

PlateSpin Migrate 12.1 Release Notes

May 2016



PlateSpin Migrate version 12.1 includes new features, enhancements, and bug fixes.

Many of these improvements were made in direct response to suggestions from our customers. We thank you for your time and valuable input. We hope you continue to help us ensure that our products meet all your needs. You can post feedback in the [PlateSpin Migrate forum](#) on NetIQ Communities, our online community that also includes product information, blogs, and links to helpful resources.

The documentation for this product is available on the NetIQ Website in HTML and PDF formats on a page that does not require you to log in. If you have suggestions for documentation improvements, click **comment on this topic** at the bottom of any page in the HTML version of the PlateSpin Migrate 12.1 documentation posted at the [NetIQ Documentation Website](#).

This product contains undocumented utilities that the Technical Support team might use to diagnose or correct problems.

For documentation that accompanied earlier releases, visit the [PlateSpin Migrate 12.1 Documentation Web Site](#) and scroll to *Previous Releases*.

- ◆ [Section 1, "What's New?," on page 1](#)
- ◆ [Section 2, "Installing PlateSpin Migrate 12.1," on page 4](#)
- ◆ [Section 3, "Upgrading to PlateSpin Migrate 12.1," on page 5](#)
- ◆ [Section 4, "Software Fixes," on page 5](#)
- ◆ [Section 5, "Known Issues," on page 8](#)
- ◆ [Section 6, "Legal Notice," on page 27](#)

1 What's New?

The following sections outline the key features and functions provided in this release:

1.1 Support for Migration of Workloads to Cloud

PlateSpin Migrate 12.1 enhances the Web Interface to let you migrate the following Windows and Linux workloads to Microsoft Azure:

Windows:

- ◆ Microsoft Windows Server 2012 R2
- ◆ Microsoft Windows Server 2012
- ◆ Microsoft Windows Server 2008 R2

Linux:

- ◆ Red Hat Enterprise Linux (RHEL) 7.1
- ◆ Red Hat Enterprise Linux (RHEL) 6.7

- ◆ SUSE Linux Enterprise Server (SLES) 11 SP4
- ◆ SUSE Linux Enterprise Server (SLES) 11 SP3

NOTE

- ◆ Migration of Windows cluster workloads is not supported because Microsoft Azure does not support Windows clusters.
 - ◆ Migration of UEFI workloads is not supported.
 - ◆ The PlateSpin Migrate Client does not support migration of workloads to Microsoft Azure. You can use only the PlateSpin Migrate Web Interface to migrate the workloads to Microsoft Azure.
 - ◆ Test Cutover of workloads is not supported. You can perform only Run Cutover of workloads.
 - ◆ PlateSpin Migrate supports Azure VM sizes with up to 64 data disks. For the maximum instance size in a selected Azure Region, Migrate will use one data disk for the OS disk replication in the PlateSpin Replication Environment. After migration, this disk becomes the OS disk, and you can add a data disk.
Each data disk must have a maximum size of 1TB (1024 GB).
 - ◆ Migrate recommends an Azure VM instance size that meets or exceeds the source workload's settings for cores, memory, data disks, and NICs. However, you can choose a smaller or larger instance size based on your requirements for the target workload, as limited by the maximum instance sizes available in the selected Azure Region.
 - ◆ The size of the disk created on the Azure VM is the size of the source disk partition plus about 1 GB because of the granularity of available disk sizes on Azure.
 - ◆ You need an OS license for the migrated target workload. For Azure target workloads, you must provide Azure with the license information or Microsoft will charge you for the OS license.
 - ◆ For each target Azure subscription, you must enable programmatic deployment for the PlateSpin Migrate Replication Environment VM. See [Enabling an Azure Subscription to Deploy the Replication Environment VM](#).
 - ◆ Currently, when the time on the PlateSpin Server goes out of sync, the cutover will fail with a 403 forbidden error. Can we either detect that the root cause is the time problem, and state that as the error message or add that to it - or otherwise build in some sort of very visible warning that the time is out of sync
 - ◆ Ensure that the PlateSpin Server host displays the correct time for the time zone it is in. If the time on the PlateSpin Server host is incorrect, the cutover process fails with a 403 forbidden error.
-

1.2 Support for Migration of Workloads to Target VMs on Microsoft Hyper-V Hosts Using Migrate Command Line Interface

PlateSpin Migrate 12.1 enhances the Migrate Command Line Interface to let you migrate workloads to VMs on target Microsoft Hyper-V hosts in addition to the existing support for migration of workloads to target VMs on VMware host. See the [Using the PlateSpin Migrate Command Line Interface](#) section in the *PlateSpin Migrate 12.1 User Guide*.

1.3 Support for Synchronization of Discovered Workloads and Target Hosts between the Migrate Client and the Web Interface

PlateSpin Migrate 12.1 enhances the Migrate Client to automatically synchronize workloads and target hosts it discovers to the Web Interface.

1.4 Support for Ensuring Data Consistency between Source and Target at Cutover

PlateSpin Migrate 12.1 introduces an option that lets you to permanently stop Windows services on the source workload throughout the cutover process to ensure application data consistency between the source workload and the target machine. These services are not restored even after the cutover process is complete. See the [Configuring the Workload for Migration](#) section in the *PlateSpin Migrate 12.1 User Guide*.

1.5 Workload Support

PlateSpin Migrate 12.1 includes support for the following workloads and containers:

- ◆ **Windows Workloads**
 - ◆ Windows Server 2012 R2 Cluster
- ◆ **Linux Workloads**
 - ◆ CentOS 7
 - ◆ Red Hat Enterprise Linux (RHEL) 7.2, 7.1 (Only BIOS-based workloads are supported)
 - ◆ Red Hat Enterprise Linux (RHEL) 6.7 (Only BIOS-based workloads are supported)
 - ◆ SUSE Linux Enterprise Server (SLES) 11 SP4 (Only BIOS-based workloads are supported)

For more information about the supported workloads and containers, see the “[Supported Configurations](#)” section in the *PlateSpin Migrate 12.1 User Guide*.

1.6 Utility to Install Block-Based Transfer Drivers on a Windows Machine

PlateSpin Migrate 12.1 introduces a new command line utility (`MigrateAgent.cli.exe`) that lets you install, upgrade, query, or uninstall the block-based transfer drivers.

Although a reboot is always required when you install, uninstall, or upgrade drivers, this utility allows you to better control when the action occurs and therefore, when the server reboots. For example, you can use the utility to install the drivers during scheduled down time, instead of during the first replication.

For information about the utility, see the “[MigrateAgent Utility](#)” section in the *PlateSpin Migrate 12.1 User Guide*.

1.7 Hyper-V Enhancements

PlateSpin Migrate 12.1 includes the following enhancements for target VMs on Microsoft Hyper-V hosts:

- ◆ [Section 1.7.1, “Option to Define Hyper-V Target Virtual Machine Generation Type,”](#) on page 4
- ◆ [Section 1.7.2, “Option to Define the Virtual Network ID for a Target VM on a Hyper-V Host,”](#) on page 4
- ◆ [Section 1.7.3, “Option to Change the Adapter Type Used during the Target Take Control Process of a Workload Migration to a Target VM on a Hyper-V Host,”](#) on page 4

1.7.1 Option to Define Hyper-V Target Virtual Machine Generation Type

PlateSpin Migrate 12.1 introduces a Virtual Machine Generation Type configuration option for VMs on Hyper-V hosts that lets you select one of the following generation types for the new virtual machine:

- ♦ **Generation 1:** Allows you deploy the target virtual machine with Hyper-V BIOS architecture.
- ♦ **Generation 2:** Allows you to deploy the target virtual machine with Hyper-V UEFI architecture.

See the [Virtual Machine Configuration: Microsoft Hyper-V](#) section in the *PlateSpin Migrate 12.1 User Guide*.

