MICRO FOCUS

PlateSpin® Migrate 2020.2

User Guide

September 2020

Legal Notice

Contents

About This Guide

This guide provides information about using PlateSpin Migrate.

Audience

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

Additional Documentation

This guide is part of the PlateSpin Migrate documentation set. For a complete list of publications supporting this release, visit the PlateSpin Migrate Documentation website (https://www.microfocus.com/documentation/platespin/platespin-migrate-2020-2).

Documentation Updates

The most recent version of this guide can be found at the PlateSpin Migrate Documentation website (https://www.microfocus.com/documentation/platespin/platespin-migrate-2020-2).

Contacting Micro Focus

For specific product issues, contact Micro Focus Support at https://support.microfocus.com/contact/.

Additional technical information or advice is available from several sources:

- Product information and resources: https://www.microfocus.com/products/platespin-migrate/
- Micro Focus Customer Center: https://www.microfocus.com/customercenter/
- Product knowledge base and videos: https://www.microfocus.com/support-and-services/
- Micro Focus Communities for PlateSpin: https://community.microfocus.com/t5/PlateSpin/ct-p/PlateSpin/
- PlateSpin Idea Exchange: https://community.microfocus.com/t5/PlateSpin-Idea-Exchange/idb-p/PlateSpin_Ideas/tab/most-recent/

Overview and Planning

PlateSpin Migrate enables you to migrate heterogeneous workloads across x86-based physical, virtual, and cloud infrastructures in your data center. It decouples the workload infrastructure from its software (operating system, applications, and data) to allow any-to-any migrations. Migrate provides tools to easily discover workloads and hosts in your environment. You can efficiently configure, execute, and test workload even before the actual cutover, and also monitor the status of workload migration. With Migrate, you can dramatically increase the migration speed and success ratios, which help reduce the costs for your migration projects.

1 Overview of Workload Migration

This section provides an overview of the workload migration scenarios and helps you understand the workload migration.

Workload Migration Scenarios

PlateSpin Migrate is designed to be used for the following scenarios:

- **Consolidation:** Automating large-scale migrations of physical machines to virtual machines, accelerating consolidation projects, and reducing administrative effort and errors.
- **Continuous Workload Optimization:** Moving workloads to and from any geographical location, onto any platform, in any direction. Workloads can be virtualized or de-virtualized during ongoing and continuous optimization of resources.
- **Migration:** Moving fully configured workloads from old hardware to new hardware without rebuilding the entire software stack.
- **Maintenance and Support Agreement Integrity:** De-virtualizing workloads along with the applications installed on them and moving them back to physical machines over the network so that the support agreements can remain valid.
- **Migration to Cloud:** Moving workloads to cloud platforms such as Amazon Web Services (AWS), Microsoft Azure, Oracle Cloud Infrastructure, and VMware vCloud Director.
- **Data Center Relocation:** Relocating data center from one geographical location to another.
- **Test Lab Deployment:** Consolidating test lab workloads by running multiple virtual machines on a single VM host, quickly deploying virtual test lab environments with ease, and replicating an entire production environment in matter of hours or days.

Understanding Workload Migration

PlateSpin Migrate automates the migration of workloads among physical, virtual machine, and cloud. The supported cloud platforms include Amazon Web Services (AWS), Microsoft Azure, Oracle Cloud Infrastructure, and VMware vCloud Director.

Figure 1-1 *Workload Migration*

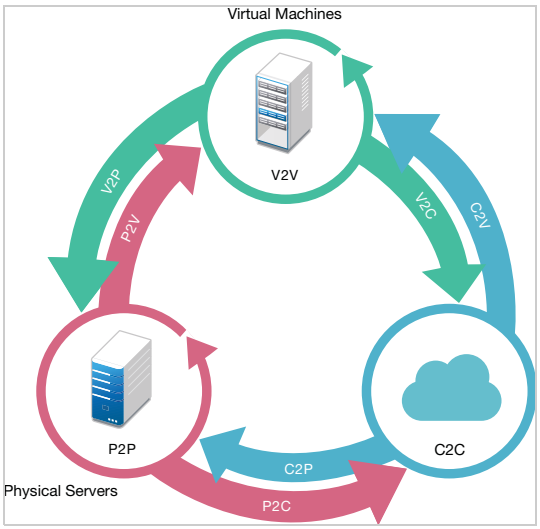

Table 1-1 *Workload Migration Operations*

Category of Operation	Migration Infrastructures
Peer-to-peer	◆ Physical to Virtual (P2V) ◆ Virtual to Virtual (V2V) ◆ Virtual to Physical (V2P) ◆ Physical to Physical (P2P)
Cloud	◆ Physical to Cloud (P2C) ◆ Virtual to Cloud (V2C) ◆ Cloud to Physical (C2P) ◆ Cloud to Virtual (C2V) ◆ Cloud to Cloud (C2C) **NOTE:** Supported Cloud platforms include Amazon Web Services (AWS), Microsoft Azure, Oracle Cloud Infrastructure, and VMware vCloud Director.

PlateSpin Migrate supports multiple workload types and virtualization platforms. For a more detailed list of supported workloads and infrastructures, see “Supported Configurations” on page 23.

2 Planning Your Workload Migrations

This section describes the configuration requirements and setup for PlateSpin Migrate. Use the information in this section to plan your migration environment.

Supported Configurations

Supported Source Workloads For Migration to Non-Cloud Platforms

PlateSpin Migrate supports the migration of the following Windows and Linux workloads to non-cloud platforms, such as physical machines and virtual machines on supported hypervisors. See "Supported Target Virtualization Platforms" on page 40.

The following migration features are supported for migration to non-cloud platforms:

- Peer-to-peer migrations (P2V, V2V, V2P, P2P).
- Peer-to-peer workload synchronization (P2V, V2V, P2P, V2P).

Review the following sections:

Requirements for Migration to Non-Cloud Platforms

- Not all workloads are supported on all target virtualization platforms. Migration of workloads to a target virtualization platform is subject to the support of the guest operating system on the target host by the host vendor.
- Before you install block-based transfer drivers on source Windows workloads, ensure that you have applied the latest Windows updates on the workload.
- BIOS workloads must have at least one partition in the boot disk and a boot loader installed in the MBR (Master Boot Record).
- Conversion of BIOS based Linux system to UEFI based is not supported.
- Conversion of a Linux UEFI source workload as a Linux BIOS target requires that a `/boot` partition is available on the source workload.

Supported Windows Workloads For Migration to Non-Cloud Platforms

PlateSpin Migrate supports the following Microsoft Windows platforms for migration to virtual machines on virtualization hosts or to physical machines, except as noted in Table 2-1. See also "Supported Workload Storage" on page 35 and "Supported Workload Architectures" on page 38.

NOTE: PlateSpin Migrate does not support migration of Active Directory domain controller servers with Flexible Single Master Operation (FSMO) roles on them. For information, see *Best Practice Tips for Active Directory Domain Controller Conversions* (KB Article 7920501).

Table 2-1 *Non-Cloud Platforms: Supported Windows Workloads*

Operating System	Remarks
Servers	
Windows Server 2019	
Windows Server 2016	Migration to a VMware VM requires VMware vCenter 6.0 or later.
Windows Server 2012 R2 Windows Server 2012	
Windows Server 2008 R2	Includes Small Business Server (SBS) editions.

Operating System	Remarks
Clusters	
Windows Server 2016 and 2019 Cluster Supports quorum models: ◆ Node and Disk Majority ◆ No Majority: Disk Only	Both Migrate Client and Web Interface support automated migration of Windows Clusters to VMware vCenter target virtualization platforms. Migrate Client also supports semi-automated migration of Windows Clusters to physical machines by using the X2P workflow. See "Preparing for Migration of Windows Clusters" on page 327. Migration of Windows Server 2016 Clusters to VMware requires VMware 6.0 or later. PlateSpin Migrate does not support migration of Windows Server clusters to the following target infrastructures: ◆ Cloud ◆ Virtualization hypervisors other than VMware PlateSpin Migrate supports only block-level replication for clusters. File-level replication is not supported. For clusters, PlateSpin Migrate provides driverless block-based transfer and driver-based block-based transfer methods. See "Block-Based Transfer for Clusters" on page 329.
Windows Server 2012 R2 Cluster Windows Server 2012 Cluster Supports quorum models: ◆ Node and Disk Majority ◆ No Majority: Disk Only	
Windows Server 2008 R2 Cluster Supports quorum models: ◆ Node and Disk Majority ◆ No Majority: Disk Only	

Supported Linux Workloads For Migration to Non-Cloud Platforms

PlateSpin Migrate supports the following Linux platforms for migration to virtual machines on virtualization hosts or to physical machines, except as noted in Table 2-2. See also "Supported Workload Storage" on page 35 and "Supported Workload Architectures" on page 38.

NOTE: To install Migrate Agent Utility for Linux, your source machine must has GNU C Library (glibc) 2.11.3 or higher installed.

Table 2-2 *Non-Cloud Platforms: Supported Linux Workloads*

Linux Distribution	Versions	Remarks
Red Hat Enterprise Linux (RHEL)	5.0 to 5.11, 6.0 to 6.10, and 7.0 to 7.7	For Red Hat Enterprise Linux 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version `2.6.32-642.13.1.el6`) for the 6.7 distribution. For Red Hat Enterprise Linux 6.8, Oracle Linux 6.8, and CentOS 6.8 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version `2.6.32-696.20.1.el6.x86_64`) for the 6.8 distribution. For Red Hat Enterprise Linux 6.10, Oracle Linux 6.10, and CentOS 6.10 workloads, incremental replication is not supported for kernel (version `2.6.32-754.6.3.el6.x86_64`) for the 6.10 distribution. See "RHEL 6.10 Source Workloads Crash During Incremental Replication" on page 568. Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for RHEL 5.*x* and 6.0. See "Paravirtualized Source Workloads" on page 40.
SUSE Linux Enterprise Server (SLES)	12 SP4 without BTRFS file systems 11 (SP1, SP2, SP3, SP4)	The SLES 11 SP2 (32-bit) with kernel `3.0.13-0.27-pae` is not supported. The kernel for this version of SLES must be upgraded to `3.0.51-0.7.9-pae` so that the conversion works. Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for SLES 11. See "Paravirtualized Source Workloads" on page 40.
CentOS	See Red Hat Enterprise Linux.	Migration of CentOS 7.*x* to VMware requires VMware vCenter 5.5 or later.
Oracle Linux (OL) (formerly Oracle Enterprise Linux)	See Red Hat Enterprise Linux.	Same level of support for standard kernels as that for workloads running RHEL. Same level of support for Unbreakable Enterprise Kernel (UEK) kernels on supported RHEL distributions for OL 6.7 and later.

Supported Source Workloads for Migration to Cloud Platforms

Use the PlateSpin Migrate Web Interface to migrate the workloads to Amazon Web Services, Microsoft Azure, and VMware vCloud Director.

Migrate supports P2C and V2C migrations to target cloud platforms. Migrate supports C2C migrations of source workloads between supported cloud platforms.

Review the following sections:

Requirements for Migration to Cloud Platforms

- Not all workloads are supported on all target cloud platforms. Migration of workloads to a cloud platform is subject to the support of the guest operating system on the target cloud platform by the cloud provider.
- Before you install block-based transfer drivers on source Windows workloads, ensure that you have applied the latest Windows updates on the workload.
- If the boot (`/boot`) partition is on a different disk than the root (`/`) partition, PlateSpin Migrate migrates them both to the first disk on the target VM.
- BIOS workloads must have at least one partition in the boot disk and a boot loader installed in the MBR (Master Boot Record).
- Windows and Linux UEFI workloads are migrated as UEFI workloads to the target vCloud platforms. However, for other target cloud platforms such as Azure and AWS that do not support UEFI workloads, Windows and Linux UEFI workloads are migrated as BIOS workloads.
- Conversion of a Linux UEFI source workload as a Linux BIOS target requires a `/boot` partition to be available on the source workload.
- Before you migrate a paravirtualized Linux source workload running on Citrix XenServer or KVM to a target platform as fully virtualized guest, see "Paravirtualized Source Workloads" on page 40.
- To install Migrate Agent Utility for Linux, your source machine must have GNU C Library (glibc) 2.11.3 or higher installed.

Supported Workloads For Migration to Amazon Web Services

PlateSpin Migrate supports the following platforms for migration to Amazon Web Services, as specified in Amazon Web Services (AWS) in Table 2-15, "Supported Target Cloud Platforms for the Migrate Web Interface," on page 43. See also "Supported Workload Storage" on page 35 and "Supported Workload Architectures" on page 38.

For information about migrating workloads to Microsoft Amazon Web Services, see:

- “Prerequisites for C2C Migration from Azure to AWS” on page 222
- “Prerequisites for C2C Migration from vCloud to AWS” on page 237
- Chapter 25, “Migration to Amazon Web Services,” on page 431

NOTE: PlateSpin Migrate does not support migration of Active Directory domain controller servers with Flexible Single Master Operation (FSMO) roles on them. For information, see *Best Practice Tips for Active Directory Domain Controller Conversions* (KB Article 7920501).

Table 2-3 *AWS: Supported Windows Platforms*

Operating System	Remarks
Windows Server 2019	For migration of Windows Server 2019 workload to AWS cloud, only t2 instances are supported.
Microsoft Windows Server 2016	Nitro System based Instance types support Windows Server 2008 R2 and later. ◆ AWS requires a minimum of Microsoft .NET Framework 4.5 for migration of Windows workloads to Nitro System based VM instances. To use Nitro System based instances, ensure that .NET Framework 4.5 or later is installed on the source Windows workload prior to Migrate discovery. ◆ Migration of Windows Domain Controller servers to AWS requires Nitro System based instance types such as c5, m5, r5, t3, c5d, m5d, and z1d. Ensure that you select a Nitro System based instance when you configure the migration.
Microsoft Windows Server 2012 R2	
Microsoft Windows Server 2012	
Microsoft Windows Server 2008 R2	

Table 2-4 *AWS: Supported Linux Platforms*

Linux Distribution	Versions	Remarks
Red Hat Enterprise Linux (RHEL)	7.0 to 7.7 6.1 to 6.10 5.1 to 5.11	Migration of RHEL, CentOS, and Oracle Linux workloads using Nitro System based instance types such as c5, m5, r5, t3, c5d, m5d, or z1d is supported for versions 7.4 and later. For Red Hat Enterprise Linux 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version `2.6.32-642.13.1.el6`) for the 6.7 distribution. For Red Hat Enterprise Linux 6.8, Oracle Linux 6.8, and CentOS 6.8 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version `2.6.32-696.20.1.el6.x86_64`) for the 6.8 distribution. For Red Hat Enterprise Linux 6.10, Oracle Linux 6.10, and CentOS 6.10 workloads, incremental replication is not supported for kernel (version `2.6.32-754.6.3.el6.x86_64`) for the 6.10 distribution. See "RHEL 6.10 Source Workloads Crash During Incremental Replication" on page 568. Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for RHEL 5.*x* and 6.0. See "Paravirtualized Source Workloads" on page 40.
SUSE Linux Enterprise Server (SLES)	12 SP4 without BTRFS file systems 11 (SP1, SP2, SP3, SP4)	Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for SLES 11. See "Paravirtualized Source Workloads" on page 40.
CentOS	See Red Hat Enterprise Linux.	Same level of support as that for workloads running RHEL.
Oracle Linux (OL) (formerly Oracle Enterprise Linux)	See Red Hat Enterprise Linux.	Same level of support for standard kernels as that for workloads running RHEL. Same level of support for Unbreakable Enterprise Kernel (UEK) kernels on supported RHEL distributions for OL 6.7 and later.

Supported Workloads For Migration to Microsoft Azure

PlateSpin Migrate supports the following platforms for migration to Microsoft Azure Cloud, as specified in Microsoft Azure in Table 2-15, "Supported Target Cloud Platforms for the Migrate Web Interface," on page 43. See also "Supported Workload Storage" on page 35 and "Supported Workload Architectures" on page 38.

NOTE: For information about installing Azure VM Agent on target workloads, see "Prerequisites for Installing Azure VM Agent" on page 175.

For information about migrating workloads to Microsoft Azure, see:

- Chapter 9, "Prerequisites for Migration to Microsoft Azure," on page 167
- "Prerequisites for C2C Migration from Azure to Azure" on page 219
- "Prerequisites for C2C Migration from AWS to Azure" on page 220
- "Prerequisites for C2C Migration from vCloud to Azure" on page 229
- Chapter 26, "Migration to Microsoft Azure," on page 453

NOTE: PlateSpin Migrate does not support migration of Active Directory domain controller servers with Flexible Single Master Operation (FSMO) roles on them. For information, see *Best Practice Tips for Active Directory Domain Controller Conversions* (KB Article 7920501).

Table 2-5 *Azure: Supported Windows Platforms*

Operating System	Remarks
Microsoft Windows Server 2019	
Microsoft Windows Server 2016	Microsoft Azure supports Azure VM Agent for Windows Server 2008 R2 and higher.
Microsoft Windows Server 2012 R2	
Microsoft Windows Server 2012	
Microsoft Windows Server 2008 R2	
Microsoft Windows Server 2008 (64-bit)	Microsoft .NET Framework 3.5 SP1 or higher is required on the source Windows Server 2008 workload. PlateSpin Migrate does not support Windows Server 2008 32-bit OS versions for migration to Azure. PlateSpin Migrate does not support installation of Azure VM Agent for Windows Server 2008 (non-R2).

Table 2-6 *Azure: Supported Linux Platforms*

Linux Distribution	Versions	Remarks
Red Hat Enterprise Linux (RHEL)	7.1 to 7.7 6.7 to 6.10	For Red Hat Enterprise Linux 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version `2.6.32-642.13.1.el6`) for the 6.7 distribution. For Red Hat Enterprise Linux 6.8, Oracle Linux 6.8, and CentOS 6.8 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version `2.6.32-696.20.1.el6.x86_64`) for the 6.8 distribution. For Red Hat Enterprise Linux 6.10, Oracle Linux 6.10, and CentOS 6.10 workloads, incremental replication is not supported for kernel version `2.6.32-754.6.3.el6.x86_64` for the 6.10 distribution. See "RHEL 6.10 Source Workloads Crash During Incremental Replication" on page 568.
SUSE Linux Enterprise Server (SLES)	12 SP4 without BTRFS file systems 11 SP4	Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported. See "Paravirtualized Source Workloads" on page 40.
CentOS	See Red Hat Enterprise Linux.	Same level of support as that for workloads running RHEL.
Oracle Linux (OL) (formerly Oracle Enterprise Linux)	See Red Hat Enterprise Linux.	Same level of support for standard kernels as that for workloads running RHEL. Same level of support for Unbreakable Enterprise Kernel (UEK) kernels on supported RHEL distributions for OL 6.7 and later.

Supported Workloads For Migration to VMware vCloud Director

PlateSpin Migrate supports the following platforms for migration to VMware vCloud Director, as specified in VMware vCloud Director in Table 2-15, "Supported Target Cloud Platforms for the Migrate Web Interface," on page 43. See also "Supported Workload Storage" on page 35 and "Supported Workload Architectures" on page 38.

For information about migrating workloads to VMware vCloud Director, see:

- Chapter 10, "Prerequisites for Migration to VMware vCloud Director," on page 201
- "Prerequisites for C2C Migration from Azure to vCloud" on page 226
- "Prerequisites for C2C Migration from AWS to vCloud" on page 233
- Chapter 27, "Migration to VMware vCloud Director," on page 469

NOTE: PlateSpin Migrate does not support migration of Active Directory domain controller servers with Flexible Single Master Operation (FSMO) roles on them. For information, see *Best Practice Tips for Active Directory Domain Controller Conversions* (KB Article 7920501).

Table 2-7 *vCloud: Supported Windows Platforms*

Operating System	Remarks
Microsoft Windows Server 2019	
Microsoft Windows Server 2016	Requires vCloud 8.20 or higher. The hosts backing the VMware resource pool must support VMs with Hardware Version 10 or higher. The Provider VDC policy for the highest supported hardware version must be set to at least Hardware Version 10.
Microsoft Windows Server 2012 R2	
Microsoft Windows Server 2012	
Microsoft Windows Server 2008 R2	

Table 2-8 *vCloud: Supported Linux Platforms*

Linux Distribution	Versions	Remarks
Red Hat Enterprise Linux (RHEL)	7.0 to 7.7 6.0 to 6.10 5.0 to 5.11	Migrate supports XFS v5 file system on source Linux UEFI and BIOS workloads for migrations using the vCloud PRE based on SLES 12 SP3. For Red Hat Enterprise Linux 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version `2.6.32-642.13.1.el6`) for the 6.7 distribution. For Red Hat Enterprise Linux 6.8, Oracle Linux 6.8, and CentOS 6.8 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version `2.6.32-696.20.1.el6.x86_64`) for the 6.8 distribution. For Red Hat Enterprise Linux 6.10, Oracle Linux 6.10, and CentOS 6.10 workloads, incremental replication is not supported for kernel (version `2.6.32-754.6.3.el6.x86_64`) for the 6.10 distribution. See "RHEL 6.10 Source Workloads Crash During Incremental Replication" on page 568. Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for RHEL 5.*x* and 6.0. See "Paravirtualized Source Workloads" on page 40. Migration of Red Hat Enterprise Linux 7.x workloads is supported only for VMware vCloud Director 9.1.
SUSE Linux Enterprise Server (SLES)	12 SP4 without BTRFS file systems 11 (SP1, SP2, SP3, SP4)	Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for SLES 11. See "Paravirtualized Source Workloads" on page 40.
CentOS	See Red Hat Enterprise Linux.	Same level of support as that for workloads running RHEL.
Oracle Linux (OL) (formerly Oracle Enterprise Linux)	See Red Hat Enterprise Linux.	Same level of support for standard kernels as that for workloads running RHEL. Same level of support for Unbreakable Enterprise Kernel (UEK) kernels on supported RHEL distributions for OL 6.7 and later.

Supported Workloads For Migration to Oracle Cloud Infrastructure

PlateSpin Migrate supports the following platforms for migration to Compute - Virtual Machine Instances on Oracle Cloud Infrastructure. See also "Supported Workload Storage" on page 35 and "Supported Workload Architectures" on page 38.

For information about migrating workloads to Oracle Cloud Infrastructure, see:

- Chapter 11, "Prerequisites for Migration to Oracle Cloud Infrastructure," on page 209
- Chapter 28, "Migration to Oracle Cloud Infrastructure," on page 481

NOTE: PlateSpin Migrate does not support migration of Active Directory domain controller servers with Flexible Single Master Operation (FSMO) roles on them. For information, see *Best Practice Tips for Active Directory Domain Controller Conversions* (KB Article 7920501).

Table 2-9 *Oracle Cloud: Supported Windows Platforms*

Operating System	Remarks
Microsoft Windows Server 2019	
Microsoft Windows Server 2016	
Microsoft Windows Server 2012 R2	
Microsoft Windows Server 2012	
Microsoft Windows Server 2008 R2	

Table 2-10 *Oracle Cloud: Supported Linux Workloads*

Linux Distribution	Versions	Remarks
Red Hat Enterprise Linux (RHEL)	5.5, 5.6, 5.9, 5.11, 6.5, 6.9, and 7.0 to 7.7	Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for RHEL 5.*x*. See "Paravirtualized Source Workloads" on page 40.
SUSE Linux Enterprise Server (SLES)	12 SP4 without BTRFS file systems 11 (SP1, SP2, SP3, SP4)	The SLES 11 SP2 (32-bit) with kernel `3.0.13-0.27-pae` is not supported. The kernel for this version of SLES must be upgraded to `3.0.51-0.7.9-pae` so that the conversion works. Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for SLES 11. See "Paravirtualized Source Workloads" on page 40.
CentOS	5.11, 6.9, and 7.0 to 7.7	Migration of CentOS 7.*x* to VMware requires VMware vCenter 5.5 or later.
Oracle Linux (OL) (formerly Oracle Enterprise Linux)	5.8, 5.11, 6.2, 6.5, and 7.0 to 7.7	

Supported Workload Storage

The following workload storage guidelines apply to all migrations:

Partitioning Schemes

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

Windows File Systems

PlateSpin Migrate supports only the NTFS file system on any supported Windows system. It does not support Windows FAT or ReFS file systems for migration.

NOTE: If the volumes are encrypted with the BitLocker disk encryption feature, they must be unlocked (decrypted) for the migration.

Linux File Systems

PlateSpin Migrate supports EXT2, EXT3, EXT4, REISERFS, and XFS file systems.

NOTE

- PlateSpin Migrate supports the XFS version 5 (v5) file system on RHEL 7.3 and later, and on distributions based on those versions.
- Migration of encrypted volumes is not supported. If the volumes are encrypted, they must be unlocked (decrypted) for the migration.

Disks

PlateSpin Migrate supports several types of storage disks, including basic disks, source Windows dynamic disks, LVM2, hardware RAID, NAS, and SAN.

NOTE: The following caveats apply for storage disks:

- **Windows Dynamic Disks:** PlateSpin Migrate does not support Windows dynamic disks at the target.

  For dynamic disks, the target storage does not follow the Same as Source mapping strategy. Dynamic Volumes (Simple, Spanned, Striped, Mirrored, or RAIDed) will be migrated as Simple Volumes on basic disks on the target workload. The target disk is partitioned as GPT if the total combined size of the dynamic volume's member disks exceeds MBR partition size limits. For more information, see *Microsoft TechNet: Understanding the 2 TB limit in Windows Storage* (https://blogs.technet.microsoft.com/askcore/2010/02/18/understanding-the-2-tb-limit-in-windows-storage/).

- **Software RAID:** PlateSpin Migrate supports hardware RAID; however, PlateSpin Migrate does not support software RAID. This is applicable for both Windows and Linux workloads.

Windows Disks, Partitions, and Volumes

- iSCSI disks mounted on a source workload are migrated as local disks to the target workload.
- Ensure that the free space on the source Windows workload volumes is as follows:

Workload Volume	Free Space	Remarks
C:	200 MB or 10% of the total volume, which ever is higher	PlateSpin Migrate uses the free space on this volume to: • Install the required PlateSpin components • Take snapshot of this volume
Other Volumes	10% of the total volume	PlateSpin Migrate uses the free space on this volume to take snapshot of the volume.

Linux Disks, Partitions, and Volumes

- Migrate supports GRUB and GRUB 2 boot loaders for Linux workloads.
- Migrate supports Linux workloads with `/boot` on the first disk (`sda`).
- iSCSI disks mounted on a source workload are migrated as local disks to the target workload.
- Ensure that the free space on the source Linux workload volumes is as follows:

Workload Volume	Free Space	Remarks
/boot	600 MB or 10% of the total volume, which ever is higher	PlateSpin Migrate uses the free space on this volume to: ◆ Create a new `initrd` image with all the required drivers to make the machine ready for the initial boot process ◆ Take snapshot of this volume
/usr	600 MB or 10% of the total volume, which ever is higher	PlateSpin Migrate uses the free space on this volume to: ◆ Install the required PlateSpin components ◆ Take snapshot of this volume
/	700 MB or 10% of the total volume, which ever is higher	PlateSpin Migrate uses the free space on this volume to: ◆ Install the required PlateSpin components ◆ Take snapshot of this volume
Other Volumes	10% of the total volume	PlateSpin Migrate uses the free space on this volume to take snapshot of the volume.

- Non-volume storage, such as a swap partition that is associated with the source workload, is recreated in the migrated workload.
- The layout of volume groups and logical volumes for LVM2 is preserved in the Same as Source mapping strategy so that you can re-create it during migration.
- LVM raw disk volumes are supported in the Same as Source configurations on Linux workloads.

Linux Block-Based Transfer

For Linux workloads using the Block-Based Transfer method, PlateSpin Migrate requires block-based Linux Kernel drivers, called *block watch* (`blkwatch`) drivers, to perform the block-based data transfer. The driver must be built for the specific kernel running on the source Linux workload. See "Block-Based Transfer for Linux" on page 46.

FC SANs

PlateSpin Migrate supports the Fibre Channel (FC) SAN communications protocol.

FCoE SANs

Fibre Channel over Ethernet (FCoE) is supported for P2P and P2V migrations for workloads listed in Table 2-11. Migration has been tested using FCoE devices from Qlogic.

Table 2-11 *Supported Source Workloads for FCoE*

Source Workloads with FCoE	Version	Remarks
Windows Server	2012 R2 2008 R2	Standalone servers only; no clusters.
SUSE Linux Enterprise Server	11 SP4	

FCoE drivers and support functionality are available in the PlateSpin ISO image. See “Downloading the PlateSpin ISO Images” on page 373.

Multipath I/O

PlateSpin Migrate supports migration of a source workload with multipath I/O (MPIO) enabled. However, MPIO on the migration target is not supported. So, we recommend that you perform migration with a single path, and then enable MPIO on the target workload after cutover.

Supported Workload Architectures

The following workload architecture guidelines apply to all migrations:

Protocols

- Linux source workloads must be running a Secure Shell (SSH) server.

Processors

PlateSpin Migrate supports migration of x86-based physical and virtual workloads in your data center:

- 64-bit
- 32-bit

Cores and Sockets for Target VMs

For VM virtualization platforms using VMware 5.1, 5.5, and 6.0 with a minimum VM hardware Level 8, PlateSpin Migrate enables you to specify the number of sockets and the number of cores per socket for the target workload. It automatically calculates the total cores. This parameter applies on the initial setup of a workload with an initial replication setting of **Full Replication**.

NOTE: The maximum number of cores the workload can use is subject to external factors such as the guest operating system, the VM hardware version, VMware licensing for the ESXi host, and ESXi host compute maximums for vSphere (see *ESXi/ESX Configuration Maximums (VMware KB 1003497)* (https://kb.vmware.com/kb/1003497)).

Some distributions of a guest OS might not honor the cores and cores per socket configuration. Instead they will retain their original cores and sockets settings as installed.

UEFI and BIOS Firmware

Migration of UEFI-based Windows and Linux source workloads is supported for all target platforms. The target workload is configured as UEFI or BIOS, as supported by the target platform vendor. For example:

- For target vCloud Cloud Director platforms, Windows and Linux UEFI workloads are migrated as UEFI workloads to the target vCloud platforms.
- For target cloud platforms such as Azure and AWS that do not support UEFI workloads, Windows and Linux UEFI workloads are migrated as BIOS workloads.

Migrate transfers workloads from source to target while enforcing the supported firmware for the respective source and target operating systems. When any migration between UEFI and BIOS systems are initiated, Migrate analyzes the transition and alerts you about its validity.

NOTE: If you are migrating UEFI-based workload onto vSphere target virtualization platform and you want to continue using the same firmware boot mode, you must target a vSphere 5.0 platform or newer.

The following are examples of Migrate behavior when doing conversion between UEFI and BIOS-based systems:

- When you migrate a UEFI-based source workload to platform that does not support UEFI, such as to a VMware vSphere 4.x, AWS, or Azure, Migrate transitions the workload's UEFI firmware to BIOS firmware.
- When you migrate a UEFI-based source workload to a BIOS-based target, Migrate converts the UEFI system's boot disks, which were GPT, to MBR disks.
- (For Windows Workloads) When you migrate a BIOS workload to a UEFI-based target, Migrate converts the BIOS system's boot disks, which are MBR, to GPT disks.

Paravirtualized Source Workloads

Paravirtualized guest VM to fully virtualized guest VM conversion is supported for the following source workloads hosted on a Citrix XenServer virtual host or KVM virtual host:

- Red Hat Enterprise Linux (RHEL) 6.0 and Linux distributions based on RHEL 6.0
- Red Hat Enterprise Linux (RHEL) 5.*x* and Linux distributions based on RHEL 5.x
- SUSE Linux Enterprise Server 10 and 11

NOTE: Only block-based data transfer is supported for paravirtualized source workloads.

Before you migrate a paravirtualized Linux source workload that is hosted on a Citrix XenServer or KVM virtual host to a target platform as a fully virtualized workload, do the following:

- Ensure that both the paravirtualized kernel and standard kernel are installed on the paravirtualized Linux source workload.
- Manually compile the block-based drivers for Xen kernel. See "Linux Distributions Supported by Migrate" on page 343.
- For semi-automated (X2P) migrations, ensure that your target is a fully virtualized (not paravirtualized) VM.

Supported Target Virtualization Platforms

PlateSpin Migrate supports the following target virtualization platforms:

- VMware vCenter and ESXi
- KVM

NOTE

- Migration of workloads to a target virtualization platform is subject to the support of the guest operating system on the target host by the host vendor.
- You need an OS license for the migrated target workload.

Review the following sections for information about supported versions and requirements.

- "Supported Target VMware Platforms" on page 40
- "Supported Target KVM Platforms" on page 42

Supported Target VMware Platforms

Table 2-12 lists supported target VMware platforms. Table 2-12 lists supported VMware datastore types.

Migrations to VMware platforms are supported using the PlateSpin Migrate Web Interface and Migrate Client. For information about migration methods, see:

- "Automated Migration to VMware Using Migrate Client" on page 489

- "Migration to VMs on VMware Using X2P Workflow" on page 506
- "Automated Migration to VMware Using Migrate Web Interface" on page 498

See also "Prerequisites for Migration to VMware" on page 247.

NOTE

- PlateSpin Migrate does not support discovery, configuration, and migration actions for a target VMware DRS Cluster where one or more hosts are in maintenance mode. See "Best Practices for Maintaining or Updating VMware Environments That Are Configured as Migration Targets" on page 265.
- For information about creating the target VM disk on VMware platforms using Raw Device Mapping (RDM), see "Migration to VMware" on page 487.

Table 2-12 *Supported Target VMware Platforms for the Migrate Web Interface and Migrate Client*

Platform	Versions	Remarks
VMware vCenter	6.7 (U1, U2, U3) 6.5 (U1 with latest patches, U2, U3) 6.0 (U1, U2, and U3) 5.5 (U1, U2, and U3)	Raw Device Mapping (RDM) for target VMs is supported using the X2P workflow. See also Table 2-13, "Supported VMware Datastore Types," on page 42.
VMware vSAN (virtual SAN)	• vSAN 6.7 U1 on vCenter 6.7 U1 platforms • vSAN 6.7 on vCenter 6.7 platforms • vSAN 6.5 U2 on vCenter 6.5 U2 platforms • vSAN 6.6 on vCenter 6.5 (U1) platforms • vSAN 6.2 on vCenter 6.0 platforms • vSAN 5.5 on vCenter 5.5 platforms	VMware vSAN storage is supported on vCenter target virtualization platforms as noted.
VMware ESXi	6.7 (U1, U2, U3) 6.5 (U1 with latest patches, U2, U3) 6.0 (U1, U2, and U3) 5.5 (U1, U2, and U3)	All ESXi versions must have a paid license; migration is unsupported with these systems if they are operating with a free license. Raw Device Mapping (RDM) for target VMs is supported using the X2P workflow. See also Table 2-13, "Supported VMware Datastore Types," on page 42.

Table 2-13 *Supported VMware Datastore Types*

Datastore Type	Supported Configurations
VMFS	Supported for all supported versions of VMware vCenter and ESXi platforms.
NFS	◆ **NFS v3**: For all supported versions of VMware vCenter and ESXi platforms ◆ **NFS v4.1**: For all supported versions of VMware vCenter 6.*x* and ESXi 6.*x* platforms
Other	Other datastore types are not supported, such as Virtual Volumes, and vFlash.

Supported Target KVM Platforms

Table 2-14 lists supported target KVM platforms. Migrations to KVM are supported only for the PlateSpin Migrate Client using the semi-automated X2P workflow. For information, see "Migration to Virtual Machines on KVM" on page 511.

See also "Prerequisites for Migration to VMs on KVM" on page 267.

Table 2-14 *Supported Target KVM Platforms for the Migrate Client Only*

Platform	Versions	Remarks
SUSE Linux Enterprise Server (SLES) with KVM	11 SP4 and 12 SP1	Fully virtualized guests are supported. Virtio devices are supported. Supported through the X2P workflow. See Migration to Virtual Machines on KVM. See also "Prerequisites for Migration to VMs on KVM" on page 267.
Red Hat Enterprise Linux (RHEL) with KVM	7.4	Fully virtualized guests are supported. Virtio devices are supported. Supported through the X2P workflow. See Migration to Virtual Machines on KVM. See also "Prerequisites for Migration to VMs on KVM" on page 267.

Supported Target Cloud Platforms

PlateSpin Migrate supports migration of workloads to target cloud platforms in the Migrate Web Interface.

Table 2-15 *Supported Target Cloud Platforms for the Migrate Web Interface*

Platform	Versions	Remarks
Amazon Web Services (AWS)	◆ AWS Global ◆ AWS China ◆ AWS Government	See also Chapter 8, " Prerequisites for Migration to Amazon Web Services," on page 145. A Migrate server can have multiple AWS Cloud target platforms. You specify the AWS Cloud environment and Regions when you create the target platform
Microsoft Azure	◆ Azure Global	See also Chapter 9, "Prerequisites for Migration to Microsoft Azure," on page 167. A Migrate server can have multiple Azure Cloud target platforms. You specify the Azure Cloud environment and Location when you create the target platform.
VMware vCloud Director	9.7 9.5 9.1 8.20	See also "Prerequisites for Migration to VMware vCloud Director" on page 201. Download the PlateSpin Replication Environment for vCloud from the Download Site for PlateSpin Migrate 2020.2. See "Understanding PlateSpin Replication Environment Used for Migration of Workloads to vCloud" on page 204.
Oracle Cloud Infrastructure		See also Chapter 11, "Prerequisites for Migration to Oracle Cloud Infrastructure," on page 209.

Table 2-16 Supported Target Cloud Platforms for the Migrate Client

Platform	Remarks
Oracle Cloud Infrastructure	See also Chapter 11, "Prerequisites for Migration to Oracle Cloud Infrastructure," on page 209.

Supported Web Browsers

The PlateSpin Migrate Web Interface, PlateSpin Configuration options, and Help files are available from a supported web browser:

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

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

Supported Data Transfer Methods

Depending on the selected workload and the migration type, PlateSpin Migrate supports file-level or block-level data transfer methods for transferring workload data from the source to the target. Source servers are online during the data transfer.

For information on how to select a transfer method, see "Conversion (Data Transfer Method)" on page 395.

- "File-Level Data Transfer" on page 44
- "Block-Level Data Transfer" on page 45

File-Level Data Transfer

The File-Based Transfer (FBT) method copies data and replicates changes at the file level. File-based transfer might be appropriate for moderately active systems. It also provides the capability to resize your volumes on the target workload.

- "File-Based Transfer for Windows" on page 44
- "File-Based Transfer for Linux" on page 45

File-Based Transfer for Windows

For Windows workloads, PlateSpin Migrate supports file-based transfer.

NOTE: PlateSpin Migrate does not support file-based transfer for Windows Clusters. See "Block-Based Transfer for Windows Clusters" on page 46.

To ensure data consistency, the file-based transfer method leverages the Microsoft Volume Shadow Copy Service (VSS) if it is available. For enterprise applications that are integrated with VSS, VSS captures and copies data in a consistent state while services are running. For applications that are not integrated with VSS, PlateSpin Migrate provides the capability to briefly pause those services while it captures the VSS snapshot to ensure that their data is captured in a consistent state.

If VSS unavailable, PlateSpin Migrate monitors source volumes for file changes while it transfers data. When the initial transfer is complete, migrate re-sends any files that have changed. If the rate of file system changes is consistently high, data transfer is stopped and a job progress warning is shown.

You can configure your migration job to stop high-transaction services, such as Microsoft SQL Server or Microsoft Exchange Server, during the transfer. See "Services or Daemons to Stop before Replication or Cutover" on page 400. Pausing services has two benefits:

- It ensures that the databases of these applications are transferred in a more consistent state.
- It reduces the rate of file system changes so that PlateSpin Migrate is able to keep up with them and complete the transfer.

File-Based Transfer for Linux

For Linux workloads, PlateSpin Migrate supports file-based transfer only in the Migrate Client.

Block-Level Data Transfer

The Block-Based Transfer (BBT) method enables PlateSpin Migrate to transfer data at the block level, providing an exact copy of the source workload. Migrate supports block-based transfer for Windows workloads, Windows Clusters, and Linux workloads.

NOTE: The Block-Based Transfer method is the preferred data transfer method for both Windows and Linux workloads.

- "Block-Based Transfer for Windows" on page 45
- "Block-Based Transfer for Windows Clusters" on page 46
- "Block-Based Transfer for Linux" on page 46

Block-Based Transfer for Windows

For Windows workloads, PlateSpin Migrate supports block-based data transfer.

To ensure data consistency, PlateSpin Migrate leverages the Microsoft Volume Snapshot Service (VSS) with applications and services that support VSS. BBT drivers can be used with standalone and cluster workloads.

NOTE: Before you install block-based transfer drivers on source Windows workloads, ensure that you have applied the latest Windows updates on the workload.

Block-Based Transfer for Windows Clusters

For Windows Clusters, PlateSpin Migrate provides driverless block-based transfer and driver-based block-based transfer methods. These methods work differently than for standalone workloads. For information about the purpose and usage of driverless and driver-based block-based transfer, see "Block-Based Transfer for Clusters" on page 329.

Block-Based Transfer for Linux

For Linux workloads, PlateSpin Migrate supports block-based transfer by using block-based Linux Kernel drivers, called *block watch* (`blkwatch`) drivers.

To ensure data consistency, the `blkwatch` driver leverages LVM snapshots if they are available. Copying data from the snapshot helps avoid potential open file conflicts. See Knowledgebase Article 7005872 *Using LVM Snapshots for Migrating and Protecting Linux Workloads* (https://support.microfocus.com/kb/doc.php?id=7005872). If LVM snapshots are not available, Migrate locks and releases each block in turn for data transfer.

The driver must be built for the specific kernel running on the source Linux workload.

- Migrate provides precompiled `blkwatch` drivers for many non-debug Linux distributions (both 32-bit and 64-bit). See "List of Distributions" on page 346.
- If your Linux workloads have a non-standard, customized, or newer kernel, you can build a custom `blkwatch` driver for your specific kernel. See Knowledgebase Article 7005873 *How to Build a Custom Block-Based Linux Kernel Driver* (https://support.microfocus.com/kb/doc.php?id=7005873).
- Some supported Linux distributions use the drivers for a supported Red Hat Enterprise Linux distribution. "Other Linux Distributions That Use blkwatch Drivers" on page 355.
- Precompiled blkwatch drivers are available for the standard kernel and Unbreakable Enterprise Kernel (UEK) as noted in the "List of Distributions" on page 346. For other Oracle Linux distributions, precompiled drivers are available only for the corresponding Red Hat Compatible Kernel (RHCK).

NOTE: Deployment or removal of the `blkwatch` driver is transparent, has no continuity impact, and requires no intervention and no reboot.

Security and Privacy

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

- “Security of Credentials” on page 49
- “User Authorization and Authentication” on page 50

Security Best Practices

As a security best practice, you should apply patches that address security vulnerabilities to your PlateSpin Server host and PlateSpin Migrate Client host, as you would for other Windows servers in your enterprise.

Micro Focus is aware of the side-channel analysis vulnerabilities described in CVEs 2017-5715, 2017-5753 and 2017-5754, known as Meltdown and Spectre. The current recommended actions have been applied on the PlateSpin Server images in the cloud.

We strongly recommend that you apply security updates that address such threats as recommended by Microsoft for the Windows operating system for the PlateSpin hosts. Consult the vendor documentation for information. See *Protect Your Windows Devices Against Spectre and Meltdown* (https://support.microsoft.com/en-us/help/4073757/protect-your-windows-devices-against-spectre-meltdown) on the Microsoft Support website.

PlateSpin Migrate and Anti-Virus Applications

A PlateSpin Migrate server stores log files and database files in the PlateSpin Migration installation folder. While migration jobs are running, the PlateSpin Migrate server will update these files frequently. Anti-virus applications either block these updates or interrupt them, which impacts the PlateSpin Migrate server performance. Anti-virus applications should either not be installed on the PlateSpin Migrate server host, or the PlateSpin Migrate installation folder must be added to the anti-virus application exclusion list.

Workloads and Anti-Virus Applications

For source workloads, an anti-virus application might block discovery or replication actions if they improperly quarantine PlateSpin Migrate tools, such as the OFX Controller.

If an anti-virus application is running on the source Windows workload, add the following folder and executable file to the anti-virus application’s exclusion list:

- `C:\Windows\PlateSpin`
- `C:\Windows\PlateSpin\Controller\ofxcontroller.exe`

For Windows Workloads, we recommend that all the non-VSS compliant services and anti-virus applications are stopped temporarily on the source while the VSS snapshot is being captured on the source. When you configure a migration, select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. Migrate restarts these services automatically as soon as the VSS snapshot creation completes.

NOTE: Some anti-virus applications will restart itself automatically if an application attempts to stop it. If the anti-virus application restarts before the VSS snapshot creation completes, the data transfer does not begin. If this conflict occurs, you must stop the anti-virus manually, modify the workload configuration to not stop the anti-virus application during replication, then start the replication. After the VSS snapshot is captured for the replication, you can manually restart the anti-virus.

An anti-virus application can also prevent successful replication. See the following information:

- "Configuring Applications Known to Cause Boot Failure on Windows Target Workloads" on page 116
- "Replication Cannot Complete If an Anti-Virus Update Is Pending a Restart on the Source" on page 566

Configuring Source Workloads to Connect Using TLS 1.2

PlateSpin Migrate server supports connections using Transport Layer Security (TLS) 1.0, 1.1, or 1.2 protocol, according to the protocols enabled on its host operating system. PlateSpin Migrate server uses TLS 1.2 protocol by default for connections with source workloads if TLS 1.2 is enabled on the underlying OS and Microsoft .NET Framework on both the Migrate server host and the source workload. Migrate does not have a setting that forces clients to use TLS 1.2 to connect.

NOTE: Older Windows operating systems, such as Windows Server 2008, do not support TLS 1.2. You must enable TLS 1.0 or TLS 1.1 protocols in the Windows Registry settings on the Migrate server host to migrate these source workloads. See "Configuring TLS Protocols for Migrate Hosts" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

To connect a source workload to the Migrate server using TLS 1.2:

- **Source workloads:** Both the Windows operating system and Microsoft .NET Framework version must support TLS 1.2 or must be updated to support TLS 1.2, and the TLS 1.2 protocol must be enabled in the Windows Registry settings.

  For Windows operating systems that do not support TLS 1.2 by default:

  1. A Microsoft update for .NET Framework might be required on the source workload in order to add support for TLS System Default Version settings. A reboot is required.
  2. Use Microsoft Windows Registry settings to force .NET Framework to choose TLS 1.2 when the workload connects with Migrate server.

  For information and configuration instructions, see "Support for TLS 1.2" in *Transport Layer Security (TLS) Best Practices with the .NET Framework* (https://docs.microsoft.com/en-us/dotnet/framework/network-programming/tls) in Microsoft Documentation.
- **Migrate server:** The Windows Registry settings for the TLS 1.2 protocol must be enabled on the Migrate server host. See "Configuring TLS Protocols for Migrate Hosts" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

Security of Workload Data in Transmission

To make the transfer of your workload data more secure, you can configure your migration jobs to encrypt the data in transit to the target. When encryption is enabled, over-the-network data transfer from the source to the target is encrypted by using 128-bit Advanced Encryption Standard (AES).For information about how to enable encryption during data transfer for a migration job, see "Encrypt Data Transfer" on page 396.

You can configure your PlateSpin Server to use a data encryption algorithm that is compliant with FIPS (Federal Information Processing Standards, Publication 140-2). If compliance with FIPS is required, it must be set up on your system prior to the PlateSpin Server installation. See "Enabling Support for FIPS-Compliant Data Encryption Algorithms (Optional)" in your *Installation Guide*.

If FIPS is enabled in a source workload, ensure that the `EnforceFIPSCompliance` parameter is enabled on the PlateSpin Migrate server before you discover the source workload. See "Enforcing FIPS Compliance for FIPS-Enabled Source Workloads" on page 101.

Security of Client-Server Communications

Data transmission between the PlateSpin Server and the PlateSpin Migrate Client can be configured to use either HTTP (default) or HTTPS (Secure Hypertext Transfer Protocol). To secure data transmission between the client and the server, enable SSL on your PlateSpin Server host and use HTTPS when specifying the server URL. See "Connecting to a PlateSpin Migrate Server" on page 67.

Security of Credentials

Credentials that you use to access sources and targets in workload migration jobs are secured by the following measures:

- Each PlateSpin Migrate server has a unique, randomly generated encryption key that it uses to encrypt credentials for source workloads and target platforms.
- Migrate uses the server's encryption key with industry-standard security algorithms to encrypt passwords for source and target credentials, and stores them encrypted in the PlateSpin database.
- Credential passwords can be stored encrypted in exported data by using a user-supplied encryption password with the Import/Export utility.
- The PlateSpin Migrate database is covered by the same security safeguards that you have in place for PlateSpin Server host (or for the PlateSpin database host if you use an external database).

NOTE: To improve security of communications between the Migrate Server host and an external PlateSpin database, you can configure the host operating systems to use the Transport Layer Security (TLS) 1.2 protocol for secure communications. See "Database Server" in "System Requirements for PlateSpin Server" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

- Passwords might be included in diagnostics, which are accessible to accredited users. You should ensure workload migration projects are handled by authorized staff.
- PlateSpin Migrate Client can store credentials locally on the Migrate Client host. The passwords are cached, encrypted, and securely stored by the PlateSpin Migrate Client, by using operating system APIs.

User Authorization and Authentication

PlateSpin Migrate provides a role-based user authorization and authentication mechanism. See "Configuring User Authorization and Authentication" on page 89.

Performance

Performance for migrations using PlateSpin Migrate depends on many factors. Use the guidelines in this section to understand those factors and better plan your migration projects.

Performance Characteristics

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

- Hardware and software profiles of your source and target
- Hardware and software profiles of your PlateSpin Server host
- Hardware and software profiles of your target virtualization host or cloud host environment as VMs compete for resources
- The specifics of your network bandwidth, configuration, and conditions
- The number of your source workloads' volumes and their sizes
- File density (number of files per unit of capacity) on your source workloads' volumes
- Source I/O levels (how busy your workloads are)
- The number of concurrent migrations and the number and type of the targets
- Whether data encryption is enabled or disabled
- Whether data compression is enabled or disabled

For planning large-scale workload migrations, you should perform a test migration of an average workload and use the result as a benchmark, fine-tuning your metrics regularly throughout the project. In addition to the data transfer process, also consider the other phases that a migration job goes through, as applicable to your project:

- Preparation and network setup

- Source workload and target machine discovery
- Target configuration

Scalability

You can set up multiple workload migrations and run them concurrently. See "Performance Characteristics" for information about the many factors that impact performance of PlateSpin Migrate in your migration environment.

- "Concurrent Replications and Migrations" on page 51
- "Workload Discovery and Inventory" on page 51

Concurrent Replications and Migrations

Performance for concurrent replications and concurrent migrations depends on the resources on the PlateSpin Migrate server and the target environment, as well as the available bandwidth. We recommend that you begin with a low load, then increase it and see how the migrations perform in your environment. Use the scheduled start dates to control when migrations begin and how many migration jobs are scheduled to run concurrently.

The available hardware resources on your Migrate server impact the number of managed workloads and concurrent replications the server can handle. Generally, the higher the load is for concurrent replication and migration, the more resources it consumes.

Scalability testing performed with VMware ESX hosts suggests the following benchmark recommendations:

- Multiple migrations to a single VMware ESX Host server: no more than 10
- Multiple migrations against multiple VMware ESX Host servers: no more than 40

In a VMware Cluster, ensure that you balance migrations across multiple hosts in the cluster for best performance.

Workload Discovery and Inventory

We recommend that you keep no more than 50 discovered workloads at a time in the inventory for your PlateSpin Migrate server, depending on its available hardware resources. As you complete workload migrations, you can remove workloads and add others.

You cannot necessarily concurrently run replications and migrations for all the workloads in your inventory. Use the scheduled start dates to control when migrations begin and how many migration jobs are scheduled to run concurrently. See "Concurrent Replications and Migrations".

PlateSpin Migrate provides three discovery tools:

- **Migrate Web Interface:** Discover one workload at a time.
- **Migrate Client:** Discover one workload at a time, multiple workloads at a time, or all workloads in a domain.
- **Mass Discover CLI:** Discover one or multiple workloads from a CSV file.

For more information, see "About Source Workload Discovery" on page 297.

Data Compression

If necessary, PlateSpin Migrate can compress the workload data before transferring it over the network. This enables you to reduce the overall amount of data transferred during a workload migration job.

Compression ratios depend on the type of files on a source workload's volumes, and might vary from approximately 0.9 (100MB of data compressed to 90 MB) to approximately 0.5 (100MB compressed to 50MB).

NOTE: Data compression utilizes the source workload's processor power.

Data Compression can be configured per migration job. You can also use the PlateSpin Migrate Client to specify a default compression value to be applied globally. See "Configuring Job Values Defaults" on page 124.

Bandwidth Throttling

PlateSpin Migrate enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. You can specify a throughput rate for each migration job. You can specify whether to throttle at all times or on specific days of the week and times of day. This provides a way to prevent migration traffic from congesting your production network and reduces the overall load of your PlateSpin Server.

Bandwidth throttling is a parameter of a workload migration job's configuration properties. To apply bandwidth throttling for the migration job, see "Bandwidth Throttling during Data Transfer" on page 394.

Blackout Window

PlateSpin Migrate Web Interface enables you to specify a blackout window for replication. The blackout window suspends scheduled replications from starting during a specified period of time and pattern. It helps you to reserve network bandwidth for users or mission critical communications during peak traffic periods. You can also use it to prevent conflicts for other data backup or snapshot activities.

Database Server

PlateSpin Migrate includes Microsoft SQL Server Express Edition. The capabilities of SQL Server Express are sufficient for the scalability characteristics described in "Scalability" on page 51.

NOTE: Microsoft SQL Server Express has a database size limit of 10 GB and can use only one CPU core at a time and 1 GB memory. For more information about requirements and limitations for SQL Server Express, see the Microsoft SQL Server 2017 Express documentation (https://www.microsoft.com/en-us/download/details.aspx?id=55994).

For large scale migrations where you want to preserve migration reports for longer time, it is recommended to use enterprise version or keep archiving data to make room for new reporting data.

We recommend that you configure the PlateSpin Server to use a database instance on your existing Microsoft SQL Server Standard Edition or Enterprise Edition database server in the following environments:

- Deployments of multiple PlateSpin Servers that use the same remote Microsoft SQL Server database server for their database instances
- Deployments where keeping all history of the reporting data is important

While multiple PlateSpin Migrate servers can use the same remote database server, each Migrate server requires a separate database instance.

Access and Communication Requirements across Your Migration Network

Ensure that your network environment meets the following requirements for access, discovery, and migration.

NOTE: Refer to the deployment diagrams based on your migration target to understand the ports and flow of information between the various migration components. See Part III, “Preparing Your Migration Environment,” on page 143.

Requirements for Discovery

Table 2-17 lists software, network, and firewall requirements that systems in your environment must meet for the discovery and inventory process. For information about discovery procedures, see Part IV, “Discovering and Preparing Workloads and Targets,” on page 277.

Table 2-17 *Network Communication Prerequisites for Discovery Operations*

System	Prerequisites
All workloads	Ping (ICMP echo request and response) support
All source workloads in AWS	◆ PowerShell 2.0 or higher

System	Prerequisites
All Windows sources	◆ Microsoft .NET Framework version 2.0 SP2, 3.5 SP1, or 4.0 ◆ Requires credentials for built-in Administrator, domain Administrator account with access to Admin$ share, or local Administrator account with administrator privileges. **NOTE:** To allow a local Administrator account that has administrator privileges to discover Windows workloads, you must enable remote access permissions for this account. ◆ **For local Administrator account that is a member of Windows domain account:** Do one of the following: ◆ Create a Group Policy to enable remote access for the account. ◆ Ensure that the `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System` registry setting has a registry key named `LocalAccountTokenFilterPolicy` of type 'DWORD' and value '1'. ◆ **For local Administrator account that is not a member of Windows domain account:** Ensure that the `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System` registry setting has a registry key named `LocalAccountTokenFilterPolicy` of type 'DWORD' and value '1'. For more information about creating a Group Policy or using the registry settings to enforce remote access for the account, see "*Enforce local account restrictions for remote access*" in Microsoft Documentation (https://docs.microsoft.com/en-us/windows/security/identity-protection/access-control/local-accounts#enforce-local-account-restrictions-for-remote-access). ◆ The Windows Firewall configured to allow **File and Printer Sharing**. Use one of these options: ◆ **Option 1, using Windows Firewall:** Use the basic **Windows Firewall** Control Panel item (`firewall.cpl`) and select **File and printer Sharing** in the list of exceptions. - OR - ◆ **Option 2, using Windows Firewall with Advanced Security:** Use the **Windows Firewall with Advanced Security** utility (`wf.msc`) with the following **Inbound Rules** enabled and set to `Allow`: ◆ **File and Printer Sharing (Echo Request - ICMPv4In)** ◆ **File and Printer Sharing (Echo Request - ICMPv6In)** ◆ **File and Printer Sharing (NB-Datagram-In)** ◆ **File and Printer Sharing (NB-Name-In)** ◆ **File and Printer Sharing (NB-Session-In)** ◆ **File and Printer Sharing (SMB-In)** ◆ **File and Printer Sharing (Spooler Service - RPC)** ◆ **File and Printer Sharing (Spooler Service - RPC-EPMAP)** ◆ The Windows Firewall configured to allow Windows Management Instrumentation (WMI-In). ◆ (Conditional) If the volumes are encrypted with the BitLocker disk encryption feature, they must be unlocked.

System	Prerequisites
All Linux sources Linux KVM servers	◆ Secure Shell (SSH) server ◆ Open port 22 (TCP) ◆ Custom SSH ports are supported; specify the port number during discovery: <*hostname* | *IP_address*>:<*port_number*>. ◆ Root-level access. For information on using an account other than `root`, see KB Article 7920711 (https://support.microfocus.com/kb/doc.php?id=7920711). **NOTE:** For source Linux workloads in Amazon Web Services, AMI templates automatically create a default non-`root` system user account that is enabled for `sudo`. The user name for this account varies by AMI provider. For Amazon Linux images, the non-`root` user name is `ec2-user` for most Linux distributions. It is `centos` for CentOS AMIs. For more information, refer to your AMI provider documentation. In AWS, a non-`root` user must run the `sudo -i` command to access the `root` shell and then run the Migrate Agent commands. Typing `sudo` in each Migrate Agent Utility command might result in a failure on some source workloads.
VMware ESX/ESXi Servers	◆ VMware account with an Administrator role ◆ VMware Web services API and file management API (HTTPS / port 443 TCP)
VMware vCenter Servers	The user with access must be assigned the appropriate roles and permissions. Refer to the pertinent release of VMware documentation for more information.
Cloud-based targets: ◆ Amazon Web Services ◆ Microsoft Azure ◆ VMware vCloud ◆ Oracle Cloud Infrastructure	Open port 443 (TCP) for HTTPS communications with the target management portal.

Requirements for Workload Registration

You can use Migrate Agent to register and inventory workloads instead of using Migrate discovery. Table 2-18 lists software, network, and firewall requirements that systems in your environment must meet for the registration and inventory process using Migrate Agent. For information about registration procedures, see "Registering Workloads and Discovering Details with Migrate Agent" on page 303. See also Appendix F, "Migrate Agent Utility," on page 359.

Table 2-18 *Network Communication Prerequisites for Migrate Agent Registration Operations*

System	Prerequisites
PlateSpin Server hosts	◆ Open port 443 (TCP) for HTTPS communications with source workloads. ◆ Open port 22 (TCP) for SSH communications with Linux source workloads. ◆ A public IP address is required for PlateSpin Server host. ◆ In PlateSpin Configuration, set the **AlternateServerAddress** parameter to the Migrate server's public IP address. The setting is configured automatically for Migrate servers available in Azure Marketplace and for Migrate servers deployed using the AWS Quick Start.
All source workloads	◆ Open port 443 (TCP) for HTTPS communications with Migrate server. ◆ A public IP address is required for source workloads.
All Windows source workloads	◆ The user who executes Migrate Agent commands must have `Administrator` privileges. ◆ For remote connections to the source workload, open port 3389 (TCP) for RDP access to the machine to install Migrate Agent.
All Linux source workloads	◆ Root-level access. For information on using an account other than `root`, see KB Article 7920711 (https://support.microfocus.com/kb/doc.php?id=7920711). **NOTE:** For source Linux workloads in Amazon Web Services, AMI templates automatically create a default non-`root` system user account that is enabled for `sudo`. The user name for this account varies by AMI provider. For Amazon Linux images, the non-`root` user name is `ec2-user` for most Linux distributions. It is `centos` for CentOS AMIs. For more information, refer to your AMI provider documentation. In AWS, a non-`root` user must run the `sudo -i` command to access the `root` shell and then run the Migrate Agent commands. Typing `sudo` in each Migrate Agent Utility command might result in a failure on some source workloads. ◆ For remote connections to the source Linux workload: ◆ Secure Shell (SSH) server ◆ Open port 22 (TCP) ◆ Custom SSH ports are supported; specify the port number during discovery: <*hostname* | *IP_address*>:<*port_number*>.

Requirements for Migration

Table 2-19 lists software and firewall requirements that systems in your environment must meet for problem-free operation during workload migration jobs.

Table 2-19 *Network Communication Prerequisites for Workload Migration*

System	Open Port (Default)	Remarks
PlateSpin Server hosts	Either TCP 80 or TCP 443	◆ Port 80 (TCP) is required for HTTP communication among the PlateSpin Server, sources, and targets. ◆ Port 443 (TCP) is required for HTTPS communication (if SSL is used) between the PlateSpin Server and the source or target machines.
All source workloads	TCP 3725	The source workload connects to the target workload on port 3725 (TCP) by default. This default connection can be reversed (target workload connects to source workload) by changing the **SourceListensForConnection** parameter setting from `False` to `True`, in which case the target workload will connect to the source workload on port 3725 (TCP). See "Configuring the Contact Direction for the Replication Port" on page 114. For Server Sync jobs, this port is required for both sources and targets. The port number is configurable by setting the **FileTransferPort** parameter in the PlateSpin Configuration settings for the Migrate server.
All targets	TCP 3725	Required for source workload to initiate communication with the target workload and for Server Sync jobs.
All Windows sources and targets	NetBIOS 137 - 139	Required for NetBIOS communications.
All Windows Server Cluster workloads. See "Clusters" on page 25.		Ensure that the PlateSpin Server can resolve DNS forward lookup and reverse lookup for the IP addresses of the Windows Server Cluster and its cluster nodes. You can update the DNS server or update the local `hosts` file (`%systemroot%\system32\drivers\etc\hosts`) on the PlateSpin Server.
All Windows sources	SMB (TCP 139, 445 and UDP 137, 138)	
All Linux sources Linux KVM servers	TCP 22	

System	Open Port (Default)	Remarks
PlateSpin Server hosts; All Windows sources	TCP 135/445	For DCOM/RPC communication between PlateSpin Server and a source for taking control of and rebooting the workload through WMI. **NOTE:** WMI (RPC/DCOM) can use TCP ports 135 and 445 as well as random/dynamically assigned ports above 1024.
PlateSpin Server hosts Windows Cluster source and target workloads	TCP 5986, outbound for host; inbound for workloads	Required for HTTPS transport for PowerShell remoting commands to shut down the non-active nodes of a Windows Cluster as appropriate for migration of a Windows Cluster to VMware.
AWS requires a minimum of Microsoft .NET Framework 4.5 for migration of Windows workloads to Nitro System based VM instances		To use Nitro System based instances, ensure that Microsoft .NET Framework 4.5 or later is installed on the source Windows workload before Migrate discovery.

Requirements for Migration of Workloads Registered Using Migrate Agent

Table 2-20 lists firewall, network, and software requirements that systems in your environment must meet for problem-free operation during migration of workloads that have been registered with the PlateSpin Server host using Migrate Agent. See also “Requirements for Migrate Agent Utility” on page 359.

Table 2-20 *Network Communication Prerequisites for Migration of Workloads Registered Using Migrate Agent*

System	Open Port (Default)	Remarks
PlateSpin Server hosts	TCP 443	Required for HTTPS communications with source and target workloads. A public IP address is required for PlateSpin Server host.
	TCP 22	Required for SSH communications with Linux workloads.

System	Open Port (Default)	Remarks
PlateSpin Configuration settings		Configuration requirements in PlateSpin Configuration for the Migrate server: ◆ Set the **AlternateServerAddress** parameter to the Migrate server's public IP address. The setting is configured automatically for Migrate servers available in Azure Marketplace and for Migrate servers deployed using the AWS Quick Start. See "Configuring Alternate IP Addresses for PlateSpin Server" on page 111. ◆ Ensure that the **SourceListensForConnection** parameter is set to `False`, which is the default value. ◆ For cloud-based Migrate servers, the server is configured by default for migration to the target type that matches its parent cloud environment. If the source workloads are in the parent cloud environment for migration to a different target, you must remove the default value (leave the field blank) for the **ServerIsHostedInCloud** parameter to allow all target types to be available in the Add Target dialog.
PlateSpin replication network		When you configure the workload migration, ensure that you enable a public IP address for the PlateSpin replication network.
All source and target workloads	TCP 443	Required for HTTPS communications with PlateSpin server.
	TCP 3725	Required for Migrate communications between the source and target machines and for data transfer from the source machine to the target machine. The port number is configurable by setting the **FileTransferPort** parameter in the PlateSpin Configuration settings for the Migrate server. When you use the Migrate Agent on the source workload, the source workload contacts the target workload for data transfers. Ensure that the **SourceListensForConnection** parameter is set to False, which is the default value.

System	Open Port (Default)	Remarks
All Linux target workloads	TCP 22	Required for SSH communications from the PlateSpin server in the PlateSpin Replication Environment.
All target workloads		Public IP addresses are required for target machines to enable source workloads to contact them over port 3725 to begin replications. Migrate sets public IP addresses on target machines during migration.

Migrations Across Public and Private Networks through NAT

In some cases, a source, a target, or PlateSpin Migrate itself, might be located in an internal (private) network behind a network address translator (NAT) device, unable to communicate with its counterpart during migration.

PlateSpin Migrate enables you to address this issue, depending on which of the following hosts is located behind the NAT device:

- **PlateSpin Server:** In your server's *PlateSpin Server Configuration* tool, record the additional IP addresses assigned to that host:
  1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

     `https://Your_PlateSpin_Server/PlateSpinConfiguration/`
  2. Locate the **AlternateServerAddresses** server parameter, click **Edit**, then add additional IP addresses, delimited by a semicolon (;), for example:

     `10.50.186.147;10.50.186.148`
- **Source:** As part of that specific migration job, record the additional IP addresses assigned to that workload. See "Network Identification (Network Connections)" on page 411.
- **Target:** When you are attempting to discover a target, such as VMware ESX, specify the public (or external) IP address in the discovery parameters.

Deciding on the Migration Interface

PlateSpin Migrate includes the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface to allow you to efficiently plan, configure, execute, and test migrations. The PlateSpin Migrate Web Interface supports large scale migration of workloads to VMware and cloud platforms such as Microsoft Azure, Amazon Web Services, Oracle Cloud Infrastructure, and VMware vCloud Director. The PlateSpin Migrate Client supports migration of workloads to physical machines, VMware, and KVM platforms.

Use the PlateSpin Migrate Web Interface when you want to concurrently migrate a large number of workloads.

The decision to use a particular migration interface depends on the migration operations or the migration tasks you have to perform.

For example:

- Migration to Amazon Web Services, Microsoft Azure, and VMware vCloud Director is possible only from the PlateSpin Migrate Web Interface.
- Migration to VMware is possible from both the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface.

For a list of migration operations that you can perform using the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface, see "Migration Operations Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface" on page 82.

For a list of migration tasks that you can perform using the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface, see "Migration Tasks Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface" on page 84.

IMPORTANT: Do not use the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface interchangeably to perform the migration tasks throughout the migration cycle of a workload. Select the appropriate tool for the workload, and use it consistently for that migration effort.

A Frequently Asked Questions

This section provides answers to frequently asked questions.

What are the performance and scalability characteristics of my PlateSpin Migrate product?

Your PlateSpin Migrate product's overall performance, including data transfer speeds and scalability, depend on a variety of factors in your specific environment. See "Performance" on page 50.

How secure is my PlateSpin Migrate product?

PlateSpin Migrate provides several features to help you safeguard your data and increase security. See "Security and Privacy" on page 46.

Does PlateSpin Migrate support my workload's data storage technology?

PlateSpin Migrate products support a number of data storage and management technologies, including Windows dynamic disks, Linux logical volumes, RAID (Redundant Array of Independent Disks) systems, and SAN (Storage Area Network) systems.

Can I use custom SSH ports to communicate with my workloads?

Yes. See "Target Discovery in the Migrate Client" on page 283.

Can multiple migrations run simultaneously?

Yes. See "Performance" on page 50.

II Working With Your PlateSpin Migrate Server

This section provides information on typical, usually one-time configuration tasks following product installation. For installation information, see the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

3 Using the PlateSpin Migrate Tools

This section introduces the PlateSpin Migrate tools and how you use them to carry out workload migration and management tasks. To interact with the product and perform tasks such as discovery of source workloads and target hosts; setting up, executing, and monitoring jobs; managing license keys; and configuring the default behavior of the server, use either the PlateSpin Migrate Client or the browser-based PlateSpin Migrate Web Interface. To decide which interface to use, see "Deciding on the Migration Interface" on page 60.

IMPORTANT: To migrate a workload, you should either use the PlateSpin Migrate Client or the PlateSpin Migrate Web Interface throughout the migration cycle of the workload.

Connecting to a PlateSpin Migrate Server

PlateSpin Server Access Using the Migrate Client

Every time you start the PlateSpin Migrate Client, it performs the following actions:

- Performs authentication of the specified user account with the PlateSpin Server.

  See "Configuring User Authorization and Authentication" on page 89.
- Connects to a specified PlateSpin Server.
- Loads a specified PlateSpin Migrate Network, a collection of discovered source workloads and targets that you work with at one time.

You specify your connection credentials, the PlateSpin Server instance, and the required PlateSpin Migrate Network in the PlateSpin Server settings.

1. In the PlateSpin Migrate Client, click **Tools** > **PlateSpin Server Settings**.

   or

Double-click one of the following three areas in PlateSpin Migrate Client status bar at the bottom: **Server**, **Network**, or **User**.

The PlateSpin Server Settings dialog box opens.

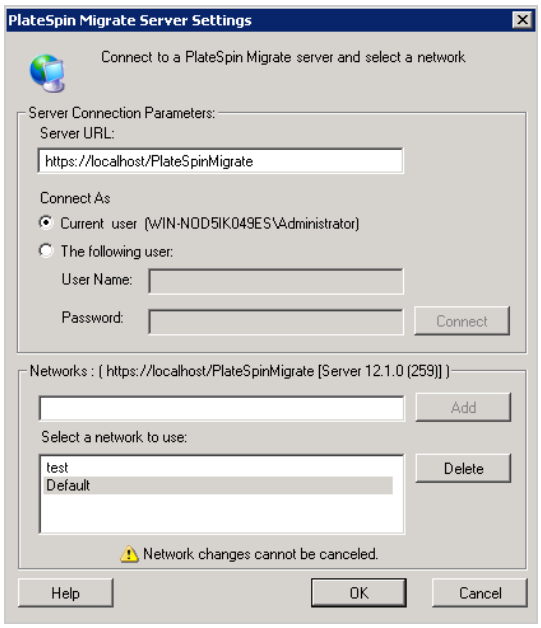

2 Specify the required PlateSpin Server URL, user, and network parameters as required:

Interface Element	Description
Server URL	Type the PlateSpin Server URL in the following format: `http://<server_host>/platespinmigrate` If SSL is enabled on the PlateSpin Server host, replace `http` in the URL with `https`. We recommend that you specify the fully qualified domain name (FQDN) if you are using a domain user account to log into Migrate Server.
Connect As	To connect to a PlateSpin Server, you must have administrative access to the PlateSpin Server host or be a member of one of the PlateSpin Migrate roles. See "Configuring User Authorization and Authentication" on page 89.
Networks	To familiarize yourself with PlateSpin Migrate features, use the **Sample Environment** network. To work with actual source workloads and targets, use the **Default** network or create your own. To add a network, type the name, then click **Add**. To remove a network, select it, then click **Delete**.

3 When you have finished, click **OK**.

PlateSpin Server Access Using the Migrate Web Interface

To access the PlateSpin Migrate Web Interface, use one of the following web browsers:

- **Google Chrome:** Version 34.0 and later
- **Microsoft Internet Explorer:** Version 11.0 and later
- **Mozilla Firefox:** Version 29.0 and later

NOTE: You must ensure that JavaScript (Active Scripting) is enabled in the browser.

To launch PlateSpin Migrate Web Interface:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`

   Replace *Your_PlateSpin_Server* with the DNS host name or IP address of your PlateSpin Migrate Server.

2. Log in using the local Administrator user credentials for the PlateSpin Server host or as an authorized user.

   For information about setting up additional users for PlateSpin, see "Configuring User Authorization and Authentication" on page 89.

About the PlateSpin Migrate Client User Interface

The PlateSpin Migrate Client provides a management tool to manage migrations to a variety of virtual host targets, physical targets, and server-sync.

For information about installing the Migrate Client, see "System Requirements for PlateSpin Migrate Client" and "Installing the PlateSpin Migrate Client" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

For information about configuration options for the Migrate Client, see Chapter 6, "Configuring PlateSpin Migrate Client," on page 123.

Use the information in this section to familiarize yourself with the Migrate Client.

Navigating the Client Interface

The PlateSpin Migrate Client window consists of the following elements:

- **Menu bar:** Reflects the current view and provides command groups for accessing program features and operations.
- **Toolbar:** Reflects the current view and provides visual shortcuts to program features and operations.
- **Servers View:** The Servers view is the main visual interface to your discovered source workloads and targets. See "Servers View" on page 70.
- **Jobs View:** The Jobs view displays all jobs, such as discovery and migration. See "Jobs View" on page 73.
- **Current view:** The main work area of the interface; lists either machines (when in Servers view mode), or jobs (when in Jobs view mode).
- **Panes:** Vertically aligned at the left side of the window, panes facilitate the selection of the current view (View pane) or a migration job (Tasks pane). A Details pane reflects the current view and provides summary information about an item selected in the current view.
- **Tasks Pane:** The Tasks pane of the PlateSpin Migrate Client window contains most essential migration actions. Clicking a task opens the Action window, which you can use to select the migration source, target, and setup method.
- **Status bar:** At the bottom of the PlateSpin Migrate Client window, the status bar displays the PlateSpin Server that the client is currently connected to, the PlateSpin Migrate Network you are currently working with, the name and role of the current user logged in, and the status of the Automatic Network Discovery feature. See "Status Bar" on page 73.

Servers View

The Servers view is the main visual interface to your discovered source workloads and targets. It consists of two panes that you can customize to suit your needs.The hierarchical display of items in the Servers view reflects the organization of items on their respective platforms; for example: VMs are shown nested beneath their VM hosts.

In addition, the **Group By** bar enables you to group machines by affiliation to a domain or to a vCenter Server (for VMware ESX server systems). See "Organizing the Servers View" on page 72.

NOTE: The Servers view hierarchy does not reflect advanced VM resource management hierarchies and structures, such as membership in resource pools or affiliation with ESX Distributed Resource Scheduler (DRS) clusters. You can view such information in an item's properties. See "Viewing the Properties of Source Workloads and Targets" on page 72.

Distinguishing Target Machines for Semi-Automated (X2P) Workflow

When you use the semi-automated (X2P) workflow, the host name displayed for the target workload in the Servers view is the registration name you provided during discovery with PlateSpin Boot OFX ISO. Additional information helps to distinguish it from the source workload:

- **If no OS is present:** The **Host Name** column displays only the registered host name. The **Operating System** column displays information from the LRD, with the annotation **Under Control.**
- **If an OS is present:** The **Host Name** column displays the registered host name followed by the host name of its operating system. The **Operating System** column displays the operating system information, with the annotation **Under Control**.

Figure 3-1 provides an example of X2P host names for target workloads with and without an operating system present. Workloads X2P-HV-LX-VM3 and X2P-HV-WIN-VM1 do not have an underlying operating system. The LRD information is displayed as the operating system.

Figure 3-1 *X2P Host Name and Operating System Displayed in the Hosts List*

In the Properties dialog for the target workload, the displayed host name is the operating system host name. The registered host name displays at the bottom of the General tab as the ***Hostname** value. The OS value displays the **Under Control** annotation.

Figure 3-2 *Properties Dialog for an X2P Target Workload*

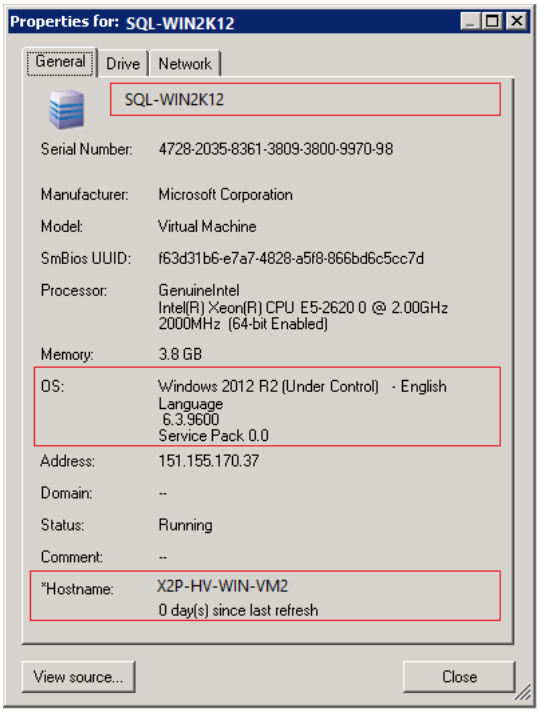

Organizing the Servers View

You can filter source workloads and targets based on operating system, domain, name, and type by using the **Group By** and **Show** drop-down menus. You can use the **Group By** drop-down menu to group the items in the Servers view by:

- Domain affiliation
- Hostname
- Affiliation to a VMware vCenter Server

To further control the scope of items shown in either pane of the view, you can also use the **Show** drop-down menu to filter machines by workload type; for example, Windows Server 2008 R2, Red Hat Linux, and so on:

Viewing the Properties of Source Workloads and Targets

In the Servers view, you can access the essential properties of your discovered source workloads and targets by right-clicking an item and selecting **Properties**.

For each machine, the system provides information about the selected system's:

- Hardware, operating system, and network profile
- Volumes, partitions, and disk usage
- Programs and services

A virtual machine's properties provide information related to the machine's environment on its corresponding virtualization platform, including information about the host, and the amount of allocated memory and processing power.

The properties for virtual machine hosts provide information specific to the selected system. For example, you can view what virtual machines are running on a selected VMware ESX server, what virtual network adapters are in use, and what resource pools are configured on them.

VMware ESX servers that are assigned to a Distributed Resource Scheduler (DRS) cluster provide information about the name of the cluster and the DRS automation level (full, manual, or partially automated). The properties for VMware ESX servers that are part of VMware vCenter platforms also indicate this.

List of Machine-Specific Icons in the Servers View

Discovered source workloads and targets are associated with unique icons to help identify the type of workload or workload host.

Table 3-1 *Machine-Specific Icons in the Servers View*

	Physical machine
	Physical machine with workload license
	Virtual machine server
	Virtual machine
	Virtual machine with workload license
	Undiscovered virtual machine
	Virtual machine - Server Sync target
	Virtual machine - Server Sync target with workload license

Jobs View

The Jobs view displays all jobs, such as discovery and migration, organized into two tabs:

- **Jobs:** All jobs submitted for execution.
- **Saved Jobs:** All saved jobs not yet submitted for execution. See “Using the Migrate Client” on page 539.

You can limit the scope of jobs displayed in the view. Use the **Job Type** and **Jobs Status** menus to specify filters for the view:

- **Job Type:** To view discovery, migration, or all other job types.
- **Job Status:** To view failed, currently running, and completed jobs.

Tasks Pane

The Tasks pane of the PlateSpin Migrate Client window contains most essential migration actions. Clicking a task opens the Action window, which you can use to select the migration source, target, and setup method.

Status Bar

The status bar of the PlateSpin Migrate Client window displays information about:

- The PlateSpin Server that you are currently connected to.
- The PlateSpin Migrate Network that you are currently working with.

- The User that you are logged in as, and the PlateSpin Migrate role assigned to your user account.
- The status of the Automatic Network Discovery feature.

Figure 3-3 *Status Bar of the PlateSpin Migrate Client Window*

Double-clicking any of the first three status items opens the PlateSpin Server Settings window. See "Connecting to a PlateSpin Migrate Server" on page 67.

Double-clicking the Network Discovery status item turns Automatic Windows Network Discovery on or off. See "Discovering Target VMs for Server Sync Jobs" on page 293.

Workload Migration Tasks

PlateSpin Migrate Client enables you to define, save, schedule, execute, and monitor the following migration tasks.

Task	Description
Copy Workload	Results in a virtual or physical duplicate of a selected physical or virtual workload, except that the new workload is assigned a new network identity. Use this migration task when you intend to keep the source workload operational.
Move Workload	Results in an exact virtual or physical duplicate of a selected physical or virtual workload. Use this migration task when you intend to retire or repurpose the original infrastructure.
Server Sync	Synchronizes a virtual or physical workload with another virtual or physical workload without transferring the entire source volume data over the network.

About the PlateSpin Migrate Web Interface

The PlateSpin Migrate Web Interface provides a web-browser-based management tool to manage automated migrations to target virtual machines on VMware host targets and cloud-based targets. No client installation is required. For information about configuration options for the Web Interface, see Chapter 7, "Configuring PlateSpin Migrate Web Interface," on page 131.

The Web Interface offers the highest levels of automation, with scheduled incremental replications, block change tracking, one-time configuration, and one-click pre-cutover testing and workload cutover.

Use the information in this section to familiarize yourself with the Migrate Web Interface.

- "Dashboard" on page 82
- "Reports" on page 82

Navigating the Web Interface

The Web Interface displays a navigation bar with the following options:

Table 3-2 *Navigation Options in the PlateSpin Migrate Web Interface*

Navigation Options	Description
Dashboard	Displays the default Dashboard page that provides information about the Migrate licenses, latest tasks, running events, upcoming events, and past events. See "Dashboard" on page 82.
Workloads	Displays the Workloads page that lists all the discovered workloads. To add or discover a workload, click **Add Workload** option on the Dashboard page or the Workloads page. For more information about adding or discovering a workload, see "Workload Discovery in the Migrate Web Interface" on page 302. You can perform various other tasks such as configuring a workload, preparing a workload for migration, and migrating a workload. See "Workloads" on page 76.
Targets	Displays the Targets page that lists the already added target platforms and lets you add new targets. For more information about adding or discovering a workload, see "Target Discovery in the Web Interface" on page 284. See "Targets" on page 81.
Tasks	Displays the Tasks page that lists items requiring user intervention. See "Tasks" on page 81.
Reports	Displays the Reports page. See "Generating Workload and Workload Migration Reports" on page 545.
Settings	Displays the Settings page that lets you configure the following: • **Licensing:** See "License Activation Using the Web Interface" on page 96 and "Viewing Workload License Designations Using Migrate Web Interface" on page 101. • **Permissions:** See "Managing Security Groups and Workload Permissions" on page 131. • **General Notification Settings:** See "Setting Up Event Notifications by Email" on page 103. • **Report Notification Settings:** See "Setting Up Replication Report Notifications by Email" on page 105. • **SMTP:** See "Setting up the SMTP Server" on page 103. • **Advanced Server Settings:** See "PlateSpin Configuration" on page 86. • **Workload Tags:** See "Using Tags to Track Logical Associations of Workloads" on page 309.

Navigation Options	Description
Downloads	Displays a page that lets you download the following: • **Migrate Agent:** Allows you to download and install the Migrate Agent utility for Windows or Linux. For information about working with the Migrate Agent Utility, see Appendix F, "Migrate Agent Utility," on page 359. • **Migrate Client Setup:** Allows you to download and install the PlateSpin Migrate Client. For information about the PlateSpin Migrate Client, see "About the PlateSpin Migrate Client User Interface" on page 69. You can also install the PlateSpin Migrate Client using the PlateSpin Migrate Installer. For more information, see Installing the PlateSpin Migrate Client in the PlateSpin Migrate 2020.2 Installation and Upgrade Guide.
About	Displays information such as the product version, copyright information, license information and also provides links to the Downloads page and the product home page.
Help	Displays the online documentation page.

Workloads

The Workloads page displays information about source Windows and Linux workloads. You can also add (discover) a new workload, remove (undiscover) a workload migration that is managed in the Web Interface, and configure migration jobs for discovered workloads.

Status for Workloads Managed in Migrate Web Interface

The Workloads page displays information for each workload you manage in the Migrate Web Interface.

Figure 3-4 *Migrate Web Interface Workloads Page*

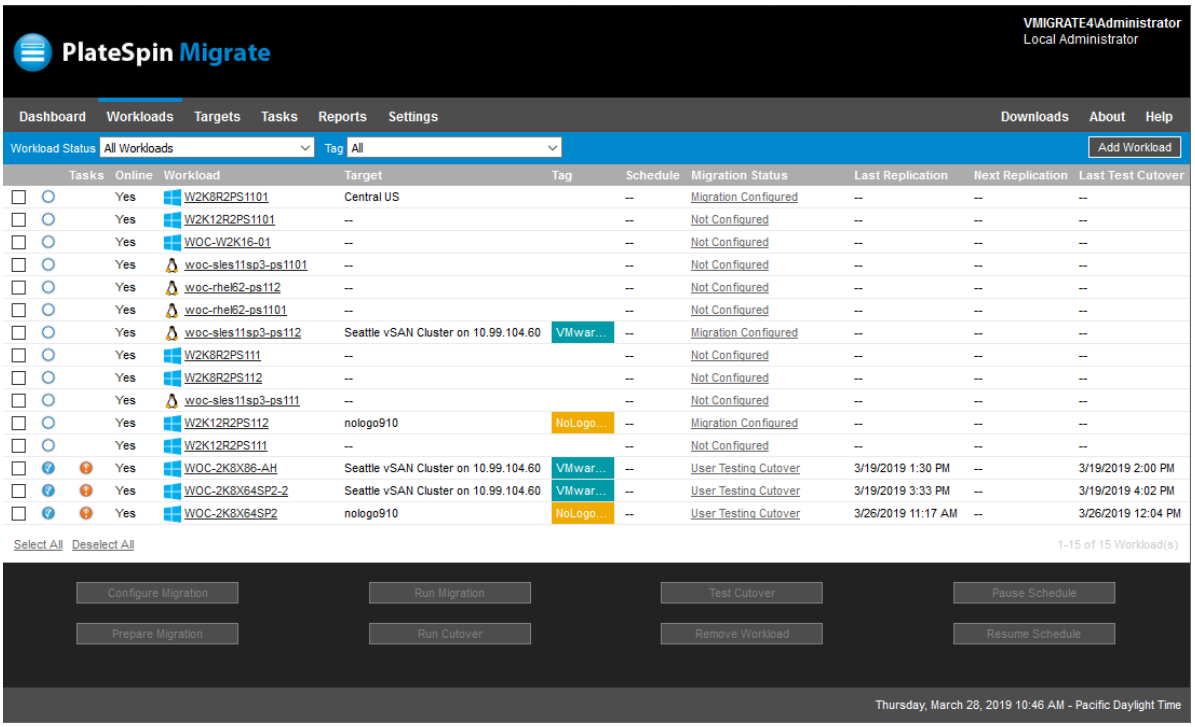

Item	Description
Tasks	Displays a warning icon for a task that might require user attention. For example: if a workload goes offline, then a warning icon displays. Pause over the icon to see more details.
Online	Displays one of the following: ◆ **Yes:** If the workload is online. ◆ **No:** If the workload is offline.
Workload	Displays the workload name. Click the workload name to configure the workload for migration.
Target	Displays the target platform name if the workload is configured for migration, such as an Azure location, AWS region, or VMware cluster or host.
Tag	Displays the tag associated with the workload. For more information about the tags, see “Managing Workload Tags” on page 133 and “Using Tags to Track Logical Associations of Workloads” on page 309.

Item	Description
Schedule	Displays the status of the schedule if you have configured a schedule for the workload migration. For example: if the schedule is configured, it displays **Active** after you have prepared the workload for migration until the end of the migration cycle, unless you pause the schedule. If you click **Pause Schedule**, then the **Paused** status displays. If you click **Resume Schedule**, then **Active** displays again.
Migration Status	Displays the current status of the workload. For example: ◆ **Adding Workload:** The process of adding or discovering a workload is in progress. ◆ **Not Configured:** The workload has been discovered but is not yet configured. ◆ **Migration Configured:** The workload has been configured for migration. ◆ **Preparing Migration:** The source workload is being prepared for migration and the target workload is being prepared for replication. ◆ **Running First Replication:** The workload is being replicated for the first time. Click the **Migration Status** link to view information about the related events.
Last Replication	Displays the date when the workload was last replicated.
Next Replication	Displays the date when the workload is scheduled for the next replication.
Last Test Cutover	Displays the date when the target workload was last tested.

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

Status for Workloads Managed in Migrate Client

The Workloads page displays read-only status for migration jobs managed in the Migrate Client.

After you discover details for a workload in the Migrate Client, the Web Interface displays it in the Workloads list with a status of Unconfigured. At this point, you can proceed to manage the workload migration in either the Migrate Client or the Web Interface, depending on your migration goals. See “Migration Operations Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface” on page 82.

After you initiate a copy job or migration job in the Migrate Client, the Web Interface displays read-only status for the Migrate Client, as described in Table 3-3. You can use the filter on the Workloads page to show **Client Managed Workloads**.

Table 3-3 *Read-Only Status for Migrate Client Copy or Move Migration Jobs*

Migrate Client Job Status	Description
Not Configured	The source workload has been added and details have been discovered, but no configuration has been attempted. The workload can be managed by either client at this point.
Client migration in progress	A Copy or Move migration job for the source workload has been initiated in the Migrate Client. The migration is in progress.
Client migration stuck	A recoverable error occurred during replication for a Copy or Move migration job. User intervention is required in the Migrate Client.
Client migration failed	A non-recoverable error occurred during replication for a Copy or Move migration job. User intervention is required in Migrate Client.
Client copy successful	A Copy migration job has ended successfully. After a typical Copy migration job, both the source workload and target workload are up and running.
Client migration successful	A Move migration job has ended successfully. After a typical Move migration job, the source workload is shut down and the target workload is up and running.

Filtering or Organizing Workloads in the Workloads View

On the Workloads page, you can filter the display of the discovered workloads. For example:

- To display all the workloads that are not yet configured, select the **Workload Status** option as **Not Configured** and the **Tag** option as **All**.
- To display all the failed Windows workloads, select the **Workload Status** option as **Failed Workloads** and the **Tag** option as **Windows.**

  For information about how to create tags and associate them with workloads, see "Using Tags to Track Logical Associations of Workloads" on page 309.

You can sort on values in any column by click the column heading.

To filter the listing of workloads:

1. In the **Workload Status** menu, select one of the following:
   - All workloads
   - Replicated
   - Scheduled
   - Running Cutover
   - Running Test Cutover

- Running Replication
- Failed Workloads
- Running Workloads
- Not configured
- Ready for Replication
- Cut Over

2 (Optional) In the **Tag** menu, select the tag that is associated with the workloads you want to list, or select **All**.

For information about how to create tags and associate them with workloads, see "Using Tags to Track Logical Associations of Workloads" on page 309.

Viewing Details for a Source Workload

After you discover a source workload, you can view its Discovery Details. After you begin configuring its migration, you can view its Migration Details.

1 On the Workloads page, click the **Name** link of the workload of interest.

2 View the Discovery Details or Migration Details, depending on where it is in the migration lifecycle.

3 (Optional) Select the Command Details tab to view information about events for the last command executed on the workload.

Viewing the Command Details for a Source Workload

After discover a source workload, you can view its Command Details to learn more about related events.

1 On the Workloads page, click the **Migration Status** link of the workload of interest.

2 On the Command Details page, view information about events for the last command executed on the workload.

3 (Optional) If Workload Commands are active for the workload, you can initiate a follow-on action for the migration by clicking the appropriate action.

OS Icons in the Workloads View

Migrate Web interface does not distinguish source workloads by the source origin of physical, virtual, or cloud. Discovered source workloads are associated with unique icons to help identify the type of workload operating system.

Table 3-4 *Operating System Icons in the Workloads View*

	Windows operating systems
	Linux operating systems

Targets

The Targets page displays the target platforms available for the migration jobs to VMware and cloud targets. You can add a new target platform in the Web Interface for VMware and cloud infrastructure-as-a-service (IaaS) platforms. See

- "Supported Target VMware Platforms" on page 40
- "Supported Target Cloud Platforms" on page 43

Each platform is identified by the cloud provider or the specific operating system installed on the VMware host server.

For information about adding target platforms, see Chapter 17, "Discovering Target Platforms," on page 279.

Tasks

The Tasks page displays the most recent tasks, the most recent events, and the upcoming events.

Events are logged whenever some action related to the system or the workload occurs. For example, an event could be the addition of a new workload, the replication of a workload starting or failing, or the detection of the failure of a migrated workload. Some events also email automatic notifications if SMTP is configured. For more information, see "Notification Service Using Migrate Web Interface" on page 103.

Tasks are special operations tied to events that require user intervention. For example, upon completion of a Test Cutover operation, the system generates an event associated with two tasks: **Mark Test as Success** and **Mark Test as Failure**. When you click either of the tasks, the Test Cutover operation is canceled and a corresponding event is logged.

The Tasks and Events panel on the dashboard displays a maximum of three entries. To see all tasks or to see past and upcoming events, click **View All** in the appropriate section.

Dashboard

The Dashboard page provides information about the Migrate licenses, tasks, running events, upcoming events, and past events.

The left pane of the Dashboard page provides a high-level view of the overall state of the PlateSpin Migrate workload inventory, summary of the license information and also lets you add or discover a new workload. For more information about adding or discovering a workload, see "Workload Discovery in the Migrate Web Interface" on page 302.

The right pane of the Dashboard page provides information about events and tasks that requires user attention.

Reports

You can generate reports that provide analytical insight into your workload migration contracts over time.

The following report types are supported:

- **Workload Migration:** Reports replication events for all workloads over a selectable time window.
- **Migration History:** Reports replication size, time, and transfer speed per selectable workload over a selectable time window.
- **Replication Statistics:** Reports the dynamics of full and incremental replications that can be summarized by **Average**, **Most Recent**, **Sum**, and **Peak** perspectives.
- **Current Migration Status:** Displays the migration status such last test cutover, last replication date, and the test age (elapsed time since the last test cutover).
- **Events:** Reports system events for all workloads over a selectable time window.
- **Scheduled Events:** Reports only upcoming workload migration events.
- **Running Events:** Reports only workload migration events that are currently in progress.
- **Resource Usage:** Displays the resources configured for the target workload.

Migration Operations Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface

Migration Operation	PlateSpin Migrate Client	PlateSpin Migrate Web Interface
Migration to Amazon Cloud		
Physical to Amazon Cloud	✘	✓
Virtual to Amazon Cloud	✘	✓

Migration Operation	PlateSpin Migrate Client	PlateSpin Migrate Web Interface
Migration to Microsoft Azure		
Physical to Microsoft Azure	✘	✓
Virtual to Microsoft Azure	✘	✓
Migration to VMware vCloud Director		
Physical to VMware vCloud Director	✘	✓
Virtual to VMware vCloud Director	✘	✓
Migration to Oracle Cloud Infrastructure		
Physical to Oracle Cloud Infrastructure	✓	✓
Virtual to Oracle Cloud Infrastructure	✓	✓
Cloud-to-Cloud Migration		
Amazon Cloud to Microsoft Azure	✘	✓
Amazon Cloud to VMware vCloud	✘	✓
Amazon Cloud to Amazon Cloud (different regions in the same AWS EC2 Subscription)	✘	✓
Microsoft Azure to Amazon Cloud	✘	✓
Microsoft Azure to VMware vCloud	✘	✓
Microsoft Azure to Microsoft Azure (different locations in the same Azure subscription)	✘	✓
VMware vCloud to Amazon Cloud	✘	✓
VMware vCloud to Microsoft Azure	✘	✓
Migration to VMware Hosts		
Physical to VMware (P2V)	✓	✓
Virtual to VMware (V2V)	✓	✓
Migration to KVM		
Physical to Virtual (P2V)	✓	✘

Migration Operation	PlateSpin Migrate Client	PlateSpin Migrate Web Interface
Virtual to Virtual (V2V)	✓	✘
Migration to Physical Hosts		
Physical to Physical (P2P)	✓	✓
Virtual to Physical (V2P)	✓	✓

Migration Tasks Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface

To migrate a workload, you should either use the PlateSpin Migrate Client or the PlateSpin Migrate Web Interface throughout the migration cycle of the workload.

The following table lists the tasks that you can perform using the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface:

Tasks	PlateSpin Migrate Client	PlateSpin Migrate web Interface
Monitor workload migration workflow	✘	✓
Discover Windows standalone workloads	✓	✓
Discover Windows cluster workloads	✓	✓
Discover Linux standalone workloads	✓	✓
Discover Linux cluster workloads	✘	✘
Discover target VMware hosts	✓	✓
Discover target KVM hosts	✓	✘
Discover target cloud platforms	✘	✓
Migrate to physical machines	✓	✓
Migrate to VMware hosts	✓	✓
Migrate to KVM hosts	✓	✘
Migrate to Azure Cloud	✘	✓

Tasks	PlateSpin Migrate Client	PlateSpin Migrate web Interface
Migrate to Amazon Web Services	✗	✓
Migrate to VMware vCloud Director	✗	✓
Migrate to Oracle Cloud	✓	✓
Migrate Windows workloads with block-based transfer	✓	✓
Migrate Linux workloads with block-based transfer	✓	✓
Migrate Windows workloads with file-based transfer	✓	✓
Migrate Linux workloads with file-based transfer	✓	✗
Migrate Windows clusters with block-based transfer	✓	✓
Schedule incremental replication	✗	✓
Support post migration scripts	✓	✓
Add new disks during migration	✓	✗
Change disk volume mapping for target workload	✓	✗
Migrate a VM to a vCenter folder	✗	✓
Move a VM to a resource pool	✓	✓
Set compression level	✓	✓
Throttle bandwidth	✓	✓
Set encryption for data transfer	✓	✓
Create tags	✗	✓
View workload migration report	✓	✓
View workload migration status reports	✗	✓
Add or remove licenses	✓	✓
Check licenses status	✓	✓

Tasks	PlateSpin Migrate Client	PlateSpin Migrate web Interface
Use security groups	✘	✓
Set global defaults for source service	✓	✘
Set global defaults for target service	✓	✘
Set global defaults for migration job values	✓	✘

Other PlateSpin Server Management Tools

PlateSpin Migrate provides additional tools you can use to help customize your migration efforts.

PlateSpin Configuration

Some aspects of your PlateSpin Server's behavior are controlled by configuration parameters that you set on a configuration web page residing on your PlateSpin Server host:

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

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

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

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`
2. Locate the required server parameter and change its value.
3. Save your settings and exit the page.

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

Table 3-5 lists common advanced settings that you might encounter in various migration scenarios. See the related documentation for more information.

Table 3-5 *Common Advanced Settings in PlateSpin Configuration*

Common Issue or Concern	Reference
Increasing the upload size limit on post-migration actions	"Increasing the Upload Size Limit for Post-Migration Actions" on page 109
Optimizing data transfer over WAN connections	"Optimizing Data Transfer over WAN Connections" on page 117
Enabling source workloads to initiate contact with target workloads for replications when sources are registered using Migrate Agent	"Configuring the Contact Direction for the Replication Port" on page 114
Adding or removing applications from a list of applications that will be temporarily disabled for conversions	"Configuring Applications Known to Cause Boot Failure on Windows Target Workloads" on page 116
Rebranding some features in the Web Interface	Appendix B, "Rebranding the UI for PlateSpin Migrate Web Interface," on page 137
Configuring parameters that control migration behavior for migrations to Amazon Web Services	"Configuring Advanced PlateSpin Settings for AWS" on page 155
Configuring Azure tags used by all migrations to Azure Cloud on the Migrate server	"Using Azure Cloud Tags for Azure Migrations" on page 194
Configuring parameters that control migration behavior for migrations to VMware vCloud Director	"Configuring Advanced PlateSpin Settings for vCloud" on page 206

PlateSpin Migrate Client Command Line Interface

The PlateSpin Migrate Client installation includes a command line interface (CLI) tool to help you perform common migrations tasks. Conversion jobs using `.ini` files is supported to VMware targets only. See Appendix I, "Using the PlateSpin Migrate Client Command Line Interface," on page 573

PlateSpin Analyzer

PlateSpin Migrate Client provides the PlateSpin Analyzer to determine whether discovered Windows machines are suitable for migration jobs. Before you begin any large-scale migration projects, you should identify potential migration problems and correct them beforehand. See "Analyzing Suitability of Discovered Windows Workloads For Conversion to Physical Machines" on page 320.

Migrate Agent Utility

The Migrate Agent utility is a command line utility that you can use to install, upgrade, query, or uninstall the block-based transfer drivers. The utility also enables you to register source workloads with PlateSpin Migrate servers and send details about the workloads to the server via HTTPS (TCP/

443). Registration allows you to add workloads that cannot be discovered, such as for Migrate Servers in Microsoft Azure when no VPN is configured between the Migrate server and the source workloads.

A reboot is not required for source Linux workloads. Although a reboot of the source Windows workload is always required when you install, uninstall, or upgrade drivers, the Migrate Agent utility allows you to better control when the action occurs and therefore, when the server reboots. For example, you can use the Migrate Agent utility to install the drivers during scheduled down time, instead of during the first replication. See Appendix F, “Migrate Agent Utility,” on page 359.

PlateSpin ISO

The PlateSpin ISO file enables you to register target physical machines and target virtual machines with PlateSpin Migrate servers and send details about them to the server via HTTPS (TCP/443). Registration allows you to add target machines that cannot be discovered because they have no operating system installed. See Appendix G, “PlateSpin ISO Image,” on page 373.

4 Configuring PlateSpin Users and Access

Users have privileges to perform tasks in PlateSpin Migrate based on their assigned PlateSpin user roles: Administrator, Power User, and Operator.

NOTE: In your VMware environment, you can configure custom PlateSpin user roles to set minimal permissions for Migrate access for a single user or multitenancy environment. See "Configuring a PlateSpin User with Minimal Permissions on VMware" on page 250.

This section explains the various PlateSpin user roles, role-based privileges, and how to assign users to the roles.

Configuring User Authorization and Authentication

PlateSpin Migrate's user authorization and authentication mechanism is based on user roles, and controls application access and operations that users can perform. The mechanism is based on Integrated Windows Authentication (IWA) and its interaction with Internet Information Services (IIS).

PlateSpin Migrate's user auditing functionality is provided through the capability to log user actions. See "Managing Migrate Client User Activity Log" on page 127.

PlateSpin Migrate Roles

A PlateSpin Migrate role is a collection of PlateSpin Migrate privileges that entitle a particular user to perform specific actions. During installation, the PlateSpin Migrate installation program creates the following three local Windows groups on the PlateSpin Server host that map directly to the three PlateSpin Migrate roles that control user authorization and authentication

Group for PlateSpin Migrate Client Users	Group for PlateSpin Migrate Web Interface Users	Description
PlateSpin Administrators	Workload Conversion Administrators	Have unlimited access to all features and functions of the application. A local administrator is implicitly part of this group.
PlateSpin Power Users	Workload Conversion Power Users	Have access to most features and functions of the application with some limitations, such as restrictions in the capability to modify system settings related to licensing and security.
PlateSpin Operators	Workload Conversion Operators	Have access to a limited subset of system features and functions, sufficient to maintain day-to-day operation.

When a user attempts to connect to a PlateSpin Server, the credentials provided through the PlateSpin Migrate Client or Web Interface are validated by IIS. If the user is not a member of one of the PlateSpin Migrate roles, connection is refused. If the user is a local administrator on the PlateSpin Server host, that account is implicitly regarded as a PlateSpin Migrate Administrator.

The Permission details for the PlateSpin Migrate roles depends on whether you use the PlateSpin Migrate Client or the PlateSpin Migrate Web Interface for migrating the workloads:

- For information on PlateSpin Migrate Roles and permission details when you use PlateSpin Migrate Client to perform the workload migration, see Table 4-1 on page 90.
- For information on PlateSpin Migrate Roles and permission details when you use PlateSpin Migrate Web Interface to perform the workload migration, see Table 4-2 on page 91.

Table 4-1 *PlateSpin Migrate Roles and Permission Details For PlateSpin Migrate Client Users*

Role Details	Administrators	Power Users	Operators
Licensing: Add, delete licenses; transfer workload licenses	Yes	No	No
Machines: Discover, undiscover	Yes	Yes	No
Machines: Delete virtual machines	Yes	Yes	No
Machines: View, refresh, export	Yes	Yes	Yes
Machines: Import	Yes	Yes	No
Machines: Export	Yes	Yes	Yes
PlateSpin Migrate Networks: Add, delete	Yes	No	No
Jobs: Create new job	Yes	Yes	No
Jobs: View, abort, change start time	Yes	Yes	Yes
Block-Based Transfer Components: Install, upgrade, remove	Yes	Yes	No

Role Details	Administrators	Power Users	Operators
Device Drivers: View	Yes	Yes	Yes
Device Drivers: Upload, delete	Yes	Yes	No
PlateSpin Server access: View Web services, download client software	Yes	Yes	Yes
PlateSpin Server settings: Edit settings that control user activity logging and SMTP notifications	Yes	No	No
PlateSpin Server settings: Edit all server settings except those that control user activity logging and SMTP notifications	Yes	Yes	No
Run Diagnostics: Generate detailed diagnostic reports on jobs.	Yes	Yes	Yes
Post-conversion Actions: Add, update, delete	Yes	Yes	No

Table 4-2 *PlateSpin Migrate Roles and Permission Details For PlateSpin Migrate Web Interface Users*

Role Details	Administrators	Power Users	Operators
Add Workload	Yes	Yes	No
Remove Workload	Yes	Yes	No
Configure Migration	Yes	Yes	No
Prepare Migration	Yes	Yes	No
Run Full Replication	Yes	Yes	Yes
Run Incremental Replication	Yes	Yes	Yes
Pause/Resume Schedule	Yes	Yes	Yes
Test Cutover	Yes	Yes	Yes
Cutover	Yes	Yes	Yes
Abort	Yes	Yes	Yes
Settings (All)	Yes	No	No
Run Reports/Diagnostics	Yes	Yes	Yes

Assigning PlateSpin Migrate Roles to Windows Users

To allow specific Windows domain or local users to carry out specific PlateSpin Migrate operations according to designated role, add the required Windows domain or user account to the applicable Windows local group (PlateSpin Administrators, PlateSpin Power Users, or PlateSpin Operators) on the PlateSpin Server host. For more information, see your Windows documentation.

Configuring Permissions for Workload Access in PlateSpin Migrate Web Interface

PlateSpin Migrate Web Interface enables you to set permissions for workload migration management. You configure security groups and assign users and workloads to it. Only members of the security group are permitted to manage the member workloads in that group. See “Managing Security Groups and Workload Permissions” on page 131.

5 Configuring PlateSpin Migrate Server

Use the information in this section to configure your PlateSpin Migrate Server.

PlateSpin Migrate Product Licensing

This section provides information about licensing and activating your PlateSpin Migrate product and managing your license keys.

NOTE: You cannot use the Licenses that you purchased for PlateSpin Migrate 9.3 and later versions with PlateSpin Migrate 9.2 and prior versions.

Activating Your Product License

For product licensing, you must have a license activation code. If you do not have a license activation code, request one through the Customer Center (https://www.microfocus.com/customercenter/). A Micro Focus representative will contact you and provide the license activation code.

NOTE: If you are an existing PlateSpin customer and you do not have a Customer Center account, you must first create an account using the same email address as specified in your purchase order. See Create Account (https://www.microfocus.com/selfreg/jsp/createAccount.jsp).

Before you activate a license, consider whether you want to split the license for different migration scenarios.

- “License Splitting” on page 94
- “License Activation Using Migrate Client” on page 94
- “License Activation Using the Web Interface” on page 96

License Splitting

A license entitles you to one instance of PlateSpin Migrate per workload. Depending on the license you purchased, you can split a license either on a per-migration or a per-workload basis.

You can only split a license that has not yet been activated. For example, you can split a per-workload license of 1000 workloads into one license covering 400 workloads and another covering 600 workloads. You can split a per-migration license for 3000 migrations into one license for 1200 migrations and one license for 1800 migrations.

For assistance with multi-license scenarios, especially if you are uncertain how to utilize licenses across your network environment, see KB Article 7920876 (https://support.microfocus.com/kb/doc.php?id=7920876).

License Activation Using Migrate Client

When you launch the PlateSpin Migrate Client for the first time after installation, the License Activation Wizard opens and prompts you to activate your product license.

Figure 5-1 *License Activation Wizard*

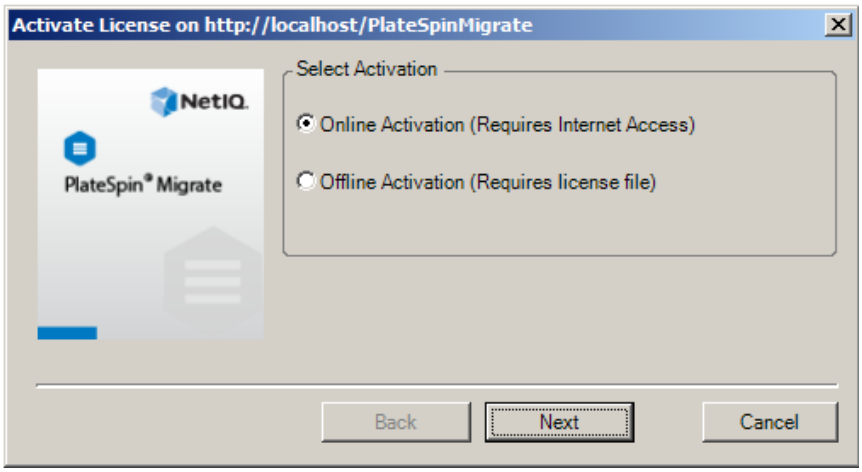

You have two options for activating your product license: online or offline.

- "Online License Activation" on page 95
- "Offline License Activation" on page 95

Online License Activation

Online activation requires that your PlateSpin Migrate Client have Internet access.

NOTE: HTTP proxies might cause failures during online activation. If you are using an HTTP proxy server and are having problems with online activation, try the offline activation method.

---

1. In the License Wizard, select the **Online Activation** option and click **Next**.
2. Enter the e-mail address that you provided when placing your order, and the activation code you received.

   The PlateSpin Migrate Client obtains the required license over the Internet and activates the product.

Offline License Activation

For offline activation, you obtain a license key over the Internet by using a machine that has Internet access.

1. In the License Wizard, select the **Offline Activation** option and click **Next**.

   The Activate License dialog box is displayed:

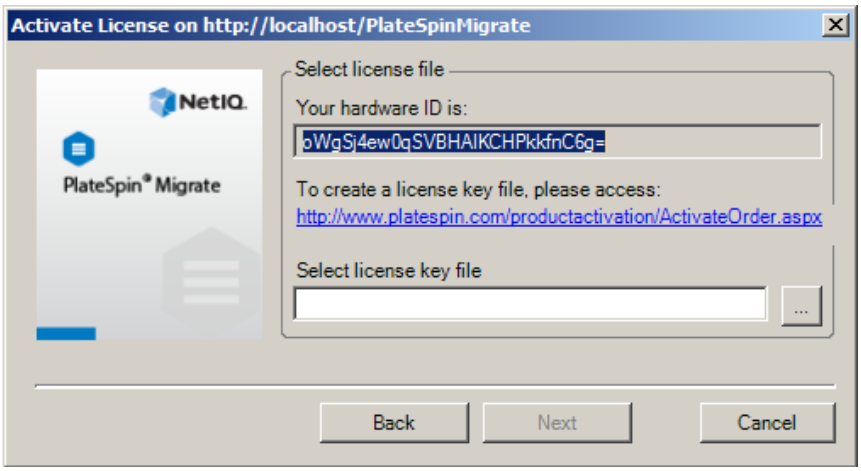

2 Save your hardware ID for use in the next steps.

3 Use a computer with Internet access to obtain a license key through the Web-based license activation utility (http://www.platespin.com/productactivation/ActivateOrder.aspx).

   To obtain a license key, you must have a Customer Center account. If you are an existing PlateSpin customer and you don't have a Customer Center account, you must first create one. (See Create Account.) Use your existing PlateSpin username (a valid e-mail address registered with PlateSpin) as input for your Customer Center account username.

4 Save your new license key in a location accessible to your PlateSpin Migrate Client.

5 In the License Wizard, type the full path to, or browse to and select, the PlateSpin Migrate license file, then click **Next**.

   The product is activated based on the selected license.

License Activation Using the Web Interface

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

Figure 5-2 *License Activation Using Migrate Web Interface*

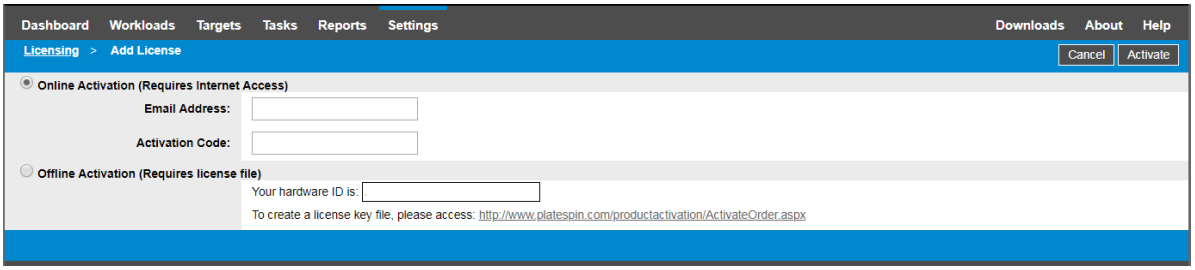

- "Online License Activation" on page 96
- "Offline License Activation" on page 97

Online License Activation

Online activation requires that your PlateSpin Migrate Web Interface has Internet access.

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

To set up online license activation:

1. In the PlateSpin Migrate Web Interface, click **Settings > Licensing**, then click **Add license**.
2. Click **Online Activation**.
3. Specify the email address that you provided when you placed your order and the activation code you received, then click **Activate**.

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

Offline License Activation

For offline activation, you must first use a computer that has Internet access to obtain a PlateSpin Migrate license key.

1. In the PlateSpin Migrate Web Interface, click **Settings > Licensing**, then click **Add license**.
2. Click **Offline Activation** and copy the hardware ID displayed in the interface.
3. Use a web browser on a computer that has Internet access to navigate to the PlateSpin Product Activation website (http://www.platespin.com/productactivation/ActivateOrder.aspx). Log in with your Customer Center user name and password.
4. Open the PlateSpin Activate Order page to generate a license key file. You need the following information:
   - activation code that you received
   - email address that you provided when you placed your order
   - hardware ID that you copied in Step 2
5. Save the generated license key file, transfer it to the product host that does not have Internet connectivity, and use it to activate the product.
6. In the PlateSpin Migrate Web Interface on the License Activation page, browse to the location of the license key file, then click **Activate**.

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

How Migration Licensing Works

PlateSpin Migrate licenses are sold on a per-workload basis. A license entitles you to an unlimited number of migrations on a specific number of workloads. With every migration, a workload unit of the license is assigned to either the source or the target. The machine that has the workload unit assigned to it can subsequently be migrated an unlimited number of times. Each time a workload is assigned, the **Workloads remaining** number is decremented.

Following is a summary of the workload assignment behavior based on the end state of a source workload after a successful cutover of the workload. This license assignment is applicable only for workloads migrated with the transfer scope as full migration. For information about the end state, see “Post-Cutover End States for Source and Target Workloads” on page 409.

Table 5-1 *PlateSpin Migrate Workload License Assignment*

Source Workload End State	Workload Assignment Behavior
Not shut down	A workload license remains with the source.
Shut down	A workload license is transferred from the source to the target.

Managing License Keys for Workload Migrations

You can add, delete, and monitor your PlateSpin licenses in the Migrate Client or Web Interface. Licenses can be used from migrations managed in either tool.

- “License Key Management with Migrate Client” on page 98
- “License Key Management Using Migrate Web Interface” on page 99

License Key Management with Migrate Client

You can manage available license keys on the License Manager’s **Available License Keys** tab.

1 In PlateSpin Migrate Client, click **Tools > License Manager** > **Available License Keys**.

Figure 5-3 *Available License Keys*

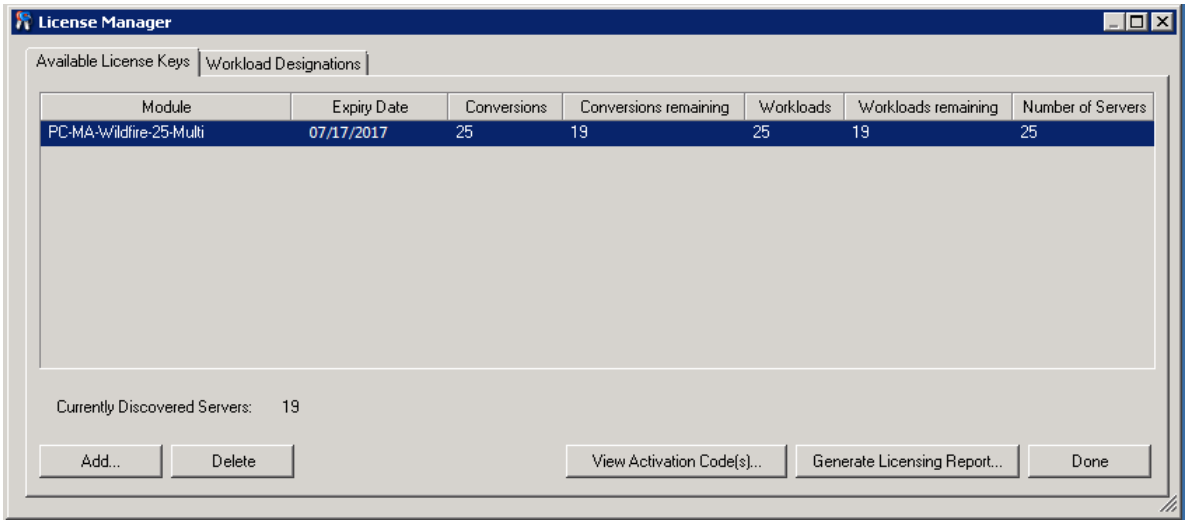

The tab displays the license name (**Module**) along with its expiry date and entitlements. These depend on the license type.

The **Number of Servers** column indicates the number of machines you can discover. This is generally the same as the number of machines that you can migrate.

Use the buttons at the bottom for related license management tasks:

Table 5-2 *License Manager Command Buttons*

Command	Description
Add	Adds licenses.
Delete	Deletes expired licenses.
View Activation Code(s)	Select a license and click this button to see the activation code and the date it was activated.
Generate Licensing Report	Creates a `*.psl` file that is used by Technical Support to troubleshoot licensing issues.

License Key Management Using Migrate Web Interface

You can manage available license keys on the **Licensing** tab in the Web Interface settings. In addition, the License Summary on the Web Interface Dashboard shows the total number of licenses and the number currently available.

1 In PlateSpin Migrate Web Interface, click **Settings > Licensing** > **Available Licenses**.

Figure 5-4 *Available License Keys*

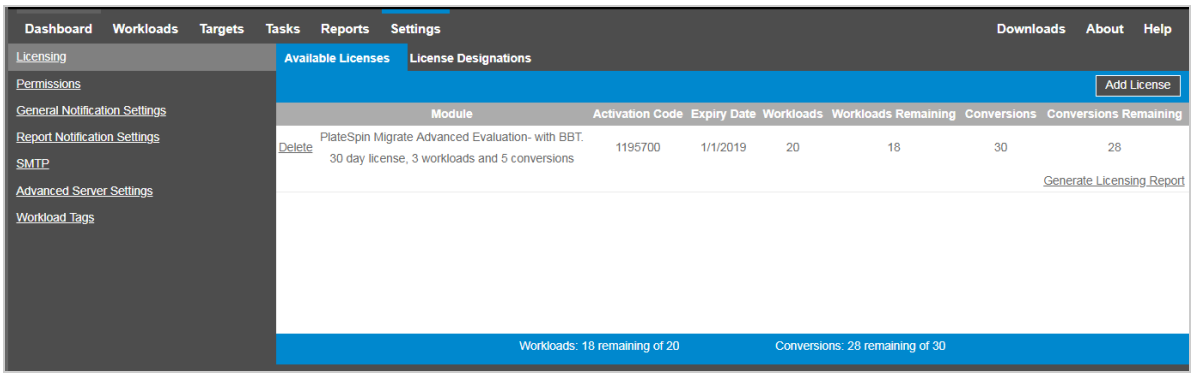

The Licensing tab displays the license name (**Module**) along with its activation code, expiry date, and the number entitlements (workload licenses available, workload licenses used, workload licenses remaining, conversions available, conversions used, and conversions remaining) for workload migration.

The sum total of all workload licenses available and remaining is displayed at the bottom of the window.

Use the options for related license management tasks:

Table 5-3 *Licensing Tab Options*

Command	Description
Add License	Adds a new license.
Delete	Deletes expired licenses.
Generate Licensing Report	Creates a `LicenseReport.txt` file that is used by Technical Support to troubleshoot licensing issues.

Managing Workload Designations

You can view license allocations for workloads in the Migrate Client or Web Interface. However, the PlateSpin Migrate Client lets you manage license allocations also.

- "Managing Workload Designations Using Migrate Client" on page 100
- "Viewing Workload License Designations Using Migrate Web Interface" on page 101

Managing Workload Designations Using Migrate Client

In the PlateSpin Migrate Client, you can view and manage license allocations on the License Manager's **Workload Designations** tab.

1 In PlateSpin Migrate Client, click **Tools > License Manager** > **Workload Designations**.

Figure 5-5 *License Manager Workload Designations*

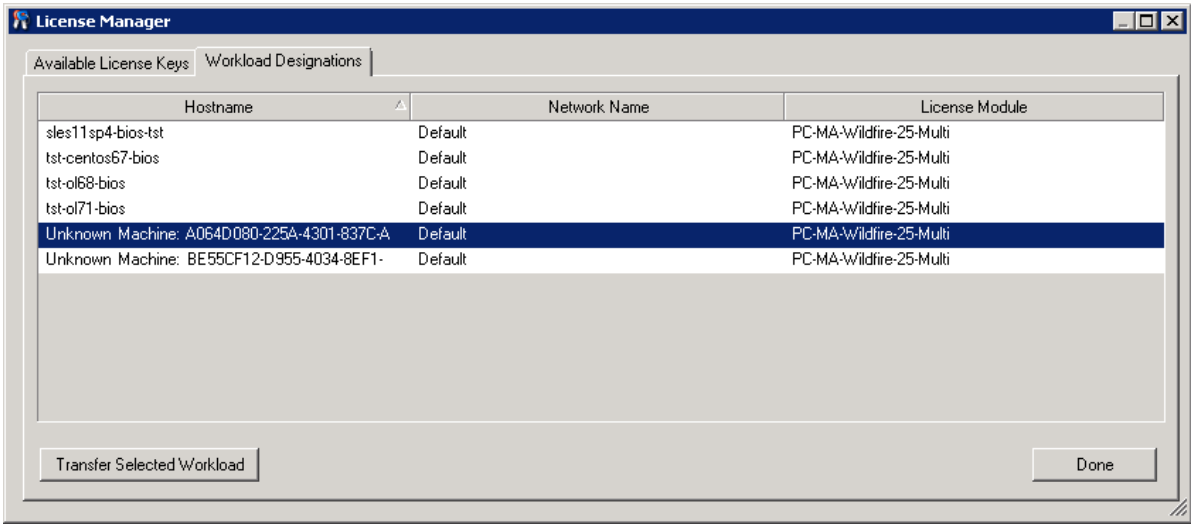

The tab lists workloads with assigned licenses. In the PlateSpin Migrate Client Servers view, each of these servers has a key icon adjacent to it.

