MICRO FOCUS

PlateSpin® Transformation Manager

User Guide

December 2016

Legal Notice

Contents

About This Book and the Library

The *User Guide* provides information about using PlateSpin Transformation Manager to manage your large-scale data center transformation and migration projects. It includes conceptual information, an overview of the user interface, and step-by-step guidance for common tasks.

Intended Audience

This document is intended for IT staff, such as data center administrators and operators, who use the product to plan, execute, and track transformation projects.

Information in the Library

The documentation library for this product is available in HTML and PDF formats on the PlateSpin Transformation Manager Documentation website (https://www.netiq.com/documentation/platespin-transformation-manager-1-1/).

Additional Resources

We encourage you to use the following additional resources online:

- Micro Focus Workload Migration and Disaster Recovery channel on YouTube.com (https://www.youtube.com/channel/UChuzpo3HbYpPl93icqeOzJQ): A channel that offers product webcasts, demos, and training.
- PlateSpin Transformation Manager Product Resources (https://www.netiq.com/products/platespin-transformation-manager/resources/): A product website that offers white papers and other technical information.
- User Community (https://www.netiq.com/communities/): A web-based community with a variety of discussion topics.
- Support Knowledgebase (https://www.netiq.com/support/kb/): A collection of in-depth technical articles.

Customer Care

For specific product issues, contact Customer Care (http://www.netiq.com/support).

Documentation Support

Our goal is to provide documentation that meets your needs. If you have suggestions for documentation improvements, click **comment on this topic** at the bottom of any page in the HTML version of the documentation. You can also email Documentation-Feedback@netiq.com. We value your input and look forward to hearing from you.

1 Overview of PlateSpin Transformation Manager

PlateSpin Transformation Manager is a planning, tracking, and automation solution for data center transformation projects. It features familiar project roles, progress visualization with a dashboard, and support for all workload transformation methods. Import with automatic discovery simplifies and standardizes the setup of workloads and target hosts for planning. In Automated Mode, you can control the transformation workflow from import to cutover from a single point of control across large server farms of PlateSpin Migrate servers. Your team can increase project predictability, transformation speed, and success ratios, which helps reduce overall project costs and time to completion.

PlateSpin Migrate Connector supports workload and host discovery, load-balances the assignment of migration jobs to PlateSpin Migrate servers, and manages communications for the execution and monitoring of transformation plans for your VMware migration projects.

Inherent Challenges for Workload Transformation

As your business evolves, the data center can expand unevenly or in very dissimilar ways through mergers and acquisitions. Legacy and new technologies coexist. Your IT staff maintains a heterogeneous mix of hardware architectures, operating systems, and applications. This workload diversity increases the stress on your IT staff as well as the likelihood of human error. Older hardware consumes more power and requires more cooling.

Benefits of Transformation

You want to transform your workloads to achieve these benefits:

- Optimize workload diversity to better meet your current and future business needs
- Simplify daily operations
- Improve overall efficiency
- Reduce operational costs
- Reduce risks in the IT environment

Goals for Transformation

The purpose of any workload transformation or migration is to change workloads from their current modes of operation to appropriate future modes of operation. How you achieve the change depends on the types of workloads you manage and your business needs. Typical project objectives include the following:

- Migrate workloads between physical, virtual, and cloud infrastructures.
- Upgrade workloads to newer hardware, different hardware vendors, or hosted provider hardware.
- Consolidate workloads on virtualization host servers or to cloud infrastructures.
- Move virtual files to newer virtualization host servers, running the same or different virtualization hypervisors.
- Lift and shift equipment from location A to location B.
- Decommission old workloads as you retire software and services.

Your transformation or migration project might be a combination of any of these goals, or thousands of instances of the same one.

Challenges for Transformation

Migration planning is unwieldy. Migrating workloads from one place to another is easy to do if you have a few servers, or even 100 servers. Large-scale migration project might have thousands, or even hundreds of thousands of workloads. The information and planning requirements are not easily captured in a spreadsheet.

Business takes priority. Transformations require minimal downtime for mission critical applications and services. Each workload transformation has different priorities and windows of opportunity based on business demands. Schedules must consider the availability of target facilities, network resources, equipment, and the IT staff needed to plan and execute the transformation. Management and organization stakeholders want to track the progress and status of your projects.

Assessment is tedious. Defining the original state of a workload can be tedious. You create a profile of the workload that includes information about its compute infrastructure, operating system, applications, data, and configuration. Because workloads might be upgraded or repurposed over time, the profile might need to be augmented or updated before you execute the transformation. The related proposed workload profile might also need to change as appropriate to the revisions to the original workload.

The process seems never-ending. Large-scale IT transformation projects typically occur over an extended period in a production environment that might span multiple locations. Complex projects with massive numbers of workloads might take months or even years to complete. It might be possible to plan details only a few months in advance. Projects require multiple phases, not a one-time effort.

Benefits of Using Transformation Manager for Large-Scale Transformations

PlateSpin Transformation Manager brings together all aspects of planning into a transformation methodology that is:

- Consistent

- Reliable
- Repeatable

Plan and Track Events through the Full Transformation Life-Cycle

Each workload transformation plan identifies the current and future environment for each workload, including the hardware, applications, and other dependent resources that must be in place for a successful cutover to the target workload. You manage and track progress for each workload independently through all phases from import to completion. You can also track cumulative metrics for each batch, wave, and project.

Automate Migrations across Large Server Farms of PlateSpin Migrate Servers

In a PlateSpin Migration Environment, you can plan, execute, and monitor workload migrations through PlateSpin Transformation Manager. The PlateSpin Migrate Connector load-balances thousands of migration jobs across large farms of PlateSpin Migrate servers. Transformation Manager automates each workload's migration, according to its transformation plan, through the Migrate Connector.

Planning Features Are Flexible

You can create custom resources for each project, and add custom fields to track information important to you. With powerful forms for search and bulk actions, you can efficiently organize workloads into batches and waves, and apply the same settings on multiple workloads at a time.

Control Access and Visibility through Role-Based Permissions

Planning involves more than the IT staff who performs the technical work. You can assign permissions for key stakeholders to monitor the project status and reports. Keeping interested parties involved with data migrations helps minimize or eliminate potential conflicts for the execution of transformations.

Performance and Scalability Are Built-In

The Web Interface and database are designed to manage up to about 500,000 transformations in a single project.

Role-Based Multi-Tenancy Supports Multiple Customers and Projects

Role-based multitenancy allows you to manage multiple organizations and projects, while protecting the security and confidentiality of their data. Organizations can be the end customers of providers, or different departments in an enterprise. User roles and their related permissions control the actions and visibility of information for users.

Your Transformation Goals Are Achievable

Using PlateSpin Transformation Manager to plan and monitor your transformation projects allows you to achieve your transformation goals:

- Dramatically increases project predictability
- Increases transformation speed, reducing the time to completion
- Improves the success ratio and reduces the likelihood of human error
- Reduces the costs

PlateSpin Migration Environment

The PlateSpin Migration Environment allows you to automate many tasks for workload migration by combining PlateSpin Transformation Manager with PlateSpin Migrate Connector and one or more PlateSpin Migrate servers. You can plan and execute migrations of physical and virtual workloads to target VMs on VMware Cluster hosts. The transformation workflow and schedule determine when migration tasks are executed. Transformation Manager can pause automation to allow the Migration Specialist to manually perform some tasks. Migration Specialists can monitor the workload migration and respond to exceptions, freeing them to handle more migrations in less time.

PlateSpin Migrate Connector integrates activities between Transformation Manager and Migrate servers. It load-balances the migration jobs across large farms of PlateSpin Migrate servers in the project. Migrate Connector listens for migration events from Transformation Manager and delivers commands to the appropriate Migrate servers. Migrate Connector listens for migration status events from the various PlateSpin Migrate servers and delivers events only to the appropriate project and workload transformation plans.

Figure 1-1 illustrates the deployment environment for automated PlateSpin migration. See Table 1-1 for a description of how automated migration works in a PlateSpin Migration Environment.

Figure 1-1 *PlateSpin Migration Environment*

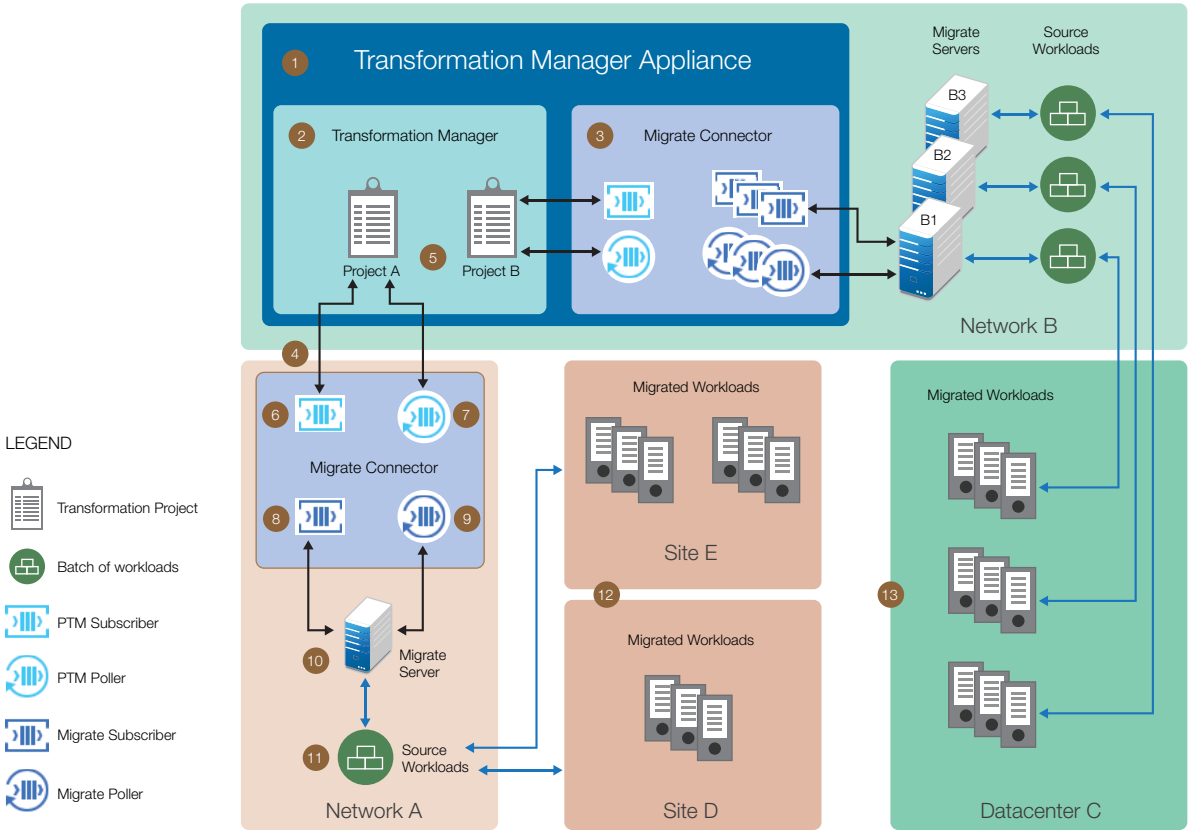

Table 1-1 *How Automated Migration Works in a PlateSpin Migration Environment*

PlateSpin Migration Environment	Description
1. PlateSpin Transformation Manager Appliance	The appliance VM hosts the PTM Server and an instance of the Migrate Connector.

PlateSpin Migration Environment	Description
2. PlateSpin Transformation Manager Server (PTM Server)	A single PTM Server manages one or more Migrate Connector instances.
3. PlateSpin Migrate Connector on the appliance	The Migrate Connector instance on the Appliance is preconfigured to work with the PTM Server. This instance can integrate events for one or more Migrate servers in the same network as the Appliance.
4. PlateSpin Migrate Connector deployed in other networks	You need a Migrate Connector instance in the same network as the PlateSpin Migrate servers and the source workloads to be migrated. Install the Connector on your own servers running SUSE Linux Enterprise Server (SLES) 11 SP4.
5. Transformation projects	A Migrate Connector works with all projects unless you configure it for a single project. A project-based connector ensures the privacy and security of an organization's data in a multi-tenant environment.
6. PlateSpin Transformation Manager Subscriber	Each Connector has one PTM Subscriber. The subscriber listens for events pushed dynamically from its assigned PTM Server. If you assign the Connector to a project, the subscriber listens only for events from its project.
7. PlateSpin Transformation Manager Poller	Each Connector has one PTM Poller. The poller periodically polls its assigned PTM Server to check that it has received all events since the last poll. If you assign the Connector to a project, the poller checks only for events from its project.
8. PlateSpin Migrate Subscriber	Each Connector has a separate Migrate Subscriber for each connected Migrate server. Each subscriber listens for events pushed dynamically from its assigned Migrate server. If you assign the Connector to a project, the subscriber listens only for events for its project.
9. PlateSpin Migrate Poller	Each Connector has a separate Migrate Poller for each connected Migrate server. Each poller periodically polls its assigned Migrate server to check that it has received all events since the last poll. If you assign the Connector to a project, the poller checks only for events for its project.
10. PlateSpin Migrate servers	For a project, you create a Migration Server resource for each PlateSpin Migrate server that you will use to execute workload migrations. When migration jobs begin, the Connector initiates a subscriber and poller for the specified Migrate server and starts listening and polling for migration state events.

PlateSpin Migration Environment	Description
11. Source workloads	For a project, you import basic information about the source workloads that you plan to migrate, then an automatic discovery process adds the details. You can manually assign a Migration Server resource to a source workload, or you can allow the Connector to auto-assign a Migration Server resource. Auto-assignment ensures that workload migrations are load-balanced across a server farm of Migrate servers. After you submit a workload, the migration workflow progresses according to the workload's transformation plan through the Migrate server.
12. Workloads migrated to hosts in different sites	Each workload's transformation plan defines the proposed workload and its target VMware cluster and network. You organize the workload migrations into waves and batches, and schedule them according to your business needs. In this example, you plan to migrate workloads to multiple sites in the same or different network. Workloads in a batch have the same destination site. The Migration Specialist at each site manages the migrations to the site.
13. Workloads migrated to different hosts in the same site	In this example, you plan to migrate workloads to different VMware clusters in a data center. Workloads in a batch have the same destination VMware cluster. The Migration Specialist responsible for the target VMware cluster manages the migrations to that cluster.

For deployment information, see "Deployment Requirements" in the *PlateSpin Migrate Connector Quick Start*.

PlateSpin Discovery Environment

In a PlateSpin Discovery Environment, PlateSpin Transformation Manager works with the PlateSpin Migrate Connector to provide automated discovery on import. You can discover source physical machines, source VMware virtual machines, and target VMware Cluster hosts in a project.

NOTE: PlateSpin Migrate servers are not required for discovery. They can be set up later in your VMware migration projects.

Import with automatic discovery simplifies and standardizes the setup of workloads and target hosts for planning. You provide minimal connection information and logon credentials for the machines. The discovery process retrieves details about each machine, populates properties for the related object in the planning database, and creates a proposed workload based on those settings.

- **Source workload discovery:** Transformation Manager provides three methods of import and automatic discovery of workloads:
  - Spreadsheet

- Range of IPv4 addresses (0 to 255)
- Single IPv4 address

See “About Automated Workload Discovery” on page 137.

- **Target host discovery:** Transformation Manager provides automatic discovery for target VMware Cluster hosts. Discovery adds the Host resource and adds resources for its discovered networks, datastores, and resource pools. See “About Automated Host Discovery” on page 199.

You can retry failed discoveries for a single workload or for multiple workloads. You can also rediscover modified workloads or hosts if needed.

Planning Mode and Automated Mode

PlateSpin Transformation Manager supports the full transformation lifecycle. For each project, you can choose whether to work in Planning Mode or Automated Mode. You can switch between the two modes if needed.

- **Planning Mode:** Planning Mode allows you to plan and track a variety of workload migrations that you execute manually. You can also configure individual workloads for automated migration.
- **Automated Mode:** Automated Mode allows you to plan, execute, and track automated workload migrations in your PlateSpin Migration Environment. You can also configure individual workloads for manual migration.

  Automated Mode supports automated execution of transformations from physical machines and VMware virtual machines to target virtual machines on VMware Cluster hosts.

  For more information, see “About Automated Migration” on page 145.

Transformation Methods

PlateSpin Transformation Manager supports planning for any-to-any transformation and migration methods as well as virtual file move, decommission, and lift and shift.

Figure 1-2 *Any-to-Any Workload Transformation Methods*

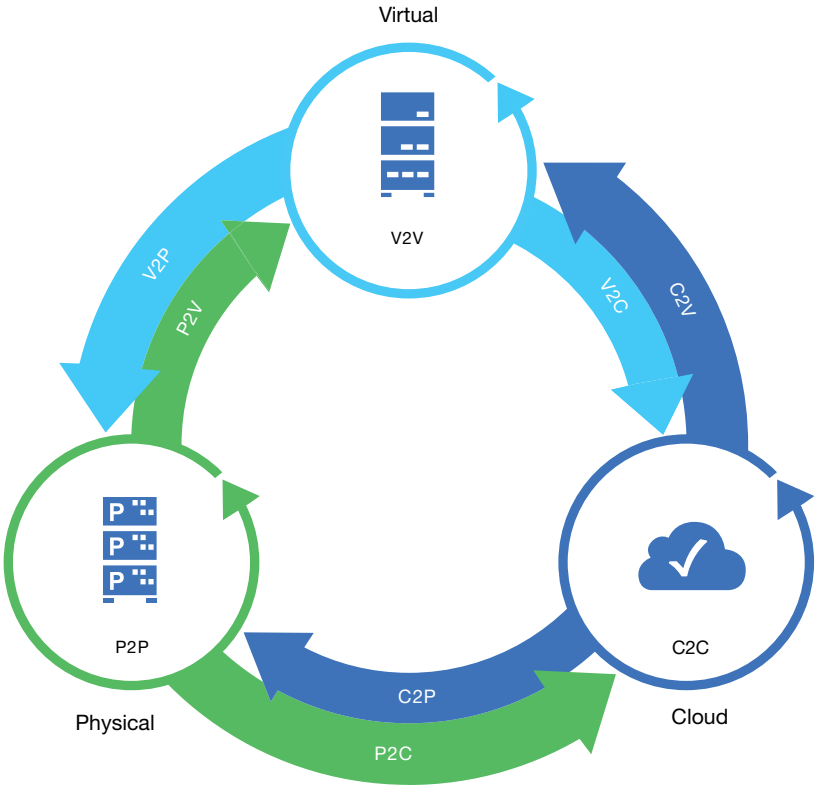

Transformation Planning Workflow

PlateSpin Transformation Manager supports transformation workflow planning that is compatible with any migration solution. Figure 1-3 shows the workflow for a transformation project as you track it using the Web Interface.

Figure 1-3 *Transformation Workflow*

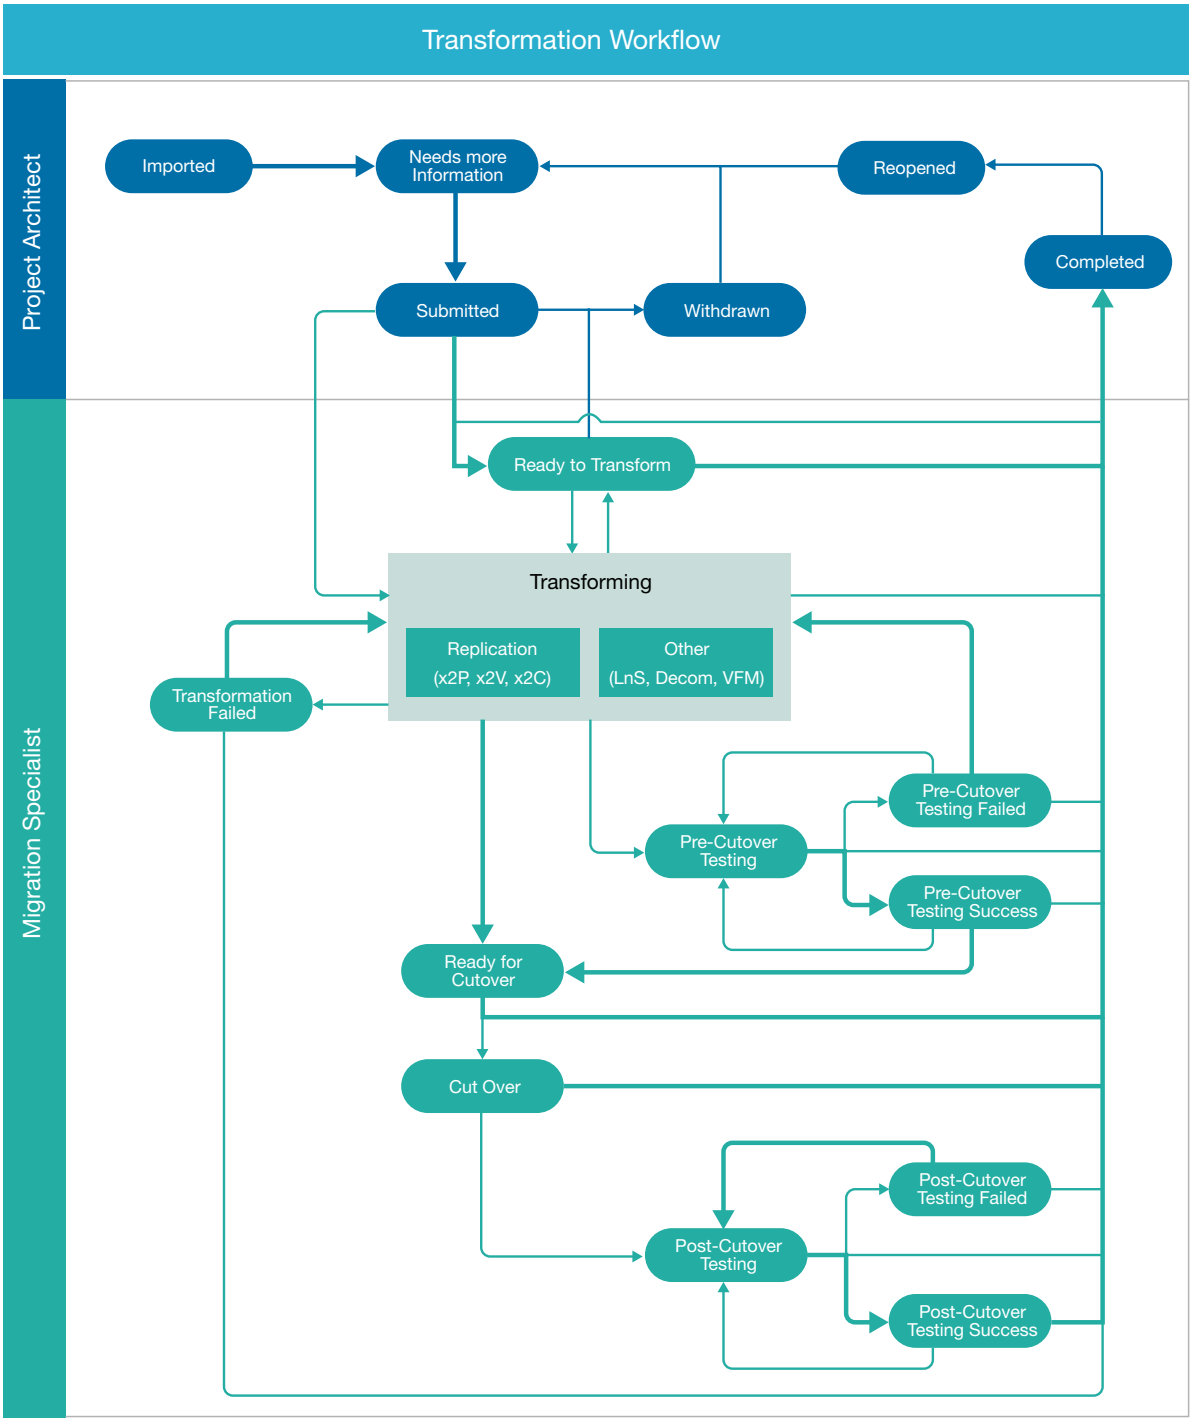

You begin by setting up components to represent different logical and physical aspects of your IT environment. Global components provide predefined resources.

As the Project Manager or Project Architect, you import information about the source workloads for transformation. You can define details about the workload, including hardware, applications, NICs, and disks. You can import details incrementally to add information, evolve the information as the project matures, or fix errors.

NOTE: If you use PlateSpin Migrate Connector and PlateSpin Migrate servers, import methods use discovery capabilities to collect information about the source workloads you add to a project. See "PlateSpin Migration Environment" on page 16.

The import process creates a target workload based on the settings for the original workload. When you re-import the source workload information, the updates apply to the target workload until you begin to plan the transformation.

As Project Architect, you define the future workload environments and refine details about the target workloads as they are known. You assign the workloads to scheduled waves and batches.

The Migration Specialists can track the transformation phases by entering different state information about the work in progress. See Table 19-1, "State Descriptions and Next States," on page 150.

Transformation Manager tracks the health of the project and provides progress metrics in the dashboard. As Dashboard Viewers, your project stakeholders can view metrics and related reports.

Key Components and Capabilities

PlateSpin Transformation Manager provides an appliance-based installation that is simple to deploy and maintain, and a web interface for users that is easily accessible on standard web browsers and is easy to use.

Appliance

The PlateSpin Transformation Manager Appliance hosts the Transformation Manager software, an instance of the PlateSpin Migrate Connector software, and the database for your transformation projects.

- **The appliance is ready to deploy in your virtualization environment.** The virtual machine includes a fully tested software stack that PlateSpin Transformation Manager needs, and omits unneeded applications and services that can consume system resources.
- **Web-based administration tools simplify server setup and maintenance.** You do not need to understand the underlying operating system, software, or databases.
- **PlateSpin Migrate Connector is ready to go.** PlateSpin Transformation Manager Appliance includes an instance of PlateSpin Migrate Connector that is automatically installed and configured to work with the PlateSpin Transformation Manager server running on the Appliance.

- **Update Channel makes appliance maintenance easy.** You can use the Online Update option to manage post-release product updates for Transformation Manager server software, product updates for the PlateSpin Migrate Connector software for the installed instance, and security updates for the software and operating system on the Appliance host VM. See "Online Update" in the *PlateSpin Transformation Manager Appliance Guide*.

For more information, see the *PlateSpin Transformation Manager Appliance Guide*.

Web Interface

The PlateSpin Transformation Manager Web Interface allows access to project information from anywhere at any time on a range of devices.

- **Securely access project information from anywhere at any time.** Visibility and actions for both internal and external stakeholders are appropriate to their user role.
- **Prioritize, organize, and schedule tasks.** Break down transformation goals into manageable chunks by project, wave, and batch.
- **Plan and track workload transformations.** Capture details for source workloads and target workloads in physical, virtual, and cloud infrastructures, with any-to-any migration, lift and shift, virtual file move, and decommission.
- **Define custom components for each project.** Associate workloads with applications, credentials, hosts, migration servers, networks, datastores, resource pools, and environments.
- **Achieve fast and efficient workload manipulations of multiple workloads at a time.** Powerful forms for Advanced Search, Bulk Edit, and Bulk Status Change allow you to apply the same settings on selected workloads.
- **Monitor project status using the real-time visual dashboard and status reports.** Internal and external stakeholders can track progress by project, wave, and batch, and view warnings for missed deadlines and schedule deviations.
- **Manage multiple concurrent projects for different organizations.** Enterprises, managed service providers, and system integrators can leverage role-based access and multi-tenancy to keep end customer data secure and confidential.
- **Effectively handle massive planning efforts.** Projects can scale from thousands to tens of thousands of workloads with minimal impact on performance.
- **Automatically discover source workloads and target VMware hosts.** You can automatically discover details about a workload or host on import. See "PlateSpin Discovery Environment" on page 18.
- **Automatically execute and track your VMware migration projects.** In a PlateSpin Migration Environment, you can automate the execution of migrations for your VMware migration projects. See "PlateSpin Migration Environment" on page 16 and "About Automated Migration" on page 145.

For more information, see "Planning Projects in the Web Interface" on page 24 and "Using the Web Interface" on page 31.

PlateSpin Migrate Connector

PlateSpin Migrate Connector supports automated discovery and migration by integrating Transformation Manager with the PlateSpin Migrate servers in a project. It provides several advantages for planning, managing, and executing workload transformation projects:

- **Integrates PlateSpin Transformation Manager and PlateSpin Migrate servers.** Migrate Connector integrates Transformation Manager and your PlateSpin Migrate servers by using event messaging and secure REST API communications.
- **Allows multiple Connector instances.** You can have multiple Connector instances for a single Transformation Manager server, where each instance is associated with a separate project.
- **Configure global settings for Migrate Connectors.** Global configuration settings on the Migrate Connector page in PlateSpin Transformation Manager apply to all Connector instances across all projects. See Chapter 5, “Configuring PlateSpin Migrate Connector,” on page 51.
- **Provides automated discovery of source workloads and target hosts.** Migrate Connector works with import options in Transformation Manager to automate discovery of source Windows and Linux workloads and target VMware Cluster hosts for your project.
- **Load-balances migration jobs across available Migrate servers.** Migrate Connector uses round-robin load-balancing to distribute workload migration jobs evenly across large server farms of the PlateSpin Migrate servers in your project.
- **Drives the automated migration on Migrate Servers.** In Automated Mode, Migrate Connector drives the automated execution of workload migrations for your VMware migration projects in a PlateSpin Migration Environment based on each workload's transformation plan. Global settings control when submitted migration jobs are set up before migration and when the jobs are removed after cutover.
- **Coordinates communications in the PlateSpin Migration Environment.** Migrate Connector supports polling and eventing types of communications in a PlateSpin Migration Environment.
  - Migrate Connector listens for migration events from Transformation Manager and delivers commands to the appropriate Migrate servers.
  - Migrate Connector listens for migration status events from the PlateSpin Migrate servers and delivers them only to the appropriate project and workloads.
- **Supports user-provided callouts.** Migrate Connector supports user-provided callouts during the transformation workflow that integrate Transformation Manager with your internal systems.

For more information, see the *PlateSpin Migrate Connector Quick Start*.

Planning Projects in the Web Interface

Most of your interaction with the PTM Server occurs through the Web Interface. This web-based planning tool allows access to project information from anywhere on a range of devices. You use the following key Web Interface components to efficiently manage your transformation projects for one or more organizations.

Dashboard

- Key stakeholders view statistics by project, wave, and batch, according to their assigned roles.
- Workload Status shows the total number of workloads and their current status as imported, defining, submitted, in progress, completed, and warning. Each category provides a link to a list of workloads in that category.
- Workload Breakdown shows core statistics by the transformation method and operating system.

For more information, see “Using the Dashboard” on page 113.

Planning

- Create projects for one or more organizations. Set each project to Planning Mode or Automated Mode.
- Plan large-scale transformation projects in manageable chunks by grouping the workload transformations in waves and batches.
- Set the start and end dates for projects, waves, and batches. Child objects automatically inherit dates from their parents. You can also set dates manually.
- Set the global project selector to automatically filter lists and dialogs for a specific project.
- Import the source workload information for your project. You can leverage automated discovery to import details for workloads.
- Assign workloads to batches through individual or bulk edits.
- Define custom workload resources to use for each project, such as custom fields, applications, credentials, hosts, migration servers, networks, datastores, resource pools, and environments.
- Define the appropriate transformation configuration for each workload.
- Use powerful forms for advanced search, bulk status change, and bulk edit to achieve fast and efficient workload manipulations of multiple workloads at a time.
- Use the workload information to execute their assigned workload transformations and to set the workflow status.
- Use custom fields to add project-specific details to all workloads, such as budget IDs, contact information, and tags to track logical or business associations among workloads. Custom fields are available in Advanced Search and Bulk Edit forms when the global project selector is set.

For more information, see “Overview of Project Planning” on page 91. See also the following related topics:

- “Managing Projects” on page 95
- “Managing Waves” on page 101
- “Managing Batches” on page 105
- “Managing Applications” on page 109
- “Configuring Workload Transformations” on page 145
- “Resources” on page 189

Users

- Define organizations, users, and groups.
- Assign users to the Administrators group for elevated privileges and responsibilities.

- Assign users to project roles:
  - Project Managers
  - Project Architects
  - Migration Specialists
  - Dashboard Viewers

  The multi-tenant architecture controls access for each user account based on the assigned roles.

For more information, see “Overview of PlateSpin User Management” on page 63. See also the following related topics:

- “Managing Organizations” on page 69
- “Managing Users” on page 73
- “Managing Groups” on page 77
- “Roles and Permissions” on page 81

Resources

Define custom resources that can be associated with workloads in a project:

- Credentials
- Hosts
- Migration servers
- Networks
- Datastores
- Resource pools
- Environments

For more information, see “Resources” on page 189.

Configuration

Configure global settings that apply across all projects:

- **Licenses:** View license information, including the start date, expiration date, and the number of licenses remaining in each license block.

  For more information, see “Managing Licenses” on page 45.
- **Migrate Connector:** Configure default global settings for PlateSpin Migrate Connectors used with projects in your transformation environment.

  For more information, see “Configuring PlateSpin Migrate Connector” on page 51.
- **Operating Systems:** Configure default global operating system types used by all projects in your transformation environment.

  For more information, see “Configuring Operating Systems” on page 57.

What’s Next

Use the Getting Started Checklist to learn more about the Web Interface, and to set up your first project.

2 Getting Started

After you install and configure your PlateSpin Transformation Manager Appliance, you can use this checklist to get acquainted with the Web Interface and set up a transformation project for your organization.

Table 2-1 *Getting Started Checklist*

Status	Task	For information, see
❑	1. Log in to the Web Interface using the System Administrator user account that you created during the appliance setup.	Accessing the Web Interface (page 33)
❑	2. Familiarize yourself with the Web Interface.	Web Interface Toolbar (page 33) Project Selector (page 35) Bulk Actions (page 35) Show Link for Navigation URLs (page 36) Custom Display and Filters for Lists (page 36) Multiple Item Selection in Lists (page 39)
❑	3. Create an organization.	About Organizations (page 69) Creating an Organization (page 70) (0ptional) Uploading an Organization Logo (page 71)
❑	4. Create a transformation project.	About Projects (page 95) Creating a Project (page 97)
❑	5. Create one or more Waves for the project.	About Waves (page 101) Creating a Wave (page 102)
❑	6. For each Wave, create one or more batches.	About Batches (page 105) Creating a Batch (page 106)
❑	7. Create user accounts for your project, and assign each user to project roles.	Roles (page 63) Creating a User (page 74)
❑	8. Prepare one or more Import Spreadsheets for the project, or use Import with auto-discovery.	Bulk Import Spreadsheet (page 173) Importing Workloads with a Spreadsheet (page 139) About Automated Workload Discovery (page 137) Importing Workloads with Auto-Discovery (page 140)

Status	Task	For information, see
❑	9. Create applications that are available to proposed workloads as you define transformations.	About Applications (page 109) Creating Applications (page 110) Creating Applications during Spreadsheet Import (page 110)
❑	10. Create target resources that are available to proposed workloads as you define transformations.	Resources (page 189) ◆ Credentials ◆ Hosts ◆ Migration Servers ◆ Networks ◆ Datastores ◆ Resource Pools ◆ Environments
❑	11. For each imported workload, use bulk edit or individual edit to define the transformation method and details for the proposed workload.	Overview of Workloads (page 131) Editing the Workload Transformation (page 162) Bulk Edit for Multiple Proposed Workloads (page 165)
❑	12. As a Project Architect, submit one or more workloads for transformation.	Status and Retry (page 149) Figure 19-1, “The Transformation Workflow,” on page 149 Bulk Status Change for Multiple Proposed Workloads (page 169)
❑	13. (Automated Migration) As a Migration Specialist, monitor the workload transformations.	Automated Mode (page 131) Status and Retry (page 149) Migration Sub Status (page 125)
❑	14. (Manual Migration) As a Migration Specialist, execute the workload transformations, and update their transformation status in PTM.	Status and Retry (page 149) Figure 19-1, “The Transformation Workflow,” on page 149 Bulk Status Change for Multiple Proposed Workloads (page 169)
❑	15. Monitor the project status by viewing core metrics and related reports in the real-time visual dashboard.	Using the Dashboard (page 113) Viewing Counts and Status for a Project, Wave, or Batch (page 116)

3 Using the Web Interface

Most of your interactions with PlateSpin Transformation Manager occur in the Web Interface.

Prerequisites for Using the Web Interface

The PTM Server hosts the Web Interface, using Jetty web services. After you install the PlateSpin Transformation Manager Appliance, ensure that your network environment meets the requirements in this section for using the Web Interface.