1.7.2 Option to Define the Virtual Network ID for a Target VM on a Hyper-V Host

PlateSpin Migrate 12.1 introduces a new VLAN ID option that allows you to specify the virtual network ID to be used on the target VM on a Hyper-V host. If you do not specify this ID, then the virtual network ID of the source machine is used by default.

See the [Post-Migration Networking for Virtual Network Interfaces \(Windows and Linux\)](#) section in the *PlateSpin Migrate 12.1 User Guide*.

1.7.3 Option to Change the Adapter Type Used during the Target Take Control Process of a Workload Migration to a Target VM on a Hyper-V Host

PlateSpin Migrate 12.1 allows you to edit the adapter type used during the Target Take Control process of workload migration to a target VM on a Hyper-V Host.

See the [Changing the Adapter Type Used during the Target Take Control Process of Workload Migration to a Target VM on a Hyper-V Host](#) section in the *PlateSpin Migrate 12.1 User Guide*.

1.8 VMware Enhancements

PlateSpin Migrate 12.1 includes the following enhancements for a target VM on a VMware host:

- ♦ Enhanced support for migration of workloads to a VMware DRS cluster.
- ♦ Support for specifying the number of CPU sockets and the number of CPU cores per socket for the target VM.

1.9 Security

The GLIBC version in this release addresses vulnerability [CVE 2015-7547](https://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2015-7547) (<https://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2015-7547>), a stack-based buffer overflow in the `getaddrinfo()` function in the `glibc` DNS client-side.

2 Installing PlateSpin Migrate 12.1

To install PlateSpin Migrate 12.1, see “[Installing PlateSpin Migrate](#)” in the *PlateSpin Migrate 12.1 Installation and Upgrade Guide*.

3 Upgrading to PlateSpin Migrate 12.1

To upgrade your PlateSpin Server to PlateSpin Migrate 12.1, you must have an existing installation of PlateSpin Migrate 12.0 Hotfix 1. See “[Upgrading PlateSpin Migrate](#)” in the *PlateSpin Migrate 12.1 Installation and Upgrade Guide*.

Other direct upgrades are not supported. For information about upgrading from PlateSpin Migrate 12.0 to PlateSpin Migrate 12.0 Hotfix 1, see the [PlateSpin Migrate 12.0 Hotfix 1 Release Notes](#).

4 Software Fixes

The following is a list of bugs that were fixed for this release:

- ◆ [Section 4.1, “Migration and Server Sync Jobs Might Fail If the Distributed Resource Scheduler Is Enabled in vCenter Server,” on page 5](#)
- ◆ [Section 4.2, “Migrated Linux Targets Have Partitions Created As Separate Disks,” on page 6](#)
- ◆ [Section 4.3, “Install Utility’s Launcher Needs to Refresh after Completion of “Install PlateSpin Server,”” on page 6](#)
- ◆ [Section 4.4, “Unable to Execute Post-Migration Scripts on a Linux Workload,” on page 6](#)
- ◆ [Section 4.5, “Unable to Add Any Workload or Target Due to Continuously Refreshing Web Interface,” on page 6](#)
- ◆ [Section 4.6, “Job Stuck at Copy Data with Source behind NAT,” on page 6](#)
- ◆ [Section 4.7, “Discovering Objects Through Web Interface Might Fail With A Warning Message,” on page 6](#)
- ◆ [Section 4.8, “Synchronization Issue between the Source Workloads and Target Hosts Discovered Using the Migrate Client and the Migrate Web Interface,” on page 7](#)
- ◆ [Section 4.9, “Performing a Test Cutover of a Workload Resulted in a Run Cutover,” on page 7](#)
- ◆ [Section 4.10, “Changes to the VM Memory in the Target Workload Settings and Target Workload Test Settings Is Not Effective,” on page 7](#)
- ◆ [Section 4.11, “Scheduled Full Replication Does Not Start When the Time Occurs,” on page 8](#)
- ◆ [Section 4.12, “Web UI Throws Exception When Saving Config Page After Re-Selecting a Volume Group on Linux Configuration Page,” on page 8](#)
- ◆ [Section 4.13, “Failure to Mount NSS Volumes,” on page 8](#)
- ◆ [Section 4.14, “\(VMware 4.1\) Poor Networking Performance By Traffic-Forwarding VMs,” on page 8](#)
- ◆ [Section 4.15, “Failure With Access Denied Error During Replication to an Image Stored on a Network Share,” on page 8](#)

4.1 Migration and Server Sync Jobs Might Fail If the Distributed Resource Scheduler Is Enabled in vCenter Server

Issue: If the Distributed Resource Scheduler (DRS) is enabled in a vCenter Server or at the cluster level, then the migration and server sync jobs might fail with some object reference errors.

Fix: When you migrate workloads to VMware cluster, the VMware DRS and VMware HA is set as `disabled`. You must not change the state of the VMware DRS or the VMware HA for the target VM during the entire migration process.

4.2 Migrated Linux Targets Have Partitions Created As Separate Disks

Issue: If you migrate a Linux workload having partitions, each partition is created as a separate disk on the migrated Linux target.

Fix: When you migrate a Linux workload having partitions, the migrated Linux target contains the same partitions as on the source workload.

4.3 Install Utility's Launcher Needs to Refresh after Completion of “Install PlateSpin Server”

Issue: After a successful installation of the PlateSpin Migrate software, the PlateSpin Migrate Installation Launcher is not refreshed, and the **Install PlateSpin Server** button is not dimmed/deactivated to confirm the detection of the installed software. (Bug 969435)

Fix: The **Install PlateSpin Server** button is automatically dimmed after a successful install.

4.4 Unable to Execute Post-Migration Scripts on a Linux Workload

Issue: Post-migration scripts failed to execute on a Linux workload. (Bug 895957)

Fix: Post-migration scripts are now successfully executed on a Linux workload.

4.5 Unable to Add Any Workload or Target Due to Continuously Refreshing Web Interface

Issue: The Web Interface refresh interval for the Workloads page and the Targets page was too short to allow a workload or target to be added. (Bug 971850)

Fix: The default refresh interval for the Workloads page and the Targets page was changed from 15 seconds to 30 seconds. The interval is now configurable. For a custom interval, modify the value for the following setting in the `\Program Files\PlateSpin Migrate Server\PlateSpin Forge\web\web.config` file:

```
<add key="WorkloadTargetsUpdateIntervalSeconds" value="30" />
```

4.6 Job Stuck at Copy Data with Source behind NAT

Issue: A source workload in a NAT environment was added using its NAT public IP address; however, the workload's NICs were mapped only to private IP addresses and its NAT public IP address was unknown to the source operating system. (Bug 961985)

Fix: When a source workload is in a NAT environment, you can configure the target workload to use the source machine's NAT public IP address as the first address to try in a NAT IP-pinning scenario when connecting to the source machine for replication.

4.7 Discovering Objects Through Web Interface Might Fail With A Warning Message

Issue: When you use the Web Interface to discover workloads and targets, the discovery might fail with a warning message. (Bugs 946132 and 970592)

Fix: A default heartbeat delay of 15 seconds (15000 ms) is set on the controller.

To enable a heartbeat delay of shorter or longer duration, do the following:

- 1 On the source workload, open the registry editor.
- 2 Go to the following location in the registry editor, depending on the operating system architecture on the source workload:

Path for a 64-bit source workload:

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

Path for a 32-bit source workload:

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

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

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

This registry key is not configured by default.

- 4 Restart the source workload.

4.8 Synchronization Issue between the Source Workloads and Target Hosts Discovered Using the Migrate Client and the Migrate Web Interface

Issue: The source workloads and target hosts that you discovered using the PlateSpin Migrate Client did not display in the PlateSpin Migrate Web Interface.

Fix: The source workloads and targets that you discover in the default network using the Migrate Client are automatically synchronized and displayed in the Web Interface.