You can reset workload licensing so that a license is no longer assigned to a particular machine. For example, you might want to do this when decommissioning servers that are already in the inventory of the PlateSpin Server.

To reset workload licensing:

1 On the License Manager's **Workload Designations** tab, select the required workload and click **Transfer Selected Workload**.

The Transfer License dialog box is displayed.

2 Use the displayed **Workload Transfer Request** string to obtain a workload transfer code from the License Entitlement Web portal (http://www.platespin.com/entitlementmgr/). Log in with credentials associated with your purchase order.

You must have a Customer Center account. If you are an existing PlateSpin customer and you don't have a Customer Center account, you must first create one. (See Create Account.) Use your existing PlateSpin username (a valid e-mail address registered with PlateSpin) as input for your Customer Center account username.

3 Return to the License Manager and specify the newly obtained transfer code. Click **Next**.

PlateSpin Migrate resets the selected workload.

Viewing Workload License Designations Using Migrate Web Interface

In the PlateSpin Migrate Web Interface, click **Settings > Licensing > License Designations** to view license allocations for workloads.

Figure 5-6 *License Designations*

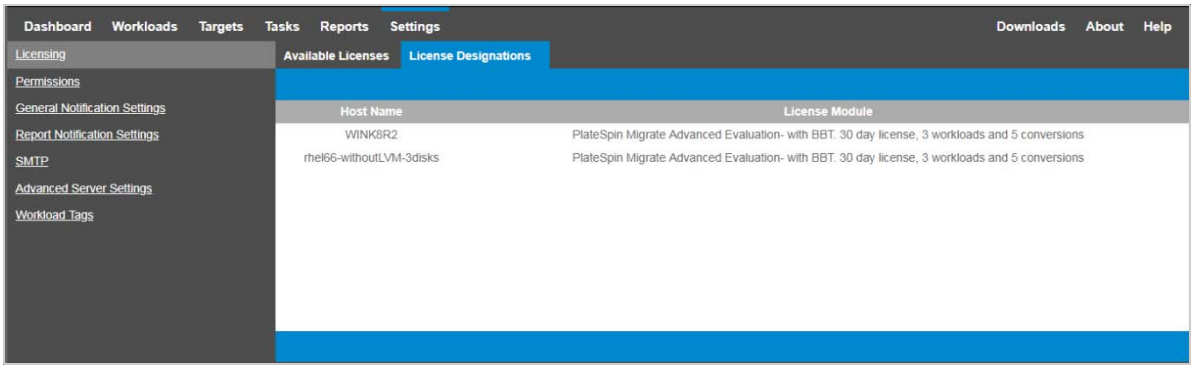

Enforcing FIPS Compliance for FIPS-Enabled Source Workloads

If FIPS is enabled in the source workload, you must enable the **EnforceFIPSCompliance** parameter before you discover the source workload:

1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2 Locate the **EnforceFIPSCompliance** parameter and click **Edit** to change its value to `True`.

3 Click **Save**.

After you modify the settings in the configuration tool, it might take up to 30 seconds for the change to take reflect on the interface. You need not reboot or restart the services.

4 Discover the FIPS enabled source workload.

Configuring the Notification Service

You can configure PlateSpin Migrate to automatically send notifications of events and replication reports to specified email addresses. This functionality requires that you first specify a valid Simple Mail Transfer Protocol (SMTP) server for PlateSpin Migrate to use.

- "Notification Service Using Migrate Client" on page 102
- "Notification Service Using Migrate Web Interface" on page 103

Notification Service Using Migrate Client

The PlateSpin Migrate Client enables you specify Simple Mail Transfer Protocol (SMTP) server settings for event and job progress notifications.

To configure the SMTP settings for the Notification Service:

1 Launch the PlateSpin Migrate Client.

2 Click **Tools > Options**.

3 Click the **Notification Service** tab.

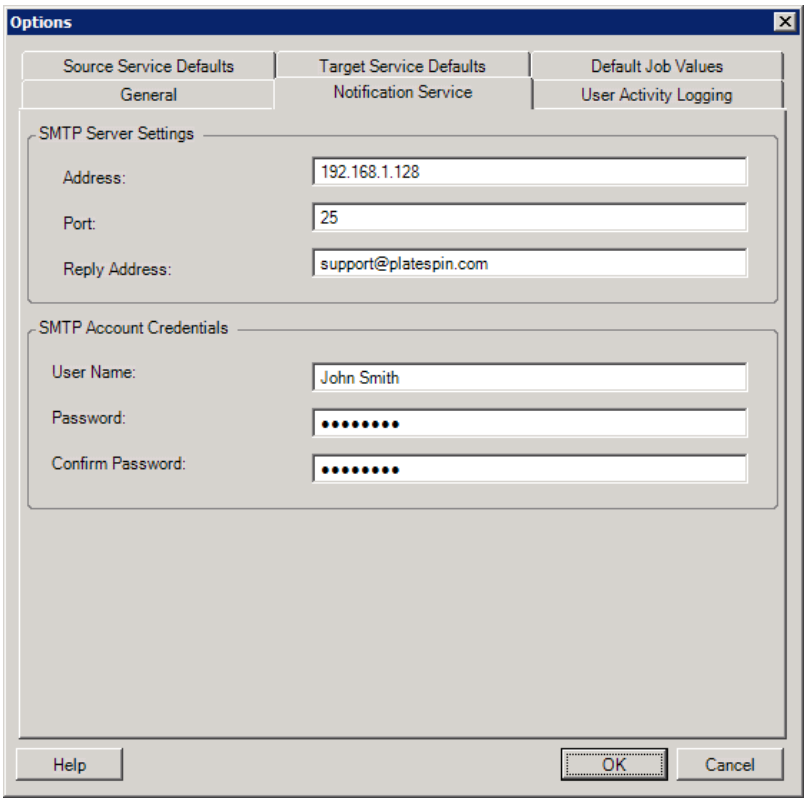

SMTP Server Settings: Specify your SMTP server's IP address, port, and a reply address for e-mail event and progress notifications.

SMTP Account Credentials: Provide valid credentials if your SMTP server requires authentication.

You can also configure migration progress notifications on a per-migration basis. See "Notifications Using the Migrate Client" on page 105.

Notification Service Using Migrate Web Interface

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

- "Setting up the SMTP Server" on page 103
- "Setting Up Event Notifications by Email" on page 103
- "Setting Up Replication Report Notifications by Email" on page 105

Setting up the SMTP Server

1. In the Migrate Web Interface, click **Settings > SMTP**.
2. Specify the following:
   - **SMTP Server Address:** The address of the SMTP server.
   - **Port:** The port at which the SMTP server is listening. By default, it is 25.
   - **Reply Address:** The address from which you want to send email event and progress notifications.
   - **Username and Password:** Provide valid credentials if your SMTP server requires authentication.
3. Click **Save**.

Setting Up Event Notifications by Email

To set up event notifications:

1. Configure an SMTP server for PlateSpin Migrate to use. See "Setting up the SMTP Server" on page 103.
2. In the PlateSpin Migrate Web Interface, select **Settings > General Notification Settings**.
3. Select the **Enable Notifications** check box.
4. Click **Edit Recipients**, specify the required email addresses separated by commas and click **OK**.
5. Click **Save**.

   To delete an email address, click **Remove** next to the address that you want to delete.

The following event types triggers email notifications if notification is configured. The events are always added to the System Application Event Log according to the log entry types such as Warning, Error, and Information.

Event Types	Remarks
Log Entry Type: Warning	
`IncrementalReplicationMissed`	Generates when any of the following applies: ◆ A replication is manually paused when a scheduled incremental replication is due. ◆ The system attempts to carry out a scheduled incremental replication when a manually-triggered replication is in progress. ◆ The system determines that the target has insufficient free disk space.
`FullReplicationMissed`	Similar to `IncrementalReplicationMissed` event.
`WorkloadOfflineDetected`	Generated when the system detects that a previously online workload is now offline. Applies to workloads whose migration state is not **Paused**.
Log Entry Type: Error	
`FailoverFailed`	Generates when a workload cutover action fails.
`FullReplicationFailed`	Generates when a full replication of the workload begins but is not able to complete successfully.
`IncrementalReplicationFailed`	Generates when an incremental replication of the workload begins but is not able to complete successfully.
`PrepareFailoverFailed`	Generates when the preparation for workload cutover fails.
Log Entry Type: Information	
`FailoverCompleted`	Generates when workload cutover completes successfully.
`FullReplicationCompleted`	Generates when workload full replication completes successfully
`IncrementalReplicationCompleted`	Generates when workload incremental replication completes successfully.
`PrepareFailoverCompleted`	Generates when the preparation for workload cutover completes successfully.
`TestFailoverCompleted`	Generates upon manually marking a Test Cutover operation a success or a failure.
`WorkloadOnlineDetected`	Generates when the system detects that a previously offline workload is now online. Applies to workloads whose migration state is not **Paused**.

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

Setting Up Replication Report Notifications by Email

1. Set up an SMTP server for PlateSpin Migrate to use. See "Setting up the SMTP Server" on page 103.
2. In the PlateSpin Migrate Web Interface, select **Settings > Report Notification Settings**.
3. Select the **Enable Report Notifications** check box.
4. In the **Report Recurrence** section, click **Edit** and specify the required recurrence pattern for the reports.
5. In the **Recipients** section, click **Edit Recipients** to specify the required email addresses separated by commas and click **OK**.
6. (Optional) In the **Migrate Access URL** section, specify a non-default URL for your PlateSpin Server.

   For example, if your PlateSpin Server host has more than one NIC or is located behind a NAT server. This URL affects the title of the report and the functionality of accessing relevant content on the server through hyperlinks within emailed reports.
7. Click **Save**.

For information on other types of reports that you can generate and view on demand, see "Generating Workload and Workload Migration Reports" on page 545.

Configuring Notifications for Events and Migrations

After you specify a valid Simple Mail Transfer Protocol (SMTP) server for PlateSpin Migrate to use, you can configure PlateSpin Migrate to automatically send notifications of events and replication reports to specified email addresses.

Notifications Using the Migrate Client

You can set up a migration job to automatically send email notifications about status and progress to a specified address:

- **Job events:** Job status messages such as `Completed`, `Recoverable Error`, and `Failed`.
- **Job progress:** Detailed job progress messages at configurable intervals.

You specify SMTP server and email account details globally. You can also specify job-specific email addresses. See "Configuring the Notification Service" on page 102.

To set up email notifications:

1. In the PlateSpin Migrate Client, configure information for the SMTP server for PlateSpin Migrate to use. See "Notification Service Using Migrate Client" on page 102.

2 Start the migration job. For information about starting a migration job, see "Initiating a Migration Job" on page 386.

3 In the Job Configuration section of the Migration Job window, click **Alerts** and configure the required options.

3a Select **Receive Event Notifications** to receive notifications for Completed, Recoverable Error, and Failed conditions for migration jobs.

3b Select **Receive Progress Notifications** to receive progress notifications via email. Specify the frequency with which you want to receive notifications for the job.

3c (Optional) In **Send To Addresses,** add or remove job-specific email addresses that will receive notifications.

4 Click **OK**.

Notifications Using the Web Interface

To set up a list of recipients for event notifications:

1 In the PlateSpin Migrate Web Interface, configure information for the SMTP server for PlateSpin Migrate to use. See "Setting up the SMTP Server" on page 103.

2 Select **Settings > General Notification Settings**.

3 Select the **Enable Notifications** check box.

4 In the **Recipients** section, click **Edit Recipients** to specify the required email addresses separated by commas and click **OK**.

5 Click **Save**.

To set up a list of recipients for report notifications:

1 In the PlateSpin Migrate Web Interface, set up an SMTP server for PlateSpin Migrate to use. See "Setting up the SMTP Server" on page 103.

2 Select **Settings > Report Notification Settings**.

3 Select the **Enable Report Notifications** check box.

4 In the **Report Recurrence** section, click **Edit** and specify the required recurrence pattern for the reports.

5 In the **Recipients** section, click **Edit Recipients** to specify the required email addresses separated by commas and click **OK**.

6 (Optional) In the **Migrate Access URL** section, specify a non-default URL for your PlateSpin Server.

For example, if your PlateSpin Server host has more than one NIC or is located behind a NAT server. This URL affects the title of the report and the functionality of accessing relevant content on the server through hyperlinks within emailed reports.

7 Click **Save**.

For information on other types of reports that you can generate and view on demand, see "Generating Workload and Workload Migration Reports" on page 545.

Managing Post-Migration Actions

PlateSpin Migrate supports the use of scripts to automatically execute custom post-migration tasks on the target workload. You configure the post-migration action in a batch file; a shell script; or a program executable, save the action and all its dependencies in a dedicated folder without any sub folders, and then add it to the PlateSpin Server library. For information about adding the action to the PlateSpin Server library, see "Adding Post-Migration Actions to PlateSpin Server Custom Actions Library" on page 107.

You can then associate this custom action with migration jobs you configure using the PlateSpin Migrate Client or the PlateSpin Migrate Web Interface. Towards the end of the migration process, PlateSpin Migrate uploads the specified action and its dependencies to the target and executes the action on the target.

Adding Post-Migration Actions to PlateSpin Server Custom Actions Library

PlateSpin Migrate enables you to execute a custom action on your target workload as a part of the migration job. You must define and save your custom actions and all their dependencies within a dedicated folder (without any subfolders) in advance and add them to the PlateSpin Server library.

Create the action, test it on a sample workload, and save it together with all its dependencies in a dedicated folder (without any subfolders) that the PlateSpin Server can access. The maximum size of the folder you upload must not exceed 64 MB. For information about increasing this limit, see "Increasing the Upload Size Limit for Post-Migration Actions" on page 109.

Take special care when developing post-migration actions for Linux workloads, which allow different characters in file names and support different ACL (Access Control List) permissions.

For Linux workloads, use tar (or a similar tool) to amalgamate the action's directory structure into a single file. See KB Article 7970214 (https://support.microfocus.com/kb/doc.php?id=7970214).

Adding Custom Actions Using the Migrate Client

1. In the PlateSpin Migrate Client, click **Tools > Manage Actions**.
2. Click **Add**.

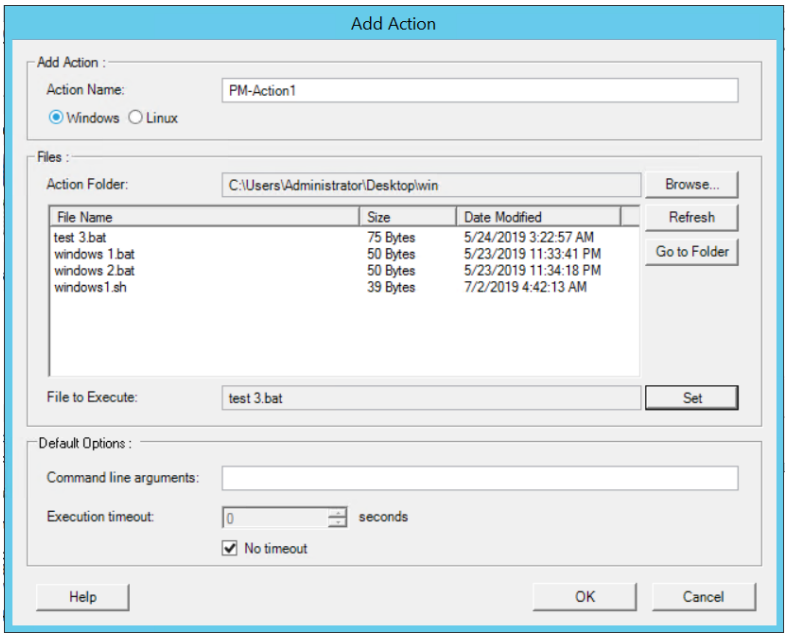

3 In the Add Action window, type a name for your custom action, select the target operating system type, then browse to and select the folder that contains the required action with all its dependencies. PlateSpin Migrate populates the list with the contents of the selected folder.

IMPORTANT: Ensure that the selected folder contains the action and all its dependencies without any sub folders.

---

4 In the **File Name** column, select the required executable, then click **Set**.

5 (Optional) In the **Default Options** section, specify the following values for the action:

- **Command line arguments:** Specify any required command line arguments that you want to be associated with this action.
- **Execution timeout:** Deselect the **No timeout** option to specify the maximum time to wait before timing out this action.

The command line parameters and timeout value that you specify is displayed when you select this action during the configuration of a migration job.

6 Click **OK**.

PlateSpin Migrate packages the action and its dependencies and uploads it to the library.

The uploaded action is then available for selection in migration jobs. See "Custom Post-Migration Actions" on page 398.

Adding Custom Actions Using the Migrate Web Interface

1 In the PlateSpin Migrate Web Interface, click **Settings > Post-Migration Actions**.

2 Click **Add Action**.

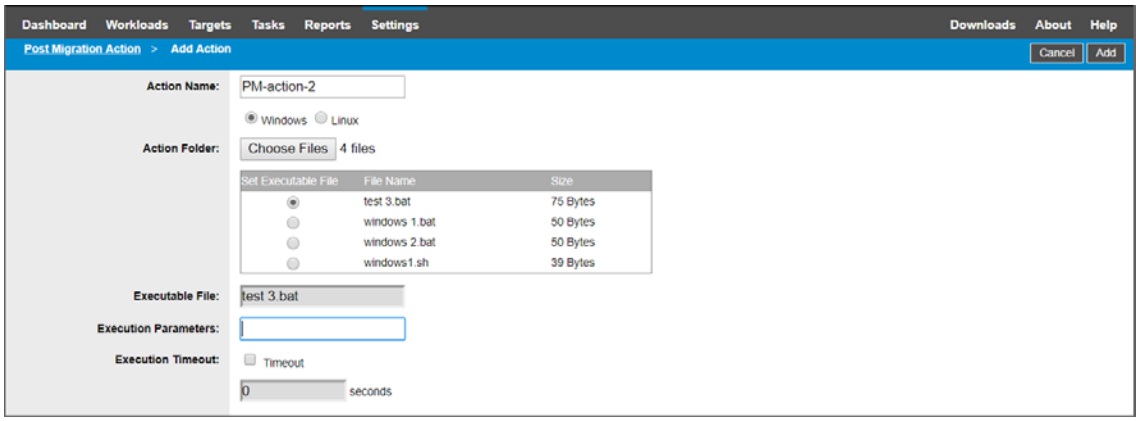

3 In the Add Action window, type a name for your custom action, select the target operating system type, then browse to and select the folder that contains the required action with its dependencies. PlateSpin Migrate populates the list with the contents of the selected folder.

IMPORTANT: Ensure that the selected folder contains the action and all its dependencies without any sub folders.

---

4 In the **Set Executable File** column, select the action executable to display the name of the selected executable in the **Executable File** option.

5 (Optional) Specify the following values for the action:

- **Execution Parameters:** Specify any required command line arguments that you want to be associated with this action.

  NOTE: Ensure that you enclose any argument that contains special characters within double quotes for Windows targets and within single quotes for Linux targets.

- **Execution Timeout:** Select the **Timeout** option to specify the maximum time to wait before timing out this action.

The command line parameters and timeout value that you specify is displayed when you select this action during the configuration of a migration job.

6 Click **OK**.

PlateSpin Migrate packages the action and its dependencies and uploads it to the library.

The uploaded action is then available for selection in migration jobs. See "Custom Post-Migration Actions" on page 398.

Increasing the Upload Size Limit for Post-Migration Actions

PlateSpin Migrate enables you to create custom scripts for post-migration actions and upload them to the PlateSpin Library. You can then associate them with certain migration jobs you configure in the PlateSpin Migrate Client. See "Managing Post-Migration Actions" on page 107.

By default, PlateSpin Migrate sets a 64 MB upload size limit for each individual post-migration action and its dependencies.

You can increase the upload size limit by modifying some attribute values in the `web.config` file on the PlateSpin Server host.

IMPORTANT: Decreasing the maximum upload size limit below the default of 64 MB might have a negative impact on the stability of your PlateSpin Server.

Modifying the Upload Size Limit for Custom Actions Using the Migrate Client

1. Close the PlateSpin Migrate Client.
2. Log in as Administrator to the PlateSpin Migrate Server host.
3. Browse to the `..\Program Files\PlateSpin Migrate Server\Web\` directory.
4. In a text editor, open the `web.config` file.
5. Locate the setting for the `httpRuntime` element with the `maxRequestLength` attribute:

   ```
   <httpRuntime maxRequestLength="65536" />
   ```

6. Replace the existing maximum upload size value of 65536 with the required new value in kilobytes.

   For example, to increase the maximum size from 64 MB to 128 MB, replace 65536 with 131072.

   ```
   <httpRuntime maxRequestLength="131072" />
   ```

7. Save the file, then restart the Migrate Client.

Modifying the Upload Size Limit for Custom Actions Using the Migrate Web Interface

1. Log in as Administrator to the PlateSpin Migrate Server host.
2. Browse to the `..\Program Files\PlateSpin Migrate Server\PlateSpin Forge\Web\` directory.
3. In a text editor, open the `web.config` file.
4. Locate the following setting for the `httpRuntime` element with the `maxRequestLength` attribute and replace the existing value of *xxxxx* with the required new value in kilobytes.

   ```
   <httpRuntime maxRequestLength="xxxxx" />
   ```

   For example, to increase the maximum size from 64 MB to 145 MB, set the value of `maxRequestLength` to 148480.

5. Locate the following setting for the `requestLimits` element with the `maxAllowedContentLength` attribute and replace the existing value of *yyyyy* with the required new value in bytes.

   ```
   <requestLimits maxAllowedContentLength="yyyyyy" />
   ```

For example, to increase the maximum size from 64 MB to 145 MB, set the value of `maxAllowedContentLength` to 152043520.

6 Save the file and refresh the PlateSpin Migrate Web Interface.

Configuring Alternate IP Addresses for PlateSpin Server

You can add alternate IP addresses to the PlateSpin Configuration **AlternateServerAddresses** parameter in order to enable the PlateSpin Server to function across NAT-enabled environments.

To add alternate IP addresses for PlateSpin Server:

1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2 Search to locate the **AlternateServerAddresses** parameter and add IP addresses for the PlateSpin Server.

3 Save your settings and exit the page.

A reboot or restart of PlateSpin services is not required to apply the changes.

Configuring the Maximum Number of IP Addresses Allowed for Each Network Interface on Target Workloads

You can specify multiple IP addresses per NIC on target workloads for migrations to VMware and the cloud (AWS, Azure, vCloud). By default, PlateSpin Migrate lets you specify up to four (4) IP addresses per NIC on the workload. PlateSpin Migrate controls this behavior by using the **IpAddressesPerNetworkInterface** parameter that is set to 4 by default. You can edit the value of this parameter to set the maximum number of IP addresses you want to specify per NIC.

The target platform might support more or fewer IP addresses per NIC than are allowed by this global limit. The lower limit is honored. For example, the maximum number of IP addresses per NIC that AWS supports depend on the AWS instance type. See AWS Documentation about the Elastic Network Interfaces (https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#AvailableIpPerENI).

To configure multiple IP addresses for NIC on the workload:

1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2 Search to locate the **IpAddressesPerNetworkInterface** parameter and set the value of this parameter to the maximum number of IP addresses you want to specify per NIC.

3 Save your settings and exit the page.

A reboot or restart of PlateSpin services is not required to apply the changes.

Enabling Public IP Address for Target Instances on Cloud

PlateSpin Migrate lets you specify whether or not a public IP address should be assigned to the primary network interface of AWS and Azure instances. PlateSpin Migrate controls this behavior by using the **EnablePublicIPForCloudInstance** parameter that is set to **False** by default. To enable public IP address to be assigned to the primary NIC on AWS and Azure Instances, set the value of **EnablePublicIPForCloudInstance** parameter to **True**.

NOTE: Before you configure the **EnablePublicIPForCloudInstance** parameter to enable public IP address for target instances on cloud, refer to the cloud-specific documentation for information about their public IP pricing.

To assign public IP address to the primary NIC on AWS and Azure Instances:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`
2. Search to locate the **EnablePublicIPForCloudInstance** parameter and set its value to **True**.
3. Save your settings and exit the page.

   A reboot or restart of PlateSpin services is not required to apply the changes.

When you configure migrations to AWS or Azure, the behavior of the Network Connections configuration option that enables AWS or Azure to automatically assign a public IP address to primary NIC on the cloud instance depends on the value set for the **EnablePublicIPForCloudInstance** parameter:

- **When EnablePublicIPForCloudInstance is True:**
  - **For AWS:** The **Auto-assign Public IP** option is selected by default and cannot be deselected.
  - **For Azure:** The **Create Public IP** option is selected by default and cannot be deselected.
- **When EnablePublicIPForCloudInstance is False:**
  - **For AWS:** The **Auto-assign Public IP** option is selected by default and you can deselect this option if you do not want AWS to automatically assign a public IP address to primary NIC on the target instance.
  - **For Azure:** The **Create Public IP** option is deselected by default and you can select this option if you want Azure to automatically assign a public IP address to primary NIC on the target instance.

For information about the **Auto-assign Public IP** option, see Network Connections for **Target Workload Settings** and **Target Workload Test Settings** in Configuring Migration of a Workload to Amazon Web Services.

For information about the **Create Public IP** option, see Network Connections for **Target Workload Settings** and **Target Workload Test Settings** in Configuring Migration of a Workload to Microsoft Azure.

Using Public IP to Register the Target Cloud Instance With Migrate Server

PlateSpin Migrate first uses a public IP address to register the target AWS and Azure instance with the Migrate server during migration. If the public IP address is not reachable, the connection attempt fails back to using a private IP address, and vice versa. PlateSpin Migrate controls this behavior by using the **UsePublicIPForRegistration** parameter that is set to **True** by default.

NOTE: Before you configure **UsePublicIPForRegistration** to enable public IP address for target instances on cloud, refer to the cloud-specific documentation for information about their public IP pricing.

To use a public IP for registering the target AWS and Azure instance with the Migrate server during migration:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2. Search to locate the **UsePublicIPForRegistration** parameter and ensure that its value is **True**:
3. Save your settings and exit the page.

   A reboot or restart of PlateSpin services is not required to apply the changes.

Setting Reboot Method for the Configuration Service

During a cutover action, the Configuration Service optimizes reboots by minimizing the number of reboots and controlling when they occur. If you experience a Configuration Service hang during a cutover action for a Windows workload with an error `Configuration Service Not Started`, you might need to allow reboots to occur as they are requested during the configuration. You can configure the single affected workload to skip reboot optimization, or configure a global **SkipRebootOptimization** parameter on the PlateSpin Server to skip reboot optimization for all Windows workloads.

To skip reboot optimization for a single Windows workload:

1. Log on as an Administrator user on the source workload.
2. Add a file at the root of the system drive (usually `C:`) called `PlateSpin.ConfigService.LegacyReboot` with no file extension. From a command prompt, enter

   ```
   echo $null >> %SYSTEMDRIVE%\PlateSpin.ConfigService.LegacyReboot
   ```

3. Run the failed Test Cutover or Cutover action again.

To skip reboot optimization for all Windows workloads:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2 Search for the **ConfigurationServiceValues** parameter, then click **Edit** for the parameter.

3 Change the setting **SkipRebootOptimization** from `False` to `True`.

4 Click **Save**.

5 Run the failed Test Cutover or Cutover again for affected Windows workloads.

Configuring the Contact Direction for the Replication Port

By default, the source workload contacts the target workload to initiate the replication data transfer. The direction is controlled at the server level by the **SourceListensForConnection** parameter that is set to `False` by default. To enable the target workload to contact the source workload in a VPN setup, reconfigure the replication port direction on the Migrate Server by setting the **SourceListensForConnection** parameter to `True` on the PlateSpin Configuration page.

To configure the direction of contact for replication traffic:

1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2 Locate the **SourceListensForConnection** parameter and edit its value as `True` or `False`, depending on your migration environment.

- **False:** (Default) The source workload contacts the target workload to initiate replication. The target listens for traffic on the replication port (default TCP/3725). The replication port must be open for inbound traffic on the target workload.
- **True:** The target workload contacts the source workload to initiate replication. The source listens for traffic on the replication port (default TCP/3725). The replication port must be open for inbound traffic on the source workload.

3 Save your settings and exit the page.

Configuring Behavior for Installing Network Drivers on Target Windows Workloads

When PlateSpin Migrate executes the Configuration Service on a target machine, Migrate by default performs the following networking tasks during the second reboot:

- Scans the network adapters and removes problematic ones.
- Uninstalls existing network drivers.
- Installs appropriate network drivers.
- Configures the network adapters according to the migration job configuration settings.

The normal networking tasks can be problematic in the following scenarios:

- If the target machine has the same network adapter hardware and networking drivers as the source machine.

  The network drivers that the target machine requires are the same as those already installed on the source machine being migrated. It is not necessary to re-install drivers. In some scenarios, removing and re-installing drivers can result in the target machine becoming unbootable.

- If the target machine is booting from SAN.

  If a target machine boots from SAN, Migrate installs drivers before the first boot. If the Configuration Service removes these newly installed drivers during the second reboot, the target machine becomes unbootable. It is necessary to avoid the driver install tasks on the second reboot.

You can configure the Migrate server to use a light networking approach in which Migrate does not perform the rescan, old driver uninstall, and new driver install during the second boot on target Windows workloads, including Windows Cluster workloads. It will perform customization as configured for the migration.

Using light networking to avoid the unneeded tasks optimizes the network configuration process and helps avoid situations that cause a target machine to become unbootable. Light networking is useful for P2P, V2V, and C2C migrations as well as for X2V semi-automated migrations where the networking hardware on the target VM is manually configured to match the source machine.

- "Understanding Light Networking Parameters" on page 115
- "Configuring Light Networking Parameters" on page 116

Understanding Light Networking Parameters

PlateSpin Configuration provides two light networking parameters to control whether or not PlateSpin Migrate should perform the networking driver tasks for specified target Windows workloads in any target platform. These parameters have no effect on Linux workloads.

EnableLightNetworking

If the **EnableLightNetworking** parameter is enabled, Migrate will not perform the following networking tasks on second reboot for specified target Windows workloads: rescan network adapters, uninstall old drivers, and install new network drivers. It will perform customization as configured for the migration. Avoiding the unneeded tasks optimizes the network configuration process for the target Windows workloads.

To take advantage of this light networking approach, set **EnableLightNetworking** to `True`, and then specify the host names of appropriate target Windows workloads in the **HostNamesForLightNetworking** parameter.

HostNamesForLightNetworking

The **HostNamesForLightNetworking** parameter is used to specify the target Windows workloads for which light networking rules should apply when **EnableLightNetworking** is set to `True`. Enable or disable the **EnableLightNetworking** parameter to control whether light networking is active for specified target Windows workloads.

Add the host names of target Windows machines in the following scenarios:

- If the source machine and target machine have the same networking hardware
- If the target machine boots from SAN

NOTE: If the target workload has different host names for test cutover and cutover, both host names must be listed in **HostNamesForLightNetworking**.

Valid values for the **HostNamesForLightNetworking** parameter are:

NONE

You can specify a value of NONE to enable all target Windows machines for light networking when the **EnableLightNetworking** parameter is set to `True`.

<FQDN>

Each value set for this parameter represents the FQDN (host name) of a target Windows workload for which light networking rules should apply when the **EnableLightNetworking** parameter is set to `True`.

If **EnableLightNetworking** value is set to `False`, the values in **HostNamesForLightNetworking** have no impact.

Configuring Light Networking Parameters

To configure the light networking parameters:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration`

2. Locate the **HostNamesForLightNetworking** parameter and edit its value as **NONE** or list one or more host names of target machines for which light networking should apply when the **EnableLightNetworking** parameter is set to `True`.
3. Locate the **EnableLightNetworking** parameter and edit its value as `True` or `False`, depending on your light networking needs.
   - **False:** (Default) Disable light networking for this Migrate server. The values set for the **HostNamesForLightNetworking** parameter have no impact.
   - **True:** Enable light networking for target machines, according to the values set in the **HostNamesForLightNetworking** parameter.
4. Save your settings and exit the page.

Configuring Applications Known to Cause Boot Failure on Windows Target Workloads

Some applications, such as backup and antivirus software, when installed on a source workload are likely to cause boot failure on the target workload if the corresponding application services are not disabled during the conversion.

The following parameters on the PlateSpin Server Configuration page helps you configure applications known to cause boot failures on the target:

- **ApplicationsKnownForBootFailuresOnTarget:** Lists some common applications such as Symantec, Kaspersky Antivirus, Backup Assist, and Carbon Black that are known to cause boot failure on the target. To edit the list of the applications, see "Editing the List of Applications Known to Cause Boot Failure on Windows Target" on page 117.
- **ApplicationsKnownForBootFailuresOnTargetDefaultValue:** Sets whether or not all the applications on the Windows source that are known to cause boot failure on the target be automatically selected for disabling during the conversion. The default value is `False` indicating that the applications are not selected by default.

When you configure the start-up mode of Windows services on the target, PlateSpin Migrate reviews the existing applications on the source to check if any of the applications listed in the **ApplicationsKnownForBootFailuresOnTarget** configuration parameter is installed on the source. PlateSpin Migrate lists all such source workload applications, which are known to cause boot failure on the target during conversion, on the user interface that you use to configure the start-up mode. These applications are selected by default if the value of the **ApplicationsKnownForBootFailuresOnTargetDefaultValue** parameter is set to `True`. However, you can review the listed applications and deselect the applications that you do not want to be disabled on the target during conversion.

For information about configuring the start-up mode of Windows services on the target, see "Service States on Target Windows Workloads" on page 402.

Editing the List of Applications Known to Cause Boot Failure on Windows Target

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`
2. Locate the **ApplicationsKnownForBootFailuresOnTarget** parameter and click **Edit**.
3. The **Values** option lists applications known to cause boot failure on target. Based on your requirement, add applications or remove existing applications whose boot services you do not want to disable during the conversion.
4. Save your settings and exit the page.

Optimizing Data Transfer over WAN Connections

You can optimize data transfer performance and fine tune it for WAN connections. You do this by modifying configuration parameters that the system reads from settings you make in a configuration tool residing on your PlateSpin Server host. For the generic procedure, see "PlateSpin Configuration" on page 86.

- "Tuning Parameters" on page 117
- "Tuning FileTransferSendReceiveBufferSize" on page 119

Tuning Parameters

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

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

Table 5-4 lists the configuration parameters on the PlateSpin Configuration page (`https://Your_PlateSpin_Server/PlateSpinConfiguration/`) that control file transfer speeds with the defaults and maximum values. You can modify these values through trial-and-error testing in order to optimize operation in a high-latency WAN environment.

Table 5-4 *Default and Optimized File Transfer Configuration Parameters*

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

Parameter	Default Value	Maximum Value
FileTransferSendReceiveBufferSize Defines the maximum size (in bytes) of the send and receive buffers for TCP connections in the replication network. The buffer size affects the TCP Receive Window (RWIN) size, which sets the number of bytes that can be sent without TCP acknowledgment. This setting is relevant for both file-based and block-based transfers. Tuning the buffer size based on your network bandwidth and latency improves throughput and reduces CPU processing. When the value is set to zero (off), the default TCP window size is used (8 KB). For custom sizes, specify the size in bytes. Use the following formula to determine the proper value: ((LINK_SPEED in Mbps / 8) * DELAY in sec)) * 1000 * 1024 For example, for a 100 Mbps link with 10 ms latency, the proper buffer size would be: `(100/8)* 0.01 * 1000 * 1024 = 128000 bytes` For tuning information, see "Tuning FileTransferSendReceiveBufferSize" on page 119.	0 (8192 bytes)	max 5242880 (5 MB)
FileTransferSendReceiveBufferSizeLinux Specifies the TCP/IP Receive Window (RWIN) Size setting for file transfer connections for Linux. It controls the number of bytes sent without TCP acknowledgment, in bytes. When the value is set to zero (off), the TCP/IP window size value for Linux is automatically calculated from the **FileTransferSendReceiveBufferSize** setting. If both parameters are set to zero (off), the default value is 248 KB. For custom sizes, specify the size in bytes. **NOTE:** In previous release versions, you were required to set this parameter to 1/2 the desired value, but this is no longer required.	0 (253952 bytes)	
FileTransferShutDownTimeOutInMinutes	1090	
FileTransferTCPTimeOutMilliSec Sets both the TCP Send and TCP Receive Timeout values.	30000	
PostFileTransferActionsRequiredTimeInMinutes	60	

Tuning FileTransferSendReceiveBufferSize

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

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

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

To calculate the optimum buffer size:

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

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

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

   1b Enter the following at a command prompt:

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

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

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

   ```
   The packet needs to be fragmented.
   ```

   1d Note the latency in seconds.

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

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

   ```
   Buffer Size = (Bandwidth in Mbps / 8) * Latency in seconds * 1000 * 1024
   ```

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

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

   ```
   Buffer Size = (10240 / 8) * 0.035 * 1000 * 1024 = 45875200 bytes
   ```

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

   3a Determine the MSS:

   ```
   MSS = MTU Size in bytes - (IP Header Size + TCP Header Size)
   ```

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

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

   ```
   MSS = 1470 bytes - (20 bytes + 20 bytes) = 1430 bytes
   ```

   3b Calculate the optimal buffer size:

   ```
   Optimal Buffer Size = (roundup( Buffer Size / MSS )) * MSS
   ```

To continue the example:

```
Optimal Buffer Size = (roundup(45875200 / 1430)) * 1430
                    = 32081 * 1430
                    = 45875830
```

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

```
Non-optimal Buffer Size = 32080 * 1430 = 45874400
```

Other Use Cases for Custom PlateSpin Server Settings (Advanced)

Table 5-5 lists configuration keys and values that might address various environmental or functional issues.

IMPORTANT: Do not use the settings in Table 5-5 unless you are advised to do so by PlateSpin Support.

Table 5-5 *List of Common Use Cases for Changing Settings in the Web Configuration Tool*

Issue or Use Case	Value Shown in the Config Tool

Discovery/Inventory issues	`<add key="UseServiceForCommonInventory" value="true" />`
	`<add key="UseServiceForMigrateInventory" value="false" />`
	`<add key="EarliestWindowsVersionForCommonInventory" value="5.2" />`

Target boot issues related to drivers	`<add key="TargetMachineRegistryCleanupLevel" value="None" />`
Controller installation issues on sources (mainly due to environmental constraints)	`<add key="InstallControllerUsingService" value="true" />` `<add key="RunControllerUnderSameCredsAsRemoteAccess" value="false" />`

Issues related to database size growth	`<add key="PowerConvertDBSizeLimitInBytes" value="4294967296" />`
	`<add key="PlateSpinDBCleanupThresholdPercent" value="80" />`
	`<add key="OFXDBCleanupThresholdInBytes" value="4294967296" />`

6 Configuring PlateSpin Migrate Client

PlateSpin Migrate Client enables you to configure global default settings that the Client uses for migration jobs, the source service, and the target service. These capabilities are available only for migration jobs configured and executed by using the Migrate Client. Use the information in this section to configure your Migrate Client.

Configuring General Options

PlateSpin Migrate Client enables you to restore default settings, clear saved credentials, and specify the locations of executable files for external applications that you can launch from within the Client.

To configure these general options:

1. Launch the PlateSpin Migrate Client.
2. Click **Tools > Options**.
3. Click the **General** tab.

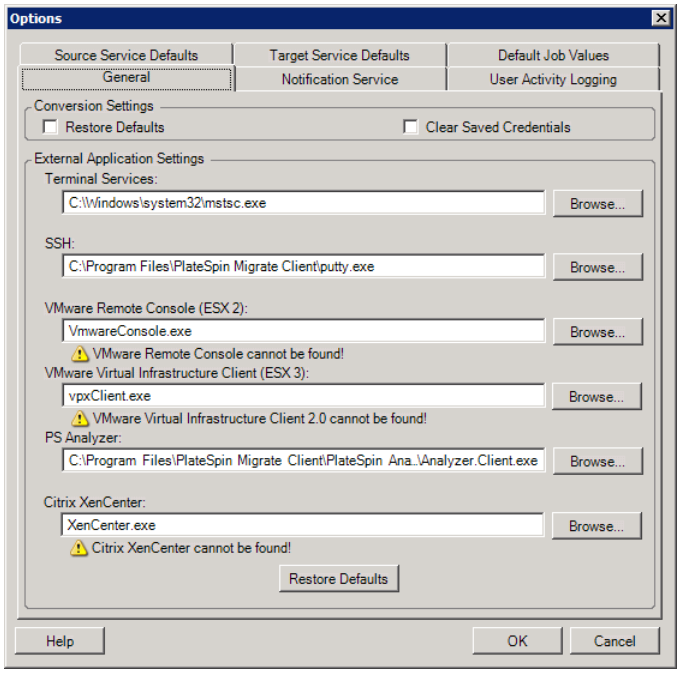

Restore Defaults: When this option is selected, PlateSpin Migrate resets the job configuration method (launches the Actions dialog box after a drag-and-drop) and resumes checking for software updates on the Client startup.

Clear Saved Credentials: Removes stored user names and passwords for source and target machines.

External Application Settings: Use the adjacent **Browse** buttons to locate application executables.

Restore Defaults: Resets the paths to their defaults.

Configuring Job Values Defaults

PlateSpin Migrate Client enables you specify default migration job values specific to the target virtualization platform.

To configure the default job values:

1. Launch the PlateSpin Migrate Client.
2. Click **Tools > Options**.
3. Click the **Default Job Values** tab.
4. In the Target Container Name and Path Defaults section, expand the required variable set (ESX Variables) and click a variable to edit its value. You can edit the following variables:

Variable Name	Variable Value	Remarks
ESX Variables		where: *%SOURCE_HOSTNAME%* is host name of the source computer. *%TARGET_DISK_EXTENSION%* is extension (`.vmdk` or `.vhd`) of the disk on the target workload.
Config Path	`/root/vmware/` *%SOURCE_HOSTNAME%*	
Disk Name	*%SOURCE_HOSTNAME%*`_#.`*%TARGET_DISK_EXTENSION%*	
Display Name	*%SOURCE_HOSTNAME%*	
ESX Config Path Within Datastore	%SOURCE_HOSTNAME%	
Config File Name	*%SOURCE_HOSTNAME%*`.vmx`	

5 In the Job Conversion Defaults section, set a default value for the following parameters that affect all migration jobs. The settings that you configure during the actual workload migration job overrides these default values.

Name	Value	Remarks
Encrypt File Transfer	◆ Yes ◆ No	See "Security of Workload Data in Transmission".
Take Control Network Settings	◆ Static ◆ DHCP	
Take Control Duplex Settings	◆ Auto-Negotiate ◆ 100 MB Full Duplex ◆ 1000 MB Full Duplex	
Install VMware Tools for ESX	◆ Yes ◆ No	See Virtualization Enhancement Software.
Virtual Disk Sizing Mode	◆ Fixed ◆ Dynamic	This setting is for ESX only. ◆ **Fixed**: Space is pre-allocated for the virtual disk ◆ **Dynamic**: The virtual disk is assigned a minimum amount of space, which grows when needed.
Compression Level	◆ None ◆ Fast ◆ Optimal ◆ Maximum	See Data Compression.
Reset	Restores default job values	
Update Defaults from Server	Retrieves defaults from the PlateSpin Server if available.	

Configuring Source Service Defaults

PlateSpin Migrate Client enables you to select Windows services and Linux daemons to stop on the source workload during migration. Because the source server is active during migration, this option helps ensure data consistency between the source and target workloads. See "Services or Daemons to Stop before Replication or Cutover" on page 400.

To configure the default services on the source:

1 Launch the PlateSpin Migrate Client.

2 Click **Tools > Options**.

3 Click the **Source Service Defaults** tab.

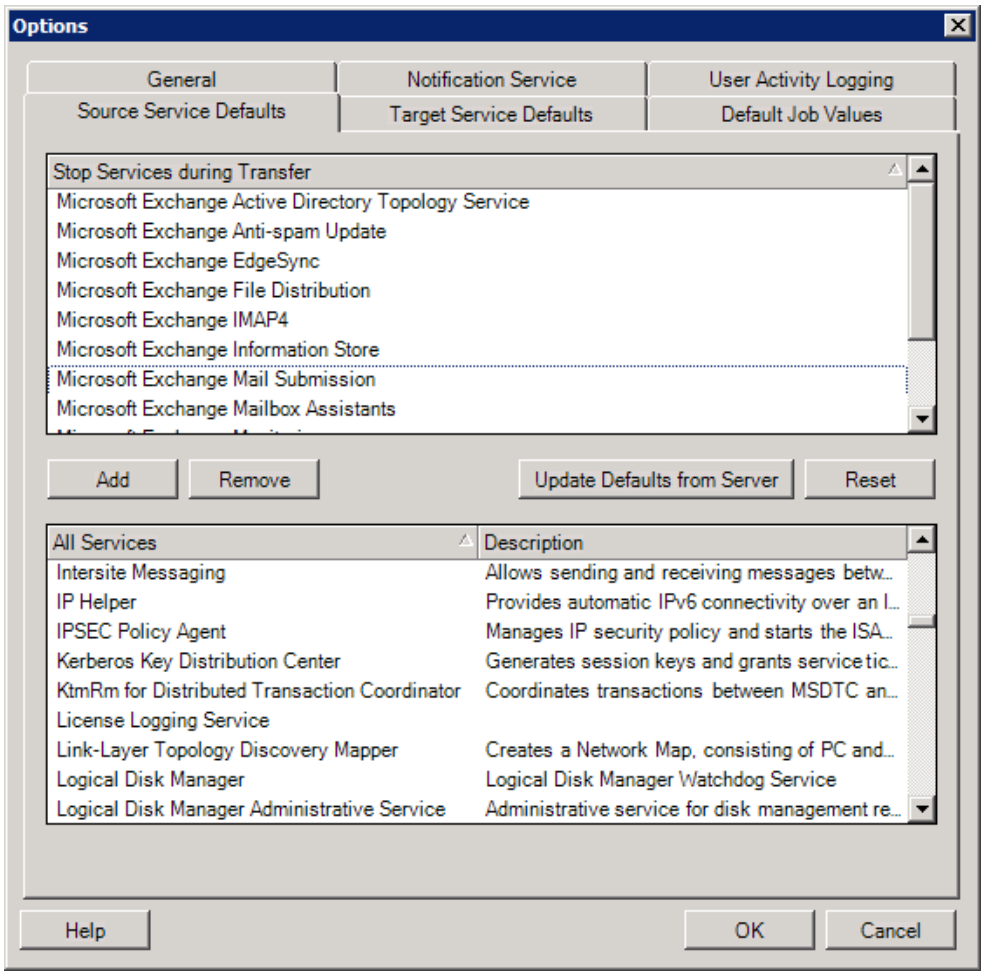

Stop Services during Transfer section: Lists services that are stopped by default. To stop a service during data transfer that uses a specific transfer method by default, select the corresponding check box. A deselected check box means the service remains active during data transfer.

All Services section: Lists unique services on all discovered machines. Click **Add** to add a selected service from the lower section to the upper section and set it to stop during the migration.

Update Defaults from Server: Retrieves defaults from PlateSpin Server.

Configuring Target Service Defaults

PlateSpin Migrate Client enables you to select Windows services whose mode on the target is to be different from that of the source. See Service States on Target Windows Workloads.

To configure the default services on the target:

1. Launch the PlateSpin Migrate Client.
2. Click **Tools > Options**.
3. Click the **Target Service Defaults** tab.

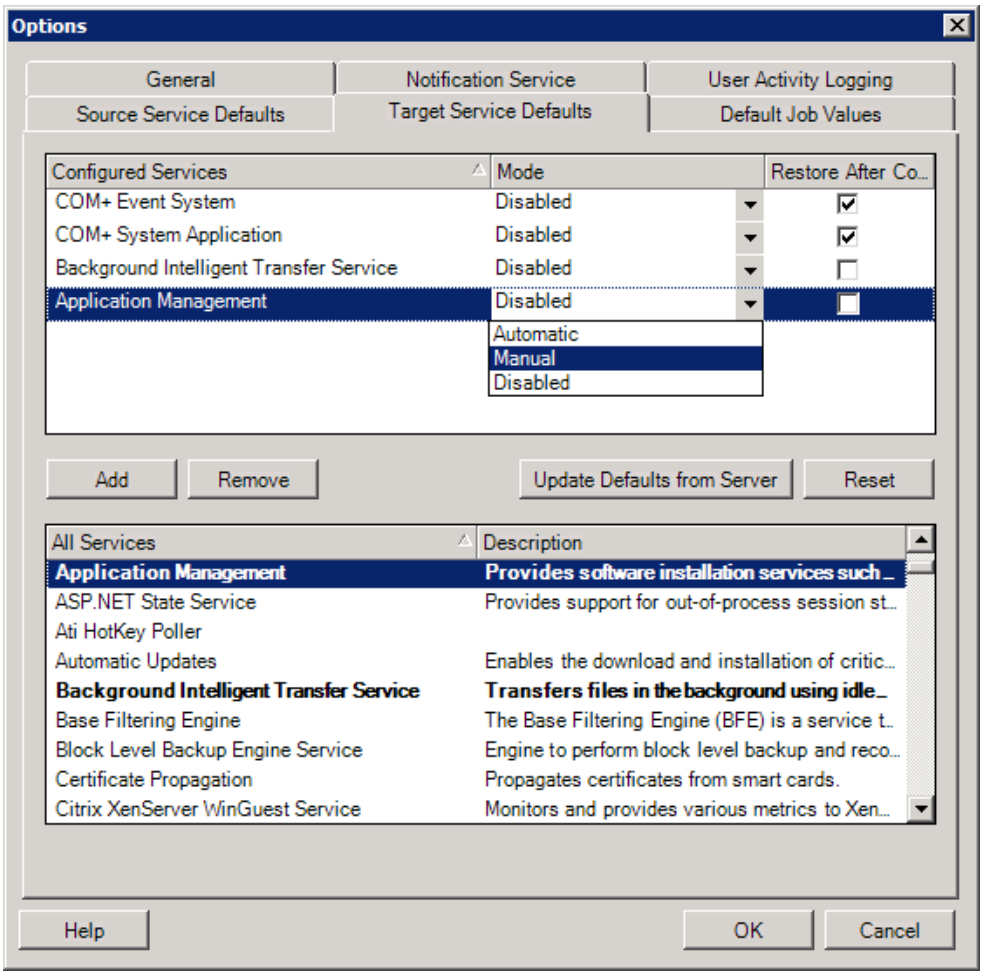

Configure Services section: Lists services and their target startup modes. Select the **Restore After Conversion** check box to use the selected mode during the migration. The service is then restored to match the source after the migration is complete and the target machine is ready to run.

All Services section: Lists unique services on all discovered machines. Click **Add** to add a service to the upper section. Use the **Mode** drop-down list to select the service state for the target. This is set during the configuration step of the job.

Remove: Removes a service.

Reset: Clears the upper section. The modes of all services in the target will match those on the source.

Managing Migrate Client User Activity Log

By default, PlateSpin Migrate Client logs all user activities that are performed in the Client. Actions logged include security, license management, target and workload discovery operations, and workload migration operations.

- "About the Migrate Client User Activity Log" on page 128
- "Configuring Migrate Client User Activity Logging" on page 128
- "Viewing Migrate Client User Activity Log" on page 129

About the Migrate Client User Activity Log

When User Activity Logging is enabled in PlateSpin Migrate Client, user actions performed in the Migrate Client are written in a User Activity Log file (`PlateSpin.UserActivityLogging.log`), located on your PlateSpin Server host, in the `..\PlateSpin Migrate Server\logs` directory.

The format of an individual log entry is:

```
date|Category|description|user|details1|details2
```

The Category element describes the functional area applicable to a particular action:

- `Security`
- `LicenseManagement`
- `Inventory` (discovery operations for workloads and targets)
- `Migration` (workload migration operations)

Elements `details1` and `details2` depend on the category and provide additional information if applicable.

The following is an example log entry that records the login action of a user with the domain account `MyDomain\John.Smith`. It has no details.

```
2017-09-02 14:14:47|Security|User logged in|MyDomain\John.Smith
```

The log file rolls over when the size of a log file reaches a specified maximum file size. The default maximum file size for the `PlateSpin.UserActivityLogging.log` file is 2 MB.

A sequential number is appended to the log file name for the rollover file. You can specify the maximum number of rollover files to keep. The default is 5.

```
PlateSpin.UserActivityLogging.log.1
PlateSpin.UserActivityLogging.log.2
PlateSpin.UserActivityLogging.log.3
```

Configuring Migrate Client User Activity Logging

The PlateSpin Migrate Client enables you to turn off or on (default) the User Activity Logging. You can configure the maximum size allowed for the User Activity Log file and how many rollover files to maintain for user activity logging.

To configure User Activity Logging:

1. Launch the PlateSpin Migrate Client.
2. Click **Tools > Options**.

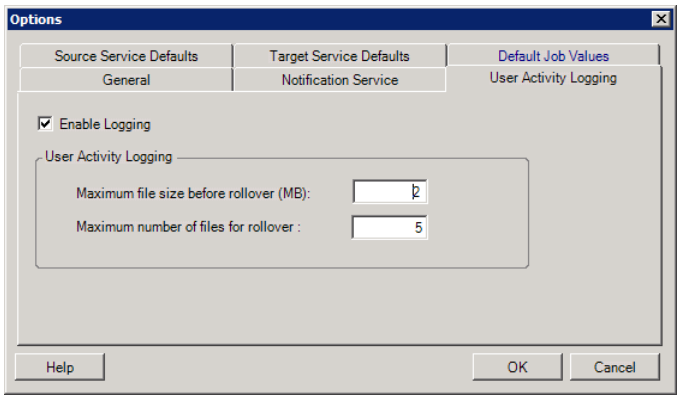

3 Click the **User Activity Logging** tab.

4 Specify the following options:

Option	Description
Enable Logging	When this option is selected, PlateSpin Migrate logs all user activities performed using the Migrate Client.
Maximum file size before rollover (MB)	When the size of a log file reaches the specified value, it is rolled over to a new file with a sequential number appended to the name.
Maximum number of files for rollover	When the number of log files reaches the specified value, the system starts overwriting the oldest file each time a rollover is performed.

5 Click **OK**.

Viewing Migrate Client User Activity Log

1 Log in to the PlateSpin Migrate Server host as the Administrator.

2 Go to the `..\PlateSpin Migrate Server\logs` directory.

3 Make a copy of the `PlateSpin.UserActivityLogging.log` file, then open the copy in a text editor.

You can also open any of its rollover files in a text editor.

7 Configuring PlateSpin Migrate Web Interface

PlateSpin Migrate Web Interface enables you to configure tags to use to track logical associates among workloads. In addition, you can control screen refresh rates for several pages. These capabilities are available only for migration jobs configured and executed by using the Migrate Web Interface. Use the information in this section to configure your Migrate Web Interface.

Managing Security Groups and Workload Permissions

PlateSpin Migrate Web Interface provides a granular application-level access mechanism that permits only specific users to carry out workload migration tasks on specified workloads. This is accomplished by setting up *security groups* and assigning users and workloads to them.

NOTE: Security group permissions apply only to migrations performed using the Web Interface.

Prerequisites for Security Groups

The default users created during Migrate installation are added to every security group you create, by default. For efficient separation of permissions, you must create additional users and assign them to appropriate Workload Migration Roles (Administrator, Power User, or Operator) with permissions that best suit their role in your organization. For more information about the workload migration roles and how to configure them, see "PlateSpin Migrate Roles" on page 89.

You must also discover workloads to be migrated by using the PlateSpin Migrate Web Interface. After discovery, you can add the workloads to an appropriate security group to be processed by its members for migration configuration and execution, according to the permissions allowed by each user's assigned roles. See "Workload Discovery in the Migrate Web Interface" on page 302.

1 Assign one or more PlateSpin Migrate users to a Workload Migration Role whose permissions best suit that role in your organization.

2 Discover workloads for migration.

Creating Security Groups for Migrate Web Interface

1 In the PlateSpin Migrate Web Interface, click **Settings** > **Permissions**.

2 On the Security Groups page, click **Create Security Group**.

3 In the **Security Group Name** field, specify a name for the security group.

4 (Optional) Click **Add Users** to select the users you want to grant access to this security group and click **OK**.

A PlateSpin Migrate user you recently added to the PlateSpin Server host might not immediately list in the user interface. To list such newly added users, click **Refresh User Accounts**.

5 (Optional) In the Migrate Web Interface, add workloads to PlateSpin Migrate that you want to add to the security group.

See "Discovering Details for Source Workloads" on page 301.

6 (Optional) Click **Assign Workloads**, select the workloads you want to include in this group, then click **OK**.

Only the users who are members of this security group have access to these workloads.

7 Click **Create** to add the new group to the security groups list on the Security Groups page.

Modifying Security Group Members or Workloads

1 In the Migrate Web Interface, select **Settings** > **Permissions**.

2 On the Security Groups page, click the security group name, then edit the group information as required:

- **Add Users**
- **Remove** assigned users

  You cannot remove the default users who were created during Migrate installation.
- **Refresh User Accounts**
- **Assign Workloads**
- **Remove** assigned workloads

3 Click **Save**.

Deleting a Security Group

1 In the Migrate Web Interface, select **Settings** > **Permissions**.

2 On the Security Groups page, click **Delete** next to the name of the security group you want to delete.

You cannot delete the default All Workloads security group with the default Migrate users.

3 Click **OK** to confirm the deletion.

Managing Workload Tags

In the PlateSpin Migrate Web Interface, the Workloads page might display a long list of workloads. Searching through these workloads to manage operations for similar workloads can be time-consuming. To overcome this issue, you can create tags for various workload categories, departments, or other logical associations appropriate to your environment. The tags you create can be associated with any workload that you manage in the Web Interface.

Creating a Workload Tag

The Workload Tags page (**Settings** > **Workload Tags**) displays all the available tags. You can create new tags and edit or delete any existing tags.

To create workload tags:

1. In the Migrate Web Interface, click **Settings** > **Workload Tags**, then click **Create Workload Tag**.
2. On the Workload Tag Creation page, specify a tag name (25-character limit) and select a color to associate with the tag.
3. Click **Save** to list the tag on the Workload Tags page.

Using Workload Tags

After you create tags, they are available on the Edit Target Details page where you can associate a tag to the appropriate workloads. Use the Tags column on the Workloads view to visually group similar workloads so that you can easily manage operations on these workloads. For information about associating tags with workloads, see "Using Tags to Track Logical Associations of Workloads" on page 309.

Modifying a Workload Tag

You can modify the name or color associated with a workload tag. Its associations with workloads is not affected.

To modify a workload tag:

1. In the Migrate Web Interface, click **Settings** > **Workload Tags**.
2. On the Create Workload Tag page, specify a different tag name or color for the tag.
3. Click **Save** to list the tag on the Workload Tags page.

Deleting a Workload Tag

You can delete tags when you no longer need it, such as when the logically associated workloads have been successfully cut over and the migration jobs cleaned up. You can also edit the migration configuration to remove or change the tags associated with workloads. You cannot delete a tag if it is associated with any workload in the list.

To delete a workload tag:

1. In the Migrate Web Interface, click **Settings** > **Workload Tags**.
2. Locate the tag of interest, then click Delete next to the tag name.
3. Click **OK** to confirm the delete.

Configuring the Refresh Rates for PlateSpin Migrate Web Interface

Several pages in the PlateSpin Migrate Web Interface have configurable refresh intervals, as shown in Table 7-1. You can modify the interval setting to meet the needs of your PlateSpin environment.

Table 7-1 *Web Interface Default Refresh Intervals*

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

1. Open the following file in a text editor:

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

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

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

3. Save the file.

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

Customizing the UI for PlateSpin Migrate Web Interface

You can modify the appearance of PlateSpin Migrate Web Interface to match the look and feel of your corporate identity. You can modify colors, logo, and product name. For more information, see Appendix B, “Rebranding the UI for PlateSpin Migrate Web Interface,” on page 137.

B Rebranding the UI for PlateSpin Migrate Web Interface

You can modify the appearance of PlateSpin Migrate Web Interface to match the look and feel of your corporate identity. You can modify colors, logo, and product name. You can even eliminate the links to **About** tab and **Help** tab in the product interface. Use the information in this section to rebrand elements in the Migrate Web Interface.

Rebranding the UI Using PlateSpin Configuration Parameters

You can change the look and feel of the Web Interface to match the proprietary look of your organization websites.

To customize the branding of the Web Interface, modify the configurable UI elements of your PlateSpin Server host:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`
2. Locate the required PlateSpin Server Configuration parameter and click **Edit** to change its value.
3. Click **Save**.

   After you modify the settings in the configuration tool, it might take up to 30 seconds for the change to take reflect on the interface. You need not reboot or restart the services.

The following sections provide information about the configurable elements in the UI for PlateSpin Migrate Web Interface.

About Configurable UI Elements for PlateSpin Migrate Web Interface

The look and feel of the PlateSpin Migrate Web Interface is consistent across the various pages. The illustration in Figure B-1 of the PlateSpin Migrate Dashboard identifies the elements that you can modify in the Web Interface UI with numbered callouts. For information about the related parameter for each element, see "Modifying PlateSpin Configuration Settings for Configurable UI Elements" on page 138.

Figure B-1 *Configurable UI Elements in the PlateSpin Migrate Web Interface*

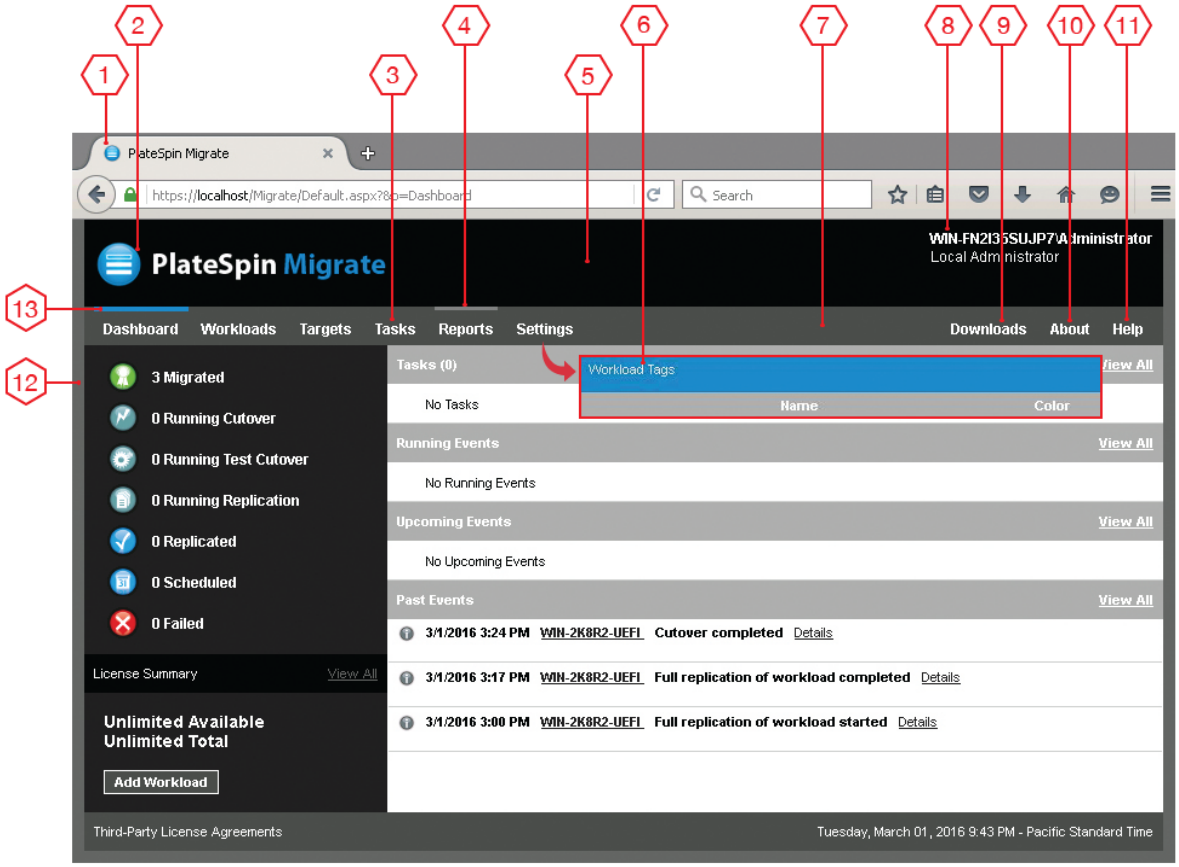

Modifying PlateSpin Configuration Settings for Configurable UI Elements

Table B-1 provides information about the setting you must use to modify the corresponding interface element. The ID column in the table lists the ID of the interface element identified in Figure B-1 provided in “About Configurable UI Elements for PlateSpin Migrate Web Interface” on page 137.

Table B-1 *Parameters for Configurable UI Elements in the PlateSpin Migrate Web Interface*

ID	Setting Name and Description	Default Value

1	`WebUIFaviconUrl`
	Location of a valid `.ico` graphic file. Specify one of the following:
	◆ A valid URL to the appropriate `.ico` file on a different machine.
	For example: https://myserver.example.com/dir1/dir2/icons/mycompany_favicon.ico
	◆ A relative path below the root of the local web server where you have uploaded the appropriate `.ico` file.
	For example, if you create a path called `mycompany\images\icons` at the root of the web server to store your custom icon graphics:
	`~/mycompany/images/icons/mycompany_favicon.ico`
	In this example, the actual file system path that contains the file is `C:\Program Files (x86)\PlateSpin Migrate Server\PlateSpin Forge\web\mycompany\images\icons\mycompany_favicon.ico`.
		`~/doc/en/favicon.ico`[1]

2	`WebUILogoUrl`
	Location of product logo graphic file. Specify one of the following:
	◆ A valid URL to the appropriate graphics file on a different machine.
	For example: https://myserver.example.com/dir1/dir2/logos/mycompany_logo.png
	◆ A relative path below the root of the local web server where you have uploaded the appropriate graphics file.
	For example, if you create a path called `mycompany\images\logos` at the root of the web server to store your custom logo images:
	`~/mycompany/images/logos/mycompany_logo.png`
	In this example, the actual file system path that contains the file is `C:\Program Files (x86)\PlateSpin Migrate Server\PlateSpin Forge\web\mycompany\images\logos\mycompany_logo.png`.
		`~/Resources/protectLogo.png`[2]

3	`WebUISiteNavigationFontColor` Color of site navigation link font color in Web UI (RGB hex value)	`#FFFFFF`

ID	Setting Name and Description	Default Value
4	`WebUISiteNavigationLinkHoverBackgroundColor` Color of site navigation link background in hover state (RGB hex value)	#808080
5	`WebUISiteHeaderBackgroundColor` Site header background color (RGB hex value)	#000000
6	`WebUISiteAccentFontColor` Font color to display with accent color in Web UI (RGB hex value)	#FFFFFF
7	`WebUISiteNavigationBackgroundColor` Color of site navigation background in Web UI (RGB hex value)	#4D4D4D
8	`WebUISiteHeaderFontColor` Site header font color in Web UI (RGB hex value)	#FFFFFF
9	`WebUIShowDownloadsTab` Toggles the visibility of the **Downloads** tab: ◆ **True:** The **Downloads** tab is visible on the interface. ◆ **False:** The **Downloads** tab is not visible on the interface.	True
10	`WebUIShowAboutTab` Toggles the visibility of the **About** tab: ◆ **True:** The **About** tab is visible on the interface. ◆ **False:** The **About** tab is not visible on the interface.	True
11	`WebUIShowHelpTab` Toggle the visibility of the **Help** tab: ◆ **True:** The **Help** tab is visible on the interface. ◆ **False:** The **Help** tab is not visible on the interface.	True
12	`WebUISiteBackgroundColor` Site background color (RGB hex value)	#666666
13	`WebUISiteAccentColor` Accent color (RGB hex value)	#0088CE

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

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

Rebranding the Product Name in the Windows Registry

The masthead at the top of the product interface provides space for the corporate logo and the product name. To change the logo, which commonly includes the product name, see "Rebranding the UI Using PlateSpin Configuration Parameters" on page 137.

To edit or eliminate the product name in a browser tab, do the following:

1. Log in to the PlateSpin Migrate Server host as the Administrator.
2. On the PlateSpin Migrate Server host, run `regedit`.
3. In the Windows Registry Editor, navigate to the following registry key:

   `HKEY_LOCAL_MACHINE\SOFTWARE\PlateSpin\MigrateServer\ProductName`

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

   `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\PlateSpin\MigrateServer`
4. Double-click the `ProductName` key and change the **Value data** for the key as required and then click **OK**.
5. Restart the IIS Server.

III Preparing Your Migration Environment

Before you discover targets and workloads, you should prepare your target migration environment. Each section describes common deployment scenarios, required settings, and a checklist for migration to the target platform.

8 Prerequisites for Migration to Amazon Web Services

PlateSpin Migrate Web Interface supports automated migration to Amazon Web Services (AWS) environments such as AWS Global, AWS China, AWS Government. This section describes the required AWS configuration that you must prepare, such as an AWS account, before you can discover an AWS target cloud platform and configure migrations to them.

Deployment for Migration to Amazon Web Services

You can deploy a PlateSpin Migrate server on premise in your data center with the source workloads or create a Migrate server in the appropriate AWS Cloud environment: AWS Global, AWS China, or AWS Government.

NOTE: For information about cloud-to-cloud (C2C) deployment requirements for migration to Azure, see Chapter 12, "Prerequisites for Cloud-to-Cloud Migrations," on page 217.

Migrate Server Deployed on Premise

For an on-premise Migrate server deployment, a site-to-site VPN connection is required between the data center and your account in the AWS cloud. Figure 8-1 shows the location of various components in your AWS migration environment and the communications between them. See "AWS Prerequisites for Using an AWS-Based Migrate Server" on page 150.

NOTE: Figure 8-1 depicts automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443). See “Requirements for Workload Registration” on page 56.

Figure 8-1 *On-Premise Migrate Server for Automated Migration to AWS*

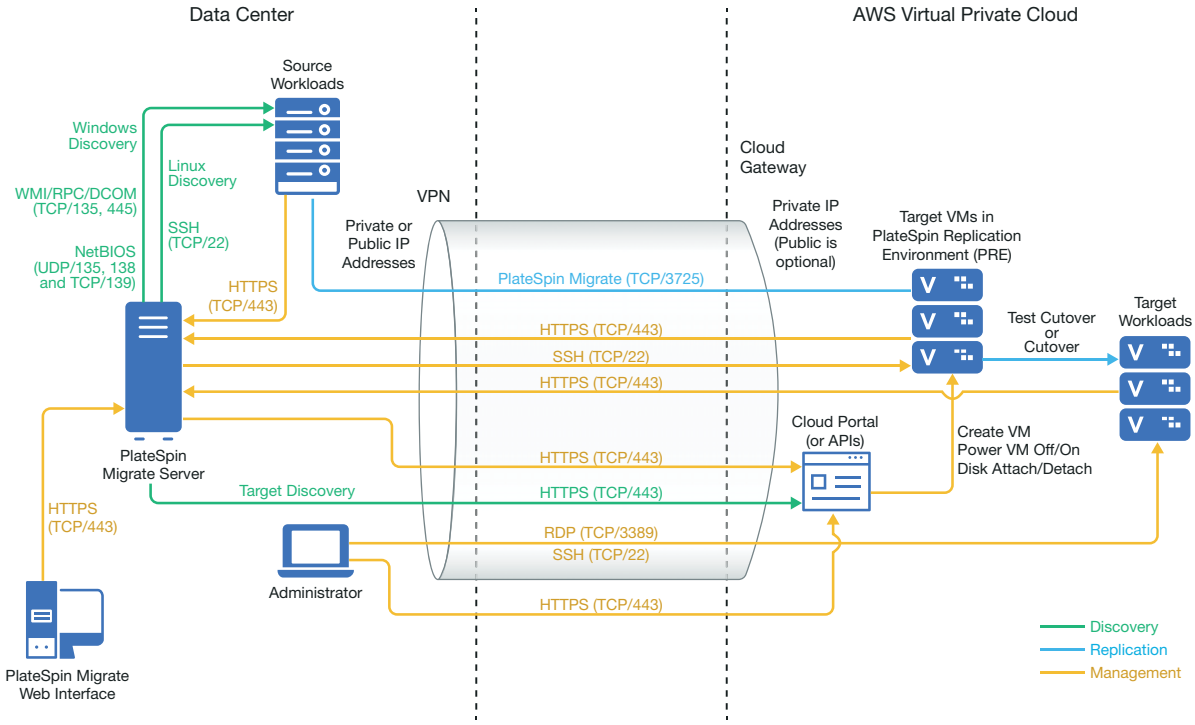

Migrate Server Deployed in AWS Cloud

AWS Quick Start lets you quickly and easily deploy Migrate server in the AWS Cloud.

Figure 8-2 shows the location of various components in your AWS migration environment without a VPN and the communications between them. See “AWS Prerequisites for Using an AWS-Based Migrate Server” on page 150.

Figure 8-2 *Cloud-Based Migrate Server for Automated Migration to AWS*

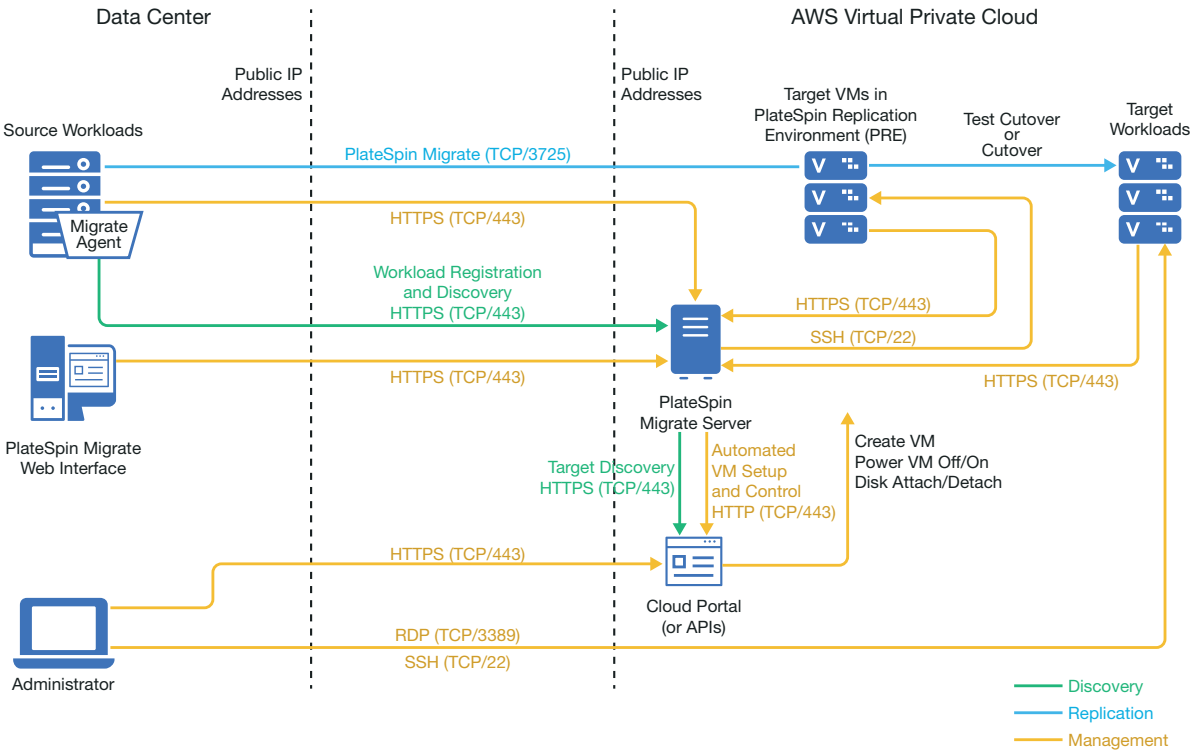

Requirements for Migrating Workloads to Amazon Web Services

Before you can migrate workloads to AWS with PlateSpin Migrate, you must set up your cloud environment. The PlateSpin Migrate server can be installed on-premise where the source workloads reside, or it can be installed in your AWS account.

- "Minimum AWS Prerequisites" on page 147
- "AWS Prerequisites for Using an On Premise Migrate Server" on page 148
- "AWS Prerequisites for Using an AWS-Based Migrate Server" on page 150

Minimum AWS Prerequisites

Before you use PlateSpin Migrate to migrate workloads to AWS, ensure that the following cloud access prerequisites are correctly configured and available:

Table 8-1 *Minimum Required Configuration for Your AWS Account*

AWS Configuration	Description
AWS Account	To create an AWS account, go to Amazon Web Services Console (http://aws.amazon.com).
AWS EC2 Subscription	PlateSpin supports only Amazon Virtual Private Cloud (VPC).

AWS Configuration	Description
Amazon Virtual Private Cloud (VPC)	Create an AWS VPC to launch AWS resources into your virtual network. See Amazon Virtual Private Cloud Documentation.
AWS user credentials	You need an AWS Identity and Access Management (IAM) user in your AWS account, with an appropriate IAM role to perform migrations into the VPC using the AWS APIs. PlateSpin Migrate provides an AWS Role Tool to enable an administrative user to create a new IAM policy based on a default policy and assign an IAM user to the policy. See "Creating an IAM Policy and Assigning an IAM User to the Policy" on page 159 Enable Programmatic Access for the IAM user to generate an access key and a secret access key. AWS Management Console Access is optional, but it can be useful for troubleshooting. See Access Keys (Access Key ID and Secret Access Key) (https://docs.aws.amazon.com/general/latest/gr/aws-sec-cred-types.html#access-keys-and-secret-access-keys). **NOTE:** We recommend that administrators regularly rotate access keys for IAM users. However, the keys must be rotated only after ensuring that no migration workflow is in progress. See "Rotating Access Keys" in the *AWS Identity and Access Management User Guide*. For information about setting up the migration user group, policy, and user, see "Creating an IAM Policy and Assigning an IAM User to the Policy" on page 159.

AWS Prerequisites for Using an On Premise Migrate Server

Before you use an on-premise PlateSpin Migrate server to migrate workloads to AWS, ensure that the following prerequisites are correctly configured and available:

- A PlateSpin Migrate license.
- PlateSpin Migrate server installed on premise in a network that can properly access the source workloads.
- A site-to-site VPN connection connecting the AWS gateway to your on-premise gateway. A public IP address for Migrate server is optional when you use a VPN.

  For information, see the following AWS resources:

  - *VPN Connections (http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpn-connections.html)*
  - *AWS-Managed VPN Connections* (http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html)
- The minimum network-related prerequisites for a successful migration are described in Table 8-2.

For information about creating and configuring an AWS Security Group, refer to *Security Groups for Your VPC* (https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_SecurityGroups.html) in the Amazon Web Services EC2 Documentation.

For detailed access and communication requirements across your migration network, see "Access and Communication Requirements across Your Migration Network" on page 53.

Table 8-2 *Port Requirements for Migrate Server on Premise for Migrations to AWS*

Location	Port	Protocol	Remarks
On-premise source workload Cloud-based target workload	TCP 443, outbound	HTTPS	The on-premise source workload and the cloud-based target workload must be able to communicate with the PlateSpin Migrate server through HTTPS (TCP/port 443) over the site-to-site VPN connection. The target workload is the replica of the source workload that will reside in AWS.
On-premise Migrate Server	TCP 443, outbound	HTTPS	The on-premise PlateSpin Migrate server must be able to communicate with the AWS API endpoint.
On-premise source workloads	TCP 22 TCP 135, 445 UDP 135, 138 and TCP 39	SSH (Linux) WMI/RPC/DCCOM NetBIOS	The PlateSpin Migrate server must be able to communicate with the source workloads on the ports that are used for discovery. See "Requirements for Discovery" on page 53 and "Discovering Details for Source Workloads" on page 301
On-premise source workloads using Migrate Agent	TCP 22 TCP 443	SSH (Linux) HTTPS	Instead of discovery, you can alternatively use the Migrate Agent utility to register source workloads with the Migrate server. See "Requirements for Workload Registration" on page 56 and "Registering Workloads and Discovering Details with Migrate Agent" on page 303.

Location	Port	Protocol	Remarks
On-premise source workload Cloud-based target workload	TCP 3725/custom	Migrate	The cloud-based target workload must be able to communicate (target to source) with the on-premise source workload across the VPN. The source workload must be able to send data to the target workload during replication across the VPN. The port number is configurable. See port 3725 in "Requirements for Migration" on page 57. If you use Migrate Agent for registration and discovery, the default direction of the replication connection must be reversed (source to target) by changing advanced settings on the Migrate server. See "Configuring the Contact Direction for the Replication Port" on page 114.
AWS Security Group for the cloud-based target workloads VPC Gateway	TCP 3725/custom, inbound and outbound	Migrate	Provide an address range covering all source workloads.
	TCP 22, inbound	SSH (Linux)	Provide the IP address of the PlateSpin Migrate server.
	TCP 3389, inbound	RDP (Windows)	Provide the IP address of the machine you plan to use to launch an RDP connect to target workloads.
	TCP 443, outbound	HTTPS	Provide the IP address of the PlateSpin Migrate server.
	TCP 123, outbound	Network Time Protocol (NTP)	Provide the IP address of your NTP server.

AWS Prerequisites for Using an AWS-Based Migrate Server

Before you use PlateSpin Migrate to migrate workloads to AWS, ensure that the following cloud access prerequisites are correctly configured and available:

- A PlateSpin Migrate license.
- Use the AWS Quick Start to deploy PlateSpin Migrate server on the AWS Cloud. See "Deploying PlateSpin Migrate Server in the Cloud" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*

NOTE: The cloud-based Migrate server does not require a site-to-site VPN connection between your local data center and AWS Portal. When no VPN is provided between the source network and the cloud-based Migrate server, you can use Migrate Agent to register workloads with the cloud-based Migrate server using secure communications over the public Internet. Internet access and public IP addresses are required. For deployment information, see Figure 8-2, "Cloud-Based Migrate Server for Automated Migration to AWS," on page 147.

- Configure migrations to AWS with a public IP address for the replication network.
- (For non-VPN setup) In the PlateSpin Configuration settings on the Migrate server, ensure that the **SourceListensForConnection** parameter is set to False, which is the default value. See "Configuring the Contact Direction for the Replication Port" on page 114.
- Allocate a Elastic IP address for the Migrate server to ensure that the IP address does not change when the server is restarted.

NOTE: A change in IP address on the PlateSpin Server breaks the heartbeat communications with source workloads.

- Ensure that workloads can reach the public IP address for Migrate server. Set the **AlternateServerAddress** parameter to the Migrate server's public IP address on the PlateSpin Configuration page. See "Configuring Alternate IP Addresses for PlateSpin Server" on page 111.
- Install the Migrate Agent on the source workload, then register the workload with the cloud-based PlateSpin Migrate server. See "Requirements for Workload Registration" on page 56 and "Registering Workloads and Discovering Details with Migrate Agent" on page 303.

  To download the Migrate Agent, launch the PlateSpin Migrate Web Interface and click the **Downloads** tab. For information about installing and using the Migrate Agent, see "Migrate Agent Utility" on page 359.
- The minimum network-related prerequisites for a successful migration when the Migrate Server is in Amazon Web Services are described in Table 8-3.

  For information about creating and configuring an AWS Security Group, refer to *Security Groups for Your VPC* (https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_SecurityGroups.html) in the Amazon Web Services EC2 Documentation.

Table 8-3 *Port Requirements for Migrate Server in AWS*

Location	Port	Protocol	Remarks
Source workload Network firewall	TCP 443, outbound	HTTPS	The source workload must be able to register (using the Migrate Agent utility) and communicate with the cloud-based PlateSpin Migrate server through HTTPS (TCP/port 443). The PlateSpin Migrate Server uses secure SSL for communications with the workloads you want to migrate.

Location	Port	Protocol	Remarks
Source workload Network firewall	TCP 3725/custom, outbound	Migrate	The source workload must be able to connect to the cloud-based target workload on TCP port 3725. The PlateSpin Migrate Server uses secure SSL for communications with the workloads you want to migrate. Port 3725 is the default port number for data transfer. By default, the data transfer is initiated from the target workload to the source workload. The port number and direction for initiating the connection are configurable. For information about changing the default port setting, see port 3725 in "Requirements for Migration" on page 57.
AWS Security Group for the Migrate Server in AWS	TCP 3725/custom, inbound and outbound	Migrate	Provide an address range covering all source workloads.
	TCP 22, outbound	SSH (Linux)	This port allows outbound communications from the Migrate server to Linux workloads. Provide the IP address of the PlateSpin Migrate server.
	TCP 3389, inbound	RDP (Windows)	Allow inbound connections in the AWS Security Group for the cloud-based Migrate server. Provide the IP address of the machine you plan to use to launch an RDP connect to target workloads.
	TCP 443, inbound	HTTPS	Provide the IP address of the PlateSpin Migrate server.
	TCP 123, outbound	Network Time Protocol (NTP)	Add this port setting to the security group if you are using an NTP service outside the virtual network where you deploy the Migrate server.

Planning for Migrating Workloads to Amazon Web Services

PlateSpin Migrate allows you to use the PlateSpin Migrate Web Interface to migrate Windows and Linux workloads to AWS. For a list of supported workloads, see "Supported Workloads For Migration to Amazon Web Services" on page 27.

NOTE: Migration of Windows Cluster workloads to AWS is not supported.

Target AWS Cloud

- Supports the following targets platforms:
  - AWS China
  - AWS Global
  - AWS Government

AWS Subscription

- PlateSpin supports only Amazon Virtual Private Cloud (VPC).

PlateSpin Server Host

- 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 access error.

OS License for Target Workload

- You select the OS licensing model on the target workload when you configure migration for the workload. See "**License Type" on page 435**.
- For cutover of workloads to AWS, PlateSpin Migrate selects an AMI based on the target workload OS licensing model that you configure in the migration job. See "Understanding PlateSpin AMIs Used for Replication and Cutover of Workloads" on page 157.

Target Workload

Consider the following guidelines before you use the PlateSpin Migrate Web Interface to migrate workloads to AWS:

- **Migrate Web Interface**

  You can use only the PlateSpin Migrate Web Interface to migrate the workloads to AWS. The PlateSpin Migrate Client does not support migration of workloads to AWS.

- **UEFI Source Workloads**

  Windows and Linux UEFI workloads are migrated as BIOS workloads.

- **Multiple NICs**

  Migration of workloads with multiple NICs to AWS is supported for Windows workloads up to the number of NICs supported by the AWS Instance size.

  NOTE: When you configure migration for a workload, PlateSpin Migrate lets you specify four IP addresses per NIC on the Windows workload by default. However, you can configure the number of IP addresses that you want to specify per NIC on the workload. See "Configuring the Maximum Number of IP Addresses Allowed for Each Network Interface on Target Workloads" on page 111.

- **Maximum Number of Data Disks**

  PlateSpin Migrate supports AWS target instances with up to 26 disks (EBS volumes) for Windows and 40 disks (EBS volumes) for Linux with each disk not exceeding 15 file system volumes.

- **Disk Size on the Target VM**

  The size of the disk created on the AWS instance is the size of the source disk partition plus about 1 GB.

- **About AWS Instance Types and Sizes**
  - Migrate recommends an AWS instance size that meets or exceeds the source workload's settings for cores, memory, volumes, 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 AWS region.
  - Migrate neither discovers or migrates ephemeral disks.
  - Migrate supports AWS instance types based on x86 and x86_64 processor architectures.
  - AWS requires a minimum of Microsoft .NET Framework 4.5 for migration of Windows workloads to Nitro System based VM instances. To use Nitro System based instances, ensure that .NET 4.5 Framework or later is installed on the source Windows workload prior to Migrate discovery.
  - Migration of Windows Domain Controller servers to AWS requires Nitro System based instance types such as c5, m5, r5, t3, c5d, m5d, and z1d. Ensure that you select a Nitro System based instance when you configure the migration.

    NOTE: Nitro System based Instance types support Windows Server 2008 R2 and later. Therefore, migration of Windows Domain Controller servers to AWS is supported only for Windows Server 2008 R2 and later.

  - For each release, PlateSpin Migrate adds support for new cloud instances and instance families that AWS has released up to a stipulated date prior to the PlateSpin Migrate release. After that date, any new AWS instances are not available by default in the Cloud Instance Size list. You can add definitions to the PlateSpin Server for new AWS Nitro Based instances and instance families to make the instances available in the AWS Cloud Instance Sizes list. See “Displaying the AWS Nitro System Based Cloud Instance Sizes That AWS Recently Supported in the AWS Cloud Instance Sizes list” on page 163.

Deploying a Migrate Server Image in AWS

Use the AWS Quick Start to quickly and easily deploy PlateSpin Migrate server on the AWS Cloud. See "Deploying PlateSpin Migrate Server in the Cloud" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

Using Enhanced Networking with ENA on Linux Distributions

Using Enhanced networking with Elastic Network Adapter (ENA) capability on a Linux workload requires ENA drivers on the workload. RHEL 7.4 and later kernel versions have built-in drivers for ENA.

PlateSpin Migrate provides precompiled ENA Linux kernel drivers for the following versions.

Linux Distribution	Precompiled Driver
RHEL 7.0	◆ `3.10.0-123.20.1.el7.x86_64-x86_64` ◆ `3.10.0-123.el7.x86_64-x86_64`
RHEL 7.1	◆ `3.10.0-229.el7.x86_64-x86_64`
RHEL 7.2	◆ `3.10.0-327.el7.x86_64-x86_64`
RHEL 7.3	◆ `3.10.0-514.el7.x86_64-x86_64`

To create custom ENA drivers for AWS enhanced networking support, follow the steps documented in the KB Article 7023023 (https://support.microfocus.com/kb/doc.php?id=7023023).

Configuring Advanced PlateSpin Settings for AWS

Some aspects of your PlateSpin Server behavior are controlled by configuration parameters that you set on a PlateSpin Configuration web page residing on your PlateSpin Server host at `https://Your_PlateSpin_Server/PlateSpinConfiguration/`.

To edit the value of the configuration parameters:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at: `https://Your_PlateSpin_Server/PlateSpinConfiguration/`.
2. Search the parameter you want to edit and make the required changes.
3. Save your settings and exit the page.

Advanced PlateSpin settings for AWS apply globally to all AWS target platforms that you define on the Migrate server.

- "Configuring the AWS Region Price List Endpoint To Be Used for Discovering Supported AWS Instance Types" on page 156
- "Configuring Target Instance Logging With Key Pair or Source Credentials" on page 156

- "Configuring PlateSpin Migrate Server to Use Public IP Address for AWS Migrations" on page 156
- "Configuring OS License Activation on Windows Targets Migrated to AWS" on page 157
- "Configuring the Number of Connection Attempts for a SSH Session from AWS Cloud-Based Migrate Server to Target VMs in PlateSpin Replication Environment" on page 157

Configuring the AWS Region Price List Endpoint To Be Used for Discovering Supported AWS Instance Types

By default, PlateSpin Migrate Server is preconfigured to use the AWS price list endpoint in the `us-east-1` region for discovering the AWS supported instance types. However, if the instance type that you want to use is not listed in the price list endpoint of the configured region, set the value of **AWSPriceListRegion** parameter to the name of region that has a price list endpoint listing the desired instance type.

Configuring Target Instance Logging With Key Pair or Source Credentials

By default, PlateSpin Migrate Server allows you to log in to an AWS target instance only by using the key pair configured in the migration job. PlateSpin Migrate controls this behavior by using the **AWSEnableSourceCredentialsForLinuxWithKeypair** parameter that is set to `False` by default. To enable logging into AWS Linux target instances either by using the key pair configured in the migration job or the source credentials, set the **AWSEnableSourceCredentialsForLinuxWithKeypair** parameter to `True`.

Configuring PlateSpin Migrate Server to Use Public IP Address for AWS Migrations

By default, PlateSpin Migrate Server is preconfigured to allow private IP addresses for communications during migrations to AWS. If the source workload cannot connect to the private IP address of the AWS target, then you require a public IP address for communications during migrations to AWS. To ensure that only public IP is used during migration:

- Set the value of the **UseOnlyPublicIPForAWS** parameter as `True`.
- Ensure that the value of the **SourceListensForConnection** parameter setting is set to False, which is the default value. See "Configuring the Contact Direction for the Replication Port" on page 114.
- Set the **AlternateServerAddress** parameter to the Migrate server's public IP address. See "Configuring Alternate IP Addresses for PlateSpin Server" on page 111.

Configuring OS License Activation on Windows Targets Migrated to AWS

PlateSpin Migrate provides the following parameters to configure KMS server for Windows OS activation on the target workload:

- **AWSKMSServers**: This parameter lets you set the AWS KMS Server information that Windows instances use for activation. The target KMS Server should be in the same AWS Region where the Windows instance is running.
- **KMSClientSetupKeys**: This parameter lists the commonly used OS version-based Microsoft KMS client setup keys that are used for activating Windows through KMS server. If the key for a particular OS is not listed, you can add an entry in the following format:

  `OperatingSystemTypeandBranding="Microsoft provided KMS Key"`

  Example: For a Windows server with OS type as `Windows 2016` and branding as `Standard Server`, the format is `Windows2016StandardServer="WC2BQ-8NRM3-FDDYY-2BFGV-KHKQY"`

Configuring the Number of Connection Attempts for a SSH Session from AWS Cloud-Based Migrate Server to Target VMs in PlateSpin Replication Environment

By default, PlateSpin Migrate Server attempts to establish a SSH connection with the target VMs in PlateSpin Replication Environment 40 times before timing out the session. PlateSpin Migrate controls this behavior by using the **AwsMaxRepEnvConnectionAttempts** parameter that is set to 40 by default. If the session times out, edit the value of the **AwsMaxRepEnvConnectionAttempts** parameter to increase the connection attempts to more than the set value.

Understanding PlateSpin AMIs Used for Replication and Cutover of Workloads

PlateSpin Migrate leverages the following PlateSpin AMIs uploaded in the Community AMI section of Amazon Web Services Console to perform replications and cutover of workloads to AWS. For cutover of workloads to AWS, PlateSpin Migrate selects an AMI based on the target workload OS licensing model that you configure in the migration job.

The AMIs are listed only for your information and you are not required to perform any action with these AMIs.

AMI Name	Description
PlateSpin Replication Environment	Used for the following: ◆ Replication of all 32-bit Windows and Linux workloads. ◆ Cutover of all Linux workloads. AWS allows you to bring your own license (BYOL) for all Linux workloads and does not bill you for the OS license on the target workload.
PlateSpin Replication Environment (64-bit Replications)	Used for the replications of 64-bit Windows and Linux workloads.
PlateSpin Template - Windows	Used during the cutover of the Windows workloads for which AWS manages the Microsoft software licensing compliance on the target workload and bills you for the license.
PlateSpin Template - Windows (BYOL)	Used during the cutover of the Windows workloads for which AWS allows you to bring your own license (BYOL) that you have already purchased from Microsoft and does not bill you for the license. You are solely responsible for complying with Microsoft licensing.

AWS Networking Guidelines

Consider the following guidelines when you are migrate workloads to AWS:

- "Private and Public IP Addresses for Workloads Connected on an AWS VPN" on page 158

Private and Public IP Addresses for Workloads Connected on an AWS VPN

Each AWS VM has both a public IP address and a private IP address for communications from machines outside the AWS environment. AWS automatically associates these IP addresses with the primary network interface for the VM.

AWS provides public IP addresses for the target instance only in case of workloads with single NIC. For workloads with multiple NICs, AWS provides only private IP addresses for the target instance and so you can connect to the target instance using only the private IP addresses. If the `UseOnlyPublicIPForAWS` PlateSpin Configuration parameter is set to `True` and you choose to migrate a source workload with multiple NICs, then you must include only one NIC for migration when you configure the migration job.

You can use the Microsoft Remote Desktop client or SSH to remotely connect to the AWS VM. Specify the IP address as follows:

- **Private IP address:** Use the VM's private IP address if your machine is part of the address space for the AWS VPN.

- **Public IP address:** Use the VM's public IP address if your machine is not part of the address space for the AWS VPN. A public IP address is not set on the target workload that has multiple NICs.

Creating an IAM Policy and Assigning an IAM User to the Policy

To migrate workloads to AWS with PlateSpin Migrate, you require an AWS Identity and Access Management (IAM) user in your AWS account with an appropriate IAM role and the required permissions to perform migrations in to the AWS VPC. You also need the AWS Access Key and AWS Secret Access Key for this user.

You can create a new IAM policy by using one of the following:

- **PlateSpin AWS Role Tool:** See "Using the AWS Role Tool to Create a New IAM Policy" on page 159.
- **AWS Management Console:** See "Using the AWS Management Console to Create an IAM Policy" on page 160.

Using the AWS Role Tool to Create a New IAM Policy

PlateSpin Migrate provides an AWS Role Tool (`AWSRoleTool.exe`) to enable an administrative user to create a new IAM policy based on a default policy (`PolicyJSON.txt`) that PlateSpin Migrate defines and assign an IAM user (either existing user or new user) to the policy. The PlateSpin Migrate AWS Role Tool (`AWSRoleTool.exe`) is included in the `Migrate-Install-folder\PlateSpin Migrate Server\bin\AWSRolesTool` directory.

By default, the `PolicyJSON.txt` file that PlateSpin Migrate defines contain the minimum permissions required for an IAM user to migrate workloads to AWS with PlateSpin Migrate. For information about the minimum permissions defined for an IAM user in the default policy, see "Defining Minimum Permissions for an IAM User" on page 160.

When you use the AWS Role Tool to create a new policy, the new policy is created as a replica of this default policy and has all the permissions that are listed in the default policy. However, you can choose to create a new policy with modified permissions than what is listed in the default policy. To create a new policy with modified permissions, you must edit the `PolicyJSON.txt` file to list only those permissions that you want to list in the new policy and then create the policy.

NOTE: If you have edited the `PolicyJSON.txt` file and want to restore the default policy that PlateSpin Migrate defines, delete the edited `PolicyJSON.txt` file. The `PolicyJSON.txt` file is recreated with the default permissions in the `Migrate-Install-folder\PlateSpin Migrate Server\bin\AWSRolesTool` directory when the AWS role tool runs.

---

1. Log in as an Administrator on your PlateSpin Migrate Server host.
2. Open a command prompt and navigate to the location that has the AWS role tool, and run the following command:

   ```
   AWSRoleTool.exe
   ```

NOTE: If the default policy (`PolicyJSON.txt`) is not available in the `Migrate-Install-folder\PlateSpin Migrate Server\bin\AWSRolesTool` directory, the tool recreates the `PolicyJSON.txt` file with the default permissions that PlateSpin Migrate recommends.

3. Enter the AWS Access Key and AWS Secret Access Key of an AWS user who has permissions to create IAM policy and users.
4. Enter a name for the AWS policy you want to create.
5. Enter the name of a new or an existing user to whom you want to assign this policy. The tool creates the new policy as a replica of the `PolicyJSON.txt` file, assigns the policy to the specified user, and provides the Access Key and Secret Key credentials for the user.
6. You can choose to save the credentials to a file or display the credentials in the command prompt:
   - To save the credentials to a file, enter y. The path of the file that contains the credentials is displayed.
   - To display the credentials in the command prompt, enter n and take a note of the displayed credentials.
7. (Optional) To restore the default policy that PlateSpin Migrate defines, delete the edited `PolicyJSON.txt` file and run the AWS Role Tool to recreate the `PolicyJSON.txt` file with the default permissions.

Using the AWS Management Console to Create an IAM Policy

You can use the AWS Management Console to create or edit an IAM policy and define user permissions by assigning the user to a policy. See Creating IAM Policies (https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_create.html).

PlateSpin Migrate provides a default policy (`PolicyJSON.txt`)that contain the minimum permissions required for an IAM user to migrate workloads to AWS with PlateSpin Migrate. For information about the minimum permissions defined for an IAM user in the default policy file, see "Defining Minimum Permissions for an IAM User" on page 160.

You can use the AWS Management Console to create a new policy with the recommended permissions included in this default policy.

Defining Minimum Permissions for an IAM User

PlateSpin Migrate provides a `PolicyJSON.txt` file that by default contains the minimum permissions required for an IAM user to migrate workloads to AWS with PlateSpin Migrate. When you use the AWS Role Tool to create a new policy, the new policy is created as a replica of this default policy and has all the permissions that are listed in the default policy.

The contents of the `PolicyJSON.txt` file is as follows:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "ec2:TerminateInstances",
                "ec2:DeleteTags",
                "ec2:StartInstances",
                "ec2:CreateTags",
                "ec2:RunInstances",
                "kms:DescribeKey",
                "ec2:StopInstances"
            ],
            "Resource": [
                "arn:aws:ec2:*:*:subnet/*",
                "arn:aws:ec2:*:*:instance/*",
                "arn:aws:ec2:*:*:volume/*",
                "arn:aws:ec2:*:*:security-group/*",
                "arn:aws:ec2:*:*:network-interface/*",
                "arn:aws:ec2:*::image/*",
                "arn:aws:kms:*:*:key/*"
            ]
        },
        {
            "Sid": "VisualEditor1",
            "Effect": "Allow",
            "Action": [
                "ec2:DeregisterImage",
                "ec2:DeleteSnapshot",
                "ec2:DescribeInstances",
                "ec2:DescribeRegions",
                "ec2:CreateImage",
                "iam:ListRoles",
                "ec2:DescribeSnapshots",
                "ec2:DescribePlacementGroups",
                "pricing:GetProducts",
                "ec2:DescribeSecurityGroups",
                "ec2:DescribeHosts",
                "ec2:DescribeImages",
                "ec2:DescribeAvailabilityZones",
                "ec2:DescribeVpcs",
                "kms:ListAliases",
                "ec2:DescribeVolumes",
                "ec2:DescribeAccountAttributes",
                "ec2:DescribeReservedInstances",
                "ec2:ModifyInstanceAttribute",
                "ec2:DescribeSubnets",
                "ec2:DescribeKeyPairs",
                "ec2:DescribeInstanceStatus"
            ],
            "Resource": "*"
        },
        {
```

```
            "Sid": "VisualEditor2",
            "Effect": "Allow",
            "Action": [
                "kms:Decrypt",
                "kms:Encrypt",
                "ec2:CreateVolume"
            ],
            "Resource": [
                "arn:aws:ec2:*:*:volume/*",
                "arn:aws:kms:*:*:key/*"
            ]
        },
        {
            "Sid": "VisualEditor3",
            "Effect": "Allow",
            "Action": [
                "ec2:AttachVolume",
                "kms:CreateGrant"
            ],
            "Resource": [
                "arn:aws:kms:*:*:key/*",
                "arn:aws:ec2:*:*:instance/*",
                "arn:aws:ec2:*:*:volume/*"
            ]
        },
        {
            "Sid": "VisualEditor4",
            "Effect": "Allow",
            "Action": "ec2:DetachVolume",
            "Resource": [
                "arn:aws:ec2:*:*:instance/*",
                "arn:aws:ec2:*:*:volume/*"
            ]
        },
        {
            "Sid": "VisualEditor5",
            "Effect": "Allow",
            "Action": "ec2:DeleteVolume",
            "Resource": "arn:aws:ec2:*:*:volume/*"
        },
        {
            "Sid": "VisualEditor6",
            "Effect": "Allow",
            "Action": "ec2:RunInstances",
            "Resource": [
```

```
                "arn:aws:ec2:*:*:subnet/*",
                "arn:aws:ec2:*:*:key-pair/*",
                "arn:aws:ec2:*:*:instance/*",
                "arn:aws:ec2:*::snapshot/*",
                "arn:aws:ec2:*:*:launch-template/*",
                "arn:aws:ec2:*:*:volume/*",
                "arn:aws:ec2:*:*:security-group/*",
                "arn:aws:ec2:*:*:placement-group/*",
                "arn:aws:ec2:*:*:network-interface/*",
                "arn:aws:ec2:*::image/*"
            ]
        }
    ]
}
```

Displaying the AWS Nitro System Based Cloud Instance Sizes That AWS Recently Supported in the AWS Cloud Instance Sizes list

For each release, PlateSpin Migrate adds support for new cloud instances and instance families that AWS has released up to a stipulated date prior to the PlateSpin Migrate release. After that date, any new AWS instances are not available by default in the Cloud Instance Size list option when you configure migration to AWS.

The AWS Cloud Instance Size list displays only instances that are defined in the `AWS-Instances.ini` file in the following package found on the PlateSpin Server:

`<`*`install-drive`*`>:\Program Files\PlateSpin Migrate Server\Packages\B9059B80-55E4-4B9C-B85B-EFEFC85EAAE2\1.package\`

You can add definitions to this file for new AWS Nitro Based instances and instance families to make the instances available in the AWS Cloud Instance Sizes list when you configure workloads for migration to AWS.

1. Log in to the PlateSpin Migrate Server host machine as the Administrator user.
2. Navigate to the PlateSpin Migrate installation directory, such as `C:\Program Files\PlateSpin Migrate Server\`.
3. Open the `.\Packages\B9059B80-55E4-4B9C-B85B-EFEFC85EAAE2\` folder and copy the `1.package` file out of this folder to a temporary directory.
4. Rename `1.package` with the `.zip` file extension (`1.zip`), then extract the contents of the file to the temporary directory.
5. Edit the file named `AWS-Instances.ini` to add an entry for the new instance or instance family that you want to support by using one of the following applicable formats:
   - **If the instance or the instance family supports both NVMe EBS volumes and enhanced networking:**

     ```
     [Instance name or Instance family name]
     ```

```
Device.Storage = nvme

Device.Network = ena
```

- **If the instance or the instance family supports only enhanced networking:**

```
[Instance name or Instance family name]

