

PlateSpin® Migrate 12.2.1

User Guide

September 2018

Legal Notice

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

Copyright © 2017 NetIQ Corporation, a Micro Focus Company. All rights reserved.

License Grant

License bought for PlateSpin Migrate 9.3 and later versions cannot be used with PlateSpin Migrate 9.2 and prior versions.

Contents

About This Guide	15
Part I Overview and Planning	17
1 Overview of Workload Migration	19
1.1 Workload Migration Scenarios	19
1.2 Understanding Workload Migration	19
1.3 Large-Scale Migration Planning and Automation	21
2 Planning Your Workload Migrations	23
2.1 Supported Configurations	23
2.1.1 Supported Source Workloads For Migration to Non-Cloud Platforms	23
2.1.2 Supported Workloads for Migration to the Cloud	27
2.1.3 Supported Workload Storage	32
2.1.4 Supported Workload Architectures	36
2.1.5 Supported Target Virtualization Platforms	38
2.1.6 Supported Target Cloud Platforms	40
2.1.7 Supported International Languages	41
2.1.8 Supported Web Browsers	41
2.2 Supported Data Transfer Methods	42
2.2.1 File-Level Transfer (Live)	42
2.2.2 Block-Level Transfer (Live)	42
2.2.3 Offline Transfer with Temporary Boot Environment	43
2.3 Security and Privacy	43
2.3.1 Secure Connections Using TLS 1.2	44
2.3.2 Security of Workload Data in Transmission	46
2.3.3 Security of Client-Server Communications	46
2.3.4 Security of Credentials	46
2.3.5 User Authorization and Authentication	46
2.4 Performance	47
2.4.1 About Product Performance Characteristics	47
2.4.2 Scalability	47
2.4.3 Data Compression	48
2.4.4 Bandwidth Throttling	48
2.4.5 Blackout Window	48
2.5 Database Server	48
2.6 Access and Communication Requirements across your Migration Network	49
2.6.1 Requirements for Discovery	49
2.6.2 Requirements for Migration	50
2.6.3 Requirements for Event Messaging	52
2.6.4 Migrations Across Public and Private Networks through NAT	53
2.7 Deciding on the Migration Interface	53

A	Frequently Asked Questions	55
Part II	Working With Your PlateSpin Server	57
3	Using the PlateSpin Migrate Tools	59
3.1	Connecting to a PlateSpin Migrate Server	59
3.1.1	PlateSpin Server Access Using the Migrate Client	59
3.1.2	PlateSpin Server Access Using the Migrate Web Interface	61
3.2	About the PlateSpin Migrate Client User Interface.	61
3.2.1	Navigating the Client Interface.	61
3.2.2	Servers View	62
3.2.3	Jobs View.	67
3.2.4	Tasks Pane	67
3.2.5	Status Bar	68
3.2.6	Workload Migration Tasks	68
3.3	About the PlateSpin Migrate Web Interface	68
3.3.1	Navigating the Web Interface.	69
3.3.2	Workloads	70
3.3.3	Targets	74
3.3.4	Tasks	74
3.3.5	Dashboard	75
3.3.6	Reports	75
3.4	Migration Operations Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface.	75
3.5	Migration Tasks Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface	76
3.6	Other PlateSpin Server Management Tools.	78
3.6.1	PlateSpin Configuration	78
3.6.2	PlateSpin Migrate Client Command Line Interface	79
3.6.3	PlateSpin Analyzer.	79
3.6.4	Migrate Agent Utility.	79
3.6.5	PlateSpin ISO	79
4	Configuring PlateSpin Users and Access	81
4.1	Configuring User Authorization and Authentication	81
4.1.1	PlateSpin Migrate Roles	81
4.1.2	Assigning PlateSpin Migrate Roles to Windows Users	83
4.2	Configuring PlateSpin Migrate Multitenancy on VMware	84
4.2.1	Defining VMware Roles for Multitenancy.	84
4.2.2	Assigning Roles In vCenter	87
4.3	Configuring Permissions for Workload Access in PlateSpin Migrate Web Interface	90
5	Configuring PlateSpin Migrate Server	91
5.1	PlateSpin Migrate Product Licensing	91
5.1.1	Activating Your Product License	91
5.1.2	How Migration Licensing Works.	95
5.1.3	Managing License Keys for Workload Migrations	95
5.1.4	Managing Workload Designations	98
5.2	Configuring Language Settings for International Versions.	99
5.2.1	Setting the Language on the Operating System	99
5.2.2	Setting the Language in Your Web Browser	99
5.3	Enforcing FIPS Compliance for FIPS-Enabled Source Workloads	100
5.4	Configuring the Notification Service.	100
5.4.1	Notification Service Using Migrate Client.	100
5.4.2	Notification Service Using Migrate Web Interface	101
5.5	Configuring Notifications for Events and Migrations	103