Port Requirements

During the appliance installation, PlateSpin Transformation Manager automatically configures the default ports shown in Table 3-1 on the PlateSpin Server. Ensure that you open the ports in any firewalls in your network between the PlateSpin Server and the computers you use to access the Web Interface.

NOTE: For remote communications, Micro Focus recommends that you use the secure port and SSL options for accessing the Web Interface.

Table 3-1 *Default Network Ports*

Network Port	Security	Description
8183	Secure (SSL)	HTTPS port for the Web Interface. Allow TCP and UDP traffic, incoming and outgoing.
8182	Not secure	HTTP traffic for the Web Interface. This port is disabled by default. If you enable this port, allow TCP and UDP traffic, incoming and outgoing.
5432		PostgreSQL port for a remote PTM database. Allow TCP traffic, incoming and outgoing. This port is closed by default if the PostgreSQL database is installed on the appliance.

For information about modifying the port setting for the Web Interface on the PlateSpin Server, see "Web Server Configuration" in the *PlateSpin Transformation Manager Appliance Guide*.

Supported Browsers for the Web Interface

You can access the PlateSpin Transformation Manager Web Interface using any of the following supported browsers:

- Mozilla Firefox, latest release
- Microsoft Internet Explorer 11
- Google Chrome, latest release

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

Supported Languages

PlateSpin Transformation Manager Web Interface supports English (En) and Japanese (Ja) languages in your web browser. Modify the Language setting in your web browser with your preferred language as the first in the list.

SSL Certificate for Secure Communications

The installation of the PlateSpin Transformation Manager Appliance generates and installs a self-signed certificate for SSL (Secure Sockets Layer) communications. It uses the DNS name that you specify for the PlateSpin Transformation Manager Appliance. The certificate applies to the appliance and the software.

For higher security, Micro Focus recommends that you use a trusted server certificate that is signed by a trusted certificate authority (CA) such as VeriSign or Equifax. You can use your own existing signed certificate, or you can use the Digital Certificate tool on the appliance to create a certificate, have it signed by a trusted certificate authority, and then add it to the appliance. The DNS name of the server must match the subject of the security certificate. To import your signed certificate, you must provide the certificate and key, as described in "Digital Certificates" in the *PlateSpin Transformation Manager Appliance Guide*.

Session Timeout

The default session timeout occurs after 30 minutes of inactivity. The timeout interval is configurable with a global setting on the PlateSpin Transformation Manager Appliance. Contact your System Administrator if you would like to specify a shorter or longer interval. See "Web Interface Session Timeout" in the *PlateSpin Transformation Manager Appliance Guide*.

Accessing the Web Interface

Most of your interaction with PlateSpin Transformation Manager takes place through the browser-based Web Interface. You use it to plan and manage your transformation projects, as well as to configure the product to meet your needs.

During the Appliance installation, you set up the default System Administrator user account for the Web Interface. This user has system-wide permissions. Log in as the System Administrator to create accounts for other users and assign roles to them. See Section 7, "Overview of PlateSpin User Management," on page 63.

To log in to the Web Interface:

1. In a supported web browser, launch the PTM Web Interface:

   https://<*ptm-server-dns-name_or_ipaddress*>:8183 (secure, default)
   http://<*ptm-server-dns-name_or_ipaddress*>:8182 (not secure, disabled by default)

2. Specify the email address and password of your user account, then click **Log In**.

Web Interface Toolbar

The PlateSpin Transformation Manager Web Interface toolbar gives you access to the key product features described in Table 3-2.

Figure 3-1 *PlateSpin Web Interface Toolbar*

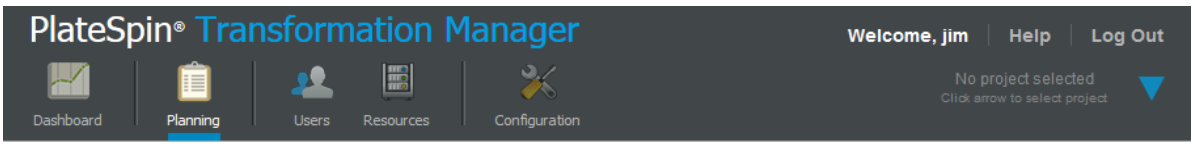

Table 3-2 *Toolbar Options*

Option	Description
Dashboard	The Dashboard tab provides status and health information about workload transformations. See "Using the Dashboard" on page 113.

Option	Description
Planning	The Planning tab allows you to configure the following elements for your transformation project: ◆ Projects ◆ Waves ◆ Batches ◆ Applications ◆ Workloads See Part III, “Planning,” on page 89.
Users	The Users tab allows you to configure the following elements for your transformation project: ◆ Organizations ◆ Users ◆ Groups See Part II, “Users,” on page 61.
Resources	The Resources tab allows you to configure the following elements for your transformation project: ◆ Credentials ◆ Hosts ◆ Migration Servers ◆ Networks ◆ Datastores ◆ Resource Pools ◆ Environments See Part V, “Resources,” on page 189.
Configuration	The Configuration tab allows the System Administrator to perform the following tasks: ◆ Licenses ◆ Migrate Connector ◆ Operating Systems See Part I, “Configuration,” on page 43.
Project Selector	The Project Selector narrows the scope globally to a single project for the information displayed and acted on. When it is enabled, the dialogs automatically complete fields for the project and its parent organization. See “Project Selector” on page 35.
Help	The Help menu provides links to help pages for your current location, help for common tasks, a link to online documentation, and information about the product.

Project Selector

You can use the Project Selector to set the global scope to a single project for the information displayed and acted on. When it is enabled, the dialogs throughout the product automatically complete fields for the project and its parent organization. It filters all tables to display only the components assigned to or associated with the selected project. It also adds the custom fields that are defined for the project to the Advanced Search form and the Bulk Edit form.

When a project is selected, the Project Selector displays the name of the project and the logo for the project's parent organization, if it is available.

To add the project-level filter:

1. Click the Project Selector arrow to open a list of projects.
2. Locate and select the appropriate project. You can scroll or filter the list to locate the project of interest.

   The selected page refreshes the list to display objects only for the selected project.

To remove the project-level filter:

1. Mouse over the displayed project name and logo in Project Selector area, then click the **X** to remove the project filter.

   The selected page refreshes the list to display objects for all projects.

Bulk Actions

PlateSpin Transformation Manager provides bulk actions to help you manipulate and perform actions on multiple objects at a time.

- "Spreadsheet Import" on page 35
- "Range Import" on page 35
- "Bulk Edit" on page 36
- "Bulk Status Change" on page 36

Spreadsheet Import

You can import source workload information by importing or re-importing the machine information using the Bulk Import spreadsheet. See Appendix C, "Bulk Import Spreadsheet," on page 173.

In the PlateSpin Migration Environment, the PlateSpin Migrate Connector provides automated discovery of details for machines in the spreadsheet after a successful spreadsheet import.

Range Import

You can import a range of 256 IPv4 addresses using the Import option. The PlateSpin Migrate Connector provides automated discovery of details for any workloads discovered in a specified address range.

Bulk Edit

You can apply settings to multiple proposed workloads at a time by setting values in the Bulk Edit dialog. Use the Project Selector, Filter, Advanced Search, sort, and multiple select functions to refine the workloads list and select the items for action. Use the Bulk Edit dialog to specify a value for one or more parameters in the selected workloads. Each specified value is set independently for a selected workload only if the value is a valid setting for it. See “Bulk Edit for Multiple Proposed Workloads” on page 165.

Bulk Status Change

You can apply state change settings to multiple proposed workloads at a time by using the Bulk Status Change tool. Use the advanced search and sort options to refine the list and select the items for action. Use the Bulk Status Change form to specify the appropriate state for the selected workloads. See “Bulk Status Change for Multiple Proposed Workloads” on page 169.

Show Link for Navigation URLs

The **Show Link** icon for a dialog provides a navigation URL that takes you directly to the dialog for the selected object. You can paste the link in the Location bar of a supported web to return directly to the location in the Web Interface. You can also paste the link in other applications, such as email, Excel, and Word, to share with stakeholders.

Users who follow the link will be prompted to log in to the Web Interface. The user must have a PlateSpin Transformation Manager user account to log in. To see the target dialog and its contents, the user account must be assigned to a role that has permission to access the page and the objects displayed on it.

To access a dialog’s Navigation URL for a selected object:

1. On the dialog, click the **Show Link** icon in the upper right corner to show the URL to the dialog for the selected object.
2. (Optional) Click the link to open the dialog in a new tab, then log in again to the Web Interface to gain access.
3. (Optional) Copy the URL to the clipboard, then do any of the following. Users of the link must log in to access the dialog.
   - Paste the link in the Location bar of any supported web browser to open the dialog in a different browser window.
   - Paste the link in an email to share it with stakeholders.
   - Paste the link in other applications to provide convenient access to the object’s information.

Custom Display and Filters for Lists

In the Web Interface, you can personalize the display for lists by using display tools integrated throughout the product.

- “Refresh List Items” on page 37
- “Number of Items in a List” on page 37
- “Show More Data in a Cell” on page 37

Refresh List Items

Lists poll to refresh the data about once per minute. Double-click the tab title to force any list to refresh immediately, regardless of its normal polling cycle.

Number of Items in a List

Mouse over the tab title to view the total number of items in the list and the total number of selected items currently selected in the list.

Show More Data in a Cell

If a list table cell contains more information than can be displayed, the text ends in an ellipsis (…). You can mouse over the cell to show the additional information in a pop-up infotip.

Sort Data

You can sort list data in ascending order (A to Z) or in descending order (Z to A), based on values in the selected column. The sort function treats numbers as text values and sorts them alphabetically, not numerically. When you sort data, an arrow icon in the column heading indicates the that column is the key for the sort. An Up arrow indicates an ascending sort order. A Down arrow indicates a descending sort order.

Most columns are available as data sort keys. If a column is not available as a sort key, the toggle sort does not work, and its Columns menu does not contain the sort options.

To sort data in a list:

1. Use either of the following methods to effect a sort:
   - **Toggle Sort:** Click a column heading to sort entries in ascending order based on values in that column. Click the column heading again to sort in descending order.
   - **Menu Sort:** Mouse over a column heading to activate its options, then click the arrow on the right column edge to access the menu for that column. Select **Sort Ascending** (A to Z) or **Sort Descending** (Z to A) to specify the preferred sort order.

Show/Hide Columns in a List

You can show or hide data in a list by specifying which parameters to display. Although the hidden data is not displayed, any filter or advanced search action considers the values. For example, in the Workloads list, hidden data includes location (site, enclosure, slot), custom fields, IP address, MAC address, workload type, and virtualization technology.

NOTE: Your column display preferences for each personalized list persists across your sessions.

To personalize the columns in a list:

1. Mouse over a column heading to activate its options, then click the arrow on the heading's right edge to access the menu for that column.
2. Select **Columns** to display the parameters available for the list.

   The menu lists parameters in the left-to-right display order in the list. Selected parameters show in the list. Deselected parameters are hidden.
3. Select the check box next to the parameter you want to show. The list updates to immediately add the column.
4. Deselect the check box next to the parameter you want to hide. The list updates immediately to remove the column.
5. When you are done, click anywhere on the page to exit the menu.

Filter Data in a List

Most lists in the Web Interface have a Filter option available to help you search the list and locate the information of interest. The filter applies to all searchable fields in the list, including any hidden columns. For example, on the Workloads list, the default hidden columns include custom fields, IP addresses, MAC addresses, and workload type.

NOTE: When you filter a list for a status condition, you must replace hyphens and spaces in the status name with underscores. Examples:

```
in_progress
pre_cutover_testing
```

To filter the objects in a list:

1. In the **Filter** field, begin typing a sequence of characters to display only the entries with values that match.

   The search and filter apply to the list almost immediately as it matches entries.
2. Click the **Filter** icon to clear the **Filter** field.

Advanced Search of Data

The Advanced Search option for lists helps you to locate multiple objects that you want to perform the same action on. Advanced Search is not available for every list.

The Advanced Search dialog is a multiple-option form that allows you to search on any combination of the following parameters for Workloads that make sense for the target of your search:

Project	Workload Type	Total Storage	Network
Wave	Transform Method	Single Disk Size	VLAN ID
Batch	Environment	Cores per Socket	Status
Hostname	Site	Amount of Memory	Sub Status
OS Type	Enclosure	On Hold	Health
Application	Custom 1	Custom 2	Custom 3
Custom 4	Custom 5	Custom 6	Custom 7

NOTE: The Project Selector limits the workloads in the Workloads list to the specified project. It also adds the custom fields you defined for the project to the Advanced Search form and Bulk Edit form.

You can use the **Health** > **All Warning State** option on the Workloads list Advanced Search function to see all workloads in the Warning state. For objects with warnings, you can mouse over the Status cell to view the condition that triggered the warning.

You can use the **Status** option to find all objects in a given transformation state, such as Imported, Ready to Submit, Transforming, Ready to Cutover, and so on.

To search:

1. (Optional) If you have permissions on multiple projects, click the **Global Project Selector**, then select the project of interest.
2. Click the **Advanced Search** icon ... to the right of the Filter field.
3. In the Advanced Search dialog, specify a value for any search parameter.

   The search filters the list to show matches to that setting.
4. (Optional) For combination searches, specify a value for additional parameters until you locate the objects of interest.
5. (Optional) Click **Clear** to reset the Advanced Search and try again with different fields.
6. After you have narrowed the list to the items of interest, click in the list to exit the Advanced Search dialog.

Scrolling Up and Down in Lists

The Web Interface is optimized to navigate lists using the scroll bar. Lists have special handling to provide responsive display and scrolling of list items, even for lists containing thousands of items. A list can display up to 25 items at a time. It caches about 200 nearby items in the browser to accommodate scrolling up and down the list. As you scroll, the Web Interface loads nearby items to the cache and releases items further away from the currently displayed items.

NOTE: Ensure that you use the scroll bar to move up and down lists in the Web Interface. Do not use the Up Arrow or Down Arrow keys to navigate lists.

Multiple Item Selection in Lists

In the Web Interface, lists have special handling to provide responsive display and scrolling of list items, even for lists containing thousands of items. A list can display up to 25 items at a time. It caches about 200 nearby items in the browser to accommodate scrolling up and down the list. As you scroll, the Web Interface loads nearby items to the cache and releases items further away from the currently displayed items.

The Web interface supports the familiar keyboard shortcuts for item selection: Shift+Click (consecutive items), Ctrl+Click (non-consecutive items), and Ctrl+A (all items). When you select an item, the Web Interface adds its information to a separate Selected Items cache. Selected rows are shaded light blue. Actions performed on the list apply only to items in the Selected Items cache.

Before you select items, use the Filter and Advanced Search options to reduce the list to the items of interest. Use the Sort function to group like items in the list to accommodate consecutive selection.

You can mouse over the list's Tab title to show the following:

- **List Size:** The total number of items in the list
- **Selected:** The current number of selected items in the list

Use the following instructions to navigate the list and select items:

- "Selecting Consecutive Items" on page 40
- "Selecting Non-Consecutive Items" on page 40
- "Selecting All Items" on page 40

Selecting Consecutive Items

Use the Shift+Click action to select consecutive items in a list. Before you begin, use the sort option in column headers to group the items of interest and facilitate consecutive selection.

To select consecutive items in a list:

1. Click the first item of interest, press and hold the Shift key, and then click the last item of interest in the displayed list to add those consecutive items to the cached list of selected items.
2. (Optional) Scroll up or down to show the next set of deselected items, press and hold the Shift key, then make your next last item selection to extend the consecutive item selection. Repeat the scroll and selection process as needed.
3. (Optional) Mouse over the list's Tab title to view the current count for the number of selected items.

Selecting Non-Consecutive Items

Use the Ctrl+Click action to select non-consecutive items in a list.

To select non-consecutive items in a list:

1. Press and hold the Ctrl key, and then click each item in the displayed items that you want to add to the cached list of selected items.
2. (Optional) Scroll up or down to show the next set of deselected items, press and hold the Ctrl key, then make your next non-consecutive item selection. Repeat the scroll and selection process as needed.
3. (Optional) Mouse over the list's Tab title to view the current count for the number of selected items.

Selecting All Items

The Ctrl+A action selects all of the currently loaded items in the list cache instead of selecting only the currently displayed items. After you select the loaded items, you will scroll through 140 to 200 selected items before you see the next set of deselected items.

To select all items in a list:

1. Mouse over the list's Tab title to view the total number of items in the list.

   The item count gives you an idea of how much scrolling is needed to load and select items.
2. Press Ctrl+A to select the currently loaded set of items and add them to the cached list of selected items.

3 Scroll up or down until you see the next set of deselected items.

4 Press Ctrl-A to add the currently loaded items to the cached list of selected items.

5 Repeat Step 3 and Step 4 until all items are selected.

6 Mouse over the list's Tab title to see the total item count and total selected item count. If they are the same, then all items in the list are selected.

Configuration

PlateSpin Transformation Manager allows you to configure licenses and operating system types for the PlateSpin Server. Modifications to the settings in the Configuration tool apply to all transformation projects.

4 Managing Licenses

PlateSpin Transformation Manager provides full and evaluation license options. You must have a license for each workload transformation.

About Licenses

PlateSpin Transformation Manager offers full and evaluation license options.

Full License

A full license allows you to plan and execute a workload transformation for each license instance. You purchase workload licenses in blocks from the Customer Center (https://www.netiq.com/customercenter/). Each block has a fulfillment window with an activation date that you specify.

You can install the appliance and begin setting up the projects before you need to activate licenses. You do not need a workload license until you start planning a workload transformation. Because transformation planning for a workload consumes a license, you cannot edit imported workloads until you enter a valid license key and licenses are available in an active license block.

Evaluation License

An evaluation license allows you a trial period to work with the your own data in the product. Evaluation mode allows you to use the product with an Evaluation key and 100 workload licenses.

After the Evaluation license expires, any unused licenses cannot be used to begin new workload transformations. You can continue to work in the product with licensed workloads through their entire workflow until the transformations are completed.

You can preserve the data you enter during the evaluation period by entering a Full license key to your evaluation installation. Your existing licensed workloads are not impacted. Consumed licenses do not expire. However, after you add the full license key, all unused evaluation workload licenses are automatically removed.

How Workloads Consume Licenses

You purchase workload licenses for PlateSpin Transformation Manager in blocks, based on the number of workloads you plan to manage for all of your projects. A workload license is not required to configure users, projects, and resources. You can import and re-import workload data multiple times without consuming a license.

NOTE: PlateSpin Transformation Manager must be able to communicate across the public Internet with the Micro Focus License Server to confirm the availability of a license before it can be consumed. Processes that trigger a license to be consumed will fail if the License Server is not accessible at that time. An error message will identify the cause of the outage.

Each workload will consume an available license when you first modify its data after import. The edit can be an individual edit action or a bulk edit action. Acquiring the license can take several seconds, resulting in a delayed response for the first edit on a workload. If no license is available for the workload at that time, the save action fails. You cannot modify the transformation plan for an unlicensed workload.

For a bulk edit, each unlicensed workload in the pre-selected list consumes an available license. If there are insufficient licenses available, the license and edit actions fail for any subsequent unlicensed workloads in the list.

NOTE: If licenses are not available, notify the System Administrator for PlateSpin Transformation Manager about the licenses you need.

When a license is consumed, the licenses remaining count decreases by one. If multiple license blocks are available, it modifies counts for the block with the next closest expiration date.

A consumed license never expires. Deleting a workload does not free the consumed license.

Adding a License Key

On the **Configuration** > **Licenses** page, the System Administrator can add the full license key or evaluation license key and view license status for current blocks of licenses for this PTM Server. The Licenses Blocks list shows each block with its purchase date, start date, expiration date, count for total licenses in the block, and count for remaining unused licenses. The list displays only license blocks with active fulfillment periods and unused licenses available. It does not include historical information about license blocks after the fulfillment period has expired or when the number of remaining unused licenses is zero (0).

1. (Optional) Log in to the Customer Center (https://www.netiq.com/customercenter) and acquire an evaluation license or full license key for PlateSpin Transformation Manager.
2. Log in to the Web Interface as a user with the System Administrator role:

   https://<*ptm-ip-address-or-dns-name*>:8183
3. Select **Configuration** > **Licenses**.
4. Complete the following information:

   License Key: The 14-character hexadecimal key provided by the Customer Center.

   License Type: Full or Eval

   Remaining Licenses: This value is automatically calculated from the license blocks that you purchase for this PlateSpin Server and that have active fulfillment periods.

License Blocks: This list is automatically retrieved from the Customer Center. It displays the license blocks that you have purchased for this PlateSpin Server that have active fulfillment periods with unused licenses available or that have future fulfillment periods. It does not display expired blocks or empty blocks.

5 Click **Save**.

6 Click **Close** to exit the System Configuration dialog.

Adding a License Block

You can purchase license blocks for PlateSpin Transformation Manager in the Customer Center (https://www.netiq.com/customercenter). You do not manually enter the block license information. PTM automatically retrieves information about license blocks from the Customer Center and displays them in the License Blocks list on the Licenses page. In addition to updates for user actions that involve licenses, PTM refreshes the license information with the Customer Center daily at about 2:00 a.m. local time.

Figure 4-1 *License Management*

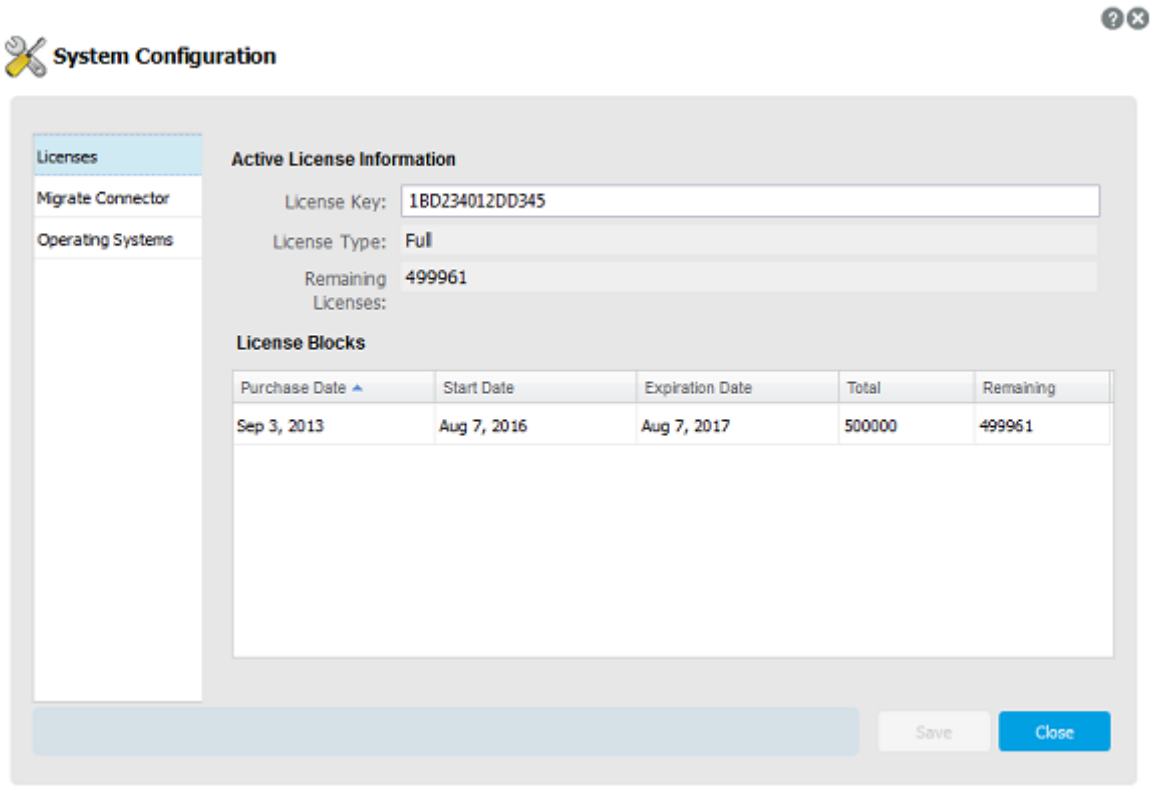

Each license block has an independent fulfillment period defined by the following dates:

- **Start Date:** The fulfillment date at which the block of licenses becomes active and you can begin to use its licenses.
- **Expiration Date:** The end of the fulfillment period. Any unused licenses in the block expire at that time.

A license block no longer displays in the list when either of the following conditions occurs:

- After the expiration date for the license block fulfillment period, even if unused licenses remain.

  NOTE: If unused licenses are due to expire soon, you can quickly consume multiple licenses by applying a Bulk Edit change to the next planned set of workloads for one or more projects. A consumed license never expires.

- After all licenses in the block are consumed, even if the block fulfillment period has not expired.

Viewing License Status for a Workload

The Transformation dialog reports the license state for the currently selected workload in the lower right of the header area. See Table 4-1 for information about license states for workloads.

Table 4-1 *Workload License States*

Icon	License State	Description
	Not licensed	The workload has been imported, and you have not yet edited the transformation plan information for the workload.
	Licensed	You edited some information about the imported workload by using the Transformation dialog or a Bulk Edit. A license is assigned permanently to this workload. Mouse over the icon to see the date the license was consumed. The license never expires.
	License is not available	You have not yet edited information for the imported workload. Because no workload licenses are currently available, any edits you attempt for the workload cannot be saved. **NOTE:** Notify the System Administrator that licenses are not available.

License Warnings

A license must be available to be consumed before you begin to plan a workload's transformation. A license is available if it is unused and its fulfillment period is active. License blocks with future start dates for their fulfillment periods are currently inactive, which makes them currently unavailable. The Web Interface notifies the user if licenses are unavailable, running low, or expiring soon.

PlateSpin Transformation Manager checks the count and expiration dates of currently available licenses when the following key actions occur in the workflow.

- At login
- When the Transformation dialog opens or closes
- When the Bulk Edit dialog opens or closes
- When the Configuration dialog closes

The License Warning message displays in a banner immediately below the main header of the Web Interface for the following license conditions:

- The number of currently available licenses is 50 or fewer.
- The next set of currently available licenses will expire within 7 days.

If no licenses are available for a currently selected unlicensed workload, the Transformation dialog displays the No License Is Available icon  in the header area, and edits cannot be saved.

NOTE: Notify the System Administrator for PlateSpin Transformation Manager about the reported license conditions.

If licenses are expiring soon, you can quickly consume multiple licenses by applying a Bulk Edit change to the next planned set of workloads for one or more projects. A consumed license never expires.

5 Configuring PlateSpin Migrate Connector

The PlateSpin Transformation Manager server works with one or more instances of PlateSpin Migrate Connector that are deployed in the networks where you have source workloads. If you associate a Connector with the PTM Server, the Connector is available to all projects that have source workloads in that network. You can also configure a separate Connector instance for each project.

NOTE: Only users assigned to the Transformation Manager System Administrator role can modify the global settings for PlateSpin Migrate Connector. If you need to modify the default workflow settings, contact the System Administrator.

For more information about the Migrate Connector, see the *PlateSpin Migrate Connector Quick Start*.

About Global Options for Connectors

On the **System Configuration** > **Migrate Connector** page in PlateSpin Transformation Manager, a System Administrator can customize the default workflow of migration projects that leverage one or more PlateSpin Migrate servers to execute workload migrations.

NOTE: You must restart PlateSpin Migrate Connector after you modify settings in order to apply the changes. See "Starting, Restarting, or Stopping Transformation Manager" in the *PlateSpin Migrate Connector Quick Start*.

General Settings

Missed Event Poll Interval

Specify the number of seconds between polls for workload migration events.

The default value is 300 seconds (5 minutes). A lower value puts more stress on your PTM server and Migrate servers.

Reconnect Retry Interval

Specify the number of seconds to wait after a connection failure before PTM tries to reconnect to a PlateSpin Migrate server.

The default value is 1500 seconds (25 minutes).

Pause for Migrate Configuration

Specify whether to pause the Transformation Workflow until a Migrate user manually configures the workload migration on a Migrate server.

The default is to disable manual migrate configuration.

Pause for Manual Pre-Cutover Testing

Specify whether to pause the Transformation Workflow in a Transforming / Incremental Replication state until the Migrate user manually triggers Pre-Cutover Testing.

The default is to disable manual pre-cutover testing.

Migrate Server Settings

Add Workload to Migrate

Specify the number of days before the start date to add the workload migration job to an auto-assigned PlateSpin Migrate server.

This option is disabled by default with a value of 0 (zero). When you submit a transformation plan, the Migrate Connector immediately auto-assigns a Migrate server, adds a workload migration job, then waits until the start date to execute the migration. The job consumes capacity and a Migrate license while it waits for start date. Consuming capacity before it is needed might block migration of workloads with earlier start dates.

Set a value of 1 or greater to enable automation control to wait until the specified pre-start-date interval to begin the preparation for migration. You can submit the workload transformation plans as they are ready without immediately consuming capacity or a license on a Migrate server.

Maximum Workloads

Specify the maximum number of workloads to allow for a PlateSpin Migrate server at a time.

The default value is 100. To determine an appropriate value for your PlateSpin Migration Environment, see “Performance” in the *PlateSpin Migrate 12.2 User Guide*.

NOTE: Capacity to add more workloads can be regained by doing the following:

- Wait until a pre-start-date interval before the start date to add the workload to a Migrate server. See Add Workload to Migrate.
- Remove the workload information after a successful cutover. See Remove Workload After Cutover.

Maximum Targets

Specify the maximum number of discovered targets for a PlateSpin Migrate server.

The default value is 27.

Verify SSL Certificate

Specify whether to enable the validation of SSL certificates for connections to the PlateSpin Migrate servers.

The default is to disable certificate validation. Select the check box to enable it.

Remove Workload After Cutover

Specify whether to clean up the workload information from the PlateSpin Migrate servers after a cutover completes.

The default is enabled. Deselect the check box to disable it.

Remove after (days)

Specify the number of days after a workload is cut over to clean up the workload information from the PlateSpin Migrate server.

The default value is 3 days.

Customer-Provided Scripts

Run Custom Import Script

Specify whether to automatically execute the Custom Import callout script after initial workload discovery.

The default is disabled.

Run Submit Validation Script

Specify whether to automatically execute the Submit Validation callout script before adding the workload to Migrate.

The default is disabled.

Run Pre-Cutover Testing Script

Specify whether to automatically execute the Pre-Cutover Testing callout script after workload replication.

The default is disabled.

Run Post-Cutover Testing Script

Specify whether to automatically execute the Post-Cutover Testing callout script after workload cutover.

The default is disabled.

Viewing Global Migrate Connector Settings

The System Administrator, Project Manager, and Project Architect can view the global settings for the PlateSpin Migrate Connector instance associated with the PTM Server.

To view the connector settings:

1. In the Web Interface toolbar, select **Configuration**.
2. In the System Configuration dialog, select **Migrate Connector**.
3. Scroll to view the settings:
   - General Settings
   - Migrate Server Settings
   - Customer-Provided Scripts
4. When you are done, click **Close** to exit the System Configuration dialog.

Editing Global Migrate Connector Settings

The System Administrator can modify any of the global settings for the PlateSpin Migrate Connector instances that are assigned to the PTM Server. The changes apply throughout the product for all transformation projects.

To edit connector settings:

1. In the Web Interface toolbar, select **Configuration**.
2. In the System Configuration dialog, select **Migrate Connector**.
3. Click **Edit**.
4. In the Edit Connector Settings dialog, specify appropriate option settings for the following:
   - General Settings
   - Migrate Server Settings
   - Customer-Provided Scripts
5. Click **Save**.
6. Click **Close** to exit the System Configuration dialog.
7. For each PlateSpin Migrate Connector server that is connected to your PTM Server, log in to the Connector server as the `root` user and restart the Connector service.

Assigning a Migrate Connector Instance

You can assign a Migrate Connector to the PlateSpin Transformation Manager server, or to a specific project.

The PlateSpin Transformation Manager Appliance includes an instance of the PlateSpin Migrate Connector that is automatically installed and configured to work with the Transformation Manager server on the Appliance. It works with all projects by default. After you set up projects in Transformation Manager, you can add a `ptm_project_id` to the `/opt/microfocus/migrate_connector/config/settings.cfg` file on the Appliance to assign the Connector instance to a specific project.

You must manually configure each PlateSpin Migrate Connector instance that you deploy on your own SUSE Linux Enterprise Server 11. SP4 server.

To configure the Connector instance for a PTM server or for a specific project:

1. Log in to the SLES 11 SP4 server with a `root` user account.
2. In a text editor, open the `/opt/microfocus/migrate_connector/config/settings.cfg` file.
3. Configure the values in the top section called `[PlateSpin Transformation Manager Server]`.

Option	Description
`ptm_host`	Specify the hostname or IP address of the Transformation Manager server. For example, the Migrate Connector instance on the Appliance is `localhost`. `ptm_host=localhost`

Option	Description
`ptm_port`	Specify the port used for the Transformation Manager server. The default port for HTTPS is 8183. The default port for HTTP is 8182. `ptm_port=8183`
`ptm_ssl`	Specify a value of `true` to use SSL to connect to the Transformation Manager server. Valid values are `true` and `false`. `ptm_ssl=true`
`verify_ptm_ssl_cert`	Specify whether you want to require the certificate to be validated for connections to the PlateSpin Transformation Manager server. Valid values are `false` or `true`. The default is to disable validation (false). `verify_ptm_ssl_cert=false`
`ptm_username`	Specify the email address of a valid user account on your PlateSpin Transformation Manager server that has been assigned a System Administrator role. `ptm_username=john.doe@example.com` **NOTE:** We recommend that you create a dedicated user account in PlateSpin Transformation Manager for the connector to use. We recommend System Administrator privileges for the account in order to allow Migrate Connector to execute on any project. The user account must have at least Project Architect level privileges for the projects on which the Connector should execute.
`ptm_password`	Specify the password for the user account. `ptm_password=yourpassword`
`ptm_project_id`	(Optional) Specify the numeric project ID for the project in Transformation Manager to restrict the Connector to a single project instead of making it available for all projects. The Connector processes only events and actions within the assigned project. `ptm_project_id=1234` To discover the numeric ID associated with a project in PlateSpin Transformation Manager: 1. In the Web Interface, go to **Planning** > **Projects**. 2. Select the project, then click **Edit**. 3. In the Edit Project dialog, mouse over the project name in the dialog title area. A tooltip displays the numeric ID of the project.

4 Save the file, then exit the text editor.

5 Start or restart PlateSpin Migrate Connector. In the terminal console, enter

```
rcps_migrate_connector restart
```

You can also restart the Connector from the Appliance Management UI.

6 Configuring Operating Systems

The System Administrator manages a list of available Operating System Types (OS Types) for the product in the System Configuration settings. The OS Types are available for use by Project Managers and Project Architects as they import original workloads or configure the proposed workloads.

NOTE: Only the System Administrator can create, edit, and delete OS Types. If you need additional OS Types for your transformation projects, contact the System Administrator.

About Operating System Types

Each OS Type uniquely represents a distribution of an operating system. The OS objects you create can be as general or as specific as necessary to meet your needs.

Name

Specify a textual name for the operating system that is unique in your PlateSpin Transformation Manager environment.

Description

Specify a brief description of the operating system.

Family

Select the appropriate operating system family from the following available options:

CentOS
Citrix
Linux
NetWare/OES
Other
Red Hat Linux
Solaris
SUSE Linux
Ubuntu
Unknown
VMware ESX
Windows

Architecture

Select the appropriate processor architecture from the following available options:

x32 (32 bit)
x64 (64 bit)

Viewing the List of Operating System Types

The System Administrator, Project Manager, and Project Architect can view the master Operating Systems list for the product. The OS Type values are available for all transformation projects.

To view a list of OS Types:

1. In the Web Interface toolbar, select **Configuration**.
2. In the System Configuration dialog, select **Operating Systems**.
3. Scroll to view the list of OS Types.
4. When you are done, click **Close** to exit the System Configuration dialog.

Creating an Operating System Type

The System Administrator can create new operating system components to make them available for all transformation projects.

To create an OS Type:

1. In the Web Interface toolbar, select **Configuration**.
2. In the System Configuration dialog, select **Operating Systems**.
3. Click **Create**.
4. In the Create Operating System dialog, specify the following information:
   - Name
   - Description
   - Family
   - Architecture
5. Click **Save** to create the operating system component.
6. Click **Close** to exit the System Configuration dialog.

Editing an Operating System Type

The System Administrator can modify any operating system component, including the predefined OS Types. The changes apply throughout the product for all transformation projects.

