AUTHORIZED DOCUMENTATION

User Guide

Novell® PlateSpin® Protect

10

July 28, 2010

www.novell.com

Legal Notices

Novell, Inc., makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc., makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. See the Novell International Trade Services Web page (http://www.novell.com/info/exports/) for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
U.S.A.
www.novell.com

Online Documentation: To access the latest online documentation for this and other Novell products, see the Novell Documentation Web page (http://www.novell.com/documentation).

Novell Trademarks

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

Third-Party Materials

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

Contents

About This Guide

This text provides information about using PlateSpin Protect 10.

Audience

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

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation, or submit your comments through the Novell Documentation Feedback site (http://www.novell.com/documentation/feedback.html).

Additional Documentation

This text is part of the PlateSpin Protect documentation set.

For a complete list of publications supporting this release, visit the PlateSpin Protect 10 Online Documentation Web Site (http://www.novell.com/documentation/platespin_protect_10).

Documentation Updates

For the most recent version of this text, visit the product’s Online Documentation Web Site (see Additional Documentation).

Additional Resources

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

- Novell User Forum (http://forums.novell.com): A Web-based community with a variety of discussion topics.
- Novell Knowledge Base (http://www.novell.com/support): A collection of in-depth technical articles.

Technical Support

- Telephone (North America): +1-877-528-3774 (1 87 PlateSpin)
- Telephone (global): +1-416-203-4799
- E-mail: support@platespin.com

You can also visit the PlateSpin Technical Support Web site (http://www.platespin.com/support).

Product Overview 1

1.1 About PlateSpin Protect

PlateSpin Protect is software that replicates and rapidly recovers workloads (operating systems, middleware, and data). In the event of a production server outage or disaster, workloads can be rapidly powered on and continue to run as normal until the production environment is restored.

PlateSpin Protect features two different mechanisms for workload protection:

- **Virtualization:** This mechanism provides you the capability to rapidly recover a workload and requires that you have an existing VM host (a VM *container*).

  In this scenario, PlateSpin Protect creates a failover workload (a virtual replica of your production workload) and regularly updates it at configurable intervals. If your production workload goes offline, you can fail it over to the VM replica, which takes over the business services of the failed workload. You can then fail this workload back to either its original or completely new infrastructure, physical or virtual.
- **Imaging:** This mechanism provides you with the capability to recover a workload using a protected image of its volumes. Recovering a workload from a protected image takes longer than recovery from a virtual replica. However, imaging requires no VM hosts; a workload’s image is captured, stored, and regularly updated on almost any host that you designate as an image server. If your production workload goes offline, you can deploy the captured image to run on either physical hardware or, if required and available, also on a VM host.

1.2 Supported Configurations

1.2.1 Supported Workloads in VM Containers

Operating systems:

- SUSE Linux Enterprise Server (SLES) 10, 11
- Red Hat Enterprise Linux (RHEL) 4, 5
- Windows Server 2008 (including domain controller (DC) and Small Business Server (SBS) systems)
- Windows Server 2003 (including DC and SBS systems)
- Windows Vista
- Windows Server 2000

- Windows XP
- Windows clusters (supported only to targets on VMware ESX 3.0.2 and later). See "Protecting Windows Clusters" on page 31.

Supported international versions:

French, German, Japanese, Chinese Traditional, and Chinese Simplified

Supported VM Containers

The following virtualization platforms are supported as VM containers:

- VMware ESX 3.0.2
- VMware ESX 3i
- VMware ESX 4i
- VMware ESX 3.5.*x*
- VMware ESX 4

1.2.2 Supported Workloads in Image Containers

Operating Systems:

- Windows Server 2008 (including DC and SBS systems)
- Windows Vista
- Windows Server 2003 (including DC and SBS systems)
- Windows 2000
- Windows XP

Supported international versions:

French, German, Japanese, Chinese Traditional, and Chinese Simplified

Supported Image Container Hosts

- Windows Server 2008
- Windows Server 2003
- Windows Server 2000

1.3 RPO, RTO, and TTO Specifications

- **Recovery Point Objective (RPO):** Describes the acceptable amount of data loss measured in time. The RPO is determined by the time between incremental replications of a protected workload and is affected by current utilization levels of PlateSpin Protect, the rate and scope of changes on the workload, and your network speed.
- **Recovery Time Objective (RTO):** Describes the time required for a failover operation (bringing a workload replica online to temporarily replace a protected production workload).

  The RTO in the process of failing a workload over to its virtual replica is affected by the time it takes to configure and execute the failover operation (10 to 45 minutes). See "Failover" on page 25.

  The RTO in the process of deploying a protected image as a bootable workload is affected by the target infrastructure (physical or virtual) and the time required for the corresponding image deployment procedure. See "Deploying a Workload Image" on page 34.
- **Test Time Objective (TTO):** Describes the time required for testing disaster recovery with some confidence of service restoration.

  Use the *Test Failover* feature to run through different scenarios and generate benchmark data.

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 than multiple failed-over workloads, which share the resources of their underlying infrastructure.

You should get average failover times for workloads in your environment by doing test failovers at various times and use them as benchmark data in your overall data recovery plans. See "Workload and Workload Protection Reports" on page 19.

Getting Started with PlateSpin Protect

2

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

2.1 Working with the PlateSpin Protect User Interface

2.1.1 Launching the PlateSpin Protect Web Client

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

The supported browsers are:

- Microsoft Internet Explorer 7, 8
- Mozilla Firefox 3.6

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

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

To use the PlateSpin Protect Web Client and integrated help in one of the supported languages, see “Language Setup for International Versions of PlateSpin Protect” in your *Application Configuration Guide*.

To launch the PlateSpin Protect Web Client:

1 Open a Web browser and go to:

`http://<`*hostname* | *IP_address*`>/Protect`

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

If SSL is enabled, use `https` in the URL.

2.1.2 Elements of the PlateSpin Protect Web Client

The default interface of the PlateSpin Protect Web Client is the Dashboard page, which contains elements for navigating to different functional areas of the interface and carrying out workload protection and recovery tasks.

Figure 2-1 *The Default Dashboard Page of the PlateSpin Protect Web Client*

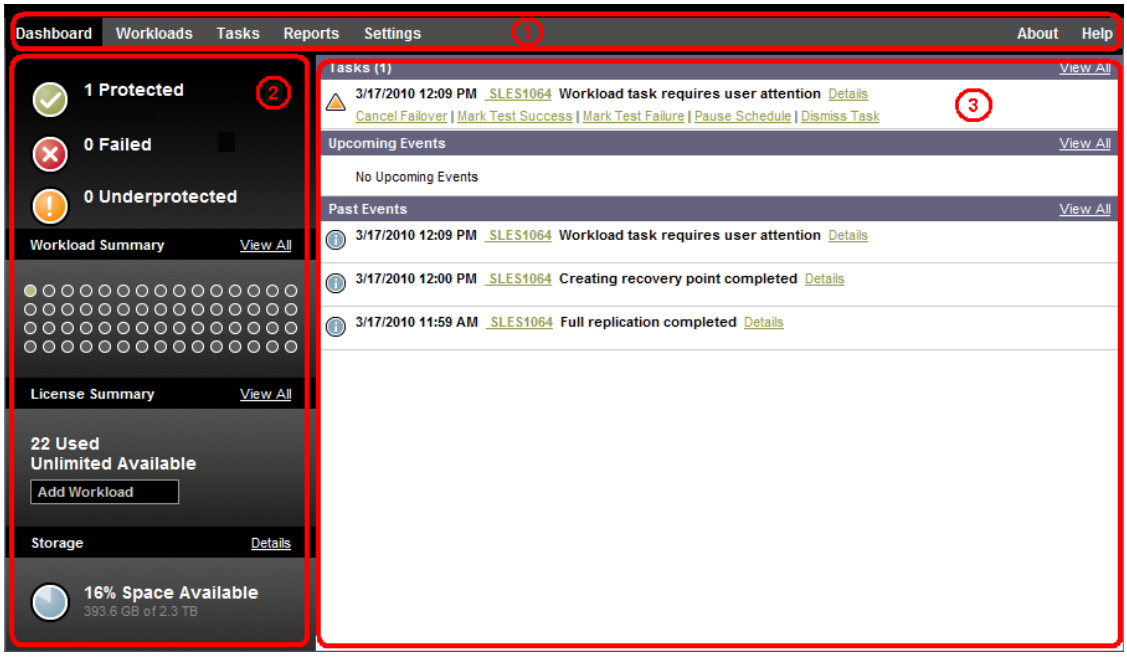

The Dashboard page consists of the following elements:

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

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 14.
- **Tasks:** Displays the Tasks page, which lists items requiring user intervention.
- **Reports:** Displays the Reports page. See "Workload and Workload Protection Reports" on page 19.
- **Settings:** Displays the Settings page, which provides access to the following configuration options:
  - **Protection Tiers:** See "Protection Tiers" on page 46.
  - **Permissions:** See "Setting Up User Authorization and Authentication" in your *Application Configuration Guide*.
  - **Containers:** See "Adding Containers" on page 18.
  - **Email/SMTP:** See "Setting Up E-Mail Notifications" in your *Application Configuration Guide*.
  - **Licenses/License Designations:** See "Product Licensing" in your *Application Configuration Guide*.

Visual Summary Panel

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

Inventoried workloads are represented by three categories:

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

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

Table 2-1 *Dot Icon Workload Representation*

	*Unprotected*		*Underprotected*

	Unprotected – Error		*Failed*
	Protected		*Expired*
	Unused		

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

Storage provides information about storage space available to PlateSpin Protect.

Tasks and Events Panel

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

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.