5.5.1	Notifications Using the Migrate Client	104
5.5.2	Notifications Using the Web Interface	104
5.6	Enabling Event Messaging for PlateSpin Migration Factory	105
5.7	Configuring Alternate IP Addresses for PlateSpin Server	106
5.8	Setting Reboot Method for the Configuration Service	106
5.9	Configuring the Contact Direction for the Replication Port.	107
5.10	Configuring Behavior for Installing Network Drivers on Target Windows Workloads	107
5.10.1	Understanding Light Networking Parameters	108
5.10.2	Configuring Light Networking Parameters	109
5.11	Specifying the Network Adapter Type to Use for Migrations to Hyper-V during Target Take-Control	109
5.12	Configuring Applications Known to Cause Boot Failure on Windows Target	110
5.12.1	Editing the List of Applications Known to Cause Boot Failure on Windows Target	110
5.13	Optimizing Data Transfer over WAN Connections	111
5.13.1	Tuning Parameters	111
5.13.2	Tuning FileTransferSendReceiveBufferSize	113
5.14	Increasing the Upload Size Limit for Post-Migration Actions	114
5.15	Other Use Cases for Custom PlateSpin Server Settings (Advanced)	114
6	Configuring PlateSpin Migrate Client	117
6.1	Configuring General Options	117
6.2	Configuring Job Values Defaults	118
6.3	Configuring Source Service Defaults	119
6.4	Configuring Target Service Defaults	120
6.5	Managing Post-Migration Actions (Windows and Linux)	121
6.6	Managing Migrate Client User Activity Log	123
6.6.1	About the Migrate Client User Activity Log	123
6.6.2	Configuring Migrate Client User Activity Logging	124
6.6.3	Viewing Migrate Client User Activity Log	124
7	Configuring PlateSpin Migrate Web Interface	125
7.1	Managing Security Groups and Workload Permissions	125
7.1.1	Prerequisites for Security Groups	125
7.1.2	Creating Security Groups for Migrate Web Interface	126
7.1.3	Modifying Security Group Members or Workloads	126
7.1.4	Deleting a Security Group	126
7.2	Managing Workload Tags	127
7.2.1	Creating a Workload Tag	127
7.2.2	Using Workload Tags	127
7.2.3	Modifying a Workload Tag	127
7.2.4	Deleting a Workload Tag	127
7.3	Configuring the Refresh Rates for PlateSpin Migrate Web Interface	128
7.4	Customizing the UI for PlateSpin Migrate Web Interface	128
B	Rebranding the UI for PlateSpin Migrate Web Interface	129
B.1	Rebranding the UI Using PlateSpin Configuration Parameters	129
B.1.1	About Configurable UI Elements for PlateSpin Migrate Web Interface	129
B.1.2	Modifying PlateSpin Configuration Settings for Configurable UI Elements	130
B.2	Rebranding the Product Name in the Windows Registry	133

Part III Preparing Your Migration Environment	135
8 Prerequisites for Migration to Amazon Web Services	137
8.1 Deployment for Migration to Amazon Web Services	137
8.2 Requirements for Migrating Workloads to Amazon Web Services	138
8.2.1 Minimum AWS Prerequisites	138
8.2.2 AWS Prerequisites for Using an On Premise Migrate Server	139
8.3 Planning For Migrating Workloads to Amazon Web Services	141
8.4 Configuring OS License Activation on Windows Targets Migrated to AWS	142
8.5 Using the PlateSpin Replication Environment AMIs	142
8.5.1 For Full Replication	143
8.5.2 For Server Sync	143
8.6 AWS Networking Guidelines	143
8.6.1 Workloads with Multiple NICs	143
8.6.2 Private and Public IP Addresses for Workloads Connected on an AWS VPN	144
8.7 Best Practices For Configuring a Migration Job to Amazon Web Services	144
8.8 Checklist for Automated Migration to AWS	145
9 Prerequisites for Migration to Microsoft Azure	147
9.1 Deployment for Migration to Azure	147
9.2 Requirements for Migrating Workloads to Azure	149
9.2.1 Minimum Azure Prerequisites	150
9.2.2 Azure Prerequisites for Using an On Premise Migrate Server	151
9.2.3 Azure Prerequisites for Using an Azure-Based Migrate Server	152
9.3 Planning For Migrating Workloads to Azure	153
9.4 Azure Networking Guidelines	154
9.4.1 Private or Public IP Addresses for Azure Migration	155
9.4.2 Windows Workloads in Azure with Multiple NICs	155
9.4.3 Private and Public IP Addresses for Workloads Connected on an Azure VPN	155
9.5 Registering a Azure Application to Represent PlateSpin Migrate	156
9.6 Configuring an Azure User for PlateSpin Migrate to Use	157
9.7 Enabling Programmatic Deployment of PlateSpin Migrate Replication Environment	157
9.8 Deploying a Migrate Server Image in Azure	158
9.9 Configuring the Target IaaS Environment for Azure Cloud	159
9.10 Managing the Azure User Password for Azure Target Containers	160
9.11 Checklist for Automated Migration to Azure	161
10 Prerequisites for Migration to VMware vCloud Director	163
10.1 Deployment for Migration to VMware vCloud	163
10.2 Planning For Migrating Workloads to VMware vCloud Director	165
10.3 Setting up vCloud Organization	165
10.4 Understanding PlateSpin Replication Environment	166
10.4.1 Resources Used in the PlateSpin Replication Environment	166
10.4.2 Creating the PlateSpin Virtual Appliance in the vCloud Organization	166
10.5 Configuring Advanced PlateSpin Settings for vCloud	167
10.5.1 Retaining the Cloud Resources For Troubleshooting Migration Errors	167
10.5.2 Setting the PlateSpin Replication Environment Password in Clear Text	167
10.6 Checklist for Automated Migration to vCloud	168
11 Prerequisites for Migration to VMware	169
11.1 Deployment for Migration to VMware	169
11.2 Planning for Migration to VMware	171

