder:

```
# ./rebuildiso.sh -i bootofx.x2p.iso -m32 -d oem-drivers
```
- ♦ To inject 64-bit drivers from an /oem-drivers folder and also inject a customized ConfigureTakeControl.xml file:

```
# ./rebuildiso.sh -i bootofx.x2p.iso -m64 -c ConfigureTakeControl.xml -d oem-drivers
```

6.11.3 Registering Physical Machines as Failback Targets with PlateSpin Forge

- 1 Burn the PlateSpin boot ISO image on a CD or save it to media from which your target can boot.
- 2 Ensure that the network switch port connected to the target is set to **Auto Full Duplex**.
- 3 Use the boot CD to boot the target physical machine, then wait for the command prompt window to open.
- 4 (Linux only) For 64-bit systems, at the initial boot prompt, type the following:

```
ps64
```
- 5 Press Enter.
- 6 When you are prompted, enter the 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 (yes) if you have enabled SSL, and N (no) if you have not.

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

6.12 Protecting Windows Clusters

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

- ♦ **Windows Server 2012 R2:** Server-based Microsoft Failover Cluster (*Node and Disk Majority Quorum* and *No Majority: Disk Only Quorum* models)
- ♦ **Windows Server 2008 R2:** Server-based Microsoft Failover Cluster (*Node and Disk Majority Quorum* and *No Majority: Disk Only Quorum* models)
- ♦ **Windows Server 2003 R2:** Server-based Windows Cluster Server (*Single-Quorum Device Cluster* model)

You can enable or disable Windows cluster discovery for your PlateSpin environment. See [Section 6.12.2, "Enabling or Disabling Windows Cluster Discovery," on page 101](#).

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

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

- ♦ [Section 6.12.1, "Planning Your Cluster Workload Protection," on page 99](#)
- ♦ [Section 6.12.2, "Enabling or Disabling Windows Cluster Discovery," on page 101](#)
- ♦ [Section 6.12.3, "Adding Resource Name Search Values," on page 102](#)
- ♦ [Section 6.12.4, "Quorum Arbitration Timeout," on page 103](#)
- ♦ [Section 6.12.5, "Setting Local Volume Serial Numbers," on page 103](#)
- ♦ [Section 6.12.6, "PlateSpin Failover," on page 103](#)
- ♦ [Section 6.12.7, "PlateSpin Failback," on page 103](#)

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

- ♦ ["Prerequisites" on page 100](#)
- ♦ ["Block-Based Transfer" on page 100](#)

- ♦ [“Cluster Node Failover during the First Full Replication” on page 100](#)
- ♦ [“Cluster Node Failover during Replication” on page 101](#)
- ♦ [“Cluster Node Failover between Replications” on page 101](#)
- ♦ [“Protection Setup” on page 101](#)

Prerequisites

The scope of support for cluster protection is subject to the following conditions:

- ♦ **Active node hostname or IP address:** You must specify the hostname or IP address of the cluster’s active node when you perform an Add Workload operation. Because of security changes made by Microsoft, Windows clusters can no longer be discovered by using the virtual cluster name (that is, the shared cluster IP address).
- ♦ **Active node discovery:** Ensure that the PlateSpin global configuration setting `DiscoverActiveNodeAsWindowsCluster` is set to `True` on the PlateSpin Server Configuration page. This is the default setting. See [Section 6.12.2, “Enabling or Disabling Windows Cluster Discovery,” on page 101](#).
- ♦ **Resource name search values:** You must specify search values to use that can help PlateSpin Forge differentiate the name of the shared Cluster IP Address resource from the name of other IP address resources on the cluster. See [Section 6.12.3, “Adding Resource Name Search Values,” on page 102](#).
- ♦ **Resolvable hostname:** The PlateSpin Server must be able to resolve the hostname of each of the nodes in the cluster.

NOTE: The hostname must be resolvable by the IP address. That is, both hostname lookup and reverse lookup are required.

- ♦ **Quorum resource:** A cluster’s quorum resource must be co-located on the node with the cluster’s resource group (service) being protected.
- ♦ **PowerShell 2.0:** Windows PowerShell 2.0 Engine must be installed on each node of the cluster.

Block-Based Transfer

When you use block-based transfer for cluster workloads, the block-based driver components are not installed on the cluster nodes. The block-based transfer occurs using a driverless synchronization 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.

Cluster Node Failover during the First Full Replication

A cluster workload requires that the first full replication completes successfully without a cluster node failover. If a cluster node failover occurs prior to the completion of the first full replication, you must remove the existing workload, re-add the cluster using the active node, and try again.

Cluster Node Failover during Replication

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

Cluster Node Failover between Replications

The nodes must have similar profiles to prevent interruptions in the replication process. If a cluster node failover occurs between the 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 for the next incremental replication. Otherwise, the next incremental replication command fails.

The profiles of cluster nodes are considered similar if all of the following conditions are met:

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

NOTE: Use the customized *Volume Manager* utility to change the local volume serial numbers to match each node of the cluster. See [“Synchronizing Serial Numbers on Cluster Node Local Storage” on page 147](#).

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

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

Protection Setup

To configure protection for a Windows cluster, follow the normal workload protection workflow. Ensure that you provide the hostname or IP address of the cluster's active node. See [“Basic Workflow for Workload Protection and Recovery” on page 69](#).

6.12.2 Enabling or Disabling Windows Cluster Discovery

The PlateSpin Forge Server can discover and inventory Windows Server failover clusters in your PlateSpin environment based on the active node in each cluster. Alternatively, it can treat all active and non-active cluster nodes as standalone machines.

To enable cluster discovery for all Windows clusters, ensure that the parameter `DiscoverActiveNodeAsWindowsCluster` is set to `True`. This is the default setting. Cluster discovery, inventory, and workload protection use the hostname or IP address of a cluster's active node, instead of using its cluster name and an administration share. You do not configure separate workloads for the cluster's non-active nodes. For other cluster workload protection requirements, see [“Prerequisites” on page 100](#).

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

To enable or disable cluster discovery:

- 1 Go to the PlateSpin Server configuration page at `https://<platespin-server-ip-address>/PlatespinConfiguration`
- 2 Search for `DiscoverActiveNodeAsWindowsCluster`, then click **Edit**.
- 3 In the **Value** field, select **True** to enable cluster discovery, or select **False** to disable cluster discovery.
- 4 Click **Save**.

6.12.3 Adding Resource Name Search Values

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

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

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

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

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

To add search values in the `MicrosoftClusterIPAddressNames` list:

- 1 Go to the PlateSpin Server configuration page at `https://<platespin-server-ip-address>/PlatespinConfiguration`
- 2 Search for `MicrosoftClusterIPAddressNames`, then click **Edit**.
- 3 In the **Value** field, add one or more search values to the list.
- 4 Click **Save**.

6.12.4 Quorum Arbitration Timeout

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

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

- 1 Go to the PlateSpin Server configuration page at <https://<platespin-server-ip-address>/PlatespinConfiguration>
- 2 Search for `FailoverQuorumArbitrationTimeout`, then click **Edit**.
- 3 In the **Value** field, specify the maximum number of seconds to allow for quorum arbitration.
- 4 Click **Save**.

6.12.5 Setting Local Volume Serial Numbers

You can use the *Volume Manager* utility to change the local volume serial numbers to match in each node of the cluster. See [“Synchronizing Serial Numbers on Cluster Node Local Storage” on page 147](#).

6.12.6 PlateSpin Failover

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

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

6.12.7 PlateSpin Failback

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

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

- ♦ Map all disks on the PlateSpin virtual one-node cluster to a single local disk on the failback target.
- ♦ Add another disk (`Disk 2`) to the physical failback machine. You can then configure the PlateSpin failback operation to restore the failover'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 a PlateSpin failback is complete, you must reattach the shared storage and rebuild the cluster environment before you can rejoin additional nodes to the newly restored cluster.

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

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

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

7 Auxiliary Tools for Working with Physical Machines

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

- ♦ [Section 7.1, “Managing Device Drivers,” on page 105](#)

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 Forge driver database.

The following sections provide more details:

- ♦ [Section 7.1.1, “Packaging Device Drivers for Windows Systems,” on page 105](#)
- ♦ [Section 7.1.2, “Packaging Device Drivers for Linux Systems,” on page 106](#)
- ♦ [Section 7.1.3, “Uploading Drivers to the PlateSpin Device Driver Database,” on page 106](#)
- ♦ [Section 7.1.4, “Using the Plug and Play \(PnP\) ID Translator Feature,” on page 108](#)

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 [“Uploading Drivers to the PlateSpin Device Driver Database” on page 106](#).

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

- ♦ All 64-bit Windows systems
 - ♦ 32-bit versions of Windows Server 2008 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 your PlateSpin boot ISO image.

- 1 On a Linux workstation, create a directory for your device driver files. All the drivers in the directory must be for the same kernel and architecture.
- 2 Download the boot image and mount it.

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

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

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

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

```
tar -xvzf packageModules.tar.gz
```

- 4 Enter the working directory and execute the following command:

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

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

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

For example, `bnx2x-1.48.107-RHEL4-2.6.9-11.EL-i686.pkg`

The package is now ready for uploading. See [“Uploading Drivers to the PlateSpin Device Driver Database” on page 106](#).