4.9 Performing a Test Cutover of a Workload Resulted in a Run Cutover

Issue: If you chose to perform a Test Cutover on a workload with **Perform Incremental Replication** option selected, it resulted in a Run Cutover. (Bug 940244)

Fix: Performing Test Cutover of a workload with **Perform Incremental Replication** option selected no longer results in a Run Cutover.

4.10 Changes to the VM Memory in the Target Workload Settings and Target Workload Test Settings Is Not Effective

Issue: If you use the PlateSpin Migrate Web Interface to configure the migration settings for a workload and specify a value for **VM Memory** in the Target Workload Settings and the Target Workload Test Settings sections, the specified value is not effective. However, the default source value continues to apply. (Bug 940013)

Fix: The value you specify for **VM Memory** in the Target Workload Settings and Target Workload Test Settings sections when configuring the migration settings is now effective.

4.11 Scheduled Full Replication Does Not Start When the Time Occurs

Issue: In PlateSpin Migrate 12.0, a schedule set for a full replication might not be honored in some circumstances. (Bug 971849)

Fix: Schedules set in PlateSpin 12.1 will be honored. However, some existing schedules might be in a broken state after you upgrade from version 12.0 Hotfix1 to version 12.1. If **Next Replication** shows an empty schedule, you must re-configure the schedule for that workload.

4.12 Web UI Throws Exception When Saving Config Page After Re-Selecting a Volume Group on Linux Configuration Page

Issue: On a Linux workload configured with volume groups and logical volumes, if you save the Config page with an LVM **Volume Group** deselected, and then edit the Config page again to re-select the **Volume Group**, the Web Interface throws an exception when you save the change. (Bug 970767)

Fix: You can now edit the Config page to re-select the **Volume Group** without any issues.

4.13 Failure to Mount NSS Volumes

Issue: After migration is completed, NSS volumes with snapshots enabled are not automatically mounted as expected. (Bug 655828)

Fix: See [KB Article 7008773](#).

4.14 (VMware 4.1) Poor Networking Performance By Traffic-Forwarding VMs

Issue: In some scenarios, the replica of a workload that is forwarding network traffic (for example, if the workload's purpose is to serve as a network bridge for NAT, VPN, or a firewall) might show significant network performance degradation. This is related to a problem with VMXNET 2 and VMXNET 3 adapters that have LRO (Large Receive Offload) enabled. (Bug 680259)

Fix: Disable LRO on the virtual network adapter. See [KB Article 7005495](#).

4.15 Failure With Access Denied Error During Replication to an Image Stored on a Network Share

Issue: The Controller service on Image servers that use network shares for storage does not preserve the service `Log On As` credentials after an upgrade. Image operations fail with an Access Denied message until the controller service is updated with the correct `Log On As` credentials. (Bug 685509)

Fix: See [KB Article 7008772](#).

5 Known Issues

- ◆ [Section 5.1, "General Issues," on page 9](#)
- ◆ [Section 5.2, "Known Issues for Upgrade," on page 20](#)
- ◆ [Section 5.3, "Rediscovering vCenter Server in Migrate Web Interface Post Upgrade Fails With An Exception," on page 20](#)

- ◆ Section 5.4, “Known Issues For Migration to Azure,” on page 20
- ◆ Section 5.5, “Known Issues For Migration to VMware,” on page 25

5.1 General Issues

The following issues are being researched:

- ◆ Section 5.1.1, “Remove Target Does Not Display a Confirmation Dialog in a German Language Browser,” on page 11
- ◆ Section 5.1.2, “Help On Chinese Locales Displays in English Language,” on page 11
- ◆ Section 5.1.3, “Test Cutover Fails When the Hostname Setting on the Configuration Page of a Target Linux Workload Contains an Underscore,” on page 11
- ◆ Section 5.1.4, “Incremental File-Based Replication Does Not Complete with Encryption Enabled,” on page 11
- ◆ Section 5.1.5, “Web Interface: Validation Message Is Not Displayed after Executing the “Remove Workload” with “Preserve Source”,” on page 11
- ◆ Section 5.1.6, “Installation Fails When the OS Language and OS Locale Do Not Match on the Computer,” on page 12
- ◆ Section 5.1.7, “Option to Install Block-Based Drivers During Prepare Replication Is Enabled When Configuring Workload Migration Even If the Block-Based Drivers Are Already Installed,” on page 12
- ◆ Section 5.1.8, “Inconsistent Behavior between Azure and VMware Targets when Deselecting LVM Volume Group but Not the Corresponding Logical Volume,” on page 12
- ◆ Section 5.1.9, “Mapping Volumes Not Supported When Migrating Linux Workloads,” on page 12
- ◆ Section 5.1.10, “Unable to Migrate Linux workloads Having Volumes Created on Raw Disks Without Partitions,” on page 13
- ◆ Section 5.1.11, “Unable to Migrate a Workload to Hitachi LPAR,” on page 13
- ◆ Section 5.1.12, “Warning Message Displayed When You Migrate a Workload To Hitachi LPAR,” on page 13
- ◆ Section 5.1.13, “Unable to Install PlateSpin Migrate on a Windows Server 2012 and Windows Server 2012 R2 Computer,” on page 13
- ◆ Section 5.1.14, “Discovered Hyper-V Container Displays As a Workload in the PlateSpin Migrate Web Interface,” on page 13
- ◆ Section 5.1.15, “Unable to Migrate a Linux Workload to a Container That Does Not Support the Source Workload Firmware,” on page 13
- ◆ Section 5.1.16, “(Windows Sources) Non-Default Per-Volume VSS Settings Not Preserved After Migration,” on page 14
- ◆ Section 5.1.17, “(ESX 4.1) No Warning or Error Message On Wrong vCPU Selection,” on page 14
- ◆ Section 5.1.18, “Special Character in Datastore Name Causes Migration Problems,” on page 14
- ◆ Section 5.1.19, “Preserving Boot Partition Causes Migration Problems,” on page 14
- ◆ Section 5.1.20, “(Linux to ESX 4) Problem Completing Migration When the Source OS Has Autologin or CD Automount Features Enabled,” on page 14
- ◆ Section 5.1.21, “Executing a Post-Migration Script with Unicode Characters in the Filename Fails,” on page 14
- ◆ Section 5.1.22, “VSS Snapshots Are Not Preserved,” on page 14