11.3	Checklist for Automated Migration to VMware	172
11.4	Checklist for Semi-Automated Migration to Target VMs on VMware	172
12	Prerequisites for Migration to Microsoft Hyper-V	175
12.1	Deployment for Migration to Microsoft Hyper-V	175
12.2	Planning for Migration to Microsoft Hyper-V	177
12.3	Checklist for Automated Migration to Hyper-V	178
12.4	Checklist for Semi-Automated Migration to Target VMs on Hyper-V	178
13	Prerequisites for Migration to VMs on Citrix XenServer	179
13.1	Deployment for Migration to Citrix XenServer	179
13.2	Planning for Migration to VMs on Citrix XenServer	180
13.3	Checklist for Semi-Automated Migration to Target VMs on Citrix XenServer	181
14	Prerequisites for Migration to VMs on Xen	183
14.1	Deployment for Migration to Xen	183
14.2	Planning for Migration to VMs on Xen	184
14.3	Checklist for Semi-Automated Migration to Target VMs on Xen	184
15	Prerequisites for Migration to VMs on KVM	185
15.1	Deployment for Migration to KVM	185
15.2	Planning for Migration to VMs on KVM	186
15.3	Checklist for Semi-Automated Migration to Target VMs on KVM	186
16	Prerequisites for Migration to Physical Machines	187
16.1	Deployment for Migration to Physical Machines	187
16.2	Planning for Migration to Physical Machines	188
16.3	Best Practices (X2P)	189
16.4	Checklist for Semi-Automated Migration to Physical Machines	189
17	Prerequisites for Migration to an Image	191
18	Preparing for Synchronization of Workloads with Server Sync	193
Part IV	Discovering and Preparing Workloads and Targets	195
19	Discovering Target Hosts	197
19.1	About Target Host Discovery	197
19.2	Network Access Requirements for Target Host Discovery	198
19.3	Discovery Guidelines for Target Hosts	198
19.3.1	Target Host Discovery Parameters for Migrate Web Interface	199
19.3.2	Target Host Discovery Parameters for Migrate Client	199
19.4	Discovering Details for Target Containers	200
19.4.1	Target Discovery in the Migrate Client	200
19.4.2	Target Discovery in the Web Interface	202
19.5	Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO	204
19.5.1	Prerequisites for Discovering Target VMs	205

19.5.2	Registering and Discovering Target VMs on Virtual Hosts	205
19.5.3	Configuration Information	206
19.6	Registering and Discovering Details for Target Physical Machines with PlateSpin ISO	206
19.6.1	Prerequisites for Discovering Target Physical Machines	206
19.6.2	Registering Target Physical Machines Using the PlateSpin ISO	207
19.6.3	Configuration Information	207
19.7	Discovering Target VMs for Server Sync Jobs	208
19.8	Refreshing Target Host Details	208
19.8.1	Refresh Target Details in the Web Interface	208
19.8.2	Refresh Target Details in Migrate Client	209
19.8.3	Refresh ESX Host Details in Migrate Client	209
19.9	Removing Target Containers	210

20 Discovering Source Workloads 211

20.1	About Source Workload Discovery	211
20.2	Network Access Requirements for Workload Discovery	212
20.3	Discovery Guidelines for Source Workloads	212
20.4	Populating the Servers View with a List of Windows Computers in a Domain	213
20.5	Discovering Details for All Windows Workloads in a Domain	213
20.6	Discovering Details for Source Workloads	213
20.6.1	Workload Discovery in the Migrate Client	214
20.6.2	Workload Discovery in the Migrate Web Interface	215
20.7	Registering Workloads and Discovering Details with Migrate Agent	215
20.7.1	Windows Workload Registration and Discovery with Migrate Agent	216
20.7.2	Linux Workload Registration and Discovery with Migrate Agent	217
20.8	Refreshing Source Workload Details	217
20.8.1	Refresh Workload Details in Migrate Client	218
20.8.2	Removing and Re-Adding Workloads in the Web Interface	218
20.9	Using Tags to Track Logical Associations of Workloads	218
20.10	Undiscovering or Removing Source Workloads	219

21 Preparing Device Drivers 221

21.1	Packaging Device Drivers for Windows Systems	221
21.2	Packaging Device Drivers for Linux Systems	221
21.3	Uploading Drivers to the PlateSpin Migrate Device Driver Database	222
21.3.1	Device Driver Upload Procedure (Windows)	222
21.3.2	Device Driver Upload Procedure (Linux)	223
21.4	Using the Plug and Play (PnP) ID Translator Feature	224
21.5	Analyzing Suitability of Discovered Windows Workloads For Conversion to Physical Machines	229
21.5.1	About PlateSpin Analyzer Tests and Results	230
21.5.2	PlateSpin Analyzer in the Migrate Client	231

22 Preparing Linux Workloads for Migration 233

22.1	Verifying Block-Based Drivers for Linux	233
22.2	Adding Drivers to the PlateSpin ISO Image	233
22.3	Configuring LVM Snapshots for Linux Volume Replication	233
22.4	Using Custom Freeze and Thaw Scripts for Linux Block-Level Migrations	234
22.5	Preparing Paravirtualized Linux Source Workload	235

23 Preparing for Migration of Windows Clusters 237

23.1	Planning Your Cluster Workload Migration	237
23.1.1	Requirements for Cluster Migration	238

23.1.2	Block-Based Transfer for Clusters	239
23.1.3	Impact of Cluster Node Failover on Replication	240
23.1.4	Cluster Node Similarity	241
23.1.5	Migration Setup	242
23.2	Configuring Windows Active Node Discovery	242
23.3	Configuring the Block-Based Transfer Method for Clusters	243
23.4	Adding Resource Name Search Values	243
23.5	Quorum Arbitration Timeout	244
23.6	Setting Local Volume Serial Numbers	244
23.7	PlateSpin Cutover	244
23.8	PlateSpin Cluster Migration	245
23.9	Migrating Windows Clusters with the Web Interface	245
23.10	Migrating Windows Clusters with the Migrate Client	245
C	Troubleshooting Discovery	247
C.1	Common Discovery Issues and Solutions	247
C.2	Modifying the OFX Controller Heartbeat Startup Delay (Windows Workloads)	249
D	Linux Distributions Supported by Migrate	251
D.1	Analyzing Your Linux Workload	251
D.1.1	Determining the Release String	251
D.1.2	Determining the Architecture	251
D.2	Pre-compiled blkwatch Drivers for Linux Distributions	252
D.2.1	List Item Syntax	252
D.2.2	List of Distributions	252
D.2.3	Other Linux Distributions That Use blkwatch Drivers	264
E	Synchronizing Serial Numbers on Cluster Node Local Storage	267
F	Migrate Agent Utility	269
F.1	Requirements for Migrate Agent Utility	269
F.2	Migrate Agent Utility for Windows	270
F.3	Migrate Agent Utility for Linux	272
F.4	Using Migrate Agent with Block-Based Transfer Drivers	275
F.5	Using Migrate Agent to Register Workloads	277
G	PlateSpin ISO Image	279
G.1	Downloading the PlateSpin ISO Images	279
G.2	Preparing the PlateSpin ISO Image for Target Registration and Discovery	280
G.3	Injecting Additional Device Drivers into the PlateSpin ISO Image	280
G.4	Adding Registration Information to the PlateSpin ISO for Unattended Registration of Physical or Virtual Machines	281
G.5	Using PlateSpin ISO	281
Part V	Configuring Workloads	283
24	Prerequisites for Automated Migrations	285
24.1	Supported Source Workloads for Automated Migration	285
24.2	Supported Target Platforms for Automated Migrations	286