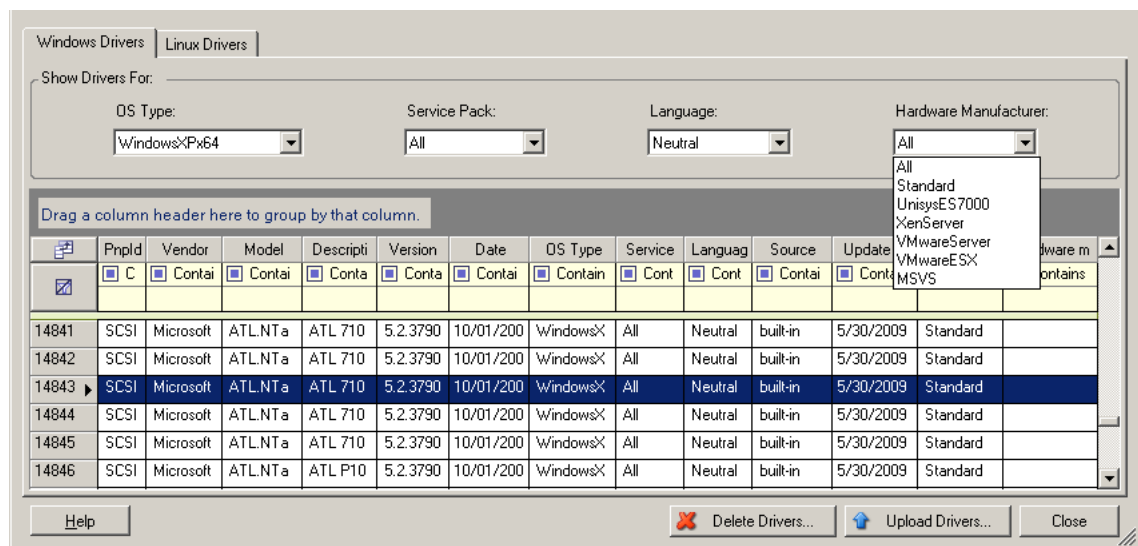
7.1.3 Uploading Drivers to the PlateSpin Device Driver Database

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

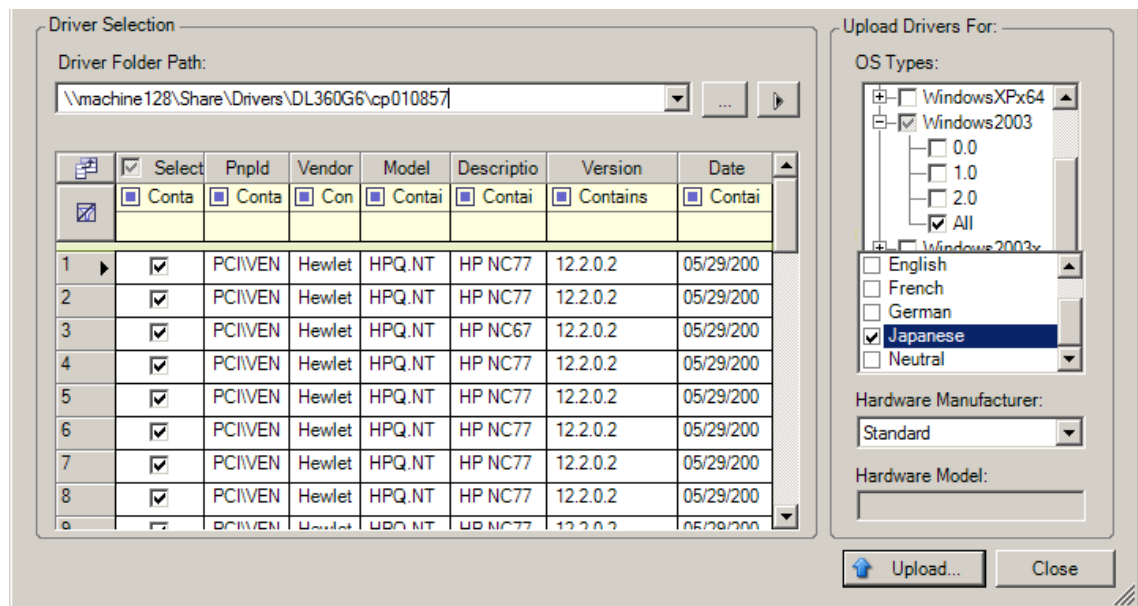
NOTE: On upload, PlateSpin Forge does not validate drivers against selected operating system types or their bit specifications; ensure that you only upload drivers that are appropriate for your target infrastructure.

Device Driver Upload Procedure (Windows)

- 1 Obtain and prepare the required device drivers. See [Packaging Device Drivers for Windows Systems](#).
- 2 On your 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.

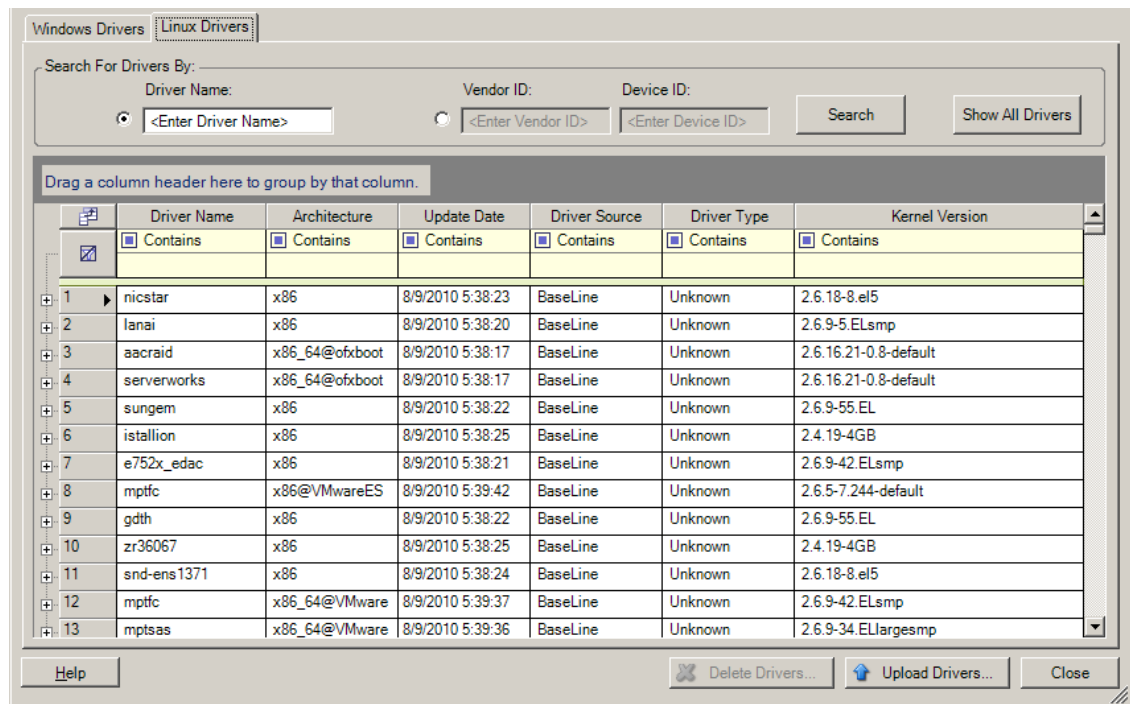


Select **Standard** as the **Hardware Manufacturer** option, unless your drivers are designed specifically for any of the target environments listed.

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

Device Driver Upload Procedure (Linux)

- Obtain and prepare the required device drivers. See [Packaging Device Drivers for Linux Systems](#).
- 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 Forge Linux RAM disk contains Linux-style PnP IDs.


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

Follow these steps to use the PnP ID Translation feature:


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




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



- 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.



- 4b2** On the dialog, 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.



- 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.


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

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


or

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

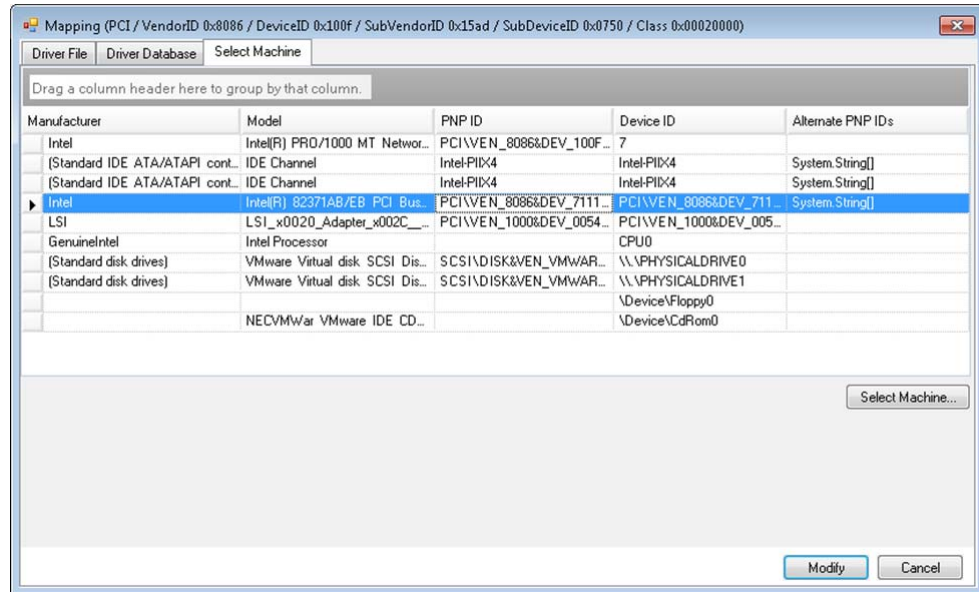
- ◆ Under the **Driver File** tab, browse to and select a Windows driver file (that is, a file with the *.inf extension), select the desired PnP ID, then click **Modify**.



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




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



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

- In the Create PnP Id Mapping dialog, 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).

To modify,

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

7b Repeat [Step 5 on page 111](#) to modify the Windows PnP ID.

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

8 ProtectAgent Utility

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

The syntax of the ProtectAgent utility is:

```
ProtectAgent.cli.exe [Option] [/psserver=%IP%]
```

Table 8-1 describes the options and switch available for the `ProtectAgent.cli.exe` command.

Table 8-1 ProtectAgent Command Options and Switch

Usage	Description
Options	
<code>h ? help</code>	Displays usage and options for the command.
<code>logs view-logs</code>	Opens the application log directory.
<code>status</code>	Shows installation status for the PlateSpin controller and drivers.
<code>din driver-install</code>	Installs the PlateSpin drivers.
<code>dup driver-upgrade</code>	Upgrades the PlateSpin drivers.
<code>dun driver-uninstall</code>	Uninstalls the PlateSpin drivers.
Switch	
<code>/psserver=%IP%</code>	Downloads the block-based transfer drivers from the specified server when you invoke the <code>status</code> , <code>driver-install</code> , or <code>driver-upgrade</code> options.