To edit an OS Type:

1. In the Web Interface toolbar, select **Configuration**.
2. In the System Configuration dialog, select **Operating Systems**.
3. Click **Edit**.

4 In the Edit Operating System dialog, specify the following information:
   - Name
   - Description
   - Family
   - Architecture
5 Click **Save** to update the operating system component.
6 Click **Close** to exit the System Configuration dialog.

Deleting an Operating System Type

The System Administrator can delete any operating system from the list of available Operating Systems. The deletion removes the OS setting for all transformation projects.

NOTE: Deleting an OS can affect the readiness of planned workloads that have this OS Type assigned to them.

To delete an OS Type:

1 In the Web Interface toolbar, select **Configuration**.
2 In the System Configuration dialog, select **Operating Systems**.
3 Select the appropriate OS Type from the list of operating systems.
4 Click **Delete**, then click **Yes** to confirm the deletion.
5 Click **Close** to exit the System Configuration dialog.

II Users

Access to PlateSpin Transformation Manager requires a user account. Through the user account, a user receives permissions to perform tasks for one or more assigned transformation projects. The default System Administrator user and other users assigned to the System Administrator role have the permissions necessary to create, manage, and delete users, groups, and organizations. This section describes common user management tasks.

7 Overview of PlateSpin User Management

For PlateSpin Transformation Manager, a user is any individual who can access Transformation Manager to plan, monitor, or execute transformation projects. Transformation Manager creates a user account during the installation process, and assigns the user to the System Administrator role. This default user is initially responsible for creating user accounts and assigning roles to them, as well as creating organizations and groups.

Transformation Manager stores user account information and authenticates the user to allow access. Access controls govern the information users can see and the actions they can perform.

System Users and Organization Users

As the default System Administrator user, you must add user accounts and assign them to project roles to enable other users to manage or view project information. When you create a user or group, you can set the user's scope at one of two levels:

- **System:** System users and groups have only the privileges associated with their assigned roles. You can assign system users to the following:
  - System Administrator role (members of the Administrators group)
  - Project Manager
  - Project Architect
  - Migration Specialist
  - Dashboard Viewer role
  - System groups

  You can also assign system groups to the roles, and then manage group membership to make management easier.
- **Organization:** Organization users and groups have only the privileges associated with their assigned roles. You can assign organization users to the following:
  - Dashboard Viewer role

Roles

You can assign users or groups to roles that let them plan, monitor, and execute transformation projects. Transformation Manager provides five roles: System Administrator, Project Manager, Project Architect, Migration Specialist, and Dashboard Viewer. Each role carries its own set of responsibilities in the PTM environment. For detailed information about the permissions for each role, see Appendix A, "Roles and Permissions," on page 81.

Roles can be assigned directly or inherited. Inherited roles can be set for system users or groups at the System, Organization, Project, Wave, or Batch level. Inherited roles can be set for organization users or groups at the Organization, Project, Wave, or Batch level. The inherited roles apply across all components in that level for existing and new components, as illustrated in Figure 7-1. For example, if you assign the system user account for John as the Project Manager for an organization, the organization's existing and new projects automatically inherit the setting.

Figure 7-1 *Scope of Permissions for Inherited Roles*

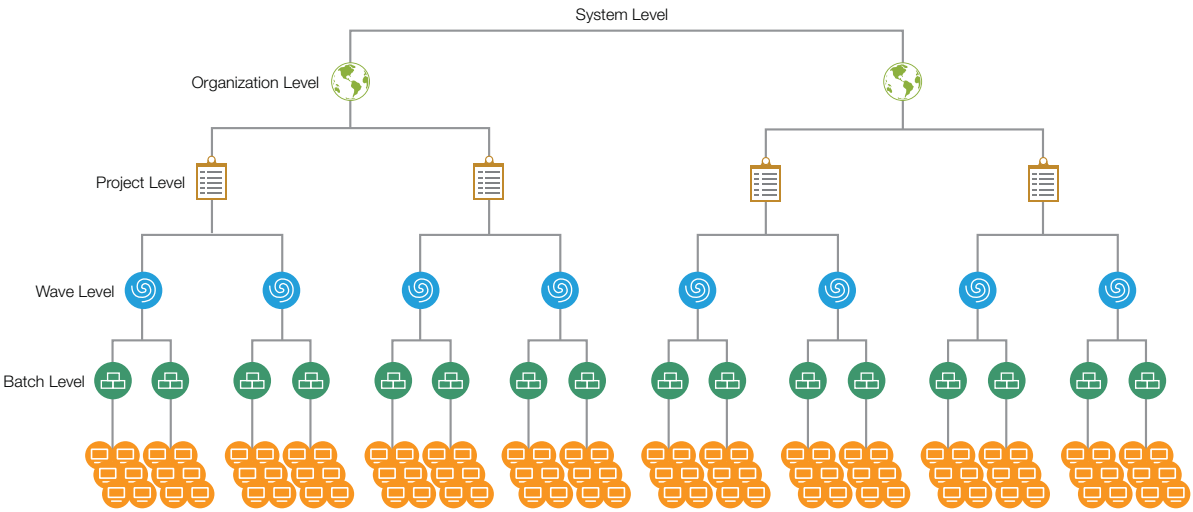

- System Administrator Role
- Project Manager Role
- Project Architect Role
- Migration Specialist Role
- Dashboard Viewer Role

System Administrator Role

The System Administrator role has full privileges in Transformation Manager. The initial user account that you create during the installation automatically has this role. You can add system users or system groups to the Administrators group to assign this role. The System Administrator typically performs the following tasks:

- Configures, maintains, and monitors the health of the PTM Server.
- Has all privileges throughout the product.
- Has exclusive privileges to perform the following tasks:
  - Create and delete organizations.
  - Create and delete projects.
  - Create and delete Operating System types.
  - Assign users and groups to roles at the Organization level.
  - Assign users and groups to the Project Manager role.
- Can perform all tasks for every role in any project.

Project Manager Role

The Project Manager role can be a user or group. For an assigned project, this role has the permissions necessary to perform the following tasks:

- Manages the project.
- Creates and deletes users.
- Creates and deletes non-administrator groups, and assigns members to them.
- Assigns users or groups to the Project Architect, Migration Specialist, and Dashboard Viewer roles.
- Tracks project progress and core statistics, using the dashboard.
- Performs any of the Project Architect tasks.
    - Creates and deletes waves, batches, and applications.
    - Bulk imports project workloads.
    - Creates and deletes resources.
    - Defines proposed workloads.
    - Submits workloads that are ready for transformation, or withdraws them if transformation changes are needed.

Project Architect Role

The Project Architect role can be user or group. For an assigned project, this role has the permissions necessary to perform the following tasks:

- Views all information for the project.
- Creates and deletes waves, batches, and applications.
- Assigns users or groups to the Migration Specialist role for waves and batches.
- Bulk imports project workloads.
- Creates and deletes resources.
- Defines proposed workloads.
- Submits workloads that are ready for transformation, or withdraws them if transformation changes are needed.
- Tracks project progress and core statistics, using the dashboard.
- Can execute the individual migrations, according to the project plan.

Migration Specialist Role

The Migration Specialist role can be a user or group. For an assigned project, wave or batch, this role has the permissions necessary to perform the following tasks:

- Views information for the project's waves, batches, and workloads.
- Views information for the project's resources.
- Executes the individual migrations, according to the project plan.
- Tracks project progress and core statistics, using the dashboard.

Dashboard Viewer Role

The Dashboard Viewer role can be a user or group. The Dashboard Viewer role has the permissions necessary to view the dashboard information only for an assigned organization, project, wave, or batch. Inherited permissions apply to this role in the child containers if you assign this role at the system, organization, project, wave, or batch level.

Example: Digital Airlines Users

Each project typically has one manager. In this example, the Digital Airlines organization has multiple projects. They want three users to share the project manager role across all of their projects. You want to manage their role assignments in a group instead of individually. The group will receive access rights at the Organization level so that its members can manage any of the organization's projects.

As a System Administrator user (a member of the Administrators group), you perform the following tasks, as shown in Figure 7-2.:

1. Create an organization account for Digital Airlines.
2. Create system user accounts for workers Alice, Ben, and Jacob.
3. Create a system group `DigAir PM`.
4. Assign the three system users to this group.
5. Assign the `DigAir PM` group to the role of Project Manager for the `Digital Airlines` organization at the Organization level.
6. Create projects for Digital Airlines.
7. Each project for the organization automatically inherits the group in the Project Manager role.

   With this configuration, any member of the `DigAir PM` group can manage any project for the `Digital Airlines` organization.

Figure 7-2 *Assigning a Group to the Project Manager Role*

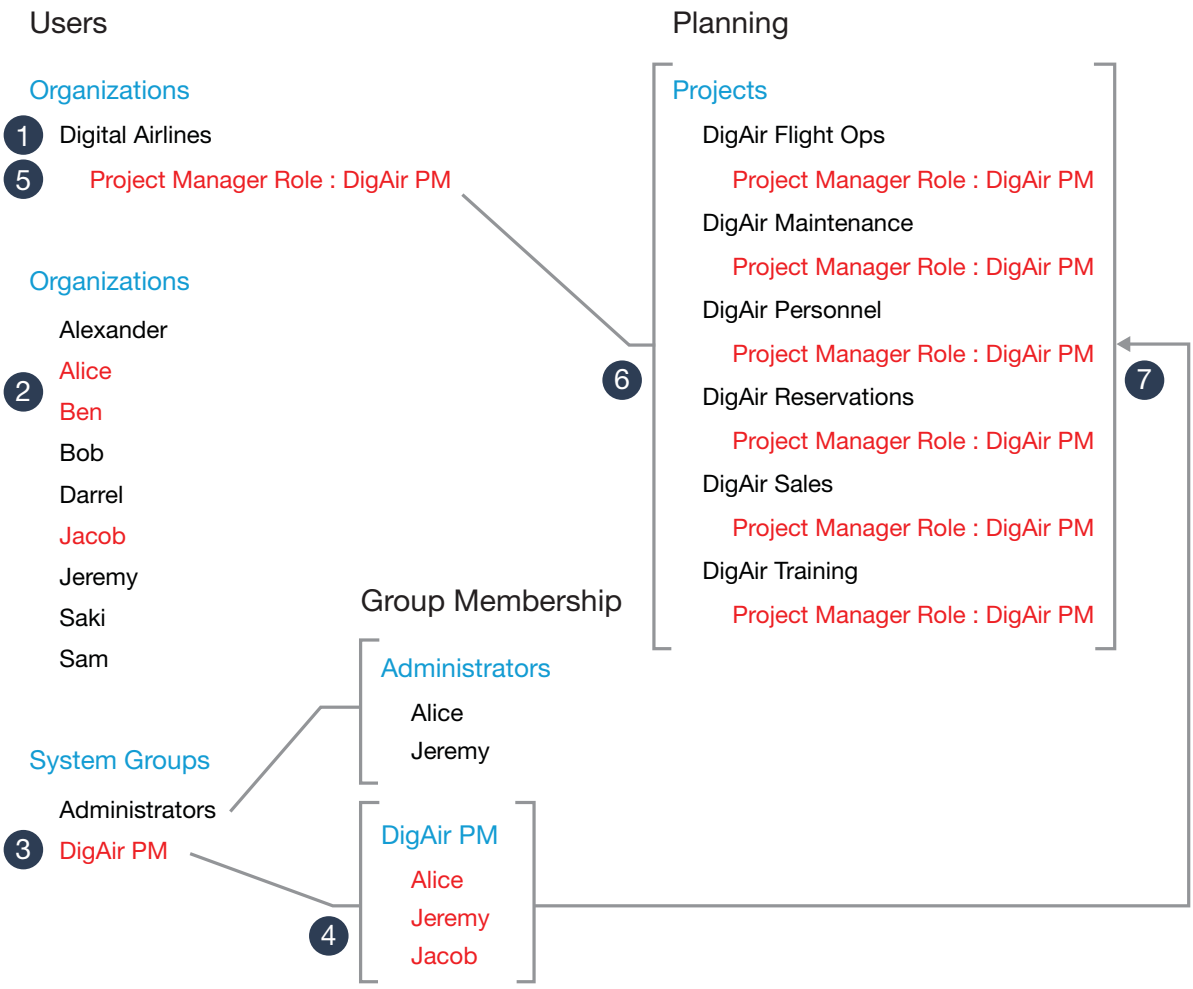

You can assign user permissions directly, or indirectly through group membership. If you assign a user to multiple roles, the user's permissions in the Web Interface are cumulative.

For example, Alice is a system user who is currently assigned to the `DigAir PM` group. You also want her to perform the System Administrator role for the Web Interface. You can add Alice as a member of the `Administrators` group to grant her the associated permissions for the System Administrator role. See the `Administrators` group membership in Figure 7-2.

The Project Manager role provides a user the ability to create, modify, and delete users and groups, without granting the additional authority for tasks associated with the System Administrator role. For example, you can assign Alice the Project Manager role. After Alice creates users and groups, the project managers for the Digital Airlines projects (that is, member in the `DigAir PM` group) can assign subordinate project roles to them.

8 Managing Organizations

An organization represents the logical association to some organizational entity in your business environment. For example, if you are data center provider, each organization might represent a company. If you are an enterprise IT department, an organization might represent a company site, a business unit or department, or a cost center within your company.

Users Who Can Perform These Tasks: System Administrator or System Administrator role

About Organizations

Transformation Manager allows you to track the following information for the organization accounts:

Name: Specify a name for the organization that is unique in the PTM system. The name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Description: (Optional) Specify a brief description of the organization. This text appears by default in the Organizations list.

Image: (Optional) Upload an image to represent the organization.

Associations: The following tabs allow you to view or manage the components associated with the organization. Some options might not be available, depending on the user's assigned role.

- **Planning**
  - Projects (Create, Edit, View, Delete) - See Chapter 12, "Managing Projects," on page 95.
- **Resources**
  - Credentials (Create, Edit, View, Delete) - See Section 21, "Managing Credentials Resources," on page 193.
  - Hosts - See Section 22, "Managing Host Resources," on page 199.
    - Source (View)
    - Target (Create, Edit, View, Delete)
  - Networks - See Section 24, "Managing Network Resources," on page 211.
    - Source (View)
    - Target (Create, Edit, View, Delete)

- Datastores - See Section 25, “Managing Datastore Resources,” on page 215.
  - Source (View)
  - Target (Create, Edit, View, Delete)
- Environments (Create, Edit, View, Delete) - See Section 27, “Managing Environment Resources,” on page 223.
- Applications (Create, Edit, View, Delete) - See Chapter 15, “Managing Applications,” on page 109.
- Migration Servers (Create, Edit, View, Delete) - See Section 23, “Managing Migration Server Resources,” on page 207.
- Resource Pools (Create, Edit, View, Delete) - See Section 26, “Managing Resource Pool Resources,” on page 219.

- **Users**

  Project role assignments at the project level are automatically inherited by their child components.
  - Dashboard Viewer (View, Add, Remove) - See “Dashboard Viewer Role” on page 66.
  - Migration Specialist (View, Add, Remove) - See “Migration Specialist Role” on page 65.
  - Project Architect (View, Add, Remove) - See “Project Architect Role” on page 65.
  - Project Manager (View, Add, Remove) - See “Project Manager Role” on page 65.

Creating an Organization

1. In the Web Interface toolbar, select **Users**.
2. Select the **Organizations** tab.
3. Click **Create**.
4. In the Create Organization dialog, specify a name for the organization that is unique within the PTM system.
5. (Optional) Specify a brief textual description of the organization.
6. (Optional) Upload an image to represent the organization. Mouse-over the image pane, click **Upload New Image**, browse to locate and select the image that you want to use for the organization, then click **Open**.
7. Click **Save** to create the organization instance and enable the **Associations** pane.
8. (Optional) In the **Associations** pane, define the associations for the organization.

   Planning
   - Projects (Create, Edit, View, Delete) - See Chapter 12, “Managing Projects,” on page 95.

   Resources
   - Credentials (Create, Edit, View, Delete) - See Section 21, “Managing Credentials Resources,” on page 193.
   - Hosts - See Section 22, “Managing Host Resources,” on page 199.
     - Source (View)
     - Target (Create, Edit, View, Delete)

- Networks - See Section 24, “Managing Network Resources,” on page 211.
  - Source (View)
  - Target (Create, Edit, View, Delete)
- Datastores - See Section 25, “Managing Datastore Resources,” on page 215.
  - Source (View)
  - Target (Create, Edit, View, Delete)
- Environments (Create, Edit, View, Delete) - See Section 27, “Managing Environment Resources,” on page 223.
- Applications (Create, Edit, View, Delete) - See Chapter 15, “Managing Applications,” on page 109.
- Migration Servers (Create, Edit, View, Delete) - See Section 23, “Managing Migration Server Resources,” on page 207.
- Resource Pools (Create, Edit, View, Delete) - See Section 26, “Managing Resource Pool Resources,” on page 219.

Users

- Dashboard Viewer (View, Add, Remove) - See “Dashboard Viewer Role” on page 66.
- Migration Specialist (View, Add, Remove) - See “Migration Specialist Role” on page 65.
- Project Architect (View, Add, Remove) - See “Project Architect Role” on page 65.
- Project Manager (View, Add, Remove) - See “Project Manager Role” on page 65.

9 Click **Save**.

10 Click **Close**.

Uploading an Organization Logo

To add a logo for an organization:

1 In the Web Interface toolbar, select **Users**.

2 Select the **Organizations** tab.

3 Select the organization, then click **Edit**.

4 Upload an image to represent the organization. Mouse-over the image pane, click **Upload New Image**, browse to locate and select the image that you want to use for the organization, then click **Open**.

5 Click **Save**.

6 Click **Close**.

Editing an Organization

1 In the Web Interface toolbar, select **Users**.

2 Select the **Organizations** tab.

3 In the Organizations list, double-click the organization, or select the organization and click **Edit**.

4 (Optional) Specify a brief textual description of the organization.

5 (Optional) Upload an image to represent the organization. Mouse-over the image pane, click **Upload New Image**, browse to locate and select the image that you want to use for the organization, then click **Open**.

6 In the **Associations** pane, view or modify the Associations information.

- Projects - See Chapter 12, “Managing Projects,” on page 95.
- Users
  - Dashboard Viewer - See “Dashboard Viewer Role” on page 66.
  - Migration Specialist - See “Migration Specialist Role” on page 65.
  - Project Architect - See “Project Architect Role” on page 65.
  - Project Manager - See “Project Manager Role” on page 65.
- Credentials - See Section 21, “Managing Credentials Resources,” on page 193.
- Hosts - See Section 22, “Managing Host Resources,” on page 199.
- Networks - See Section 24, “Managing Network Resources,” on page 211.
- Datastores - See Section 25, “Managing Datastore Resources,” on page 215.
- Environments - See Section 27, “Managing Environment Resources,” on page 223.
- Applications - See Chapter 15, “Managing Applications,” on page 109.
- Migration Servers - See Section 23, “Managing Migration Server Resources,” on page 207.
- Resource Pools - See Section 26, “Managing Resource Pool Resources,” on page 219.

7 If you modified information, click **Save**.

8 Click **Close**.

Removing an Organization

When you remove an organization, the associations set up for the organization are automatically deleted.

1 In the Web Interface toolbar, select **Users**.

2 Select the **Organizations** tab.

3 In the Organizations list, select the organization and click **Delete**.

4 Click **Yes** to confirm the deletion, or click **No** to keep the organization.

9 Managing Users

Transformation Manager creates a user account during the installation process, and assigns this user to the System Administrator role. This default System Administrator initially creates user accounts, as well as organizations and groups. You cannot delete the default System Administrator user.

Users Who Can Perform These Tasks: System Administrator, or System Administrator role

About Users

PlateSpin Transformation Manager allows you to track the following information for users:

Full Name: Specify a first and last name for the user. Names can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Email Address: Specify an email address for the user that is unique in the PTM system.

Phone Number: (Optional) Specify a contact phone number for the user.

Password: Specify a password for the user account in the PTM system. Type it again to confirm. The password length must be a minimum of 5 characters. After you set up a role for the user account, send the initial password to the user. The user logs in with the initial password, and then sets their preferred password.

NOTE: Passwords are local to the product. They are stored securely in the PTM database.

Scope: Specify whether the user's privileges apply at the system or organization level in the PTM system. After the user's scope is set, it cannot be modified. Organization scope is used only for users in a Dashboard Viewer role for their organization's projects.

Membership and Access: The following tabs allow you to view or manage the user's access to information in the transformation environment. Some options might not be available, depending on the user's assigned role.

- All Roles
- System (Add, Remove)
- Organization (Add, Remove)
- Project (Add, Remove)
- Wave (Add, Remove)

- Batch (Add, Remove)
- Group Membership (View, Add, Remove)

Viewing Users

1. In the Web Interface toolbar, select **Users**.
2. Select the **Users** tab.
3. In the Users list, double-click the user, or select the user and click **View**.
4. In the View User dialog, view the User and the Membership and Access information.
5. Click **Close**.

Creating a User

1. In the Web Interface toolbar, select **Users**.
2. Select the **Users** tab.
3. Click **Create**.
4. In the User pane of the Create User dialog, specify the following information for the user account:
   - First and last name
   - Email address
   - Phone number (optional)
   - Password and Confirm Password
   - Scope (System or Organization)
5. Click **Save** to create the user account instance and enable the **Membership and Access** pane.
6. In the **Membership and Access** pane, define the permissions and roles assignments for the user account.
7. Click **Save**.
8. Click **Close**.

Editing a User

1. In the Web Interface toolbar, select **Users**.
2. Select the **Users** tab.
3. In the Users list, double-click the user, or select the user and click **Edit**.
4. In the Edit User dialog, view or modify the User information or the Membership and Access information.
5. If you modified information, click **Save**.
6. Click **Close**.

Changing a User Password

A System Administrator can change the password for any user account. A non-administrator user can change the password associated with the user account assigned to them.

The user logs in with the initial password sent to them by the System Administrator, and then sets a preferred password. The new password is stored, and takes effect the next time the user logs in for a session.

1. In the Web Interface toolbar, select **Users**.
2. Select the **Users** tab.
3. In the Users list, double-click the user name, or select the user name and click **Edit**.
4. In the Edit User dialog, type the preferred password, then type it again to confirm the change.
5. Click **Save**.
6. Click **Close**.

Removing a User

When you remove a user, the roles and permissions set up for the user are automatically removed.

1. In the Web Interface toolbar, select **Users**.
2. Select the **Users** tab.
3. In the Users list, select the user and click **Delete**.
4. Click **Yes** to confirm the deletion, or click **No** to keep the user.

10 Managing Groups

You can also associate users with groups to more efficiently manage access. Members of a group inherit the access controls assigned to the group.

About Groups

Transformation Manager allows you to track the following information for groups:

Full Name: Specify a name for the group account that is unique in the PTM system. Names can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Email Address: (Optional) Specify an email address for the group that is unique in the PTM system.

Scope: Specify whether the user’s privileges apply at the system or organization level in the PTM system. After the user’s scope is set, it cannot be modified. Organization scope is used only for groups in a Dashboard Viewer role for their organization’s projects.

Membership and Access: The following tabs allow you to view or manage the group’s access to information in the transformation environment. Some options might not be available, depending on the group’s assigned role.

- All Roles
- System (Add, Remove)
- Organization
  - Migration Specialist (Add, Remove)
  - Project Architect (Add, Remove)
  - Project Manager (Add, Remove)
- Project
  - Migration Specialist (Add, Remove)
  - Project Architect (Add, Remove)
  - Project Manager (Add, Remove)
- Wave
  - Migration Specialist (Add, Remove)
  - Project Architect (Add, Remove)

- Batch
  - Migration Specialist (Add, Remove)
  - Project Architect (Add, Remove)
- Members (View, Add, Remove)

Viewing Groups

1. In the Web Interface toolbar, select **Users**.
2. Select the **Group** tab.
3. In the Group list, double-click the group, or select the group and click **View**.
4. In the View Group dialog, view the Group and the Membership and Access information.
5. Click **Close**.

Creating a Group

1. In the Web Interface toolbar, select **Users**.
2. Select the **Groups** tab.
3. Click **Create**.
4. In the Group pane of the Create Group dialog, specify the following information for the group account:
   - Full Name
   - Email address
   - Scope (System or Organization)
5. Click **Save** to create the group account instance and enable the **Membership and Access** pane.
6. In the **Membership and Access** pane, define the roles and member assignments for the group account.
7. Click **Save**.
8. Click **Close**.

Editing a Group

1. In the Web Interface toolbar, select **Users**.
2. Select the **Group** tab.
3. In the Group list, double-click the group, or select the group and click **Edit**.
4. In the Edit Group dialog, view or modify the Group and the Membership and Access information.
5. If you modified information, click **Save**.
6. Click **Close**.

Removing a Group

When you remove a group, the roles and membership set up for the group are automatically removed.

1. In the Web Interface toolbar, select **Users**.
2. Select the **Groups** tab.
3. In the Groups list, select the group and click **Delete**.
4. Click **Yes** to confirm the deletion, or click **No** to keep the group.

A Roles and Permissions

PlateSpin Transformation Manager provides role-based access control to restrict access in the Web Interface to authorized users. It associates permissions with roles rather than with individual users and groups. A role grants a user or group the authority to perform specific actions in the Web Interface. Permissions determine the various job functions that the role can perform. Roles in the Web Interface include the following:

- System Administrator
- Project Manager
- Project Architect
- Migration Specialist
- Dashboard Viewer

Users and groups have no permissions in Transformation Manager except those you grant through the roles you assign them. Because members of a group automatically inherit the role assignments for the group, you can assign appropriate roles to groups to simplify the management of user authorizations. You can easily move users in and out of various groups to assign roles to them.

The scope of responsibility for each project role depends on the level of the role assignment: system, organization, project, wave, or batch. Only system users can be assigned roles at the system level or to multiple organizations.

The roles and permissions matrixes identify the various actions and access control permissions that users and groups have based on their assigned roles.

User Roles

Table A-1 *Legend for the Roles and Permissions Matrixes*

Column Heading	User or Role	Description
NA	Network Administrator	An IT administrator who has sufficient permissions and authority in the data center to install and configure the PlateSpin Transformation Manager Appliance. After the initial setup, this user (or another designated IT member) monitors and manages the appliance by logging in to the Appliance Management tool with either of the following appliance user identities: vaadmin `root`
SA	System Administrator role	Any Transformation Manager system user who is a member of the `Administrators` group.
PM	Project Manager role	A Transformation Manager user who is assigned the Project Manager role.

Column Heading	User or Role	Description
PA	Project Architect role	A Transformation Manager user who is assigned to the Project Architect role.
MS	Migration Specialist role	A Transformation Manager user who is assigned to the Migration Specialist role.
DV	Dashboard Viewer role	A Transformation Manager user who is assigned to the Dashboard Viewer role.

Annotations

N No

Y Yes

Y^U Restricted to user's own user account.

Y^A Restricted to the projects assigned to the user.

Y^N Restricted to non-Administrator groups.

Roles and Permissions Matrixes

Table A-2 *Roles and Permissions for the Appliance and Application*

Tasks and Actions	NA	SA	PM	PA	MS	DV
Appliance Installation and Setup						
Download software	Y	N	N	N	N	N
Prepare VM Environment						
Hypervisor Host	Y	N	N	N	N	N
Network	Y	N	N	N	N	N
External Storage	Y	N	N	N	N	N
Install VM	Y	N	N	N	N	N
Configure VM (vaadmin or `root`)						
Modify VM Settings	Y	N	N	N	N	N
Monitor VM (Ganglia metrics and reports)	Y	N	N	N	N	N
Add System Administrator User	Y	N	N	N	N	N
Change Password for System Administrator User	Y	N	N	N	N	N
Upgrade VM	Y	N	N	N	N	N
License Key/Blocks (Customer Center account owner)	Y	N	N	N	N	N

Tasks and Actions	NA	SA	PM	PA	MS	DV
Manage PlateSpin Server (vaadmin or `root`)						
Initial Configuration	Y	N	N	N	N	N
Web Server	Y	N	N	N	N	N
Add System Administrator User	Y	N	N	N	N	N
Change Password for System Administrator User	Y	N	N	N	N	N
Web Interface Session Timeout	Y	N	N	N	N	N
Web Interface Configuration						
Licenses						
Add license key	N	Y	N	N	N	N
Edit license key	N	Y	N	N	N	N
Remove license key	N	Y	N	N	N	N
View active license key and unused blocks	N	Y	N	N	N	N
Migrate Connector						
General Settings	N	Y	N	N	N	N
Migrate Server Settings	N	Y	N	N	N	N
Customer-Provided Scripts	N	Y	N	N	N	N
View	N	Y	N	N	N	N
Operating Systems						
Create	N	Y	N	N	N	N
Edit	N	Y	N	N	N	N
Delete	N	Y	N	N	N	N
View	N	Y	N	N	N	N

Table A-3 *Roles and Permissions for User Management*

Tasks and Actions	SA	PM	PA	MS	DV
Users					
Organizations					
Create	Y	N	N	N	N
Edit	Y	N	N	N	N
Delete	Y	N	N	N	N
View	Y	Y[U]	Y[U]	N	N

Tasks and Actions	SA	PM	PA	MS	DV
Users					
Create	Y	Y	N	N	N
Edit	Y	Y[U]	Y[U]	Y[U]	Y[U]
Delete (except the last member of the Administrators group)	Y	Y	N	N	N
View	Y	Y	Y	Y	N
Change password	Y	Y[U]	Y[U]	Y[U]	Y[U]
Groups					
Create	Y	Y[N]	N	N	N
Edit	Y	Y[N]	N	N	N
Delete (except the Administrators group)	Y	Y[N]	N	N	N
View	Y	Y[N]	Y[N]	N	N
Role-Based Access Control (who can assign roles)					
System Administrator Role (Administrators Group)					
Add members	Y	N	N	N	N
Remove members (except the last member)	Y	N	N	N	N
View members	Y	N	N	N	N
Project Manager Role					
Add	Y	N	N	N	N
Remove	Y	N	N	N	N
Project Architect Role					
Add	Y	Y	N	N	N
Remove	Y	Y	N	N	N
Migration Specialist Role					
Add	Y	Y	Y	N	N
Remove	Y	Y	Y	N	N
Dashboard Viewer Role					
Add	Y	Y	N	N	N
Remove	Y	Y	N	N	N

Table A-4 *Roles and Permissions for Transformation Management*

Tasks and Actions	SA	PM	PA	MS	DV
Dashboard					
View Warning and Error Statistics	Y	N	N	N	N
View Transformation Statistics	Y	Y[A]	Y[A]	Y[A]	Y[A]
Planning					
Projects					
View	Y	Y[A]	Y[A]	Y[A]	Y[A]
Create	Y	Y[A]	N	N	N
Edit	Y	Y[A]	Y[A]	N	N
Delete	Y	Y[A]	N	N	N
Waves					
View	Y	Y[A]	Y[A]	Y[A]	Y[A]
Create	Y	Y[A]	Y[A]	N	N
Edit	Y	Y[A]	Y[A]	N	N
Delete	Y	Y[A]	Y[A]	N	N
Batches					
View	Y	Y[A]	Y[A]	Y[A]	Y[A]
Create	Y	Y[A]	Y[A]	N	N
Edit	Y	Y[A]	Y[A]	N	N
Delete	Y	Y[A]	Y[A]	N	N
Applications					
View	Y	Y[A]	Y[A]	Y[A]	N
Create	Y	Y[A]	Y[A]	N	N
Edit	Y	Y[A]	Y[A]	N	N
Delete	Y	Y[A]	Y[A]	N	N

Tasks and Actions	SA	PM	PA	MS	DV
Workloads					
View	Y	Y[A]	Y[A]	N	N
Import (Single or Range)	Y	Y[A]	Y[A]	N	N
Spreadsheet Import	Y	Y[A]	Y[A]	N	N
Retry (Migration action)	Y	Y[A]	Y[A]	N	N
Rediscover (retry auto-discovery if initial discovery failed)	Y	Y[A]	Y[A]	N	N
Bulk Edit	Y	Y[A]	Y[A]	N	N
Edit	Y	Y[A]	Y[A]	N	N
Bulk Status Change	Y	Y[A]	Y[A]	N	N
Delete	Y	Y[A]	Y[A]	N	N
Transformations					
View	Y	Y[A]	Y[A]	Y[A]	N
View Statistics	Y	Y[A]	Y[A]	Y[A]	Y[A]
Submit	Y	Y[A]	Y[A]	N	N
Withdraw	Y	Y[A]	Y[A]	N	N
Reopen	Y	Y[A]	Y[A]	N	N
Bulk Status Change	Y	Y[A]	Y[A]	Y[A]	N
Execute Transformation	Y	N	Y[A]	Y[A]	N
Resources					
Credentials					
View	Y	Y[A]	Y[A]	Y[A]	N
Create	Y	Y[A]	Y[A]	N	N
Edit	Y	Y[A]	Y[A]	N	N
Delete	Y	Y[A]	Y[A]	N	N
Hosts (Source)					
View	Y	Y[A]	Y[A]	Y[A]	N
Delete	Y	Y[A]	Y[A]	N	N

Tasks and Actions	SA	PM	PA	MS	DV
Hosts (Target)					
View	Y	Y[A]	Y[A]	Y[A]	N
Create	Y	Y[A]	Y[A]	N	N
Edit	Y	Y[A]	Y[A]	N	N
Delete	Y	Y[A]	Y[A]	N	N
Migration Servers					
View	Y	Y[A]	Y[A]	Y[A]	N
Create	Y	Y[A]	Y[A]	N	N
Edit	Y	Y[A]	Y[A]	N	N
Delete	Y	Y[A]	Y[A]	N	N
Networks (Source)					
View	Y	Y[A]	Y[A]	Y[A]	N
Delete	Y	Y[A]	Y[A]	N	N
Networks (Target)					
View	Y	Y[A]	Y[A]	Y[A]	N
Create	Y	Y[A]	Y[A]	N	N
Edit	Y	Y[A]	Y[A]	N	N
Delete	Y	Y[A]	Y[A]	N	N
Datastores (Source)					
View	Y	Y[A]	Y[A]	Y[A]	N
Delete	Y	Y[A]	Y[A]	N	N
Datastores (Target)					
View	Y	Y[A]	Y[A]	Y[A]	N
Create	Y	Y[A]	Y[A]	N	N
Edit	Y	Y[A]	Y[A]	N	N
Delete	Y	Y[A]	Y[A]	N	N

Tasks and Actions	SA	PM	PA	MS	DV
Resource Pools					
View	Y	Y[A]	Y[A]	Y[A]	N
Create	Y	Y[A]	Y[A]	N	N
Edit	Y	Y[A]	Y[A]	N	N
Delete	Y	Y[A]	Y[A]	N	N
Environments					
View	Y	Y[A]	Y[A]	Y[A]	N
Create	Y	Y[A]	Y[A]	N	N
Edit	Y	Y[A]	Y[A]	N	N
Delete	Y	Y[A]	Y[A]	N	N

III Planning

PlateSpin Transformation Manager provides the tools you need to create elements and organize them to define your transformation projects.

11 Overview of Project Planning

PlateSpin Transformation Manager provides planning tools that allow you to schedule the transformations for large projects with thousands to tens of thousands of workloads. You can schedule the transformations to occur in waves with smaller groupings of batches in each wave.

Planning Waves and Batches

Large-scale IT transformation projects typically occur over an extended period in a production environment that might span multiple locations. Future network activities and conditions can be difficult to predict and details are as yet unknown. Complex projects with massive numbers of workloads might take months or even years to complete. It might be possible to plan details only a few months in advance.

PlateSpin Transformation Manager supports rolling wave planning to accommodate the near-term and long-term planning of your transformation tasks. You can organize the project’s transformation tasks in waves, as shown in Figure 11-1. For near-term waves, you can plan the tasks in detail. You can refine the transformation plan for future waves as newer and better information becomes available.