24.3	Preparing Targets for Automated Migration	286
24.4	Network Connections and Bandwidth	287
24.5	Automated Workflow	287
25	Prerequisites for Semi-Automated (X2P) Migrations	289
25.1	Supported Source Workloads for X2P Migrations	289
25.2	Supported Target Platforms for X2P Migrations	289
25.3	X2P Workflow for VMs	289
26	Configuration Essentials	291
26.1	Configuration Workflows	291
26.1.1	Configuration Workflows Using Migrate Client	292
26.1.2	Configuring Workflows Using Migrate Web Interface	292
26.2	Initiating a Migration Job	292
26.2.1	Prerequisites for Migration Jobs	292
26.2.2	Initiate a Migration Job Using Migrate Client	293
26.2.3	Initiate a Migration Job Using the Migrate Web Interface	294
26.3	Saving a Migration Configuration	295
26.3.1	Using the Migrate Client	295
26.3.2	Using the Migrate Web Interface	295
26.4	Editing a Migration Job	295
26.4.1	Edit Migration Job Using Migrate Client	295
26.4.2	Edit Migration Job Using Migrate Web Interface	296
26.5	Migrate License Key	296
26.5.1	License Key in Migrate Client	296
26.5.2	License Key in Migrate Web Interface	296
26.6	Credentials for Source Workloads and Target Hosts	297
26.6.1	About Credentials	297
26.6.2	Credentials in Migrate Client	297
26.6.3	Credentials in Migrate Web Interface	297
26.7	Migration Schedule	298
26.7.1	Migration Schedule Using Migrate Client	298
26.7.2	Migration Schedule Using Migrate Web Interface	298
26.8	Blackout Window for Data Transfer	298
26.8.1	Blackout Window Using the Migrate Client	298
26.8.2	Blackout Window Using the Migrate Web Interface	299
26.9	Compression during Data Transfer	299
26.9.1	Compression Using Migrate Client	299
26.9.2	Compression Using Migrate Web Interface	299
26.10	Bandwidth Throttling during Data Transfer	300
26.10.1	Bandwidth Throttling Using Migrate Client	300
26.10.2	Bandwidth Throttling Using Migrate Web Interface	300
26.11	Conversion (Data Transfer Method)	300
26.11.1	Conversion Using Migrate Client	301
26.11.2	Data Transfer Using Migrate Web Interface	301
26.12	Encrypt Data Transfer	301
26.12.1	Encrypt Data Transfer Using Migrate Client	301
26.12.2	Encrypt Data Transfer Using Migrate Web Interface	301
26.13	Virtualization Enhancement Software	302
26.13.1	Replace VMware Tools using Migrate Client	302
26.13.2	Replace VMware Tools using Migrate Web Interface	303
26.14	Custom Post-Migration Actions	303
26.15	Services or Daemons to Stop before Replication or Cutover	304
26.15.1	Services and Daemons to Stop Using Migrate Client	304
26.15.2	Services and Daemons to Stop using Migrate Web Interface	305

26.16	Service States on Target Windows Workloads	306
26.16.1	Service States using Migrate Client	306
26.16.2	Service States using Migrate Web Interface	308
26.17	Daemon States on Target Linux Workloads	310
26.17.1	Daemon States using Migrate Client	311
26.17.2	Daemon States using Migrate Web Interface	311
26.18	Windows HAL or Kernel File Replacements	312
26.19	Post-Cutover End States for Source and Target Workloads	313
26.19.1	Workload End States Using the Migrate Client	313
26.19.2	Workload End States Using the Migrate Web Interface	313
26.20	Target Workload Settings for VMs	314
26.20.1	Target VM Configuration in Migrate Client	314
26.20.2	Target VM Configuration in Migrate Web Interface	314
26.21	Network Identification (Network Connections)	315
26.21.1	Network Identification Using Migrate Client	315
26.21.2	Network Connections Using Migrate Web Interface	317
26.22	Migration Network (Replication Network)	318
26.22.1	Migration Network Using Migrate Client	318
26.22.2	Replication Network Using Migrate User Interface	324
26.23	Storage Disks and Volumes	326
26.23.1	Storage Disks and Volumes Using Migrate Client	326
26.23.2	Storage and Volume Using Migrate Web Interface	330
27	Migration to Amazon Web Services	331
27.1	Planning for Migration to Amazon Web Services	331
27.2	Configuring Migration of a Workload to Amazon Web Services	332
27.3	Configuring Full Migration of a Workload to Amazon Web Services	332
28	Migration to Microsoft Azure	335
28.1	Planning for Migration to Microsoft Azure	335
28.2	Configuring Migration of a Workload to Microsoft Azure	336
29	Migration to VMware vCloud Director	349
29.1	Planning for Migration to VMware vCloud Director	349
29.2	Configuring Migration of a Workload to VMware vCloud Director	349
30	Migration to VMware	359
30.1	Planning for Migration to VMware	359
30.2	Automated Migration to VMware Using Migrate Client	360
30.2.1	Target VM Configuration: VMware ESXi 5 and Later	368
30.2.2	Target VM Configuration: VMware ESX 4.1	369
30.2.3	Drive Configuration: VMware ESX	371
30.3	Migration to VMs on VMware Using X2P Workflow	371
30.3.1	Downloading and Saving the PlateSpin ISO Image (VMware)	372
30.3.2	Creating and Configuring the Target Virtual Machine (VMware)	372
30.3.3	Setting Up VMware Tools for the Target Workload	372
30.3.4	Registering the Virtual Machine with PlateSpin Server (VMware)	373
30.3.5	Migrating Your Source Workload to the Target Virtual Machine (VMware)	374
30.4	Automated Migration to VMware Using Migrate Web Interface	374