A copy of the block-based transfer drivers is bundled with the ProtectAgent utility. You can alternatively specify the `/psserver=` command line switch in order to download the drivers from the PlateSpin Server when you invoke the `status`, `driver-install`, or `driver-upgrade` options. This is useful when the server is patched with a new driver package, but the ProtectAgent command line utility is not patched.

NOTE: To avoid confusion, the recommended method of using the ProtectAgent is to install, uninstall, or upgrade the drivers and then reboot prior to doing a replication.

You should reboot the system each time that you install, upgrade, or uninstall the drivers. The reboot forces the running driver to stop and the new driver to be applied on system restart. If you do not reboot the system prior to replication, the source continues to act as if the operation has not been completed. For example, if you install drivers without rebooting the system, the source acts as if no driver is installed during replication. Similarly, if you upgrade the drivers without rebooting, the source continues to use the running driver during replication until you reboot the system.

If the version of the installed driver is different than the version of the running driver, the status option will remind the user to reboot. For example:

```
C:\ProtectAgent\ProtectAgent.cli.exe status
Step 1 of 2: Querying the PlateSpin controller service
Done
Step 2 of 2: Querying the installed PlateSpin driver version
Done

The task completed successfully
PlateSpin Controller Service Status
Status: Running
Version: 9.9.9.9
Last Successful Contact: 1/5/2015 12:14:25 PM

PlateSpin Driver Status
Installed Driver Version: 8.0.0.11
Running Driver Version: Not running. Reboot to load the driver.
Upgrade Available: No
```

PlateSpin creates a task to warn the user that a reboot is necessary in order to complete the driver installation or upgrade. The notification appears in the Tasks list (Figure 8-1). During replication, the notification appears on the Command Details page (Figure 8-2).

Figure 8-1 Reboot Notification Task


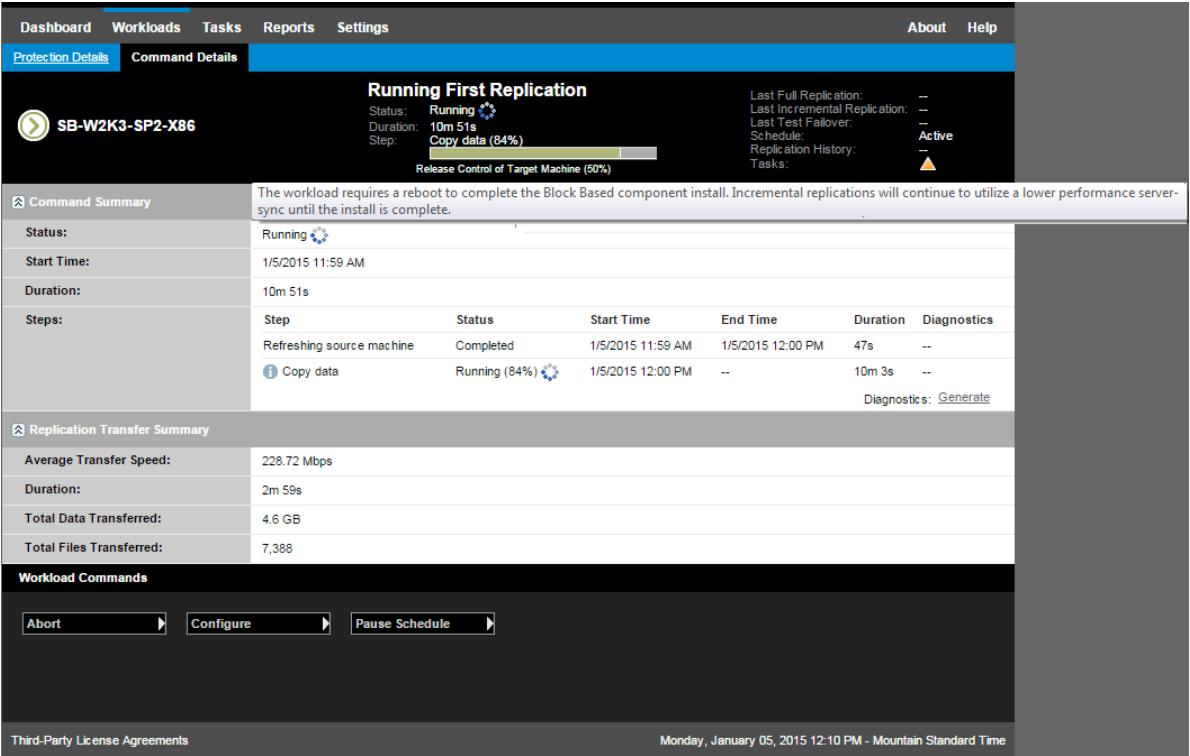


Figure 8-2 Reboot Notification During Replication



Rebooting the source machine applies and starts the installed or upgraded drivers. If the driver was recently installed, after the reboot, one full replication or a server-sync replication is required in order to ensure that all of a source's changes are captured. This server-sync replication will be represented to the user in the Status field as a warning (Figure 8-3). Subsequent incremental replications will complete without warning.

Figure 8-3 Server-Sync Required Notification

SB-W2K3-SP2-X86

Running Incremental

Status: Running

Duration: 7m 38s

Step: Copy data (8%)

Copying Volume Data from Source to Target (5%)

Last Full Replication: 1/5/2015 12:11 PM

Last Incremental Replication: 1/5/2015 12:29 PM

Last Test Failover: --

Schedule: Active

Replication History: View

Tasks: --

Command Summary

Events:

Event	Details	User	Date
Incremental replication started		scb-pelennor\Scott	1/5/2015 12:37 PM

Status:

Running

The Block Based component has recently completed the install process. This replication requires a server-sync to be performed.

Start Time:

1/5/2015 12:37 PM

Duration:

7m 38s

Steps:

Step	Status	Start Time	End Time	Duration	Diagnostics
Refreshing source machine	Completed	1/5/2015 12:37 PM	1/5/2015 12:38 PM	51s	--
Revert to snapshot	Completed	1/5/2015 12:38 PM	1/5/2015 12:38 PM	30s	--
Copy data	Running (8%)	1/5/2015 12:38 PM	--	6m 17s	--

Diagnostics: [Generate](#)

Replication Transfer Summary

Average Transfer Speed:

1.51 Mbps

Duration:

37s

Total Data Transferred:

6.2 MB

Total Files Transferred:

103

Workload Commands

Abort

Configure

Pause Schedule

Third-Party License Agreements

Monday, January 05, 2015 12:45 PM - Mountain Standard Time

9 Troubleshooting

This section includes the following information:

- ♦ [Section 9.1, “Troubleshooting Workload Inventory \(Windows\),” on page 119](#)
- ♦ [Section 9.2, “Troubleshooting Workload Inventory \(Linux\),” on page 123](#)
- ♦ [Section 9.3, “Troubleshooting Problems during the Prepare Replication Command \(Windows\),” on page 123](#)
- ♦ [Section 9.4, “Troubleshooting Workload Replication,” on page 124](#)
- ♦ [Section 9.5, “Troubleshooting the Configuration Service,” on page 125](#)
- ♦ [Section 9.6, “Troubleshooting Traffic-forwarding Workloads,” on page 130](#)
- ♦ [Section 9.7, “Troubleshooting Online Help,” on page 130](#)
- ♦ [Section 9.8, “Generating and Viewing Diagnostic Reports,” on page 131](#)
- ♦ [Section 9.9, “Removing Workloads,” on page 131](#)
- ♦ [Section 9.10, “Post-Protection Workload Cleanup,” on page 131](#)
- ♦ [Section 9.11, “Shrinking the PlateSpin Forge Databases,” on page 133](#)
- ♦ [Section 9.12, “Active Directory Domain Services Are Not Available After a Failback \(Windows\),” on page 134](#)

9.1 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 121 again. If the test succeeds, try adding the workload again.</p> <ul style="list-style-type: none">♦ “Troubleshooting DCOM Connectivity” on page 121♦ “Troubleshooting RPC Service Connectivity” on page 121

Problems or Messages	Solutions
Unable to connect to Windows server...The network path was not found	Network connectivity failure. Perform the tests in “Performing Connectivity Tests” on page 120 . If a test fails, ensure that PlateSpin Forge and the workload are on the same network. Reconfigure the network and try again.
"Discover Server Details {hostname}" Failed Progress: 0% Status: NotStarted	This error can occur for several reasons and each has a unique solution: <ul style="list-style-type: none"> ♦ For environments using a local proxy with authentication, bypass the proxy or add the proper permissions. See Knowledgebase Article 7920339 (https://www.netiq.com/support/kb/doc.php?id=7920339) for more details. ♦ If local or domain policies restrict required permissions, follow the steps outlined in Knowledgebase Article 7920862 (https://www.netiq.com/support/kb/doc.php?id=7920862).
Workload Discovery fails with error message	There are several possible reasons for the Could not find file output.xml error:
Could not find file output.xml	<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 122.
or	
Network path not found	<ul style="list-style-type: none"> ♦ File and Printer Sharing for Microsoft Networks might not be enabled. Enable it under the Network Interface Card properties.
or (upon attempting to discover a Windows cluster)	<ul style="list-style-type: none"> ♦ 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 122.
Inventory failed to discover. Inventory result returned nothing.	<ul style="list-style-type: none"> ♦ The Server or the Workstation service might not be running. If this is the case, enable them and set the startup mode to automatic. ♦ The Windows remote registry service is disabled. Start the service and set the startup type to automatic.

The following sections provide more troubleshooting information on Windows workloads:

- ♦ [Section 9.1.1, “Performing Connectivity Tests,” on page 120](#)
- ♦ [Section 9.1.2, “Disabling Antivirus Software,” on page 122](#)
- ♦ [Section 9.1.3, “Enabling File/Share Permissions and Access,” on page 122](#)

9.1.1 Performing Connectivity Tests

- ♦ [“Network Connectivity Test” on page 120](#)
- ♦ [“WMI Connectivity Test” on page 121](#)
- ♦ [“Troubleshooting DCOM Connectivity” on page 121](#)
- ♦ [“Troubleshooting RPC Service Connectivity” on page 121](#)

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 and Installing the vSphere Client Program” on page 48](#).