- ◆ Section 5.1.23, "Migration Over WAN Taking a Long Time When the Target VM Host Has a High Number of Datastores," on page 15
- ◆ Section 5.1.24, "Unmapped Home Directory is Disabled And Unmounted After One Time Server Sync," on page 15
- ◆ Section 5.1.25, "VMware tools Are Not Installed during Windows Server 2012 Core Conversion," on page 15
- ◆ Section 5.1.26, "Network Card Not Initialized on SLES 11 Target VM Hosted on a Windows Server 2008 Hyper-V Host," on page 15
- ◆ Section 5.1.27, "Target VM Does Not Boot After Migration from VMware ESX to Citrix Xen If Boot Files Are Located in Second Disk," on page 15
- ◆ Section 5.1.28, "XenServer Tools Not Removed After Conversion," on page 15
- ◆ Section 5.1.29, "Post Migration the Primary Partition Is Converted to a Logical Partition on the Target," on page 16
- ◆ Section 5.1.30, "Machine Node on ESX Host Is Not Undiscovered When Migrate Undiscovers a Machine," on page 16
- ◆ Section 5.1.31, "Attempt to Retrieve Data from VMware vCenter Server Fails," on page 16
- ◆ Section 5.1.32, "V2P Conversion Hangs at the Configuring Operating System Step," on page 16
- ◆ Section 5.1.33, "UEFI to BIOS Conversion for Windows 8.1 Workload Fails at the Sending Files Step," on page 16
- ◆ Section 5.1.34, "Conversion Fails While Downgrading from UEFI to BIOS Firmware," on page 17
- ◆ Section 5.1.35, "File-Based Transfer Breaks for Windows Server 2012 R2 UEFI Workload," on page 17
- ◆ Section 5.1.36, "Image Capture of a Windows 32-Bit OS Fails," on page 17
- ◆ Section 5.1.37, "Source Machine Stays in an Under Control State After Offline Conversion," on page 17
- ◆ Section 5.1.38, "Source Machine Boot Configuration Not Restored After Offline Conversion," on page 17
- ◆ Section 5.1.39, "Creating and Moving a VM Under Resource Pool as a Setting Is Not Supported in the CLI Tool," on page 17
- ◆ Section 5.1.40, "Partitions Not Mounted to Drive Letters After Conversion," on page 18
- ◆ Section 5.1.41, "Adding Disk and Volume Mapping Does Not Work Properly for a Conversion of a Workload to Windows Server 2012 R2 Hyper-V," on page 18
- ◆ Section 5.1.42, "Redundant Disks Are Present After a RHEL 6.2 x64 Block Migration to Windows Server 2012 R2 Hyper-V," on page 19
- ◆ Section 5.1.43, "Requirements for VMware DRS Cluster support," on page 19
- ◆ Section 5.1.44, "Installation of the PlateSpin Image Server Fails When the Fully Qualified Image File Path is More Than 248 Characters," on page 19
- ◆ Section 5.1.45, "Conversion Job Fails to Configure NICs on a Windows 2000 Server Target Machine," on page 19
- ◆ Section 5.1.46, "Rediscovering ESX Server in Migrate Client Displays a Duplicate Server Entry With Add Failed Error Message in the Web Interface," on page 19

5.1.1 Remove Target Does Not Display a Confirmation Dialog in a German Language Browser

Issue: If the Languages option of your web browser is set to German, the Web Interface does not display a confirmation dialog when you click **Remove** next to a target in the Targets list. This issue applies for both VMware targets and Azure targets. The Web Interface displays the Remove Target confirmation dialog for web browsers with Languages set to English or the other supported national languages (French, Japanese, Chinese Simplified, and Chinese Traditional). (Bug 978490)

Workaround: If you use a German language browser, use caution when you select to remove targets. You can alternatively modify the web browser's Languages setting to use a different supported language.

5.1.2 Help On Chinese Locales Displays in English Language

On Chinese locales, help is displayed in the English language.

5.1.3 Test Cutover Fails When the Hostname Setting on the Configuration Page of a Target Linux Workload Contains an Underscore

Issue: For target Linux workloads, the Test Cutover fails with the following message if the specified hostname on the Configuration page contains an underscore:

```
Failed to configure virtual machine
```

(Bug 975854)

Workaround: Using an underscore in the hostname is generally not supported by Linux platforms. Modify the hostname to use only supported characters a to z, 0 to 9, and hyphen (-), then try again.

5.1.4 Incremental File-Based Replication Does Not Complete with Encryption Enabled

Issue: After you enable Encryption for a Windows workload that is configured for file-based data transfer, the Windows receiver might hang at the end of the transfer for incremental replications. The hang occurs if the last byte read of the transfer is incorrectly set by the encryption process to a non-zero value, indicating that more files are being transferred and to continue reading from the stream.

(Bug 944559)

Workaround: You can use block-based data transfer for Windows workloads if you want to enable Encryption for replication data transfers.

5.1.5 Web Interface: Validation Message Is Not Displayed after Executing the "Remove Workload" with "Preserve Source"

Issue: After you remove a workload with the options **Remove Workload** and **Preserve Source**, the discovered workload is in a Not Configured state, and the target location is cleaned up. The validation history relates only to the most recent operation, which was **Remove Workload**. The validation history for the previous workload discovery is no longer available. (Bug 971118)

Workaround: Before you remove the workload, note the information from the discovery validation by taking a screen shot or copying the message to another location.

5.1.6 Installation Fails When the OS Language and OS Locale Do Not Match on the Computer

Issue: If you choose to install PlateSpin Migrate on a computer where the OS Language setting does not match the OS Locale setting, the installation fails. (Bug 939805)

Workaround: To successfully install PlateSpin Migrate, ensure that the OS Language setting matches the OS Locale setting on the computer. You can change the locale of the computer as per your requirement after the installation is complete.

For example, if the OS Language setting is English, then you must ensure that the OS Locale setting is also English when you install the English or localized version of PlateSpin Migrate. After the installation is complete, you can change the locale as per your requirement.

5.1.7 Option to Install Block-Based Drivers During Prepare Replication Is Enabled When Configuring Workload Migration Even If the Block-Based Drivers Are Already Installed

Issue: When you configure migration for a workload that has block-based drivers already installed, the **Install During Prepare Replication** option in the Transfer Method setting is enabled and selected by default. This option must be disabled because the block-based drivers are already installed. However, there is no loss in the functionality. (Bug 967018)

Workaround: Ignore the option.

5.1.8 Inconsistent Behavior between Azure and VMware Targets when Deselecting LVM Volume Group but Not the Corresponding Logical Volume

Issue: For a Linux workload with LVM, if you deselect the volume group but do not deselect the corresponding logical volume, the behavior is inconsistent for the following targets:

- ♦ **VMware:** A validation error/message is displayed and user intervention is required to fix the incorrect configuration.

```
The LVM volume group <vg_name> assigned to <logical_volume_name> is missing in the target.
```

- ♦ **Azure:** The corresponding logical volume is automatically deselected upon save.

(Bug 973926)

Workaround: If you deselect a volume group, you should also deselect its corresponding logical volume.

5.1.9 Mapping Volumes Not Supported When Migrating Linux Workloads

Issue: When you use the PlateSpin Migrate Client to migrate Linux workloads, the following are not supported: (Bug 930355)

- ♦ Mapping the boot volume to LVM
- ♦ Mapping any volume to an existing Volume Group
- ♦ Mapping any volume to new Volume Group
- ♦ Re-mapping Volume group to disk

5.1.10 Unable to Migrate Linux workloads Having Volumes Created on Raw Disks Without Partitions

Issue: PlateSpin Migrate does not support migrating Linux workloads that have volumes created on raw disks without partitions. (Bug 937071)

5.1.11 Unable to Migrate a Workload to Hitachi LPAR

Issue: When you migrate a workload to Hitachi LPAR that has an operating system running on it, the migration might not complete. This is because the migration job waits for user intervention during the **Configure Target Machine** step of migration. (Bug 902489)

Workaround: Modify the UEFI Boot Order of Hitachi LPAR to enable it to boot from the hard disk instead of the ISO image.

5.1.12 Warning Message Displayed When You Migrate a Workload To Hitachi LPAR

Issue: When you migrate a workload to Hitachi LPAR, a warning message similar to the following might get displayed: (Bug 917209)

```
Device 'Unassigned Hitachi Shared FC Device 3017' is not supported by .....
```

Workaround: Ignore the message.

5.1.13 Unable to Install PlateSpin Migrate on a Windows Server 2012 and Windows Server 2012 R2 Computer

Issue: On a Windows Server 2012 or Windows Server 2012 R2 computer, if you disable UAC through the Control Panel and then install PlateSpin Migrate on the computer, the prerequisites check utility displays an error that the UAC is still enabled. This is because when we disable UAC from the Control Panel, the change is not reflected in the corresponding registry key. (Bug 929511)

Workaround: To disable UAC on a Windows Server 2012 or a Windows Server 2012 R2 computer, see “Windows Server 2012: Deactivating UAC” in the [Microsoft TechNet Wiki](#).

5.1.14 Discovered Hyper-V Container Displays As a Workload in the PlateSpin Migrate Web Interface

Issue: If you use the PlateSpin Migrate Web Interface to discover a Hyper-V container, the Hyper-V container is listed as a workload in the interface. You must not migrate this Hyper-V container. (Bug 929978)

Workaround: You must not migrate this Hyper-V container.