31 Migration to Microsoft Hyper-V	383
31.1 Planning for Migration to Hyper-V	383
31.2 Automated Migration to Hyper-V	384
31.2.1 Target VM Configuration: Microsoft Hyper-V	390
31.2.2 Drive Configuration: Hyper-V	392
31.3 Migration to VMs on Hyper-V Using X2P Workflow	393
31.3.1 Downloading and Saving the PlateSpin ISO Image (Hyper-V)	393
31.3.2 Creating and Configuring the Target Virtual Machine (Hyper-V)	393
31.3.3 Registering the Virtual Machine with PlateSpin Server (Hyper-V)	394
31.3.4 Migrating Your Source Workload to the Target Virtual Machine (Hyper-V)	394
31.3.5 Post-Migration Steps (Hyper-V)	394
32 Migration to Virtual Machines on Citrix XenServer	395
32.1 Planning for Migration to Citrix XenServer	395
32.2 Configuring Migration to a VM on a Citrix XenServer Virtual Host	396
32.2.1 Downloading and Preparing the PlateSpin ISO Image (Citrix XenServer)	396
32.2.2 Creating and Configuring the Target Virtual Machine (Citrix XenServer)	396
32.2.3 Registering the Virtual Machine with PlateSpin Server (Citrix XenServer)	397
32.2.4 Migrating Your Source Workload to the Target Virtual Machine (Citrix XenServer)	397
32.2.5 Target VM Configuration: Citrix XenServer	398
33 Migration to Virtual Machines on Xen	399
33.1 Planning for Migration to Xen	399
33.2 Configuring Migration to a VM on a Xen Virtual Host	400
33.2.1 Downloading and Preparing the PlateSpin ISO Image (Xen on SLES)	400
33.2.2 Creating and Configuring the Target Virtual Machine (Xen on SLES)	400
33.2.3 Registering the Virtual Machine with PlateSpin Server (Xen on SLES)	401
33.2.4 Migrating Your Source Workload to the Target Virtual Machine (Xen on SLES)	401
33.2.5 Post-Migration Steps (Xen on SLES)	401
34 Migration to Virtual Machines on KVM	403
34.1 Planning for Migration to KVM	403
34.2 Configuring Migration to a VM on a KVM Virtual Host	404
34.2.1 Downloading and Preparing the PlateSpin ISO Image (KVM)	404
34.2.2 Creating and Configuring the Target Virtual Machine (RHEL KVM)	404
34.2.3 Registering the Virtual Machine with PlateSpin Server (RHEL KVM)	405
34.2.4 Migrating Your Source Workload to the Target Virtual Machine (RHEL KVM)	405
35 Migration to Physical Machines	407
35.1 Planning for Migration to Physical Machines	407
35.2 Configuring Migration to a Physical Target (P2P, V2P)	408
36 Workload Migration with a PlateSpin Image	415
36.1 About PlateSpin Images	415
36.2 Designating a PlateSpin Image Server	415
36.3 Capturing a Workload to a PlateSpin Image	417
36.4 Deploying a PlateSpin Image	418
36.5 Managing PlateSpin Images	419
36.5.1 Moving Images from One PlateSpin Image Server to Another	420
36.5.2 Automating Image Operations	420
36.5.3 Browsing and Extracting Image Files	420

37 Synchronizing Workloads with Server Sync	423
37.1 Server Sync to a Virtual Target	423
37.2 Server Sync to a Physical Target	426
37.3 Server Sync to Amazon Web Services	426
37.4 Server Sync Volume Mapping	427
37.4.1 Server Sync Volume Configuration (Windows)	428
37.4.2 Server Sync Volume Configuration (Linux)	429
 Part VI Executing Migrations	 431
 38 Executing Workload Migrations	 433
38.1 Preparing a Migration	433
38.1.1 Using the Migrate Client	433
38.1.2 Using the Migrate Web Interface	434
38.2 Starting Migration Execution (First Replication)	434
38.2.1 Using the Migrate Client	435
38.2.2 Using the Migrate Web Interface	435
38.3 Scheduling Migration Execution (First Replication)	435
38.3.1 Using the Migrate Client	436
38.3.2 Using the Migrate Web Interface	436
38.4 Starting Incremental Replications	436
38.4.1 Using the Migrate Web Interface	436
38.5 Scheduling Incremental Replications	437
38.5.1 Using the Migrate Web Interface	437
38.6 Viewing Properties for an In-Progress or Completed Migration	438
38.6.1 Using the Migrate Client	438
38.6.2 Using the Migrate Web Interface	438
38.7 Canceling an In-Progress Migration	438
38.7.1 Using the Migrate Client	438
38.7.2 Using the Migrate Web Interface	438
38.8 Restarting or Shutting Down the Source Workload	439
 39 Generating Reports	 441
39.1 Generating Workload and Workload Migration Reports	441
39.1.1 Generate Reports using the Migrate Client	441
39.1.2 Generate Reports using the Web Interface	442
39.2 Generating Diagnostic Reports	442
39.2.1 Using the Migrate Client	442
39.2.2 Using the Migrate Web Interface	443
 40 Post-Migration Cleanup of Source Workloads	 445
40.1 Cleaning Up Windows Workloads	445
40.2 Cleaning Up Linux Workloads	446
 H Troubleshooting PlateSpin Migrate	 447
H.1 Migration of Workloads to Azure Cloud	447
H.1.1 Assigning a Reserved IP Address to a Migrate Server in Azure	447
H.1.2 Computer Name Not Available for Azure VM in Azure Portal	448
H.1.3 Outbound Email Stuck after Migrating Microsoft Exchange Server 2016 to Azure Cloud	448
H.2 Migration of Workloads to VMware	449

