PlateSpin Migrate 12.2 Release Notes

September 2017

PlateSpin Migrate version 12.2 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 12.2 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 12.2 Documentation Web Site and scroll to *Previous Releases.*

- Section 1, “What’s New (September 2017),” on page 1
- Section 2, “What’s New for PlateSpin Migrate 12.2 Patch 1,” on page 1
- Section 3, “What’s New for PlateSpin Migrate 12.2 RabbitMQ Fix,” on page 2
- Section 4, “What’s New for PlateSpin Migrate 12.2,” on page 2
- Section 5, “Installing PlateSpin Migrate 12.2,” on page 4
- Section 6, “Upgrading to PlateSpin Migrate 12.2,” on page 4
- Section 7, “Software Fixes,” on page 4
- Section 8, “Known Issues,” on page 9
- Section 9, “Supplemental Documentation,” on page 32
- Section 10, “Legal Notice,” on page 32

1 What’s New (September 2017)

This Release Notes document provides supplemental documentation for PlateSpin Migrate 12.2. See “Linux Boot Partition Requirement” under Supplemental Documentation.

2 What’s New for PlateSpin Migrate 12.2 Patch 1

PlateSpin Migrate 12.2 Patch 1 is mandatory fix that you must apply on a base installation of PlateSpin Migrate 12.2 (with or without field patches applied) to update your existing installations of the PlateSpin Server and Migrate Client. It resolves several issues. For more information, see the *PlateSpin Migrate 12.2 Patch 1 Readme*.

3 What's New for PlateSpin Migrate 12.2 RabbitMQ Fix

PlateSpin Migrate 12.2 RabbitMQ Fix is a patch that addresses an issue where the Event Messaging components RabbitMQ and Erlang are not installed and configured properly after you have installed or upgraded PlateSpin Migrate 12.2 Server on a drive that is not the `C:` drive. PlateSpin Transformation Manager 1.1 is the only consumer of the Event Messaging feature. For more information, see the Knowledge Base article "The Rabbit MQ server used by PlateSpin Transformation Manager 1.2 is unable to connect to PlateSpin Migrate 12.2 (KB 7018515)" (https://www.netiq.com/support/kb/doc.php?id=7018515).

You can download the PlateSpin Migrate 12.2 RabbitMQ Fix at Micro Focus Patch Finder (https://download.microfocus.com/Download?buildid=YEjOyq56-cA~).

4 What's New for PlateSpin Migrate 12.2

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

- Section 4.1, "Support for Migration of Workloads to Cloud," on page 2
- Section 4.2, "Supported Configurations," on page 3
- Section 4.3, "Migration Web Interface Enhancements," on page 4
- Section 4.4, "Large-Scale Migration Planning and Automation," on page 4
- Section 4.5, "Security Enhancements," on page 4

4.1 Support for Migration of Workloads to Cloud

PlateSpin Migrate 12.2 introduces support for migration of workloads to VMware vCloud Director and Amazon Web Services (AWS). It also enhances the support for migration of workloads to Azure Cloud.

- Section 4.1.1, "Migration of Workloads to VMware vCloud Director," on page 2
- Section 4.1.2, "Migration of Workloads to Amazon Web Services," on page 2
- Section 4.1.3, "Migration of Workloads to Azure Cloud," on page 2

4.1.1 Migration of Workloads to VMware vCloud Director

PlateSpin Migrate 12.2 enhances the Web Interface to let you migrate Windows and Linux workloads to VMware vCloud Director. For a list of the supported workloads, see "Supported Workloads For Migration to VMware vCloud Director" in the PlateSpin Migrate 12.2 User Guide.

4.1.2 Migration of Workloads to Amazon Web Services

PlateSpin Migrate 12.2 enhances the PlateSpin Migrate Client to let you migrate Windows and Linux workloads to AWS. For a list of the supported workloads, see "Supported Workloads For Migration to Amazon Web Services" in the PlateSpin Migrate 12.2 User Guide.

4.1.3 Migration of Workloads to Azure Cloud

PlateSpin Migrate 12.2 includes the following support:

- Publishing the PlateSpin Migrate Server VM in Azure Market place. See "Azure Cloud Based Migrate Server" in the *PlateSpin Migrate 12.2 User Guide*.

- Perform incremental replication of workloads to Azure Cloud.
- Use the enhanced Migrate Agent utility to migrate workloads to Azure Cloud over the Internet even when no VPN is configured. See "Migrate Agent Utility" in the *PlateSpin Migrate 12.2 User Guide*

4.2 Supported Configurations

PlateSpin Migrate 12.2 includes support for the following:

- Section 4.2.1, "Workload Support," on page 3
- Section 4.2.2, "LVM Raw Disks," on page 3
- Section 4.2.3, "VMware Containers," on page 3
- Section 4.2.4, "VMware vSAN Support," on page 4
- Section 4.2.5, "PlateSpin ISO (Linux Ram Disk)," on page 4

4.2.1 Workload Support

PlateSpin Migrate 12.2 includes support for the following operating systems:

- Windows Server 2012, 2008, 2003 Cluster
- Red Hat Enterprise Linux (RHEL) 6.8
- Oracle Enterprise Linux (OEL) 6.8
- CentOS 6.8
- In addition to the existing BIOS-based support, PlateSpin Migrate 12.2 includes UEFI-based support for the following:
  - Red Hat Enterprise Linux (RHEL) 7.2, 7.1, 7.0, 6.7
  - Oracle Enterprise Linux (OEL) 7.2, 7.1, 7.0, 6.7
  - CentOS 7.2, 7.1, 7.0, 6.7
  - SLES 11 SP4

See the "Supported Configurations" section in the *PlateSpin Migrate 12.2 User Guide*.

4.2.2 LVM Raw Disks

PlateSpin Migrate 12.2 adds support for LVM raw disk volumes for Same as Source storage configurations on Linux workloads.

4.2.3 VMware Containers

PlateSpin Migrate 12.2 includes support for the following VM containers:

- VMware vCenter 5.1 Update 3, 5.5 Update 3, 6.0 Update 1 and 2
- VMware ESXi 5.1 Update 3, 5.5 Update 3, 6.0 Update 1 and 2
- SUSE Linux Enterprise Server (SLES) 12 SP1 KVM
- Red Hat Enterprise Linux (RHEL) 7.2 KVM

See the "Supported Configurations" section in the *PlateSpin Migrate 12.2 User Guide*.