Device.Network = ena
```

NOTE: If the network and storage capabilities of an instance is not the same as that of the family it belongs to, you must have a separate entry for the instance and the support details listed within this entry takes precedence for the instance instead of its instance family entry.

6 Repackage the modified file as `1.package`.

Best Practices for Configuring a Migration Job to Amazon Web Services

To help prevent the failure of a migration job to AWS, you must adopt the following best practices when you configure migration jobs:

- If you use a static IP address for the network, ensure that the address is unique within the supported subnet range.
- The number of target instances running at any point of time must not exceed the instance limit applicable for your subscription.
- You must select a subnet such that the replication, run cutover, and test cutover instances are all in the same availability zone.

Checklist for Automated Migration to AWS

Task	Description
1. Prepare your AWS migration environment.	Figure 8-1, “On-Premise Migrate Server for Automated Migration to AWS,” on page 146 Figure 8-2, “Cloud-Based Migrate Server for Automated Migration to AWS,” on page 147 “Planning for Migrating Workloads to Amazon Web Services” on page 152
2. Discover target cloud platforms.	“Target Discovery in the Web Interface” on page 284
3. Discover source workloads.	“Workload Discovery in the Migrate Web Interface” on page 302 -OR- “Registering Workloads and Discovering Details with Migrate Agent” on page 303

Task	Description
4. Configure target workload migration.	“Configuring Migration of a Workload to Amazon Web Services” on page 432
5. Execute migration.	Chapter 33, “Executing Workload Migrations,” on page 537

9 Prerequisites for Migration to Microsoft Azure

PlateSpin Migrate Web Interface supports automated migration to Microsoft Azure Cloud and Azure-connected Azure Stack environments.

This section describes the required Azure configuration that you must prepare in the appropriate environment, such as an Azure account, subscriptions, and services, before you can discover Azure target cloud platforms and configure migrations to them.

Deployment for Migration to Azure

You can deploy a PlateSpin Migrate server on premise in your data center with the source workloads or in the appropriate Microsoft Azure Cloud environment.

NOTE: For information about cloud-to-cloud (C2C) deployment requirements for migration to Azure, see Chapter 12, "Prerequisites for Cloud-to-Cloud Migrations," on page 217.

Migrate Server Deployed on Premise

For an on-premise Migrate server deployment, a site-to-site VPN connection is required between the data center and your account in the Azure environment. Figure 9-1 and Figure 9-2 shows the location of various components in your Azure migration environment and the communications between them. See “Azure Prerequisites for Using an On-Premise Migrate Server” on page 175.

NOTE: You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443). See “Requirements for Workload Registration” on page 56 and “Registering Workloads and Discovering Details with Migrate Agent” on page 303.

Figure 9-1 *On-Premise Migrate Server for Automated Migration to Azure Cloud*

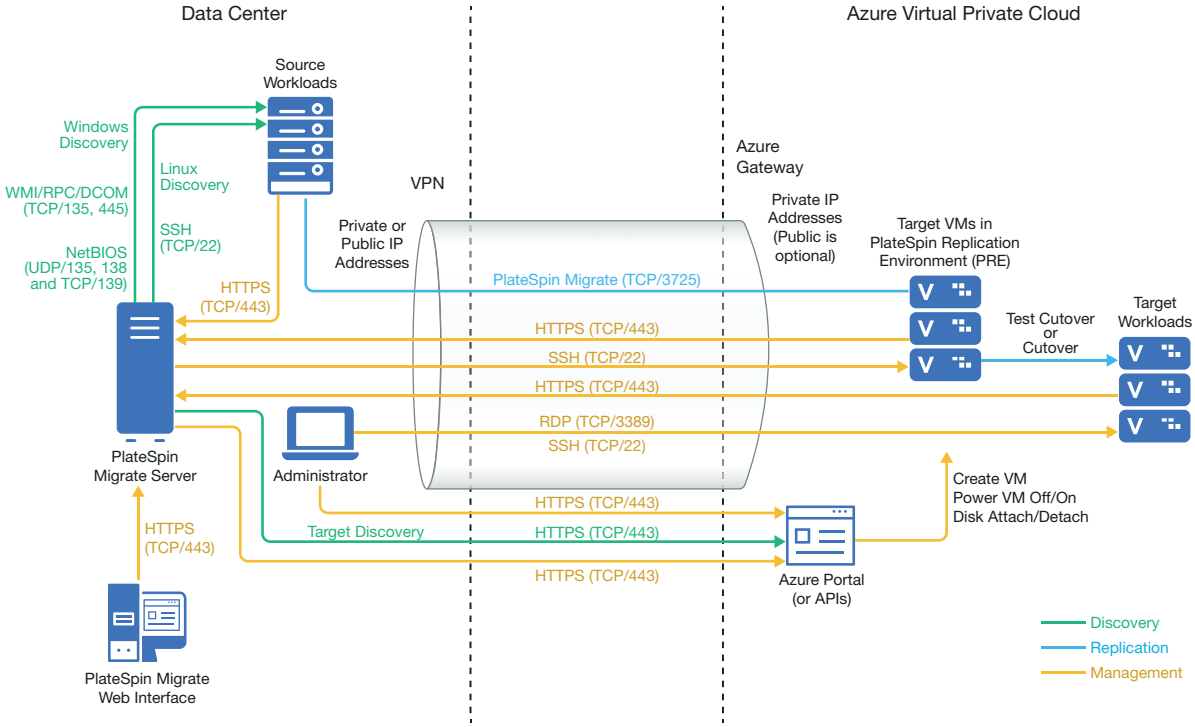

Figure 9-2 *On-Premise Migrate Server for Automated Migration to Azure Stack*

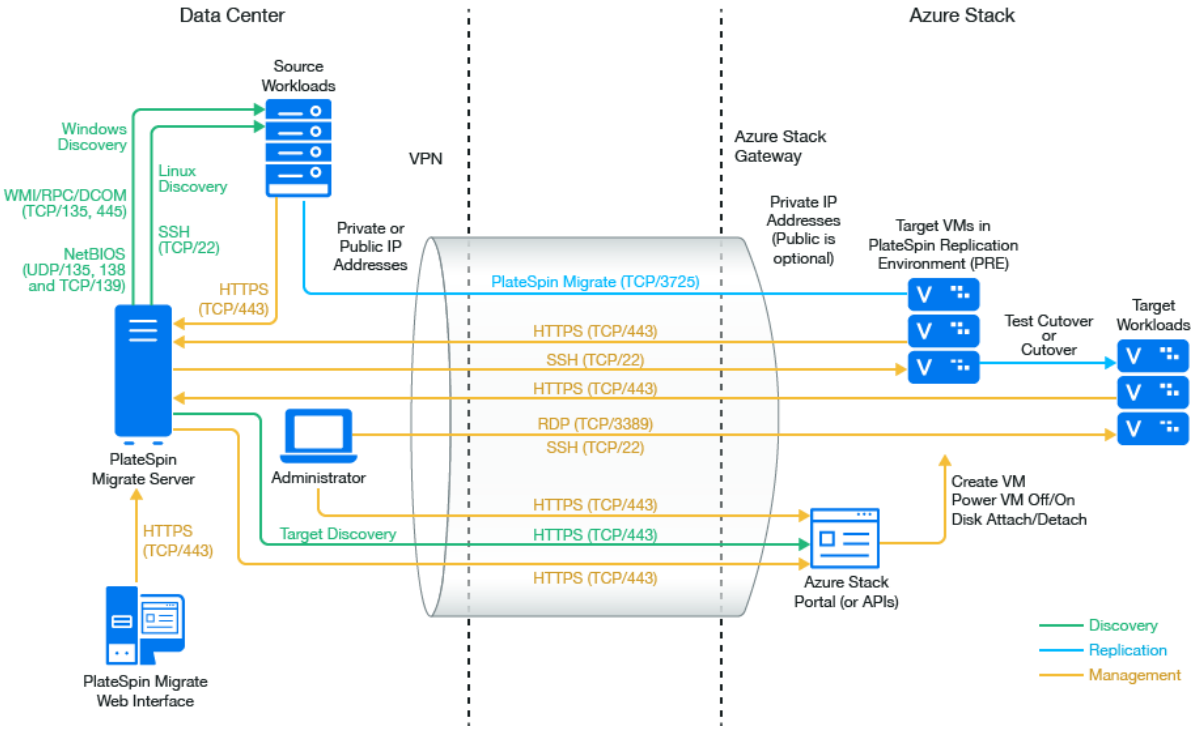

Migrate Server Deployed in Azure Cloud

For a cloud-based Migrate server deployment, the Azure Marketplace in the target Azure environment offers a PlateSpin Migrate Server image that is preconfigured to support its host IaaS environment. Figure 9-3 and Figure 9-4 shows the location of various components in your Azure migration environment and the communications between them. See “Azure Prerequisites for Using an Azure-Based Migrate Server” on page 177.

Figure 9-3 *Cloud-Based Migrate Server for Automated Migration to Azure*

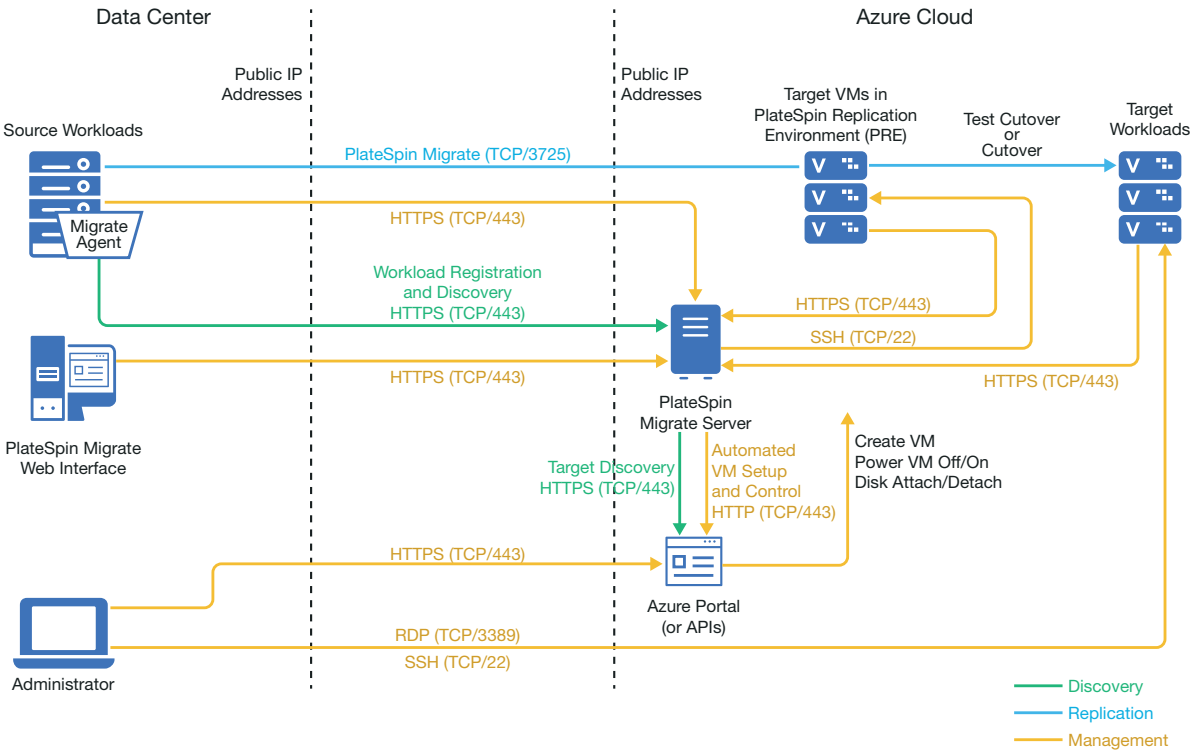

Figure 9-4 *Cloud-Based Migrate Server for Automated Migration to Azure Stack*

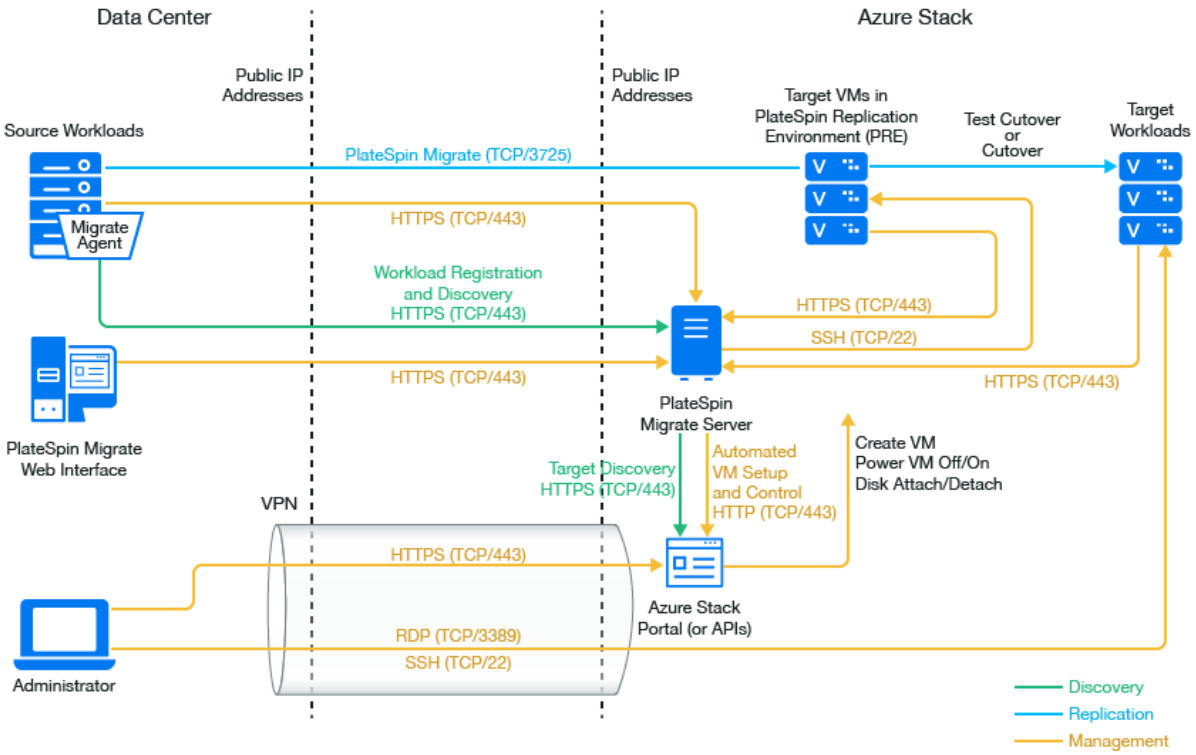

Requirements for Migrating Workloads to Azure

Based on the location of your PlateSpin Migrate server, review the following sections for information about the requirements for migrating workloads to Azure Cloud and Azure-connected Azure Stack environments. For information about Azure-connected Azure Stack environments, see Azure Stack documentation (https://docs.microsoft.com/en-us/azure-stack/?view=azs-1910).

Minimum Azure Prerequisites

PlateSpin Migrate requires the use of Microsoft Azure Resource Management for migrating workloads to the Microsoft Azure environments. For migrations to Microsoft Azure, you must prepare your Azure account, subscriptions, and services in the desired Azure environment.

Table 9-1 describes the minimum configuration you must perform in the appropriate Azure environment before you can migrate workloads to Azure.

Table 9-1 *Minimum Required Configuration for Your Azure Account*

Azure Configuration	Description
Microsoft Azure Account	Use the Azure Global Portal (https://portal.azure.com/) to create an account in the Azure Global or Azure Stack environment where you will migrate workloads. An administrator on the account is required to perform the Application setup, to enable PRE programmatic access, and to create a Contributor user that is to be used by Migrate.
Azure Subscription ID	The ID for the Azure Subscription in the specified Azure account that you want to bill for Azure-related costs for migrations using PlateSpin Migrate. An account can have multiple subscriptions. You must provide the Subscription ID when you add a Microsoft Azure Location as a migration target in PlateSpin Migrate. To view a Subscription ID, log in to the Azure portal with your Azure user account that you use to manage your subscriptions, then go to the Subscriptions page and select the subscription.

Azure Configuration	Description
Contributor user for the subscription created in Azure Active Directory	A special-purpose user identity for PlateSpin Migrate that you create in Azure Active Directory. You add a Contributor role to the user account for the specified subscription. Using this Contributor user only for Migrate helps to uniquely identify actions performed by Migrate in Azure for the subscription. In Migrate, you use the Contributor user credentials to add Azure as a target platform. Migrate uses the credentials for this user when it accesses the Migrate Azure API through the related subscription. See “Configuring a Contributor User for PlateSpin Migrate to Use” on page 188.
Application (client) ID	An ID that represents PlateSpin Migrate as it makes use of the Microsoft Azure API when it replicates or migrates workloads on your behalf to VMs in a target Azure location for a subscription in your Azure account. See “Configuring an Application in Azure to Represent PlateSpin Migrate” on page 182. **NOTE:** For migrations to Azure Stack, you must not create an Azure application ID for PlateSpin Migrate. Instead, use the following ID as the Azure application ID for PlateSpin Migrate: `872cd9fa-d31f-45e0-9eab-6e460a02d1f1`
Azure Virtual Network and Subnet	You must create least one Virtual Network with a Subnet in the specified Subscription. If you have an site-to-site VPN set up, the subnet must be different than the default Gateway Subnet. Network resources are never created automatically by PlateSpin Migrate, so they always must be set up manually in advance. For instructions, refer to Azure documentation.
Azure Storage Account	For each target VM, PlateSpin Migrate supports using Azure Managed Disks or using Azure Storage Accounts and unmanaged disks. The setting applies at the VM level. **NOTE:** At least one Azure storage account is required when you use unmanaged disks. Managed Disks do not require a storage account. See “Prerequisites for Azure Storage” on page 173.
Azure Availability Sets	PlateSpin Migrate supports using pre-defined Azure Availability Sets for target VMs. See “Prerequisites for Using Azure Availability Sets” on page 174.
Azure tags	PlateSpin Migrate provides the ability to define tags for migrations to Azure that it applies to objects it creates in the target Azure subscription. See “Using Azure Cloud Tags for Azure Migrations” on page 194.

Prerequisites for Azure Storage

For each target VM, PlateSpin Migrate supports using Azure Managed Disks or using Azure Storage Accounts and unmanaged disks.

- "Azure Managed Disks" on page 173
- "Azure Storage Accounts and Unmanaged Disks" on page 174

Azure Managed Disks

Azure Managed Disks is an Azure service that provides replicas of your data to help ensure persistence of your data and high tolerance against failures. Azure automatically creates and manages the placement of the disks in the target Azure location. You do not need to specify a storage account or location in a storage account. The specified **Cloud Instance Size** for the target VM determines how many data disks you can attach to the VM and the type of storage you can use to host the managed disks. The Azure Managed Disks setting applies to all disks for the target VM.

Using Azure Managed Disks is optional. It is enabled by default in keeping with Azure default VM settings.

NOTE

- To verify the availability in the target Azure location of Managed Disks service, refer to the *Microsoft Azure Products Available by Region*.
- For information about disk types, sizes, and performance characteristics for Azure Manage Disks, refer to *Introduction to Azure Managed Disks* in Microsoft Azure Documentation.
- Azure automatically and transparently provides Azure Storage Service Encryption for all storage, including Azure Managed Disks.
- PlateSpin Migrate does not support configuration of Azure Disk Encryption for the target workload.

  Azure Disk Encryption provides OS-based encryption for OS and data volumes at rest by using BitLocker for Windows VMs or DM-Crypt for Linux VMs. After the cutover is complete, you can use Azure tools to enable Azure Disk Encryption for the cutover workload.

When you configure migrations to Azure, specify one of the following storage types for hosting the managed disks for the target VM:

- Standard HDD (hard disk drives)
- Standard SSD (solid state drives)

  NOTE: Standard SSD option is not applicable for migrations to Azure Stack.

- Premium SSD

A storage type of Standard HDD allows you to choose VMs with HDD or SDD storage, whereas Standard SSD and Premium SSD are restricted to VMs that support the appropriate level of SSD storage. For information about disk performance characteristics, refer to *What Disk Types Are Available in Azure?* in the Microsoft Azure Documentation.

You can use Azure Managed Disks in combination with a predefined Azure Availability Set that has been configured to use managed disks.

Azure Storage Accounts and Unmanaged Disks

If you do not use Managed Disks for a VM, the VM disks are *unmanaged* by the Azure Managed Disks server. The VM disks will use the Azure page blob type of general-purpose storage, which can run on HDD (Standard) or SSD (Premium) storage media. You can use Azure General Purpose Storage V1 or V2, according to your needs.

If you plan to use unmanaged disks, you must create at least one Azure Storage account in the specified Subscription for the target Azure platform. For information about Azure Storage Accounts, refer to *Introduction to Azure Storage* in Microsoft Azure documentation.

NOTE: A Standard Storage Account can be used for Azure VM sizes that use HDD or SDD storage media. A Premium Storage Account can be used only for Azure VM sizes that use SDD storage media.

If no Azure Storage Account is associated with a subscription, PlateSpin Migrate sets up a Standard general-purpose storage account to use as the datastore for the target VM. The datastore name is based on the Azure Resource Group for the Subscription.

For information, see *Azure Page Blobs* (unmanaged disks) in Microsoft Azure Documentation.

Prerequisites for Using Azure Availability Sets

An Azure Availability Set is an Azure service that provides high availability to two or more member VMs by deploying them in different *fault domains* and *update domains*. Fault domains ensure that the VMs run on different host servers, compute racks, storage units, and network switches. Update domains ensure that maintenance outages for host hardware do not occur at the same time.

NOTE: VMs must be created in the Availability Set. You cannot add an existing VM to an Availability Set.

Using Azure Availability Sets is optional. No availability set is configured by default for the target VM, even if the selected Resource Group includes availability sets.

PlateSpin Migrate supports using pre-defined Azure Availability Sets for target VMs. You cannot create Availability Sets in Migrate.

Before you configure workloads for the target VMs that you want to add to an Availability Set:

- Use the Azure portal to create and define the Availability Set in a resource group for the specified target location.
- You can use an Availability Set in combination with Azure Managed Disks only if you enable the Use Managed Disks option in the Availability Set definition.

When the first VM is created in the Availability Set, Azure selects the physical hardware that will be used in the target Azure location. Thereafter, Azure limits the cloud instance sizes and networks that can be used for future member VMs accordingly.

NOTE: If you have multiple VMs configured for the same empty Availability Set, the first cutover VM determines the Azure limits based on the configuration characteristics of that VM. PlateSpin Migrate might prompt you to choose different cloud instance sizes or networks after the Azure limitations are enacted.

Prerequisites for Installing Azure VM Agent

PlateSpin Migrate provides the ability to automatically install Microsoft Azure Virtual Machine Agent on target Windows and Linux workloads at cutover and test cutover. The minimum requirements for the source workloads are:

- **Windows:** Microsoft Azure supports Azure VM Agent only for workloads running Windows Server 2008 R2 and higher.
- **Linux:** Before you add the source Linux workload to Migrate, ensure that it meets the following requirements:
  - PlateSpin Migrate requires Python 2.7 or higher to be installed on the source workload.
  - The source Linux workload must also meet minimum Azure Linux Agent requirements and system dependencies. See *Understanding and Using the Azure Linux Agent* (https://docs.microsoft.com/en-us/azure/virtual-machines/extensions/agent-linux) in the Microsoft Azure Documentation.

Azure Prerequisites for Using an On-Premise Migrate Server

If you set up an Azure site-to-site VPN (or an Azure Express Route connection) between the premises where your source workloads reside and the target Azure environment, you can deploy your PlateSpin Migrate server on-premises. Before you use PlateSpin Migrate to migrate workloads to Microsoft Azure, ensure that the following cloud access prerequisites are correctly configured and available:

- A PlateSpin Migrate license.
- A PlateSpin Migrate server deployed on-premise.
- A site-to-site VPN connection between your local data center and Microsoft Azure Portal.

  For information, see the following Microsoft resources:
  - *Create a Site-to-Site Connection in the Azure Portal* (https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-howto-site-to-site-resource-manager-portal)
  - *Create VNet with Site-to-Site VPN Connection Using PowerShell* (https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-create-site-to-site-rm-powershell)
- A default Gateway Subnet.
- The minimum network-related prerequisites for a successful migration are described in Table 9-2.

Table 9-2 *Port Requirements for Migrate Server on Premise*

Location	Port	Protocol	Remarks
On-premise source workload Cloud-based target workload	TCP 443, outbound	HTTPS	The on-premise source workload and the cloud-based target workload must be able to communicate with the PlateSpin Migrate server through HTTPS (TCP/port 443) over the site-to-site VPN connection.

Location	Port	Protocol	Remarks
On-premise Migrate Server	TCP 443, outbound	HTTPS	The on-premise PlateSpin Migrate server must be able to communicate with the Microsoft Azure API endpoint.
On-premise source workloads	TCP 22 TCP 135, 445 UDP 135, 138 and TCP 39	SSH (Linux) WMI/RPC/DCCOM NetBIOS	The PlateSpin Migrate server must be able to communicate with the source workloads on the ports that are used for discovery. See "Requirements for Discovery" on page 53 and "Discovering Details for Source Workloads" on page 301.
On-premise source workloads using Migrate Agent	TCP 22 TCP 443	SSH (Linux) HTTPS	Instead of discovery, you can use the Migrate Agent utility to register source workloads with the Migrate server. See "Requirements for Workload Registration" on page 56 and "Registering Workloads and Discovering Details with Migrate Agent" on page 303.
On-premise source workload Cloud-based target workload	TCP 3725	Migrate	The cloud-based target workload must be able to communicate (target to source) with the on-premise source workload across the VPN. The source workload must be able to send data to the target workload during replication across the VPN. The port number is configurable. See port 3725 in "Requirements for Migration" on page 57. If you use Migrate Agent for registration and discovery, the default direction of the replication connection must be reversed (source to target) by changing advanced settings on the Migrate server. See "Configuring the Contact Direction for the Replication Port" on page 114.

Location	Port	Protocol	Remarks
Network Security Group in Azure for the cloud-based target workloads	TCP 443, inbound TCP 3389, inbound TCP 22, inbound	HTTPS RDP (Windows) SSH (Linux)	Allow inbound connections in the Network Security Group for the cloud-based target workloads. For information about creating and configuring a Network Security Group in Azure, refer to *Create, Change, or Delete a Network Security Group* (https://docs.microsoft.com/en-us/azure/virtual-network/manage-network-security-group) in Microsoft Azure Documentation.

Azure Prerequisites for Using an Azure-Based Migrate Server

Before you use PlateSpin Migrate to migrate workloads to Microsoft Azure, ensure that the following cloud access prerequisites are correctly configured and available:

- A PlateSpin Migrate license.
- Deploy an Azure Marketplace image of the PlateSpin Migrate server in the target Azure environment. See "Deploying PlateSpin Migrate Server in the Cloud" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

  NOTE: The cloud-based Migrate server does not require a site-to-site VPN connection between your local data center and Microsoft Azure Portal. When no VPN is provided between the source network and the cloud-based Migrate server, you can use Migrate Agent to register workloads with the cloud-based Migrate server using secure communications over the public Internet. Internet access and public IP addresses are required. For deployment information, see Figure 8-2, "Cloud-Based Migrate Server for Automated Migration to AWS," on page 147.

- 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.

  NOTE: A change in IP address on the PlateSpin Server breaks the heartbeat communications with source workloads.

  You cannot specify the actual IP address assigned to the public IP resource. Azure allocates and reserves an IP address from a pool of its available IP addresses in the Azure location where you deploy the Migrate server. The address persists through server restarts. Azure releases the IP address only when you delete the resource or change the resource's allocation method to **Dynamic**.

- Install the Migrate Agent on the source workload, then register the workload with the cloud-based PlateSpin Migrate server. See "Registering Workloads and Discovering Details with Migrate Agent" on page 303.

To download the Migrate Agent, launch the PlateSpin Migrate Web Interface and click the **Downloads** tab. For information about installing and using the Migrate Agent, see "Migrate Agent Utility" on page 359.

- The minimum network-related prerequisites for a successful migration when the Migrate Server is in Azure are described in Table 9-3.

Table 9-3 *Port Requirements for Migrate Server in Azure*

Location	Port	Protocol	Remarks
Source workload Network firewall	TCP 443, outbound	HTTPS	Required to allow the source workload to register (using the Migrate Agent utility) and communicate with the cloud-based PlateSpin Migrate server. The PlateSpin Migrate Server uses secure SSL for communications with the workloads you want to migrate.
Source workload Network firewall Network Security Group (NSG) in Azure	TCP 3725, outbound	Migrate	Required to allow communications with the target machine and to transfer data from the source to the target during replication. The direction of the communication (source to target) is automatic, but the port number is configurable. For information about changing the default port setting, see port 3725 in "Requirements for Migration" on page 57. For information about creating and configuring a Network Security Group in Azure, refer to *Create, Change, or Delete a Network Security Group* (https://docs.microsoft.com/en-us/azure/virtual-network/manage-network-security-group) in Microsoft Azure Documentation.
NSG in Azure for the Migrate Server	TCP 443, inbound TCP 3389, inbound	HTTPS RDP	Allow inbound connections in the Network Security Group for the cloud-based Migrate server. The <*Migrate-server-name*>-nsg is created automatically when you deploy the Migrate server in Azure.
NSG in Azure for the Migrate Server	TCP 123, outbound	Network Time Protocol (NTP)	Add this port setting to the security group if you are using an NTP service outside the virtual network where you deploy the Migrate server.
NSG in Azure for the Migrate Server	TCP 22, outbound	SSH	This port allows outbound communications from the Migrate server to Linux workloads.

Planning For Migrating Workloads to Azure

PlateSpin Migrate enables you to use the PlateSpin Migrate Web Interface to migrate Windows and Linux workloads to Microsoft Azure. For a list of supported workloads, see "Supported Workloads For Migration to Microsoft Azure" on page 29.

NOTE: Migration of Windows Cluster workloads to Azure is not supported.

Target Azure IaaS Environment

- Each PlateSpin Migrate server can support migration to multiple Azure global and sovereign environments. Set the appropriate Azure environment when you configure a target Azure platform:
  - Azure Global
  - Azure Stack

Azure Subscription

- Provide valid credentials for the Azure subscription. See "Managing the Azure User Password for Azure Target Cloud Platforms" on page 197.

PlateSpin Server Host

- 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 access error.

OS License for Target Workload

- 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.

Target Workload

Consider the following guidelines before you migrate workloads to Azure:

- **Migrate Web Interface**

  You can use only the PlateSpin Migrate Web Interface to migrate the workloads to Microsoft Azure. The PlateSpin Migrate Client does not support migration of workloads to Microsoft Azure.

- **Migrate API**

  PlateSpin Migrate API supports headless migrations only to Azure Cloud. See *PlateSpin Migrate API Getting Started Reference*.

  NOTE: PlateSpin Migrate API currently does not support headless migrations to Azure Stack.

- **UEFI Source Workloads**

  Windows and Linux UEFI workloads are migrated as BIOS workloads.

- **Multiple NICs**

  Migration of workloads with multiple NICs to Azure is supported for Windows workloads, up to the number of NICs supported by the Azure VM size.

NOTE: When you configure migration for a workload, PlateSpin Migrate lets you specify four IP addresses per NIC on the Windows workload by default. However, you can configure the number of IP addresses that you want to specify per NIC on the workload. See "Configuring the Maximum Number of IP Addresses Allowed for Each Network Interface on Target Workloads" on page 111.

- **Maximum Number of Data Disks**

  PlateSpin Migrate supports Azure VM sizes with up to 64 data disks. For the maximum VM size in a selected Azure Region, Migrate will use one data disk for the OS disk replication in the PlateSpin Replication Environment. After cutover, this disk becomes the OS disk, and you can add a data disk.

- **Maximum Disk Size for Data Disks**

  Data disks can have a maximum size of 4 TB (4092 GB), depending on the maximum size allowed for the target VM instance size and storage type. This data disk size limit applies to Azure managed disks and unmanaged disks.

- **Disk Size on the Target VM**

  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 disk space on Azure.

- **Cloud Instance Size**

  Migrate initially identifies an Azure VM size in the specified target location that meets or exceeds the source workload's settings for cores, memory, data disks, and NICs. However, you can choose a smaller or larger VM size based on your requirements for the target workload, as limited by the maximum VM sizes available in the selected Azure Region.

- **Temporary Storage on the Target VM**

  In Azure, every VM contains a temporary disk that provides short-term storage for applications and processes. It can also be used to store temporary data such as a page file or swap file. PlateSpin Migrate configures the page file for a target Windows VM to use the Azure temporary storage drive and mounts it to a drive letter.

  WARNING: Do not use this temporary storage to store data that you are not willing to lose.

  The temporary storage size depends on the VM type and size. See *Sizes for Windows Virtual Machines in Azure* (https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sizes/) in the Microsoft Azure Documentation.

Azure Networking Guidelines

You can create a virtual machine with multiple NICs in Azure virtual networks. Each NIC must be located in one subnet; one subnet can be assigned to multiple NICs. Each NIC has an IP address consistent with its subnet assignment. The IP address and MAC pairing for each NIC persists, even if the order of the NICs changes.

Consider the following guidelines when you are migrating workloads to Microsoft Azure.

- "Private or Public IP Addresses for Azure Migration" on page 181
- "Windows Workloads in Azure with Multiple NICs" on page 181
- "Private and Public IP Addresses for Workloads Connected on an Azure VPN" on page 181

Private or Public IP Addresses for Azure Migration

You can use private IP addresses for workload migration if you have configured an Azure VPN to connect your premise network with your Azure cloud environment. Otherwise, you must enable a public IP address to be assigned to the replication network, cutover network, and test cutover network. If the VM has multiple NICs, only the primary NIC can have a public IP address. The assigned public IP addresses will be in the address space of the specified network and subnet for the designated NIC in each network.

NOTE: PlateSpin requires a public IP address only if a site-to-site Azure VPN is not available.

If you enable a public IP address for the primary NIC, Azure assigns the NIC both a public IP address and a private IP address. For more information about connecting to the Azure VM, see "Private and Public IP Addresses for Workloads Connected on an Azure VPN" on page 181.

Windows Workloads in Azure with Multiple NICs

Azure configures the VM with a default gateway that is associated with the primary network interface. Azure removes the gateway information for all secondary NICs, which limits their communications to the same subnet as the primary interface.

For Windows workloads with multiple NICs, you can enable a secondary NIC to communicate outside its own subnet. Use the Windows `route add` command to add a different gateway entry for the secondary NIC in the routing table. See "Configure Windows VMs" in *Create a VM with Multiple NICs* (https://azure.microsoft.com/en-us/documentation/articles/virtual-networks-multiple-nics/) on the Microsoft Azure website (https://azure.microsoft.com/).

Private and Public IP Addresses for Workloads Connected on an Azure VPN

An Azure VM can have one or more NICs attached to it. The primary NIC for the VM can have both a public and private IP address. A private IP address is used for communications from other resources in a virtual network and from machines inside the address space for the Azure VPN that connects your premise network to your Azure cloud environment. A public IP address can be used to communicate with the Internet and with machines outside the Azure cloud environment. Azure automatically associates these IP addresses with the primary network interface for the VM.

You can use the Microsoft Remote Desktop client to connect remotely to the Azure VM. Specify the IP address as follows:

- **Private IP address:** Use the VM's private IP address if your machine is part of the address space for the Azure VPN.
- **Public IP address:** Use the VM's public IP address if your machine is not part of the address space for the Azure VPN.

You can alternatively use the **Connect** option in the Microsoft Azure portal (https://azure.microsoft.com/en-us/features/azure-portal/) from a machine with an address space that is not part of the Azure VPN. This option automatically launches the Microsoft Remote Desktop client configured to connect to the VM's public IP address for the primary NIC.

NOTE: This portal operation fails if your machine is in the address space of the Azure VPN.

Configuring an Application in Azure to Represent PlateSpin Migrate

PlateSpin Migrate uses the Microsoft Azure API to automate workload migrations to Azure. You need to create an Azure application ID for PlateSpin Migrate to use when it uses the Azure API for replicating and migrating workloads to your Azure account.

NOTE: For migrations to Azure Stack, you must not create an Azure application ID for PlateSpin Migrate. Instead, use the following ID as the Azure application ID for PlateSpin Migrate:

```
872cd9fa-d31f-45e0-9eab-6e460a02d1f1
```

When you register an application, Azure creates a central identity for it in Azure Active Directory and assigns it an Application ID. The registration stores basic information that Migrate uses to authenticate to and use Azure APIs. It also records delegated permissions that Migrate needs when it accesses the APIs on behalf of the Contributor user for the subscription.

NOTE: For more information about application registration and permissions. see the following Microsoft resources:

- "Register Application" in *Use Resource Manager Authentication API to Access Subscriptions* in the Azure Resource Manager documentation
- *Quick Start: Register an Application with the Microsoft Identity Platform* in the Azure Active Directory documentation

- "Registering an Application in Azure" on page 182
- "Granting Admin Consent for Application Permissions" on page 186
- "Verifying Admin Consent for the Application" on page 187

Registering an Application in Azure

To register an application in Azure for PlateSpin Migrate to use:

1. Use the Azure Global Portal (https://portal.azure.com/) to log in to your Azure account as a user with rights necessary to create and manage applications.
2. In the Portal sidebar, select **Azure Active Directory**, then in the left pane under **Manage**, select **App registrations**.
3. Register PlateSpin Migrate for your account:
   - **3a** On the App registrations page, click **New Registration** in the command bar.
   - **3b** On the **Register an application** page, provide the following information:

Register an application	
Name	Specify friendly name for the application, such as `PlateSpin Migrate Integration`. The name must be unique in your Azure Active Directory. This is the name that appears in the Applications list.
Supported account types	Select **Accounts in any organizational directory**. This option enables the application for user accounts in your default directory in Azure Active Directory.
Redirect URI	Specify **Web**, then leave the **Redirect URI** field empty.

3c Read the Microsoft Platform Policies, then click **Register**.

The working pane opens to the Overview page of the newly registered application.

4 On the Overview page of the new application, copy the **Application (client) ID** value to the clipboard and paste it in a text document where you can access it when you set up the target cloud platforms for subscriptions in this account.

An Application (client) ID is a value in the format of: abc12b34-c5df-6e78-f9a0-bc123456d789.

You must provide the Application (client) ID when you add a Microsoft Azure Location as a migration target in PlateSpin Migrate.

5 Configure Authentication settings for the application that you registered in Step 3.

The Public Client setting allows the Migrate server to authenticate programmatically to Azure when it uses the Azure APIs.

5a Continuing on the Overview page of the new application, in the left pane under **Manage**, select **Authentication**.

To access the Overview page: In the Portal sidebar, select **Azure Active Directory**, select **Manage** > **App registrations**, then under **Owned applications**, select the application by its name.

5b On the Authentication page, scroll down to **Advanced Settings** > **Default Client Type**.

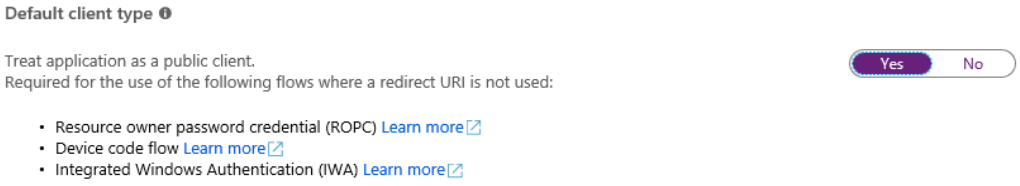

5c Click **Yes** to treat the application as a public client.

5d In the command bar, click **Save**.

5e Dismiss the confirmation message for the updated Authentication settings.

5f In the left pane, click **Overview** to return to the Overview page for the application.

6 Configure API permissions for the application that you registered in Step 3.

6a Continuing on the Overview page for the new application, in the left pane under **Manage**, select **API permissions**. You can also click **View API Permissions** on the Overview page.

To access the Overview page: In the Portal sidebar, select **Azure Active Directory**, select **Manage** > **App registrations**, then under **Owned applications**, select the application by its name.

6b **Microsoft Graph:**Modify this default permission to add the User.ReadBasic.All permission.

User.Read grants permission to read the profile of the signed-in user. User.ReadBasic.All constrains application access to a limited set of properties known as the basic profile (display name, given name, mail, photo, surname, and user principal name).

6b1 On the API Permissions page, click **Microsoft Graph**.

Its current permission type is `Delegated` and the permission granted is `User.Read`.

6b2 On the Request API Permissions page under **Select Permissions**, scroll to the **User** section and select **User.ReadBasic.All**.

6b3 On the bottom of the page, click **Update Permissions**.

6b4 Dismiss the confirmation message for the updated permissions.

6b5 On the API Permissions page, visually verify that Microsoft Graph now has Delegated permissions for `User.Read` and `UserReadBasic.All`.

6c **Azure Service Management:**Add a new API permission for Azure Service Management.

Migrate uses the Azure Resource Manager Authentication API to access subscriptions, which requires the following permissions for the registered application: Azure Service Management, Delegated Permissions, and user_impersonation.

6c1 On the API Permissions page, click **Add a Permission**.

6c2 On the Request API Permissions page, provide the following information:

Request API Permissions	
API Category	Select **Microsoft APIs**.
API	Select **Azure Service Management**.
Type of Permissions	Select **Delegated Permissions**. When PlateSpin Migrate accesses Azure APIs, the permissions are delegated to an authorized Contributor User that PlateSpin Migrate logs in as to perform migration actions. See "Configuring a Contributor User for PlateSpin Migrate to Use" on page 188.
Permissions	Select **user_impersonation**.

6c3 At the bottom of the page, click **Add permissions**.

6c4 Dismiss the confirmation message for the newly added permissions.

6c5 On the API Permissions page, visually verify that Azure Service Management is listed in the API permissions list with Delegated permissions for user_impersonation.

Although permissions are configured, they are not yet effective for the application. An Azure global administrator must grant admin consent for the permissions in the Default Directory to ensure that consent is not required when PlateSpin Migrate accesses Azure APIs.

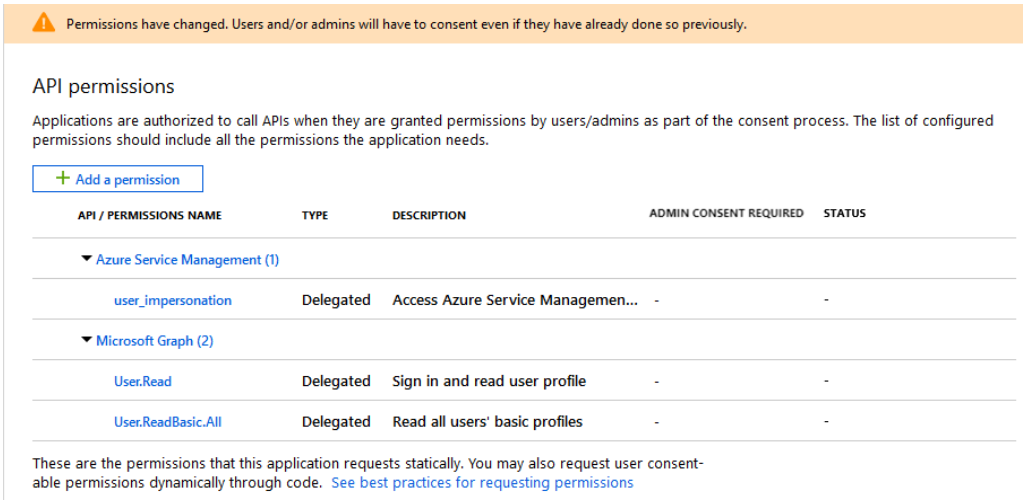

7 [Azure Global Administrator rights for the Default Directory are required for this step.] Grant admin consent for the Default Directory for the requested permissions.

Do one of the following:

- If you have admin rights, you can grant consent now from the API Permissions page. Continue with "Grant Admin Consent from App Registrations" on page 186.

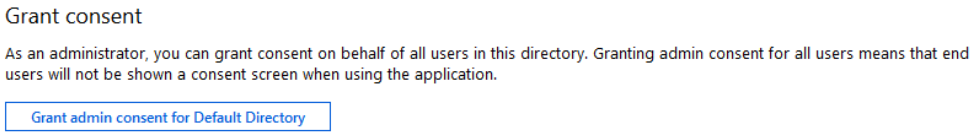

- If you are not currently logged in as a user with admin rights, the option to grant admin consent is disabled on the API Permissions page.

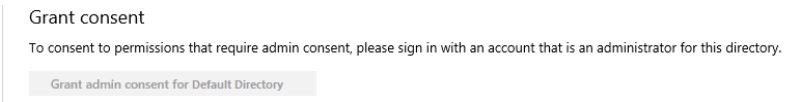

Use either of these methods to grant consent:

- You can log in as a user with rights from the application Permissions page in Enterprise Applications. Continue with "Grant Admin Consent from Enterprise Applications" on page 187.
- You can log out of the portal and log in as a user with the administrator rights, then return to the API Permissions page to grant consent. Continue with "Grant Admin Consent from App Registrations" on page 186.

Granting Admin Consent for Application Permissions

A user with Azure Global Administrator rights for the Default Directory must grant consent for the permissions you request for the registered application that PlateSpin Migrate will use. The consent on behalf of all users will ensure that consent is not required when PlateSpin Migrate accesses Azure APIs. After consent is granted, the application is ready for use by PlateSpin Migrate.

IMPORTANT: Ensure that you are logged in with an Azure user account that has Global Administrator rights for the Default Directory where the application will be used.

- "Grant Admin Consent from App Registrations" on page 186
- "Grant Admin Consent from Enterprise Applications" on page 187

Grant Admin Consent from App Registrations

To grant admin consent from App Registrations:

1. In the Portal sidebar, select **Azure Active Directory**, select **Manage** > **App registrations**, then under **Owned applications**, select the application by its name.
2. Under Manage, select **API Permissions**.
3. On the API Permissions page for the application, click **Grant admin consent for <*Default Directory*>**.

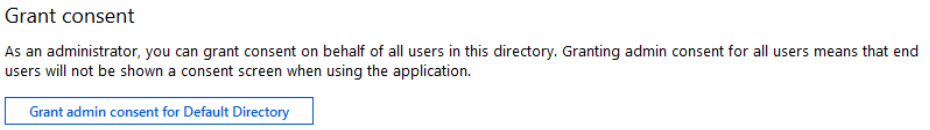

4. When you are prompted to grant consent for the requested permissions for all accounts in the Default Directory, click **Yes**.
5. Dismiss the confirmation message that you successfully granted admin consent for the requested permissions.
6. On the API Permissions page, visually verify that the Status shows that the permissions are Granted for Default Directory.

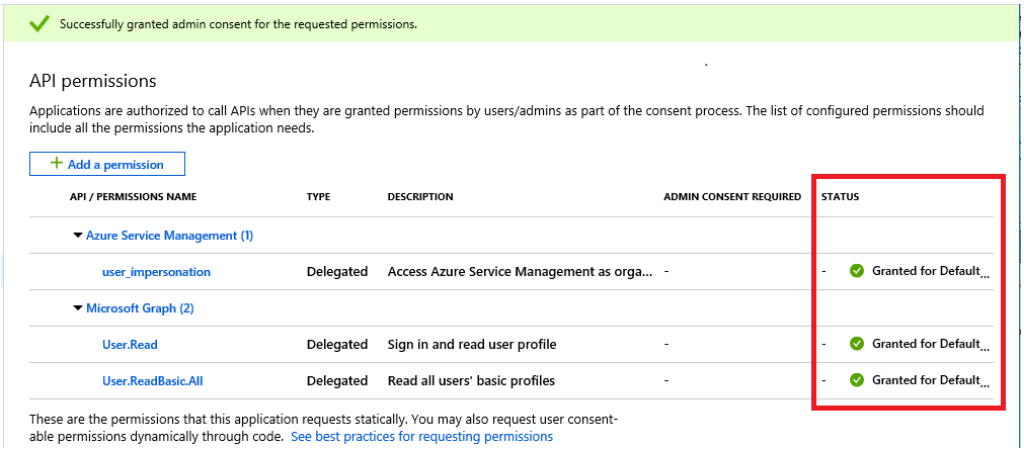

7 Continue with "Verifying Admin Consent for the Application" on page 187.

Grant Admin Consent from Enterprise Applications

1 In the Portal sidebar, select **Azure Active Directory**, then In the left pane under **Manage**, select **Enterprise Applications**.

2 On the Enterprise applications - All applications page, locate and select the new application for PlateSpin Migrate that was created in Step 3 of "Registering an Application in Azure" on page 182.

3 In the left pane under **Security**, select **Permissions**.

Initially, there are no permissions listed in the **Admin Consent** section for the application.

4 Click **Grant admin consent for <*Default Directory*>**.

A separate browser window opens for the authentication and permission flow for the selected application.

5 When you are prompted to sign in to administer the application, sign in using an Azure global administrator account that has rights to grant admin consent for the Default Directory.

6 After authentication succeeds, the **Permissions requested - Accept for your organization** dialog prompts you to consent to the requested application permissions. Click **Accept**, then wait for the browser to refresh its content.

7 After the permissions are successfully granted, close the pop-up browser window.

8 Continue with "Verifying Admin Consent for the Application" on page 187.

Verifying Admin Consent for the Application

To verify the application and its permissions:

1 In Portal sidebar, select **Azure Active Directory**, then select **Manage** > **Enterprise Applications**.

2 Locate and select the new application that you created in Step 3.

3 Under **Security**, click **Permissions**.

4 Verify that required permissions are listed in the **Admin Consent** section and that consent is granted for the Default Directory.

The permissions should match the permissions you configured in Step 6 of "Registering an Application in Azure" on page 182.

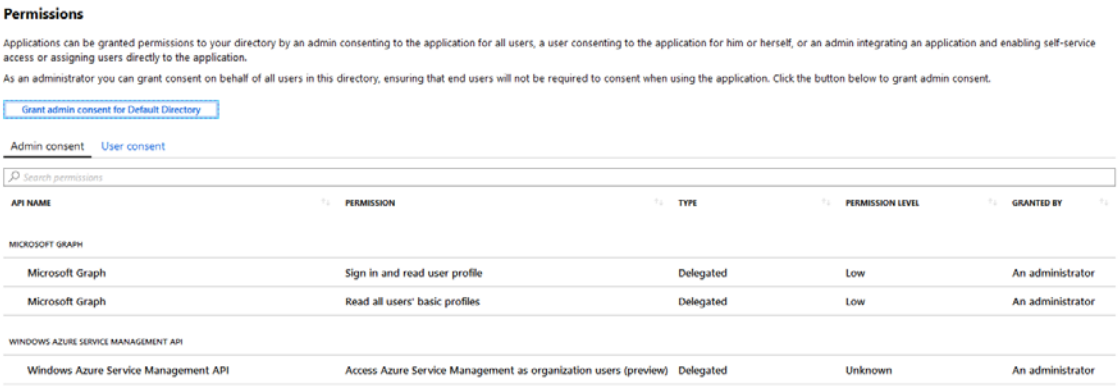

Configuring a Contributor User for PlateSpin Migrate to Use

PlateSpin Migrate uses the Microsoft Azure API to automate workload migrations to Azure. A user with administrative privileges on the Azure subscription must create a user for PlateSpin Migrate in your Azure Active Directory, then assign the Contributor role to the user for the subscription. You use this Contributor user credentials to add Azure as a target in Migrate. Migrate uses the credentials for this user when it accesses the Migrate Azure API through the related subscription.

To configure a Contributor user for a subscription:

1 In Azure Active Directory, create a new special-purpose user for PlateSpin Migrate.

1a Use the Azure Global Portal (https://portal.azure.com/) to log in as a user with administrative privileges on the Azure Subscription in your account.

1b In the Portal sidebar, select **Azure Active Directory**.

1c In the left pane under **Manage**, select **Users**, then select **All Users**.

1d In the Command bar, select **New User** to create a new user.

1e On the User page, specify the following information:

Parameter	Description
Name	Specify the name of the new user that you will use for PlateSpin Migrate. For example: `PlateSpin Migrate`.

Parameter	Description

User name	Specify the user name in the following format:
	`<username>@<default_directory>.onmicrosoft.com`
	The user must be a member of the Default Directory (federation is not supported). To view the name of the default directory for your account, pause over the Account area in the upper right corner of the Azure portal and view the Domain information related to your Azure administrative account. The default directory name is the first part of the Domain name. For example, if the Azure administrative account is jim@acme.com, the default directory name is `acme` and the Domain name is `acme.onmicrosoft.com`.
	For example:
	`platespinmigrate@acme.onmicrosoft.com`

Profile	Not configured
Properties	Default
Groups	None
Directory role	User
Password	Click **Show Password** and note down the auto-generated password provided in the **Password** box. You need this information to change the password after the account is created.

1f Click **Create** to create the new user.

The newly created user appears in the Users - All Users list.

2 Log in to the Azure portal as the new user for PlateSpin Migrate and change the user password.

2a Launch a separate web browser window.

2b Log in to the Azure portal by using the credentials for the new user for PlateSpin Migrate.

Use the user name and auto-generated password that you noted when you created the user.

2c When you are prompted to **Update Your Password**, specify the auto-generated password, then specify a new secure password.

You need this new password when you add a Microsoft Azure Location as a migration target in PlateSpin Migrate.

2d Click **Sign In**.

2e After a successful log in, log out of the Azure portal and close this separate browser window.

3 Continuing your Azure session, assign the Contributor role to the new user for the appropriate Azure Subscription.

3a Log in to the Azure portal with your Azure user account that you use to manage your subscriptions.

3b In the Portal sidebar, select **All services**, then select **Subscriptions**.

3c In the list of your subscriptions, select the subscription that you want to use for your migrations.

3d Copy the **Subscription ID** value to the clipboard and paste it in a text document where you can access it when you set up the target cloud platforms for this account.

You must provide the Subscription ID when you add a Microsoft Azure Location as a migration target in PlateSpin Migrate.

3e In the subscription detail view, select **Access control (IAM)**, select **Add**, then select **Add role assignment** to assign a new role for the user.

3f Specify the following information:

Parameter	Description
Role	Select **Contributor**.
Assign access to	Select **Azure AD user, group, or service principal**.
Select	Locate the user by typing the first characters in the name, and then select the user.

3g Click **Save.**

Enabling PlateSpin Replication Environment for Azure Subscriptions

PlateSpin Migrate must be able to programmatically deploy a PlateSpin Replication Environment (PRE) VM during the replication of workloads to Azure. The required VM image is available in the Azure Marketplace. You must enable the PRE image for each Azure Subscription that you plan to use as a migration target. You must have at least the Contributor role for a subscription to manage programmatic access.

IMPORTANT: Ensure that you enable the PRE and accept the Azure terms of use for the target Azure Subscription before you attempt to migrate workloads to Azure. Otherwise, migrations will fail with the following error:

```
User failed validation to purchase resources. Legal terms have not been
accepted for this item on this subscription.
```

- "Enabling Programmatic Deployment of PlateSpin Replication Environment" on page 190
- "Verifying the Programmatic Deployment of PRE for a Subscription" on page 192

Enabling Programmatic Deployment of PlateSpin Replication Environment

Before you attempt to migrate workloads to Azure Cloud, you must enable PlateSpin Replication Environment image for programmatic access.

To enable the PlateSpin Replication Environment for an Azure Subscription:

1. Use the Azure Global Portal (https://portal.azure.com/) to log in to your Azure account that has administrator rights for the Azure Subscriptions that you will use for migrations. You must have at least the Contributor role for a subscription manage programmatic access.
2. In the portal menu, click **Create a Resource**.
3. In the Marketplace Search, type `platespin`, then press Enter.

   The Search Results shows one or more PlateSpin offerings.
4. Select the **PlateSpin Replication Environment** offering by Micro Focus.

   The Micro Focus PlateSpin Replication Environment runs the SUSE Linux Enterprise 12 SP3 operating system.

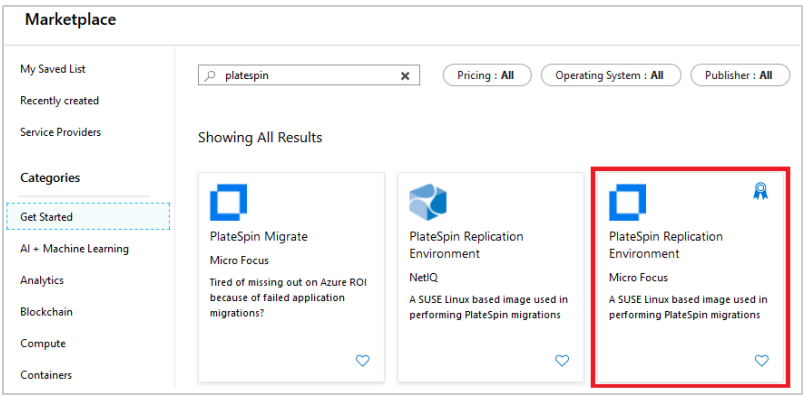

5. On the PlateSpin Replication Environment page under **Create** next to **Want to deploy programmatically?**, click **Get Started**.

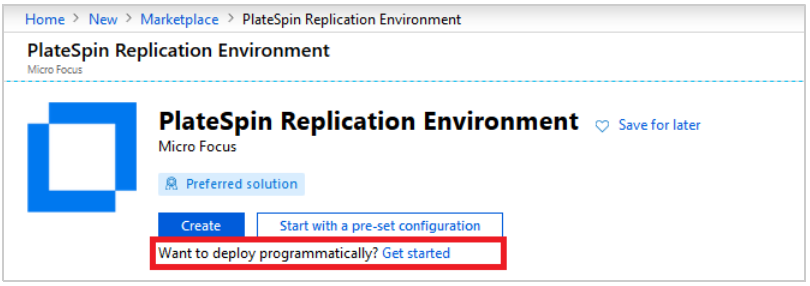

6. On the Configure Programmatic Deployment page, read the *Offer Details* for PlateSpin Replication Environment and the *Terms of Use* for programmatic purchases.
7. Scroll down to **Choose the subscriptions** to view a list of Azure Subscriptions to which you have access.
8. Click **Enable** for each Azure Subscription that you will configure as a migration target in PlateSpin Migrate.

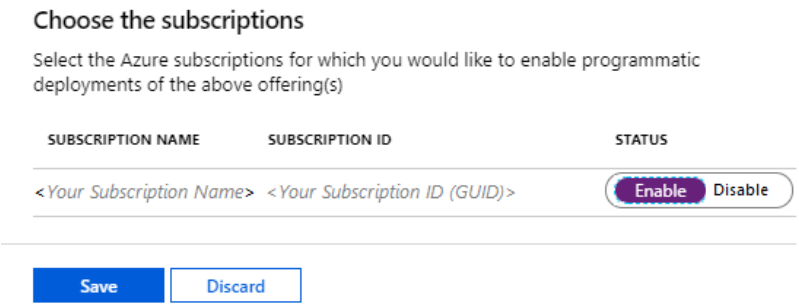

9 Click **Save**.

10 Click X in the top right to close the page.

11 Continue with "Verifying the Programmatic Deployment of PRE for a Subscription" on page 192.

Verifying the Programmatic Deployment of PRE for a Subscription

Verify the programmatic deployment of PlateSpin Replication Environment for the Azure Subscription.

1 In the Portal sidebar, select **All Services,** then search for and select **Subscriptions**.

2 On the Subscriptions page, select a Subscription that you enabled for programmatic deployment of PlateSpin Replication Environment.

3 On the Overview page for the Subscription in the left pane under **Settings**, select **Programmatic deployment**.

4 On the <*Your Subscription Name*> - Programmatic Deployment page, visually verify that PlateSpin Replication Environment is listed as an offering and its status is **Enabled**.

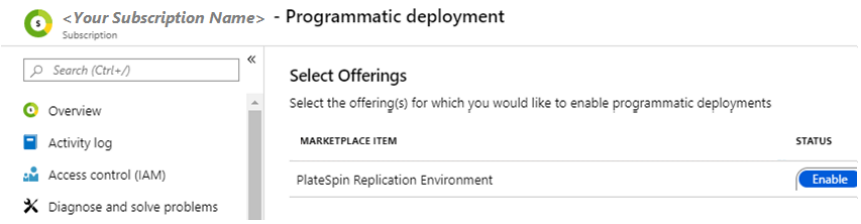

Configuring the Azure Instance Type Used For the Azure Replication Environment Virtual Machine

For the Azure Replication Environment VM, PlateSpin Migrate Server by default uses the smallest instance size that is required for the workload migration based on the number of disk drives you configure for the target workload.

To change the Azure instance type used during replication, set the value of the **AzureInstanceSizeForReplicationEnvironment** parameter to the Azure instance type you want to use for the Replication Environment Virtual Machine based on the number of disk drives you configure for the target workload.

To specify an instance size for the Azure Replication Environment VM:

1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2 Search to locate the AzureInstanceSizeForReplicationEnvironment parameter and set its value to the required instance size.

IMPORTANT: To specify the instance size, replace all the spaces in the instance size name with underscore(_). For example, "Standard A4m v2" must be specified as "Standard_A4m_v2".

---

3 Save your settings and exit the page.

A reboot or restart of PlateSpin services is not required to apply the changes.

Deploying a Migrate Server Image in Azure

PlateSpin Migrate offers a PlateSpin Migrate Server image in Azure through the Azure Marketplace in each of the supported Azure environments. See “Deploying PlateSpin Migrate Server in the Cloud” in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*:

Manually Adding PlateSpin Replication Environment Image to Azure Stack

When you perform a first-time discovery of an Azure Stack, the discovery might take some considerable time (20-30 mins) because PlateSpin Migrate downloads the `PlateSpin Replication Environment(PRE)` image from the Azure Storage and uploads it to the Azure Stack. If the `PlateSpin Replication Environment` image is not successfully uploaded to Azure Stack, then the following messages are logged and you must manually upload the PRE image to the Azure Stack before you configure migrations to it.

```
Create UserImage has failed With Message: {0}
```

```
Create a User Image manually in the portal with Name {0} before continuing
with the migration
```

For information about discovery, see “Target Discovery in the Web Interface” on page 284.

To manually upload the PRE image to Azure Stack:

1 Download the PlateSpin Replication Environment disk from the following location to a temporary folder on your computer:

```
https://platespinmigration2.blob.core.windows.net/vhds/azure-
platespin-replication-environment-image-12.2.2.vhd?st=2018-02-
09T23%3A00%3A00Z&se=2021-03-31T22%3A00%3A00Z&sp=rl&sv=2016-05-
31&sr=b&sig=yb0xWeriJh7AJHEnm1LQaQQBtIYraGNnK4NL3auldD0%3D
```

2 Upload the PlateSpin Replication Environment disk from the temporary folder on your computer to a storage account in Azure Stack. Make a note of the blob storage URI where you upload the image. See Azure documentation (https://docs.microsoft.com/en-us/azure-stack/operator/azure-stack-add-vm-image?view=azs-1910).

3 Create a Linux image named `PlateSpinReplicationEnvironment`.

   3a Log in to the Azure Stack Portal (https://portal.local.azurestack.external/).

   3b Click **Images** in the portal menu, then click Add in the working pane.

   3c Specify the name of the image you want to create as **PlateSpinReplicationEnvironment**.

   3d Select the subscription, resource group and location as applicable.

   3e Set the **OS disk** as **Linux**.

   3f In the **Storage blob** option, click **Browse** and browse to select the blob storage location where you previously uploaded the PRE image.

   3g Select the account type and host caching policy.

   3h Click **Create**.

4 After the PRE image is added to the Azure Stack, refresh the discovered Azure Stack by using the PlateSpin Migrate Web Interface. See "Refreshing Target Host Details" on page 293.

See Azure documentation (https://docs.microsoft.com/en-us/azure-stack/?view=azs-1910).

Using Azure Cloud Tags for Azure Migrations

PlateSpin Migrate enables you to configure Azure Cloud tags that it applies to VMs, resource groups, and resources that it creates on your behalf in target Azure platforms (for Azure objects that accept tags). Each tag consists of a name and value pair. In Azure, you can use the tags to logically organize your assets and track their related costs, and to monitor resource usage and retention.

About Cloud Tags for Azure

Azure Cloud tags are *key name* and *value* pairs. On your PlateSpin Server, you can define tags that apply globally for workload migrations to all target Azure Cloud platforms on the PlateSpin Server. These tag key names and values (if defined) are pre-populated on the Workload Configuration page for migrations to Azure. You can specify appropriate tag values for that workload. You can also specify other tags for an individual workload.

If Azure tags are configured for a workload, Migrate applies them to the target VM and its related resource objects that Migrate creates in the target Azure environment during the first full replication.

Azure tags for a workload are configurable for workloads before the first full replication. If you add or modify global tag definitions, updated tag definitions will apply automatically to new workload configurations. For configured workloads, only changes for required global tags will be applied automatically.

NOTE: After a successful full replication, a workload's tags and their values can no longer be modified. You cannot modify its tag settings on the Workload Configuration page and modifications to globally defined tags will not be applied to the workload.

PlateSpin allows up to 8 (eight) tags for a workload. You can define up to 8 global Azure Cloud tags that will be available automatically in the Workload Configuration for migrations to Azure. Use the `DefaultCloudTagsAzure` parameter on the PlateSpin Configuration page to specify definitions for your global Azure Cloud tags. Examples of tag names include: Cost Center, Business Service, Created By, Authorized By, Environment.

Table 9-4 describes the usage and format for global Azure Cloud tag definitions.

Table 9-4 *Format for Global Azure Cloud Tag Definitions*

Tag Behavior	Format	Example Tag Entry	Workload Configuration View
Optional	`key`	Service	Service : *<value>*
Optional, suggested value	`key%value`	Service%Engineering	Service : *<Engineering>*
Required, user must specify a value	`%key`	%Service	**Service : *<required>***
Required, fixed value	`%key%value`	%Service%Engineering	**Service : Engineering**

Consider the following requirements when you configure global tags:

- Define each key name and value pair on a separate line.
- Use the percent (%) character to delimit the key string and value string in the tag entry.
- Start an entry with the % character to indicate that the specified tag is required. Required tags will always be included as defined in the workload configuration.
- Entries that do not begin with the % character are optional. You can remove or modify them in the workload configuration.
- Tag names can be up to 512 alphanumeric and Unicode characters. Tag names must not contain the following special characters: less than (<), greater than (>), percent (%), ampersand (&), backslash (\), forward slash (/), question mark (?), and period (.). The names are case insensitive.
- Tag values can be up to 256 alphanumeric and Unicode characters. Tag values must not contain the less than (<) character. The values are case insensitive.
- If a tag value needs the percent (%) character, you must escape it with a preceding backslash. For example, to set a global Discount tag with a value of 5 percent, your definition needs a backslash only before the % character in the value field:

  ```
  Discount%5\%
  %Discount%5\%
  ```

Defining Global Azure Cloud Tags on a PlateSpin Server

PlateSpin supports up to 8 Azure tags for a workload. Tags defined in the **DefaultCloudTagsAzure** parameter will be available automatically in the Workload Configuration page for all migrations to Azure on the PlateSpin Server.

To define global Azure tags on a PlateSpin Server:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2. Search to locate the **DefaultCloudTagsAzure** parameter, then click **Edit**.
3. Define up to 8 global Azure tags that you want to make available during configuration of migrations to Azure.

   Add each tag definition on a separate line in the following format:

Tag Behavior	Format
Optional	`key`
Optional, suggested value	`key%value`
Required, user must specify a value	`%key`
Required, fixed value	`%key%value`

4. Save your settings and exit the page.

   A reboot or restart of PlateSpin services is not required to apply the changes.

Configuring Azure Cloud Tags for a Workload

For workload migrations to Azure, you configure Azure tags on the Workload Configuration page. In the Azure Tags field, click **Edit Tags** to access the dialog where you can configure the tag names and values as appropriate for the workload.

You can specify up to 8 tags for the workload. Global tags are pre-populated automatically. Required global tags with fixed values are applied automatically to the workload and cannot be modified. You must specify a value for a required tag if its value is not fixed. For optional global tags, you can modify the tag name or value, or you can remove them. You can optionally specify other tags, up to the maximum of 8 tags per workload.

You cannot modify tags and values for the workload after a successful full replication occurs. Changes to the global settings for tag names and values will not apply to migrations in this state.

Viewing Resources with Tags in the Azure Portal

You can use the Azure portal to view Azure tags that Migrate creates for resource groups and resources. See “Portal” in *Use Tags to Organize Your Azure Resources* in the Microsoft Azure Documentation.

To view all resources with a Migrate-related tag:

1. Log in to the Azure Portal.
2. In the sidebar, select **All Services**, then select **General** > **Tags** in the working pane.
3. (Optional) If you have multiple subscriptions and want to narrow your search, select the appropriate subscription.
4. In the list of Azure tags, select the Migrate-related tag of interest to view a list of all of the resources with that tag.

Managing the Azure User Password for Azure Target Cloud Platforms

Provide a valid password for the Microsoft Azure user when you add the Azure target cloud platform. Ensure that you update the password for the cloud platform in PlateSpin Migrate if you modify it in Azure.

Workload migrations can fail in the following conditions:

- **Invalid password:** If the stored password for the Azure user is invalid, an authentication error occurs when the next connection to Azure is requested.

  If the Azure user modifies the password in the Microsoft Azure portal while migration tasks are running, the tasks will fail with an authentication error when the next connection to Azure is requested.
- **Expired password:** If the stored password for the Azure user expires in Microsoft Azure, a `Password Is Expired` error occurs when the next connection to Azure is requested.

  If the password expires while migration tasks are running, the tasks will fail with a `Password Is Expired` error when the next connection to Azure is requested.

To resolve failed migrations to Azure for password issues:

1. (Conditional) If the Azure user's password is expiring or has expired, log in to the user account in the Microsoft Azure portal, then set a new user password. See *Reset Your Work or School Password* (https://docs.microsoft.com/en-us/azure/active-directory/user-help/active-directory-passwords-update-your-own-password) in the Microsoft Azure Documentation.
2. Log in to the PlateSpin Migrate Web Interface, then go to the Targets page.
3. Update the password that is stored for the Azure user for any affected Azure target cloud platforms.
   - **3a** Click the name of a target platform to access the target platform settings, then click **Edit**.
   - **3b** Specify a valid password.
   - **3c** (Optional) Click **Test Credentials**.
   - **3d** Click **Save**.
4. Rerun any failed workload migrations to the affected Azure target cloud platforms.

Checklist for Automated Migration to Azure

Task	Description
1. Prepare your Azure account for Migrate.	"Configuring an Application in Azure to Represent PlateSpin Migrate" on page 182 "Enabling PlateSpin Replication Environment for Azure Subscriptions" on page 190 (Non-VPN deployment) "Deploying a Migrate Server Image in Azure" on page 193
2. Prepare your Azure migration environment.	Figure 9-1, "On-Premise Migrate Server for Automated Migration to Azure Cloud," on page 168 "Planning For Migrating Workloads to Azure" on page 179

3. Discover target cloud platform.	"Target Discovery in the Web Interface" on page 284
	
	IMPORTANT: (For Azure Stack only) The first-time discovery of an Azure Stack might take some considerable time (20-30 mins) because PlateSpin Migrate downloads the `PlateSpin Replication Environment(PRE)` image from the Azure Marketplace and uploads it to the Azure Stack. If the `PlateSpin Replication Environment` does not get successfully uploaded to Azure Stack, then the following messages are logged and you must manually upload the PRE image to the Azure Stack before you configure migrations to it.
	
	`Create UserImage has failed With Message: {0}`
	
	`Create a User Image manually in the portal with Name {0} before continuing with the migration`
	
	To manually upload the PRE image to the discovered Azure Stack, see "Manually Adding PlateSpin Replication Environment Image to Azure Stack" on page 193.

4. Discover source workloads.	"Workload Discovery in the Migrate Web Interface" on page 302 -OR- "Registering Workloads and Discovering Details with Migrate Agent" on page 303
5. Configure target workload migration.	"Configuring Migration of a Workload to Microsoft Azure" on page 454

Task	Description
6. Execute migration.	Chapter 33, “Executing Workload Migrations,” on page 537

10 Prerequisites for Migration to VMware vCloud Director

PlateSpin Migrate Web Interface supports automated migration to VMware vCloud Director environments. This section describes the required vCloud configuration that you must prepare in the appropriate environment, such as a vCloud Organization, before you can discover vCloud target cloud platforms and configure migrations to them.

Deployment for Migration to VMware vCloud

You can deploy a PlateSpin Migrate server on premise in your data center with the source workloads or in the appropriate VMware vCloud Organization.

For an on-premise Migrate server deployment, a site-to-site VPN connection is required between the data center and your account in the vCloud cloud. Figure 10-1 shows the location of various components in your vCloud migration environment and the communications between them. See "Planning For Migrating Workloads to VMware vCloud Director" on page 203.

NOTE: Figure 10-1 depicts automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443). See "Requirements for Workload Registration" on page 56 and "Registering Workloads and Discovering Details with Migrate Agent" on page 303.

Figure 10-1 *On-Premise Migrate Server for Automated Migration to vCloud*

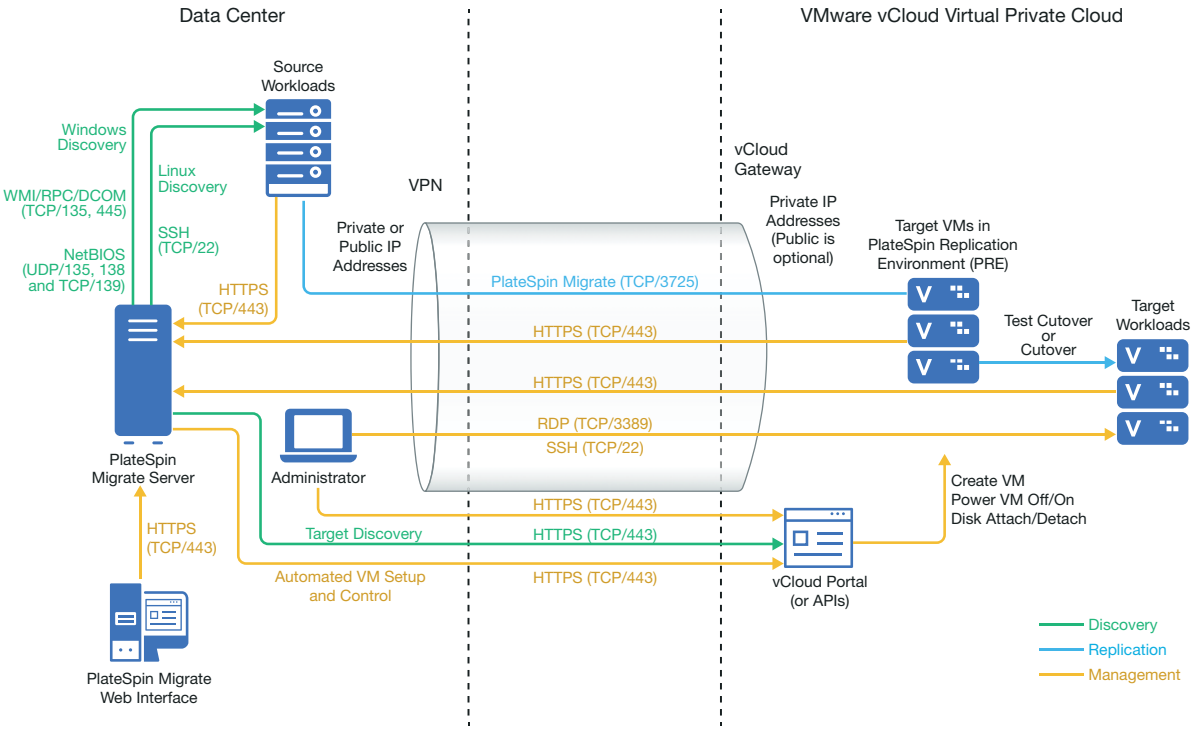

For a cloud-based Migrate server deployment, the PlateSpin Migrate Server is available. Figure 10-2 shows the location of various components in your vCloud migration environment and the communications between them. See “Planning For Migrating Workloads to VMware vCloud Director” on page 203.

Figure 10-2 *Cloud-Based Migrate Server for Automated Migration to vCloud*

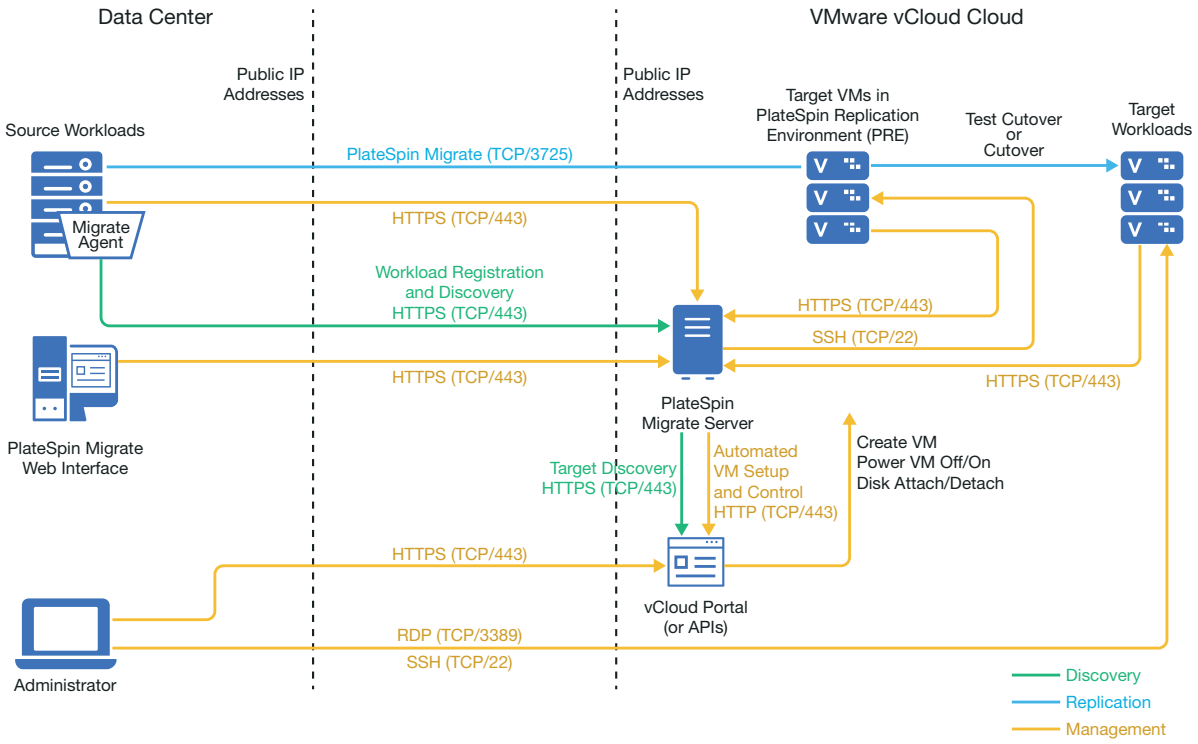

Planning For Migrating Workloads to VMware vCloud Director

- PlateSpin Migrate uses the VMware vCloud Director for migrating workloads to VMware vCloud. For a list of supported workloads, see "Supported Workloads For Migration to VMware vCloud Director" on page 31.
- Set up the vCloud Organization with at least the minimum required set of resources. See "Setting up vCloud Organization" on page 204.
- PlateSpin requires a replication environment to migrate workloads to a vCloud Organization. The replication environment is a virtual appliance based on a SLES operating system and contains all the required PlateSpin tools. It also contains a OVF PlateSpin Package that you must upload to the vCloud organization before you migrate workloads to a vCloud Organization. See "Understanding PlateSpin Replication Environment Used for Migration of Workloads to vCloud" on page 204.
- Migration of workloads with multiple NICs to vCloud is supported for Windows workloads up to the number of NICs supported by the VMware vcloud Director.

NOTE: When you configure migration for a workload, PlateSpin Migrate lets you specify four IP addresses per NIC on the Windows workload by default. However, you can configure the number of IP addresses that you want to specify per NIC on the workload. See "Configuring the Maximum Number of IP Addresses Allowed for Each Network Interface on Target Workloads" on page 111.

Setting up vCloud Organization

You must set up a vCloud Organization with at least the following minimum set of resources:

- ❒ Define one or more Organization Virtual Data Center (Org vDC).
- ❒ Define one or more Org vDC Network for the target VM.
- ❒ Create a private Catalog and grant full access permissions for the organization users to access the contents and settings of the catalog.
- ❒ Use Administrator level credentials for discovering the vCloud Organization and performing migrations.
- ❒ Define policies that apply to the target VMs in your Org vDC and ensure the following:
  - ❒ The lease period of the vCloud Organization resources should not expire during migration.
  - ❒ No restrictions are set on VM quota.
  - ❒ No restrictions are set on the number of connections or operations to the vCloud organization.
- ❒ The VDC Hardware Version policy limits the maximum hardware version for VMs that Migrate will create for the vCloud platform.
- ❒ Migration of Windows Server 2016 workloads to vCloud 9.1 requires Hardware Version 10 or higher support by the underlying VMware platform. The Hardware Version policy for the VDC must be set to at least Hardware Version 10.

NOTE: During the Test Cutover, there is cloning for the target VM that consumes twice the storage resources than you need for Cutover. Ensure that the storage quotas for the Org vDC support that need. The additional resources used is temporary, and will be released after the Test Cutover.

For more information, see VMware vCloud Director Documentation (https://www.vmware.com/support/pubs/vcd_pubs.html).

Understanding PlateSpin Replication Environment Used for Migration of Workloads to vCloud

PlateSpin requires a replication environment to migrate workloads to a vCloud Organization. The replication environment is a virtual appliance based on a SLES operating system and contains all the required PlateSpin tools. It also contains a OVF PlateSpin Package that you must upload to the vCloud organization before you migrate workloads to a vCloud Organization.

The following PREs are available on Micro Focus Download site:

Name	Description
`PlateSpin_Replication_Environment-<x>.zip` where <*x*> is the product release version.	This replication environment is a virtual appliance based on a SLES 11 operating system and is required for migration of 32-bit BIOS workloads to vCloud.
`PlateSpin_Replication_Environment_BIOS_x64-<x>.zip` where <*x*> is the product release version.	This replication environment is a virtual appliance based on a SLES 12 operating system and is required for migration of 64-bit BIOS workloads to vCloud.
`PlateSpin_Replication_Environment_UEFI-<x>.zip` where <*x*> is the product release version.	This replication environment is a virtual appliance based on a SLES 12 operating system and is required for migration of UEFI workloads to vCloud.

Depending on whether you want to migrate 32-bit BIOS, 64-bit BIOS, or UEFI workloads to vCloud, you are required to upload the corresponding PlateSpin Replication Environment OVF Package to the vCloud organization. You can download this package from the Micro Focus Download site for this PlateSpin Migrate release. For more information about downloading the OVF package and uploading to vCloud, see "Creating the PlateSpin Virtual Appliance in the vCloud Organization" on page 206.

Review the following sections:

- "Resources Used in the PlateSpin Replication Environment" on page 205
- "Creating the PlateSpin Virtual Appliance in the vCloud Organization" on page 206

Resources Used in the PlateSpin Replication Environment

PlateSpin uses the following minimum resources for the Replication Environment Virtual Machine:

Hardware Resource	Details
Virtual CPUs	1
Cores Per Socket	1
RAM	1 GB
Disk	4 GB (for 32-bit and BIOS PREs) 7 GB (for UEFI PRE)
Network Adapter of type E1000	1
Virtual Hardware Version	7 (for 32-bit PRE) 9 (for BIOS and UEFI PREs)

Creating the PlateSpin Virtual Appliance in the vCloud Organization

1. Ensure that you have set up a vCloud Organization with at least the minimum set of resources. See "Setting up vCloud Organization" on page 204.
2. Download one of the following PlateSpin Replication Environment file from the Micro Focus Download site (https://www.microfocus.com/support-and-services/download/) for this PlateSpin Migrate release, depending on whether you want to migrate 32-bit BIOS, 64-bit BIOS, or UEFI workloads:
   - **`PlateSpin_Replication_Environment-<x>.zip`**: For migration of 32-bit BIOS workloads.
   - **`PlateSpin_Replication_Environment_BIOS-<x>.zip`:** For migration of 64-bit BIOS workloads.
   - **`PlateSpin_Replication_Environment_UEFI-<x>.zip`:** For migration of UEFI workloads.
3. Unzip the `.zip` file that you downloaded and extract the contents to a temporary directory. For example, `C:\PlateSpin_Replication_Environment`.
4. Use the vCloud Director Web Console to upload the OVF PlateSpin package, which you extracted in the previous step, as a vApp Template to a Catalog, such as `PlateSpin Catalog`. Sample listing of the replication environment in the vCloud Director Web console is as follows:
   - Catalogs
     - *<Catalog Name>*
   - vApp Templates
     - PlateSpin Replication Environment
     - PlateSpin Replication Environment - UEFI
   - VMs
     - PlateSpin Virtual Appliance

Configuring Advanced PlateSpin Settings for vCloud

Some aspects of your PlateSpin Server behavior are controlled by configuration parameters that you set on a PlateSpin Configuration page (`https://Your_PlateSpin_Server/PlateSpinConfiguration/`) using the Migrate Web Interface. You can also access this page from the Web Interface, by selecting **Settings** > **Advanced Server Settings**, then click the link provided. On the PlateSpin Server Configuration Settings page, search for the parameter of interest.

- "Configuring vCloud vApp Template Name Used for Replication Environment" on page 207
- "Retaining the Cloud Resources For Troubleshooting Migration Errors" on page 207
- "Configuring a Custom Password for the PlateSpin Replication Environment for vCloud" on page 207

Configuring vCloud vApp Template Name Used for Replication Environment

The **VCloudAppTemplateName** parameter in PlateSpin Configuration lets you configure the name of the vApp template used for the Replication Environment during vCloud replications.

By default, the value of this parameter is **PlateSpin Replication Environment**. However, if you have edited the name of the vApp Template to which you uploaded the OVF PlateSpin package, then you must set the value of the **VCloudAppTemplateName** parameter to the new name of the vApp Template.

Retaining the Cloud Resources For Troubleshooting Migration Errors

When an error occurs during a migration, cloud resources are either deleted or retained based on the setting for the **LeaveCloudResourcesOnError**parameter in PlateSpin Configuration.

By default, this parameter is set to `False` and PlateSpin deletes the target VM and its associated resources if there is an error during migration. If you need PlateSpin to retain these resources for troubleshooting and do not want to delete them, set the **LeaveCloudResourcesOnError**setting to `True`.

Configuring a Custom Password for the PlateSpin Replication Environment for vCloud

The **vCloudReplicationEnvironmentPassword** parameter in PlateSpin Configuration lets you configure a custom password for the `root` user of the PlateSpin Replication Environment (PRE) for vCloud.

By default, the `root` user password required to access the PlateSpin Replication Environment is randomly generated and encrypted. To troubleshoot replication failures, you must access the PRE as the `root` user. Use the **vCloudReplicationEnvironmentPassword** parameter to set a custom password that overrides the default password. You can then access the PlateSpin Replication Environment as a `root` user with the newly set password.

Checklist for Automated Migration to vCloud

Task	Description
1. Prepare your vCloud migration environment.	Figure 10-1, “On-Premise Migrate Server for Automated Migration to vCloud,” on page 202 Figure 10-2, “Cloud-Based Migrate Server for Automated Migration to vCloud,” on page 203 “Planning For Migrating Workloads to VMware vCloud Director” on page 203
2. Discover target cloud platform.	“Target Discovery in the Web Interface” on page 284

Task	Description
3. Discover source workloads.	“Workload Discovery in the Migrate Web Interface” on page 302 -OR- “Registering Workloads and Discovering Details with Migrate Agent” on page 303
4. Configure target workload migration.	“Configuring Migration of a Workload to VMware vCloud Director” on page 470
5. Execute migration.	Chapter 33, “Executing Workload Migrations,” on page 537

11 Prerequisites for Migration to Oracle Cloud Infrastructure

PlateSpin Migrate supports semi-automated migration of Windows and Linux workloads to Compute - Virtual Machine Instances on your Oracle Cloud Infrastructure environment. This section describes the required Oracle Cloud configuration that you must prepare, such as an Oracle Cloud account, before you can register and discover details for Oracle Cloud target VMs and configure migrations to them.

Deployment for Migration to Oracle Cloud Infrastructure

You must manually install a PlateSpin Migrate server with a public IP address on a virtual host on the Oracle Cloud Infrastructure. See "Installing a PlateSpin Migrate Server in Oracle Cloud" on page 214.

Figure 8-1 shows the location of various components in your semi-automated migration environment and the communications between them. See "Prerequisites for Using a Migrate Server on Oracle Cloud" on page 211.

Figure 11-1 *Cloud-Based Migrate Server for Semi-Automated Migration to Oracle Cloud*

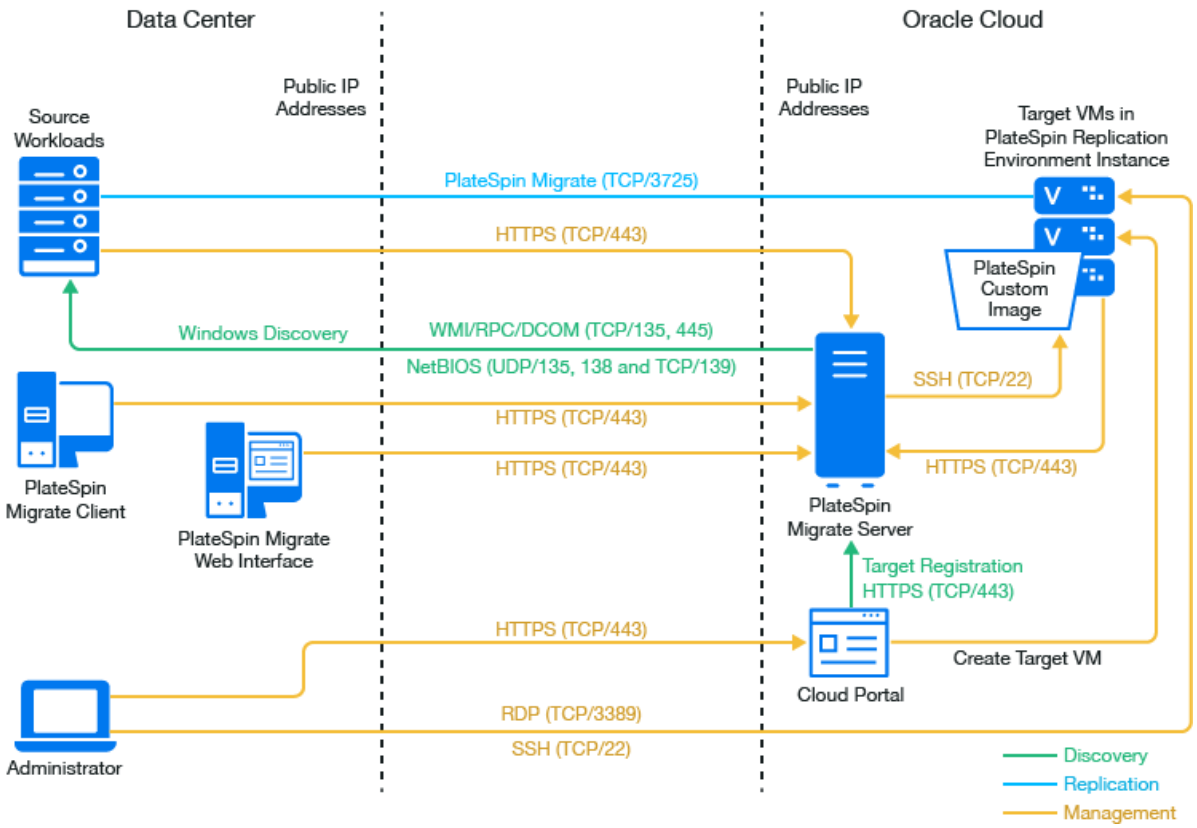

Requirements for Migrating Workloads to Oracle Cloud Infrastructure

Before you use PlateSpin Migrate to migrate workloads to Oracle Cloud Infrastructure, you must set up your cloud environment. Create a Windows Instance in Oracle Cloud Infrastructure and manually install a PlateSpin Migrate server with a public IP address on that instance. For information about manually installing the Migrate server on Oracle Cloud, see "Installing PlateSpin Migrate Server in Oracle Cloud Infrastructure" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

- "Minimum Oracle Cloud Prerequisites" on page 211
- "Prerequisites for Using a Migrate Server on Oracle Cloud" on page 211

Minimum Oracle Cloud Prerequisites

Before you use PlateSpin Migrate to migrate workloads to Oracle Cloud Infrastructure, you must ensure that you have a Oracle Cloud Account with all the required permissions for launching Oracle Cloud Compute Instances. See Oracle Cloud Infrastructure Documentation (https://docs.cloud.oracle.com/iaas/Content/Compute/Tasks/launchinginstance.htm).

Prerequisites for Using a Migrate Server on Oracle Cloud

Before you use PlateSpin Migrate to migrate workloads to Oracle Cloud Infrastructure, ensure that the following prerequisites are correctly configured and available:

- A PlateSpin Migrate license.
- Create a Windows instance in the target Oracle Cloud Infrastructure environment and install a Migrate server with a public IP address. See "Installing PlateSpin Migrate Server in Oracle Cloud Infrastructure" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

  NOTE: The cloud-based Migrate server does not require a site-to-site VPN connection between your local data center and Oracle Cloud Portal. When no VPN is provided between the source network and the cloud-based Migrate server, a public IP address is required for the source workload to register the workload with the cloud-based Migrate server using secure communications over the public Internet. For deployment information, see Figure 11-1, "Cloud-Based Migrate Server for Semi-Automated Migration to Oracle Cloud," on page 210.

- You can use PlateSpin Migrate Web Interface (recommended) or PlateSpin Migrate Client to migrate workloads to Oracle Cloud instances.

  You can install PlateSpin Migrate Client on the PlateSpin Migrate server or on another machine in the same network as the PlateSpin Migrate server. Migrate Client must be able to connect to the PlateSpin Migrate server.
- For semi-automated migrations to Oracle Cloud, you must set **SourceListensForConnection** to `True`. See "Configuring the Contact Direction for the Replication Port" on page 114.
- To enable the cloud-based PlateSpin Server to function across NAT-enabled environments and discover source workloads by using secure communications over the public Internet, you must set the AlternateServerAddress PlateSpin Configuration parameter to the public IP addresses of the Migrate server. For information about configuring the AlternateServerAddress parameter, see "Configuring Alternate IP Addresses for PlateSpin Server" on page 111.

  IMPORTANT: You must set the AlternateServerAddress PlateSpin Configuration parameter to the public IP addresses of the Migrate server only after you prepare the target VM Instance in the Oracle Cloud Infrastructure and use PlateSpin Custom Image to register it with PlateSpin Migrate Server. For information about registering the target VM instance and configuring migration of workload to Oracle Cloud, see "Configuring Full Migration of a Workload to Oracle Cloud VM Instance Using the X2P Workflow" on page 482.

- Allocate a Public IP address for the Migrate server to ensure that the IP address does not change when the server is restarted.

  NOTE: A change in IP address on the PlateSpin Server breaks the heartbeat communications with source workloads.

- A Security Group that provides the following inbound and outbound rules.

Table 11-1 *Port Requirements for Migrate Server in Oracle Cloud infrastructure*

Location	Port	Protocol	Remarks
Source workload Network firewall	TCP 443, outbound	HTTPS	When no VPN is provided between the source network and the cloud-based Migrate server, a public IP address is required for the source workload to allow the Migrate server to discover the workload using secure communications over the public Internet. The PlateSpin Migrate server must be able to communicate with the source workloads on the ports that are used for discovery. See "Requirements for Discovery" on page 53 and "Discovering Details for Source Workloads" on page 301
Source workload Network firewall	TCP 3725/custom, outbound	Migrate	The source workload must be able to connect to the cloud-based target workload on TCP port 3725. The PlateSpin Migrate Server uses secure SSL for communications with the workloads you want to migrate. Port 3725 is the default port number for data transfer. By default, the data transfer is initiated from the target workload to the source workload. The port number and direction for initiating the connection are configurable. For information about changing the default port setting, see port 3725 in "Requirements for Migration" on page 57.

Location	Port	Protocol	Remarks
Security Group for the Migrate Server in Oracle Cloud	TCP 3725/custom, inbound and outbound	Migrate	Provide an address range covering all source workloads.
	TCP 22, outbound	SSH (Linux)	This port allows outbound communications from the Migrate server to target workloads. Provide the IP address of the PlateSpin Migrate server.
	TCP 3389, inbound	RDP (Windows)	Allow inbound connections in the AWS Security Group for the cloud-based Migrate server. Provide the IP address of the machine you plan to use to launch an RDP connect to target workloads.
	TCP 443, inbound	HTTPS	Provide the IP address of the PlateSpin Migrate server. The source and the target workload must be able to communicate with the PlateSpin Migrate server on port 443. The PlateSpin Migrate Client and the PlateSpin Migrate Web Interface must be able to connect to the PlateSpin Migrate Server over HTTPS (TCP/443).
	TCP 123, outbound	Network Time Protocol (NTP)	Add this port setting to the security group if you are using an NTP service outside the virtual network where you deploy the Migrate server.

Planning For Migrating Workloads to Oracle Cloud Infrastructure

PlateSpin Migrate allows you to migrate workloads to Oracle Cloud infrastructure. For a list of supported workloads, see "Supported Workloads For Migration to Oracle Cloud Infrastructure" on page 34.

NOTE: Migration of Windows Cluster workloads to Oracle Cloud Infrastructure is not supported.

PlateSpin Server Host

- 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 access error.

Installing a PlateSpin Migrate Server in Oracle Cloud

You must manually install Migrate server on your own virtual host in Oracle Cloud. See "Installing PlateSpin Migrate Server in Oracle Cloud Infrastructure" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

Importing PlateSpin Custom Images to Oracle Cloud Infrastructure

PlateSpin Migrate provides customized image that you can import from the Oracle Cloud Object Storage provided you have read access to the Object Storage object containing the image. You use this image to create a target VM instance (Windows or Linux) on Oracle Cloud to which you can migrate the source workload.

For information about importing the PlateSpin image from the Object Storage, see Importing an Image in Oracle Cloud Infrastructure Documentation.

NOTE: To import the PlateSpin custom image:

- In the **OBJECT STORAGE URL** option, specify the following URL that contains the PlateSpin custom image:

  ```
  https://objectstorage.us-ashburn-1.oraclecloud.com/p/
  kGD8AQf5DraYKVCvhAgqjj_Fu6FMMpgLLXMKLuj8ZQ0/n/platespin/b/platespin/o/
  Platespin_Replication_Evironment_2019_8.vmdk
  ```

- Set the **IMAGE TYPE** as **VMDK** and the **LAUNCH MODE** as **EMULATED MODE**.

Checklist for Semi-Automated Migration to Oracle Cloud Infrastructure

Task	Description
1. Prepare your Oracle Cloud migration environment.	Figure 11-1, "Cloud-Based Migrate Server for Semi-Automated Migration to Oracle Cloud," on page 210 "Planning For Migrating Workloads to Oracle Cloud Infrastructure" on page 213
2. Discover target containers.	"Creating and Connecting to the Target Virtual Machine Instance in Oracle Cloud Infrastructure Environment" on page 482
3. Discover source workloads.	"Workload Discovery in the Migrate Client" on page 301 -OR- "Workload Discovery in the Migrate Web Interface" on page 302

Task	Description
4. Configure target workload migration.	“Configuring Full Migration of a Workload to Oracle Cloud VM Instance Using the X2P Workflow” on page 482
5. Execute migration.	Chapter 33, “Executing Workload Migrations,” on page 537

12 Prerequisites for Cloud-to-Cloud Migrations

PlateSpin Migrate Web Interface supports automated cloud-to-cloud (C2C) migration of workloads. For migrations using a cloud-based PlateSpin Migrate Server and public IP addresses, Migrate does not require site-to-site VPN connections between any of the participating locations: source cloud, target cloud, and data center.

To plan your cloud-to-cloud migrations, use the following information about supported C2C deployment scenarios, required configurations, and checklists for migration.

Requirements for C2C Non-VPN Migrations

A cloud-based PlateSpin Migrate server does not require a site-to-site VPN connection between your local data center and the target cloud platform.

To use a cloud-based Migrate server without a VPN:

- Internet access is required.
- Deploy a Migrate server in the source cloud or target cloud, as appropriate for your deployment scenario.
  - **For Azure:** Use the Azure Marketplace template.
  - **For AWS:** Use the AWS Quick Start.
  - **For Oracle Cloud Infrastructure:** Create a Windows instance in the target Oracle Cloud Infrastructure environment and install a Migrate server with a public IP address.

- **For VMware vCloud:** Create a Windows instance in the target vCloud environment and install a Migrate server with a public IP address.

See "Deploying PlateSpin Migrate Server in the Cloud" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

- Public IP addresses are required for the PlateSpin Migrate server, the replication network, and target machines. A public IP address is not required for the source machine when you use the Migrate Agent. If you do not use the Migrate Agent, then all components need public IP addresses.
- In the PlateSpin Configuration settings on the cloud-based Migrate server:
  - **AlternateServerAddress:** Set the **AlternateServerAddress** parameter to the public IP address of the Migrate server. For Migrate servers deployed from Azure Marketplace or using the AWS Quick Start, Migrate automatically adds the public IP address to this parameter. See "Configuring Alternate IP Addresses for PlateSpin Server" on page 111.
  - **SourceListensForConnection:** Ensure that the **SourceListensForConnection** parameter is set to `False`, which is the default value. See "Configuring the Contact Direction for the Replication Port" on page 114.
- (Migrate Discovery) If the Migrate server is in the same cloud network as the source workloads, you can use Migrate discovery to add workloads to the Migrate server. Ensure that your network security groups for the source network and target network allow traffic for ports required for discovery and migration. See:
  - "Requirements for Discovery" on page 53.
  - "Requirements for Migration" on page 57.
- (Migrate Agent registration) If the Migrate server is in the target cloud network, ensure that your network security groups for the source network and target network allow traffic for ports required for registration with Migrate Agent and migration over the public Internet. You might also use Migrate Agent to register workloads if the Migrate server is in a different network security group than the source workloads, or if you do not want to enable discovery ports on source workloads. See:
  - "Requirements for Workload Registration" on page 56.
  - "Requirements for Migration of Workloads Registered Using Migrate Agent" on page 58.
- When you configure a workload migration:
  - Enable a public IP address for the replication network.
  - Ensure that you enable **Encrypt Data Transfer** to transfer data securely between the source workload in AWS and the PlateSpin Replication Environment in vCloud over the public Internet. See "Encrypt Data Transfer Using Migrate Web Interface" on page 396.
- (Migrate Agent) Install the Migrate Agent on the source workload, then register the workload with the cloud-based PlateSpin Migrate server. See "Registering Workloads and Discovering Details with Migrate Agent" on page 303.

  To download the Migrate Agent, launch the PlateSpin Migrate Web Interface and click the **Downloads** tab. For information about installing and using the Migrate Agent, see "Migrate Agent Utility" on page 359.

Prerequisites for C2C Migration from AWS to AWS

PlateSpin Migrate supports migration of workloads between AWS regions in the same AWS EC2 subscription.

- **PlateSpin Migrate Server**
  - Deploy the PlateSpin Migrate server in the target AWS region.

    See "Deploying PlateSpin Migrate Server in the Cloud" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.
  - Internet access and a public IP address are required for the Migrate Server.
  - No VPN is required between the participating regions.
  - A Migrate license is required to perform the migrations.
- **Target AWS Region**

  In the Migrate Web Interface, add a region for the same AWS subscription that contains the source workload as the target location. See "Target Discovery in the Web Interface" on page 284 and Table 17-4, "Options for Amazon Cloud Region," on page 285.
- **Source Workload**
  - In AWS, configure the required ports in the source workload's security group on the NIC that you will use to discover details for the source workload and migrate its data.

    See "Access and Communication Requirements across Your Migration Network" on page 53.
  - Discover source workloads by using their public IP addresses.
  - Enable the **Public IP** option for the Replication Environment.
  - Enable the **Public IP** option for the primary IP address for Cutover and Test Cutover.

    NOTE: If you have a VPN between your site and the target location, it is not necessary to use a public IP address for the cutover VM and test cutover VM.

Prerequisites for C2C Migration from Azure to Azure

PlateSpin Migrate supports migration of workloads between Azure regions in the same Azure subscription.

- **PlateSpin Migrate Server**
  - Deploy the PlateSpin Migrate server in the target Azure region.

    See "Deploying PlateSpin Migrate Server in the Cloud" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.
  - Internet access and a public IP address are required for the Migrate Server.
  - No VPN is required between the participating regions.
  - A Migrate license is required to perform the migrations.

- **Target Azure Location**

  In the Migrate Web Interface, add the region as the target location for the same Azure subscription that contains the source workload. See "Target Discovery in the Web Interface" on page 284 and Table 17-5, "Options for Microsoft Azure Location Target," on page 286.

- **Source Workload**
  - In Azure, configure the required ports in the source workload's security group on the NIC that you will use to discover details for the source workload and migrate its data.

    See "Access and Communication Requirements across Your Migration Network" on page 53.
  - Discover source workloads by using their public IP addresses.
  - Enable the **Public IP** option for the Replication Environment.
  - Enable the **Public IP** option for the primary IP address for Cutover and Test Cutover.

  NOTE: If you have a VPN between your site and the target location, it is not necessary to use a public IP address for the cutover VM and test cutover VM.

Prerequisites for C2C Migration from AWS to Azure

PlateSpin Migrate supports migration of workloads from Amazon Web Services EC2 Cloud to Microsoft Azure Cloud.

Deployment for C2C Migration from AWS to Azure

For migration of workloads from Amazon Web Services EC2 Cloud to Microsoft Azure Cloud, deploy the PlateSpin Migrate server in the target Azure environment. No VPN is required between the participating sites. Internet access and public IP addresses are required. Figure 12-1 shows the location of various components in your AWS, Azure, and data center migration environments and the communications between them.

You must also enable the application use of PlateSpin Replication Environment from the Azure Marketplace in the target Azure environment.

You use Migrate Agent to register workloads with the cloud-based Migrate server using secure communications over the public Internet. Enable data transfer encryption to transfer data securely between the source workload in AWS and the PlateSpin Replication Environment in Azure over the public Internet.

NOTE: A reboot of the source Windows workload is required when you install, uninstall, or upgrade block-based transfer drivers. A reboot is not required for source Linux workloads.

Figure 12-1 *Cloud-Based Migrate Server for Automated Migration from AWS to Azure with No VPNs*

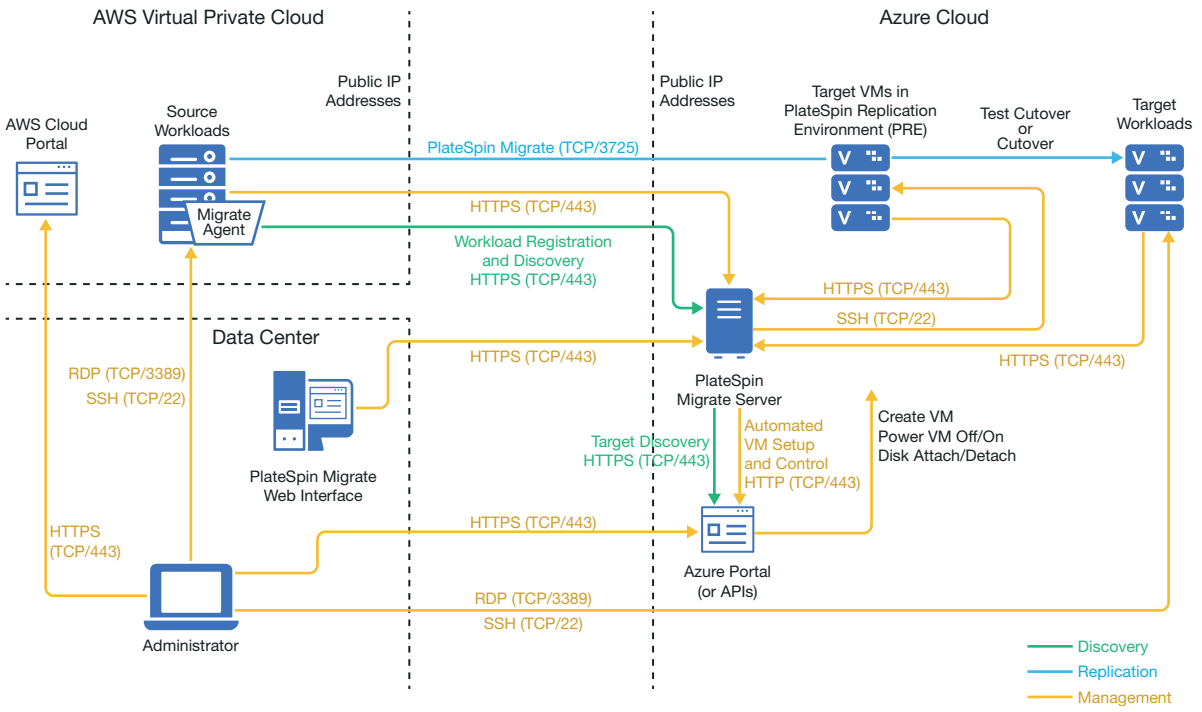

Requirements for Migrating Workloads to Azure

To prepare your target Azure environment, review the following information in "Requirements for Migrating Workloads to Azure" on page 171:

- "Minimum Azure Prerequisites"
- "Azure Prerequisites for Using an Azure-Based Migrate Server"

Ensure that the source workload is supported by the target Azure configuration.

Requirements for Migrating Workloads from AWS to Azure

Deploy a PlateSpin Migrate server in the target Azure network environment. Ensure that your non-VPN migration environment meets the "Requirements for C2C Non-VPN Migrations" on page 217.

For source workloads in AWS:

- AWS automatically adds the Remote Desktop Protocol (RDP) port (TCP/3389) and Secure Shell (SSH) port (TCP/22) in the AWS Security Group for the source workload VMs. You must manually add other ports to the source workload's AWS Security Group that are required by PlateSpin Migrate to provide migration services, such as Port 3725 for replication traffic and Port 443 for HTTPS traffic.
- For Windows workloads, use a user name and password.

- For Linux workloads, use the root user or root equivalent user.

  In AWS, Amazon Linux AMIs by default enable the `ec2user` user name and PEM key credentials, and disable the `root` user name and password credentials. To use Migrate discovery to inventory workloads, you must enable `root` user access for the AWS source Linux workload. See "Enabling Root User Credentials for Source Linux Workloads in AWS" on page 242.

Checklist for Automated Migration from AWS to Azure

Task	Description
1. Prepare your network resources.	Figure 12-1, "Cloud-Based Migrate Server for Automated Migration from AWS to Azure with No VPNs," on page 221 "Deployment for C2C Migration from AWS to Azure" on page 220
2. Prepare your Azure migration environment.	"Requirements for Migrating Workloads to Azure" on page 221
3. Prepare your AWS source workloads for PlateSpin Migrate.	"Requirements for Migrating Workloads from AWS to Azure" on page 221
4. Discover target cloud platform.	"Target Discovery in the Web Interface" on page 284
5. Register source workloads with the cloud-based Migrate server by using Migrate Agent.	"Registering Workloads and Discovering Details with Migrate Agent" on page 303
6. Configure target workload migration.	"Configuring Migration of a Workload to Microsoft Azure" on page 454
7. Execute migration.	Chapter 33, "Executing Workload Migrations," on page 537

Prerequisites for C2C Migration from Azure to AWS

PlateSpin Migrate supports migration of workloads from Microsoft Azure Cloud to Amazon Web Services EC2 Cloud.

Deployment for C2C Migration from Azure to AWS

For migration of workloads from Microsoft Azure Cloud to Amazon Web Services EC2 Cloud, you can deploy a cloud-based PlateSpin Migrate server in Azure or in AWS.

Migrate Server in Azure

Deploy PlateSpin Migrate server from the Azure Marketplace in the source Azure environment. The Migrate server image in Azure Marketplace is preconfigured to support its host Azure Cloud environment. When the Migrate Server and source workloads are in the same network security group, you can use Migrate discovery to add workload details to Migrate. Figure 12-2 shows the location of various components in your AWS, Azure, and data center migration environments and the communications between them.

NOTE: Figure 12-2 depicts source workloads and the Migrate server in the same network security group. If they are in different security groups, use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443). See "Requirements for Workload Registration" on page 56 and "Registering Workloads and Discovering Details with Migrate Agent" on page 303.

Figure 12-2 *Migrate Server in Azure for Automated Migration from Azure to AWS with No VPNs*

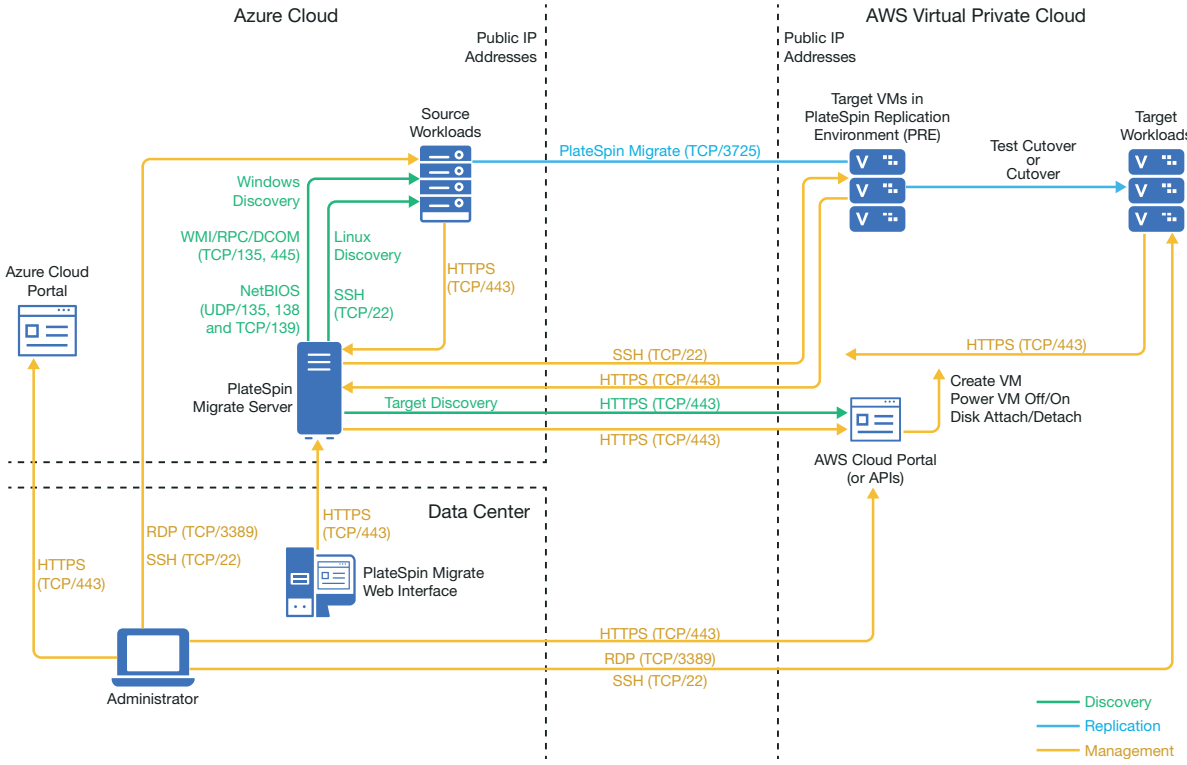

Migrate Server in AWS

Deploy PlateSpin Migrate server in the target AWS environment. See "Deploying a Migrate Server Image in AWS" on page 155.

You use Migrate Agent to register workloads with the cloud-based Migrate server using secure communications over the public Internet. Internet access and public IP addresses are required. Figure 12-3 show the location of various components in your AWS, Azure, and data center migration environments and the communications between them

NOTE: A reboot of the source Windows workload is required when you install, uninstall, or upgrade block-based transfer drivers. A reboot is not required for source Linux workloads.

Enable data transfer encryption to transfer data securely between the source workload in Azure and the PlateSpin Replication Environment in AWS over the public Internet.

Figure 12-3 *Migrate Server in AWS for Automated Migration from Azure to AWS with No VPNs*

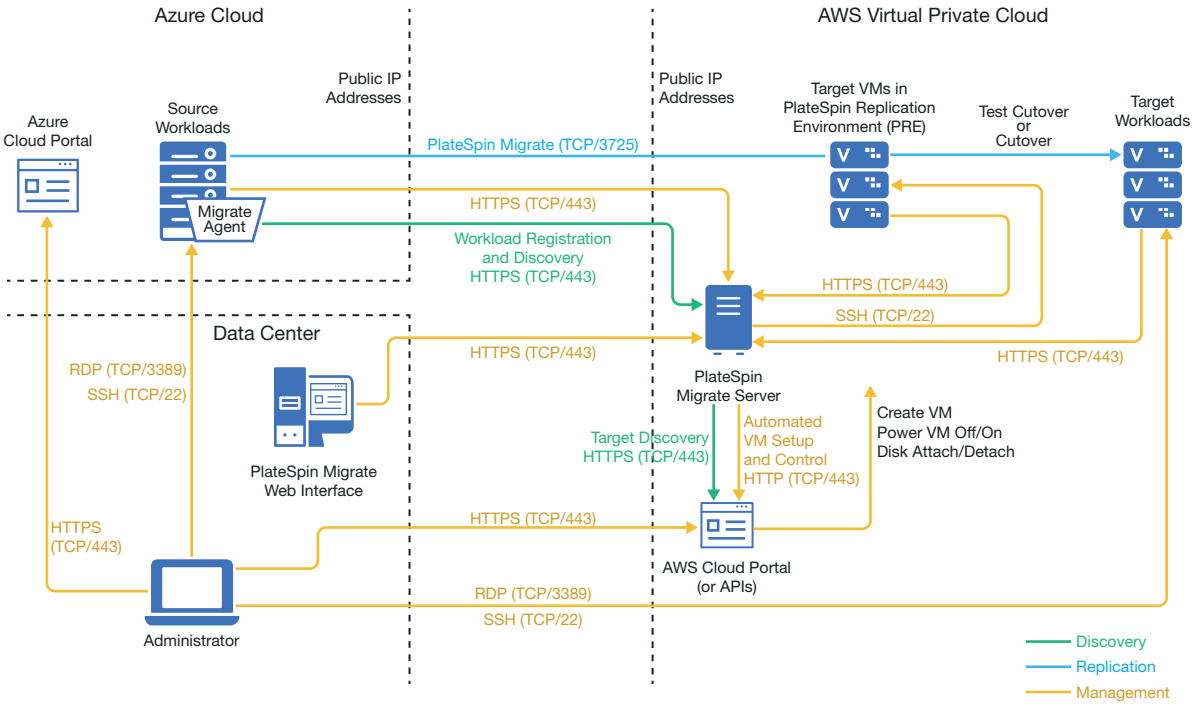

Requirements for Migrating Workloads to AWS

To prepare your target AWS environment, review the following information in "Requirements for Migrating Workloads to Amazon Web Services" on page 147:

- "Minimum AWS Prerequisites" on page 147
- "AWS Prerequisites for Using an AWS-Based Migrate Server" on page 150

Ensure that the source workload is supported by the target AWS configuration.

Requirements for Migrating Workloads from Azure to AWS

Deploy a PlateSpin Migrate server in the source Azure network environment or the target AWS network environment. Ensure that your non-VPN migration environment meets the "Requirements for C2C Non-VPN Migrations" on page 217.

Ensure that your migration environment meets these additional requirements:

- In the PlateSpin Configuration settings on the Migrate server:
  - **(Migrate Server in Azure) ServerIsHostedInCloud:** Remove the value of `azure` from the **ServerIsHostedInCloud** parameter to enable the Add Target dialog to provide all target types for selection. When you set up the AWS target, select **Amazon Cloud Region** as the target type.
- Azure automatically adds the Remote Desktop Protocol (RDP) port (TCP/3389) and Secure Shell (SSH) port (TCP/22) in the Azure Security Group for the source workload VMs. You must manually add other ports to the source workload's Azure Security Group that are required by PlateSpin Migrate to provide migration services, such as Port 3725 for replication traffic and Port 443 for HTTPS traffic.
- For information about workload login requirements for migration, see the Windows and Linux source workload login requirements in Table 18-2, "Guidelines for Machine Type and Credentials for Source Workloads," on page 299.

Checklist for Automated Migration from Azure to AWS

Task	Description
1. Prepare your network resources.	Figure 12-2, "Migrate Server in Azure for Automated Migration from Azure to AWS with No VPNs," on page 223 Figure 12-3, "Migrate Server in AWS for Automated Migration from Azure to AWS with No VPNs," on page 224 "Deployment for C2C Migration from Azure to AWS" on page 222
2. Prepare your AWS migration environment.	"Requirements for Migrating Workloads to AWS" on page 224
3. Prepare your Azure source workloads for PlateSpin Migrate.	"Requirements for Migrating Workloads from Azure to AWS" on page 224
4. Discover target cloud platform.	"Target Discovery in the Web Interface" on page 284
5. Discover source workloads. You can optionally register source workloads with the cloud-based Migrate server in AWS using Migrate Agent.	"Workload Discovery in the Migrate Web Interface" on page 302 -OR- "Registering Workloads and Discovering Details with Migrate Agent" on page 303
6. Configure target workload migration.	"Configuring Migration of a Workload to Amazon Web Services" on page 432
7. Execute migration.	Chapter 33, "Executing Workload Migrations," on page 537

Prerequisites for C2C Migration from Azure to vCloud

PlateSpin Migrate supports migration of workloads from Microsoft Azure to VMware vCloud Director.

Deployment for C2C Migration from Azure to vCloud

For migration of workloads from Microsoft Azure to VMware vCloud Director, deploy a PlateSpin Migrate server on premise in your source network. With an on-premise Migrate server, site-to-site VPN gateways are required between the data center and Azure and between the data center and vCloud. Figure 12-4 shows the location of various components in your Azure, vCloud, and data center migration environments and the communications between them.

Figure 12-4 *Migrate Server on Premise for Migration from Azure to vCloud*

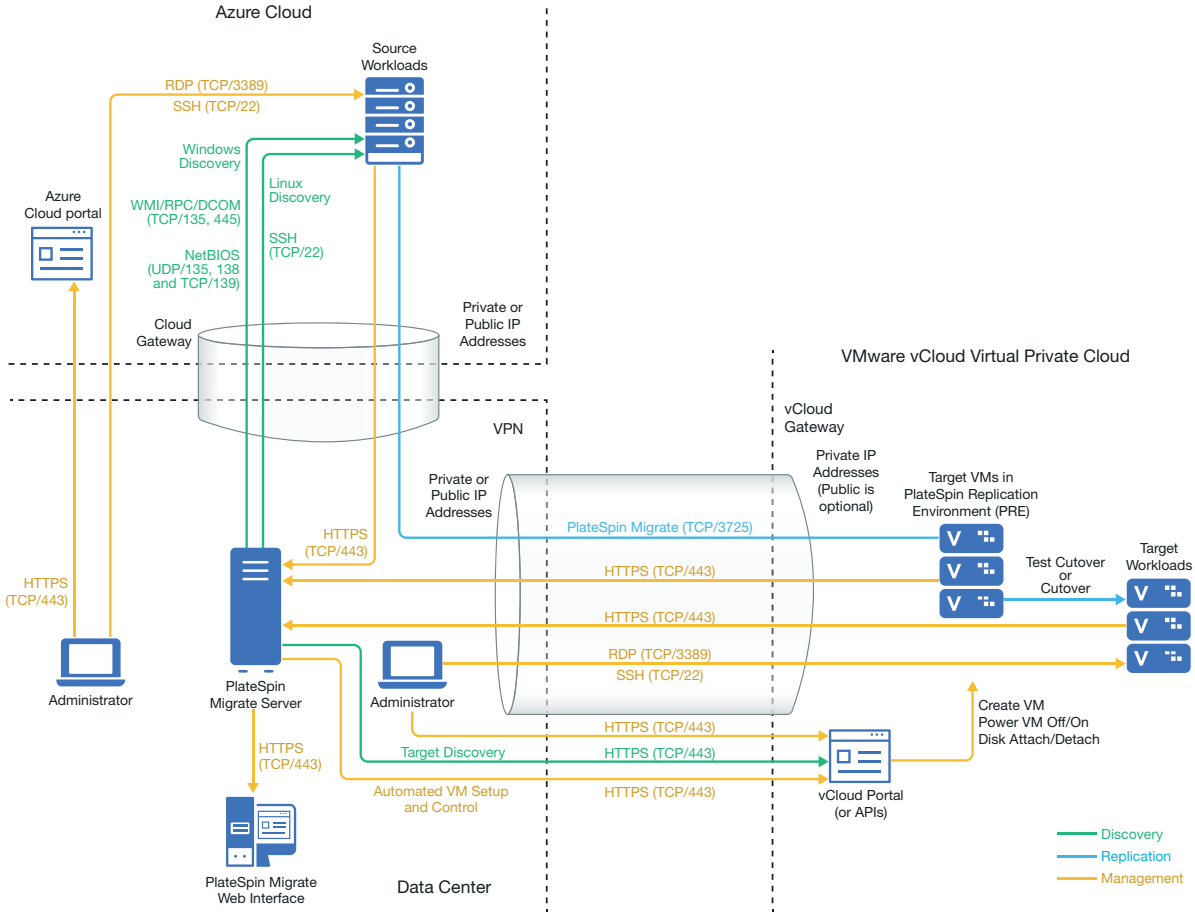

You can alternatively deploy the PlateSpin Migrate server from the Azure Marketplace in the source Azure environment. No VPN is required. With the Azure server in the same network security group as the source workloads, you can use discovery to add workloads to Azure. Use data encryption to secure data for replications over the public Internet. Figure 12-5 shows the location of various components in your Azure, vCloud, and data center migration environments and the communications between them.

NOTE: Figure 12-5 depicts source workloads and the Migrate server in the same network security group. If they are in different security groups, use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443). See "Requirements for Workload Registration" on page 56 and "Registering Workloads and Discovering Details with Migrate Agent" on page 303.

Figure 12-5 *Migrate Server in Azure for Migration from Azure to vCloud with No VPNs*

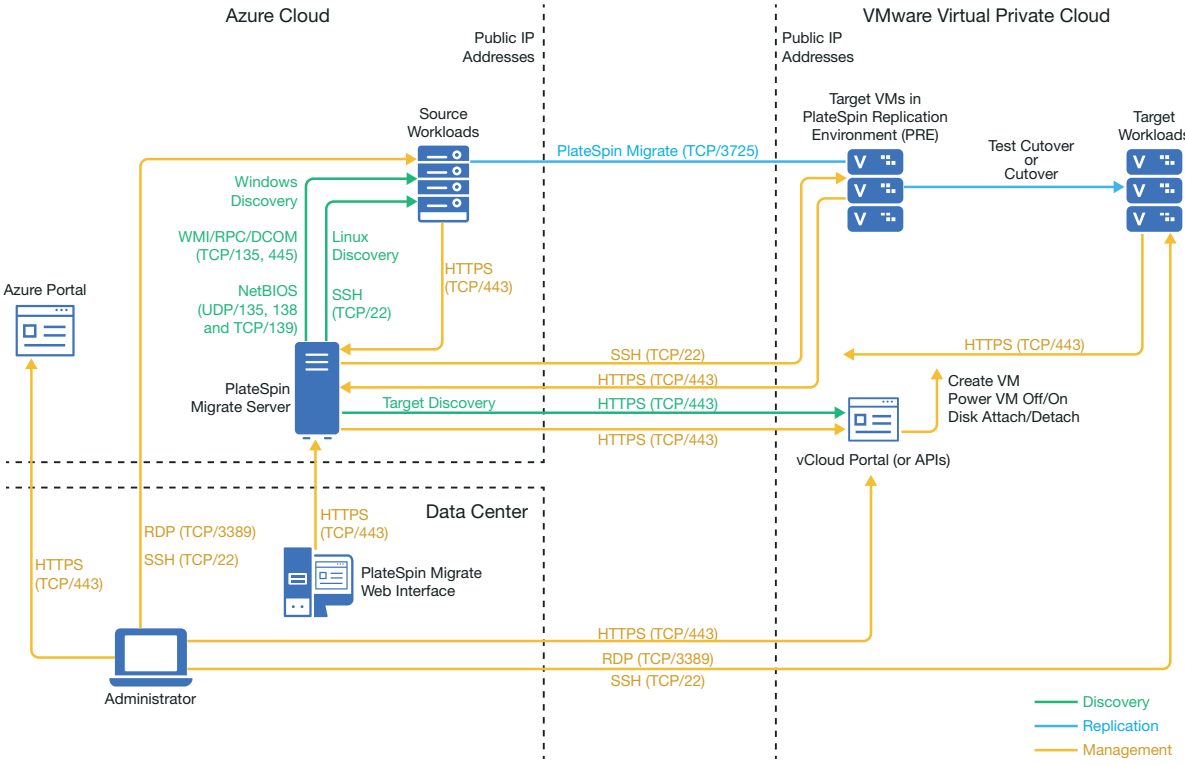

Requirements for Migration to vCloud

To prepare your target vCloud environment, review the information in "Planning For Migrating Workloads to VMware vCloud Director" on page 203.

Ensure that the source workload is supported by the target vCloud configuration.

Requirements for Migrating Workloads from Azure to vCloud

For source workloads in Azure:

- Azure automatically adds the Remote Desktop Protocol (RDP) port (TCP/3389) and Secure Shell (SSH) port (TCP/22) in the Azure Security Group for the source workload VMs. You must manually add other ports to the source workload's Security Group that are required by PlateSpin Migrate to provide migration services, such as Port 3725 for replication traffic and Port 443 for HTTPS traffic.
- For Windows workloads, use a user name and password.
- For Linux workloads, use the root user or root equivalent user.

To use an on-premise Migrate server for migration of workloads from Azure to vCloud:

- Deploy a site-to-site VPN between your data center and your Azure environment.
- Deploy a site-to-site VPN between your data center and your VMware vCloud Virtual Private Cloud.
- Because you are using VPNs with an on-premise Migrate server, you can use a private IP address for the Migrate server.
- Ensure that your source and target network meet the following requirements.
  - "Requirements for Discovery" on page 53.
  - "Requirements for Migration" on page 57.
- Migrate Agent is not required because a VPN is available, but it would also work. For network ports and firewall requirements for registration, see "Requirements for Workload Registration" on page 56.

To use a cloud-based Migrate server for migration of workloads from Azure to vCloud without a VPN:

- Deploy a PlateSpin Migrate server in the source Azure network environment. Ensure that your non-VPN migration environment meets the "Requirements for C2C Non-VPN Migrations" on page 217.
- In the PlateSpin Configuration settings on the Migrate server:
  - **(Migrate Server in Azure) ServerIsHostedInCloud:** Remove the value of `azure` from the **ServerIsHostedInCloud** parameter to enable the Add Target dialog to provide all target types for selection. When you set up the vCloud target, select the **VMware vCloud Organization** option.

Checklist for Automated Migration from Azure to vCloud

Task	Description
1. Prepare your network resources.	Figure 12-4, “Migrate Server on Premise for Migration from Azure to vCloud,” on page 226 Figure 12-5, “Migrate Server in Azure for Migration from Azure to vCloud with No VPNs,” on page 227 “Deployment for C2C Migration from Azure to vCloud” on page 226
2. Prepare your vCloud migration environment.	“Requirements for Migration to vCloud” on page 227
3. Prepare your Azure source workloads for PlateSpin Migrate.	“Requirements for Migrating Workloads from Azure to vCloud” on page 228
4. Discover target cloud platform.	“Target Discovery in the Web Interface” on page 284
5. Discover source workloads in Azure.	“Workload Discovery in the Migrate Web Interface” on page 302
6. Configure target workload migration.	“Configuring Migration of a Workload to VMware vCloud Director” on page 470
7. Execute migration.	Chapter 33, “Executing Workload Migrations,” on page 537

Prerequisites for C2C Migration from vCloud to Azure

PlateSpin Migrate supports migration of workloads from VMware vCloud Director to Microsoft Azure.

- “Deployment for C2C Migration from vCloud to Azure” on page 229
- “Requirements for Migrating Workloads to Azure” on page 231
- “Requirements for Migrating Workloads from vCloud to Azure” on page 231
- “Checklist for Automated Migration from vCloud to Azure” on page 232

Deployment for C2C Migration from vCloud to Azure

For migration of workloads from VMware vCloud Director to Microsoft Azure, deploy a PlateSpin Migrate server on premise in your source network. With an on-premise Migrate server, site-to-site VPN gateways are required between the data center and Azure and between the data center and vCloud. Figure 12-6 shows the location of various components in your Azure, vCloud, and data center migration environments and the communications between them.

Figure 12-6 *Migrate Server on Premise for Migration from vCloud to Azure*

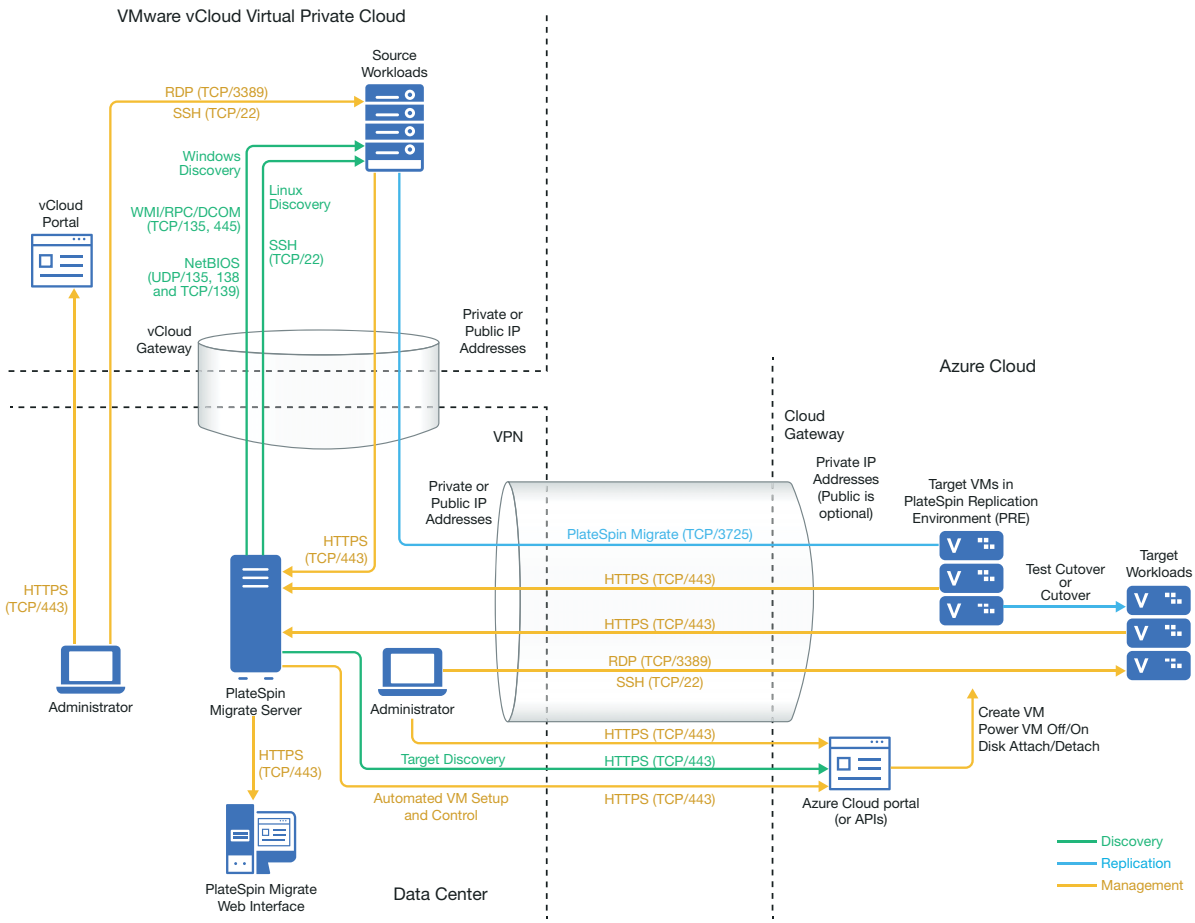

You can alternatively deploy the PlateSpin Migrate server from the Azure Marketplace in the target Azure environment. No VPN is required. You use Migrate Agent to register workloads with the cloud-based Migrate server using secure communications over the public Internet. Use data encryption to secure data for replications over the public Internet. Internet access and public IP addresses are required. Figure 12-7 shows the location of various components in your Azure, vCloud, and data center migration environments and the communications between them.

NOTE: A reboot of the source Windows workload is required when you install, uninstall, or upgrade block-based transfer drivers. A reboot is not required for source Linux workloads.

Figure 12-7 *Migrate Server in Azure for Migration from vCloud to Azure with No VPNs*

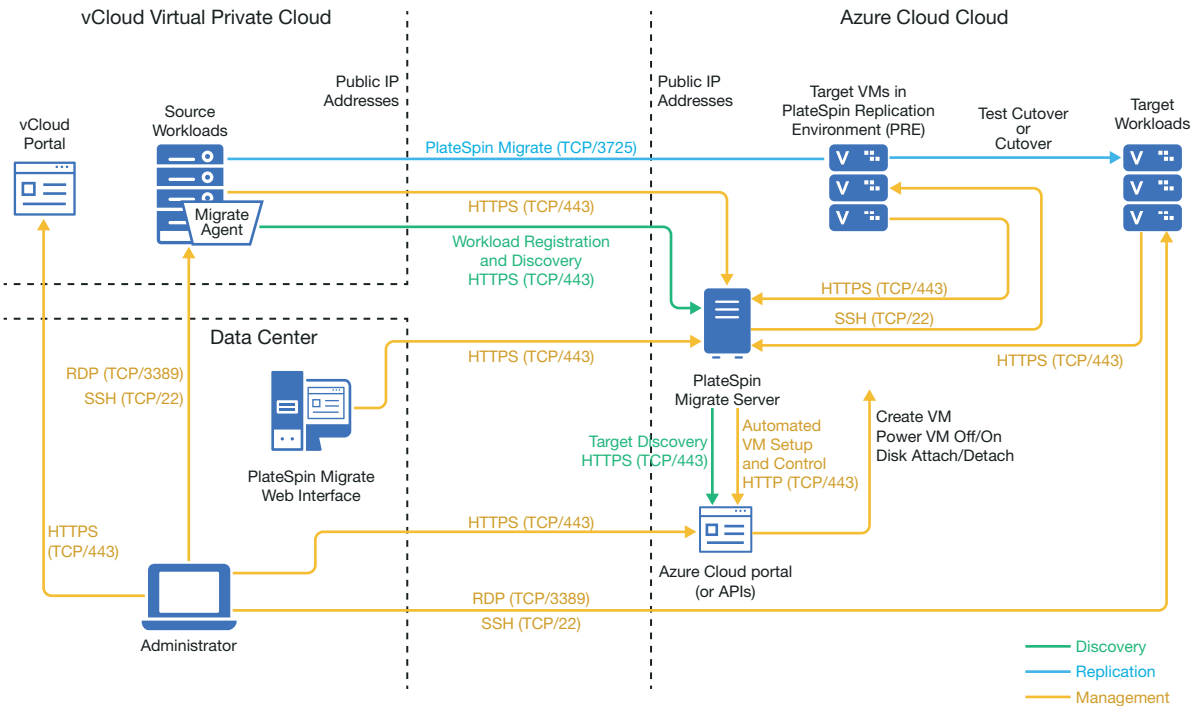

Requirements for Migrating Workloads to Azure

To prepare your target Azure environment, review the following information in "Requirements for Migrating Workloads to Azure" on page 171:

- "Minimum Azure Prerequisites" on page 171
- "Azure Prerequisites for Using an On-Premise Migrate Server" on page 175
- "Azure Prerequisites for Using an Azure-Based Migrate Server" on page 177

Ensure that the source workload is supported by the target Azure configuration.

Requirements for Migrating Workloads from vCloud to Azure

To use an on-premise Migrate server for migration of workloads from vCloud to Azure:

- Deploy a site-to-site VPN between your data center and your Azure environment.
- Deploy a site-to-site VPN between your data center and your VMware vCloud Virtual Private Cloud.
- Because you are using a VPN Gateway between the data center and Azure, you can use a private IP address for the Migrate server.
- Migrate Agent is not required because a VPN is available, but it would also work. For network ports and firewall requirements for registration, see "Requirements for Workload Registration" on page 56.

- (Migrate Discovery) Ensure that your source and target network meet the following requirements. See also Figure 12-6, “Migrate Server on Premise for Migration from vCloud to Azure,” on page 230.
  - “Requirements for Discovery” on page 53.
  - “Requirements for Migration” on page 57.

To use a cloud-based Migrate server for migration of workloads from vCloud to Azure without a VPN:

- Deploy a PlateSpin Migrate server in the target Azure network environment. Ensure that your non-VPN migration environment meets the “Requirements for C2C Non-VPN Migrations” on page 217.
- Azure automatically adds the Remote Desktop Protocol (RDP) port (TCP/3389) and Secure Shell (SSH) port (TCP/22) in the Azure Security Group for the source workload VMs. You must manually add other ports to the source workload’s Azure Security Group that are required by PlateSpin Migrate to provide migration services, such as Port 3725 for replication traffic and Port 443 for HTTPS traffic.
- For information about workload login requirements for migration, see the Windows and Linux source workload login requirements in Table 18-2, “Guidelines for Machine Type and Credentials for Source Workloads,” on page 299.

Checklist for Automated Migration from vCloud to Azure

Task	Description
1. Prepare your network resources.	Figure 12-6, “Migrate Server on Premise for Migration from vCloud to Azure,” on page 230 Figure 12-7, “Migrate Server in Azure for Migration from vCloud to Azure with No VPNs,” on page 231 “Deployment for C2C Migration from vCloud to Azure” on page 229
2. Prepare your vCloud migration environment.	“Requirements for Migrating Workloads to Azure” on page 231
3. Prepare your Azure source workloads for PlateSpin Migrate.	“Requirements for Migrating Workloads from vCloud to Azure” on page 231
4. Discover target cloud platform.	“Target Discovery in the Web Interface” on page 284
5. Discover source workloads in vCloud. You can optionally register source workloads with the cloud-based Migrate server in Azure using Migrate Agent.	“Workload Discovery in the Migrate Web Interface” on page 302 -OR- “Registering Workloads and Discovering Details with Migrate Agent” on page 303
6. Configure target workload migration.	“Configuring Migration of a Workload to VMware vCloud Director” on page 470

Task	Description
7. Execute migration.	Chapter 33, “Executing Workload Migrations,” on page 537

Prerequisites for C2C Migration from AWS to vCloud

PlateSpin Migrate supports migration of workloads from Amazon Web Services EC2 Cloud to VMware vCloud Director.

Deployment for C2C Migration from AWS to vCloud

For migration of workloads from Amazon Web Services EC2 Cloud to VMware vCloud Director, deploy a PlateSpin Migrate server on premise in your source network. VPN gateways are required between the data center and AWS and between the data center and vCloud. Figure 12-8 shows the location of various components in your AWS, vCloud, and data center migration environments and the communications between them.

Figure 12-8 *Migrate Server on Premise for Migration from AWS to vCloud*

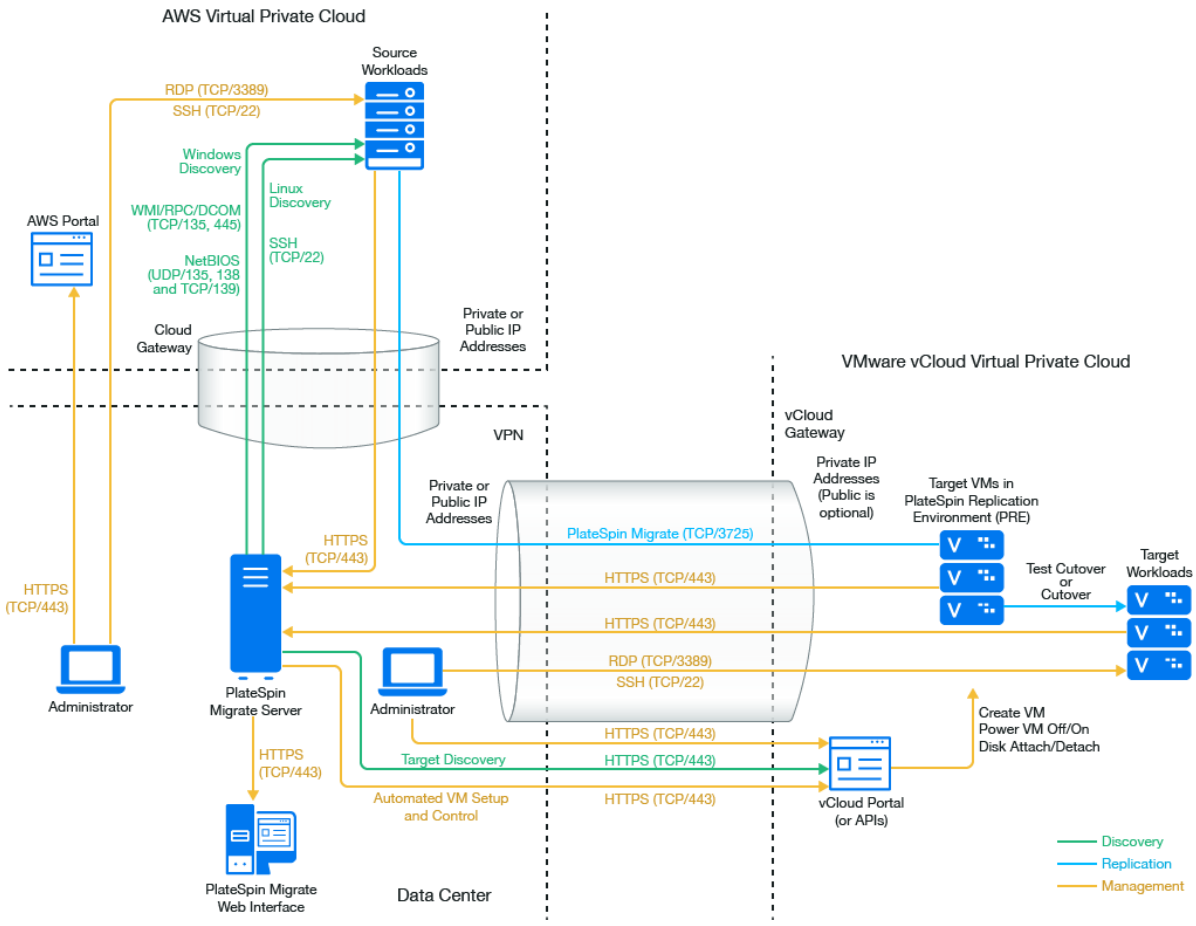

You can alternatively deploy the PlateSpin Migrate server using the AWS Quick Start in the source AWS environment. No VPN is required. With the AWS server in the same network security group as the source workloads, you can use discovery to add workloads to AWS. Use data encryption to secure data for replications over the public Internet. Figure 12-9 shows the location of various components in your AWS, vCloud, and data center migration environments and the communications between them.

NOTE: Figure 12-9 depicts source workloads and the Migrate server in the same network security group. If they are in different security groups, use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443). See "Requirements for Workload Registration" on page 56 and "Registering Workloads and Discovering Details with Migrate Agent" on page 303.

Figure 12-9 *Migrate Server in AWS for Migration from AWS to vCloud with No VPNs*

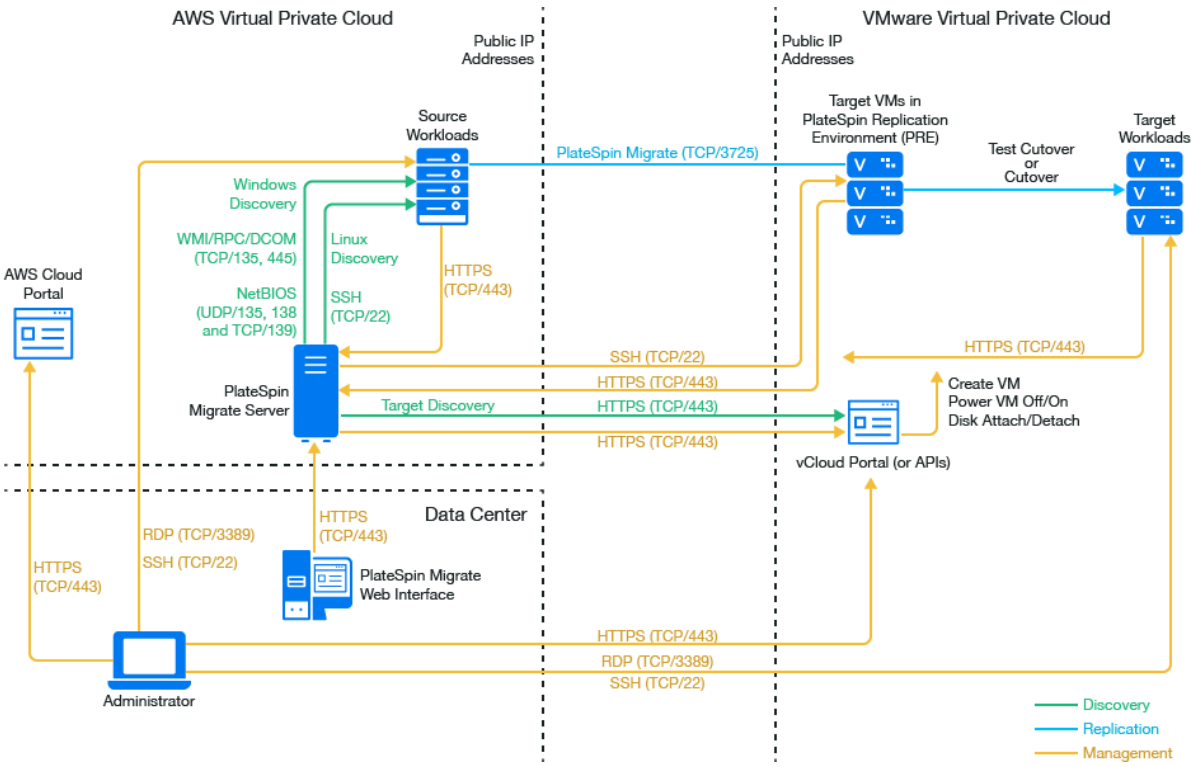

Requirements for Migration to vCloud

To prepare your target vCloud environment, review the information in "Planning For Migrating Workloads to VMware vCloud Director" on page 203.

Ensure that the source workload is supported by the target vCloud configuration.

Requirements for Migrating Workloads from AWS to vCloud

For source workloads in AWS:

- AWS automatically adds the Remote Desktop Protocol (RDP) port (TCP/3389) and Secure Shell (SSH) port (TCP/22) in the AWS Security Group for the source workload VMs. You must manually add other ports to the source workload's AWS Security Group that are required by PlateSpin Migrate to provide migration services, such as Port 3725 for replication traffic and Port 443 for HTTPS traffic.
- For Windows workloads, use a user name and password.
- For Linux workloads, use the root user or root equivalent user.

  In AWS, Amazon Linux AMIs by default enable the `ec2user` user name and PEM key credentials, and disable the `root` user name and password credentials. To use Migrate discovery to inventory workloads, you must enable `root` user access for the AWS source Linux workload. See "Enabling Root User Credentials for Source Linux Workloads in AWS" on page 242.

To use an on-premise Migrate server for migration of workloads from AWS to vCloud:

- Deploy a site-to-site VPN between your data center and your AWS environment.
- Deploy a site-to-site VPN between your data center and your VMware vCloud Virtual Private Cloud.
- Because you are using a VPN Gateway between the data center and AWS, you can use a private IP address for the Migrate server.
- Migrate Agent is not required because a VPN is available, but it would also work. For network ports and firewall requirements for registration, see "Requirements for Workload Registration" on page 56.

To use a cloud-based Migrate server for migration of workloads from AWS to vCloud without a VPN:

- Deploy a PlateSpin Migrate server in the source AWS network environment. Ensure that your non-VPN migration environment meets the "Requirements for C2C Non-VPN Migrations" on page 217.

Checklist for Automated Migration from AWS to vCloud

Task	Description
1. Prepare your network resources.	Figure 12-8, "Migrate Server on Premise for Migration from AWS to vCloud," on page 234 Figure 12-9, "Migrate Server in AWS for Migration from AWS to vCloud with No VPNs," on page 235 "Deployment for C2C Migration from AWS to vCloud" on page 233
2. Prepare your vCloud migration environment.	"Requirements for Migration to vCloud" on page 235
3. Prepare your AWS source workloads for PlateSpin Migrate.	"Requirements for Migrating Workloads from AWS to vCloud" on page 235
4. Discover target cloud platform.	"Target Discovery in the Web Interface" on page 284
5. Discover source workloads in AWS.	"Workload Discovery in the Migrate Web Interface" on page 302
6. Configure target workload migration.	"Configuring Migration of a Workload to VMware vCloud Director" on page 470
7. Execute migration.	Chapter 33, "Executing Workload Migrations," on page 537

Prerequisites for C2C Migration from vCloud to AWS

PlateSpin Migrate supports migration of workloads from VMware vCloud Director to Amazon Web Services EC2 Cloud.

Deployment for C2C Migration from vCloud to AWS

For migration of workloads from VMware vCloud Director to Amazon Web Services EC2 Cloud, deploy a PlateSpin Migrate server on premise in your source network. VPN gateways are required between the data center and AWS and between the data center and vCloud. Figure 12-10 shows the location of various components in your AWS, vCloud, and data center migration environments and the communications between them.

Figure 12-10 *Migrate Server on Premise for Migration from vCloud to AWS*

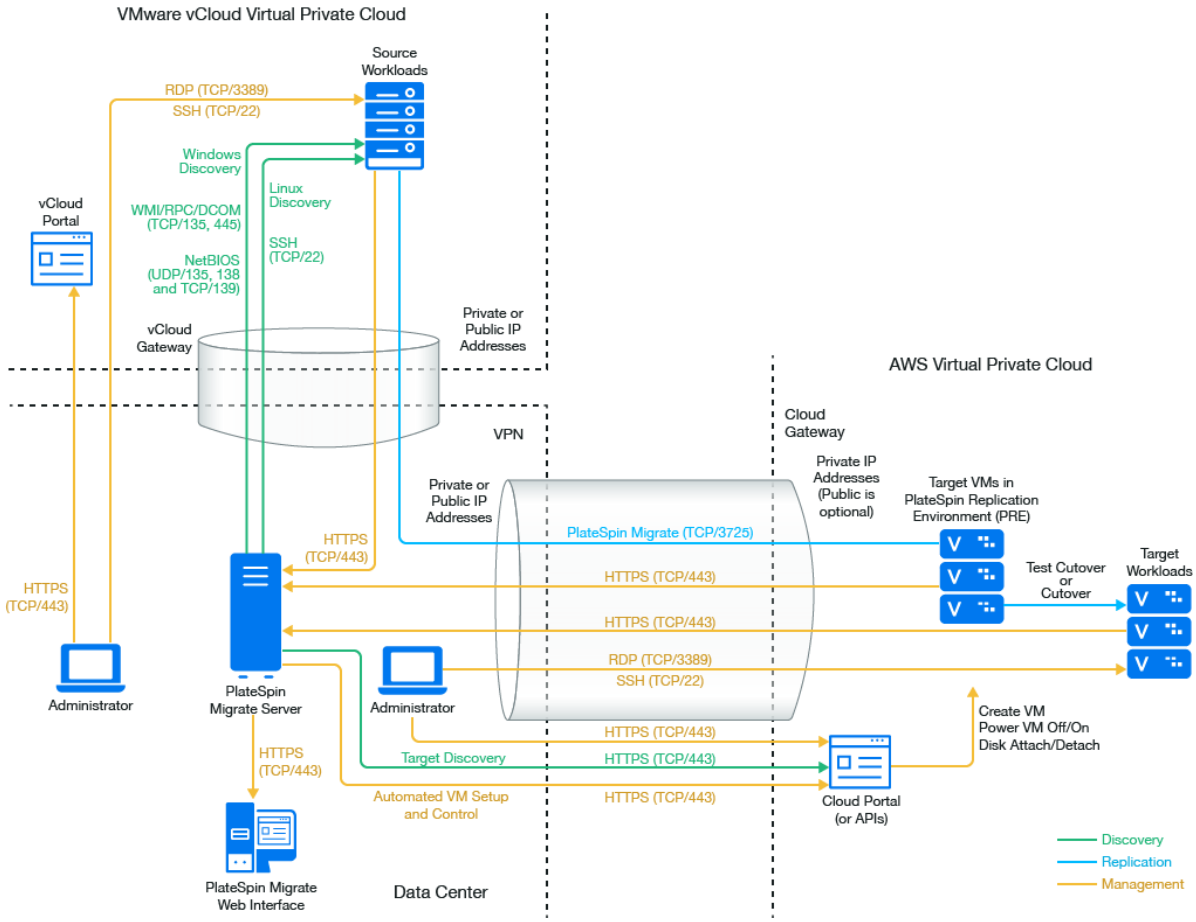

You can alternatively deploy the PlateSpin Migrate server using the AWS Quick Start in the target AWS environment. No VPN is required. You use Migrate Agent to register workloads with the cloud-based Migrate server using secure communications over the public Internet. Use data encryption to secure data for replications over the public Internet. Internet access and public IP addresses are required. Figure 12-11 shows the location of various components in your AWS, vCloud, and data center migration environments and the communications between them.

NOTE: A reboot of the source Windows workload is required when you install, uninstall, or upgrade block-based transfer drivers. A reboot is not required for source Linux workloads.

Figure 12-11 *Migrate Server in AWS for Migration from vCloud to AWS with No VPNs*

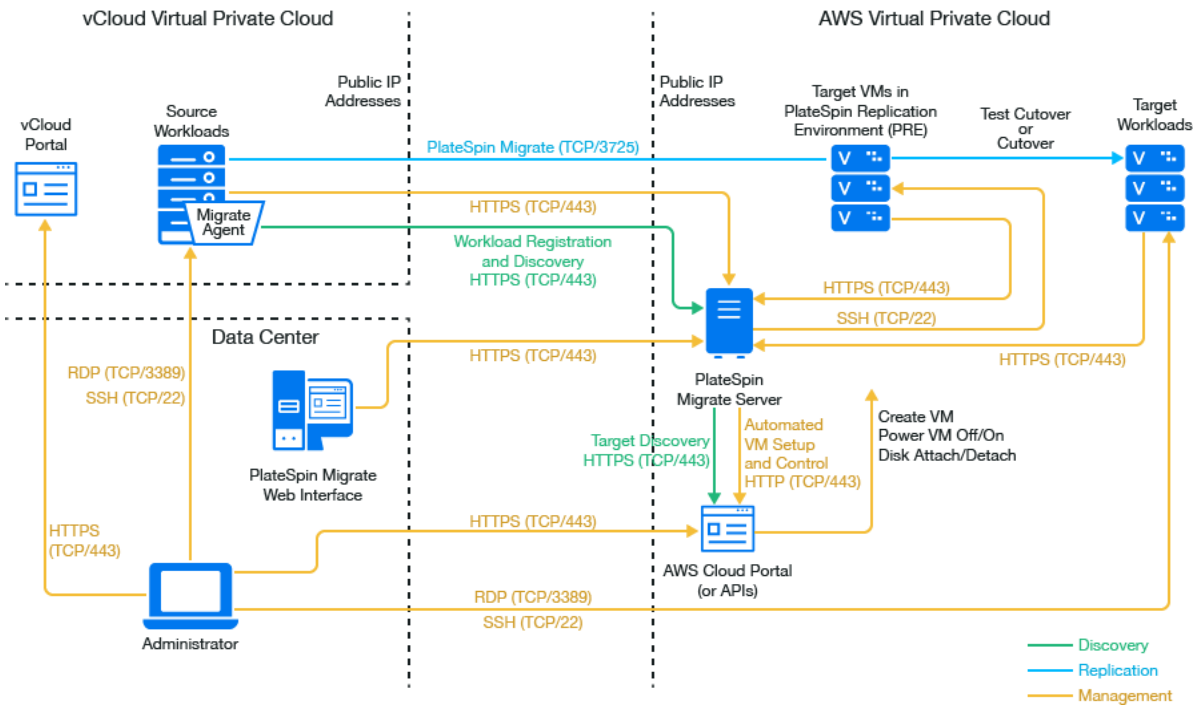

Requirements for Migrating Workloads to AWS

To prepare your target AWS environment, review the following information in "Requirements for Migrating Workloads to Amazon Web Services" on page 147:

- "Minimum AWS Prerequisites" on page 147
- "AWS Prerequisites for Using an AWS-Based Migrate Server" on page 150

Ensure that the source workload is supported by the target AWS configuration.

Requirements for Migrating Workloads from vCloud to AWS

To use an on-premise Migrate server for migration of workloads from vCloud to AWS:

- Deploy a site-to-site VPN between your data center and your AWS environment.
- Deploy a site-to-site VPN between your data center and your VMware vCloud Virtual Private Cloud.

- Because you are using a VPN Gateway between the data center and AWS, you can use a private IP address for the Migrate server.
- Migrate Agent is not required because a VPN is available, but it would also work. For network ports and firewall requirements for registration, see "Requirements for Workload Registration" on page 56.
- (Migrate Discovery) Ensure that your source and target network meet the following requirements. See also Figure 12-10, "Migrate Server on Premise for Migration from vCloud to AWS," on page 237.
  - "Requirements for Discovery" on page 53.
  - "Requirements for Migration" on page 57.

To use a cloud-based Migrate server for migration of workloads from vCloud to AWS without a VPN:

- Deploy a PlateSpin Migrate server in the target AWS network environment. Ensure that your non-VPN migration environment meets the "Requirements for C2C Non-VPN Migrations" on page 217.

Checklist for Automated Migration from vCloud to AWS

Task	Description
1. Prepare your network resources.	Figure 12-10, "Migrate Server on Premise for Migration from vCloud to AWS," on page 237 Figure 12-11, "Migrate Server in AWS for Migration from vCloud to AWS with No VPNs," on page 238 "Deployment for C2C Migration from vCloud to AWS" on page 237
2. Prepare your vCloud migration environment.	"Requirements for Migrating Workloads to AWS" on page 238
3. Prepare your AWS source workloads for PlateSpin Migrate.	"Requirements for Migrating Workloads from vCloud to AWS" on page 238
4. Discover target cloud platform.	"Target Discovery in the Web Interface" on page 284
5. Discover source workloads in vCloud. You can optionally register source workloads with the cloud-based Migrate server in AWS using Migrate Agent.	"Workload Discovery in the Migrate Web Interface" on page 302 -OR- "Registering Workloads and Discovering Details with Migrate Agent" on page 303
6. Configure target workload migration.	"Configuring Migration of a Workload to VMware vCloud Director" on page 470
7. Execute migration.	Chapter 33, "Executing Workload Migrations," on page 537

Prerequisites for C2C Migration from AWS to Oracle Cloud infrastructure

PlateSpin Migrate supports migration of workloads from Amazon Web Services EC2 Cloud to Oracle Cloud Infrastructure.

Deployment for C2C Migration from AWS to Oracle Cloud Infrastructure

For migration of workloads from Amazon Web Services EC2 Cloud to Oracle Cloud Infrastructure, deploy the PlateSpin Migrate server in the target Oracle Cloud environment. No VPN is required between the participating sites. Internet access and public IP addresses are required. Figure 12-1 shows the location of various components in your AWS, Oracle Cloud, and data center migration environments and the communications between them.

A public IP address is required for the source workload to register the workload with the cloud-based Migrate server using secure communications over the public Internet. For deployment information, see Figure 12-1, "Cloud-Based Migrate Server for Automated Migration from AWS to Azure with No VPNs," on page 221

Figure 12-12 *Cloud-Based Migrate Server for Automated Migration from AWS to Oracle Cloud with No VPNs*

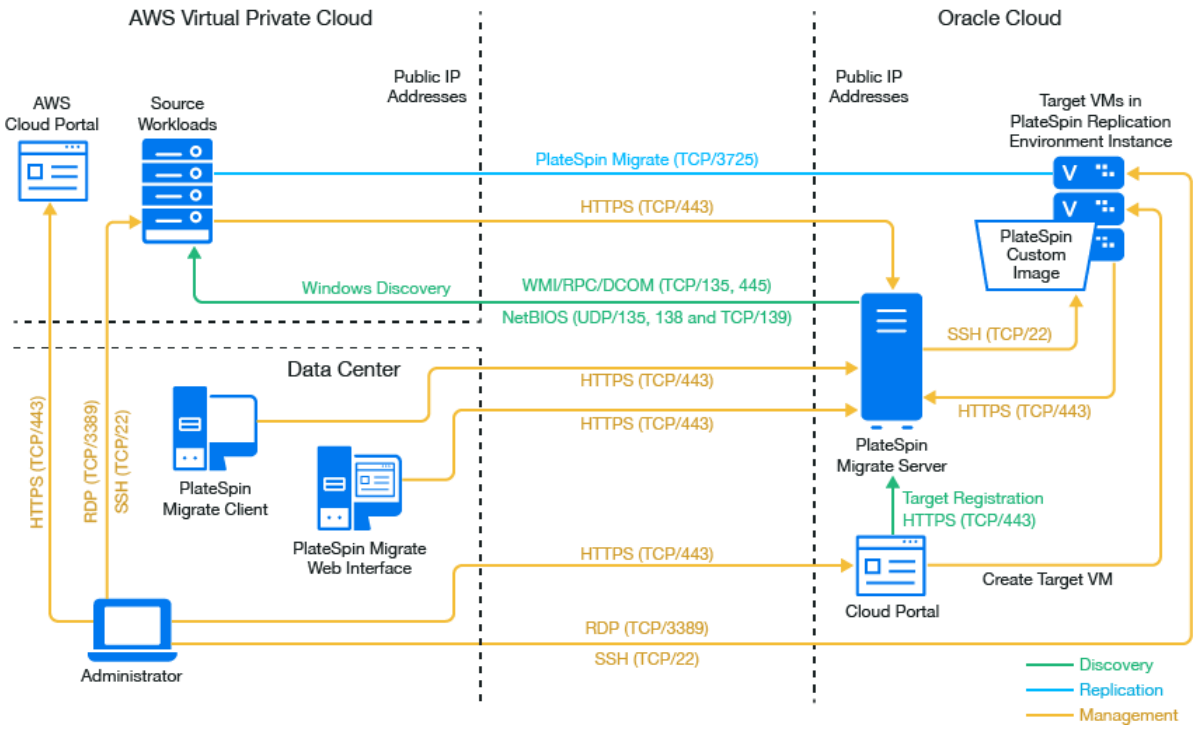

Requirements for Migrating Workloads to Oracle Cloud

To prepare your target Oracle Cloud Infrastructure environment, review the following information in "Requirements for Migrating Workloads to Oracle Cloud Infrastructure" on page 210:

- "Minimum Oracle Cloud Prerequisites"
- "Prerequisites for Using a Migrate Server on Oracle Cloud"

Ensure that the source workload is supported by the target Oracle Cloud configuration.

Requirements for Migrating Workloads from AWS to Oracle Cloud

Deploy a PlateSpin Migrate server in the target Oracle Cloud infrastructure environment. Ensure that your non-VPN migration environment meets the "Requirements for C2C Non-VPN Migrations" on page 217.

For source workloads in AWS:

- AWS automatically adds the Remote Desktop Protocol (RDP) port (TCP/3389) and Secure Shell (SSH) port (TCP/22) in the AWS Security Group for the source workload VMs. You must manually add other ports to the source workload's AWS Security Group that are required by PlateSpin Migrate to provide migration services, such as Port 3725 for replication traffic and Port 443 for HTTPS traffic.
- For Windows workloads, use a user name and password.
- For Linux workloads, use the root user or root equivalent user.

  In AWS, Amazon Linux AMIs by default enable the `ec2user` user name and PEM key credentials, and disable the `root` user name and password credentials. To use Migrate discovery to inventory workloads, you must enable `root` user access for the AWS source Linux workload. See "Enabling Root User Credentials for Source Linux Workloads in AWS" on page 242.

Checklist for Automated Migration from AWS to Oracle Cloud

Task	Description
1. Prepare your network resources.	Figure 12-12, "Cloud-Based Migrate Server for Automated Migration from AWS to Oracle Cloud with No VPNs," on page 240 "Deployment for C2C Migration from AWS to Azure" on page 220
2. Prepare your Oracle Cloud migration environment.	"Requirements for Migrating Workloads to Azure" on page 221

Task	Description
3. Discover source workloads.	"Workload Discovery in the Migrate Client" on page 301 -OR- "Workload Discovery in the Migrate Web Interface" on page 302
4. Prepare your AWS source workloads for PlateSpin Migrate.	"Requirements for Migrating Workloads from AWS to Azure" on page 221
5. Discover target cloud platform.	"Creating and Connecting to the Target Virtual Machine Instance in Oracle Cloud Infrastructure Environment" on page 482
6. Configure target workload migration.	"Configuring Full Migration of a Workload to Oracle Cloud VM Instance Using the X2P Workflow" on page 482
7. Execute migration.	Chapter 33, "Executing Workload Migrations," on page 537

Enabling Root User Credentials for Source Linux Workloads in AWS

PlateSpin Migrate requires `root` user credentials for discovery of Linux workloads. To use Migrate discovery instead of the Migrate Agent to inventory source workloads in AWS, you must enable `root` user access for the workload. In AWS, Amazon Linux AMIs by default enable the `ec2user` user name and PEM key credentials, and disable the `root` user name and password credentials.

NOTE: If the Migrate Server resides on premise in the data center, you must have a site-to-site VPN between the AWS account and the data center in order to use Migrate discovery for the inventory.

To enable root user credentials on an AWS source Linux workload:

1. Use SSH tool (such as Putty) to connect to the source Linux workload in AWS, and log in with the `ec2user` user name and PEM key credentials.
2. Run `sudo su`.
3. Create a password for the `root` user by running the `passwd` command.
4. In a text editor, edit the `/etc/ssh/sshd_config` file. Ensure that the directive "`PasswordAuthentication no`" is uncommented and set to `yes`.

   ```
   PasswordAuthentication yes
   ```

5. Run the `/etc/init.d/sshd reload` command, or reboot the workload to apply the changes.

   On Red Hat Enterprise Linux 7.*x*, use the following command:

```
/bin/systemctl restart sshd.service
```

Reloading or restarting the SSH daemon might not work on some Linux distributions, In this case, a reboot is required to apply the settings.

Configuring Advanced Settings for a Cloud-Based Migrate Server

Cloud-based PlateSpin Migrate servers, such as Migrate images deployed in Azure Marketplace or using AWS Quick Start, have pre-configured PlateSpin advanced settings for workload migrations to the parent cloud, as described in Table 12-1. If you intend to use the cloud-based Migrate server to migrate workloads from the parent cloud environment, you must modify the settings.

IMPORTANT: PlateSpin Migrate supports semi-automated migration of workloads to Virtual Machine Instances on your Oracle Cloud Infrastructure environment. You must manually install the PlateSpin Migrate server in the Oracle Cloud Infrastructure environment. Migrate Agent is not supported for registering source workloads that you want to migrate to Oracle Cloud Infrastructure environment.

Table 12-1 *PlateSpin Configuration Settings for PlateSpin Migrate Server in the Cloud*

Parameter	Migrations to Cloud	Migrations from Cloud	Remarks
SourceListensForConnection	`False` Assumes that Migrate Agent is used to register workloads.	`False`	The source and target both must have public IP addresses accessible to each other. Registration of source workloads using Migrate Agent is not supported for semi-automated migrations to Oracle Cloud Infrastructure. Instead, the source workloads must be discovered by the Migrate server. For semi-automated migrations to Oracle Cloud, set SourceListensForConnection to `True`, then use either of the following methods to allow source discovery: ◆ Provide a VPN connection between the source network and the Migrate server deployed in Oracle Cloud. ◆ Ensure that source workloads have public IP addresses to allow discovery by the Migrate server deployed in Oracle Cloud using secure communications over the public Internet.

Parameter	Migrations to Cloud	Migrations from Cloud	Remarks
AlternateServerAddress	Migrate server's public IP address	Migrate server's public IP address	If you use Migrate Agent to register source workloads, the public IP address is set automatically for this parameter when you register the source. (For non-VPN setup) For semi-automated migrations to Oracle Cloud Infrastructure, you must manually set the value of AlternateServerAddress to the Migrate server's public IP address after you have registered the target VM instances. See "Configuring Alternate IP Addresses for PlateSpin Server" on page 111. **IMPORTANT:** Before you set the AlternateServerAddress parameter, ensure that you first prepare the target VM instances in Oracle Cloud and use PlateSpin Custom Image to register them with the PlateSpin Migrate server deployed in Oracle Cloud.
ServerIsHostedInCloud	Depending on the cloud location: ◆ `azure` -OR- ◆ `aws`	(no value, empty field)	This parameter limits the type of targets available in the Add Targets dialog. When it is empty, all target types are available. This parameter is not applicable for semi-automated migrations to Oracle Cloud Infrastructure. See "Enabling a Cloud-Based Migrate Server to Handle Migrations to Other Target Platforms" on page 245.

Enabling a Cloud-Based Migrate Server to Handle Migrations to Other Target Platforms

For Migrate servers deployed from Azure Marketplace or using the AWS Quick Start, the ServerIsHostedInCloud parameter is set to the parent cloud value, such as `azure` or `aws`. This setting determines what target types are available to you in the Add Target dialog in the Migrate Web Interface, as described in Table 12-2.

NOTE: The settings for the **ServerIsHostedInCloud** parameter are valid only for automated migrations using PlateSpin Migrate Web Interface. The settings do not apply for semi-automated migrations to Oracle Cloud Infrastructure.

Table 12-2 *Target Types Allowed for Cloud-Based Migrate Servers*

ServerIsHostedInCloud Value	**Target Type in Add Target**	**Description**
`azure`	**Microsoft Azure Location**	Default setting for Migrate servers deployed from Azure Marketplace.
`aws`	**Amazon Cloud Region**	Default setting for Migrate servers deployed using the AWS Quick Start.
No value	All target types	Remove the pre-assigned value if you are using the cloud-based Migrate server to migrate workloads from the parent cloud environment to a different target type.

If you are migrating workloads from the parent cloud of a cloud-based Migrate server to a different target type, you must remove the default value (leave the field blank) for the **ServerIsHostedInCloud** parameter. After you remove the value, all target types are available in the Add Target dialog in the Migrate Web Interface.

To enable migrations from the source cloud using a cloud-based Migrate server:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://`*`Your_PlateSpin_Server`*`/PlateSpinConfiguration/`

2. Search to locate the **ServerIsHostedInCloud** parameter and remove the pre-configured cloud setting. Leave the field blank.
3. Save your settings and exit the page.

   A reboot or restart of PlateSpin services is not required to apply the changes.

13 Prerequisites for Migration to VMware

PlateSpin Migrate supports automated or semi-automated migration to your VMware environment. This section describes the required VMware configuration that you must prepare before you can discover VMware target virtualization platforms (for automated migration) or target VMs (for semi-automated migrations) and configure migrations to them.

Deployment for Migration to VMware

Figure 13-1 shows the location of various components in your automated VMware migration environment and the communications between them. Both PlateSpin Migrate Client and PlateSpin Migrate Web Interface supports automated migrations to VMware target virtualization platforms.

NOTE: Figure 13-1 and Figure 13-2 depict automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443). See "Requirements for Workload Registration" on page 56 and "Registering Workloads and Discovering Details with Migrate Agent" on page 303.

Figure 13-1 *Automated Migration to VMware*

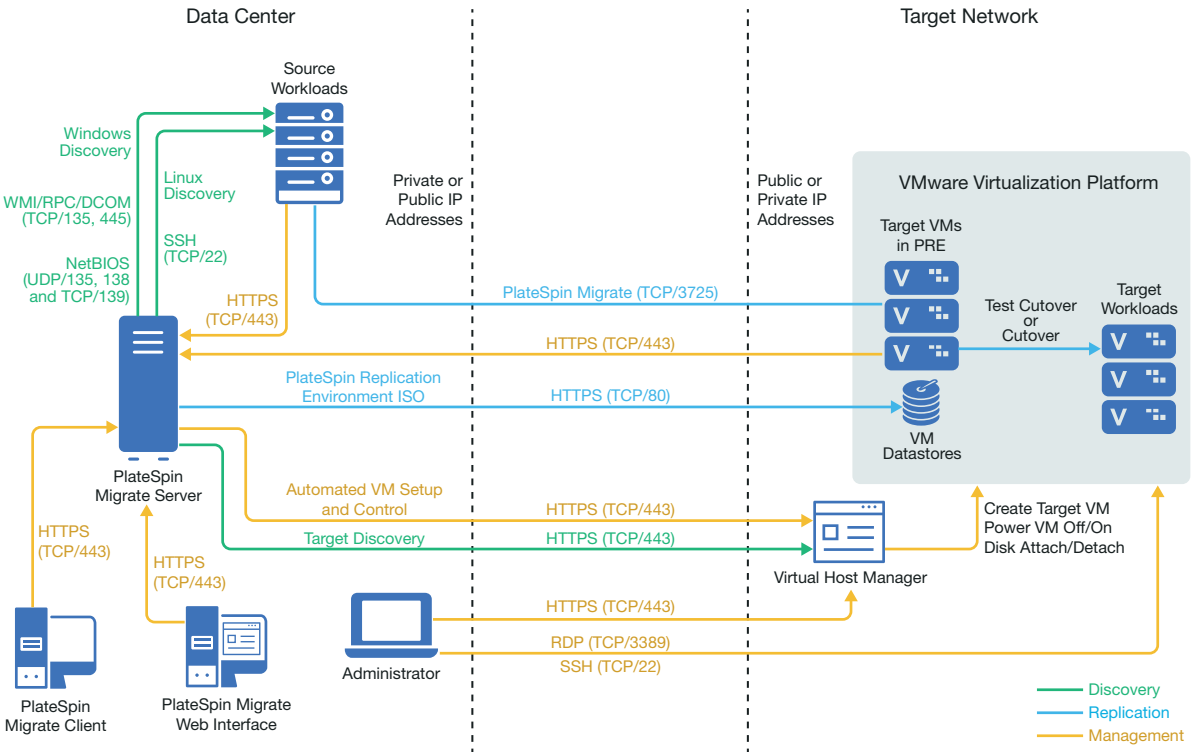

Figure 13-2 shows the location of various components in your semi-automated VMware migration environment and the communications between them. Both PlateSpin Migrate Client and PlateSpin Migrate Web Interface supports semi-automated migrations to VMware target virtualization platforms.

Figure 13-2 *Semi-Automated Migration to VMs on VMware*

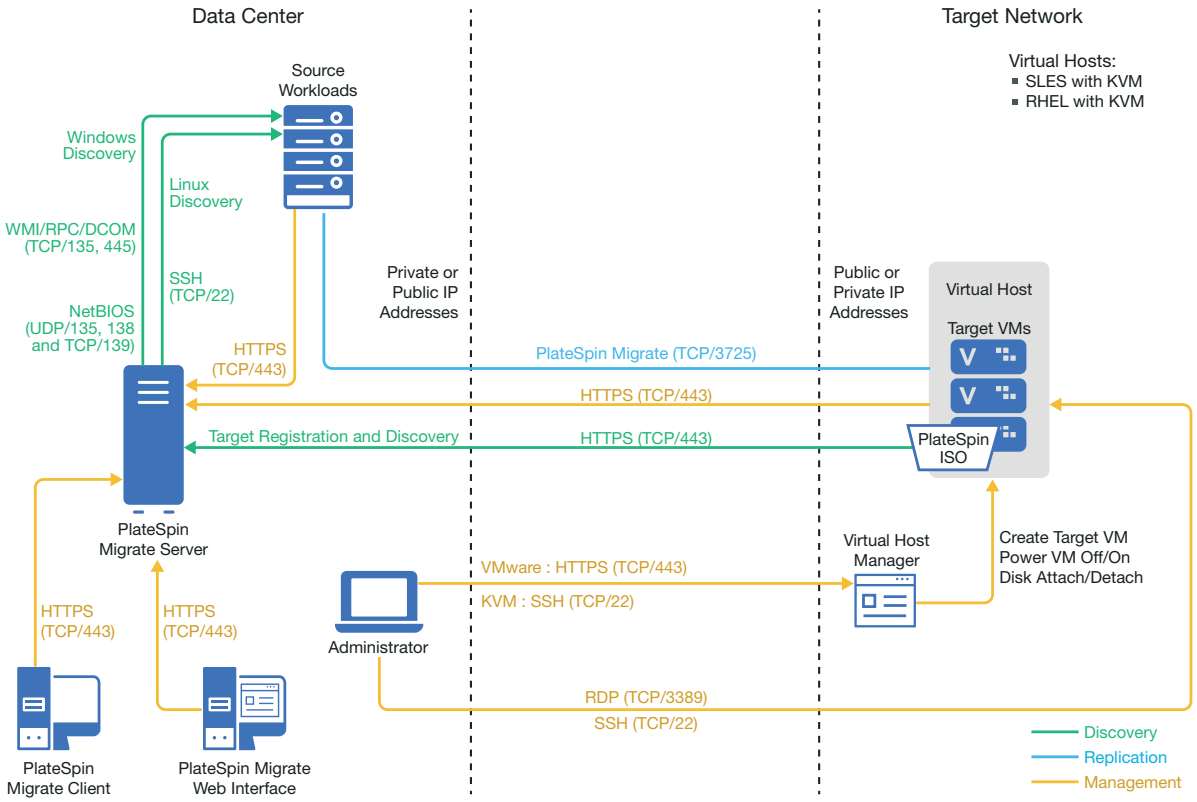

Planning for Migration to VMware

Ensure that your VMware environment meets the following prerequisites for migration to VMware:

- Use PlateSpin Migrate Client or PlateSpin Migrate Web Interface to migrate workloads to VMware.

  See Table 2-12, “Supported Target VMware Platforms for the Migrate Web Interface and Migrate Client,” on page 41.

- Your source workload must be supported by PlateSpin Migrate and VMware.

  See “Supported Source Workloads For Migration to Non-Cloud Platforms” on page 23.

- Your network environment must meet the requirements for access, discovery, and migration described in “Access and Communication Requirements across Your Migration Network” on page 53.

- You can optionally set up a PlateSpin Virtual Machine Manager role on your VMware vCenter server that Migrate will use for migrations instead of the vCenter administrator user. See “Configuring a PlateSpin User with Minimal Permissions on VMware” on page 250.

- For semi-automated migrations, ensure that you configure volumes on the target disks with about 50 MB of additional storage space than the source disks.

- Raw Device Mapping (RDM) for target VMs on VMware is supported only by using the X2P workflow.

- Installing VMware Tools on the target workload:
  - When you use the X2P workflow for migrating a workload to VMware, you must manually set up the VMware Tools for the target workload before you perform the conversion. See "Setting Up VMware Tools for the Target Workload" on page 507.
  - Before you migrate a Linux workload, ensure that Perl module is installed on the source Linux workload to enable PlateSpin Migrate to install the VMware tools on the target workload during conversion. Alternatively, you can manually install the VMware tools after the migration is completed on the cutover workload.
- Migration of workloads with multiple NICs is supported for Windows workloads up to the number of NICs supported by VMware.

NOTE: When you configure migration for a workload, PlateSpin Migrate lets you specify four IP addresses per NIC on the Windows workload by default. However, you can configure the number of IP addresses that you want to specify per NIC on the workload. See "Configuring the Maximum Number of IP Addresses Allowed for Each Network Interface on Target Workloads" on page 111.

For information about configuring the migration, see "Migration to VMware" on page 487.

Configuring a PlateSpin User with Minimal Permissions on VMware

PlateSpin Migrate does not require administrative access to a VMware environment to perform workload migrations. Migrate requires permissions only for the VMware resources it must access to execute tasks in the migration workflow and functions.

To establish minimal permissions for a Migrate Server in a target VMware environment:

1. Create a special-purpose PlateSpin user to represent the Migrate server.
2. In vSphere, create predefined PlateSpin roles, and configure each role with the predefined minimum required VMware privileges.
3. Assign to the PlateSpin user an appropriate role for each type of VMware resource that you want Migrate to be able to access for migration to the VMware environment.

Review the information in this section to understand the PlateSpin roles for VMware and how to assign them in your VMware environment.

- "About PlateSpin VMware Roles and Permissions" on page 251
- "Assigning VMware Roles in a Single User Environment" on page 252
- "Assigning VMware Roles in a Multitenancy Environment" on page 255
- "Configuring VMware Roles for PlateSpin Migrate" on page 259

About PlateSpin VMware Roles and Permissions

PlateSpin Migrate provides custom VMware roles that you can use to control access for Migrate in your VMware environment. Each role is a set of predefined minimal permissions needed to perform actions in the migration workflow. For each Migrate server, you create a special-purpose PlateSpin user that represents the Migrate server in vCenter. To permit required actions, you associate the PlateSpin user with the appropriate role on various VMware resources that Migrate needs to access.

The `PlateSpinRole.xml` file on the Migrate server defines the minimum required privileges for each role. After you associate the roles for a special PlateSpin user, PlateSpin Migrate will have all the necessary VMware permissions to complete a migration.

- PlateSpin Virtual Infrastructure Manager
- PlateSpin Virtual Machine Manager
- PlateSpin User

For more granular control, you can establish the following four roles for a PlateSpin user to filter out resources for which the user does not have sufficient privileges to perform migrations. Use equivalent roles to more narrowly apply permissions to different VMware resource types. The PlateSpin VMware Role Tool cannot create these roles for you.

- PlateSpin Datastore Manager
- PlateSpin Network Manager
- PlateSpin Cluster Manager
- PlateSpin Virtual Machine User

For your convenience, PlateSpin Migrate provides the PlateSpin VMware Role Tool and role definitions that you can use to configure the required PlateSpin roles and permissions in your VMware environment. Table 13-1 describes these resources, which are available in the `<Migrate-install-folder>\PlateSpin Migrate Server\bin\VMwareRolesTool` directory.

Table 13-1 *Migrate Resources for Setting VMware Roles and Permissions*

File Name	Description
`PlateSpinRole.xml`	The PlateSpin Role XML file defines VMware custom roles and the minimum required VMware privileges for each role.
`PlateSpin.VMwareRoleTool.exe`	PlateSpin VMware Role Tool creates certain VMware roles in a VMware environment and automatically sets the minimum privileges for them.

Assigning VMware Roles in a Single User Environment

To grant permissions for the PlateSpin user in a single-user environment, you assign the appropriate role for the PlateSpin user for a specific vCenter inventory object. The user can have different roles for different objects.

About Role Assignments in a Single-User Environment

As you set up a single-user environment, you create a PlateSpin user to represent the Migrate server in the VMware environment. You assign the PlateSpin VMware roles to the user for the VMware resources, according to the resources that user must be able to access, and only those resources. In the VMware environment, it is this PlateSpin user that creates the Migrate target and performs actions on it during the migration.

Table 13-2 describes the role assignments across your VMware environment that are required to perform a migration. Assign the roles to the PlateSpin user for the VMware resources that you want the Migrate Server to access. Ensure that you do not extend permissions to the PlateSpin user for resources that you do not want the Migrate Server to access.

IMPORTANT: Propagating any permission has security implications. For security reasons, propagating privileges is not recommended for role assignments at the root of the Inventory tree and for DataCenter objects. Otherwise, propagation settings are at the discretion of the VMware administrator, except as noted.

Table 13-2 *PlateSpin User: Role Assignments for VMware Resource Types*

PlateSpin Role (or Equivalent Custom Role)	Propagation Instructions	More information
Root of the vCenter Inventory tree		
PlateSpin Virtual Infrastructure Manager	Non-propagation recommended	This setting enables Migrate to monitor tasks being performed by Migrate software and to end any stale VMware sessions.
Each DataCenter object that contains target clusters and hosts		
PlateSpin Virtual Infrastructure Manager	Non-propagation recommended	This setting enables Migrate to access the data center's datastores for file upload and download.

PlateSpin Role (or Equivalent Custom Role)	Propagation Instructions	More information
Target cluster and its member hosts		
PlateSpin Virtual Infrastructure Manager	You can propagate permission from the Cluster object to Host objects, or create an additional permission on each of its member Host objects. If you assign the role on the cluster object and propagate it, no further changes are necessary when you add a new host to the cluster. However, propagating this permission has security implications.	Configure the role on each target cluster (and its member hosts) that you will specify as a migration target.
Resource pool		
PlateSpin Virtual Machine Manager		Configure the PlateSpin Virtual Machine Manager role on at least one resource pool. You can grant permissions to one or multiple resource pools that are available to the enabled clusters.
PlateSpin Virtual Machine User		(Optional) Configure the PlateSpin Virtual Machine User role on resource pools you want to filter out.
VM folder		
PlateSpin Virtual Machine Manager		Configure the PlateSpin Virtual Machine Manager role on at least one VM folder. You can grant permissions to one or multiple VM folders.
PlateSpin Virtual Machine User		(Optional) Configure the PlateSpin Virtual Machine User role on VM folders you want to filter out.
Network (dvSwitch or vNet)		
PlateSpin Virtual Machine Manager	To assign the correct role to a dvSwitch, propagate the role on the data center (resulting in an additional object receiving the role), or place the dvSwitch in a folder and assign the role on that folder.	Configure the PlateSpin Virtual Machine Manager role on at least one Network. You can grant permissions to one or multiple networks that are available to the enabled clusters. For a standard portgroup to be listed as an available network in the Migrate UI, create a definition for it on every host in the cluster.

PlateSpin Role (or Equivalent Custom Role)	Propagation Instructions	More information
PlateSpin Network Manager		(Optional) Configure the Network Manager role on networks you want to filter out. For a standard portgroup to be listed as an available network in the Migrate UI, create a definition for it on every host in the cluster.
Datastore / Datastore cluster		
PlateSpin Virtual Machine Manager	For Datastore Clusters, the permission must be propagated to the contained datastores. Not providing access to an individual member of the cluster causes both Prepare and Full Replication to fail.	Configure the PlateSpin Virtual Machine Manager role on at least one datastore or datastore cluster. You can grant permissions to one or multiple datastores or datastore clusters that are available to the enabled clusters.
PlateSpin Datastore Manager		(Optional) Configure the PlateSpin Datastore Manager role on datastores or datastore clusters you want to filter out.

Security Implications of Assigning VMware Roles

When you assign VMware roles to the PlateSpin user, the key security implications include:

- With the *PlateSpin Virtual Infrastructure Manager* role assigned to the vCenter object, the PlateSpin user can see (but not affect) the tasks performed by every other user.
- Because there is no way to set permissions on datastore folders and subfolders, the PlateSpin user with permissions on a datastore has access to all disks stored on that datastore.
- With the *PlateSpin Virtual Infrastructure Manager* role assigned to a Cluster object, the PlateSpin user is able to turn off/on HA or DRS on the entire cluster.
- Setting the *PlateSpin Virtual Infrastructure Manager* role on the DRS Cluster object and propagating this role allows the PlateSpin user to see all VMs placed in the default resource pool and/or default VM folder. Also, propagation requires the administrator to explicitly set the PlateSpin user to have a “No Access” role on every resource pool and VM folder that the user should not be able to access.
- Setting the *PlateSpin Virtual Infrastructure Manager* role on the vCenter object allows the PlateSpin user to end sessions of any other user connected to the vCenter.

Assigning VMware Roles in a Multitenancy Environment

In a multitenancy environment, the custom VMware roles for PlateSpin makes it possible for you, as a service provider, to segment your VMware cluster to allow multitenancy: where multiple Migrate targets are instantiated in your data center to accommodate Migrate customers or "tenants" who want to keep their data and evidence of their existence separate from and inaccessible to other customers who also use your data center.

- "About Role Assignments in a Multitenancy Environment" on page 255
- "Security Implications of Assigning VMware Roles in a Multitenancy Environment" on page 258

About Role Assignments in a Multitenancy Environment

As you set up a multitenancy environment, you need to provision a single Migrate server per customer or "tenant." For each tenant, you create a PlateSpin user to represent that tenant's Migrate server in the VMware environment. You assign the PlateSpin VMware roles to the user for the VMware resources, according to the resources that user must be able to access, and only those resources. In the VMware environment, it is the tenant-based PlateSpin user that creates the Migrate target and performs actions on it during the migration. As service provider, you maintain this user's credentials and do not disclose them to your tenant customer.

Table 13-3 describes the role assignments across your VMware environment that are required to perform a migration. Assign the roles to each of the tenant-based PlateSpin users for the VMware resources that you want the tenant's Migrate Server to access. Ensure that you do not extend permissions to the tenant-based PlateSpin user for the tenant's resources that you do not want the Migrate Server to access.

IMPORTANT: Propagating any permission has security implications. For security reasons, propagating privileges is not recommended for role assignments at the root of the Inventory tree and for DataCenter objects. Otherwise, propagation settings are at the discretion of the VMware administrator, except as noted.

Table 13-3 *Tenant-Based PlateSpin User: Role Assignments for the Tenant's VMware Resource Types*

PlateSpin Role	Propagation Instructions	More information
Root of the vCenter Inventory tree		
PlateSpin Virtual Infrastructure Manager	Non-propagation recommended	This setting enables Migrate to monitor tasks being performed by Migrate software and to end any stale VMware sessions.
Each DataCenter object that contains target clusters and hosts		
PlateSpin Virtual Infrastructure Manager	Non-propagation recommended	This setting enables Migrate to access the data center's datastores for file upload and download.

PlateSpin Role	Propagation Instructions	More information
Target cluster and its member hosts		
PlateSpin Virtual Infrastructure Manager	You can propagate permission from the Cluster object to Host objects, or create an additional permission on each of its member Host objects. If you assign the role on the cluster object and propagate it, no further changes are necessary when you add a new host to the cluster. However, propagating this permission has security implications.	Configure the role on each target cluster (and its member hosts) that you will specify as a migration target.
Resource pool		
PlateSpin Virtual Machine Manager		Configure the PlateSpin Virtual Machine Manager role on at least one resource pool. You can grant permissions to one or multiple resource pools that are available to the enabled clusters.
PlateSpin Virtual Machine User		(Optional) Configure the PlateSpin Virtual Machine User role on resource pools you want to filter out.
VM folder		
PlateSpin Virtual Machine Manager		Configure the PlateSpin Virtual Machine Manager role on at least one VM folder. You can grant permissions to one or multiple VM folders.
PlateSpin Virtual Machine User		(Optional) Configure the PlateSpin Virtual Machine User role on VM folders you want to filter out.
Network (dvSwitch or vNet)		
PlateSpin Virtual Machine Manager	To assign the correct role to a dvSwitch, propagate the role on the data center (resulting in an additional object receiving the role), or place the dvSwitch in a folder and assign the role on that folder.	Configure the PlateSpin Virtual Machine Manager role on at least one Network. You can grant permissions to one or multiple networks that are available to the enabled clusters. For a standard portgroup to be listed as an available network in the Migrate UI, create a definition for it on every host in the cluster.

PlateSpin Role	Propagation Instructions	More information
PlateSpin Network Manager		(Optional) Configure the Network Manager role on networks you want to filter out. For a standard portgroup to be listed as an available network in the Migrate UI, create a definition for it on every host in the cluster.
Datastore / Datastore cluster		
PlateSpin Virtual Machine Manager	For Datastore Clusters, the permission must be propagated to the contained datastores. Not providing access to an individual member of the cluster causes both Prepare and Full Replication to fail.	Configure the PlateSpin Virtual Machine Manager role on at least one datastore or datastore cluster. You can grant permissions to one or multiple datastores or datastore clusters that are available to the enabled clusters.
PlateSpin Datastore Manager		(Optional) Configure the PlateSpin Datastore Manager role on datastores or datastore clusters you want to filter out.

Table 13-4 describes the role you can assign to the customer or tenant user.

Table 13-4 *Tenant-Based PlateSpin User: Role Assignments for the Tenant User*

PlateSpin Role	Propagation Instructions	More information
Each Resource Pool and Folder where the tenant's VMs will be created		
PlateSpin User role		This tenant is a member of the PlateSpin Administrators group on the PlateSpin Migrate server and is also on the vCenter Server. If the tenant will be granted the ability to change the resources used by the VM (that is, networks, ISO images, and so forth), grant this user the necessary permissions on those resources. For example, if want to you allow the customer to change the network where their VM is attached, this user should be assigned the Read-only role (or better) on all of the networks being made accessible to the customer.

Figure 13-3 illustrates a Virtual Infrastructure in the vCenter console. The objects labeled in blue are assigned the Infrastructure Manager role. The objects labeled in green are assigned the Virtual Machine Manager role. The tree does not show VM Folders, Networks, and Datastores. Those objects are assigned the *PlateSpin Virtual Machine Manager* role.

Figure 13-3 *Roles assigned in vCenter*

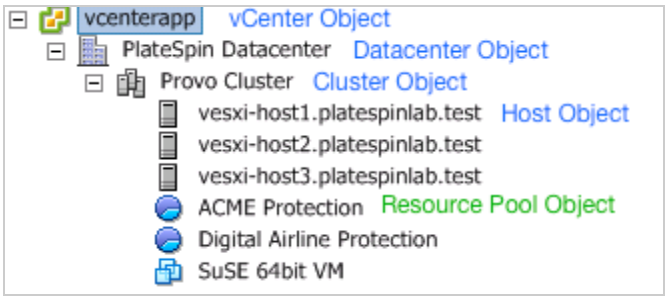

Security Implications of Assigning VMware Roles in a Multitenancy Environment

PlateSpin Migrate software uses a tenant-based PlateSpin user only to perform actions in the migration workflow. From your perspective as a service provider, an end user never has access to the tenant-based PlateSpin user's credentials and is unable to access the same set of VMware resources. In an environment where multiple Migrate servers are configured to use the same vCenter environment, Migrate prevents possibilities for cross-client access.

The major security implications include:

- With the *PlateSpin Virtual Infrastructure Manager* role assigned to the vCenter object, each tenant-based PlateSpin user can see (but not affect) the tasks performed by every other user.
- Because there is no way to set permissions on datastore folders/subfolders, each tenant-based PlateSpin user with permissions on a datastore has access to all other tenant users' disks stored on that datastore.
- With the *PlateSpin Virtual Infrastructure Manager* role assigned to the Cluster object, each tenant-based PlateSpin user is able to turn off/on HA or DRS on the entire cluster.
- Setting the *PlateSpin Virtual Infrastructure Manager* role on the DRS Cluster object and propagating this role allows the tenant-based PlateSpin user to see all VMs placed in the default resource pool and/or default VM folder. Also, propagation requires the administrator to explicitly set the tenant-based PlateSpin user to have a "no-access" role on every resource pool/ VM folder that he or she should not have access to.
- Setting the *PlateSpin Virtual Infrastructure Manager*role on the vCenter object allows the enabled user to end sessions of any other user connected to the vCenter.

NOTE: Remember, in these scenarios, different tenant-based PlateSpin users are actually different instances of the PlateSpin software.

Configuring VMware Roles for PlateSpin Migrate

To migrate workloads to a VMware environment, PlateSpin Migrate Server must have permissions on multiple types of resources on the vCenter Server. You grant permissions to a PlateSpin user that represents the PlateSpin Server by assigning custom roles to the user on different resources.

Migrate provides information about the custom VMware roles and the minimum required privileges for each role in the `<Migrate-install-folder>\PlateSpin Migrate Server\bin\VMwareRolesTool` directory on your Migrate Server. For your convenience, it also provides the PlateSpin VMware Role Tool that you can use to easily create and configure the essential roles. Table 13-1 describes the files provided for this purpose.

Table 13-5 *Migrate Resources for Configuring Custom VMware Roles and Permissions*

File Name	Description
`PlateSpinRole.xml`	The PlateSpin Role XML file defines custom VMware roles and the minimum required privileges for each role.
`PlateSpin.VMwareRoleTool.exe`	The PlateSpin VMware Role Tool creates essential custom VMware roles in a VMware environment and sets the minimum required privileges for each role.

Use the information in this section to create essential custom roles on a target vCenter Server and set the minimum set of privileges for each role.

- “Using the PlateSpin VMware Role Tool to Define PlateSpin VMware Roles in vCenter” on page 259
- “(Optional) Manually Defining the PlateSpin Roles in vCenter” on page 262
- “Using vCenter to View Privileges for PlateSpin Custom Roles” on page 262

Using the PlateSpin VMware Role Tool to Define PlateSpin VMware Roles in vCenter

You can use the PlateSpin VMware Role Tool to conveniently establish the following three roles for a PlateSpin user in a VMware environment:

- PlateSpin Virtual Machine Manager
- PlateSpin Virtual Infrastructure Manager
- PlateSpin User

These roles will contain all the necessary VMware permissions to complete a migration. You will assign the roles to appropriate resource types for the PlateSpin user:

- Table 13-2, “PlateSpin User: Role Assignments for VMware Resource Types,” on page 252
- Table 13-3, “Tenant-Based PlateSpin User: Role Assignments for the Tenant’s VMware Resource Types,” on page 255
- Table 13-4, “Tenant-Based PlateSpin User: Role Assignments for the Tenant User,” on page 257

This section provides usage guidelines for the PlateSpin VMware Role Tool.

Location

The PlateSpin VMware Role Tool (`PlateSpin.VMwareRoleTool.exe` file) is available in the `<Migrate-install-folder>\PlateSpin Migrate Server\bin\VMwareRolesTool` directory on the PlateSpin Migrate Server.

Syntax

From the location where the role tool is installed, run the tool from the command line, using this basic syntax:

```
PlateSpin.VMwareRoleTool.exe /host=<hostname-or-IP-address> /
user=<vCenter-admin-user-name> /role=<the role definition file name and
location> /create
```

Parameters

/host=<*hostname-or-IP-address*>

Specifies the host name or IP address of the vCenter Server where the custom roles will be created or updated.

/user=<*vCenter-admin-user-name*>

Specifies the user name of an administrator account on the vCenter Server.

/role=<*the role definition file name and location*>

Specifies the file name of the file that defines the custom roles and their required minimum privileges. Use the `PlateSpinRole.xml` file unless you are advised to use a different file.

/help

Displays information about parameters, actions, and optional flags.

Actions

Apply the following action parameters as needed when you use `PlateSpin.VMwareRoleTool.exe` to create or update custom roles in vCenter.

/create

(Mandatory) Creates the roles defined by the file you specified for the `/role` parameter.

The default roles defined in the `PlateSpinRole.xml` file are:

- PlateSpin Virtual Machine Manager role

- PlateSpin Virtual Infrastructure Manager role
- PlateSpin User role

/get_all_privileges

Display all server-defined privileges.

/get_compatible_roles

Display all roles that are compatible to the role defined by `/role`.

/check_role=<*role name*>

Check the given role for compatibility with the role defined by `/role`.

Options

Apply the following optional flags as needed when you use `PlateSpin.VMwareRoleTool.exe` to create or update custom roles in vCenter.

/interactive

Run the tool with interactive options that allow you to choose to create individual roles, check role compatibility, or list all compatible roles.

For information about using the tool in interactive mode, see *VMware Role Tool to Verify Permissions to the Roles* (KB 7018547) (https://support.microfocus.com/kb/doc.php?id=7018547).

/password=<*password*>

Provide the VMware password (bypasses the password prompt).

/verbose

Display detailed information.

Tool Usage Example

Usage:

```
PlateSpin.VMwareRoleTool.exe /host=houston_sales /user=pedrom /
role=PlateSpinRole.xml /create
```

Resulting Actions:

1. The PlateSpin VMware Role Tool runs on the `houston_sales` vCenter Server, which has an administrator with the user name `pedrom`.
2. In the absence of the `/password` parameter, the tool prompts for the user password, which you enter.
3. The tool accesses the role definition file, `PlateSpinRole.xml`, which is located in the same directory as the tool executable (there was no need to further define its path).
4. The tool locates the definition file and is instructed (`/create`) to create the roles defined in the contents of that file in the vCenter environment.
5. The tool accesses the definition file and creates the new roles (including the appropriate minimum privileges for defined, limited access) inside vCenter.

You will later assign the roles to appropriate resource types for the PlateSpin user:

- Table 13-2, “PlateSpin User: Role Assignments for VMware Resource Types,” on page 252
- Table 13-3, “Tenant-Based PlateSpin User: Role Assignments for the Tenant’s VMware Resource Types,” on page 255
- Table 13-4, “Tenant-Based PlateSpin User: Role Assignments for the Tenant User,” on page 257

Additional Information

For information about using the tool, see *VMware Role Tool to Verify Permissions to the Roles* (KB 7018547) (https://support.microfocus.com/kb/doc.php?id=7018547).

(Optional) Manually Defining the PlateSpin Roles in vCenter

You can use the vCenter client to manually create and assign the PlateSpin custom roles. This requires creating the roles with the enumerated privileges as defined in `PlateSpinRole.xml`. When you manually create roles:

- You can use the PlateSpin-defined names or custom names for the roles.
- Each role must have the required appropriate minimum privileges from the PlateSpin definition file.

For more information about how to create custom roles in vCenter, see *Using Roles to Assign Privileges* in the VMware vSphere 6.7 Documentation.

Using vCenter to View Privileges for PlateSpin Custom Roles

You use the vCenter client to view the minimal privileges set for the PlateSpin custom roles.

1. In vCenter, select a custom role:
   - PlateSpin Virtual Machine Manager
   - PlateSpin Virtual Infrastructure Manager
   - PlateSpin User
   - PlateSpin Datastore Manager
   - PlateSpin Network Manager
   - PlateSpin Cluster Manager
   - PlateSpin VM User
2. Click **Edit** to view the privileges settings in the Edit Role dialog.

For example, the following figure shows some of the privileges set for the PlateSpin Virtual Machine Manager role.

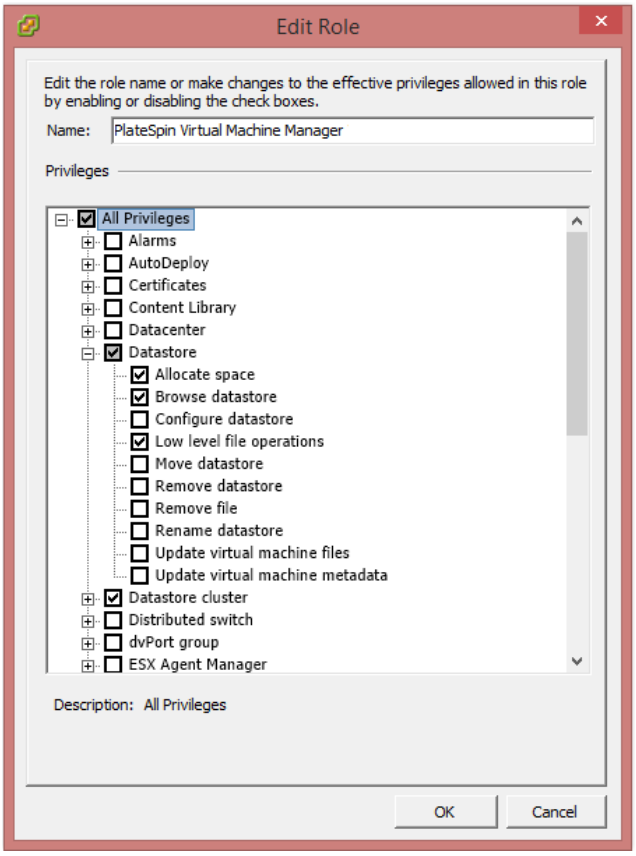

Configuring the SCSI Adapter and VMware Tools for the Target Workload

By default, PlateSpin Migrate configures the SCSI adapter and VMware tools that are the best choice for the OS on the target workload. PlateSpin Migrate controls this behavior by using the **UseSameAsSourceVMwareSCSIAdapterType** parameter that is set to **False** by default. To enable PlateSpin Migrate to configure the same SCSI adapter as on the source workload and retain the source workload VMware tools, set the value of **UseSameAsSourceVMwareSCSIAdapterType** parameter to **True**.

NOTE: iSCSI disks mounted on a source workload are migrated as local disks to the target workload. See "Supported Workload Storage" on page 35.

To use the same SCSI adapter and retain the VMware tools as on the source workload:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2. Search to locate the **UseSameAsSourceVMwareSCSIAdapterType** parameter and set its value to True.

3 Save your settings and exit the page.

A reboot or restart of PlateSpin services is not required to apply the changes.

Retaining the Source VMware Tools on the Target Workload

For workloads configured with the **Install VM Tools** option, the default behavior of PlateSpin Migrate is to remove the source workload VMware tools and install target workload-specific VMware tools on the target workload. PlateSpin Migrate controls this behavior by using the **RetainSourceVMTools** parameter that is set to **False** by default. To enable PlateSpin Migrate to retain the source VMware tools on the target workload, set the value of **RetainSourceVMTools** to **True**.

To retain the source VMware tools on the target workload:

1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2 Search to locate the **RetainSourceVMTools** parameter and set its value to **True**.

3 Save your settings and exit the page.

A reboot or restart of PlateSpin services is not required to apply the changes.

Checklist for Automated Migration to VMware

Task	Description
1. Prepare your VMware migration environment.	Figure 13-1, “Automated Migration to VMware,” on page 248. “Planning for Migration to VMware” on page 249
2. Discover target virtualization platform.	“Discovering Details for Target Platforms” on page 283
3. Discover source workloads.	“Workload Discovery in the Migrate Client” on page 301 -OR- “Workload Discovery in the Migrate Web Interface” on page 302 -OR- “Registering Workloads and Discovering Details with Migrate Agent” on page 303

Task	Description
4. Configure target workload migration.	"Automated Migration to VMware Using Migrate Client" on page 489 -OR- "Automated Migration to VMware Using Migrate Web Interface" on page 498
5. Execute migration.	Chapter 33, "Executing Workload Migrations," on page 537

Checklist for Semi-Automated Migration to Target VMs on VMware

Task	Description
1. Prepare your VMware migration environment.	Figure 13-2, "Semi-Automated Migration to VMs on VMware," on page 249 "Planning for Migration to VMware" on page 249
2. Discover target virtualization platform.	"Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO" on page 288
3. Discover source workloads.	"Workload Discovery in the Migrate Client" on page 301 -OR- "Workload Discovery in the Migrate Web Interface" on page 302
4. Configure target workload migration.	"Migration to VMs on VMware Using X2P Workflow" on page 506
5. Execute migration.	Chapter 33, "Executing Workload Migrations," on page 537

Best Practices for Maintaining or Updating VMware Environments That Are Configured as Migration Targets

Use the following best practices for update and maintenance of VMware DRS clusters and its member hosts that are configured as target platforms in PlateSpin Migrate.

Before you begin VMware maintenance or update:

1. In PlateSpin Migrate, do the following for all workloads associated with the target VMware host:
   - **1a** Pause all scheduled migrations.
   - **1b** Wait for any in-progress full replications and incremental replications to complete, or abort the replications.
   - **1c** Wait for any in-progress cutovers or test cutovers to complete.
2. As a precaution, back up the PlateSpin migration database by using the PlateSpin Import/Export utility (`ImportExportAll.bat`).

   See "Exporting Workload Migration Data" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

During the VMware maintenance or update:

- In PlateSpin Migrate:
  - Do not refresh, modify, or delete the target VMware DRS cluster or host.
  - Do not refresh, re-configure, or delete source workloads that are already configured for migration to the target VMware DRS cluster or host.
  - Do not configure additional migrations to the target VMware DRS cluster or host.
- In the VMware environment:
  - Ensure that IP addresses, hostname, number of NICs, and so on for the target VMware DRS cluster and hosts do not change.
  - Follow VMware best practices for maintenance or update with regard to VM handling in your VMware environment. You might need to relocate target VMs to alternate hosts or power off all VMs on the host.
  - As you complete the maintenance or update, ensure that target VMs are returned to their prior host and power on state.

After you complete the maintenance or update:

1. In the PlateSpin Migrate Web Interface, refresh the VMware target.
2. Unpause migrations to resume scheduled migrations for all workloads associated with the target VMware host.
3. If you need to recover migration data, import the PlateSpin migration database by using the PlateSpin Import/Export utility (`ImportExportAll.bat`).

   See "Importing Workload Migration Data" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.

14 Prerequisites for Migration to VMs on KVM

This section describes the required KVM configuration that you must prepare before you can discover target VMs and configure migrations to them.

Deployment for Migration to KVM

Figure 14-1 shows the location of various components in your semi-automated KVM migration environment and the communications between them.

NOTE: Figure 14-1 depicts automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443). See "Requirements for Workload Registration" on page 56 and "Registering Workloads and Discovering Details with Migrate Agent" on page 303.

Figure 14-1 *Semi-Automated Migration to VMs on KVM*

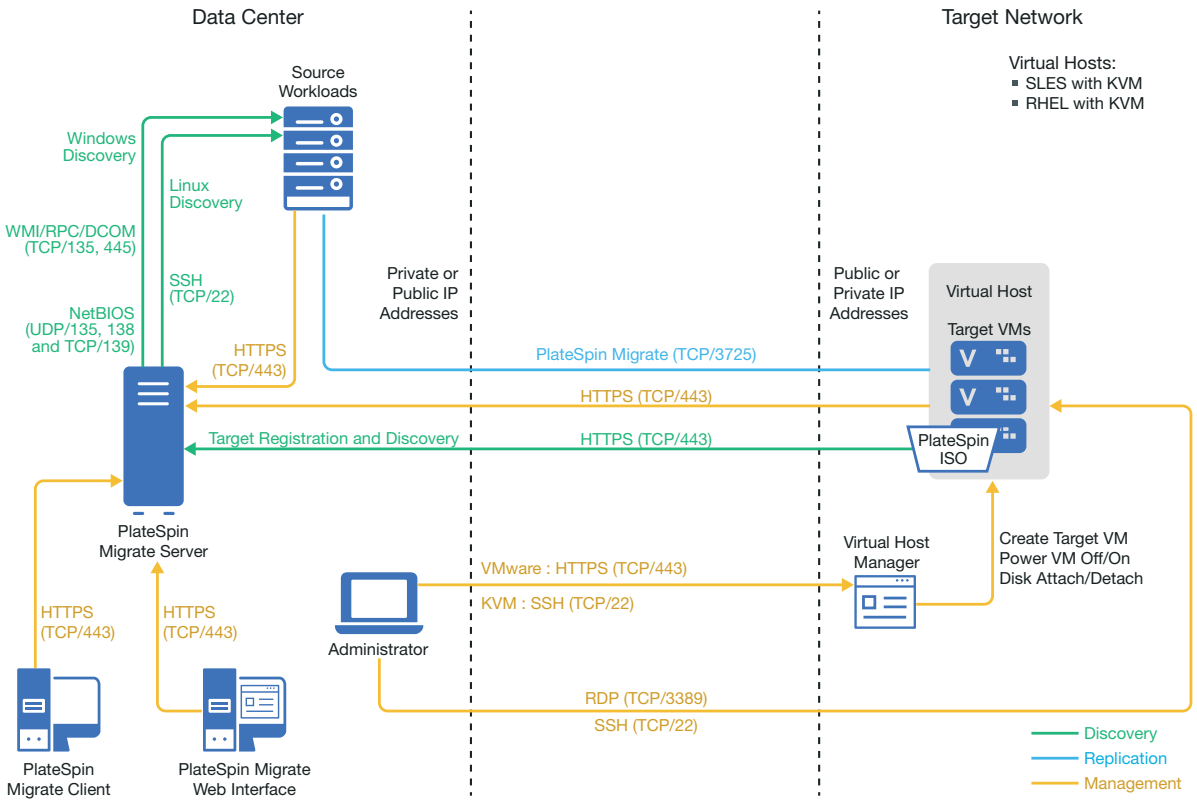

Planning for Migration to VMs on KVM

Ensure that your KVM environment meets the following prerequisites for migration to VMs on KVM:

- Use PlateSpin Migrate Client or PlateSpin Migrate Web Interface to migrate workloads to virtual machines on KVM virtual hosts.
- You can use KVM as the target virtualization platform in a semi-automated workload migration.
- Your target must be a fully virtualized (not paravirtualized) VM.
- Your source workload must be supported by PlateSpin Migrate and KVM.

  See the following information in Table 2-14, "Supported Target KVM Platforms for the Migrate Client Only," on page 42.

  - "SUSE Linux Enterprise Server (SLES) with KVM"
  - "Red Hat Enterprise Linux (RHEL) with KVM"
- Your network environment must meet the requirements for access, discovery, and migration described in "Access and Communication Requirements across Your Migration Network" on page 53.
- Configure volumes on the target disks with about 50 MB of additional storage space than the source disks.
- When you use Virtio disks in the target VM on a KVM host, ensure that you configure the target VM with the appropriate disk type as the boot disk:
  - **Virtio and IDE disks:** Configure the IDE disk as the boot disk and the Virtio disk as the data disk.
  - **Virtio and non-IDE disks:** Configure the Virtio disk as the boot disk and a non-IDE disk such as SATA or SCSI disk as the data disk.

For information about configuring semi-automated migration to a virtual machine on KVM, see "Migration to Virtual Machines on KVM" on page 511.

Checklist for Semi-Automated Migration to Target VMs on KVM

Task	Description
1. Prepare your KVM migration environment.	Figure 14-1, "Semi-Automated Migration to VMs on KVM," on page 267 "Planning for Migration to VMs on KVM" on page 268
2. Discover target virtualization platform.	"Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO" on page 288

Task	Description
3. Discover source workloads.	“Workload Discovery in the Migrate Client” on page 301 -OR- “Workload Discovery in the Migrate Web Interface” on page 302
4. Configure target workload migration.	Chapter 30, “Migration to Virtual Machines on KVM,” on page 511
5. Execute migration.	Chapter 33, “Executing Workload Migrations,” on page 537

15 Prerequisites for Migration to Physical Machines

This section describes the required configuration for migrations to physical machines.

Deployment for Migration to Physical Machines

Figure 15-1 shows the location of various components in your semi-automated physical machine migration environment and the communications between them.

NOTE: Figure 15-1 depicts automated discovery and the network requirements for Windows and Linux workloads.

Figure 15-1 *Semi-Automated Migration to Physical Machines*

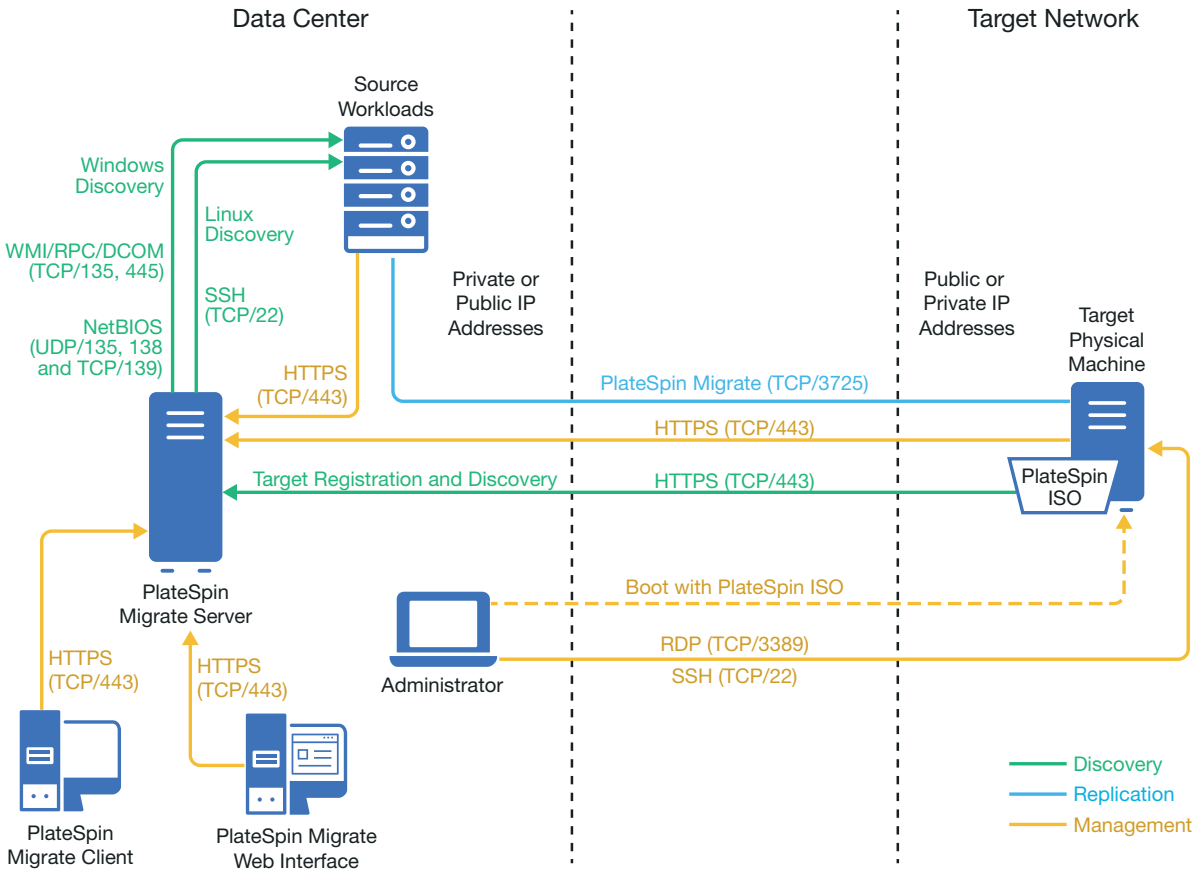

Planning for Migration to Physical Machines

Ensure that your environment meets the following prerequisites for migration to physical machines:

- Use PlateSpin Migrate Client or PlateSpin Migrate Web Interface to migrate workloads to a target physical machine.
- Your physical hardware must be supported by PlateSpin Migrate. See the following information in "Supported Configurations" on page 23:
  - Supported Workload Storage
  - Supported Workload Architectures
- Your network environment must meet the requirements for access, discovery, and migration described in "Access and Communication Requirements across Your Migration Network" on page 53.
- Configure volumes on the target disks with about 50 MB of additional storage space than the source disks.

For information about configuring semi-automated migration to a physical machine, see "Migration to Physical Machines" on page 515.

Best Practices (X2P)

- When you are migrating a workload from one vendor to a target hardware infrastructure from another vendor (for example, from HP to Dell), or if your source is a virtual machine, ensure that you disable vendor-specific or VM-specific services during the transfer. For example, disable the HP Insight service and the VMware Tools service.

  See "Windows HAL or Kernel File Replacements" on page 408.

- Ensure that vendor partitions are not being copied from the source.

  See "Storage Disks and Volumes" on page 422.

Checklist for Semi-Automated Migration to Physical Machines

Task	Description
1. Prepare your physical migration environment.	Figure 15-1, "Semi-Automated Migration to Physical Machines," on page 272 "Planning for Migration to Physical Machines" on page 272
2. Discover target physical platforms.	"Registering and Discovering Details for Target Physical Machines with PlateSpin ISO" on page 291
3. Discover source workloads.	"Workload Discovery in the Migrate Client" on page 301 -OR- "Workload Discovery in the Migrate Web Interface" on page 302
4. Configure target workload migration.	"Configuring Migration to a Physical Target (P2P, V2P) Using Migrate Client" on page 516 -OR- "Configuring Migration to a Physical Target (P2P, V2P) Using Migrate Web Interface" on page 521
5. Execute migration.	Chapter 33, "Executing Workload Migrations," on page 537

16 Preparing for Synchronization of Workloads with Server Sync

For information about synchronizing workloads to synchronize just the data that is different between the source and a target, see Chapter 32, “Synchronizing Workloads with Server Sync,” on page 527.

IV Discovering and Preparing Workloads and Targets

Before you can configure migrations, you must identify your planned target platforms and source workloads. You get details about targets and workloads through a discovery and inventory process.

17 Discovering Target Platforms

Discovery refers to the process of adding unmanaged workloads and platforms in your network and retrieving information about them. For any workload migration, you must have a discovered source and a discovered target platform. For semi-automated migrations, the target is a virtual machine or a physical machine. A target discovery operation populates the PlateSpin Migrate database with detailed inventory information about the target host and its resources. The inventory provides the data necessary to determine the host's use and to properly configure one or more migrations to the target host.

About Target Discovery

PlateSpin Migrate Web Interface and PlateSpin Migrate Client provide automated discovery and inventory of supported target host platforms. See Table 8-1 for an overview of the target host discovery capabilities of each tool.

Table 17-1 *Supported Target Host Discovery Capabilities*

Target Host Discovery	Migrate Client	Web Interface
Cloud Targets		
Amazon Web Services (Cloud Region)	✘	✓
Microsoft Azure (Cloud Location)	✘	✓
VMware vCloud Director (Organization)	✘	✓
Oracle Cloud infrastructure	✓	✓

Target Host Discovery	Migrate Client	Web Interface
VMware Targets		
VMware DRS Cluster (A vCenter Cluster is the target; any available node might be used for the VM.)	✓	✓
VMware DRS Cluster as Hosts (Each VMware ESX host in a vCenter Cluster is a potential target.)	✗	✓
VMware ESX Server	✓	✓
Other Targets		
Linux KVM virtual host	✓	✗
Physical host	✓	✓
Discovery Capabilities		
An individual host server	✓	✓
Multiple virtual host servers at a time	✓	✓
All hosts in a domain	✓	✗
Refresh Target Discovery	✓	✓

You can view discovered target platforms in the Targets list in either tool:

- **Web Interface:** The Targets list includes:
  - All cloud and VMware hosts discovered using the Web Interface
  - All VMware hosts in the default network discovered using Migrate Client

    NOTE: Use the Web Interface to discover target cloud and VMware hosts in non-default networks if you plan to use the Web Interface for migrations to those locations.

  All target hosts displayed in the Web Interface Targets list are supported as migration targets using the Web Interface. See Table 17-1, “Supported Target Host Discovery Capabilities,” on page 279.
- **Migrate Client:** The Targets list includes all the discovered VMware target hosts, no matter where you initiated discovery.

For information about the target hosts that the Web interface and the Migrate Client supports, see Table 17-1, “Supported Target Host Discovery Capabilities,” on page 279.

Network Access Requirements for Target Host Discovery

For information about network access requirements for discovery of target hosts, see “Requirements for Discovery” on page 53.

Discovery Guidelines for Target Hosts

For information about the software, network, and firewall requirements that systems in your environment must meet for the discovery and inventory process, see "Requirements for Discovery" on page 53.

Target Host Discovery Parameters for Migrate Web Interface

Table 17-2 provides guidelines for target type selection, credential format, and syntax for discovery parameters for target hosts using the Migrate Web Interface.

Table 17-2 *Guidelines for Migrate Web Interface Target Type and Credentials for Target Hosts*

To Discover	Target Type	Credentials	Remarks
Amazon Cloud Region	**Amazon Cloud Region**	IAM role or Access Key ID and Secret Key ID	If you are using a AWS-based Migrate server that has an IAM role attached, PlateSpin Migrate by default uses the attached IAM role for accessing the AWS account. However, you can override this default behavior and use the Access Key ID and Secret Key ID credentials for accessing the AWS account. See Table 17-4, "Options for Amazon Cloud Region," on page 285.
Azure Cloud Location	**Microsoft Azure Location**	Subscription ID Application ID Azure user with Subscription administrator role	You must create an Azure Application ID for PlateSpin Migrate to use when it uses the Azure API for replicating and migrating workloads to your Azure Global. See "Configuring an Application in Azure to Represent PlateSpin Migrate" on page 182.
Azure Stack Location	**Microsoft Azure Location**	Subscription ID Application ID Management Endpoint Storage Endpoint Suffix Azure user with Subscription administrator role	You need NOT create an Azure Application ID for PlateSpin Migrate to use when it uses the Azure API for replicating and migrating workloads to your Azure Stack account. Specify the following value for the Application ID: `872cd9fa-d31f-45e0-9eab-6e460a02d1f1`

To Discover	Target Type	Credentials	Remarks
VMware vCenter Cluster	**VMware DRS Cluster**	VMware vCenter Web service credentials (user name and password) To configure minimal permissions needed, see "Configuring a PlateSpin User with Minimal Permissions on VMware" on page 250.	All subsequent communications with ESX hosts in the Cluster take place through the vCenter Server. VMware high availability and DRS rules apply for a target VM except during replications. The VM can reside on any available node.
VMware ESXi Hosts managed in a VMware vCenter Cluster	**VMware DRS Cluster as Hosts**	VMware vCenter Web service credentials (user name and password) To configure minimal permissions needed, see "Configuring a PlateSpin User with Minimal Permissions on VMware" on page 250.	Each host in the vCenter Cluster appears as a separate potential target in the Web Interface. All subsequent communications with each ESX host take place through the vCenter Server. High availability and DRS rules apply for a target VM except during replications. The VM must reside on the designated host for prepare, replication, test cutover, and cutover actions.
VMware ESXi host	**VMware ESX Server**	ESX account with administrator role To configure minimal permissions needed, see "Configuring a PlateSpin User with Minimal Permissions on VMware" on page 250.	
vCloud Organization	**VMware vCloud Organization**	Organization Administrator credentials (user name and password)	

Target Host Discovery Parameters for Migrate Client

Table 17-3 provides guidelines for machine type selection, credential format, and syntax for discovery parameters for target hosts using the Migrate Client.

Table 17-3 *Guidelines for Migrate Client Machine Type and Credentials for Target Hosts*

To Discover	Machine Type	Credentials	Remarks
VMware ESX hosts affiliated with a VMware vCenter Server	**VMware vCenter**	VMware vCenter Web service credentials (user name and password) To configure minimal permissions needed, see "Configuring a PlateSpin User with Minimal Permissions on VMware" on page 250.	
VMware ESX hosts	**VMware ESX**	ESX account with administrator role To configure minimal permissions needed, see "Configuring a PlateSpin User with Minimal Permissions on VMware" on page 250.	
All Linux KVM virtual hosts	**Linux**	Root-level username and password	Non-root accounts must be properly configured to use `sudo`. See KB Article 7920711 (https://support.microfocus.com/kb/doc.php?id=7920711).

Discovering Details for Target Platforms

Before you configure a migration job, you must discover and perform an inventory of the target platform. The inventory collects information about the host platform and its resources, such as the amount of RAM, number of cores and processors, datastores, networks, and resource groups.

- "Target Discovery in the Migrate Client" on page 283
- "Target Discovery in the Web Interface" on page 284

Target Discovery in the Migrate Client

In Migrate Client, you can discover:

- An individual virtual machine host server
- Multiple virtual machine host servers
- All VMware ESX hosts affiliated with a VMware vCenter Server

Before you begin discovery operations, ensure that PlateSpin Server can communicate with your source workloads and targets. See "Requirements for Discovery" on page 53.

To discover targets using Migrate Client:

1 In the Migrate Client toolbar, click **Discover Details**.

or

In the Servers view, right-click in a blank area, then select **Discover Details**.

2 In the **Discover Details** dialog box, type the host name or IP address of the target.

To discover multiple machines, specify multiple host names or IP addresses separated by semicolons. If the target is behind a NAT device, specify its public (external) IP address.

See "Migrations Across Public and Private Networks through NAT" on page 60.

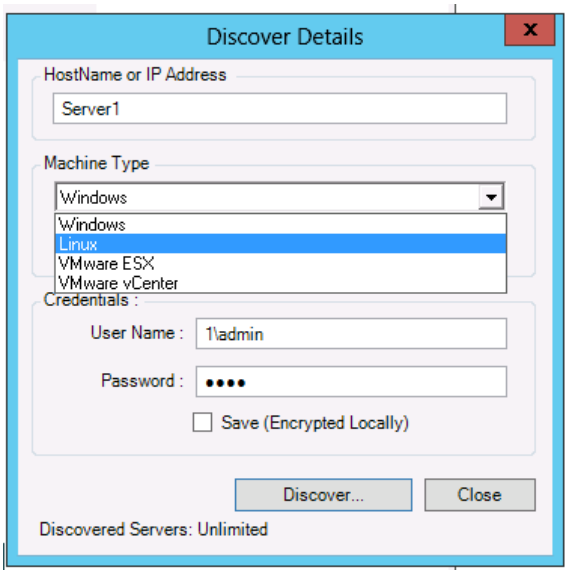

3 Select the machine type for the target platform. If you select VMware vCenter, also provide the name of the vCenter cluster.

- Windows
- Linux
- VMware ESX
- VMware vCenter

See "Discovery Guidelines for Target Hosts" on page 281.

4 Provide administrator credentials for the machine you are discovering.

See "Discovery Guidelines for Target Hosts" on page 281.

5 (Optional) If you want to store these credentials for use during future jobs, enable the **Save (Encrypted Locally)** option.

6 Click **Discover** and wait for the process to complete.

7 (Optional) If you want to monitor the progress of the job, switch to the Jobs view.

Target Discovery in the Web Interface

To migrate a workload through the Web Interface, you must first add or discover the intended target platform and its resources.

PlateSpin Migrate Web Interface supports discovery of virtual and cloud target platforms:

- Amazon Cloud Region
- Microsoft Azure Location

- VMware DRS Cluster (The cluster appears in the Targets list.)
- VMware DRS Cluster as Hosts (Each host in the cluster appears in the Targets list, but not their parent cluster.)
- VMware ESX Server
- VMware vCloud Organization

When you add the target, its associated resources are automatically discovered. You can add one platform at a time. All available target platforms are listed on the Targets page.

Before you begin discovery operations, ensure that PlateSpin Server can communicate with your source workloads and targets. See section "Requirements for Discovery" on page 53.

To add a target platform:

1. In the Migrate Web Interface, click **Targets** > **Add Target**.
2. Select one of the following target types:
   - Amazon Cloud Region
   - Microsoft Azure Location
   - VMware DRS Cluster
   - VMware DRS Cluster as Hosts
   - VMware ESX Server
   - VMware vCloud Organization
3. Depending on the type of targets you selected in the previous step, specify the appropriate access information.
   - **Amazon Cloud Region:** See Table 17-4.
   - **Microsoft Azure Location:** See Table 17-5.
   - **VMware DRS Cluster:** See Table 17-6.
   - **VMware DRS Cluster as Hosts:** See Table 17-7.
   - **VMware ESX Server:** See Table 17-8.
   - **VMware vCloud Organization:** See Table 17-9.

Table 17-4 *Options for Amazon Cloud Region*

Option	Description
AWS Cloud	Select one of the following appropriate AWS environment for the target AWS platform. By default, AWS Global Cloud is selected. • AWS China Cloud • AWS Global Cloud • AWS Government Cloud

Option	Description
This Migrate Server instance has an IAM role attached. Use the IAM role to access Amazon EC2 Region	When you use an AWS-based Migrate server that has an IAM role attached, this option displays in the user interface and is selected by default. PlateSpin Migrate uses the attached IAM role for accessing the AWS account. However, to override this default behavior and use the Access Key ID and Secret Key ID credentials for accessing the AWS account, you must deselect this option.
Access Key ID	Specify the access key ID for your AWS account. This option is not displayed if the **This Migrate Server instance has an IAM role attached. Use the IAM role to access Amazon EC2 Region** option is selected.
Secret Key ID	Specify the secret key ID required to access your AWS account. This option is not displayed if **This Migrate Server instance has an IAM role attached. Use the IAM role to access Amazon EC2 Region** option is selected.
Region Name	Select the region for the Amazon target.

Table 17-5 *Options for Microsoft Azure Location Target*

Option	Description
Azure Cloud	Select one of the following appropriate Azure environment for the target Azure platform. By default, Azure Global is selected. ◆ Azure Global ◆ Azure Stack
Management Endpoint	(Applicable for Azure Stack only) Specify the Azure Resource Manager endpoint that you configured for the Azure Stack. For Azure Stack Development kit, this value is set to `https://management.local.azurestack.external`.
Storage Endpoint Suffix	(Applicable for Azure Stack only) Specify the Storage endpoint that you configured for the Azure Stack. For Azure Stack Development kit, this value is set to `local.azurestack.external`.
Subscription Id	Specify the subscription ID for your Microsoft Azure account.

Option	Description
Application Id	Use either of the following to specify your Azure Application ID required to enable PlateSpin Migrate to use the Azure APIs when it replicates or migrates workloads on your behalf to VMs in the target Azure account: ◆ **For Azure Cloud:** Create an Azure Application ID for PlateSpin Migrate to use when it uses the Azure API for replicating and migrating workloads to your Azure Global account. See "Configuring an Application in Azure to Represent PlateSpin Migrate" on page 182. ◆ **For Azure Stack:** Specify the following value for the Application ID: `872cd9fa-d31f-45e0-9eab-6e460a02d1f1`
Username and Password	Specify administrator-level credentials for accessing the parent Microsoft Azure account.
Location Name	Select the location for the Microsoft Azure target. Click **Update Location List** to refresh the list of available locations in the menu. For predefined Azure Cloud environments, locations are sorted by the geographical region and alphabetically. The mapping is fixed and is based on the current categories that Azure uses. If Microsoft Azure adds new locations after the current release, Migrate displays them dynamically and alphabetically in the **Recently Added** category.

Table 17-6 *Options for VMware DRS Cluster Target*

Option	Description
vCenter Hostname or IP	Specify the host name or IP address of the vCenter server.
Cluster Name	Specify the name of the DRS cluster. This is applicable only for VMware DRS Cluster.
Username and Password	Specify administrator-level credentials for accessing the target host.

Table 17-7 *Options for VMware DRS Cluster as Hosts Target*

Option	Description
vCenter Hostname or IP	Specify the host name or IP address of the vCenter server.
Cluster Name	Specify the name of the DRS cluster. This is applicable only for VMware DRS Cluster.
Username and Password	Specify administrator-level credentials for accessing the target host.

Table 17-8 Options for VMware ESX Server Target

Option	Description
Hostname or IP	Specify the host name or IP address of the VMware ESX server.
Username and Password	Specify administrator-level credentials for accessing the target host.

Table 17-9 Options for VMware vCloud Organization Target

Option	Description
vCloud Director Server Address	Specify the server host name or the IP address of the vCloud Director server. For example: `cloud.example.com` or `10.10.10.101`
Organization Name	Specify the name of the organization in the vCloud Director server. The name is case sensitive in vCloud. Type the name exactly as you created it. For example: `DemoOrg001`
Username and Password	Specify the organization-level administrator credentials for accessing the target host. For example: `demouser1` and `demopwd`

4 Click **Test Credentials** to validate the credential values you specified.

5 Click **Add** to add and discover details about the target and list it on the Targets page.

IMPORTANT: (For Azure Stack only) The first-time discovery of an Azure Stack might take some considerable time (20-30 mins) because PlateSpin Migrate downloads the `PlateSpin Replication Environment(PRE)` image from the Azure Marketplace and uploads it to the Azure Stack. If the `PlateSpin Replication Environment` does not get successfully uploaded to Azure Stack, then the following messages are logged and you must manually upload the PRE image to the Azure Stack before you configure migrations to it.

```
Create UserImage has falied With Message: {0}
```

```
Create a User Image manually in the portal with Name {0} before continuing with the migration
```

To manually upload the PRE image to the discovered Azure Stack, see "Manually Adding PlateSpin Replication Environment Image to Azure Stack" on page 193.

Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO

PlateSpin Migrate enables you to migrate a source workload to a target virtual machine on a virtual host, where the VM is regarded as a target physical machine:

- VMware

Semi-automated migration to VMs on VMware can be done, but fully automated migration to target VMware platforms is preferred. Discovery for target VMware platforms is available both in the Migrate Client and the Migrate Web Interface. See “Discovering Details for Target Platforms”.

- KVM
- Oracle Cloud Infrastructure

For information about supported virtual host platforms, see “Supported Target Virtualization Platforms” on page 40.

PlateSpin ISO registers the target physical machine with the PlateSpin Migrate server and performs an inventory of the machine to collect information about it, such as the amount of RAM, number of cores and processors, storage disks, and NICs.

Prerequisites for Discovering Target VMs

PlateSpin Migrate does not automatically build the target VM for you on the target virtual host. You must manually set up the target virtual machine with guest operating system type and version settings that match your source workload, in accordance with the features and capabilities of the virtualization platform. You must also prepare the PlateSpin ISO file and attach it as a boot CD for the VM.

1. Download the PlateSpin ISO image for use with the target VM.

   See “Downloading the PlateSpin ISO Images” on page 373.
2. Prepare the PlateSpin ISO image for use with the target VM. Attended and unattended registration options are possible.

   See “Preparing the PlateSpin ISO Image for Target Registration and Discovery” on page 374.
3. Use the native interface of the virtualization platform to create a virtual machine.

   See “Creating and Configuring the Target Virtual Machine (RHEL KVM)” on page 512:
4. Ensure that the VM is configured to restart on reboot and that you attach the PlateSpin ISO file as a boot CD for the VM.

Registering and Discovering Target VMs on Virtual Hosts

After you create and prepare the virtual machine to boot with the PlateSpin ISO, you are ready to register it as a target VM with your PlateSpin Server.

1. From the Virtual Machine Manager, power on (or reboot) the virtual machine, then launch the virtual machine console and monitor the boot process.

   When the virtual machine completes the boot process, it prompts you for parameters that control the registration of the machine and its profile with PlateSpin Migrate. If you are using the unattended registration process, the required parameters are read from an answer file.

2 At the initial boot prompt, type one of the following options, then press Enter:

Boot Option	Boot Action
`ps`	PlateSpin Linux for taking control You can also press Enter to select this option.
`fcoe`	PlateSpin Linux for taking control with FCoE support
`next`	Boot from the next boot device set in the BIOS

If no key is pressed for 20 seconds, the workload boots from the next boot device set in the BIOS.

3 At the command line, provide the required information at each individual prompt:

- **PlateSpin Server:** Enter the PlateSpin Server URL, using the following format:

  `http://`*Your_PlateSpin_Server*`/platespinmigrate`

  Replace *Your_PlateSpin_Server* with the host name or the IP address of your PlateSpin Server host.

- **Credentials (User Name/Password):** Enter the name of an administrator-level user on the PlateSpin Server host, including the domain or machine name. For example: *domain**username*, or *localhost*\Administrator. Provide a valid password for the specified user.

- **Network Card:** Select the network card that is active, then either enter a temporary static IP address for this NIC or press Enter to dynamically obtain an IP address from a DHCP server.

- **Temporary hostname:** Provide a temporary VM name for PlateSpin Migrate Client to use to list the newly registered VM. The workload's target host name you select in the migration job overwrites this name.

- **SSL encryption:** If your PlateSpin Migrate is installed on a host with SSL encryption enabled, enter `Yes`. If not, enter `No`.

- **PlateSpin Migrate Network:** Unless you have defined your own PlateSpin Migrate Network in PlateSpin Migrate Client, press Enter. If you are working with a non-default PlateSpin Migrate Network, type its name, then press Enter.

  A controller on your target virtual machine communicates with PlateSpin Server and registers the virtual machine as a physical target for a migration job.

After a few moments, PlateSpin Migrate Client displays the target virtual machine in the Servers view.

NOTE: If registration fails with an authorization error, you might need to synchronize the clocks of the source and the target, modify the LAN Manager Authentication Level on the target, or both. See Table C-1, "Common Issues and Solutions Related to Discovery Operations," on page 337.

Configuration Information

For information about configuring migration for target VMs on virtual hosts, see the following:

- “Migration to VMs on VMware Using X2P Workflow” on page 506
- “Migration to Virtual Machines on KVM” on page 511

Registering and Discovering Details for Target Physical Machines with PlateSpin ISO

To discover a physical target and inventory its hardware components, you must boot the target machine with the PlateSpin ISO image on a CD or other media from which your target can be booted.

PlateSpin ISO registers the target physical machine with the PlateSpin Migrate server and performs an inventory of the machine to collect information about it, such as the amount of RAM, number of cores and processors, storage disks, and NICs.

- “Prerequisites for Discovering Target Physical Machines” on page 291
- “Registering and Discovering Target Physical Machines” on page 292
- “Configuration Information” on page 293

Prerequisites for Discovering Target Physical Machines

You must prepare the PlateSpin ISO file and attach it as a boot CD for the physical machine.

1. Download the PlateSpin ISO image for use with the target VM.

   See “Downloading the PlateSpin ISO Images” on page 373.
2. Prepare the PlateSpin ISO image for use with the physical machine. Attended and unattended registration options are possible.

   See “Preparing the PlateSpin ISO Image for Target Registration and Discovery” on page 374.
3. Ensure that the physical machine is configured to restart on reboot and that you attach the PlateSpin ISO file as a boot CD.

Registering and Discovering Target Physical Machines

After you create and prepare the physical machine to boot with the PlateSpin ISO, you are ready to register the target machine with your PlateSpin Server.

1. Boot the target machine from the PlateSpin ISO image.
2. At the initial boot prompt, type one of the following options, then press Enter:

Boot Option	Boot Action
ps	PlateSpin Linux for taking control You can also press Enter to select this option.
fcoe	PlateSpin Linux for taking control with FCoE support
next	Boot from the next boot device set in the BIOS

   If no key is pressed for 20 seconds, the workload boots from the next boot device set in the BIOS.

3. At the command line, provide the required information at each individual prompt:
   - **PlateSpin Server:** Enter the PlateSpin Server URL, using the following format:

     `http://`*Your_PlateSpin_Server*`/platespinmigrate`

     Replace *Your_PlateSpin_Server* with the host name or the IP address of your PlateSpin Server host.
   - **Credentials (User Name/Password):** Enter the name of an administrator-level user on the PlateSpin Server host, including the domain or machine name. For example: *domain**username*, or *localhost*\Administrator. Provide a valid password for the specified user.
   - **Network Card:** Select the network card that is active, then either enter a temporary static IP address for this NIC or press Enter to dynamically obtain an IP address from a DHCP server.
   - **Temporary hostname:** Provide a temporary VM name for PlateSpin Migrate Client to use to list the newly registered VM. The workload's target host name you select in the migration job overwrites this name.
   - **SSL encryption:** If your PlateSpin Migrate is installed on a host with SSL encryption enabled, enter `Yes`. If not, enter `No`.
   - **PlateSpin Migrate Network:** Unless you have defined your own PlateSpin Migrate Network in PlateSpin Migrate Client, press Enter. If you are working with a non-default PlateSpin Migrate Network, type its name, then press Enter.

     A controller on your target virtual machine communicates with PlateSpin Server and registers the virtual machine as a physical target for a migration job.

After a few moments, PlateSpin Migrate Client displays the physical target in the Servers view.

NOTE: If registration fails with an authorization error, you might need to synchronize the clocks of the source and the target, modify the LAN Manager Authentication Level on the target, or both. See Table C-1, "Common Issues and Solutions Related to Discovery Operations," on page 337.

Configuration Information

For information about configuring migration to physical machines, see "Migration to Physical Machines" on page 515.

Discovering Target VMs for Server Sync Jobs

If you want to synchronize two workloads, and if your synchronization target is a virtual machine, you must discover and register an appropriate virtual machine first. For information about the Server Sync feature, see "Synchronizing Workloads with Server Sync" on page 527.

1. On your virtual machine host, create a virtual machine with the desired specifications and install the operating system that matches the intended source workload, including the exact service pack.
2. Discover the virtual machine host or refresh its details.
3. In the Servers view, right-click the newly created virtual machine underneath the virtual machine server, then select **Prepare for synchronization**.
4. Specify administrator credentials for the virtual machine server.
5. (Optional) If you want to store these credentials for use during future jobs, enable the **Save (Encrypted Locally)** option.
6. (Optional) To configure the temporary (Take Control) network settings, such as choosing which virtual network to use from those available on the virtual machine server and configuring TCP/IP settings, click **Configure**, then configure the network settings as required.
7. Click **Prepare** and wait for the job to complete.

   On completion, the Servers view lists a new Server Sync target underneath the VM host: 

Refreshing Target Host Details

You should routinely refresh details about your target platforms before setting up or executing a migration job.

- "Refresh Target Details in the Web Interface" on page 293
- "Refresh Target Details in Migrate Client" on page 294

Refresh Target Details in the Web Interface

PlateSpin Migrate Web Interface enables you to refresh the discovered resources for virtual and cloud target platforms:

- Amazon Cloud Region
- Microsoft Azure Location
- VMware DRS Cluster
- VMware DRS Cluster as Hosts
- VMware ESX Server
- VMware vCloud Organization

When you refresh the target, its associated resources are automatically rediscovered and updated. You can refresh one target platform at a time.

To refresh details for a target platform:

1. In the PlateSpin Migrate Web Interface, click **Targets**.
2. Select a target.
3. Click **Refresh**.
4. Expand the panels for the associated resources to view the changes.

Refresh Target Details in Migrate Client

Migrate Client allows you to refresh target details for VMware ESX servers platforms discovered using the Migrate Client:

To refresh target details:

1. In the Servers view, right-click the required item, then select **Refresh Details**.

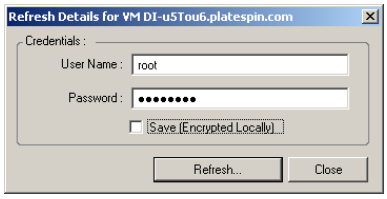

2. Specify the credentials appropriate for the system being refreshed, then click **Refresh**.

   PlateSpin Migrate starts a discovery job, which you can monitor in the Jobs view.

Removing (Undiscovering) Target Platforms

After you complete all migration jobs for a target platforms, you can remove (undiscover) the target platform. You might also remove a target that will not be used.

IMPORTANT

- If an object is listed both in the Migrate Client and the Migrate Web Interface, then you must use the Web Interface to remove the object.
- Before you delete a target platform that is in use for configured jobs, you must ensure that all the affected jobs are completed.
- For potential clean-up of files that might have been copied during discovery on the target platform, ensure that the platform is up and running and that it is reachable before you attempt to remove or undiscover the target.

  NOTE: If this step cannot be attempted, the process reports a failure even though the target platform is successfully removed (undiscovered) from the database and is no longer available in the Migrate Client or Migrate Web Interface.

To undiscover a workload through the Migrate Client:

1. On the Workloads page, right-click the target and select **Undiscover Target**.

To remove a target through the Migrate Web Interface:

1. On the Targets page, click **Remove** next to the target you want to remove from Migrate.

18 Discovering Source Workloads

Discovery refers to the process of adding unmanaged workloads and platforms in your network and retrieving information about them. For any workload migration, you must have a discovered source and a discovered target. A workload discovery operation populates the PlateSpin Migrate database with detailed inventory information about a workload that you want to migrate. The workload inventory provides the data necessary to determine the machine's use and to properly configure its migration.

About Source Workload Discovery

PlateSpin Migrate Web Interface and PlateSpin Migrate Client provide automated discovery and inventory of supported source workloads. See Table 9-1 for an overview of the workload discovery capabilities of each tool.

IMPORTANT

- Before you discover a source workload, you must ensure that the source workload has an active partition. If you discover a source workload that does not have an active partition, the discovery fails. See "The workload cannot be migrated because it has 0 active partitions. Ensure that the workload has exactly 1 active partition and try again" on page 337.
- Discovery of source Windows workloads in AWS requires PowerShell 2.0 or higher on the source workload.

Table 18-1 *SUPPORTED SOURCE WORKLOAD DISCOVERY CAPABILITIES*

Source Workload Discovery Capabilities	Migrate Client	Web Interface
Windows standalone workloads	✓	✓
Windows cluster workloads (to target VMware host)	✓	✓
Linux standalone workloads	✓	✓
Linux cluster workloads	✘	✘
Multiple machines at a time	✓	✘
All machines in a domain	✓	✘
Refresh Source Discovery	✓	✘

The Mass Discover CLI enables you to discover workloads from a CSV file. The related migration jobs start according to the schedules you set for them. See “massdiscover” in “Using the PlateSpin Migrate Client Command Line Interface” on page 573.

As an alternative to Migrate discovery, you can use Migrate Agent to register a workload with the Migrate Server and inventory its details. See Appendix F, “Migrate Agent Utility,” on page 359.

You can view discovered source workloads in the Workloads list in either tool:

- **Web Interface:** The Workloads list includes:
  - All source workloads discovered using the Web Interface
  - Source workloads in the default network discovered using Migrate Client

    NOTE: Use the Web Interface to discover source workloads in non-default networks if you plan to migrate them using the Web Interface.
  - All source workloads registered using the Migrate Agent utility

  All workloads displayed in the Web Interface Workloads list are supported for migration using the Web Interface. See Table 18-1 and “Migration Operations Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface” on page 82.
- **Migrate Client:** The Workloads list includes all discovered source workloads, no matter where you initiated discovery.

  Some workloads in the Migrate Client Workloads list might not be supported for some migration targets using the Migrate Client. See Table 18-1 and “Migration Operations Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface” on page 82.

Network Access Requirements for Workload Discovery

For information about network access requirements for gathering details about source Windows and Linux workloads, see the following as appropriate:

- **Discovery and inventory process:** “Requirements for Discovery” on page 53

  -OR-

- **Registration using Migrate Agent:** “Requirements for Workload Registration” on page 56

Discovery Guidelines for Source Workloads

For information about the software, network, and firewall requirements that systems in your environment must meet before you add workloads to Migrate, see the following information as appropriate:

- **Discovery and inventory process:** “Requirements for Discovery” on page 53

  -OR-

- **Registration using Migrate Agent:** “Requirements for Workload Registration” on page 56

Table 18-2 provides guidelines for machine type selection, credential format, and syntax for discovery parameters for workloads.

Table 18-2 *Guidelines for Machine Type and Credentials for Source Workloads*

To Discover	Machine Type	Credentials	Remarks
All Windows workloads	**Windows**	Local or domain administrator credentials.	For the username, use this format: ◆ For domain member machines: *authority**principal* ◆ For workgroup member machines: *hostname**principal*
All Linux workloads	**Linux**	Root-level user name and password	Non-`root` user accounts must be properly configured to use `sudo`. See KB Article 7920711 (https://support.microfocus.com/kb/doc.php?id=7920711).
Windows workloads in AWS (no VPN connection, C2C migration from AWS to Azure or to vCloud)	**Windows**		For C2C migrations from AWS, log in to the source Windows workload in AWS with RDP, then use Migrate Agent Utility to register the workload. See “Windows Workload Registration and Discovery with Migrate Agent” on page 304.

To Discover	Machine Type	Credentials	Remarks
Linux workloads in AWS (no VPN connection, C2C migration from AWS to Azure or to vCloud)	**Linux**	User name with root-level access and the private key file you created for your AWS EC2 Key Pair	For C2C migrations from AWS, log in to the source Linux workload in AWS with SSH, then use Migrate Agent Utility to register the workload. See "Windows Workload Registration and Discovery with Migrate Agent" on page 304. Non-`root` user accounts must be properly configured to use `sudo`. See KB Article 7920711 (https://support.microfocus.com/kb/doc.php?id=7920711). **NOTE:** For AMI images in AWS, use the default non-`root` user system account that is automatically configured to use `sudo`. To run Migrate Agent commands, run the `sudo -i` command to access the `root` shell, and then run the Migrate Agent commands.

Populating the Servers View with a List of Windows Computers in a Domain

In the PlateSpin Migrate Client, the Network Discovery feature populates the Server view with all Windows physical machines and virtual machines that are online in a specified domain. PlateSpin Migrate uses the standard Windows network browser function for discovery. Because Linux workloads and virtual machine servers do not advertise to the Windows network browser, they are not automatically detected and do not appear in the list.

Unlike a full discovery with inventory, Network Discovery lists the Windows machines but does not inventory each workload to gather its details. A workload inventory is required for migration jobs. You can use either of the following methods to inventory the workloads:

- Use **Discover All Servers** to discover details for each of the listed Windows workloads. See "Discovering Details for All Windows Workloads in a Domain" on page 301.
- Use **Discover Details** to discover details a specific workload. See "Workload Discovery in the Migrate Client" on page 301.

Network Discovery is enabled by default. The option is a toggle between enabled and disabled modes.

To enable or disable Network Discovery:

1. In the Migrate Client, double-click **Network Discovery** at the bottom right corner of the Migrate Client window.

Discovering Details for All Windows Workloads in a Domain

You can use the **Discover All Servers** option in the Servers view to discover and perform an inventory of all Windows workloads in a specified domain. The Network Discovery option must be enabled to detect the Windows servers in the network.

1. In Migrate Client, enable the Network Discovery feature.

   See "Populating the Servers View with a List of Windows Computers in a Domain" on page 300.
2. Expand the list of domains that contain the machines to be inventoried.
3. Right-click the domain name, then select **Discover All Servers**.
4. Specify domain-level administrator credentials.
5. Click **Discover** and wait for the process to complete.
6. (Optional) If you want to monitor the progress of the discovery job, switch to the Jobs view.

Discovering Details for Source Workloads

Before you configure a migration job, you must discover and perform an inventory of the workload. The inventory collects information about the workload such as the server host name, amount of RAM, number of cores and processors, storage disks and volumes, NICs and applications and their start states.

- "Workload Discovery in the Migrate Client" on page 301
- "Workload Discovery in the Migrate Web Interface" on page 302

Workload Discovery in the Migrate Client

In the PlateSpin Migrate Client, you can use the **Discover Details** option in the Servers view to discover and perform an inventory for physical or virtual machines:

- An individual Windows workload
- An individual Linux workload
- Multiple Windows or Linux workloads at a time

Before starting discovery operations, ensure that PlateSpin Server can communicate with your source workloads. See "Requirements for Discovery" on page 53.

To discover workloads using Migrate Client:

1. On the Migrate Client toolbar, click **Discover Details**.

   or

   In the Servers view, right-click in a blank area, then select **Discover Details**.

   or

   In the Servers view, right-click a Windows workload that has been populated through network discovery. then select **Discover Details**.
2. In the **Discover Details** dialog box, type the host name or IP address of the source workload.

To discover multiple machines at a time, specify multiple host names or IP addresses separated by semicolons. If the machine is behind a NAT device, specify its public (external) IP address.

See "Migrations Across Public and Private Networks through NAT" on page 60.

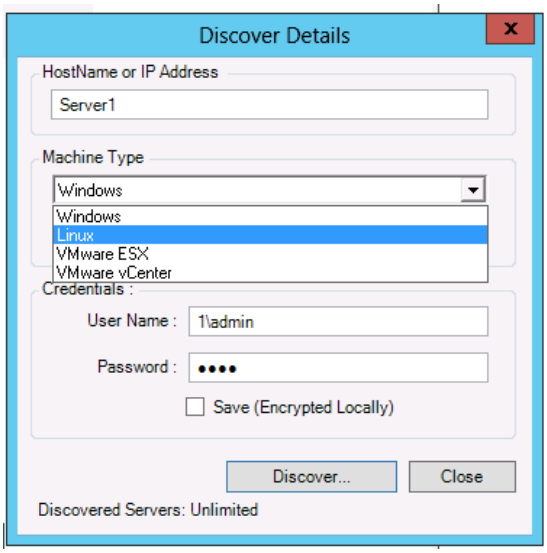

3 Select the machine type for the source workload

- Windows
- Linux

4 Provide administrator credentials for the machine you are discovering.

See "Discovery Guidelines for Source Workloads" on page 299.

5 (Optional) If you want to store these credentials for use during future jobs, enable the **Save (Encrypted Locally)** option.

6 Click **Discover** and wait for the process to complete.

7 (Optional) If you want to monitor the progress of the job, switch to the Jobs view.

Workload Discovery in the Migrate Web Interface

To migrate a workload through the Web Interface, you must first add (or *discover*) the workload.

PlateSpin Migrate Web Interface supports discovery of a physical, virtual, or cloud-based machine:

- An individual Windows workload
- An individual Linux workload

Before you discover a workload, ensure that PlateSpin Server can communicate with your source workloads. See "Requirements for Discovery" on page 53.

To discover a workload:

1 In the PlateSpin Migrate Web Interface, click **Workloads** > **Add Workload**.

Alternatively, you can click the **Add Workload** option on the Dashboard page.

2 Specify the host name or the IP address of the workload you want to add.

3 Select the type of workload.

4 Specify the credentials to connect to the workload.

5 Click **Add Workload** to discover the workload and list it on the Workloads page.

Registering Workloads and Discovering Details with Migrate Agent

Migrate Agent is a command line utility that enables you to register source workloads with PlateSpin Migrate servers and send details about the workloads to the server via HTTPS (TCP/443). Registration allows you to add workloads that cannot be discovered, such as:

- When you deploy Migrate server in the cloud without a site-to-site VPN
- When corporate network or policy restrictions prohibit opening ports for automated discovery

Migrate Agent enables you to migrate a Windows workload without opening any inbound ports, such as SMB or NetBIOS. Only HTTPS (TCP/443) and a replication port (TCP/3725 is the default) are needed outbound for the source Windows workloads. For source Linux workloads, you also need to open the SSH port (TCP/22). See "Requirements for Workload Registration" on page 56.

When you use the Migrate Agent on the source workload, the source workload contacts the target workload for data transfers. The direction is controlled at the server level. Ensure that the value of the `SourceListensForConnection` parameter setting is set to False, which is the default value See "Configuring the Contact Direction for the Replication Port" on page 114.

You must install Migrate Agent on each source workload. When you use the `register` option, Migrate Agent performs discovery locally on the workload and sends its details to the Migrate Server through HTTPS (TCP/443). After you register the workload, use the Migrate Web Interface to configure the workload migration to the target cloud where the Migrate Server instance is deployed.

Registered workloads differ from discovered workloads in the following ways:

- Registered source workloads do not store the source credentials on the Migrate Server.
- You must use Migrate Agent to install, upgrade, and remove the Windows PlateSpin drivers from registered source workloads.
- After you delete the contract for a registered source workload, use the `cleanup` command to remove the OFX controller from the workload. See clu | cleanup for Windows workload and cleanup for Linux workload

For information about the Migrate Agent commands, see "Migrate Agent Utility" on page 359.

- "Windows Workload Registration and Discovery with Migrate Agent" on page 304
- "Linux Workload Registration and Discovery with Migrate Agent" on page 305
- "Linux Workload Registration and Discovery with Migrate Agent for Workloads in AWS" on page 306

Windows Workload Registration and Discovery with Migrate Agent

Before you begin, ensure that your source Windows workload and network settings meet the "Requirements for Migrate Agent Utility". For Windows workloads, Migrate Agent Utility requires `Administrator` privileges to execute commands.

1. Log in as Administrator to the source Windows workload.
2. Ensure that TCP port 443 is open on the workload.
3. Download Migrate Agent Utility for Windows. Save the `MigrateAgent.cli.exe` file to a convenient location on the workload.

   See "Migrate Agent Utility for Windows" on page 361.
4. In an Administrator Prompt, navigate to the location where you saved the file, then view the command Help by entering:

   ```
   MigrateAgent.cli.exe help
   ```
5. Register the workload with the appropriate Migrate Server cloud instance. Enter

   ```
   MigrateAgent.cli.exe /register /psserver=ps_dns_or_ipaddr <username> /
   password=<password>
   ```

   Provide the credentials for an administrator-level user of the PlateSpin Migrate Server who has the permissions needed to add a workload. You can use the `/password=` option with the password, use the `-pwdfile=` option with a path to a file that contains the password, or do not specify the password in the command sequence. If you exclude the password from the command line, the script will prompt for it. The password is obscured as you type it and it does not appear in the process list.

   For example:

   ```
   Migrate.Agent.cli.exe /register /psserver=10.10.10.101 /username=jsmith
   /password=jspwd
   ```

   NOTE: If you modify the public IP address of the Migrate Server, you must run the following command on each of the source Windows workloads that are configured for the server to modify the IP address.

   ```
   MigrateAgent.cli.exe /config /setting=psserver:<new-ps-dns-or-ipaddr>
   ```

   For example:

   ```
   MigrateAgent.cli.exe /config /setting=psserver:10.10.20.202
   ```

   ---
6. Verify that the PlateSpin Controller is running. Enter

   ```
   MigrateAgent.cli.exe /status
   ```

   If the controller is running, the status reports results similar to the following:

   ```
   The PlateSpin Controller daemon is running and registered to server
   10.165.x.x
   The PlateSpin blockwatch driver is not installed.
   ```

Linux Workload Registration and Discovery with Migrate Agent

Before you begin, ensure that your source workload and network settings meet the "Requirements for Migrate Agent Utility". Key Linux considerations are:

- The Migrate Agent Utility for Linux requires the source machine to have GNU C Library (glibc) 2.11.3 or higher installed.
- Migrate Agent requires root-level access to execute commands. A non-`root` user must be an authorized `sudo` user.

  For a non-`root` user, type `sudo` in the Migrate Agent commands to execute them with `root` privileges. For example:

  ```
  sudo ./MigrateAgent -h
  ```

  If you are prompted for a password, provide the password of the non-`root` system user name you logged in as.

NOTE: In AWS, you must run `sudo -i` and execute commands in a root shell. Use the registration procedure in "Linux Workload Registration and Discovery with Migrate Agent for Workloads in AWS" on page 306.

To register source Linux workloads:

1. Log in to the source Linux workload as the `root` user or as a non-`root` user with root level access.
2. Ensure that TCP port 443 is open on the workload.
3. Download the Migrate Agent Utility for Linux. Extract the downloaded file to the `/MigrateAgent` directory,

   See "Migrate Agent Utility for Linux" on page 364.
4. In a terminal, navigate to the `/MigrateAgent` directory, then view the command Help by entering:

   ```
   ./MigrateAgent -h
   ```
5. Register the workload with the appropriate Migrate Server cloud instance. Enter

   ```
   ./MigrateAgent register [-h] <ps_dns_or_ipaddr> <ps_username> [[-p
   <user_password>] | [-pf <passwordfile_path>]]
   ```

   Specify the IP address or DNS name of the PlateSpin Migrate Server instance in the cloud. Provide the credentials for an administrator-level user of the PlateSpin Migrate Server who has the permissions needed to add a workload. You can use the `-p` option with the password, use the `-pf` option with a path to a file that contains the password, or do not specify the password in the command sequence. If you exclude the password from the command line, the script will prompt for it. The password is obscured as you type it and it does not appear in the process list.

   For example:

   ```
   ./MigrateAgent register 10.10.10.101 jsmith -p jspwd
   ```

NOTE: If you modify the public IP address of the Migrate Server, you must run the following command on each of the source Linux workloads that are configured for the server to modify the IP address.