Figure 11-1 *Wave Planning*

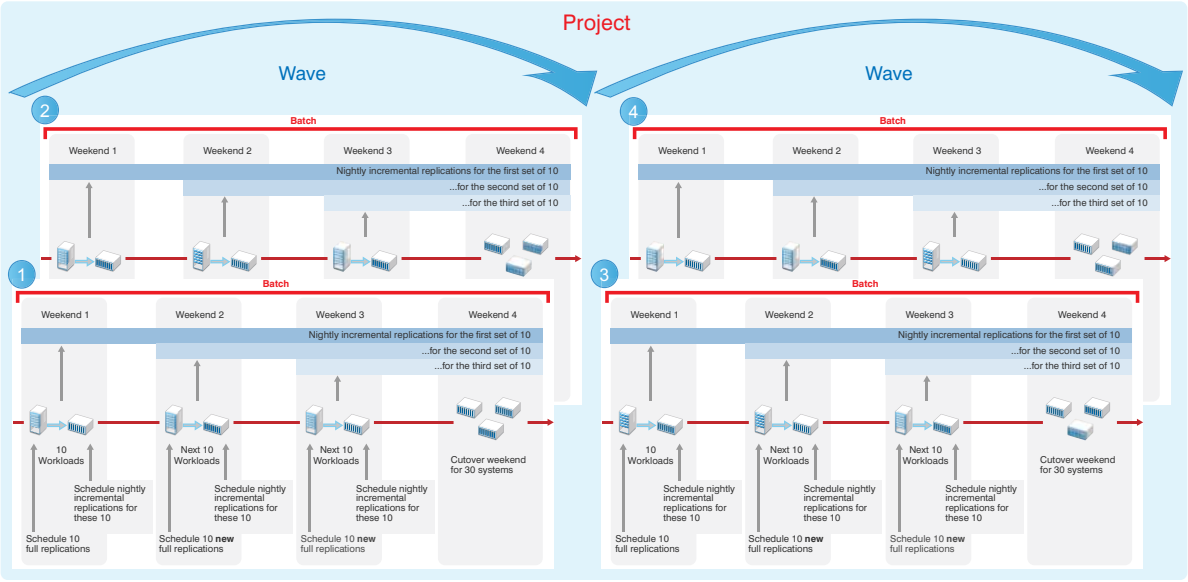

Within a wave, batches group like workloads that you want to transform together. Batches can be more easily scheduled during intervals when network resources are available. You can deliver valuable results in each batch and wave. Schedules are flexible. You can coordinate the start dates and cutover dates with stakeholders to work around planned events that are critical to the business.

Figure 11-2 *Batch Planning*

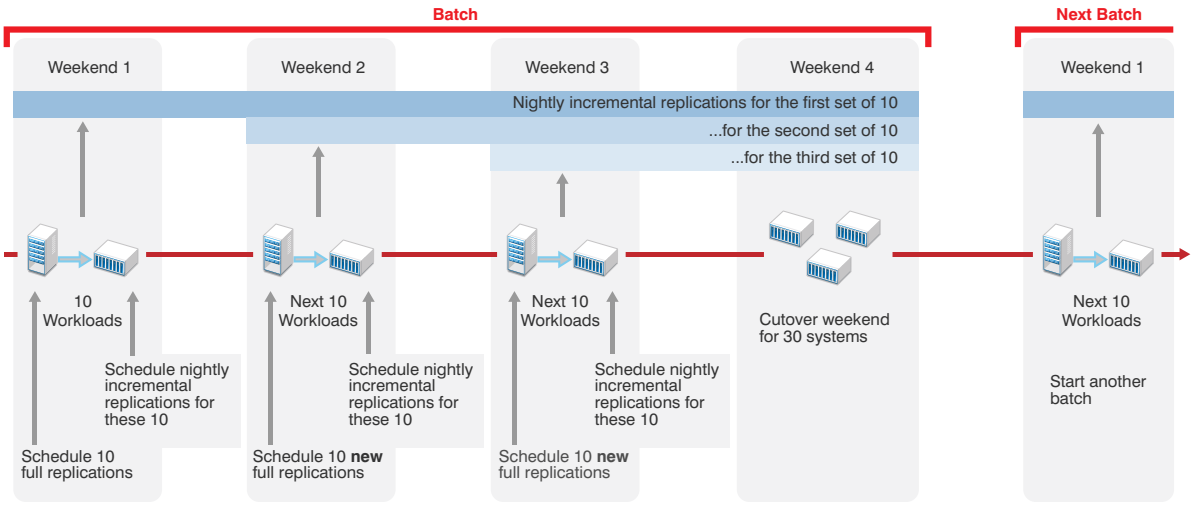

Prerequisites for Planning

When you create a project, you must associate it with a specific organization. Before you can add a project, you must create the parent organization to ensure that it is available when you create its transformation projects. Waves, batches, applications, and workloads are all child containers of their parent project.

Granting Access

PlateSpin Transformation Manager provides the following roles for managing your project:

- **Project Manager:** The Project Manager is responsible for managing all aspects of the project.
- **Project Architect:** The Project Architect is responsible for configuring transformation plans for the workloads, scheduling the transformations, and monitoring the health of transformations.
- **Migration Specialists:** The Migration Specialists are responsible for executing the transformations.
- **Dashboard Viewer:** The Dashboard Viewers are stakeholders who want to observe the progress and metrics for projects.

Transforming Workloads

Transforming workloads from their current operational mode to a future operational mode is the fundamental management goal for your transformation project. A transformation plan includes the following information:

- Dependencies between workloads
- Data about each original workload and its proposed workload

- The sequence for workload tasks
- When tasks need to be executed

Figure 11-3 shows the parent-child relationships between planning objects: Organizations, Projects, Waves, Batches, and Workloads.

Figure 11-3 *Parent-Child Relationships of Planning Objects*

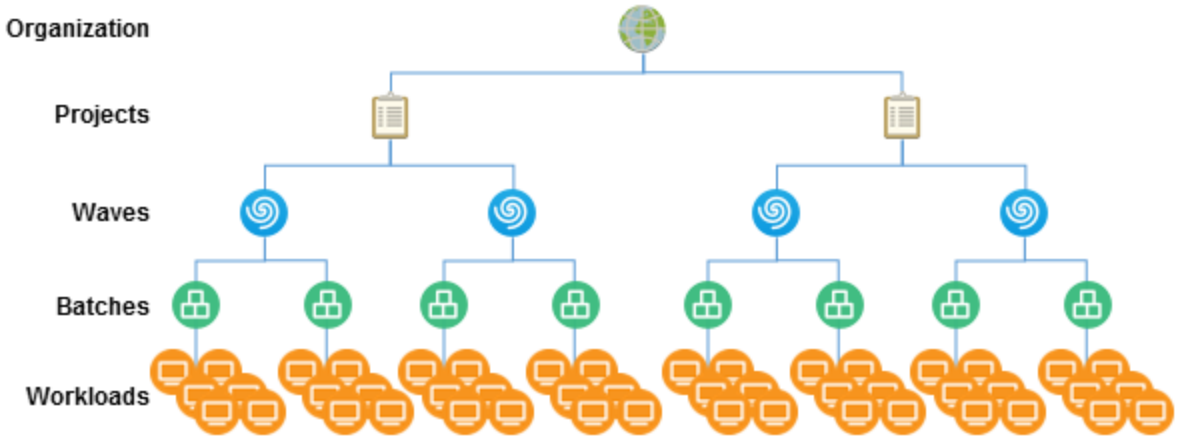

Scheduling Dates

Project dates can be set according to your business and network needs. Business factors include your project priorities, task dependencies, and the availability of resources. Network environmental factors include available bandwidth, connection speeds, and the amount of data being transformed.

You might have a specific cutover date in mind, or simply want to create a plan to get the work accomplished as efficiently and quickly as you can. External events might determine when certain tasks must be completed. IT and human resources availability also constrain your schedule.

Dates for waves and batches can be set manually or inherited from their parent objects. Dates of child objects must fall within the execution window of its parent object.

Figure 11-4 demonstrates how the dates flow down from parent to child. Manually set dates appear in a black font in fields and tables. Inherited dates appear in a gray font in fields and tables.

If you set dates manually, the date setting overrides inheritance rules. If you move a workload with a manual date from one batch to another, the manual date setting does not change, but the inherited date setting changes automatically.

Figure 11-4 shows how the manual date remains the same after you move the workloads from Batch 1 to Batch 3. If the manual date falls outside the execution window for the new batch, the date appears in a red font in fields and tables. You must re-configure the dates for workloads if their execution windows extend outside their new parent window.

Figure 11-4 *Inherited Start Date and End Date*

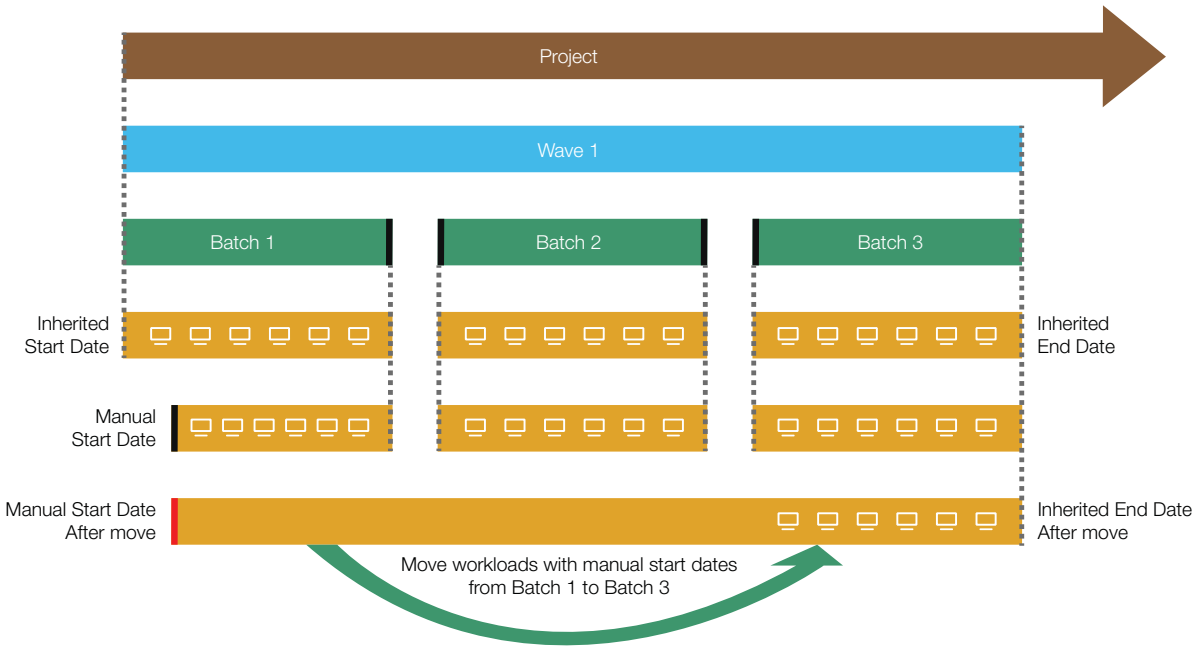

NOTE: Dates display in the format of your computer browser's Locale setting.

PlateSpin Transformation Manager displays an object's dates in the following font colors, depending on how the object obtained the setting:

Font Color for Dates	Condition
Gray	The date is inherited from its parent object.
Black	The date has been set directly on the object. The new date is automatically inherited by child objects.
Red	The date does occurs before or after the execution window set on the parent object: ◆ The object's start date occurs before the parent's start date. ◆ The object's end or cutover date occurs after the parent's end or cutover date.

12 Managing Projects

PlateSpin Transformation Manager allows you to analyze and organize information about workload transformations into projects based on your business needs.

About Projects

The Project allows you to track the following information for your transformation project:

Name: A friendly name for the project that is unique in your organization. The name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Organization: The name of the parent organization for the project.

Description: (Optional) A brief description of the project.

Start Date: The date that the transformation of workloads is planned to begin.

End Date: The date that the transformation of all workloads must be completed successfully.

Mode: Specify a mode for the project. You can switch between the modes if needed.

- **Planning:** Planning Mode allows the widest variety of supported transformation methods and planning options.

  Use Planning Mode to plan transformations that you execute manually or using third-party migration tools. You can also configure individual workloads for automation.

- **Automated:** Automated Mode allows you to plan and execute workload migrations in your PlateSpin Migration Environment. It provides additional settings that apply only to migration with PlateSpin Migrate servers, while limiting some options to those supported by automation. See “PlateSpin Migration Environment” on page 16.

  Use Automated Mode to plan transformations that you execute using automated or semi-automated options in a PlateSpin Migration Environment. You can also configure individual workloads for manual or third-party migration.

Custom Field Names: Additional information types that you want to track for a project. You can define 1 to 7 custom fields. Each custom field is independent, and can be used for any purpose. Uses might include a new attribute, a logical tag, a priority system, contact information, and so on.

For example, if you want to identify the workload with its day-to-day IT administrator, you might define the **Custom 1** field name to be **Contact** for the project. Values might be the contact person's name, user name, or email address, as appropriate for your project. You specify a workload's value for the Contact field in the Transformation dialog.

Associations: Associations define the relationship between a variety of components and the project. You can also perform the same tasks on each tab that you can on their primary tab, with the exception of the Workloads tab. For workloads, you can view information about the original and proposed workloads associated with the project.

- **Planning**
  - Waves (Create, Edit, View, Delete) - See "Managing Waves" on page 101.
  - Batches (Create, Edit, View, Move, Delete) - See "Managing Batches" on page 105.
  - Applications (Create, Edit, View, Delete) - See "Managing Applications" on page 109.
  - Workloads (Original, Proposed) - See "Configuring Workload Transformations" on page 145.
- **Resources**
  - Credentials (Create, Edit, View, Delete) - See "Managing Credentials Resources" on page 193.
  - Hosts - See "Managing Host Resources" on page 199.
    - Source (View)
    - Target (Create, Edit, View, Delete)
  - Migration Servers (Create, Edit, View, Delete) - See "Managing Migration Server Resources" on page 207.
  - Networks - See "Managing Network Resources" on page 211.
    - Source (View)
    - Target (Create, Edit, View, Delete)
  - Datastores - See "Managing Datastore Resources" on page 215.
    - Source (View)
    - Target (Create, Edit, View, Delete)
  - Resource Pools (Create, Edit, View, Delete) - See "Managing Resource Pool Resources" on page 219.
  - Environments (Create, Edit, View, Delete) - See "Managing Environment Resources" on page 223.
- **Users**

  Project role assignments at the project level are automatically inherited by their child components.
  - Dashboard Viewer (View, Add, Remove) - See "Dashboard Viewer Role" on page 66.
  - Migration Specialist (View, Add, Remove) - See "Migration Specialist Role" on page 65.
  - Project Architect (View, Add, Remove) - See "Project Architect Role" on page 65.
  - Project Manager (View, Add, Remove) - See "Project Manager Role" on page 65.

Prerequisites for Projects

When you create a project, you must associate it with a specific organization. Before you can create a project, you must create the parent organization to ensure that it is available when you create the project.See “Creating an Organization” on page 70.

Viewing a Project

The View option allows users with the View permissions to view the project information.

To view project information:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Projects** tab.
3. In the **Projects** list, search for and locate the appropriate project.
4. Select the project, then click **View**.
5. Click **Close** when you are done.

Creating a Project

Only the System Administrator user can create projects.

To create a project:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Projects** tab.
3. Click **Create**.
4. In the **Project** pane, specify the following information:
   - Name
   - Organization
   - (Optional) Description
   - Start Date
   - End Date
   - Mode (Planning or Automated)
   - (Optional) Custom Field Names (Field 1 to Field 7)
5. Click **Save** to create the project object and activate the **Associations** pane.
6. In the **Associations** pane, click each tab to configure settings for the associated components. See “Configuring Project Associations” on page 99.
7. (Optional) Click **Set to Completed**, or click **Reopen**.
8. Click **Save**.
9. Click **Close**.

Editing a Project

You might need to modify dates, custom fields, and associations for a project as you configure the project and as the project matures. The System Administrator and Project Manager can modify the project settings.

To edit a project:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Projects** tab.
3. Double-click the project to open the Edit Project dialog, then view the project details.

   You can alternatively select the project, then click **Edit**.
4. Modify the values as appropriate:
   - Name
   - Organization
   - (Optional) Description
   - Start Date
   - End Date
   - Mode (Planning or Automated)
   - (Optional) Custom Field Names (Field 1 to Field 7)
5. (Optional) View or modify Associations. See "Configuring Project Associations" on page 99.
6. (Optional) Click **Set to Completed**, or click **Reopen**.
7. Click **Save**.
8. Click **Close**.

Configuring Custom Field Names for a Project

The custom fields defines additional information that you want to track for a project. The custom fields apply project-wide. The System Administrator and Project Manager can configure 1 to 7 custom fields names for a project. Each field name must be unique in the project. Only defined custom fields are available in the Transformation dialog. The Project Architect sets the values appropriate for the workload.

To define custom field names:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Projects** tab.
3. Do one of the following:
   - **3a** Click **Create** to open the Create Project dialog, then configure the required fields.
   - **3b** Double-click the project to open the Edit Project dialog, then view the project details.
4. Under **Custom Field Names**, specify a name for up to seven custom fields.
5. (Optional) Click **Set to Completed**, or click **Reopen**.
6. Click **Save**.
7. Click **Close**.

Configuring Project Associations

The Associations pane for a project shows all of the possible associations. The System Administrator user, Project Manager user, and Project Architect user can configure associations for a project.

Depending on your permissions, you can perform the same actions from the tabs under Associations as you can from the main tabs for these components. Some actions might not be available at this time. You can save the project and return later to complete information.

To associate components with a project:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Projects** tab.
3. Double-click the project to open the Edit Project dialog, then view the project details.
4. Under Associations, click each tab to view or modify settings for the associated components:

   Planning

   - Wave - See Chapter 13, “Managing Waves,” on page 101.
   - Batches - See Chapter 14, “Managing Batches,” on page 105.
   - Applications - See Chapter 15, “Managing Applications,” on page 109.
   - Workloads - See Chapter 19, “Configuring Workload Transformations,” on page 145.

   Resources

   - Credentials - See Section 21, “Managing Credentials Resources,” on page 193.
   - Hosts - See Section 22, “Managing Host Resources,” on page 199.
   - Migration Servers - See Section 23, “Managing Migration Server Resources,” on page 207.
   - Networks - See Section 24, “Managing Network Resources,” on page 211.
   - Datastores - See Section 25, “Managing Datastore Resources,” on page 215.
   - Resource Pools - See Section 26, “Managing Resource Pool Resources,” on page 219.
   - Environments - See Section 27, “Managing Environment Resources,” on page 223.

   Users

   - Dashboard Viewer - See “Dashboard Viewer Role” on page 66.
   - Migration Specialist - See “Migration Specialist Role” on page 65.
   - Project Architect - See “Project Architect Role” on page 65.
   - Project Manager - See “Project Manager Role” on page 65.
5. (Optional) Click **Set to Completed**, or click **Reopen**.
6. Click **Save**.
7. Click **Close**.

Deleting a Project

Only the System Administrator user can delete a project.

NOTE: Deleting a project deletes all data for the project from the PTM database.

To delete a project:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Projects** tab.
3. Select the appropriate project, click **Delete**, then click **Yes** to confirm the deletion.

13 Managing Waves

PlateSpin Transformation Manager supports rolling wave planning to accommodate the near-term and long-term planning of your transformation tasks. Waves represent a major effort for workload transformation. For near-term waves, you can plan the tasks in detail. You can refine the transformation plan for future waves as newer and better information becomes available.

About Waves

The Wave allows you to track the following information for your transformation activities:

Name: A friendly name for the wave that is unique in your transformation project. The name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Organization: The name of the parent organization for the wave.

Project: The name of the parent project for the wave.

Description: (Optional) A brief description of the wave.

Start Date: The date that the transformation of workloads is planned to begin. By default, the start date is inherited from its parent project’s start date. Dates with a gray font indicate inherited values.

End Date: The date that the transformation of all workloads must be completed successfully. By default, the end date is inherited from its parent project’s end date. Dates with a gray font indicate inherited values.

Associations: Associations define the relationship between a variety of component types and the wave.

- Batches (Create, Edit, View, Delete) - See “Managing Batches” on page 105.
- Workloads - See “Configuring Workload Transformations” on page 145.
  - Original (Edit, View, Delete)
  - Proposed (Edit, View, Delete)
- Users
  - Dashboard Viewer (View) - See “Dashboard Viewer Role” on page 66.
  - Migration Specialist (View, Add, Remove) - See “Migration Specialist Role” on page 65.

- Project Architect (View) - See “Project Architect Role” on page 65.
- Project Manager (View) - See “Project Manager Role” on page 65.

Prerequisites for Waves

When you create a wave, you must associate it with a specific organization and project. Before you can create a wave, you must create the parent organization and project to ensure that they are available when you create the wave.

For instructions, see the following:

- “Creating an Organization” on page 70.
- “Creating a Project” on page 97.

Viewing Waves

The View option allows users with only the View permissions to view the wave information.

To view wave information:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Waves** tab.
3. In the **Waves** list, search for and locate the appropriate wave.
4. Select the wave, then click **View**.
5. Click **Close** when you are done.

Creating a Wave

A System Administrator user and a Project Manager can create waves.

To create a wave:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Waves** tab.
3. Click **Create**.
4. In the **Waves** pane, specify the following information:
   - Name
   - Organization
   - Project
   - (Optional) Description
   - Start Date
   - End Date
5. Click **Save** to create the wave object and activate the **Associations** pane.
6. (Optional) In the **Associations** pane, click each tab to configure settings for the associated components. See “Configuring Wave Associations” on page 103.
7. (Optional) Click **Set to Completed**.

8 Click **Save**.

9 Click **Close**.

Editing a Wave

You might need to modify dates and associations for a wave as you configure the wave and as its parent project matures. The System Administrator and Project Manager can modify the wave.

To edit a wave:

1 In the Web Interface toolbar, select **Planning**.

2 Select the **Waves** tab.

3 Double-click the wave to open the Edit Wave dialog, then view the wave details.

You can alternatively select the wave, then click **Edit**.

4 Modify the values as appropriate.

5 (Optional) View or modify the Associations. See “Configuring Wave Associations” on page 103.

6 (Optional) Click **Set to Completed**.

7 Click **Save**.

8 Click **Close**.

Configuring Wave Associations

The Associations pane for a wave shows all of the associations. The System Administrator user, Project Manager user, and Project Architect user can configure associations for a wave.

Depending on the status of workloads, some actions might not be available at this time. You can save the wave and return later to complete information.

To associate components with a wave:

1 In the Web Interface toolbar, select **Planning**.

2 Select the **Waves** tab.

3 Double-click the wave to open the Edit Wave dialog, then view the wave details.

4 Under Associations, click each tab to view or modify settings for the associated components:

- Batches - See Chapter 14, “Managing Batches,” on page 105.
- Workloads - See Chapter 19, “Configuring Workload Transformations,” on page 145.
- Users
  - Dashboard Viewer - See “Dashboard Viewer Role” on page 66.
  - Migration Specialist - See “Migration Specialist Role” on page 65.
  - Project Architect - See “Project Architect Role” on page 65.
  - Project Manager - See “Project Manager Role” on page 65.

5 (Optional) Click **Set to Completed**.

6 Click **Save**.

7 Click **Close**.

Deleting a Wave

Only the System Administrator user and Project Manager can delete a wave. Deleting a wave deletes all data for the wave from the PTM database. It moves the member workloads to an Imported state, but it does not delete the associated workload information.

To delete a wave:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Waves** tab.
3. Select the appropriate wave, click **Delete**, then click **Yes** to confirm the deletion.

14 Managing Batches

Batches are the component efforts for waves in the workload transformation. A batch groups like workloads that you want to transform together. Batches can be more easily scheduled during intervals when network resources are available. You can coordinate the start dates and cutover dates with stakeholders to work around planned events that are critical to the business.

About Batches

The Batch allows you to track the following information for your transformation activities:

Name: A friendly name for the batch that is unique in your transformation project. The name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Organization: The name of the parent organization for the batch.

Project: The name of the parent project for the batch.

Wave: The name of the parent wave for the batch.

Description: (Optional) A brief description of the batch.

Start Date: The date that the transformation of its member workloads is planned to begin. By default, the start date is inherited from its parent wave’s start date. Dates with a gray font indicate inherited values.

End Date: The date that the transformation of all of its member workloads must be completed successfully. By default, the end date is inherited from its parent wave’s end date. Dates with a gray font indicate inherited values.

Associations: Associations define the relationship between a variety of component types and the batch.

- Workloads - See “Configuring Workload Transformations” on page 145.
  - Original (Edit, View, Delete)
  - Proposed (Edit, View, Delete)
- Users
  - Dashboard Viewer (View) - See “Dashboard Viewer Role” on page 66.

- Migration Specialist (View, Add, Remove) - See “Migration Specialist Role” on page 65.
- Project Architect (View) - See “Project Architect Role” on page 65.
- Project Manager (View) - See “Project Manager Role” on page 65.

Prerequisites for Batches

When you create a batch, you must associate it with a specific organization, project, and wave. Before you can create a batch, you must create the parent organization, project, and wave to ensure that they are available when you create the batch.

For instructions, see the following:

- “Creating an Organization” on page 70.
- “Creating a Project” on page 97.
- “Creating a Wave” on page 102.

Viewing a Batch

The View option allows users with only the View permissions to view the batch information.

To view batch information:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Batches** tab.
3. In the **Batches** list, search for and locate the appropriate batch.
4. Select the batch, then click **View**.
5. Click **Close** when you are done.

Creating a Batch

A System Administrator user and a Project Manager can create batches.

To create a batch:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Batches** tab.
3. Click **Create**.
4. In the **Batches** pane, specify the following information:
   - Name
   - Organization
   - Project
   - Wave
   - (Optional) Description
   - Start Date
   - End Date
5. Click **Save** to create the batch object and activate the **Associations** pane.

6 (Optional) In the **Associations** pane, click each tab to configure settings for the associated components. See “Configuring Batch Associations” on page 108

7 (Optional) Click **Set to Completed**.

8 Click **Save**.

9 Click **Close**.

Editing a Batch

You might need to modify dates and associations for a batch as you configure the batch and as its parent project matures. The System Administrator and Project Manager can modify the batch.

To edit a batch:

1 In the Web Interface toolbar, select **Planning**.

2 Select the **Batches** tab.

3 Double-click the batch to open the Edit Batch dialog, then view the batch details.

You can alternatively select the batch, then click **Edit**.

4 Modify the values as appropriate.

5 (Optional) View or modify Associations. See “Configuring Batch Associations” on page 108

6 (Optional) Click **Set to Completed**.

7 Click **Save**.

8 Click **Close**.

Moving a Batch to a Different Wave

1 In the Web Interface, select **Planning** in the toolbar, then select the **Batches** tab.

2 Select one or more batches from the list that you want to move to a different wave. The batches can be in the same or different wave. Press Shift and click the appropriate batches to select multiple batches.

3 Click **Move**.

4 In the Move dialog, select the appropriate destination wave, then click **OK**.

5 If there are date conflicts between the dates set for the batches and their new parent wave, specify your preference for how you want to handle the dates.

For example, you might choose to allow the batch dates to change to fit into the window scheduled for the new parent wave, or keep the existing dates and fix them later. Conflicted dates display in a red font.

If the batches’ dates are inherited, the batches automatically inherit the dates from the new parent wave.

6 Click **Save**.

7 Click **Close**.

8 For each batch that has dates in conflict, select the batch, click **Edit**, modify the dates, then click **Save**.

Configuring Batch Associations

The Associations pane for a batch shows all of the associations. The System Administrator user, Project Manager user, and Project Architect user can configure associations for a batch.

Depending on the status of workloads, some actions might not be available at this time. You can save the batch and return later to complete information.

To associate components with a batch:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Batches** tab.
3. Double-click the batch to open the Edit Batch dialog, then view the batch details.
4. Under Associations, click each tab to view or modify settings for the associated components:
   - Workloads - See Chapter 19, “Configuring Workload Transformations,” on page 145.
   - Users
     - Dashboard Viewer - See “Dashboard Viewer Role” on page 66.
     - Migration Specialist - See “Migration Specialist Role” on page 65.
     - Project Architect - See “Project Architect Role” on page 65.
     - Project Manager - See “Project Manager Role” on page 65.
5. (Optional) Click **Set to Completed**.
6. Click **Save**.
7. Click **Close**.

Deleting a Batch

Only the System Administrator user and Project Manager can delete a batch. Deleting a batch deletes all schedule data for the batch from the PTM database. It moves the member workloads to an Imported state, but it does not delete the associated workload information.

To delete a batch:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Batches** tab.
3. Select the appropriate batch, click **Delete**, then click **Yes** to confirm the deletion.

15 Managing Applications

PlateSpin Transformation Manager allows you to define the application types to track for workload transformations.

About Applications

The Application allows you to track the following information for your transformation activities:

Name: A friendly name for the application that is unique in your transformation project. The name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Organization: The name of the parent organization for the project.

Project: The name of the parent project for the application component.

Description: (Optional) A brief description of the application.

Prerequisites for Applications

When you create an application component, you must associate it with a specific organization and project. Before you can create an application, you must create the parent organization and project to ensure that they are available when you create the application.

Viewing Applications

The View option allows users with only the View permissions to view the application information.

To view application information:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Applications** tab.
3. In the **Applications** list, search for and locate the appropriate application.

4 Select the application, then click **View**.

5 Click **Close** when you are done.

Creating Applications

A System Administrator user and a Project Manager can create application components.

To create an application:

1 In the Web Interface toolbar, select **Planning**.

2 Select the **Applications** tab.

3 Click **Create**.

4 Specify the following information:

- Name
- Organization
- Project
- (Optional) Description

5 Click **Save** to create the batch object and activate the **Associations** pane.

6 Click **Close**.

Creating Applications during Spreadsheet Import

You can also create Application objects during the Spreadsheet Import. For each workload, add one or more **Application** columns. If an application does not match an existing Application object for the project, the import creates it. The specified applications are automatically associated with the workload. See Table C-8, “Application Parameter,” on page 185.

Editing Applications

The System Administrator and Project Manager can modify the application.

To edit an application:

1 In the Web Interface toolbar, select **Planning**.

2 Select the **Applications** tab.

3 Double-click the application to open the Edit Application dialog, then view the application details.

You can alternatively select the application, then click **Edit**.

4 Modify the values as appropriate.

5 Click **Save**.

6 Click **Close**.

Associating Applications and Workloads

The Applications panel in the Transformation dialog shows all of the applications associated with the workload. The System Administrator user, Project Manager user, and Project Architect can associate an application with a workload.

To associate an application with a workload:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Workloads** tab.
3. Double-click the workload to open the Edit Workload dialog, then scroll down to the Applications panel to view the applications associated with the workload.
4. Click **Add** to add an application association.
5. Click **Remove** to remove an application association.
6. Click **Save**.
7. Click **Close**.

In Spreadsheet Import, the applications you specify for the workload are automatically associated with the workload. See also Table C-8, “Application Parameter,” on page 185.

Deleting an Application

Only the System Administrator user and Project Manager can delete an application.

NOTE: If you delete an application, it automatically removes the associations for it, including those assigned to workloads that are in progress and completed. If the same associations apply, consider whether it is more appropriate to edit the application name and settings.

To delete an application:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Applications** tab.
3. Select the appropriate application, click **Delete**, then click **Yes** to confirm the deletion.

16 Monitoring Projects

PlateSpin Transformation Manager provides the ability to view status and metrics for the project and individual workloads.

Using the Dashboard

PlateSpin Transformation Manager provides the ability to view summary metrics for your transformation project in the Dashboard.

- Shows users metrics only for their assigned projects, waves, or batches.
- Shows the progress of projects, waves, and batches.
- Shows metrics for workloads by state, operating system, and transformation method.
- Runs status and count calculation jobs every 5 minutes. Displays **Needs Update** if workload changes occur while the calculation jobs run.

About the Dashboard

The PlateSpin Transformation Manager Dashboard displays statistics about your assigned transformation projects, waves, or batches. The

Figure 16-1 *Project View of the Dashboard*

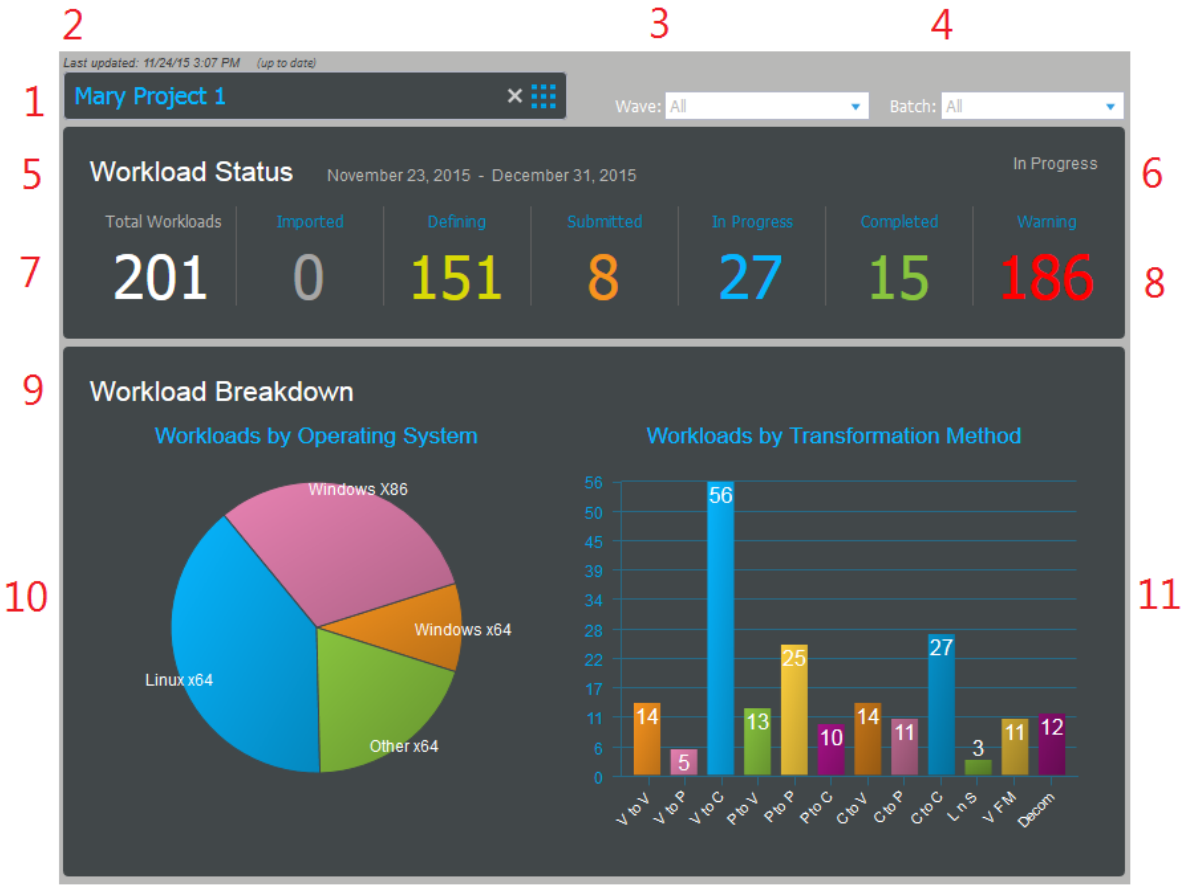

Table 16-1 *Dashboard Information*

Key	Metric	Description
1	Project	If you have roles on multiple projects, the Project field allows you select which project's information to display. If you have a role on a single project, the Project field and Project Selector settings are static.
2	Last Updated	The PTM Server compiles status and counts for each project every 5 minutes. This timestamp tells you how current are the displayed statistics. If changes occur while you view the dashboard, it also displays the `Needs update` message. If no changes have occurred, it displays the `Up to date` message.
3	Waves	If a project has multiple waves, the Waves field allows you to select which wave's information to display. The default is all waves.

Key	Metric	Description
4	Batches	If a project has multiple batches, the Batches field allows you to select which batch's information to display. The default is all batches.
5	Schedule	At the top of the Workload Status pane, PTM displays the start date and cutover date of the project, wave, or batch currently being displayed.
6	Project Status	The current state of the project. For example: In Progress, Closed
7	Transformation Status	The total number of workloads for the project, wave, or batch currently being displayed, and the counts for workloads by category: Total Workloads, Imported, Defining, Submitted, In Progress, Completed, Warning. Click the category title or value to open a Workloads list with the appropriate member items. You can alternatively use the **Health** > **All Warning State** option on the Advanced Search dialog to see all workloads in the Warning state.
8	Warning	The count includes any workload transformation that has one or more of the following conditions: ◆ Transformation is in a failed state. ◆ Transformation is not in progress, but the transformation start date has passed. ◆ Transformation has not been competed, but the transformation cutover date has passed. ◆ Transformation start date is before its parent's start date (batch, wave, project). ◆ Transformation cutover date is after its parent's cutover date.
9	Workload Breakdown	The graphical presentation of workloads based on workload characteristics.
10	Workloads by Operating System	A pie chart showing the percentage of workloads for the project, wave, or batch currently being displayed in each of the operating system categories. This breakout depends on the mix of operating systems allocated in the parent schedule container.