H.2.1	Outbound Email Stuck after Migrating Microsoft Exchange Server 2016 to VMware	449
H.3	Peer-to-Peer Migrations (Windows)	449
H.4	PlateSpin Images	450
H.5	Shrinking the PlateSpin Migrate Databases	451
H.6	Troubleshooting the Configuration Service	451
H.6.1	Understanding What Is Causing the Problem	452
H.6.2	What Can Be Done to Resolve the Problem	452
H.6.3	Additional Troubleshooting Tips	455
H.7	Validation Warning for Bandwidth Throttling	456
H.8	PlateSpin OFX Controller Does Not Start on a Virtual Machine Source	456
H.9	Target Windows Machine Becomes Unbootable on Second Boot	456
H.10	Two or More Volumes Have the Same Volume Serial Number	457
Part VII Additional PlateSpin Tools		459
I	Using the PlateSpin Migrate Client Command Line Interface	461
I.1	Where Is the Tool Located?	461
I.2	Before You Use the Tool	461
I.2.1	Pre-configuring the Migrate Server Values for CLI	462
I.2.2	Becoming Familiar with the Commands	462
I.3	Configurable .ini Files (Jobs) You Can Use with the Tool	464
I.3.1	Conversion Jobs	465
I.3.2	ServerSync Jobs	472
I.3.3	Imaging Jobs	476
J	Using the iPerf Network Test Tool to Optimize Network Throughput for PlateSpin Products	487
J.1	Introduction	487
J.2	Calculations	488
J.3	Setup	489
J.4	Methodology	490
J.5	Expectations	491
Part VIII Documentation Updates		493
K	Documentation Updates	495
K.1	October 2018	495
K.2	September 2018	495
K.3	March 2018	495
K.4	February 2018	496
K.5	January 2018	497

About This Guide

This guide provides information about using PlateSpin Migrate.

- ♦ [Part I, “Overview and Planning,” on page 17](#)
- ♦ [Part II, “Working With Your PlateSpin Server,” on page 57](#)
- ♦ [Part III, “Preparing Your Migration Environment,” on page 135](#)
- ♦ [Part IV, “Discovering and Preparing Workloads and Targets,” on page 195](#)
- ♦ [Part V, “Configuring Workloads,” on page 283](#)
- ♦ [Part VI, “Executing Migrations,” on page 431](#)
- ♦ [Part VII, “Additional PlateSpin Tools,” on page 459](#)
- ♦ [Part VIII, “Documentation Updates,” on page 493](#)

Audience

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

Additional Documentation

This guide is part of the PlateSpin Migrate documentation set. For a complete list of publications supporting this release, visit the [PlateSpin Migrate Documentation website \(https://www.netiq.com/documentation/platespin-migrate-12-2-1/\)](https://www.netiq.com/documentation/platespin-migrate-12-2-1/).

Documentation Updates

The most recent version of this guide can be found at the [PlateSpin Migrate Documentation website \(https://www.netiq.com/documentation/platespin-migrate-12-2-1/\)](https://www.netiq.com/documentation/platespin-migrate-12-2-1/).

Contacting Micro Focus

For specific product issues, contact Micro Focus Support at <https://www.microfocus.com/support-and-services/>.

Additional technical information or advice is available from several sources:

- ♦ Product documentation, Knowledge Base articles, and videos: <https://www.microfocus.com/support-and-services/>
- ♦ The Micro Focus Community pages: <https://www.microfocus.com/communities/>

Contacting 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.

Overview and Planning

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

- ♦ [Chapter 1, “Overview of Workload Migration,” on page 19](#)
- ♦ [Chapter 2, “Planning Your Workload Migrations,” on page 23](#)
- ♦ [Appendix A, “Frequently Asked Questions,” on page 55](#)

1 Overview of Workload Migration

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

- ♦ [Section 1.1, “Workload Migration Scenarios,” on page 19](#)
- ♦ [Section 1.2, “Understanding Workload Migration,” on page 19](#)
- ♦ [Section 1.3, “Large-Scale Migration Planning and Automation,” on page 21](#)

1.1 Workload Migration Scenarios

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

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

1.2 Understanding Workload Migration

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

Figure 1-1 Workload Migration




Table 1-1 Workload Migration Operations

Category of Operation	Migration Infrastructures
Peer-to-peer	<ul style="list-style-type: none">Physical to Virtual (P2V)Virtual to Virtual (V2V)Virtual to Physical (V2P)Physical to Physical (P2P)
Imaging	<ul style="list-style-type: none">Physical to Image (P2I)Virtual to Image (V2I)Image to Virtual (I2V)Image to Physical (I2P)
Cloud	<ul style="list-style-type: none">Physical to Cloud (P2C)Virtual to Cloud(V2C)Cloud to Physical (C2P)Cloud to Virtual (C2V)
NOTE <ul style="list-style-type: none">Supported Cloud platforms include Microsoft Azure, VMware vCloud Director, or Amazon Web Services (AWS).	

PlateSpin Migrate supports multiple workload types and virtualization platforms. Imaging is supported for workloads with Microsoft Windows operating systems. For a more detailed list of supported workloads and infrastructures, see [“Supported Configurations” on page 23](#).

1.3 Large-Scale Migration Planning and Automation

PlateSpin Migration Factory is a planning, scheduling, migration execution, and visualization solution that streamlines the execution of large-scale cloud and data center migration projects. PlateSpin Transformation Manager and PlateSpin Migrate Connector work with multiple PlateSpin Migrate servers to handle the full migration lifecycle—from planning to fully automated or semi-automated migration activities to successful cutover.