2.1.3 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 2-2 *The Workloads Page*

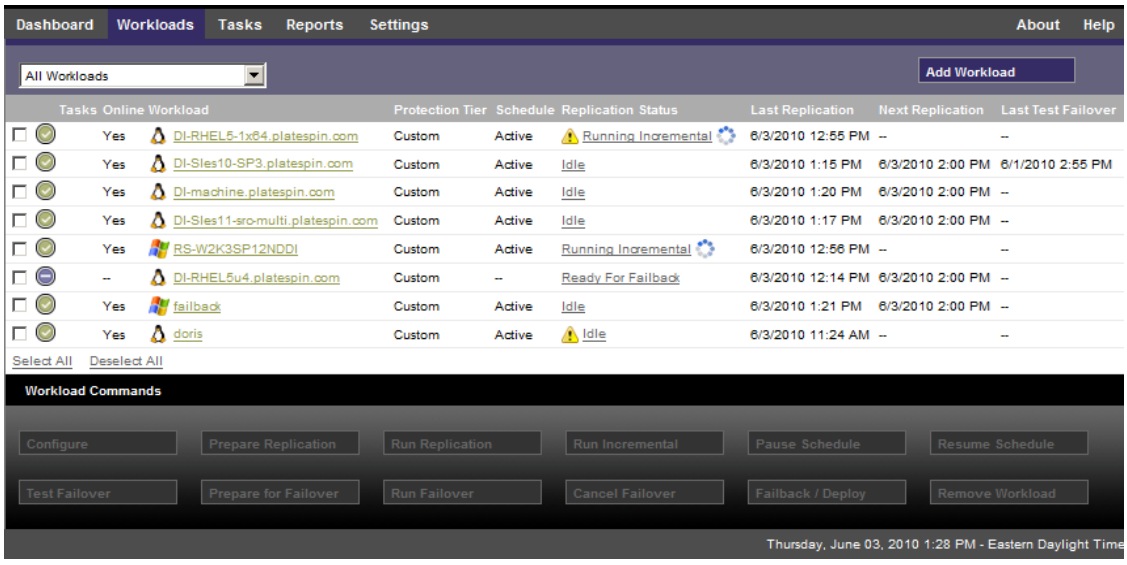

NOTE: All time stamps reflect the time zone of the PlateSpin Protect Server host. This might be different from the time zone of the protected workload or the time zone of the host on which you are running the PlateSpin Protect Web Client. A display of the server date and time appears at the bottom right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

Workload Protection and Recovery Commands

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

Figure 2-3 *Workload Commands*

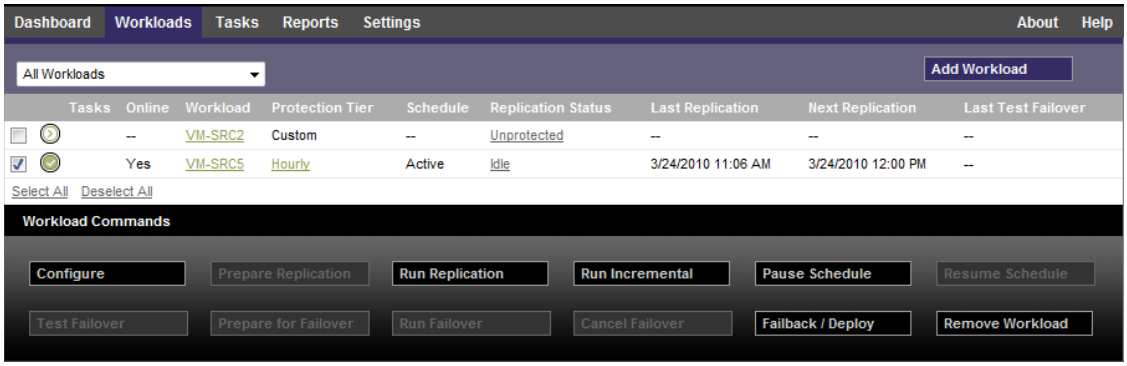

The following is a summary of workload commands along with their functional descriptions.

Table 2-2 *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 VM in preparation of workload replication.
Run Replication	Starts replicating the source workload and establishes the workload protection contract according to specified parameters.
Run Incremental	Performs an individual transfer of changed data from the source to the target outside the workload protection schedule.
Pause Schedule	Suspends the protection and pauses data transfers from the protected workload.
Resume Schedule	Resumes the protection according to saved protection settings.
Test Failover	Brings the recovery workload online in an isolated environment within the container for testing purposes.
Prepare for Failover	Boots the recovery workload in preparation for a failover operation.
Run Failover	Boots and configures the recovery workload, which takes over the business services of a failed workload.
Cancel Failover	Aborts the failover process.

Workload Command	Description
Failback / Deploy	Following a failover operation, fails the recovery workload back to its original infrastructure or to a new infrastructure (virtual or physical).
Remove Workload	Removes a workload from the inventory.

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

You can use workload protection functionality programmatically, using the `protection.webservices` API from within your applications. You can use any programming or scripting language that supports Web services.

`http://`<*hostname* | *IP_address*>`/protection.webservices`

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

To script common workload protection operations, use the referenced sample written in Python as guidance.

2.3 Managing Multiple Instances of PlateSpin Protect

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

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

- Section 2.3.1, “Using the PlateSpin Protect Management Console,” on page 16
- Section 2.3.2, “About PlateSpin Protect Management Console Cards,” on page 17
- Section 2.3.3, “Adding Instances of PlateSpin Protect to the Management Console,” on page 18
- Section 2.3.4, “Managing Cards on the Management Console,” on page 18

2.3.1 Using the PlateSpin Protect Management Console

1. Open a Web browser on a machine that has access to your PlateSpin Protect instances and navigate to the following URL:

   `http://`<*IP_address* | *hostname*>`/console`

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

2. Log in with your username and password.

   The console’s default Dashboard page is displayed.

Figure 2-4 *The Management Console's Default Dashboard Page*

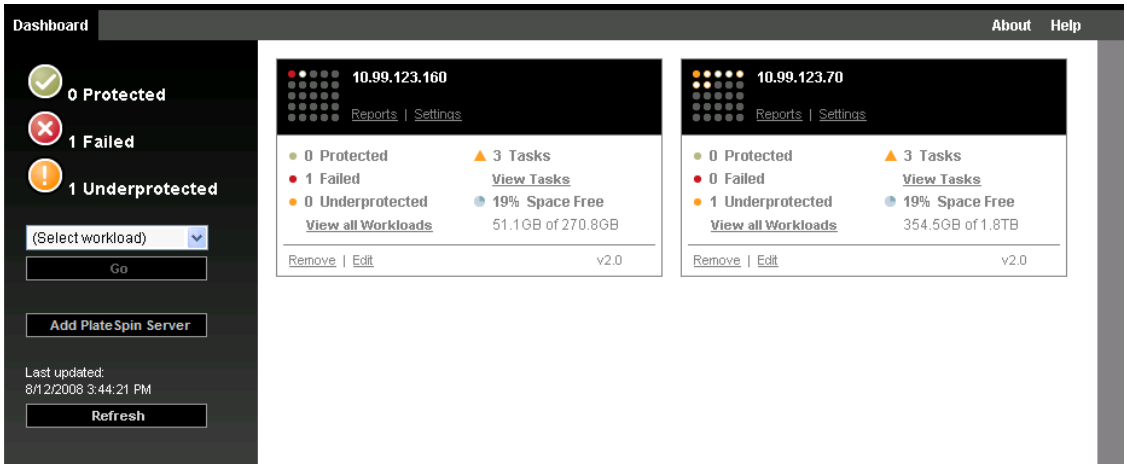

2.3.2 About PlateSpin Protect Management Console Cards

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

Figure 2-5 *PlateSpin Protect Instance Card*

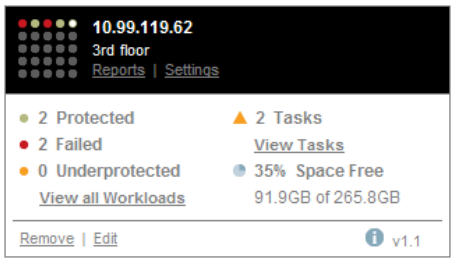

A card displays basic information about the specific instance of PlateSpin Protect, such as:

- 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.

2.3.3 Adding Instances of PlateSpin Protect to the Management Console

Adding a PlateSpin Protect 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 on a PlateSpin Protect instance, that instance is not automatically added to the console. It must be manually added to the console.

To add a PlateSpin Protect instance to the console:

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

   The *Add/Edit* page is displayed.
2. Specify the URL of the PlateSpin Protect Server host. Both HTTP and HTTPS protocols are supported.
3. (Optional) Enable the *Use Management Console Credentials* check box to use the same credentials as those used by the console. When it is selected, the console automatically populates the *Domain\Username* field.
4. In the *Domain\Username* field, type a domain name and a username valid for the PlateSpin Protect instance that you are adding. In the *Password* field, type the corresponding password.
5. (Optional) Specify a descriptive or identifying *Display Name* (15 characters max), a *Location* (20 characters max), and any *Notes* you might require (400 characters max).
6. Click *Add/Save*.

   A new card is added to the dashboard.

2.3.4 Managing Cards on the Management Console

You can modify the details of a PlateSpin Protect 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 PlateSpin Protect 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 appliance card is removed from the dashboard.

2.4 Adding Containers

A container is either of the two workload protection infrastructures supported by PlateSpin Protect: a VM host or an image server.

To add a container:

1 In your PlateSpin Protect Web Client, click *Settings* > *Containers* > *Add Container*.

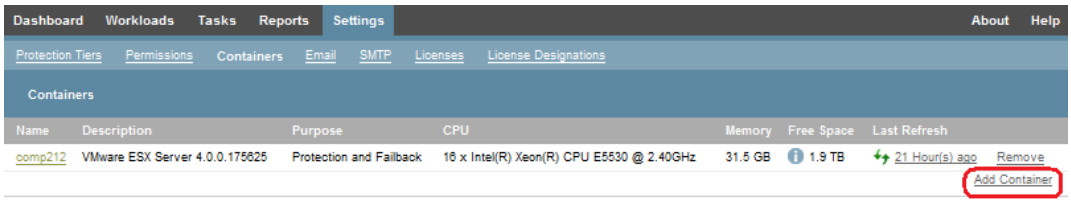

2 Specify the following parameters:

- **Type:** Select the type of the container (ESX Server or PlateSpin Image Server). Make sure the container is supported. See:
  - "Supported VM Containers" on page 8
  - "Supported Image Container Hosts" on page 8
- **Hostname or IP:** Type the container's hostname or IP address.
- **Username/Password:** Provide admin-level credentials for accessing the required host. See "Guidelines for Workload and Container Credentials" on page 45
- **Purpose:** (applicable to ESX containers only) Select the required item (*Protection*, *Failback*, or both). Selecting both results in that container being available for selection as a target in both protection and failback operations.

3 Click *Add*.

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

To refresh a container, click the *Refresh* icon next to the container you want to refresh.

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

2.5 Workload and Workload Protection Reports

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

The following report types are supported:

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

Figure 2-6 *Options for a Replication History Report*

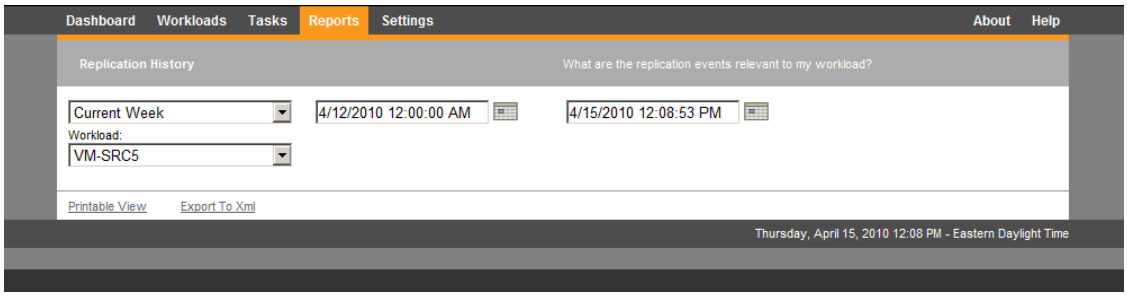

To generate a report:

1. In your PlateSpin Protect Web Client, click *Reports*.

   A list of the report types is displayed.
2. Click the name of the required report type.