4.2.4 VMware vSAN Support

PlateSpin Migrate 12.2 includes support for VMware vSAN.

4.2.5 PlateSpin ISO (Linux Ram Disk)

PlateSpin Migrate 12.2 provides an updated PlateSpin ISO (Linux Ram Disk) based on SLES 11 SP4 kernel for X2P migrations.

4.3 Migration Web Interface Enhancements

- **Replication Networks for Source.** This option was renamed from **Networks Allowed for Replication** and moved to follow the **Storage** options under Migration Settings.
- **Replication Network for Target.** This option was renamed from **Replication Network** and moved to follow the **Storage** options under Migration Settings.

4.4 Large-Scale Migration Planning and Automation

You can now plan large scale workload migrations, execute and monitor the workload migrations through PlateSpin Transformation Manager 1.1. PlateSpin Transformation Manager is a planning, tracking, and automation solution for data center transformation projects.

See the *PlateSpin Transformation Manager User Guide* on the Documentation site (https://www.netiq.com/documentation/platespin-transformation-manager-1-1/).

4.5 Security Enhancements

PlateSpin Migrate 12.2 updates OpenSSL to version 1.0.2h to address vulnerability issues in OpenSSL. For more information see the OpenSSL project (http://openssl.org).

5 Installing PlateSpin Migrate 12.2

To install PlateSpin Migrate 12.2, see “Installing PlateSpin Migrate” in the *PlateSpin Migrate 12.2 Installation and Upgrade Guide*.

6 Upgrading to PlateSpin Migrate 12.2

To upgrade your PlateSpin Server to PlateSpin Migrate 12.2, you must have an existing installation of PlateSpin Migrate 12.1 or PlateSpin Migrate 12.1 Patch Update 1 (P1). See “Upgrading PlateSpin Migrate” in the *PlateSpin Migrate 12.2 Installation and Upgrade Guide*.

Other direct upgrades are not supported. For information about upgrading from PlateSpin Migrate 12.1 to PlateSpin Migrate 12.1 P1, see the PlateSpin Migrate 12.1 Patch Update 1 (P1) Release Notes.

7 Software Fixes

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

- Section 7.1, “VMware Roles Tool Is Missing Some Permissions for the PlateSpin Virtual Machine Manager,” on page 5
- Section 7.2, “Remove Target Does Not Display a Confirmation Dialog in a German Language Browser,” on page 5

- Section 7.3, “Replication Environment Blobs Were Not Automatically Cleaned Up after ‘Cutover’ or 'Remove Workload' or 'Abort',” on page 6
- Section 7.4, “Unable to Discover the Cores Per Socket on a Windows Source Workload When the Target Is Set as VMware,” on page 6
- Section 7.5, “Inconsistent Behavior between Azure and VMware Targets when Deselecting LVM Volume Group but Not the Corresponding Logical Volume,” on page 6
- Section 7.6, “Full Replication Breaks Differential Backup for SQL Server,” on page 6
- Section 7.7, “Cutover Stalls with a Configuration Service Not Started Error,” on page 7
- Section 7.8, “Partitions Not Mounted to Drive Letters After Conversion,” on page 7
- Section 7.9, “Discovery Fails on Windows Server 2012 without .NET 3.5 Feature Installed,” on page 7
- Section 7.10, “Azure to Azure Migration Might Delete the Source Workload If You Select Same Target and Datastore,” on page 7
- Section 7.11, “Test Cutover Failed When the Hostname Setting on the Configuration Page of a Target Linux Workload Contained an Underscore,” on page 7
- Section 7.12, “Two Disks Share the Same Disk Path on Config Page for Dynamic Disk Workloads,” on page 8
- Section 7.13, “Web Interface: Validation Message Is Not Displayed after Executing the "Remove Workload" with "Preserve Source",” on page 8
- Section 7.14, “Migrate Client Migration of a Windows Server 2008 R2 Workload to VMware 5.5 or 6.0 Shows a Non-Critical Error for Installing VMware Tools,” on page 8
- Section 7.15, “Migrate Client Warns That a vSAN Datastore Is Not Shared Across a VMware Cluster When It Is,” on page 8
- Section 7.16, “Cutover Hangs with VMware vCDROM Locked Message in vSphere Client; User Intervention Required,” on page 9
- Section 7.17, “After Successful Cutover, Azure Portal Web UI Does Not Show Target VM's Computer Name and DNS Name,” on page 9

7.1 VMware Roles Tool Is Missing Some Permissions for the PlateSpin Virtual Machine Manager

Issue: The PlateSpin.VMwareRoleTool.exe reported exceptions for some permissions that were missing for the PlateSpin Virtual Machine Manager in the `PlateSpinRole.xml file`. `(Bugs 962265, 958576, 969197)`

Fix: The PlateSpinRole.xml file now includes all required minimum permissions.

7.2 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)`

Fix: A confirmation dialog is available in all supported languages.

7.3 Replication Environment Blobs Were Not Automatically Cleaned Up after ‘Cutover’ or 'Remove Workload' or 'Abort'

Issue: For migrations to Microsoft Azure, the Azure Blob Service created 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 did not automatically remove the artifacts. Microsoft Azure charged you to store these unneeded data files. (Bug 977308)

Fix: After a workload migration is completed, PlateSpin Migrate now automatically removes the artifacts.

7.4 Unable to Discover the Cores Per Socket on a Windows Source Workload When the Target Is Set as VMware

Issue: If you discover a Windows workload with multiple sockets and cores, and then select a VMware container for the migration target, the values for the **Sockets** and the **Cores Per Socket** settings on the Migration Settings page in the Web Interface display as 1. (Bug 999444)

Fix: Migrate correctly detects the CPU and sets the **Sockets** and the **Cores Per Socket** values properly.