PlateSpin Migration Factory provides several benefits:

- ♦ Helps project managers create realistic project plans
- ♦ Gives project architects insights into environmental challenges
- ♦ Enables migration specialists to execute server migrations on time, with more automation, and less room for human error

PlateSpin Transformation Manager uses import with automated discovery to simplify and standardize the setup of migration workload and target containers 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. In Manual Mode, you can plan migrations and monitor semi-automated migration activities across the project. 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.

PlateSpin Migrate servers provide migration capabilities needed to execute and monitor defined migration jobs.

For more information about PlateSpin Transformation Manager and PlateSpin Migrate Connector, see the [PlateSpin Transformation Manager User Guide](https://www.netiq.com/documentation/platespin-transformation-manager-1-1/) on the [Documentation website \(https://www.netiq.com/documentation/platespin-transformation-manager-1-1/\)](https://www.netiq.com/documentation/platespin-transformation-manager-1-1/).

2 Planning Your Workload Migrations

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

- ♦ [Section 2.1, “Supported Configurations,” on page 23](#)
- ♦ [Section 2.2, “Supported Data Transfer Methods,” on page 42](#)
- ♦ [Section 2.3, “Security and Privacy,” on page 43](#)
- ♦ [Section 2.4, “Performance,” on page 47](#)
- ♦ [Section 2.5, “Database Server,” on page 48](#)
- ♦ [Section 2.6, “Access and Communication Requirements across your Migration Network,” on page 49](#)
- ♦ [Section 2.7, “Deciding on the Migration Interface,” on page 53](#)

2.1 Supported Configurations

- ♦ [Section 2.1.1, “Supported Source Workloads For Migration to Non-Cloud Platforms,” on page 23](#)
- ♦ [Section 2.1.2, “Supported Workloads for Migration to the Cloud,” on page 27](#)
- ♦ [Section 2.1.3, “Supported Workload Storage,” on page 32](#)
- ♦ [Section 2.1.4, “Supported Workload Architectures,” on page 36](#)
- ♦ [Section 2.1.5, “Supported Target Virtualization Platforms,” on page 38](#)
- ♦ [Section 2.1.6, “Supported Target Cloud Platforms,” on page 40](#)
- ♦ [Section 2.1.7, “Supported International Languages,” on page 41](#)
- ♦ [Section 2.1.8, “Supported Web Browsers,” on page 41](#)

2.1.1 Supported Source Workloads For Migration to Non-Cloud Platforms

PlateSpin Migrate supports the migration of the following Windows and Linux workloads to non-cloud platforms, such as physical machines and virtual machines on supported hypervisors. See [“Supported Target Virtualization Platforms” on page 38](#).

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

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

NOTE: Not all workloads are supported on all target virtualization platforms. Migration of workloads to a target virtualization platform is subject to the support of the guest operating system on the target host by the host vendor.

Review the following sections:

- ♦ “Supported Microsoft Windows Workloads For Migration to Non-Cloud Platforms” on page 24
- ♦ “Supported Linux Workloads For Migration to Non-Cloud Platforms” on page 25

Supported Microsoft Windows Workloads For Migration to Non-Cloud Platforms

PlateSpin Migrate supports the following Microsoft Windows platforms for migration to virtual machines on virtualization hosts or to physical machines, except as noted in [Table 2-1](#). See also “Supported Workload Storage” on page 32 and “Supported Workload Architectures” on page 36.

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

Operating System	Remarks
Servers	
Windows Server 2016	Migration to a VMware VM requires VMware vCenter 6.0 or later.
Windows Server 2012 R2 Windows Server 2012	
Windows Server 2008 R2 Windows Server 2008	Includes domain controller (DC) systems and Small Business Server (SBS) editions. Does not support migration of Windows Server 2008 R2 SP0 to Hyper-V because Microsoft no longer supports it. See Microsoft TechNet Website (https://technet.microsoft.com/library/dn792027.aspx) .
Windows Server 2003 R2 Windows Server 2003 SP 1 and later	

Operating System	Remarks
Clusters	
Windows Server 2016 Cluster Supports quorum models: <ul style="list-style-type: none">◆ Node and Disk Majority◆ No Majority: Disk Only	Both Migrate Client and Web Interface support automated migration of Windows Clusters to VMware vCenter containers. Migrate Client also supports semi-automated migration of Windows Clusters to physical machines by using the X2P workflow. See Chapter 23, “Preparing for Migration of Windows Clusters,” on page 237. Migration of Windows Server 2016 Clusters to VMware requires VMware 6.0 or later. PlateSpin Migrate does not support migration of Windows Server clusters to the following target infrastructures: <ul style="list-style-type: none">◆ Images◆ Cloud◆ Virtualization hypervisors other than VMware PlateSpin Migrate supports only block-level replication for clusters. File-level replication is not supported. Support includes block-based data transfer with a driver (Fibre Channel SANs only) or without a driver for incremental replications for clusters. WARNING: Do not attempt to use the block-based driver on clusters with shared iSCSI drives. It renders the cluster unusable.
Windows Server 2012 R2 Cluster Windows Server 2012 Cluster Supports quorum models: <ul style="list-style-type: none">◆ Node and Disk Majority◆ No Majority: Disk Only	
Windows Server 2008 R2 Cluster Windows Server 2008 Cluster Supports quorum models: <ul style="list-style-type: none">◆ Node and Disk Majority◆ No Majority: Disk Only	
Windows Server 2003 R2 Cluster Windows Server 2003 Cluster Supports quorum model: <ul style="list-style-type: none">◆ Single-Quorum Device Cluster	
Desktops	
Windows 8 and 8.1	Requires the High Performance Power Plan.
Windows 7	Supports only Professional, Enterprise, and Ultimate.

Supported Linux Workloads For Migration to Non-Cloud Platforms

PlateSpin Migrate supports the following Linux platforms for migration to virtual machines on virtualization hosts or to physical machines, except as noted in [Table 2-2](#). See also [“Supported Workload Storage”](#) on page 32 and [“Supported Workload Architectures”](#) on page 36.