- 2 Open a command prompt and ping your workload:

```
ping workload_ip
```

WMI Connectivity Test

- 1 Go to your Forge VM.
See [“Downloading and Installing the vSphere Client Program” on page 48](#)“Downloading and Installing the vSphere Client Program” on page 48.
- 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 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 [“Configuring Access and Communication Settings across your Protection Network” on page 28.](#)

9.1.2 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.

9.1.3 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 [“Downloading and Installing the vSphere Client Program” on page 48.](#)
- 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.

9.2 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 "Configuring Access and Communication Settings across your Protection Network" on page 28.</p>
Access denied	<p>This authentication problem indicates either an invalid user name or password. For information on proper workload access credentials, see "Guidelines for Workload and Container Credentials" on page 86.</p>

9.3 Troubleshooting Problems during the Prepare Replication Command (Windows)

Problems or Messages	Solutions
Authentication error when verifying the controller connection while setting up the controller on the source.	<p>The account used to add a workload needs to be allowed by this policy. See "Group Policy and User Rights" on page 123.</p>
Failure to determine whether .NET Framework is installed (with exception The trust relationship between this workstation and the primary domain failed).	<p>Check whether the Remote Registry service on the source is enabled and started. See also "Troubleshooting Workload Inventory (Windows)" on page 119.</p>

9.3.1 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`.

9.4 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 <code>ip</code> 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 37.
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 Knowledgebase Article 7920339 (https://www.netiq.com/support/kb/doc.php?id=7920339). ♦ If local or domain policies restrict required permissions, follow the steps outlined in Knowledgebase Article 7920862 (https://www.netiq.com/support/kb/doc.php?id=7920862). <p>This is a common issue when Forge VM is affiliated with a domain and the domain policies are applied with restrictions. See “Group Policy and User Rights” on page 123.</p>
After a Windows Update, some files in the C:\Windows\SoftwareDistribution folder are not transferred to the target machine during incremental file-based replication.	<p>This is a Microsoft Windows common practice: For optimization purposes, some files are marked for deletion in the HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\BackupRestore\FilesNotToSnapshot registry key to prevent them from being included in VSS snapshots. See the Microsoft Developer Network article, Excluding Files from Shadow Copies (http://msdn.microsoft.com/en-us/library/aa819132.aspx) for more information.</p> <p>Generally, these files are used to install Windows updates before they are deleted and are no longer necessary after the update. If you choose to restore these files, run Windows Update on the target machine after failover to repopulate the SoftwareDistribution folder.</p>

9.5 Troubleshooting the Configuration Service

After Test Failover or Failover, an error occurs on the target VM because of non-specific Configuration Service issues. The common error message is:

Configuration service in the target machine does not seem to have started

Troubleshooting tips in this section explain common Configuration Service issues and some alternative ways to resolve them.

- ♦ [Section 9.5.1, “Understanding What Is Causing the Problem,” on page 126](#)
- ♦ [Section 9.5.2, “What Can Be Done to Resolve the Problem,” on page 126](#)
- ♦ [Section 9.5.3, “Additional Troubleshooting Tips,” on page 129](#)

9.5.1 Understanding What Is Causing the Problem

The Configuration Service error indicates that the PlateSpin Server is unable to communicate with the Configuration Service on the Target VM. Analyze your system to determine the possible root cause of the problem.

- ♦ [“Target VM Fails to Boot” on page 126](#)
- ♦ [“Network Is Not Set Up Correctly” on page 126](#)
- ♦ [“Unable to Read or Write Status Messages to Floppy Devices” on page 126](#)

Target VM Fails to Boot

The operating system must be loaded in the target VM in order for the Configuration Service to start up normally. A failure to boot indicates that there could be a driver conflict, a boot loader error, or possible disk corruption.

We recommend that you open a service ticket with Micro Focus Customer Care if the operating system fails to boot on the target VM.

Network Is Not Set Up Correctly

The network must be set up correctly in order for the Configuration Service on the target workload to communicate with the PlateSpin Server.

Ensure that you have configured your network in a way that the target workload can communicate with the PlateSpin Server.

Unable to Read or Write Status Messages to Floppy Devices

The Configuration Service must be able to communicate with the floppy devices for VMware VMs in order to read and write status messages for the PlateSpin Server.

On the target VM, verify that the machine is able to communicate with the floppy devices:

- 1 On the VM, open the log file (C:\windows\platespin\configuration\data\log.txt).
- 2 Any of the following messages might be an indication that the floppy is inaccessible:

```
Failed (5) to write to file \\?\Volume{<guid-number>}\log.zip

CopyFile \\?\Volume{<guid-number>}\windows\platespin\configuration\data\result.txt
to \\?\Volume{<guid-number>}\result.txt failed

The output floppy was not accessible after the timeout period
```

9.5.2 What Can Be Done to Resolve the Problem

To resolve a Configuration Service error, you can try any of the solutions in this section.

NOTE: Some settings in these solutions require Forge 11.2 with the latest patches applied.

- ♦ [“Skip the Target VM Reboot Optimizations” on page 127](#)
- ♦ [“Reduce the Read/Write Traffic to Floppy Devices” on page 127](#)

- “Change the Startup Type to Increase the Delay” on page 128
- “Configure Conflicting Services to Not Run Automatically at Startup” on page 129

Skip the Target VM Reboot Optimizations

Forge tries to minimize the number of reboots that occur on the target VM by default in order to speed up the Failover process. It is possible that allowing the additional reboots will improve the target VM’s ability to communicate with the PlateSpin Server.

To skip reboot optimizations:

- 1 Log in to the PlateSpin Server, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/platespinconfiguration/`
- 2 Search for the parameter **ConfigurationServiceValues**.
- 3 Edit the **ConfigurationServiceValues** parameter and set the **SkipRebootOptimization** option to true.
- 4 Click **Save**.
- 5 Run an incremental or full replication.
The replication also propagates the modified configuration settings to the target VM.
- 6 Run the Test Failover or Failover again for affected workloads.

Reduce the Read/Write Traffic to Floppy Devices

You can decrease the number of times the PlateSpin Server attempts to read from and write to the VMware input or output floppy devices if the diagnostic log shows the following error:

```
Information:1:Attempting floppy download
```

followed by

```
Verbose:1:Failed to copy file from remote URL
```

-or-

```
Exception: The remote server returned an error: (500) Internal Server Error
```

This error is caused by VMware locking the resource. It indicates that the PlateSpin Server is detaching and reattaching the floppy each time it checks the status. Locking can cause the target VM to fail to read and write to the floppy device. See [Using the VMware vCenter Server 4.x, 5.x and 6.0 Datastore Browser to Download or Copy a Powered-On Virtual Machine's .vmx and .nvram Files Fails \(1019286\)](https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1019286) (https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1019286).

If you experience floppy device locking issues, you can increase values for the Configuration Service polling settings on the PlateSpin Server:

vmwareConfigServicePollStartDelay

This parameter determines how long to wait before the PlateSpin Server starts polling for target workload status. The default value is 120 seconds (2 minutes).

vmwareConfigServicePollIntervalInMilliseconds

This parameter determines how frequently the PlateSpin Server attempts to communicate with the target workload and to read or write to the VMware floppy devices. The poll interval default is 30000 ms (30 seconds).

vmwareConfigServicePollStartTimeout

This parameter determines how long the PlateSpin Server waits after it starts the target VM before it displays an error in the Web Interface. The default value is 420 seconds (7 minutes).

vmwareConfigServicePollUpdateTimeout

This parameter determines how long the PlateSpin Server waits after each polling interval before displaying an error in the Web Interface. The default value is 300 seconds (5 minutes).

Higher values for these parameters reduce the frequency that the PlateSpin Server attempts to read from and write to the VMware floppy devices on target VMs.

To reduce read and write traffic for VMware floppy devices:

- 1 Log in to the PlateSpin Server, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/platespinconfiguration/`
- 2 Search for the Configuration Service polling parameters, modify their settings as appropriate, then click **Save**.

For example:

```
vmwareConfigServicePollStartDelay = 180 (3 minutes)
vmwareConfigServicePollIntervalInMilliseconds = 300000 (5 minutes)
vmwareConfigServicePollStartTimeout = 1200 (20 minutes)
vmwareConfigServicePollUpdateTimeout = 900 (15 minutes)
```

or

```
vmwareConfigServicePollStartDelay = 300 (5 minutes)
vmwareConfigServicePollIntervalInMilliseconds = 480000 (8 minutes)
vmwareConfigServicePollStartTimeout = 1200 (20 minutes)
vmwareConfigServicePollUpdateTimeout = 900 (15 minutes)
```

- 3 Run an incremental or full replication.
The replication also propagates the modified configuration settings to the target VM.
- 4 Run the Test Failover or Failover again for affected workloads.

Change the Startup Type to Increase the Delay

The Configuration Service might be coming up before resources are accessible. You can change the Configuration Service startup type to have increase the delay.

To change the startup type:

- 1 Log in to the PlateSpin Server, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/platespinconfiguration/`
- 2 Search for the parameter **windowsConfigServiceStartType**.
- 3 Change the **windowsConfigServiceStartType** value to **AutoDelay**.

Options for **windowsConfigServiceStartType** are:

- ♦ **GroupDelay** is the default value and adds the Configuration Service to the end of the **ServiceGroupOrder** in the registry.

- ♦ **AutoDelay** will maximize the amount of time the service waits before starting (2 minutes after boot). Also modify the **ServicesPipeTimeoutForWindowsConfigService** parameter value in [Step 4](#).
 - ♦ **NoDelay** is the most efficient option and starts the service as soon as Windows can. However, it is not recommended because of the potential issues connecting to resources.
- 4 (AutoDelay) Change the **ServicesPipeTimeoutForWindowsConfigService** parameter setting to 180 seconds to account for the 120 seconds that the service will take to start up after boot when AutoDelay is set for **windowsConfigServiceStartType** in [Step 3](#).
 - 5 Click **Save**.
 - 6 Run an incremental or full replication.
The replication also propagates the modified configuration settings to the target VM.
 - 7 Run the Test Failover or Failover again for affected workloads.

Configure Conflicting Services to Not Run Automatically at Startup

During a failover action, a Windows service interferes with the mounting of floppy drivers.