Workload Protection 3

PlateSpin Protect creates a replica of your production workload and regularly updates that replica based on changes that the protected workload undergoes over time.

The replica, or the *failover workload*, is a virtual machine in the VM container of PlateSpin Protect and takes over the business function of your production workload in case of a disruption at the production site.

In addition to workload protection through virtualization, PlateSpin Protect provides workload image protection through volume imaging. See “Workload Image Protection” on page 33.

3.1 Basic Workflow for Workload Protection and Recovery

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

1. Preparatory step:
   - **1a** Make sure that PlateSpin Protect supports your workload. See “Supported Configurations” on page 7.
   - **1b** Make sure that your workloads and containers meet access and network prerequisites. See “Access and Communication Requirements across your Protection Network” in your *Application Configuration Guide*.
   - **1c** (Linux only)
     - (Conditional) If you plan to protect a supported Linux workload that has a non-standard, customized, or newer kernel, rebuild the PlateSpin `blkwatch` module, which is required for block-level data replication. See KB Article 7005873 (http://www.novell.com/support/viewContent.do?externalId=7005873).
     - (Recommended) Prepare LVM snapshots for block-level data transfer. See KB Article 7005872 (http://www.novell.com/support/viewContent.do?externalId=7005872).
     - (Optional) Determine and prepare any custom scripts that you want to execute on your source workload upon each replication. See “Automatically Executing Custom Scripts upon Every Replication (Linux)” on page 49.
   - **1d** (Optional) Define a replication blackout window if required. See *Parameters for Imposing a Replication Blackout Window* in your *Application Configuration Guide*.
2. Add a container. See “Adding Containers” on page 18.
3. Add a workload. See “Adding a Workload for Protection” on page 22.
4. Configure protection details and prepare the replication. See “Configuring Protection Details and Preparing the Replication” on page 23.
5. Start the workload protection schedule. See “Starting the Workload Protection” on page 25.
6. (Optional) Manually run an incremental.
7. (Optional) Test the failover functionality. See Testing the Recovery Workload and the Failover Functionality

8 Perform a failover. See “Failover” on page 25

9 Perform a failback. See “Failback” on page 28.

10 (Optional) Reprotect a workload after failback.

Except for Steps 1, 8, and 9, these are represented by workload commands on the Workloads page. See “Workloads and Workload Commands” on page 14.

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

3.2 Adding a Workload for Protection

1 Follow the required preparatory steps. See Step 1 in “Basic Workflow for Workload Protection and Recovery” on page 21.

2 Add a VM container. See “Adding Containers” on page 18.

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

The PlateSpin Protect Web Client displays the Add Workload page.

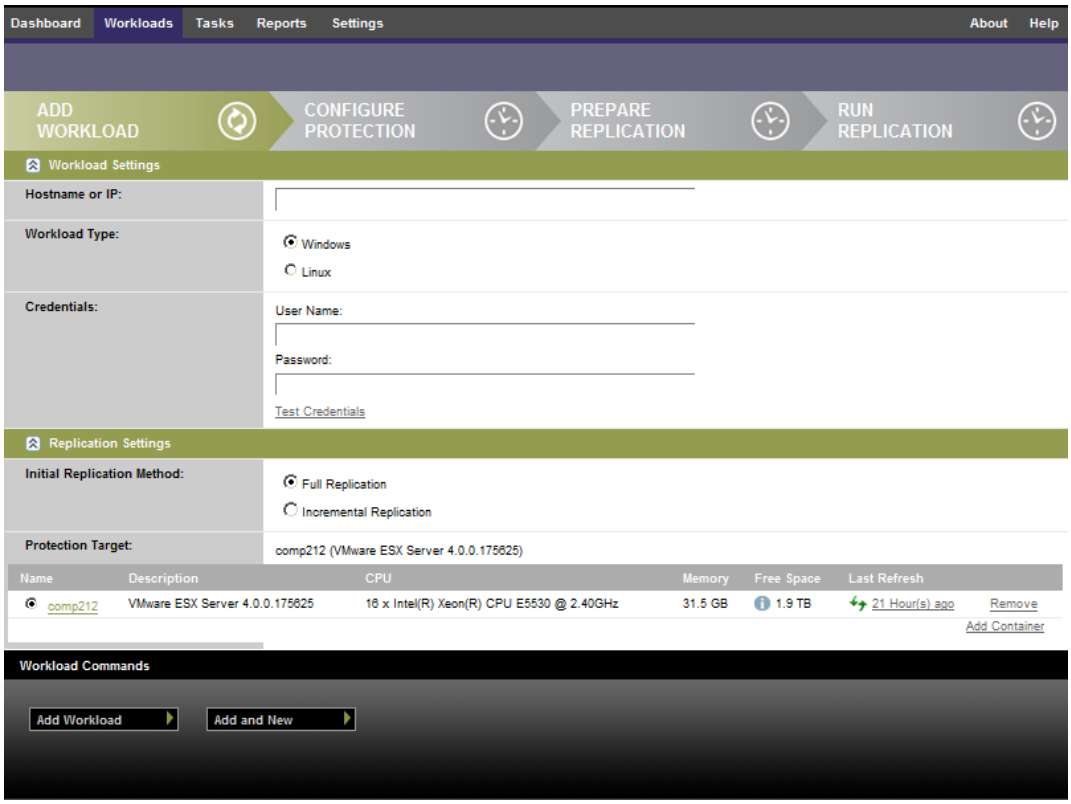

4 Specify the required workload details:

- **Workload Settings:** Specify your workload’s hostname or IP address, the operating system, and admin-level credentials. Use the required credential format (see “Guidelines for Workload and Container Credentials” on page 45. To make sure that PlateSpin Protect can access the workload, click *Test Credentials*.
- **Replication Settings:** Select the required replication settings. See “Initial Replication Method (Full and Incremental)” on page 47.

- **Protection Target:** Select the required protection target. This is either the target container or, if you have selected *Incremental Replication* as the initial replication method, a prepared workload. See “Initial Replication Method (Full and Incremental)” on page 47.

5 Click *Add Workload*.

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

3.3 Configuring Protection Details and Preparing the Replication

Protection details control the workload protection and recovery settings and behavior over the entire life cycle of a workload under protection. At each phase of the protection and recovery workflow (see “Basic Workflow for Workload Protection and Recovery” on page 21), relevant settings are read from the protection details.

To configure your workload’s protection details:

1 Add a workload. See “Adding a Workload for Protection” on page 22.

2 On the Workloads page, select the required workload and click *Configure*.

The PlateSpin Protect Web Client displays the workload’s Protection Details page.

3 Configure the protection details in each set of settings as dictated by your business continuity needs. See “Workload Protection Details” on page 23.

4 Correct any validation errors.

5 Click *Save*.

Alternatively, click *Save & Prepare*. This saves the settings and simultaneously executes the *Prepare Replication* command (installing data transfer drivers on the source workload if required 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.

3.3.1 Workload Protection Details

Workload protection details are represented by five sets of parameters:

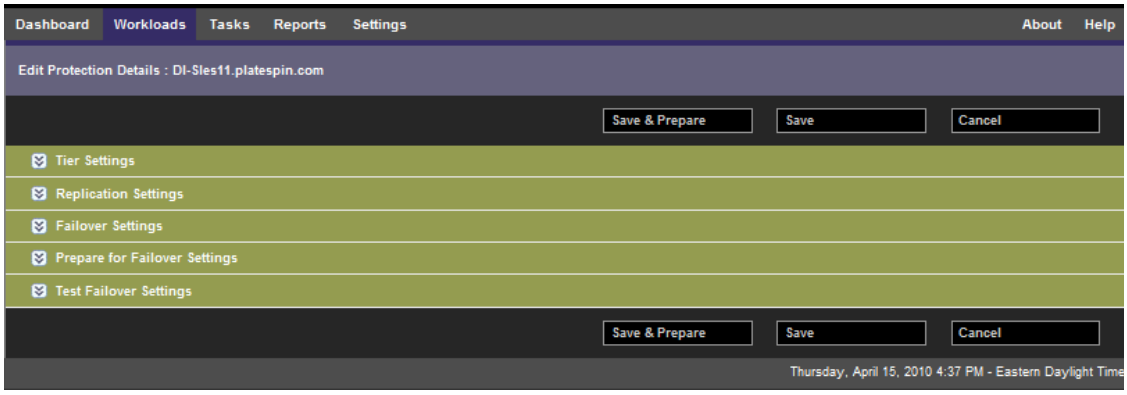

You can expand or collapse each parameter set by clicking the icon at the left.

The following are the details of the five parameter sets:

Table 3-1 *Workload Protection Details*

Parameter Set (Settings)	Details
Tier	Indicates the Protection Tier that the current protection contract follows. See “Protection Tiers” on page 46.
Replication	**Transfer Method:** (Windows) Enables you to select a data transfer mechanism and security through encryption. See “Transfer Methods and Data Transfer Security” on page 46.
	Source Credentials: Required for accessing the workload. See “Guidelines for Workload and Container Credentials” on page 45.
	Number of CPUs: Enables you to specify the required number of vCPUs assigned to the recovery workload.
	Replication Network: Enables you to separate replication traffic based on virtual networks defined on your VM container. See “Networking” on page 51.
	Recovery Point Datastore: Enables you to select a datastore associated with your VM container for storing Recovery Points. See “Recovery Points” on page 47.
	Protected Volumes: Use these options to select volumes for protection and to assign their replicas to specific datastores on your VM container. For Linux, you can also select logical volumes and volume groups for protection. See “Volumes” on page 50.
	Services/Daemons to stop: Enables you to select Windows services or Linux Daemons that are automatically stopped during the replication. See “Service and Daemon Control” on page 49.
Failover	**VM Memory:** Enables you to specify the amount of memory allocated to the failover VM.
	Hostname and Domain/Workgroup affiliation: Use these options to control the identity and domain/workgroup affiliation of the failover workload when it is live. For domain affiliation, domain admin credentials are required.
	Network Connections: Use these options to control the LAN settings of the failover workload. See “Networking” on page 51.
	Service States to Change: Enables you to control the startup state of specific application services (Windows) or daemons (Linux). See “Service and Daemon Control” on page 49.
Prepare for Failover	Enables you to control the temporary network settings of the failover workload during the optional Prepare for Failover operation. See “Networking” on page 51.

Parameter Set (Settings)	Details
Test Failover	**VM Memory:** Enables you to assign the required RAM to the temporary workload. **Hostname:** Enables you to assign a hostname to the temporary workload. **Domain/Workgroup:** Enables you to affiliate the temporary workload with a domain or a workgroup. For domain affiliation, domain admin credentials are required. **Network Connections:** Controls the LAN settings of the temporary workload. See “Networking” on page 51. **Service States to Change:** Enables you to control the startup state of specific application services (Windows) or daemons (Linux). See “Service and Daemon Control” on page 49.