Key	Metric	Description
11	Workloads by Transformation Method	A bar chart showing the count of workloads for the project, wave, or batch currently being displayed in the transformation method categories: ◆ Virtual to virtual (V to V) ◆ Virtual to physical (V to P) ◆ Virtual to cloud (V to C) ◆ Physical to virtual (P to V) ◆ Physical to physical (P to P) ◆ Physical to cloud (P to C) ◆ Cloud to virtual (C to V) ◆ Cloud to physical (C to P) ◆ Cloud to cloud (C to C) ◆ Lift and shift (LnS) ◆ Virtual file move (VFM) ◆ Decommission (Decom)

Viewing Counts and Status for a Project, Wave, or Batch

The Projects, Waves, and Batches lists include the following summary health and status counts in the **Counts** column:

Total

Total number of workloads in the project.

Def

Total number of workloads currently in the defining state, typically with a status of Needs Additional Info or Ready to be Submitted.

Submit

Total number of workloads that have been defined and are waiting for transformation to start.

InProg

Total number of workloads currently in the transforming state.

Done

Total number of workload transformations currently in the Completed state.

Warn

Total number of workloads in the project that currently have one or more warning conditions.

The Projects, Waves, Batches, and Workloads lists report the current status in the Status column.

Viewing Workload Status

On the Workloads tab, the Status column shows the transformation workflow state and the migration sub state. You can mouse over the state or sub state message to view details in a tooltip.

B Global Planning Objects

In PlateSpin Transformation Manager, global planning objects are available to all projects.

User Roles

- System Administrator
- Project Manager
- Project Architect
- Migration Specialist
- Dashboard Viewer

See “Roles” on page 63.

Workload Types

- Cloud
- Physical
- Virtual

Workload Transformation Methods

- Cloud to cloud (C to C)
- Cloud to physical (C to P)
- Cloud to virtual (C to V)
- Decommission (Decom)
- Lift and shift (LnS)
- Physical to cloud (P to C)
- Physical to physical (P to P)
- Physical to virtual (P to V)
- Virtual to cloud (V to C)
- Virtual file move (VFM)
- Virtual to physical (V to P)
- Virtual to virtual (V to V)

See “Editing the Workload Transformation” on page 162.

Operating System Families

- CentOS
- Citrix
- Linux
- NetWare/OES
- Other
- Red Hat Linux
- Solaris
- SUSE Linux
- Ubuntu
- Unknown
- VMware ESX
- Windows

See Section 6, “Configuring Operating Systems,” on page 57.

Operating System Architectures

- x32 (32 bit)
- x64 (64 bit)

See Section 6, “Configuring Operating Systems,” on page 57.

Cloud Providers

The Cloud Provider options represent the top providers by market share in the industry.

- Amazon Web Services
- Google
- Microsoft Azure
- None
- Rackspace
- vCloud

See “About Host Resources” on page 201.

Virtualization Technologies (Hypervisors)

The hypervisor options represent the top virtualization technologies by market share in the industry.

- Citrix XenServer
- Linux KVM
- Microsoft Hyper-V
- None
- SUSE Xen
- VMware

See “About Host Resources” on page 201.

Volume Types

- Local
- NAS (network attached storage)
- SAN (storage area network)
- System

Storage Mapping Strategies

- Same as Source
- Custom
- All Volumes on a Single Disk
- One Volume per Disk

Storage Objects

- Disk
- Dynamic Mirrored

- Dynamic RAID 5
- Dynamic Simple
- Dynamic Spanned
- Dynamic Striped
- Partition
- Swap
- Volume
- Volume Group

Disk Copy Methods

Disk copy methods are used for data transfer between the source workload and target workload.

- Block with Driver
- Block without Driver
- File

Compression Levels

Compression levels represent the available compression methods for data transfer during replication.

- Fast
- Maximum
- None
- Optimal

Migration Server Types

- PlateSpin Migrate
- Other (any migration tool can be defined using this category)

Host Discovery States

The Hosts list displays the Host discovery state in the Status column.

- Discovering
- Discovery Failed
- Discovery Succeeded
- Discovered

Project Phases

The Projects list displays the Project Phase in the Status column.

- Planning
- Ready to Transform
- In Progress
- Completed

Transformation Phases

The Transformation Phases are options in the Transformation Plan workflow. The Workloads list displays the Transformation Phase in the Status column. You can set the transformation phase manually by using the Transformation dialog for a single workload or the Bulk Status Change dialog for one or more selected workloads.

- Ready to be Imported
- Imported
- Needs Additional Info
- Ready to Submit
- Submitted, Waiting for Start Date
- Ready to Transform
- Transforming
- Transformation Failed
- Pre-Cutover Testing
- Pre-Cutover Testing Failed
- Pre-Cutover Testing Succeeded
- Ready to Cut Over, Waiting for Cutover Date
- Ready to Cut Over
- Cut Over
- Post-Cutover Testing
- Post-Cutover Testing Failed
- Post-Cutover Testing Succeeded
- Completed
- Reopened
- Withdrawn
- Error

Workload Health

Workload health reports error conditions for workload transformation. You can search on Health objects in the Advanced Search for workloads.

- All

- OK
- Warning
- Error
- All Warning states
- Other warnings
- State warnings
- Sub State warnings
- All Planning warnings
- Start date is before batch start date
- Cutover date is after batch end date
- Invalid date range
- Not started, but start date has passed
- Not cut over, but cutover date has passed

Workflow Status

Status reflects the transformation workflow for the workload. You can search on Status objects in the Advanced Search for workloads.

- All
- All Defining states
- All Failed states
- All In-progress states
- All Submitted states
- Completed
- Cut Over
- Error
- Imported
- Importing
- Needs Additional Info
- Post-cutover Testing
- Post-cutover Testing Failed
- Post-cutover Testing Succeeded
- Pre-cutover Testing
- Pre-cutover Testing Failed
- Pre-cutover Testing Succeeded
- Ready to Cut Over
- Ready to Cut Over, waiting for Cutover Date
- Ready to Import
- Ready to Submit
- Ready to Transform

- Reopened
- Submitted, waiting for Start Date
- Transformation Failed
- Transforming
- Withdrawn

Migration Sub Status

In a PlateSpin Migration Environment, the PlateSpin Migrate Server reports events in the migration process. You can search on Sub Status objects in the Advanced Search for workloads.

- All
- Adding workload failed
- Adding workload to migration server
- Automated testing
- Automated testing failed
- Automated testing stalled. Attempting auto-recovery.
- Automated testing succeeded
- Automatic migrations not supported for this workload
- Canceling create test cutover environment
- Canceling cutover
- Canceling incremental replication
- Canceling initial replication
- Canceling reverting test cutover
- Canceling target VM prepare
- Canceling target host discovery
- Custom import scripts failed
- Custom import scripts succeeded
- Cutover canceled
- Cutover in progress
- Discovering target host
- Discovery succeeded, ready to plan
- Incremental replication
- Incremental replication canceled
- Incremental replication failed
- Incremental replication paused while testing
- Incremental replication stalled. Attempting auto-recovery.
- Incremental replication succeeded
- Initial replication
- Initial replication canceled
- Initial replication failed

- Initial replication stalled. Attempting auto-recovery.
- Initial replication succeeded
- Initial workload discovery
- Initial workload discovery failed
- Initial workload discovery stalled. Attempting auto-recovery.
- Initial workload discovery succeeded
- Manual tests failed
- Manual tests in progress
- Manual tests succeeded
- Migrate Server insufficient licenses
- Migrate Server not configured
- Migration cutover process failed
- Migration cutover progress stalled. Attempting auto-recovery.
- Migration environment prepare failed
- Migration environment prepare stalled. Attempting auto-recovery.
- Post-cutover testing failed
- Post-cutover testing in progress
- Post-cutover testing stalled. Attempting auto-recovery.
- Post-cutover testing succeeded
- Pre-migration validation
- Pre-migration validation failed
- Pre-migration validation stalled. Attempting auto-recovery.
- Pre-migration validation succeeded
- Preparing target VM
- Ready to replicate
- Removing the completed workload details failed
- Removing the completed workload details succeeded
- Removing the completed workload from the migration server
- Reverting test cutover
- Reverting test cutover failed
- Reverting test cutover stalled. Attempting auto-recovery.
- Running custom import scripts
- Target VM prepare canceled
- Target host discovered
- Target host discovery canceled
- Target host discovery failed
- Target host discovery stalled. Attempting auto-recovery.
- Test cutover failed
- Test cutover in progress
- Test cutover stalled. Attempting auto-recovery.

- Test cutover succeeded
- Unhandled error
- User-provided note
- Waiting for physical host to register
- Waiting for user to complete manual tests
- Waiting for user to initiate prepare
- Withdrawing workload
- Workload add stalled. Attempting auto-recovery.
- Workload added to migration server
- Workload being removed from migration server
- Workload configure failed
- Workload configure stalled. Attempting auto-recovery.
- Workload not configured
- Workload removed from migration server
- Workload target removed from migration server and host

IV Workloads

PlateSpin Transformation Manager allows you to import information about the cloud, physical, virtual, machines for each transformation project. You can use each workload's original configuration to define its appropriate destination workload and target network.

17 Overview of Workloads

Workloads are the cloud, physical, or virtual machines that you want to transform from their current mode of operation to a future mode of operation. You can use a single import or bulk import method to add information about a project's original workloads to the transformation database. You can modify the original workload information by re-importing the machine information.

NOTE: If you use PlateSpin Migrate Connector and PlateSpin Migrate servers for a project, a successful import triggers an automated discovery process that adds more information about each of the imported source workloads.

On import of an original workload, PlateSpin Transformation Manager automatically creates a matching proposed workload based on the original configuration. You can use individual edit and bulk edit to customize the proposed workload.

Automated Mode

Automated Mode allows you to plan workload migrations and to automate execution of supported migrations on PlateSpin Migrate servers in your PlateSpin Migration Environment. You can begin planning without setting up the entire migration environment. PlateSpin Migrate servers must be deployed in the network before you begin automated migration tasks. See the *PlateSpin Migrate 12.2 Installation and Upgrade Guide* for installation requirements and instructions.

If you switch from Planning Mode to Automated Mode for a project, the proposed workload settings for Storage and NICs revert to the settings for the original workload. Otherwise, all previously imported workloads settings are unchanged. For subsequently imported workloads, the Migration Server is automatically set to Auto Assign.

Workload Types

PlateSpin Transformation Manager supports the following workload types:

- **Cloud:** The workload is hosted in a cloud provider's infrastructure-as-a-service environment, such as Amazon Web Services, Google Cloud Platform, Microsoft Azure, Rackspace, or VMware vCloud.
- **Physical:** The workload is hosted on a physical machine, such as a tower, rack, or blade server.
- **Virtual:** The workload is a virtual machine hosted on a virtualization host server running a hypervisor, such as Citrix XenServer, Linux KVM, Microsoft Hyper-V, SUSE Xen, or VMware.

Workload Transformation Methods

PlateSpin Transformation Manager supports planning for the following workload transformation methods:

- Cloud to cloud (C to C)
- Cloud to physical (C to P)
- Cloud to virtual (C to V)
- Physical to cloud (P to C)
- Physical to physical (P to P)
- Physical to virtual (P to V)
- Virtual to cloud (V to C)
- Virtual to physical (V to P)
- Virtual to virtual (V to V)
- Lift and shift (LnS)
- Virtual file move (VFM)
- Decommission (Decom)

  For Decommission, the Transformation dialog displays the Original configuration, project, wave, batch, and schedule. PTM does not display the Proposed workload configuration, but it stores the values in the PTM database. It includes proposed values in Advanced Search, Bulk Edit, and Bulk Status Change actions and tracks indirect edits. Thus, you can move a workload in and out of a decommission state.

Automated Mode supports automated execution for two transformation methods where the virtualization hypervisor is VMware and the virtual machines are on VMware Cluster hosts:

- Physical to virtual
- Virtual to virtual

All other transformation methods are automatically set up in Planning Mode because they are not supported for automated migration.

OS Types

The System Administrator manages a list of available Operating System Types for the product in the System Configuration settings.

Contact your Transformation Manager administrator if you need additional OS Types for your transformation projects. See Section 6, “Configuring Operating Systems,” on page 57.

Custom Fields

You can define custom fields for each project. The Workloads list can include custom field names as column headers if a project is selected in the Project Selector. PTM also adds the custom fields that you defined for the project to the Advanced Search form and the Bulk Edit form. See “Configuring Custom Field Names for a Project” on page 98.

Network Connectivity and Access Requirements

PlateSpin Migrate Connector requires network connectivity to the following resources, based on its assignment to the PlateSpin Transformation Manager server or to a specific project:

- Its assigned PlateSpin Transformation Manager server
- Source workloads
- Target VMware cluster hosts
- PlateSpin Migrate servers

Your environment must meet the following requirements for network connectivity and access:

- “Event Messaging” on page 133
- “PlateSpin Discovery Environment” on page 134
- “PlateSpin Migration Environment” on page 134
- “Secure HTTPS Setup for PlateSpin Transformation Manager” on page 135

Event Messaging

PlateSpin Transformation Manager publishes workload workflow state change messages for its registered listeners. Each Migrate Connector instance registers with its assigned Transformation Manager server or project and listens for events and performs the appropriate actions for them on Migrate servers in a project.

In a PlateSpin Migration Environment, each PlateSpin Migrate server publishes workload migration state change messages for its registered listeners. A Migrate Connector instance registers with the Migrate servers in a project, then listens for messages and delivers them to the appropriate project and workload in Transformation Manager.

These message queues are pre-configured and start automatically when you start the PlateSpin service for a PlateSpin Transformation Manager server and a PlateSpin Migrate server. They do nothing unless you open the necessary port to allow registration and a PlateSpin Migrate Connector registers for the messages. The messaging function starts, stops, and restarts automatically with its parent service. Do not modify the default settings.

Table 17-1 shows the protocol and port required for messaging for discovery and automated migration. These messages reflect events and state changes and do not contain sensitive information.

Table 17-1 *Event Messaging Requirements for Network Protocols and Ports*

Traffic	Network Protocol and Port	Other Requirements
Event Messaging	61613 TCP, incoming (not secure)	This port is open by default on the Transformation Manager Appliance. For automated migration, open this port on the PlateSpin Migrate servers assigned to the project.

PlateSpin Discovery Environment

Workload discovery requires that you enable incoming ping (ICMP echo reply and ICMPv4-In echo request) traffic for source workloads and firewalls. PlateSpin supports only IPv4. For information about required software, network, and port settings for workload discovery, see Table 17-2.

Table 17-2 *Workload Discovery Requirements for Network Access and Communications*

Discovery Target	Network Protocols and Ports	Other Requirements
Windows workloads	◆ ICMP, incoming ◆ SMB (TCP 445 or 139)	◆ Microsoft .NET Framework 2.0 SP2, 3.5 SP1 or 4.0 ◆ Credentials with Domain Admin or built-in Administrator privileges
Linux workloads	◆ ICMP, incoming ◆ SSH (TCP 22, incoming)	Root-level access. For information on using an account other than `root`, see KB Article 7920711 (https://www.netiq.com/support/kb/doc.php?id=7920711).

Host discovery requires that you enable incoming ping (ICMP echo reply and ICMPv4-In echo request) traffic for target VMware hosts and firewalls. PlateSpin supports only IPv4. For information about required software, network, and port settings for host discovery, see Table 17-3.

Table 17-3 *Host Discovery Requirements for Network Access and Communications*

Discovery Target	Network Protocols and Ports	Other Requirements
VMware Cluster hosts	◆ ICMP, incoming ◆ SMB (TCP 445 or 139, incoming)	VMware account with an Administrator role

Ensure that you open the port required for event messaging on the PlateSpin Transformation Manager server and on each PlateSpin Migrate Connector. See “Event Messaging” on page 133.

PlateSpin Migration Environment

Table 17-4 provides the ports to open in the firewall and on each of the Migrate servers in order for PlateSpin Transformation Manager to use the Migrate REST APIs for automated migration. In addition, the Migration Server resource for Migrate server must provide a valid Credentials resource for the Migrate Administrator user.

Table 17-4 *REST API Requirements for Network Access and Communications*

REST API Traffic	Network Protocol and Port	Access
HTTPS (secure)	Port 443, TCP, incoming and outgoing	Administrator login credentials for the Migrate server
HTTP (non-secure)	Port 80, TCP, incoming and outgoing	Administrator login credentials for the Migrate server

Ensure that you open the port required for event messaging on the PlateSpin Transformation Manager server and on each PlateSpin Migrate server. See “Event Messaging” on page 133.

PlateSpin Transformation Manager requires that your migration environment meets the PlateSpin Migrate requirements for network communications. See “Requirements for Migration” in the *PlateSpin Migrate 12.2 User Guide*.

Secure HTTPS Setup for PlateSpin Transformation Manager

For secure HTTPS connections between PlateSpin Migrate Connector and PlateSpin Transformation Manager, the Jetty SSL settings on the PlateSpin Transformation Manager Appliance VM are configured with the latest recommended security settings.

18 Importing and Discovering Workloads

Large-scale transformation projects typically involve hundreds, or even thousands, of workloads. It is time-consuming to define each workload individually in your Transformation Manager system. PlateSpin Transformation Manager provides tools to help simplify that process. You can import workload details from a spreadsheet, or leverage an automated discovery process that retrieves details about the source workloads.

About Automated Workload Discovery

If a PlateSpin Migrate Connector is available to the project, PlateSpin Transformation Manager provides automated discovery of source workloads as you import them into the project. You can also rediscover details for a workload if necessary.

Discovery Requirements for Source Workloads

Before you attempt an import with auto-discovery, you must configure your PlateSpin Discovery Environment:

- **PlateSpin Transformation Manager 1.1**

  Configure PlateSpin Transformation Manager and set up a project for the workloads you want to discover.

- **PlateSpin Migrate Connector 1.1**

  Deploy a Migrate Connector in the same network as the source workloads and assign it to the PTM Server or to your project on the server. For deployment information, see "Deployment Requirements" in the *PlateSpin Migrate Connector Quick Start*.

  PlateSpin Transformation Manager Appliance includes a Migrate Connector instance that is pre-configured to work with the PTM Server. You can also deploy Migrate Connector instances separately on your own SUSE Linux Enterprise Server 11 SP4 servers. See the *PlateSpin Migrate Connector Quick Start*.

NOTE: Ensure that the Migrate Connector is up and running before you attempt discovery or rediscovery or workloads.

- **Source Workloads**
  - The Connector supports automated discovery of Windows and Linux source workloads for the following platforms:
    - Physical machines with x86 and x64 architectures
    - Virtual machines (VMs) with x86 and x64 architectures on VMware Cluster hosts

    For cloud or other virtual workload types, you can use the Spreadsheet Import method to add workload details.
  - Ensure that the source workloads are up and running.
  - Ensure that the network connection is working between the PlateSpin Migrate Connector and the source workloads.
- **Credentials**

  Create a Credentials resource to use for source workload discovery. The Credentials resource is also used later for replication to the target VM on the host.
- **Network Connectivity and Access**

  For information about network connectivity and access information for automated discovery, see “PlateSpin Discovery Environment” on page 134

Import and Discovery Process

The import and auto-discovery process performs the following actions:

1. The import creates a new Workload object unless it matches an existing workload.
2. The Workloads list displays the new workloads at the next page refresh, using the IP address as the workload name if the FQDN is not provided.
3. Auto-discovery begins for each new workload, using the provided Credentials resources to log in to the workload. Discovery status displays in the Workloads list and Transformation dialog.
4. On successful workload discovery, the Workloads list or Transformation dialog displays the newly discovered details at the next page refresh. The discovered hostname replaces the IP address or FQDN as the workload name, as appropriate.

Troubleshooting Discovery Failures

Discovery fails for any of the following conditions:

- A Migrate Connector is not available.
- The source workload is not supported for automated discovery.
- The necessary ports are not open in the source workload’s operating system firewall.
- The source workload Credentials resources are not provided.
- The Credentials resource has an invalid value for the user name and password.
- The source workload is not running at the time of discovery.
- The provided IP address or FQDN is used by a non-computer device such as a UPS, printer, or network equipment. The workload is automatically removed from the database.

After you resolve Migrate Connector, network, or credentials issues, you can retry discovery for supported source workloads. See “Retrying Workload Discovery” on page 142.

If you retry discovery for a workload that has failed discovery and the workload goes into a transformation workflow phase of **Imported** with a sub status of **Workload Not Configured**, it might be an indication that there is no valid Credential resource assigned to the workload. In order for the initial discovery to proceed, go to the Transformation dialog and assign a Credential resource in the Workload panel. See “Associating Credentials and Workloads” on page 195.

Viewing Workloads

You manage workloads through their transformation lifecycle from the Workloads tab. You import information about the source workloads, then track key settings and status for them from this page.

1. Select your project in the **Project Selector**.

   If you have rights in only in a single project, the project is selected automatically.
2. In the Web Interface, select **Planning** in the toolbar.
3. Select the **Workloads** tab.
4. Use the Filter to perform a simple search on the list. Use the Advanced Search ... to search on a variety of workload settings, custom fields, status, and health. See “Locating Workloads with the Filter and Advanced Search” on page 161.

   

5. View the **Status** to follow the workload state.

Importing Workloads with a Spreadsheet

Large-scale transformation projects typically involve hundreds, or even thousands, of workloads. It is time-consuming to define each workload individually in your Transformation Manager system. You can use the Spreadsheet Import dialog on the Workloads page to import information about the workloads for a single project from a specially formatted spreadsheet. For detailed information about preparing and validating the spreadsheet, see Appendix C, “Bulk Import Spreadsheet,” on page 173.

NOTE: If a PlateSpin Migrate Connector is available to the project, the successful import of workloads initiates an automated discovery process that retrieves details about the source workloads. For more information about auto-discovery, see “Import and Discovery Process” on page 138 and “Discovery Requirements for Source Workloads” on page 137.

Users Who Can Perform This Task: System Administrator, Project Manager, Project Architect

---

To use a spreadsheet to import multiple workloads for a project:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. Click **Spreadsheet Import**.
3. From the **Project** list, select the project to use for the workloads in the spreadsheet.

   If the global scope of operations is set to a project, the **Project** field is pre-populated with that value. To select a different project, exit the Spreadsheet Import dialog, deselect the global scope or reassign the scope to the appropriate project, and then return to Step 2 and try again.
4. Click **Browse**, then locate and select the Bulk Import spreadsheet (`.xls` file) that contains the workload information that you want to add to the specified project.

5 Click **Validate**, then view the **Validation Results** to verify that all required information is present and in the appropriate format.

   For example, it reports if an entry is missing required information, uses an unknown option in a field, or uses wrong the format for values.

6 If the spreadsheet contains errors:

   6a Click **Get Results** at the bottom of the dialog to generate an annotated spreadsheet, and save the file to your computer.

   6b Correct the errors in your Bulk Import spreadsheet, and save the file.

   6c Return to Step 5, and repeat the validation and correction process until there are no reported errors.

7 After the spreadsheet is validated, click **Import**.

8 After a successful import, click **Get Results** at the bottom of the dialog, then click **Save**.

   PlateSpin Transformation Manager adds the assigned workload IDs to the spreadsheet, as well as the final status of `Imported`. If you provided NIC and Disk information, it also adds the UUIDs assigned for those components.

9 In the Web Interface, select to **Planning** > **Workloads** to view the list of workloads.

10 (Connector) If the spreadsheet included credential information for one or more workloads, view **Status** to confirm the automated discovery for the workloads.

Importing Workloads with Auto-Discovery

If a PlateSpin Migrate Connector is available to the project, you can use the **Import** option to import workloads and leverage an auto-discovery process to add details about them:

- Specify an IPv4 address or FQDN to import and discover details for a single workload.
- Specify a range of IPv4 addresses to import and discover details for up to 256 workloads at a time.

A workload can have only one instance in a transformation project. If a workload already exists in the project, the discovery process is not performed for it. The auto-discovery also ignores non-computer networked devices.

Users Who Can Perform This Task: System Administrator, Project Manager, Project Architect

---

- “About Import” on page 140
- “Importing a Single Workload by IP Address or FQDN” on page 141
- “Importing Workloads in an IP Address Range” on page 142

About Import

The Import dialog for single or range import requires the following information:

Organization: Select the parent organization for the project.

Project: Select the parent project for the workload.

Description: (Optional) Specify a brief description of the workload. This text appears by default in the Workloads list.

IP or FQDN: Specify a single address or a range of addresses:

- **Single:** Specify the IP address or fully qualified domain name for the workload. PTM uses this information to discover and communicate with the workload. A workload can have only one instance in a transformation project.

  Examples:

  ```
  sqlserver5.paris.digitalair.com
  192.168.1.1
  ```

- **Range:** Specify the first IPv4 address in the workload range, then specify the fourth octet value for the last IPv4 address in the range. Valid values for the range octet are 0 to 255. PTM uses this information to discover and communicate with the workloads.

  Examples:

  ```
  10.10.1.0 – 255
  10.10.10.16 – 48
  ```

Credential: PTM uses the credentials to discover and communicate with the workload. You can specify 1 to 4 Credentials resources to try. The import tries each specified Credentials resource. At least one of the specified Credentials resources must be valid in order for discovery to succeed.

Importing a Single Workload by IP Address or FQDN

You can specify a single IPv4 address or FQDN to import and discover details for a single workload.

To import a single workload:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. Select **Import**.
3. In the Import a Workload dialog, specify the following information:
   - Organization
   - Project
   - Description
   - IP or FQDN

     Examples:

     ```
     sqlserver5.paris.digitalair.com
     192.168.1.1
     ```

   - Credentials

     Specify 1 to 4 Credentials resources to try.
4. Click **Import**.
5. Click **Close** to return to the Workloads list.
6. In the Workloads list, monitor the discovery status for the workload.
7. (Discovery Failed) If discovery fails, mouse over the workload status to understand the problem, then resolve the issues and retry discovery.

Importing Workloads in an IP Address Range

You can specify a range of IPv4 addresses to import and discover details for up to 256 workloads at a time.

To import workloads in a specified IP address range:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. Select **Import**.
3. In the Import a Range of Workloads dialog, specify the following information:
   - Organization
   - Project
   - Description
   - IP or FQDN

     Specify the first IP address in the range and last octet in the range.

     Examples:

     ```
     10.10.1.0 – 255
     10.10.10.16 – 48
     ```
   - Credentials

     Specify 1 to 4 Credentials resources to try against each workload in the range.
4. Click **Import**.
5. Click **Close** to return to the Workloads list.
6. In the Workloads list, monitor the discovery process for each workload in the range.
7. (Discovery Failed) If discovery fails, mouse over the workload status to understand the problem, then resolve the issues and retry discovery for the failed workload.

Retrying Workload Discovery

You might want to retry discovery for a workload if the initial discovery fails. For information about possible causes of discovery failure, see “Troubleshooting Discovery Failures” on page 138.

To retry discovery for one or more workloads:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. Use the Advanced Search to select one or more workloads where the initial discovery failed. In the **Sub Status** list, select **Initial workload discovery failed**.
3. Click in the Workloads list to exit the Advanced Search.
4. For each workload, mouse over the **Initial workload discovery failed** status to view details about the reason the discovery failed.
5. Select one or multiple workloads in the refined list.
6. In the actions above the list, click **Retry** to initiate discovery for each of the selected workloads.

To retry discovery for a single workload:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. Use the Filter and Advanced Search to locate and select the workload, then click **Edit**.

3 Mouse over **Retry** in the header area for information about why the initial discovery failed.

4 Click **Retry** to initiate discovery.

Rediscovering Workloads

You might want to rediscover a workload if you modify the source workload before you submit the transformation for execution. After you edit a proposed workload by using the Transformation dialog or the Bulk Edit options, any rediscovered workload information applies only to the original workload.

To rediscover one or more workloads:

1 In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.

2 Use the Advanced Search to select one or more workloads where the initial discovery failed. In the **Sub Status** list, select **Discovery succeeded, ready to plan**.

3 Click in the Workloads list to exit the Advanced Search.

4 Select one or multiple workloads in the refined list.

5 In the actions above the list, click **Rediscover** to initiate rediscovery for each of the selected workloads.

To rediscover a single workload:

1 In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.

2 Use the Filter and Advanced Search to locate and select the workload, then click **Edit**.

3 In the header area, note the Sub Status **Discovery succeeded, ready to plan**.

4 Click **Rediscover** to initiate rediscovery.

19 Configuring Workload Transformations

After you import workloads and discover information about them, you are ready to configure the proposed transformation for each workload.

About Automated Migration

In a PlateSpin Migration Environment, PlateSpin Transformation Manager provides automated migration for physical machines and VMware virtual machines to target virtual machines on VMware Cluster hosts.

Automation Deployment Requirements

Before you attempt to execute an automated migration, you must configure your PlateSpin Migration Environment:

- **PlateSpin Transformation Manager 1.1**

  Configure PlateSpin Transformation Manager and set up a project and discover the workloads and hosts.

- **PlateSpin Migrate Connector 1.1**

  A PlateSpin Migrate Connector must be available to the project. For deployment information, see "Deployment Requirements" in the *PlateSpin Migrate Connector Quick Start*.

- **PlateSpin Migrate 12.2**

  PlateSpin Migrate servers must be deployed in the network. For deployment information, see "Planning Your PlateSpin Migration Environment" in the *PlateSpin Migrate 12.2 User Guide*.

  You can create Migration Server resources without the detailed information for initial planning. Select Auto Assign in the workload's Transformation Plan to allow the PlateSpin Migrate Connector to manage PlateSpin Migrate server assignments from among the Migration Server resources that you create for the project.

- **Supported Transformations**

  The Connector supports automated migration only for the following source workloads and target VM hosts:

  - Physical machines with x86 and x64 architectures
  - Virtual machines (VMs) on VMware Cluster hosts
  - Target VMware Cluster hosts

  For other workload types, use manual or third-party migration tools.

- **Source Workloads**

  Ensure that the source workloads are up and running.

  NOTE: Operating systems supported for automated execution are subject to those supported by PlateSpin Migrate. See "Supported Source Workloads" in the *PlateSpin Migrate 12.2 User Guide*.

- **Target Hosts**

  VMware Cluster hosts must be available as target hosts. You can create Host resources without the detailed information for initial planning. For host discovery information, see "About Host Resources" on page 201.

- **Network Connectivity and Access**

  Ensure that the network connections are working:

  - Between the PlateSpin Migrate Connector and the source workloads
  - Between the PlateSpin Migrate Connector and the PlateSpin Migrate servers
  - Between the source network and target network

  For information about network connectivity and access information for automated migration, see "PlateSpin Migration Environment" on page 134.

Automation Process

For information about how the automated migration works in the PlateSpin Migration Environment, see "PlateSpin Migration Environment" on page 16.

Troubleshooting Automated Migration Failures

Migration fails for any of the following conditions:

- A Migrate Connector is not available.
- The assigned Migrate server is not available.

- The target VMware Cluster host is not available.
- The source workload is not supported for automated migration.
- The target transformation method is not supported for automated migration.
- The target host is not supported for automated migration.
- The necessary ports are not open in the source workload's operating system firewall.
- The source workload Credentials resource is not provided.
- The Migrate server Credentials resource is not provided.
- The target VMware Cluster host Credentials resource is not provided.
- The Credentials resource has an invalid value for the user name and password.
- The source workload is not running at the time of migration execution.

About the Transformation Dialog

You plan each transformation through the Transformation dialog. It contains information about the original workload, the proposed workload, and the transformation status. This dialog allows you to modify the proposed workload and the transformation state for the workload.

In the Transformation dialog, the font colors for Proposed values indicate states for the data:

- **Light Gray:** The Proposed field values are light gray (dimmed) if the section has not yet been edited directly. Each Proposed field value is inherited from its source workload's imported or discovered value. In the Transformation Plan section, each field value is inherited from its parent batch, wave, or project value.
- **Gray:** The Proposed field values are gray if the section has been edited directly and the value matches the displayed Original value. The Proposed values can be modified if the Edit action is enabled for the section.
- **Blue italics:** The field value has been modified from its imported or discovered value.

  If the section in the Transformation dialog has not been edited directly, any changes to the Proposed default or imported values have been made indirectly. Otherwise, the changes might have been made directly or indirectly. The source and history of changes to field values are not tracked.
- **Red:** The section heading is in a red boldface font if errors or warnings exist for any of its fields. The red error condition displays next to the field label. For example, an error exists if a required value is missing.

The dialog allows you to view or modify the following information for the workload transformation:

- Header
- Workload IDs and Discovery Address
- URL for the Transformation dialog
- Errors and Warnings
- Status and Retry
- Migration Sub Status
- Discovery Status and Rediscovery
- Last Modified
- License State
- Transformation Plan

- Migration Settings
- Workload
- Location
- Network Interfaces
- Testing Network Interfaces
- Storage
- Applications
- Custom Fields
- Transformation History

Header

The Transformation dialog header provides the following information about the workload:

- FQDN (fully qualified domain name)
- Project name
- Wave name or <wave not assigned>
- Batch name or <batch not assigned>

Workload IDs and Discovery Address

Mouse over any of the objects in the Transformation dialog header area to view the following IDs for the selected workload:

- Transformation ID
- Original Workload ID
- Proposed Workload ID
- Migration Server Workload ID

  This value is known after a PlateSpin Migrate server has been assigned. It is used for matching workloads in the database of the specified PlateSpin Migrate server.
- Discovery Address

  This is the IP address used for discovery and communications with the workload.

URL for the Transformation dialog

Click the Navigation URL icon in the upper right corner of the Transformation dialog to show the URL to the selected workload transformation. You can open the link in a new tab, or copy the URL to the clipboard. You can also email the copied URL to stakeholders to allow them to go directly to the selected page. Log in is required for those users to gain access.

Errors and Warnings

If errors or warnings exist, the Transformation dialog reports the total number that are associated with the workload transformation. For sections that contain errors or warnings, the section title and the field values of missing or invalid information appear in a red font.

	Errors exist
	Warnings exist

NOTE: A **Retry** option is available for failed states that are recoverable conditions.

Status and Retry

The Status indicates where the workload is in its transformation workflow. Figure 19-1 shows the user responsible for setting the status, the various workload states, and the workflow for the transformation process. The thick lines show the default next state available for each state. See Table 19-1 for a description of each workload state and its possible next states.

Some states are associated with the start date or cutover date for the workload's assigned batch. The state reports a Warning if that date has passed, but the expected action has not yet occurred.

NOTE: A **Retry** option is available for failed states that are recoverable conditions.