Determine which Windows Services are configured to start up at reboot. Some services are known to interfere with the Configuration Service writing to a floppy, such as Wireless Configuration and some Anti-virus software. You should configure these services to not run automatically on Test Failover or Failover, then run the Test Failover or Failover again.

You can also try to disable all non-essential services for Test Failover and Failover on the Configuration page, then run the Test Failover or Failover again.

9.5.3 Additional Troubleshooting Tips

If the Configuration Service cannot contact the PlateSpin Server, diagnostics will tell only part of the picture. You must also get logs from the target VM:

- ♦ **Windows workloads:** The Configuration Service logs are found in the `C:\windows\platespin\configuration\data` folder.
 - ♦ The `log.txt` file contains all of the logging information, but the `Config.ini` file is useful in understanding what is to be configured.
 - ♦ The `result.txt` file contains the status of the Configuration Service run.
 - ♦ If the target VM cannot read from the input floppy device, it will not have the merged `Config.ini` file, which might include custom network configuration information for the test failover network environment.
 - ♦ If the `Config.ini` file has no network related information (such as a `[NIC0]`), the target VM network adapter might have special characters in the name.
It is a known issue that the `Config.ini` file might not be accurate until it is merged with the one from the floppy device.
 - ♦ The target VM tries a reboot if it cannot connect to either the output floppy or input floppy (one time only). You will see a `config.ini.floppyreboot` file if this is the case.
- ♦ **Linux workloads:** The Configuration Service logs are found in the `/tmp` folder.
 - ♦ The main log files are named `file*.platespin.fileLogger`.
We recommend examining any configuration folders in `/tmp`. Tar the configuration folders along with the `file*.platespin.fileLogger` files to send to Micro Focus Customer Care.
 - ♦ Other config files to check for include the following:

```
/tmp/Ofx.RunCommand.Output*  
/tmp/*DiskHelper*  
/tmp/*VmTools*
```

- ♦ The configuration file is `/usr/lib/psconfigservice/data/config.conf`.
- ♦ The end result log file is `/usr/lib/psconfigservice/data/result.txt`.

9.6 Troubleshooting Traffic-forwarding Workloads

In some scenarios, the replica of a workload that is forwarding network traffic (for example, if the workload's purpose is to serve as a network bridge for NAT, VPN, or a firewall) might show significant network performance degradation. This is related to a problem with VMXNET 2 and VMXNET 3 adapters that have LRO (large receive offload) enabled.

To work around this issue, you need to disable LRO on the virtual network adapter. For more information, see [Knowledgebase Article 7005495 \(https://www.netiq.com/support/kb/doc.php?id=7005495\)](https://www.netiq.com/support/kb/doc.php?id=7005495).

9.7 Troubleshooting Online Help

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

To enable JavaScript:

- ♦ **Chrome:**

1. From the Chrome menu, select **Settings**, then scroll to and click **Show advanced settings**.
2. Under **Privacy**, click **Content Settings**.
3. Scroll to **JavaScript**, then select **Allow all sites to run JavaScript**.
4. Click **Done**.

- ♦ **Firefox:**

1. In the Location bar, type `about:config` and press Enter.
2. Click **I'll be careful, I promise!**
3. In the **Search** bar, type `javascript.enabled`, then press Enter.
4. In the search results, view the value for the `javascript.enabled` parameter. If its value is `false`, right-click `javascript.enabled` and select **Toggle** to set its value to `true`.

- ♦ **Internet Explorer:**

1. From the Tools menu, select **Internet Options**.
2. Select **Security**, then click **Custom level**.
3. Scroll to **Scripting > Active scripting**, then select **Enable**.
4. Click **Yes** at the warning dialog box, then click **OK**.
5. Click **Apply > OK**.

9.8 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** link in the lower right of the panel.
After a few moments, the page refreshes and displays a **Download** link above the **Generate** link.
- 2 Click **Download**.
A .zip file contains the comprehensive diagnostic information about the current command.
- 3 Save the file, then extract the diagnostics to view them.
- 4 Have the .zip file ready if you need to contact Technical Support.

9.9 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. Ensure that 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.

9.10 Post-Protection Workload Cleanup

Use these steps to clean up your source workload from all PlateSpin software components when required, such as following an unsuccessful or problematic protection.

The following sections include more information:

- ♦ [Section 9.10.1, “Cleaning Up Windows Workloads,” on page 131](#)
- ♦ [Section 9.10.2, “Cleaning Up Linux Workloads,” on page 132](#)

9.10.1 Cleaning Up Windows Workloads

Component	Removal Instructions
PlateSpin Block-Based Transfer Component	See Knowledgebase Article 7005616 (https://www.netiq.com/support/kb/doc.php?id=7005616) .

Component	Removal Instructions
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 v7 2. 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	<p>In the workload's Windows directory:</p> <ul style="list-style-type: none"> ♦ Remove all files named <code>machinediscovery*</code>. ♦ Remove the subdirectory named <code>platespin</code>.
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"> ♦ <code>\Program Files\platespin*</code> (32-bit systems) ♦ <code>\Program Files (x86)\platespin*</code> (64-bit systems) 2. Run the following command: <pre>ofxcontroller.exe /uninstall</pre> 3. Remove the <code>platespin*</code> directory

9.10.2 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>kill -9 ofxcontrollerd</code> ♦ <code>kill -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.

Component	Removal Instructions
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>
LVM snapshots	<p>LVP snapshots used by ongoing replications are named according to a <code>volume_name-PS-snapshot</code> convention. For example, a snapshot of a <code>LogVol01</code> 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	For each volume under protection, at the root of the volume, remove the corresponding <code>.blocks_bitmap</code> file.
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>

9.11 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. [Knowledgebase 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.

9.12 Active Directory Domain Services Are Not Available After a Failback (Windows)

Active Directory domain services might not come up after a Failover, if `chkdsk` errors occur. Two avoidable causes of `chkdsk` errors are:

- ♦ Log files related to Microsoft Updates if the source machine is not up-to-date with all Microsoft recommended patches or updates when you perform the first full replication.
- ♦ System files and folders that should be excluded from your anti-virus software.

To avoid these issues, NetIQ recommends the following best practices before you run the first full replication:

- ♦ Ensure that you update Windows (Windows Update) on your source system before you run the first full replication. If the Windows machine is a Domain Controller, ensure that you also disable anti-virus software on the system during the replication.
- ♦ Ensure that you set up your anti-virus software with the recommended file and folder exclusions described in the [Microsoft Knowledge Article: Virus scanning recommendations for Enterprise computers that are running currently supported versions of Windows \(KB: 822158\) \(https://support.microsoft.com/en-us/kb/822158\)](https://support.microsoft.com/en-us/kb/822158).

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 135](#)
- ♦ [Section A.2, “PlateSpin Forge Pre-compiled “blkwatch” Driver \(Linux\),” on page 136](#)

A.1 Analyzing Your Linux Workload

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

- ♦ [Section A.1.1, “Determining the Release String,” on page 135](#)
- ♦ [Section A.1.2, “Determining the Architecture,” on page 135](#)

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 PlateSpin Forge Pre-compiled “blkwatch” Driver (Linux)

Following is a list of non-debug Linux distributions for which PlateSpin 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 of Distributions](#). If you find your release string and architecture, PlateSpin Forge has a pre-compiled version of the `blkwatch` driver.

If your search is unsuccessful, you can create a custom `blkwatch` driver by following the steps found in the [Knowledgebase Article 7005873 \(https://www.netiq.com/support/kb/doc.php?id=7005873\)](https://www.netiq.com/support/kb/doc.php?id=7005873). Self-compiled drivers are supported only for the Linux major and minor kernel versions that appear in the [List of Distributions](#), or a patched version thereof. If the major and minor kernel version in the release string of your Linux workload kernel matches a major and minor kernel version in the list, your self-compiled driver will be supported.

- ♦ [Section A.2.1, “List Item Syntax,” on page 136](#)
- ♦ [Section A.2.2, “List of Distributions,” on page 136](#)
- ♦ [Section A.2.3, “Other Linux Distributions That Use blkwatch Drivers,” on page 145](#)

A.2.1 List Item Syntax

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

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

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

`SLES9-SP1-2.6.5-7.139-bigsmpt-x86`

A.2.2 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-U11-2.6.18-398.el5-x86
RHEL5-U11-2.6.18-398.el5-x86_64
RHEL5-U11-2.6.18-398.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
RHEL6-U6-2.6.32-504.el6-x86
RHEL6-U6-2.6.32-504.el6-x86_64
RHEL7-GA-3.10.0-123.el7.x86_64-x86_64
SLES10-GA-2.6.16.21-0.8-bigsmpp-x86
SLES10-GA-2.6.16.21-0.8-default-x86
SLES10-GA-2.6.16.21-0.8-default-x86_64
SLES10-GA-2.6.16.21-0.8-smp-x86

SLES10-GA-2.6.16.21-0.8-smp-x86_64
SLES10-GA-2.6.16.21-0.8-xen-x86
SLES10-GA-2.6.16.21-0.8-xen-x86_64
SLES10-GA-2.6.16.21-0.8-xenpae-x86
SLES10-SP1-2.6.16.46-0.12-bigsmp-x86
SLES10-SP1-2.6.16.46-0.12-default-x86
SLES10-SP1-2.6.16.46-0.12-default-x86_64
SLES10-SP1-2.6.16.46-0.12-smp-x86
SLES10-SP1-2.6.16.46-0.12-smp-x86_64
SLES10-SP1-2.6.16.46-0.12-xen-x86
SLES10-SP1-2.6.16.46-0.12-xen-x86_64
SLES10-SP1-2.6.16.46-0.12-xenpae-x86
SLES10-SP2-2.6.16.60-0.21-bigsmp-x86
SLES10-SP2-2.6.16.60-0.21-default-x86
SLES10-SP2-2.6.16.60-0.21-default-x86_64
SLES10-SP2-2.6.16.60-0.21-smp-x86
SLES10-SP2-2.6.16.60-0.21-smp-x86_64
SLES10-SP2-2.6.16.60-0.21-xen-x86
SLES10-SP2-2.6.16.60-0.21-xen-x86_64
SLES10-SP2-2.6.16.60-0.21-xenpae-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-bigsmp-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-default-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-default-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-smp-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-smp-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xen-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xen-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xenpae-x86
SLES10-SP3-2.6.16.60-0.54.5-bigsmp-x86
SLES10-SP3-2.6.16.60-0.54.5-default-x86
SLES10-SP3-2.6.16.60-0.54.5-default-x86_64
SLES10-SP3-2.6.16.60-0.54.5-smp-x86
SLES10-SP3-2.6.16.60-0.54.5-smp-x86_64
SLES10-SP3-2.6.16.60-0.54.5-xen-x86
SLES10-SP3-2.6.16.60-0.54.5-xen-x86_64
SLES10-SP3-2.6.16.60-0.54.5-xenpae-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-bigsmp-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-default-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-default-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-smp-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-smp-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xen-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xen-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xenpae-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-bigsmp-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-default-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-default-x86_64