3.4 Starting the Workload Protection

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

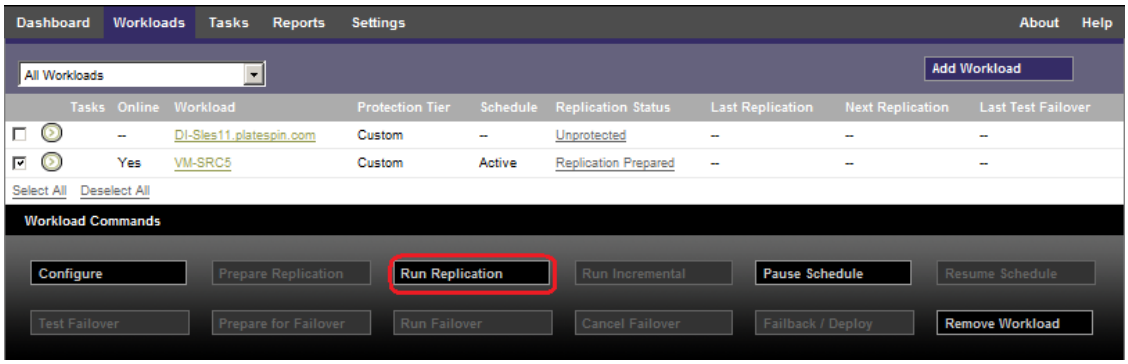

You can execute the Run Replication command after:

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

When you are ready to proceed:

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

   PlateSpin Protect starts the execution and displays a process indicator for the *Copy data* step  .

3.5 Failover

Failover is when the business function of a failed workload is taken over by a recovery workload within a PlateSpin Protect VM container.

- Section 3.5.1, “Failure Detection,” on page 26

- Section 3.5.2, “Performing a Failover,” on page 26
- Section 3.5.3, “Testing the Recovery Workload and the Failover Functionality,” on page 27

3.5.1 Failure Detection

If an attempt to detect a workload fails for a predefined number of times, PlateSpin Protect generates a *Workload is offline* event. Criteria that determine and log a workload failure are part of a workload protection contract’s Tier settings (see the Tier row in “Workload Protection Details” on page 23).

If notifications are configured along with SMTP settings, PlateSpin Protect simultaneously sends a notification e-mail to the specified recipients. See *Setting Up E-Mail Notifications* in your *Application Configuration Guide*.

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, and then proceed to the *Run Failover* command. See “Performing a Failover” on page 26.

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

Figure 3-1 *The Dashboard Page upon Workload Failure Detection*

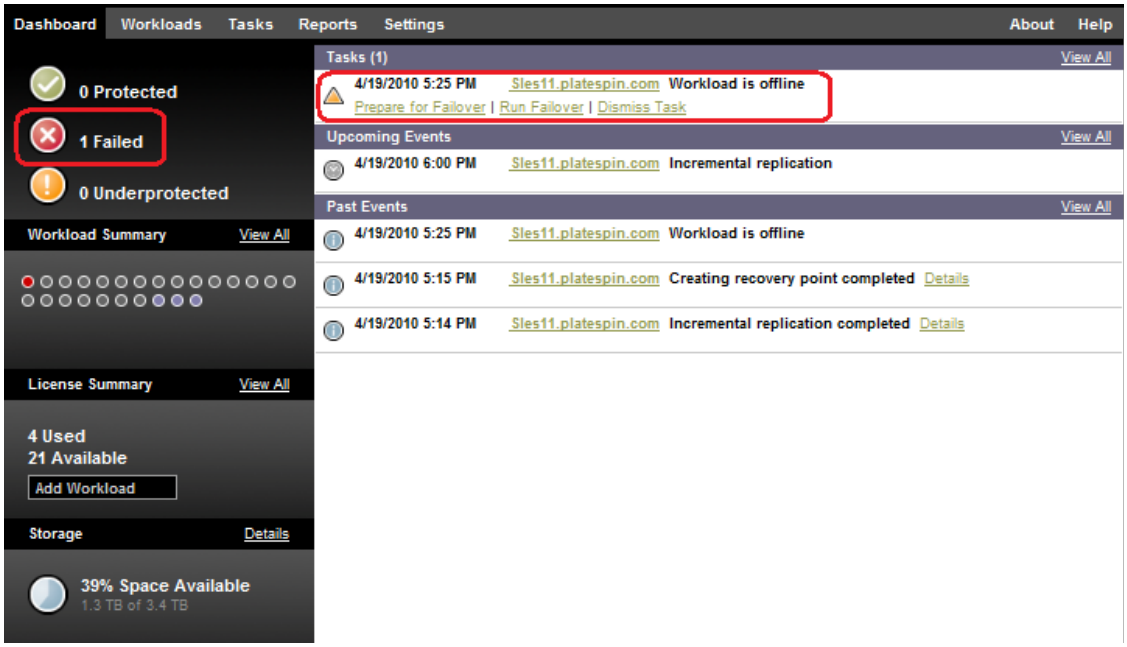

3.5.2 Performing a Failover

Failover settings, including the recovery workload’s network identity and LAN settings, are saved together with the workload’s protection details at configuration time. See the Failover row in “Workload Protection Details” on page 23.

You can use the following methods to perform a failover:

- Selecting the required workload on the Workloads page and clicking *Run Failover*. You can use the optional *Prepare for Failover* command for applying your saved failover settings to the recovery workload and booting it in advance of a full failover. Consider a separate *Prepare for Failover* operation to make sure that your production workload has indeed failed. This saves time when running a full *Failover* command.
- Clicking the appropriate command hyperlink of the *Workload is offline* event in the Tasks and Events pane. See Figure 3-1.
- Manually booting the recovery workload by using the VMware Infrastructure Client (VIC). When using this method, use the VIC's Snapshot Manager to select a snapshot (a recovery point).

  See your VMware documentation.

  NOTE: When performing a failover manually, the system applies failover settings as saved upon the workload's replication.

Use one of these methods to start the failover process and select a recovery point to apply to the recovery workload (see "Recovery Points" on page 47). Click *Execute* and monitor the progress. Upon completion, the replication status of the workload should indicate *Live*.

For testing the recovery workload or testing the failover process as part of a planned disaster recovery exercise, see "Testing the Recovery Workload and the Failover Functionality" on page 27.

3.5.3 Testing the Recovery Workload and the Failover Functionality

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

When you execute the command, PlateSpin Protect applies the Test Failover Settings, as saved in the workload protection details, to the recovery workload (see the Test Failover row in "Workload Protection Details" on page 23.

1. Define an appropriate time window for testing and make sure that there are no replications underway. The replication status of the workload must be *Idle*.
2. On the Workloads page, select the required workload, click *Test Failover*, select a recovery point (see "Recovery Points" on page 47), and the click *Execute*.

   Upon completion, PlateSpin Protect generates a corresponding event and a task with a set of applicable commands:

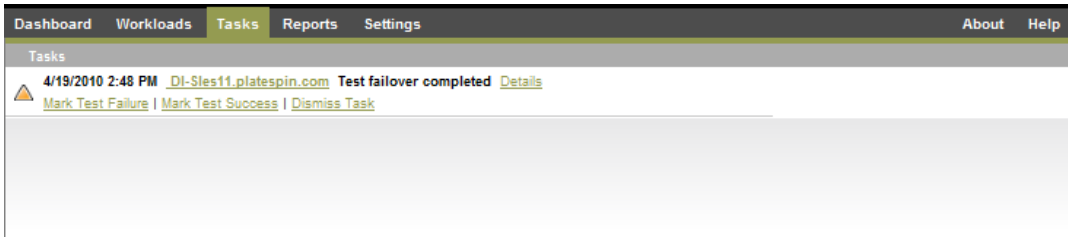

3 Verify the integrity and business functionality of the recovery workload. Use the VMware vSphere Client to access the recovery workload in the VM container.

4 Mark the test as a failure or a success. Use the corresponding commands in the task (*Mark Test Failure*, *Mark Test Success*). The selected action is saved in the history of events associated with the workload. *Dismiss Task* discards the task and the event.

Upon completion of the *Mark Test Failure* or *Mark Test Success* tasks, PlateSpin Protect discards temporary settings that were applied to the recovery workload, and the protection contract returns to its pre-test state.

3.6 Failback

A Failback operation is the next logical step after a failover; it transfers the failover workload to either its original physical or virtual infrastructure, or a new one.

3.6.1 Workload Failback to a VM Container

Use these steps to fail a workload back to a virtual machine. The VM host might be either the original infrastructure or a new one.

1 Following a failover, select the workload on the Workloads page and click *Failback / Deploy*.

2 Specify the following sets of parameters:

- **Workload Settings:** Specify the recovery workload's hostname or IP address and provide admin-level credentials. Use the required credential format (see "Guidelines for Workload and Container Credentials" on page 45.
- **Failback Target Settings -** Specify the 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 47.
  - ***Target Type*:** Select *Virtual Target*. If you don't yet have a failback container, click *Add Container* and inventory a supported VM host using root-level credentials.

3 Click *Save and Prepare* and monitor the progress on the Command Details screen.

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

4 Configure the failback details. See "Failback Details (Workload to VM)" on page 29.

5 Click *Save and Failback* and monitor the progress on the Command Details page. See Figure 3-2.

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

Figure 3-2 *Failback Command Details*

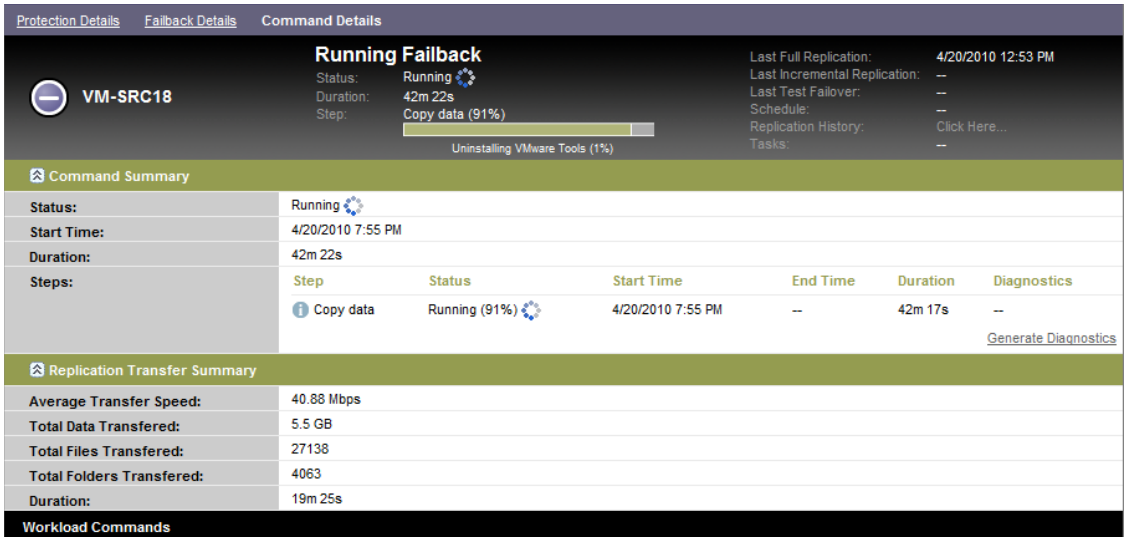

Failback Details (Workload to VM)

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

Table 3-2 *Failback Details (VM)*

Parameter Set (Settings)	Details
Failback	**Transfer Method:** (Windows) Enables you to select a data transfer mechanism and security through encryption. See "Transfer Methods and Data Transfer Security" on page 46. **Failback Network:** Enables you to direct failback traffic over a dedicated network based on virtual networks defined on your VM container. See "Networking" on page 51. **VM Datastore:** Enables you to select a datastore associated with your failback container for the target workload. **Volumes to Copy:** Enables you to select the volumes for re-creating on the target and assigning to a specific datastore. **Services/Daemons to stop:** Enables you to select Windows services or Linux daemons that are automatically stopped during the failback. See "Service and Daemon Control" on page 49.