5.1.15 Unable to Migrate a Linux Workload to a Container That Does Not Support the Source Workload Firmware

Issue: The migration of a Linux workload fails in the following scenarios because UEFI to BIOS conversion and vice versa is not supported: (Bug 937070)

- ♦ Migrating a Linux workload with UEFI firmware to a container that supports BIOS firmware.
- ♦ Migrating a Linux workload with BIOS firmware to a container that supports UEFI firmware.

5.1.16 (Windows Sources) Non-Default Per-Volume VSS Settings Not Preserved After Migration

Issue: The non-default per-volume VSS Settings are not preserved post migration of a Windows workload. (Bug 493589)

Workaround: After the migration, you must re-configure your custom VSS settings.

5.1.17 (ESX 4.1) No Warning or Error Message On Wrong vCPU Selection

Issue: If the number of the requested vCPUs exceeds the number of physical CPUs on the ESX 4.1 host, the requested number is ignored and the target VM is created with a single vCPU without a warning. (Bug 505426)

Workaround: Ensure that your vCPU selection does not exceed the number of physical CPUs on the ESX 4.1 host server.

5.1.18 Special Character in Datastore Name Causes Migration Problems

Issue: Migration operations might fail when they are attempted on ESX datastores that have the “+” or other special characters in the datastore name. (Bug 506154)

See [KB Article 7009373](#).

5.1.19 Preserving Boot Partition Causes Migration Problems

Issue: In some migration scenarios, the system improperly allows you to preserve your boot partition on the target, preventing the proper workload from booting. (Bug 595490)

Workaround: Do not opt to preserve your boot partition on the target.

5.1.20 (Linux to ESX 4) Problem Completing Migration When the Source OS Has Autologin or CD Automount Features Enabled

Issue: When the Source OS has Autologin or CD Automount features enabled, the migration is affected. The migration is also affected if you log in to the target during the job’s Configuration step. (Bug 604320)

Workaround: Disable the autologin and CD automount features on the source; avoid logging in to the target workload prior to the completion of the migration.

5.1.21 Executing a Post-Migration Script with Unicode Characters in the Filename Fails

Issue: If you use Unicode characters in the filename of your post-migration script, the script fails to execute. (Bug 619942)

Workaround: Use only ASCII characters when naming a post-migration action.

5.1.22 VSS Snapshots Are Not Preserved

Issue: VSS snapshots taken by third-party applications on the source workload are not replicated to the target upon migration. (Bug 692680)

5.1.23 Migration Over WAN Taking a Long Time When the Target VM Host Has a High Number of Datastores

Issue: Under some circumstances, when your Migrate server is connected to the VM host over a WAN, and if your VM host has a high number of datastores, the process of locating the appropriate ISO image required for booting the target might take longer than expected. (Bug 702152)

Workaround: To optimize the bandwidth available, plan the migration to occur during non-peak traffic hours.

5.1.24 Unmapped Home Directory is Disabled And Unmounted After One Time Server Sync

Issue: If you perform a server sync and then unmap the `/home` partition to `none`, the partition `/home` directory should be mounted and enabled on the target server, instead it is disabled and unmounted. (Bug 779194)

Workaround: Following the Server Sync, uncomment the appropriate line in the `/etc/fstab` file of the target server. See [KB Article 7014638](#).

5.1.25 VMware tools Are Not Installed during Windows Server 2012 Core Conversion

Issue: VMware tools are not installed during the conversion of a Windows Server 2012 Core (Bug 810460)

Workaround: Install the VMware tools manually after the conversion.

5.1.26 Network Card Not Initialized on SLES 11 Target VM Hosted on a Windows Server 2008 Hyper-V Host

Issue: If you perform a SLES 11 workload (cloned VM) migration using the semi-automated method to a target VM (faked physical) on a Windows Server 2008 Hyper-V host, the process freezes at the **Configuring OS** step. (Bug 822601)

Workaround: See [KB 7012911](#).

5.1.27 Target VM Does Not Boot After Migration from VMware ESX to Citrix Xen If Boot Files Are Located in Second Disk

Issue: When a VM is converted from VMware ESX to Citrix Xen and its boot files are allocated in second disk, the VM does not boot and manual intervention is requested. This is because the Citrix XEN VM tries to boot with disk 0 rather than with the bootfiles allocated to disk 2. (Bug 824724)

Workaround: To resolve this problem, rearrange the virtual-disk position in XenCenter so that the virtual machine boots from the virtual disk containing the operating system. [The knowledge article at the Citrix Web site \(http://support.citrix.com/servlet/KbServlet/download/32320-102-691310/xcm-10-guide.pdf\)](http://support.citrix.com/servlet/KbServlet/download/32320-102-691310/xcm-10-guide.pdf) includes information about how to change the position of the virtual disk containing the operating system. See [KB Article 7012906](#).

5.1.28 XenServer Tools Not Removed After Conversion

Issue: XenServer tools on a Windows VM in a Citrix XenServer hypervisor environment are not removed when the VM is converted to a VMware container or a physical container. (Bug 825016)

Workaround: Manually uninstall the XenServer tools after conversion.

5.1.29 Post Migration the Primary Partition Is Converted to a Logical Partition on the Target

Issue: Consider the following scenario:

Scenario: Moving or copying a Windows OS machine with more than three primary partitions to a physical machine where a Windows OS has been installed with minimum 3 primary partitions. At least one primary partition is preserved in the target machine. (Bug 825434)

Effect: After the migration, the Windows OS machine is unable to boot.

Example: The following error occurs when Windows Server 2003 machine is converted to Physical machine:

```
Windows could not start because the following file is missing or corrupt:  
<Windows root>\system32\ntoskrnl.exe.Please re-install a copy of the above file.
```

Workaround: See [KB Article 7012913](#).

5.1.30 Machine Node on ESX Host Is Not Undiscovered When Migrate Undiscovers a Machine

Issue: When you undiscover a workload, it displays as such in the Migrate client, but the ESX host shows that the node is not undiscovered (Bug 826545)

Workaround: Undiscover the workload on the ESX host, then refresh the ESX host.

5.1.31 Attempt to Retrieve Data from VMware vCenter Server Fails

Issue: An attempt to retrieve data from VMware vCenter Server fails with the following exception:
Permission to perform this operation was denied. (Bug 839329)

Workaround: This problem can be corrected by following the procedures to define VMware Roles with tools as described in “[Using Tools to Define VMware Roles](#)” in the *PlateSpin Migrate 12.1 User Guide*.

5.1.32 V2P Conversion Hangs at the Configuring Operating System Step

Issue: When there are multiple boot options in the firmware and the hard disk is not the first boot device in the boot options list, the target machine does not boot from hard disk and conversion hangs. (Bug 859440)

Workaround: In the boot options of the physical machine, change the boot order so that **Hard Drive** is the first option, then restart the machine. See also [KB Article 7014623](#).

5.1.33 UEFI to BIOS Conversion for Windows 8.1 Workload Fails at the Sending Files Step

Issue: The default OEM installation of Windows 8.1 (UEFI) creates a recovery partition with insufficient free space, making it impossible to create a Volume Shadow Copy (VSS) for the partition. (Bug 864325)

Workaround: Remove or expand the recovery partition. See [KB Article 7014696](#).

5.1.34 Conversion Fails While Downgrading from UEFI to BIOS Firmware

Issue: The conversion of a UEFI workload (Windows 6.2 and above kernel versions) to BIOS-based machine fails at the **Preparing OS** step because the active partition cannot be found to update boot parameters. (Bug 864326)

Workaround: Update the partition type of **Disk as MBR** where the system volume is present in either the source workload or the image. Use Export and Import of UI options or OFX Browser to edit the XML. For a complete list of steps, see [KB Article 7014637](#).