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

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

SLES11SP2-GA-3.0.13-0.27-xen-x86_64
SLES11SP2-LTSS_U1-3.0.101-0.7.19-default-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-default-x86_64
SLES11SP2-LTSS_U1-3.0.101-0.7.19-pae-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-xen-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-xen-x86_64
SLES11SP2-LTSS_U2-3.0.101-0.7.21-default-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-default-x86_64
SLES11SP2-LTSS_U2-3.0.101-0.7.21-pae-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-xen-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-xen-x86_64
SLES11SP2-U1-3.0.26-0.7-default-x86
SLES11SP2-U1-3.0.26-0.7-default-x86_64
SLES11SP2-U1-3.0.26-0.7-pae-x86
SLES11SP2-U1-3.0.26-0.7-xen-x86
SLES11SP2-U1-3.0.26-0.7-xen-x86_64
SLES11SP2-U10-3.0.74-0.6.8-default-x86
SLES11SP2-U10-3.0.74-0.6.8-default-x86_64
SLES11SP2-U10-3.0.74-0.6.8-pae-x86
SLES11SP2-U10-3.0.74-0.6.8-xen-x86
SLES11SP2-U10-3.0.74-0.6.8-xen-x86_64
SLES11SP2-U11-3.0.74-0.6.10-default-x86
SLES11SP2-U11-3.0.74-0.6.10-default-x86_64
SLES11SP2-U11-3.0.74-0.6.10-pae-x86
SLES11SP2-U11-3.0.74-0.6.10-xen-x86
SLES11SP2-U11-3.0.74-0.6.10-xen-x86_64
SLES11SP2-U12-3.0.80-0.5-default-x86
SLES11SP2-U12-3.0.80-0.5-default-x86_64
SLES11SP2-U12-3.0.80-0.5-pae-x86
SLES11SP2-U12-3.0.80-0.5-xen-x86
SLES11SP2-U12-3.0.80-0.5-xen-x86_64
SLES11SP2-U13-3.0.80-0.7-default-x86
SLES11SP2-U13-3.0.80-0.7-default-x86_64
SLES11SP2-U13-3.0.80-0.7-pae-x86
SLES11SP2-U13-3.0.80-0.7-xen-x86
SLES11SP2-U13-3.0.80-0.7-xen-x86_64
SLES11SP2-U14-3.0.93-0.5-default-x86
SLES11SP2-U14-3.0.93-0.5-default-x86_64
SLES11SP2-U14-3.0.93-0.5-pae-x86
SLES11SP2-U14-3.0.93-0.5-xen-x86
SLES11SP2-U14-3.0.93-0.5-xen-x86_64
SLES11SP2-U15-3.0.101-0.5-default-x86
SLES11SP2-U15-3.0.101-0.5-default-x86_64
SLES11SP2-U15-3.0.101-0.5-pae-x86
SLES11SP2-U15-3.0.101-0.5-xen-x86
SLES11SP2-U15-3.0.101-0.5-xen-x86_64
SLES11SP2-U16-3.0.101-0.7.15-default-x86

SLES11SP2-U16-3.0.101-0.7.15-default-x86_64
SLES11SP2-U16-3.0.101-0.7.15-pae-x86
SLES11SP2-U16-3.0.101-0.7.15-xen-x86
SLES11SP2-U16-3.0.101-0.7.15-xen-x86_64
SLES11SP2-U17-3.0.101-0.7.17-default-x86
SLES11SP2-U17-3.0.101-0.7.17-default-x86_64
SLES11SP2-U17-3.0.101-0.7.17-pae-x86
SLES11SP2-U17-3.0.101-0.7.17-xen-x86
SLES11SP2-U17-3.0.101-0.7.17-xen-x86_64
SLES11SP2-U2-3.0.31-0.9-default-x86
SLES11SP2-U2-3.0.31-0.9-default-x86_64
SLES11SP2-U2-3.0.31-0.9-pae-x86
SLES11SP2-U2-3.0.31-0.9-xen-x86
SLES11SP2-U2-3.0.31-0.9-xen-x86_64
SLES11SP2-U3-3.0.34-0.7-default-x86
SLES11SP2-U3-3.0.34-0.7-default-x86_64
SLES11SP2-U3-3.0.34-0.7-pae-x86
SLES11SP2-U3-3.0.34-0.7-xen-x86
SLES11SP2-U3-3.0.34-0.7-xen-x86_64
SLES11SP2-U4-3.0.38-0.5-default-x86
SLES11SP2-U4-3.0.38-0.5-default-x86_64
SLES11SP2-U4-3.0.38-0.5-pae-x86
SLES11SP2-U4-3.0.38-0.5-xen-x86
SLES11SP2-U4-3.0.38-0.5-xen-x86_64
SLES11SP2-U5-3.0.42-0.7-default-x86
SLES11SP2-U5-3.0.42-0.7-default-x86_64
SLES11SP2-U5-3.0.42-0.7-pae-x86
SLES11SP2-U5-3.0.42-0.7-xen-x86
SLES11SP2-U5-3.0.42-0.7-xen-x86_64
SLES11SP2-U6-3.0.51-0.7.9-default-x86
SLES11SP2-U6-3.0.51-0.7.9-default-x86_64
SLES11SP2-U6-3.0.51-0.7.9-pae-x86
SLES11SP2-U6-3.0.51-0.7.9-xen-x86
SLES11SP2-U6-3.0.51-0.7.9-xen-x86_64
SLES11SP2-U7-3.0.58-0.6.2-default-x86
SLES11SP2-U7-3.0.58-0.6.2-default-x86_64
SLES11SP2-U7-3.0.58-0.6.2-pae-x86
SLES11SP2-U7-3.0.58-0.6.2-xen-x86
SLES11SP2-U7-3.0.58-0.6.2-xen-x86_64
SLES11SP2-U8-3.0.58-0.6.6-default-x86
SLES11SP2-U8-3.0.58-0.6.6-default-x86_64
SLES11SP2-U8-3.0.58-0.6.6-pae-x86
SLES11SP2-U8-3.0.58-0.6.6-xen-x86
SLES11SP2-U8-3.0.58-0.6.6-xen-x86_64
SLES11SP2-U9-3.0.74-0.6.6-default-x86
SLES11SP2-U9-3.0.74-0.6.6-default-x86_64
SLES11SP2-U9-3.0.74-0.6.6-pae-x86

SLES11SP2-U9-3.0.74-0.6.6-xen-x86
SLES11SP2-U9-3.0.74-0.6.6-xen-x86_64
SLES11SP3-GA-3.0.76-0.11-default-x86
SLES11SP3-GA-3.0.76-0.11-default-x86_64
SLES11SP3-GA-3.0.76-0.11-pae-x86
SLES11SP3-GA-3.0.76-0.11-xen-x86
SLES11SP3-GA-3.0.76-0.11-xen-x86_64
SLES11SP3-U1-3.0.82-0.7-default-x86
SLES11SP3-U1-3.0.82-0.7-default-x86_64
SLES11SP3-U1-3.0.82-0.7-pae-x86
SLES11SP3-U1-3.0.82-0.7-xen-x86
SLES11SP3-U1-3.0.82-0.7-xen-x86_64
SLES11SP3-U2-3.0.93-0.8-default-x86
SLES11SP3-U2-3.0.93-0.8-default-x86_64
SLES11SP3-U2-3.0.93-0.8-pae-x86
SLES11SP3-U2-3.0.93-0.8-xen-x86
SLES11SP3-U2-3.0.93-0.8-xen-x86_64
SLES11SP3-U3-3.0.101-0.8-default-x86
SLES11SP3-U3-3.0.101-0.8-default-x86_64
SLES11SP3-U3-3.0.101-0.8-pae-x86
SLES11SP3-U3-3.0.101-0.8-xen-x86
SLES11SP3-U3-3.0.101-0.8-xen-x86_64
SLES11SP3-U4-3.0.101-0.15-default-x86
SLES11SP3-U4-3.0.101-0.15-default-x86_64
SLES11SP3-U4-3.0.101-0.15-pae-x86
SLES11SP3-U4-3.0.101-0.15-xen-x86
SLES11SP3-U4-3.0.101-0.15-xen-x86_64
SLES11SP3-U5-3.0.101-0.21-default-x86
SLES11SP3-U5-3.0.101-0.21-default-x86_64
SLES11SP3-U5-3.0.101-0.21-pae-x86
SLES11SP3-U5-3.0.101-0.21-xen-x86
SLES11SP3-U5-3.0.101-0.21-xen-x86_64
SLES11SP3-U6-3.0.101-0.29-default-x86
SLES11SP3-U6-3.0.101-0.29-default-x86_64
SLES11SP3-U6-3.0.101-0.29-pae-x86
SLES11SP3-U6-3.0.101-0.29-xen-x86
SLES11SP3-U6-3.0.101-0.29-xen-x86_64
SLES11SP3-U7-3.0.101-0.31-default-x86
SLES11SP3-U7-3.0.101-0.31-default-x86_64
SLES11SP3-U7-3.0.101-0.31-pae-x86
SLES11SP3-U7-3.0.101-0.31-xen-x86
SLES11SP3-U7-3.0.101-0.31-xen-x86_64
SLES11SP3-U8-3.0.101-0.35-default-x86
SLES11SP3-U8-3.0.101-0.35-default-x86_64
SLES11SP3-U8-3.0.101-0.35-pae-x86
SLES11SP3-U8-3.0.101-0.35-xen-x86
SLES11SP3-U8-3.0.101-0.35-xen-x86_64