7.5 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 VMware and Azure. (Bugs 973926 and 970584)

Fix: If you deselect a volume group, 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.
```

7.6 Full Replication Breaks Differential Backup for SQL Server

Issue: After a full replication for a Microsoft Windows Server 2012 workload running Microsoft SQL Server 2012, the SQL Server’s differential backup fails until a full backup with Microsoft tools runs again. This issue might occur if Microsoft VSS Writer for SQL Server (SqlServerWriter) is running during the replication. (Bug 983567)

Fix: The issue with SqlServerWriter has been resolved. SQL server differential backups are not broken.

As a best practice, ensure that no backups run during the narrow window when the Microsoft VSS Snapshot is created. This practice applies for any software that uses Microsoft's Volume Shadow Copy Service to create snapshots, including anti-virus, SQL backups, and so on.

7.7 Cutover Stalls with a Configuration Service Not Started Error

Issue: During a cutover job, the target VM fails to boot. The cutover process stalls at configuring the services with a `Configuration Service Not Started` error. The configuration service on the target workload appears to hang, when it actually has completed successfully but could not report back to the PlateSpin server. A hang might also occur if software on the cutover workload keeps processes from continuing until the workload reboots and reboot optimization is enabled (the default) for the Configuration Service. `(Bugs 994314, 991068, 991089, 988132, 987156, 986629, 984153, and 982362)`

Fix: Cutover jobs perform as expected for most workloads. If the Configuration Service hangs for a Windows workload during a cutover action, you can disable optimized reboots for the Configuration Service, which allows the reboots to occur as requested by software. You can skip optimized reboots for a specific Windows workload or for all Windows workloads See Setting Reboot Method for the Configuration Service.

7.8 Partitions Not Mounted to Drive Letters After Conversion

Issue: Following a conversion to Windows Server 2012 R2 Hyper-V, only the "C" drive displayed. Other partitions were not mounted to drive letters.`(Bug 894623/915807)`

Fix: The partitions are now mounted to the drives post conversion.

7.9 Discovery Fails on Windows Server 2012 without .NET 3.5 Feature Installed

Issue: Discovery fails for Windows Server 2012 workloads where only Microsoft .NET 4.0 is installed. The Discovery expected `.NET 2.0 SP2` or `.NET 3.5 SP1` to be installed. `(Bug 998423)`

Fix: Migrate discovery supports `Microsoft .NET Framework version 2.0 SP2`, `3.5 SP1`, or `4.0` installed on Windows sources and Hyper-V hosts.

7.10 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)`

Fix: The option to Remove Workload is disabled in the UI so that you cannot inadvertently delete the source workload.

7.11 Test Cutover Failed When the Hostname Setting on the Configuration Page of a Target Linux Workload Contained an Underscore

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

```
Failed to configure virtual machine
```

`(Bug 975854)`

Fix: Test Cutover is successful even if the hostname contains an underscore.

7.12 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)

Fix: Target disks are named sequentially by their order on the target.

7.13 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)

Fix: Validation of the discovered machine is re-run when the workload is returned to an Unprotected state during a **Delete with preserve source** operation.

7.14 Migrate Client Migration of a Windows Server 2008 R2 Workload to VMware 5.5 or 6.0 Shows a Non-Critical Error for Installing VMware Tools

Issue: When using the PlateSpin Migrate Client to migrate a Windows Server 2008 R2 workload to VMware 5.5 or 6.0, a non-critical error occurs for installing tools when Install VMware Tools is enabled.

```
Installing Tools [Failed: Non-critical error]
```

However, the VMware tools are properly installed. (Bug 992705)

Fix: The installation of VMware Tools on a Windows Server 2008 R2 workload completes without errors.

7.15 Migrate Client Warns That a vSAN Datastore Is Not Shared Across a VMware Cluster When It Is

Issue: PlateSpin Migrate Client displays the following warning for a vSAN datastore:

The selected datastore <*datastore*> is not shared across VMware cluster.

However, the selected datastore is shared across the VMware cluster. The warning is not seen in the Web Interface. (Bug 992643)

Fix: The vSAN datastores are listed as non-local so that the Migrate Client validators correctly identify them as shared storage in a VMware cluster.

7.16 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 and you are prompted to override the vCDROM lockout. This issue did not occur on all target Linux workloads or on all VMware container versions. (Bug 975853)

Fix: Manual intervention is no longer required.

7.17 After Successful Cutover, Azure Portal Web UI Does Not 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)

Fix: The Azure Portal does not show this information. 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.

8 Known Issues

- Section 8.1, "General Issues," on page 9
- Section 8.2, "Known Issues For Install," on page 18
- Section 8.3, "Known Issues For Uninstall," on page 19
- Section 8.4, "Known Issues for Upgrade," on page 20
- Section 8.5, "Rediscovering vCenter Server in Migrate Web Interface Post Upgrade Fails With An Exception," on page 21
- Section 8.6, "Known Issues For Migration to Azure," on page 21
- Section 8.7, "Known Issues For Migration to Amazon Web Services," on page 28
- Section 8.8, "Known Issues For Migration to vCloud," on page 29
- Section 8.9, "Known Issues For Migration to VMware," on page 29

8.1 General Issues

The following issues are being researched:

- Section 8.1.1, "Unable to configure migration job for source Linux workloads having hostname and LVM name exceeding 63 characters," on page 11
- Section 8.1.2, "Test Cutover and Cutover Confirmation Warns That the Last Replication Is Older Than 24 Hours but the Elapsed Time Is Less Than 24 Hours," on page 11
- Section 8.1.3, "Quorum Disk Converted as a Local Disk When Migrating a Cluster Node from A Virtual Source Machine to A Physical Target Machine," on page 11
- Section 8.1.4, "V2P Windows Cluster Conversion Creates a Copy of the Quorum Disk on a Local Disk," on page 12
- Section 8.1.5, "Incremental Replication Might Take a Long Time When an MD5 Checksum or Server-Sync Is Needed," on page 12
- Section 8.1.6, "Help On Chinese Locales Displays in English Language," on page 12

- Section 8.1.7, “Incremental File-Based Replication Does Not Complete with Encryption Enabled,” on page 12
- Section 8.1.8, “Installation Fails When the OS Language and OS Locale Do Not Match on the Computer,” on page 13
- Section 8.1.9, “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 13
- Section 8.1.10, “Mapping Volumes Not Supported When Migrating Linux Workloads,” on page 13
- Section 8.1.11, “Unable to Migrate Linux workloads Having Volumes Created on Raw Disks Without Partitions,” on page 13
- Section 8.1.12, “Unable to Migrate a Workload to Hitachi LPAR,” on page 13
- Section 8.1.13, “Warning Message Displayed When You Migrate a Workload To Hitachi LPAR,” on page 14
- Section 8.1.14, “Discovered Hyper-V Container Displays As a Workload in the PlateSpin Migrate Web Interface,” on page 14
- Section 8.1.15, “Unable to Migrate a Linux Workload to a Container That Does Not Support the Source Workload Firmware,” on page 14
- Section 8.1.16, “(Windows Sources) Non-Default Per-Volume VSS Settings Not Preserved After Migration,” on page 14
- Section 8.1.17, “Special Character in Datastore Name Causes Migration Problems,” on page 14
- Section 8.1.18, “VSS Snapshots Are Not Preserved,” on page 14
- Section 8.1.19, “Migration Over WAN Taking a Long Time When the Target VM Host Has a High Number of Datastores,” on page 15
- Section 8.1.20, “Network Card Not Initialized on SLES 11 Target VM Hosted on a Windows Server 2008 Hyper-V Host,” on page 15
- Section 8.1.21, “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 8.1.22, “XenServer Tools Not Removed After Conversion,” on page 15
- Section 8.1.23, “Post Migration the Primary Partition Is Converted to a Logical Partition on the Target,” on page 15
- Section 8.1.24, “Machine Node on ESX Host Is Not Undiscovered When Migrate Undiscovers a Machine,” on page 16
- Section 8.1.25, “V2P Conversion Hangs at the Configuring Operating System Step,” on page 16
- Section 8.1.26, “UEFI to BIOS Conversion for Windows 8.1 Workload Fails at the Sending Files Step,” on page 16
- Section 8.1.27, “Conversion Fails While Downgrading from UEFI to BIOS Firmware,” on page 16
- Section 8.1.28, “Source Machine Stays in an Under Control State After Offline Conversion,” on page 16
- Section 8.1.29, “Source Machine Boot Configuration Not Restored After Offline Conversion,” on page 16
- Section 8.1.30, “Creating and Moving a VM Under Resource Pool as a Setting Is Not Supported in the CLI Tool,” on page 17
- Section 8.1.31, “Partitions Not Mounted to Drive Letters After Conversion,” on page 17
- Section 8.1.32, “Redundant Disks Are Present After a RHEL 6.2 x64 Block Migration to Windows Server 2012 R2 Hyper-V,” on page 17

- Section 8.1.33, “Requirements for VMware DRS Cluster support,” on page 17
- Section 8.1.34, “Installation of the PlateSpin Image Server Fails When the Fully Qualified Image File Path is More Than 248 Characters,” on page 17
- Section 8.1.35, “Conversion Job Fails to Configure NICs on a Windows 2000 Server Target Machine,” on page 17
- Section 8.1.36, “Rediscovering ESX Server in Migrate Client Displays a Duplicate Server Entry With Add Failed Error Message in the Web Interface,” on page 18
- Section 8.1.37, “Block-based Full Replication Fails for Windows Server 2012 R2 Server with Complex Partitioning,” on page 18
- Section 8.1.38, “Unable to Validate the Transfer Settings in the INI File When Using CLI Tool For Conversion or ServerSync Jobs,” on page 18

8.1.1 Unable to configure migration job for source Linux workloads having hostname and LVM name exceeding 63 characters

Issue: If you choose to configure a migration job for source Linux workloads with the combination of hostname and LVM name exceeding 63 characters, the job configuration fails. `(Bug 978065)`

Workaround: Ensure that the combination of the hostname and LVM name on your Linux workload does not exceed 63 characters.

8.1.2 Test Cutover and Cutover Confirmation Warns That the Last Replication Is Older Than 24 Hours but the Elapsed Time Is Less Than 24 Hours

Issue: The confirmation dialog for the Test Cutover or Cutover actions displays a yellow **Warning** icon next to the **Last Replication** date and time with the following message:

```
The last replication is older than 24 hours.
```

However, the interval between the completion time and the current time is less than 24 hours. All times represent the time zone of the Migrate server. `(Bug 1020776)`

Workaround: Take one of the following actions in the Test Cutover or Cutover confirmation dialog, as appropriate for your data integrity needs:

- Enable **Perform Incremental Replication**, then click **Execute**. An incremental replication runs before the cutover as part of the cutover process.
- Ignore the warning and click **Execute**. Data changes made to the source workload since the last replication will not be there on the target workload.

8.1.3 Quorum Disk Converted as a Local Disk When Migrating a Cluster Node from A Virtual Source Machine to A Physical Target Machine

Issue: When you migrate a cluster node from a virtual source machine to a physical target machine and the number of disks that you choose to migrate are less than the number of disks on the target machine, the quorum disk from the source machine is converted as local disk on the target machine irrespective of whether you selected the quorum disk for migration during JOB configuration or not.

Workaround: Rebuild the cluster environment. For information about rebuilding the cluster environment after the migration, see KB Article 7016770.

8.1.4 V2P Windows Cluster Conversion Creates a Copy of the Quorum Disk on a Local Disk

Issue: During a virtual-to-physical Windows Cluster conversion of the active node using the Migrate Client, a copy of the Quorum disk is created on the second local disk on the target workload. The conversion completes successfully, but Cluster Services are not running. The local copy of the Quorum disk is created whether you select or deselect the Quorum disk during the job setup. (Bug 1017154)

Workaround: Delete the copy of Quorum disk on the target workload's local disk, and then start the Windows Cluster service.

8.1.5 Incremental Replication Might Take a Long Time When an MD5 Checksum or Server-Sync Is Needed

Issue: An incremental replication can take twice as long as a full replication if the workload needs an MD5 checksum or server-sync action. For example:

- **Windows:** If a failover occurs in Windows with a warning stating "`switching to server sync mode`", the transfer method fails back to perform an MD5 checksum.
- **Linux:** Incremental replications for Linux workloads take twice the time of a full replication if the workload needs to perform an MD5 checksum or server-sync.
- **Linux (Azure):** After a successful Abort action for an incremental replication of a Linux workload to Azure, the subsequent incremental replication appears to hang at Copy Data at 36% with no error message. An MD5 checksum is being performed on the backend that takes about an hour to complete.

The incremental replications eventually complete successfully. (Bug 1016246)

Workload: None.

8.1.6 Help On Chinese Locales Displays in English Language

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

8.1.7 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.

8.1.8 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.

8.1.9 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 lose in the functionality. (Bug 967018)

Workaround: Ignore the option.

8.1.10 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 a existing Volume Group
- Mapping any volume to new Volume Group
- Re-mapping Volume group to disk

8.1.11 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)

8.1.12 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.

8.1.13 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.

8.1.14 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.

8.1.15 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.

8.1.16 (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.

8.1.17 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.

8.1.18 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)

8.1.19 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.

8.1.20 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.

8.1.21 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) includes information about how to change the position of the virtual disk containing the operating system. See KB Article 7012906.

8.1.22 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.

8.1.23 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.

8.1.24 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.

8.1.25 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.

8.1.26 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.

8.1.27 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.

8.1.28 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.

8.1.29 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.

8.1.30 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.

8.1.31 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.

8.1.32 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.

8.1.33 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. `(Bug 896598)`

See "Discovery Guidelines for Machine Types and Credentials" in your *User Guide*.

8.1.34 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.

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

Issue: If you choose to migrate a Window 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.

8.1.36 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.

8.1.37 Block-based Full Replication Fails for Windows Server 2012 R2 Server with Complex Partitioning

Issue: Block-based full replication fails for Windows Server 2012 R2 servers with complex partitioning of more than 44 partitions or volumes. High disk utilization can prevent creation of VSS snapshots for any additional partitions or volumes. (`Bug 1007833`)

Workaround: Ensure that workloads you attempt to replicate have no more than 44 partitions or volumes.

8.1.38 Unable to Validate the Transfer Settings in the INI File When Using CLI Tool For Conversion or ServerSync Jobs

Issue: Migrating Linux workloads using file-based transfer is possible only through offline conversion. If you choose to use the CLI tool to migrate a Linux workload using file-based transfer and edit the settings in `[Transfer]` section of the `INI` file to set the `TransferType` setting to `filebased`, but incorrectly set the `LiveTransferEnabled` setting to `True` instead of `False`, the job fails to display a validation error. Also, the migration continues using the block-based transfer method. (`Bug 994513`)

Workaround: To migrate Linux workloads using File-based transfer, you must ensure that the `LiveTransferEnabled` setting is set to `False`.

8.2 Known Issues For Install

The following issue is being researched:

- Section 8.2.1, “Installation Fails on a Windows Server 2012 R2 Domain Computer If the User Performing the Installation Does Not Have Administrator Privileges,” on page 18
- Section 8.2.2, “Silent Install Is Not Supported for Migrate Servers Used with PlateSpin Transformation Manager,” on page 19
- Section 8.2.3, “Event Messaging Components Are Not Properly Configured,” on page 19

8.2.1 Installation Fails on a Windows Server 2012 R2 Domain Computer If the User Performing the Installation Does Not Have Administrator Privileges

Issue: If a domain user chooses to install PlateSpin Migrate on a Windows Server 2012 R2 domain computer, the installation fails and an `Operating system error` message is logged.(`Bug 1011557`)

Workaround: The user must have administrator privileges to install PlateSpin Migrate on a Windows Server 2012 R2 domain computer.

8.2.2 Silent Install Is Not Supported for Migrate Servers Used with PlateSpin Transformation Manager

Issue: The components needed for Event Messaging are not properly installed and configured if you use a batch file to silently install PlateSpin Migrate as described in "Installing the PlateSpin Migrate Software Components Using a Batch File" in the *PlateSpin Migrate 12.2 Installation and Upgrade Guide*. PlateSpin Transformation Manager 1.2 is the only consumer of the Event Messaging feature. (Bug 1020689)

Workaround: Do not use silent install to set up Migrate servers that you plan to use with PlateSpin Transformation Manager 1.2.

8.2.3 Event Messaging Components Are Not Properly Configured

Issue: Event Messaging components are not configured properly if you install or upgrade PlateSpin Migrate Server on a drive that is not the C: drive. PlateSpin Transformation Manager 1.2 is the only consumer of the Event Messaging feature. (Bug 1021229)

Workaround: To avoid this problem, ensure that you install PlateSpin Migrate Server on the C: drive for Migrate servers that you plan to use with PlateSpin Transformation Manager 1.2. If you encounter this issue and need to install PlateSpin Migrate Server on a different drive, contact Customer Support for assistance.

8.3 Known Issues For Uninstall

The following issue is being researched:

- Section 8.3.1, "After Uninstalling PlateSpin Server, the Reinstall Stalls at 'Installing OFX: Enabling STOMP Protocol'," on page 19

8.3.1 After Uninstalling PlateSpin Server, the Reinstall Stalls at 'Installing OFX: Enabling STOMP Protocol'

Issue: After an uninstall of PlateSpin Server, if you reinstall without rebooting, the installation might stall with the following message:

```
Status: Installing OFX
        Enabling STOMP protocol
```

The Migrate Event Messaging function uses RabbitMQ software, Erlang programming language, and STOMP. The uninstall removes the installed files, but does not automatically stop the related processes. If the processes are still running when you attempt to reinstall PlateSpin Server, the installation stalls. (Bug 999112)

Workaround: After the uninstall of PlateSpin Server, ensure that you stop the following services for the Event Messaging function (if they are running) before you perform the next install:

- `erlsrv.exe`
- `erl.exe`
- `epmd.exe`

8.4 Known Issues for Upgrade

The following issues are being researched:

- Section 8.4.1, “After Upgrade, There Is an Extra ‘Default’ Workload Tag,” on page 20
- Section 8.4.2, “Migration Contracts in the Configured State Are Reset to Not Configured,” on page 20
- Section 8.4.3, “Linux Workloads Prepared For Migration Before Upgrade Should be Prepared Again Post Upgrade,” on page 20
- Section 8.4.4, “Upgrade Fails If Migrate Client Is Running,” on page 20
- Section 8.4.5, “Web Interface Does Not List Any of The Target Virtual Machines When You Configure A Workload For Incremental Replication For the First Time Post Upgrade,” on page 21
- Section 8.4.6, “The Web Interface Displays More Than One Instance of the Same Discovered Object Post Upgrading to Migrate 12.2,” on page 21

8.4.1 After Upgrade, There Is an Extra ‘Default’ Workload Tag

Issue: After you upgrade the Migrate server from version 12.1 to version 12.2, the Web Interface shows an extra workload tag named **Default** that is assigned the color gray. (`Bug 1018730`)

Workaround: In the Web Interface, select **Settings** > **Workload Tags**, then delete the **Default** workload tag.

8.4.2 Migration Contracts in the Configured State Are Reset to Not Configured

Issue: The upgrade to Migrate 12.2 resets any migration contracts that are in the Configured state to the Not Configured state. You must reconfigure the contracts after the upgrade. (`Bug 1018109`)

Workaround: Do the following:

- **Windows Workloads:** Before you begin the upgrade, run the **Prepare Migration** operation for each Windows migration contract that is in the Configured state.
- **Linux Workloads:** After the upgrade, reconfigure the migration contract and run **Prepare Migration**. See also “Linux Workloads Prepared For Migration Before Upgrade Should be Prepared Again Post Upgrade” on page 20.

8.4.3 Linux Workloads Prepared For Migration Before Upgrade Should be Prepared Again Post Upgrade

Issue: If you have Linux workloads prepared for synchronization or incremental replication and you upgrade the Migrate server to version 12.2, the prepared state of the workloads is reset. (`Bug 1018112`)

Workaround: Prepare the Linux workloads again post upgrade.

8.4.4 Upgrade Fails If Migrate Client Is Running

Issue: The upgrade fails if the Migrate Client is running because the installation cannot update the open files. (`Bug 992912`)

Workaround: Close the Migrate Client before you attempt to run the upgrade for the PlateSpin Server or Migrate Client.

8.4.5 Web Interface Does Not List Any of The Target Virtual Machines When You Configure A Workload For Incremental Replication For the First Time Post Upgrade

Issue: After you upgrade the Migrate server to version 12.2, if you use the Web Interface to configure incremental replication of a workload for the first time and select an ESX target, the **Target Virtual Machines** option does not list any virtual machines. (Bug 1017328)

Workaround: Refresh the target ESX computer to enable the listing of the target VMs.

8.4.6 The Web Interface Displays More Than One Instance of the Same Discovered Object Post Upgrading to Migrate 12.2

Issue: If you use the PlateSpin Migrate Client to discover workloads and targets and have discovered the same object in different networks of the Migrate Client, then only the object that you discovered in the default network displays on the PlateSpin Migrate Web Interface. However, on upgrading to Migrate 12.2, more than one instance of the discovered object displays on the Web Interface depending on the number of times it has been discovered in the various networks. (Bug 977577)

Workaround: Delete all the extra instances of the object listed in the Web Interface.

8.5 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 12.2. If you now upgrade to Migrate 12.2, 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)

8.6 Known Issues For Migration to Azure

The following issues are being researched:

- Section 8.6.1, “Hostname and Domain Membership Changes Are Not Honored at Cutover,” on page 22
- Section 8.6.2, “Cutover to Azure Fails for a Workload with a Data Disk Partition Greater Than 1 TB in Size,” on page 22
- Section 8.6.3, “For Linux Migration to Azure, the Ethernet Interfaces Have Different Names on the Target VM,” on page 23
- Section 8.6.4, “Azure-to-Azure Migration Might Delete the Source Workload If You Select the Same VM and Datastore as Source and Target,” on page 23
- Section 8.6.5, “Assign a Reserved Virtual IP Address to a Migrate Server in Azure,” on page 23
- Section 8.6.6, “Failed Workload Migration to Azure with Password Is Expired Error,” on page 23
- Section 8.6.7, “Cutover Fails During Partitioning Operations,” on page 24
- Section 8.6.8, “Network Connections on the Target VM Might Not Be Mapped to the Correct Azure Virtual Network Or Subnet,” on page 24
- Section 8.6.9, “Azure Target VMs Are Created with One Additional Hard Disk,” on page 24
- Section 8.6.10, “For the Maximum Azure Instance Size, Need a UI Check to Limit the Data Disks for Replication to 63 Data Disks,” on page 25

- Section 8.6.11, “Disk Numbers and DiskIndex Numbers Are Not Sequential for Discovered Dynamic Disk Workloads,” on page 25
- Section 8.6.12, “Dynamic Disk Workload Default Cloud Instance Size Is Too Large,” on page 25
- Section 8.6.13, “Virtual Disk Ordering Incorrectly Changes When a Disk Is Excluded for Migration on the Config Page,” on page 25
- Section 8.6.14, “Migration to Azure failed for SLES 11 SP 4 with 3 Disks and No LVM,” on page 25
- Section 8.6.15, “Linux Disks or Partitions Are Migrated on to Target in a Different Order than that of the Linux Source,” on page 26
- Section 8.6.16, “LVM Volume Groups Are Created on Opposite Partitions within the Same Disk on Linux Target VM,” on page 26
- Section 8.6.17, “Linux Partitions Are Created on Opposite Partitions within the Same Disk on Linux Target VM,” on page 26
- Section 8.6.18, “Azure Target VM Launched in Safe Mode After Successful Cutover of a Workload,” on page 26
- Section 8.6.19, “The Virtual Machine Settings Page of the Target VM in the Azure Portal Does Not Display the Size of the VM,” on page 27
- Section 8.6.20, “SLES 11 SP4 with /boot on sdb - Migration Fails with Configuration Service Not Started Error,” on page 27
- Section 8.6.21, “Abort of First Full Replication Does Not Clean Up root-PS-snapshot on Linux Source Workload,” on page 27
- Section 8.6.22, “Incremental Replication Does Not Transfer Data from Source to Target for LVM Volumes on RHEL 6.7, Oracle Linux 6.7, and CentOS 6.7,” on page 27
- Section 8.6.23, “Setting 'Bandwidth Throttling' on RHEL 6.7/7.1 Triggers a Validation Warning that Should Not Be Triggered,” on page 28