5.1.35 File-Based Transfer Breaks for Windows Server 2012 R2 UEFI Workload

Issue: X2P File-based transfer of Windows 6.2 and above kernel versions fails during the sending and receiving files stage. (Bug 865570)

Workaround: To force file transfer to work in this X2P scenario, disable the CPU advanced flags in the firmware: `VT-d, VT-s, Execute Disable Bit`. See [KB Article 7014698](#).

5.1.36 Image Capture of a Windows 32-Bit OS Fails

Issue: Migrate expects a folder named `C:\Windows\Boot\EFI` to be present in the source server for exporting content for future use. The folder is not present in Windows 32-bit operating systems earlier than Windows Server 2008 or Windows Vista, so when Migrate exports BCD information to the folder, the operation fails with the error:

```
Error message: Failed: C:\Windows\Boot\EFI
```

(Bug 866467)

Workaround: Create the `C:\Windows\Boot\EFI` folder, then create a Directory Junction under `C:\Windows` for `C:\Windows\System32`. See [KB Article 7014710](#).

5.1.37 Source Machine Stays in an Under Control State After Offline Conversion

Issue: If you configure the `End State` setting of an offline conversion job as `Restart`, the source machine remains in an `under control` state after the job completes successfully. (Bug 875562)

Workaround: Manually restart the source when the conversion completes.

5.1.38 Source Machine Boot Configuration Not Restored After Offline Conversion

Issue: The boot menu of a Windows source machine is not restored after offline conversion. (Bug 878043)

Workaround: After the conversion, the source boot menu displays two options: the Linux RAM disk (LRD) and the Operating System (OS). When you boot for the first time after the conversion, manually select the OS option. This action purges the boot menu of the LRD boot option in future boot operations.

5.1.39 Creating and Moving a VM Under Resource Pool as a Setting Is Not Supported in the CLI Tool

Issue: The command line interface (CLI) tool does not currently support moving or creating a *VM under resource pool* as a setting in the `conversion.ini` file. (Bug 891690)

Workaround: After the conversion, manually move the new machine to the resource pool you want.

5.1.40 Partitions Not Mounted to Drive Letters After Conversion

Issue: Following a conversion to Windows Server 2012 R2 Hyper-V, only the "C" drive is visible. Other partitions are not mounted to drive letters. (Bug 894623)

Workaround: After the conversion, go to disk management and manually assign the drive letters to the partitions.

5.1.41 Adding Disk and Volume Mapping Does Not Work Properly for a Conversion of a Workload to Windows Server 2012 R2 Hyper-V

Issue: Booting the Windows Server 2012 R2 Hyper-V VM with LRD returns randomly listed devices in the Hard Disk Devices List, whether IDEs, SCSIs, or a mix of both. (Bug 896584)

Workaround: The list should contain IDE disks at the top followed by SCSI disks. Use the Migrate Client to customize the list.

The following scenarios provide examples of the list behavior. **Assumptions in these scenarios:** The target VM is Generation 1. You need to create three or more virtual disk drives:

Scenario 1-- IDE to SCSI Behavior

Given initial setting:

Disk2: IDE

Disk3: IDE

- ◆ If Disk2 changes to SCSI, Disk3 changes to SCSI. List settings after the modification display as:

Disk2: SCSI

Disk3: SCSI

- ◆ If Disk3 changes to SCSI, Disk2 does not change. List settings after the modification display as:

Disk2: IDE

Disk3: SCSI

Scenario 2-- SCSI to IDE Behavior

Given initial setting:

Disk2: SCSI

Disk3: SCSI

- ◆ If Disk2 changes to IDE, Disk3 does not change. List settings after the modification display:

Disk2: IDE

Disk3: SCSI

- ◆ If Disk3 changes to IDE, Disk2 changes to IDE. List settings after the modification display:

Disk2: IDE

Disk3: IDE

5.1.42 Redundant Disks Are Present After a RHEL 6.2 x64 Block Migration to Windows Server 2012 R2 Hyper-V

Issue: After performing a successful RHEL 6.2 x64 block-based migration with the `Install Integration Services` option selected, running the `fdisk -l` command shows redundant disks. That is, a single disk is displayed twice as `sda` and `sdb`. (Bug 896598)

Workaround: This is a known Microsoft issue and is being addressed.

5.1.43 Requirements for VMware DRS Cluster support

Issue: PlateSpin Migrate supports VMware Clusters with and without DRS enabled, and with any level of DRS (`Manual`, `Partially Automated`, or `Fully Automated`). However, to be a valid migration target, your VMware Cluster must be discovered via vCenter and not by directly inventorying individual ESX servers.

See “[Discovery Guidelines for Machine Types and Credentials](#)” in your *User Guide*.

5.1.44 Installation of the PlateSpin Image Server Fails When the Fully Qualified Image File Path is More Than 248 Characters

Issue: If you choose to designate a machine as a PlateSpin Image Server and specify the fully qualified path of an image file that is more than 248 characters, the installation of the image server fails. (Bug 967414)

Workaround: Ensure that the fully qualified path of the specified image file is less than or equal to 248 characters.

5.1.45 Conversion Job Fails to Configure NICs on a Windows 2000 Server Target Machine

Issue: If you choose to migrate a Windows 2000 Server source that has one or more NICs, the conversion job fails to configure the NICs on the target machine. (Bug 971414)

Workaround: Manually configure the NICs after the conversion job is complete.

5.1.46 Rediscovering ESX Server in Migrate Client Displays a Duplicate Server Entry With Add Failed Error Message in the Web Interface

Issue: When you use the Migrate Client to discover an ESX Server either directly or via vCenter server, the discovered server is listed in the Web Interface. If you rediscover the same ESX server again through the Migrate Client, a duplicate entry of the server with an `Add Failed` error is displayed in the Web Interface. (Bug 975870)

Workaround: Delete the duplicate server entry from the Web Interface.

5.2 Known Issues for Upgrade

The following issue is being researched:

- ◆ [Section 5.2.1, “Scheduled Full Replication Does Not Start When the Time Occurs,” on page 20](#)

5.2.1 Scheduled Full Replication Does Not Start When the Time Occurs

Issue: In PlateSpin Migrate 12.0, a schedule set for a full replication might not be honored in some circumstances. (Bug 971849)

Workaround: After you upgrade from version 12.0 Hotfix 1 to version 12.1, some existing schedules might be in a broken state. If **Next Replication** shows an empty schedule, you must re-configure the schedule for that workload. Schedules set in PlateSpin 12.1 will be honored.

5.3 Rediscovering vCenter Server in Migrate Web Interface Post Upgrade Fails With An Exception

Issue: Consider that you have used the Migrate Client to discover a vCenter server in more than one network prior to upgrading to Migrate 12.1. If you now upgrade to Migrate 12.1, the vCenter server discovered in the Migrate Client is synchronized with the Web Interface. If you now choose to add the same vCenter target using the Web Interface, the operation fails with a `Could not add target` exception. (Bug 977577)

5.4 Known Issues For Migration to Azure

The following issues are being researched:

- ◆ [Section 5.4.1, “Cutover Fails During Partitioning Operations,” on page 21](#)
- ◆ [Section 5.4.2, “Network Connections on the Target VM Might Not Be Mapped to the Correct Azure Virtual Network Or Subnet,” on page 21](#)
- ◆ [Section 5.4.3, “After Successful Cutover, Azure Portal Web UI Should Show Target VM's Computer Name and DNS Name,” on page 22](#)
- ◆ [Section 5.4.4, “Azure Target VMs Are Created with One Additional Hard Disk,” on page 22](#)
- ◆ [Section 5.4.5, “For the Maximum Azure Instance Size, Need a UI Check to Limit the Data Disks for Replication to 63 Data Disks,” on page 22](#)
- ◆ [Section 5.4.6, “Azure to Azure Migration Might Delete the Source Workload If You Select Same Target and Datastore,” on page 22](#)
- ◆ [Section 5.4.7, “Disk Numbers and DiskIndex Numbers Are Not Sequential for Discovered Dynamic Disk Workloads,” on page 22](#)
- ◆ [Section 5.4.8, “Dynamic Disk Workload Default Cloud Instance Size Is Too Large,” on page 23](#)
- ◆ [Section 5.4.9, “Virtual Disk Ordering Incorrectly Changes When a Disk Is Excluded for Migration on the Config Page,” on page 23](#)
- ◆ [Section 5.4.10, “Migration to Azure failed for SLES 11 SP 4 with 3 Disks and No LVM,” on page 23](#)
- ◆ [Section 5.4.11, “Linux Disks or Partitions Are Migrated on to Target in a Different Order than that of the Linux Source,” on page 23](#)
- ◆ [Section 5.4.12, “LVM Volume Groups Are Created on Opposite Partitions within the Same Disk on Linux Target VM,” on page 23](#)