SLES9-GA-2.6.5-7.97-bigsmp-x86
SLES9-GA-2.6.5-7.97-default-x86
SLES9-GA-2.6.5-7.97-default-x86_64
SLES9-GA-2.6.5-7.97-smp-x86
SLES9-GA-2.6.5-7.97-smp-x86_64
SLES9-SP1-2.6.5-7.139-bigsmp-x86
SLES9-SP1-2.6.5-7.139-default-x86
SLES9-SP1-2.6.5-7.139-default-x86_64
SLES9-SP1-2.6.5-7.139-smp-x86
SLES9-SP1-2.6.5-7.139-smp-x86_64
SLES9-SP2-2.6.5-7.191-bigsmp-x86
SLES9-SP2-2.6.5-7.191-default-x86
SLES9-SP2-2.6.5-7.191-default-x86_64
SLES9-SP2-2.6.5-7.191-smp-x86
SLES9-SP2-2.6.5-7.191-smp-x86_64
SLES9-SP3-2.6.5-7.244-bigsmp-x86
SLES9-SP3-2.6.5-7.244-default-x86
SLES9-SP3-2.6.5-7.244-default-x86_64
SLES9-SP3-2.6.5-7.244-smp-x86
SLES9-SP3-2.6.5-7.244-smp-x86_64
SLES9-SP4-2.6.5-7.308-bigsmp-x86
SLES9-SP4-2.6.5-7.308-default-x86
SLES9-SP4-2.6.5-7.308-default-x86_64
SLES9-SP4-2.6.5-7.308-smp-x86
SLES9-SP4-2.6.5-7.308-smp-x86_64

A.2.3 Other Linux Distributions That Use blkwatch Drivers

CentOS: PlateSpin Forge supports workloads for a CentOS version if it is based on a supported Red Hat Enterprise Linux distribution. See the RHEL entries in the [List of Distributions](#).

Open Enterprise Server: PlateSpin Forge supports workloads for an Open Enterprise Server 11 SPx version if it is based on a supported SUSE Linux Enterprise Server distribution, for kernel versions 3.0.27 or later. See the SLES entries in the [List of Distributions](#).

Oracle Enterprise Linux: PlateSpin Forge supports workloads for an Oracle Enterprise Linux version if it is based on a supported Red Hat Enterprise Linux distribution, except for workloads using the Unbreakable Enterprise Kernel. See the RHEL entries in the [List of Distributions](#).

B Synchronizing Serial Numbers on Cluster Node Local Storage

This section details the procedure you can use to change local volume serial numbers to match each node of the Windows cluster that you want to protect. The information includes the use of the Volume Manager utility (`VolumeManager.exe`) to synchronize serial numbers on cluster node local storage.

To download and run the utility:

- 1 From the [NetIQ Downloads site](#), search for the PlateSpin Forge product, then click **Submit Query**.
- 2 On the Products tab, select PlateSpin Forge 11.2 to go to the product-specific download page, then click **proceed to download**.
- 3 On the download page, click **download** on the `VolumeManager.exe` line or select the comparable download manager link.
- 4 Download the utility, then copy it to an accessible location on each cluster node.
- 5 On the active node of the cluster, open an administrative command prompt, navigate to the location of the downloaded utility, and run the following command:

```
VolumeManager.exe -l
```

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

```
Volume Listing:
```

```
-----
```

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

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

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

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

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

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

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

- 6d** When you have successfully changed all of the volume serial numbers on a node of the cluster, you need to restart that node.
- 6e** Repeat [Step 6a](#) through [Step 6d](#) for each node of the cluster.
- 7** (Conditional) If the cluster has already been protected in a PlateSpin environment, we recommend running a full replication on the active node to ensure that any changes are propagated to the database.

C

Rebranding the PlateSpin Forge Web Interface

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

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

- ♦ [Section C.1, “Rebranding the Interface By Using Configuration Parameters,” on page 149](#)
- ♦ [Section C.2, “Rebranding the Product Name in the Windows Registry,” on page 152](#)

C.1 Rebranding the Interface By Using Configuration Parameters

As with [other aspects of the PlateSpin Server's behavior](#), you control the appearance of its Web Interface can be controlled by using configuration parameters that you set on a configuration web page residing your Forge VM (https://Your_PlateSpin_Server/platespinconfiguration/). Using these parameters, you can give the Web Interface a “look and feel” that is proprietary to your own organization. This section includes information that you can use to set up custom branding.

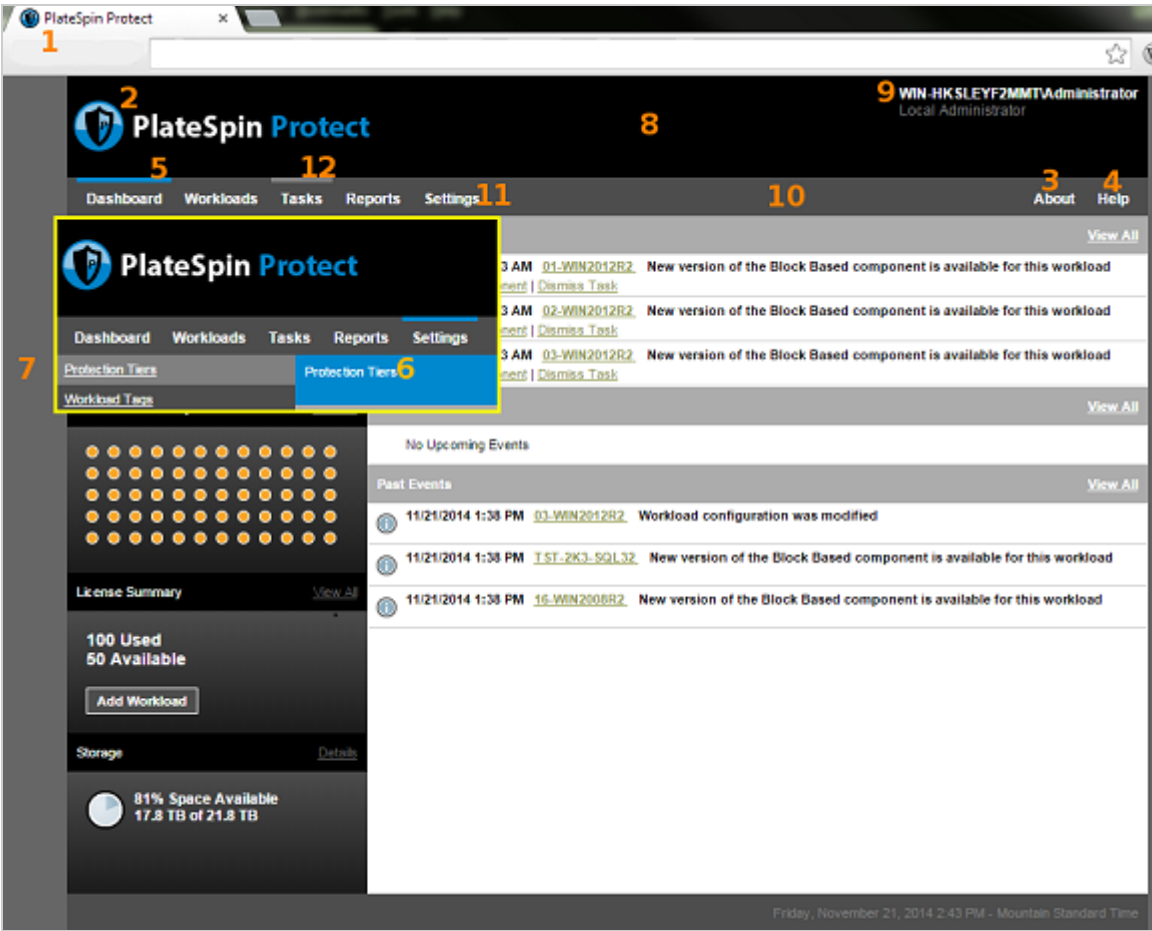
Use the following procedure for changing and applying any configuration parameters:

- 1 From any web browser, open https://Your_PlateSpin_Server/platespinconfiguration/, then log in as Administrator.
- 2 Locate the required server parameter, click **Edit**, then change its value.
For more information, see [Figure C-1](#) and the setting name, description, and default value information for each modifiable element.
- 3 Save and your settings and exit the page.

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

The Web Interface has some common “look and feel” elements throughout its various pages. The illustration of the PlateSpin Forge Dashboard in [Figure C-1](#) identifies the elements you can modify with numbered callouts.

Figure C-1 PlateSpin Forge Web Interface with Configurable Elements Labeled (inset added)



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

ID	Setting Name and Description	Default Value
1	<p>WebUIFaviconUrl</p> <p>Location of a valid .ico graphic file. Specify one of the following:</p> <ul style="list-style-type: none"> ♦ A valid URL to the appropriate .ico file on a different machine. <p>For example: <code>https://myserver.example.com/dir1/dir2/icons/mycompany_favicon.ico</code></p> ♦ A relative path below the root of the local web server where you have uploaded the appropriate .ico file. <p>For example, if you create a path called <code>mycompany\images\icons</code> at the root of the web server to store your custom icon graphics:</p> <p><code>~/mycompany/images/icons/ mycompany_favicon.ico</code></p> <p>In this example, the actual file system path that contains the file is <code>C:\Program Files (x86)\PlateSpin Protect Server\PlateSpin Forge\web\mycompany\images\icons\mycompany_favicon.ico</code>.</p>	<p>~/doc/en/favicon.ico ¹</p>
2	<p>WebUILogoUrl</p> <p>Location of product logo graphic file. Specify one of the following:</p> <ul style="list-style-type: none"> ♦ A valid URL to the appropriate graphics file on a different machine. <p>For example: <code>https://myserver.example.com/dir1/dir2/logos/mycompany_logo.png</code></p> ♦ A relative path below the root of the local web server where you have uploaded the appropriate graphics file. <p>For example, if you create a path called <code>mycompany\images\logos</code> at the root of the web server to store your custom logo images:</p> <p><code>~/mycompany/images/logos/mycompany_logo.png</code></p> <p>In this example, the actual file system path that contains the file is <code>C:\Program Files (x86)\PlateSpin Protect Server\PlateSpin Forge\web\mycompany\images\logos\mycompany_logo.png</code>.</p>	<p>~/Resources/protectLogo.png ²</p>
3	<p>WebUIShowAboutTab</p> <p>Toggle the visibility of the About tab on (True) or off (False).</p>	<p>True</p>
4	<p>WebUIShowHelpTab</p> <p>Toggle the visibility of the Help tab on (True) or off (False).</p>	<p>True</p>
5	<p>WebUISiteAccentColor</p> <p>Accent color (RGB hex value)</p>	<p>#0088CE</p>

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

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

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

C.2 Rebranding the Product Name in the Windows Registry

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

To change the product name:

- 1 At the PlateSpin Server, run `regedit`.
- 2 In the Windows Registry Editor, navigate to the following registry key:
`HKEY_LOCAL_MACHINE\SOFTWARE\PlateSpin\ForgeServer\ProductName`

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

`HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\PlateSpin\Forge`

- 3 Double-click the `ProductName` key and change the **Value data** for the key as you prefer, then click **OK**.
- 4 Restart the IIS Server for the interface change to take effect.