8.6.1 Hostname and Domain Membership Changes Are Not Honored at Cutover

Issue: A hostname change or domain membership made to the migration contract configuration at Cutover is not honored on the target VM. A replication is required to propagate the hostname or domain membership change to the target VM. A hostname or domain membership change made at Test Cutover is honored because the test cutover process is handled differently than Cutover. (Bug 1019651)

Workaround: After you change the hostname or domain membership in the migration contract configuration, ensure that you run an incremental or full replication before you attempt Cutover.

8.6.2 Cutover to Azure Fails for a Workload with a Data Disk Partition Greater Than 1 TB in Size

Issue: Cutover to Azure does not complete for a workload that has a GPT data disk greater than 1 TB in size. If the source machine is BIOS with MBR as the boot disk, you see the following error message:

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

(Bug 977922)

Workaround: Azure supports a maximum disk size of 1 TB, whether it is MBR or GPT.

8.6.3 For Linux Migration to Azure, the Ethernet Interfaces Have Different Names on the Target VM

Issue: Because the MAC addresses for NICs on a target Linux VM in Azure are different than the MAC addresses on the source Linux workload, Linux assigns different names to them. For example, the eth0 and eth1 Ethernet interfaces on a source Linux workload become eth2 and eth3 on the target Linux VM in Azure. `(Bug 967468)`

Workaround: You can manually rename the NICs in the target VM after cutover.

8.6.4 Azure-to-Azure Migration Might Delete the Source Workload If You Select the Same VM and Datastore as Source and Target

Issue: In an Azure-to-Azure cutover, if you inadvertently select a target VM and datastore that are actually the source objects, Migrate destroys the source workload if you select **Abort** before the **Cutover** operation fails on its own. `(Bug 1011726)`

Workaround: Let the **Cutover** action fail until the **Abort** button is replaced by the **Configure Migration**, **Run Cutover**, and **Remove Workload** buttons, then click **Remove Workload**.

8.6.5 Assign a Reserved Virtual IP Address to a Migrate Server in Azure

Issue: The IP address for the Migrate Server in Azure can change on shutdown or reboot unless it has a Reserved Virtual IP Address. A change in IP address on the PlateSpin Server breaks the heartbeat communications with source workloads. `(Bug 1015893)`

Workaround: When you provision the Migrate Server in Azure, specify **Static** as the allocation method for the public IP address of the Migrate Server to ensure that the IP address does not change when the server is restarted. See Azure Cloud Based Migrate Server and Deploying a Migrate Server in the Azure Cloud in the *PlateSpin Migrate User Guide*.

If you set up the Migrate Server in Azure with a Dynamic assignment method for the public IP address, you should modify the setting to use the Static assignment method. See Assigning a Reserved IP Address to a Migrate Server in Azure in the *PlateSpin Migrate User Guide*.

8.6.6 Failed Workload Migration to Azure with Password Is Expired Error

Issue: When you attempt to migrate a workload to an Azure target container, it fails with a `Password Is Expired` error message from Azure. On the Azure Target Container page, a tooltip displays the `Password Is Expired` error message for the password.

This error can occur if you allow the password to expire for the Azure user that you created in the Microsoft Azure portal to use for PlateSpin Migrate migrations.

Workaround: Resolve the password issue in Microsoft Azure, update the password stored for the Azure user for any affected Azure target containers, then rerun the failed workload migrations.

1. Log in to the user account in the Microsoft Azure portal and use the Azure Self-Service Password Reset (https://azure.microsoft.com/en-us/documentation/articles/active-directory-passwords-getting-started/#step-3-reset-your-azure-ad-password-as-a-user) to set a new password.
2. Log in to the PlateSpin Migrate Web Interface.
3. Update the password stored for the Azure user in the Azure target container.
4. Rerun any failed workload migrations to the Azure target container.

8.6.7 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.

8.6.8 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.

8.6.9 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.

8.6.10 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.

8.6.11 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.

8.6.12 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.

8.6.13 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.

8.6.14 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.

8.6.15 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.

8.6.16 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.

8.6.17 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.

8.6.18 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.

8.6.19 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 DS*X*_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 DS*X*_v2 series.

8.6.20 SLES 11 SP4 with /boot on sdb - Migration Fails with Configuration Service Not Started Error

Issue: Migration to Azure fails for a SUSE Linux Enterprise Server (SLES) 11 SP4 workload that has /boot on the second drive (sdb). The migration job fails with a Configuration Not Started error. (Bug 978115)

Workaround: This configuration is not supported. Migrate supports Linux workloads with /boot on the first disk (sda).

8.6.21 Abort of First Full Replication Does Not Clean Up root-PS-snapshot on Linux Source Workload

Issue: After a successful Abort action during cutover of a source Linux VM on VMware to a target Linux VM in Azure, a subsequent cutover attempt fails with an error:

Under-control conversion of a Linux source with LVM snapshots is not supported: See /dev/<*source-hostname*>/root-PS-snapshot

This error occurs because the root-PS-snapshot symbolic link was not removed during the clean up process for the Abort action. (Bug 1016619)

Workaround: Manually delete root-PS-snapshot symbolic link.

8.6.22 Incremental Replication Does Not Transfer Data from Source to Target for LVM Volumes on RHEL 6.7, Oracle Linux 6.7, and CentOS 6.7

Issue: On Red Hat Enterprise Linux (RHEL) 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads running kernel version 2.6.32-573.el6, incremental replications do not transfer data from source to target for LVM volumes. Full replications and server sync work with this kernel version. (Bug 1018176)

Workaround: Do one of the following for RHEL 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads if they have LVM volumes:

- Run only full replications or server sync for workloads running the RHEL kernel version 2.6.32-573.el6.
- Register with the Red Hat network and update the workload by using the Red Hat Package Manager (yum update) to get the latest available kernel fixes for RHEL 6.7. The workload is updated to use the same kernel version 2.6.32-642.13.1.el6 that is used by the RHEL 6.8 distribution. Both full and incremental replications perform as expected with this kernel version.

  The following blkwatch drivers are available for RHEL 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads running kernel 2.6.32-642:

  ```
  RHEL6-RHSA201700361-2.6.32-642.13.1.el6.i686-x8
  RHEL6-RHSA201700361-2.6.32-642.13.1.el6.x86_64-x86_64
  ```

8.6.23 Setting 'Bandwidth Throttling' on RHEL 6.7/7.1 Triggers a Validation Warning that Should Not Be Triggered

Issue: On a Red Hat 6.7 or 7.1 workload, after you configure migration for a workload with no warnings or validation errors, if you enable **Bandwidth Throttling** and save the configuration, a validation warning pops up:

```
This configuration contains validation warnings. You can save the configuration or
review the warnings.
```

The warning is triggered even if the **Bandwidth Throttling** setting is valid. (`Bug 1021521`)

Workaround: If you set a valid value, you can save the configuration and continue.

8.7 Known Issues For Migration to Amazon Web Services

The following issues are being researched:

- Section 8.7.1, “Migration of a RHEL 7.2 Source Workload Installed With a Non-Default Disk Layout Fails,” on page 28
- Section 8.7.2, “Migration of Source Workload With Volumes Labels Containing Special Characters Fails,” on page 28
- Section 8.7.3, “Migration of RHEL 7.X BIOS Workloads From AWS to VMware ESX Fails,” on page 28
- Section 8.7.4, “Discovery Fails for Windows Workloads in AWS for Private IP Address or WMI Method,” on page 29

8.7.1 Migration of a RHEL 7.2 Source Workload Installed With a Non-Default Disk Layout Fails

Issue: If you choose to migrate a RHEL 7.2 source workload that is installed with a non-default disk layout, the migration fails. (`Bug 1032569`)

Workaround: None.

8.7.2 Migration of Source Workload With Volumes Labels Containing Special Characters Fails

Issue: If you choose to migrate a source workload that has volume labels containing special characters, the migration fails at the Creating and Partitioning Volumes step. (`Bug 1017270`)

Workaround: Ensure that the volume labels on the source workload does not contain any special characters.

8.7.3 Migration of RHEL 7.*X* BIOS Workloads From AWS to VMware ESX Fails

Issue: If you migrate a RHEL 7.*x* BIOS workload having GPT partitioning scheme, the migration fails if the non-volume storage setting in the **Drive Configuration** > **Hard Drive** section of the migration job window is not selected. (`Bug 1022235`)

Workaround: Before migrating a RHEL 7.*x* BIOS workload to VMware ESX, ensure that you have selected the non-volume storage, which must be recreated during the migration, in the **Drive Configuration** > **Hard Drive** section of the migration job window.

8.7.4 Discovery Fails for Windows Workloads in AWS for Private IP Address or WMI Method

Issue: Discovery of a Windows workload in AWS fails if you use a private IP address or if you use the WMI discovery method. Possible error messages include:

```
Cannot connect to named pipe

Workload authentication failed

(Bug 1012001)
```

Workaround: For successful discovery of Windows workloads in AWS, use the workload's public IP address and the Common Inventory discovery method.

8.8 Known Issues For Migration to vCloud

The following issue is being researched:

- Section 8.8.1, "A Linux Workload Migrated to vCloud Target Does Not Boot If the Boot Partition is Not in the First Disk of the Source Workload," on page 29
- Section 8.8.2, "Discovered Cores and Sockets Settings Are Not Displayed Correctly If the Target Container Is vCloud," on page 29

8.8.1 A Linux Workload Migrated to vCloud Target Does Not Boot If the Boot Partition is Not in the First Disk of the Source Workload

Issue: On migrating a Linux workload that does not have the boot partition in the first disk to a vCloud target, the target workload fails to boot and the migration job moves to recoverable error. `(Bug 1006972)`

Workaround: None.

8.8.2 Discovered Cores and Sockets Settings Are Not Displayed Correctly If the Target Container Is vCloud

Issue: After discovery, if you select vCloud as the target container for a discovered source workload, the discovered source values for the cores and sockets settings are not correctly displayed on the Configuration page. For example, the Configuration page displays 1 core and 1 socket for a workload with discovered values of 2 cores and 2 sockets. `(Bug 1020264)`

Workaround: Configure the cores and sockets settings to values that you want to implement on the vCloud target.

8.9 Known Issues For Migration to VMware

The following issues are being researched:

- Section 8.9.1, "Disk Numbers and DiskIndex Numbers Are Not Sequential for Discovered Dynamic Disk Workloads," on page 30
- Section 8.9.2, "Virtual Disk Ordering Incorrectly Changes When a Disk Is Excluded for Migration on the Config Page," on page 30
- Section 8.9.3, "LVM Volume Groups Are Created on Opposite Partitions within the Same Disk on Linux Target VM," on page 30

- Section 8.9.4, “Linux Partitions Are Created on Opposite Partitions within the Same Disk on Linux Target VM,” on page 30
- Section 8.9.5, “Incremental Replication Does Not Transfer Data from Source to Target for LVM Volumes on RHEL 6.7, Oracle Linux 6.7, and CentOS 6.7,” on page 31
- Section 8.9.6, “Configure Fails for Migration of Oracle Linux 5.x with Paravirtual Kernel to VMware,” on page 31

8.9.1 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.

8.9.2 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.

8.9.3 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.

8.9.4 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.

8.9.5 Incremental Replication Does Not Transfer Data from Source to Target for LVM Volumes on RHEL 6.7, Oracle Linux 6.7, and CentOS 6.7

Issue: On Red Hat Enterprise Linux (RHEL) 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads running kernel version `2.6.32-573.el6`, incremental replications do not transfer data from source to target for LVM volumes. Full replications and server sync work with this kernel version. (`Bug 1018176`)

Workaround: Do one of the following for RHEL 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads if they have LVM volumes:

- Run only full replications or server sync for workloads running the RHEL kernel version `2.6.32-573.el6`.
- Register with the Red Hat network and update the workload by using the Red Hat Package Manager (`yum update`) to get the latest available kernel fixes for RHEL 6.7. The workload is updated to use the same kernel version `2.6.32-642.13.1.el6` that is used by the RHEL 6.8 distribution. Both full and incremental replications perform as expected with this kernel version.

  The following blkwatch drivers are available for RHEL 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads running kernel `2.6.32-642`:

  ```
  RHEL6-RHSA201700361-2.6.32-642.13.1.el6.i686-x8
  RHEL6-RHSA201700361-2.6.32-642.13.1.el6.x86_64-x86_64
  ```

8.9.6 Configure Fails for Migration of Oracle Linux 5.x with Paravirtual Kernel to VMware

Issue: Configuration fails for migration of a Oracle Linux 5.x with a Paravirtual kernel under Citrix XenServer to a target VM under VMware. One of the following errors might be seen in the Migrate Client or the Web Interface:

64-bit

```
PlateSpin.Forge.DataAccess.Connector.PowerConvert.ConnectorException: Wrapped
IndexOutOfRangeException --->
System.IndexOutOfRangeException: Index was outside the bounds of the array.
```

32-bit

```
Error, (Source) The boot loader unknown is not supported
```

In Migrate Client, no NIC displays in the **Source Take Control** adapter field. (`Bugs 1001424, 1001433, 1001436`)

Workaround: None. You cannot migrate an Oracle Linux 5.x (32-bit or 64-bit) workload that is paravirtualized under Citrix XenServer to VMware.

The VMware Paravirtual SCSI (PVSCSI) adapter does not support boot disks for guest operating systems running Red Hat Enterprise Linux 5.x (32-bit and 64-bit). This restriction applies to distributions based on RHEL 5.x such as Oracle Linux 5.x and CentOS 5.x. See the VMware Knowledge Base article "Configuring disks to use VMware Paravirtual SCSI (PVSCSI) adapters (1010398)" (https://kb.vmware.com/selfservice/search.do?cmd=displayKC&docType=kc&docTypeID=DT_KB_1_1&externalId=1010398).

9 Supplemental Documentation

9.1 Linux Boot Partition Requirement

For Linux workloads, the boot partition of the source workload must have a minimum of 100 MB free space. PlateSpin Migrate creates a new `initrd` image with required drivers on the source `/boot` directory for the initial boot of the target workload.

10 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 © 2017 NetIQ Corporation. All Rights Reserved.