```
./MigrateAgent configure <ps_dns_or_ipaddr> <new-ps-dns-or-ipaddr>
```

For example:

```
./MigrateAgent configure 10.10.10.101 10.10.20.202
```

6 Verify that PlateSpin Controller is running. Enter

```
./MigrateAgent status
```

If the controller is running, the status reports results similar to the following:

```
The PlateSpin Controller daemon is running and registered to server
10.165.x.x
The PlateSpin blockwatch driver is not installed.
```

Linux Workload Registration and Discovery with Migrate Agent for Workloads in AWS

PlateSpin Migrate Web Interface supports migration of Amazon Web Services EC2 VM instances to Microsoft Azure, without requiring a VPN. The source workload operating system and architecture of the workload must be supported for VMs in Azure. For migration requirements for this scenario, see Chapter 12, “Prerequisites for Cloud-to-Cloud Migrations,” on page 217.

Before you begin, ensure that your source Linux workload and network settings meet the “Requirements for Migrate Agent Utility”. Key Linux considerations for Linux workloads in AWS are:

- The Migrate Agent Utility for Linux requires the source machine to have GNU C Library (glibc) 2.11.3 or higher installed.
- Migrate Agent requires root-level access to execute commands. A non-`root` user must be an authorized `sudo` user.

  NOTE: For source Linux workloads in Amazon Web Services, AMI templates automatically create a default non-`root` system user account that is enabled for `sudo`. The user name for this account varies by AMI provider. For Amazon Linux images, the non-`root` user name is `ec2-user` for most Linux distributions. It is `centos` for CentOS AMIs. For more information, refer to your AMI provider documentation.

  In AWS, a non-`root` user must run the `sudo -i` command to access the `root` shell and then run the Migrate Agent commands. Typing `sudo` in each Migrate Agent Utility command might result in a failure on some source workloads.
- AWS login for SSH requires the local path of the private key file that you created for the AWS EC2 Key Pair.

To register a source workload in AWS with your Migrate server:

1 Log in to the source Linux workload in AWS by using a system user name with `root`-level access and the local path of the private key file.

2 Ensure that TCP port 443 is open on the workload.

3 Download the Migrate Agent Utility for Linux. Extract the downloaded file to the `/MigrateAgent` directory,

See "Migrate Agent Utility for Linux" on page 364.

4 In a terminal, navigate to the `/MigrateAgent` directory.

5 (Non-`root` user) At the server console, run `sudo -i`. Enter

```
sudo -i
```

This command puts you in a root shell where commands are executed as the `root` user. The terminal prompt now shows `root` instead of your non-`root` user name, such as `ec2-user`.

If you are prompted by Linux for a password, provide the password of the user name you logged in as.

6 View the Migrate Agent command Help by entering:

```
./MigrateAgent -h
```

7 Register the workload with the appropriate Migrate Server cloud instance. Enter

```
./MigrateAgent register [-h] <ps_dns_or_ipaddr> <ps_username> [[-p
<user_password>] | [-pf <passwordfile_path>]]
```

Specify the IP address or DNS name of the PlateSpin Migrate Server instance in the cloud. Provide the credentials for an administrator-level user of the PlateSpin Migrate Server who has the permissions needed to add a workload. You can use the `-p` option with the password, use the `-pf` option with a path to a file that contains the password, or do not specify the password in the command sequence. If you exclude the password from the command line, the script will prompt for it. The password is obscured as you type it and it does not appear in the process list.

For example:

```
./MigrateAgent register 10.10.10.101 jsmith -p jspwd
```

NOTE: If you modify the public IP address of the Migrate Server, you must run the following command on each of the source Linux workloads that are configured for the server to modify the IP address.

```
./MigrateAgent configure <ps_dns_or_ipaddr> <new-ps-dns-or-ipaddr>
```

For example:

```
./MigrateAgent configure 10.10.10.101 10.10.20.202
```

---

8 Verify that PlateSpin Controller is running on the source workload. Enter

```
./MigrateAgent status
```

If the controller is running, the status reports results similar to the following:

```
The PlateSpin Controller daemon is running and registered to server
10.165.x.x
The PlateSpin blockwatch driver is not installed.
```

9 (Non-`root` user) Exit the `sudo -i` root shell. Press Ctrl+D, or enter

```
exit
```

The terminal prompt now shows your non-`root` user name, such as `ec2-user`.

Refreshing Source Workload Details

If you make changes on the source workload before the migration begins, you might need to re-discovery the workload details. In the Migrate Client, you can refresh discovery details. In the Migrate Web Interface, you must remove and re-add the workload.

- "Refresh Workload Details in Migrate Client" on page 308
- "Removing and Re-Adding Workloads in the Web Interface" on page 308

Refresh Workload Details in Migrate Client

PlateSpin Migrate Client allows you to refresh workload details. You should routinely refresh your source workloads and targets before setting up a migration job.

To refresh a source workload details:

1 In the Servers view, right-click the required item, then select **Refresh Details**.

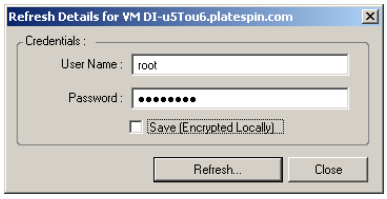

2 Specify the credentials appropriate for the system being refreshed, then click **Refresh**.

PlateSpin Migrate starts a discovery job, which you can monitor in the Jobs view.

Removing and Re-Adding Workloads in the Web Interface

PlateSpin Migrate Web Interface does not support refreshing details for the discovered workloads. To update details about a discovered workload, you must remove the workload, and then add and discover its details again. For example, if you modify the host name of the discovered workload or add or remove volumes, you must remove and re-add the workload to capture the new information. For information about removing the workload, see "Undiscovering or Removing Source Workloads" on page 310.

On removing a workload:

- Configuration details of the workload are lost if the workload is in a configured state.
- Migration license allocated to the workload post the successful completion of its first full replication will be lost if the workload is removed with the **Preserve Source** option unselected.

IMPORTANT: To retain the migration license of a workload that you want to remove from the Web Interface, ensure that the **Preserve Source** option is selected when you remove the workload.

Using Tags to Track Logical Associations of Workloads

In the PlateSpin Migrate Web Interface, the Workloads page might display a long list of workloads. Searching through these workloads to manage operations for similar workloads can be time-consuming. To overcome this issue, you can create tags for various workload categories, departments, or other logical associations appropriate to your environment. A tag can be associated with any workload that you manage in the Web Interface.

For information about creating, modifying, or deleting workload tags, see "Managing Workload Tags" on page 133.

After you create tags, they are available at the bottom of the Edit Target Details page where you can assign a tag to the appropriate workloads. The Workloads page includes a **Tag** column where the single tag you associate with a workload is displayed. You can sort on this column to group similar workloads together. This enables you to easily locate and run operations on the tagged workloads at the same time.

NOTE: When you export a workload with a tag setting to a new server, the tag settings persist.

To associate a tag with a workload during Configure Migration:

1. In the Migrate Web Interface, click **Workloads.**
2. In the workload list, select the workload you want to tag and click **Configure Migration**.
3. Configure the workload.
4. In the Tag section at the bottom of the Edit Target Details page, select the tag name you want to associate with the workload
5. Click **Save**.

To add or modify a tag associated with configured workload:

1. In the Migrate Web Interface, click **Workloads.**
2. In the workload list, click the workload you want to tag to open the Target Details page.
3. Click **Edit**.
4. In the Tag section at the bottom of the Edit Target Details page, select the tag name you want to associate with the workload.
5. Click **Save**.

To disassociate a tag from a workload:

1. In the Migrate Web Interface, click **Workloads.**
2. In the workload list, select the workload for which you want to remove the tag and click **Configure Migration**.
3. In the Tag section of the configuration page, select the empty string and click **Save**.

Undiscovering or Removing Source Workloads

After you complete all migration jobs for a source workload and the cutover completes successfully, you can remove (undiscover) the source workload.

IMPORTANT

- Before you delete an object that is in use for configured jobs, you must ensure that all the affected jobs are completed.
- If block-level transfer is enabled, remove the block-based transfer driver from the source workload:
  - **Windows:** Select to uninstall the block-based transfer driver.

    A reboot of the source workload is required after the driver is removed.
  - **Linux:** Manually uninstall the blkwatch driver from the source. See Block-level data transfer software in Cleaning Up Linux Workloads.
- For potential cleanup of files copied during discovery to the target platform, ensure that the target platform is reachable before you remove (undiscover) the target platform.

To undiscover a workload through the Migrate Client:

1. On the Workloads page, right-click the workload object and select **Undiscover Server**.
2. (Block-level transfer) Remove the block-based driver from the source workload.
3. (Windows) Reboot the source workload.

To remove a workload through the Migrate Web Interface:

1. On the Workloads page, select the workload, then click **Remove Workload**.
2. (Block-level transfer) Remove the block-based driver from the source workload.
3. (Windows) Reboot the source workload.

19 Preparing Device Drivers

PlateSpin Analyzer ships with a library of device drivers, and during migration jobs it installs the appropriate drivers for the target. If you require specific drivers for your target infrastructure, you might need to add (upload) drivers to the PlateSpin Migrate driver database.

To determine if the required drivers are available for conversion of Windows workloads to physical machines, you can use the PlateSpin Analyzer function in PlateSpin Migrate Client. PlateSpin Analyzer can help identify missing or incompatible drivers. See "Analyzing Suitability of Discovered Windows Workloads For Conversion to Physical Machines" on page 320.

Packaging Device Drivers for Windows Systems

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

1. Prepare all interdependent driver files (`*.sys`, `*.inf`, `*.dll`, etc.) for your target infrastructure and device. If you have obtained manufacturer-specific drivers as a `.zip` archive or an executable, extract them first.
2. Save the driver files in separate folders, with a discrete folder per device.

The drivers are now ready for upload. See "Uploading Drivers to the PlateSpin Migrate Device Driver Database" on page 312.

NOTE: For problem-free operation of your migration job and the target workload, upload *only digitally signed* drivers for all 64-bit Windows systems.

Packaging Device Drivers for Linux Systems

To package your Linux device drivers for uploading to the PlateSpin Migrate driver database, you can use a custom utility included in your Linux ISO boot image.

1. Find a Linux workstation that has the same kernel version as the source machine. Source machine itself is one of the best choices. On the Linux workstation, create a directory for your device driver files. All the drivers in the directory must be for the same kernel and architecture.
2. Download the boot image and mount it.

For example, assuming that the ISO has been copied under the `/root` directory, issue these commands:

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

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

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

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

4 Enter the working directory and execute the following command:

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

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

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

For example, `bnx2x-1.48.107-RHEL7-U6-3.10.0-957.el7.x86_64-x86_64.pkg`

The package is now ready for upload. See "Uploading Drivers to the PlateSpin Migrate Device Driver Database" on page 312.

Uploading Drivers to the PlateSpin Migrate Device Driver Database

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

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

Device Driver Upload Procedure (Windows)

1 Obtain and prepare the required device drivers.

See Packaging Device Drivers for Windows Systems.

2 Click **Tools** > **Manage Device Drivers** and select the **Windows Drivers** tab:

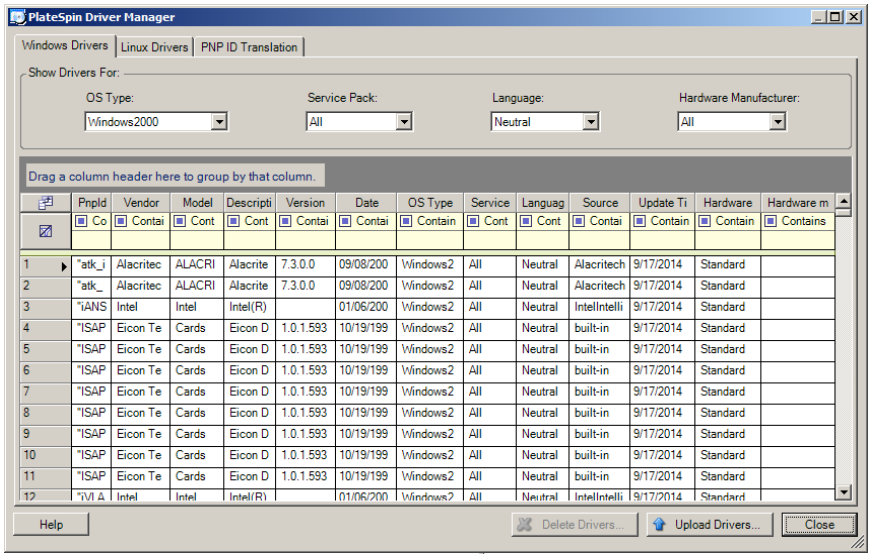

3 Click **Upload Drivers**.

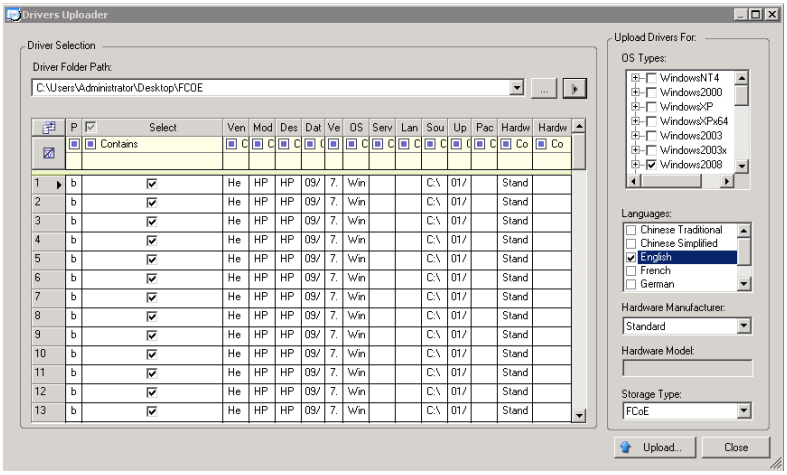

4 Select the Hardware Manufacturer.

For most X2P migrations, select **Standard** as the **Hardware Manufacturer** option, unless your drivers are designed specifically for any of the target environments listed.

5 Select the Storage Type.

IMPORTANT: If you select the **Storage Type** as **FCoE**, then you must ensure that all the drivers applicable for the FCoE storage device are in the same folder.

---

6 Browse to the folder that contains the required driver files, and select applicable OS type, language, and hardware manufacturer options.

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

The system uploads the selected drivers to the driver database.

Device Driver Upload Procedure (Linux)

1 Obtain and prepare the required device drivers.

See Packaging Device Drivers for Linux Systems.

2 Click **Tools** > **Manage Device Drivers** and select the **Linux Drivers** tab:

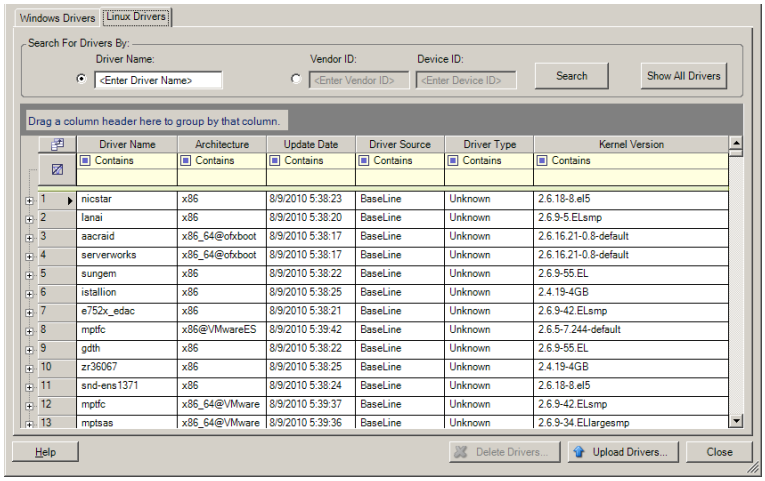

3 Click **Upload Drivers**, browse to the folder that contains the required driver package (`*.pkg`), and click **Upload All Drivers**.

The system uploads the selected drivers to the driver database.

Using the Plug and Play (PnP) ID Translator Feature

"Plug and Play" (PnP) refers to Windows operating system functionality that supports connectivity, configuration, and management with native plug and play devices. In Windows, the feature facilitates discovery of PnP compliant hardware devices attached to a PnP compliant bus. PnP compliant devices are assigned a set of Device Identification Strings by their manufacturer. These strings are programmed into the device when it is built. These strings are fundamental to how PnP works: they are part of the Windows' information source used to match the device with a suitable driver.

When the PlateSpin Server discovers workloads and their available hardware, the discovery includes these PnP IDs and the storage of that data as part of the workload's details. PlateSpin uses the IDs to determine which, if any, drivers need to be injected during a conversion operation. The PlateSpin

Server maintains a database of PnP IDs for the associated drivers of each of the supported operating systems. Because Windows and Linux use different formats for PnP IDs, a Windows workload discovered by the Migrate Linux RAM disk contains Linux-style PnP IDs.

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

Follow these steps to use the PnP ID Translation feature:

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

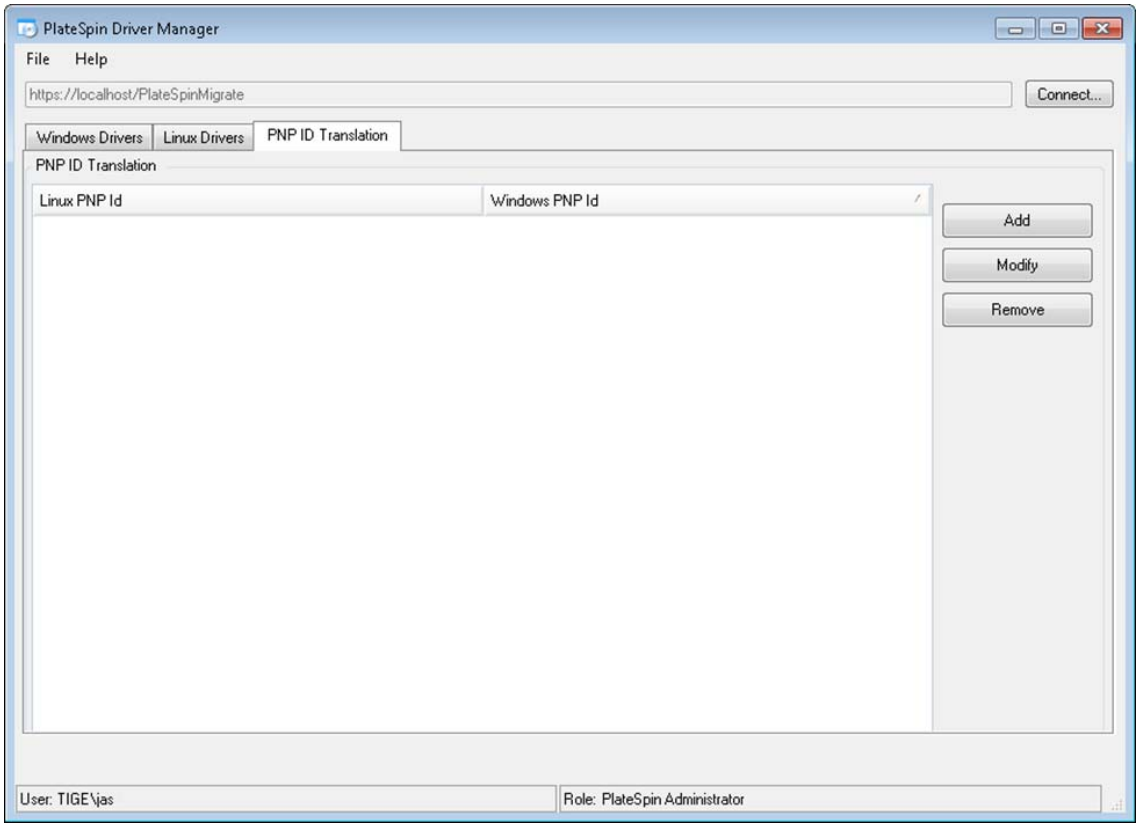

3. On the list page, click **Add** to display the Create PNP ID Mapping dialog box.

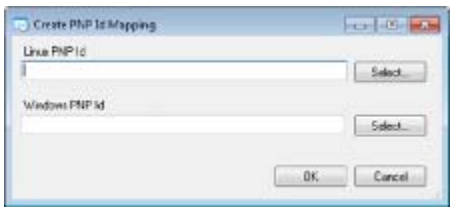

4 In the **Linux PNP ID** field, add a Linux PnP ID.

4a (Conditional) If you know it, type the Linux PnP ID you want to use.

or

4b (Conditional) Select an ID from a previously discovered workload:

4b1 Adjacent to the **Linux PnP ID** field, click **Select** to open the Select Linux PnP ID dialog box.

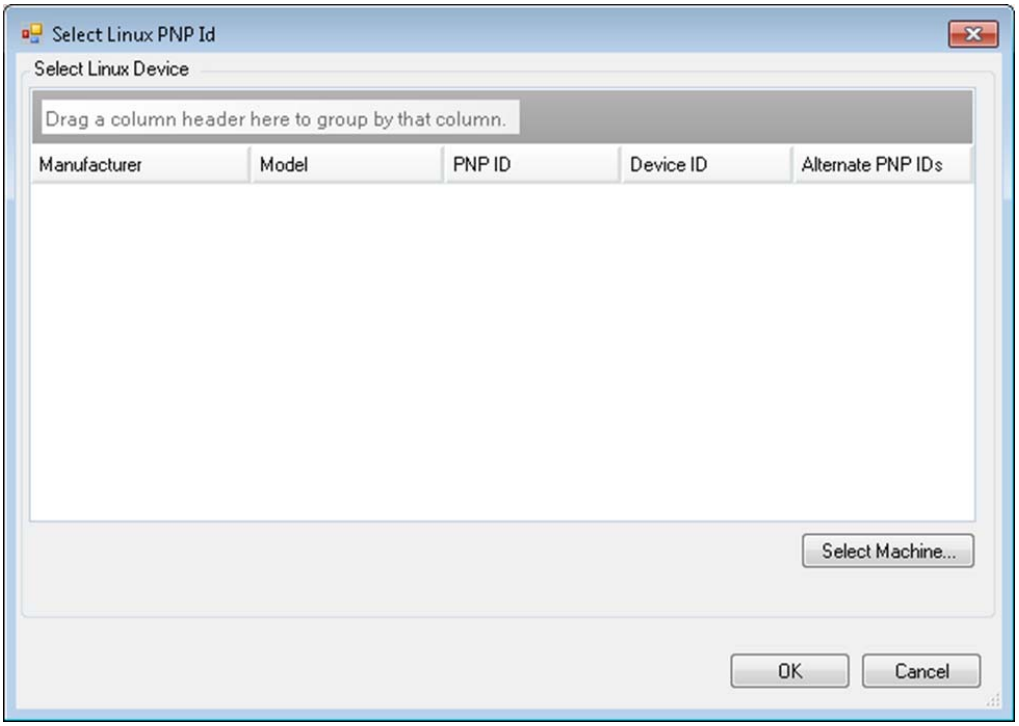

4b2 On the dialog box, click **Select Machine** to display a list of the machines previously discovered by the PlateSpin Linux RAM disk.

4b3 Highlight one of the devices in the list, then click **Select** to populate the list in the Select Linux PnP ID dialog box.

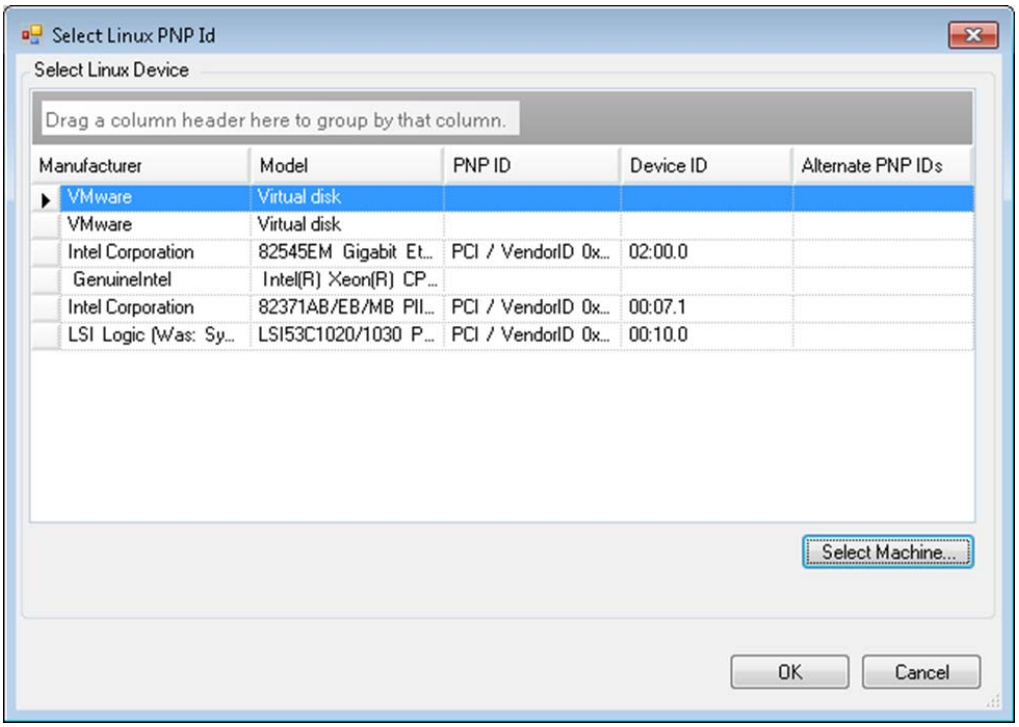

4b4 Select a device on the list, then click **OK** to apply the standard transformation to the PnP ID and display it in the Create PnP ID Mapping dialog box.

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

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

or

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

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

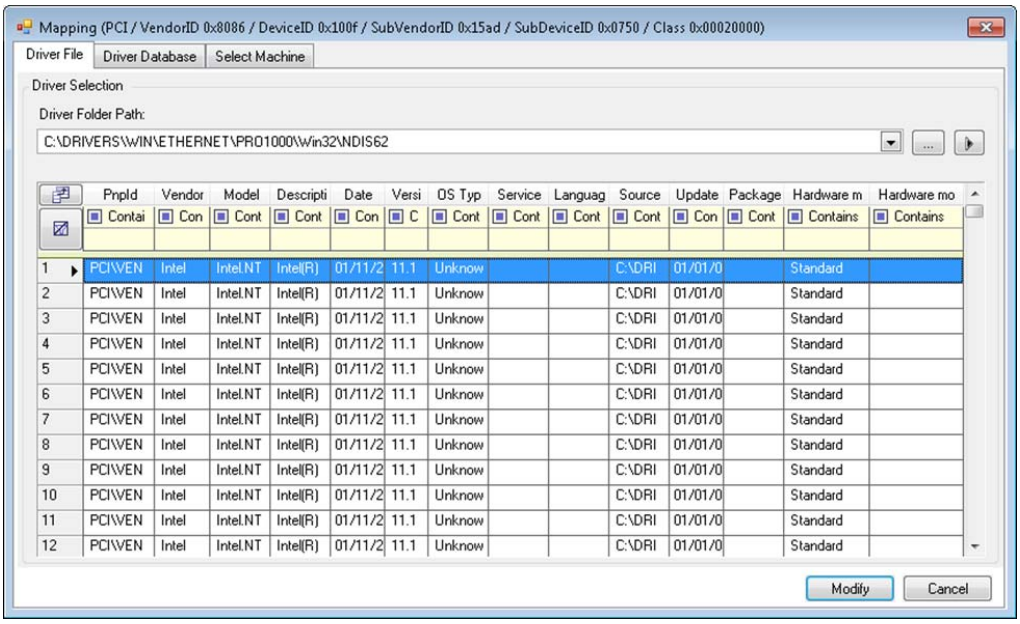

- Under the **Driver Database** tab, browse to and select the existing driver database, select the correct PnP ID, then select **Modify**.

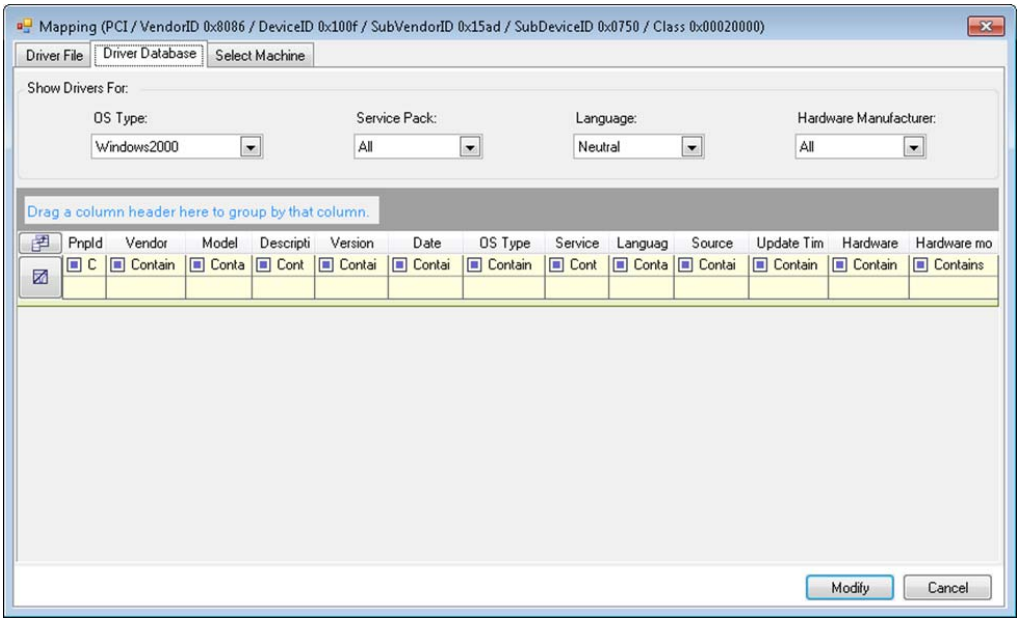

- Under the **Select Machine** tab, click **Select Machine**, then, from the list of Windows machines discovered using live discovery, select a machine, click **OK** to display its devices, select the desired PnP ID, then click Modify.

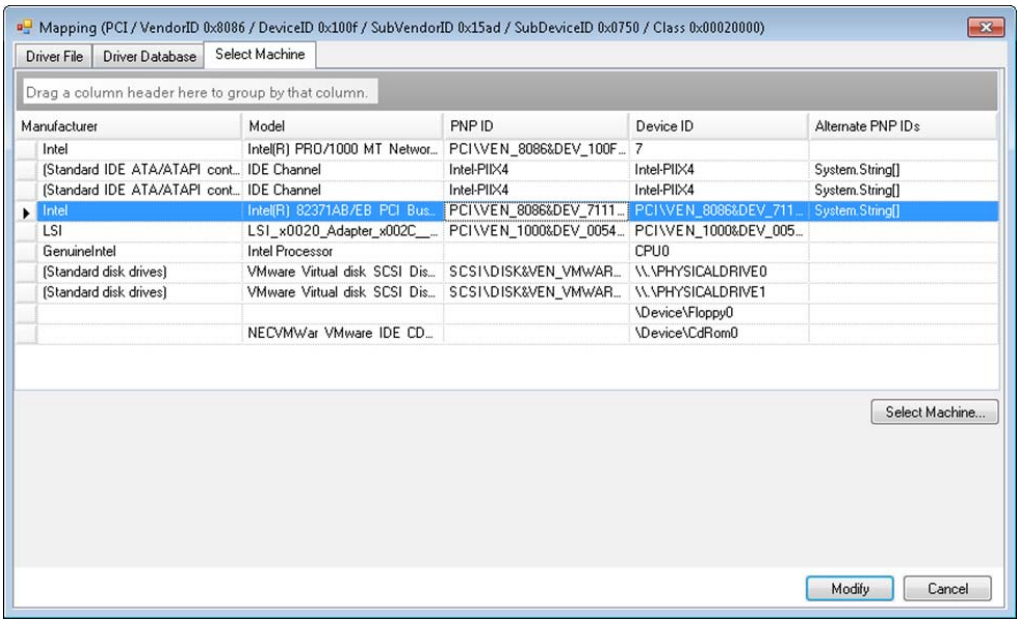

IMPORTANT: Selecting a Windows PnP ID that does not have an associated driver package installed might result in a failure at conversion time.

6 In the Create PnP Id Mapping dialog box, confirm that the correct Linux PnP ID and the correct Windows PnP are selected, then click **OK** to display the PNP ID Translation page of the PlateSpin Driver Manager.

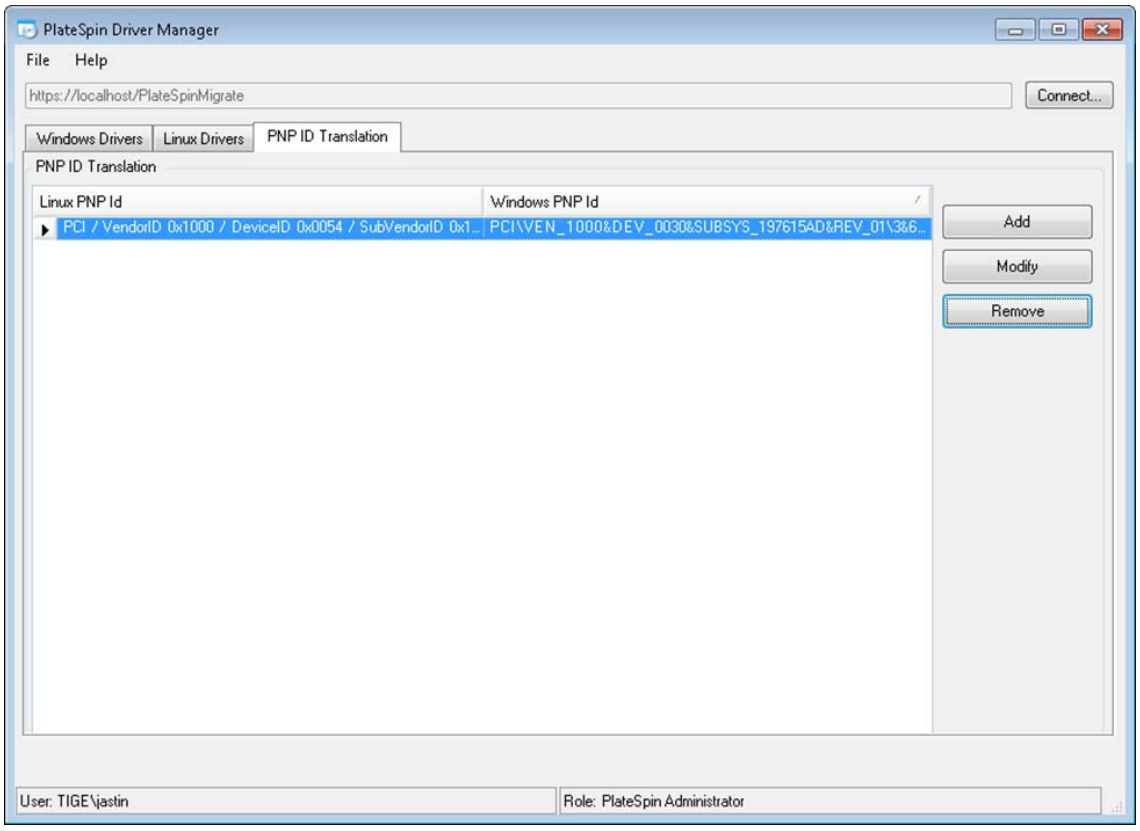

7 (Optional) To modify or remove the mapping in the PNP ID Translation list, select the mapping pattern, then click **Remove** or **Modify**, depending on the operation you want to perform.

Remove simply deletes the mapping (after displaying a confirmation dialog box).

To modify,

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

7b Repeat Step 5 to modify the Windows PnP ID.

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

Analyzing Suitability of Discovered Windows Workloads For Conversion to Physical Machines

Before you begin any large-scale migration projects, you should identify potential migration problems and correct them beforehand. The PlateSpin Migrate Client provides the PlateSpin Analyzer utility to validate the following:

- Compatibility of target hardware for migration to physical targets
- Availability of drivers in the driver database for the physical server hardware

NOTE: PlateSpin Analyzer currently supports only Windows workloads.

About PlateSpin Analyzer Tests and Results

For target hardware support, PlateSpin Analyzer checks whether hardware drivers are in the driver repository for the following conversion types:

- Physical to physical (P2P)
- Virtual to physical (V2P)

Table 19-1 describes the purpose of each test.

Table 19-1 *PlateSpin Analyzer Tests*

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

Table 19-2 describes the status messages in the test results.

Table 19-2 *Status Messages in PlateSpin Analyzer Test Results*

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

For more information about using PlateSpin Analyzer and understanding the results, see KB Article 7920478 (https://support.microfocus.com/kb/doc.php?id=7920478).

PlateSpin Analyzer in the Migrate Client

To open PlateSpin Analyzer:

1. On the **Tools** menu, click **Analyze Servers**.

   The PlateSpin Analyzer window opens.
2. Select the required PlateSpin Migrate Network and the required machines to analyze.
3. (Optional) To reduce the analysis time, limit the scope of machines to a specific language.
4. (Optional) To analyze machines in the inventory of a different PlateSpin Server, click **Connect**, then specify the required PlateSpin Server URL and valid credentials.
5. Click **Analyze**.

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

Analyzed servers are listed in the left pane. Select a server to view test results in the right pane.

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

The **Test Results** tab provides the test results about a selected machine.

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

20 Preparing Linux Workloads for Migration

Perform the tasks in this section to prepare your Linux workloads for migration using PlateSpin Migrate

Verifying Block-Based Drivers for Linux

Verify that a blkwatch module is available for the workload's Linux distribution. For a list of preconfigured drivers, see Appendix D, "Linux Distributions Supported by Migrate," on page 343.

If you plan to protect a supported Linux workload that has a non-standard, customized, or newer kernel, rebuild the PlateSpin `blkwatch` module, which is required for block-level data replication.

See Knowledgebase Article 7005873 (https://support.microfocus.com/kb/doc.php?id=7005873).

Adding Drivers to the PlateSpin ISO Image

The PlateSpin ISO image contains a large library of device drivers sufficient to boot most common targets. However, occasionally you might want to use your own, such as lesser-known, vendor-specific, or custom-developed drivers for Linux workloads.

You can modify the PlateSpin ISO image to add your vendor-specific or custom-developed drivers. See "Injecting Additional Device Drivers into the PlateSpin ISO Image" on page 374.

Configuring LVM Snapshots for Linux Volume Replication

We recommend that you prepare snapshots for block-level data transfer. Ensure that each volume group has sufficient free space for snapshots (at least 10% of the sum of all partitions). If snapshots are not available, PlateSpin Migrate locks and releases each block in turn on the source workload for data transfer.

The `blkwatch` driver leverages LVM snapshots if they are available. Copying blocks from the snapshot helps avoid potential open file conflicts.

For LVM storage, see Knowledgebase Article 7005872 (https://support.microfocus.com/kb/doc.php?id=7005872).

Using Custom Freeze and Thaw Scripts for Linux Block-Level Migrations

For Linux workload migrations, PlateSpin Migrate supports the use of *freeze* and *thaw* shell scripts to provide an additional means of control over your Linux block-level migration process.

Migrate executes these scripts during Linux workload migrations, at the beginning and end of block-level data transfer sessions. Specifically, they interject in the migration process in the following fashion:

1. First pass of all volumes without snapshots:
   - Regular (non-LVM) volumes
   - LVM without enough space to take a snapshot
2. *Freeze script*
3. Take snapshots
4. Second pass of all non-snapshot volumes
5. *Thaw script*
6. Transfer volume snapshots

You can use this capability to complement the automated daemon control feature provided through the user interface. See "Services or Daemons to Stop before Replication or Cutover" on page 400.

For example, you might want to use this feature to cause an application to flush its data to disk so that the workload remains in a more consistent state during migration.

To use the feature, do the following before setting up your migration job:

1 Create the following files:
   - `platespin.freeze.sh` is a shell script to contain the `freeze` logic.
   - `platespin.thaw.sh` is a shell script to contain the `thaw` logic.
   - `platespin.conf` is a text file that defines any required arguments, along with a timeout value.

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

     `[ServiceControl]`

     (optional) `FreezeArguments=`*<arguments>*

     (optional) `ThawArguments=`*<arguments>*

     (optional) `TimeOut=`*<timeout>*

     Replace *<arguments>* with the required command arguments, separated by a space, and *<timeout>* with a timeout value in seconds. If unspecified, the default timeout is used (60 seconds).

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

`/etc/platespin/`

Preparing Paravirtualized Linux Source Workload

Before you migrate a paravirtualized Linux source workload that is hosted on a Citrix XenServer or KVM virtual host to a target platform as a fully virtualized workload, do the following:

- Ensure that both the paravirtualized kernel and the standard kernel are installed on the paravirtualized Linux source workload.
- Manually compile the block-based drivers for Xen kernel. See "Linux Distributions Supported by Migrate" on page 343.
- Use block-based data transfer for the migration.
- For semi-automated (X2P) migrations, ensure that your target is a fully virtualized (not paravirtualized) VM.

See "Paravirtualized Source Workloads" on page 40.

21 Preparing for Migration of Windows Clusters

You can migrate Microsoft Windows Cluster business services to a target VMware vCenter virtualization platform or to a physical machine. For information about supported Microsoft Windows Clusters, see "Clusters" in "Supported Source Workloads For Migration to Non-Cloud Platforms" on page 23.

You can use PlateSpin Migrate Client or PlateSpin Migrate Web Interface to migrate Windows Clusters to VMware vCenter virtualization platforms. You can also use PlateSpin Migrate Client to migrate Windows Clusters to physical machines. The prerequisites for migration are the same.

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

Planning Your Cluster Workload Migration

When active node discovery is enabled (the default) for the PlateSpin environment, migration of a Windows cluster is achieved through incremental replications of changes on the active node streamed to a virtual one node cluster. If you disable active node discovery, each node of a Windows cluster can be discovered and migrated as a standalone node.

Before you configure Windows clusters for migration, ensure that your environment meets the prerequisites and that you understand the conditions for migrating cluster workloads.

- "Cluster Node Similarity" on page 332
- "Migration Setup for the Active Node" on page 332

Requirements for Cluster Migration

The scope of support for cluster migration is subject to the conditions described in Table 21-1. Consider these requirements when you configure migration for clusters in your PlateSpin environment.

Table 21-1 *Cluster Migration Requirements*

Requirement	Description
Discover the active node as a Windows Cluster	The PlateSpin global configuration setting `DiscoverActiveNodeAsWindowsCluster` determines whether Windows clusters are migrated as clusters or as separate standalone machines: • **True (Default):** The active node is discovered as a Windows cluster. • **False:** Individual nodes can be discovered as standalone machines. See "Configuring Windows Active Node Discovery" on page 332.
Resource name search values	The PlateSpin global configuration setting `MicrosoftClusterIPAddressNames` determines the cluster resource names that can be discovered in your PlateSpin environment. You must configure search values that help to differentiate the name of the shared Cluster IP Address resource from the name of other IP address resources on the cluster. See "Adding Resource Name Search Values" on page 333.
Windows Cluster Mode	The PlateSpin global configuration setting `WindowsClusterMode` determines the method of block-based data transfer for incremental replications: • **Default:** Driverless synchronization. • **SingleNodeBBT:** Driver-based block-based transfer. See the following: • "Block-Based Transfer for Clusters" on page 329 • "Configuring the Block-Based Transfer Method for Clusters" on page 333
Active node host name or IP address	You must specify the host name or IP address of the cluster's active node when you perform an **Add Workload** operation. Because of security changes made by Microsoft, Windows clusters can no longer be discovered by using the virtual cluster name (that is, the shared cluster IP address).
Resolvable host name	The PlateSpin Server must be able to resolve the host name of each of the nodes in the cluster by their IP address. **NOTE:** DNS forward lookup and reverse lookup are required to resolve the host name by its IP address.

Requirement	Description
Quorum resource	A cluster's quorum resource must be co-located on the node with the cluster's resource group (service) being migrated.
Similarity of cluster nodes	In the default Windows Cluster Mode, driverless sync can continue from any node that becomes active if the nodes are similar. If they do not match, replications can occur only on the originally discovered active node. See "Cluster Node Similarity" on page 332.
PowerShell 2.0	Windows PowerShell 2.0 must be installed on each node of the cluster.

Block-Based Transfer for Clusters

Block-based transfer for clusters works differently than for standalone servers. The initial replication either makes a complete copy (full) or use a driverless block-based synchronization method performed on the active node of the cluster. Subsequent incremental replications can use a driverless synchronization method or driver-based synchronization method for block-based data transfer.

NOTE: PlateSpin Migrate does not support file-based transfer for Windows Clusters.

The PlateSpin global configuration setting `WindowsClusterMode` determines the method of block-based data transfer for incremental replications:

- **Default:** Driverless synchronization using an MD5-based replication on the currently active node.
- **SingleNodeBBT:** Driver-based synchronization using a BBT driver installed on the originally discovered active node.

Both methods support block-level replication of local storage and shared storage on Fibre Channel SANs and iSCSI SANs.

Table 21-2 describes and compares the two methods.

Table 21-2 *Comparison of Block-Based Data Transfer Methods for Incremental Replication*

Consideration	Default BBT	Single-Node BBT
Data transfer method	Uses driverless synchronization with an MD5-based replication on the currently active node.	Uses a BBT driver installed on the originally discovered active node.
Performance	Potentially slow incremental replications.	Significantly improves performance for incremental replications.
Supported Windows Clusters	Works with any supported Windows Server clusters.	Works with Windows Server 2008 R2 and later clusters.

Consideration	Default BBT	Single-Node BBT
Drivers	◆ Driverless; no BBT driver to install. ◆ No reboot is required on the source cluster nodes.	◆ Use the Migrate Agent utility to install a BBT driver on the originally discovered active node of the cluster. ◆ Reboot the node to apply the driver. This initiates a failover to another node in the cluster. After the reboot, make the originally discovered node the active node again. ◆ The same node must remain active for replications to occur and to use single-node block-based transfer. ◆ After you install the BBT driver, either a full replication or a driverless incremental replication must occur before the driver-based incremental replications can begin.
First incremental replication	Uses driverless sync on the active node.	Uses driver-based block-based transfer on the originally discovered active node if a full replication was completed after the BBT driver was installed. Otherwise, it uses driverless sync on the originally discovered active node.
Subsequent incremental replication	Uses driverless sync on the active node.	Uses driver-based block-based transfer on the originally discovered active node. If a cluster switches nodes, the driverless sync method is used for the first incremental replication after the originally active node becomes active again. See “Impact of Cluster Node Failover on Replication” on page 330.

Impact of Cluster Node Failover on Replication

Table 21-3 describes the impact of cluster node failover on replication and the required actions for the Migrate administrator.

Table 21-3 *Impact of Cluster Node Failover on Replication*

Cluster Node Failover or Failback	Default BBT	Single-Node BBT
Cluster node failover occurs during the first full replication	Replication fails. The first full replication must complete successfully without a cluster node failover. 1. Remove the cluster from Migrate. 2. (Optional) Make the originally discovered active node the active node again. 3. Re-add the cluster using the active node. 4. Re-run the first full replication.	
Cluster node failover occurs during a subsequent full replication or a subsequent incremental replication	The replication command aborts and a message displays indicating that the replication needs to be re-run. If the new active node's profile is similar to the failed active node, the migration contract remains valid. 1. Re-run the replication on the now-active node. If the new active node's profile is not similar to the failed active node, the migration contract is valid only on the originally active node. 1. Make the originally discovered active node the active node again. 2. Re-run the replication on the active node.	The replication command aborts and a message displays indicating that the replication needs to be re-run. The migration contract is valid only on the originally discovered active node. 1. Make the originally discovered active node the active node again. 2. Re-run the replication on the active node. This first incremental replication after a cluster failover/failback event automatically uses driverless sync. Subsequent incremental replications will use the block-based driver as specified by single-node BBT.
Cluster node failover occurs between replications	If the new active node's profile is similar to the failed active node, the migration contract continues as scheduled for the next incremental replication. Otherwise, the next incremental replication command fails. If a scheduled incremental replication fails: 1. Make the originally discovered active node the active node again. 2. Run an incremental replication.	Incremental replication fails if the active node switches between replications. 1. Ensure that the originally discovered active node is again the active node. 2. Run an incremental replication. This first incremental replication after a cluster failover/failback event automatically uses driverless sync. Subsequent incremental replications will use the block-based driver as specified by single-node BBT.

Cluster Node Similarity

In the default Windows Cluster Mode, the cluster nodes must have similar profiles to prevent interruptions in the replication process. The profiles of cluster nodes are considered similar if all of the following conditions are met:

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

  NOTE: Use the customized *Volume Manager* utility to change the local volume serial numbers to match each node of the cluster. See "Synchronizing Serial Numbers on Cluster Node Local Storage" on page 357.

  If the local volumes on each node of the cluster have different serial numbers, you cannot run a replication after a cluster node failover occurs. For example, during a cluster node failover, the active node Node 1 fails, and the cluster software makes Node 2 the active node. If the local drives on the two nodes have different serial numbers, the next replication command for the workload fails.
- The nodes must have the same number of volumes.
- Each volume must be exactly the same size on each node.
- The nodes must have an identical number of network connections.

Migration Setup for the Active Node

To configure migration for a Windows cluster, follow the normal workload migration workflow.

Ensure that you provide the host name or IP address of the cluster's active node.

Configuring Windows Active Node Discovery

You can discover Windows Server clusters as clusters or as individual standalone machines, depending on the PlateSpin global configuration setting `DiscoverActiveNodeAsWindowsCluster`.

To discover Windows clusters as clusters, set the `DiscoverActiveNodeAsWindowsCluster` parameter to `True`. This is the default setting. Cluster discovery, inventory, and workload migration use the host name or IP address of a cluster's active node, instead of using its cluster name and an administration share. You do not configure separate workloads for the cluster's non-active nodes. For other cluster workload migration requirements, see "Requirements for Cluster Migration" on page 328.

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

To enable or disable cluster discovery:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`

   Replace *Your_PlateSpin_Server* with the DNS host name or IP address of your PlateSpin Migrate Server.
2. Search for `DiscoverActiveNodeAsWindowsCluster`, then click **Edit**.
3. In the **Value** field, select **True** to enable cluster discovery, or select **False** to disable cluster discovery.
4. Click **Save**.

Configuring the Block-Based Transfer Method for Clusters

Incremental replications for Windows clusters can use a driverless method (Default) or driver-based method (SingleNodeBBT) for block-based data transfer, depending on the PlateSpin global configuration setting `WindowsClusterMode`. For more information, see "Block-Based Transfer for Clusters" on page 329.

To configure WindowsClusterMode:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`

   Replace *Your_PlateSpin_Server* with the DNS host name or IP address of your PlateSpin Migrate Server.
2. Search for `WindowsClusterMode`, then click **Edit**.
3. In the **Value** field, select **Default** to use driverless synchronization for incremental replication, or select **SingleNodeBBT** to use block-based drivers for incremental replication.
4. Click **Save**.

Adding Resource Name Search Values

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

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

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

The default name for the shared Cluster IP Address resource is `Cluster IP Address` in English, or the equivalent if the cluster node is configured in another language. The default search values in the `MicrosoftClusterIPAddressNames` list include the resource name `Cluster IP Address` in English and examples of Microsoft translations of that name in other languages. Ensure that you use the Microsoft translation of the phrase for the language configured on your cluster.

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

To add search values in the `MicrosoftClusterIPAddressNames` list:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`

   Replace *Your_PlateSpin_Server* with the DNS host name or IP address of your PlateSpin Migrate Server.
2. Search for `MicrosoftClusterIPAddressNames`, then click **Edit**.
3. In the **Value** field, add one or more search values to the list.
4. Click **Save**.

Quorum Arbitration Timeout

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

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

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`

   Replace *Your_PlateSpin_Server* with the DNS host name or IP address of your PlateSpin Migrate Server.
2. Search for `FailoverQuorumArbitrationTimeout`, then click **Edit**.
3. In the **Value** field, specify the maximum number of seconds to allow for quorum arbitration.
4. Click **Save**.

Setting Local Volume Serial Numbers

In the default Windows Cluster Mode, replication of the currently active in the Windows cluster fails if the serial numbers for the nodes' local volumes (System volume and System Reserved volume) is not the same on each cluster node. See "Cluster Node Similarity" on page 332.

You can use the *Volume Manager* utility to change the local volume serial numbers to match in each node of the cluster. See "Synchronizing Serial Numbers on Cluster Node Local Storage" on page 357.

Guidelines for PlateSpin Cutover

- When the PlateSpin cutover operation is complete and the virtual one-node cluster comes online, you see a multi-node cluster with one active node (all other nodes are unavailable).
- To perform a PlateSpin cutover (or to test the PlateSpin cutover on) a Windows cluster, the cluster must be able to connect to a domain controller. To leverage the test failover functionality, you need to migrate the domain controller along with the cluster. During the test, bring up the domain controller, followed by the Windows cluster workload (on an isolated network).

Guidelines for PlateSpin Cluster Migration

- A PlateSpin cluster migration operation requires a full replication for Windows Cluster workloads.
- (P2P migrations) After PlateSpin cluster migration is complete, you must reattach the shared storage and rebuild the cluster environment before you can rejoin additional nodes to the newly restored cluster.

  For information about rebuilding the cluster environment after a PlateSpin migration, see *Rebuilding a Windows Server 2012 R2 Cluster* (KB 7016770).

Migrating Windows Clusters with the Web Interface

After you prepare your environment for migrating the Windows cluster, you can use the PlateSpin Migrate Web Interface to migrate the essential services of a cluster that results in a functional single-node cluster in a virtual machine in VMware. The workflow of migrating the Windows cluster is similar to that of migrating a standalone server, except that you migrate the active node.

1. In the Web Interface, add the active node by specifying the IP address of the active node.
2. Configure migration for the active node to VMware.
3. Run the migration.

   See "Guidelines for PlateSpin Cluster Migration" on page 335.
4. Perform cutover.

   See "Guidelines for PlateSpin Cutover" on page 335.

Migrating Windows Clusters with the Migrate Client

In the PlateSpin Migrate Client, you can use a **Move** job to migrate the essential services of a cluster that results in a functional single-node cluster in a virtual machine in VMware or a physical machine.

The workflow of migrating a Windows cluster is similar to that of migrating a standalone server:

1. Discover the active node by specifying the IP address of the active node.
2. In the Servers view, use drag-and-drop to start a migration job, then configure the job's parameters.
3. (Conditional: successful migration) If the migration job completes successfully, perform a Server Sync operation on the active node.

   NOTE: If the active node in the cluster fails over before you can perform a Server Sync operation, perform a full migration using the *new* active node, and then perform a Server Sync on this new node.

4. (Conditional: failover prior to migration) If a cluster failover occurs prior to the completion of file transfer, the migration job aborts. If this happens, refresh the source and retry the migration job.

NOTE: If you select **Shut down** for the source's post-migration end state, a shutdown of all source nodes of the cluster results.

C Troubleshooting Discovery

Table C-1 provides information to help you troubleshoot common problems that might occur during workload discovery or target discovery.

Common Discovery Issues and Solutions

Table C-1 *Common Issues and Solutions Related to Discovery Operations*

Problems or Messages	Solutions
Unable to Configure X2P Migration for a Workload Discovered Using Migrate Agent	Use the PlateSpin Migrate Client or PlateSpin Migrate Web Interface to discover the source workload and then migrate the workload. You can alternatively use a different replication method for the migration.
Discovering a Source Workload By Host Name Fails When a Discovered Under Control Target Has the Same Host Name As the Source	Use the IP Address of the source workload instead of the host name to discover it.
The workload cannot be migrated because it has 0 active partitions. Ensure that the workload has exactly 1 active partition and try again	This error occurs if there is no active partition on the source workload. Use the `diskpart` `SELECT` and `ONLINE` commands to make a partition active: 1. Open a command prompt as an Administrator and run `diskpart`. 2. Enter `list volume` and make a note of the volume number that you want to make active. 3. Enter `select volume <volume_number>` 4. Enter `online volume` and then `exit`.
`Application has generated an error` occurs during registration of physical server	This error occurs if the physical server is unable to contact the PlateSpin Server. A common cause is incorrect information entered during the registration process. To restart the registration process: 1. Enter `RegisterMachine.bat`. 2. Ping to confirm basic connectivity with the PlateSpin Server.

Problems or Messages	Solutions
My physical server has completed the registration process, but is not seen in PlateSpin Migrate Client.	The full registration process can take some time to complete. After the second command prompt window has closed on the physical server, wait a few minutes before clicking the **Refresh** button in PlateSpin Migrate Client.
Problems discovering source and target servers	KB Article 7920291 (https://support.microfocus.com/kb/doc.php?id=7920291) contains troubleshooting checklists for discovering the following: ◆ Linux servers and VMware ESX Servers ◆ Windows-based source and target servers The article also has instructions for troubleshooting WMI connections and checking if DCOM is enabled.
`Package <…> Not Found` occurs during discovery of existing Windows servers	Check for required IIS configuration and network settings. See "Installing Prerequisite Software" in the *PlateSpin Migrate 2020.2 Installation and Upgrade Guide*.
`Unable to connect neither to the SSH server running on <IP_address> nor to VMware Virtual Infrastructure web-services at <ip_address>/sdk`	This message has a number of possible causes: ◆ The workload is unreachable. ◆ The workload does not have SSH running. ◆ The firewall is on and the required ports have not been opened. ◆ The workload's specific operating system is not supported. For network and access requirements for a workload, see "Access and Communication Requirements across Your Migration Network" on page 53
`Access denied`	This authentication problem indicates either an invalid user name or password. For information on proper workload access credentials, see Table 18-2, "Guidelines for Machine Type and Credentials for Source Workloads," on page 299. Access can be denied for SSH connections if the key algorithm or ciphers settings in the `/etc/ssh/sshd_config` file on the source Linux workload are missing or are incompatible with the settings used by Migrate server. See "Test Credentials or Discovery Fails with Access Denied Error" on page 339.

Related KB Articles are listed in Table C-2.

Table C-2 *KB Articles for Discovery Issues*

ID	Description
7920339 (https://support.microfocus.com//kb/doc.php?id=7920339)	ERRMSG: Discovery fails with `The request failed with HTTP status 407` message
7920862 (https://support.microfocus.com/kb/doc.php?id=7920862)	ERRMSG: Recoverable Error: `ControllerConnectionBroken during discovery`

ID	Description
7920291 (https://support.microfocus.com/kb/doc.php?id=7920291)	ERRMSG: `Server details discovery problems`
7021574 (https://support.microfocus.com/kb/doc.php?id=7021574)	ERRMSG: X2P Target Discovery Failed: `Linux job did not complete successfully`
For more discovery-related TIDs in the Knowledgebase	Search on “discovery” for the PlateSpin Migrate product (https://support.microfocus.com/kb/?&q=discovery&product=PlateSpin_Migrate).

Test Credentials or Discovery Fails with Access Denied Error

Issue: Test Credentials, Add workload, or Discover workload actions for a source Linux workload fails with the following error:

`Access denied. The root credentials provided cannot be used to connect to the server <source-Linux-workload-IP-address>. Please ensure that the password is correct, and that root has not been blocked from using SSH.`

Workaround: Access can be denied for SSH connections if the key algorithm or ciphers settings in the `/etc/ssh/sshd_config` file on the source Linux workload are missing or are incompatible with the settings used by Migrate server.

1. Verify the following are working correctly:
   - You correctly specified the source Linux workload’s IP address, user name, and password.
   - On the source Linux workload, the SSH service is enabled and running; and the firewall (if any) allows inbound SSH traffic on TCP port 22.
   - You can log in successfully to this Linux Workload as `root` user from a remote machine using an SSH client such as Putty.
2. On the source Linux workload, log in as the `root` user, then view the log file (`/var/log/messages`) or check the status of the SSH daemon (`systemctl status sshd`) to search for error messages for the Migrate server IP address.
3. Add or discover the source Linux workload again.
   - **3a** Verify that Test Credential is successful.
   - **3b** Verify that the workload is added successfully.

See also the following related KB Articles:

- *Discovering Linux workload states access denied* (KB 7018214)
- *Linux discovery failure with access denied error* (KB 7018128)

Modifying the OFX Controller Heartbeat Startup Delay (Windows Workloads)

To avoid discovery failures caused by timing issues, a default heartbeat startup delay of 15 seconds (15000 ms) is set on the OFX Controller. The setting is configurable by adding the `HeartbeatStartupDelayInMS` registry key on the source workload. This registry key is not configured by default.

To enable a heartbeat delay of shorter or longer duration:

1. On the source workload, open the Windows 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"
   ```

4. Restart the source workload.

Web Interface Does Not Display the Edited Host Name of a Discovered Workload

Issue: If you edit the host name of a discovered workload, the new host name displays in the Migrate Client, but not in the Web Interface. (`Bug 1042869`)

Workaround: A discovery refresh option is not available in the Migrate Web Interface. For workload migrations that you manage in the Web Interface, if you modify information about the workload, such as changing its host name or adding or removing volumes, you must undiscover the workload and then rediscover it.

PlateSpin Migrate Client

Use information in this section to help troubleshoot common problems that might occur when working with workloads in PlateSpin Migrate Client.

- “Deleting a Workload from the PlateSpin Migrate Client Does Not Remove the Workload from the Web Interface” on page 341
- “Discovering a Workload in the PlateSpin Migrate Client Fails When a Migration License Is Not Available” on page 341

Deleting a Workload from the PlateSpin Migrate Client Does Not Remove the Workload from the Web Interface

Issue: When you use the Migrate Client to delete a workload, the workload is deleted from the Migrate Client but is not automatically removed from the Web Interface. Consequently, rediscovering the same workload fails.

Workaround: Manually remove the workload from the Web Interface before you rediscover the workload.

Discovering a Workload in the PlateSpin Migrate Client Fails When a Migration License Is Not Available

Issue: If you choose to discover a workload in the PlateSpin Migrate Client when there are no remaining migrations available for the license, the workload discovery job in the Migrate Client fails to initiate and no message is displayed. However, the workload is added in the Web Interface, which prevents rediscovery.

Workaround: None. You can discover the workload in the Migrate Client only after the migration license is available. Ensure that you manually remove the workload from the Web Interface before you attempt to rediscover the workload.

D Linux Distributions Supported by Migrate

PlateSpin Migrate uses an exactly matching blkwatch driver for block-based replication of Linux workloads. The PlateSpin Migrate software by default includes precompiled block-based Linux kernel drivers, called block watch (blkwatch) drivers, for many non-debug Linux distributions (32-bit and 64-bit).

If the PlateSpin Migrate software does not include the precompiled blkwatch drivers for your workload, you must create custom blkwatch drivers before you migrate the workload. The custom blkwatch driver must exactly match the specific kernel running on your Linux workload. You can create custom drivers either manually or by using the PlateSpin Migrate Agent Utility.

You can also copy custom drivers from one Migrate Server to other Migrate Servers. This helps you to build the custom drivers library on the Migrate Servers without having to recreate them.

Review the following sections for information about the PlateSpin Migrate precompiled blkwatch drivers, how to create the custom drivers, and how to copy custom drivers from one Migrate Server to another Migrate Server:

Create Custom Blkwatch Drivers Manually

You must first determine whether your Linux workload kernel matches a supported precompiled blkwatch driver. If no match is found, manually create the custom blkwatch driver. For more information, review the following sections:

Analyze Your Linux Workload

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

Determine the Release String

You can determine the release string of the kernel of your Linux workload by running the following command at the workload's Linux terminal:

```
uname -r
```

For example, if you run `uname -r`, you might see the following output:

```
3.0.76-0.11-default
```

If you search the list of distributions, you see there are two entries that match this string:

- `SLES11SP3-GA-3.0.76-0.11-default-x86`
- `SLES11SP3-GA-3.0.76-0.11-default-x86_64`

The search results indicate that the product has drivers for both 32-bit (x86) and 64-bit (x86_64) architectures.

Determine the Architecture

You can determine the architecture of your Linux workload by running the following command at the workload's Linux terminal:

```
uname -m
```

For example, if you run `uname -m`, you might see the following output:

```
x86_64
```

With this information, you can determine that the workload has 64-bit architecture.

Search the Linux Driver Distibution List

You can search the List of Distributions to determine if the release string and architecture of your Linux workload kernel matches a supported distribution in the list. If you find your release string and architecture, PlateSpin Migrate has a precompiled version of the `blkwatch` driver.

If your search is unsuccessful, you must create a custom `blkwatch` driver before you migrate the workload.

Create Custom Blkwatch Drivers

To manually create a custom `blkwatch` driver, follow the steps in the KB Article 7005873 (https://support.microfocus.com/kb/doc.php?id=7005873).

IMPORTANT: The folder that contains the custom driver must follow the naming convention: `CustomKernel-<release string>-<architecture>`

Self-compiled drivers are supported only for the Linux major and minor kernel versions that appear in the List of Distributions, or a patched version thereof. If the major and minor kernel version in the release string of your Linux workload kernel matches a major and minor kernel version in the list, your self-compiled driver will be supported.

Create Custom Drivers Using the Migrate Agent Utility

If your Linux source workload is registered to the PlateSpin Migrate Server using the PlateSpin Migrate Agent Utility, run the `driver-install` command to automatically install the custom drivers on your source workload.

When you run the `driver-install` command, the utility checks if precompiled drivers for the kernel running on your Linux workload are available. If no precompiled drivers match the source workload kernel, the utility automatically compiles the required drivers and installs them on the source workload.

The folder that contains the custom driver is named based on the following naming convention: `CustomKernel-<Kernel_Release_String>-<Kernel_Architecture>`. The utility then uploads the custom drivers to the PlateSpin Server library. See driver-install command in Table F-2 on page 366.

Copying Custom Drivers to Migrate Servers

You can copy custom drivers from one Migrate Server to other Migrate Servers. This helps you to build the custom drivers library on many Migrate Servers without having to recreate them.

Consider that you want to copy the custom drivers from the Migrate Server M1 to the Migrate Server M2.

To copy the drivers from M1:

1. Log in to the PlateSpin Migrate Server M1 host machine as the Administrator user.
2. Navigate to the PlateSpin Migrate installation directory such as `<install-folder>\PlateSpin Migrate Server`.
3. Open the `\Packages\9D470BFF-E30F-4015-83D8-4D308A8C8B1B\` folder and copy the latest `.package` (`x.package`) file out of this folder to a temporary directory dir1.

   For example, if the folder contains `1.package`, `2.package`, and `3.package`, copy the `3.package` file to the temporary directory `dir1`.
4. Rename `x.package` with the `.zip` file extension (`x.zip`) and extract the contents of the file to this temporary directory `dir1`.
5. Copy all the custom drivers (named `CustomKernel*`) from `dir1` to another temporary location `dir2`.

To copy the drivers to M2:

1. Log in to the PlateSpin Migrate Server M2 host machine as the Administrator user.
2. Navigate to the PlateSpin Migrate installation directory such as `<install-folder>\PlateSpin Migrate Server`.
3. Open the `\Packages\9D470BFF-E30F-4015-83D8-4D308A8C8B1B\` folder and copy the latest `.package` (`y.package`) file out of this folder to a temporary directory `dir3`.

   For example, if the folder contains `1.package`, `2.package`, and `3.package`, copy the `3.package` file to a temporary directory `dir3`.
4. Rename `y.package` with the `.zip` file extension (`y.zip`) and extract the contents of the file to this temporary directory `dir4`.

5 Copy all the custom drivers from the temporary location `dir2` on the PlateSpin Migrate Server M1 to the temporary location `dir4` on the PlateSpin Migrate Server M2.

6 Zip all the files in the folder `dir4` on the PlateSpin Migrate Server M2 to recreate the `y.zip` file.

7 Rename `y.zip` to `y.package`.

This `y.package` on the PlateSpin Migrate Server M2 now contains the custom drivers copied from M1 in addition to all the drivers that existed on M2 before copying the drivers from M1.

Precompiled blkwatch Drivers for Linux Distributions

List Item Syntax

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

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

So, for a SLES 11 SP1 distribution with a kernel release string of `2.6.32.12-0.6-pae` for 32-bit (x86) architecture, the item is listed in a format like this:

`SLES11-SP1-2.6.32.12-0.6-pae-x86`

List of Distributions

Oracle Linux 5 U11 GA

`OEL511-GA-2.6.18-398.el5-x86`

`OEL511-GA-2.6.18-398.el5-x86_64`

`OEL511-GA_UEK-2.6.39-400.215.10.el5uek-x86`

`OEL511-GA_UEK-2.6.39-400.215.10.el5uek-x86_64`

Oracle Linux 6 U7

NOTE: Blkwatch drivers for kernel version 2.6.32-573 do not support incremental replication for workloads with LVM volumes. Update the kernel, then see RHEL 6 U7 for drivers for kernel 2.6.32-642.

`OEL6-U7-2.6.32-573.el6.i686-x86`

`OEL6-U7-2.6.32-573.el6.x86_64-x86_64`

`OEL6-U7_UEK-2.6.39-400.250.7.el6uek.i686-x86`

`OEL6-U7_UEK-3.8.13-68.3.4.el6uek.x86_64-x86_64`

Oracle Linux 6 U8

NOTE: Blkwatch drivers for kernel version 2.6.32-642 on RHEL 6 U8 do not support incremental replication for workloads with LVM volumes. Update the kernel, then see RHEL 6.8 for drivers for kernel 2.6.32-696.20.1.

`OEL6-U8-2.6.32-642.el6.i686-x86`

`OEL6-U8-2.6.32-642.el6.x86_64-x86_64`

OEL6-U8_UEK-2.6.39-400.278.2.el6uek.i686-x86
OEL6-U8_UEK-4.1.12-37.4.1.el6uek.x86_64-x86_64

Oracle Linux 6 U9

OEL6-U9-2.6.32-696.el6.i686-x86
OEL6-U9-2.6.32-696.el6.x86_64-x86_64
OEL6-U9_UEK-2.6.39-400.294.3.el6uek.i686-x86
OEL6-U9_UEK-4.1.12-61.1.28.el6uek.x86_64-x86_64

Oracle Linux 6 U10

OEL6-U10-2.6.32-754.el6.i686-x86
OEL6-U10-2.6.32-754.el6.x86_64-x86_64
OEL6-U10_UEK-2.6.39-400.299.3.el6uek.i686-x86
OEL6-U10_UEK-4.1.12-124.16.4.el6uek.x86_64-x86_64

Oracle Linux 7 GA

OEL7-GA-3.10.0-123.el7.x86_64-x86_64
OEL7-GA_UEK-3.8.13-35.3.1.el7uek.x86_64-x86_64

Oracle Linux 7 U1

OEL7-U1-3.10.0-229.el7.x86_64-x86_64
OEL7-U1_UEK-3.8.13-55.1.6.el7uek.x86_64-x86_64

Oracle Linux 7 U2

OEL7-U2-3.10.0-327.el7.x86_64-x86_64
OEL7-U2_UEK-3.8.13-98.7.1.el7uek.x86_64-x86_64

Oracle Linux 7 U3

OEL7-U3-3.10.0-514.el7.x86_64-x86_64
OEL7-U3_UEK-4.1.12-61.1.18.el7uek.x86_64-x86_64

Oracle Linux 7 U4

OEL7-U4-3.10.0-693.el7.x86_64-x86_64
OEL7-U4_UEK-4.1.12-94.3.9.el7uek.x86_64-x86_64

Oracle Linux 7 U5

OEL7-U5-3.10.0-862.el7.x86_64-x86_64
OEL7-U5_UEK-4.1.12-112.16.4.el7uek.x86_64-x86_64

Oracle Linux 7 U6

OEL7-U6-3.10.0-957.el7.x86_64-x86_64
OEL7-U6_UEK-4.14.35-1818.3.3.el7uek.x86_64-x86_64

Oracle Linux 7 U7

OEL7U7-U7-3.10.0-1062.el7.x86_64-x86_64
OEL7U7-U7_UEK-4.14.35-1902.3.2.el7uek.x86_64-x86_64

Red Hat Enterprise Linux 5 GA

RHEL5-GA-2.6.18-8.el5-x86
RHEL5-GA-2.6.18-8.el5-x86_64
RHEL5-GA-2.6.18-8.el5PAE-x86

Red Hat Enterprise Linux 5 U1

RHEL5-U1-2.6.18-53.el5-x86
RHEL5-U1-2.6.18-53.el5-x86_64
RHEL5-U1-2.6.18-53.el5PAE-x86

Red Hat Enterprise Linux 5 U2

RHEL5-U2-2.6.18-92.el5-x86

RHEL5-U2-2.6.18-92.el5-x86_64

RHEL5-U2-2.6.18-92.el5PAE-x86

Red Hat Enterprise Linux 5 U3

RHEL5-U3-2.6.18-128.el5-x86

RHEL5-U3-2.6.18-128.el5-x86_64

RHEL5-U3-2.6.18-128.el5PAE-x86

Red Hat Enterprise Linux 5 U4

RHEL5-U4-2.6.18-164.el5-x86

RHEL5-U4-2.6.18-164.el5-x86_64

RHEL5-U4-2.6.18-164.el5PAE-x86

Red Hat Enterprise Linux 5 U5

RHEL5-U5-2.6.18-194.el5-x86

RHEL5-U5-2.6.18-194.el5-x86_64

RHEL5-U5-2.6.18-194.el5PAE-x86

Red Hat Enterprise Linux 5 U6

RHEL5-U6-2.6.18-238.el5-x86

RHEL5-U6-2.6.18-238.el5-x86_64

RHEL5-U6-2.6.18-238.el5PAE-x86

Red Hat Enterprise Linux 5 U7

RHEL5-U7-2.6.18-274.el5-x86

RHEL5-U7-2.6.18-274.el5-x86_64

RHEL5-U7-2.6.18-274.el5PAE-x86

Red Hat Enterprise Linux 5 U8

RHEL5-U8-2.6.18-308.el5-x86

RHEL5-U8-2.6.18-308.el5-x86_64

RHEL5-U8-2.6.18-308.el5PAE-x86

Red Hat Enterprise Linux 5 U9

RHEL5-U9-2.6.18-348.el5-x86

RHEL5-U9-2.6.18-348.el5-x86_64

RHEL5-U9-2.6.18-348.el5PAE-x86

Red Hat Enterprise Linux 5 U10

RHEL5-U10-2.6.18-371.el5-x86

RHEL5-U10-2.6.18-371.el5-x86_64

RHEL5-U10-2.6.18-371.el5PAE-x86

Red Hat Enterprise Linux 5 U11

RHEL5-U11-2.6.18-398.el5-x86

RHEL5-U11-2.6.18-398.el5-x86_64

RHEL5-U11-2.6.18-398.el5PAE-x86

Red Hat Enterprise Linux 6 GA

RHEL6-GA-2.6.32-71.el6.i686-x86

RHEL6-GA-2.6.32-71.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U1

RHEL6-U1-2.6.32-131.0.15.el6.i686-x86

RHEL6-U1-2.6.32-131.0.15.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U2

RHEL6-U2-2.6.32-220.el6.i686-x86

RHEL6-U2-2.6.32-220.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U3

RHEL6-U3-2.6.32-279.el6.i686-x86

RHEL6-U3-2.6.32-279.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U4

RHEL6-U4-2.6.32-358.el6.i686-x86

RHEL6-U4-2.6.32-358.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U5

RHEL6-U5-2.6.32-431.el6.i686-x86

RHEL6-U5-2.6.32-431.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U6

RHEL6-U6-2.6.32-504.el6-x86

RHEL6-U6-2.6.32-504.el6-x86_64

Red Hat Enterprise Linux 6 U7

NOTE: Blkwatch drivers for kernel version 2.6.32-573 do not support incremental replication for workloads with LVM volumes. Update the kernel, then use drivers for kernel 2.6.32-642.

RHEL6-U7-2.6.32-573.el6.i686-x86

RHEL6-U7-2.6.32-573.el6.x86_64-x86_64

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

RHEL6-RHSA201700361-2.6.32-642.13.1.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U8

NOTE: Blkwatch drivers for kernel version 2.6.32-642 on RHEL 6 U8 do not support incremental replication for workloads with LVM volumes. Update the kernel, then use drivers for kernel 2.6.32-696.20.1.

RHEL6-U8-2.6.32-642.el6.i686-x86

RHEL6-U8-2.6.32-642.el6.x86_64-x86_64

RHEL6-RHSA20180169-2.6.32-696.20.1.el6.i686-x86

RHEL6-RHSA20180169-2.6.32-696.20.1.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U9

RHEL6-U9-2.6.32-696.el6.i686-x86

RHEL6-U9-2.6.32-696.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U10

RHEL6-U10-2.6.32-754.el6.i686-x86

RHEL6-U10-2.6.32-754.el6.x86_64-x86_64

Red Hat Enterprise Linux 7 GA

RHEL7-GA-3.10.0-123.el7.x86_64-x86_64

Red Hat Enterprise Linux 7 U1

RHEL7-U1-3.10.0-229.el7.x86_64-x86_64

Red Hat Enterprise Linux 7 U2

RHEL7-U2-3.10.0-327.el7.x86_64-x86_64

Red Hat Enterprise Linux 7 U3

RHEL7-U3-3.10.0-514.el7.x86_64-x86_64

Red Hat Enterprise Linux 7 U4

RHEL7-U4-3.10.0-693.el7.x86_64-x86_64

RHEL7-RHSA2018015101-3.10.0-693.17.1.el7.x86_64-x86_64

Red Hat Enterprise Linux 7 U5

RHEL7-U5-3.10.0-862.el7.x86_64-x86_64

Red Hat Enterprise Linux 7 U6

RHEL7-U6-3.10.0-957.el7.x86_64-x86_64

Red Hat Enterprise Linux 7 U7

RHEL7U7-U7-3.10.0-1062.el7.x86_64-x86_64

SUSE Linux Enterprise Server 11 GA

SLES11-GA-2.6.27.19-5-default-x86

SLES11-GA-2.6.27.19-5-default-x86_64

SLES11-GA-2.6.27.19-5-pae-x86

SUSE Linux Enterprise Server 11 SP 1

SLES11-SP1-2.6.32.12-0.6-default-x86

SLES11-SP1-2.6.32.12-0.6-default-x86_64

SLES11-SP1-2.6.32.12-0.6-pae-x86

SUSE Linux Enterprise Server 11 SP 1 U14

SLES11-SP1_U14-2.6.32.54-0.3-default-x86

SLES11-SP1_U14-2.6.32.54-0.3-default-x86_64

SLES11-SP1_U14-2.6.32.54-0.3-pae-x86

SUSE Linux Enterprise Server 11 SP 1 U15

SLES11-SP1_U15-2.6.32.59-0.3-default-x86

SLES11-SP1_U15-2.6.32.59-0.3-default-x86_64

SLES11-SP1_U15-2.6.32.59-0.3-pae-x86

SUSE Linux Enterprise Server 11 SP 1 U16

SLES11-SP1_U16-2.6.32.59-0.7-default-x86

SLES11-SP1_U16-2.6.32.59-0.7-default-x86_64

SLES11-SP1_U16-2.6.32.59-0.7-pae-x86

SUSE Linux Enterprise Server 11 SP 1 LTSS U1

SLES11-SP1_LTSS_U1-2.6.32.59-0.9-default-x86

SLES11-SP1_LTSS_U1-2.6.32.59-0.9-default-x86_64

SLES11-SP1_LTSS_U1-2.6.32.59-0.9-pae-x86

SUSE Linux Enterprise Server 11 SP 1 LTSS U2

SLES11-SP1_LTSS_U2-2.6.32.59-0.13-default-x86

SLES11-SP1_LTSS_U2-2.6.32.59-0.13-default-x86_64

SLES11-SP1_LTSS_U2-2.6.32.59-0.13-pae-x86

SUSE Linux Enterprise Server 11 SP 2 GA

SLES11SP2-GA-3.0.13-0.27-default-x86

SLES11SP2-GA-3.0.13-0.27-default-x86_64

SLES11SP2-GA-3.0.13-0.27-pae-x86

SLES11SP2-GA-3.0.13-0.27-xen-x86

SLES11SP2-GA-3.0.13-0.27-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U1

SLES11SP2-U1-3.0.26-0.7-default-x86
SLES11SP2-U1-3.0.26-0.7-default-x86_64
SLES11SP2-U1-3.0.26-0.7-pae-x86
SLES11SP2-U1-3.0.26-0.7-xen-x86
SLES11SP2-U1-3.0.26-0.7-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U2

SLES11SP2-U2-3.0.31-0.9-default-x86
SLES11SP2-U2-3.0.31-0.9-default-x86_64
SLES11SP2-U2-3.0.31-0.9-pae-x86
SLES11SP2-U2-3.0.31-0.9-xen-x86
SLES11SP2-U2-3.0.31-0.9-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U3

SLES11SP2-U3-3.0.34-0.7-default-x86
SLES11SP2-U3-3.0.34-0.7-default-x86_64
SLES11SP2-U3-3.0.34-0.7-pae-x86
SLES11SP2-U3-3.0.34-0.7-xen-x86
SLES11SP2-U3-3.0.34-0.7-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U4

SLES11SP2-U4-3.0.38-0.5-default-x86
SLES11SP2-U4-3.0.38-0.5-default-x86_64
SLES11SP2-U4-3.0.38-0.5-pae-x86
SLES11SP2-U4-3.0.38-0.5-xen-x86
SLES11SP2-U4-3.0.38-0.5-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U5

SLES11SP2-U5-3.0.42-0.7-default-x86
SLES11SP2-U5-3.0.42-0.7-default-x86_64
SLES11SP2-U5-3.0.42-0.7-pae-x86
SLES11SP2-U5-3.0.42-0.7-xen-x86
SLES11SP2-U5-3.0.42-0.7-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U6

SLES11SP2-U6-3.0.51-0.7.9-default-x86
SLES11SP2-U6-3.0.51-0.7.9-default-x86_64
SLES11SP2-U6-3.0.51-0.7.9-pae-x86
SLES11SP2-U6-3.0.51-0.7.9-xen-x86
SLES11SP2-U6-3.0.51-0.7.9-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U7

SLES11SP2-U7-3.0.58-0.6.2-default-x86
SLES11SP2-U7-3.0.58-0.6.2-default-x86_64
SLES11SP2-U7-3.0.58-0.6.2-pae-x86
SLES11SP2-U7-3.0.58-0.6.2-xen-x86
SLES11SP2-U7-3.0.58-0.6.2-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U8

SLES11SP2-U8-3.0.58-0.6.6-default-x86
SLES11SP2-U8-3.0.58-0.6.6-default-x86_64
SLES11SP2-U8-3.0.58-0.6.6-pae-x86
SLES11SP2-U8-3.0.58-0.6.6-xen-x86
SLES11SP2-U8-3.0.58-0.6.6-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U9

SLES11SP2-U9-3.0.74-0.6.6-default-x86
SLES11SP2-U9-3.0.74-0.6.6-default-x86_64
SLES11SP2-U9-3.0.74-0.6.6-pae-x86
SLES11SP2-U9-3.0.74-0.6.6-xen-x86
SLES11SP2-U9-3.0.74-0.6.6-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U10

SLES11SP2-U10-3.0.74-0.6.8-default-x86
SLES11SP2-U10-3.0.74-0.6.8-default-x86_64
SLES11SP2-U10-3.0.74-0.6.8-pae-x86
SLES11SP2-U10-3.0.74-0.6.8-xen-x86
SLES11SP2-U10-3.0.74-0.6.8-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U11

SLES11SP2-U11-3.0.74-0.6.10-default-x86
SLES11SP2-U11-3.0.74-0.6.10-default-x86_64
SLES11SP2-U11-3.0.74-0.6.10-pae-x86
SLES11SP2-U11-3.0.74-0.6.10-xen-x86
SLES11SP2-U11-3.0.74-0.6.10-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U12

SLES11SP2-U12-3.0.80-0.5-default-x86
SLES11SP2-U12-3.0.80-0.5-default-x86_64
SLES11SP2-U12-3.0.80-0.5-pae-x86
SLES11SP2-U12-3.0.80-0.5-xen-x86
SLES11SP2-U12-3.0.80-0.5-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U13

SLES11SP2-U13-3.0.80-0.7-default-x86
SLES11SP2-U13-3.0.80-0.7-default-x86_64
SLES11SP2-U13-3.0.80-0.7-pae-x86
SLES11SP2-U13-3.0.80-0.7-xen-x86
SLES11SP2-U13-3.0.80-0.7-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U14

SLES11SP2-U14-3.0.93-0.5-default-x86
SLES11SP2-U14-3.0.93-0.5-default-x86_64
SLES11SP2-U14-3.0.93-0.5-pae-x86
SLES11SP2-U14-3.0.93-0.5-xen-x86
SLES11SP2-U14-3.0.93-0.5-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U15

SLES11SP2-U15-3.0.101-0.5-default-x86
SLES11SP2-U15-3.0.101-0.5-default-x86_64

SLES11SP2-U15-3.0.101-0.5-pae-x86
SLES11SP2-U15-3.0.101-0.5-xen-x86
SLES11SP2-U15-3.0.101-0.5-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U16

SLES11SP2-U16-3.0.101-0.7.15-default-x86
SLES11SP2-U16-3.0.101-0.7.15-default-x86_64
SLES11SP2-U16-3.0.101-0.7.15-pae-x86
SLES11SP2-U16-3.0.101-0.7.15-xen-x86
SLES11SP2-U16-3.0.101-0.7.15-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U17

SLES11SP2-U17-3.0.101-0.7.17-default-x86
SLES11SP2-U17-3.0.101-0.7.17-default-x86_64
SLES11SP2-U17-3.0.101-0.7.17-pae-x86
SLES11SP2-U17-3.0.101-0.7.17-xen-x86
SLES11SP2-U17-3.0.101-0.7.17-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 LTSS U1

SLES11SP2-LTSS_U1-3.0.101-0.7.19-default-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-default-x86_64
SLES11SP2-LTSS_U1-3.0.101-0.7.19-pae-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-xen-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 LTSS U2

SLES11SP2-LTSS_U2-3.0.101-0.7.21-default-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-default-x86_64
SLES11SP2-LTSS_U2-3.0.101-0.7.21-pae-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-xen-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 GA

SLES11SP3-GA-3.0.76-0.11-default-x86
SLES11SP3-GA-3.0.76-0.11-default-x86_64
SLES11SP3-GA-3.0.76-0.11-pae-x86
SLES11SP3-GA-3.0.76-0.11-xen-x86
SLES11SP3-GA-3.0.76-0.11-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U1

SLES11SP3-U1-3.0.82-0.7-default-x86
SLES11SP3-U1-3.0.82-0.7-default-x86_64
SLES11SP3-U1-3.0.82-0.7-pae-x86
SLES11SP3-U1-3.0.82-0.7-xen-x86
SLES11SP3-U1-3.0.82-0.7-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U2

SLES11SP3-U2-3.0.93-0.8-default-x86
SLES11SP3-U2-3.0.93-0.8-default-x86_64
SLES11SP3-U2-3.0.93-0.8-pae-x86
SLES11SP3-U2-3.0.93-0.8-xen-x86

SLES11SP3-U2-3.0.93-0.8-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U3

SLES11SP3-U3-3.0.101-0.8-default-x86
SLES11SP3-U3-3.0.101-0.8-default-x86_64
SLES11SP3-U3-3.0.101-0.8-pae-x86
SLES11SP3-U3-3.0.101-0.8-xen-x86
SLES11SP3-U3-3.0.101-0.8-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U4

SLES11SP3-U4-3.0.101-0.15-default-x86
SLES11SP3-U4-3.0.101-0.15-default-x86_64
SLES11SP3-U4-3.0.101-0.15-pae-x86
SLES11SP3-U4-3.0.101-0.15-xen-x86
SLES11SP3-U4-3.0.101-0.15-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U5

SLES11SP3-U5-3.0.101-0.21-default-x86
SLES11SP3-U5-3.0.101-0.21-default-x86_64
SLES11SP3-U5-3.0.101-0.21-pae-x86
SLES11SP3-U5-3.0.101-0.21-xen-x86
SLES11SP3-U5-3.0.101-0.21-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U6

SLES11SP3-U6-3.0.101-0.29-default-x86
SLES11SP3-U6-3.0.101-0.29-default-x86_64
SLES11SP3-U6-3.0.101-0.29-pae-x86
SLES11SP3-U6-3.0.101-0.29-xen-x86
SLES11SP3-U6-3.0.101-0.29-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U7

SLES11SP3-U7-3.0.101-0.31-default-x86
SLES11SP3-U7-3.0.101-0.31-default-x86_64
SLES11SP3-U7-3.0.101-0.31-pae-x86
SLES11SP3-U7-3.0.101-0.31-xen-x86
SLES11SP3-U7-3.0.101-0.31-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U8

SLES11SP3-U8-3.0.101-0.35-default-x86
SLES11SP3-U8-3.0.101-0.35-default-x86_64
SLES11SP3-U8-3.0.101-0.35-pae-x86
SLES11SP3-U8-3.0.101-0.35-xen-x86
SLES11SP3-U8-3.0.101-0.35-xen-x86_64

SUSE Linux Enterprise Server 11 SP 4 GA

SLES11SP4-GA-3.0.101-63-default-x86
SLES11SP4-GA-3.0.101-63-default-x86_64
SLES11SP4-GA-3.0.101-63-pae-x86
SLES11SP4-GA-3.0.101-63-xen-x86
SLES11SP4-GA-3.0.101-63-xen-x86_64

SUSE Linux Enterprise Server 11 SP 4 U1

SLES11SP4-U1-3.0.101-65-default-x86
SLES11SP4-U1-3.0.101-65-default-x86_64
SLES11SP4-U1-3.0.101-65-pae-x86
SLES11SP4-U1-3.0.101-65-xen-x86
SLES11SP4-U1-3.0.101-65-xen-x86_64

SUSE Linux Enterprise Server 11 SP 4 U2

SLES11SP4-U2-3.0.101-68-default-x86
SLES11SP4-U2-3.0.101-68-default-x86_64
SLES11SP4-U2-3.0.101-68-pae-x86
SLES11SP4-U2-3.0.101-68-xen-x86
SLES11SP4-U2-3.0.101-68-xen-x86_64

SUSE Linux Enterprise Server 12 SP 4 GA

SLES12SP4-GA-4.12.14-94.41-default-x86_64

Other Linux Distributions That Use blkwatch Drivers

PlateSpin Migrate supports other Linux distributions listed in Table D-1 if the distribution is based on a supported release version of Red Hat Enterprise Linux. You can use the precompiled blkwatch driver for the supported Linux Distribution.

Table D-1 *Blkwatch Driver Support for Other Linux Distributions*

Other Linux Distribution	Based on a Supported Release Version for RHEL or SLES	Notes
CentOS	Red Hat Enterprise Linux	Use RHEL blkwatch drivers. Precompiled RHEL drivers are available as noted in “List of Distributions” on page 346.
Oracle Linux (OL) (formerly Oracle Enterprise Linux (OEL))	Red Hat Enterprise Linux	Blkwatch drivers are available for the standard kernel and the Unbreakable Enterprise Kernel (UEK) as noted in the “List of Distributions” on page 346. For other Oracle Linux distributions, precompiled drivers are not available for UEK kernels. Use RHEL blkwatch drivers for the corresponding Red Hat Compatible Kernel (RHCK). Workloads using the Oracle Linux Unbreakable Enterprise Kernel are not supported in Migrate 12.1 and lower versions.

For a list of supported kernel distributions, see “List of Distributions” on page 346.

E Synchronizing Serial Numbers on Cluster Node Local Storage

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

To download and run the utility:

1. From the Micro Focus Downloads (https://www.microfocus.com/support-and-services/download/) site, search for the PlateSpin Migrate product, then click **Submit Query**.
2. On the Products tab, select PlateSpin Migrate 2020.2 to go to the release-specific download page, then click **proceed to download**.
3. On the download page, click **download** on the `VolumeManager.exe` line or select the comparable download manager link, then save the file.
4. Copy the downloaded file to an accessible location on each cluster node.
5. On the active node of the cluster, open an administrative command prompt, navigate to the location of the downloaded utility, and run the following command:

   ```
   VolumeManager.exe -l
   ```

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

   ```
   Volume Listing:

   --------------------

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

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

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

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

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

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

6d When you have successfully changed all of the volume serial numbers on a node of the cluster, you need to restart that node.

6e Repeat Step 6a through Step 6d for each node of the cluster.

7 (Conditional) If the cluster has already been migrated in a PlateSpin environment, we recommend running a full replication on the active node to ensure that any changes are propagated to the database.

F Migrate Agent Utility

Migrate Agent is a command line utility that you can use to install, upgrade, query, or uninstall the block-based transfer drivers. The utility also enables you to register source workloads with PlateSpin Migrate servers and send details about the workloads to the server via HTTPS (TCP/443, outbound). See “Using Migrate Agent to Register Workloads” on page 369.

Requirements for Migrate Agent Utility

Ensure that your source workloads and network environment meets the following requirements for using the Migrate Agent Utility.

Supported Migrations for Migrate Agent

- Migrate Agent requires that the source workload is active during the migration.
- Migrate Agent is supported for automated migrations performed using Migrate Client and Migrate Web Interface.
- Migrate Agent is not supported for semi-automated (X2P) migrations performed using Migrate Client and Migrate Web Interface.

Deployment Requirements for Migrate Agent

When you use the Migrate Agent for workload registration and migration, ensure that your migration environment meets the following requirements:

- Public IP addresses are required for the PlateSpin Migrate server host, replication network, and target machines. In some deployment scenarios, public IP addresses are also required for source machines.
  - Ensure that workloads can reach the public IP address for Migrate server.

    Set the **AlternateServerAddress** parameter to the Migrate server's public IP address on the PlateSpinConfiguration page. For Migrate servers deployed from Azure Marketplace or deployed using the AWS Quick Start, Migrate automatically adds the public IP address to this parameter. See "Configuring Alternate IP Addresses for PlateSpin Server" on page 111.
  - Enable a public IP address for the replication network when you configure the migration for a workload.
  - Migrate automatically configures public IP addresses on target machines during migration.
- For information about network requirements for registration and migration, see
  - "Requirements for Workload Registration" on page 56
  - "Requirements for Migration of Workloads Registered Using Migrate Agent" on page 58

  NOTE: Refer to the deployment diagrams based on your migration target to understand the ports and flow of information between the various migration components. See Part III, "Preparing Your Migration Environment," on page 143.

- Ensure that you configure source workloads to support outbound traffic for the following ports:
  - HTTPS port (TCP/443)
  - Replication port (TCP/3725 is the default)

    The replication port is configurable. If you modify the **FileTransferPort** parameter on the PlateSpin Configuration page, you must modify your firewall settings accordingly.
- When you use the Migrate Agent on the source workload, the source workload contacts the target workload for data transfers. Ensure that the value of the `SourceListensForConnection` parameter setting is set to False, which is the default value See "Configuring the Contact Direction for the Replication Port" on page 114.
- For cloud-based Migrate servers, the server is configured by default for migration to the target type that matches its parent could environment. If the source workloads are in the parent cloud environment for migration to a different target, you must remove the default value (leave the field blank) for the **ServerIsHostedInCloud** parameter to allow all target types to be available in the Add Target dialog.

Usage Requirements for Migrate Agent Utility

- **Software Prerequisites**

  Migrate Agent Utility for Linux requires the source machine to have GNU C Library (glibc) 2.11.3 or higher installed.
- **Reboot**

A reboot of the source Windows workload is required when you install, uninstall, or upgrade block-based transfer drivers. A reboot is not required for source Linux workloads.

Although a reboot is always required for Windows workloads, using the Migrate Agent utility allows you to better control when the action occurs and therefore, when the server reboots. For example, you can use the Migrate Agent utility to install the drivers during scheduled down time, instead of during the first replication.

- **Credentials**
  - For Windows workloads, Migrate Agent Utility requires `Administrator` privileges to execute commands.
  - For Linux workloads, Migrate Agent Utility requires `root`-level access to execute the commands. A non-`root` user account must be authorized to use the `sudo` command. That is, the user name must be listed as an authorized user in the `/etc/sudoers` configuration file. For information on using an account other than `root`, see KB Article 7920711 (https://support.microfocus.com/kb/doc.php?id=7920711).

NOTE: For source Linux workloads in Amazon Web Services, AMI templates automatically create a default non-`root` system user account that is enabled for `sudo`. The user name for this account varies by AMI provider. For Amazon Linux images, the non-`root` user name is `ec2-user` for most Linux distributions. It is `centos` for CentOS AMIs. For more information, refer to your AMI provider documentation.

In AWS, a non-`root` user must run the `sudo -i` command to access the `root` shell and then run the Migrate Agent commands. Typing `sudo` in each Migrate Agent Utility command might result in a failure on some source workloads.

Migrate Agent Utility for Windows

Downloading and Installing Migrate Agent on a Source Windows Workload

To download and install the Migrate Agent utility for Windows to the source workload:

1. Log in to the source Windows machine as the Administrator user.
2. In a web browser, launch the PlateSpin Migrate Web Interface and log in.
3. Click the **Downloads** tab.
4. Click the Migrate Agent application link for the Windows target platform, then save the compressed `MigrateAgent.cli.exe` file.
5. Extract the contents of the file to access the executable file.
6. (Optional) View the Migrate Agent Help by entering

   ```
   MigrateAgent.cli.exe -h
   ```

Migrate Agent Commands for Windows

The syntax for running the Migrate Agent utility for Windows is:

`MigrateAgent.cli.exe {command} [command_option] [/psserver=%IP%]`

Table F-1 describes the commands, command options, and switch available for the `MigrateAgent.cli.exe` command on Windows.

Table F-1 *Migrate Agent Utility for Windows Commands, Command Options, and Switch*

Usage	Description
Commands	
`h | ? | help`	Displays usage and options for the command.
`logs | view-logs`	Opens the application log directory.

`reg | register`
 	
`/reg /psserver=%IP% /username=<username> [[ /password=<password>] | [/pwdfile=<path-to-password-file>]] [/proxyurl=<your-proxy-server-url>]`
 	
If you do not specify the password or a path to a file that contains the password, you will be prompted for the password. The password is obscured as you type it and it does not appear in the process list.
 	
Example:
 	
`MigrateAgent.cli.exe /register /psserver=10.10.10.101 /username=jsmith /password=jspwd`
	Registers this machine as a workload on the specified server. It also checks for driver upgrades from the specified PlateSpin Server.
	
	Enables you to add workloads that cannot be discovered. Registered workloads differ from discovered workloads in the following ways:
	
	◆ Registered source workloads do not store the source credentials.
	◆ You must use Migrate Agent to install, upgrade, and remove Block-based Transfer (BBT) drivers from registered source workloads.
	◆ After you delete the contract for a registered source workload, use the clu | cleanup command to remove the OFX controller from the workload.

`status`
 	
`/status [/psserver=%IP%]`
	Shows installation status for the PlateSpin controller and drivers on this workload.
	
	If you specify the PlateSpin Server, it checks for driver upgrades from the server.

`din | driver-install`
 	
`/din [/psserver=%IP%] [/proxyurl=<your-proxy-server-url>]`
	Installs the PlateSpin drivers.
	
	NOTE: Before you install block-based transfer drivers on source Windows workloads, ensure that you have applied the latest Windows updates on the workload.
	
	If you specify the PlateSpin Server, it checks for driver upgrades from the server.

`dup | driver-upgrade`
 	
`/dup [/psserver=%IP%] [/proxyurl=<your-proxy-server-url>]`
	Upgrades the PlateSpin drivers.
	
	If you specify the PlateSpin Server, it checks for driver upgrades from the server.

Usage	Description
`dun | driver-uninstall` `[/dun /psserver=%IP%]`	Uninstalls the PlateSpin drivers.

`con | config`
`/con /setting=<setting_name>:<value>`
Example:
`migrateagent.cli.exe /config /setting=psserver:10.10.10.202`
	Specifies the name of the setting and its value to change in the configuration file on this workload.
	The `psserver` option stops the OFX Controller (`ofxcontroller`) service, modifies the `OfxController.exe.config` file with the new IP address, and restarts the service. If you modify the public IP address of the PlateSpin Server, you must run this command on each of the source workloads that are configured for the server.

`clu | cleanup`	Removes PlateSpin drivers, the OFX controller software, and Migrate Agent related files.
Switch	
`/psserver=%IP%`	Specifies the IPv4 address of the PlateSpin Server. Downloads the block-based transfer drivers from the specified server when you invoke the `status`, `driver-install`, or `driver-upgrade` options.
Command Options	
`username` `/username=value`	Specifies the PlateSpin Server user name for an administrator-level user with rights to add a workload.
`password | pwd | p` `/password=value`	Specifies the password for the specified PlateSpin Server user name. If you exclude the password from the command line, the script will prompt for it. The password is obscured as you type it and it does not appear in the process list. Do not combine this option with the `pwdfile` option.
`pwdfile | pf` `/pwdfile=value`	Specifies the path to a file that contains the password for the specified PlateSpin Server user name. Do not combine this option with the `password` option.

`setting`
`/setting=<setting_name>:<value>`
	Specifies the setting name and value of the configuration setting to modify.
	Supported setting names are:
	`psserver`
	`altAddress`
	`heartbeat`

Usage	Description

`proxyurl | pr`
 	
`/proxyurl=<your-proxy-server-url>`
 	
-OR-
 	
`/pr=<your-proxy-server-url>`
 	
Specify the HTTP proxy URL in the following format to send all HTTP and HTTPS traffic to Migrate Server through the proxy server:
 	
http://<*proxy-server-fqdn*>:<*port-number*>
 	
Example:
 	
`MigrateAgent.cli.exe /register /psserver=10.10.10.101 /username=jsmith /password=jspwd /proxyurl=http://myproxy.example.com:8088`
	Proxy URL to use for connecting to the PlateSpin server.
	
	If HTTP and HTTPS communications with the PlateSpin Migrate Server are required to go through your proxy server, you can use the `proxyurl` option to specify the URL of your proxy server. Include the option with the following commands on source Windows workloads:
	
	`reg|register`
	`din|driver-install`
	`dup|driver-upgrade`
	
	The `MigrateAgent.cli.exe status` command reports whether a proxy is set for the PlateSpin Controller Service. If a proxy is set, the Proxy URL parameter displays the proxy URL.
	
	NOTE: PlateSpin Migrate requires that you open port 3725 in your firewall for data transfer and communications between workloads. Migrate will not send data traffic through your proxy server.

Migrate Agent Utility for Linux

Before you install or use Migrate Agent, ensure that your system satisfies the Requirements for Migrate Agent Utility.

- "Downloading and Installing Migrate Agent on a Source Linux Workload" on page 364
- "Migrate Agent Commands for Linux" on page 365

Downloading and Installing Migrate Agent on a Source Linux Workload

Before you install Migrate Agent Utility for Linux, ensure that the source machine has GNU C Library (glibc) 2.11.3 or higher installed.

Ensure that you download the application with the appropriate architecture for your source Linux machines. The file name is case sensitive.

- **64-bit:** `MigrateAgent-x86_64.tar.gz`
- **32-bit:** `MigrateAgent-x86.tar.gz`

To download and install the Migrate Agent utility for Linux on the source workload:

1. Log in to the source Linux workload as the `root` user.
2. Use either of the following methods to get the `MigrateAgent-`*`arch`*`.tar.gz` file.

Replace *arch* with the appropriate architecture (x86_64 or x86).

- Download the zipped file from the Web Interface:

  1. In a web browser, launch the PlateSpin Migrate Web Interface and log in.

     https://<*Your_PlateSpin_Server* >/Migrate

     Replace *Your_PlateSpin_Server* with the DNS name or IP address of your PlateSpin Migrate server.

  2. Click the **Downloads** tab.
  3. Click the Migrate Agent application link for the appropriate Linux platform (x86_64 or x86), then save the `MigrateAgent-arch.tar.gz` file.

  -OR-

- Use the `wget` command to copy the file from the PlateSpin Server.

  NOTE: If the operating system on the PlateSpin Server host accepts only TLS 1.2 connections, use `wget` version 1.16.1 or higher on your source Linux workload.

  ---

  1. Launch a terminal, then enter

     ```
     wget --no-check-certificate --http-user=<username> --http-
     password=<password> https://<Your_PlateSpin_Server>/Migrate/Downloads/
     MigrateAgent-<arch>.tar.gz
     ```

     Replace *Your_PlateSpin_Server* with the DNS name or IP address of your PlateSpin Migrate server. Replace *arch* with x86_64 or x86.

3 Open the `MigrateAgent-arch.tar.gz` file in Archive Manager, then extract the `MigrateAgent` directory and its contents to the `root` directory (`/`).

  Alternatively, in a shell prompt, enter

  ```
  tar xvf MigrateAgent-<arch>.tar.gz
  ```

  Replace *arch* with x86_64 or x86.

4 Change directory to the `/MigrateAgent` directory, then list its contents. In a terminal, enter:

  ```
  cd MigrateAgent
  ls
  ```

  The directory contains a `commands` file and the `MigrateAgent` script file.

5 View the command Help by entering:

  ```
  ./MigrateAgent -h
  ```

Migrate Agent Commands for Linux

The syntax for running the Migrate Agent utility is:

```
./MigrateAgent [Command] [-h]
```

Table F-2 describes the options and arguments available for the `MigrateAgent` command on Linux.

Table F-2 *Migrate Agent Utility for Linux Command Options and Arguments*

Usage	Description
Commands	

`register <server> <user> [-pr <your-proxy-server-url>] [[-p password] | [-pf <password-file-path>]]`
 	
For *server*, specify the DNS name or IP address of your PlateSpin Migrate server.
 	
For *user*, specify a valid PlateSpin Server user name for an administrator-level user with rights to add a workload.
 	
For the password, do one of the following:
◆ Use the `-p` option and type the password in the command for the specified PlateSpin user name.
`-p mypassword`
◆ Use the -pf option to specify the path to a file that contains the password for the specified PlateSpin user name.
`-pf /tmp/jsmith-password-file.txt`
◆ Do not specify the password in the command. You will be prompted to enter the password at the command line.
 	
Example:
`./MigrateAgent register 10.10.10.101 jsmith -p jspwd`
	Registers this machine as a workload on the specified server. It also checks for driver upgrades from the specified PlateSpin Server.
	
	Enables you to add workloads that cannot be discovered. Registered workloads differ from discovered workloads in the following ways:
	◆ Registered source workloads do not store the source credentials.
	◆ You must use Migrate Agent to install, upgrade, and remove the Linux blkwatch drivers from registered source workloads.
	◆ After you delete the contract for a registered source workload, use the cleanup command to remove the OFX controller from the workload.

`status <server>` For *server*, specify the DNS name or IP address of your PlateSpin Migrate server.	Shows installation status for the PlateSpin controller and drivers. It also checks for driver upgrades from the specified PlateSpin Server.

Usage	Description
`driver-install` <*server*> [-pr <*your-proxy-server-url*>] For *server*, specify the DNS name or IP address of your PlateSpin Migrate server.	Installs the appropriate PlateSpin blkwatch drivers. It also checks for driver upgrades from the specified PlateSpin Server. If precompiled drivers for the kernel running on your Linux workload are not available, the utility does the following: 1. Compiles the required drivers and installs them on the source workload. 2. Uploads the compiled drivers to the following PlateSpin Server library that packages the precompiled drivers and custom blkwatch drivers. `<install-drive>:\Program Files\PlateSpin Migrate Server\Packages\9D470BFF-E30F-4015-83D8-4D308A8C8B1B` By default, PlateSpin Migrate packages the precompiled drivers in a `1.package` file. The PlateSpin Migrate Agent Utility creates a new *x*`.package` file that is a copy of the latest available package file and adds the compiled custom drivers to this newly created package file. The value of *x* is the numeric value in the latest package filename incremented with 1. For example: If the latest package file is `1.package`, the utility creates a `2.package` file that contains the following ◆ The existing drivers from `1.package` (default precompiled drivers and custom blkwatch drivers). ◆ The utility compiled new custom drivers. **NOTE:** PlateSpin Migrate uses the latest package file containing the drivers.
`driver-upgrade` <*server*> [-pr <*your-proxy-server-url*>] For *server*, specify the DNS name or IP address of your PlateSpin Migrate server.	Upgrades the installed PlateSpin blkwatch driver. It also checks for driver upgrades from the specified PlateSpin Server.
`driver-uninstall`	Uninstalls the installed PlateSpin blkwatch driver from the source Linux workload.
`configure <server> <new-server>` [-pr <*your-proxy-server-url*>] For *server*, specify the DNS name or IP address of your PlateSpin Migrate server. For *new-server*, specify the new DNS name or IP address of the PlateSpin Migrate server. Example: `./MigrateAgent configure 10.10.10.10 10.10.20.20`	Stops the OFX Controller (`ofxcontroller`) service, modifies the OFX Controller configuration file with the new address, and restarts the service. If you modify the public IP address of the PlateSpin Server, you must run this command on each of the source workloads that are configured for the server.

Usage	Description
`cleanup`	Removes PlateSpin drivers and the OFX controller.
Command Options	
`server`	Specifies the DNS name or IP address of the PlateSpin Migrate Server. Downloads the blkwatch drivers from the specified server when you invoke the `status`, `driver-install`, or `driver-upgrade` options.
`user`	Specifies the PlateSpin Server user name for an administrator-level user with rights to add a workload.
Options	
`-h`, `--help`	Displays usage and options for the command.

`-pr, --proxyurl`
 	
`--proxyurl<`*`your-proxy-server-url`*`>`
 	
-OR-
 	
`-pr <`*`your-proxy-server-url`*`>`
 	
Specify the HTTP proxy URL in the following format to send HTTP and HTTPS traffic to Migrate Server through the proxy server:
 	
http://<*proxy-server-fqdn*>:<*port-number*>
 	
Example:
 	
`./MigrateAgent register 10.10.10.101 jsmith -p jspwd -pr http://myproxy.example.com:8088`
	Proxy URL to use for connecting to the PlateSpin server.
	
	If HTTP and HTTPS communications with the PlateSpin Migrate Server are required to go through your proxy server, you can use the `proxyurl` option to specify the URL of your proxy server. Include the option with the following commands on source Linux workloads:
	
	`reg|register`
	`din|driver-install`
	`dup|driver-upgrade`
	`configure`
	
	The `./MigrateAgent status` command reports whether a proxy is set for the PlateSpin Controller Service. If a proxy is set, the Proxy URL parameter displays the proxy URL.
	
	NOTE: PlateSpin Migrate requires that you open port 3725 in your firewall for data transfer and communications between workloads. Migrate will not send data traffic through your proxy server.

`-p, --password`
 	
`-p <`*`user_password`*`>`
	Specifies the password for the PlateSpin Server user name.
	
	If you exclude the password from the command line, the script will prompt for it. The password is obscured as you type it and it does not appear in the process list.
	
	Do not combine this option with the `passwordfile` option.

`-pf, --passwordfile`
 	
`-pf <`*`passwordfile_path`*`>`
	Specifies the path to a file that contains the password for the specified PlateSpin Server user name.
	
	Do not combine this option with the `password` option.

Usage	Description
Logging	

`logging.json`	Contains the logging configuration settings in JSON format for logging Migrate Agent utility actions.
	
	To view logging settings, use the `cat` command:
	
	`cat MigrateAgent/logging.json`
	
	You can edit the file in a text editor. Set the level of the logging by changing the `"level:"` value from `"DEBUG"` to `"INFO"` or `"ERROR"`. For example:
	
	`"level": "DEBUG"`
	
	or
	
	`"level": "INFO"`
	
	or
	
	`"level": "ERROR"`
	
	Logged messages are written by default to the `MigrateAgent.log` file in the `MigrateAgent` directory. You can modify the log file name setting in the `logging.json` file.

`MigrateAgent.log`	Contains the logged messages for the `MigrateAgent` command. To view the log, use the `cat` command. `cat MigrateAgent.log`

Using Migrate Agent to Register Workloads

You can use the Migrate Agent utility for registration and discovery instead of automated discovery. Using Migrate Agent is required to register and discover details about source workloads in scenarios where automated discovery is not possible, such as:

- When you deploy Migrate server in the cloud without deploying a site-to-site VPN between your network and your cloud environment.
- When you plan cloud-to-cloud migrations without deploying a site-to-site VPN between the participating locations: your network, your source cloud environment, and your target cloud environment.
- When corporate network or policy restrictions prohibit opening inbound ports on source workloads.

  For information about inbound ports required for automated discovery of Windows and Linux workloads, see “Requirements for Discovery” on page 53.

Migrate Agent enables you to migrate a Windows workload without opening any inbound ports, such as SMB or NetBIOS. Only HTTPS (TCP/443) and a replication port (TCP/3725 is the default) are needed outbound for the source Windows workloads. For source Linux workloads, you also need to open the SSH port (TCP/22). See “Requirements for Workload Registration” on page 56.

When you use the Migrate Agent on the source workload, the source workload contacts the target workload for data transfers. The direction is controlled at the server level. Ensure that the value of the `SourceListensForConnection` parameter setting is set to False, which is the default value See "Configuring the Contact Direction for the Replication Port" on page 114.

You must install Migrate Agent on each source workload. When you use the `register` option, Migrate Agent performs discovery locally on the workload and sends its details to the Migrate Server through HTTPS (TCP/443). After you register the workload, use the Migrate Web Interface to configure the workload migration to the target cloud where the Migrate Server instance is deployed.

Registered workloads differ from discovered workloads in the following ways:

- Registered source workloads do not store the source credentials on the Migrate Server.
- You must use Migrate Agent to install, upgrade, and remove the Windows PlateSpin drivers from registered source workloads.
- After you delete the contract for a registered source workload, use the `cleanup` command to remove the OFX controller from the workload. See clu | cleanup for Windows workload and cleanup for Linux workload

See the following procedures in "Registering Workloads and Discovering Details with Migrate Agent" on page 303:

- Windows Workload Registration and Discovery with Migrate Agent
- Linux Workload Registration and Discovery with Migrate Agent

Using Migrate Agent with Block-Based Transfer Drivers

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

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

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

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

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

The task completed successfully
PlateSpin Controller Service Status
    The PlateSpin Controller service is not installed

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

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

Figure F-1 *Reboot Notification Task*

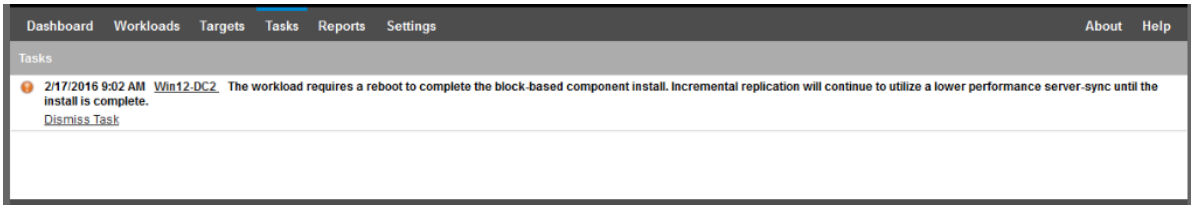

Figure F-2 *Reboot Notification During Replication*

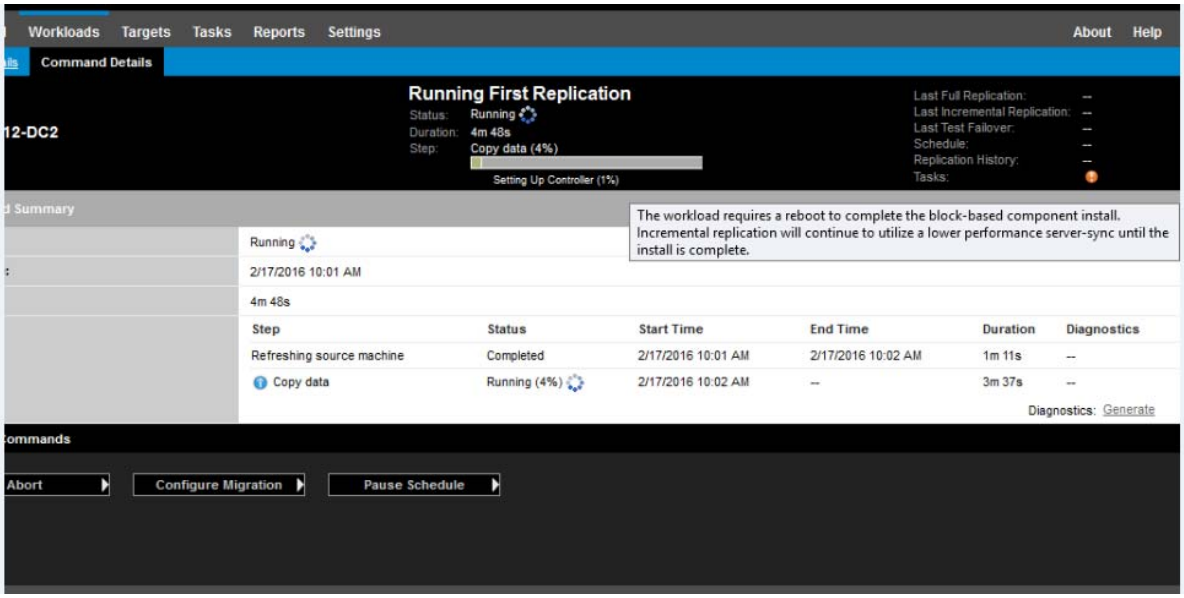

Rebooting the source machine applies and starts the installed or upgraded drivers. If the driver was recently installed, after the reboot, one full replication or a server-sync replication is required in order to ensure that all of a source's changes are captured. This server-sync replication will be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

Figure F-3 *Server-Sync Required Notification*

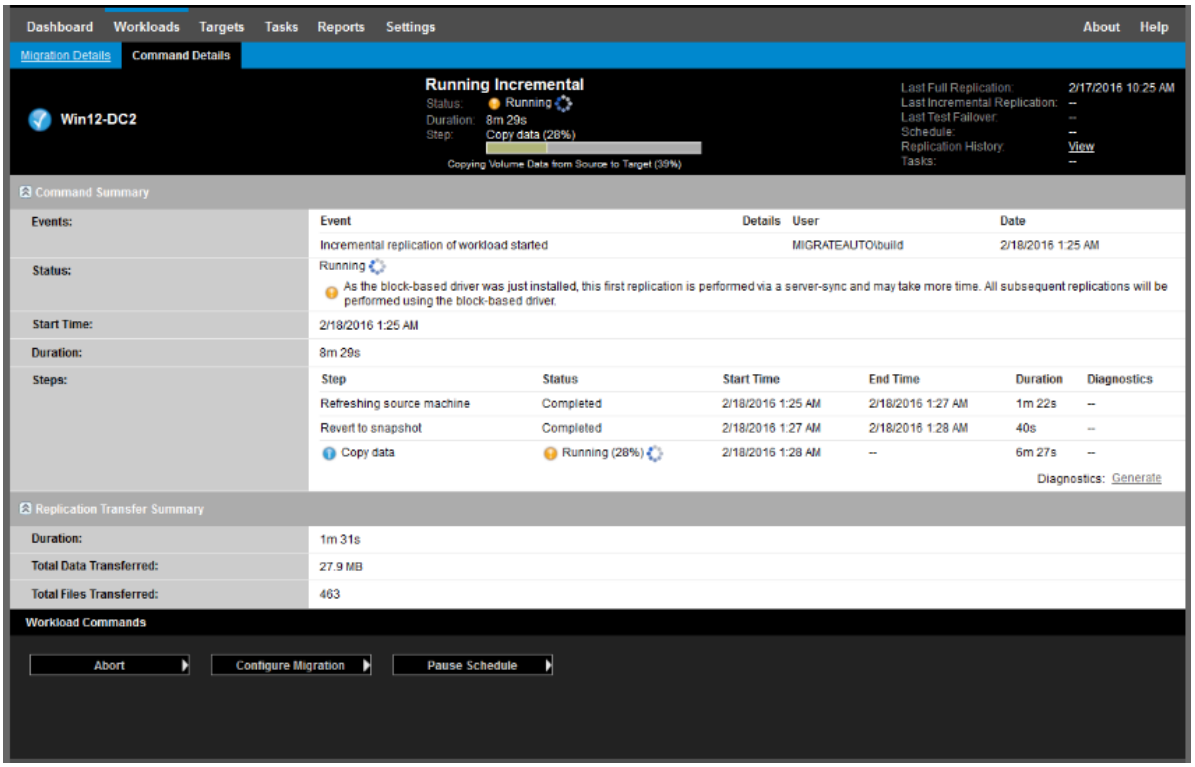

G PlateSpin ISO Image

The PlateSpin ISO image file enables you to boot BIOS or UEFI firmware-based target physical machines and virtual machines during semi-automated migrations and semi-automated Server Sync operations. The semi-automated migration is used to transfer the workload to a physical machine or virtual machine that has been registered in PlateSpin Migrate. This registration occurs when you boot the target machine with the PlateSpin ISO image and register with the PlateSpin Server by following the prompts. It also discovers the target's hardware details and sends them to the server.

Downloading the PlateSpin ISO Images

You can download the PlateSpin ISO image from the PlateSpin Migrate software download page at Micro Focus Downloads (https://www.microfocus.com/support-and-services/download/). Search for downloads for the current product and version:

Product: PlateSpin Migrate
Version:2020.2
Dates: All dates

The compressed `.iso` files is contained in `PhysicalTarget-2020_2.zip` at the download site. The ISO file uses the SUSE Linux Enterprise Server (SLES) operating system for the Linux RAMDisk (LRD). The LRD contains a minimal set of system files, drivers, and executables, sufficient for an initial, temporary boot. See Table G-1 for information about the operating system version used for the LRD and boot options.

Table G-1 *PlateSpin ISO Image File*

PlateSpin ISO Image File	LRD OS	Workload Architecture	FCoE
`bootofx.x2p.iso`	SLES 12 SP3	64-bit	Optional
`bootofx.x2p.sles11sp4.iso`	SLES 11 SP4	32-bit	No

Preparing the PlateSpin ISO Image for Target Registration and Discovery

1. Download the PlateSpin ISO image from Micro Focus Downloads and extract the contents. See Downloading the PlateSpin ISO Images.
2. (Optional) Inject additional device drivers for Linux workloads into the PlateSpin ISO image, complete the steps in Injecting Additional Device Drivers into the PlateSpin ISO Image.
3. (Optional) For an unattended registration, modify the PlateSpin ISO to provide the appropriate responses from an answer file. See Adding Registration Information to the PlateSpin ISO for Unattended Registration of Physical or Virtual Machines.
4. Save the PlateSpin ISO image:
   - **Physical Machine:** Burn the PlateSpin ISO image on a CD or save it to the required media, from which your target can boot.
   - **Virtual Machine:** Save the PlateSpin ISO image on the virtual host for a target VM in a location where you can use it to boot the target machine.
5. Use native tools to prepare the target machine to boot from the PlateSpin ISO image.

   Ensure that the machine is configured to restart on reboot and that you attach the PlateSpin ISO file as a boot CD for the VM.

For information about registering the target machine, see the following:

- “Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO” on page 288
- “Registering and Discovering Details for Target Physical Machines with PlateSpin ISO” on page 291

Injecting Additional Device Drivers into the PlateSpin ISO Image

The PlateSpin ISO image contains a large library of device drivers sufficient to boot most common targets. However, occasionally you might want to use your own, such as lesser-known, vendor-specific or custom-developed drivers for Linux workloads.

The `rebuildiso.sh` script that helps you rebuild the ISO file has different options and kernel version requirements, as shown in Table G-2.

Table G-2 *Comparison of rebuildiso.sh for the PlateSpin ISO*

PlateSpin ISO Image File	LRD OS	Kernel Version	Bit Switch
`bootofx.x2p.iso`	SLES 12 SP3	`4.4.73-5-default`	None, assumes 64-bit
`bootofx.x2p.sles11sp4.iso`	SLES 11 SP4	`3.1.101-63-pae`	`-m32` for 32-bit `-m64` for 64-bit

To inject drivers into the PlateSpin ISO image for Linux workloads:

1. Download and extract the PlateSpin ISO images. See Downloading the PlateSpin ISO Images.
2. Obtain or compile the required `*.ko` driver files.

   IMPORTANT: Ensure that the drivers are valid for the kernel version included with the ISO file you are trying to rebuild. See Table G-2, “Comparison of rebuildiso.sh for the PlateSpin ISO,” on page 374.

3. Mount the ISO image in any Linux machine (`root` credentials required). Use the following command syntax:

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

4. Copy the `rebuildiso.sh` script, located in the `/tools` subdirectory of the mounted ISO file, into a temporary working directory.
5. Create another working directory for the required driver files and save them in that directory.
6. In the directory where you saved the `rebuildiso.sh` script, run the following command as `root`, according to the ISO file you are rebuilding.

   For the PlateSpin ISO for SLES 12 SP3:

   ```
   ./rebuildiso.sh –i <ISO_file> -d <driver_dir>
   ```

   For the PlateSpin ISO for SLES 11 SP4:

   ```
   ./rebuildiso.sh –i <ISO_file> -d <driver_dir> -m32
   ```

   ```
   ./rebuildiso.sh –i <ISO_file> -d <driver_dir> -m64
   ```

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

   NOTE: To rebuild Migrate LRD ISO, a minimum of `genisoimage 1.1.11` is required. By default, operating systems such as RHEL 7 and CentOS 7 have the required `genisoimage` version.

7. Unmount the ISO file (execute the command `unmount` *<mount_point>*).

Adding Registration Information to the PlateSpin ISO for Unattended Registration of Physical or Virtual Machines

PlateSpin Migrate provides a mechanism for automating the registration and discovery of details for a target physical or virtual machine. You must first update the PlateSpin ISO image with specific registration information before booting the target.

For details, see KB Article 7013485 (https://support.microfocus.com/kb/doc.php?id=7013485).

Using PlateSpin ISO

After you have prepared the PlateSpin ISO for your environment, you can use the file to register and discover target physical machines or target virtual machines in a semi-automated migration or Server Sync operation. See the following procedures in "Discovering Target Platforms":

- "Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO" on page 288
- "Registering and Discovering Details for Target Physical Machines with PlateSpin ISO" on page 291

V Configuring Workloads

After you discover targets and workloads, you are ready to prepare for migration by configuring migration jobs for your workloads.

22 Prerequisites for Automated Migrations

PlateSpin Migrate Client and PlateSpin Migrate Web Interface enable you to automate migration of workloads to target virtualization platforms and target cloud platforms.

Supported Source Workloads for Automated Migration

In an automated migration, PlateSpin Migrate builds the target virtual machine on the destination platform based on the target workload details you configure for the conversion. Automation supports source workloads based on the destination target platform. For information about source workloads for supported virtualization and cloud platforms, see Table 22-2.

Table 22-1 *Supported Source Workloads for Automated Migrations*

Target Platform	Migrate Client	Migrate Web Interface
Amazon Web Services	Not supported	Table 2-3, "AWS: Supported Windows Platforms," on page 28 Table 2-4, "AWS: Supported Linux Platforms," on page 29
Microsoft Azure	Not supported	Table 2-5, "Azure: Supported Windows Platforms," on page 30 Table 2-6, "Azure: Supported Linux Platforms," on page 31
VMware vCloud Director	Not supported	Table 2-7, "vCloud: Supported Windows Platforms," on page 32 Table 2-8, "vCloud: Supported Linux Platforms," on page 33
VMware	Table 2-1, "Non-Cloud Platforms: Supported Windows Workloads," on page 24 Table 2-2, "Non-Cloud Platforms: Supported Linux Workloads," on page 26	Table 2-1, "Non-Cloud Platforms: Supported Windows Workloads," on page 24 Table 2-2, "Non-Cloud Platforms: Supported Linux Workloads," on page 26

Supported Target Platforms for Automated Migrations

In an automated migration, PlateSpin Migrate prepares the virtual machine on the target platform before the replications begin. You can schedule when the first full replication begins. The Prepare Workload step must be executed prior to the scheduled start time. For information about supported virtualization and cloud platforms, see Table 22-2.

Table 22-2 *Supported Target Platforms for Automated Migrations*

Target Platform	Migrate Client	Migrate Web Interface
Amazon Web Services	Not supported	See “Amazon Web Services” in Table 2-15, “Supported Target Cloud Platforms for the Migrate Web Interface,” on page 43
Microsoft Azure	Not supported	See “Microsoft Azure” in Table 2-15, “Supported Target Cloud Platforms for the Migrate Web Interface,” on page 43
VMware vCloud Director	Not supported	See “VMware vCloud Director” in Table 2-15, “Supported Target Cloud Platforms for the Migrate Web Interface,” on page 43
VMware	Table 2-12, “Supported Target VMware Platforms for the Migrate Web Interface and Migrate Client,” on page 41	Table 2-12, “Supported Target VMware Platforms for the Migrate Web Interface and Migrate Client,” on page 41

Preparing Targets for Automated Migration

In an automated migration, PlateSpin requires information about the target platform where it will create the virtual machines. You must prepare your target environment for discovery, and discover the target. For information about configuring the target platform environment for use with PlateSpin Migrate, see Table 22-3. For discovery of target platforms, see “Discovering Details for Target Platforms” on page 283.

Table 22-3 *Prerequisites for Target Platforms*

Target Platform	Migrate Client	Migrate Web Interface
Amazon Web Services	Not supported	Chapter 8, “ Prerequisites for Migration to Amazon Web Services,” on page 145
Microsoft Azure	Not supported	“Prerequisites for Migration to Microsoft Azure” on page 167
VMware vCloud Director	Not supported	“Prerequisites for Migration to VMware vCloud Director” on page 201

Target Platform	Migrate Client	Migrate Web Interface
Cloud-to-Cloud	Not supported	"Prerequisites for Cloud-to-Cloud Migrations" on page 217
VMware	"Prerequisites for Migration to VMware" on page 247	"Prerequisites for Migration to VMware" on page 247

Network Connections and Bandwidth

Before you execute replications for an automated migration:

- Ensure that your network access and ports are properly configured. See "Requirements for Migration" on page 57.

  If you are using Migrate Agent, see "Requirements for Migration of Workloads Registered Using Migrate Agent" on page 58.
- Ensure that you test the connection to see if there are any connection or bandwidth issues, and resolve them. For information about optimizing throughput on the connection, see "Using the iPerf Network Test Tool to Optimize Network Throughput for PlateSpin Products" on page 589.

Automated Workflow

Refer to the checklist to understand the automated workflow:

- "Checklist for Automated Migration to AWS" on page 164
- "Checklist for Automated Migration to Azure" on page 198
- "Checklist for Automated Migration to vCloud" on page 207
- "Checklist for Automated Migration to VMware" on page 264
- "Checklist for Automated Migration from AWS to Azure" on page 222
- "Checklist for Automated Migration from Azure to AWS" on page 225
- "Checklist for Automated Migration from Azure to vCloud" on page 229
- "Checklist for Automated Migration from vCloud to Azure" on page 232
- "Checklist for Automated Migration from AWS to vCloud" on page 236
- "Checklist for Automated Migration from vCloud to AWS" on page 239

For information about configuring automated migration to a target platform, see:

- "Configuring Migration of a Workload to Amazon Web Services" on page 432
- "Configuring Migration of a Workload to Microsoft Azure" on page 454
- "Configuring Migration of a Workload to VMware vCloud Director" on page 470
- "Automated Migration to VMware Using Migrate Client" on page 489
- "Automated Migration to VMware Using Migrate Web Interface" on page 498

23 Prerequisites for Semi-Automated (X2P) Migrations

To migrate workloads to physical machines (X2P), use either the PlateSpin Migrate Client or the PlateSpin Migrate Web Interface. Use PlateSpin ISO to register the target physical machine with the PlateSpin Migrate server and report details about it. This manual process of target preparation and discovery is referred to as the *X2P workflow*.

Supported Source Workloads for X2P Migrations

You can also use the X2P workflow to migrate workloads to a virtual machines that you set up on a supported virtual host. You configure the VM with guest operating system type and version settings that match your source workload, in accordance with the features and capabilities of the target virtualization platform. For information about source workloads for supported virtualization platforms, see:

- Table 2-1, "Non-Cloud Platforms: Supported Windows Workloads," on page 24
- Table 2-2, "Non-Cloud Platforms: Supported Linux Workloads," on page 26

Supported Target Platforms for X2P Migrations

PlateSpin Migrate Client and PlateSpin Migrate Web Interface supports using the X2P workflow for migrations to physical machines and to any supported virtual host, even if an automated alternative is available. For information about supported virtualization platforms, see "Supported Target Virtualization Platforms" on page 40.

X2P Workflow for VMs

To migrate a workload to a VM on a virtual host:

1. Use the native interface of the required virtualization platform to set up the target virtual machine with guest operating system type and version settings that match your source workload, in accordance with the features and capabilities of the target virtualization platform.
2. Begin booting the newly created virtual machine by using the appropriate PlateSpin ISO image, load the appropriate driver, if needed, then continue the boot process.

   This special boot process discovers and registers the target virtual machine as a PlateSpin Migrate physical machine target. See "Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO" on page 288.

3 Use the PlateSpin Migrate Client or PlateSpin Migrate Web Interface to create and execute an X2P migration job.

4 Upon completion of the migration job, install virtualization enhancement software specific to the target virtualization platform.

For information about configuring semi-automated migration to a virtual machine running on virtualization hosts that PlateSpin Migrate regards as a physical machine:

- “Migration to VMs on VMware Using X2P Workflow” on page 506
- “Migration to Virtual Machines on KVM” on page 511
- “Migration to Oracle Cloud Infrastructure” on page 481

24 Configuration Essentials

When you configure a workload for migration, the workload type and target determine the configuration options available. This section describes the essentials for configuration of each parameter.

Configuration Workflows

Refer to the migration configuration sections for a step-by-step walk through the migration configuration for the various migration job types.

Configuration Workflows Using Migrate Client

The PlateSpin Migrate Client supports migration of workloads to VMware platforms, physical machines, and server-sync.

- Migration to VMware
- Migration of Windows Clusters
- Migration to Virtual Machines on KVM
- Migration to Physical Machines
- Synchronizing Workloads with Server Sync

Configuring Workflows Using Migrate Web Interface

The PlateSpin Migrate Web Interface supports large scale migration of workloads to VMware and cloud platforms such as Amazon Web Services, Microsoft Azure, and VMware vCloud Director.

- Migration to Amazon Web Services
- Migration to Microsoft Azure
- Migration to VMware vCloud Director
- Automated Migration to VMware Using Migrate Web Interface
- Migration to VMs on VMware Using X2P Workflow
- Preparing for Migration of Windows Clusters

Initiating a Migration Job

After workload discovery, the migration job for the workload is in a unconfigured state. Migration jobs are not automatically initiated with default settings. You must initiate the migration job by starting configuration for the migration.

- “Prerequisites for Migration Jobs” on page 386
- “Initiate a Migration Job Using Migrate Client” on page 387
- “Initiate a Migration Job Using the Migrate Web Interface” on page 388

Prerequisites for Migration Jobs

For any migration job, ensure that you have completed the following tasks:

- You must have discovered details for the source workload and the target host. See Part IV, “Discovering and Preparing Workloads and Targets,” on page 277.
- Ensure that the credentials for the source workload and target host are valid.

Initiate a Migration Job Using Migrate Client

To start a migration job for a Workload:

1 In the Migrate Client, open the Action window. Use any of the following methods:

- Drag a discovered source and drop it on a discovered target.
- Click a task in the Tasks pane.
- Click the New Job toolbar.
- In the Jobs view, right-click a source and select a command from the context menu. Available commands depend on the type of source.

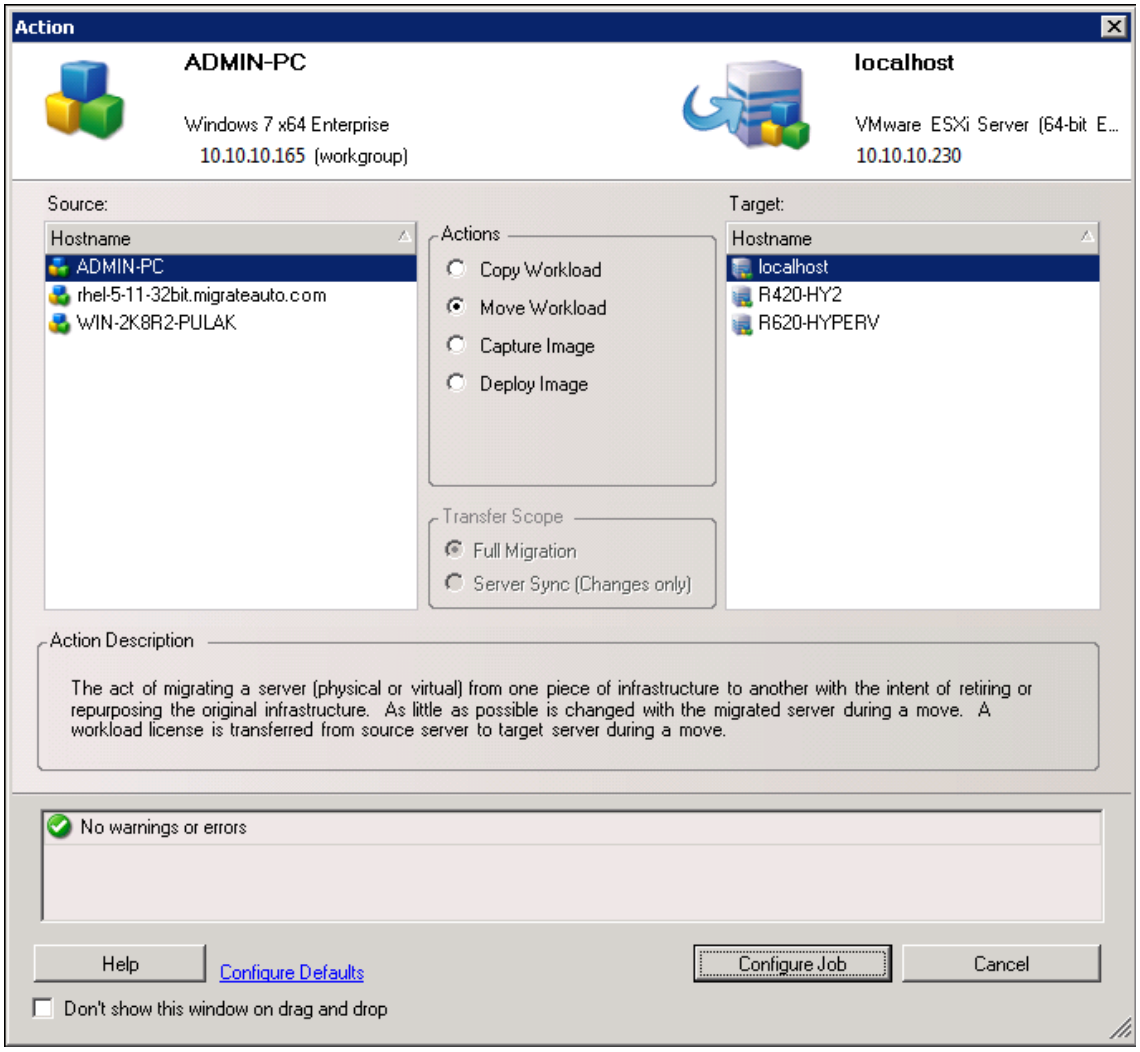

The Source and Target panes display workloads and targets applicable to the selected type of a migration job under **Actions**:

- Copy Workload
- Move Workload

For **Transfer Scope**, the **Full Transfer** and **Server Sync** options are enabled in the following circumstances:

- The system detects an existing operating system on the target
- The operating system profile of the target matches that of the source workload

See "Synchronizing Workloads with Server Sync" on page 527.

2. Check the validation messages at the bottom of the window.
3. To start configuring your migration job, click **Configure Job**.
4. (Optional) For convenience, to avoid displaying the **Action** window on drag-and-drop, select **Don't show this window on drag and drop** before you proceed. Subsequent drag-and-drops actions bypass the Action window and directly open a Conversion Job window.

   To restore the job migration startup behavior, restore the application defaults. See "Configuring General Options" on page 123.
5. Configure the migration as appropriate for the workload and target host.
   - Automated Migration to VMware Using Migrate Client
   - Preparing for Migration of Windows Clusters
   - Migration to Virtual Machines on KVM
   - Migration to Physical Machines
   - Synchronizing Workloads with Server Sync

Initiate a Migration Job Using the Migrate Web Interface

1. In the PlateSpin Migrate Web Interface, click **Workloads**.
2. On the Workloads page, select the workload to migrate.
3. Click **Configure Migration**.
4. Specify the **Initial Transfer Method** for replication based on the scope of data you want to transfer from the source to the target:
   - **Full Replication:** Migrate replicates the full volume from the source to the target.
   - **Incremental Replication:** Migrate replicates only differences in data from the source to the target, provided the workloads have similar operating system and volume profiles.
5. Select a discovered target host, then click **Configure Migration**.
6. Configure the Target Workload Details as appropriate for the workload and target host.
   - Migration to Amazon Web Services
   - Migration to Microsoft Azure
   - Migration to VMware vCloud Director
   - Automated Migration to VMware Using Migrate Web Interface
   - Preparing for Migration of Windows Clusters
7. Click one of the following:
   - **Save & Prepare**
   - **Save**
   - **Cancel**

Saving a Migration Configuration

After you configure a workload for migration, you can save the migration configuration for execution at a later time.

- "Using the Migrate Client" on page 389
- "Using the Migrate Web Interface" on page 389

Using the Migrate Client

To save a migration configuration:

1. Set up a migration job and configure the options.
2. On the Edit Migration Details page, click the arrow at the right side of the **Save** button to expand the Save menu, then select **Save as** or **Save with NTFS Encryption**.

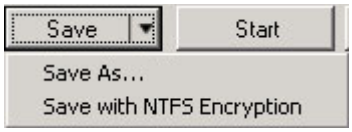

Using the Migrate Web Interface

To save a migration configuration:

1. Set up a migration job and configure the options.
2. Do one of the following:
   - Click **Save & Prepare** to save the migration and to begin preparations for the target VM replication environment on the target host.
   - Click **Save** to save the migration for subsequent changes or later execution.

Editing a Migration Job

You can save an incomplete configuration for a migration job, then add or change settings later.

- "Edit Migration Job Using Migrate Client" on page 389
- "Edit Migration Job Using Migrate Web Interface" on page 390

Edit Migration Job Using Migrate Client

1. In the Jobs view, locate the required job.
2. Open the Migration Job window.
3. Modify the settings as appropriate.
4. Click **OK**.

Edit Migration Job Using Migrate Web Interface

1. On the Workloads page, click the name link of the workload to migrate.
2. On the Migration Details page, click **Edit**.
3. Modify the settings as appropriate.
4. Click **Save**.

Migrate License Key

By default, PlateSpin Migrate automatically selects the best license key for a particular migration job. For information about product licensing and license key management, see "PlateSpin Migrate Product Licensing" on page 93.

License Key in Migrate Client

If you have multiple license keys, PlateSpin Migrate client enables you to select a specific license key to apply to a particular migration job, assuming its workload licenses are available (neither expired nor exhausted). Certain licenses cannot be selected if they are invalid for the current migration. Licenses can be invalid for reasons such as:

- There are no remaining migrations for the license.
- The license does not allow X2V migrations and the current migration is a P2V.

To view or modify the license key selected for a migration job:

1. Start the migration job. For information about starting a migration job, see "Initiating a Migration Job" on page 386.
2. In the Job Configuration section of the Migration Job window, click **License.**
3. To manually choose a different key, deselect **Automatically select the best license key during the conversion**, then select the required license key from the menu.
4. Click **OK**.

   The selected license key is displayed on the **Licenses** tab and the description is updated accordingly.

License Key in Migrate Web Interface

If multiple license keys are available, PlateSpin Migrate Web Interface consumes workload licenses associated with the license keys in order of their start date until all workloads associated with the key are consumed. You cannot specify the key to be used by each workload.

Network Options

Network options are settings for security, performance, and connectivity, and enable you to specify:

- Whether you want the system to compress workload data that is being transferred over the network.

  See “Data Compression” on page 52.

  Fast consumes the least CPU resources on the source but yields a lower compression ratio, **Maximum** consumes the most, but yields a higher compression ratio. **Optimal**, the middle ground, is the recommended option.

- Whether to encrypt the data transferred from source to target.

  See “Security and Privacy” on page 46.

- Whether you want to apply bandwidth throttling for the current migration job.

  See “Bandwidth Throttling” on page 52.

  To throttle this job’s bandwidth to a specified rate, select the **Enable Bandwidth Throttling** option, specify the required maximum value in Mbps, and optionally a time period during which to enforce the throttling. If specified, the **from** and **to** time values are based on the source workload’s system time.

  If no time interval is defined, bandwidth is throttled to the specified rate at all times by default. If time interval is defined and the migration job executes outside this interval, data is transferred at full speed.

- Additional IP addresses for source workloads to enable communication in environments that use network address translation (NAT).

  For information on how to specify additional IP addresses for your PlateSpin Server, see “Migrations Across Public and Private Networks through NAT” on page 60.

Credentials for Source Workloads and Target Hosts

When you configure a migration job, you can validate the provided credentials and save them for future migration jobs that use the same source and target. If you modify the password on the workload or target host, you must also modify the credentials stored in PlateSpin Migrate.

About Credentials

For a migration job to execute properly, you must provide valid credentials for your source and target. For more information about credentials format, see:

- “Discovery Guidelines for Target Hosts” on page 281
- “Discovery Guidelines for Source Workloads” on page 299

Credentials in Migrate Client

To modify source and target credentials:

1. In the Jobs view, select the required workload or target.
2. In the Job Configuration section of the Migration Job window, click **Access.**
3. Specify the credentials.
4. Click **OK**.

Credentials in Migrate Web Interface

To modify target credentials:

1. In the Migrate Web Interface, click **Targets**, then click the target name.
2. On the Target Details page, click **Edit**.
3. On the Edit Target Details page, specify the new user name and password.
4. Click **Save**.

To modify source workload credentials:

1. In the Migrate Web Interface, click **Workloads**, then click the workload name.
2. On the Workload Details page, click **Edit**.
3. On the Edit Target Workload Details page, go to **Migration Settings** > **Source Credentials**.
4. Specify the new user name and password for the source workload.
5. Click **Save**.

Migration Schedule

The migration schedule enables you to specify whether to start the first replication manually or on a specific date and a specific time.

- "Migration Schedule Using Migrate Client" on page 392
- "Migration Schedule Using Migrate Web Interface" on page 393

Migration Schedule Using Migrate Client

To schedule the migration start date and time:

1. In the Jobs view, locate the required job.
2. In the Job Configuration section of the Migration Job window, click **Schedule.**
3. Select **Run at a later time**, then specify the date and start time for the first replication.
4. Click **OK**.

Migration Schedule Using Migrate Web Interface

To schedule the migration start date and time:

1. On the Edit Migration Details page, go to **Schedule Settings** > **Full Replication**, then click **Edit**.
2. Click **Start**, then set the date and time when you want to start the first full replication.

   You can type the date (dd/mm/yyyy) or click the Calendar icon to select the date. The default run time is 12:00:00 a.m. (hh:mm:ss a.m. or p.m.).
3. Click **Close** to return to the Edit Migration Details page.
4. Click **Save**.

Blackout Window for Data Transfer

The blackout window suspends scheduled replications from starting during a specified period of time and pattern. It helps you to reserve network bandwidth for users or mission critical communications during peak traffic periods. You can also use it to prevent conflicts for other data backup or snapshot activities. For example, suspend replications during peak network utilization hours or to prevent conflicts between VSS-aware software and the PlateSpin VSS block-level data transfer component.

The default setting is None. No blackout window is scheduled.

NOTE: The blackout start and end times are based on the system clock on the PlateSpin Server.

- "Blackout Window Using the Migrate Client" on page 393
- "Blackout Window Using the Migrate Web Interface" on page 393

Blackout Window Using the Migrate Client

PlateSpin Migrate Client does not provide and option to configure a blackout window for data transfer.

Blackout Window Using the Migrate Web Interface

To set or modify a blackout window:

1. On the Edit Migration Details page, go to **Schedule Settings** > **Blackout Window**, then click **Edit**.
2. Specify the start and end time for the blackout period.

   The blackout start and end times are based on the system clock on the PlateSpin Server.
3. Select **Daily**, **Weekly**, or **Monthly** to enable a blackout window, then set the recurrence pattern.
4. Click **Close** to return to the Edit Migration Details page.
5. Click **Save**.

Compression during Data Transfer

The Compression Level setting controls whether data is compressed during transmission between the source and target workloads, and the level of data compression applied.See “Data Compression” on page 52.

Select one of the following options:

- **None:** No compression.
- **Fast:** Consumes the least CPU resources on the source, but yields a lower compression ratio.
- **Optimal:** (Default) Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option.
- **Maximum:** Consumes the most CPU resources on the source, but yields a higher compression ratio.

- “Compression Using Migrate Client” on page 394
- “Compression Using Migrate Web Interface” on page 394

Compression Using Migrate Client

To enable and use compression for data transfer:

1. In the Jobs view, locate the required job.
2. In the Network Configuration section of the Migration Job window, select **Enable Compression.**
3. Specify the appropriate compression level: **Fast**, **Optimal**, or **Maximum**.
4. Click **OK**.

Compression Using Migrate Web Interface

To enable and use compression for data transfer:

1. On the Edit Migration Details page, go to **Schedule Settings** > **Compression Level**.
2. Specify the appropriate compression level: **Fast**, **Optimal**, or **Maximum**.
3. Click **Save**.

Bandwidth Throttling during Data Transfer

Bandwidth throttling enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. Throttling helps to prevent migration traffic from congesting your production network and to reduce the overall load of your PlateSpin Server. You can specify a throughput rate for each migration job. See “Bandwidth Throttling” on page 52.

NOTE: Throttling time is local to the source workload.

- "Bandwidth Throttling Using Migrate Client" on page 395
- "Bandwidth Throttling Using Migrate Web Interface" on page 395

Bandwidth Throttling Using Migrate Client

To enable and use bandwidth throttling for data transfer:

1. In the Jobs view, locate the required job.
2. In the Network Configuration section of the Migration Job window, view **Bandwidth Throttling.**
3. Select the **Enable Throttling** option, specify the required maximum value in Mbps, and optionally a time period during which to enforce the throttling.

   If no time interval is defined, bandwidth is throttled to the specified rate at all times by default. If time interval is defined and the migration job executes outside this interval, data is transferred at full speed.
4. Click **OK**.

Bandwidth Throttling Using Migrate Web Interface

To enable and use bandwidth throttling for data transfer:

1. On the Edit Migration Details page, go to **Schedule Settings** > **Bandwidth Throttling**.
2. Specify the maximum bandwidth to consume in Mbps as the **Throttling Rate**.

   A value of Off disables bandwidth throttling.
3. Specify one of the following throttling patterns:
   - **Always:** Always throttle data transfer for the replications. No throttling pattern is specified.
   - **Custom:** Specify the start and end time and days of the week to throttle data transfer for the replications running in that window.
4. Click **Save**.

Conversion (Data Transfer Method)

Conversion or Data Transfer options enable you to specify:

- How data is transferred from source to target.

  PlateSpin Migrate supports file-level data transfer for Windows workloads and block-level data transfer for Windows and Linux workloads. The file-level method enables you to specify the free space available for the target volume when you configure the migration. See "Supported Data Transfer Methods" on page 44.
- The scope of workload data to transfer from the source to the target.

  PlateSpin Migrate Client supports **Full Migration** for any workload and **Server Sync (Changes only)** for Server Sync jobs. See "Synchronizing Workloads with Server Sync" on page 527.

Conversion Using Migrate Client

To specify the transfer options for a migration job:

1. In the Jobs view, locate the required job.
2. In the Job Configuration section of the Migration Job window, click **Conversion.**
3. Select the scope and method of data transfer.
4. Click **OK**.

Data Transfer Using Migrate Web Interface

1. On the Edit Migration Details page, go to **Migration Settings** > **Transfer Method**.
2. Specify the appropriate data transfer method.
3. Click **Save**.

Encrypt Data Transfer

The Encrypt Data Transfer option determines whether to encrypt the data for transmission from the source workload to the target workload.See "Security and Privacy" on page 46.

- "Encrypt Data Transfer Using Migrate Client" on page 396
- "Encrypt Data Transfer Using Migrate Web Interface" on page 396

Encrypt Data Transfer Using Migrate Client

To enable and use encryption for data transfer:

1. In the Jobs view, locate the required job.
2. In the Network Configuration section of the Migration Job window, click **Encryption.**
3. Select **Encrypt Data Transfer**.
4. Click **OK**.

Encrypt Data Transfer Using Migrate Web Interface

To enable and use encryption for data transfer for Windows workloads:

1. On the Edit Migration Details page, go to **Migration Settings** > **Data Transfer**.
2. Select **Encrypt Data Transfer**.
3. Click **Save**.

To enable and use encryption for data transfer for Linux workloads:

1. On the Edit Migration Details page, go to **Migration Settings** > **Transfer Encryption**.
2. Select **Encrypt Data Transfer**.
3. Click **Save**.

Virtualization Enhancement Software

For migrations between different virtualization hosts, PlateSpin Migrate provides a mechanism to automatically uninstall virtualization enhancement software, such as VMware Tools.

When converting a workload on a VMware platform that has an earlier version of VMware Tools installed, PlateSpin Migrate identifies the presence of obsolete software and adds a VMware Tools Cleanup step in the migration job.

You must provide administrator credentials to uninstall VMware Tools. The credentials provided must match the administrator-level user account that was logged in during the installation of VMware Tools.

When the earlier version is uninstalled, PlateSpin Migrate proceeds with the installation of the new version of VMware Tools.

NOTE: If you are downgrading a virtual machine that has VMware Tools installed, or if you are converting a virtual machine to another VMware target that has an older version of VMware Tools, the installation of VMware Tools during the configuration of the target will fail.

Replace VMware Tools using Migrate Client

To configure a job to remove or replace VMware Tools during the migration:

1. In the Jobs view, select the required workload.
2. In the Operating System and Application Configuration section of the Migration Job window, click **Clean up VMware Tools**.

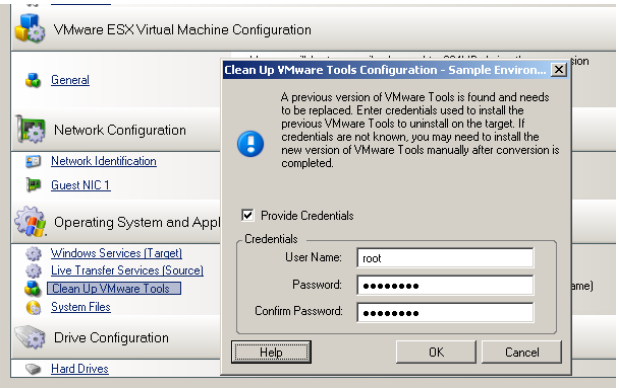

3. Depending on the target, PlateSpin Migrate identifies existing instances of VMware Tools and prompts to either replace or remove them, as applicable:
   - **For non-VMware targets:** The job configuration interface prompts you to uninstall VMware Tools. Provide the same administrator-level credentials used to install the software. If the credentials are unknown, VMware Tools remains on the target machine after migration.

- **For VMware targets:** The job configuration interface prompts you to replace VMware Tools. Provide the same administrator-level credentials used to install the obsolete version of VMware Tools. If the credentials are unknown, install the new version of VMware Tools manually after the migration completes.

4 Click **OK**.

Replace VMware Tools using Migrate Web Interface

To remove or replace VMware Tools during a migration:

1 On the Edit Target Workload Details page, go to **Target Workload Settings** > **VM Tools**.

2 To install the VM tools, select the **Install VM Tools option**. This option is selected by default.

3 On the Edit Target Workload Details page, go to **Target Workload Test Settings** > **VM Tools**.

4 To install the VM tools, select the **Install VM Tools option**. This option is selected by default.

5 Click **Save**.

Custom Post-Migration Actions

PlateSpin Migrate enables you to execute a custom action on your target. You must define and save your custom actions and their dependencies in advance. See "Managing Post-Migration Actions" on page 107.

You can then associate this custom action with migration jobs you configure using the PlateSpin Migrate Client or the PlateSpin Migrate Web Interface. Towards the end of the migration process, PlateSpin Migrate uploads the specified action and its dependencies to the target and executes the action on the target.

IMPORTANT: Consider that you have associated a custom post-migration action with a migration job and have run a full replication job. Before you perform a test cutover or cutover of the replicated workload, you must perform an incremental replication of the workload in the following scenarios:

- The custom action associated with the replicated workload is either deleted or modified after a full replication of the workload.
- The migration job is reconfigured to execute a different custom action on the workload.

Specify Custom Action Using Migrate Client

To specify a custom post-migration action for your migration job:

1 Start the migration job. For information about starting a migration job, see "Initiating a Migration Job" on page 386.

2 In the Virtual Machine Configuration section of the Migration Job window, click **Post-Migration.**

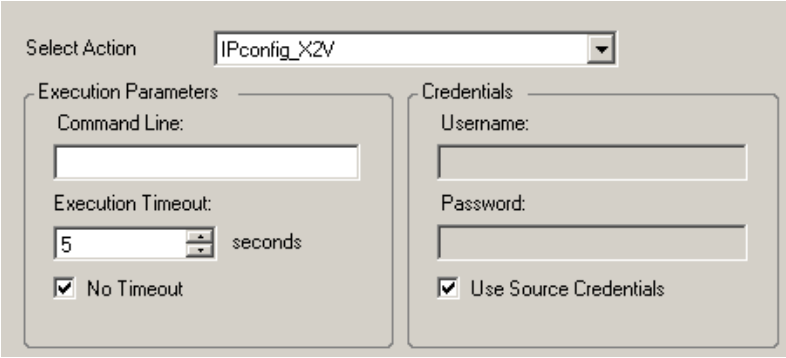

3 Specify the following options:

- **Select Action:** From the drop-down list, select a custom action previously saved in your library of post-migration actions.
- **Execution Parameters:** By default, it displays the command line parameters and timeout value that you specified when you added this action to the PlateSpin Server library. You can edit the displayed command line parameters and the timeout value.
- **Credentials:** Provide administrator credentials for the target. If they are the same as those for the source, and if they have been saved, select **Use Source Credentials**.

Specify Custom Action Using Migrate Web Interface

1 On the Edit Migration Details page, go to **Target Workload Settings** > **Post-Migration Action**.

2 Click **Configure Action**.

3 Specify the following options:

- **Action Name:** Specify a name for the action.
- **Execution Parameters:** Specify the following execution parameters for the action:
  - **Command Line:** By default, it displays the command line parameters that you specified when you added this action to the PlateSpin Server library. You can edit the displayed command line parameters if required.
  - **Execution Timeout:** By default, it displays the timeout value that you specified when you added this action to the PlateSpin Server library. You can edit the timeout value if required.

    If a timeout value is not specified, then PlateSpin Migrate by default waits for 7200 seconds (120 mins) before timing out the action. However, the value that you specify for the timeout takes precedence over this default value of 7200 seconds.

4 Click **Apply**.

For non-cloud target platforms (VMware), the post-migration action that you configure for the Target Workload Settings is automatically applied for the Target Workload Test Settings.

5 (For cloud target platforms) Go to **Target Workload Test Settings** > **Post-Migration Action** and perform Step 2 through Step 4 to configure the action.

Services or Daemons to Stop before Replication or Cutover

Because source servers are active during migrations, PlateSpin Migrate provides a mechanism to stop selected services or daemons during the migration. This ensures that data on your source is captured in a consistent state.

If your source workload is running Microsoft SQL Server or Microsoft Exchange Server software, you can configure your migration job to automatically copy the database files of these servers. If you do not require the migration to include the volume containing the databases, consider not stopping these services.

If your source workload includes I/O-intensive application services that might inhibit the ability of the file transfer process to keep up with the changes, consider stopping them during the migration.

After the completion of the migration, services that you select to stop are automatically restarted on the source, unless you explicitly configure your migration job to power off the source on completion.

For Linux systems, consider using the custom `freeze` and `thaw` scripting capability. See "Using Custom Freeze and Thaw Scripts for Linux Block-Level Migrations" on page 324.

TIP: You can globally configure your preferences for stopping selected Windows services during VSS File-based or VSS Block-based Transfer performed using the PlateSpin Migrate Client. See "Configuring Source Service Defaults" on page 125.

Services and Daemons to Stop Using Migrate Client

To specify which services or daemons you want the system to stop during data transfer:

1. In the Jobs view, select the required workload.
2. In the Operating System and Application Configuration section of the Migration Job window, click **Live Transfer Services/Daemons (Source)**.
3. To indicate that you want SQL Server and Exchange Server database files copied during the migration, click **Advanced** (applicable to Windows systems only).

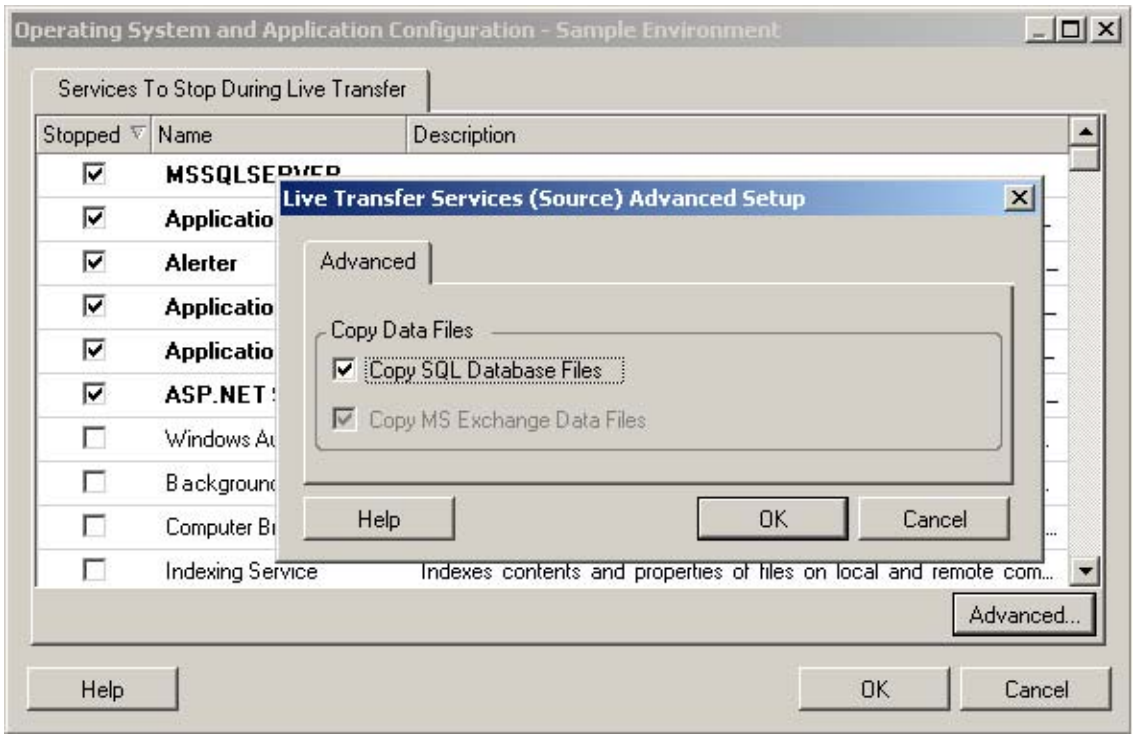

4 Click **OK**.

Services and Daemons to Stop using Migrate Web Interface

To stop Windows services:

1 On the Edit Target Workload Details page, go to **Migration Settings** > **Services to Stop before Any Replication**.

2 Select the services to stop for replication.

We recommend that all the non-VSS compliant services or antivirus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.

3 On the Edit Target Workload Details page, go to **Migration Settings** > **Services to Stop before Cutover with Replication**.

4 Select the Windows services that should be permanently stopped on the source workload for cutover with any replication. The services stopped on the source workload during the replication process are not restored after Cutover. The stopped services are restored after a Test Cutover.

5 Click **Save**.

To stop Linux Daemons:

1 On the Edit Target Workload Details page, go to **Migration Settings** > **Daemons to Stop before Any Replication**.

2 Select the Linux daemons that you want to be temporarily stopped on the source workload before replication. These daemons will be restored after replication completes.

3 On the Edit Target Workload Details page, go to **Migration Settings** > **Daemons to Stop before Cutover with Replication**.

4 Select the Linux daemons that should be permanently stopped on the source workload for Cutover with any Replication. The daemons stopped on the source workload during the replication process are not restored after Cutover. The stopped daemons are restored after a Test Cutover.

5 Click **Save**.

Service States on Target Windows Workloads

In scenarios such as the following, you might want to change the start-up mode of the services on target Windows workloads:

- If you do not want a certain Windows service to continue running on a virtualized workload, then configure the job to disable the service on the target workload.
- If you require that a service on the target starts based on a request from some other service, you can set the start-up mode of the required service to manual.
- If you want to configure a job to restore the original start-up mode of the service post the migration. For example, you might want to disable a virus scanner during the migration, but restore the start-up mode of the scanner after the migration completes.
- Some applications on a source workload are known to cause boot failure on the target workload if the corresponding application services are not disabled during the conversion. The **ApplicationsKnownForBootFailuresOnTarget** parameter on the PlateSpin Server Configuration page lists such applications that are likely to cause boot failure on target workload. You can edit this list to add or remove the applications from the list.

  A global setting, **ApplicationsKnownForBootFailuresOnTargetDefaultValue,** on the PlateSpin Server Configuration page sets whether the services of all such applications listed in the **ApplicationsKnownForBootFailuresOnTarget** parameter must be selected by default so that the corresponding application services can be disabled on the target during the conversion.

  For information about the configuring applications known to cause boot failure on Windows target, see “Configuring Applications Known to Cause Boot Failure on Windows Target Workloads” on page 116.

For information on modifying or disabling the service state on the target, review the following sections:

- “Service States using Migrate Client” on page 402
- “Service States using Migrate Web Interface” on page 404

Service States using Migrate Client

You can specify the preferred run states for services on target Windows workloads that will be enabled after cutover or test cutover. Windows service states options are:

- Automatic

- Automatic (Delayed Start)
- Manual
- Disabled

Modifying the Windows Service State on the Target Post Migration

To configure post-migration startup mode of Windows services:

1. Start the migration job. For information about starting a migration job, see "Initiating a Migration Job" on page 386.
2. In the Operating System and Application Configuration section of the Migration Job window, click **Windows Services (Target)** and then click an item in the **Start Mode** column.

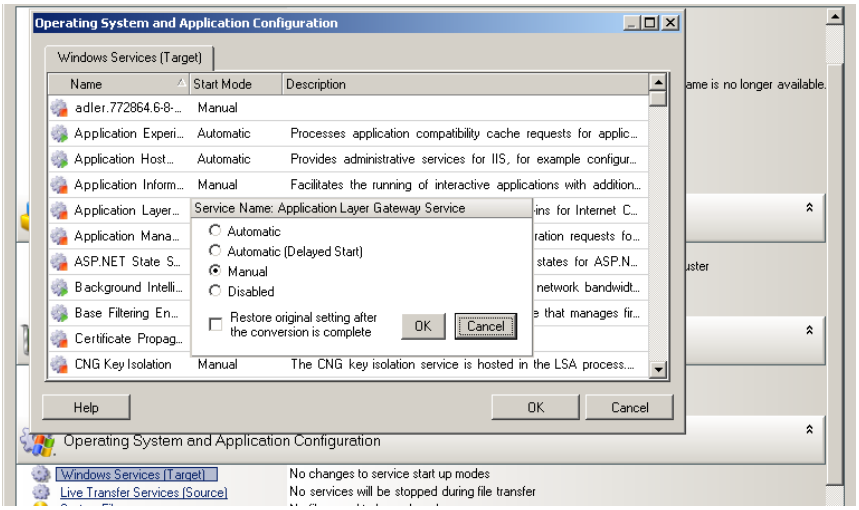

3. Select the desired startup mode.
4. To restore the original setting after conversion is complete, select the check box.
5. Click **OK**.

Disabling the Windows Boot Service State on the Target Post Migration

1. Start the migration job. For information about starting a migration job, see "Initiating a Migration Job" on page 386.
2. In the Operating System and Application Configuration section of the Migration Job window, click **Windows Services (Target)** and then click **More Options**.

   PlateSpin Migrate reviews the existing applications on the source to check if any of the applications listed in the **ApplicationsKnownForBootFailuresOnTarget** configuration parameter is installed on the source. PlateSpin Migrate lists all such applications, which are known to cause boot failure on the target during conversion in the Application Known For Boot Failure panel.

   These applications are selected by default if the value of the **ApplicationsKnownForBootFailuresOnTargetDefaultValue** parameter on the PlateSpin Configuration page is set to **true**.

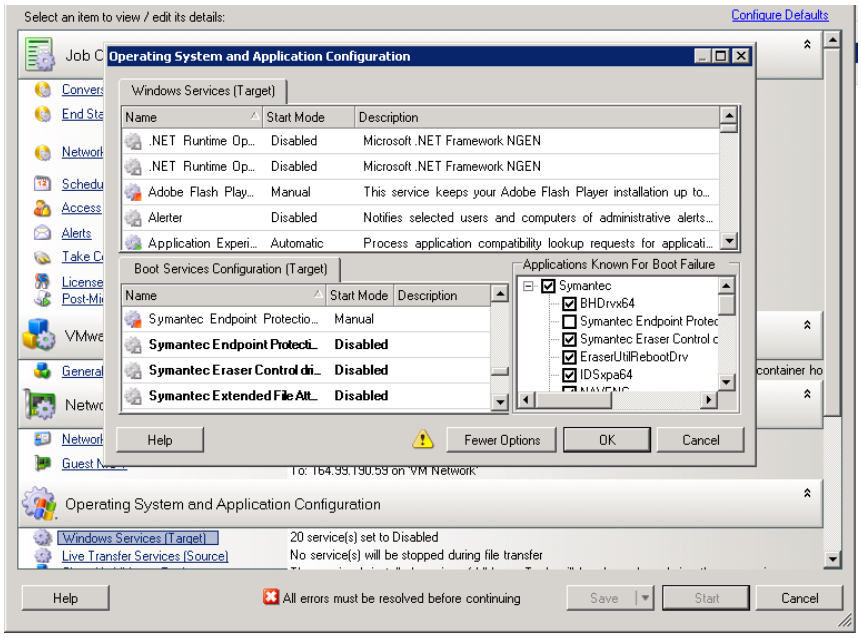

3 Modify the selection of the applications in the Application Known For Boot Failure panel depending on whether or not you want to disable the boot services of the applications on the target. Selecting an application sets the start-up mode of the corresponding boot service on the target as **Disabled**.

4 In the Boot Services Configuration (Target) panel, review the modified boot services configuration. Ensure that the settings are correctly configured to prevent any operating system issues.

5 Click **OK**.

Service States using Migrate Web Interface

You can specify the preferred run states for services on target Windows workloads that will be enabled after cutover or test cutover. Windows service states options are:

- Automatic
- Manual
- Disabled
- Automatic (Delayed Start)
- Boot
- System

Modifying the Windows Service State on the Target Post Migration

1 On the Edit Migration Details page, go to **Target Workload Settings** > **Service States on Target VM**.

2 Click **Add Services**.

3 Select the start-up mode of the Windows service on the target VM.

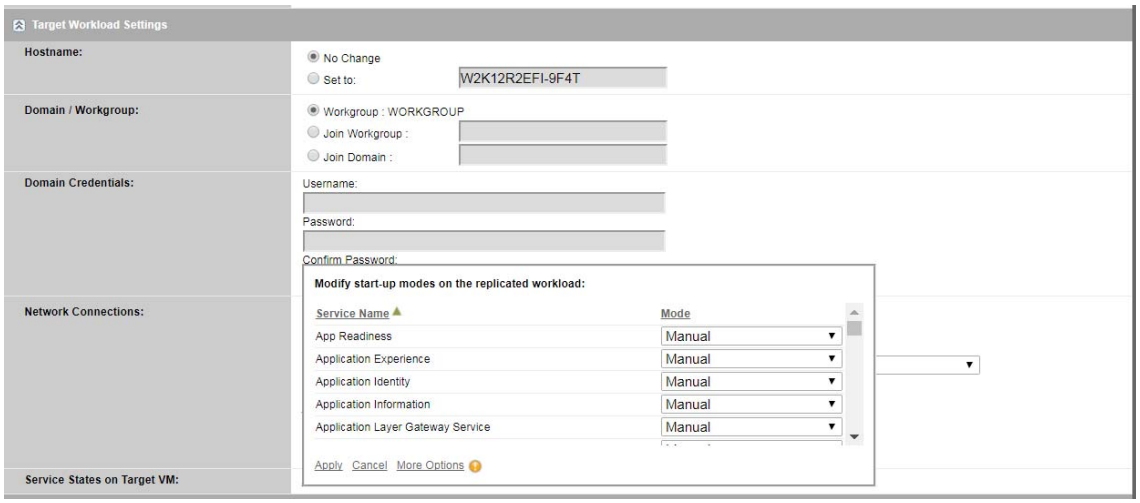

4 Click **Apply**.

Disabling the Windows Boot Service State on the Target Post Migration

1 On the Edit Migration Details, go to **Migration Settings** > **Boot Services to Disable on Target**.

2 Click **Add Services**.

PlateSpin Migrate reviews the existing applications on the source to check if any of the applications listed in the **ApplicationsKnownForBootFailuresOnTarget** configuration parameter is installed on the source. PlateSpin Migrate lists all such applications, which are known to cause boot failure on the target during conversion in the Application Known For Boot Failure panel.

These applications are selected by default if the value of the **ApplicationsKnownForBootFailuresOnTargetDefaultValue** parameter on the PlateSpin Configuration page is set to **true**.

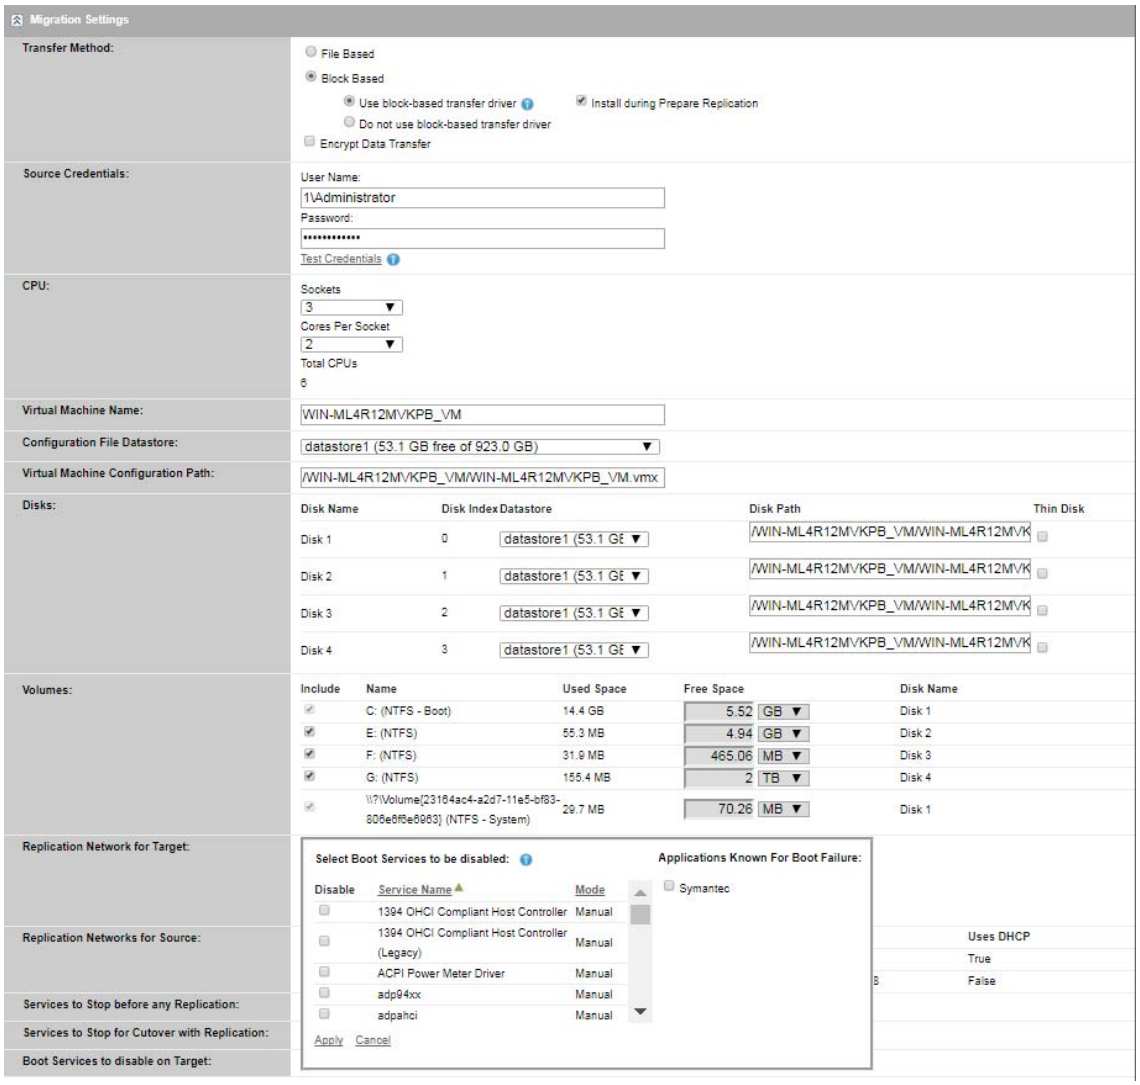

3 Modify the selection of the applications in the Application Known For Boot Failure panel depending on whether or not you want to disable the boot services of the applications on the target. Selecting an application sets the start-up mode of the corresponding boot service on the target as **Disabled**.

4 In the Select Boot Services to be disabled panel, review the modified boot services configuration. Ensure that the settings are correctly configured to prevent any operating system issues.

5 Click **Apply**.

Daemon States on Target Linux Workloads

You can specify the preferred run states for daemons on target Linux workloads that will be enabled after cutover or test cutover. Linux daemons states options are enabled or disabled at the following runlevels and system boot:

0	Shutdown

1	Single-user mode
2	Unused (user-defined)
3	Full multi user-mode (no GUI)
4	Unused (user-defined)
5	Full multi-user mode with display manager (GUI)
6	Reboot
Boot	Start at power on

Daemon States using Migrate Client

To configure the post-migration run level of Linux daemons:

1. Start the migration job. For information about starting a migration job, see "Initiating a Migration Job" on page 386.
2. In the Operating System and Application Configuration section of the Migration Job window, click **Linux Daemons (Target)**, and then click an item in the **Run Level** column

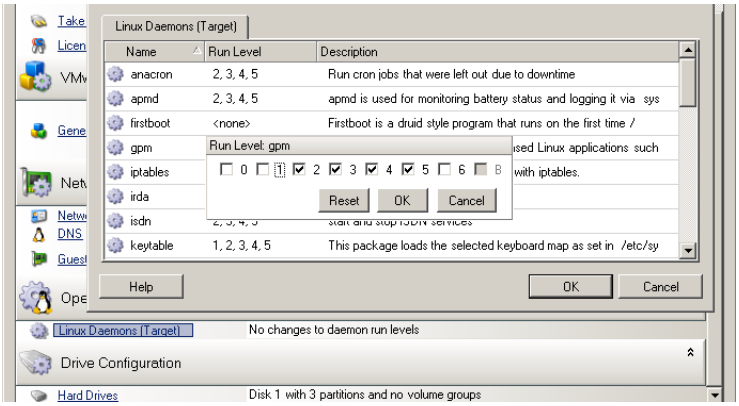

3. Select the desired run levels. Click **OK**.

Daemon States using Migrate Web Interface

To set start states for Linux daemons on the target VM:

1. On the Edit Target Workload Details page, go to **Target Workload Settings** > **Daemon States on Target VM**.
2. Select Linux daemons' start conditions on the target VM. Enable the daemon to start by selecting the check boxes at the appropriate runlevels (0 to 6) and Boot.
3. Click **Save**.

Windows HAL or Kernel File Replacements

When you use PlateSpin Migrate Client to migrate Windows workloads with system files (such as a HAL or kernel files) that are incompatible with the target infrastructure, PlateSpin Migrate uses an appropriate file from its library and saves a backup copy of the source file (`*.bak`) on the target, in the same system directory.

You can use Migrate Client to view the HAL or kernel files that PlateSpin Migrate identifies as those requiring replacement.

To view the files selected for replacement during migration:

1. In the Jobs view, select the required workload.
2. In the Operating System and Application Configuration section of the Migration Job window, click **System Files**.

   Files selected for replacement during migration are listed.

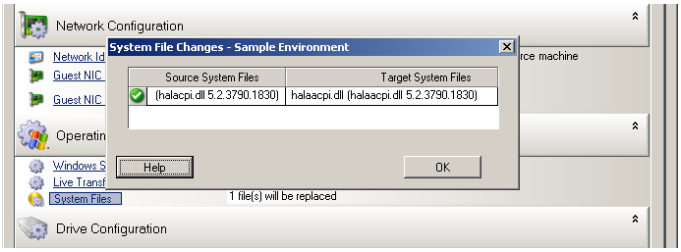

3. Click **OK**.

The following warnings might display at the bottom of the dialog box:

Driver Cache is empty	Indicates that you might need to place the necessary files into the local driver cache on the source Windows server (`..\Windows\Driver Cache`).
The driver cache contains a higher version	PlateSpin Migrate has a partial match with its matrix but the driver cache contains a later version of one or more system files than the one that PlateSpin Migrate will use.
File <filename> will be replaced with lower version	PlateSpin Migrate has not found a match for the system files in its matrix. It will replace the system files with a version that is earlier than the ones that were discovered as the source machine's original system files.
File <filename> will be replaced with higher version	PlateSpin Migrate has not found a match for the system files in its matrix. It will replace the system files with a version that is later than the ones that were discovered as the source machine's original system files.

If warnings appear on the screen, click **More Help** (only available if warnings exist) to learn more.

See also KB Article 7920815 *FAQ: Understanding the System Files Information Screen* (https://support.microfocus.com/kb/doc.php?id=7920815).

Post-Cutover End States for Source and Target Workloads

After a successful cutover, PlateSpin Migrate shuts down or starts the source workload and target workload, depending on the nature of the migration. For example, if the migration goal is to copy the workload, you might want both the source and target workload to be running after cutover. If you are moving a workload, you might want to stop the source workload after cutover and leave the target workload running.

- "Workload End States Using the Migrate Client" on page 409
- "Workload End States Using the Migrate Web Interface" on page 409

Workload End States Using the Migrate Client

To specify non-default post-cutover end states for your source and target:

1. In the Jobs view, select the required workload.
2. In the Job Configuration section of the Migration Job window, click **End States.**
3. Configure the appropriate settings:
   - **Source Machine End State:** Specify whether to shut down the source workload after a successful cutover. For a workload move, the shut down is selected by default.
   - **Target Machine End State:** Specify whether to power on, power off, or suspend the target workload after a successful cutover.
4. Click **OK**.

Workload End States Using the Migrate Web Interface

To specify post-cutover end states for the source and target workloads after a cutover with replication:

1. On the Workloads page, select the prepared workload that you want to migrate.
2. Click **Run Migration**.
3. On the Workload Commands page, specify the full or incremental replication method.
4. For Post-Replication Cutover, enable **Run cutover after successful replication**.
5. Specify the appropriate run state for the source and target workload by enabling or disabling the following settings:
   - Shut down source after cutover
   - Shut down target after cutover
6. Click **Execute**.

   PlateSpin Migrate starts the replication for the workload, executes the cutover, then shuts down the source or target as configured.

NOTE: For workloads in the cloud, the end-state shutdown stops the VM and deallocates its resources.

Target Workload Settings for VMs

For jobs that involve workload virtualization, PlateSpin Migrate provides a mechanism for specifying target VM configuration options, such as providing a target VM name and a configuration file path, selecting a datastore to use, and allocating virtual memory, in accordance with the features and capabilities of the selected virtualization platform.

If you have resource pools configured on your target virtualization platform, you can select a resource pool for your VM to be assigned to.

NOTE: If your target VMware ESX server is part of a fully automated Distributed Resource Scheduler (DRS) cluster (a cluster with its VM migration automation level set to **Fully Automated)**, the newly created target VM's automation level is changed to **Partially Automated** for the duration of the migration. This means that your target VM might power up on a different ESX server from the one initially selected, but migration is prevented from automatic execution.

- "Target VM Configuration in Migrate Client" on page 410
- "Target VM Configuration in Migrate Web Interface" on page 410

Target VM Configuration in Migrate Client

To modify target VM configuration options:

1. In the Jobs view, select the required workload.
2. In the Virtual Machine Configuration section of the Migration Job window, click **General**.
3. Specify the values for the configuration options and click **OK**.

PlateSpin Migrate displays target virtual machine configuration options specific to the selected target and also provides access to advanced configuration options. See "Target VM Configuration: VMware ESXi" on page 496.

Target VM Configuration in Migrate Web Interface

Migrate Web Interface displays target virtual machine configuration options specific to the selected target. You can specify different values as needed for the target workload test settings.

1. On the Edit Target Workload Details page, go to **Target Workload Settings**.
2. Modify the target VM settings as appropriate for the target platform:
   - **AWS:** Target Workload Settings
   - **Azure:** Target Workload Settings
   - **vCloud:** Target Workload Settings
   - **VMware:** Target Workload Settings
3. (Optional) Go to **Target Workload Test Settings**, then modify the target VM test settings as appropriate for the target platform:
   - **AWS:** Target Workload Settings
   - **Azure:** Target Workload Test Settings

- **vCloud:** Target Workload Test Settings
- **VMware:** Target Workload Test Settings

4 Click **Save**.

Network Identification (Network Connections)

PlateSpin Migrate enables you to manage the network identity and domain registration of your migration target workload and specify related preferences as part of a migration job. By default, a job is configured to preserve a source workload's network identity and domain registration. You can modify the default configuration to suit the objectives of your migration job.

Proper configuration of migration target's network identity is especially important when you are migrating a workload to a different domain, planning to take it off a domain, or if you intend to change the host name of a workload while it is in the domain.

- "Network Identification Using Migrate Client" on page 411
- "Network Connections Using Migrate Web Interface" on page 413

Network Identification Using Migrate Client

To configure a target workload's network identity options:

1 In the Jobs view, select the required workload.

2 In the Network Configuration section of the Migration Job window, click **Network Identification.**

3 Specify the options and then click **OK**.

Configuration options vary depending on whether the target machine is Windows or Linux. For information about the configuration options, see the following sections:

- "Managing the Identity of Windows Workloads" on page 411
- "Managing the Network Identity of Linux Workloads" on page 413

Managing the Identity of Windows Workloads

Use these settings to configure the network identity of your target Windows workload.

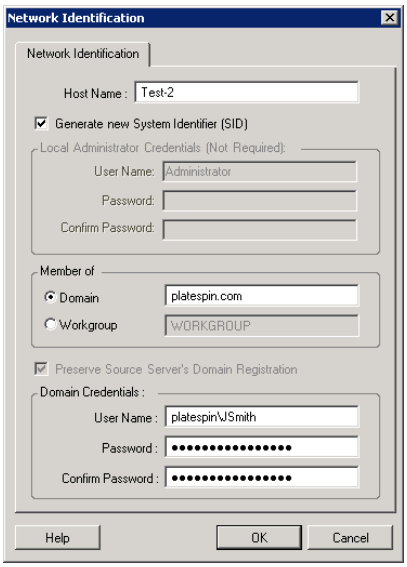

Host Name: Specify the desired host name for the target machine.

Member of (Domain / Workgroup): Select the required option and type the name of the domain or workgroup that you want the target machine to join.

Preserve Source Server's Domain Registration: Preserves domain registration and ensures that the source server domain registration remains intact during migration. If you disable this option, the source machine's domain account is transferred to the target machine. The source server still appears to be on the domain, but does not have a valid connection.

Domain Credentials: If the target machine is to be part of a domain, specify valid credentials for a user account with permission to add servers to the domain, such as a member of the Domain Admins group or Enterprise Admins group.

Managing the Network Identity of Linux Workloads

Use these settings to configure the network identity of your target Linux workload and DNS server addresses as required.

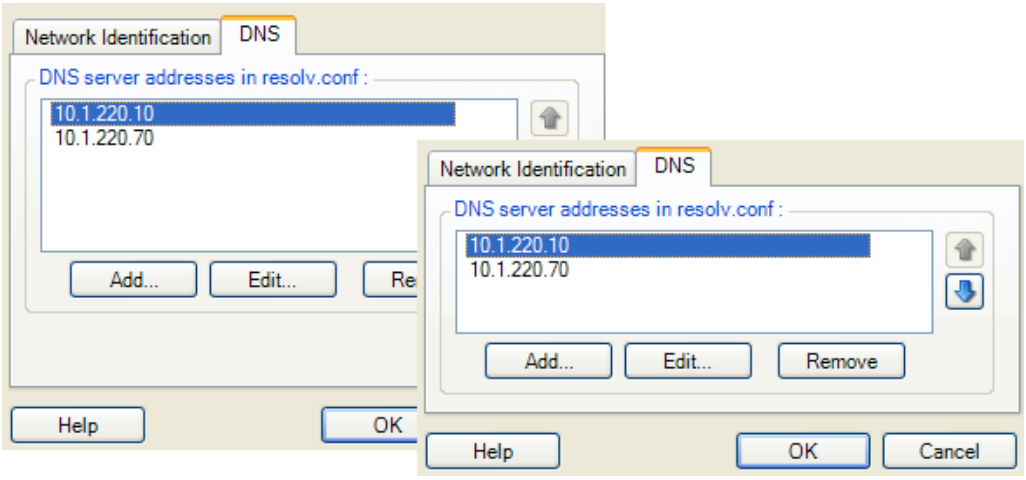

Network Identification tab: Specify the desired host name for the target server.

DNS tab: Use the **Add**, **Edit**, and **Remove** buttons to manage DNS server entries for the new virtual machine.

Network Connections Using Migrate Web Interface

Migrate Web Interface displays target network configuration options specific to the selected target. You can specify different network values as needed for the target workload test settings.

1. On the Edit Target Workload Details page, go to **Target Workload Settings** > **Network Connections**
2. Modify the Network Connections settings as appropriate for the target workload on the target platform:

Parameter	Description
IP Address	Specify DHCP, or select Static and provide an IP address for each network connection. For Windows workloads, click **Advanced** to specify multiple IP addresses for the NIC. For target AWS and Azure cloud platforms, the network mask and gateway address is auto-populated based on the subnet selected. For VMware or vCloud, specify the appropriate network mask for each IP address. You can optionally specify IP addresses for multiple default gatways.
DNS Servers	If you choose static, specify information about your DNS servers.

- **AWS:** Target Workload Settings > Network Connection
- **Azure:** Target Workload Settings > Network Connections

  For Azure, configure these additional settings:

Parameter	Description
Include	If the workload has multiple NICs, select **Include** for each NIC to be migrated. ◆ At least one NIC is required. ◆ The number of NICs to migrate cannot exceed the maximum number of NICs supported by the selected cloud instance. The available NICs apply to the NICs in Target Workload Test Settings.
Network and Subnet	For each NIC, specify the network to use and a subnet in that network.
Primary Connection	If you have multiple NICs, specify one of the included NICs to use as the primary connection. The default **Primary Connection** is the first NIC in the list.
Public IP	If you do not use an Azure VPN, the primary NIC requires a public IP address that is automatically assigned by a Azure.
Resource Group	Type or select a resource group to use for the NIC. The **Azure Resource Group** setting is the default.

- **vCloud:** Target Workload Settings > Network Connection
- **VMware:** Target Workload Settings > Network Connections

3 (Optional) Go to **Target Workload Test Settings** > **Network Connections**, then modify the target VM test settings as appropriate for the target platform:

- **AWS:** Target Workload Test Settings > Network Connection
- **Azure:** Target Workload Test Settings > Network Connections
- **vCloud:** Target Workload Test Settings > Network Connections
- **VMware:** Target Workload Test Settings > Network Connections

4 Click **Save**.

Migration Network (Replication Network)

For each workload migration job, you must properly configure workload networking to enable communications between the source workloads and the target workloads or PlateSpin Replication Environment during the migration process. The network configuration of a target workload must be appropriate for its end state.

- "Migration Network Using Migrate Client" on page 415
- "Replication Network Using Migrate Web Interface" on page 420

Migration Network Using Migrate Client

Temporary Networking: Also called *Take Control Network Settings*; they apply to target workloads booted into a temporary pre-execution environment.

- "Temporary (Take Control) Network Settings" on page 415
- "TCP/IP and Advanced Network Settings" on page 419

Temporary (Take Control) Network Settings

Temporary (Take Control) Network Settings control how source workloads, target workloads, and the PlateSpin Server communicate among each other during the migration. If required, you can manually specify a temporary network address to your target workload, or configure it to use a DHCP-assigned IP address during the migration.

During Windows and Linux workload migrations, the Temporary Network Settings control the PlateSpin Server's communication with the target workloads that are booted into a temporary pre-execution environment.

To configure Temporary (Take Control) network settings:

1. Start the migration job. For information about starting a migration job, see "Initiating a Migration Job" on page 386.
2. In the Job Configuration section of the Migration Job window, click **Take Control**.
3. To access network interface mapping and TCP/IP settings, click **Configure** in the target area.
4. Click **OK**.

Configuration options for the Temporary networking vary and depend on whether the network interface is virtual or physical, and whether it is connecting a Windows or a Linux workload.

- "Temporary (Take Control) Network Settings: Physical Network Interfaces" on page 416
- "Temporary (Take Control) Network Settings: Virtual Network Interfaces" on page 417
- "Target Post-Migration Networking" on page 417

Temporary (Take Control) Network Settings: Physical Network Interfaces

These settings apply only to source physical machines. For target physical machines, Temporary (Take Control) network settings are configured during the boot process that uses the PlateSpin ISO image. See "Registering and Discovering Details for Target Physical Machines with PlateSpin ISO" on page 291.

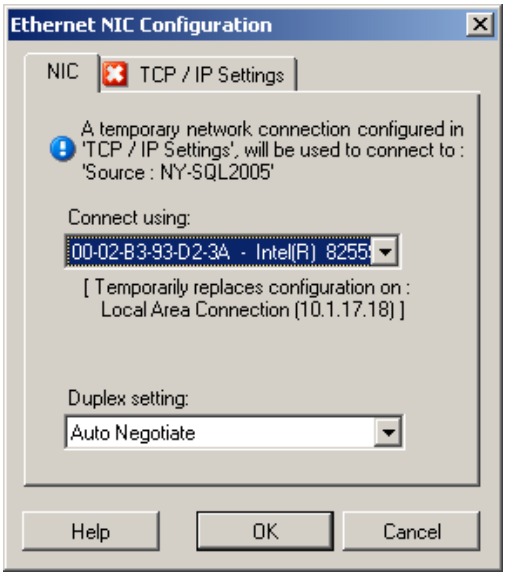

Connect using: If multiple network adapters are present, select the adapter that can communicate with both the PlateSpin Server and the target.

Duplex setting: Use the drop-down list to select network card duplexing. It must match the duplex setting for the switch to which the network interface is connected. When the source is connected to switch ports that are set to 100 Mbps full duplex and cannot be changed to auto negotiation, select **Force NIC to Full Duplex**.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See "TCP/IP and Advanced Network Settings" on page 419.

Temporary (Take Control) Network Settings: Virtual Network Interfaces

These settings apply to target Take Control network settings.

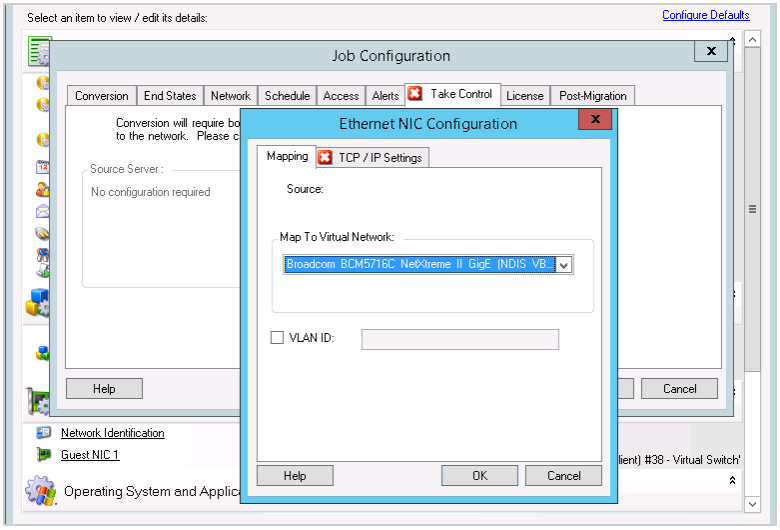

Map to Virtual Network: From the drop-down list, select the virtual switch or network to use for communications. If multiple virtual network adapters are present, select the adapter that can communicate with both the PlateSpin Server and the source machine. This network can differ from the network on which the target virtual machine will run after the migration.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See "TCP/IP and Advanced Network Settings" on page 419.

Target Post-Migration Networking

Target post-migration network settings defined in a migration job control the network configuration of a target after the migration is complete. This applies to both physical and virtual network interfaces.

During workload migration, the target workload's post-migration network settings are configured while the workload is booted into a pre-execution environment.

To configure target post-migration network settings:

1. Start the migration job. For information about starting a migration job, see "Initiating a Migration Job" on page 386.
2. In the Network Configuration section of the Migration Job window, do one of the following:.
   - **For target virtual machines:** click **Guest NIC**.
   - **For target physical machines:** click **Network Connection**.
3. Configure the options as required and click **OK**.

   The Configuration options for the target post-migration network settings vary and depend on whether the network interface is virtual or physical, and whether it is connecting a Windows or a Linux workload. For more information about the options, review the following sections:

- "Post-Migration Networking for Physical Network Interfaces (Windows and Linux)" on page 418
- "Post-Migration Networking for Virtual Network Interfaces (Windows and Linux)" on page 418

Post-Migration Networking for Physical Network Interfaces (Windows and Linux)

Use these settings to configure the post-migration network settings of a workload being migrated to physical hardware.

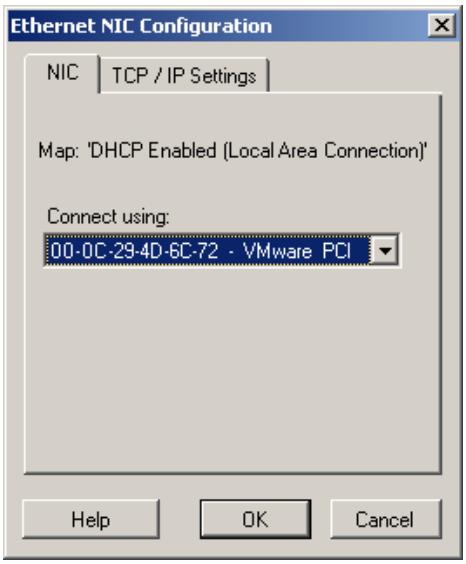

Connect using: If multiple network adapters are present, select the adapter that can communicate with the PlateSpin Server.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See "TCP/IP and Advanced Network Settings" on page 419.

Post-Migration Networking for Virtual Network Interfaces (Windows and Linux)

By default, PlateSpin Migrate configures a migration job to create a virtual NIC for each NIC found on the source. For post-migration connectivity, ensure that the target virtual NIC is mapped to the appropriate virtual network on the target virtualization platform.

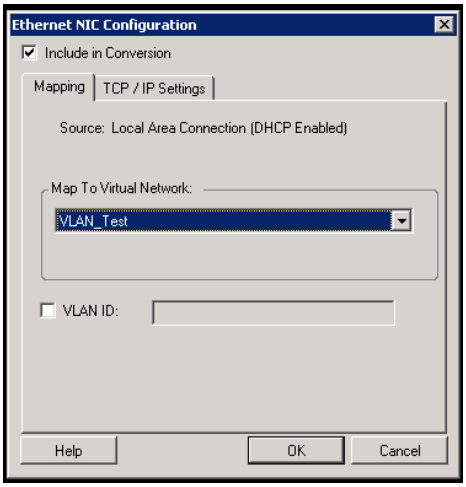

Include in Conversion: When this option is selected, PlateSpin Migrate creates a virtual NIC for a source NIC.

Map to Virtual Network: Select the virtual network that will be used on the target VM. Choose a virtual network that allows the target VM to communicate with the server.

Start connected: Enable this option to connect the virtual network interface when starting the ESX target machine.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See "TCP/IP and Advanced Network Settings" on page 419.

TCP/IP and Advanced Network Settings

PlateSpin Migrate provides a standard network configuration interface to both source and target network settings, and for both Temporary and target post-migration networking. Configuration settings vary slightly, depending on the operating system.

TCP/IP and Advanced Network Settings (Windows)

The following are standard TCP/IP and advanced network settings for Windows workloads:

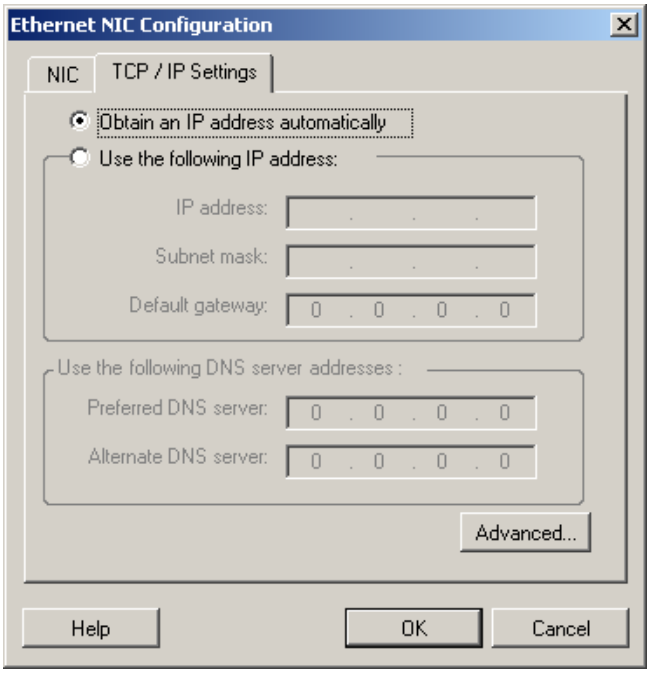

Obtain an IP address automatically: When this option is selected, the workload uses an IP address automatically assigned by a DHCP server during the migration process.

Use the following IP address: Select this option to specify a static IP address.

Use the following DNS server addresses: If required, specify preferred and alternative DNS server addresses.

Advanced: Click this button to access advanced TCP/IP configuration settings, then specify or edit default gateway, DNS server, and WINS server information as required.

TCP/IP and Advanced Network Settings (Linux)

The following are standard TCP/IP and advanced network settings for Linux workloads:

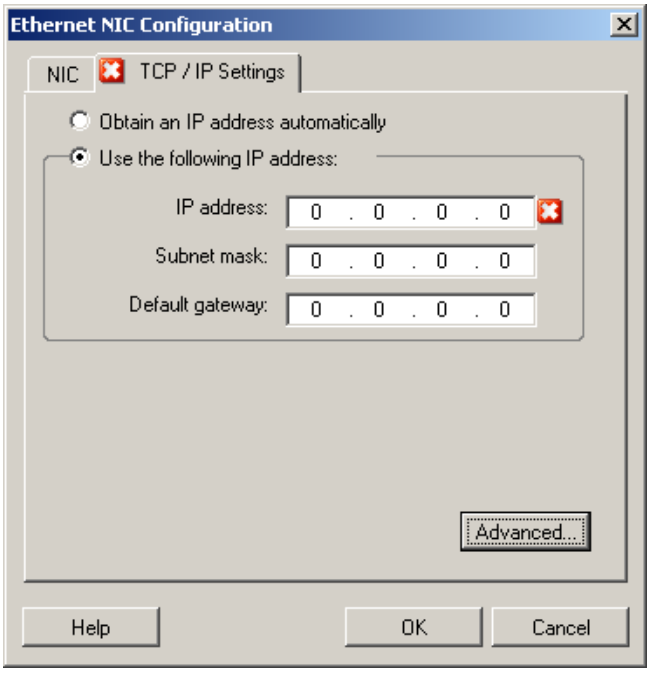

Obtain an IP address automatically: When this option is selected, the workload uses an IP address automatically assigned by a DHCP server during the migration process.

Use the following IP address: Select this option to specify a static IP address.

Advanced: Click this button to access DNS configuration settings, then specify preferred and alternate DNS server addresses as required. You can also indicate whether you want DNS addresses copied to the `resolv.conf` file located in your target's `/etc` directory.

Replication Network Using Migrate Web Interface

Replication Network for Migration to AWS

To specify the Replication Network for migration to Amazon Web Services:

1. In the Web Interface, select the Workload to go to the Target Configuration page, then click **Edit**.
2. Navigate to **Target Workload Settings** > **Network Connections**, then specify the Primary NIC. Migrate uses the Primary NIC as the Replication NIC.

3 Under **Migration Settings** in **Replication Network for Target**, specify the replication network settings:

3a Select a network and subnet to use for replication traffic.

3b If you do not use an AWS VPN, the replication NIC requires a public IP address that is automatically assigned by AWS. To enable AWS to automatically assign the public IP, select **Auto-assign Public IP**.

3c Select one of the following:

- **DHCP:** Obtain an IP address automatically assigned by a DHCP server.
- **Static:** Specify a static private IP address.

3d Click **Add Security Groups** to add one or more security groups to be used for the replication network. See "Create a Security Group" in the *Best Practices for Migrating Servers to Amazon Web Services with PlateSpin Migrate* white paper.

4 In **Replication Networks for Source**, specify one or more network interfaces (NIC or IP address) on the source workload to use for replication traffic that are valid for communications with the replication environment. If the network for the NIC you specify is not part of your AWS VPN, ensure that the NIC has a public IP address.

Replication Network for Migration to Azure

To specify the Replication Network for migration to Azure:

1 In the Web Interface, select the Workload to go to the Target Configuration page, then click **Edit**.

2 Navigate to **Target Workload Settings** > **Network Connections**, then specify the Primary NIC.

Migrate uses the Primary NIC as the Replication NIC.

3 Under **Migration Settings** in **Replication Network for Target**, specify the replication network settings:

3a Select a network and subnet to use for replication traffic.

3b If you do not use an Azure VPN, click **Edit**, then select **Create Public IP**.

When no VPN is present in the deployment, the replication NIC requires a public IP address that is automatically assigned by Azure.

3c Specify a resource group to use for the replication network.

The **Azure Resource Group** setting is the default. To specify a different resource group, click **Edit** and do one of the following:

- Type the name to use when PlateSpin creates a new resource group.
- Select an existing resource group from the list.

3d Select one of the following:

- **DHCP:** Obtain an IP address automatically assigned by a DHCP server.
- **Static:** Specify a static private IP address.

4 In **Replication Networks for Source**, specify one or more network interfaces (NIC or IP address) on the source workload to use for replication traffic that are valid for communications with the replication environment.

Replication Network for Migration to vCloud

To specify the Replication Network for migration to vCloud:

1. In the Web Interface, select the Workload to go to the Target Configuration page, then click **Edit**.
2. Under **Migration Settings** in **Replication Network for Target**, specify a network interface (NIC or IP address) on the target to use for replication traffic.
3. Under **Migration Settings** in **Replication Networks for Source**, specify one or more network interfaces (NIC or IP address) on the source to use for replication traffic.
   - **DHCP:** Obtain an IP address automatically assigned by a DHCP server.
   - **Static - Manual:** Specify a static IP address.
   - **Static - IP Pool:** Select this option to automatically issue IP address from the IP pool.

   For Windows workloads that have more than one NIC, select the connection for each NIC.

   For this setting, you can also specify an MTU value that the PlateSpin Migrate Linux RAM Disk (LRD) replication network can use. Setting a low value helps to avoid jabber over networks. For example: a VPN.

   The default value is an empty string. When networking is configured in the LRD, it allows the network device to set its own default, which is usually 1500. However, if you specify a value, PlateSpin Migrate adjusts the MTU when it configures the network interface.

Replication Network for Migration to VMware

To specify the Replication Network for migration to VMware:

1. In the Web Interface, select the Workload to go to the Target Configuration page, then click **Edit**.
2. Under **Migration Settings** in **Replication Network for Target**, specify a network interface (NIC or IP address) on the target to use for replication traffic.
3. Under **Migration Settings** in **Replication Networks for Source**, specify one or more network interfaces (NIC or IP address) on the source to use for replication traffic.

Storage Disks and Volumes

PlateSpin Migrate provides mechanisms for configuring your migration job to handle your workload volumes and their physical or virtual layout in the target infrastructure. For information about the supported storage, see “Supported Workload Storage” on page 35.

Storage layout and volume configuration settings depend on the job configuration mode (Advanced or Wizard), migration type, target virtualization platform, and source operating system.

The following topics provide additional information:

- “Storage Disks and Volumes Using Migrate Client” on page 423
- “Storage Disks and Volumes Using Migrate Web Interface” on page 428

Storage Disks and Volumes Using Migrate Client

To access drive configuration options:

- In the **Drive Configuration** of the Migration Job window, click **Hard Drives**.

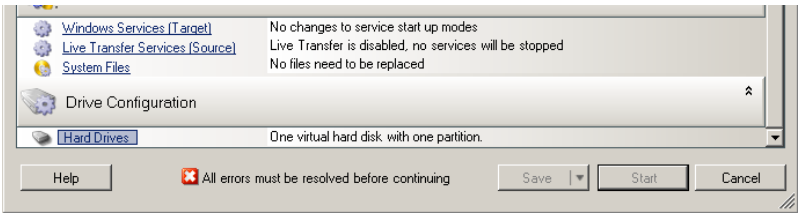

Settings vary depending on the target system.

- "Windows Drive Configuration" on page 423
- "Linux Drive and LVM Volume Configuration" on page 424
- "Target VM-Specific P2V/V2V Drive Configuration" on page 426
- "Volume Mapping in Server Sync" on page 427

Windows Drive Configuration

Use these settings to select the volumes to copy during the migration:

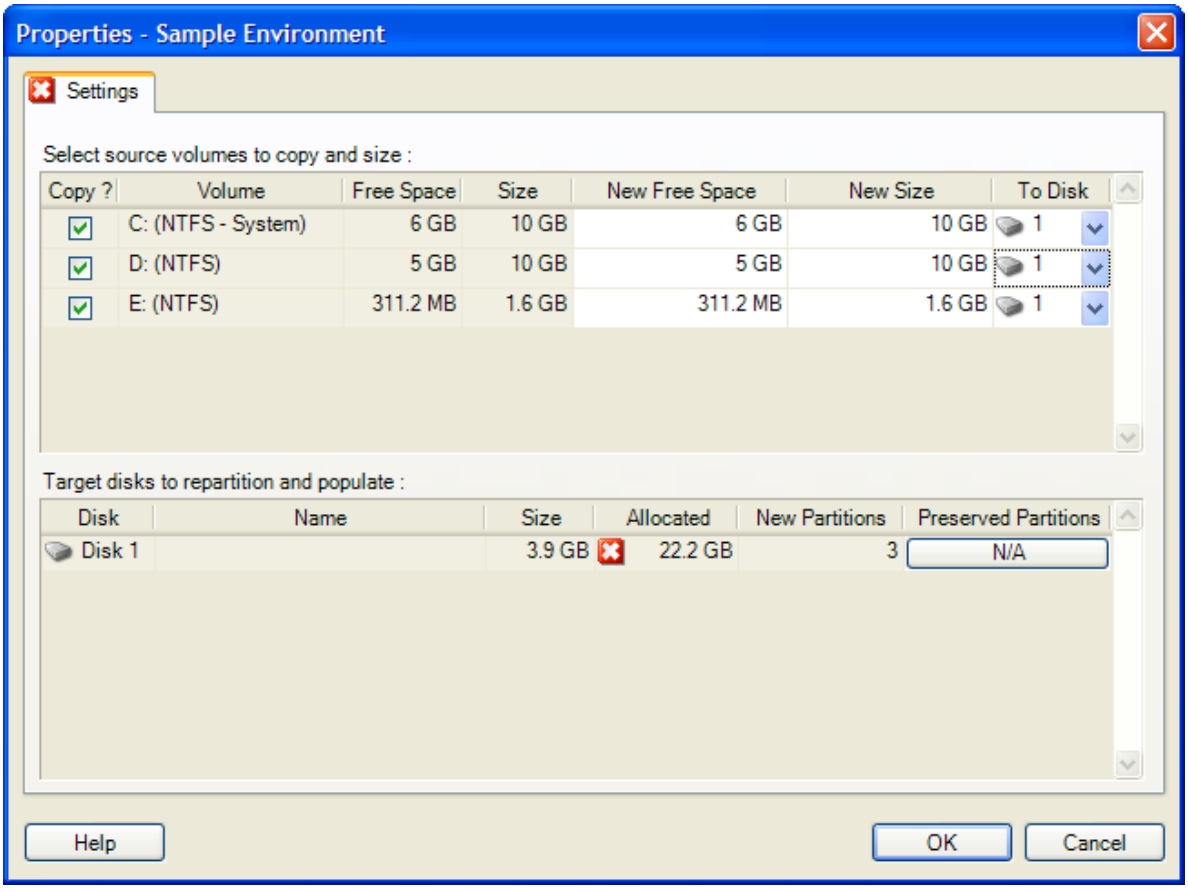

Copy: Select the volumes to be copied during the migration.

New Free Space: To resize the volume during the migration, specify the desired amount of free space. PlateSpin Migrate automatically adjusts **New Size**.

New Size: To resize the volume during the migration, specify the desired size. PlateSpin Migrate automatically adjusts **New Free Space**.

To Disk: Select which hard drive the volume will be copied to on the physical target machine.

Preserve Partitions: Click this column to determine if an existing vendor partition should remain intact during the migration. If the partitions are not selected, PlateSpin Migrate permanently removes the partitions from the server.

Linux Drive and LVM Volume Configuration

Use these settings to select the volumes and non-volume source spaces to copy and size during the migration. If LVM is installed on the source, a **Volume Group** tab provides you with corresponding options.

- "Handling Linux Disks and Volume Groups" on page 424
- "Linux Drive and LVM Volume Configuration (Settings Tab)" on page 425
- "Linux Drive and LVM Volume Configuration (Volume Groups Tab)" on page 426

Handling Linux Disks and Volume Groups

The PlateSpin Migrate Client provides you with Linux-specific user interface elements that provide you with options to properly handle your Linux storage.

Note the following sequence of steps that you must take for properly configuring and mapping newly-added disks and volume groups.

1. After adding a new disk, go to the **Volume Groups** tab and map the required volume group name by selecting the **Include** option.

   See Linux Drive and LVM Volume Configuration (Volume Groups Tab).
2. Specify Size in Allocation for Volume Group Box.
3. For each added disk, specify the required size in the corresponding **Allocation for Volume Group** field.

   After the system focus shifts away from the field, the size of the newly-added disk is updated dynamically.

Linux Drive and LVM Volume Configuration (Settings Tab)

Use these settings to select source volumes to copy, non-volume source spaces to re-create and size, and target disks to repartition and populate.

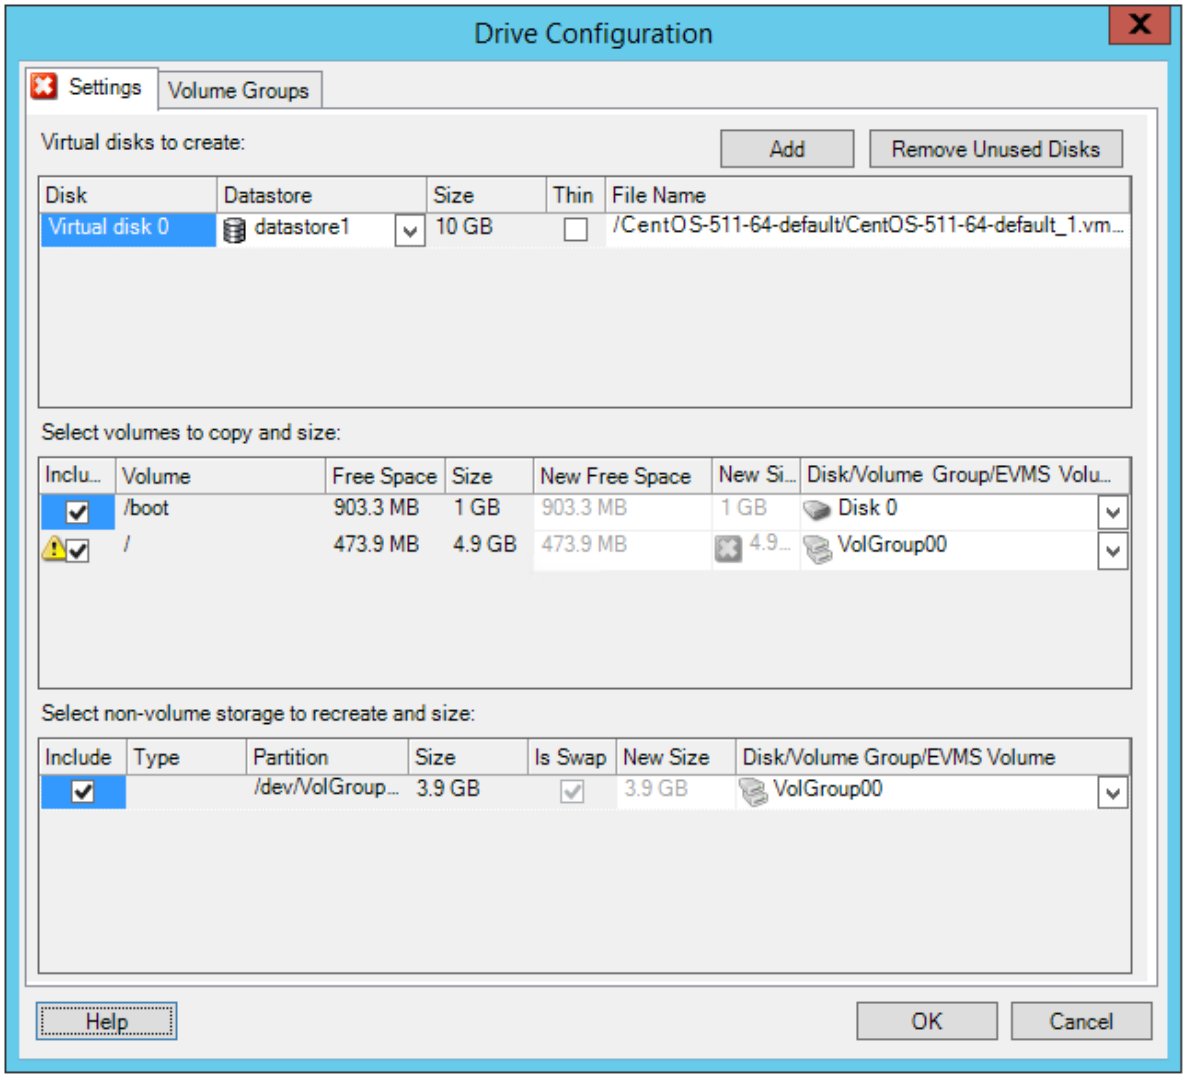

Include: Select the volumes or non-volume source spaces to be copied or re-created and sized during the migration.

New Free Space: To resize the volume during the migration, enter the desired amount of free space. PlateSpin Migrate automatically adjusts **New Size**.

New Size: To resize the volume during the migration, enter the desired size. PlateSpin Migrate automatically adjusts **New Free Space**.

Disk/Volume Group: Select which hard drive or volume group the volume will be copied to on the physical target machine.

Preserve Partitions: For each disk, click the corresponding cell in this column to select existing vendor partitions to preserve during the migration. If the partitions are not selected, PlateSpin Migrate permanently removes them from the server.

Linux Drive and LVM Volume Configuration (Volume Groups Tab)

Use these settings to manage volume groups.

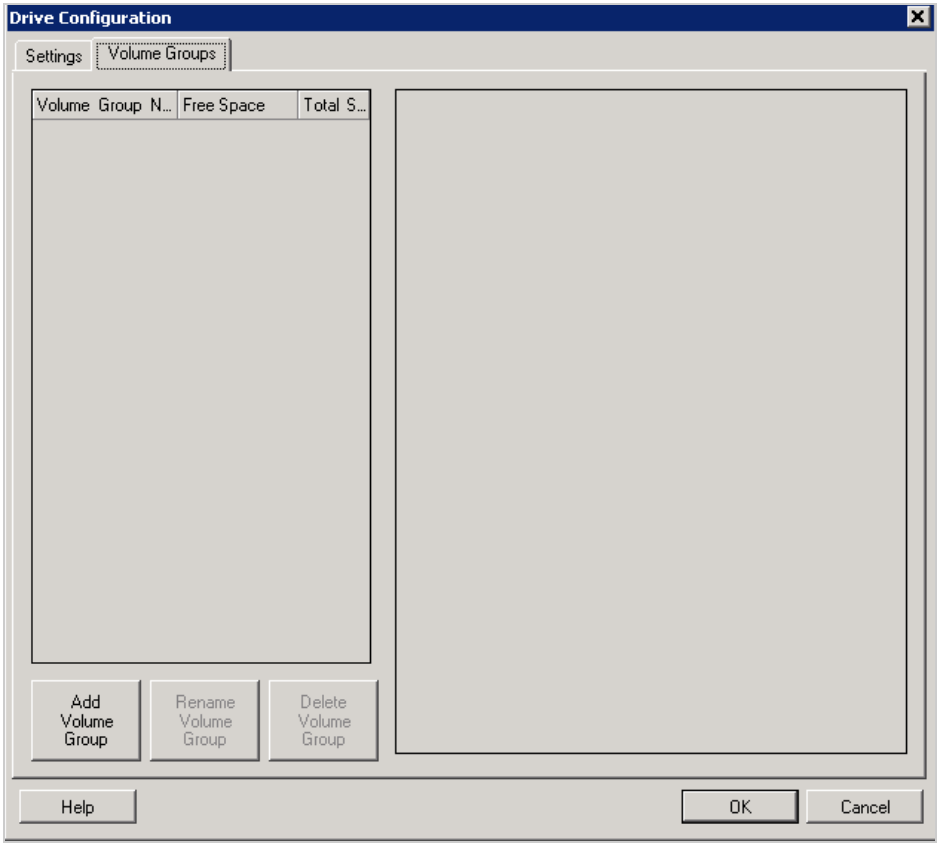

Add Volume Group: Creates a volume group on the target machine that is not present on the source machine.

Rename Volume Group: Renames a volume group that is being copied from the source to the target.

Delete Volume Group: Deletes a volume group so that it is not created on the target machine. The volumes assigned to the volume group can be reassigned to other locations by using the **Settings** tab (by default, they are assigned to disk).

Allocation for Volume Group: To allocate space on disks to a volume group, select the volume group, then select the disks to include in it. Specify the amount of space to be allocated to it on each included disk.

Target VM-Specific P2V/V2V Drive Configuration

When you configure a peer-to-peer virtualization job, the job configuration window provides access to settings specific to the target virtualization platform.

PlateSpin Migrate displays target virtual machine drive configuration settings specific to the selected target. See "Drive Configuration: VMware ESX" on page 497.

Volume Mapping in Server Sync

When you are using Server Sync to synchronize two Windows or Linux workloads, PlateSpin Migrate Client provides you with the capability to specify the required mapping between source volumes and existing volumes on the target. See “Server Sync Volume Mapping” on page 532.

Storage Disks and Volumes Using Migrate Web Interface

1 On the Edit Target Workload Details page, go to **Target Workload Settings** > **Migration Settings**.

2 Configure the following options:

Setting Name	Description
(Azure) Use Managed Disks	
(Azure) When you configure migrations to Azure, specify one of the following storage types for hosting the managed disks for the target VM: ◆ Standard HDD (hard disk drives) ◆ Standard SSD (solid state drives) ◆ Premium SSD A storage type of Standard HDD allows you to choose VMs with HDD or SDD storage, whereas Standard SSD and Premium SSD are restricted to VMs that support the appropriate level of SSD storage. For more information, see "Azure Managed Disks" on page 173.	
Storage Layout	
PlateSpin Migrate offers storage layout options of Same as Source or One Volume Per Disk. The option affects the layout of data volumes on the target workload. ◆ Source-based The Source-based storage layout assigns data volumes to disks according to their layout on the source workload. ◆ **Single disk** The Single disk storage layout assigns all the data volumes or LVM volume groups to a single disk on the target workload. An LVM volume group might contain multiple volumes. ◆ **One volume per disk** The One volume per disk storage layout assigns each data volume or LVM volume group to a separate disk on the target workload. An LVM volume group might contain multiple volumes. The following target storage configurations are the same for both layout options: ◆ **Windows Workloads** ◆ The first disk on the target workload will be the system disk that contains the operating system and the system reserve partition. ◆ Dynamic disks are resolved as Simple Basic Volume disks on the target workload. See "Disks" on page 35. ◆ **Linux Workloads** ◆ The first disk on the target workload will be the system disk that contains the `/boot` and root (`/`) volume. If the root (`/`) volume is in an LVM volume group (VG), then the VG is on the same disk as `/boot`. ◆ Spanned disks are resolved as a single disk on the target workload. See "Disks" on page 35.	

Setting Name	Description
Disks	
NOTE: Disks that host both regular volumes (such as `/boot`) and volume groups are configured in the Disks section. Configuration settings affect both the regular volumes and volume groups on the disk. **AWS** ◆ Select a disk type for each disk. The **Disk Type** option lists the type of disks that AWS supports. See Amazon EBS Volume Types (https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html). ◆ Select an encryption key to enable encryption of AWS target instance disks. Ensure that the currently logged in IAM user has sufficient permissions to use this encryption key. For information about creating the encryption key, see Creating Keys (https://docs.aws.amazon.com/kms/latest/developerguide/create-keys.html). **Azure** ◆ For Azure Managed Disks, you can optionally modify the name of the disk resource that represents the actual disks being managed by Azure. ◆ For unmanaged disks, specify the Storage Account to use and the path to the disk. You can optionally modify the name of the disk. **VMware or vCloud** ◆ Specify the path to the hard disk on the target virtual machine.	
Volumes	
Select volumes to be included in the target for migration. **NOTE:** To migrate volumes that are unmounted on the source workload, you must specify a mount path. For Windows file-based migrations, you can optionally adjust the size of the target volume by modifying the **Size** value in the row for that volume. **NOTE:** The ability to adjust the size of the target volume is currently not available in the Disk Layout view of storage in the Migrate Web Interface. See "Cannot Adjust the Target Volume Size for File-Based Migrations in the Web Interface" on page 569.	
NTFS Cluster Size	
(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see the Microsoft Support KB Article 140365.	
Non-volume Storage	
(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.	
Disks For Volume Groups	(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default. **NOTE:** Disks that host both regular volumes (such as `/boot`) and volume groups are configured in the Disks section.
Volume Groups	(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the **Converted Logical Volumes** section of the settings.

Setting Name	Description
Converted Logical Volumes	(For Linux Workloads) Specify one or more LVM logical volumes to be migrated for a Linux workload.

3 Click **Save**.

25 Migration to Amazon Web Services

Planning for Migration to Amazon Web Services

Before you begin migrations to your cloud environment in Amazon Web Services (AWS), ensure that your migration environment meets the following guidelines:

Supported Cloud Platforms

- See "Amazon Web Services (AWS)" in "Supported Target Cloud Platforms" on page 43.

Supported Workloads

- See "Supported Workloads For Migration to Amazon Web Services" on page 27, as appropriate for the target AWS environment.

Network Access and Communications

- See "Access and Communication Requirements across Your Migration Network" on page 53.

Prerequisites

- See Chapter 8, " Prerequisites for Migration to Amazon Web Services," on page 145.
- See Chapter 12, "Prerequisites for Cloud-to-Cloud Migrations," on page 217.

Targets and Workloads

- **Target AWS EC2 cloud account (automated):** See "Target Discovery in the Web Interface" on page 284 using Table 17-4, "Options for Amazon Cloud Region," on page 285.
- **Source Workloads:** Use either of the following discovery methods:
  - "Workload Discovery in the Migrate Web Interface" on page 302
  - "Registering Workloads and Discovering Details with Migrate Agent" on page 303

Additional Information

- See the Amazon Elastic Compute Cloud Documentation (https://aws.amazon.com/documentation/ec2/).
- See AWS Managed VPN (http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html) in the Amazon Virtual Private Cloud User Guide.
- See Your Customer Gateway (http://docs.aws.amazon.com/AmazonVPC/latest/NetworkAdminGuide/Introduction.html) in the Amazon Virtual Private Cloud Network Administrator Guide.

Configuring Migration of a Workload to Amazon Web Services

When you add or discover a workload, the workload is listed on the Workloads page and the status is set as **Not Configured**. Before you migrate the workload, you must configure the workload for migration:

1. Launch the PlateSpin Migrate Web Interface.
2. If you have not configured a Amazon Cloud Region as a migration target, click **Targets** > **Add Target**, and then configure the target AWS cloud platform.

   See Table 17-4, “Options for Amazon Cloud Region,” on page 285.
3. On the Workloads page, select the workload you want to configure.
4. Click **Configure Migration**.
5. Specify the **Initial Transfer Method** for replication based on the scope of data you want to transfer from the source to the target:
   - **Full Replication:** Migrate replicates the full volume from the source to the target.
   - **Incremental Replication:** Migrate replicates only differences in data from the source to the target, provided the workloads have similar operating system and volume profiles.

   NOTE: PlateSpin Migrate does not support Incremental Replication for the initial replication of data to existing target workloads in Amazon Cloud. However, you can schedule Incremental Replications for subsequent replication of data. See **Incremental Recurrence** in Step 8.

   ---
6. Select an existing Amazon Cloud Region target to which you want to migrate the source workload.
7. Click **Configure Migration**.

8 Configure the following settings:

Schedule Settings
Incremental Recurrence
Specify the time and pattern when you want to run incremental replications after the first full replication, or start each incremental replication manually. The default setting is None. The incremental replications are unscheduled. **To set or modify the incremental recurrence time and pattern:** 1. Click **Edit**. 2. For **Begin the recurrence schedule**, set the date and time when you want to begin the scheduled incremental replications. You can type the date (dd/mm/yyyy) or click the Calendar icon to select the date. By default, the run time is 12:00:00 a.m. (hh:mm:ss a.m. or p.m.). 3. For **Recurrence run setting**, set the pattern to follow for scheduled incremental replications: ◆ **Daily:** The replication takes place on the specified daily intervals or on weekdays every week for a period of 60 days from the time the replication starts. ◆ **Weekly:** The replication takes place at specified intervals for a period of 8 weeks from the time the replication starts. ◆ **Monthly:** The replication takes place at specified intervals for a period of 2 months from the time the replication starts. **NOTE:** ◆ Scheduled incremental replications are skipped until the first full replication is complete. ◆ Scheduled incremental replications take place for a maximum period of 60 days from the time that the scheduled incremental replication runs begin.
Full Replication
Specify when you want the first full replication to run, or start the first full replication manually. The first full replication is a one-time event, but the run is attempted daily as scheduled until the first replication begins and completes successfully. The default setting is None. The first full replication is unscheduled. **NOTE:** You must prepare the workload prior to the scheduled time or the manual start.The full replication cannot run unless the target VM exists and the workload preparation is complete. If they are not ready, Migrate skips the scheduled full replication and retries it at the scheduled time on the next day. **To set or modify the schedule for the first full replication:** 1. Click **Edit**. 2. Click **Start**, then set the date and time when you want to start the first full replication. You can type the date (dd/mm/yyyy) or click the Calendar icon to select the date. By default, the run time is 12:00:00 a.m. (hh:mm:ss a.m. or p.m.).

Blackout Window

Specify a replication blackout window that suspends scheduled replication activities for a specified period of time and pattern. For example, suspend replications during peak network utilization hours or to prevent conflicts between VSS-aware software and the PlateSpin VSS block-level data transfer component.

The default setting is None. No blackout window is scheduled.

To set or modify a blackout window:

1. Click **Edit**.
2. Specify the start and end time for the blackout period.

   The blackout start and end times are based on the system clock on the PlateSpin Server.
3. Select **Daily**, **Weekly**, or **Monthly** to enable a blackout window, then set the recurrence pattern.

Compression Level

This setting controls whether data is compressed during transmission between the source and target workloads, and the level of data compression applied.See "Data Compression" on page 52.

Select one of the following options:

- **None:** No compression.
- **Fast:** Consumes the least CPU resources on the source, but yields a lower compression ratio.
- **Optimal:** (Default) Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option.
- **Maximum:** Consumes the most CPU resources on the source, but yields a higher compression ratio.

Bandwidth Throttling

Bandwidth throttling enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. Throttling helps to prevent migration traffic from congesting your production network and to reduce the overall load of your PlateSpin Server. You can specify a throughput rate for each migration job.

Throttling is disabled by default with a Throttling Rate value of Off.

To throttle replications to a specified rate:

1. Specify a maximum throughput value in Mbps for data transfer for the workload.
2. Specify the throttling pattern:
   - **Always:** Always throttle data transfer for the replications.
   - **Custom:** Specify the time and days to throttle data transfer for the replications running in that window.

     Throttling time is local to the source workload.

Migration Settings

Transfer Method

(For Windows Workloads) Select a data transfer mechanism and security through encryption.See "Supported Data Transfer Methods" on page 44.

To enable encryption, select the **Encrypt Data Transfer** option. See "Security and Privacy" on page 46.

Transfer Encryption

(For Linux Workloads) To enable encryption, select the **Encrypt Data Transfer** option.See "Security and Privacy" on page 46.

Source Credentials

Specify the credentials required for accessing the workload. See "Discovery Guidelines for Source Workloads" on page 299.

Virtual Machine Name	

Specify a display name for the new virtual machine.

License Type

Select the OS licensing model on the target workload.

- **Auto:** (For Windows Workloads) Enables PlateSpin Migrate to decide whether to allow AWS to activate Windows license on the target Windows workload or allow users to bring their own licenses.
- **AWS:** (For Windows Workloads) Enables AWS to activate Windows license on the target Windows workload.
- **BYOL:** Enables you to bring your own Microsoft licenses (BYOL) and AWS does not bill you for the license. You are responsible for complying with Microsoft licensing and activating the OS license on the target workload. This option is applicable both for Windows and Linux workloads.

NOTE

- For AWS to activate the Windows license on the target workload, it is required that the KMS server is configured for Windows OS activation on the target workload. See "Configuring OS License Activation on Windows Targets Migrated to AWS" on page 157
- Based on the selected OS licensing model, PlateSpin Migrate uses one of the PlateSpin AMIs uploaded in the AWS community during the cutover of workloads to AWS. For information about the PlateSpin AMIs, see "Understanding PlateSpin AMIs Used for Replication and Cutover of Workloads" on page 157.
- If you choose to migrate a Windows workload to a dedicated host, the OS licensing model on the target workload is always set to BYOL irrespective of the licensing model you choose.

Storage Layout

Specify the desired layout of data volumes (and volume groups) on the target workload disks:

- **Source-based**

  The Source-based storage layout assigns data volumes to disks according to their layout on the source workload.

- **Single disk**

  The Single disk storage layout assigns all the data volumes or LVM volume groups to a single disk on the target workload. An LVM volume group might contain multiple volumes.

- **One volume per disk**

  The One volume per disk storage layout assigns each data volume or LVM volume group to a separate disk on the target workload. An LVM volume group might contain multiple volumes.

Storage Options > Disks
Select a disk type for each disk. The **Disk Type** option lists the type of disks that AWS supports. See Amazon EBS Volume Types (https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html). Select an encryption key to enable encryption of AWS target instance disks. Ensure that the currently logged in IAM user has sufficient permissions to use this encryption key. For information about creating the encryption key, see Creating Keys (https://docs.aws.amazon.com/kms/latest/developerguide/create-keys.html).
Storage Options > Volumes
Select volumes to be included in the target for migration. For Windows file-based migrations, you can optionally adjust the size of the target volume by modifying the **Size** value in the row for that volume. **NOTE:** The ability to adjust the size of the target volume is not currently available in the Disk Layout view of storage in the Migrate Web Interface. See "Cannot Adjust the Target Volume Size for File-Based Migrations in the Web Interface" on page 569.
NTFS Cluster Size
(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see the Microsoft Support KB Article 140365.
Non-volume Storage
(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups
(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.
Volume Groups
(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the **Converted Logical Volumes** section of the settings. **NOTE:** Disks that host both regular volumes (such as `/boot`) and volume groups are configured in the Disks section.
Converted Logical Volumes
(For Linux Workloads) Select LVM logical volumes to be included in the target for migration.

Replication Network for Target

The replication NIC is the primary NIC that you specify in **Target Workload Settings**> **Network Connections**.

1. Select a network and subnet to use for replication traffic.
2. If the workload is not part of the address space for the AWS VPN, the replication NIC requires a public IP address. Select **Auto-assign Public IP** to enable AWS to automatically assign the public IP.
3. Select one of the following:
   - **DHCP:** Obtain an IP address automatically assigned by a DHCP server.
   - **Static:** Specify a static private IP address that is unique within the supported subnet.
4. Click **Add Security Groups** to add one or more security groups. See "Create a Security Group" in the *Best Practices for Migrating Servers to Amazon Web Services with PlateSpin Migrate* white paper.

Replication Networks for Source

Specify one or more network interfaces (NIC or IP address) on the source workload to use for replication traffic that are valid for communications with the replication environment.

If the network for the NIC you specify is not part of your AWS VPN, ensure that the NIC has a public IP address.

Services to Stop Before Any Replication

(For Windows Workloads) We recommend that all the non-VSS compliant services or antivirus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.

Services to Stop for Cutover with Replication

(For Windows Workloads) Select the Windows services that should be permanently stopped on the source workload for cutover with any replication. The services stopped on the source workload during the replication process are not restored after Cutover. The stopped services are restored after a Test Cutover.

Daemons to Stop before Any Replication

(For Linux Workloads) Select the Linux services that you want to be temporarily stopped on the source workload before replication. These services will be restored back after replication completes.

Daemons to Stop for Cutover with Replication

(For Linux Workloads) Select the Linux services that should be permanently stopped on the source workload for Cutover with any Replication. The services stopped on the source workload during the replication process are not restored after Cutover. The stopped services are restored after a Test Cutover.

Boot Services to Disable on Target

(For Windows Workloads) Select the boot services that you want to disable on the target workload.

Target Workload Settings

(These settings are applied during the Run Cutover)

AWS Agent

For collecting metrics and logs from Amazon EC2 Instance on the target workload, select the **Install CloudWatch Agent** option to install the AWS CloudWatch agent on the target workload. This option is displayed for the following workloads that supports CloudWatch agent:

- Red Hat Enterprise Linux (RHEL) versions 7.6, 7.5, 7.4, 7.2, 7.0, and 6.5
- CentOS versions 7.6, 7.2, 7.0, 6.8, and 6.5
- 64-bit versions of Windows server

If the AWS CloudWatch Agent is already installed on the source workload, then this option is disabled.

Tenancy

Select one of the following options to specify whether your instance should run on a shared or a dedicated hardware:

- **Run a shared hardware instance:** Your instance runs on a shared hardware and this is selected by default.
- **Run a dedicated instance:** Your instance runs on a single-tenant hardware.
- **Launch this instance on a dedicated host:** Your instance runs on a dedicated host, which is an isolated server already allocated for use in your account.

  NOTE: If you choose to launch the instance on a dedicated host, the OS licensing model on the target workload is always set to BYOL irrespective of the licensing model you selected.

  Set the following options based on your requirement:

  - **Host:** Select a specific host to launch the instance or select **Use auto-placement** to allow the instance to launch on to any host that has a matching instance type and auto-placement enabled. The **Use auto-placement** option is selected by default if any of the available dedicated hosts supports auto-placement.
  - **Affinity:** For a specific dedicated host, the affinity is always **Host**. However, if you set the **Host** option to **Use auto-placement**, then select one of the following:
    - **Off:** Restarts a stopped instance on any available host. This option is selected by default.
    - **Host:** Restarts a stopped instance on the same host where it was launched.

Cloud Instance Size

Click **Change Cloud Instance Size** to select a supported cloud instance size appropriate for your workload.

- For migration of Windows Server 2019 workload to AWS cloud, only `t2` instances are supported.
- If an instance type that AWS supports is not listed, then you can configure the `AWSPriceListRegion` PlateSpin Configuration parameter to set its value to the region name that has a price list endpoint listing the desired instance type. See "Configuring the AWS Region Price List Endpoint To Be Used for Discovering Supported AWS Instance Types" on page 156.
- For Linux workloads, you must ensure that the instance type you select for the Target Workload and Target Workload Test Settings must both be of the same type: Nitro or non-Nitro. Using a combination of Nitro and non-Nitro instances during Test Cutover and Run Cutover (that is, Nitro instance during Test Cutover and a non-Nitro instance during Run Cutover or vice versa) results in migration failure.
- As AWS adds support for new instance types, Migrate detects them dynamically and displays them with an asterisk symbol. PlateSpin Migrate might not have tested such instance types that AWS recently added and support for such AWS instance types is experimental. For a list of the new instance types that Migrate supports in this release, see "What's New" in the *PlateSpin Migrate 2020.2 Release Notes*.
- For each release, PlateSpin Migrate adds support for new cloud instances and instance families that AWS has released up to a stipulated date prior to the PlateSpin Migrate release. For a list of the new instance types that Migrate supports in this release, see "What's New" in the *PlateSpin Migrate 2020.2 Release Notes*.

  You can add definitions to the PlateSpin Server for new AWS Nitro Based instances and instance families to make the instances available in the AWS Cloud Instance Sizes list. See "Displaying the AWS Nitro System Based Cloud Instance Sizes That AWS Recently Supported in the AWS Cloud Instance Sizes list" on page 163.
- The m5a, r5a, and F1 instance types are supported only for Linux workloads.
- For migration of a source Windows Domain Controller workload, select a Nitro System based instance type such as c5, m5, r5, t3, c5d, m5d, or z1d.
- AWS requires a minimum of Microsoft .NET Framework 4.5 for migration of Windows workloads to Nitro System based VM instances. To use Nitro System based instances, ensure that .NET 4.5 Framework or later is installed on the source Windows workload prior to Migrate discovery.

By default, Migrate selects a cloud instance size that most closely matches your source workload for the following components:

- Total number of cores
- Amount of memory
- Number of NICs
- Network Performance
- AWS Instance Family

The default instance either meets or exceed the settings for each of these components on the source workload. However, you can choose a smaller instance size based on your requirements:

- The target VM uses the allowed CPU and memory for the instance size. To reduce the number of CPUs or amount of memory on the target workload:
  1. Select a smaller cloud instance size with fewer CPUs or less memory that best fits your needs.
- The target VM uses up to the maximum allowed number of NICs for the instance size. To migrate only some of the NICs:
  1. Select a cloud instance size with fewer NICs that best fits your needs. At least one NIC is required.
  2. Under **Target Workload Settings**, deselect the NICs that should not be migrated until the number of NICs for migration fits the selected instance.

NOTE: The `i3.16xlarge` cloud instance size is not supported for migration of Windows Server 2008 R2 Workload to AWS. Use a supported cloud instance size other than `i3.16xlarge`.

AWS Instance Tags

AWS allows you to assign metadata to their resources in the form of tags thereby making it easy to manage, search for, and filter resources. To add tags, do the following:

1. Click **Add/Edit** Tags and then click **Create Tag**.
2. Specify a key and value for the tag.
3. Click Apply.

You can edit tags key and value, and also remove tags.

Placement Groups

This setting is applicable only if you set the **Tenancy** to run your instance as a shared instance.

Select a placement group where you want to launch your instance.

IMPORTANT: Placement Group configuration in Migrate is limited to cloud instance types supported by Amazon EC2. Refer to AWS EC2 Documentation for the latest information about placement groups and AWS rules and limitations for using them:

- "Placement Groups" in the *AWS EC2: User Guide for Windows* Instances (https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/placement-groups.html).
- "Placement Groups" in the *AWS EC2: User Guide for Linux Instances* (https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html#concepts-placement-groups).

IAM Role

Select an AWS Identity and Access Management (IAM) user in your AWS account, with an appropriate IAM role for the user to perform migrations into the VPC using the AWS APIs.

Key Pair

Select the AWS EC2 Key Pair that you want to use for logging in to your AWS target instance. However, if you do not want to use a key pair, select **Proceed without a key pair** to use only the source credentials for logging in to your AWS target instance.

NOTE: When you select a key pair, PlateSpin Migrate by default allows you to log in to the AWS target instance only by using the selected key pair. To enable logging into AWS Linux target instance either by using the key pair configured in the migration job or the source credentials, see "Configuring Target Instance Logging With Key Pair or Source Credentials" on page 156.

For information about creating the key pair, see:

- **For Windows:** *Amazon EC2 Key Pairs and Windows Instances* (https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/ec2-key-pairs.html)
- **For Linux:** *Amazon EC2 Key Pairs and Linux Instances* (https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-key-pairs.html).

Hostname

Do one of the following:

- To retain the same host name, select **No Change**.
- To change the host name, select **Set To** and specify the new name.

  NOTE: An incremental replication is required if you change the host name at cutover.

Domain / Workgroup
(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays: ◆ **Workgroup:** *Workgroup_name* where *Workgroup_name* is the workgroup name to which the source belongs. ◆ **Domain:** *Domain_name* where *Domain_name* is the domain name to which the source belongs. **NOTE:** An incremental replication is required if you change the domain or workgroup name at cutover. Do one of the following depending on where you want the target workload to join: ◆ **When the source workload belongs to a workgroup:** Assume that the source workload belongs to a workgroup named **WorkGroup1**. ◆ For the target workload to join the same workgroup (**WorkGroup1**), retain the following existing selection: **Workgroup: Workgroup1** ◆ For the target workload to join a different workgroup (say **WorkGroup2)**, select **Join Workgroup** and specify the name as **WorkGroup2**. ◆ For the target workload to join a domain, select **Join Domain** and specify the domain name you want the target to join. ◆ **When the source workload belongs to a domain:** Assume that the source workload belongs to a domain named **Domain1**. ◆ For the target workload to join a workgroup, click **Join Workgroup** and specify the name of the workgroup you want the target to join. ◆ For the target workload to join the same domain (**Domain1**) with the domain registration settings preserved, retain the following existing selection: **Domain: Domain1** ◆ For the target workload to join the same domain (**Domain1)** without preserving the domain registration settings, select **Join Domain** and specify the domain name as **Domain1**. ◆ For the target workload to join a different domain, select **Join Domain** and specify the domain name you want the target to join.
Domain Credentials
(For Windows Workloads) If you select **Join Domain**, specify the domain administrator credentials.

Network Connections

1. To provide high-performance networking capabilities on the workload, PlateSpin Migrate selects the **Enable Enhanced Networking** option by default if the selected instance type supports only ENA adapter. However, if the selected instance type supports both ENA and Intel adapters, then select the **Enable Enhanced Networking** option if you want to use ENA adapter.

   IMPORTANT

   - AWS supports enhanced networking capabilities on selected instance types. If you select this option to enable enhanced networking for an unsupported instance type, you receive a validation error. To see the list of supported instances, refer to the following topics in the AWS Documentation:
     - Enhanced Networking on Windows
     - Enhanced Networking on Linux
   - (For Linux workloads) Using Enhanced networking with Elastic Network Adapter (ENA) capability on a Linux workload requires ENA drivers on the workload. See "Using Enhanced Networking with ENA on Linux Distributions" on page 155.

2. For workloads that have more than one NIC, select **Include** for each NIC to be migrated. Deselect **Include** to exclude a NIC.
   - At least one NIC is required.
   - The number of NICs to migrate cannot exceed the maximum number of NICs supported by the selected cloud instance.
   - If the source workload is not part of the address space for the AWS VPN, then a public IP address is required for migration. To enable AWS to automatically assign a public IP address, you must include only one NIC for migration. This is because AWS supports assigning public IP address only to instances with a single network interface. To ensure that only public IP is used during migration, configure the `UseOnlyPublicIPForAWS` parameter in the PlateSpin Configuration settings for the Migrate server as `True`. See "Configuring PlateSpin Migrate Server to Use Public IP Address for AWS Migrations" on page 156.
3. For each included NIC, select a network and subnet.
4. (For primary NIC) Select **Auto-assign Public IP** to enable AWS to automatically assign a public IP address.

   NOTE: PlateSpin Migrate uses the `EnablePublicIPForCloudInstance` configuration parameter to control whether or not a public IP address should be assigned to the primary network interface of AWS instances. If the value of `EnablePublicIPForCloudInstance` is true, then **Auto-assign Public IP** option is selected by default and you cannot deselect it. However, if the value of `EnablePublicIPForCloudInstance` is false, then also **Auto-assign Public IP** option is selected by default, but you can deselect it if you do not want AWS to automatically assign a public IP address to primary NIC on the target instance. See "Enabling Public IP Address for Target Instances on Cloud" on page 112.
5. For each included NIC, select one of the following:
   - **DHCP:** Obtain an IP address automatically assigned by a DHCP server.
   - **Static:** Specify one or more static IP addresses. The network mask and gateway are auto-populated based on the selected subnet. Click **Advanced** to add multiple IP addresses, edit an IP address, or remove an IP address. The IP addresses must be unique within the supported subnet.

DNS Servers

Specify the DNS Servers for the target workloads. This is applicable only if you select **Static** in the **Network Connections** option:

- **Primary DNS server:** Specify the primary DNS server address.
- **Alternative DNS server:** Specify an alternate DNS server address.
- **Additional DNS server:** To specify additional DNS server addresses:
  1. Click **Advanced**.
  2. Specify the DNS server address.
  3. Click **Add** to add the server in the DNS Server Addresses list.
  4. Click **OK**.

Services States on Target VM

(For Windows Workloads) Select Windows services' start conditions on the target VM. Start options are **Automatic**, **Manual**, **Disabled**, and **Automatic (Delayed Start).**

D**aemons States to Change**

(For Linux Workloads) Select Linux daemons' start conditions on the target VM. Enable the daemon to start by selecting the check boxes at the appropriate runlevels (0 to 6) and Boot.

Post-Migration Action

Click **Configure Action** to select a previously configured custom action that you want to execute on your target workload as a part of the migration job. The command line parameters and timeout value that you previously configured for the selected action is displayed by default. You can edit these values as required. If you choose to timeout the action and have not specified any timeout value, then PlateSpin Migrate waits for a maximum period of 7200 seconds before timing out this action. For information about managing the post-migrations, see "Managing Post-Migration Actions" on page 107.

Target Workload Test Settings

(These settings are applied during the Test Cutover)

Copy Target Workload Settings

Click the **Copy Target Workload Settings** option to automatically copy the workload settings from **Target Workload Settings** section to **Target Workload Test Settings** section.

AWS Agent

For collecting metrics and logs from Amazon EC2 Instance on the target workload, select the **Install CloudWatch Agent** option to install the AWS CloudWatch agent on the target workload. This option is displayed for the following workloads that supports CloudWatch agent:

- Red Hat Enterprise Linux (RHEL) versions 7.6, 7.5, 7.4, 7.2, 7.0, and 6.5
- CentOS versions 7.6, 7.2, 7.0, 6.8, and 6.5
- 64-bit versions of Windows server

If the AWS CloudWatch Agent is already installed on the source workload, then this option is disabled.

Tenancy
Select one of the following options to specify whether your instance should run on a shared or a dedicated hardware: ◆ **Run a shared hardware instance:** Your instance runs on a shared hardware and this is selected by default. ◆ **Run a dedicated instance:** Your instance runs on a single-tenant hardware. ◆ **Launch this instance on a dedicated host:** Your instance runs on a dedicated host, which is an isolated server already allocated for use in your account. **NOTE:** If you choose to launch the instance on a dedicated host, the OS licensing model on the target workload is always set to BYOL irrespective of the licensing model you selected. Set the following options based on your requirement: ◆ **Host:** Select a specific host to launch the instance or select **Use auto-placement** to allow the instance to launch on to any host that has a matching instance type and auto-placement enabled. The **Use auto-placement** option is selected by default if any of the available dedicated hosts supports auto-placement. ◆ **Affinity:** For a specific dedicated host, the affinity is always **Host**. However, if you set the **Host** option to **Use auto-placement**, then select one of the following: ◆ **Off:** Restarts a stopped instance on any available host. This option is selected by default. ◆ **Host:** Restarts a stopped instance on the same host where it was launched.

Cloud Instance Size

Click **Change Cloud Instance Size** to select a supported cloud instance size appropriate for your workload.

- For migration of Windows Server 2019 workload to AWS cloud, only `t2` instances are supported.
- If an instance type that AWS supports is not listed, then you can configure the `AWSPriceListRegion` PlateSpin Configuration parameter to set its value to the region name that has a price list endpoint listing the desired instance type. See "Configuring the AWS Region Price List Endpoint To Be Used for Discovering Supported AWS Instance Types" on page 156.
- For Linux workloads, you must ensure that the instance type you select for the Target Workload and Target Workload Test Settings must both be of the same type: Nitro or non-Nitro. Using a combination of Nitro and non-Nitro instances during Test Cutover and Run Cutover (that is, Nitro instance during Test Cutover and a non-Nitro instance during Run Cutover or vice versa) results in migration failure.
- As AWS adds support for new instance types, Migrate detects them dynamically and displays them with an asterisk symbol. PlateSpin Migrate might not have tested such instance types that AWS recently added and support for such AWS instance types is experimental. For a list of the new instance types that Migrate supports in this release, see "What's New" in the *PlateSpin Migrate 2020.2 Release Notes*.
- For each release, PlateSpin Migrate adds support for new cloud instances and instance families that AWS has released up to a stipulated date prior to the PlateSpin Migrate release. For a list of the new instance types that Migrate supports in this release, see "What's New" in the *PlateSpin Migrate 2020.2 Release Notes*.

  You can add definitions to the PlateSpin Server for new AWS Nitro Based instances and instance families to make the instances available in the AWS Cloud Instance Sizes list. See "Displaying the AWS Nitro System Based Cloud Instance Sizes That AWS Recently Supported in the AWS Cloud Instance Sizes list" on page 163.
- The m5a, r5a, and F1 instance types are supported only for Linux workloads.
- For migration of a source Windows Domain Controller workload, select a Nitro System based instance type such as c5, m5, r5, t3, c5d, m5d, or z1d.
- AWS requires a minimum of Microsoft .NET Framework 4.5 for migration of Windows workloads to Nitro System based VM instances. To use Nitro System based instances, ensure that .NET 4.5 Framework or later is installed on the source Windows workload prior to Migrate discovery.

By default, Migrate selects a cloud instance size that most closely matches your source workload for the following components:

- Total number of cores
- Amount of memory
- Number of NICs
- Network Performance
- AWS Instance Family

The default instance either meets or exceed the settings for each of these components on the source workload. However, you can choose a smaller instance size based on your requirements:

- The target VM uses the allowed CPU and memory for the instance size. To reduce the number of CPUs or amount of memory on the target workload:
  1. Select a smaller cloud instance size with fewer CPUs or less memory that best fits your needs.
- The target VM uses up to the maximum allowed number of NICs for the instance size. To migrate only some of the NICs:
  1. Select a cloud instance size with fewer NICs that best fits your needs. At least one NIC is required.
  2. Under **Target Workload Settings**, deselect the NICs that should not be migrated until the number of NICs for migration fits the selected instance.

NOTE: The `i3.16xlarge` cloud instance size is not supported for migration of Windows Server 2008 R2 Workload to AWS. Use a supported cloud instance size other than `i3.16xlarge`.

AWS Instance Tags

AWS allows you to assign metadata to their resources in the form of tags thereby making it easy to manage, search for, and filter resources. To add tags, do the following:

1. Click **Add/Edit** Tags and then click **Create Tag**.
2. Specify a key and value for the tag.
3. Click Apply.

You can edit tags key and value, and also remove tags.

Placement Groups

This setting is applicable only if you set the **Tenancy** to run your instance as a shared instance.

Select a placement group where you want to launch your instance.

IMPORTANT: Placement Group configuration in Migrate is limited to cloud instance types supported by Amazon EC2. Refer to AWS EC2 Documentation for the latest information about placement groups and AWS rules and limitations for using them:

- "Placement Groups" in the *AWS EC2: User Guide for Windows Instances* (https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/placement-groups.html).
- "Placement Groups" in the *AWS EC2: User Guide for Linux Instances* (https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html#concepts-placement-groups).

IAM Role

Select an AWS Identity and Access Management (IAM) user in your AWS account, with an appropriate IAM role for the user to perform migrations into the VPC using the AWS APIs.

Key Pair

Select the AWS EC2 Key Pair that you want to use for logging in to your AWS target instance. However, if you do not want to use a key pair, select **Proceed without a key pair** to use only the source credentials for logging in to your AWS target instance.

NOTE: When you select a key pair, PlateSpin Migrate by default allows you to log in to the AWS target instance only by using the selected key pair. To enable logging into AWS Linux target instance either by using the key pair configured in the migration job or the source credentials, see "Configuring Target Instance Logging With Key Pair or Source Credentials" on page 156.

For information about creating the key pair, see:

- **For Windows:** Amazon EC2 Key Pairs and Windows Instances (https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/ec2-key-pairs.html)
- **For Linux:** Amazon EC2 Key Pairs and Linux Instances (https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-key-pairs.html).

Hostname

Do one of the following:

- To retain the same host name, select **No Change**.
- To change the host name, select **Set To** and specify the new name.

  NOTE: An incremental replication is not required if you change the host name at Test Cutover.

Domain / Workgroup

(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays:

- **Workgroup:** *Workgroup_name*

  where *Workgroup_name* is the workgroup name to which the source belongs.

- **Domain:** *Domain_name*

  where *Domain_name* is the domain name to which the source belongs.

NOTE: An incremental replication is not required if you change the domain or workgroup name at Test Cutover.

Do one of the following depending on where you want the target workload to join:

- **When the source workload belongs to a workgroup:** Assume that the source workload belongs to a workgroup named **WorkGroup1**.
  - For the target workload to join the same workgroup (**WorkGroup1**), retain the following existing selection:

    Workgroup: Workgroup1
  - For the target workload to join a different workgroup (say **WorkGroup2)**, select **Join Workgroup** and specify the name as **WorkGroup2**.
  - For the target workload to join a domain, select **Join Domain** and specify the domain name you want the target to join.
- **When the source workload belongs to a domain:** Assume that the source workload belongs to a domain named **Domain1**.
  - For the target workload to join a workgroup, click **Join Workgroup** and specify the name of the workgroup you want the target to join.
  - For the target workload to join the same domain (**Domain1**) with the domain registration settings preserved, retain the following existing selection:

    Domain: Domain1
  - For the target workload to join the same domain (**Domain1)** without preserving the domain registration settings, select **Join Domain** and specify the domain name as **Domain1**.
  - For the target workload to join a different domain, select **Join Domain** and specify the domain name you want the target to join.

Domain Credentials

(For Windows Workloads) If you select **Join Domain**, specify the domain administrator credentials.

Network Connections

1. To provide high-performance networking capabilities on the workload, PlateSpin Migrate selects the **Enable Enhanced Networking** option by default if the selected instance type supports only ENA adapter. However, if the selected instance type supports both ENA and Intel adapters, then select the **Enable Enhanced Networking** option if you want to use ENA adapter.

   IMPORTANT

   - AWS supports enhanced networking capabilities on selected instance types. If you select this option to enable enhanced networking for an unsupported instance type, you receive a validation error. To see the list of supported instances, refer to the following topics in the AWS Documentation:
     - Enhanced Networking on Windows
     - Enhanced Networking on Linux
   - (For Linux workloads) Using Enhanced networking with Elastic Network Adapter (ENA) capability on a Linux workload requires ENA drivers on the workload. See "Using Enhanced Networking with ENA on Linux Distributions" on page 155.

2. For workloads that have more than one NIC, select **Include** for each NIC to be migrated. Deselect **Include** to exclude a NIC.
   - At least one NIC is required.
   - The number of NICs to migrate cannot exceed the maximum number of NICs supported by the selected cloud instance.
   - If the source workload is not part of the address space for the AWS VPN, then a public IP address is required for migration. To enable AWS to automatically assign a public IP address, you must include only one NIC for migration. This is because AWS supports assigning public IP address only to instances with a single network interface. To ensure that only public IP is used during migration, configure the `UseOnlyPublicIPForAWS` parameter in the PlateSpin Configuration settings for the Migrate server as `True`. See "Configuring PlateSpin Migrate Server to Use Public IP Address for AWS Migrations" on page 156.
3. For each included NIC, select a network and subnet.
4. (For primary NIC) Select **Auto-assign Public IP** to enable AWS to automatically assign a public IP address.

   NOTE: PlateSpin Migrate uses the `EnablePublicIPForCloudInstance` configuration parameter to control whether or not a public IP address should be assigned to the primary network interface of AWS instances. If the value of `EnablePublicIPForCloudInstance` is true, then **Auto-assign Public IP** option is selected by default and you cannot deselect it. However, if the value of `EnablePublicIPForCloudInstance` is false, then also **Auto-assign Public IP** option is selected by default, but you can deselect it if you do not want AWS to automatically assign a public IP address to primary NIC on the target instance. See "Enabling Public IP Address for Target Instances on Cloud" on page 112.
5. For each included NIC, select one of the following:
   - **DHCP:** Obtain an IP address automatically assigned by a DHCP server.
   - **Static:** Specify one or more static IP addresses. The network mask and gateway are auto-populated based on the selected subnet. Click **Advanced** to add multiple IP addresses, edit an IP address, or remove an IP address. The IP addresses must be unique within the supported subnet.

DNS Servers
Specify the DNS Servers for the target workloads. This is applicable only if you select **Static** in the **Network Connections** option: ◆ **Primary DNS server:** Specify the primary DNS server address. ◆ **Alternative DNS server:** Specify an alternate DNS server address. ◆ **Additional DNS server:** To specify additional DNS server addresses: 1. Click **Advanced**. 2. Specify the DNS server address. 3. Click **Add** to add the server in the DNS Server Addresses list. 4. Click **OK**.
Services States on Target VM
(For Windows Workloads) Select Windows services that must be automatically stopped on the target VM.
Daemons States to Change
(For Linux Workloads) Select Linux daemons that must be automatically stopped on the target VM.
Tag
Post-Migration Action
Click **Configure Action** to select a previously configured custom action that you want to execute on your target workload as a part of the migration job. The command line parameters and timeout value that you previously configured for the selected action is displayed by default. You can edit these values as required. If you choose to timeout the action and have not specified any timeout value, then PlateSpin Migrate waits for a maximum period of 7200 seconds before timing out this action. For information about managing the post-migrations, see "Managing Post-Migration Actions" on page 107.
Tag
Select a tag to assign to the workload. For more information about tags, see "Using Tags to Track Logical Associations of Workloads" on page 309.

9 (Optional) To change the target, click **Change Target**.

NOTE: If you change the target, all the settings you specified will be cleared.

10 Do one of the following:

- Click **Save** to save the settings.
- Click **Save and Prepare** to save the settings and start preparing the workload migration.
- Click **Cancel** to exit.

26 Migration to Microsoft Azure

Planning for Migration to Microsoft Azure

Before you begin migrations to your cloud environment in Microsoft Azure, ensure that your migration environment meets the following guidelines:

Supported Cloud Platforms

- See "Microsoft Azure" in "Supported Target Cloud Platforms" on page 43.

Supported Workloads

- See "Supported Workloads For Migration to Microsoft Azure" on page 29, as appropriate for the target Azure cloud environment.

Network Access and Communications

- See "Access and Communication Requirements across Your Migration Network" on page 53.

Prerequisites

- See Chapter 9, "Prerequisites for Migration to Microsoft Azure," on page 167.
- See Chapter 12, "Prerequisites for Cloud-to-Cloud Migrations," on page 217.

Targets and Workloads

- **Target Azure cloud subscription (automated):** See "Target Discovery in the Web Interface" on page 284 using Table 17-5, "Options for Microsoft Azure Location Target," on page 286.
- **Source Workloads:** Use either of the following discovery methods:
  - "Workload Discovery in the Migrate Web Interface" on page 302
  - "Registering Workloads and Discovering Details with Migrate Agent" on page 303

Additional Information

- See "Create a Site-to-Site Connection in the Azure Portal" in the *Microsoft Azure VPN Gateway Documentation*.
- See "Create a VNet with a Site-to-Site VPN Connection Using PowerShell" in the *Microsoft Azure VPN Gateway Documentation*.

Configuring Migration of a Workload to Microsoft Azure

When you add or discover a workload, the workload is listed on the Workloads page and the status is set as **Not Configured**. Before you migrate the workload, you must configure the workload for migration:

1. Launch the PlateSpin Migrate Web Interface.
2. If you have not configured a Microsoft Azure Location as a migration target, click **Targets** > **Add Target**, and then configure the target Azure cloud platform.

   See Table 17-5, “Options for Microsoft Azure Location Target,” on page 286.
3. On the Workloads page, select the workload you want to configure.
4. Click **Configure Migration**.
5. Specify the **Initial Transfer Method** for replication based on the scope of data you want to transfer from the source to the target:
   - **Full Replication:** Migrate replicates the full volume from the source to the target.
   - **Incremental Replication:** Migrate replicates only differences in data from the source to the target, provided the workloads have similar operating system and volume profiles.

   NOTE: PlateSpin Migrate does not support Incremental Replication for the initial replication of data to existing target workloads in Azure Cloud. However, you can schedule Incremental Replications for subsequent replication of data. See **Incremental Recurrence** in Step 8.

   ---
6. Select an existing Microsoft Azure Location target platform to which you want to migrate the source workload.

   To verify the availability in the target location of Premium Blob Storage for unmanaged disks or of Managed Disks, refer to the *Microsoft Azure Products Available by Region*.
7. Click **Configure Migration**.

8 Configure the following settings:

Schedule Settings
Incremental Recurrence
Specify the time and pattern when you want to run incremental replications after the first full replication, or start each incremental replication manually. The default setting is None. The incremental replications are unscheduled. **To set or modify the incremental recurrence time and pattern:** 1. Click **Edit**. 2. For **Begin the recurrence schedule**, set the date and time when you want to begin the scheduled incremental replications. You can type the date (dd/mm/yyyy) or click the Calendar icon to select the date. By default, the run time is 12:00:00 a.m. (hh:mm:ss a.m. or p.m.). 3. For **Recurrence run setting**, set the pattern to follow for scheduled incremental replications: ◆ **Daily:** The replication takes place on the specified daily intervals or on weekdays every week for a period of 60 days from the time the replication starts. ◆ **Weekly:** The replication takes place at specified intervals for a period of 8 weeks from the time the replication starts. ◆ **Monthly:** The replication takes place at specified intervals for a period of 2 months from the time the replication starts. **NOTE:** ◆ Scheduled incremental replications are skipped until the first full replication is complete. ◆ Scheduled incremental replications take place for a maximum period of 60 days from the time that the scheduled incremental replication runs begin.
Full Replication
Specify when you want the first full replication to run, or start the first full replication manually. The first full replication is a one-time event, but the run is attempted daily as scheduled until the first replication begins and completes successfully. The default setting is None. The first full replication is unscheduled. **NOTE:** You must prepare the workload prior to the scheduled time or the manual start. The full replication cannot run unless the target VM exists and the workload preparation is complete. If they are not ready, Migrate skips the scheduled full replication and retries it at the scheduled time on the next day. **To set or modify the schedule for the first full replication:** 1. Click **Edit**. 2. Click **Start**, then set the date and time when you want to start the first full replication. You can type the date (dd/mm/yyyy) or click the Calendar icon to select the date. By default, the run time is 12:00:00 a.m. (hh:mm:ss a.m. or p.m.).

Blackout Window

Specify a replication blackout window that suspends scheduled replication activities for a specified period of time and pattern. For example, suspend replications during peak network utilization hours or to prevent conflicts between VSS-aware software and the PlateSpin VSS block-level data transfer component.

The default setting is None. No blackout window is scheduled.

To set or modify a blackout window:

1. Click **Edit**.
2. Specify the start and end time for the blackout period.

   The blackout start and end times are based on the system clock on the PlateSpin Server.
3. Select **Daily**, **Weekly**, or **Monthly** to enable a blackout window, then set the recurrence pattern.

Compression Level

This setting controls whether data is compressed during transmission between the source and target workloads, and the level of data compression applied.See "Data Compression" on page 52.

Select one of the following options:

- **None:** No compression.
- **Fast:** Consumes the least CPU resources on the source, but yields a lower compression ratio.
- **Optimal:** (Default) Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option.
- **Maximum:** Consumes the most CPU resources on the source, but yields a higher compression ratio.

Bandwidth Throttling

Bandwidth throttling enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. Throttling helps to prevent migration traffic from congesting your production network and to reduce the overall load of your PlateSpin Server. You can specify a throughput rate for each migration job.

Throttling is disabled by default with a Throttling Rate value of Off.

To throttle replications to a specified rate:

1. Specify a maximum throughput value in Mbps for data transfer for the workload.
2. Specify the throttling pattern:
   - **Always:** Always throttle data transfer for the replications.
   - **Custom:** Specify the time and days to throttle data transfer for the replications running in that window.

     Throttling time is local to the source workload.

Migration Settings

Transfer Method

(For Windows Workloads) Select a data transfer mechanism and security through encryption.See "Supported Data Transfer Methods" on page 44.

To enable encryption, select the **Encrypt Data Transfer** option. See "Security and Privacy" on page 46.

Transfer Encryption

(For Linux Workloads) To enable encryption, select the **Encrypt Data Transfer** option.See "Security and Privacy" on page 46.

Source Credentials

Specify the credentials required for accessing the workload. See "Discovery Guidelines for Source Workloads" on page 299.

Azure Resource Group

Specify a resource group to use for the target VM resources. Do one of the following:

- Select an existing resource group from the list.

  NOTE: An existing resource group is required for target VMs that you will add to an Azure Availability Set. Before you configure the workload, use the Azure portal to create an Availability Set in a resource group for the specified target location. The Availability Set definition determines whether you can use the set in combination with Azure Managed Disks.

- Allow PlateSpin to create a new resource group with the default name:

  `<hostname>-Resources`

- Type the name to use when PlateSpin creates a new resource group.

Azure Tags

Click **Edit Tags** to access the dialog where you can configure the tag names and values as appropriate for the workload.

You can specify up to 8 tags for the workload. Global tags are pre-populated automatically. Required global tags with fixed values are applied automatically to the workload and cannot be modified. You must specify a value for a required tag if its value is not fixed. For optional global tags, you can modify the tag name or value, or you can remove them. You can optionally specify other tags, up to the maximum of 8 tags per workload.

You cannot modify tags and values for the workload after a successful full replication occurs. Changes to the global settings for tag names and values will not apply to migrations in this state.

See "Using Azure Cloud Tags for Azure Migrations" on page 194.

Virtual Machine Name

Specify a display name for the new virtual machine.

Storage Layout

Specify the desired layout of data volumes (and volume groups) on the target workload disks:

- **Source-based**

  The Source-based storage layout assigns data volumes to disks according to their layout on the source workload.

- **Single disk**

  The Single disk storage layout assigns all the data volumes or LVM volume groups to a single disk on the target workload. An LVM volume group might contain multiple volumes.

- **One volume per disk**

  The One volume per disk storage layout assigns each data volume or LVM volume group to a separate disk on the target workload. An LVM volume group might contain multiple volumes.

Storage Options > Disks

Specify the information for disks, depending on whether the target VM uses Azure Managed Disks or uses Storage Accounts and unmanaged disks.

Azure Managed Disks

Managed disks are resource objects that represent the actual replicated disks being managed. The Azure Managed Disks setting applies to all disks for the target VM. Using Azure Managed Disks is optional and is enabled by default. Deselect the option to specify an Azure Storage Account and use unmanaged disks in Azure. See "Prerequisites for Azure Storage" on page 173.

To use Azure Managed Disks for the Target VM:

1. Select **Use Managed Disks** to enable Azure Managed Disks for the VM. This is the default setting.
2. Specify the **Storage Type** to use for hosting the managed disks for the target VM:
   - Standard HDD (hard disk drives)
   - Standard SSD (solid state drives)
     NOTE: Standard SSD option is not applicable for migrations to Azure Stack.
   - Premium SSD
3. (Optional) Under **Disks**, specify a custom name for one or more managed disk. You do not need to specify a storage account or a location in a storage account.

Managed Disk Names:

You can optionally provide a custom name for each managed disk. You cannot rename the managed disk after it is created. PlateSpin Migrate applies a default name for each managed disk in the following format:

```
<parent VM name>-<context>-<numeric suffix>
```

Valid characters for the resource name are alphanumeric, underscore (_), hyphen (-), and period (.). The name should start with a letter or number, and end with a letter, number, or underscore. The maximum name length is 80 characters. Managed disk names do not have a `.vhd` extension.

Unmanaged Disks:

If you disabled **Use Managed Disks**, specify the Datastore (Azure Storage Account) and the disk path to use in the datastore.

To use unmanaged disks for the Target VM:

1. Deselect **Use Managed Disks**.
2. Under **Disks**, specify the Datastore (Azure Storage Account) and Disk Path to use for the disk for the target VM.
3. (Optional) Specify a custom name for one or more disk.

Unmanaged Disk Names:

You can optionally provide a custom name for each unmanaged disk. Names must comply with the naming conventions for disks. Disk names have a .vhd extension.

Storage Options > Volumes

Select volumes to be included in the target for migration.

For Windows file-based migrations, you can optionally adjust the size of the target volume by modifying the **Size** value in the row for that volume.

NOTE: The ability to adjust the size of the target volume is not currently available in the Disk Layout view of storage in the Migrate Web Interface. See "Cannot Adjust the Target Volume Size for File-Based Migrations in the Web Interface" on page 569.

NTFS Cluster Size

(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see the Microsoft Support KB Article 140365.

Non-volume Storage

(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.

Disks For Volume Groups

(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.

NOTE: Disks that host both regular volumes (such as `/boot`) and volume groups are configured in the Disks section.

Volume Groups

(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the **Converted Logical Volumes** section of the settings.

Converted Logical Volumes

(For Linux Workloads) Select LVM logical volumes to be included in the target for migration.

Replication Network for Target

The replication NIC is the primary NIC that you specify in **Target Workload Settings**> **Network Connections**.

1. Select a network and subnet to use for replication traffic.
2. If you do not use an Azure VPN, the replication NIC requires a public IP address that is automatically assigned by Azure. Click **Edit**, then select **Create Public IP**.
3. Specify a resource group to use for the replication network. The **Azure Resource Group** setting is the default. To specify a different resource group, click **Edit** and do one of the following:
   - Type the name to use when PlateSpin creates a new resource group.
   - Select an existing resource group from the list.
4. Select one of the following:
   - **DHCP:** Obtain an IP address automatically assigned by a DHCP server.
   - **Static:** Specify a static private IP address that is unique within the supported subnet.

Replication Networks for Source

Specify one or more network interfaces (NIC or IP address) on the source workload to use for replication traffic that are valid for communications with the replication environment.

If the network for the NIC you specify is not part of your Azure VPN, ensure that the NIC has a public IP address.

Services to Stop Before Any Replication

(For Windows Workloads) We recommend that all the non-VSS compliant services or antivirus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.

Services to Stop for Cutover with Replication
(For Windows Workloads) Select the Windows services that should be permanently stopped on the source workload for cutover with any replication. The services stopped on the source workload during the replication process are not restored after Cutover. The stopped services are restored after a Test Cutover.
Daemons to Stop before Any Replication
(For Linux Workloads) Select the Linux services that you want to be temporarily stopped on the source workload before replication. These services will be restored back after replication completes.
Daemons to Stop for Cutover with Replication
(For Linux Workloads) Select the Linux services that should be permanently stopped on the source workload for Cutover with any Replication. The services stopped on the source workload during the replication process are not restored after Cutover. The stopped services are restored after a Test Cutover.
Boot Services to Disable on Target
(For Windows Workloads) Select the boot services that you want to disable on the target workload.
Target Workload Settings (These settings are applied during the Run Cutover)
Azure Agent
Specify whether to install the latest Azure Virtual Machine Agent on the target VM. If the check box is selected (enabled), Migrate installs the agent on the target workload at Cutover. If the Azure VM Agent is installed on the source, the agent will be uninstalled on the target workload and reinstalled using the latest version of the agent.
Azure Availability Set
An Azure Availability Set is an Azure service that provides high availability to two or more member VMs by deploying them in different *fault domains* and *update domains*. VMs must be created in the Availability Set. You cannot add an existing VM to an Availability Set. See "Prerequisites for Using Azure Availability Sets" on page 174. Using an Availability Set is optional. The default setting is `No Availability Set`. Click **Change Availability Set** to select an Availability Set from a list of sets available within the specified **Azure Resource Group** in Migration Settings. **NOTE:** After Azure assigns the physical hardware for the Availability Set, Azure limits the cloud instance sizes and networks that can be used for future member VMs. PlateSpin Migrate might prompt you to choose different cloud instance sizes or networks accordingly.

Cloud Instance Size

Select the cloud instance size appropriate for your workload and the storage account type for the target platform.

IMPORTANT: The Cloud Instance Size must be of the same storage type as the target account: HDD or SDD. Otherwise, you receive a validation error. To verify the availability in the target location of Premium Blob Storage for unmanaged disks or of Managed Disks, refer to the *Microsoft Azure Products Available by Region*.

By default, PlateSpin selects a cloud instance size that supports the same Storage account type and that most closely matches your source workload for the following components:

- Total number of cores
- Amount of memory
- Number of data disks
- Number of NICs

The default instance either meets or exceed the settings for each of these components on the source workload. However, you can choose a smaller instance size based on your requirements:

- The target VM uses the allowed CPU and memory for the instance size. To reduce the number of CPUs or amount of memory on the target workload:
  1. Select a smaller cloud instance size with fewer CPUs or less memory that best fits your needs.
- The target VM uses up to the maximum allowed number of data disks for the instance size. To migrate only some of the data disks:
  1. Select a smaller cloud instance size with fewer data disks that best fits your needs.
  2. Deselect the volumes that should not be migrated until the number of disks for migration fits the selected instance.
- The target VM uses up to the maximum allowed number of NICs for the instance size. To migrate only some of the NICs:
  1. Select a cloud instance size with fewer NICs that best fits your needs. At least one NIC is required.
  2. Under **Target Workload Settings**, deselect the NICs that should not be migrated until the number of NICs for migration fits the selected instance.

     NOTE: The number of data disks consumed by volumes on the target VM cannot exceed the maximum number of data disks supported by the selected cloud instance.
- In the **Cloud Instance Size** list, the **Supports SSD Storage** column indicates the storage account type of the instance: HDD Storage (No) or SDD Storage (Yes). Ensure that your new instance size supports the same storage account type as the target platform.

Hostname

Do one of the following:

- To retain the same host name, select **No Change**.
- To change the host name, select **Set To** and specify the new name.

  NOTE: An incremental replication is required if you change the host name at cutover.

Domain / Workgroup

(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays:

- **Workgroup:** *Workgroup_name*

  where *Workgroup_name* is the workgroup name to which the source belongs.

- **Domain:** *Domain_name*

  where *Domain_name* is the domain name to which the source belongs.

NOTE: An incremental replication is required if you change the domain or workgroup name at cutover.

Do one of the following depending on where you want the target workload to join:

- **When the source workload belongs to a workgroup:** Assume that the source workload belongs to a workgroup named **WorkGroup1**.
  - For the target workload to join the same workgroup (**WorkGroup1**), retain the following existing selection:

    Workgroup: Workgroup1
  - For the target workload to join a different workgroup (say **WorkGroup2)**, select **Join Workgroup** and specify the name as **WorkGroup2**.
  - For the target workload to join a domain, select **Join Domain** and specify the domain name you want the target to join.
- **When the source workload belongs to a domain:** Assume that the source workload belongs to a domain named **Domain1**.
  - For the target workload to join a workgroup, click **Join Workgroup** and specify the name of the workgroup you want the target to join.
  - For the target workload to join the same domain (**Domain1**) with the domain registration settings preserved, retain the following existing selection:

    Domain: Domain1
  - For the target workload to join the same domain (**Domain1)** without preserving the domain registration settings, select **Join Domain** and specify the domain name as **Domain1**.
  - For the target workload to join a different domain, select **Join Domain** and specify the domain name you want the target to join.

Domain Credentials

(For Windows Workloads) If you select **Join Domain**, specify the domain administrator credentials.

Network Connections

1. For workloads that have more than one NIC, select **Include** for each NIC to be migrated. Deselect **Include** to exclude a NIC.
   - At least one NIC is required.
   - The number of NICs to migrate cannot exceed the maximum number of NICs supported by the selected cloud instance.
2. For each included NIC, select a network and subnet.
3. Ensure that the Primary NIC is properly configured for its role as Primary. The default **Primary Connection** is the first NIC in the list. For more information, see "Azure Networking Guidelines" on page 180.
4. If you do not use an Azure VPN, the primary NIC requires a public IP address that is automatically assigned by a Azure. For the primary NIC, click **Edit**, then select **Create Public IP**.

   NOTE: PlateSpin Migrate uses the `EnablePublicIPForCloudInstance` configuration parameter to control whether or not a public IP address should be assigned to the primary network interface of Azure instances. If the value of `EnablePublicIPForCloudInstance` is true, then **Create Public IP** option is selected by default and you cannot deselect it. However, if the value of `EnablePublicIPForCloudInstance` is false, then **Create Public IP** option is deselected by default, but you can select it if you want Azure to automatically assign a public IP address to primary NIC on the target instance. See "Enabling Public IP Address for Target Instances on Cloud" on page 112.
5. For each included NIC:
   a. Specify a resource group to use for the NIC. The **Azure Resource Group** setting is the default. To specify a different resource group, click **Edit** and do one of the following:
      - Type the name to use when PlateSpin creates a new resource group.
      - Select an existing resource group from the list.
   b. Select one of the following:
      - **DHCP:** Obtain an IP address automatically assigned by a DHCP server.
      - **Static:** Specify one or more static IP addresses. The network mask and gateway are auto-populated based on the selected subnet. Click **Advanced** to add multiple IP addresses, edit an IP address, or remove an IP address. The IP addresses must be unique within the supported subnet.

DNS Servers

Specify the DNS Servers for the target workloads. This is applicable only if you select **Static** in the **Network Connections** option:

- **Primary DNS server:** Specify the primary DNS server address.
- **Alternative DNS server:** Specify an alternate DNS server address.
- **Additional DNS server:** To specify additional DNS server addresses:
  1. Click **Advanced**.
  2. Specify the DNS server address.
  3. Click **Add** to add the server in the DNS Server Addresses list.
  4. Click **OK**.

Services States on Target VM

(For Windows Workloads) Select Windows services' start conditions on the target VM. Start options are **Automatic**, **Manual**, **Disabled**, and **Automatic (Delayed Start).**

Daemons States to Change
(For Linux Workloads) Select Linux daemons' start conditions on the target VM. Enable the daemon to start by selecting the check boxes at the appropriate runlevels (0 to 6) and Boot.
Post-Migration Action
Click **Configure Action** to select a previously configured custom action that you want to execute on your target workload as a part of the migration job. The command line parameters and timeout value that you previously configured for the selected action is displayed by default. You can edit these values as required. If you choose to timeout the action and have not specified any timeout value, then PlateSpin Migrate waits for a maximum period of 7200 seconds before timing out this action. For information about managing the post-migrations, see "Managing Post-Migration Actions" on page 107.
Target Workload Test Settings (These settings are applied during the Test Cutover)
Copy Target Workload Settings
Click the **Copy Target Workload Settings** option to automatically copy the workload settings from **Target Workload Settings** section to **Target Workload Test Settings** section.
Azure Agent
Specify whether to install the latest Azure Virtual Machine Agent on the target VM. If the check box is selected (enabled), Migrate installs the agent on the target workload at Test Cutover. If the Azure VM Agent is installed on the source, the agent will be uninstalled on the target workload and reinstalled using the latest version of the agent.
Azure Availability Set
An Azure Availability Set is an Azure service that provides high availability to two or more member VMs by deploying them in different *fault domains* and *update domains*. VMs must be created in the Availability Set. You cannot add an existing VM to an Availability Set. See "Prerequisites for Using Azure Availability Sets" on page 174. Using an Availability Set is optional. The default setting is `No Availability Set`. Click **Change Availability Set** to select an Availability Set from a list of sets available within the specified **Azure Resource Group** in Migration Settings. **NOTE:** After Azure assigns the physical hardware for the Availability Set, Azure limits the cloud instance sizes and networks that can be used for future member VMs. PlateSpin Migrate might prompt you to choose different cloud instance sizes or networks accordingly.

Cloud Instance Size

Select the cloud instance size appropriate for your workload and the storage account type for the target platform.

IMPORTANT: The Cloud Instance Size must be of the same storage type as the target account: HDD or SDD. Otherwise, you receive a validation error. To verify the availability in the target location of Premium Blob Storage for unmanaged disks or of Managed Disks, refer to the *Microsoft Azure Products Available by Region*.

By default, PlateSpin selects a cloud instance size that supports the same Storage account type and that most closely matches your source workload for the following components:

- Total number of cores
- Amount of memory
- Number of data disks
- Number of NICs

The default instance either meets or exceed the settings for each of these components on the source workload. However, you can choose a smaller instance size based on your requirements:

- The target VM uses the allowed CPU and memory for the instance size. To reduce the number of CPUs or amount of memory on the target workload:
  1. Select a smaller cloud instance size with fewer CPUs or less memory that best fits your needs.
- The target VM uses up to the maximum allowed number of data disks for the instance size. To migrate only some of the data disks:
  1. Select a smaller cloud instance size with fewer data disks that best fits your needs.
  2. Deselect the volumes that should not be migrated until the number of disks for migration fits the selected instance.
- The target VM uses up to the maximum allowed number of NICs for the instance size. To migrate only some of the NICs:
  1. Select a cloud instance size with fewer NICs that best fits your needs. At least one NIC is required.
  2. Under **Target Workload Settings**, deselect the NICs that should not be migrated until the number of NICs for migration fits the selected instance.

     NOTE: The number of data disks consumed by volumes on the target VM cannot exceed the maximum number of data disks supported by the selected cloud instance.
- In the **Cloud Instance Size** list, the **Supports SDD Storage** column indicates the storage account type of the instance: HDD Storage (No) or SDD Storage (Yes). Ensure that your new instance size supports the same storage account type as the target platform.

Hostname

Do one of the following:

- To retain the same host name, select **No Change**.
- To change the host name, select **Set To** and specify the new name.

  NOTE: An incremental replication is not required if you change the host name at test cutover.

Domain / Workgroup
(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays: ◆ **Workgroup:** *Workgroup_name* where *Workgroup_name* is the workgroup name to which the source belongs. ◆ **Domain:** *Domain_name* where *Domain_name* is the domain name to which the source belongs. **NOTE:** An incremental replication is not required if you change the domain or workgroup name at test cutover. Do one of the following depending on where you want the target workload to join: ◆ **When the source workload belongs to a workgroup:** Assume that the source workload belongs to a workgroup named **WorkGroup1**. ◆ For the target workload to join the same workgroup (**WorkGroup1**), retain the following existing selection: **Workgroup: Workgroup1** ◆ For the target workload to join a different workgroup (say **WorkGroup2)**, select **Join Workgroup** and specify the name as **WorkGroup2**. ◆ For the target workload to join a domain, select **Join Domain** and specify the domain name you want the target to join. ◆ **When the source workload belongs to a domain:** Assume that the source workload belongs to a domain named **Domain1**. ◆ For the target workload to join a workgroup, click **Join Workgroup** and specify the name of the workgroup you want the target to join. ◆ For the target workload to join the same domain (**Domain1**) with the domain registration settings preserved, retain the following existing selection: **Domain: Domain1** ◆ For the target workload to join the same domain (**Domain1)** without preserving the domain registration settings, select **Join Domain** and specify the domain name as **Domain1**. ◆ For the target workload to join a different domain, select **Join Domain** and specify the domain name you want the target to join.
Domain Credentials
(For Windows Workloads) If you select **Join Domain**, specify the domain administrator credentials.

Network Connections

Available NICs match the included NICs in **Target Workload Settings** > **Network Connections**.

1. For each included NIC, select a network and subnet.
2. Ensure that the Primary NIC is properly configured for its role as Primary. The default **Primary Connection** is the first NIC in the list. For more information, see "Azure Networking Guidelines" on page 180.
3. If you do not use an Azure VPN, the primary NIC requires a public IP address that is automatically assigned by a Azure. For the primary NIC, click **Edit**, then select **Create Public IP**.

   NOTE: PlateSpin Migrate uses the `EnablePublicIPForCloudInstance` configuration parameter to control whether or not a public IP address should be assigned to the primary network interface of Azure instances. If the value of `EnablePublicIPForCloudInstance` is true, then **Create Public IP** option is selected by default and you cannot deselect it. However, if the value of `EnablePublicIPForCloudInstance` is false, then **Create Public IP** option is deselected by default, but you can select it if you want Azure to automatically assign a public IP address to primary NIC on the target instance. See "Enabling Public IP Address for Target Instances on Cloud" on page 112.
4. For each included NIC:
   a. Specify a resource group to use for the NIC. The **Azure Resource Group** setting is the default. To specify a different resource group, click **Edit** and do one of the following:
      - Type the name to use when PlateSpin creates a new resource group.
      - Select an existing resource group from the list.
   b. Select one of the following:
      - **DHCP:** Obtain an IP address automatically assigned by a DHCP server.
      - **Static:** Specify one or more static IP addresses. The network mask and gateway are auto-populated based on the selected subnet. Click **Advanced** to add multiple IP addresses, edit an IP address, or remove an IP address. The IP addresses must be unique within the supported subnet.

DNS Servers

Specify the DNS Servers for the target workloads. This is applicable only if you select **Static** in the **Network Connections** option:

- **Primary DNS server:** Specify the primary DNS server address.
- **Alternative DNS server:** Specify an alternate DNS server address.
- **Additional DNS server:** To specify additional DNS server addresses:
  1. Click **Advanced**.
  2. Specify the DNS server address.
  3. Click **Add** to add the server in the DNS Server Addresses list.
  4. Click **OK**.

Services States on Target VM

(For Windows Workloads) Select Windows services that must be automatically stopped on the target VM.

Daemons States to Change

(For Linux Workloads) Select Linux daemons that must be automatically stopped on the target VM.

Post-Migration Action
Click **Configure Action** to select a previously configured custom action that you want to execute on your target workload as a part of the migration job. The command line parameters and timeout value that you previously configured for the selected action is displayed by default. You can edit these values as required. If you choose to timeout the action and have not specified any timeout value, then PlateSpin Migrate waits for a maximum period of 7200 seconds before timing out this action. For information about managing the post-migrations, see "Managing Post-Migration Actions" on page 107.
Tag
Tag
Select a tag to assign to the workload. For more information about tags, see "Using Tags to Track Logical Associations of Workloads" on page 309.

9 (Optional) To change the target, click **Change Target**.

NOTE: If you change the target, all the settings you specified will be cleared.

10 Do one of the following:

- Click **Save** to save the settings.
- Click **Save and Prepare** to save the settings and start preparing the workload migration.
- Click **Cancel** to exit.

27 Migration to VMware vCloud Director

Planning for Migration to VMware vCloud Director

Before you begin migrations to your cloud environment in VMware vCloud Director, ensure that your migration environment meets the following guidelines:

Supported Cloud Platforms

- See "VMware vCloud Director" in "Supported Target Cloud Platforms" on page 43.

Supported Workloads

- See "Supported Workloads For Migration to VMware vCloud Director" on page 31, as appropriate for the target VMware vCloud platform.

Network Access and Communications

- See "Access and Communication Requirements across Your Migration Network" on page 53.

Prerequisites

- See Chapter 10, "Prerequisites for Migration to VMware vCloud Director," on page 201.
- See Chapter 12, "Prerequisites for Cloud-to-Cloud Migrations," on page 217.

Targets and Workloads

- **Target VMware vCloud Organization (automated):** See "Target Discovery in the Web Interface" on page 284 using Table 17-9, "Options for VMware vCloud Organization Target," on page 288.
- **Source Workloads:** Use either of the following discovery methods:
  - "Workload Discovery in the Migrate Web Interface" on page 302
  - "Registering Workloads and Discovering Details with Migrate Agent" on page 303

Additional Information

- "Working with Virtual Machines" in the *VMware vCloud Director 5.6 Documentation Center*.

Configuring Migration of a Workload to VMware vCloud Director

When you add or discover a workload, the workload is listed on the Workloads page and the status is set as **Not Configured**. Before you migrate the workload, you must configure the workload for migration:

1. Launch the PlateSpin Migrate Web Interface.
2. If you have not configured a VMware vCloud Organization as a migration target, click **Targets** > **Add Target**, and then configure the target vCloud platform.

   See Table 17-9, “Options for VMware vCloud Organization Target,” on page 288.
3. On the Workloads page, select the workload you want to configure.
4. Click **Configure Migration**.
5. Select one of the following based on the scope of data you want to transfer from the source to the target:
   - **Full Replication:** A full volume of data transfer takes place from the source to the target.
   - **Incremental Replication:** Only differences are transferred from the source to the target, provided they have similar operating system and volume profiles.

   NOTE: PlateSpin Migrate does not support Incremental Replication for the initial replication of data to existing target workloads in VMware vCloud Director. However, you can schedule Incremental Replications for subsequent replication of data. See **Incremental Recurrence** in Step 8.

   ---
6. Select an existing vCloud target platform to which you want to migrate the source workload.
7. Click **Configure Migration**.

8 Configure the following settings. Ensure that the IP address for the source workload, the replication network for target, the cutover network, and the test cutover network are all different.

Schedule Settings
Incremental Recurrence

Specify the following:
◆ **Start of Recurrence:** The date when you want to start the replication. You can specify the date or click the calendar icon to select the date. By default, the time is 12:00 a.m.
◆ **Recurrence Pattern:** The pattern to follow for the recurrence of the replication. For example:
  ◆ To use incremental recurrence everyday, select **Daily.**
  ◆ To never use incremental recurrence, select **None.**
NOTE
◆ Scheduled incremental replications are skipped until the first full replication is complete.
◆ When you schedule incremental recurrence, the replication takes place for a maximum period of 60 days from the starting time of replication. For example:
  ◆ If you select **Daily**, then the replication takes place for 60 days from the time the replication starts.
  ◆ If you select **Weekly**, then the replication takes place for 8 weeks from the time the replication starts.
  ◆ If you select **Monthly**, then the replication takes place for 2 months from the time the replication starts.

Full Replication
Do one of the following: ◆ To specify a schedule for the replication, click **Start** and specify the date when you want to start the full replication. ◆ To start full replication manually without setting a schedule, click **None**. **NOTE:** You must prepare the workload prior to the scheduled time.The full replication cannot run unless the target VM exists and the workload preparation is complete. Migrate skips the scheduled full replication and retries it at the next scheduled time.
Blackout Window
Use these settings to force a replication blackout. The replication blackout suspends scheduled replications during peak utilization hours or prevents conflicts between VSS-aware software and the PlateSpin VSS block-level data transfer component. To specify a blackout window, click **Edit** and do the following: ◆ Specify the start and end time for the blackout period. ◆ Select one of the blackout recurrence pattern such as daily, weekly, or monthly. If you do not want to force a replication blackout, select **None**. **NOTE:** The blackout start and end times are based on the system clock on the PlateSpin Server.

Compression Level

These settings control whether data is compressed during transmission between the source and target workloads, and the level of data compression applied.See “Data Compression” on page 52. Select one of the following options:

- **Fast:** Consumes the least CPU resources on the source, but yields a lower compression ratio.
- **Optimal:** Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option.
- **Maximum:** Consumes the most CPU resources on the source, but yields a higher compression ratio.

Bandwidth Throttling

These settings control the bandwidth throttling. PlateSpin Migrate enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. You can specify a throughput rate for each migration job. Throttling provides a way to prevent migration traffic from congesting your production network and to reduce the overall load of your PlateSpin Server.

To throttle replications to a specified rate, specify the required throughput value in Mbps and the time pattern.

Migration Settings

Transfer Method

(For Windows Workloads) Select a data transfer mechanism and security through encryption.See “Supported Data Transfer Methods” on page 44.

To enable encryption, select the **Encrypt Data Transfer** option. See “Security and Privacy” on page 46.

Transfer Encryption

(For Linux Workloads) To enable encryption, select the **Encrypt Data Transfer** option.See “Security and Privacy” on page 46.

Source Credentials

Specify the credentials required for accessing the workload. See “Discovery Guidelines for Source Workloads” on page 299.

CPU

(For migration to vCloud and VM platforms using VMware 5.1, 5.5, and 6.0 with a minimum VM hardware Level 8) Specify the number of sockets and the number of cores per socket for the target workload. It automatically calculates the total cores. This parameter applies on the initial setup of a workload with an initial replication setting of **Full Replication.**

NOTE: The maximum number of cores the workload can use is subject to external factors such as the guest operating system, the VM hardware version, VMware licensing for the ESXi host, and ESXi host compute maximums for vSphere (see *ESXi/ESX Configuration Maximums (VMware KB 1003497)* (https://kb.vmware.com/kb/1003497)).

Some distributions of a guest OS might not honor the cores and cores per socket configuration. Instead they will retain their original cores and sockets settings as installed.

Organization Virtual Data Center

(For migration to vCloud) Select a virtual data center associated with your organization.

vApp

Specify a name for the VMware vApp.

Virtual Machine Name

Specify a display name for the new virtual machine.

Storage Layout

Specify the desired layout of data volumes (and volume groups) on the target workload disks:

- **Source-based**

  The Source-based storage layout assigns data volumes to disks according to their layout on the source workload.

- **Single disk**

  The Single disk storage layout assigns all the data volumes or LVM volume groups to a single disk on the target workload. An LVM volume group might contain multiple volumes.

- **One volume per disk**

  The One volume per disk storage layout assigns each data volume or LVM volume group to a separate disk on the target workload. An LVM volume group might contain multiple volumes.

Storage Options > Disks

Specify the path to the hard disk on the target virtual machine.

Storage Options > Volumes

Select volumes to be included in the target for migration.

For Windows file-based migrations, you can optionally adjust the size of the target volume by modifying the **Size** value in the row for that volume.

NOTE: The ability to adjust the size of the target volume is not currently available in the Disk Layout view of storage in the Migrate Web Interface. See "Cannot Adjust the Target Volume Size for File-Based Migrations in the Web Interface" on page 569.

NTFS Cluster Size

(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see Microsoft Support KB Article 140365.

Non-volume Storage

(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.

Disks For Volume Groups

(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.

Volume Groups

(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the **Converted Logical Volumes** section of the settings.

NOTE: Disks that host both regular volumes (such as `/boot`) and volume groups are configured in the Disks section.

Converted Logical Volumes

(For Linux Workloads) Specify one or more LVM logical volumes to be migrated for a Linux workload.

Replication Network for Target

By default, the replication NIC is the primary NIC that you specify in **Target Workload Settings**> **Network Connections**. Specify a network interface (NIC or IP address) on the target to use for replication traffic.

1. Select a network to use for replication traffic.
2. Select one of the following:
   - **DHCP:** Obtain an IP address automatically assigned by a DHCP server.
   - **Static - Manual:** Specify a static IP address that is unique within the supported subnet.
   - **Static - IP Pool:** Select this option to automatically issue IP address from the IP pool.
3. Specify an MTU value that the PlateSpin Migrate Linux RAM Disk (LRD) replication network can use. Setting a low value helps to avoid jabber over networks. For example: a VPN.

   The default value is an empty string. When networking is configured in the LRD, it allows the network device to set its own default, which is usually 1500. However, if you specify a value, PlateSpin Migrate adjusts the MTU when it configures the network interface.

Replication Networks for Source

Select one or more network interfaces (NIC or IP address) on the source workload to use for replication traffic that are valid for communications with the replication environment.

Services to Stop Before Any Replication

(For Windows Workloads) We recommend that all the non-VSS compliant services or antivirus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.

Services to Stop for Cutover with Replication

(For Windows Workloads) Select the Windows services that should be permanently stopped on the source workload for cutover with any replication. The services stopped on the source workload during the replication process are not restored after Cutover. The stopped services are restored after a Test Cutover.

Daemons to Stop Before Any Replication

(For Linux Workloads) Select the Linux services that you want to be temporarily stopped on the source workload before replication. These services will be restored back after replication completes.

Daemons to Stop for Cutover with Replication

(For Linux Workloads) Select the Linux services that should be permanently stopped on the source workload for Cutover with any Replication. The services stopped on the source workload during the replication process are not restored after Cutover. The stopped services are restored after a Test Cutover.

Boot Services to Disable on Target

(For Windows Workloads) Select the boot services that you want to disable on the target workload.

Target Workload Settings

(These settings are applied during the Run Cutover)

VM Memory

Specify the amount of memory allocated to the target workload.

VM Tools

To install the VM tools, select the **Install VM Tools option**. This option is selected by default.

Hostname

Do one of the following:

- To retain the same host name, select **No Change**.
- To change the host name, select **Set To** and specify the new name.

  NOTE: An incremental replication is required if you change the host name at cutover.

Domain / Workgroup

(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays:

- **Workgroup:** *`Workgroup_name`*

  where *Workgroup_name* is the workgroup name to which the source belongs.
- **Domain:** *Domain_name*

  where *Domain_name* is the domain name to which the source belongs.

NOTE: An incremental replication is required if you change the domain or workgroup at cutover.

Do one of the following depending on where you want the target workload to join:

- **When the source workload belongs to a workgroup:** Assume that the source workload belongs to a workgroup named **WorkGroup1**.
  - For the target workload to join the same workgroup (**WorkGroup1**), retain the following existing selection:

    Workgroup: Workgroup1
  - For the target workload to join a different workgroup (say **WorkGroup2)**, select **Join Workgroup** and specify the name as **WorkGroup2**.
  - For the target workload to join a domain, select **Join Domain** and specify the domain name you want the target to join.
- **When the source workload belongs to a domain:** Assume that the source workload belongs to a domain named **Domain1**.
  - For the target workload to join a workgroup, click **Join Workgroup** and specify the name of the workgroup you want the target to join.
  - For the target workload to join the same domain (**Domain1**) with the domain registration settings preserved, retain the following existing selection:

    Domain: Domain1
  - For the target workload to join the same domain (**Domain1)** without preserving the domain registration settings, select **Join Domain** and specify the domain name as **Domain1**.
  - For the target workload to join a different domain, select **Join Domain** and specify the domain name you want the target to join.

Domain Credentials

(For Windows Workloads) If you select **Join Domain**, specify the domain administrator credentials.

Network Connections

1. For workloads that have more than one NIC, select **Include** for each NIC to be migrated. Deselect **Include** to exclude a NIC.
   - At least one NIC is required.
   - The number of NICs to migrate cannot exceed the maximum number of NICs supported by the selected cloud instance.
2. Ensure that the Primary NIC is properly configured for its role as Primary. The default **Primary Connection** is the first NIC in the list. To set a different NIC as Primary NIC, click **Edit** for the corresponding NIC and select **Primary Connection** for that NIC.
3. For each included NIC:
   a. Select **Start Connected** to connect the virtual network interface when starting the target workload.
   b. Select a network.
   c. (Conditional) To set the NIC as Primary NIC, click **Edit** and select **Primary Connection.** This resets the **Primary Connection** for the previously set Primary NIC.
   d. Select one of the following:
      - **DHCP:** Obtain an IP address automatically assigned by a DHCP server.
      - **Static:** Specify one or more static IP addresses. Click **Advanced** to add multiple IP addresses, edit an IP address, or remove an IP address. The IP addresses must be unique within the supported network.

DNS Servers

(For Linux Workloads) Specify the DNS Servers for the target workloads. This is applicable only if you select **Static** in the **Network Connections** option:

- **Primary DNS server:** Specify the primary DNS server address.
- **Alternative DNS server:** Specify an alternate DNS server address.
- **Additional DNS server:** To specify additional DNS server addresses:
  1. Click **Advanced**.
  2. Specify the DNS server address.
  3. Click **Add** to add the server in the DNS Server Addresses list.
  4. Click **OK**.

Services States on Target VM

(For Windows Workloads) Select Windows services that must be automatically stopped on the target VM.

Daemons States on Target VM

(For Linux Workloads) Select Linux daemons that must be automatically stopped on the target VM.

Post-Migration Action
Click **Configure Action** to select a previously configured custom action that you want to execute on your target workload as a part of the migration job. The command line parameters and timeout value that you previously configured for the selected action is displayed by default. You can edit these values as required. If you choose to timeout the action and have not specified any timeout value, then PlateSpin Migrate waits for a maximum period of 7200 seconds before timing out this action. For information about managing the post-migrations, see "Managing Post-Migration Actions" on page 107.
Target Workload Test Settings (These settings are applied during the Test Cutover)
Copy Target Workload Settings
Click the **Copy Target Workload Settings** option to automatically copy the workload settings from **Target Workload Settings** section to **Target Workload Test Settings** section.
VM Memory
Specify the amount of memory allocated to the target workload.
VM Tools
To install the VM tools, select the **Install VM Tools option**. This option is selected by default.
Hostname
Do one of the following: ◆ To retain the same host name, select **No Change**. ◆ To change the host name, select **Set To** and specify the new name. **NOTE:** An incremental replication is not required if you change the host name at test cutover.

Domain / Workgroup

(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays:

- **Workgroup:** *Workgroup_name*

  where *Workgroup_name* is the workgroup name to which the source belongs.

- **Domain:** *Domain_name*

  where *Domain_name* is the domain name to which the source belongs.

NOTE: An incremental replication is not required if you change the domain or workgroup at test cutover.

Do one of the following depending on where you want the target workload to join:

- **When the source workload belongs to a workgroup:** Assume that the source workload belongs to a workgroup named **WorkGroup1**.
  - For the target workload to join the same workgroup (**WorkGroup1**), retain the following existing selection:

    Workgroup: Workgroup1
  - For the target workload to join a different workgroup (say **WorkGroup2)**, select **Join Workgroup** and specify the name as **WorkGroup2**.
  - For the target workload to join a domain, select **Join Domain** and specify the domain name you want the target to join.
- **When the source workload belongs to a domain:** Assume that the source workload belongs to a domain named **Domain1**.
  - For the target workload to join a workgroup, click **Join Workgroup** and specify the name of the workgroup you want the target to join.
  - For the target workload to join the same domain (**Domain1**) with the domain registration settings preserved, retain the following existing selection:

    Domain: Domain1
  - For the target workload to join the same domain (**Domain1)** without preserving the domain registration settings, select **Join Domain** and specify the domain name as **Domain1**.
  - For the target workload to join a different domain, select **Join Domain** and specify the domain name you want the target to join.

Domain Credentials

(For Windows Workloads) If you select **Join Domain**, specify the domain administrator credentials.

Network Connections

Available NICs match the included NICs in **Target Workload Settings** > **Network Connections**. The default **Primary Connection** is the first NIC in the list.

1. For each included NIC:
   a. Select **Start Connected** to connect the virtual network interface when starting the target workload.
   b. Select a network.
   c. (Conditional) To set the NIC as Primary NIC, click **Edit** and select **Primary Connection.** This resets the **Primary Connection** for the previously set Primary NIC.
   d. Select one of the following:
      - **DHCP:** Obtain an IP address automatically assigned by a DHCP server.
      - **Static:** Specify one or more static IP addresses. Click **Advanced** to add multiple IP addresses, edit an IP address, or remove an IP address. The IP addresses must be unique within the supported network.

DNS Servers

Specify the DNS Servers for the target workloads. This is applicable only if you select **Static** in the **Network Connections** option:

- **Primary DNS server:** Specify the primary DNS server address.
- **Alternative DNS server:** Specify an alternate DNS server address.
- **Additional DNS server:** To specify additional DNS server addresses:
  1. Click **Advanced**.
  2. Specify the DNS server address.
  3. Click **Add** to add the server in the DNS Server Addresses list.
  4. Click **OK**.

Services States on Target VM

(For Windows Workloads) Select Windows services that must be automatically stopped on the target VM.

Daemons States to Change

(For Linux Workloads) Select Linux daemons that must be automatically stopped on the target VM.

Post-Migration Action

Click **Configure Action** to select a previously configured custom action that you want to execute on your target workload as a part of the migration job. The command line parameters and timeout value that you previously configured for the selected action is displayed by default. You can edit these values as required. If you choose to timeout the action and have not specified any timeout value, then PlateSpin Migrate waits for a maximum period of 7200 seconds before timing out this action. For information about managing the post-migrations, see "Managing Post-Migration Actions" on page 107.

Tag

Tag

Select a tag to assign to the workload. For more information about tags, see "Using Tags to Track Logical Associations of Workloads" on page 309.

9 (Optional) To change the target, click **Change Target**.

NOTE: If you change the target, all the settings you specified will be cleared.

10 Do one of the following:

- Click **Save** to save the settings.
- Click **Save and Prepare** to save the settings and start preparing the workload migration.
- Click **Cancel** to exit.

28 Migration to Oracle Cloud Infrastructure

Planning for Migration to Oracle Cloud Infrastructure

Before you begin migrations to the Oracle Cloud Infrastructure, ensure that your migration environment meets the following guidelines:

Supported Cloud Platforms

- See Table 2-16, "Supported Target Cloud Platforms for the Migrate Client," on page 44 and Table 2-15, "Supported Target Cloud Platforms for the Migrate Web Interface," on page 43.

Supported Workloads

- See "Supported Workloads For Migration to Oracle Cloud Infrastructure" on page 34, as appropriate for the target Oracle Cloud environment.

Network Access and Communications

- See "Access and Communication Requirements across Your Migration Network" on page 53.

Prerequisites

- See Chapter 11, "Prerequisites for Migration to Oracle Cloud Infrastructure," on page 209.

Targets and Workloads

- **Target VM on Virtual Hosts on Oracle Cloud (semi-automated):** See "Registering the Target Virtual Machine Instance with PlateSpin Server" on page 483.
- **Source Workloads:**
  - **Using Migrate Client:** See "Workload Discovery in the Migrate Client" on page 301.
  - **Using Migrate Web Interface:** See "Workload Discovery in the Migrate Web Interface" on page 302.

Additional Information

- Oracle Cloud Infrastructure Compute Cloud Documentation (https://docs.cloud.oracle.com/iaas/Content/Compute/Concepts/computeoverview.htm).

Configuring Full Migration of a Workload to Oracle Cloud VM Instance Using the X2P Workflow

PlateSpin Migrate uses X2P workflow to support full migration of Windows and Linux workloads to Oracle Cloud Infrastructure.

IMPORTANT: We recommend that no changes are made on the source workload during the full migration process to avoid any changes that are not replicated to the target workload.

Use the guidelines in this section to configure migration to VMs on Oracle Cloud Infrastructure virtual hosts.

Discovering Source Workloads

Before you configure a migration job to migrate workloads to Oracle Cloud, you must discover and perform an inventory of the source workload you want to migrate. See "Workload Discovery in the Migrate Client" on page 301 and "Workload Discovery in the Migrate Web Interface" on page 302.

Creating and Connecting to the Target Virtual Machine Instance in Oracle Cloud Infrastructure Environment

Use the custom PlateSpin image that you imported from the Oracle Object Storage to create a target VM instance on the Oracle Cloud Infrastructure virtual host. You can then connect to this created instance and register the instance with PlateSpin Migrate server. For information about importing the custom PlateSpin image, see "Importing PlateSpin Custom Images to Oracle Cloud Infrastructure" on page 214.

1. Sign in to the Oracle Cloud Console (https://cloud.oracle.com/).
2. Open the navigation menu. Go to **Compute** and click **Instances**. Choose your **Compartment** and then click **Create Instance**.
3. In the Create Compute Instance dialog box, you specify the resources to use for your instance. By default, your instance launches in the selected compartment, and the resources you choose also come from the selected compartment.

NOTE

- For the **Choose an operating system or image source** option, select the name of the custom image you previously imported from Oracle Object Storage. See "Importing PlateSpin Custom Images to Oracle Cloud Infrastructure" on page 214.
- For the **Configure boot volume** option, you must ensure to provide the boot volume size that is large enough to recreate the volumes from the source workload.

For more information about the resources, see Oracle Cloud Infrastructure Documentation (https://docs.cloud.oracle.com/iaas/Content/Compute/Tasks/launchinginstance.htm#UsingtheConsole).

4 Click **Create**. The created instance displays in the instance list.

NOTE: The created VM instance has a single disk and no additional disks or volumes can be added until the migration has completed. You can specify the disk size, up to 16 TB, when the VM instance is deployed.

5 When the instance is fully provisioned and running, create a console connection for this created instance so that you can use SSH to connect to it. For information about creating the console connection for the instance, see Oracle Cloud Infrastructure Documentation (https://docs.cloud.oracle.com/iaas/Content/Compute/References/serialconsole.htm#InstanceConsoleConnections).

6 Register the target VM instance with PlateSpin Migrate server. See "Registering the Target Virtual Machine Instance with PlateSpin Server" on page 483.

Registering the Target Virtual Machine Instance with PlateSpin Server

After you use the PlateSpin Custom Image to create a virtual machine instance on the Oracle Cloud Infrastructure and connect to it using SSH, you are ready to register the target VM instance with your PlateSpin Server:

IMPORTANT: Use the PlateSpin Migrate Web Interface to ensure that no value is set for the `AlternateServerAddress` PlateSpin Configuration parameter. See "Configuring Alternate IP Addresses for PlateSpin Server" on page 111.

1 In the Oracle Cloud Console (https://cloud.oracle.com/), select the instance that you previously created on Oracle Cloud and click **Reboot** to shutdown the target instance and restart it. See Oracle Cloud Infrastructure Documentation (https://docs.cloud.oracle.com/iaas/Content/Compute/Tasks/restartinginstancepool.htm).

NOTE: If rebooting the VM instance using the Oracle Cloud Infrastructure Console takes too long, then the SSH session to the target VM instance might time out. In such cases, reset the VM instance using the Oracle Cloud Infrastructure Command Line Interface.

2 In the SSH session that you are using to connect to the target VM instance, type one of the following boot options when prompted and press Enter:

Boot Option	Boot Action
ps	PlateSpin Linux for taking control You can also press Enter to select this option.
fcoe	PlateSpin Linux for taking control with FCoE support
next	Boot from the next boot device set in the BIOS

If no key is pressed for 20 seconds, the workload boots from the next boot device set in the BIOS.

3 At the command line, provide the required information at each individual prompt:

- **PlateSpin Server:** Enter the PlateSpin Server URL, using the following format:

  `http://`*Your_PlateSpin_Server*`/platespinmigrate`

  Replace *Your_PlateSpin_Server* with the host name or the IP address of your PlateSpin Server host.
- **Credentials (User Name/Password):** Enter the name of an administrator-level user on the PlateSpin Server host, including the domain or machine name. For example: *domain**username*, or *localhost*\Administrator. Provide a valid password for the specified user.
- **Network Card:** Select the network card that is active, then either enter a temporary static IP address for this NIC or press Enter to dynamically obtain an IP address from a DHCP server.
- **Temporary hostname:** Provide a temporary VM name for PlateSpin Server to use to list the newly registered VM. The workload's target host name you select in the migration job overwrites this name.
- **SSL encryption:** If your PlateSpin Migrate is installed on a host with SSL encryption enabled, enter `Yes`. If not, enter `No`.
- **PlateSpin Migrate Network:** Unless you have defined your own PlateSpin Migrate Network in PlateSpin Migrate Client, press Enter. If you are working with a non-default PlateSpin Migrate Network, type its name, then press Enter.

  A controller on your target virtual machine communicates with PlateSpin Server and registers the virtual machine as a physical target for a migration job.

The registered target virtual machine displays in the Servers view of the PlateSpin Migrate Client Interface and in the PlateSpin Migrate Web Interface and you can now configure migration jobs to this target VM. See "Migrating the Source Workload to a Target Instance on Oracle Cloud Infrastructure" on page 485.

Migrating the Source Workload to a Target Instance on Oracle Cloud Infrastructure

1. (For non-VPN setup) After you register the target virtual machine with Migrate Server, you must set the `AlternateServerAddress` PlateSpin Configuration parameter to the public IP addresses of the Migrate server. For information about configuring the AlternateServerAddress parameter, see "Configuring Alternate IP Addresses for PlateSpin Server" on page 111.
2. Use either the PlateSpin Migrate Web Interface (recommended) or the PlateSpin Migrate Client to migrate the workload to the target instance:
   - **Using PlateSpin Migrate Web Interface:** See "Configuring Migration to a Physical Target (P2P, V2P) Using Migrate Web Interface" on page 521.

     NOTE: The Oracle Cloud target VM instance is already registered with your PlateSpin Server in "Registering the Target Virtual Machine Instance with PlateSpin Server" on page 483. When you use the PlateSpin Migrate Web Interface to configure migrations to the physical target, ignore the step that requires you to register your target physical machine with PlateSpin Migrate by booting it with the PlateSpin ISO.
   - **Using PlateSpin Migrate Client:** See "Configuring Migration to a Physical Target (P2P, V2P) Using Migrate Client" on page 516.

     NOTE: The Oracle Cloud target VM instance is already registered with your PlateSpin Server in "Registering the Target Virtual Machine Instance with PlateSpin Server" on page 483. When you use the PlateSpin Migrate Client to configure migrations to the physical target, ignore the step that requires you to register your target physical machine with PlateSpin Migrate by booting it with the PlateSpin ISO.

29 Migration to VMware

For migration of workloads to a VMware virtual host, PlateSpin Migrate provides automated setup of the target virtual machine on a specified ESX host, in accordance with the features and capabilities of the selected virtualization platform. In addition to the migration settings, you specify settings for the target VM that Migrate will create, such as:

- Target VM name and configuration file path
- Datastore to use from available resources on the target virtual host
- Network settings
- Virtual memory allocation

NOTE

- Raw Device Mapping (RDM) for target VMs on VMware is supported only by using the X2P workflow.
- When you use the X2P workflow for migrating a workload to VMware, you must set up the VMware Tools for the target workload before you perform the conversion. See "Setting Up VMware Tools for the Target Workload" on page 507.
- Before you migrate a Linux workload, ensure that Perl module is installed on the source Linux workload to enable PlateSpin Migrate to install the VMware tools on the target workload during conversion. Alternatively, you can manually install the VMware tools after the migration is completed on the cutover workload.
- If your target VMware ESX server is part of a fully automated Distributed Resource Scheduler (DRS) cluster (a cluster with its VM migration automation level set to **Fully Automated)**, the newly created target VM's automation level is changed to **Partially Automated** for the duration of the migration. This means that your target VM might power up on a different ESX server from the one initially selected, but migration is prevented from automatic execution.

Use the guidelines in this section to configure migration to VMware.

Planning for Migration to VMware

Before you begin migrations to virtual machines on VMware, ensure that your migration environment meets the following guidelines:

Supported VMware Platforms

- See "Supported Target VMware Platforms" on page 40. The configurations apply to VMware environments on-premises.

Supported Workloads

- See "Supported Source Workloads For Migration to Non-Cloud Platforms" on page 23, as appropriate for the target VMware platform.

Network Access and Communications

- See "Access and Communication Requirements across Your Migration Network" on page 53.

Prerequisites

- See Chapter 13, "Prerequisites for Migration to VMware," on page 247.
- See Chapter 21, "Preparing for Migration of Windows Clusters," on page 327.

Target Discovery

Using Migrate Client

- **Target VMware virtual host (automated):** See "Target Discovery in the Migrate Client" on page 283.

Using Migrate Web Interface

- **Target VMware virtual host (automated):** See "Target Discovery in the Web Interface" on page 284 using the appropriate credentials:
  - Table 17-6, "Options for VMware DRS Cluster Target," on page 287
  - Table 17-7, "Options for VMware DRS Cluster as Hosts Target," on page 287
  - Table 17-8, "Options for VMware ESX Server Target," on page 288

Using PlateSpin ISO

- **Target VMs on VMware virtual host (semi-automated):** See "Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO" on page 288.

Workload Discovery

- **Using Migrate Client:** See "Workload Discovery in the Migrate Client" on page 301.
- **Using Migrate Web Interface:** See "Workload Discovery in the Migrate Web Interface" on page 302.
- **(For automated migrations only) Using Migrate Agent:** See "Registering Workloads and Discovering Details with Migrate Agent" on page 303.

Additional Information

- *vSphere Virtual Machine Administration* (https://docs.vmware.com/en/VMware-vSphere/6.5/vsphere-esxi-vcenter-server-65-virtual-machine-admin-guide.pdf)

Automated Migration to VMware Using Migrate Client

1 In the Migrate Client, if the VMware target you need is not yet configured, click **Add Targets**, then configure the target VMware platform.

See "Target Discovery in the Migrate Client" on page 283.

2 In the Migrate Client, initiate a peer-to-peer workload migration.

2a Expand the Tasks options, then select the conversion type, depending on your goals for the migration:

- **Copy Workload**
- **Move Workload**

The Source and Target panes display workloads and targets applicable to the selected type of a migration job.

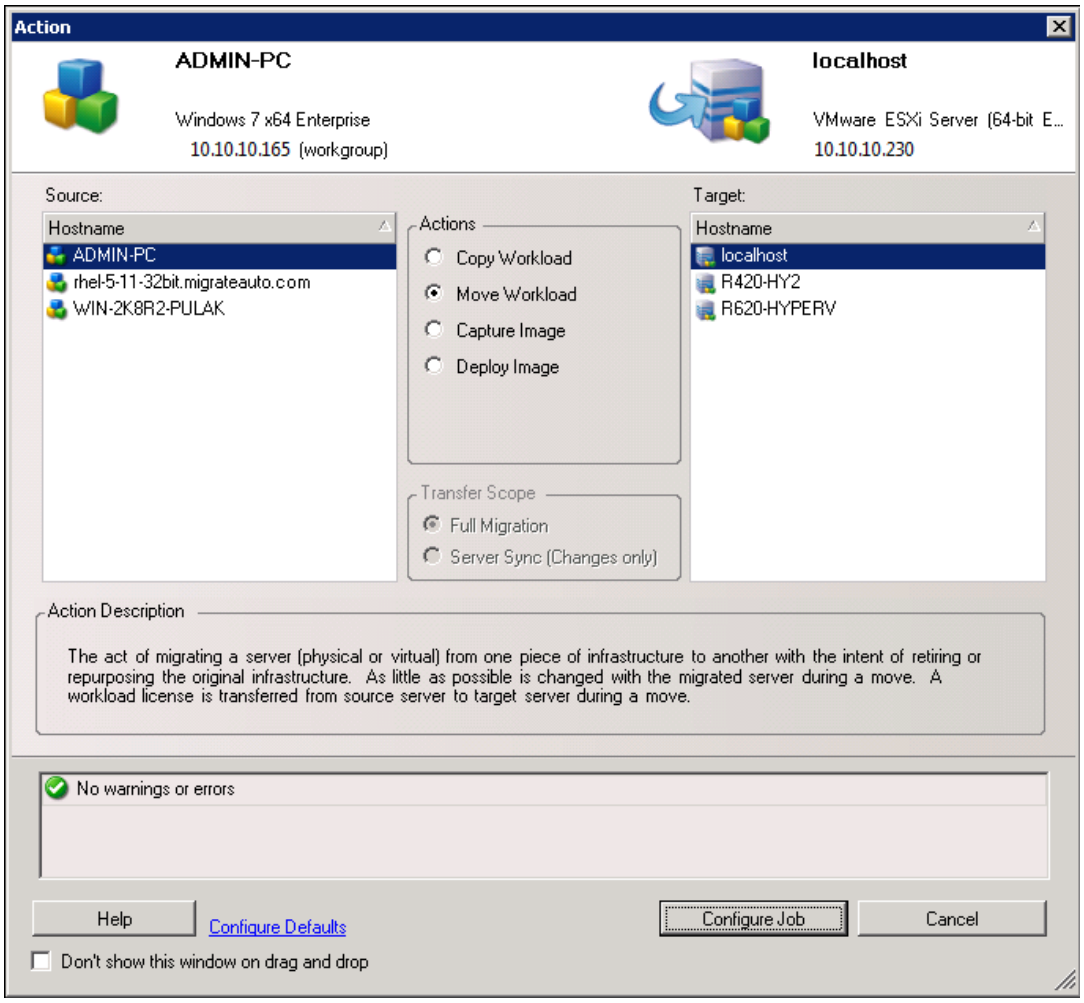

2b In the Source pane, select the workload you want to migrate.

2c In the Target pane, select target VMware platform for the migration.

2d Check the validation messages at the bottom of the window.

2e Click **Configure Job** to access the Peer-to-Peer Migration Job window.

Figure 29-1 *Peer-to-Peer Migration Job Window*

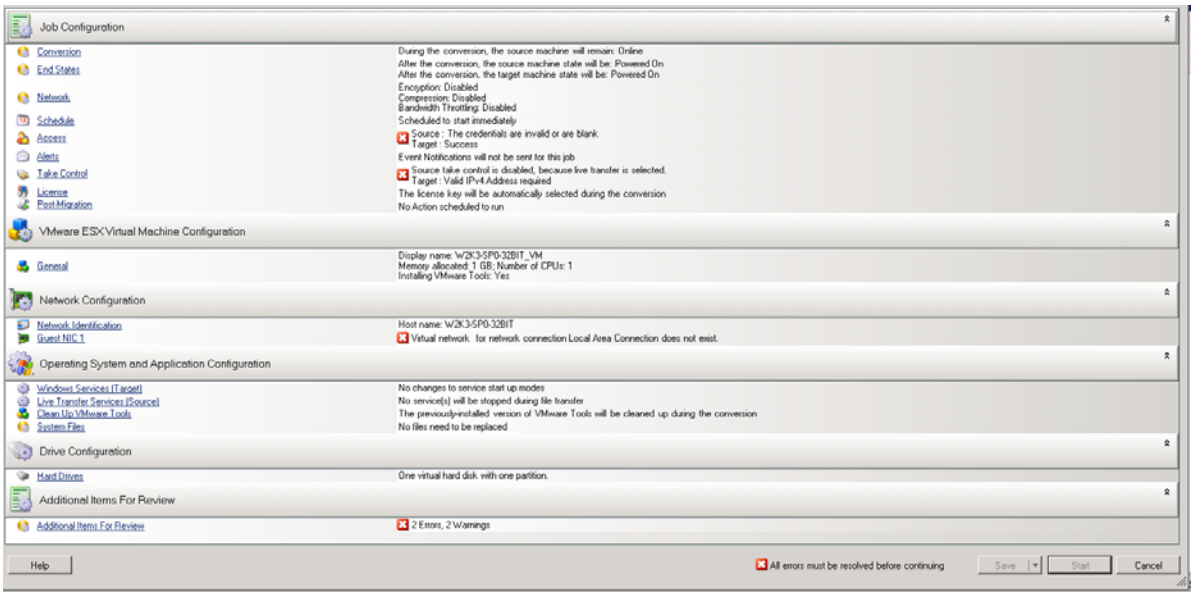

3 In the Job Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
License	
License Key	PlateSpin Migrate automatically selects the best license key for a migration job. If you have multiple license keys, you can specify the license key to use for the workload, assuming licenses are available (neither expired nor exhausted). **To specify an alternate key to use:** 1. Deselect **Automatically select the best license key during the conversion**, then select the appropriate license key from the menu. 2. Click **OK**. The selected license key is displayed on the **License** tab and its description is updated.
Conversion	
Transfer Scope	Specify the scope of workload data to transfer from the source to the target as **Full Migration** or **Server Sync (Changes Only**).
Transfer Method	Specify how data is transferred from source to target. The availability depends on your workload and migration job type.See "Supported Data Transfer Methods" on page 44.
End State	
Source Machine End State	Specify whether to shut down the source workload after a successful cutover. For a workload move, the shut down is selected by default.
Target Virtual Machine End State	Specify whether to power on, power off, or suspend the target workload after a successful cutover.

Setting Name	Description
Network	
Compression	Specify whether to compress data during transmission between the source and target workloads, and the level of data compression to apply.See "Data Compression" on page 52. Select one of the following options: ◆ **Fast:** Consumes the least CPU resources on the source, but yields a lower compression ratio. ◆ **Optimal:** Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option. ◆ **Maximum:** Consumes the most CPU resources on the source, but yields a higher compression ratio.
Encryption	Select **Encrypt Data Transfer** to encrypt the data as it is transferred from source to target.See "Security and Privacy" on page 46.
Bandwidth Throttling	Specify whether to throttle bandwidth for data transfer traffic between the source and target machines. To enable throttling, select the **Enable Bandwidth Throttling** option, specify the required maximum value in Mbps, and optionally a time period during which to enforce the throttling. If specified, the **from** and **to** time values are based on the source workload's system time. If no time interval is defined, bandwidth is throttled to the specified rate at all times by default. If time interval is defined and the migration job executes outside this interval, data is transferred at full speed.
IP Addresses	Specify additional IP addresses for source workloads to enable communication in environments that use network address translation (NAT). For information on how to specify additional IP addresses for your PlateSpin Server, see "Migrations Across Public and Private Networks through NAT" on page 60.
Schedule	
Schedule	Specify when to start the migration job: ◆ Run immediately ◆ Run at a later time Use the calendar menu to specify the date and time to begin the migration. **NOTE:** You must prepare the workload prior to the scheduled time.The full replication cannot run unless the target VM exists and the workload preparation is complete. Migrate skips the scheduled full replication and retries it at the next scheduled time.

Setting Name	Description
Access Settings	
Source Credentials	(Windows) Specify the account user name with local or domain-level administrative privileges and a valid password. Use this format: ◆ For domain member machines: *authority\principal* ◆ For workgroup member machines: *hostname\principal* (Linux) Specify the root or root-level user name and a valid password.
Target Credentials	(VMware DRS Cluster) Specify VMware vCenter Web service user name and password. (VMware ESX Server) Specify the ESX account with administrator role.
Alerts	
Receive Event Notifications	Specify whether to send email notifications for event conditions. You must configure an SMTP server to use this feature.
Receive Progress Notifications	If you enable Event notifications, you can optionally receive progress notifications at a specified interval.
Send to Addresses	Add or remove valid email addresses for recipients of the notifications.
Take Control Settings	
Target Virtual Machine	Under Target Virtual Machine, click Configure, then specify the options for the virtual network and the TCP/IP settings for the replication NIC, then click **OK**.
Post-Migration	
Action	Specify a previously configured custom action that you want to execute on your target workload as a part of the migration job. See "Managing Post-Migration Actions" on page 107.
Execution Parameters	Specify the command line command to run the selected action. You can specify a timeout for the execution.
Credentials	Specify the user name and password to use for the post-migration tasks. You can optionally use the source credentials.

4 In the Virtual Machine Configuration section of the Migration Job window, click **General**, then configure the following settings:

Setting Name	Description
VMware ESX Virtual Machine	
Virtual Machine Name	Specify a name to use for the target VM as it appears in the VMware.
Datastore	Select a datastore associated with your VM for storing VM configuration files.
Path	Type the path to use for the target VM file, including the VM file name. For example: `/hostname/hostname.vmx`

Setting Name	Description
Virtual Machine Memory Allocation	Specify the amount of virtual memory in GB.
Install VMware Tools	Specify whether to install the latest VMware tools on the target VM. If they are installed on the source, they will be uninstalled and reinstalled using the version appropriate for the platform of the VMware target host.
Virtual Devices	Specify preferences for the virtual devices.
Advanced	(For expert users) Specify preferences for the Resource Pool, number of CPUs, and CPU scheduling affinity based on their availability on the target VMware server. Each vCPU is presented to the guest OS on the VM platform as a single core, single socket. (For migration to VM platform that is part of a DRS Cluster) Specify the Resource Pool location where the migrated VM is to be created.

PlateSpin Migrate displays target virtual machine configuration options specific to the selected target and also provides access to advanced configuration options. For information about host-specific configuration options, see Target VM Configuration: VMware ESXi.

5 In the Network Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Network Configuration	
Network Identification Settings for Windows	
Host Name	Specify the desired host name for the target machine.
Member of Domain / Workgroup	Select the required option and type the name of the domain or workgroup that you want the target machine to join.
Preserve Source Server's Domain Registration	Preserves domain registration and ensures that the source server domain registration remains intact during migration. If you disable this option, the source machine's domain account is transferred to the target machine. The source server still appears to be on the domain, but does not have a valid connection.
Domain Credentials	If the target machine is to be part of a domain, specify valid credentials for a user account with permission to add servers to the domain, such as a member of the Domain Admins group or Enterprise Admins group.
Network Identification Settings for Linux	
Host Name	On the Network Identification tab, specify the desired host name for the target machine.
DNS	Use the **Add**, **Edit**, and **Remove** buttons to manage DNS server entries for the new virtual machine.

6 In the Operating System and Applications Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Operating System and Application Configuration	
Windows Services (Target)	Select Windows services' start conditions on the target VM after cutover. Start options are **Automatic**, **Manual**, **Disabled**, and **Automatic (Delayed Start).** **To modify the settings:** 1. Click the Status column for the service, then select from the Windows start options. 2. When you are done setting services start states, click **OK**.
Live Transfer Services (Source)	Specify the Windows services to stop on the source workload during data transfers. We recommend that all the non-VSS compliant services or antivirus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes. **To modify the settings:** 1. Select **Stopped** next to the service to be stopped for data transfer. 2. When you are done setting services to stop, click **OK**.
Linux Daemons (Target)	Specify the start states for daemons on the target VM after cutover. **To modify the settings:** 1. Click the **Run Level** column for the daemon, then select from run levels 0 through 6 and Boot (B), then click **OK**. 2. When you are done setting daemon start states, click **OK**.
Live Transfer Daemons (Source)	Specify the daemons to stop on the source workload during data transfers. **To modify the settings:** 1. Select **Stopped** next to the daemon to be stopped for data transfer. 2. When you are done setting daemons to stop, click **OK**.

7 In the Drive Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Drive Configuration	
Hard Drives	Specify drive and volume configurations to be migrated.
Disks	Specify the path to the hard disk on the target virtual machine.
Volumes	Select volumes to be included in the target for migration.

Setting Name	Description
NTFS Cluster Size	(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see the Microsoft Support KB Article 140365.
Non-volume Storage	(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups	(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.
Volume Groups	(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the **Converted Logical Volumes** section of the settings.
Converted Logical Volumes	(For Linux Workloads) Specify one or more LVM logical volumes to be migrated for a Linux workload.

PlateSpin Migrate displays storage configuration options specific to the selected target. For information about host-specific configuration options, see:

- Drive Configuration: VMware ESX

8 In the Additional Items for Review section of the Migration Job window, review errors and messages about the workload configuration. You must resolve errors before you can submit the migration job.

9 Click **OK**.

Target VM Configuration: VMware ESXi

The following are configuration options specific to VMware vSphere (applicable to all VMs under the containing resource pool).

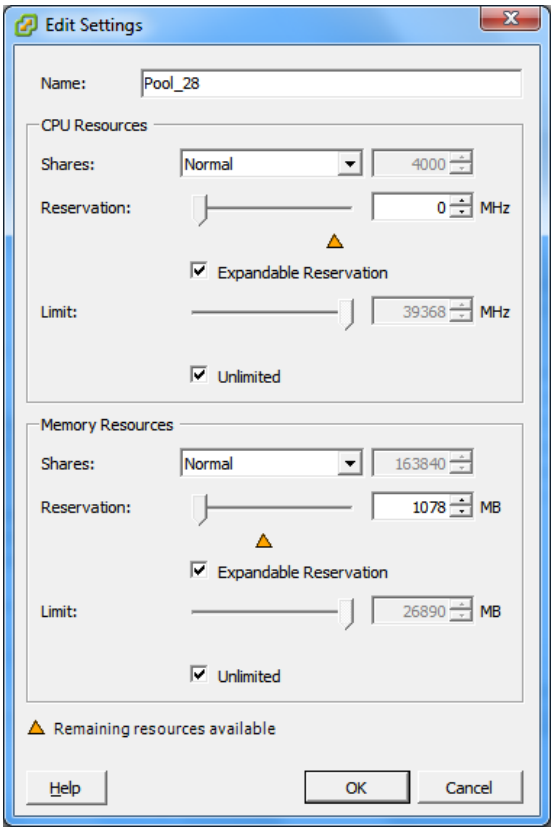

Name: Specify the display name for the new virtual machine.

CPU Resources

- **Shares:** CPU shares for this virtual machine with respect to the parent's total. Peer VMs share resources according to their relative share values bounded by the **Reservation** and **Limit**. Select **Low**, **Normal**, or **High**, which specify share values respectively in a 1:2:4 ratio. Select **Custom** to give each virtual machine a specific number of shares, which express a proportional weight.
- **Reservation:** Guaranteed CPU allocation for this VM.

  Expandable Reservation: Select this option to specify that more than the specified reservation is allocated if resources are available in a parent.
- **Limit:** Upper limit for this virtual machine's CPU allocation.

  Unlimited: Select this option to specify no upper limit.

Memory Resources: (these are similar to CPU resource settings, but apply to memory resources)

Drive Configuration: VMware ESX

The following are drive configuration settings specific to VMware ESX:

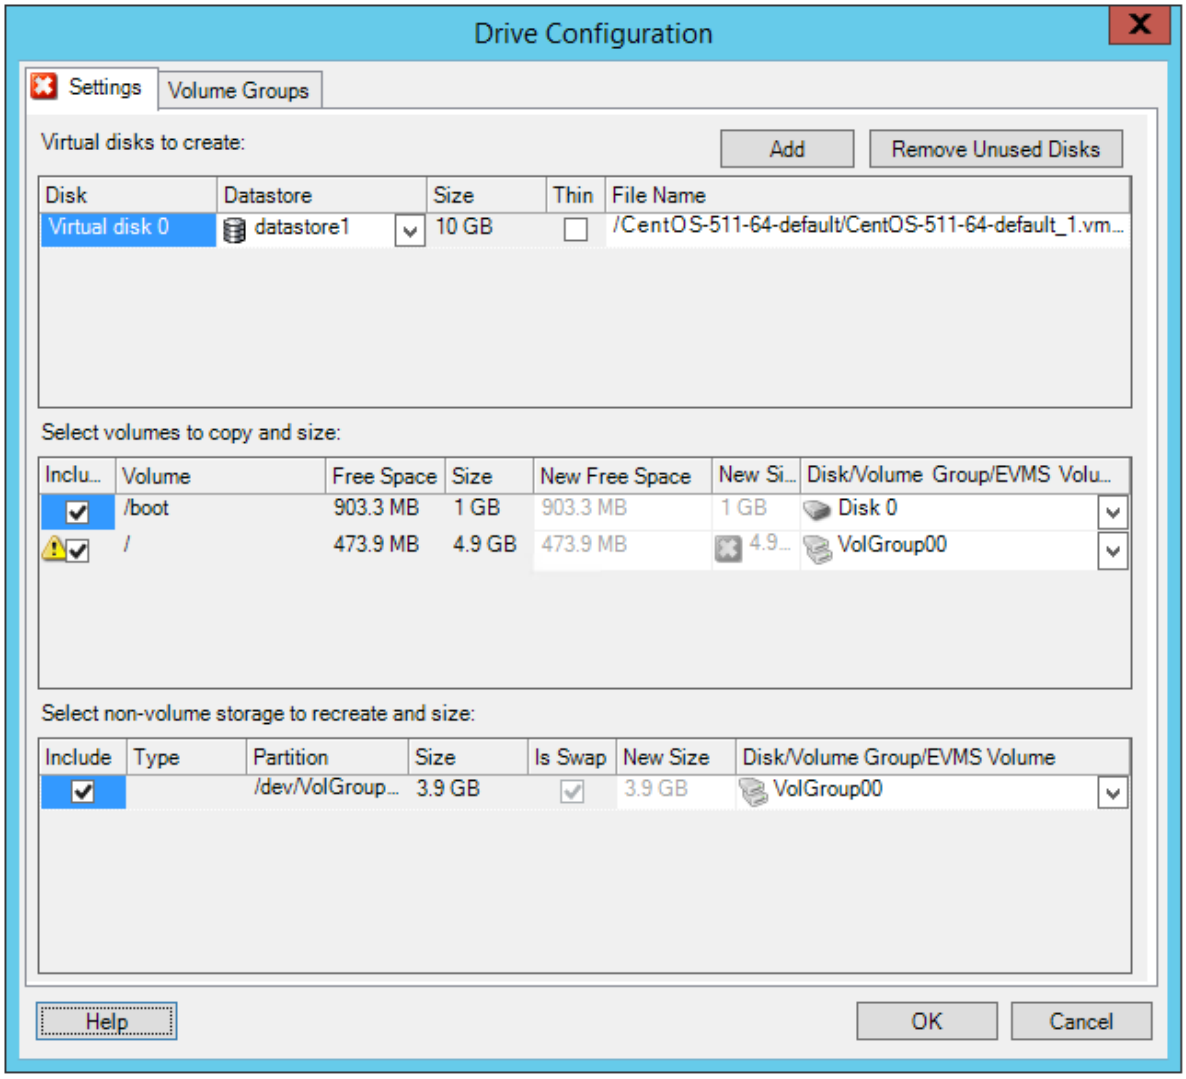

Datastore: Select the datastore volume on the ESX server where you want to place the vmdk files.

Copy: Select the volumes to be copied during the migration.

New Free Space: To resize the volume during the migration, specify the desired amount of free space. PlateSpin Migrate automatically adjusts New Size.

New Size: To resize the volume during the migration, specify the desired size. PlateSpin Migrate automatically adjusts New Free Space.

Disk/Volume Group: Assign the volume to a disk or, if LVM is enabled, to a volume group. The volume will be copied to this disk or volume group on the target machine.

Create: Select any non-volume disk partitions that should be created on the target machine (for example, a Linux swap partition).

New Size: To resize the non-volume partition during the migration, specify the desired size.

Automated Migration to VMware Using Migrate Web Interface

1. Launch the PlateSpin Migrate Web Interface.
2. If you have not configured a VMware platform as a migration target, click **Targets** > **Add Target**, and then configure the target VMware platform.

   See the following, depending on the target VMware platform:

   - Table 17-6, “Options for VMware DRS Cluster Target,” on page 287
   - Table 17-7, “Options for VMware DRS Cluster as Hosts Target,” on page 287
   - Table 17-8, “Options for VMware ESX Server Target,” on page 288
3. On the Workloads page, select the workload you want to configure.
4. Click **Configure Migration**.
5. Select one of the following based on the scope of data you want to transfer from the source to the target:
   - **Full Replication:** A full volume of data transfer takes place from the source to the target.
   - **Incremental Replication:** Only differences are transferred from the source to the target, provided they have similar operating system and volume profiles.
6. Depending on whether you want to perform full replication or incremental replication, do one of the following:
   - **Full Replication:** Select an existing VMware target platform to which you want to migrate the source workload.
   - **Incremental Replication:** Select the target type as **Virtual Targets.**From the list of the prepared target workloads, select a virtual target workload that matches the source workload you want to migrate. If the prepared target workloads list is either empty or does not list a matching target workload, then click **Prepare Workload** to prepare the target workload to which you want to migrate the source workload.
7. Click **Configure Migration**.
8. Configure the following settings:

NOTE: All of the following settings are applicable for full replication of the workload. Some of the listed settings might not display or display differently when you configure incremental replication for the workload. For example:

- The **Virtual Machine Configuration Path** setting does not display because a target workload matching the source workload must already be available to transfer only the differences from the source to this target.
- The **Volume Mapping** setting displays the volumes and their mappings. You can edit the volume mappings.

Setting Name	Description
Schedule Settings	
Incremental Recurrence	Specify the following: ◆ **Start of Recurrence:** The date when you want to start the replication. You can specify the date or click the calendar icon to select the date. By default, the time is 12:00 a.m. ◆ **Recurrence Pattern:** The pattern to follow for the recurrence of the replication. For example: ◆ To use incremental recurrence everyday, select **Daily.** ◆ To never use incremental recurrence, select **None.** **NOTE** ◆ Scheduled incremental replications are skipped until the first full replication is complete. ◆ When you schedule incremental recurrence, the replication takes place for a maximum period of 60 days from the starting time of replication. For example: ◆ If you select **Daily**, then the replication takes place for 60 days from the time the replication starts. ◆ If you select **Weekly**, then the replication takes place for 8 weeks from the time the replication starts. ◆ If you select **Monthly**, then the replication takes place for 2 months from the time the replication starts.
Full Replication	Do one of the following: ◆ To specify a schedule for the replication, click **Start** and specify the date when you want to start the full replication. ◆ To start full replication manually without setting a schedule, click **None**. **NOTE:** You must prepare the workload prior to the scheduled time.The full replication cannot run unless the target VM exists and the workload preparation is complete. Migrate skips the scheduled full replication and retries it at the next scheduled time.
Blackout Window	Use these settings to force a replication blackout. The replication blackout suspends scheduled replications during peak utilization hours or prevents conflicts between VSS-aware software and the PlateSpin VSS block-level data transfer component. To specify a blackout window, click **Edit** and do the following: ◆ Specify the start and end time for the blackout period. ◆ Select one of the blackout recurrence pattern such as daily, weekly, or monthly. If you do not want to force a replication blackout, select **None**. **NOTE:** The blackout start and end times are based on the system clock on the PlateSpin Server.

Setting Name	Description
Compression Level	These settings control whether data is compressed during transmission between the source and target workloads, and the level of data compression applied.See "Data Compression" on page 52.Select one of the following options: ◆ **Fast:** Consumes the least CPU resources on the source, but yields a lower compression ratio. ◆ **Optimal:** Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option. ◆ **Maximum:** Consumes the most CPU resources on the source, but yields a higher compression ratio.
Bandwidth Throttling	These settings control the bandwidth throttling. PlateSpin Migrate enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. You can specify a throughput rate for each migration job. Throttling provides a way to prevent migration traffic from congesting your production network and to reduce the overall load of your PlateSpin Server. To throttle replications to a specified rate, specify the required throughput value in Mbps and the time pattern.
Migration Settings	
Transfer Method	(For Windows Workloads) Select a data transfer mechanism and security through encryption.See "Supported Data Transfer Methods" on page 44. To enable encryption, select the **Encrypt Data Transfer** option. See "Security and Privacy" on page 46.
Transfer Encryption	(For Linux Workloads) To enable encryption, select the **Encrypt Data Transfer** option.See "Security and Privacy" on page 46.
Source Credentials	Specify the credentials required for accessing the workload. See "Discovery Guidelines for Source Workloads" on page 299.
CPU	(For migration to vCloud and VM platforms using supported versions of VMware 5.5 and later with a minimum VM hardware Level 8) Specify the number of sockets and the number of cores per socket for the target workload. It automatically calculates the total cores and displays the total CPUs. This parameter applies on the initial setup of a workload with an initial replication setting of **Full Replication.** **NOTE:** The maximum number of cores the workload can use is subject to external factors such as the guest operating system, the VM hardware version, VMware licensing for the ESXi host, and ESXi host compute maximums for vSphere (see *ESXi/ESX Configuration Maximums (VMware KB 1003497)* (https://kb.vmware.com/kb/1003497)). Some distributions of a guest OS might not honor the cores and cores per socket configuration. Instead they will retain their original cores and sockets settings as installed.
Resource Pool for Target VM	(For migration to VM platform that is part of a DRS Cluster) Specify the Resource Pool location where the migrated VM is to be created.

Setting Name	Description
VM Folder for Target VM	(For migration to VM platform that is part of a DRS Cluster) Specify the VM folder location where the migrated VM is to be created.
Virtual Machine Name	Specify a display name for the new virtual machine.
Configuration File Datastore	Select a datastore associated with your VM for storing VM configuration files.
Virtual Machine Configuration Path	Specify the path to the configuration file on the target virtual machine.
Storage Layout	Specify the desired layout of data volumes (and volume groups) on the target workload disks: ◆ **Source-based** The Source-based storage layout assigns data volumes to disks according to their layout on the source workload. ◆ **Single disk** The Single disk storage layout assigns all the data volumes or LVM volume groups to a single disk on the target workload. An LVM volume group might contain multiple volumes. ◆ **One volume per disk** The One volume per disk storage layout assigns each data volume or LVM volume group to a separate disk on the target workload. An LVM volume group might contain multiple volumes.
Storage Options > Disks	Specify the path to the hard disk on the target virtual machine.
Storage Options > Volumes	Select volumes to be included in the target for migration. For Windows file-based migrations, you can optionally adjust the size of the target volume by modifying the **Size** value in the row for that volume. **NOTE:** The ability to adjust the size of the target volume is not currently available in the Disk Layout view of storage in the Migrate Web Interface. See "Cannot Adjust the Target Volume Size for File-Based Migrations in the Web Interface" on page 569.
NTFS Cluster Size	(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see the Microsoft Support KB Article 140365.
Non-volume Storage	(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups	(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default. **NOTE:** Disks that host both regular volumes (such as `/boot`) and volume groups are configured in the Disks section.

Setting Name	Description
Volume Groups	(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the **Converted Logical Volumes** section of the settings.
Converted Logical Volumes	(For Linux Workloads) Specify one or more LVM logical volumes to be migrated for a Linux workload.
Replication Network for Target	Specify a network interface (NIC or IP address) on the target to use for replication traffic.
Replication Networks for Source	Specify one or more network interfaces (NIC or IP address) on the source to use for replication traffic.
Services to Stop Before Any Replication	(For Windows Workloads) We recommend that all the non-VSS compliant services or anti-virus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.
Services to Stop for Cutover with Replication	(For Windows Workloads) Select the Windows services that should be permanently stopped on the source workload for cutover with any replication. The services stopped on the source workload during the replication process are not restored after cutover. The stopped services are restored after a Test Cutover.
Daemons to Stop Before Any Replication	(For Linux Workloads) Select the Linux daemons that you want to be temporarily stopped on the source workload before replication. These daemons will be restored after replication completes.
Daemons to Stop for Cutover with Replication	(For Linux Workloads) Select the Linux daemons that should be permanently stopped on the source workload for Cutover with any Replication. The daemons stopped on the source workload during the replication process are not restored after Cutover. The stopped daemons are restored after a Test Cutover.
Boot Services to Disable on Target	(For Windows Workloads) Select the boot services that you want to disable on the target workload.
Target Workload Settings (These settings are applied during the Run Cutover)	
VM Memory	Specify the amount of memory allocated to the target workload.
VM Tools	To install the VM tools, select the **Install VM Tools option**. This option is selected by default. PlateSpin Migrate uninstalls any VMware tools that are already available on the source and reinstalls the version of VMware tools appropriate for the platform of the VMware target host.
Hostname	Do one of the following: ◆ To retain the same host name, select **No Change**. ◆ To change the host name, select **Set To** and specify the new name.

Setting Name	Description
Domain / Workgroup	(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays: ◆ **Workgroup:** *Workgroup_name* where *Workgroup_name* is the workgroup name to which the source belongs. ◆ **Domain:** *Domain_name* where *Domain_name* is the domain name to which the source belongs. Do one of the following depending on where you want the target workload to join: ◆ **When the source workload belongs to a workgroup:** Assume that the source workload belongs to a workgroup named **WorkGroup1**. ◆ For the target workload to join the same workgroup (**WorkGroup1**), retain the following existing selection: **Workgroup: Workgroup1** ◆ For the target workload to join a different workgroup (say **WorkGroup2)**, select **Join Workgroup** and specify the name as **WorkGroup2**. ◆ For the target workload to join a domain, select **Join Domain** and specify the domain name you want the target to join. ◆ **When the source workload belongs to a domain:** Assume that the source workload belongs to a domain named **Domain1**. ◆ For the target workload to join a workgroup, click **Join Workgroup** and specify the name of the workgroup you want the target to join. ◆ For the target workload to join the same domain (**Domain1**) with the domain registration settings preserved, retain the following existing selection: **Domain: Domain1** ◆ For the target workload to join the same domain (**Domain1)** without preserving the domain registration settings, select **Join Domain** and specify the domain name as **Domain1**. ◆ For the target workload to join a different domain, select **Join Domain** and specify the domain name you want the target to join.
Domain Credentials	(For Windows Workloads) If you select **Join Domain**, specify the domain administrator credentials.
Network Connections	Select the local area connection and then select one of the following: ◆ **DHCP:** Obtain an IP address automatically assigned by a DHCP server. ◆ **Static:** Specify one or more static IP addresses, subnet masks, and gateway IP addresses. Click **Advanced** to add multiple entries, edit an existing entry, or remove an existing entry. The IP addresses must be unique within the supported network. For Windows workloads that have more than one NIC, select the connection for each NIC.

Setting Name	Description
DNS Servers	Specify the DNS Servers for the target workloads. This is applicable only if you select **Static** in the **Network Connections** option: ◆ **Primary DNS server:** Specify the primary DNS server address. ◆ **Alternative DNS server:** Specify an alternate DNS server address. ◆ **Additional DNS server:** To specify additional DNS server addresses: 1. Click **Advanced**. 2. Specify the DNS server address. 3. Click **Add** to add the server in the DNS Server Addresses list. 4. Click **OK**.
Services States on Target VM	(For Windows Workloads) Select Windows services' start conditions on the target VM. Start options are Automatic, Manual, Disabled, and Automatic (Delayed Start).
Daemons States to Change	(For Linux Workloads) Select Linux daemons' start conditions on the target VM. Enable the daemon to start by selecting the check boxes at the appropriate runlevels (0 to 6) and Boot.
Post-Migration Action	Click **Configure Action** to select a previously configured custom action that you want to execute on your target workload as a part of the migration job. The command line parameters and timeout value that you previously configured for the selected action is displayed by default. You can edit these values as required. If you choose to timeout the action and have not specified any timeout value, then PlateSpin Migrate waits for a maximum period of 7200 seconds before timing out this action. For information about managing the post-migrations, see "Managing Post-Migration Actions" on page 107. **NOTE:** The post-migration action that you configure for the Target Workload Settings is automatically applicable for the Target Workload Test Settings and you cannot edit it.
Target Workload Test Settings (These settings are applied during the Test Cutover)	
Copy Target Workload Settings	
Click the **Copy Target Workload Settings** option to automatically copy the workload settings from **Target Workload Settings** section to **Target Workload Test Settings** section.	
VM Memory	Specify the amount of memory allocated to the target workload.
VM Tools	To install the VM tools, select the **Install VM Tools option**. This option is selected by default. PlateSpin Migrate uninstalls any VMware tools that are already available on the source and reinstalls the version of VMware tools appropriate for the platform of the VMware target host.
Hostname	Do one of the following: ◆ To retain the same host name, select **No Change**. ◆ To change the host name, select **Set To** and specify the new name.

Setting Name	Description
Domain / Workgroup	(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays: ◆ **Workgroup:** *Workgroup_name* where *Workgroup_name* is the workgroup name to which the source belongs. ◆ **Domain:** *Domain_name* where *Domain_name* is the domain name to which the source belongs. Do one of the following depending on where you want the target workload to join: ◆ **When the source workload belongs to a workgroup:** Assume that the source workload belongs to a workgroup named **WorkGroup1**. ◆ For the target workload to join the same workgroup (**WorkGroup1**), retain the following existing selection: **Workgroup: Workgroup1** ◆ For the target workload to join a different workgroup (say **WorkGroup2)**, select **Join Workgroup** and specify the name as **WorkGroup2**. ◆ For the target workload to join a domain, select **Join Domain** and specify the domain name you want the target to join. ◆ **When the source workload belongs to a domain:** Assume that the source workload belongs to a domain named **Domain1**. ◆ For the target workload to join a workgroup, click **Join Workgroup** and specify the name of the workgroup you want the target to join. ◆ For the target workload to join the same domain (**Domain1**) with the domain registration settings preserved, retain the following existing selection: **Domain: Domain1** ◆ For the target workload to join the same domain (**Domain1)** without preserving the domain registration settings, select **Join Domain** and specify the domain name as **Domain1**. ◆ For the target workload to join a different domain, select **Join Domain** and specify the domain name you want the target to join.
Domain Credentials	(For Windows Workloads) If you select **Join Domain**, specify the domain administrator credentials.
Network Connections	Select the network connection and then select one of the following: ◆ **DHCP:** Obtain an IP address automatically assigned by a DHCP server. ◆ **Static:** Specify one or more static IP addresses, subnet masks, and gateway IP addresses. Click **Advanced** to add multiple entries, edit an existing entry, or remove an existing entry. The IP addresses must be unique within the supported network.

Setting Name	Description
DNS Servers	Specify the DNS Servers for the target workloads. This is applicable only if you select **Static** in the **Network Connections** option: ◆ **Primary DNS server:** Specify the primary DNS server address. ◆ **Alternative DNS server:** Specify an alternate DNS server address. ◆ **Additional DNS server:** To specify additional DNS server addresses: 1. Click **Advanced**. 2. Specify the DNS server address. 3. Click **Add** to add the server in the DNS Server Addresses list. 4. Click **OK**.
Services States on Target VM	(For Windows Workloads) Select Windows services that must be automatically stopped on the target VM.
Daemons States to Change	(For Linux Workloads) Select Linux daemons that must be automatically stopped on the target VM.
Post-Migration Action	The post-migration action that you configure for the Target Workload Settings is automatically applied and cannot be changed.
Tag	
Tag	Select a tag to assign to the workload. See "Managing Workload Tags" on page 133.

9 (Optional) To change the target, click **Change Target**.

NOTE: If you change the target, all the settings you specified will be cleared.

10 Do one of the following:

- Click **Save** to save the settings.
- Click **Save and Prepare** to save the settings and start preparing the workload migration.
- Click **Cancel** to exit.

Migration to VMs on VMware Using X2P Workflow

Raw Device Mapping (RDM) for target VMs on VMware is only supported by using the X2P workflow. When you use the X2P workflow for migrating a workload to VMware, you must set up the VMware Tools for the target workload before you perform the conversion.

Use the guidelines in this section to configure migration to VMs on VMware virtual hosts.

- "Downloading and Saving the PlateSpin ISO Image (VMware)" on page 507
- "Creating and Configuring the Target Virtual Machine (VMware)" on page 507
- "Setting Up VMware Tools for the Target Workload" on page 507
- "Registering the Virtual Machine with PlateSpin Server (VMware)" on page 509

- "Using PlateSpin Migrate Client to Migrate Your Source Workload to the Target Virtual Machine (VMware)" on page 509
- "Using PlateSpin Migrate Web Interface to Migrate Your Source Workload to the Target Virtual Machine (VMware)" on page 509

Downloading and Saving the PlateSpin ISO Image (VMware)

1 Download and prepare the PlateSpin ISO image for use with the target VM. Attended and unattended registration options are possible.

See "Preparing the PlateSpin ISO Image for Target Registration and Discovery" on page 374.

2 Save the ISO image in a location that VMware server can access. For example: `c:\temp`.

This ensures that the PlateSpin ISO image is available to the target VM as a bootable CD-ROM image.

Creating and Configuring the Target Virtual Machine (VMware)

1 Log on to the VMware server using the vSphere client and then use the New Virtual Machine Wizard to create a new virtual machine with the following settings:

- **Name and Location:** Specify a name for your new target and accept the default location.
- **Operating System Type and Version:** Specify the operating system type and version settings that matches the source workload. The wizard uses this information to set appropriate default values, such as the amount of memory needed, and resource limits for the VM.
- **Assign Memory:** Assign at least 384 MB of RAM to the VM.
- **Connect Virtual Hard Disk:** Ensure that the disk size of every disk is about 50 MB more than the corresponding disk on your source workload.
- **Installation Options:** Configure the VM to boot from an ISO image file, and point the wizard to the downloaded PlateSpin ISO image.
- **Summary:** Configure the VM to not start upon creation (deselect the **Start the virtual machine after it is created** option).

2 Set up VMware tool for the target workload. See "Setting Up VMware Tools for the Target Workload" on page 507.

Setting Up VMware Tools for the Target Workload

VMware Tools setup packages are automatically copied to the target during conversion so that the configuration service can install the tools on the target VM when the target VM contacts the PlateSpin Server. However, if you choose to migrate workloads to VMware using the X2P workflow,

you must set up the VMware tools for the target workload before you perform the conversion. Perform the following steps to prepare your environment for setting the VMware tools for the target workload:

1 Retrieve the VMware Tools packages from an ESX host:

   1a Secure copy (`scp`) the `windows.iso` image from the `/usr/lib/vmware/isoimages` directory on an accessible ESX host to a local temporary folder.

   1b Open the ISO and extract its setup packages, saving them to an accessible location.

   The setup packages are `setup.exe` and `setup64.exe`.

2 Create OFX packages from the setup packages you extracted:

   2a Zip the package you want, making sure that the setup installer file is at the root of the `.zip` archive.

   2b Rename the `.zip` archive to `1.package` so that it can be used as an OFX package.

   NOTE: If you want to create an OFX package for more than one of the setup packages, remember that each setup package must have its own unique `.zip` archive.

   Because each package must have the same name (`1.package`), if you want to save multiple `.zip` archives as OFX packages, you need to save each in its own unique subdirectory.

3 Copy the appropriate OFX package (`1.package`) to the `%ProgramFiles%\PlateSpin Migrate Server\Packages\%`*`GUID`*`%` directory on the PlateSpin Server.

   The value of *%GUID%* depends on the version of your VMware ESX host and its VMware Tools architecture, as shown in Table 29-1. Use the appropriate GUID value to copy the package to the correct directory.

Table 29-1 *GUIDs for the VMware Tools Directory Names*

VMware Server Version	VMware Tools Architecture	GUID
5.5	x86	`660C345A-7A91-458b-BC47-6A3914723EF7`
5.5	x64	`8546D4EF-8CA5-4a51-A3A3-6240171BE278`
6.0	x86	`311E672E-05BA-4CAF-A948-B26DF0C6C5A6`
6.0	x64	`D7F55AED-DA64-423F-BBBE-F1215529AD03`
6.5	x86	`D61C0FCA-058B-42C3-9F02-898F568A3071`
6.5	x64	`5D3947B7-BE73-4A00-A549-B15E84B98803`
6.7	x86	`C1170EA0-B145-4E77-8167-BDB3125487B1`
6.7	x64	`BAD60698-C360-47CC-BBFF-DBF29BD24062`

Registering the Virtual Machine with PlateSpin Server (VMware)

After you create the virtual machine and prepare it to boot with the PlateSpin ISO, you are ready to register it as a target VM with your PlateSpin Server. See "Registering and Discovering Target VMs on Virtual Hosts" on page 289.

Using PlateSpin Migrate Client to Migrate Your Source Workload to the Target Virtual Machine (VMware)

1. Use PlateSpin Migrate Client to start an X2P migration job with your source workload being the job's migration source and the target being the new VM on VMware.

   See "Migration to Physical Machines" on page 515.
2. For host-specific target VM configuration options for the Virtual Machine Configuration dialog, see "Target VM Configuration: VMware ESXi" on page 496.
3. For host-specific storage configuration options, see "Drive Configuration: VMware ESX" on page 497.
4. Monitor the migration job in Jobs view in PlateSpin Migrate Client.

   When the job reaches the **Configure Target Machine** step, the virtual machine's console returns to the boot prompt of the PlateSpin ISO image.
5. Shut down the virtual machine and reconfigure it to boot from disk rather than from the boot image.
6. Power on the virtual machine.

   The migration job resumes, reboots the target, and completes the workload configuration.

Using PlateSpin Migrate Web Interface to Migrate Your Source Workload to the Target Virtual Machine (VMware)

1. Use PlateSpin Migrate Web Interface to start an X2P migration job with your source workload being the job's migration source and the target being the new VM on VMware.

   See "Migration to Physical Machines" on page 515.
2. Monitor the status of the migration job. When the job reaches the **Configure Target Machine** step, the virtual machine's console returns to the boot prompt of the PlateSpin ISO image.
3. Shut down the virtual machine and reconfigure it to boot from disk rather than from the boot image.
4. Power on the virtual machine.

   The migration job resumes, reboots the target, and completes the workload configuration.

Migration of Windows Clusters to VMware

You can migrate a Microsoft Windows cluster's business services to VMware. For information about migrating Windows clusters, see:

- Chapter 21, "Preparing for Migration of Windows Clusters," on page 327

30 Migration to Virtual Machines on KVM

For migration of workloads to a KVM virtual host, PlateSpin Migrate requires that you manually set up the target virtual machine with guest operating system type and version settings that match your source workload, in accordance with the features and capabilities of the KVM virtualization platform. Use the PlateSpin ISO to register the target machine with the PlateSpin Server and send machine details. Use PlateSpin Migrate Client or PlateSpin Migrate Web Interface to migrate workloads to a KVM virtual host.

Use the guidelines in this section to configure migration to VMs on KVM virtual hosts.

Planning for Migration to KVM

Before you begin migrations to virtual machines on KVM virtual hosts, ensure that your migration environment meets the following guidelines:

Supported KVM Platforms

- See "Supported Target KVM Platforms" on page 42.

Supported Workloads

- See "Supported Source Workloads For Migration to Non-Cloud Platforms" on page 23, as appropriate for the target KVM platform.

Network Access and Communications

- See "Access and Communication Requirements across Your Migration Network" on page 53.

Prerequisites

- See "Prerequisites for Migration to VMs on KVM" on page 267.

Targets and Workloads

- **Target VM on a KVM virtual host (semi-automated):** See "Registering and Discovering Target VMs on Virtual Hosts" on page 289.
- **Source Workload Discovery:**
  - **Using Migrate Client:** See "Workload Discovery in the Migrate Client" on page 301.
  - **Using Migrate Web Interface:** See "Workload Discovery in the Migrate Web Interface" on page 302.

Additional Information

- *SUSE Linux Enterprise Server 11 SPX Virtualization with KVM* (https://www.suse.com/documentation/sles11/singlehtml/book_kvm/book_kvm.html)
- *Red Hat Enterprise Linux 7.X Virtualization Deployment and Administration Guide* (https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Virtualization_Deployment_and_Administration_Guide/index.html)

Configuring Migration to a VM on a KVM Virtual Host

You can use KVM as the target virtualization platform in a semi-automated workload virtualization.

- "Downloading and Preparing the PlateSpin ISO Image (KVM)" on page 512
- "Creating and Configuring the Target Virtual Machine (RHEL KVM)" on page 512
- "Registering the Virtual Machine with PlateSpin Server (RHEL KVM)" on page 513
- "Migrating Your Source Workload to the Target Virtual Machine (RHEL KVM)" on page 513

Downloading and Preparing the PlateSpin ISO Image (KVM)

1 Download and prepare the PlateSpin ISO image for use with the target VM. Attended and unattended registration options are possible.

See "Preparing the PlateSpin ISO Image for Target Registration and Discovery" on page 374.

2 Save the ISO image in a location that the KVM virtual host can access.

This ensures that the PlateSpin ISO image is available to the target VM as a bootable CD-ROM image.

Creating and Configuring the Target Virtual Machine (RHEL KVM)

1 On RHEL KVM, use the Virtual Machine Manager Wizard or the Create Virtual Machines program shortcut to create a new virtual machine.

Ensure that the new virtual machine is created with the following settings:

- **Virtualization method:** Fully virtualized.
- **Operating System Type and Version:** Specify the operating system type and version settings that matches the source workload. The wizard uses this information to set appropriate default values (such as the amount of memory needed) and resource limits for the VM.
- **Memory:** Assign at least 384 MB of RAM to the VM. This ensures that the VM has sufficient resources during the migration and improves transfer speed. If the virtual machine requires less memory after the migration, reduce the assigned memory after the migration completes.
- **Disks:** Assign disks such that the disk size of every disk is about 50 MB more than the corresponding disk on your source workload. The storage can be either a raw SAN LUN or a virtual disk. Also, create a Virtual CD-ROM assigned to the downloaded PlateSpin ISO image.

2 Ensure that the VM is configured to restart on reboot.

3 From the Virtual Machine Manager, launch the virtual machine console and monitor the boot process.

When the virtual machine completes the boot process, it prompts you for parameters that control the registration of the machine and its profile with PlateSpin Migrate. If you are using the unattended registration process, the required parameters are read from an answer file.

Registering the Virtual Machine with PlateSpin Server (RHEL KVM)

After you create the virtual machine and prepare it to boot with the PlateSpin ISO, you are ready to register it as a target VM with your PlateSpin Server. See "Registering and Discovering Target VMs on Virtual Hosts" on page 289.

Migrating Your Source Workload to the Target Virtual Machine (RHEL KVM)

1 Use PlateSpin Migrate Client or the PlateSpin Migrate Web Interface to start an X2P migration job with your source workload being the job's migration source and the target being the new VM on the RHEL KVM hypervisor.

See "Migration to Physical Machines" on page 515.

2 Monitor the status of the migration job. When the job reaches the **Configure Target Machine** step, the virtual machine's console returns to the boot prompt of the PlateSpin ISO image.

3 Shut down the virtual machine, reconfigure it to boot from disk rather than from the boot image.

4 Power on the virtual machine.

The migration job resumes, reboots the target, and completes the workload configuration.

31 Migration to Physical Machines

PlateSpin Migrate supports semi-automated migration to physical machines. You prepare the target machine to meet migration needs, and then use PlateSpin Migrate to automate the data migration.

NOTE: PlateSpin Migrate Web Interface supports one-time migration of workloads to physical targets. Consequently, if the migration job of a migration prepared workload to a physical machine fails or is aborted, the status of migration is marked as Cut Over and is not reverted to the previous successful Migration Prepared state. To retry migration to the physical machine, you must use PlateSpin ISO to register the target physical machine again with the PlateSpin Migrate server.

Use the guidelines in this section to configure migration to physical machines.

Planning for Migration to Physical Machines

Before you begin migrations to physical machines, ensure that your migration environment meets the following guidelines:

Supported Physical Hardware

- See the following information in "Supported Configurations" on page 23:
  - Supported Workload Storage
  - Supported Workload Architectures

Supported Workloads

- See "Supported Source Workloads For Migration to Non-Cloud Platforms" on page 23.

Network Access and Communications

- See "Access and Communication Requirements across Your Migration Network" on page 53.

Prerequisites

- See "Prerequisites for Migration to Physical Machines" on page 271.

Targets and Workloads

- **Target physical host (semi-automated):** See "Registering and Discovering Details for Target Physical Machines with PlateSpin ISO" on page 291.

- **Source Workload Discovery:**
  - **Using Migrate Client:** See “Workload Discovery in the Migrate Client” on page 301.
  - **Using Migrate Web Interface:** See “Workload Discovery in the Migrate Web Interface” on page 302.

Configuring Migration to a Physical Target (P2P, V2P) Using Migrate Client

To initiate a peer-to-peer workload migration to a physical machine:

1. (Recommended) Use PlateSpin Analyzer to ensure that:
   - Your source operating system and hardware are supported by PlateSpin Migrate.
   - PlateSpin Migrate’s X2P device driver database contains device drivers that your target requires for the operating system being ported.

   See “Analyzing Suitability of Discovered Windows Workloads For Conversion to Physical Machines” on page 320.
2. Discover your source workload. See “Workload Discovery in the Migrate Client” on page 301.
3. (Conditional) If drivers for the physical target are not available in the PlateSpin Migrate’s X2P device driver database, upload the required drivers to the database.

   See Chapter 19, “Preparing Device Drivers,” on page 311.
4. Register your target physical machine with PlateSpin Migrate by booting it with the PlateSpin ISO.

   See “Registering and Discovering Details for Target Physical Machines with PlateSpin ISO” on page 291.
5. Launch Migrate Client, then start a peer-to-peer workload migration.

   The Source and Target panes display workloads and targets applicable to the selected type of a migration job.

   See “Initiating a Migration Job” on page 386.

   5a. Under Tasks, select the conversion type, depending on your goals for the migration:
   - **Copy Workload**
   - **Move Workload**

   In the Action dialog, the Transfer Scope is set to **Full Migration**.

   5b. In the Source pane, select the workload you want to migrate.

   5c. In the Target pane, select target physical machine for the migration.

   5d. Read the validation messages at the bottom of the window.

   5e. Click **Configure Job** to access the Peer-to-Peer Migration Job window.
6. Configure the required parameters of the job.

   See Chapter 24, “Configuration Essentials,” on page 385.

Setting Name	Description
License	
Licenses License Key	PlateSpin Migrate automatically selects the best license key for a migration job. If you have multiple license keys, you can specify the license key to use for the workload, assuming licenses are available (neither expired nor exhausted). **To specify an alternate key to use:** 1. Deselect **Automatically select the best license key during the conversion**, then select the appropriate license key from the menu. 2. Click **OK**. The selected license key is displayed on the **License** tab and its description is updated.
Conversion	
Transfer Scope	Set by default to **Full Migration**.
Transfer Method	Specify how data is transferred from source to target. The availability depends on your workload and migration job type.See "Supported Data Transfer Methods" on page 44.
End State	
Source Machine End State	Specify whether to shut down the source workload after a successful cutover. For a workload move, **Shutdown** is selected by default.
Target Virtual Machine End State	Specify whether to power on, power off, or suspend the target workload after a successful cutover.
Network	
Compression	Specify whether to compress data during transmission between the source and target workloads, and the level of data compression to apply: **Fast**, **Optimal**, or **Maximum**. Compression is disabled by default.See "Compression during Data Transfer" on page 394.
Encryption	Select **Encrypt Data Transfer** to encrypt the data as it is transferred from source to target.See "Security and Privacy" on page 46.
Bandwidth Throttling	Select **Enable Throttling** to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. Specify the required throughput value in Mbps and the time pattern. Bandwidth throttling is disabled by default. See "Bandwidth Throttling during Data Transfer" on page 394. Time-based throttling is based on the source server time.
Advanced Additional Source Machine Addresses	Specify additional IP addresses for source workloads to enable communication in environments that use network address translation (NAT). See "Migrations Across Public and Private Networks through NAT" on page 60.

Setting Name	Description
Schedule	
Schedule	Specify when to start the migration job: ◆ Run immediately ◆ Run at a later time Use the calendar menu to specify the date and time to begin the migration. **NOTE:** You must prepare the target machine prior to the scheduled time.The full replication cannot run unless the target machine is available. Migrate skips the scheduled full replication and retries it at the next scheduled time.
Access Settings	
Source Credentials	(Windows) Specify the account user name with local or domain-level administrative privileges and a valid password. Use this format: ◆ For domain member machines: *authority**principal* ◆ For workgroup member machines: *hostname**principal* (Linux) Specify the `root` or root-level user name and a valid password.
Target Credentials	
Alerts	
Receive Event Notifications	Specify whether to send email notifications for event conditions. You must configure an SMTP server to use this feature. See "Notification Service Using Migrate Client" on page 102.
Receive Progress Notifications	If you enable Event notifications, you can optionally receive progress notifications at a specified interval.
Send to Addresses	Add or remove valid email addresses for recipients of the notifications.
Take Control Settings	
Target Virtual Machine	Under Target Virtual Machine, click **Configure**, then specify the options for the virtual network and the TCP/IP settings for the replication NIC, then click **OK**.
Post-Migration	
Action	Specify a pre-configured action from the PlateSpin Migrate library.See "Managing Post-Migration Actions" on page 107.
Execution Parameters	Specify the command line command to run the selected action. You can specify a timeout for the execution.
Credentials	Specify the user name and password to use for the post-migration tasks. You can optionally use the source credentials.

7 In the Network Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Network Configuration	
Network Identification Settings for Windows	
Host Name	Specify the desired host name for the target machine.
Member of Domain / Workgroup	Select the required option and type the name of the domain or workgroup that you want the target machine to join.
Preserve Source Server's Domain Registration	Preserves domain registration and ensures that the source server domain registration remains intact during migration. If you disable this option, the source machine's domain account is transferred to the target machine. The source server still appears to be on the domain, but does not have a valid connection.
Domain Credentials	If the target machine is to be part of a domain, specify valid credentials for a user account with permission to add servers to the domain, such as a member of the Domain Admins group or Enterprise Admins group.
Network Identification Settings for Linux	
Host Name	On the Network Identification tab, specify the desired host name for the target machine.
DNS	Use the **Add**, **Edit**, and **Remove** buttons to manage DNS server entries for the new virtual machine.

8 In the Operating System and Applications Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Operating System and Application Configuration	
Windows Services (Target)	Select Windows services' start conditions on the target VM after cutover. Start options are **Automatic**, **Manual**, **Disabled**, and **Automatic (Delayed Start).** **To modify the settings:** 1. Click the Status column for the service, then select from the Windows start options. 2. When you are done setting services start states, click **OK**.

Setting Name	Description
Live Transfer Services (Source)	Specify the Windows services to stop on the source workload during data transfers. We recommend that all the non-VSS compliant services or antivirus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes. **To modify the settings:** 1. Select **Stopped** next to the service to be stopped for data transfer. 2. When you are done setting services to stop, click **OK**.
Linux Daemons (Target)	Specify the start states for daemons on the target VM after cutover. **To modify the settings:** 1. Click the **Run Level** column for the daemon, then select from run levels 0 through 6 and Boot (B), then click **OK**. 2. When you are done setting daemon start states, click **OK**.
Live Transfer Daemons (Source)	Specify the daemons to stop on the source workload during data transfers. **To modify the settings:** 1. Select **Stopped** next to the daemon to be stopped for data transfer. 2. When you are done setting daemons to stop, click **OK**.

9 In the Drive Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Drive Configuration	
Hard Drives	Specify drive and volume configurations to be migrated.
Disks	Specify the path to the hard disk on the target virtual machine.
Volumes	Select volumes to be included in the target for migration.
NTFS Cluster Size	(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see the Microsoft Support KB Article 140365.
Non-volume Storage	(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups	(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.

Setting Name	Description
Volume Groups	(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the **Converted Logical Volumes** section of the settings.
Converted Logical Volumes	(For Linux Workloads) Specify one or more LVM logical volumes to be migrated for a Linux workload.

10 In the Additional Items for Review section of the Migration Job window, review errors and messages about the workload configuration. You must resolve errors before you can submit the migration job.

11 Click **OK**.

Configuring Migration to a Physical Target (P2P, V2P) Using Migrate Web Interface

To initiate a peer-to-peer workload migration to a physical machine:

1 (Recommended) Use PlateSpin Analyzer to ensure that:

- Your source operating system and hardware are supported by PlateSpin Migrate.
- PlateSpin Migrate's X2P device driver database contains device drivers that your target requires for the operating system being ported.

See "Analyzing Suitability of Discovered Windows Workloads For Conversion to Physical Machines" on page 320.

2 Discover your source workload. See "Workload Discovery in the Migrate Web Interface" on page 302.

3 (Conditional) If drivers for the physical target are not available in the PlateSpin Migrate's X2P device driver database, upload the required drivers to the database.

See Chapter 19, "Preparing Device Drivers," on page 311.

4 Register your target physical machine with PlateSpin Migrate by booting it with the PlateSpin ISO.

See "Registering and Discovering Details for Target Physical Machines with PlateSpin ISO" on page 291.

5 Launch the PlateSpin Migrate Web Interface.

6 On the Workloads page, select the workload you want to configure.

7 Click **Configure Migration**.

8 Specify the **Initial Transfer Method** for replication based on the scope of data you want to transfer from the source to the target:

- **Full Replication:** Migrate replicates the full volume from the source to the target.
- **Incremental Replication:** Migrate replicates only differences in data from the source to the target, provided the workloads have similar operating system and volume profiles.

9 Depending on whether you want to perform full replication or incremental replication, do one of the following:

- **Full Replication:** Select the target physical machine to which you want to migrate the source workload.
- **Incremental Replication:** Select the target type as **Physical Targets.** From the list of the prepared target workloads, select a physical target that matches the source workload you want to migrate. If no matching physical target workload is listed, add a physical target and then boot up this physical target and use the PlateSpin Migrate ISO Image to register it with Migrate server.

10 Click **Configure Migration**.

11 Configure the following settings:

NOTE: All of the following settings are applicable for full replication of the workload. Some of the listed settings might not display or display differently when you configure incremental replication for the workload. For example:

- The **Virtual Machine Configuration Path** setting does not display because a target workload matching the source workload must already be available to transfer only the differences from the source to this target.
- The **Volume Mapping** setting displays the volumes and their mappings. You can edit the volume mappings.

Setting Name	Description
Schedule Settings	
Compression Level	These settings control whether data is compressed during transmission between the source and target workloads, and the level of data compression applied.See "Data Compression" on page 52.Select one of the following options: ◆ **Fast:** Consumes the least CPU resources on the source, but yields a lower compression ratio. ◆ **Optimal:** Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option. ◆ **Maximum:** Consumes the most CPU resources on the source, but yields a higher compression ratio.
Bandwidth Throttling	These settings control the bandwidth throttling. PlateSpin Migrate enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. You can specify a throughput rate for each migration job. Throttling provides a way to prevent migration traffic from congesting your production network and to reduce the overall load of your PlateSpin Server. To throttle replications to a specified rate, specify the required throughput value in Mbps and the time pattern.

Setting Name	Description
Migration Settings	
Transfer Method	(For Windows Workloads) Select a data transfer mechanism and security through encryption.See "Supported Data Transfer Methods" on page 44. To enable encryption, select the **Encrypt Data Transfer** option. See "Security and Privacy" on page 46.
Transfer Encryption	(For Linux Workloads) To enable encryption, select the **Encrypt Data Transfer** option.See "Security and Privacy" on page 46.
Source Credentials	Specify the credentials required for accessing the workload. See "Discovery Guidelines for Source Workloads" on page 299.
Volumes	Select volumes to be included in the target for migration. For Windows file-based migrations, you can optionally adjust the size of the target volume by modifying the **Size** value in the row for that volume. **NOTE:** The ability to adjust the size of the target volume is currently not available in the Disk Layout view of storage in the Migrate Web Interface. See "Cannot Adjust the Target Volume Size for File-Based Migrations in the Web Interface" on page 569.
Non-volume Storage	(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups	(For Linux Workloads) Displays the disks to which volume groups are mapped.
Volume Groups	(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the **Converted Logical Volumes** section of the settings.
Converted Logical Volumes	(For Linux Workloads) Specify one or more LVM logical volumes to be migrated for a Linux workload.
Replication Networks for Source	Specify one or more network interfaces (NIC or IP address) on the source to use for replication traffic.
Services to Stop Before Any Replication	(For Windows Workloads) We recommend that all the non-VSS compliant services or anti-virus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.
Services to Stop for Cutover with Replication	(For Windows Workloads) Select the Windows services that should be permanently stopped on the source workload for cutover with any replication. The services stopped on the source workload during the replication process are not restored after cutover. The stopped services are restored after a Test Cutover.
Daemons to Stop Before Any Replication	(For Linux Workloads) Select the Linux daemons that you want to be temporarily stopped on the source workload before replication. These daemons will be restored after replication completes.

Setting Name	Description
Daemons to Stop for Cutover with Replication	(For Linux Workloads) Select the Linux daemons that should be permanently stopped on the source workload for Cutover with any Replication. The daemons stopped on the source workload during the replication process are not restored after Cutover. The stopped daemons are restored after a Test Cutover.
Boot Services to Disable on Target	(For Windows Workloads) Select the boot services that you want to disable on the target workload.
Target Workload Settings (These settings are applied during the Run Cutover)	
Hostname	Do one of the following: ◆ To retain the same host name, select **No Change**. ◆ To change the host name, select **Set To** and specify the new name.

Setting Name	Description
Domain / Workgroup	(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays: ◆ **Workgroup:** *Workgroup_name* where *Workgroup_name* is the workgroup name to which the source belongs. ◆ **Domain:** *Domain_name* where *Domain_name* is the domain name to which the source belongs. Do one of the following depending on where you want the target workload to join: ◆ **When the source workload belongs to a workgroup:** Assume that the source workload belongs to a workgroup named **WorkGroup1**. ◆ For the target workload to join the same workgroup (**WorkGroup1**), retain the following existing selection: **Workgroup: Workgroup1** ◆ For the target workload to join a different workgroup (say **WorkGroup2)**, select **Join Workgroup** and specify the name as **WorkGroup2**. ◆ For the target workload to join a domain, select **Join Domain** and specify the domain name you want the target to join. ◆ **When the source workload belongs to a domain:** Assume that the source workload belongs to a domain named **Domain1**. ◆ For the target workload to join a workgroup, click **Join Workgroup** and specify the name of the workgroup you want the target to join. ◆ For the target workload to join the same domain (**Domain1**) with the domain registration settings preserved, retain the following existing selection: **Domain: Domain1** ◆ For the target workload to join the same domain (**Domain1)** without preserving the domain registration settings, select **Join Domain** and specify the domain name as **Domain1**. ◆ For the target workload to join a different domain, select **Join Domain** and specify the domain name you want the target to join.
Domain Credentials	(For Windows Workloads) If you select **Join Domain**, specify the domain administrator credentials.
Network Connections	Select the local area connection and then select one of the following: ◆ **DHCP:** Obtain an IP address automatically assigned by a DHCP server. ◆ **Static:** Specify one or more static IP addresses, subnet masks, and gateway IP addresses. Click **Advanced** to add multiple entries, edit an existing entry, or remove an existing entry. The IP addresses must be unique within the supported network. For Windows workloads that have more than one NIC, select the connection for each NIC.

Setting Name	Description
DNS Servers	(For Linux Workloads) Specify the DNS Servers for the target workloads. This is applicable only if you select **Static** in the **Network Connections** option: ◆ **Primary DNS server:** Specify the primary DNS server address. ◆ **Alternative DNS server:** Specify an alternate DNS server address. ◆ **Additional DNS server:** To specify additional DNS server addresses: 1. Click **Advanced**. 2. Specify the DNS server address. 3. Click **Add** to add the server in the DNS Server Addresses list. 4. Click **OK**.
Services States on Target VM	(For Windows Workloads) Select Windows services' start conditions on the target VM. Start options are Automatic, Manual, Disabled, and Automatic (Delayed Start).
Daemons States on Target VM	(For Linux Workloads) Select Linux daemons' start conditions on the target VM. Enable the daemon to start by selecting the check boxes at the appropriate runlevels (0 to 6) and Boot.
Post-Migration Action	Click **Configure Action** to select a previously configured custom action that you want to execute on your target workload as a part of the migration job. The command line parameters and timeout value that you previously configured for the selected action is displayed by default. You can edit these values as required. If you choose to timeout the action and have not specified any timeout value, then PlateSpin Migrate waits for a maximum period of 7200 seconds before timing out this action. For information about managing the post-migrations, see "Managing Post-Migration Actions" on page 107.
Tag	
Tag	Select a tag to assign to the workload. See "Managing Workload Tags" on page 133.

12 (Optional) To change the target, click **Change Target**.

NOTE: If you change the target, all the settings you specified will be cleared.

13 Do one of the following:

- Click **Save** to save the settings.
- Click **Cancel** to exit.

32 Synchronizing Workloads with Server Sync

The Server Sync feature enables you to reduce the scope of data that is transferred from your source to your target to just data that is different between a source and a target, effectively synchronizing their volume contents.

For example, when setting up a job for a workload migration operation, you can choose to update an existing physical or virtual machine to match the state of your source workload without transferring volume data in its entirety. PlateSpin Migrate compares the target physical or virtual workload with the selected source and transfers only data that is different between the two, overwriting files on the target with those on the source workload.

Server Sync is useful in situations where the size of volume data or network conditions are prohibitive for a direct source-to-target virtualization over the network.

Server Sync to a Virtual Target

1 Discover your source workload.

  See "Discovering Details for Source Workloads" on page 301.

2 Create a target virtual machine by using one of the following methods:

  - Do an initial migration of your workload to a virtual machine. See Chapter 24, "Configuration Essentials," on page 385.

    - OR -

  - Using your virtualization platform's native interface, manually install a virtual machine with the same operating system profile as that of your source.

    NOTE: When you are creating a virtual target for Server Sync, you should also manually install the appropriate virtualization enhancement tools, such as VMware Tools.

3 In the Servers view, drag your source workload and drop it on the required target (Server Sync target or discovered physical machine under control).

  The system validates the selected source and target and, if it detects matching operating systems on them, provides you with two **Transfer Scope** options, **Full Migration** and **Server Sync**:

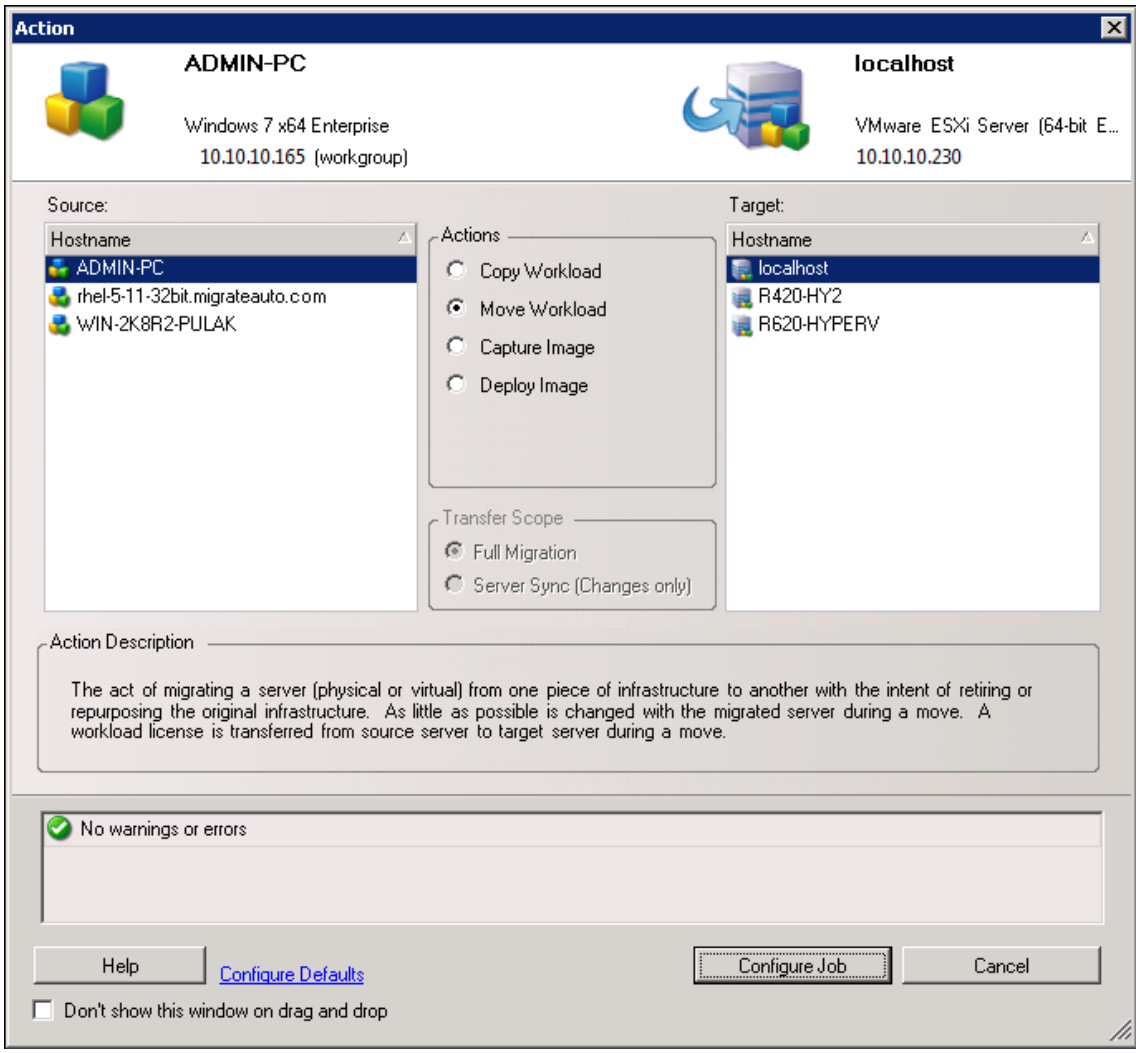

4 Select the **Server Sync** option, then click **Configure Job**.

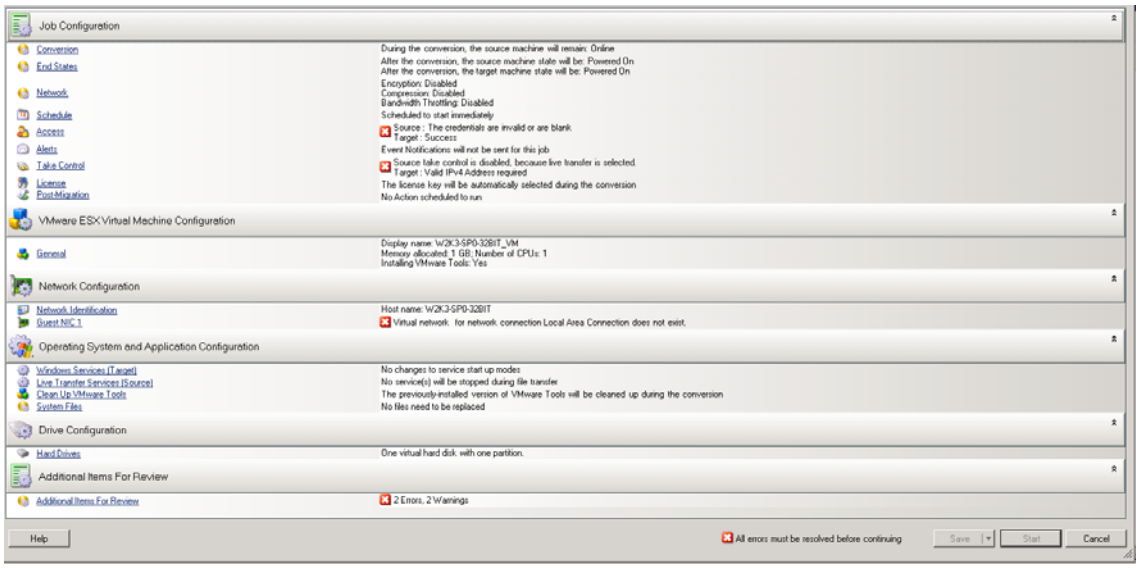

5 In the job configuration window, specify the parameters of the job as dictated by the purpose of the operation, address any warnings and errors, and ensure that you map the required volumes on the source to those on the target (see "Server Sync Volume Mapping" on page 532).

When you have finished, click **Start**.

PlateSpin Migrate starts the job and lists it in the Jobs view.

Server Sync to a Physical Target

1 Discover your source workload.

See "Discovering Details for Source Workloads" on page 301.

2 Discover your physical target by using the appropriate PlateSpin ISO boot image.

See "Registering and Discovering Details for Target Physical Machines with PlateSpin ISO" on page 291.

3 In the Servers view, drag your source workload and drop it on the required target (Server Sync target or discovered physical machine under control).

The system validates the selected source and target and, if it detects matching operating systems on them, it provides you with two **Transfer Scope** options, **Full Migration** and **Server Sync**, similar to the "Server Sync to a Virtual Target" on page 527 (see Step 4).

4 Select the **Server Sync** option, then click **Configure Job**.

5 In the job configuration window, specify the parameters of the job as dictated by the purpose of the operation, address any warnings and errors, and ensure that you map the required volumes on the source to those on the target.

6 When you have finished, click **Start**.

PlateSpin Migrate starts the job and lists it in the Jobs view.

Selective Server Sync to a Physical or Virtual Target

When you are using Server Sync to synchronize two Windows or Linux workloads, PlateSpin Migrate Client provides you with the capability to select the sources volumes that you want to synchronize with the target. Consider a scenario where only the data volumes might have changed post the replication of the workloads. In such a case, you might want to synchronize only the data volumes and exclude the boot and system volumes from synchronizing.

1 Discover your source workload.

See "Discovering Details for Source Workloads" on page 301.

2 Discover your physical or virtual target.

3 In the Servers view, drag your source workload and drop it on the required target (Server Sync target or discovered physical machine under control).

The system validates the selected source and target and, if it detects matching operating systems on them, it provides you with two **Transfer Scope** options, **Full Migration** and **Server Sync**, similar to the "Server Sync to a Virtual Target" on page 527 (see Step 4).

4 Select the **Server Sync** option, then click **Configure Job**.

5 In the job configuration window, specify the parameters of the job as dictated by the purpose of the operation, address any warnings and errors, and ensure that you map the required volumes on the source to those on the target.

6 In the **Drive Configuration** section of the Migration Job window, click the **Volume Mapping** or **Drives and Volumes** option displayed depending on the target type.

7 Configure the Server Sync volume configuration options.

The following topics provide information about how to select volume configuration options specific to Windows and Linux workloads.

- "Server Sync Volume Configuration (Windows)" on page 530.
- "Server Sync Volume Configuration (Linux)" on page 531.

8 When you have finished, click **Start**.

PlateSpin Migrate starts the job and lists it in the Jobs view.

Server Sync Volume Configuration (Windows)

A Server Sync job for Windows workloads provides detailed drive and volume information for both the source and the target, and enables you to specify the required mapping. For the volumes that you do not want to synchronize, set the mapping to **None**. For information about mapping the volumes, see "Server Sync Volume Mapping" on page 532.

NOTE

- Either include or exclude all the OS volumes (boot and system volumes) from synchronizing the changes. If you exclude an OS volume (boot or system volume), then PlateSpin Migrate Client notifies you that all the OS volumes must be excluded.
- Do not exclude the OS volumes (boot or system volumes), if you are using BBT Driver for X2P replications.
- At least one volume must be included.

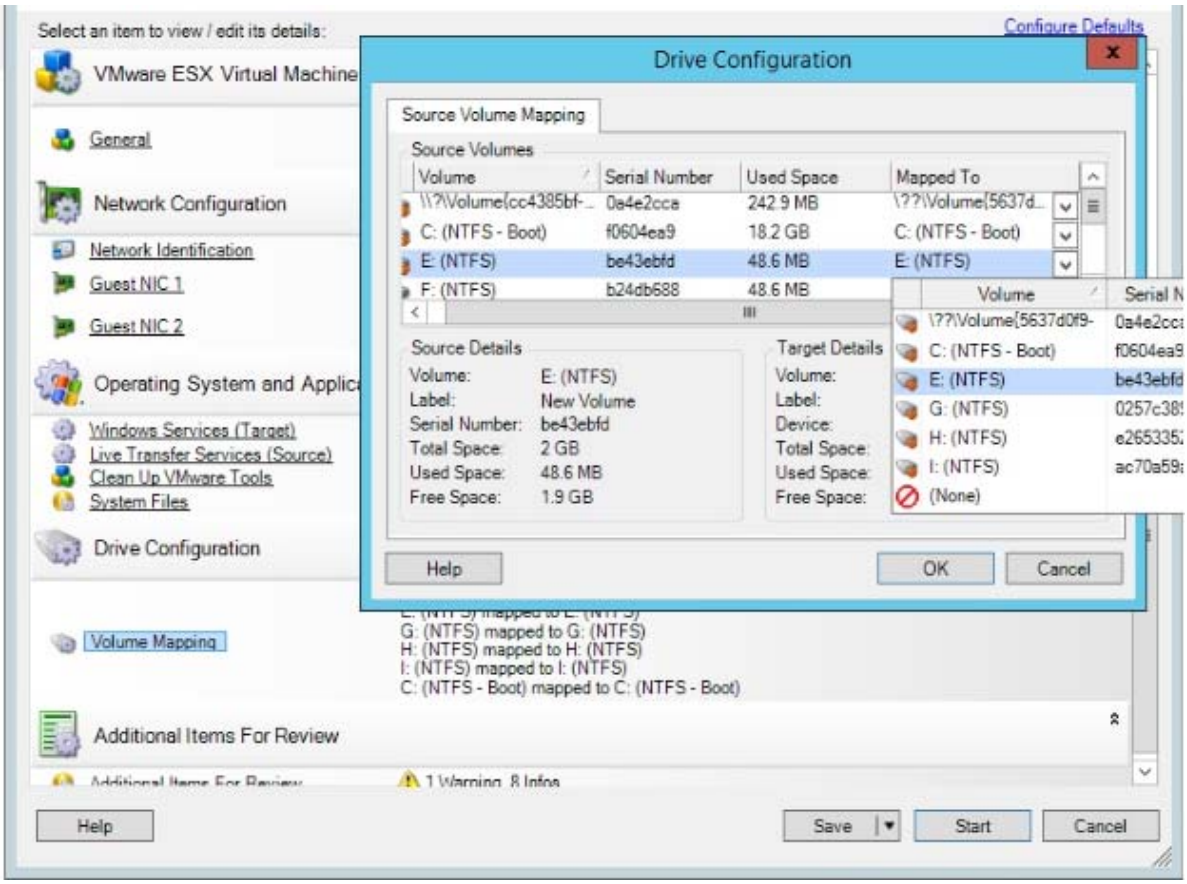

Server Sync Volume Configuration (Linux)

A Server Sync job for Linux workloads provides detailed mount point and volume information for both the source and the target, and enables you to specify the required mapping. For the volumes that you do not want to synchronize, set the mapping to **None**. For information about mapping the volumes, see "Server Sync Volume Mapping" on page 532.

NOTE

- Either include or exclude all the OS volumes (boot and system volumes) from synchronizing the changes. If you exclude an OS volume (boot or system volume), then PlateSpin Migrate Client notifies you that all the OS volumes must be excluded.
- Do not exclude the OS volumes (boot or system volumes), if you are using BBT Driver for X2P replications.
- At least one volume must be included.

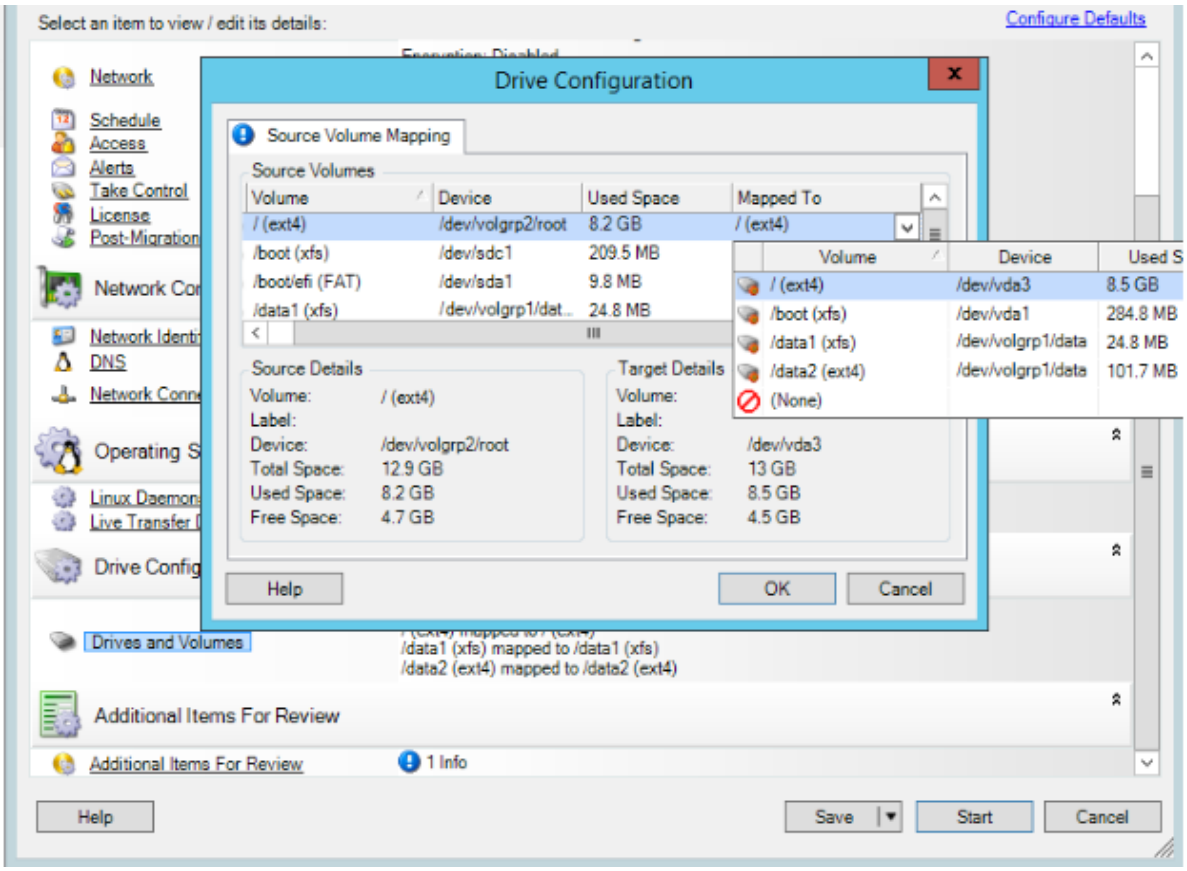

Server Sync Volume Mapping

When you are using Server Sync to synchronize two Windows or Linux workloads, PlateSpin Migrate Client provides you with the capability to specify the required mapping between source volumes and existing volumes on the target. See “Synchronizing Workloads with Server Sync” on page 527.

To access volume configuration options in a Server Sync job:

1. In the Jobs view, select the required workload.
2. In the **Drive Configuration** section of the Migration Job window, click the **Volume Mapping** or **Drives and Volumes** option displayed depending on the target type.
3. Configure the Server Sync volume configuration options.

   The following topics provide information about Server Sync volume configuration options specific to Windows and Linux workloads.

- “Server Sync Volume Configuration (Windows)” on page 533
- “Server Sync Volume Configuration (Linux)” on page 534

Server Sync Volume Configuration (Windows)

A Server Sync job for Windows workloads provides detailed drive and volume information for both the source and the target, and enables you to specify the required mapping.

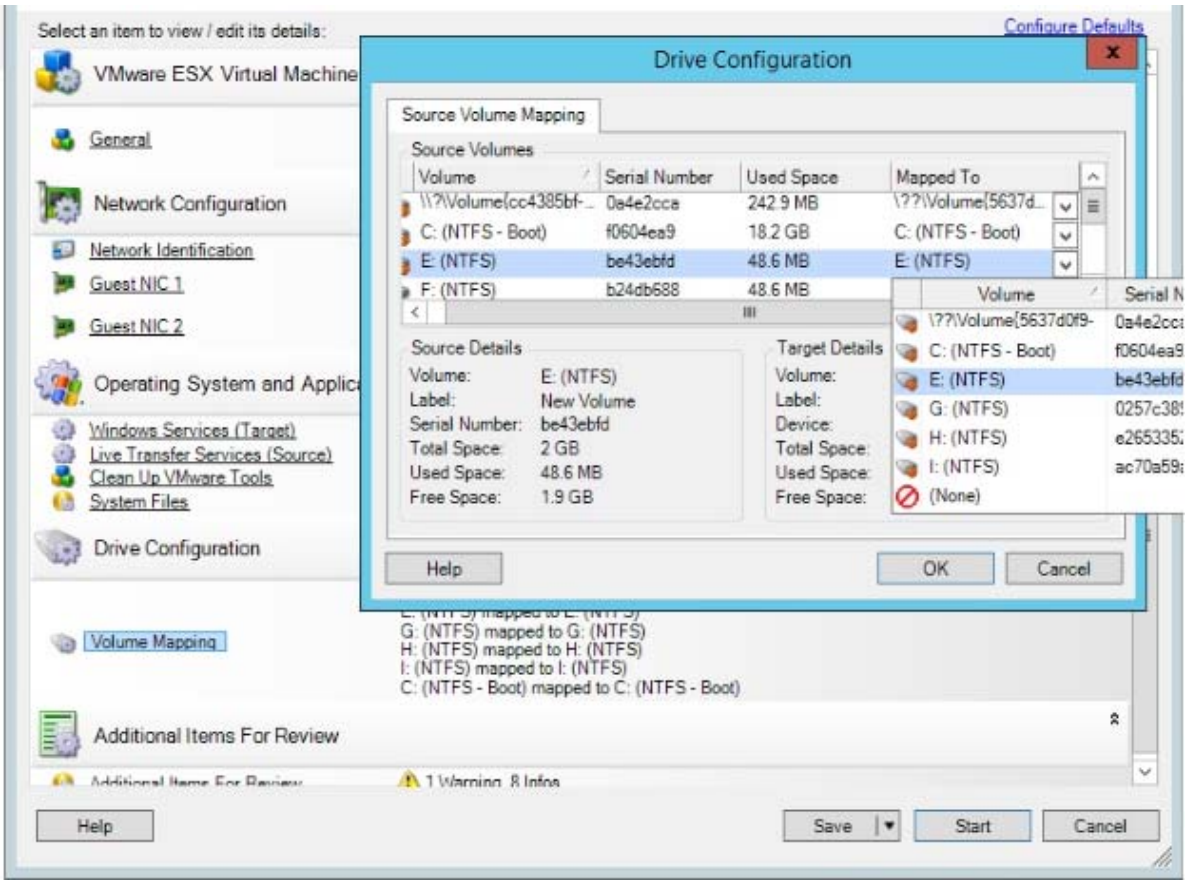

Mapped To: Map each volume on the source to an existing volume on the target.

Server Sync Volume Configuration (Linux)

A Server Sync job for Linux workloads provides detailed mount point and volume information for both the source and the target, and enables you to specify the required mapping.

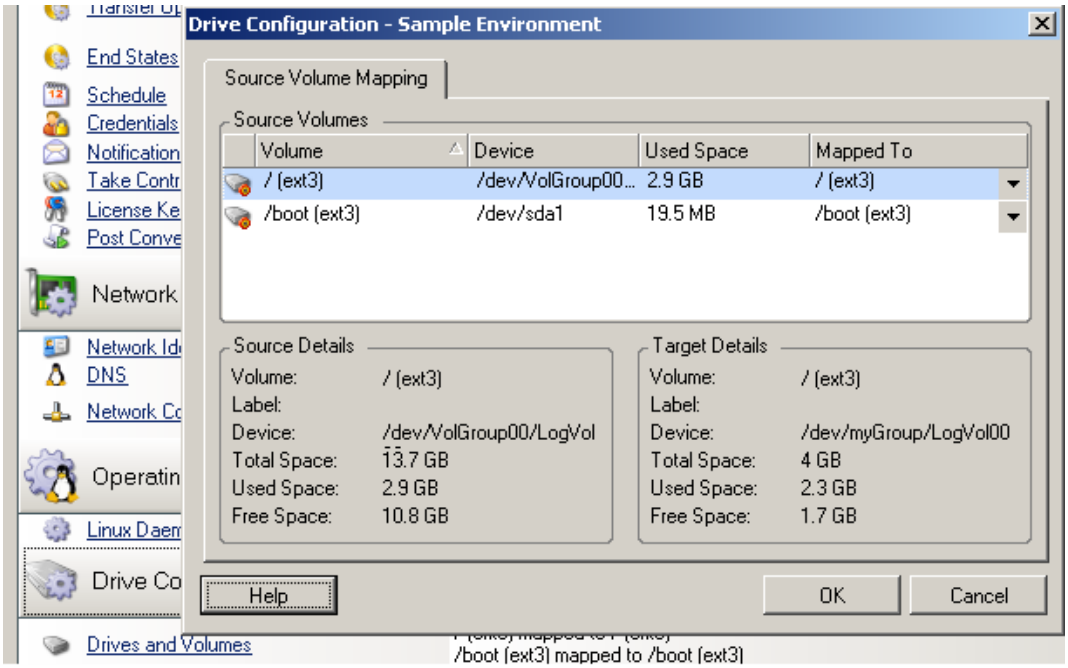

Mapped To: Map each volume on the source to an existing volume on the target.

VI Executing Migrations

After you configure migration settings for the workload, you are ready execute the migration. Ensure that the target VMs are prepared for migration, then begin replicating data to the target. You can monitor the health of migration jobs and generate reports about them.

33 Executing Workload Migrations

After you discover and configure workloads for migration, you execute and monitor the migration by performing the migration tasks described in this section. Use the PlateSpin Migrate Web Interface or PlateSpin Migrate Client as appropriate for the migration types and target platforms. See "Migration Tasks Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface" on page 84.

Preparing a Migration

After you configure a workload for migration, PlateSpin Migrate uses the migration settings to install any required data transfer software on the source workload and create a target workload on the target platform.

Using the Migrate Client

When you start a migration job from the PlateSpin Migrate Client, PlateSpin Migrate validates the job type, the source, the target, and the selected parameters, and might generate errors and warnings.

Error markers show configurations that you need to change before the migration job can start.

Warning markers alert you to settings that should be confirmed prior to starting the migration.

In a default PlateSpin Migrate configuration, validation messages display at the bottom of the Action window. However, If you have configured PlateSpin Migrate to bypass the Action window on drag-and-drop, errors and warnings are displayed in a separate window:

Figure 33-1 *Migration Validation Window*

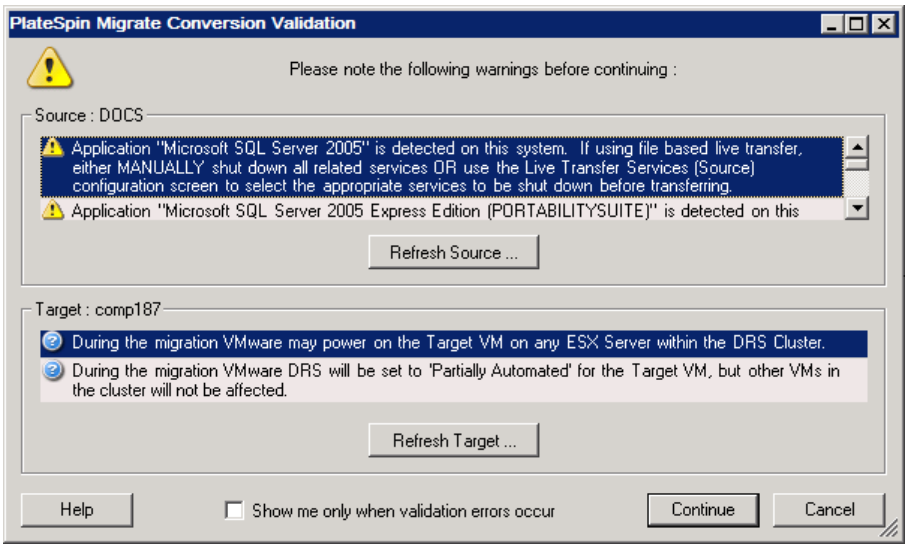

To force this window to open only on errors, select **Show me only when validation errors occur**.

Using the Migrate Web Interface

To immediately prepare the workload for migration:

1 On the Edit Migration Details page, click **Save and Prepare**.

To prepare a preconfigured workload for migration:

1 On the Workloads page, select the preconfigured workload you want to migrate.

2 Click **Prepare Migration**.

Starting Migration Execution (First Replication)

After the migration preparation completes successfully, the migration is ready for execution. Execution begins with the first replication. The first replication is a full replication with a Full Replication contract type or an incremental data synchronization for a pre-existing target workload with an Incremental Replication contract type.

The first replication is unscheduled by default. You can manually start the first replication. You can alternatively schedule the date and time to run the first replication. See "Scheduling Migration Execution (First Replication)" on page 539.

NOTE: You must prepare the source and target workload prior to the manual start. The full replication cannot run unless the target workload exists and the workload preparation is complete. See "Preparing a Migration" on page 537.

Using the Migrate Client

To manually start the first replication:

1. In the Jobs view, locate the prepared workload that you want to migrate.
2. Right-click the job and select **Start**.

   PlateSpin Migrate starts the first full replication for the workload.

Using the Migrate Web Interface

To manually start the first replication:

1. On the Workloads page, select the prepared workload that you want to migrate.
2. Click **Run Migration**.
3. On the Workload Commands page, do one of the following depending on the migration contract type you configured for the workload:
   - **Full Replication:** Select **Full Replication** as the replication method.
   - **Incremental Replication:** Select **Incremental Replication** as the replication method.
4. (Optional) Set the following options as appropriate if you want to cut over the workload after a successful manual replication:
   - Run cutover after successful replication
     - Shut down source after cutover
     - Shut down target after cutover

   NOTE: For workloads in the cloud, the end-state shutdown stops the VM and deallocates its resources.
5. Click **Execute**.

   PlateSpin Migrate starts the first replication for the workload.

Scheduling Migration Execution (First Replication)

After the migration preparation completes successfully, the migration is ready for execution. Execution begins with the first replication, which might be a full replication or a data synchronization for a pre-existing target workload.

The default schedule setting is None. The first replication is unscheduled. You can schedule the start date and time to run the first replication. You can alternatively start the first replication manually. See "Starting Migration Execution (First Replication)" on page 538.

The first replication for a scheduled migration execution is a one-time event, but the run is attempted daily as scheduled until the first replication begins and completes successfully.

NOTE: You must prepare the workload prior to the scheduled time or the manual start. The first replication cannot run unless the target workload exists and the workload preparation is complete. If they are not ready, Migrate skips the scheduled replication and retries it at the scheduled time on the next day.

- "Using the Migrate Client" on page 540
- "Using the Migrate Web Interface" on page 540

Using the Migrate Client

To modify the start date and time for the first replication:

1. In the Jobs view, locate the required job.
2. Right-click the job and select **Change Start Time** to open the Change Job Start Time dialog box.

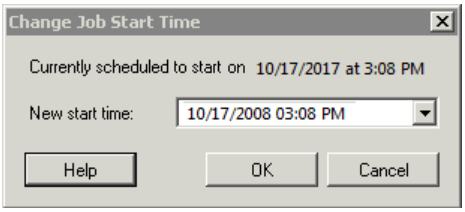

3. Specify the required start date and time, then click **OK**.

   PlateSpin Migrate reschedules the job and executes it at the specified time.

Using the Migrate Web Interface

To modify the start date and time for the first replication:

1. On the Workloads page, locate and click the workload.
2. On the Migration Details page, click **Edit**.
3. On the Edit Migration Details page, go to **Schedule Settings** > **Full Replication**, then click **Edit**.
4. Click **Start**, then set the date and time when you want to start the first full replication.

   You can type the date (dd/mm/yyyy) or click the Calendar icon to select the date. The default run time is 12:00:00 a.m. (hh:mm:ss a.m. or p.m.).
5. Click **Close** to return to the Edit Migration Details page, then click **Save**.

Starting Incremental Replications

After the first replication completes successfully, you can start each incremental replication manually. You can alternatively schedule the time and pattern to run incremental replications that occur after the first replication. See "Scheduling Incremental Replications" on page 541.

- "Using the Migrate Web Interface" on page 541

Using the Migrate Web Interface

To start an incremental replication manually:

1. On the Workloads page, locate and select the workload.
2. Click **Run Migration**.
3. On the Workload Commands page, select **Incremental Replication** as the replication method.
4. (Optional) Set the following options as appropriate if you want to cut over the workload after a successful manual replication:
   - Run cutover after successful replication
     - Shut down source after cutover
     - Shut down target after cutover

   NOTE: For workloads in the cloud, the end-state shutdown stops the VM and deallocates its resources.
5. Click **Execute**.

   PlateSpin Migrate starts the incremental replication for the workload.

Scheduling Incremental Replications

After you configure and save a workload migration, you can modify the time and pattern to run incremental replications that occur after the first replication. You can alternatively start each incremental replication manually. See "Starting Incremental Replications" on page 541.

NOTE:

- Scheduled incremental replications are skipped until the first full replication is complete.
- Scheduled incremental replications take place for a maximum period of 60 days from the time that the scheduled incremental replication runs begin.

- "Using the Migrate Web Interface" on page 541

Using the Migrate Web Interface

To schedule the incremental replication recurrence time and pattern:

1. On the Workloads page, locate and click the workload.

2 On the Migration Details page, click **Edit**.

3 On the Edit Migration Details page, go to **Schedule Settings** > **Incremental Recurrence**, then click **Edit**.

The default Incremental recurrence setting is None. The incremental replications are unscheduled.

4 For **Begin the recurrence schedule**, set the date and time when you want to begin the scheduled incremental replications.

You can type the date (dd/mm/yyyy) or click the Calendar icon to select the date. The default run time is 12:00:00 a.m. (hh:mm:ss a.m. or p.m.).

5 For **Recurrence run setting**, set the pattern to follow for scheduled incremental replications:

- **Daily:** The replication takes place on the specified daily intervals or on weekdays every week for a period of 60 days from the time the replication starts.
- **Weekly:** The replication takes place at specified intervals for a period of 8 weeks from the time the replication starts.
- **Monthly:** The replication takes place at specified intervals for a period of 2 months from the time the replication starts.

6 Click **Close** to return to the Edit Migration Details page, then click **Save**.

Performing Workload Test Cutover or Cutover

After a workload is successfully replicated, you can either test the workload even before the actual cutover or perform an actual cutover of the workload to complete the migration.

- "Using the Migrate Web Interface" on page 542

Using the Migrate Web Interface

To perform test cutover of the workload:

1 On the Workloads page, select the replicated workload that you want to migrate.

2 Click **Test Cutover**.

3 (Optional) On the Workload Commands page, set the following options as appropriate:

- Shutdown Target Workload
- Perform Incremental Replication

4 Click **Execute**.

PlateSpin Migrate starts the test cutover of the workload. On successful completion of a Test Cutover operation, the system generates an event associated with two tasks: **Mark Test as Success** and **Mark Test as Failure.**

5 Depending on the status of test cutover, click one of the appropriate tasks to cancel the Test Cutover operation and log a corresponding event.

To perform cutover of the workload:

1 On the Workloads page, select the replicated workload that you want to migrate.

2 Click **Run Cutover**.

3 (Optional) On the Workload Commands page, set the following options as appropriate:
   - Shutdown Source (Selected by default)
   - Shutdown Target Workload
   - Perform Incremental Replication (Selected by default)
4 Click **Execute** to create the cutover workload.

Viewing Properties for an In-Progress or Completed Migration

After you add a workload to PlateSpin Migrate, the Configuration page displays the properties of the workload's migration configuration throughout the migration lifecycle.

- "Using the Migrate Client" on page 543
- "Using the Migrate Web Interface" on page 543

Using the Migrate Client

To view properties for a workload migration:

1 In the Jobs view, locate the required job.
2 Right-click the job and select **View**.
   Migrate Client opens the job configuration window.
3 View the workload migration configuration parameters and settings in read-only mode.

Using the Migrate Web Interface

To view properties for a workload migration:

1 On the Workloads page, locate and click the workload.
   Migrate Web Interface opens the Migration Details page.
2 View the workload migration configuration parameters and settings in read-only mode.

Canceling an In-Progress Migration

You might need to cancel an in-progress workload migration that has become non-responsive.

- "Using the Migrate Client" on page 543
- "Using the Migrate Web Interface" on page 544

Using the Migrate Client

1 In the Jobs view, locate the required job.
2 Right-click the job and select **Abort**.

Using the Migrate Web Interface

1. On the Workloads page, locate and click the stalled workload.
2. View the replication or cutover status.
3. Click **Abort**.

Restarting or Shutting Down the Source Workload

PlateSpin Migrate Client enables you to restart or shut down a source workload if the migration job is not active.

To shut down or restart the source workload from the Migrate Client:

1. In the Jobs view, locate the required job.
2. Right-click the job and select **Restart Source** or **Shutdown Source** as applicable.

To automate the startup state of source and target workloads, specify the required post-migration state in your migration job. See "Post-Cutover End States for Source and Target Workloads" on page 409.

34 Generating Reports

You can generate reports about discovered workloads and the workload migrations by using PlateSpin Migrate Client or PlateSpin Migrate Web Interface. For information about generating a Licensing Report, see "License Key Management with Migrate Client" on page 98.

Generating Workload and Workload Migration Reports

You can generate detailed reports of running and completed migration jobs. A migration report records the tasks performed during the job.

Generate Reports using the Migrate Client

To generate a job report:

1. In the Jobs view, locate the required job.
2. Right-click the job and select **Report**.

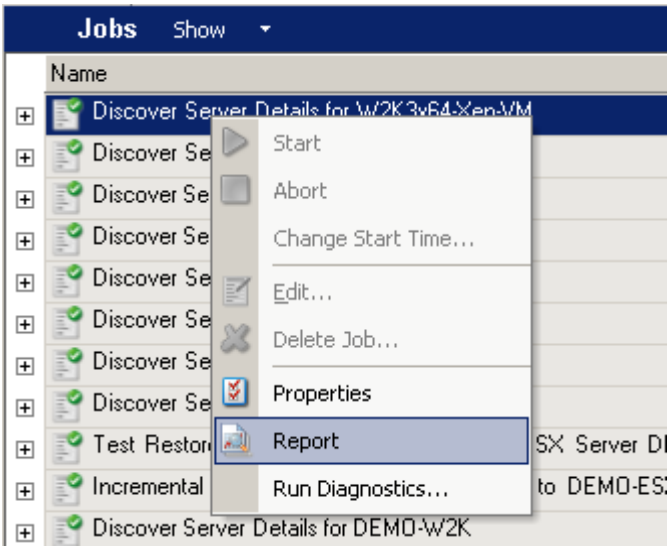

A Web browser window displays the requested report.

Generate Reports using the Web Interface

PlateSpin Migrate Web provides reports that give analytical insight into your workload migration contracts over time. See Table 34-1 for a list of available reports. The reports open in the Web Interface. You can print a report using browser options, or export it to XML.

Table 34-1 *Reports Available in PlateSpin Migrate Web Interface*

Report	Description
Workload Migration	Reports replication events for all workloads over a selectable time window.
Migration History	Reports replication size, and time per selectable workload over a selectable time window.
Replication Statistics	Reports the dynamics of full and incremental replications that can be summarized by **Average**, **Most Recent**, **Sum**, and **Peak** perspectives.
Current Migration Status	Displays the migration status such last test cutover, last replication date, and the test age (time elapsed since the last test cutover completed).
Events	Reports system events for all workloads over a selectable time window.
Scheduled Events	Reports only upcoming workload migration events.
Running Events	Reports only migration events that are running at the time that the report is generated.
Resource Usage	Displays the resources configured for the target workload.

To generate a report:

1. In your PlateSpin Migrate Interface, click **Reports**.

   A list of the report types is displayed.
2. Click the name of the required report type.
3. Select one or more workloads for which you want to create the report.
4. Configure the time period for which you want to view the report.
5. Do one of the following:
   - Click **Printable View** to view the report in your web browser.
   - Click **Export to XML**, then save the XML file to your computer.

Generating Diagnostic Reports

- "Using the Migrate Client" on page 547
- "Using the Migrate Web Interface" on page 548

Using the Migrate Client

PlateSpin Migrate provides a tool that can produce a diagnostics report for any running or completed job.

To view a diagnostics report:

1. In the Jobs view, right-click the required job and select **Run Diagnostics**.
2. Click **OK** to dismiss the notice that the Diagnostics report has started.

   The process might take a few moments.
3. The Diagnostics report is displayed in your web browser.

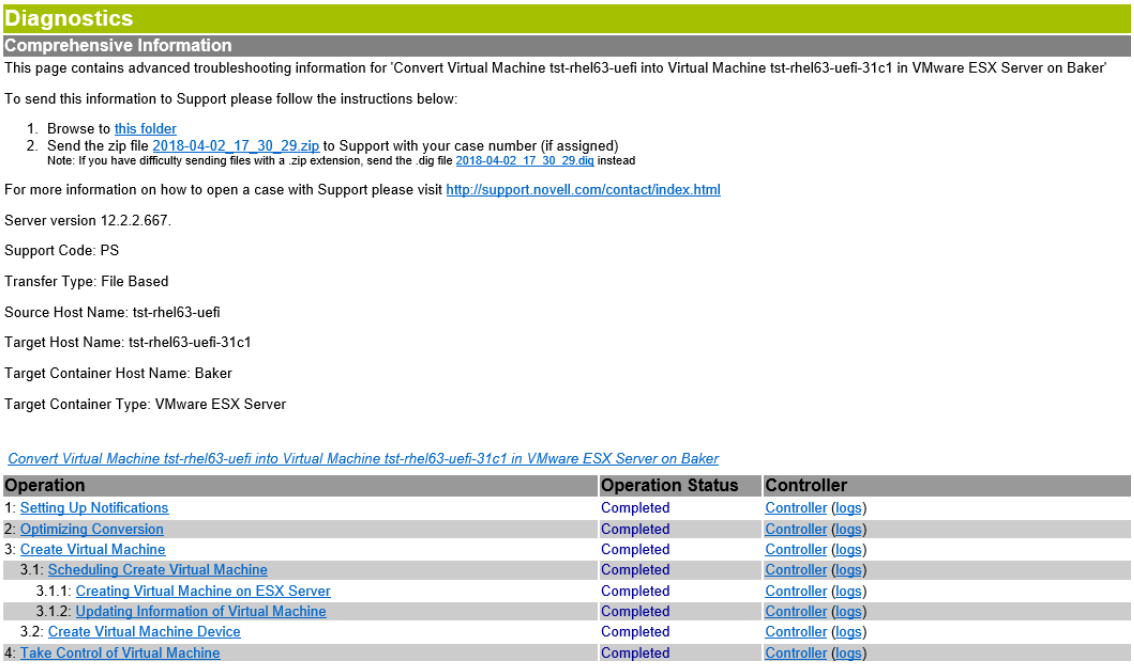

Diagnostics

Comprehensive Information

This page contains advanced troubleshooting information for 'Convert Virtual Machine tst-rhel63-uefi into Virtual Machine tst-rhel63-uefi-31c1 in VMware ESX Server on Baker'

To send this information to Support please follow the instructions below:

1. Browse to this folder
2. Send the zip file 2018-04-02_17_30_29.zip to Support with your case number (if assigned)
   Note: If you have difficulty sending files with a .zip extension, send the .dig file 2018-04-02_17_30_29.dig instead

For more information on how to open a case with Support please visit http://support.novell.com/contact/index.html

Server version 12.2.2.667.

Support Code: PS

Transfer Type: File Based

Source Host Name: tst-rhel63-uefi

Target Host Name: tst-rhel63-uefi-31c1

Target Container Host Name: Baker

Target Container Type: VMware ESX Server

Convert Virtual Machine tst-rhel63-uefi into Virtual Machine tst-rhel63-uefi-31c1 in VMware ESX Server on Baker

Operation	Operation Status	Controller
1: Setting Up Notifications	Completed	Controller (logs)
2: Optimizing Conversion	Completed	Controller (logs)
3: Create Virtual Machine	Completed	Controller (logs)
3.1: Scheduling Create Virtual Machine	Completed	Controller (logs)
3.1.1: Creating Virtual Machine on ESX Server	Completed	Controller (logs)
3.1.2: Updating Information of Virtual Machine	Completed	Controller (logs)
3.2: Create Virtual Machine Device	Completed	Controller (logs)
4: Take Control of Virtual Machine	Completed	Controller (logs)

The diagnostics report lists several statistics:

- All the operations involved in the job. Click any operation to view its XML representation.
- The status of each operation.
- The controller that ran the operation. Click the controller to view its XML representation, or click **Logs** to view its event log.

In addition, the report contains links to:

- The XML representations of the source machine, original target machine, and the target VM host.
- The root operation for the job, as well as a variety of logs and reports.

You can send diagnostics reports directly to Technical Support. Follow the instructions provided on the report.

Using the Migrate Web Interface

In the Migrate Web Interface, after you have executed a command, you can generate detailed diagnostic reports about the command's details.

1. Click **Command Details,** then click the **Generate** link in the lower right of the panel.

   After a few moments, the page refreshes and displays a **Download** link above the **Generate** link.

2. Click **Download**.

   A `.zip` file contains the comprehensive diagnostic information about the current command.

3. Save the file, then extract the diagnostics to view them.

4. Have the `.zip` file ready if you need to contact Technical Support.

35 Post-Migration Tasks

The following sections describe tasks you might need to perform after a workload migration.

Cleaning Up Windows Workloads

The following are instructions for cleaning up Windows workloads by component and use case.

Table 35-1 *Use Cases and Instructions for Cleaning Up Windows Workloads*

Component	Use Case	Removal Instructions
File-based Transfer Component	All Migrations	At root level for each volume migrated, remove all files named `PlateSpinCatalog*.dat`
Workload discovery software	All migrations	1. In the Servers view, undiscover the source (right-click, then select **Undiscover**). 2. In the source workload's `Windows` directory: ◆ Remove all files named `machinediscovery*`. ◆ Remove the subdirectory named `platespin`.
Controller software	All migrations	1. In the Servers view, undiscover the source (right-click, then select **Undiscover**). 2. Open a command prompt and change the current directory to: ◆ `\Program Files\platespin*` (32-bit systems) ◆ `\Program Files (x86)\platespin` (64-bit systems) 3. Run the following command: `ofxcontroller.exe /uninstall` 4. Remove the `platespin*` directory

Cleaning Up Linux Workloads

The following are instructions for cleaning up Linux workloads by component and use case.

Table 35-2 *Use Cases and Instructions for Cleaning Up Linux Workloads*

Component	Use Case	Removal Instructions

Controller software	All migrations	1. Stop the OFX controller process:
		`/etc/init.d/ofxcontrollerd stop`
		2. Remove the OFX controller service:
		`chkconfig --del ofxcontrollerd`
		3. Clean up the OFX controller files:
		◆ `rm -rf /usr/lib/ofx`
		◆ `rm -f /etc/init.d/ofxcontrollerd`

Block-level data transfer software	All block-level migrations	1. Check if the driver is active:
		`lsmod | grep blkwatch`
		If the driver is still loaded in memory, the result should contain a line, similar to the following:
		`blkwatch_7616  70924  0`
		2. (Conditional) If the driver is still loaded, remove it from memory:
		`rmmod blkwatch_7616`
		3. Remove the driver from the boot sequence:
		`blkconfig -u`
		4. Remove the driver files by deleting the following directory with its contents:
		`rm -rf /lib/modules/<kernel-version>/platespin`
		For example:
		`rm -rf /lib/modules/3.0.101-63-default/platespin`
		You can alternately use a variable `$(uname -r)` to dynamically retrieve the kernel version for the directory name:
		`rm -rf /lib/modules/$(uname -r)/platespin`
		5. Delete the following file:
		`/etc/blkwatch.conf`

Component	Use Case	Removal Instructions
LVM snapshots	Block-level migrations using LVM snapshots	1. In the Jobs view, generate a Job Report for the failed job, then note the name of the snapshot. 2. Remove the snapshot device by using the following command: `lvremove`*snapshot_name*

H Troubleshooting PlateSpin Migrate

This section provides a series of topics about troubleshooting PlateSpin Migrate.

For information about common problems that occur during discovery of workloads or targets, see Appendix C, "Troubleshooting Discovery," on page 337.

Migration of Workloads to Azure Cloud

Use information in this section to help troubleshoot common problems that might occur during migration of workloads to Microsoft Azure Cloud.

Assigning a Reserved IP Address to a Migrate Server in Azure

In Azure, the Dynamic assignment method is the default setting for the public IP address. The IP address can change every time the server is stopped and started. You should modify the setting to use the Static assignment method. Using a reserved IP address ensures that Azure allocates and reserves an IP address for the life of the resource.

NOTE: A change in IP address on the PlateSpin Server breaks the heartbeat communications with source workloads.

To apply a reserved IP address to an existing Migrate Server in Azure that has a dynamic IP address:

1 Specify **Static** as the assignment method for the public IP address of the Migrate Server resource:

   1a Use the Azure Global Portal (https://portal.azure.com/) to log in to your Azure account.

   1b Open Resources, select the Migrate Server resource, then select **Stop**.

   1c In the information for the Migrate Server, select the public IP address.

   1d In the **Public IP Address Configuration** panel under **Settings**, select **Configuration**.

   1e Specify **Static** as the assignment method for the Public IP address.

   1f Click **Save**.

   Azure allocates and reserves an IP address from a pool of its available IP addresses in the Azure location where you deploy the Migrate server.

   1g Start the Migrate Server resource.

   Heartbeat communications for existing migration jobs will be broken until you modify the server IP address stored in the OFX Controller configuration file on the source workload.

2 For each source workload that has already been configured for migration on the Migrate Server, use the Migrate Agent to set the new IP address:

```
migrateagent.cli.exe config /
setting=psserver:<new_ipaddress_or_dns_name>
```

   The `psserver` option stops the OFX Controller (`ofxcontroller`) service, modifies the `OfxController.exe.config` file with the new address, and restarts the service. Heartbeat communications now work with the server's new IP address.

Install Azure Agent Option Is Not Available for a Source Linux Workload

Issue: For workload migrations to Azure, the **Install Azure Agent** option is deselected and dimmed (not available) for selection when you configure a source Linux workload. The message indicates that Python 2.7 is required. You installed Python 2.7, but the option is still not available.

Fix: Python 2.7 must be installed on the source workload prior to discovery to make the information available to PlateSpin Migrate. After you update Python to version 2.7, remove the workload from PlateSpin Migrate, then re-add/rediscover the workload. The Install Azure Agent option will be available when you configure the workload migration.

Outbound Email Stuck after Migrating Microsoft Exchange Server 2016 to Azure Cloud

Issue: After you migrate a Microsoft Exchange 2016 server to Microsoft Azure, the user's outgoing messages get stuck in the `Drafts` folder of their Microsoft Outlook application.

Fix: After you migrate a Microsoft Exchange Server workload to Microsoft Azure, ensure that you modify the Exchange internal and external DNS settings to use **Microsoft Hyper-V Network Adapter**. See KB Article 7021909 (https://support.microfocus.com/kb/doc.php?id=7021909).

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.

Fix: To boot the target VM in Normal Mode:

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

Linux Disks or Partitions on the Target Are in a Different Order Than on the Source

Issue: Because Azure controls the mapping of devices, Linux disks might be created in a different order on the target workload than they are on the source workload. All migrated disks and partitions are present; only the disk order differs.

Fix: None. The target VM is fully functional.

Migration of Workloads to vCloud

Use information in this section to help troubleshoot common problems that might occur during migration of workloads to VMware vCloud Director.

- "Duplicate MAC Address Alarm for a VM Migrated to vCloud" on page 555

Duplicate MAC Address Alarm for a VM Migrated to vCloud

Issue: Alarms for duplicate MAC addresses are observed when a VM is deployed to a VMware vCenter 6.*x* Server hosted in a VMware vCloud virtual private cloud.

Fix: This is a known issue for VMware vCloud Director. See the VMware KB Article *Duplicate MAC address alarms are present when a VM is deployed in vCloud Director* (2148579) (https://kb.vmware.com/s/article/2148579).

Migration of Workloads to VMware

Use information in this section to help troubleshoot common problems that might occur during migration of workloads to VMware.

Outbound Email Stuck after Migrating Microsoft Exchange Server 2016 to VMware

Issue: After you migrate a Microsoft Exchange 2016 server to VMware, the users' outgoing messages get stuck in their `Drafts` folder.

Fix: After you migrate a Microsoft Exchange Server workload to VMware, ensure that you modify the Exchange internal and external DNS settings to use **VMXNET 3**. See KB Article 7021909 (https://support.microfocus.com/kb/doc.php?id=7021909).

Mouse Does Not Work in the VM Console Window for the Target VM

Issue: Sometimes on Test Cutover or Cutover, the mouse does not work for the VM in the vSphere Web Client. That is, when you perform **Actions** > **Open Console** to open the VMware Web Console, the mouse pointer does not function properly within the virtual machine console window.

Fix: Manually restart the VM to allow VMware Tools to recognize the USB Controller for the mouse. In vSphere, select **Actions** > **Power** > **Restart Guest OS**.

Cleaning Up Floppy Devices on the Cutover VM on VMware

During migrations to VMware, Migrate uses floppy devices on the target VM. After cutover is completed for a migration to VMware, extra floppy devices remain attached but not connected to the cutover VM. By default, Migrate cleans up the floppy devices at cutover.

For migrations to VMware, the PlateSpin Configuration parameter `LeaveVMwareDevicesAtCutover` controls whether Migrate removes (cleans up) or leaves the floppy devices on a cutover VM after a successful Cutover or Test Cutover. The parameter settings are:

- **False:** (Default) At cutover, Migrate removes the floppy devices from the cutover VM. Migrate stops the guest OS on the cutover VM, removes its floppy devices, then restarts the VM. This extra reboot occurs after every Cutover or Test Cutover for migrations to VMware. After the reboot, the cutover VM will have no floppy devices.

  NOTE: The removal action requires that VMware Tools be installed and running on the cutover VM. The **Install VMware Tools** option is enabled by default when you configure migrations to VMware.

  If VMware Tools is not installed or is not running, Migrate cannot stop the cutover VM to remove the floppy devices. See "Known Issue for Floppy Devices on Cutover VMs" on page 557.

- **True:** At cutover, Migrate leaves the unneeded floppy devices on the cutover VM. The cutover VM will have two disconnected floppy devices. There is no extra reboot.

  You can leave the floppy devices in the attached but disconnected state, or you can manually remove them by using vSphere.

The `LeaveVMwareDevicesAtCutover` parameter setting applies globally on the Migrate Server for all migrations to VMware.

Configuring the `LeaveVMwareDevicesAtCutover` Parameter

To modify the `LeaveVMwareDevicesAtCutover` setting:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2. Locate the **LeaveVMwareDevicesAtCutover** parameter and click **Edit**.
3. Change the setting.
4. Save your setting and exit the page.

Known Issue for Floppy Devices on Cutover VMs

Issue: If `LeaveVMwareDevicesAtCutover` is set to False, but VMware Tools is not available (not installed or not running) on a cutover VM, Migrate cannot stop the VM to remove its floppy devices. The extra reboot does not occur. The cutover VM will have two disconnected floppy devices.

Fix: If this issue occurs, you can leave the devices in the attached but disconnected state, or you can manually remove them by using vSphere.

To avoid this issue, ensure that you select to Install VMware Tools when you configure migrations to VMware.

vSphere Alarm: Virtual Machine Consolidation Needed

Issue: When you migrate a workload to a VMware target, the migration completes successfully. However, the following message is displayed in the vSphere Web Client:

```
vSphere Web Client Configuration Issue: Virtual Machine Disks Consolidation is needed.
```

```
vSphere Web Client Triggered Alarm: Virtual machine Consolidation Needed status
```

This error condition is caused by the state of the VMware environment when the snapshot is removed. Some virtual disk files might remain on the disk.

Fix: In the vSphere Web Client, consolidate the snapshots. For information, see the following VMware resources:

- *Consolidate Snapshots* in the VMware vSphere 6.7 Documentation
- *How to Consolidate Snapshots in vSphere 5.x/6.x (2003638)* in the VMware Knowledgebase

VMware Tools Not Installed on Target Workload When You Use the X2P Workflow for Migrating a Workload to VMware

Issue: When you use the X2P workflow to migrate a workload to VMware, you are required to set up the VMware Tools for the target workload before you perform the conversion so that the configuration service can install the tools on the target VM when the target VM contacts the PlateSpin Server. See "Setting Up VMware Tools for the Target Workload" in the *User Guide*.

However, the VMware tools does not automatically install on the target workload if the target workload already has some version of VMware tools installed.

Fix: After conversion, manually uninstall the older version of VMware tools on the target workload and install the newer version.

Migration of an OEL or CentOS 6.*x* Source Workload on a VMware Host Fails

Issue: When you migrate an OEL or a CentOS 6.*x* source workload that is on a VMware host, the migration fails if the SCSI adapter type is not `Paravirtual` and the network adapter is not `VMXNET3`.

Fix: Edit the SCSI adapter type to `Paravirtual` and the network adapter type to `VMXNET3` and then perform the migration.

Migration of Workloads Using File-Based Transfer Method

Use information in this section to help troubleshoot common problems that might occur during migration of workloads using file-based data transfer method.

- "File-Based Transfer Conversion Fails at Cutover with Kernel Panic or GRUB Rescue Mode for Older Linux Workloads with an XFS /boot Directory" on page 559

File-Based Transfer Conversion Fails at Cutover with Kernel Panic or GRUB Rescue Mode for Older Linux Workloads with an XFS /boot Directory

Issue: In the Migrate Client, file-based transfer conversions fail at cutover for older Linux workloads that have an XFS /boot directory. The replication completes normally. However, when the target workload boots at cutover, it either has a kernel panic (UEFI workloads) or fails into a GRUB rescue console with XFS errors (BIOS workloads). This issue has been observed on RHEL/CentOS/OL 7.1 and older workloads.

Fix: You can try the migration using block-based data transfer.

Peer-to-Peer Migrations (Windows)

Table H-1 provides information to help you troubleshoot common problems that might occur during Windows peer-to-peer migrations.

Table H-1 *Common Issues and Solutions Related to Peer-to-Peer Migrations (Windows)*

Problems or Messages	Solutions
File transfer hangs at 1% or progresses at a slow pace	By default, a link type of AUTO is used on the source server during a migration. If the source server is connected to a switch port that is forced to 100/FULL, the **Force Full Duplex** option must be enabled when configuring the migration. If this option is set incorrectly, a duplex mismatch occurs on the network.
Unable to determine suitable boot partition	When converting existing source servers, the boot volume must pass the following checks: • It must be on a basic disk • It must have 175 MB of free space • It must be a primary partition • If any of these are not true for the system volume, the migration fails while attempting to take control of the source server.
Job remains in a **Scheduled** state for a long period and then changes to **Recoverable error** (all sub-steps display **NotStarted** status)	There is a problem with the Operations Framework Controller on the PlateSpin Server. Use the Windows services plug-in to confirm that the Controller is running. See KB Article 7920862 (https://support.microfocus.com/kb/doc.php?id=7920862) for other troubleshooting instructions.

Problems or Messages	Solutions
Troubleshooting failures at the **Configuring Operating System** stage (also applicable to **Configure Target Machine** or **Configuring Virtual Machine** migration steps)	Generally, failures during the configuration step indicate that a time-out occurred when attempting to configure the target physical or virtual machine. Although the migration job appears to have failed, the overall migration is probably successful and the configuration service running on the target will likely continue its operations. KB Article 7920327 (https://support.microfocus.com/kb/doc.php?id=7920327) contains a detailed troubleshooting checklist and lists information required if technical support is necessary.
Live Transfer is unavailable	Either an unsupported file system or operating system exists on the server.

Related KB Articles:

ID	Description
7920862 (https://support.microfocus.com/kb/doc.php?id=7920862)	ERRMSG: PlateSpin Migrate Job remains at a "Scheduled" or "Recoverable Error" state
7920810 (https://support.microfocus.com/kb/doc.php?id=7920810)	INFO: Restore job stalls - “The configuration service in the target machine”
2790341 (https://support.microfocus.com/kb/doc.php?id=7920341)	INFO: What ports does PlateSpin Migrate use during discovery, migration and file transfer?

Shrinking the PlateSpin Migrate Databases

When the PlateSpin Migrate databases (OFX and PortabilitySuite) reach a predetermined capacity, cleanup on those databases occurs at regular intervals. If there is a need to further regulate the size or content of those databases, Migrate provides a PlateSpin Database Cleanup utility (`PlateSpin.DBCleanup.exe`) to further clean up and shrink those databases. KB Article 7006458 (https://support.microfocus.com/kb/doc.php?id=7006458) explains the location of the tool and the options available for it, should you decide to use it for offline database operations.

Troubleshooting the Configuration Service

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

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

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

- “Understanding What Is Causing the Problem” on page 561
- “What Can Be Done to Resolve the Problem” on page 561
- “Additional Troubleshooting Tips” on page 564

Understanding What Is Causing the Problem

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

- "Target VM Fails to Boot" on page 561
- "Network Is Not Set Up Correctly" on page 561
- "Unable to Read or Write Status Messages to Floppy Devices" on page 561

Target VM Fails to Boot

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

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

Network Is Not Set Up Correctly

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

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

Unable to Read or Write Status Messages to Floppy Devices

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

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

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

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

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

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

What Can Be Done to Resolve the Problem

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

- "Skip the Target VM Reboot Optimizations" on page 562
- "Reduce the Read/Write Traffic to Floppy Devices" on page 562

- "Change the Startup Type to Increase the Delay" on page 563
- "Configure Conflicting Services to Not Run Automatically at Startup" on page 564

Skip the Target VM Reboot Optimizations

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

To skip reboot optimizations:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`
2. Search for the parameter **ConfigurationServiceValues**.
3. Edit the **ConfigurationServiceValues** parameter and set the **SkipRebootOptimization** option to `true`.
4. Click **Save**.
5. Run an incremental or full replication.

   The replication also propagates the modified configuration settings to the target VM.
6. Run the Test Cutover or Cutover again for affected workloads.

Reduce the Read/Write Traffic to Floppy Devices

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

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

followed by

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

-or-

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

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

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

vmwareConfigServicePollStartDelay

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

vmwareConfigServicePollIntervalInMilliseconds

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

vmwareConfigServicePollStartTimeout

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

vmwareConfigServicePollUpdateTimeout

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

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

To reduce read and write traffic for VMware floppy devices:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2. Search for the Configuration Service polling parameters, modify their settings as appropriate, then click **Save**.

   For example:

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

   or

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

3. Run an incremental or full replication.

   The replication also propagates the modified configuration settings to the target VM.

4. Run the Test Cutover or Cutover again for affected workloads.

Change the Startup Type to Increase the Delay

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

To change the startup type:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`

2 Search for the parameter **windowsConfigServiceStartType**.

3 Change the **windowsConfigServiceStartType** value to **AutoDelay**.

Options for **windowsConfigServiceStartType** are:

- **GroupDelay** is the default value and adds the Configuration Service to the end of the **ServiceGroupOrder** in the registry.
- **AutoDelay** will maximize the amount of time the service waits before starting (2 minutes after boot). Also modify the **ServicesPipeTimeoutForWindowsConfigService** parameter value in Step 4.
- **NoDelay** is the most efficient option and starts the service as soon as Windows can. However, it is not recommended because of the potential issues connecting to resources.

4 (AutoDelay) Change the **ServicesPipeTimeoutForWindowsConfigService** parameter setting to 180 seconds to account for the 120 seconds that the service will take to start up after boot when AutoDelay is set for **windowsConfigServiceStartType** in Step 3.

5 Click **Save**.

6 Run an incremental or full replication.

The replication also propagates the modified configuration settings to the target VM.

7 Run the Test Cutover or Cutover again for affected workloads.

Configure Conflicting Services to Not Run Automatically at Startup

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

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

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

Additional Troubleshooting Tips

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

- **Windows workloads:** The Configuration Service logs are found in the `C:\windows\platespin\configuration\data` folder.
  - The `log.txt` file contains all of the logging information, but the `Config.ini` file is useful in understanding what is to be configured.
  - The `result.txt` file contains the status of the Configuration Service run.
  - If the target VM cannot read from the input floppy device, it will not have the merged `Config.ini` file, which might include custom network configuration information for the test Cutover network environment.
  - If the `Config.ini` file has no network related information (such as a `[NIC0]`, the target VM network adapter might have special characters in the name.

It is a known issue that the `Config.ini` file might not be accurate until it is merged with the one from the floppy device.

- The target VM tries a reboot if it cannot connect to either the output floppy or input floppy (one time only). You will see a `config.ini.floppyreboot` file if this is the case.

- **Linux workloads:** The Configuration Service logs are found in the `/tmp` folder.
  - The main log files are named `file*.platespin.fileLogger`.

    We recommend examining any configuration folders in `/tmp`. Tar the configuration folders along with the `file*.platespin.fileLogger` files to send to Micro Focus Customer Care.
  - Other config files to check for include the following:

    ```
    /tmp/Ofx.RunCommand.Output*
    /tmp/*DiskHelper*
    /tmp/*VmTools*
    ```
  - The configuration file is `/usr/lib/psconfigservice/data/config.conf`.
  - The end result log file is `/usr/lib/psconfigservice/data/result.txt`.

PlateSpin OFX Controller Does Not Start on a Virtual Machine Source

Issue: If you configure Migrate to install the block-based component during the first replication, PlateSpin OFX Controller might not start on the source workload during the Install Block-Based Components step. The Service Manager reports this problem if the VM is running so slowly that the OFX Controller startup event times out.

Fix: Manually start PlateSpin OFX Controller on the source workload.

To avoid the problem, for workloads with low memory and CPU resources, do either of the following to improve startup performance:

- Configure the workload to install the block-based component during Prepare Workload instead of First Replication.
- Increase the Memory and CPU resources of the source VM.

Validation Warning for Bandwidth Throttling

Issue: After you configure migration for a workload with no warnings or validation errors, you might get a warning message if you then set or modify the value for **Bandwidth Throttling**, even if the setting is valid.

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

Target Windows Machine Becomes Unbootable on Second Boot

Issue: The target Windows machine becomes unbootable during the second boot.

When PlateSpin Migrate executes the Configuration Service on a target Windows machine, the normal networking tasks performed during the second boot can be problematic in the following scenarios:

- If the target machine has the same network adapter hardware and networking drivers as the source machine.

  The network drivers that the target machine requires are the same as those already installed on the source machine being migrated. It is not necessary to re-install drivers. In some scenarios, removing and re-installing drivers can result in the target machine becoming unbootable.

- If the target machine is booting from SAN.

  If a target machine boots from SAN, Migrate installs drivers before the first boot. If the Configuration Service removes these newly installed drivers during the second reboot, the target machine becomes unbootable. It is necessary to avoid the driver install tasks on the second reboot.

Fix: PlateSpin Migrate provides two light networking configuration settings for the PlateSpin Server that optimizes the network configuration process on the target machine during the second boot and helps avoid situations that can cause a target machine to become unbootable. Light networking is useful for P2P, V2V, and C2C migrations as well as for X2V semi-automated migrations where the networking hardware on the target VM is manually configured to match the source machine. See "Configuring Behavior for Installing Network Drivers on Target Windows Workloads" on page 114.

Two or More Volumes Have the Same Volume Serial Number

Issue: When you attempt to set up a migration job for a Windows server, the following error is displayed:

```
[Source] Two or more volumes have the same serial number. Change the serial numbers so that they are unique and rediscover the machine.
```

Fix: This problem can occur if the Volume Serial Numbers for two or more volumes are the same. PlateSpin Migrate requires the serial numbers to be unique.

To resolve this issue, modify the serial numbers for the data volumes as appropriate, and then rediscover the machine. For information about how to use Windows native tools to modify the serial numbers, see KB Article 7921101.

Replication Cannot Complete If an Anti-Virus Update Is Pending a Restart on the Source

Issue: Automatic updates for anti-virus software on Windows source workloads sometimes have pending system changes that require a restart. While the required restart is pending, any replication seems to get stuck and cannot complete.

Fix: To prevent this potential replication conflict, ensure that you restart the source Windows workload after an anti-virus automatic update occurs that requires a restart. Perform the restart before the next replication begins.

To gracefully resolve this conflict for an in-progress replication:

1. Abort the replication by using the Migrate Client or Migrate Web Interface, as appropriate.
2. Reboot the source Windows workload.
3. In Migrate Client or Migrate Web Interface, initiate the replication again.

   The replication should complete successfully.

Disk Not Properly Aligned on the Target VM

Issue: One or more disks in the target workload's primary partition is misaligned with the backend storage resulting in increased I/O operations per second.

Fix: The PlateSpin Configuration parameter `PartitionAlignmentSizeInKB` controls the alignment of a target workload's primary partition that is not cylinder aligned at the beginning of a disk and rounds the offset to the closest alignment boundary. The value of this parameter is the number of kilobytes (KB) from the beginning of the disk to the closest alignment boundary. This is applicable only for workloads with MBR partitions.

To specify the disk alignment value:

1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:

   `https://Your_PlateSpin_Server/PlateSpinConfiguration/`
2. Locate the **PartitionAlignmentSizeInKB** parameter and click **Edit**.
3. Edit the value based on the following allowed values. If you specify a value other than the allowed value, then the default value is applicable.
   - **For a Windows workload:** The default value is **1024** and you can set one of the following allowed values: **1024**, **2048**, **4096**.
   - **For a Linux workload:** The default value is **64** and you can set one of the following allowed values: **64**,**128**,**256**, **512**,**1024**, **2048**.
4. Save your settings and exit the page.

Cutover Fails If `root-PS-snapshot` on the Source Linux Workload Is Not Cleaned Up Properly

Issue: A 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 after a successful Abort of the first full replication of after numerous incremental replications of the source workload.

Fix: Manually delete `root-PS-snapshot` symbolic link on the source Linux workload, then repeat the cutover. See "LVM snapshots" in Table 35-2, "Use Cases and Instructions for Cleaning Up Linux Workloads," on page 550.

Source Passive Node Does Not Shut Down at Cutover for Windows Server 2016 Cluster

Issue: When **Shut Down** is set as the post-migration end state for a Windows Server 2016 Cluster, the PlateSpin Migrate Web Interface shuts down only the active node of the cluster; the passive nodes are not shut down. Migrate Client properly shuts down all source nodes.

Fix: Manually shut down the source passive nodes if they do not automatically shut down when **Shut Down** is selected for the post-migration end state of a Windows Server 2016 Cluster.

RHEL 6.10 Source Workloads Crash During Incremental Replication

Issue: During incremental replications, a system crash occurs for source Red Hat Enterprise Linux (RHEL) 6.10 workloads running Linux kernel 2.6.32-754.6.3.el6.x86_64.

Fix: The system crash is a Red Hat known issue for RHEL 6.10 running Linux kernel version 2.6.32-754.6.3.el6.x86_64. See the following Red Hat Knowledgebase Articles:

- *System crashes at "block/blk-throttle.c:1222" after running veeam agent backup job* (KB 3658111) (https://access.redhat.com/solutions/3658111)
- *Kernel panic due to Hard LOCKUP in the function blk_throtl_drain()* (KB 3676431) (https://access.redhat.com/solutions/3676431)

To avoid this issue, do either of the following:

- Follow the Red Hat Bug Advisory RHBA-2018:3763 (https://access.redhat.com/errata/RHBA-2018:3763), which will apply Linux kernel 2.6.32-754.9.1.el6.x86_64 to the source workload.

  -OR-

- Use Linux kernel version 2.6.32-754.3.5.el6.x86_64 or earlier on the source workload.

After the crash occurs, the blkwatch driver might cause a failure in the normal boot process on the source workload. Remove the blkwatch driver until you have resolved the Red Hat issue. See *How to Disable the blkwatch Kernel Module on Source Linux Workloads (KB 7006279)* (https://support.microfocus.com/kb/doc.php?id=7006279).

Replication Fails in a FIPS-Enabled Environment for vCloud and AWS Targets

Issue: In a FIPS-enabled environment, FIPS is enabled for the PlateSpin Migrate Server and source workloads. Replication of a FIPS-enabled workload to a target vCloud or target AWS platform fails with the following error:

```
Exception has been thrown by the target of an invocation.
This implementation is not part of the Windows Platform FIPS validated
cryptographic algorithms.
```

This error is not observed for migration of FIPS-enabled source workloads to other target platforms.

Fix: You can enable PlateSpin Migrate to suppress errors for non-compliant FIPS algorithms.

1. In a text editor, open the `ofxcontrollerexecution.exe.config` file found in this folder:

   `<install folder>\PlateSpin Migrate Server\Controller\Packages\0\C863075B-8130-4d29-893B-70FF2AD9308C\1`

2. Add the following element to the `runtime` section of the file:

```
<configuration>
    <runtime>
         ...
        <enforceFIPSPolicy enabled="false"/>
    </runtime>
</configuration>
```

3. Save your changes.

Cannot Adjust the Target Volume Size for File-Based Migrations in the Web Interface

Issue: In the Migrate Web Interface, you could previously modify the **Volumes > Free Space** parameter to adjust the target volume size for Windows file-based migrations. However, the new Disk Layout view for Volumes currently does not allow you to adjust the volume size. You cannot adjust the target volume size for file-based migrations in the Web Interface when the Disk Layout view is enabled.

Fix: Disable the Disk Layout view and use the legacy storage view to configure target disks and volumes. You can modify the **Free Space** value for the target volume to adjust its volume size.

To disable the Disk Layout view and use the legacy storage view instead:

1. Go to the PlateSpin Configuration page:

   https://*Your_PlateSpin_Server*/PlateSpinConfiguration

2. Locate the **EnableDiskLayoutOptions** parameter, then click **Edit**.
3. Set the value to `False`.
4. Save your settings and exit the page.

Replication Fails for Oracle 5.11 UEK Workload Migrations to AWS

Issue: Replication fails for an Oracle 5.11 UEK with kernel version 2.6.39-400.215.10.el5uek for migrations to AWS. Diagnostics show error `mkinitrd command failed`.

Replication of this workload completes successfully for migrations to VMware and vCloud. (Oracle 5.*x* is not supported for migrations to Azure.) Replication of Oracle 5.11 RHCK with kernel version 2.6.18-398.el5 completes successfully for migrations to AWS. (`Bug 1136126`)

Fix: None. You cannot migrate a workload running Oracle 5.11 UEK with kernel version 2.6.39-400.215.10.el5uek to AWS.

VII Additional PlateSpin Tools

PlateSpin Migrate provides additional tools to support your migration efforts.

Using the PlateSpin Migrate Client Command Line Interface

The PlateSpin Migrate Client installation includes a command line interface (CLI) tool to help you perform common migrations tasks. Conversion jobs using `.ini` files is supported to VMware targets only. Using this tool, you can

- Discover and subsequently refresh a host or target server to populate the Migrate Server with server information.
- Migrate (also known as "convert") heterogeneous workloads across x86 server and desktop infrastructure in the data center.
- Prepare the target host for its new workload and then, after a conversion, synchronize the host and the target.
- Check the status of a job as it is running, and if necessary, abort it.

This section includes information that can help you use the CLI effectively.

- "Where Is the Tool Located?" on page 573
- "Before You Use the Tool" on page 573
- "Configurable .ini Files (Jobs) You Can Use with the Tool" on page 576

Where Is the Tool Located?

The CLI tool, `PlateSpin.Migrate.Console.exe`, is installed with the PlateSpin Migrate Client at the following location:

- **32-bit host:** `C:\Program Files\PlateSpin Migrate Client\CommandLine\PlateSpin.Migrate.Console.exe`
- **64-bit host:** `C:\Program Files(x86)\PlateSpin Migrate Client\CommandLine\PlateSpin.Migrate.Console.exe`

Before You Use the Tool

This section includes the following information:

- "Pre-configuring the Migrate Server Values for CLI" on page 574
- "Becoming Familiar with the Commands" on page 574

Pre-configuring the Migrate Server Values for CLI

Before you begin using the command line utility, you need to ensure that the Migrate Server is properly configured. You can check the configuration in the `PlateSpin.Migrate.Console.exe.config` file, located in the same path as the command line utility. After the Migrate installation, the following `config` file should already be populated with values.

```
<?xml version="1.0" encoding="utf-8 ?>"
<configuration>
  <appSettings>
    <add key="MigrateServerURL" value="https://localhost/PlateSpinMigrate/
"/>
    <add key="ServerDomain" value=""/>
    <add key="psuser" value="administrator"/>
    <add key="pspassword" value="encoded_password"/>
    <add key ="encoded" value="yes"/>
  </appSettings>
</configuration>
```

The tool uses these values as it executes commands. You need to reconcile the values in the file with the settings for the Migrate Server with which you want to connect.

The value for the `pspassword` key is blank by default and you must specify an encoded password as the value. To encode the password, use the `encode` command. For more information about commands, see "Becoming Familiar with the Commands" on page 574.

If you choose to provide encoded passwords for source workload and target platform, set the value for the encoded key in the following line of the `PlateSpin.Migrate.Console.exe.config` file to `yes`, otherwise set value to `no`.

```
<add key ="encoded" value="no"/>
```

Becoming Familiar with the Commands

You can display the commands supported in the tool by running it with the `Help` option or with the `?` option from the command prompt, like this:

```
C:\Program Files\PlateSpin Migrate
Client\CommandLine>PlateSpin.Migrate.Console.exe Help
```

The tool displays a matrix that includes information similar to what is included in the following table:

Table I-1 *Commands available from the Migrate CLI tool*

Command	Description
`run`	Runs a configured `.ini` file as a scheduled job. When the you add the `/wait=no` parameter and the job starts to run, its Job ID is displayed in the interface.
`query`	Runs a query on the job (when you specify a Job ID) to display its current status.
`discover`	Runs an operation that inventories the details of a supported workload or target computer in preparation for a migration or "conversion" job.

Command	Description
`refresh`	Refreshes a discovered server.
`unDiscover`	Undiscovers a server.
`abort`	Aborts a scheduled job.
`licenseInfo`	Displays the license information of the migrate server.
`serversync`	Prepares the server for the Server Sync operation and then runs a serversync job using the configuration file.
`encode`	Encodes the text input or the data in the text file.
`massdiscover`	Performs mass discovery of source workloads and targets. The discovered workloads and targets are displayed both in the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface To mass discover workloads and targets, you must first list the workloads and targets that you want to discover in a CSV file. To create this CSV file, refer to the sample CSV file located at `\PlateSpin Migrate Client\CommandLine\Sample INI\MassDiscovery.csv`.

When you run any of these commands, you must include its required parameter(s) in the command line. You can also include some optional parameters, when needed. For example, `savejob=` parameter saves the job in default location.

To display a list of these parameters at the command prompt, run the command without any parameter. For example, if you run the `discover`command without parameters, like this:

```
C:\Program Files\PlateSpin Migrate
Client\CommandLine>PlateSpin.Migrate.Console.exe discover
```

the command line interface displays these following:

```
[discover]                      discovers a server
Required Parameters:
    /machineAddress=      machine address to discover
    /userName=            the username to use
    /password=            the password to use
    /type=                type like windows,
linux,vmware_esx,vmware_vcenter,
    Optional Parameters:
    /network=             network name to connect to
    /address=             server address to connect to
    /psuser=              Username used for accessing PlateSpin Migrate
server as user different from the one logged on this computer
    /pspassword=          Password used for accessing Platespin Migrate
server for the user different from the one logged on this computer
    /wait=                wait for completion of job [yes,no]
    /clusterName=         clustername to be discovered
    /verbose=             verbose mode for output [on,off]
    /output=              the output file
    /format=              the ouptput format to display in [text,html,xml]
    /sslcertificatewarnings=      Whether to Ignore or Enforce SSL
Certificate Warnings [Ignore| Enforce]
```

NOTE: You should become familiar with the different CLI commands and their respective required and optional parameters.

Command Line Syntax

If you were to run the discover command (which is also a job), you would use a syntax similar to this example, at the command prompt:

```
C:\Program Files\PlateSpin Migrate
Client\CommandLine>PlateSpin.Migrate.Console.exe discover /
machineaddress=10.10.8.100 /username=administrator /password=password /
type=windows /wait=no
```

Note that all required parameters and one optional parameter are included in this example.

When the discover command (job) starts, the CLI tool displays its job ID, similar to this example:

```
8be8d306-7665-4869-9795-a9dbb3ce1471
```

You can leverage this ID to learn the status of the job, just by using the `query` command, like this:

```
C:\Program Files\PlateSpin Migrate
Client\CommandLine>PlateSpin.Migrate.Console.exe query /id=8be8d306-7665-
4869-9795-a9dbb3ce1471
```

The query command yields a status report that includes all of the details of the job. This is the same kind of information you might see from the Migrate Client Jobs view.

Configurable .ini Files (Jobs) You Can Use with the Tool

When you install the PlateSpin Migrate Client, the installation creates a separate directory for a number of preconfigured jobs (actually, `.ini` files) that can do the following:

- Workload conversion (that is, a migration operation)
- Server Sync

You execute a job by using the `run` command at the command line. The values in the files are the optional parameters that run along with the job. Each of these functions has a “default” `.ini` file version that runs with basic settings, and one or more “platform-specific” `.ini` file(s) that run with custom settings:

- `Conversion-Default.ini`
- `Conversion-Windows.ini` (customized)
- `Conversion-Linux.ini` (customized)
- `ServerSync-Default.ini`
- `ServerSync-Windows.ini` (customized)
- `ServerSync-Linux.ini` (customized)
- `IncrementalImaging-Default.ini`
- `IncrementalImaging.ini` (customized)

This section includes more details about these jobs.

- “Conversion Jobs” on page 577
- “ServerSync Jobs” on page 583

Conversion Jobs

The CLI tool supports converting Windows and Linux workloads (source) to vCenter or ESX servers (target). There are two types of `.ini` files, one for a basic job configuration, and one for custom configurations. While the job is running you can abort the job or check its status.

Before you start a conversion job, ensure that you run the `discover` command on the source computer and then on the target platform. The following is example syntax for running the discover command:

```
discover /machineaddress=10.10.10.10 /username=administrator /
password=anything@123 /type=vmware_vcenter
```

The tables in this section are named by the respective conversion jobs `.ini`files they represent. The table contents include the file section names within the `.ini` and the available settings you can configure according to your conversion needs:

Conversion-Default.ini

Table I-2 *Details of `Conversion-Default.ini`*

File Sections and Default Settings	Comment
`[Type]`	
`Conversion=X2V`	{required} This value must be used for every conversion.
`[JobConfig]`	
`Default=true`	
`[Source]`	
`Address=`	{required} Specify an IP address for the source workload.
`UserName=`	{required} Specify a username credential for the source workload.
`Password=`	{required} Specify a password credential for the source workload.
`TakeControlAddress=`	
`SubnetMask=`	
`DefaultGateway=`	
`DNS=`	
`[TargetContainer]`	

File Sections and Default Settings	Comment
`Address=`	{required} Specify the IP address for the target platform depending on how it is discovered. For example, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery.
`UserName=`	{required} Specify the username for the target platform depending on how it is discovered. For example: ◆ If ESX is discovered via VCenter, specify the vCenter username. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root username.
`Password=`	{required} Specify the password for the target platform depending on how it is discovered. For example: ◆ If ESX is discovered via VCenter, specify the vCenter password. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root password.
`[NewMachine]`	
`DisplayName=`	{required} Specify the name you want to display in the target platform console.
`HostName=`	{required} Host name of the target machine.

Conversion-Windows.ini

You can skip system volume.

Table I-3 *Details of* `Conversion-Windows.ini`

File Sections and Default Settings	Comment
`[Type]`	
`Conversion=X2V`	{required} This value must be used for every conversion.
`[JobConfig]`	
`Default=false`	
`[Transfer]`	
`TransferType=VSSFileBased/ VSSblockBased/FileBased`	Possible settings shown. If the Windows source machine support VSS snapshotting, use the VSS setting, if it does not support VSS, use the `Filebased` setting.

File Sections and Default Settings	Comment
`LiveTransferEnabled=true/false`	Possible settings shown. This setting is dependent on the `TransferType` setting. Offlline migrations (false) are not supported.
`[Source]`	
`Address=`	{required} Specify an IP address for the source workload.
`UserName=`	{required} Specify a username credential for the source workload.
`Password=`	{required} Specify a password credential for the source workload.
`EndState=ShutDown/Donothing/Reboot`	Possible settings shown.
`TakeControlAddress=`	
`SubnetMask=`	
`DefaultGateway=`	
`DNS=`	
`[TargetContainer]`	
`Address=`	{required} Specify the IP address for the target platform depending on how it is discovered. For example, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery.
`UserName=`	{required} Specify the username for the target platform depending on how it is discovered. For example: ◆ If ESX is discovered via VCenter, specify the vCenter username. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root username.
`Password=`	{required} Specify the password for the target platform depending on how it is discovered. For example: ◆ If ESX is discovered via VCenter, specify the vCenter password. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root password.
`VirtualNetwork=`	Specify the target platform virtual network name you want to use.
`TakeControl=static/dhcp`	Specify `static` or `dhcp`, depending on your networking configuration.

File Sections and Default Settings	Comment
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
DNS=	
[NewMachine]	
DisplayName=	{required} Specify the name you want to display in the target platform console.
DataStore=	Specify the name of datastore (datastore1) you want to use for configuration files.
ConfigPath=	Specify the complete path where you want to create the .vmx file. For example: */folder_name/vmx_file_name* The .vmxfile is created in the specified folder within the datasource.
Memory=	Specify the amount of RAM you want for the target computer. The setting can be in MB or GB and must be specified with integers (no decimal values).
InstallTools=true/false	Possible settings shown. Default is true.
NumberofCPU=	Specify the number of CPUs you want for the target computer.
HostName=	{required} Specify the target host name.
WorkGroup=	{optional} Specify the workgroup name you want to join.
Domain=	
DomainUserName=	
DomainUserPassword=	
EndState=VMPowerOFF/VMPowerON/ VMSuspend	Possible settings shown.
ScsiType=	(On VMware) Specify the SCSI Adapter type. If you do not specify a type or specify an unsupported adapter type, the default adapter type is used.
ResourcePool=	(On VMware) Specify the ResourcePool name in the vCenter. If the resource pool is nested, then use \ to separate names. For example, windows\local.
UseThinDisks=	To use thin disks, specify **true**. Else, specify **false**.

File Sections and Default Settings	Comment
[EthernetNic1]	You can repeat this section of the .ini file for every NIC at the target platform. For example, the second NIC section would be named **[EthernetNic2].** Configuration settings would be specified for each NIC section in the file.
DHCPEnabled=true/false	Specify **true** for DHCP and **false** for static IP.
VirtualNetwork=	Specify the target platform virtual network name you want to use.
Address=	Specify the IP address for the target machine.
SubnetMask=	
DefaultGateway=	
DNS=	Specify one or more DNS names separated by commas.
[DriveGeneral]	If you have multiple disks at the source, you can specify them here. You can specify as many disks as there are at the source.
DataStore1=	Specify the datastore (datastore1)on the target platform. For example:
Disk1=	Specify the path to the configuration file on the target platform. For example: /win2k8r2/ win2k8r2.vmdk
DataStore2=	
Disk2=	
[Volume1]	You can repeat this section of the .ini file for every volume at the target platform. For example, the second volume section would be named **[Volume2].** Configuration settings would be specified for each volume section in the file.
VolumeToCopy=	Specify the volume to copy to the target.
MapTo=	Specify the disk to map.
FreeSpace=	Specify the amount of free space, in MB or GB, available on the target for File-Based conversion.

Conversion-Linux.ini

The sections in the Conversion-Windows.ini and in the Conversion-Linux.inifile are identical, except for the settings in **[Transfer]** section, along with the settings for workgroup and domain configuration. The differences for the Linux source job are shown in the following table.

Table I-4 *Conversion-Linux.ini: Differences in Setting Details of the [Transfer] section*

File Sections and Default Settings (differences only)	Comment
[Transfer]	
TransferType=BlockBased/FileBased	Possible settings shown. Linux does not support VSS.
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the TransferType setting. Offline migration (false) are not supported.
[Source]	
[NewMachine]	
ScsiType=	(On VMware) Specify the Scsi Adapter type. If you do not specify a type or specify an unsupported adapter type, the default adapter type is used.
ResourcePool=	(On VMware) Specify the ResourcePool name in the vCenter. If the resource pool is nested, then use \ to separate names. For example, windows\local.
UseThinDisks=	To use thin disks, specify **true**. Else, specify **false**.
[EthernetNic1]	
DNS=	Specify one or more DNS names separated by commas.
[LVMGroup]	
Group1= Add enteries depending on the number of groups you want. If you have two groups, then add the following: Group1= Group2=	Name of the LVM group in the source.
[Volume1]	
FreeSpace=	Specify the amount of free space, in MB or GB, available on the target for File-Based conversion.

ServerSync Jobs

Use `serversync` command to perform the Server Sync operation. There are two types of `.ini` files, one for a basic job configuration, and one for custom configurations. While the job is running you can abort the job or check its status. If you specify the required settings, it will start the job. Then, when it runs, the job populates the other values with default settings.

The tables in this section are named by the respective Server Sync jobs `.ini`files they represent. The table contents include the file section names within the `.ini` and the available settings you can configure according to your conversion needs:

ServerSync-Default.ini

Table I-5 *Details of `ServerSync-Default.ini`*

File Sections and Default Settings	Comment
`[Type]`	
`Conversion=Sync2V`	{required} This value must be used for every Server Sync operation.
`[JobConfig]`	
`Default=true`	
`[Source]`	
`Address=`	{required} Specify an IP address for the source workload.
`UserName=`	{required} Specify a username credential for the source workload.
`Password=`	{required} Specify a password credential for the source workload.
`TakeControlAddress=`	
`SubnetMask=`	
`DefaultGateway=`	
`DNS=`	
`[TargetContainer]`	
`Address=`	{required} Specify the IP address for the target platform depending on how it is discovered. For example, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery.

File Sections and Default Settings	Comment
`UserName=`	{required} Specify the username for the target platform depending on how it is discovered. For example: ◆ If ESX is discovered via VCenter, specify the vCenter username. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root username.
`Password=`	{required} Specify the password for the target platform depending on how it is discovered. For example: ◆ If ESX is discovered via VCenter, specify the vCenter password. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root password.
`[ExistingTargetMachine]`	
`DisplayName=`	{required} Specify the display name of the target machine where you want to sync.
`HostName=`	{required}

ServerSync-Windows.ini

For prepare for Sync, the ServerSync command uses target platform and network details from TargetContainer and machine name from ExistingTargetMachine file sections.

Table I-6 *Details of `ServerSync-Windows.ini`*

File Sections and Default Settings	Comment
`[Type]`	
`Conversion=Sync2V`	{required} This value must be used for every Server Sync operation.
`[JobConfig]`	
`Default=false`	
`[Transfer]`	
`TransferType=VSSFileBased/ VSSblockBased/FileBased`	Possible settings shown. If the Windows source machine support VSS snapshotting, use the VSS settings, if it does not support VSS, use the `Filebased` setting.
`LiveTransferEnabled=true/false`	Possible settings shown. This setting is dependent on the `TransferType` setting. Offline migration (false) is not supported.

File Sections and Default Settings	Comment
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
EndState=ShutDown/Donothing/Reboot	Possible settings shown.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
DNS=	
[TargetContainer]	
Address=	{required} Specify the IP address for the target platform depending on how it is discovered. For example, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery.
UserName=	{required} Specify the username for the target platform depending on how it is discovered. For example: ◆ If ESX is discovered via VCenter, specify the vCenter username. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root username.
Password=	{required} Specify the password for the target platform depending on how it is discovered. For example: ◆ If ESX is discovered via VCenter, specify the vCenter password. ◆ If ESX is discovered via Direct ESX discovery, specify the ESX root password.
VirtualNetwork=	Specify the target platform virtual network name you want to use.
TakeControl=static/dhcp	Specify `static` or `dhcp` depending on your networking configuration.
TakeControlAddress=	
SubnetMask=	

File Sections and Default Settings	Comment
DefaultGateway=	
DNS=	
[ExistingTargetMachine]	
DisplayName=	{required} Specify the display name of the target machine where you want to sync.
HostName=	.
InstallTools=true/false	.
WorkGroup=	Specify the workgroup name if you want to join workgroup.
Domain=	.
DomainUserName=	.
DomainUserPassword=	.
EndState=VMPowerOFF/VMPowerON/ VMSuspend	Possible settings shown.
[EthernetNic1]	You can repeat this section of the .ini file for every NIC at the target platform. For example, the second NIC section would be named **[EthernetNic2].** Configuration settings would be specified for each NIC section in the file.
DHCPEnabled=true/false	Specify **true** for DHCP and **false** for static IP.
VirtualNetwork=	Specify the target platform virtual network name you want to use.
Address=	Specify the IP address for the target machine.
SubnetMask=	
DefaultGateway=	
DNS=	

ServerSync-Linux.ini

The sections in the `ServerSync-Windows.ini` and in the `ServerSync-Linux.ini`file are identical, except for the settings in **[Transfer]** section, along with the settings for the workgroup and domain configuration. For prepare for Sync, the ServerSync command uses target platform and network details from TargetContainer and machine name from ExistingTargetMachine file sections.

The differences for the Linux source job are shown in the following table.

Table I-7 *ServerSync-Linux.ini: Differences in Setting Details of the [Transfer] section*

File Sections and Default Settings (differences only)	Comment
[Transfer]	
TransferType=BlockBased/FileBased	Possible settings shown. Linux does not support VSS.
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the TransferType setting. Offline migration (false) is not supported.

J Using the iPerf Network Test Tool to Optimize Network Throughput for PlateSpin Products

Before you execute replication, ensure that you test the connection to see if there are any connection or bandwidth issues, and resolve them. This section describes how to use the open source iPerf Network Test tool to optimize throughput on the connection.

Introduction

In the interest of helping PlateSpin administrators in their efforts to achieve better network throughput when using PlateSpin products, the iPerf Network Test tool is provided on the PlateSpin LRD (Linux RAM Disk) take-control environment. As stated in the iPerf documentation: "The primary goal of iPerf is to help in tuning TCP connections over a particular path. The most fundamental tuning issue for TCP is the TCP window size, which controls how much data can be in the network at any one point."

The purpose of this README is to describe a basic method for network tuning and testing as it relates to using PlateSpin products. First, you calculate a theoretical optimum TCP window size. Then you use the iPerf tool to validate and fine-tune this calculated size and measure the resulting throughput. Using this method is also useful in determining the real achievable throughput for a given network.

Both the iPerf tool and PlateSpin products are actually using the *TCP send/receive buffer size* in order to affect the eventual internal choice of *TCP window size*. Going forward, these terms will be used interchangeably.

NOTE: There are many factors that affect network throughput. A wealth of information is available on the Internet that can aid in understanding. One such resource is the *Network Throughput Calculator* (http://wintelguy.com/wanperf.pl), which can help in calculating the expected maximum TCP throughput given applicable customer network characteristics. We strongly recommend that this online calculator be used in order to correctly set expectations regarding throughput.

Calculations

Tuning of the TCP window size is based on a number of factors, including network link speed and network latency. For our purposes relating to PlateSpin products, the initial choice of TCP window size for tuning is based on standard calculations (widely available on the Internet and elsewhere) as follows:

```
WinSizeInBytes=((LINK_SPEED(Mbps)/8)*DELAY(sec))*1000*1024
```

For example, for a 54 Mbps link with 150 ms latency, the proper initial window size would be:

```
(54/8)*0.15*1000*1024 = 1,036,800 bytes
```

For a 1000 Mbps link with 10 ms latency, the proper initial window size would be:

```
(1000/8)*.01*1000*1024 = 1,280,000 bytes
```

In order to get a latency value for the network, use `ping` from the command prompt (Windows) or the terminal (Linux). Although the `ping` round-trip time (RTT) is arguably different than the actual latency, the value obtained is sufficiently close for use in this method.

The following is a sample output from a Windows `ping` command, where the latency is observed to be 164 ms on average:

`ping 10.10.10.232 -n 5`

```
Pinging 10.10.10.232 with 32 bytes of data:
Reply from 10.10.10.232: bytes=32 time=154ms TTL=61
Reply from 10.10.10.232: bytes=32 time=157ms TTL=61
Reply from 10.10.10.232: bytes=32 time=204ms TTL=61
Reply from 10.10.10.232: bytes=32 time=153ms TTL=61
Reply from 10.10.10.232: bytes=32 time=153ms TTL=61

Ping statistics for 10.10.10.232:
    Packets: Sent = 5, Received = 5, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 153ms, Maximum = 204ms, Average = 164ms
```

The following is a sample output from a Linux `ping` command, where the latency is observed to be 319 ms on average:

`ping 10.10.10.232 -c 5`

```
PING 10.10.10.232 (10.10.10.232) 56(84) bytes of data.
64 bytes from 10.10.10.232: icmp_seq=1 ttl=62 time=0.328 ms
64 bytes from 10.10.10.232: icmp_seq=2 ttl=62 time=0.280 ms
64 bytes from 10.10.10.232: icmp_seq=3 ttl=62 time=0.322 ms
64 bytes from 10.10.10.232: icmp_seq=4 ttl=62 time=0.349 ms
64 bytes from 10.10.10.232: icmp_seq=5 ttl=62 time=0.316 ms

--- 10.10.10.232 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 3998ms
rtt min/avg/max/mdev = 0.280/0.319/0.349/0.022 ms
```

In practice, you should use the `-n` or `-c` option to specify a larger number of ping packets in order to more closely measure the latency value.

Setup

The iPerf tool runs in either server mode or client mode.

The basic usage syntax for `iperf` server mode is:

```
iperf -s -w <win_size>
```

The basic usage syntax for `iperf` client mode is:

```
iperf -c <server_ip> -w <win_size>
```

You can specify units such as K (kilobytes) or M (megabytes) in the value for the `-w` option. For example: `1.3M` or `1300K`.

NOTE: Linux automatically doubles the requested TCP buffer size. If you use `iperf` on a Linux server, the `win_size` value for the `-w` option should be only 1/2 of the desired value (that is, `<win_size>/2`). Otherwise, you will inadvertently test with a buffer size that is twice the desired value.

Our intent is to measure and tune the network between a source and target workload. In many cases, these can be the actual source and targets in use. It is possible to complete the testing using a different workload for either source or target, provided that the substitute has the same network characteristics as the original, such as NIC, network connection, and so on.

NOTE: Ensure that you are not testing the throughput from the PlateSpin server to either the source or the target, as this traffic is minimal, and does not represent the traffic that occurs during a migration or replication.

While it is possible to use a live workload (either Windows or Linux) as the target/`iperf` server, the following steps provide the environment most similar to what happens at migration/replication time, and is strongly recommended.

To set up and run `iperf` on the target:

1. Boot the target workload using the LRD.
2. In the LRD console, use the helper terminal (accessible via Alt-F2) to do the following:
   - **2a** Set up networking using option 5.
   - **2b** Mount the CD media using option 6.
3. In the LRD console, switch to the debug terminal (accessible via Alt-F7) to go to the location of the iPerf tool:

   ```
   cd /mnt/cdrom/LRDTools/iperf_2.0.X/linux
   ```

4. Run the iPerf tool in server mode. Enter

   ```
   ./iperf -s -w <win_size>
   ```

   NOTE: For a Linux target, remember that the TCP window size you specify for the `-w` option should be half of the desired value.

To set up and run `iperf` on the source:

1. Mount the LRD ISO by using software or physical media.
2. Open a command prompt (Windows) or terminal (Linux) and go to the location of the iPerf tool:

   ```
   cd <media>/LRDTools/iperf_2.0.X/
   ```

3. As determined by the source operating system, go to the `windows` or `linux` subdirectory:

   ```
   cd windows
   ```

   ```
   -OR-
   ```

   ```
   cd linux
   ```

4. Run the iPerf tool in client mode. Enter

   ```
   iperf -c <target_ip> -w <win_size>
   ```

   NOTE: For a Linux source, remember that the TCP window size you specify for the `-w` option should be half of the desired value.

NOTE: You can download and use `iperf3` for the calculations, which is helpful in certain scenarios where `iperf2` is unable to generate useful throughput numbers. Although the command syntax and output from `iperf3` differs slightly, it should be fairly straightforward to adapt and interpret the newer output, if necessary.

Methodology

Starting with the initial `win_size` calculated in the Calculations section, record the output from several iterations of the iPerf tool using the calculated value as well as slightly larger and smaller values. We recommend that you increase and decrease the `win_size` by increments of about 10 percent of the original value. Of course, it is assumed that only the run step is repeated for each iteration of the iPerf tool.

NOTE: For a Linux source or target, remember that the TCP window size you specify should be half of the desired value. For command syntax, see "Setup" on page 591.

Using the of 1,280,000 bytes example, you might increase or decrease `win_size` in increments of about 100,000 bytes. You can use `-w` values of `1.28M`, `1.38M`, `1.18M`, and so on as the `win_size` for the `-w` option in the `iperf` command.

Sample output from an iperf client iteration looks similar to the following:

```
iperf.exe -c 10.10.10.232 -w 1.1M

------------------------------------------------------------
Client connecting to 10.10.10.232, TCP port 5001
TCP window size: 1.10 MByte
------------------------------------------------------------
[296] local 10.10.10.224 port 64667 connected with 10.10.10.232 port 5001
[ ID] Interval       Transfer     Bandwidth
[296]  0.0-10.2 sec  11.3 MBytes  9.29 Mbits/sec
```

Sample output from the referenced target server looks similar to the following:

```
./iperf -s -w .6M

------------------------------------------------------------
Server listening on TCP port 5001
TCP window size: 1.20 MByte (WARNING: requested 614 Kbyte)
------------------------------------------------------------
[  4] local 10.10.10.232 port 5001 connected with 10.10.10.224 port 64667
[  4] 0.0-10.2 sec  11.3 MBytes  9.29 Mbits/sec
```

NOTE: The client disconnects from the server after a single iteration, while the server continues to listen until stopped by using Ctrl-C.

Use several iterations to determine the optimal value for the TCP window size.

Increased throughput indicates that you are getting closer to an optimal TCP window size. Finally, as you get closer to an optimal value, use longer iterations in order to more closely simulate real running conditions.To achieve a longer iteration, use the `-t <time_in_seconds>` option to `iperf`. This option needs to be specified only on the client side.

For example:

```
iperf.exe -c 10.10.10.232 -w 1.25M -t 60
```

After an optimal value has been determined, configure this value in the `FileTransferSendReceiveBufferSize` parameter for the appropriate PlateSpin server at:

https://<*my_ps_server*>/PlatespinConfiguration/

This global value applies to all workloads on the PlateSpin server, so care should be taken to group workloads and their respective networks in a sensible manner across available PlateSpin servers.

NOTE: The value for the `FileTransferSendReceiveBufferSize` parameter setting is the optimal value you determined for `win_size`. Migrate automatically takes the Linux behavior of halving buffer sizes into consideration when it configures your Linux workloads.

Expectations

Modifying the TCP window size indirectly with the TCP send/receive buffer size can be a very effective method for increasing network throughput in some scenarios. Two to three or even more times the original throughput can sometimes be achieved. However, it is important to remember that the network characteristics can (and often do) change over time because of changes in usage patterns, hardware, software, or other infrastructure.

We strongly recommend that you use this method to calculate the optimum value at the same time of day and under the same network usage patterns that are intended to be in use during the planned live migration or replication tasks. We also recommend that you recalculate the setting periodically in order to account for changing network conditions.

VIII Documentation Updates

This guide has been updated since the General Availability of PlateSpin Migrate 2020.2.

K Documentation Updates

This section contains information on documentation content changes that were made in the English translation of the *PlateSpin Migrate Installation and Upgrade Guide* since the General Availability of PlateSpin Migrate 2020.2.

September 2020

PlateSpin Migrate discontinues support for migration to Azure China targets. So, removed all instances of Azure China across the guide.

April 2020

Location	Update
Appendix D, "Linux Distributions Supported by Migrate," on page 343	Updated this section to restructure the information, add information about how to copy custom drivers from one Migrate Server to other Migrate Server.
"Migrate Agent Commands for Linux" on page 365	Modified the `driver-install` command to update information about the custom drivers