Figure 19-1 *The Transformation Workflow*

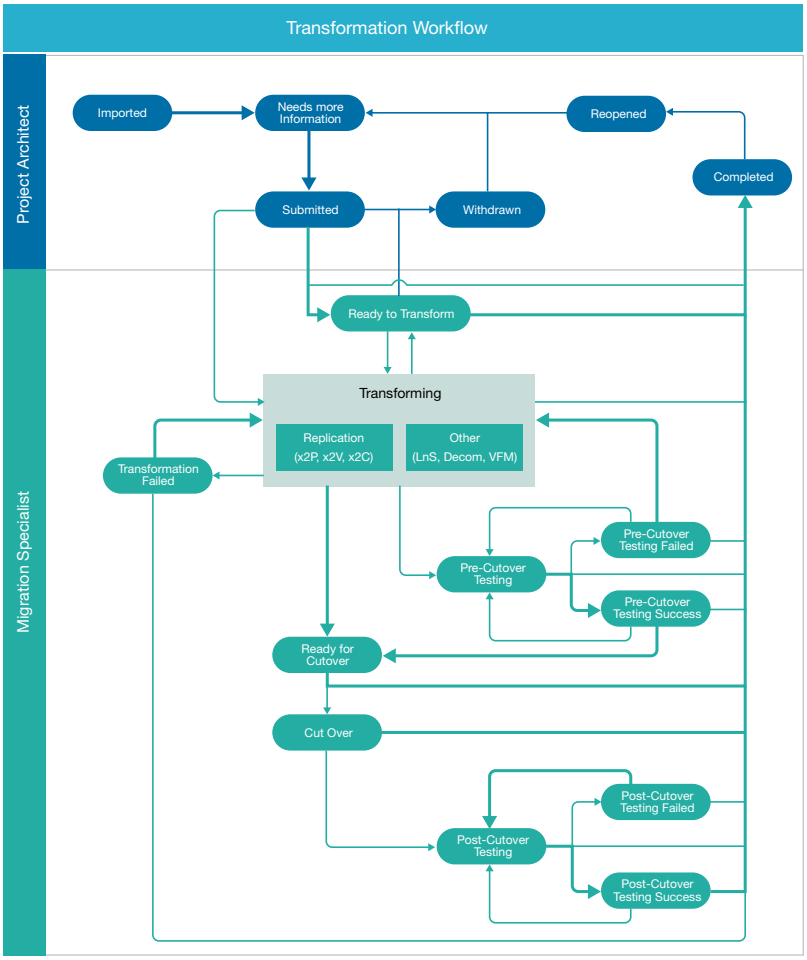

Table 19-1 *State Descriptions and Next States*

State	Description	Possible Next States
Imported	The workload information has been imported as an original workload and a proposed workload. You have not yet modified the imported information. The workload is unlicensed until you begin to edit the workload information.	Needs Additional Info (automatic on edit)
Needs Additional Info	You have modified the imported workload information directly in the Transformation dialog or with a Batch Edit action. Some required values might be missing, or there might be errors to resolve. Edits include setting the schedule for the workload, modifying its proposed workload configuration, or adding required values.	Ready to Submit (automatic on no errors) Submitted, Waiting for Start Date
Ready to Submit	All required values are present and there are no errors to resolve. A red Warning notice appears in the Status column if the start date for its assigned batch has passed.	Submitted, Waiting for Start Date
Submitted, Waiting for Start Date	The workload has been submitted for transformation, and the start date for its assigned batch is a future date. When the start date occurs, the workload's status automatically changes to Ready to Transform. The Migration Specialist can manually change the workload's status from Submitted to Transforming before the start date.	Ready to Transform Transforming Completed Withdrawn
Ready to Transform	The workload is ready for transformation, and the start date for its assigned batch has passed.	Completed Transforming Withdrawn
Withdrawn	A submitted workload has been withdrawn because it needs changes for the planned transformation before the transformation process can continue.	Needs Additional Info
Transforming	The Migration Specialist has begun the transformation process for the workload, such as replication methods (x2P, x2V, or x2C) or other methods (LnS, Decom, and VFM).	Ready to Cutover Completed Pre-Cutover Testing Transformation Failed Ready to Transform
Transformation Failed	The transformation from original workload to proposed workload has been attempted and failed.	Transforming Completed
Ready to Cutover	The transformed workload is ready, waiting for the Cutover Date.	Completed Cut Over Post-Cutover Testing
Pre-Cutover Testing	The transformed workload tests are in progress.	Pre-Cutover Testing Succeeded Pre-Cutover Testing Failed Completed

State	Description	Possible Next States
Pre-Cutover Testing Succeeded	The transformed workload test succeeded, and the workload is ready for cutover.	Ready for Cutover Completed Pre-Cutover Testing
Pre-Cutover Testing Failed	The transformed workload test failed. Repeat the transformation or the re-test the workload.	Transforming Completed Pre-Cutover Testing
Cut Over	The workload has been cut over to its target network.	Completed Post-Cutover Testing
Post-Cutover Testing	After the cutover, tests on the cutover workload are in progress.	Post-Cutover Testing Succeeded Post-Cutover Testing Failed Completed
Post-Cutover Testing Succeeded	The tests on the cutover workload succeeded.	Completed Post-Cutover Testing
Post-Cutover Testing Failed	The tests on the cutover workload failed. Re-test the workload if needed.	Post-Cutover Testing Completed
Completed	The workload transformation process is completed.	Reopened
Reopened	A completed workload transformation has been reopened because it needs changes for the planned transformation before the transformation process can continue.	Needs Additional Info

Migration Sub Status

During automated migration, the PlateSpin Migrate server reports migration state events for the workload, referred to as its Migration Sub Status. You can search on Sub Status states in the Advanced Search dialog.

For a list of reportable migration state events, see “Migration Sub Status” on page 125.

Discovery Status and Rediscovery

When a PlateSpin Migrate Connector is available to the project, the header area displays the discovery status of the workload below the Status.

After a successful discovery, a **Rediscovery** option is available. You might initiate rediscovery if you make manual changes to the source workload, such as adding NICs or modifying its assigned datastores. See “Rediscovering Workloads” on page 143.

Last Modified

The Last Modified field shows the timestamp of the last modifications made directly in the Transformation dialog, or indirectly by actions performed through import, discovery, and changes to the proposed workload or status.

License State

View the license state in the lower right corner of the header area of the Transformation dialog. Mouse over the Licensed (green) icon to view the license date.


Licensed


Not licensed


License is not available

Transformation Plan

The Transformation Plan panel allows you to specify the following information for the workload transformation:

- **Project:** The parent project for the workload.
- **Wave:** The parent wave for the specified batch.
- **Batch:** The parent batch scheduled for this workload.
- **Start Date:** The planned start date for the workload transformation. The date must fall in the execution window for its parent batch, and occur before the cutover date.
  - **Time:** Specify the time of day to begin workload replication for the migration. The time is based on the time zone of the PlateSpin Migrate server that you will use to execute the migration.
- **Cutover Date and Time:** The planned cutover date for the workload transformation. The date must fall in the execution window for its parent batch.
  - **Time:** Specify the time of day to begin the cutover execution. The time is based on the time zone of the PlateSpin Migrate server that you will use to execute the migration.
- **Transformation Method:** The transformation method to use for this workload. The options depend on the workload type of the original workload. See “Workload Transformation Methods” on page 132.
- **Migration Server:** The Migration Server resource that will be used to migrate the original workload to the proposed workload. Before you can select a Migration Server, you must first create Migration Server resources for your project in **Resources** > **Migration Servers**. See Chapter 23, “Managing Migration Server Resources,” on page 207.

  NOTE: Migration Server selection is disabled if you enable Auto Assign.

  You can save a Migration Server resource in an incomplete state without a URL and credentials. In this case, the Migration Server field displays the Migration Server resource name and the **Not configured** state in the Transformation Plan panel. When you later deploy the migration server in your environment, you should update the Migration Server resource with its URL and credentials. The resource state is automatically updated for all of its associated workloads.
- **Auto Assign:** Select the check box to enable the load-balanced automatic assignment of migration jobs to your PlateSpin Migrate servers. In Automated Mode, you can use Auto Assign as an alternative to assigning a specific Migration Server resource to each workload.

The Migrate Connector uses round-robin load-balancing to distribute workload migration jobs evenly across all of the PlateSpin Migrate servers in your project. It detects all currently running Migrate servers for the round-robin. When it receives a new workload migration request, the Connector assigns the workload to the next Migrate server in the sequence with available capacity and licenses.

After the workload is assigned to a Migrate server, a link to the workload in PlateSpin Migrate is available in the workload's Transformation dialog. You can follow the link to go directly to the workload on the assigned Migrate server.

NOTE: For Windows Cluster migration, the PlateSpin Migrate server must meet the prerequisites described in "Migrating Windows Clusters" in the *PlateSpin Migrate 12.2 User Guide*. If you use Auto Assign, ensure that all Migrate servers for the project satisfy these prerequisites.

- **On Hold:** Select the check box to put the workload in an On Hold state, where the workflow cannot progress. Deselect the option to remove the hold. Workloads in an On Hold state are dimmed in the Workloads list.

Migration Settings

The Migration Settings panel provides information about the default behaviors for the replication environment.

General Settings

- **Disk Copy Method:** The method used to create an exact copy of volume data from the source disk to the target disk. Options for the data transfer are:
  - **Block with Driver:** A data transfer method that dynamically installs a special driver on the source workload to facilitate volume data transfer at the block level. It leverages snapshots if they are available. PlateSpin Migrate requires this option for all Linux workloads.
  - **Block without Driver:** A driverless data transfer method for Windows workloads. It leverages snapshots if they are available, or it uses a driverless synchronization with an MD5-based replication.
  - **File:** A data transfer method for Windows workloads that copies data and replicates changes on a file-by-file basis.
- **Install Driver:** If you select Block with Driver as the disk copy method, you can specify when in the migration process to install the driver on the source workload.
  - **Prepare:** Installs the required data transfer software on the source workload and creates a target workload in preparation for the workload replication.
  - **First Replication:** Installs the required data transfer software on the source workload and begins the first full replication of the workload.
- **Compression Level:** Specify level of compression to apply to workload data before its transmission for replications to the target workload.
  - **Fast:** Consumes the least CPU resources on the source, but yields a lower compression ratio.
  - **Optimal:** (Recommended) Consumes optimal CPU resources on the source and yields an optimal compression ratio.
  - **Maximum:** Consumes the most CPU resources on the source, but yields a higher compression ratio.
  - **None:** No compression.

- **Perform daily incremental replication:** Select the check box to enable daily incremental replication to the target workload.
  - **Time:** Specify the time of day to run the incremental replication job. The time is based on the time zone of the PlateSpin Migrate server that you will use to execute the migration. The default run time is 1:00 a.m. in the time zone of the PlateSpin Migrate server.
- **Encrypt data:** Select the check box to enable encryption of data for transmission.
- **Pause workflow:** Select the check box enable the workflow to pause between the migration configuration and the workload prepare process.
- **Automatic cutover:** Select the check box to enable the automatic cutover of the workload when the cutover date is reached.
- **Perform replication with test cutover:** Select the check box to enable an incremental replication before executing the test cutover.
- **Perform replication with cutover:** Select the check box to enable an incremental replication before executing the cutover.
- **Shut down source during cutover:** Select the check box to enable the shut down of the source workload during the cutover.
- **Shut down target during cutover:** Select the check box to enable the shut down of the target workload during the cutover.

Replication Network

- **Network:** Select the network resource assigned to the target host machine.
- **MTU:** The *Maximum Transmission Unit* setting in bytes that represents the largest packet size that can be used on the replication network. Valid values are 68 to 65535. For example, most Ethernet networks support an MTU setting of 1500 bytes. On a VPN, the MTU value is typically smaller than 1500 bytes.

  Specify an MTU setting that matches the smallest MTU setting on the replication network path between the source workload and the target workload. If a packet is larger than the MTU value of an interface, the packet might be split into smaller packets, or it might be dropped. Both packet fragmentation and retransmission introduce latency and congestion that can dramatically increase the replication time. Tuning a workload's MTU value for its replication network helps avoid these conditions, which optimizes the transmission performance and minimizes replication time.
- **Use DHCP to get network address details:** This setting is enabled by default. Deselect the check box to specify static IP addresses, then specify the network address settings for the replication network.
  - **IP Address:** IPv4 format.
  - **Subnet Mask:** IPv4 format.
  - **Default Gateway:** IPv4 format.
  - **DNS Servers:** IPv4 format. For multiple entries, specify one DNS server per line.

Workload

The Workload panel displays the original workload configuration that was imported or re-imported using the Spreadsheet Import feature. Initially, the proposed workload has the same configuration. You can edit the configuration of the proposed workload.

Workload Details

- **Host Name:** The simple hostname (or machine name) for the workload, not its FQDN. Example: `sqlserver10`
- **Domain Name:** The name of the domain that the proposed workload will join. Examples: `digitalair.com`, `paris.digitalair.com`
- **Admin Credentials:** (Optional) Select the Credentials resource to use to authenticate to a specified source workload.
- **Operating System:** The specified value from the OS Types list.
- **Architecture:** This value is inherited from the specified operating system.
- **Description:** (Optional) A textual description of the workload, such as information about migration constraints or dependencies.
- **Workload Type:** Cloud, Physical, or Virtual. This value is inherited from the transformation method.

Workload Resources

- **Amount of RAM:** The amount of memory for the workload in MB, GB, TB, PB, EB, ZB, or YB.
- **# Sockets:** The number of sockets installed on the workload.
- **# Cores per Socket:** The number of cores in each socket. The total number of cores for the workload is the cores per socket times the number of sockets.

Windows Settings

The Windows Settings area is available for Windows virtual workloads.

- **Domain or Workgroup:** Specify whether the Windows workload is a member of a domain or a workgroup, then specify the name.
  - **Domain:** Specify its domain name, such as `PARIS` or `paris.digitalair.com`.
  - **Workload:** Specify its workgroup name, such as `WORKGROUP`.
- **Domain Credentials:** If you specified Domain, select the Credentials resource for the domain administrator account needed to join the workload to the existing domain.

Location

The Location panel provides information about where the proposed workload resides in the network:

Host

- **Environment:** The environment resource assigned to the workload.
- **Cloud Provider:** (Type: Cloud) The cloud provider that hosts the workload:
  - Amazon Web Services
  - Microsoft Azure
  - Rackspace
  - vCloud
  - Google
- **Hypervisor:** (Type: Virtual) The type of hypervisor that hosts the workload's VM:
  - Citrix XenServer
  - Linux KVM

- Microsoft Hyper-V
- SUSE Xen
- VMware

- **Host:** (Type: Virtual) The Host resource to use for the target host.
- **Site:** The geographic location or facility where the workload network resides. Example: London
- **Hardware:** The vendor or model of hardware used by the workload. Example: Dell PowerEdge
- **Enclosure:** The location of the workload in the data center. Example: Rack A-31
- **Slot:** The slot number in the rack for the host machine. Example: 8

Virtual Machine Settings

- **VM Name:** The name of the VM used for the workload. The default value is the hostname of the workload.
- **Resource Pool:** (VMware) Select the Resource Pool resource where the target VM will be created on the VMware host.
- **VM Config Datastore:** The Datastore resource that contains the VM configuration file.
- **Path on Datastore:** The path to the VM configuration file.
- **Install VMware Tools:** (VMware) Specify whether to install VMware tools on the target virtual workload running on a VMware host.

Network Interfaces

The Network Interfaces panel defines information about each NIC on the workload. You cannot delete the source NIC information provided during the import, but you can add, edit, and delete other NICs for the target workload.

You must have at least one NIC for the source workload and one NIC for the target workload. You cannot delete the last remaining NIC in the list. You must create another NIC first, and then delete the unwanted NIC.

If you have multiple NICs on the source workload, specify which NIC to use for replication.

- **Interface Name:** The name for the NIC. Example: `eth0`
- **Network:** The network resource assigned to the workload. If you have multiple NICs, this setting is required for each NIC.
- **MAC Address:** (Planning Mode) MAC address of the source workload NIC in the format of 00-00-00-00-00-0c.

  For proposed workloads, the MAC address is assigned by the target workload. The address is not tracked for planning.
- **VLAN ID:** The network VLAN that carries the workload traffic.
- **Connect on boot:** The NIC is automatically set to connect to the network on boot.
- **Use DHCP to get network address details:** This setting is enabled by default. Deselect the check box to specify a static IP address, then specify the network address settings for the target network.
  - **IP Address:** IPv4 format.
  - **Subnet Mask:** IPv4 format.
  - **Default Gateway:** IPv4 format.

NOTE: If the source workload or target workload is located in a different internal (private) network behind a network address translation (NAT) device, it might be unable to communicate with its counterpart to replicate the workload. On the affected workload, ensure that you specify a public IP address for the NIC you use for replication.

- **Use DHCP to get name servers:** Specify whether the DNS name is acquired from the DHCP server in the associated network. This option is enabled by default. If you use static IP addresses, deselect this option, then provide the following information:
  - **DNS Servers:** IPv4 format. For multiple entries, specify one DNS server per line.
  - **Search Domains:** FQDN of one or more search domains.
- **Use for Replication:** On the Network Interfaces panel, specify whether to use the NIC for replication communications. The default value is **Yes**. At least one of the NICs must be enabled for replication.

Testing Network Interfaces

The Testing Network Interfaces panel defines information about the NIC configuration during cutover testing for the target workload. You can copy or reapply the default settings from Network Interfaces, and then modify the live network settings with settings for the test network. You can also create and delete NICs for testing.

- **Interface Name:** The name for the NIC. Example: `eth0`
- **Network:** The network resource assigned to the workload. If you have multiple NICs, this setting is required for each NIC.
- **MAC Address:** (Planning Mode) MAC address of the source workload NIC in the format of 00-00-00-00-00-0c.

  For proposed workloads, the MAC address is assigned by the target workload. The address is not tracked for planning.
- **VLAN ID:** The network VLAN that carries the workload traffic.
- **Connect on boot:** Select the check box to enable the NIC to connect to the network on boot.
- **Use DHCP to get network address details:** This setting is enabled by default. Deselect the check box to specify a static IP address, then specify the network address settings for the target network.
  - **IP Address:** IPv4 format.
  - **Subnet Mask:** IPv4 format.
  - **Default Gateway:** IPv4 format.

NOTE: If the source workload or target workload is located in a different internal (private) network behind a network address translation (NAT) device, it might be unable to communicate with its counterpart to replicate the workload. On the affected workload, ensure that you specify a public IP address for the NIC you use for replication.

- **Use DHCP to get name servers:** Specify whether the DNS name is acquired from the DHCP server in the associated network. This option is enabled by default. If you use static IP addresses, deselect this option, then provide the following information:
  - **DNS Servers:** IPv4 format. For multiple entries, specify one DNS server per line.
  - **Search Domains:** FQDN of one or more search domains.

Storage

The Storage panel defines information about the storage mapping strategy and about the mounted volumes on the workload. You cannot modify information for the source workload. To update the source workload disk and volume information, you can rediscover the source workload if you are using auto-discovery, or update the Bulk Import Spreadsheet and import the workload again using the Spreadsheet Import.

Storage Mapping Strategy

The mapping strategy translates the layout of storage disks and volumes from the source workload to a proposed layout in the destination storage location. In Planning mode, you can keep the resulting proposed layout, or modify the disk details for a custom mapped solution.

- **Same as source:** Storage objects for the target volumes and disks will be configured in the destination storage location using the same layout, size, and format as the source workload. For the proposed workload, you can select a datastore and path for target VMware environments, and define a remote path for SAN and NAS devices.
- **Custom:** In Planning Mode, you can add, edit, and remove volume groups, volumes, disks, and partitions. You can also restore the settings to the Same as Source settings for the original volumes and disks.

Volumes Tab

The Volumes tab allows you to view and configure the proposed volume layout on the proposed workload from a volume perspective.

- **Volumes:** A tree-view list of volumes and volume groups on the proposed workload.
  - **Name:** The name of the volume group or its member volumes.
  - **File System:** The file system format for the partition, such as Ext3, Linux-swap, NTFS, unallocated, extended, and so on.
  - **Size:** The total amount of space in the units specified that are allocated to the volume group or its child volumes. The unused amount of space available in the volume or volume group displays below the size.
  - **Cluster Size:** (Windows) For NTFS volumes, specifies the block allocation size.
  - **Type:** The type of storage object and file system for the selected volume:

    Example Storage Types:

    - Dynamic Mirrored
    - Dynamic RAID 5
    - Dynamic Simple
    - Dynamic Spanned
    - Dynamic Striped
    - Local
    - NAS
    - SAN
    - Swap
    - System
    - Volume

Example File System Types:

- NTFS
- Ext4

- **Mount Point:** The path on the workload where the volume can be mounted, such as `K:` or `/mnt/data2`. The volume might be mounted or unmounted.
- **Datastore:** The name of the VM datastore to use for the volume. The remote path to the virtual disk file in the datastore appears under the datastore name.
- **Path on Datastore:** For a datastore, specify the path on the datastore for the volume.
- **Remote Path:** The remote path for a SAN or Network Attached Storage.
- **Devices:** The devices and partitions used by the selected volume group or volume.

Disks Tab

The Disks tab allows you to view and configure the proposed disk and partition layout on the proposed workload.

- **Disks:** A tree-view list of disks and partitions on the proposed workload.
  - **Name:** The name of the parent disk and its child partitions.
  - **Size:** The total amount of space in the units specified that is allocated to a disk and its partitions. The amount of unallocated available space displays below the size.
  - **Volume:** The name of the volume that resides on the child partition.
  - **Type:** The type of storage object as disk or partition. The Partition type includes the file system format, such as Ext3, Linux-swap, NTFS, unallocated, extended, and so on.
  - **Datastore:** The name of the VM datastore to use for the disk. The path to the virtual disk file in the datastore appears under the datastore name.
- **Disk Details:** Select a partition in the Disks list to view or modify details about its disk and volume for the proposed workload.
  - **Type:** The type of storage object (disk or partition).
  - **Size:** The total amount of space allocated to a disk or partition in the units specified.
  - **Name:** The name of the volume or volume group, such as `Network_Share_11-T1`.
  - **Mount Point:** The path on the workload where the volume can be mounted, such as `K:` or `/mnt/data2`. The volume might be mounted or unmounted.
  - **Volume:** The name of the volume that resides on the partition.
  - **Unallocated:** The amount of free space available in the disk or partition in the units specified.
- **Virtual Disk Details:** Select a partition in the Disks list to view or modify details about its virtual disk and volume for the proposed workload. Customizing the mapped disk or volume is optional.
  - **Datastore:** The name of the VMware datastore for the target VM to use.
  - **Path on Datastore:** The path to the virtual disk file (`.vmdk`) in the selected datastore.
  - **Disk Provisioning:** The virtual disk provisioning policy to use on VMware for the virtual disk.
    - **Thin Disk:** Creates a disk in a space-efficient sparse disk format that can grow to the size is provisioned at creation time.
    - **Thick Disk:** Creates a virtual disk in a type of thick virtual disk that supports clustering features such as fault tolerance. Space required for the virtual disk is allocated at creation time.

Applications

The Applications panel identifies the applications installed on the workload. You can add or remove applications for the proposed workload.

Custom Fields

The Custom Fields panel allows you to specify values for the Custom Fields 1-7 that are defined for the project. Custom Fields are available only if they are defined for the project and the project is selected in the Global Project Selector when you open the Transformation dialog.

Transformation History

The Transformation History panel displays the most recent workflow state changes, migration sub-state changes, or user-provided notes for the workload. Each state change triggers an entry. You can also enter user notes to be logged, such as manual changes and actions performed for the workload.

Viewing the Workloads List

The Workloads tab allows you to track key settings and status for each workload. Use the following tips to manipulate the page and view various status and details.

- Set the Project Selector to the project of interest to list only workloads for that project.
- View the list based on details for the Original workload or for the Proposed workload.
- Mouse over the Status or Sub Status states to view details about the event or error conditions.
- Use the Filter to perform a simple search on the list. Use the Advanced Search ... to search on a variety of workload settings, custom fields, status, and health. See “Locating Workloads with the Filter and Advanced Search” on page 161.

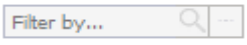

- Select a Health warning type to quickly list all workloads that are in a specific state. For example, `Not cut over, but cutover date has passed`. See “Workload Health” on page 123.
- Double-click the tab title to force any list to refresh immediately, regardless of its normal polling cycle.
- Mouse over the tab title to view the total number of items in the list and the total number of selected items currently selected in the list.
- A **Retry** option is available when you select a workload in a failed state that is a recoverable condition.
- A **Rediscover** option is available when you select a workload that was previously successfully discovered.

Viewing a Workload NIC and Volume Summary

The Workload Summary allows you to view key original and proposed workload information without opening the Edit Workload dialog. The information pops up as a tool-tip when you mouse over the workload name.

To view the workload NIC and disk summary:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. Use the **Filter** and **Advanced Search** options to locate the workload of interest.
3. (Optional) Mouse over a workload name to view a summary of its original workload and proposed workload configurations for NICs and volumes.

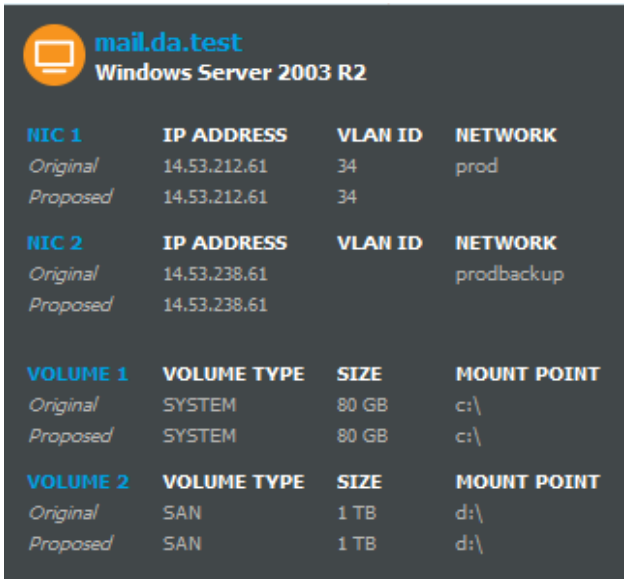

4. Move the mouse away from the workload name to exit the summary pop-up.

Locating Workloads with the Filter and Advanced Search

The Filter and Advanced Search tools allow you to locate one or more workloads that you want to perform actions on. It is an easy way to select multiple workloads to assign to batches or to modify transformation information.

1. In the Web Interface, click the **Global Project Selector**, then select the project of interest.
2. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
3. In the **Filter**, type characters to search for values in the searchable columns. The filter applies immediately to the list as you type the characters.

   You can click the **Clear Filter** icon to clear the field and remove the filter.
4. Click the **Advanced Search** icon to narrow the search by specifying your desired search parameters and values.

The Advanced Search dialog is a multiple-option form that allows you to search on any combination of the following parameters for Workloads that make sense for the target of your search:

Project	Workload Type	Total Storage	Network
Wave	Transform Method	Single Disk Size	VLAN ID
Batch	Environment	Cores per Socket	Status
Hostname	Site	Amount of Memory	Sub Status
OS Type	Enclosure	On Hold	Health
Application	Custom 1	Custom 2	Custom 3
Custom 4	Custom 5	Custom 6	Custom 7

NOTE: Each custom field that you define for a project appears in the form if the Global Project Selector is set to a specific project.

Click **Clear** to reset the Advanced Search and try again with different fields.

5. After you have narrowed the list to the items of interest, click in the list to exit the Advanced Search dialog.
6. Mouse over the Workloads tab to see the number of total items in the list.
7. In the filtered Workloads list, view the workloads listed in the search results. Use the scroll bar to navigate up and down through the items. See “Scrolling Up and Down in Lists” on page 39.
8. (Optional) Select multiple items in the list to perform Bulk Edit or Bulk Status Change actions on the selected workloads. See “Multiple Item Selection in Lists” on page 39.
9. Mouse over the Workloads tab to see the number of total items in the list and the total number of selected items.

Viewing Workload Transformation Details

The Workloads list provides key information and statistics about each workload. The **View** option on the Workloads tab allows users with only the View permissions to view workload transformation details for a selected workload.

To view workload transformation details:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. In the **Workloads** list, use the Filter and Advanced Search to locate the appropriate workload.
3. Select the workload, then click **View** to open the Transformation dialog.
4. View the workload information.
5. Click **Close** when you are done.

Editing the Workload Transformation

The initial workload import creates an original workload and a proposed workload. You use the Transformation dialog for the workload to set the appropriate configuration for the transformation and proposed workload.

Users Who Can Perform These Tasks: System Administrator, Project Manager, Project Architect

For information about the various parameters for the workload, see “About the Transformation Dialog” on page 147.

To configure the transformation and proposed workload for a single workload:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. Use the **Filter** and **Advanced Search** options to locate the workload of interest.
3. (Optional) Mouse over a workload **Name** to view a summary of its original workload and proposed workload configurations.
4. Double-click the workload you want to modify to open its Transformation dialog. You can alternatively select the workload, and then click **Edit**.
5. Modify the workload information on the Transformation dialog as appropriate:
   - Status
   - License
   - Transformation Plan
   - Migration Settings
   - Workload
   - Location
   - Network Interfaces
   - Testing Network Interfaces
   - Storage
   - Applications
   - Custom Fields
   - Transformation History
6. Use the following to save your changes for now, to submit the workload as ready for execution, or to return to a submitted workload to make changes:
   - Select **Mark In Progress** to keep the workload available for additional changes. This allows errors to exist until you return to make additional modifications.
   - When you have resolved all errors and provided the appropriate transformation details, click **Submit** to indicate that the workload is ready for transformation.

     NOTE: In Automated Mode, you cannot submit a workload transformation that has not been successfully discovered.

Configuring the Source NIC to Use for Replication

If you have multiple NICs for a source workload, you must specify which NIC to use for replication.

To enable a NIC for Replication:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. Use the **Filter** and **Advanced Search** options to locate the workload of interest.
3. Double-click the workload you want to modify to open its Transformation dialog. You can alternatively select the workload, and then click **Edit**.
4. Scroll down to view the Network Interfaces panel and locate the NIC you want to use for replication.

5 Under the Original workload, click **Enable** to set **Use for Replication** to **Yes**. At least one source NIC must be enabled.

6 Click **Close**.

Editing a Workload Transformation after Submit

You can modify the following settings for the workload transformation after you submit the workload without needing to withdraw it:

- **Transformation Plan**
  - Wave
  - Batch
  - Start Date and Time
  - Cutover Date and Time
  - On Hold
- **Transformation History**
  - User-provided notes

To modify the plan or add a note for a single workload:

1 In the Web Interface, select the project in the Global Project Selector.

2 Select **Planning** in the toolbar, then select the **Workloads** tab.

3 Use the **Filter** and **Advanced Search** options to locate the workload of interest.

4 Double-click the workload you want to modify to open its Transformation dialog. You can alternatively select the workload, and then click **Edit**.

5 In the Transformation Plan panel, click **Edit**, then modify the schedule information as needed.

6 In the Transformation History panel, click **View**, then add notes as needed.

7 Click **Close** to exit the Transformation dialog.

To modify the plan of one or more workloads using Bulk Edit:

1 In the Web Interface, select the project in the Global Project Selector.

2 Select **Planning** in the toolbar, then select the **Workloads** tab.

3 Use the **Filter** and **Advanced Search** options to locate the workloads of interest.

4 Select the specific workloads you want to modify.

5 Click **Bulk Edit**.

6 (Optional) Modify the plan settings as needed:
  - **Wave**
  - **Batch**
  - **Remove from Batch and Wave**
  - **Start Date**
  - **Cutover Date**
  - **On Hold**

7 Click **Apply**.

8 Review the change confirmation, then click **Proceed**.

Workloads dimmed in an On Hold state are dimmed in the Workloads list.

Bulk Edit for Multiple Proposed Workloads

The Bulk Edit function is a powerful editor that allows you to apply the same parameter setting to multiple selected workloads at a time. Each parameter is independent, and applies separately against a set of selected workloads. Bulk Edit validates each value against the type of workload and other characteristics to ensure that the change is made only if the setting is a valid option for that workload. This safeguard allows you to concurrently edit a set of diverse workloads that have a logical relationship that is meaningful to your network environment.

About Bulk Edit

The Bulk Edit dialog title area provides information about the number of workloads that you have preselected in the Workloads list before opening the Bulk Edit dialog. A System Administrator and users with Project Manager or Project Architect roles in multiple projects are able to select workloads across projects.

If a specified parameter value applies to one or more selected workloads, the confirmation dialog identifies the number of affected workloads, the number of skipped workloads, and the issues that resulted in skipped workloads.

NOTE: If you specify values for multiple parameters in the same bulk edit request, the confirmation dialog does not present all of the various matching conditions for the remainder of the parameters. It reports on only the first match it encounters. Thus, if you seek clear cause-and-effect reporting, you should edit only one parameter setting at a time in your bulk edit request.

Workloads might be skipped if the specified change or value is not valid for the workload type. For example, Host and Datastore changes apply only to target virtual workloads, or to workloads where the Transformation Method is unknown. Cloud Provider applies only to target cloud workloads, or to workloads where the Transformation Method is unknown.

Some actions cannot be performed against a selected workload based on the workload state:

- Values for workloads in a Submitted or later state are not editable. You must withdraw the workload to return it to an editable state.
- Values for workloads in a Completed state are not editable. You must reopen the workload to return it to an editable state. If the Completed state was initiated from a parent wave, batch, or project, you must reopen the parents first.

Table 19-2 describes wave, batch, and network parameters that apply for workloads in a single project at a time. When you select values, it narrows the action to workloads for a specific project. After you apply the change for one project, you can clear the Project field, and select different values to isolate a second Project, and so on until all appropriate changes are complete.

Table 19-2 *Bulk Edit Parameters for a Single Project at a Time*

Parameter	Description
Project	The Project value is informational if you preset the Project Selector. All selected workloads belong to that project. If you select any of the project-specific parameters, the Project value represents the project related to your selected value. All remaining project-related parameter values are restricted based on that choice.
Wave	This option reflects the parent wave of the target batch and sets the Project value if it is not set.
Batch	Set the target batch assignment for the selected workloads. It sets the parent wave and project if they are not set. Each workload inherits the dates from the newly assigned batch. Manual date settings override inheritance. If manual dates fall outside the new execution window, the affected dates display in a red font in the Workloads list and the Transformation dialog.
Remove from Batch and Wave	Move selected workloads out of their current assigned batches and waves.
VLAN ID	For any network, change the VLAN ID: ◆ From any value to the specified value. ◆ From a specified value to a new specified value. For a specified network, change the VLAN ID: ◆ From any value to the specified value. ◆ From a specified value to a new specified value.
Network resource	For any NIC, change the Network resource: ◆ From any value to the specified value. ◆ From a specified value to a new specified value. For a specified NIC, change the network resource: ◆ From any value to the specified value. ◆ From a specified value to a new specified value.
Host resource	Change the Host resource: ◆ From any value to the new specified value. ◆ From a specified value to a new specified value. The Host resource setting applies only to virtual target workloads, or to workloads where the Transformation Method is not currently set.

Parameter	Description
Datastore resource	Change the Datastore resource: ◆ From any value to the new specified value. ◆ From a specified value to a new specified value. The Datastore resource setting applies to virtual workloads, or to workloads where the Transformation Method is not currently set.
Migration Server resource	Change the Migration Server resource: ◆ From any value to the new specified value. ◆ From a specified value to a new specified value.
Replication network	Change the Replication network: ◆ From any network to the new specified network. ◆ From a specified network to a new specified network.

Table 19-3 describes parameters that apply to all selected workloads that are in an editable state.

Table 19-3 *Bulk Edit Parameters for a Single Project or Multiple Projects*

Parameter	Description
Number of workloads selected	The number of workloads that you have preselected in the Workloads list before opening the Bulk Edit dialog.
Number of projects represented	A System Administrator and users with Project Manager or Project Architect roles in multiple projects are able to select workloads across projects.
Transformation Method	Specify the transformation method. The setting applies only to workloads where the source workload type matches. For example, Physical to Cloud modifies the proposed transformation type only for physical source workloads in the selected workloads.
Disk Copy Method	Specify the disk copy method: block with driver, block without driver, and file.
Cloud Provider	Specify a cloud provider for the cloud workloads. **NOTE:** The Cloud Provider setting applies to cloud workloads, or to workloads where the Transformation Method is not currently set.
Start Date	Set the start date of the transformation execution.
Cutover Date	Set the cutover date for the workload.
On Hold	Pause the workload transformation workflow for the selected workloads.
Amount of Memory	Specify the size of the target memory.
Custom Fields	Each custom field that you define for a project appears in the form if the Global Project Selector is set to a specific project.

Selecting Workloads for Bulk Edit

Before you open the Bulk Edit dialog, use the Project Selector, Filter, Advanced Search, sort, and multiple select functions to refine the Workloads list and select the workloads for action. You can pre-select any number of workloads for concurrent bulk editing actions.

To select the workloads for a bulk edit action:

1. (Conditional) To work with workloads for a single project, click the **Project Selector**, then select the project of interest.
2. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
3. Use the Filter and Advanced Search options to refine the workloads list.
4. Use the multiple-select functions to select one or more workloads that you want to modify in the same way.
5. Mouse over the Workloads tab to view the number of workloads currently in the list, and the number of workloads that are selected.
6. Continue with "Applying a Bulk Edit" on page 168.

Applying a Bulk Edit

After your selection is ready, use the Bulk Edit dialog to specify values for one or more parameters in the selected workloads.

To perform a bulk edit:

1. After you select the workloads for action, click **Bulk Edit**.
2. Specify the information you want to change for the selected workloads. See "About Bulk Edit" on page 165.

   The following fields apply to a single selected project:

   - Project
   - Wave
   - Batch
   - Remove from Batch and Wave
   - VLAN ID (network, from VLAN ID, to VLAN ID)
   - Network resource (NIC, from Network, to Network)
   - Host resource (from Host, to Host)
   - Datastore resource (from Datastore, to Datastore)
   - Migration Server resource (from Migration Server, to Migration Server)
   - Replication Network (from Network, to Network)

   The following fields apply to all selected workloads:

   - Transformation Method
   - Disk Copy Method
   - Cloud Provider
   - Start Date
   - Cutover Date
   - On Hold

- Amount of Memory
- Custom field values (available only if the Global Project Selector is set)

3 Click **Apply**.

4 Review the confirmation message, then click **Proceed**.

Workloads placed in an On Hold status are dimmed in the Workloads list.