Parameter Set (Settings)	Details
Workload	***Number of CPUs:*** Enables you to specify the required number of vCPUs assigned to the target workload. **VM Memory:** Enables you to assign the required RAM to the target workload . **Hostname, Domain/Workgroup:** Use these options to control the identity and domain/ workgroup affiliation of the target workload. For domain affiliation, domain admin credentials are required. **Network Connections:** Use these options to specify the network mapping of the target workload based on the virtual networks of the underlying VM container. **Service States to Change:** Enables you to control the startup state of specific application services (Windows) or daemons (Linux). See "Service and Daemon Control" on page 49.
Post-Failback	**Reprotect Workload:** Use this option if you plan to re-create the protection contract for the target workload after deployment. This maintains a continuous event history for the workload and auto-assigns/designates a workload license. ◆ ***Reprotect after Failback:*** Select this option if you intend to re-create a protection contract for the target workload. ◆ ***No reprotect*:** Select this option if you don't intend to re-create a protection contract for the target workload.

3.6.2 Workload 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 Protect Server. See "Registering Physical Machines with PlateSpin Protect for Failback" on page 51.
2. Run the PS Analyzer tool to determine whether any drivers are missing. See "Analyzing Workloads with PlateSpin Analyzer" on page 41.
3. If the PS Analyzer reports missing or incompatible drivers, upload the required drivers to the PlateSpin Protect device driver database. See "Managing Device Drivers" on page 42.
4. Following a failover, select the workload on the Workloads page and click *Failback / Deploy*.
5. Specify the following sets of parameters:
   - **Workload Settings:** Specify the recovery workload's hostname or IP address and provide admin-level credentials. Use the required credential format (see "Guidelines for Workload and Container Credentials" on page 45.
   - **Failback Target Settings -** Specify the following parameters:
     - ***Replication Method*** : Select the scope of data replication. See "Initial Replication Method (Full and Incremental)" on page 47.
     - ***Target Type:*** Select the *Physical Target* option and then select the physical machine you registered in Step 1.
6. Click *Save and Prepare* and monitor the progress on the Command Details screen.

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

7 Configure the failback details, then click *Save and Failback*.

Monitor the progress on the Command Details screen.

3.7 Protecting Windows Clusters

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

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

Protection of a cluster is achieved through incremental replications of changes on the active node streamed to a virtual single-node cluster, which you can use while troubleshooting the source infrastructure.

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

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

If a node failover occurs between incremental replications of a protected cluster, PlateSpin Protect generates a protection event. If the new active node's profile is similar to the failed active node, the protection schedule continues, otherwise the command fails.

To protect a Windows cluster, follow the normal workload protection workflow (see "Basic Workflow for Workload Protection and Recovery" on page 21).

On Failback, PlateSpin Protect provides validation that helps you ensure that shared volume layouts are preserved on the target. Make sure you map the volumes correctly.

Workload Image Protection

4

Imaging is one of the two workload protection infrastructures of PlateSpin Protect. For information about the virtualization infrastructure, see “Workload Protection” on page 21.

A PlateSpin Image is a static stored copy of the state of a physical or virtual machine (including volume data and the configuration details of the workload’s hardware profile, operating system, and network identity), captured at a specific point in time and regularly updated at intervals you specify in the workload’s protection settings.

Upon failure of the protected workload, you can deploy the captured image to run on physical hardware or in a VM host.

Similar to workload protection functionality with virtualization, image deployment allows for key workload configuration options, such as those for managing the workload’s disk layout, volume sizes, network identity, and domain or workgroup affiliation

4.1 Protecting a Workload Image

4.1.1 Adding a Workload for Image Protection

Protecting a workload image captures the specified workload’s volume data in PlateSpin Flexible Image format and establishes an ongoing protection of the image through incremental updates at specified intervals.

1. Make sure your workload is supported for image protection. See “Supported Workloads in Image Containers” on page 8.
2. Add an Image Container. See “Adding Containers” on page 18.
3. In the PlateSpin Protect Web Client, click *Add Workload*.
4. Specify the required information about the workload. See “Guidelines for Workload and Container Credentials” on page 45.
5. Select the required image server as the protection target.
6. Click *Add Workload*.

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

4.1.2 Configuring Workload Image Protection Details

Image protection settings determine how often a workload image is synchronized with changes on the workload's volumes, what transfer mechanism is used during replications, and which volumes of the workload are selected for protection.

1. Add a workload for image protection. See "Adding a Workload for Image Protection" on page 33.
2. On the Workloads page, select the required workload and click *Configure*.
3. Configure the required Protection Tier settings. See "Protection Tiers" on page 46.
4. Configure the image replication settings:
   - **Transfer Method and Encryption:** See "Transfer Methods and Data Transfer Security" on page 46.
   - **Source Credentials:** Specify admin-level credentials for the source workload. See "Guidelines for Workload and Container Credentials" on page 45.
   - **Protected Volumes:** Select the volumes of the workload for which you require image protection.
   - **Datastore:** Select the datastore for workload image data.
   - **Services to Stop During Replication:** See "Service and Daemon Control" on page 49.
5. Click *Save*.
6. To start the operation, click *Run Replication*, then confirm by clicking *Execute*.

   PlateSpin Protect reloads the Workloads page and displays a process indicator for the workload being replicated  .

4.2 Deploying a Workload Image

If a workload with a protected image fails, you can deploy the protected image that is stored on your image server, as a bootable workload onto a physical or virtual infrastructure. Criteria that determine and log a workload failure are part of a workload image protection contract's Tier settings (see "Protection Tiers" on page 46).

If notifications are configured along with SMTP settings, PlateSpin Protect simultaneously sends a notification e-mail to the specified recipients. See "Setting Up E-Mail Notifications" in your *Application Configuration Gude*.

4.2.1 Deploying an Image to a Virtual Target

Use this procedure to deploy an image as a bootable workload on a virtual machine.

1. Add a failback container. See "Adding Containers" on page 18.
2. In your PlateSpin Protect Web Client, select the required workload, then click *Failback / Deploy*. Select *Virtual* as the Target Type, and select your failback container as the target.
3. Click *Save and Prepare*. When prompted, specify the complete parameters for the operation:

Parameter Set (Settings)	Details
Deployment	**Transfer Method:** (Windows) Enables you to select a data transfer mechanism and security through encryption. See “Transfer Methods and Data Transfer Security” on page 46.
	Failback Network: Enables you to direct failback traffic over a dedicated network based on virtual networks defined on your VM container. See “Networking” on page 51.
	VM Datastore: Enables you to select a datastore associated with your failback container for the target workload.
	Volumes to Copy: Enables you to select the protected volumes for deployment on the selected target and to assign them to specific datastores.
Workload	**Number of CPUs:** Enables you to specify the required number of vCPUs assigned to the target workload.
	VM Memory: Enables you to assign the required RAM to the target workload.
	Hostname, Domain/Workgroup: Use these options to control the identity and domain/workgroup affiliation of the target workload. For domain affiliation, domain admin credentials are required.
	Network Connections: Use these options to specify the network mapping of the target workload based on the virtual networks of the underlying VM container.
	Service States to Change: Enables you to control the startup state of specific application services. See “Service and Daemon Control” on page 49.
Post-Failback	**Reprotect Workload:** Use if you plan to re-create the image protection contract for the target workload after deployment. This maintains a continuous event history for the workload. ◆ *Turn off Deployed Workload*: If this option is selected, the target VM is powered off upon completion of the deployment.

4 Click *Run Deployment* and monitor the progress.

Upon completion, PlateSpin Protect reports the status of the command as *Deployment completed*.

Access the target physical machine and verify its functionality and integrity.

Deployment Details (Image to VM)

Deployment details are represented by three sets of parameters that you configure when you are performing a workload image deployment to a virtual machine.

4.2.2 Deploying an Image to a Physical Target

Use this procedure to deploy an image as a bootable workload on a physical machine.

1. Register the required physical machine with your PlateSpin Protect Server. See "Registering Physical Machines with PlateSpin Protect for Failback" on page 51. Upon completion:
   - **1a** Run the PS Analyzer tool to determine whether any drivers are missing. See "Analyzing Workloads with PlateSpin Analyzer" on page 41.
   - **1b** If the PS Analyzer reports missing or incompatible drivers, upload the required drivers to the PlateSpin Protect device driver database. See "Managing Device Drivers" on page 42.
2. In your PlateSpin Protect Web Client, select the required workload, then click *Failback / Deploy*. Select *Physical* as the Target Type, and select your physical machine as the target.
3. Click *Save and Prepare*. When prompted, specify the complete parameters for the operation:

Parameter Set (Settings)	Details
Failback	**Transfer Method:** (Windows) Enables you to select a data transfer mechanism and security through encryption. See "Transfer Methods and Data Transfer Security" on page 46.
	Volumes to Copy: Enables you to select the protected volumes for deployment on the selected target.
Workload	**Hostname, Domain/Workgroup:** Use these options to control the identity and domain/workgroup affiliation of the target workload. For domain affiliation, domain admin credentials are required.
	Network Connections: Use these options to specify the LAN settings of the target workload.
	Service States to Change: Enables you to control the startup state of specific application services. See "Service and Daemon Control" on page 49.
	Partitions to Preserve: Enables you to preserve any existing partitions on the target.
	Service States to Change: Enables you to control the startup state of specific application services. See "Service and Daemon Control" on page 49.
Post-Failback	**Reprotect Workload:** Use this option if you plan to re-create the image protection contract for the target workload after deployment. This maintains a continuous event history for the workload. ◆ *Turn off Deployed Workload*: If this option is selected, the target VM is powered off upon completion of the deployment.

4. Click *Run Deployment* and monitor the progress.

   Upon completion, PlateSpin Protect reports the status of the command as *Deployment completed*.

Access the target physical machine and verify its functionality and integrity.

Deployment Details (Image to Physical)

Deployment details are represented by three sets of parameters that you configure when you are performing a workload image deployment to a physical machine.