- ◆ [Section 5.4.13, “Linux Partitions Are Created on Opposite Partitions within the Same Disk on Linux Target VM,”](#) on page 24
- ◆ [Section 5.4.14, “Azure Target VM Launched in Safe Mode After Successful Cutover of a Workload,”](#) on page 24
- ◆ [Section 5.4.15, “The Virtual Machine Settings Page of the Target VM in the Azure Portal Does Not Display the Size of the VM,”](#) on page 24
- ◆ [Section 5.4.16, “Replication Environment Blobs Are Not Automatically Cleaned Up after ‘Cutover’ or ‘Remove Workload’ or ‘Abort’,”](#) on page 24

5.4.1 Cutover Fails During Partitioning Operations

Issue: A race condition can occur during the disk partitioning where PlateSpin Migrate attempts to read the partition table before all of the partition information has been returned by the `npart` utility. This condition causes the cutover to fail with the following message:

```
Unable to find a device to map the windows volume
(Bug 959079)
```

Workaround: Re-run the cutover.

5.4.2 Network Connections on the Target VM Might Not Be Mapped to the Correct Azure Virtual Network Or Subnet

Issue: If you migrate a Windows workload that has multiple NICs with DHCP to Azure, the network connection on the target VM might not get mapped to the correct Azure network or subnet.

For example: Consider that the source workload has three NICs with DHCP and the network connections are mapped as follows:

- ◆ Ethernet mapped to subnet4
- ◆ Ethernet2 mapped to subnet3
- ◆ Ethernet3 mapped to subnet2

After migration, the network connections on the target VM might get mapped as follows:

- ◆ Ethernet mapped to subnet3
- ◆ Ethernet2 mapped to subnet4
- ◆ Ethernet3 mapped to subnet2

(Bug 967316)

Workaround: You cannot use the target workload in this state. You must remove the faulty target workload, then repeat the migration setup and cutover. A second attempt typically succeeds in properly mapping the DHCP NICs to their correct Azure virtual network or subnet assignments.

To remove the faulty target workload and repeat the migration:

- 1 In the Web Interface, select the Migrated workload, click **Remove Workload**, select the **Preserve Source** and **Delete Target VM** options, then click **Execute**.

This action removes the target workload and leaves the source workload in the `Not Configured` state.

- 2 Configure the source workload for migration to Azure, using the same settings as in the first attempt.

3 Click **Run Cutover**.

4 Verify that the NICs are assigned to their correct Azure virtual network or subnet assignments.

5.4.3 After Successful Cutover, Azure Portal Web UI Should Show Target VM's Computer Name and DNS Name

Issue: In the Microsoft Azure Portal, the workload's **DNS name** field and **Properties > COMPUTER NAME** field are empty. (Bug 969489)

Workaround: Use Remote Desktop to log in to the computer, then open the Control Panel to the System page to view the workload's computer name and DNS name.

5.4.4 Azure Target VMs Are Created with One Additional Hard Disk

Issue: Azure automatically adds a disk to the VM that is not mounted with drive letter. The size of this disk varies, depending on the cloud instance you choose for deployment. (Bug 967314)

Workaround: You can remove the extra Azure disk from the VM configuration.

5.4.5 For the Maximum Azure Instance Size, Need a UI Check to Limit the Data Disks for Replication to 63 Data Disks

Issue: PlateSpin Migrate supports Azure VM sizes with up to 64 data disks. For the maximum instance size, Migrate will use one data disk for the OS disk replication in the PlateSpin Replication Environment. After migration, this disk becomes the OS disk, and you can add a data disk. If you submit a source workload for replication with 64 data disks, there is no warning in the UI and the replication fails. (Bug 972053)

Workaround: For the maximum Azure instance size, ensure that your source workload has 63 or fewer data disks when you submit the workload for replication.

5.4.6 Azure to Azure Migration Might Delete the Source Workload If You Select Same Target and Datastore

Issue: If you mistakenly try to migrate an Azure source workload to the same target and datastore location in Azure, the migration setup fails as expected with the error `BlobAlreadyExists`. The post-failure cleanup of the target workload deletes the source workload because they are the same location. (Bug 971998)

Workaround: Do not migrate an Azure source workload to the same target and datastore location in Azure.

5.4.7 Disk Numbers and DiskIndex Numbers Are Not Sequential for Discovered Dynamic Disk Workloads

Issue: For Windows source workloads with dynamic disk types of Simple, Spanned, Striped, Mirrored, and RAID5, the target workload configuration assigns nonsequential numbers in disk names and disk indexes. The non-sequential numbering is an artifact of the types of dynamic disks on the source workload. All necessary disks are present for the target workload. This issue occurs for target Azure workloads and target VMware workloads. (Bug 973266)

Workaround: There is no workaround.

5.4.8 Dynamic Disk Workload Default Cloud Instance Size Is Too Large

Issue: For Windows source workloads with dynamic disk types of Simple, Spanned, Striped, Mirrored, and RAID5, the default Cloud Instance Size for the target workload might be larger than is necessary. The default setting for the data disks component is based on the total disk count on the source workload rather than on the net number of disks to be created on the target workload. (Bug 973265)

Workaround: You must manually change the Cloud Instance Size to the appropriate and desired setting. A Cloud Instance Size with fewer data disks might fit your needs.

5.4.9 Virtual Disk Ordering Incorrectly Changes When a Disk Is Excluded for Migration on the Config Page

Issue: A discovered workload lists all discovered disks on the Configuration page. If you exclude a disk from the migration and save the change, the vdisk list with corresponding disk path is reordered and the expected disk might not be the one excluded. This problem is observed for target VMware and Azure VMs. (Bug 969639)

Workaround: This is a cosmetic modification of the configuration in the UI. The underlying configuration is saved correctly and there is no need for user modification of the disk paths or disk ordering.

5.4.10 Migration to Azure failed for SLES 11 SP 4 with 3 Disks and No LVM

Issue: Migration to Azure failed for a SUSE Linux Enterprise Server (SLES) 11 SP 4 workload with only non-LVM disks. Different errors occurred in different attempts:

```
Command failed. Refer to command details.
```

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

Migrations to Azure are successful in workloads running other versions of SLES and other Linux operating systems that are configured with only non-LVM disks. (Bug 972062)

Workaround: There is no workaround for SLES 11 SP4 workloads using non-LVM disks.

5.4.11 Linux Disks or Partitions Are Migrated on to Target in a Different Order than that of the Linux Source

Issue: Linux disks are created in a different order on the target workload than they are on the source workload. All disks and partitions are present; only the disk order differs. (Bug 974156)

Workaround: No workaround available. The target VM is fully functional.

5.4.12 LVM Volume Groups Are Created on Opposite Partitions within the Same Disk on Linux Target VM

Issue: On a Linux workload with multiple LVM volume groups on the same disk, the LVM volume groups are created in the opposite order on the target workload. For example, if the source volume group order is AB, the target volume group order is BA. This issue occurs for target workloads on Azure and on VMware. (Bug 973227)

Workaround: The order of the LVM volume groups on the disk does not impact functionality. The target machine works as expected.