Bulk Status Change for Multiple Proposed Workloads

You can use the Bulk Status Change tool to modify the state value for proposed workloads for multiple workloads at a time. Use the Filter and Advanced Search options to locate and select the workloads of interest. See “Status and Retry” on page 149.

The Bulk Status Change dialog indicates the number of workloads selected and the number of projects you will affect with the Bulk Status Change. The change applies only to the selected items for which the next phase selected is a valid state change.

To change the status of one or more workloads:

1 In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.

2 Select one or more workloads that you want to set to the same state.

3 Click **Bulk Status Change**.

4 Do one of the following actions for the selected workloads:

- If the **Submit** option is available, one or more of the selected workloads is eligible to move to a phase of readiness for transformation. Click **Submit** to change the workload status to Ready for Transformation.
- Specify the **Next Phase** you want to set for selected workloads, then click **Apply**.

Synchronizing Edits from the Migrate Server

In the PlateSpin Migration Environment, a Migration Specialist might make manual changes to the workload migration job on the Migrate server. Events are synchronized to Transformation Manager through the Migrate Connector if the change affects settings that Transformation Manager tracks.

Examples of workload settings that are synchronized include Number of Cores, Cores per Socket, Datastores, static IP for replication network, and so on. Because Transformation Manager does not have Blackout Window or Bandwidth Throttling settings, it does not synchronize changes for them. Also, source credentials set in Migrate are not synchronized.

The Migration Specialist might also initiate an action by clicking a command in Migrate. The state change is synchronized, and Transformation Manager displays the changed state.

Viewing the Transformation History

The Transformation History tracks the following:

- Workflow state changes

- Migration sub-state changes
- User-provided notes that capture events that occur outside the monitoring provided by Transformation Manager

On the Transformation dialog, the Transformation History panel shows the 10 most recent events or user-provided notes. All events are available in the Transformation History log for the workload.

To view the transformation history for a workload:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. Use the **Filter** and **Advanced Search** options to locate the workload of interest.
3. Double-click the workload you want to modify to open its Transformation dialog. You can alternatively select the workload, and then click **Edit**.
4. Scroll down to view the Transformation History panel to view the 10 most recent events and user-provided notes.
5. Click **View** to open the Transformation History log.
6. Scroll or use the Filter to locate the events of interest.

   You can mouse over a long user-provided note to view the full message in a tooltip.
7. Click **Close** to exit the Transformation History log.
8. Click **Close** to exit the Transformation dialog.

Adding User Notes to the Transformation History

In the Transformation History, you can add user notes that capture events that occur outside the monitoring provided by Transformation Manager.

To add user notes to the history:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. Use the **Filter** and **Advanced Search** options to locate the workload of interest.
3. Double-click the workload you want to modify to open its Transformation dialog. You can alternatively select the workload, and then click **Edit**.
4. Scroll down to view the Transformation History panel, then click **View**.
5. Under **Create Note**, type the comments you want to add to the log.

   The note length is not limited to the size of the note field. The words scroll to the left as you type.
6. Click **Add Note**.

   The new note is time-stamped and added to the Transformation History with the status of **User-provided note**. The first two lines of the note display in the Transformation History dialog. For notes longer than two lines, you can mouse over the note to view the full message in a tooltip.
7. Click **Close** to exit the Transformation History log.

   The new note displays as an entry in the most recent events in the Transformation History. The first line of user-provided note displays in the Transformation History panel. For notes longer than one line, you can mouse over the note to view the full message in a tooltip.
8. Click **Close** to exit the Transformation dialog.

Withdrawing a Submitted Workload

After you submit a workload for migration, you can withdraw it at any point in the migration. In Planning Mode, the state returns immediately to **Prepared, waiting**.

In Automated Mode, you have the option of removing the target VM and cleaning up the workload job from the Migrate server when you Withdraw. The behavior depends on where in the workflow the migration is at the time:

- If the migration is waiting to start, the state returns immediately to **Prepared, waiting**.
- If the migration has started, you have an option to automatically clean up the VM host and Migrate server:
  - If you opt to remove the target VM, the clean-up occurs before it the state returns to **Prepared, waiting**. The clean-up can take a while. You cannot modify the workload until the clean-up is done.
  - If you opt to not remove the target VM, the state returns immediately to **Prepared, waiting**. The workload job remains on the Migrate server and the target VM remains on the target host. You should manually clean up the target VM and Migrate server.

To withdraw a workload:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. Use the **Filter** and **Advanced Search** options to locate the workload of interest.
3. Double-click the workload you want to modify to open its Transformation dialog. You can alternatively select the workload, and then click **Edit**.
4. In the Transformation dialog, click **Withdraw**.
5. If you are prompted for permission to remove the target VM, do one of the following:
   - Click **Yes** to initiate and automated clean up the target VM on the VM host and the workload job on the Migrate server.
   - Click **No** if you want to manually clean up the target VM on the VM host and the workload job on the Migrate server.

To withdraw one or more workloads:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. Select one or more submitted workloads that you want to withdraw.
3. Click **Bulk Status Change**.
4. Click **Withdraw**.
5. If you are prompted for permission to remove the target VM, do one of the following:
   - Click **Yes** to initiate and automated clean up the target VM on the VM host and the workload job on the Migrate server for each of the selected workloads.
   - Click **No** if you want to manually clean up the target VM on the VM host and the workload job on the Migrate server for each of the selected workloads.

Deleting a Workload

The Delete option allows a user with Workload management permissions to delete the workload from the project. Deleting a licensed workload does not free the consumed license.

Before you execute a workload transformation or migration, you might want to delete a workload if the workload is no longer exits or is no longer a candidate for transformation or migration. For example, the workload is no longer available if the hardware fails or if you decommission the workload prior to the transformation project. The workload is no longer a candidate if you determine that some workloads will remain in the current facility.

However, if only the profile for the source workload has changed, you should re-import the source workload information instead of deleting it. You can use the same FQDN in the Bulk Import spreadsheet to re-import the source workload with the modified data. After you have edited a proposed workload, a re-import modifies only the source workload information. See “Re-Importing Workloads” on page 186.

To delete a workload from a project:

1. In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.
2. In the **Workloads** list, use the Filter and Advanced Search to locate the appropriate workload.
3. Select the workload, then click **Delete**.
4. Click **Yes** to confirm the deletion.

C Bulk Import Spreadsheet

Users Who Can Perform This Task: Project Manager, Project Architect, or any user in the Administrators group

A large-scale data center transformation or migration project can involve hundreds or thousands of workloads. Importing data from a spreadsheet allows you to leverage data you might already have in an asset database.

The information the you include in your spreadsheet depends on your import goals:

- **Planning Only:** PlateSpin Transformation Manager allows you to plan for a variety of transformation goals, migration tools, and environments. You can import or re-import data for workloads from the same or different spreadsheets to define the original workload.
- **Automated Discovery:** In a PlateSpin Discovery Environment, your import spreadsheet needs to include only the minimal access information for each workload and values for Credentials. The import begins an automated discovery to retrieve detailed information for each new workload. Before you import the spreadsheet, ensure that your setup meets the following requirements:
  - A PlateSpin Migrate Connector instance must be deployed and running in the network where the source workloads reside.
  - The Connector instances must be configured to work with your PTM Server.
  - The Connector instances must be available to your project.
  - The Connector supports automated discovery only for physical and virtual machines.

For each source workload, the import automatically creates an original workload object for your project. It also create an initial proposed workload based on those original settings.

This section describes the bulk import process, and how to prepare a specially formatted spreadsheet that contains data about the machines to import for a project.

About Spreadsheet Import

PlateSpin Transformation Manager supports the bulk import of multiple machines at a time to your transformation project. *Bulk import* refers to loading data about the machines from a specially formatted Microsoft Excel 97-2003 (`.xls`) spreadsheet file to a previously defined project.

For each newly imported machine, the import automatically configures its original workload and an initial proposed workload. The workloads are then available in the Web Interface for further association with the project resources, users, and schedules. After you edit a proposed workload by using the Transformation dialog or the Bulk Edit options, any re-imported workload information applies only to the original workload.

The spreadsheet allows you to import a variety of data about each machine. PlateSpin Transformation Manager requires only the discovery IP address or FQDN (fully qualified domain name) and the type of machine (physical, virtual, or cloud). Predefined and custom parameters allow you to specify more details about the workloads, depending on your needs.

NOTE: If a PlateSpin Migrate Connector is available to the project, the successful import of workloads initiates an automated discovery process that retrieves details about the source workloads. For more information about auto-discovery, see “Import and Discovery Process” on page 138 and “Discovery Requirements for Source Workloads” on page 137.

You use your own spreadsheet software to create the Bulk Import spreadsheet. You can leverage the data in your network inventory records by mapping and exporting data to one or more versions of the Bulk Import spreadsheet.

Spreadsheet Import Tools

The Spreadsheet Import function in the PlateSpin Transformation Manager Web Interface provides the following tools to help you validate the spreadsheet for import:

- **Validate:** The Validate tool checks the spreadsheet format to identify the types of data it contains, performs limited checks of values, and identifies errors.

  Transformation Manager restricts what values are possible for some parameters, and expects text or numbers for others. It also checks for dependencies among the parameters to ensure all expected values are present. Otherwise, it accepts the values you input for the parameters. It does not validate the values for a machine against the actual machine in your network environment.

  For re-imported machines, the Validate tool also check the current workflow state for the machine. It ignores rows for machines that have already been submitted for transformation. You can withdraw the machine from transformation to make it eligible again for re-import.

  You can re-submit the spreadsheet iteratively until all rows of data are valid, and the spreadsheet is ready for import.
- **Validation Results:** The Validation Results reporting tool provides the Validation Status for each machine in the spreadsheet. For each invalid row, it identifies the type and location of errors. It also identifies which rows are ignored. You can view results for all rows, valid rows, invalid rows, and ignored rows.

- **Get Results:** The Get Results tool allows you to download a Results spreadsheet in Microsoft Excel 97-2003 (`.xls`) format. In addition to the Validation Status for each machine, it adds the following information if the column is present:
  - Workload ID
  - NIC UUID
  - Disk UUID

Spreadsheet Import Workflow

The Spreadsheet Import workflow, shown in Figure C-1, includes the tasks you perform to prepare the data for import, as well as the tasks required to validate that data before you can import it to your PlateSpin Transformation Manager environment.

Figure C-1 *Spreadsheet Import Workflow*

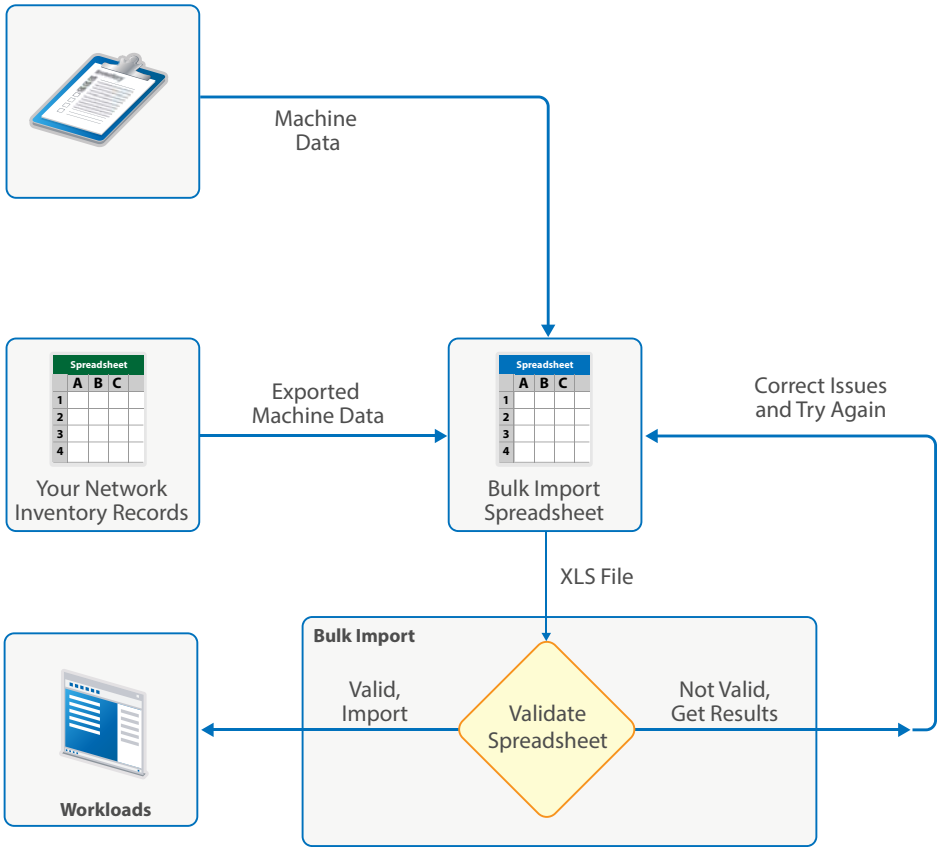

To import or update data for one or multiple machines in a project:

1. Collect data about the machines for a single transformation project in a spreadsheet, and save it in Microsoft Excel 97-2003 (`.xls`) format.

   A template for this specially formatted spreadsheet is available in the Spreadsheet Import dialog.
2. Open the Spreadsheet Import dialog.
3. Select the project, then browse and select the spreadsheet you prepared.
4. Validate the spreadsheet.
5. View the validation results by all rows, valid rows, invalid rows, or ignored rows.

6. If rows are invalid, download an annotated Results spreadsheet. PTM adds validation information that can help you to troubleshoot issues.
7. For each invalid row, correct the errors according to the various validation messages, then clear its Validation Status cell.
8. Repeat Step 2 to Step 7 iteratively until all rows are valid.
9. After all rows are valid, import the spreadsheet.

   For each machine with an empty Validation Status cell, PTM does the following:

   a. Creates or updates the original workload and the proposed workload (if it is unedited) for the machine.
   b. (Connector) If a PlateSpin Migrate Connector is enabled and active for the project, automated discovery retrieves additional information about each imported workload if you provide credentials in the spreadsheet. Discovery status is available in the Workloads list.
   c. Generates the Results spreadsheet and includes the following information, as appropriate:
      - Machine data
      - Validation Status of **Imported**
      - Workload ID of the original workload
      - NIC UUID for each NIC instance
      - Disk UUID for each disk instance
10. (Optional) Before you close the Spreadsheet Import dialog, download the Results spreadsheet.
11. (Optional) Re-import data for a machine from the same or different Bulk Import spreadsheet. A re-import can add new data or modify previously submitted data.

Spreadsheet Parameters

The Bulk Import spreadsheet contains the following information for one or more workloads in your transformation project:

- Required workload data
- Workload credentials
- Workload details
- Location
- Custom fields
- NICs
- Disks
- Applications

You can download a sample spreadsheet from the Spreadsheet Import dialog (**Planning** > **Workloads** > **Spreadsheet Import**) in the Web Interface. Mouse over the red triangle in a heading cell to display help about the purpose and format of the data in the column and whether it is required as part of the mapping process.

See the following tables for information about the expected values and format for each parameter in the Bulk Import spreadsheet:

- Table C-1, “Required Workload Parameters,” on page 177
- Table C-2, “Workload Credentials Parameters,” on page 178

Required Workload Parameters

The spreadsheet requires only three columns to set up a workload: FQDN, Type, and Validation Status. See Table C-1. This flexibility requires that you provide a header row with proper column names. These three columns must be present in all re-imports for the workloads. If you also include the Workload ID column, PTM automatically populates the IDs in the Results spreadsheet.

NOTE: If you prepare a spreadsheet without column headings, all defined columns must be present in the expected order indicated in the sample spreadsheet.

Table C-1 *Required Workload Parameters*

Parameter	Required/Optional	Description

IP or FQDN	Required for all imports.	Specify the IP address or fully qualified domain name for the workload. This parameter might be used to discover and communicate with the workload. You can specify the hostname only if you do not need discovery or validation.
		Examples:
		`sqlserver5.paris.digitalair.com`
		`sqlserver5`
		`192.168.1.1`
		A workload can have only one instance in a transformation project.

Validation Status	Required for all imports. Values are provided by PTM.	Transformation Manager writes the status for each workload to the Results spreadsheet after validation or after successful import. Clear the field for import.
Workload ID	Optional; recommended for all imports. Values are provided by PTM.	If this column is present, PTM writes the original workload’s unique workload ID to the Results spreadsheet after a workload is successfully imported to a project. Example: `1635326`
Type	Required for all imports.	Specify the workload type. Acceptable values are `CLOUD`, `PHYSICAL`, or `VIRTUAL`.

Workload Credentials

If you use PlateSpin Migrate for workload migrations, you can specify the Credentials resource to use for communications with the Migrate server. If you also use PlateSpin Migrate Connector, PTM uses Credentials resources for workload discovery, host discovery, and migration execution. See Table C-2.

NOTE: A valid Credentials resource is required for automated discovery. Ensure that you provide the Workload Credential Name of an existing, valid Credentials resource, or provide all three credentials fields to create a new one.

After you create a Credentials resource, you cannot modify the user name and password for it by modifying the value in the spreadsheet and re-importing the workload. Use the Edit Credential dialog in the Web Interface to modify the settings for the resource. See "Editing a Credential" on page 196.

Table C-2 *Workload Credentials Parameters*

Parameters	Required/Optional	Description
Workload Credential Name	Optional	Specify a name for the credential that is unique within your project. If no matching value is found in the Credentials list, PTM creates a new Credentials resource. Names are case insensitive. If you specify a workload credential user name and password without providing a credential name, PTM uses the FQDN as the credential name.

Workload Credential User Name	Optional (Required for a new Credentials resource)	Specify the logon user name of a domain administrator or local administrator to use for communications with the workload.
		
		Provide the user name exactly as it is expected by the workload operating system and environment:
		
		◆ **Windows:** Case insensitive. Use local or domain administrator credentials. For example:
		  ◆ For domain member machines: `CORPDOM\username`
		  ◆ For workgroup member machines: `WORKGROUP\username`
		  ◆ For a local user account: `username`
		◆ **Linux:** Case sensitive. Use `root` or a root-level user name for the workload that has been properly configured to use `sudo`. See *KB Article 7920711* (https://www.netiq.com/support/kb/doc.php?id=7920711).

Workload Credential Password	Optional (Required for a new Credentials resource)	Specify the password of the specified user name. You can add or modify the password for the credential later from the Web Interface. The password is case sensitive. Passwords are stored securely and are not visible in the Web interface.

Workload Details

Workload details define the operating system and hardware for the source workload. These details are automatically populated for the target workload. You can later modify the details for the target workload by using the Transformation dialog or the Bulk Edit action in the Web Interface. See Table C-3.

You can create source resources for Hosts, Networks, and Datastores by including those columns in the spreadsheet. Include the Environment column to create a resource that can apply to a source workload or target workload.

NOTE: A source resource name must be unique among source resources of the same type. The Environment resource name must be unique among environment resources for the project.

The resource name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Table C-3 *Workload Details Parameters*

Parameters	Required/Optional	Description

Operating System	Optional	Specify an operating system from the Operating Systems list configured in PTM. Names are case sensitive.
		
		Examples:
		
		`CentOS 6`
		`Citrix XEN Server`
		`Red Hat Enterprise Linux 7`
		`Solaris 11`
		`SUSE Linux Enterprise Server 12`
		`Windows Server 2003 (x32)`
		`Windows 2008`
		`Windows Server 2012 R2`
		`Windows Server 2016`
		`VMware ESX`
		
		PTM includes most operating system types and release versions in the default Operating Systems list. Contact your System Administrator to add custom operating system values to the list.

Memory (MB)	Optional	Specify the maximum memory available for the workload. Values default to MB if no unit is specified. Valid binary units are B, MB, GB, TB, PB, EB, ZB, and YB. Examples: `2 GB`, `512`, `8GB`
# of Processors	Optional	Specify the number of processors installed on the workload. Example: 2
# of Cores	Optional	Specify the number of cores per processor. Example: 8

Location

Location information can describe the logical, virtual, physical locations of the workload. See Table C-4.

Table C-4 *Location Parameters*

Parameter	Required/Optional	Description
Environment	Optional	Specify the name of the Environment resource for this workload. Values are case sensitive. If no matching value is found, PTM creates a new Environment resource for the project. Examples: `staging`, `preproduction`, and `production` **NOTE:** The Environment resource name must be unique among environment resources for the project, but it can apply to multiple workloads.
Cloud Provider	Optional; recommended for Cloud workloads	Specify the cloud provider that will host a cloud workload. Acceptable values are: ◆ `Amazon Web Services` ◆ `Microsoft Azure` ◆ `Rackspace` ◆ `vCloud` ◆ `Google`
Host	Optional	For a virtual workload, specify the name of the source Host resource for the source workload. Values are case sensitive. If no matching value is found, PTM creates a new source Host resource for the project. Examples: `mailhost`, `hyperv-cluster`, `esx2`, `citrix5.digitalair.com` **NOTE:** The source Host resource name must be unique among source Host resources for the project, but it can apply to multiple workloads.
Hypervisor	Optional	For a virtual workload, specify the hypervisor type running on its source Host resource. Names are case sensitive. Acceptable values are: ◆ `Citrix XenServer` ◆ `Linux KVM` ◆ `Microsoft Hyper-V` ◆ `SUSE Xen` ◆ `VMware`
VM Name	Optional	For a virtual workload, specify the VM name known to its virtualization host server. Examples: `vm-erm2`, `vm-msx-mailbox`, `vm-msx-web`

Parameter	Required/Optional	Description
Site	Optional	Specify the geographical or logical location of the host machine. Examples: `Paris`, `Metro General Hospital`, `South Campus`
Hardware Type	Optional	Specify the hardware for the workload's host machine (that is, the physical machine or a VM's virtualization host machine). Examples: `Dell PowerEdge`, `Levono X`, `IBM BladeCenter`
Enclosure	Optional	Specify the enclosure for the host machine. Examples: `AA-21`, `C205-R5`, `IL6-C5-R3`
Enclosure Slot	Optional	Specify the slot number in the specified enclosure. Examples: `U19`, `B07`
Is Domain?	Optional	(OS: Windows family) Boolean value must be either `TRUE` or `FALSE`. Specify `TRUE` if the Windows workload is a member of a domain. Specify `FALSE` if the Windows workload is a member of a workgroup.
Domain/Workgroup Name	Optional	Specify the name of the domain or workgroup for the Windows workload, as appropriate. Domain name examples: `CORPDOM`, `PARIS`, or `paris.digitalair.com` Workgroup name example: `WORKGROUP`

Custom Fields

The Spreadsheet Import expects the headings of the **Custom X** columns to match the custom field names you have previously defined for the project. Each project can have up to 7 user-defined fields that apply project-wide. See Table C-5.

You can associate logical names and purpose with the custom fields for each project. For example, if you want to identify the workload with its day-to-day IT administrator, you might define a **Field 1** name of **Contact** for the project. In the spreadsheet, you use **Contact** instead of **Custom 1** as the column heading. Specify values for the field as appropriate for your project, such as a contact person's name, user name, or email address.

Table C-5 *Custom Field Parameters*

Parameters	Required/Optional	Description
Custom 1	Optional	Specify the project's Field 1 name as the column heading. For each workload, specify the value based on the project's Field 1 definition.
Custom 2	Optional	Specify the project's Field 2 name as the column heading. For each workload, specify the value based on the project's Field 2 definition.

Parameters	Required/Optional	Description
Custom 3	Optional	Specify the project's Field 3 name as the column heading. For each workload, specify the value based on the project's Field 3 definition.
Custom 4	Optional	Specify the project's Field 4 name as the column heading. For each workload, specify the value based on the project's Field 4 definition.
Custom 5	Optional	Specify the project's Field 5 name as the column heading. For each workload, specify the value based on the project's Field 5 definition.
Custom 6	Optional	Specify the project's Field 6 name as the column heading. For each workload, specify the value based on the project's Field 6 definition.
Custom 7	Optional	Specify the project's Field 7 name as the column heading. For each workload, specify the value based on the project's Field 7 definition.

NICs

You can identify one or more NICs for a workload. You can specify any one of the available NIC input parameters to set up a NIC. See Table C-6. If you also include the **NIC1 UUID** column, PTM automatically populates the UUIDs in the Results spreadsheet. You can set up multiple NICs by adding additional NIC# column instances.

Table C-6 *NIC Parameters*

Parameters	Required/Optional	Description
NIC1 UUID	Optional; recommended for NIC. Values are provided by PTM.	The UUID is a set of hexadecimal numbers that uniquely identify the NIC in the project. If this column is present, PTM writes a unique ID for each NIC to the Results spreadsheet after a workload is successfully imported to a project, On a re-import, PTM keeps the same ID, and replaces the NIC's configuration information for the original workload. Example: `b916afef-ccfc-4501-aa4c-a4658cd9845d`
NIC1 Interface Name	Optional	Specify a unique interface name on the workload. Examples: `eth0`, `eth1`
NIC1 IP DHCP	Optional	Specify whether a DHCP server provides IP addresses for the NIC. Boolean value must be either `TRUE` or `FALSE`.
NIC1 IP Address	Optional	Specify an address for the NIC in IPv4 format. Example: `10.10.10.40`
NIC1 MAC	Optional	Specify the MAC address for the source workload NIC in the MM:MM:MM:SS:SS:SS format. Example: `01:23:45:67:89:ab`

Parameters	Required/Optional	Description
NIC1 Network Name	Optional	Specify the name of the source Network resource for the source workload. Values are case sensitive. If no match is found, PTM creates a new source Network resource for the project. Examples: `OrgNet-DMZ`, `OrgNet-Isolated`, `Azure-VNet-01`, `VMnet-08` **NOTE:** The source Network resource name must be unique among source Network resources for the project, but it can apply to multiple NICs on the same or different workloads.
NIC1 VLAN ID	Optional	Specify the assigned VLAN ID for this NIC. Providing the ID makes the validation and import functions run faster. Valid values are 1 to 4094.
NIC1 Network Mask	Optional	Specify the network mask in IPv4 format. Example: `255.255.252.0`
NIC1 Gateway	Optional	Specify the default gateway address in IPv4 format. Multiple entries are not supported.
NIC1 DNS DHCP	Optional	Specify whether DNS DHCP is available for the NIC. Boolean value must be either `TRUE` or `FALSE`.
NIC1 DNS Servers	Optional	Specify one or more addresses in IPv4 format. For multiple entries, place each value on a separate line.
NIC1 Search Domains	Optional	Specify one or more search domains. For multiple entries, place each value on a separate line.

Create additional NIC# sections as appropriate, then continue to the next component section.

A workload can have multiple network cards, each of which has a unique network configuration. The first defined network card is designated as NIC1, the second as NIC2, and so on.

Disks

You can identify one or more disks for a workload. See Table C-7. A disk requires only the **Disk1 Vol Type** and **Disk1 Size** columns. If you also include the **Disk1 UUID** column, PTM automatically populates the UUIDs in the Results spreadsheet. You can set up multiple disks by adding additional Disk# column instances with the minimum of **Disk# Vol Type** and **Disk# Size** columns.

Table C-7 *Disk Parameters*

Parameters	Required/Optional	Description
Disk1 UUID	Optional; recommended for Disks. Values are provided by PTM.	The UUID is a set of hexadecimal numbers that uniquely identify the disk in the project. If this column is present, PTM writes a unique ID for each disk to the Results spreadsheet after a workload is successfully imported to a project. On a re-import, PTM keeps the same ID, and replaces the disk's configuration information for the original workload. Example: `866b4209-4779-48f8-a81d-e8aa04c63d55`

Parameters	Required/Optional	Description
Disk1 Datastore Name	Optional, applies to Virtual workloads	Specify the name of the source Datastore resource that contains the disk. Values are case sensitive. If no matching value is found, PTM creates a new source Datastore resource for the project. Example: `san-dc-east` **NOTE:** The datastore name must be unique among source Datastore resources for the project. A datastore can apply to multiple disks on the same virtual workload or on different virtual workloads.
Disk1 Vol Name	Optional	Specify the unique volume name on the workload.
Disk1 Vol Type	Optional; required for Disk	Acceptable values are `System`, `Local`, `NAS`, or `SAN`.
Disk1 Mount Point	Optional	Specify the path from the root of the volume.
Disk1 Remote Path	Optional	(Volume Type: SAN or NAS) Specify the remote path for the volume. Example: `\\vol1\mnt`
Disk1 Size	Optional; required for Disk	Specify the total size of the disk. Values default to MB if no unit is specified. Valid binary units are B, MB, GB, TB, PB, EB, ZB, and YB. Examples: `2TB`, `500 GB`, `102400`
Disk1 Used	Optional	Specify the space used on the disk. Values default to MB if no unit is specified. Valid binary units are B, MB, GB, TB, PB, EB, ZB, and YB. Examples: `1.22TB`, `162 GB`, `51200`

Create additional Disk# sections as appropriate, then continue to the next component section.

A workload can have multiple disks, each of which has its own configuration. The first defined disk is designated as `Disk1`, the second as `Disk2`, and so on.

Applications

You can identify one or more applications for a workload. See Table C-8. Add an **Application#** column instance for each application you want to track for the workload.

Table C-8 *Application Parameter*

Parameter	Required	Description
Application1 Name	Optional	Only the name of the application is needed. The application must be available to the workload. Values are case sensitive. If no matching value is found, PTM creates a new Application resource for the project. Example: `Paris-Project2-Exchange` **NOTE:** An Application name must be unique among Applications for the project, but it can apply to multiple workloads.

Create additional Application# sections as appropriate.

A workload can have multiple applications or services associated with it.The first defined application is designated with a column heading of `Application1 Name`, the second column heading as `Application2 Name`, and so on.

Downloading a Sample Import Spreadsheet

To download a sample spreadsheet:

1. Log in to the Web Interface as the Project Manager user, or as any system user.
2. Select **Planning > Workloads > Spreadsheet Import** to open the Spreadsheet Import dialog.
3. Click **Get sample spreadsheet**, then save the file to your local machine.

   The spreadsheet is in `.xls` file format.

Validating a Spreadsheet

Before you submit the Bulk Import spreadsheet, clear the contents in the Validation Status cell for each machine that you want to validate for import (or re-import).

The Validate tool displays the following states for machines and adds the information to the Validation Status column of the Results spreadsheet:

- **Ready for Import:** The row contains valid information.

  The Validate tool found no errors in the data submitted for the machine. All values are valid options and in the expected format. All required and dependent values are present.
- **Invalid:** The row contains invalid information.

  The Validate tool found one or more errors in the data submitted for this machine. The cell also includes messages about the type and location of the errors.
- **Imported:** The row was successfully imported or re-imported.

On import, PTM creates (or modifies) the original and proposed workloads for this machine in the specified project. If you have edited the proposed workload by using the Web Interface, a re-import modifies data only for the original workload.

- **Ignored:** The row was previously imported for a workload that has already been submitted for transformation.

  You can use the Transformation dialog to withdraw the machine from transformation to make it eligible again for re-import.

Downloading the Results Spreadsheet

After you perform a validation or an import of the Bulk Import spreadsheet, the **Get Results** option becomes available on the Spreadsheet Import dialog.

1. On the Spreadsheet Import dialog, click **Get Results**, then save the file to your local machine.

   The spreadsheet is in `.xls` file format.

After a validation, the Results spreadsheet reports the import readiness of each workload.

After a successful import, the Results spreadsheet contains the exact results of the current state of the import process, including:

- The deletion of blank rows
- More verbose results of the validation status
- The Workload ID assigned to each workload (if the **Workload ID** column is present in the input Bulk Import spreadsheet)
- The UUID assigned to each NIC (if the **NIC# UUID** column is present in the input Bulk Import spreadsheet)
- The UUID assigned to each disk (if the **Disk# UUID** column is present in the input Bulk Import spreadsheet)

Re-Importing Workloads

You can re-import workloads multiple times to separately set up the various categories of data. This partial import capability allows you to import a workload's full details from multiple sources, each with its own import spreadsheet. Each re-import adds details for the original workload and for the proposed workload.

For example, you can import the minimal data to create the workload. Subsequent re-imports can use separate spreadsheets to add workload details, location, NICs, disks, applications, and custom fields. Alternatively, you can continue to work in the same spreadsheet, adding the appropriate columns for each re-import at the end of the spreadsheet.

NOTE: After you edit the proposed workload by using the Web Interface, the proposed workload is no longer updated by any subsequent re-imports.

If you attempt to re-import a workload that is in a workflow state of Submitted or a later phase, then the Validation Status returns a message:

```
Cannot re-import workload after transformation has been submitted.
```

The import ignores workloads in the Submitted or later state, and continues unless there are invalid rows. Rows that are in the ignored state appear in the Ignored Rows tab of the Validation Results.

V Resources

PlateSpin Transformation Manager allows you to define the resources to track for workload transformations. Resources include credentials, hosts, networks, datastores, environments, migration servers, and resource pools.

20 Overview of Resources

PlateSpin Transformation Manager allows you to define the resources to track for workload transformations. Resources include credentials, hosts, networks, datastores, environments, migration servers, and resource pools.

About Resources

Resources are elements of your IT infrastructure that are external to the workloads to be transformed. These resources represent the operational and replication environment for the workload. Resource types include:

- Credentials
- Hosts
- Migration servers
- Networks
- Datastores
- Resource pools
- Environments

You can use the Resources feature to create, manage, and delete information about resources for your project. In a Spreadsheet Import, the values for host, network, and datastore create source resources, and values for credentials and environments apply for both sources and targets.

Prerequisites for Resources

When you create a resource, you must associate it with a specific transformation project. Before you can add resources for a project, you must create the parent organization and project to ensure that they are available when you create the resources for a project.

21 Managing Credentials Resources

Credentials resources allow you to track the authentication credentials that are needed to log in as an administrator user to the related location:

- Workloads
- Target hosts
- Migration servers

You can use the Credentials page to create, manage, view, and delete the Credentials resources for the target workloads in your transformation project. Credentials do not apply to source workloads.

About Credentials Resources

The Credentials resource allows you to track the following information for authentication credentials:

Name: Specify a name for the Credentials resource that is unique to your transformation project. The name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Organization: Select the parent organization for the project.

Project: Select the parent project for the credential.

Description: (Optional) Specify a brief description of the credential. This text appears by default in the Credentials list.

Username: Specify the logon user name of a domain administrator or local administrator to use for communications with the workload.

Provide the user name exactly as it is expected by the intended workload operating system or authentication system:

- **Windows:** Case insensitive. Use local or domain administrator credentials. For example:
  - For domain member machines: `CORPDOM\username`
  - For workgroup member machines: `WORKGROUP\username`
  - For a local user account: `username`
- **Linux:** Case sensitive. Use `root` or a root-level user name for the workload that has been properly configured to use `sudo`. See *KB Article 7920711* (https://www.netiq.com/support/kb/doc.php?id=7920711).

Password: Specify the password for the user name.

Confirm Password: Re-type the password.

Credentials resources facilitate authentication for the associated components. The passwords are stored securely in the PlateSpin Transformation Manager database. Passwords are obscured in the interface.

NOTE: Credentials resource settings are not integrated with your directory password management system. If the password changes in your network environment or for accounts in cloud provider locations, you must manually update the password stored for the resource.

Viewing Credentials

All roles for a project can view the Credentials resource information. The password for the user name is always obscured.

1. In the Web Interface toolbar, select **Resources**, then select the **Credentials** tab.

   You can also view Credentials resources for a single project in the **Associations** area of the Edit Project dialog. Select **Planning** > **Projects**, double-click the project. In the Edit Project dialog under **Associations**, select the **Resources** tab, then select **Credentials**.
2. On the **Credentials** tab, view the list of resources.
3. Select a Credentials resource, then click **View** (or double-click the resource) to view the following settings for the resource:
   - Name
   - Organization
   - Project
   - Description
   - Username
   - Password
   - Confirm Password
4. (Optional) Filter entries in the Credentials list using the **Search** field to show only the resource pools of interest. Type a sequence of characters in the **Search** area to initiate filtering. The filter applies as you type. You can filter on information in any one of the columns at a time.
5. (Optional) Sort entries by the values in a column. Mouse over the column to expose the menu option, click it, then select **Sort Ascending** or **Sort Descending.**

Creating a Credential

You can use the Credentials page to create Credentials resources for the target transformation environment of a project.

1. In the Web Interface toolbar, select **Resources**.
2. Select the **Credentials** tab.
3. Click **Create** to open the Create Credential dialog.
4. Specify the following information for the Credentials resource:
   - Name

- Organization
- Project
- Description
- Username
- Password
- Confirm Password

5 Click **Save** to create the Credentials resource.

6 Click **Close** to exit the dialog.

Associating Credentials with Workloads, Hosts, or Migration Servers

After you create Credentials resources for a project, you can associate them with workloads, hosts, or migration servers, as appropriate.