4.3 Browsing and Extracting Image Files

During a disaster recovery effort or a business continuity exercise you can selectively restore files in your production server's file system by using backup versions of those files that are stored in images.

To do this, you can use the Image Browser utility, which enables you to browse, search, sort, and extract files from an image file or a specific image increment file.

You can work with both base images and image increments by loading either of the following:

- A base image's corresponding binary file (*volume-x*`.pkg`) or text configuration file (*image_name*`.xml`).
- An image increment's binary (*image_increment*`.pkg`) file. You cannot use an increment's text configuration file (*image_increment_name*`.xml`).

The utility enables you to work with image files in a Windows Explorer-like environment. A command line version enables you to extract files at the command line.

4.3.1 Starting the Image Browser and Loading Image Files

1. Start the Image Browser program (`ImageBrowser.exe`), which is located on your image server host, in the following directory:

   ```
   \Program Files\PlateSpin Image Server\ImageOperations
   ```

   The utility starts and displays the Open dialog box. At any time after the program's initial startup, you can load an image file by clicking *File* > *Open*.

2. In the Open dialog box, select the file type, navigate to and select the required image or image increment file, then click *OK*.

   The utility loads the required file and displays its contents in a two-pane interface.

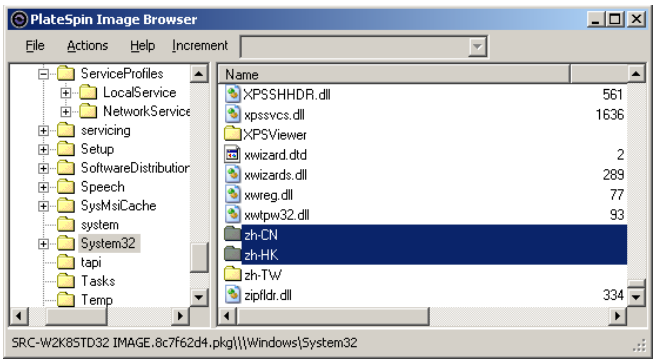

Depending on the size of the image, it might take a few seconds to several minutes for the utility to load the required file.

4.3.2 Sorting and Searching Items in the Image Browser Interface

You can sort the contents of a selected directory by name, size, type, date last modified, and by file attribute. To sort items in a selected view, click the corresponding bar at the top of the right pane.

You can search for a specific directory name or filename. You can use alphanumeric text, wildcards, and regular expressions. Regular expression search patterns that you specify must adhere to the Microsoft .NET Framework regular expression syntax requirements. See the Microsoft .NET Framework Regular Expressions page on MSDN (http://msdn.microsoft.com/en-us/library/hs600312.aspx).

To search for an item:

1. Load the required image or image increment.
2. In the left pane, select a volume or a subdirectory.
3. On the *Actions* menu, click *Search*.

   Alternatively, you can right-click the required volume or subdirectory in the left pane and click *Search* in the context menu.

   The Image Browser Search window opens.
4. Specify the name of the file you are searching. If you are using a regular expression, select the corresponding option.
5. Click *Search*.

   The results are shown in the right pane.

4.3.3 Extracting Items

1. Load the required image or image increment.
2. Locate and select the required file or directory.You can select multiple files and directories in the right pane.
3. On the *Actions* menu, click *Extract*.

   Alternatively, you can right-click the required item and click *Extract* in the context menu.

The Browse for Folder dialog box opens.

4. Browse to the required destination, then click *OK*.

   The selected items are extracted to the specified destination.

NOTE: Files that you choose to overwrite are deleted if you interrupt the extraction process.

4.3.4 Browsing and Extracting Image Files at the Command Line

To browse and extract files from images and image increments at the command line, you can use the `ImageBrowser.Console` utility.

To start the utility:

1. On your Flexible Image Server host, open a command interpreter (`cmd.exe`) and change the current directory to `\Program Files\PlateSpin Image Server\ImageOperations`.
2. At the command prompt, type `ImageBrowser.Console`, then press Enter.

   For command syntax and usage details, type `ImageBrowser.Console /help`, then press Enter.

Auxiliary Tools for Working with Physical Machines

5

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

5.1 Analyzing Workloads with PlateSpin Analyzer

Before running a workload failback or image deployment operation to a physical machine, use the PlateSpin Analyzer to identify potential driver problems and correct them beforehand.

NOTE: PlateSpin Analyzer currently supports only Windows workloads.

1. On your PlateSpin Protect Server host, start the `Analyzer.Client.exe` program, located in the following directory:

   `\Program Files\PlateSpin Protect Server\PlateSpin Analyzer`

2. Make sure that the network selection is *Default*, then select the required machine in the *All Machines* drop-down list.
3. (Optional) To reduce the analysis time, limit the scope of machines to a specific language.
4. Click *Analyze*.

   Depending on the number of inventoried workloads you select, the analysis might take a few seconds to several minutes.

Analyzed servers are listed in the left pane. Select a server to view test results in the right pane. Test results can be any combination of the following:

Table 5-1 *Status Messages in PlateSpin Analyzer Test Results*

Result	Description
Passed	The machine passed the PlateSpin Analyzer tests.
Warning	One or more tests returned warnings for the machine, indicating potential migration issues. Click the hostname to see the details.
Failed	One or more tests failed for this machine. Click the hostname to see the details and obtain more information.

The *Summary* tab provides a listing of the number of machines analyzed and not checked, as well as those that passed the test, failed the test, or were assigned a warning status.

The *Test Results* tab provides the following information:

Table 5-2 *PlateSpin Analyzer Test Results Tab*

Section	Details
System Test	Validates that the machine fulfills minimum hardware and operating system requirements.
Hardware Support	Checks the workload for hardware compatibility.
Target Hardware Support	Checks hardware compatibility for use as a target physical machine.
Software Test	Checks for applications that must be shut down for Live Transfer, and databases that should be shut down during Live Transfer to guarantee transactional integrity.
Incompatible Application Test	Verifies that applications known to interfere with the migration process are not installed on the system. These applications are stored in the Incompatible Application Database. To add, delete or edit entries in this database, select *Incompatible Application* from the *Tools* menu.

The *Properties* tab provides detailed information about a selected machine.

5.2 Managing Device Drivers

PlateSpin Protect ships with a library of device drivers and automatically installs the appropriate ones on target workloads. To determine if the required drivers are available, use the PlateSpin Analyzer utility. See "Analyzing Workloads with PlateSpin Analyzer" on page 41.

If PlateSpin Analyzer encounters missing or incompatible drivers, or if you require specific drivers for a target infrastructure, you might need to add (upload) drivers to the PlateSpin Protect driver database.

- Section 5.2.1, "Packaging Device Drivers for Windows Systems," on page 42
- Section 5.2.2, "Packaging Device Drivers for Linux Systems," on page 43
- Section 5.2.3, "Uploading Drivers to the PlateSpin Protect Device Driver Database," on page 43

5.2.1 Packaging Device Drivers for Windows Systems

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

1. Prepare the entire library of all interdependent device driver files (`*.sys`, `*.inf`, `*.dll`, etc.) for your target infrastructure.
2. Make a `.zip` archive of the collection of driver files.
3. Upload the archive to the PlateSpin Protect driver database. See "Uploading Drivers to the PlateSpin Protect Device Driver Database" on page 43.

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

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

5.2.2 Packaging Device Drivers for Linux Systems

To package your Linux device drivers for uploading to the PlateSpin Protect driver database, you can use a custom utility included in your Linux Take Control ISO boot image. See Table 6-2, "ISO Boot Images for Target Physical Machines," on page 51.

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 and mount the boot image, and from its `/tools` subdirectory, copy and extract the `ackageModules.tar.gz` archive into a another working directory.
3. 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`
4. Upload the package to the PlateSpin Protect driver database. See "Uploading Drivers to the PlateSpin Protect Device Driver Database" on page 43.

5.2.3 Uploading Drivers to the PlateSpin Protect Device Driver Database

After you have packaged the necessary device drivers in the required format, use the following procedure to upload them to the PlateSpin Protect device driver database:

1. Save the driver files to a local directory.
2. On your PlateSpin Protect Server host, start the `DriverManager.exe` program, located in the following directory:

   `\Program Files\PlateSpin Protect Server\DriverManager`

   The Driver Manager window opens.

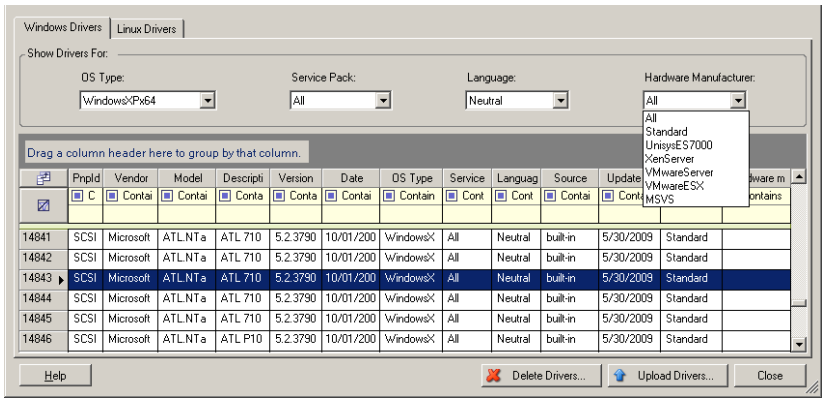

3. Select the tab corresponding to the required operating system, then select the applicable operating system and target hardware manufacturer details.
4. Click *Upload Drivers*.
5. Use the *Device Driver Uploader* dialog box to specify the path to the directory where your saved the required driver files.
6. Click *Upload*, then click *OK*.

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

Essentials of Workload Protection Details

6

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

6.1 Guidelines for Workload and Container Credentials

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

Table 6-1 *Workload and Container Credentials*

To Discover	Credentials	Remarks
All Windows workloads Image Server Containers	Local or domain admin credentials.	For the username, use this format: ◆ For domain member machines: *authority\principal* ◆ For workgroup member machines: *hostname\principal*
Windows Clusters	Domain admin credentials	Use the cluster’s virtual IP address. If you use the IP address of an individual Windows cluster node, that node is discovered as a regular (cluster-unaware) Windows workload.
All Linux workloads	Root-level username and password	Non-root accounts must be properly configured to use `sudo`. See KB Article 7920711 (http://www.novell.com/support/viewContent.do?externalId=7920711).

To Discover	Credentials	Remarks
VMware ESX hosts	Root-level username and password	

6.2 Transfer Methods and Data Transfer Security

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

- **Block-level:** Data is replicated at a volume's block level. For this transfer method, PlateSpin Protect uses a driver to monitor changes on the source workload.
  - **Windows systems:** For Windows systems, PlateSpin Protect uses a block-based component that leverages the Microsoft Volume Snapshot Service (VSS) with applications and services that support VSS. The automatic installation of the block-based component requires a reboot of the source workload. When you are configuring workload protection details, you can select the timing of the component's installation. Similarly, when removing a workload, uninstallation of the block-based component requires a reboot.
  - **Linux systems:** For the block-level transfer of Linux systems, PlateSpin Protect uses a block-based component driver that can leverage LVM snapshots (recommended). See KB Article 7005872 (http://www.novell.com/support/viewContent.do?externalId=7005872).