5.4.13 Linux Partitions Are Created on Opposite Partitions within the Same Disk on Linux Target VM

Issue: On a Linux workload with multiple Linux partitions on the same disk, the partitions are created in the opposite order on the target workload. For example, if the source partition order is AB, the target partition order is BA. This issue occurs for target workloads on Azure and on VMware. (Bug 970822)

Workaround: The order of the Linux partitions on the disk does not impact functionality. The target machine works as expected.

5.4.14 Azure Target VM Launched in Safe Mode After Successful Cutover of a Workload

Issue: If you choose to migrate a Windows Small Business Server 2011 workload to Azure, the cutover completes but the target VM in Azure is launched in Safe Mode. (Bug 978131)

Workaround: To boot the target VM in Normal Mode:

1. Run `msconfig`.
2. Uncheck the **Boot > Safe boot** option.
3. Reboot the VM.

5.4.15 The Virtual Machine Settings Page of the Target VM in the Azure Portal Does Not Display the Size of the VM

Issue: After successful cutover of a workload to Azure, the Virtual Machine Settings page of the Azure Portal does not display the size of the Azure VM if the VM belongs to `DSX_v2` series. Although the VM size is not shown on the settings page, the underlying VM configuration contains the VM size. (Bug 977497)

Workaround: You can see the VM size in the Azure CLI for VMs in the `DSX_v2` series.

5.4.16 Replication Environment Blobs Are Not Automatically Cleaned Up after 'Cutover' or 'Remove Workload' or 'Abort'

Issue: For migrations to Microsoft Azure, the Azure Blob Service creates storage artifacts (a page blob and a block blob) in the assigned datastore for the workload's replication environment. PlateSpin Migrate no longer needs the artifacts after the workload is successfully cut over, aborted, or removed; however, it does not automatically remove the artifacts. Microsoft Azure charges you to store these unneeded data files. (Bug 977308)

Workaround: After a workload migration is completed, aborted, or removed, you should manually remove the related storage artifacts in the `vhds` storage container of the datastore assigned to the migration. Do not remove any blob file where the related migration is in progress.

To remove the old blobs:

1. Log into [Azure Portal Web UI](#) and manually delete the blobs.
2. Go to **Storage Accounts > *datastore-name* > Services > Blobs > vhds**.
3. Delete the page blob and block blob for the source workload's replication environment.

`<source-hostname>-RepEnv.<guid>.status` (the Block blob)

`<source-hostname>-RepEnvOS<guid>.vhd` (the Page blob)

For example, for a source workload named `TST-2K12-SBS`, the blob file names are:

5.5 Known Issues For Migration to VMware

The following issues are being researched:

- ◆ [Section 5.5.1, “Two Disks Share the Same Disk Path on Config Page for Dynamic Disk Workloads,” on page 25](#)
- ◆ [Section 5.5.2, “Disk Numbers and DiskIndex Numbers Are Not Sequential for Discovered Dynamic Disk Workloads,” on page 25](#)
- ◆ [Section 5.5.3, “Virtual Disk Ordering Incorrectly Changes When a Disk Is Excluded for Migration on the Config Page,” on page 26](#)
- ◆ [Section 5.5.4, “LVM Volume Groups Are Created on Opposite Partitions within the Same Disk on Linux Target VM,” on page 26](#)
- ◆ [Section 5.5.5, “Linux Partitions Are Created on Opposite Partitions within the Same Disk on Linux Target VM,” on page 26](#)
- ◆ [Section 5.5.6, “Cutover Hangs with VMware vCDROM Locked Message in vSphere Client; User Intervention Required,” on page 26](#)
- ◆ [Section 5.5.7, “User Intervention Might Be Required When Using the Migrate Client To Configure OS For the Target Linux Workloads,” on page 27](#)

5.5.1 Two Disks Share the Same Disk Path on Config Page for Dynamic Disk Workloads

Issue: For Windows source workloads with dynamic disk types of Simple, Spanned, Striped, Mirrored, and RAID5, the target workload configuration might use the same DiskPath setting for two disks. The migration ignores the duplicate setting and configures unique paths for the two disks on the target workload. The migration completes successfully. (Bug 973271)

Workaround: No action is required for the configuration.

5.5.2 Disk Numbers and DiskIndex Numbers Are Not Sequential for Discovered Dynamic Disk Workloads

Issue: For Windows source workloads with dynamic disk types of Simple, Spanned, Striped, Mirrored, and RAID5, the target workload configuration assigns nonsequential numbers in disk names and disk indexes. The non-sequential numbering is an artifact of the types of dynamic disks on the source workload. All necessary disks are present for the target workload. This issue occurs for target Azure workloads and target VMware workloads. (Bug 973266)

Workaround: There is no workaround.

5.5.3 Virtual Disk Ordering Incorrectly Changes When a Disk Is Excluded for Migration on the Config Page

Issue: A discovered workload lists all discovered disks on the Configuration page. If you exclude a disk from the migration and save the change, the vdisk list with corresponding disk path is reordered and the expected disk might not be the one excluded. This problem is observed for target VMware and Azure VMs. (Bug 969639)

Workaround: This is a cosmetic modification of the configuration in the UI. The underlying configuration is saved correctly and there is no need for user modification of the disk paths or disk ordering.

5.5.4 LVM Volume Groups Are Created on Opposite Partitions within the Same Disk on Linux Target VM

Issue: On a Linux workload with multiple LVM volume groups on the same disk, the LVM volume groups are created in the opposite order on the target workload. For example, if the source volume group order is AB, the target volume group order is BA. This issue occurs for target workloads on Azure and on VMware. (Bug 973227)

Workaround: The order of the LVM volume groups on the disk does not impact functionality. The target machine works as expected.

5.5.5 Linux Partitions Are Created on Opposite Partitions within the Same Disk on Linux Target VM

Issue: On a Linux workload with multiple Linux partitions on the same disk, the partitions are created in the opposite order on the target workload. For example, if the source partition order is AB, the target partition order is BA. This issue occurs for target workloads on Azure and on VMware. (Bug 970822)

Workaround: The order of the Linux partitions on the disk does not impact functionality. The target machine works as expected.

5.5.6 Cutover Hangs with VMware vCDROM Locked Message in vSphere Client; User Intervention Required

Issue: For target Linux workloads on VMware containers, after the data copy completes and the configuration service begins, the cutover hangs with the following message in the Web Interface:

```
The ReconfigVM_Task submitted to VMware vCenter server failed:  
Connection control operation failed for disk ide0:0
```

In the vSphere Client for the target environment, the Virtual Machine Message dialog and Virtual Machine Question dialog prompt you to override the CD-ROM lock. In the Web Interface, the cutover hang continues until you manually override the vCDROM lockout in the vSphere Client for the target environment.

This issue does not occur on all target Linux workloads or on all VMware container versions. (Bug 975853)

Workaround: Log in to the vSphere Client for the target environment. When you are prompted to override the CD-ROM lock, select **Yes**, then click **OK**.

5.5.7 User Intervention Might Be Required When Using the Migrate Client To Configure OS For the Target Linux Workloads

Issue: If you use the Migrate Client to configure OS for target Linux workloads on VMware containers, the Migrate Client might not respond.

In the vSphere Client for the target environment, the Virtual Machine Message dialog and Virtual Machine Question dialog prompt you to override the CD-ROM lock. The Migrate Client might not respond until you manually override the vCDROM lockout in the vSphere Client for the target environment.

This issue does not occur on all target Linux workloads or on all VMware container versions. (Bug 975853)

Workaround: Log in to the vSphere Client for the target environment. When you are prompted to override the CD-ROM lock, select **Yes**, then click **OK**.

6 Legal Notice

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

Copyright © 2016 NetIQ Corporation. All Rights Reserved.