- "Associating Credentials and Workloads" on page 195
- "Associating Credentials and Hosts" on page 195
- "Associating Credentials and Migration Servers" on page 196

Associating Credentials and Workloads

After you create Credentials resources for a project, you can associate them with workloads.

1 In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.

2 In the **Workloads** list, use the Filter and Advanced Search to locate the appropriate workload.

3 Select the workload, then click **Edit**. You can alternatively double-click the workload of interest.

4 On the Transformation dialog, scroll to the **Workload** panel, then click **Edit**.

5 From the **Admin Credentials** list, select the Credentials resource that you want to assign to the target workload.

6 (Windows workloads) In a Domain, select **Domain** under **Windows Settings**, specify the domain name, then select the Credentials resource to use for the domain in the **Domain Credential** list.

7 (Optional) Modify other Workload information as appropriate.

8 Click **Save**, then click **Close**.

9 Click **Close** to exit the Transformation dialog.

Associating Credentials and Hosts

After you create Credentials resources for a project, you can associate them with host resources.

1 In the Web Interface toolbar, select **Resources**.

2 Select the **Hosts** tab.

3 Select the host, then click **Edit**. You can alternatively double-click the host of interest.

4 From the **Credential** list, select the Credentials resource that you want to assign to the target host.

5 (Optional) Modify other Host information as appropriate.

6 Click **Save**.

7 Click **Close** to exit the dialog.

Associating Credentials and Migration Servers

After you create Credentials resources for a project, you can associate them with migration servers.

1 In the Web Interface toolbar, select **Resources**.

2 Select the **Migration Servers** tab.

3 Select the migration server, then click **Edit**. You can alternatively double-click the migration server of interest.

4 From the **Credential** list, select the Credentials resource that you want to assign to the target migration server.

5 (Optional) Modify other Migration Server information as appropriate.

6 Click **Save**.

7 Click **Close** to exit the dialog.

Editing a Credential

You can use the Credentials page to edit a Credentials resource for the target transformation environment of a project. For example, you might need to modify the user name and password assigned to the credential, or change the password for the existing user name.

1 In the Web Interface toolbar, select **Resources**.

2 Select the **Credentials** tab.

3 In the **Credentials** list, select the credential, then click **Edit** (or double-click the credential) to open the Edit Credential dialog.

4 Modify any of the following information for the Credentials resource:

- Name
- Organization
- Project
- Description
- Username
- Password
- Confirm Password

5 Click **Save** to update the Credentials resource.

6 Click **Close** to exit the dialog.

Deleting a Credential

You can use the Credentials page to delete a Credentials resource for the target transformation environment of a project.

NOTE: If you delete a resource, it automatically removes the associations for the resource, including those assigned to workloads that are in progress and completed. If the same associations apply, consider whether it is more appropriate to edit the resource name and its user name and password.

1. In the Web Interface toolbar, select **Resources**.
2. Select the **Credentials** tab.
3. View the list of Credentials resources.
4. Select the appropriate credential, click **Delete**, then click **Yes** to confirm the deletion.
5. Click **Close** to exit the dialog.

22 Managing Host Resources

Host resources allow you to track information about the host machines for the workload VMs. You can use the Hosts page to create, edit, view, and delete host resources for the target workloads in your transformation project. You can define host resources for the source workloads when you import machine their machine data. A host resource can be used for source workloads or for target workloads, but not to both. A host resource for a physical machine can be assigned to only a single workload.

About Automated Host Discovery

If a PlateSpin Migrate Connector is available to the project, PlateSpin Transformation Manager provides automated discovery for target VMware Cluster hosts. Discovery retrieves details about the host, including its networks, datastores, and resource pools. If the networks, datastores, and resource pools do not match existing resources, it creates resources for them. You can also rediscover details for a host if necessary.

Discovery Requirements for Target Hosts

Before you attempt discovery for a VMware Cluster host, you must configure your PlateSpin Discovery Environment:

- **PlateSpin Transformation Manager 1.1**

  Configure PlateSpin Transformation Manager and set up a project for the target hosts you want to discover.

- **PlateSpin Migrate Connector 1.1**

  A PlateSpin Migrate Connector must be available to the project. For deployment information, see "Deployment Requirements" in the *PlateSpin Migrate Connector Quick Start*.

PlateSpin Transformation Manager Appliance includes a Migrate Connector instance that is pre-configured to work with the PTM Server. You can also deploy Migrate Connector instances separately on your own SUSE Linux Enterprise Server 11 SP4 servers. See the *PlateSpin Migrate Connector Quick Start*.

NOTE: Ensure that the Migrate Connector is up and running before you attempt discovery or rediscovery of target hosts.

- **Target Hosts**
  - The Connector supports automated discovery of VMware Cluster Hosts.

    For other target host types, you can use the Spreadsheet Import method to add host details.
  - Ensure that the target VMware Cluster hosts are up and running.
  - Ensure that the network connections are working between the PlateSpin Migrate Connector and the target VMware Cluster host.
- **Credentials**

  Create a Credentials resource to use for host discovery that provides the administrator-level credentials needed to access the target VMware host. The Credentials resource is also used later for automated migration setup of the target VM on the host.
- **Network Connectivity and Access for Target Hosts**

  For information about network connectivity and access requirements for target host discovery, see “PlateSpin Discovery Environment” on page 134.

Discovery Process for Hosts

The automated host discovery process performs the following actions:

1. The Create Host dialog creates a new Host object unless it matches an existing host.
2. Auto-discovery begins for the host, using the provided Credentials resources to log in to the host.
3. If discovery details for the host include networks, datastores, or resource pools, PTM creates new resources for each one in the project unless they match an existing resource.
4. On successful host discovery, the Hosts list or Edit Host dialog displays the host information.

Troubleshooting Host Discovery Failures

Discovery fails for any of the following conditions:

- A Migrate Connector is not available.
- The target host is not supported for automated discovery.
- The Credentials resource is not provided.
- The Credentials resource has an invalid value for the user name and password.
- The host is not running at the time of discovery.

After you resolve Migrate Connector, network, or credentials issues, you can retry discovery for supported hosts. See “Retrying Workload Discovery” on page 142.

About Host Resources

The host resource allows you to track the following information for a server that hosts one physical workload, or that hosts the hypervisor for multiple virtual workloads:

Status: The discovery status for the target VM host. Discovery is available only when a PlateSpin Migrate Connector is available to the project and the target VM host is a VMware Cluster.

Name: Specify a name for the server that is unique to your transformation project. The name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Organization: Select the parent organization for the project.

Project: Select the parent project for the virtualization host.

Description: (Optional) Specify a brief description of the host. This text appears by default in the Hosts list.

IP/DNS Name: (Optional, required for discovery and automated migration) Specify an IP address in IPv4 format, or the DNS name of the host machine in the FQDN format.

Credential: (Optional, required for discovery and automated migration) Expand the Credentials list, then locate and select the Credentials resource to use for accessing the host machine. You can scroll the list, or type a sequence of characters in the Filter to locate the credentials of interest.

Hypervisor: Select the hypervisor type used by the host. Options include:

`Citrix XenServer`

`Linux KVM`

`Microsoft Hyper-V`

`None` (for physical workloads)

`SUSE Xen`

`VMware`

Cluster: (Optional) The name of the vSphere Cluster.

Hardware: (Optional) Specify a textual description of the hardware architecture for the host server. If a match is not found, PTM adds a new Hardware type in the database. Options include: x32, x64.

Site: (Optional) Specify a textual description of the facility or location where the host server resides. If a match is not found, PTM adds a new Site type in the database. Examples: `HQ data center`, `West campus`, `Sydney`.

Enclosure: (Optional) Specify a textual description of the enclosure where the host server resides, such as the rack cabinet, blade server chassis, or tower cabinet. If a match is not found, PTM adds a new Enclosure type in the database. Examples: `RC-10`, `BS-2`, `TC-5`.

Slot: (Optional) Specify a textual description of the slot where the host server resides in the specified enclosure. If a match is not found, PTM adds a new Slot type in the database. Examples: `R3-2U`, `Bay-2`, `Sh-1`.

Viewing Hosts

You can view the host information for transformation workloads that you manage. For information about the parameters used to define the Host, see “About Host Resources” on page 201.

1. In the Web Interface, select **Resources** in the toolbar, then select the **Hosts** tab.

   You can also view hosts for a single project in the **Associations** area of the Edit Project dialog. Select **Planning** > **Projects**, double-click the project, then under **Associations**, select the **Hosts** tab.
2. On the **Hosts** tab, do either of the following to view a list of resources:
   - Click the **Source** link to view host information that you imported for original workloads using Spreadsheet Import.
   - Click the **Target** link to view host information for the target transformation environment.
3. Select a host resource, then click **View** (or double-click the resource) to view the following settings for the resource:
   - Status
   - Name
   - Organization
   - Project
   - Description
   - IP/DNS Name
   - Credential
   - Hypervisor
   - Cluster
   - Hardware
   - Site
   - Enclosure
   - Slot
4. (Optional) Filter entries in the Hosts list using the Search field to show only the hosts of interest.

   Type a sequence of characters in the **Search** area to initiate filtering. The filter applies as you type. You can filter on information in any one of the columns at a time.
5. (Optional) Sort entries by the values in a column. Mouse over the column to expose the menu option, click it, then select **Sort Ascending** or **Sort Descending.**

Creating Source Hosts

You can create a host resource for a source workload by specifying the host information with the machine definition in the Bulk Import spreadsheet.

Creating a Target Host

You can use the Hosts page to create host resources for the target transformation environment of a project.

NOTE: If a PlateSpin Migrate Connector is available to the project and the target VM host is a VMware Cluster, auto-discovery creates related Network, Datastore, and Resource Pool resources if they do not already exist.

1. In the Web Interface toolbar, select **Resources**.
2. Select the **Hosts** tab, then click **Target**.
3. Click **Create** to open the Create Host dialog.
4. Specify the following information for the host resource:
   - Name
   - Organization
   - Project
   - Description
   - IP/DNS Name (required for discovery and automated migration)
   - Credential (required for discovery and automated migration)
   - Hypervisor
   - Cluster
   - Hardware
   - Site
   - Enclosure
   - Slot
5. Click **Save** to create the host resource.

   If auto-discovery is available, the host's Network, Datastore, and Resource Pool resources are automatically created.
6. Click **Close** to exit the dialog.
7. (Optional) Visit the Network, Datastore, and Resource Pool tabs to visually verify information about the newly discovered resources.

Retrying Host Discovery

You might want to retry discovery for a host if the initial discovery fails. For information about possible causes of discovery failure, see "Troubleshooting Host Discovery Failures" on page 200.

To retry discovery for one or more hots:

1. In the Web Interface, select **Resources** in the toolbar, then select the **Hosts** tab.
2. Use the Advanced Search to select one or more hosts where the initial discovery failed. In the **Sub Status** list, select **Target host discovery failed**.
3. Click in the Hosts list to exit the Advanced Search.
4. For each host, mouse over the **Target host discovery failed** status to view details about the reason the discovery failed.
5. Select one or multiple hosts in the refined list.
6. In the actions above the list, click **Retry** to initiate discovery for each of the selected hosts.

To retry discovery for a single host:

1. In the Web Interface, select **Resources** in the toolbar, then select the **Hosts** tab.
2. Use the Filter and Advanced Search to locate and select the host, then click **Edit**.
3. Mouse over **Retry** in the header area for information about why the initial discovery failed.
4. Click **Retry** to initiate discovery.

Rediscovering Hosts

You might want to rediscover a host if you modify the host platform, modify the network, add or remove datastores, or add or remove resource pools. Rediscovery automatically adds any new Network, Datastore, or Resource Pool resources; however, it does not remove the old resources.

To rediscover one or more hosts:

1. In the Web Interface, select to **Resources** > **Hosts**.
2. Select one or multiple hosts in the Hosts list.
3. In the actions above the list, click **Rediscover** to initiate rediscovery for each of the selected hosts.

To rediscover a single host:

1. In the Web Interface, select to **Resources** > **Hosts**.
2. In the Hosts lists, select the host, then click **Edit**.
3. In the header area, note the Sub Status **Discovered**.
4. Click **Rediscover** to initiate rediscovery.

Editing Target Hosts

You can use the Hosts page to edit a host resource for the target transformation environment of a project. For information about the parameters used to define the Host, see “About Host Resources” on page 201.

1. In the Web Interface toolbar, select **Resources**.
2. Select the **Hosts** tab, then click **Target**.
3. In the **Hosts** list, select the host, then click **Edit** (or double-click the host) to open the Edit Host dialog.
4. Modify any information for the host resource:
   - Name
   - Organization
   - Project
   - Description
   - IP/DNS Name (required for discovery and automated migration)
   - Credential (required for discovery and automated migration)
   - Hypervisor
   - Cluster

- Hardware
- Site
- Enclosure
- Slot

5 Click **Save** to update the host resource.

6 Click **Close** to exit the dialog.

Associating Hosts and Target Workloads

After you create target host resources for a project, you can associate them with workloads. For information about the parameters used to define the Host, see “About Host Resources” on page 201.

1 In the Web Interface, select **Planning** in the toolbar, then select the **Workloads** tab.

2 Select the workload that will be transformed to a virtual workload, then click **Edit**. You can alternatively double-click the workload of interest.

3 Verify that the transformation target workload is a virtual machine.

3a On the Transformation dialog in the Transformation Plan panel, click **Edit**.

3b In the Transformation Method list, select the appropriate method.

3c Click **Save**, and then click **Close**.

4 On the Transformation dialog in the **Location** panel, click **Edit**.

5 From the **Host** list, select the host resource that you want to assign to the target workload.

6 (Optional) Modify other Location information as appropriate.

7 Click **Save**, and then click **Close**.

8 Click **Close** to exit the Transformation dialog.

Deleting a Host

You can use the Hosts page to delete a host resource for the source or target transformation environment of a project. The deleted resource is automatically removed from each of its assigned workloads.

NOTE: If you delete a resource, it automatically removes the associations for the resource, including those assigned to workloads that are in progress and completed. If the same associations apply, consider whether it is more appropriate to edit the resource name and settings.

---

1 In the Web Interface toolbar, select **Resources**.

2 Select the **Hosts** tab.

3 (Optional) Delete a source host resource:

3a On the **Hosts** tab, click **Source** to view the list of Source host resources.

3b Select the appropriate host, click **Delete**, then click **Yes** to confirm the deletion.

4 (Optional) Delete a target host resource:

4a On the **Hosts** tab, click **Target** to view the list of Target host resources.

4b Select the appropriate host, click **Delete**, then click **Yes** to confirm the deletion.

23 Managing Migration Server Resources

Migration Server resources allow you to track and associate the servers that host the migration tools used by migration specialists to transform workloads. For example, a migration server might host the migration tool PlateSpin Migrate (https://www.netiq.com/products/migrate/). You can use the Migration Servers page to create, manage, and delete the Migration Server resources for the target workloads in your transformation project. Migration Server resources do not apply to source workloads.

About Migration Server Resources

The Migration Server resource allows you to track the following information for a migration server that resides in your transformation project:

Name: Specify a name for the migration server that is unique to your transformation project. The name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Organization: Select the parent organization for the project.

Project: Select the parent project for the migration server.

Description: (Optional) Specify a brief description of the migration server. This text appears by default in the Migration Servers list.

Migration Server Type: Select one of the following:

- **PlateSpin Migrate:** Select this option to identify the migration server as a PlateSpin Migrate Server, and to allow integration with your PlateSpin Migrate Server.
- **Others:** Select this option to identify that the migration server runs any other type of migration software. PTM does not provide integration.

Server URL: Specify the URL of the login page for the migration tool hosted on the migration server. The URL must conform to all syntax rules and can be up to 1024 characters. For example:

```
https://psmigrate.example.com/Migrate
```

Credential: Expand the Credentials list, then locate and select the Credentials resource to use to log in to the migration server's web-based management tool. You can scroll the list, or type a sequence of characters in the Filter to locate the Credentials resource of interest.

Migration Server resources can be saved in an incomplete state. The Migration Server field displays the value **Not Configured** in the Transformation dialog. When you later deploy the server, you can update the Migration Server resource. The status is updated automatically for all of its associated workloads.

Viewing Migration Servers

All roles for a project can view the migration server information.

1. In the Web Interface toolbar, select **Resources**, then select the **Migration Servers** tab.

   You can also view migration servers for a single project in the **Associations** area of the Edit Project dialog. Select **Planning** > **Projects**, double-click the project. In the Edit Project dialog under **Associations**, select the **Resources** tab, then select **Migration Servers**.
2. On the **Migration Servers** tab, view the list of resources.
3. Select a migration server resource, then click **View** (or double-click the resource) to view the following settings for the resource:
   - Name
   - Organization
   - Project
   - Description
   - Migration Server Type
   - Server URL
   - Credential
4. (Optional) Select the project in the **Global Project Selector** to list only the Migration Server resources for a single project.
5. (Optional) Filter entries in the Migration Servers list using the **Search** field to show only the migration servers of interest. Type a sequence of characters in the **Search** area to initiate filtering. The filter applies as you type. You can filter on information in any one of the columns at a time.
6. (Optional) Sort entries by the values in a column. Mouse over the column to expose the menu option, click it, then select **Sort Ascending** or **Sort Descending.**

Creating a Migration Server

You can use the Migration Servers page to create Migration Server resources for the target transformation environment of a project.

1. In the Web Interface toolbar, select **Resources**.
2. Select the **Migration Servers** tab.
3. Click **Create** to open the Create Migration Server dialog.
4. Specify the following information for the Migration Server resource:
   - Name
   - Organization
   - Project
   - Description

- Migration Server Type
- Server URL
- Credential

5 Click **Save** to create the Migration Server resource.

6 Click **Close** to exit the dialog.

Associating Migration Servers and Target Workloads

After you create Migration Server resources for a project, you can associate them with workloads. You can specify the Migration Server resource for each workload, or select Auto Assign to allow the PlateSpin Migrate Connector to load-balance the migration jobs across all of the project's available Migration Server resources of type PlateSpin Migrate.

In Automated Mode, imported workloads are automatically set to Auto Assign. You can set individual workloads separately. Auto Assign is disabled for a workload if the transformation method is not supported by automated migration. See "Workload Transformation Methods" on page 132.

1 In the Web Interface toolbar, select **Planning**.

2 Select the **Workloads** tab.

3 Select the workload, then click **Edit**. You can alternatively double-click the workload of interest.

4 On the Transformation dialog in the **Transformation Plan** panel, click **Edit**.

5 Do one of the following:

- **Migration Server:** From the **Migration Server** list, select the migration server resource that you want to assign to the target workload.
- **Auto Assign:** Select the Auto Assign check box to enable automatic assignment from among the available Migration Server resources that are of type PlateSpin Migrate.

6 (Optional) Modify other Transformation Plan information as appropriate.

7 Click **Save**, and then click **Close**.

8 Click **Close** to exit the Transformation dialog.

Editing a Migration Server

You can use the Migration Servers page to edit a Migration Server resource for the target transformation environment of a project.

1 In the Web Interface toolbar, select **Resources**.

2 Select the **Migration Servers** tab.

3 In the **Migration Servers** list, select the migration server, then click **Edit** (or double-click the migration server) to open the Edit Migration Server dialog.

4 Modify any of the following information for the Migration Server resource:

- Name
- Organization
- Project
- Description

- Migration Server Type
- Server URL
- Credential

5 Click **Save** to update the Migration Server resource.

6 Click **Close** to exit the dialog.

Deleting a Migration Server

You can use the Migration Servers page to delete a Migration Server resource for the target transformation environment of a project.

NOTE: If you delete a resource, it automatically removes the associations for the resource, including those assigned to workloads that are in progress and completed. If the same associations apply, consider whether it is more appropriate to edit the resource name and settings.

---

1 In the Web Interface toolbar, select **Resources**.

2 Select the **Migration Servers** tab.

3 View the list of migration server resources.

4 Select the appropriate migration server, click **Delete**, then click **Yes** to confirm the deletion.

24 Managing Network Resources

Network resources allow you to track information about the networks for the workloads. You can use the Networks page to create, manage, and delete networks for the target workloads in your transformation project. You can define network resources for the source workloads in the machine information for a bulk import. A network resource can apply to source workloads or to target workloads, but not to both.

About Network Resources

The network resource allows you to track the following information for a network used by the one or more workloads in your transformation project:

Name: Specify a name for the network that is unique to your transformation project. The name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Organization: Select the parent organization for the project.

Project: Select the parent project for the network.

Description: (Optional) Specify a brief description of the network. This text appears by default in the Networks list.

Viewing Networks

All roles for a project can view the network information.

1. In the Web Interface toolbar, select **Resources**, then select the **Networks** tab.

   You can also view networks for a single project in the **Associations** area of the Edit Project dialog. Select **Planning** > **Projects**, double-click the project, then under **Associations**, select the **Networks** tab.
2. On the **Networks** tab, do either of the following to view a list of resources:
   - Click the **Source** link to view network information that you imported for original workloads using Spreadsheet Import.
   - Click the **Target** link to view network information for the target transformation environment.

3 Select a network resource, then click **View** (or double-click the resource) to view the following settings for the resource:
   - Name
   - Organization
   - Project
   - Description
4 (Optional) Filter entries in the Networks list using the **Search** field to show only the networks of interest. Type a sequence of characters in the **Search** area to initiate filtering. The filter applies as you type. You can filter on information in any one of the columns at a time.
5 (Optional) Sort entries by the values in a column. Mouse over the column to expose the menu option, click it, then select **Sort Ascending** or **Sort Descending.**

Creating a Source Network

You can create a network resource for a source workload by specifying the **NIC# Network Name** in the machine definition in the Bulk Import spreadsheet. The name is case sensitive. If no match is found, PTM creates a new source network resource for the project. On import, it automatically associates the network resource with the source workload.

Creating a Target Network

You can use the Networks page to create network resources for the target transformation environment of a project. Auto-discovery for target VMware Clusters hosts automatically creates new Network resources that it discovers for the target host.

1 In the Web Interface toolbar, select **Resources**.
2 Select the **Networks** tab, then click **Target**.
3 Click **Create** to open the Create Network dialog.
4 Specify the following information for the network resource:
   - Name
   - Organization
   - Project
   - Description
5 Click **Save** to create the network resource.
6 Click **Close** to exit the dialog.

Associating Networks and Target Workloads

After you create target Network resources for a project, you can associate them with workloads.

1 In the Web Interface toolbar, select **Planning**.
2 Select the **Workloads** tab.
3 Select the workload, then click **Edit**. You can alternatively double-click the workload of interest.
4 On the Transformation dialog in the **Network Interfaces** panel, click **Edit**.

5 From the **Network** list, select the network resource that you want to assign to the target workload.

6 (Optional) Modify other network interface information as appropriate.

7 Click **Save**, and then click **Close**.

8 Click **Close** to exit the Transformation dialog.

Editing a Target Network

You can use the Networks page to edit a network resource for the target transformation environment of a project.

1 In the Web Interface toolbar, select **Resources**.

2 Select the **Networks** tab, then click **Target**.

3 In the **Networks** list, select the network, then click **Edit** (or double-click the network) to open the Edit Network dialog.

4 Modify any of the following information for the network resource:

- Name
- Organization
- Project
- Description

5 Click **Save** to update the network resource.

6 Click **Close** to exit the dialog.

Deleting a Network

You can use the Networks page to delete a network resource for the source or target transformation environment of a project.

NOTE: If you delete a resource, it automatically removes the associations for the resource, including those assigned to workloads that are in progress and completed. If the same associations apply, consider whether it is more appropriate to edit the resource name and settings.

---

1 In the Web Interface toolbar, select **Resources**.

2 Select the **Networks** tab.

3 (Optional) Delete a source network resource:

3a Click **Source** to view the list of source network resources.

3b Select the appropriate network, click **Delete**, then click **Yes** to confirm the deletion.

4 (Optional) Delete a target network resource:

4a Click **Target** to view the list of target network resources.

4b Select the appropriate network, click **Delete**, then click **Yes** to confirm the deletion.

25 Managing Datastore Resources

Datastore resources allow you to track information about the datastores for virtual workloads. In a virtualization environment, the datastore represents a storage location for virtual machine files and virtual disk files.

You can use the Datastores page to create, manage, and delete datastores for the target workloads in your transformation project. You can define datastore resources for the source workloads in the machine information for a bulk import. A datastore resource can apply to source workloads or to target workloads, but not to both.

About Datastore Resources

The Datastore resource allows you to track the following information for a datastore used by the one or more workloads in your transformation project:

Name: Specify a name for the datastore that is unique to your transformation project. The name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Organization: Select the parent organization for the project.

Project: Select the parent project for the datastore.

Description: (Optional) Specify a brief description of the datastore. This text appears by default in the Datastores list.

Viewing Datastores

All roles for a project can view the datastore information.

1 In the Web Interface toolbar, select **Resources**, then select the **Datastores** tab.

You can also view datastores for a single project in the **Associations** area of the Edit Project dialog. Select **Planning** > **Projects**, double-click the project, then under **Associations**, select the **Datastores** tab.

2 On the **Datastores** tab, do either of the following to view a list of resources:
   - Click the **Source** link to view datastore information that you imported for original workloads using Spreadsheet Import.
   - Click the **Target** link to view datastore information for the target transformation environment.
3 Select a datastore resource, then click **View** (or double-click the resource) to view the following settings for the resource:
   - Name
   - Organization
   - Project
   - Description
4 (Optional) Filter entries in the Datastores list using the **Search** field to show only the datastores of interest. Type a sequence of characters in the **Search** area to initiate filtering. The filter applies as you type. You can filter on information in any one of the columns at a time.
5 (Optional) Sort entries by the values in a column. Mouse over the column to expose the menu option, click it, then select **Sort Ascending** or **Sort Descending.**

Creating a Source Datastore

You can create a datastore resource for a source workload by specifying the **Disk*x* Datastore Name** in the machine definition in the Bulk Import spreadsheet. The name is case sensitive. If no match is found, PTM creates a new source datastore resource for the project. On import, it automatically associates the datastore resource with the source workload.

Creating a Target Datastore

You can use the Datastores page to create Datastore resources for the target transformation environment of a project. Auto-discovery for target VMware Clusters hosts automatically creates new Datastore resources that it discovers for the target host.

1 In the Web Interface toolbar, select **Resources**.
2 Select the **Datastores** tab, then click **Target**.
3 Click **Create** to open the Create Datastore dialog.
4 Specify the following information for the Datastore resource:
   - Name
   - Organization
   - Project
   - Description
5 Click **Save** to create the Datastore resource.
6 Click **Close** to exit the dialog.

Associating Datastores and Target Workloads

After you create target datastore resources for a project, you can associate them with virtual workloads. You can use the same datastore or different datastores for the system and data volumes and the VM Config Datastore.

To set the target datastore resource for a storage volume:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Workloads** tab.
3. Select the workload, then click **Edit**. You can alternatively double-click the workload of interest.
4. On the Transformation dialog in the **Storage** panel, click **Edit**.
5. On the Volumes tab, select the volume.
6. From the **Datastore** list, select the datastore resource that you want to assign to the target workload.
7. (Optional) Modify other disk and volume interface information as appropriate.
8. Click **Save**, and then click **Close**.
9. Click **Close** to exit the Transformation dialog.

To set the datastore resource for a virtual workload's VM Config Datastore:

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Workloads** tab.
3. Select the workload, then click **Edit**. You can alternatively double-click the workload of interest.
4. On the Transformation dialog in the **Location** panel, click **Edit**.
5. From the **VM Config Datastore** list, select the datastore resource that you want to assign to the target workload.
6. Click **Save**, and then click **Close**.
7. Click **Close** to exit the Transformation dialog.

Editing a Target Datastore

You can use the Datastores page to edit a Datastore resource for the target transformation environment of a project.

1. In the Web Interface toolbar, select **Resources**.
2. Select the **Datastores** tab, then click **Target**.
3. In the **Datastores** list, select the datastore, then click **Edit** to open the Edit Datastore dialog. You can alternatively double-click the datastore of interest.
4. Modify any of the following information for the Datastore resource:
   - Name
   - Organization
   - Project
   - Description
5. Click **Save** to update the Datastore resource.
6. Click **Close** to exit the dialog.

Deleting a Target Datastore

You can use the Datastores page to delete a Datastore resource for the source or target transformation environment of a project.

NOTE: If you delete a resource, it automatically removes the associations for the resource, including those assigned to workloads that are in progress and completed. If the same associations apply, consider whether it is more appropriate to edit the resource name and settings.

---

1. In the Web Interface toolbar, select **Resources**.
2. Select the **Datastores** tab.
3. (Optional) Delete a source Datastore resource:
   - **3a** Click **Source** to view the list of source datastore resources.
   - **3b** Select the appropriate datastore, click **Delete**, then click **Yes** to confirm the deletion.
4. (Optional) Delete a target Datastore resource:
   - **4a** Click **Target** to view the list of target datastore resources.
   - **4b** Select the appropriate datastore, click **Delete**, then click **Yes** to confirm the deletion.

26 Managing Resource Pool Resources

Resource Pool resources allow you to track the resource pools that are used by VMs on a host server running a VMware hypervisor. You can use the Resource Pool page to create, manage, view, and delete the resource pool resources for the target workloads in your transformation project. Resource Pools do not apply to source workloads.

About Resource Pool Resources

The Resource Pool resource allows you to track the following information for a VMware resource pool where one or more virtual workloads reside in your transformation project:

Name: Specify a name for the resource pool that is unique to your transformation project. The name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Organization: Select the parent organization for the project.

Project: Select the parent project for the resource pool.

Description: (Optional) Specify a brief description of the resource pool. This text appears by default in the Resource Pools list.

Viewing Resource Pools

All roles for a project can view the resource pool information.

1. In the Web Interface toolbar, select **Resources**, then select the **Resource Pools** tab.

   You can also view resource pools for a single project in the **Associations** area of the Edit Project dialog. Select **Planning** > **Projects**, double-click the project. In the Edit Project dialog under **Associations**, select the **Resources** tab, then select **Resource Pools**.
2. On the **Resource Pools** tab, view the list of resources.
3. Select a resource pool resource, then click **View** (or double-click the resource) to view the following settings for the resource:
   - Name
   - Organization
   - Project
   - Description

4. (Optional) Filter entries in the Resource Pools list using the **Search** field to show only the resource pools of interest. Type a sequence of characters in the **Search** area to initiate filtering. The filter applies as you type. You can filter on information in any one of the columns at a time.
5. (Optional) Sort entries by the values in a column. Mouse over the column to expose the menu option, click it, then select **Sort Ascending** or **Sort Descending.**

Creating a Resource Pool

You can use the Resource Pools page to create Resource Pool resources for the target transformation environment of a project.

1. In the Web Interface toolbar, select **Resources**.
2. Select the **Resource Pools** tab.
3. Click **Create** to open the Create Resource Pool dialog.
4. Specify the following information for the Resource Pool resource:
   - Name
   - Organization
   - Project
   - Description
5. Click **Save** to create the Resource Pool resource.
6. Click **Close** to exit the dialog.

Associating Resource Pools with Workloads

After you create Resource Pools resources for a project, you can associate them with virtual workloads on VMware virtualization host servers.

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Workloads** tab.
3. Select the workload, then click **Edit**. You can alternatively double-click the workload of interest.
4. On the Transformation dialog in the **Location** panel, click **Edit**.
5. From the **Resource Pool** list, select the resource pool resource that you want to assign to the target workload.
6. (Optional) Modify other Location information as appropriate.
7. Click **Save**, and then click **Close**.
8. Click **Close** to exit the Transformation dialog.

Editing a Resource Pool

You can use the Resource Pools page to edit a Resource Pool resource for the target transformation environment of a project.

1. In the Web Interface toolbar, select **Resources**.
2. Select the **Resource Pools** tab.

3 In the **Resource Pools** list, select the resource pool, then click **Edit** (or double-click the resource pool) to open the Edit Resource Pool dialog.

4 Modify any of the following information for the Resource Pool resource:

- Name
- Organization
- Project
- Description

5 Click **Save** to update the Resource Pool resource.

6 Click **Close** to exit the dialog.

Deleting a Resource Pool

You can use the Resource Pools page to delete a resource pool resource for the target transformation environment of a project.

NOTE: If you delete a resource, it automatically removes the associations for the resource, including those assigned to workloads that are in progress and completed. If the same associations apply, consider whether it is more appropriate to edit the resource name and settings.

---

1 In the Web Interface toolbar, select **Resources**.

2 Select the **Resource Pools** tab.

3 View the list of resource pool resources.

4 Select the appropriate resource pool, click **Delete**, then click **Yes** to confirm the deletion.

27 Managing Environment Resources

Environment resources allow you to track the logical environments for workloads in your network. Examples include staging, preproduction, and production. You can use the Environments page to create, manage, and delete environments for the target workloads in your transformation project. You can define environment resources for the source workloads in the machine information for a bulk import. Environment resources can apply to source and target workloads.

About Environment Resources

The Environment resource allows you to track the following information for a logical environment where one or more workloads reside in your transformation project:

Name: Specify a name for the environment that is unique to your transformation project. The name must begin with a letter or number. It can contain only letters, numbers, and the following special characters: space, hyphen, underscore, apostrophe, percent, ampersand, comma, and period.

Organization: Select the parent organization for the project.

Project: Select the parent project for the environment.

Description: (Optional) Specify a brief description of the environment. This text appears by default in the Environments list.

Viewing Environments

All roles for a project can view the environment information.

1. In the Web Interface toolbar, select **Resources**, then select the **Environments** tab.

   You can also view environments for a single project in the **Associations** area of the Edit Project dialog. Select **Planning** > **Projects**, double-click the project, then under **Associations**, select the **Environments** tab.
2. On the **Environments** tab, view the list of resources.
3. Select an environment resource, then click **View** (or double-click the resource) to view the following settings for the resource:
   - Name
   - Organization

- Project
- Description

4. (Optional) Filter entries in the Environments list using the **Search** field to show only the environments of interest. Type a sequence of characters in the **Search** area to initiate filtering. The filter applies as you type. You can filter on information in any one of the columns at a time.
5. (Optional) Sort entries by the values in a column. Mouse over the column to expose the menu option, click it, then select **Sort Ascending** or **Sort Descending.**

Creating a Source Environment

You can create an environment resource during a Spreadsheet Import by specifying the **Environment** in the machine definition in the Bulk Import spreadsheet. The name is case sensitive. If no match is found, PTM creates a new environment resource for the project. On import, it automatically associates the environment resource with the source workload.

Creating an Environment

You can use the Environments page to create Environment resources for a project.

1. In the Web Interface toolbar, select **Resources**.
2. Select the **Environments** tab.
3. Click **Create** to open the Create Environment dialog.
4. Specify the following information for the Environment resource:
   - Name
   - Organization
   - Project
   - Description
5. Click **Save** to create the Environment resource.
6. Click **Close** to exit the dialog.

Associating Environments and Target Workloads

After you create environment resources for a project, you can associate them with workloads.

1. In the Web Interface toolbar, select **Planning**.
2. Select the **Workloads** tab.
3. Select the workload, then click **Edit**. You can alternatively double-click the workload of interest.
4. On the Transformation dialog in the **Location** panel, click **Edit**.
5. From the **Environment** list, select the environment resource that you want to assign to the target workload.
6. (Optional) Modify other Location information as appropriate.
7. Click **Save**, and then click **Close**.
8. Click **Close** to exit the Transformation dialog.

Editing an Environment

You can use the Environments page to edit an Environment resource for a project.

1. In the Web Interface toolbar, select **Resources**.
2. Select the **Environments** tab.
3. In the **Environments** list, select the environment, then click **Edit** (or double-click the environment) to open the Edit Environment dialog.
4. Modify any of the following information for the Environment resource:
   - Name
   - Organization
   - Project
   - Description
5. Click **Save** to update the Environment resource.
6. Click **Close** to exit the dialog.

Deleting an Environment

You can use the Environments page to delete an Environment resource for a project.

NOTE: If you delete a resource, it automatically removes the associations for the resource, including those assigned to workloads that are in progress and completed. If the same associations apply, consider whether it is more appropriate to edit the resource name and settings.

---

1. In the Web Interface toolbar, select **Resources**.
2. Select the **Environments** tab.
3. Click **Target** to view the list of environment resources.
4. Select the appropriate environment, click **Delete**, then click **Yes** to confirm the deletion.