    The Linux block-based component included in your PlateSpin Protect distribution is precompiled for the standard, non-debug kernels of the supported Linux distributions. If you have a non-standard, customized, or newer kernel, you can rebuild the block-based component for your specific kernel. See KB Article 7005873 (http://www.novell.com/support/viewContent.do?externalId=7005873).

    Deployment or removal of the component is transparent, has no continuity impact, and requires no intervention.
- **File-level:** Data is replicated on a file-by-file basis (Windows only). Supported with or without VSS.

  In image protection and deployment operations, data is replicated at the file level without requiring explicit selection.

To make the transfer of workload data more secure, PlateSpin Protect enables you to encrypt data replication. When encryption is enabled, over-the-network data transfer from the source to the target is encrypted by using AES (Advanced Encryption Standard) or 3DES if FIPS-compliant encryption is enabled.

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

6.3 Protection Tiers

A Protection Tier is a customizable collection of workload protection parameters that define the frequency of replications and criteria for the system to consider a workload as failed; it is an integral part of every workload protection contract.

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

You can also create custom protection tiers in advance:

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

Name	Type the name you want to use for the tier.
Workload Failure	Specify the number of workload detection attempts before it is considered failed.
Workload Detection Every	Specify the time interval (in seconds) between workload detection attempts.
Recovery Points to Keep	Specify the number of recovery points to keep for workloads that use this Protection Tier. See "Recovery Points" on page 47. A `0` value disables this feature.
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.

6.4 Recovery Points

A recovery point is a point-in-time snapshot of a workload or a workload image. It allows a replicated workload or a workload image to be restored to a specific state.

For each protected workload, you can keep up to 28 recovery points.

For each protected image, you can keep up to 100 recovery points.

Recovery points that accumulate over time might cause your PlateSpin Protect storage to run out of space.

6.5 Initial Replication Method (Full and Incremental)

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

- **Full:** A full volume transfer takes place from a production workload to its replica (the recovery workload), or from a failover workload to its original virtual or physical infrastructure.

- **Incremental:** Only differences are transferred from a selected operation's source to its target, provided that they have a similar operating system and volume profile.
  - During protection: The production workload is compared with an existing VM in the VM container. The existing VM might be:
    - A previously-protected workload's recovery VM (when a *Remove Workload* command's *Delete VM* option is deselected).
    - A VM that is manually imported into the VM container, such as a workload VM physically moved, on portable media, from the production site to a remote recovery site (for VMware ESX 3.5 and later only).

      For details, see your VMware documentation.
  - During failback to a virtual machine - the failover workload is compared with an existing VM in a failback container.
  - During failback to a physical machine - the failover workload is compared with a workload on the target physical machine, if the physical machine is registered with PlateSpin Protect (see "Workload Failback to a Physical Machine" on page 30).

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

1. Proceed with the required workload command, such as *Add Workload* or *Failback.*
2. For the *Initial Replication Method* option, select *Incremental Replication*.
3. Click *Prepare Workload*.

   The PlateSpin Protect Web Client displays the Prepare for Incremental Replication page.

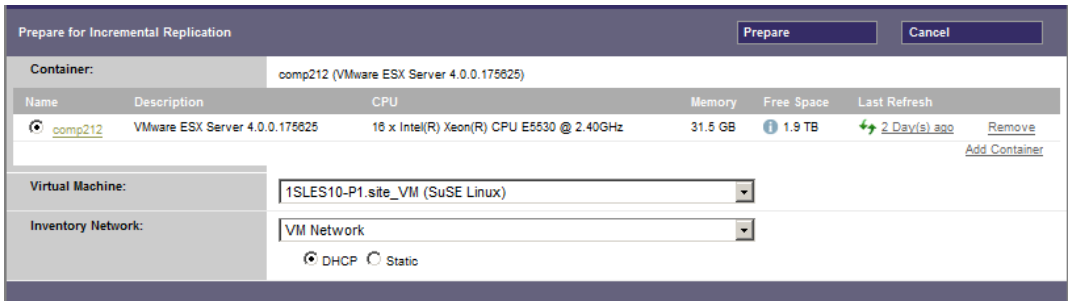

4. Select the required container, the virtual machine, and the inventory network to use for communicating with the VM.
5. Click *Prepare*.

   Wait for the process to complete and for the user interface to return to the original command, then select the prepared workload.

6.6 Service and Daemon Control

PlateSpin Protect enables you to control services and daemons:

- **Source service/daemon control:** During data transfer, you can automatically stop Windows services or Linux daemons that are running on your source workload. This ensures that the source workload is transferred to the recovery workload 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 "Automatically Executing Custom Scripts upon Every Replication (Linux)" on page 49.

- **Target startup sate/run level control:** You can select the startup state (Windows) or the run level (Linux) of services/daemons on the target workload. 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.

6.7 Automatically Executing Custom Scripts upon Every Replication (Linux)

For Linux systems, PlateSpin Protect provides you with the capability to automatically execute custom scripts: `freeze` (executed at the beginning of a replication) and `thaw` (executed at the end of a replication).

You might want to consider using this capability to complement the automated daemon control feature provided through the user interface (see "Source service/daemon control:" on page 49. 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, do the following 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 unspecified, 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`

6.8 Volumes

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

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

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

- Non-volume storage that is associated with the source workload is recreated and assigned to the recovery workload.
- The layout of volume groups and logical volumes is preserved so that you can re-create it during failback.

The following figure 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*

6.9 Networking

PlateSpin Protect enables you to control your recovery workload's network identity and LAN settings to prevent replication traffic from interfering with your main LAN or WAN traffic.

You can specify distinct networking settings in your workload protection details for use at different stages of the workload protection and recovery workflow:

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

6.10 Registering Physical Machines with PlateSpin Protect for Failback

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

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

Download the ISO image from the Novell Downloads (http://download.novell.com). Use the image appropriate for your target infrastructure:

Table 6-2 *ISO Boot Images for Target Physical Machines*

Filename	Remarks
`winperamdisk.iso`	Windows systems with 384 MB RAM or more
`winpe.iso`	Windows systems with 256 to 384 MB RAM
`bootofxx2p.iso`	All Linux systems
`winpe_cisco.iso`	Windows systems on Cisco hardware
`winpe_dell.iso`	Windows systems on Dell hardware
`winpe_fujitsu.iso`	Windows systems on Fujitsu hardware

After downloading the required file, unzip and save the extracted ISO file.

- Section 6.10.1, "Registering Target Physical Machines," on page 51

6.10.1 Registering Target Physical Machines

1. Burn the appropriate 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*.

   Because the Windows version of the boot CD image supports only *Auto Negotiate Full Duplex*, this ensures that there are no conflicts in the duplex settings.

3 Use the boot CD to boot the target physical machine, then wait for the command prompt window to open.

4 (Linux only) For 64-bit systems, at the initial boot prompt, type the following:

   - `ps64` (for systems with up to 512 MB RAM)
   - `ps64_512m` (for systems with more than 512 MB RAM)

   Press Enter.

5 When prompted, enter the following URL:

   `http://`<*hostname* | *IP_address*>`/platespinprotect`

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

6 Provide your admin-level credentials for the PlateSpin Protect Server host, specifying an authority. For the user account, use this format:

   *domain**username* or *hostname**username*

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

7 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.

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

9 Enter *Yes* if you have enabled SSL; otherwise, enter *No*.

After a few moments, the physical machine should be available in the failback/image deployment settings of the PlateSpin Protect Web Client.

Injecting Drivers into a PlateSpin Boot Image (Linux)

You can use a custom utility to package and inject additional Linux device drivers into the PlateSpin boot image (`bootofxx2p`) before burning it on a CD:

1 Obtain or compile the required `*.ko` driver filess.

   IMPORTANT: Make sure the drivers are valid for the kernel included with the ISO file (`2.6.16.21-0.8-default`) and are appropriate for the target architecture.

   ---

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 `unmount` <*mount_point*>).

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

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

`./rebuildiso.sh –i <`*ISO_file*`> -d <`*driver_dir*`> -m i586|x86_64`

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

Troubleshooting 7

7.1 Troubleshooting Workload Inventory (Windows)

The following are common problems that you might need to troubleshoot during the workload inventory.

Problems or Messages	Solutions
`The domain in the credentials is invalid or blank`	This error occurs when the Credential Format is incorrect. Try the discovery by using a local admin account with the credential format `hostname\LocalAdmin` Or try the discovery by using a domain admin account with the credential format `domain\DomainAdmin`
`Unable to connect to Windows server...Access is denied`	A non-admin account was used when trying to add a workload. Use an admin account or add the user to the administrators group and try again. This message might also indicate WMI connectivity failure. For each of the following possible resolutions, attempt the solution and then perform the “WMI Connectivity Test” on page 57 again. If the test succeeds, try adding the workload again. ◆ “Troubleshooting DCOM Connectivity” on page 57 ◆ “Troubleshooting RPC Service Connectivity” on page 57
`Unable to connect to Windows server...The network path was not found`	Network connectivity failure. Perform the “Performing Connectivity Tests” on page 56. If it fails, ensure that PlateSpin Protect and the workload are on the same network. Reconfigure the network and try again.

Problems or Messages	Solutions
`"Discover Server Details {hostname}" Failed Progress: 0% Status: NotStarted`	This error can occur for several reasons and each has a unique solution: ◆ For environments using a local proxy with authentication, bypass the proxy or add the proper permissions. See KB Article 7920339 (http://www.novell.com/support/viewContent.do?externalId=7920339) for more details. ◆ If local or domain policies restrict required permissions, follow the steps outlined in Knowledge Base article KB Article 7920862 (http://www.novell.com/support/viewContent.do?externalId=7920862).

Workload Discovery fails with error message
`Could not find file output.xml`
or
`Network path not found`
or (upon attempting to discover a Windows cluster)
`Inventory failed to discover. Inventory result returned nothing.`
	There are several possible reasons for the "Could not find file `output.xml`" error:
	◆ Anti-virus software on the source could be interfering with the discovery. Disable the anti-virus software to determine whether or not it is the cause of the problem. See "Disabling Anti-Virus Software" on page 58.
	◆ File and Printer Sharing for Microsoft Networks might not be enabled. Enable it under the Network Interface Card properties.
	◆ The C$ and/or Admin$ shares on the source might not be accessible. Ensure that PlateSpin Protect can access those shares. See "Enabling File/Share Permissions and Access" on page 58.
	◆ Change the flag `ForceMachineDiscoveryUsingService` to `true` in the `web.config` file in the `\Program Files\PlateSpin Portability Suite Server\Web` folder.
	◆ 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.