D Using Workload Protection Features through the PlateSpin Protect Server API

You can use workload protection features of PlateSpin Forge programmatically through the PlateSpin Protect Server API (`protectionsservices`) from within your applications. You can use any programming or scripting language that supports an HTTP client and JSON serialization framework.

NOTE: The Protect Server API is experimental. Information in this section is provided as a technology preview.

- ♦ [Section D.1, “API Overview,” on page 155](#)
- ♦ [Section D.2, “PlateSpin Protect Server API Documentation,” on page 155](#)
- ♦ [Section D.3, “Samples and Other References,” on page 156](#)

D.1 API Overview

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

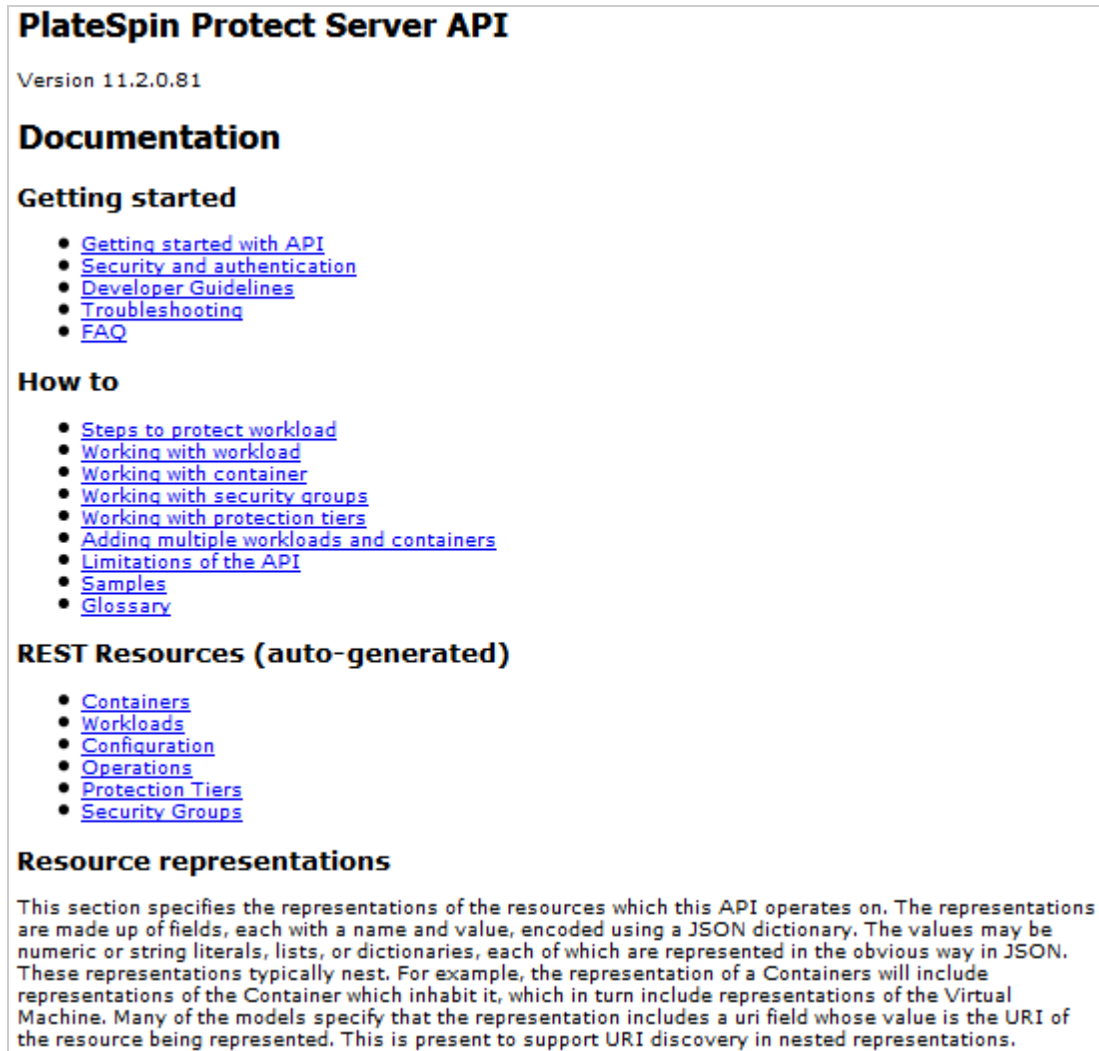
D.2 PlateSpin Protect Server API Documentation

The PlateSpin Protect Server API home page for `protectionsservices` provides documentation and samples that can be useful for developers and administrators. For information, go to the following location on your Forge VM:

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

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

Figure D-1 The Home Page of the Protect Server API



D.3 Samples and Other References

Forge administrators can leverage a JScript sample from the command line to access the product through the API. On the Forge VM, see the sample at

<https://localhost/protectionservices/Documentation/Samples/protect.js>

The sample can help you write scripts to help you work with the product. Using the command line utility, you can perform the following operations:

- ♦ Add a single workload
- ♦ Add a single container
- ♦ Run the replication, failover, and failback operations
- ♦ Add multiple workloads and containers at one time

NOTE: For more information about this operation, see the API documentation at <https://localhost/protectionservices/Documentation/AddWorkloadsAndContainersFromCsvFile.htm>

- ♦ remove all workloads at one time
- ♦ remove all container at one time

To script common workload protection operations, use the referenced samples written in Python as guidance. A Microsoft Silverlight application, along with its source code, is also provided for reference purposes.

E

Documentation Updates

This section contains information on documentation content changes that were made in this *User Guide* after the initial release of NetIQ PlateSpin Forge 11.2.

- ♦ [Section E.1, “March 2018,” on page 159](#)
- ♦ [Section E.2, “January 2018,” on page 159](#)
- ♦ [Section E.3, “October 2017,” on page 160](#)
- ♦ [Section E.4, “June 2017,” on page 160](#)
- ♦ [Section E.5, “February 2017,” on page 160](#)
- ♦ [Section E.6, “January 2017,” on page 160](#)
- ♦ [Section E.7, “November 2016,” on page 160](#)
- ♦ [Section E.8, “September 2016,” on page 161](#)
- ♦ [Section E.9, “April 2016,” on page 161](#)
- ♦ [Section E.10, “February 2016,” on page 161](#)

E.1 March 2018

Location	Update
Section 3.6, “Resetting the Forge Appliance to Factory Defaults,” on page 55	Added a step for Forge models that do not have an optical media player. Added success and failure information for the factory reset.

E.2 January 2018

Location	Update
Section 3.4, “Physically Relocating the Appliance,” on page 50	Ensure that you use a single container VMware license for the Forge Appliance. Using a multiple container license can create issues when you attempt to physically relocate the Appliance.

E.3 October 2017

Location	Update
Section 1.1.1, "Supported Windows Workloads," on page 10	Protection is no longer supported for desktop (workstation) workloads.

E.4 June 2017

Location	Update
Section 1.1.1, "Supported Windows Workloads," on page 10	Protection is no longer supported for Windows 8 and 8.1.

E.5 February 2017

Location	Update
Section 1.1.3, "Supported VM Containers," on page 13	PlateSpin Forge ships with VMware ESXi 5.5 GA2 Update 2 as the protection VM container. Previous content in this section applied only to PlateSpin Protect.

E.6 January 2017

Location	Update
Section 9.5, "Troubleshooting the Configuration Service," on page 125	This section is new.

E.7 November 2016

Location	Update
Section 3.5, "Recovering the Forge Management VM to Factory Defaults," on page 54	This section is new.

E.8 September 2016

Location	Update
Section A.2.2, "List of Distributions," on page 146	Corrected the list of supported distributions for RHEL 6 U6.

E.9 April 2016

Location	Update
Section 1.2.3, "Security of Credentials," on page 15	PlateSpin Forge protects credentials by using an SSL connection for communications and the Windows cryptographic library to encrypt passwords.
Section 2.2, "Activating Your Product License," on page 22	A Customer Care representative will contact you and provide the license activation code.
Section 2.2.3, "Managing Licenses," on page 23	This section is new.
Section 2.9.1, "Tuning Parameters," on page 38	Updated the definition for the FileTransferSendReceiveBufferSize parameter.
Section 2.9.2, "Tuning FileTransferSendReceiveBufferSize," on page 46	Updated the definition for the FileTransferSendReceiveBufferSize parameter.

E.10 February 2016

Location	Update
"Tuning Parameters" on page 38	Updated the equation for the FileTransferSendReceiveBufferSize parameter.
"Tuning FileTransferSendReceiveBufferSize" on page 46	This section is new.
"ProtectAgent Utility" on page 115	Corrected a typo in the following sample command: <code>C:\ProtectAgent\ProtectAgent.cli.exe status</code>