7.1.1 Performing Connectivity Tests

Network Connectivity Test

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

1. Go to your PlateSpin Protect Server host.
2. Open a command prompt and ping your workload:

```
ping workload_ip
```

WMI Connectivity Test

1. Go to your PlateSpin Protect Server host.
2. Click *Start* > *Run*, type `Wbemtest` and press `Enter`.
3. Click *Connect*.
4. In the *Namespace*, type the name of the workload you are trying to discover with `\root\cimv2` appended to it. For example, if the hostname is `win2k`, type:

   ```
   \\win2k\root\cimv2
   ```

5. Enter the appropriate credentials, using either the `hostname\LocalAdmin` or `domain\DomainAdmin` format.
6. Click *Connect* to test the WMI connection. If an error message is returned, a WMI connection cannot be established between PlateSpin Protect and your workload.

Troubleshooting DCOM Connectivity

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

Troubleshooting RPC Service Connectivity

There are three potential blockages for the RPC service:

- The Windows Service
- A Windows firewall
- A Hardware firewall

For the Windows Service, ensure that the RPC service is running on the workload. To access the services panel, run `services.msc` from a command prompt.

For a Windows firewall, add an RPC exception.

For hardware firewalls, you can try the following strategies:

- Putting PlateSpin Protect and the workload on the same side of the firewall
- Opening up specific ports between PlateSpin Protect and the workload (See "Access and Communication Requirements across your Protection Network" in your *Application Configuration Guide*).

7.1.2 Disabling Anti-Virus Software

Anti-virus software might occasionally block some of the PlateSpin Protect functionality related to WMI and Remote Registry.

In order to ensure that workload inventory is successful, it might be necessary to first disable the anti-virus service on a workload.

In addition, anti-virus 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 anti-virus software. These services are only disabled for the duration of the file transfer, and are restarted when the process completes. This is not necessary during block-level data replication.

7.1.3 Enabling File/Share Permissions and Access

To successfully protect a workload, PlateSpin Protect needs to successfully deploy and install the OFX Controller and, if you require block-level replication, a dedicated block-based component. Upon deployment of these components to a workload, as well as during the Add Workload process, PlateSpin Protect uses the workload's administrative shares. PlateSpin Protect needs administrative access to the shares, using either a local administrator account or a domain admin 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 C$ and Admin$, among other shares.

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

1. Go to your PlateSpin Protect Server host.
2. Click *Start* > *Run*, type \\<*server_host*>\C$, then click *OK*.
3. If prompted, use the same credentials as those you will use to add the workload to the PlateSpin Protect workload inventory.

   The directory is opened and you should be able to browse and modify its contents.
4. Repeat the process for all shares with the exception of the IPC$ share.

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

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

7.2 Troubleshooting Workload Inventory (Linux)

Problems or Messages	Solutions
`Unable to connect neither to the SSH server running on <IP_address> nor to VMware Virtual Infrastructure web-services at <ip_address>/sdk`	This message has a number of possible causes: ◆ The workload is unreachable. ◆ The workload does not have SSH running. ◆ The firewall is on and the required ports have not been opened. ◆ The workload's specific operating system is not supported. For network and access requirements for workload, see "Access and Communication Requirements across your Protection Network" in your *Application Configuration Guide*.
`Access denied`	Authentication problem: either invalid username or password. For information on proper workload access credentials, see "Guidelines for Workload and Container Credentials" on page 45.

7.3 Troubleshooting Problems during the Prepare Replication Command (Windows)

Problems or Messages	Solutions
Authentication error when verifying the controller connection while setting up the controller on the source.	The account used to add a workload needs to be allowed by this policy. See "Group Policy and User Rights" on page 59.

7.3.1 Group Policy and User Rights

Refresh the policy immediately by using `gpudate /force` (for Windows 2003/XP) or `secedit /refreshpolicy machine_policy /enforce` (for Windows 2000). Because of the way that PlateSpin Protect interacts with the source workload's operating system, it requires the administrator account used to add a workload 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 alternatively `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.

7.4 Troubleshooting Workload Replication

Problems or Messages	Solutions
Workload issue requires user intervention 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. The solution is to wait until the replication is complete.
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 from the source workload: `ping ip –t` (replace *ip* with the IP address of your PlateSpin Protect Server host). Allow it to run for 50 iterations and the average indicates the latency. Also see *Parameters for Optimizing Transfers over WAN Connections* in your *Application Configuration Guide*.

`The file transfer cannot begin - port 3725 is already in use`
 	
or
 	
`3725 unable to connect`
	Ensure that the port is open and listening:
	
	Run `netstat -ano` on the workload.
	
	Check the firewall.
	
	Retry the replication.

`Controller connection not established` Replication fails at the *Take Control of Virtual Machine* step.	This error occurs when the replication networking information is invalid. Either the DHCP server is not available or the replication virtual network is not routable to the PlateSpin Protect Server host. Change the replication IP to a static IP or enable the DHCP server. Ensure that the virtual network selected for replication is routable to the PlateSpin Protect Server host.

Problems or Messages	Solutions
Replication job does not start (stuck at 0%)	This error can occur for different reasons and each has a unique solution: • For environments using a local proxy with authentication, bypass the proxy or add proper permissions to resolve this problem. See Knowledge Base article KB Article 20339 (http://www.novell.com/support/viewContent.do?externalId=7920339) for more details. • If local or domain policies restrict required permissions, to resolve this problem follow the steps outlined in KB Article 7920862 (http://www.novell.com/support/viewContent.do?externalId=7920862). This is a common issue when PlateSpin Protect Server host is affiliated with a domain and the domain policies are applied with restrictions. See "Group Policy and User Rights" on page 59.

7.5 Generating and Viewing Diagnostic Reports

After you execute a command, you can generate detailed diagnostic reports about the command's details.

1 Click *Command Details*, then click the *Generate Diagnostics* link.

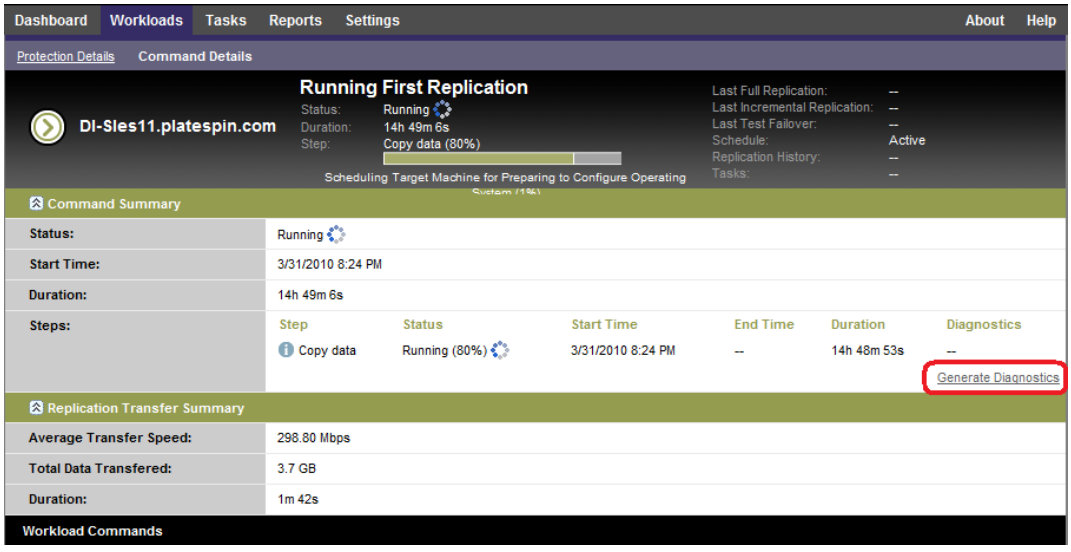

After a few moments, the page refreshes and displays a *View* link above the *Generated Diagnostics* link.

2 Click *View*.

A new page opens with comprehensive diagnostic information about the current command.

3 Save the diagnostics page and have it ready when seeking technical support.

7.6 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.

7.6.1 Cleaning Up Windows Workloads

Component	Removal Instructions
Third-party Block-based Transfer Component (discontinued)	1. Use the Windows Add/Remove Programs applet (run `appwiz.cpl`) and remove the component. Depending on the source, you might have either of the following versions: ◆ SteelEye Data Replication for Windows v6 Update2 ◆ SteelEye DataKeeper For Windows v7 2. Reboot the machine.
File-based Transfer Component	At root level for each volume under protection, remove all files named `PlateSpinCatalog*.dat`
Workload Inventory software	In the workload's `Windows` directory: ◆ Remove all files named `machinediscovery*`. ◆ Remove the subdirectory named `platespin`.
Controller software	1. Open a command prompt and change the current directory to: ◆ `\Program Files\platespin*` (32-bit systems) ◆ `\Program Files (x86)\platespin` (64-bit systems) 2. Run the following command: `ofxcontroller.exe /uninstall` 3. Remove the `platespin*` directory

7.6.2 Cleaning Up Linux Workloads

Component	Removal Instructions

Controller software	◆ In the source workload's file system, under `/boot`, remove the `ofx` directory with its contents.
	◆ Kill the OFX controller process if running:
	`pkill -9 ofxcontrollerd`
	◆ remove the OFX controller rpm package:
	`rpm -e ofxcontrollerd`

Component	Removal Instructions
Block-level data transfer software	In the source workload's file system: ◆ Under `/lib/modules/`*`kernel_version`*, remove the `platespin` directory with its contents ◆ Under `/etc`, remove the `blkwatch.conf` file
LVM snapshots	1. In the Jobs view, generate a Job Report for the failed job, then note the name of the snapshot. 2. Remove the snapshot device by using the following command: `lvremove` *snapshot_name*
Bitmap files	For each volume under protection, at the root of the volume, remove the corresponding `.blocks_bitmap` file.
Tools	On the source workload, under `/sbin`, remove the following files: ◆ `bmaputil` ◆ `blkconfig`

Glossary

container
Either of the two workload protection infrastructures supported by PlateSpin Protect (a VM host or an image server).

Deploy
PlateSpin Protect command that recovers a failed workload by deploying its image, which is under protection on an image server, to run on physical hardware or a VM container.

Failback
The restoration of the business function of a failed workload in its original environment when the business function of a temporary recovery workload within PlateSpin Protect is no longer required.

Failover
The taking over of the business function of a failed workload by a recovery workload within a PlateSpin Protect VM container.

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

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

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

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

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

recovery point objective (RPO)
Tolerable data loss measured in time and defined by a configurable interval between incremental replications of a protected workload or a protected workload image.

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

recovery workload
A protected workload's bootable virtual replica.

replication
1. The creation of an initial base copy of a workload (*initial replication*). 2. Any transfer of changed data from a protected workload to its replica in the container.

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

Reprotect
A PlateSpin Protect command that reestablishes a protection contract for a workload following the Failover and Failback operations.

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

In image deployment, it is the image of a workload on a designated image server.

See also target.

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

In image deployment, it is the infrastructure onto which a protected image is deployed to boot.

See also source.

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

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

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