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

Linux Distribution	Versions	Remarks
Red Hat Enterprise Linux (RHEL)	AS/ES/WS 4.x, 5.0 to 5.11, 6.0 to 6.9, and 7.0 to 7.3	<p>PlateSpin Migrate does not support the XFS version 5 (v5) file system on RHEL 7.3 and on distributions based on RHEL 7.3.</p> <p>For Red Hat Enterprise Linux 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads with LVM volumes, PlateSpin Migrate supports incremental replication only for the latest available kernel (version 2.6.32-642.13.1.el6) for the 6.7 distribution.</p> <p>For Red Hat Enterprise Linux 6.8, Oracle Linux 6.8, and CentOS 6.8 workloads with LVM volumes, PlateSpin Migrate supports incremental replication only for the latest available kernel (version 2.6.32-696.20.1.el6.x86_64) for the 6.8 distribution.</p> <p>Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for RHEL 5. See “Paravirtualized Source Workloads” on page 37.</p>
SUSE Linux Enterprise Server (SLES)	9, 10, 11 (SP1, SP2, SP3, and SP4)	<p>The SLES 11 SP2 (32-bit) with kernel 3.0.13-0.27-pae is not supported. The kernel for this version of SLES must be upgraded to 3.0.51-0.7.9-pae so that the conversion works.</p> <p>Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for SLES 10 and 11. See “Paravirtualized Source Workloads” on page 37.</p>
CentOS	See Red Hat Enterprise Linux.	<p>Same level of support as that for workloads running RHEL except that CentOS 4.x is not supported for Hyper-V.</p> <p>Migration of CentOS 7.x to VMware requires VMware vCenter 5.5 or later.</p>

Linux Distribution	Versions	Remarks
Oracle Linux (OL) (formerly Oracle Enterprise Linux)	See Red Hat Enterprise Linux.	<p>Same level of support for standard kernels as that for workloads running RHEL except that OEL 4.x is not supported for Hyper-V.</p> <p>Same level of support for Unbreakable Enterprise Kernel (UEK) kernels on supported RHEL distributions for OL 6.7 and later.</p>

NOTE: The following guidelines apply to Linux migrations to non-cloud target platforms:

- ♦ Workload imaging is not supported in Linux workloads.
- ♦ Conversion of BIOS based Linux system to UEFI based is not supported.

2.1.2 Supported Workloads for Migration to the Cloud

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

Only P2C and V2C migrations are supported.

NOTE

- ♦ Not all workloads are supported on all target cloud platforms. Migration of workloads to a cloud container is subject to the support of the guest operating system on the target cloud platform by the cloud provider.
- ♦ Windows and Linux UEFI workloads are migrated as BIOS workloads to the target cloud platform.
- ♦ Before you migrate a paravirtualized Linux source workload running on Citrix XenServer or KVM to a target platform as fully virtualized guest, see [“Paravirtualized Source Workloads” on page 37](#).

Review the following sections:

- ♦ [“Supported Workloads For Migration to Microsoft Azure” on page 27](#)
- ♦ [“Supported Workloads For Migration to VMware vCloud Director” on page 29](#)
- ♦ [“Supported Workloads For Migration to Amazon Web Services” on page 31](#)

Supported Workloads For Migration to Microsoft Azure

PlateSpin Migrate supports the following platforms for migration to Microsoft Azure Cloud for the global environment and the sovereign Azure China environment. See also [“Supported Workload Storage” on page 32](#) and [“Supported Workload Architectures” on page 36](#).

For information about migrating workloads to Microsoft Azure, see [“Prerequisites for Migration to Microsoft Azure” on page 147](#) and [“Migration to Microsoft Azure” on page 335](#).

Table 2-3 Azure: Supported Windows Platforms

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

Table 2-4 Azure: Supported Linux Platforms

Linux Distribution	Versions	Remarks
Red Hat Enterprise Linux (RHEL)	6.7 to 6.9 and 7.1 to 7.3	<p>PlateSpin Migrate does not support the XFS v5 file system on RHEL 7.3 and on distributions based on RHEL 7.3.</p> <p>For Red Hat Enterprise Linux 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version 2.6.32-642.13.1.el6) for the 6.7 distribution.</p> <p>For Red Hat Enterprise Linux 6.8, Oracle Linux 6.8, and CentOS 6.8 workloads with LVM volumes, PlateSpin Migrate supports incremental replication only for the latest available kernel (version 2.6.32-696.20.1.el6.x86_64) for the 6.8 distribution.</p>
SUSE Linux Enterprise Server (SLES)	11 (SP3 and SP4)	Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for SLES 11. See “Paravirtualized Source Workloads” on page 37.
CentOS	See Red Hat Enterprise Linux.	Same level of support as that for workloads running RHEL.
Oracle Linux (OL) (formerly Oracle Enterprise Linux)	See Red Hat Enterprise Linux.	<p>Same level of support for standard kernels as that for workloads running RHEL.</p> <p>Same level of support for Unbreakable Enterprise Kernel (UEK) kernels on supported RHEL distributions for OL 6.7 and later.</p>

NOTE: PlateSpin Migrate does not support migration for source Linux workloads to Azure if the boot (/boot) partition is on a different disk than the root (/) partition.

Supported Workloads For Migration to VMware vCloud Director

PlateSpin Migrate supports the following platforms for migration to VMware vCloud Director. See also [Section 2.1.3, “Supported Workload Storage,” on page 32](#) and [Section 2.1.4, “Supported Workload Architectures,” on page 36](#).

For information about migrating workloads to Microsoft VMware Cloud Director, see [Chapter 10, “Prerequisites for Migration to VMware vCloud Director,” on page 163](#) and [Chapter 29, “Migration to VMware vCloud Director,” on page 349](#).

Table 2-5 vCloud: Supported Windows Platforms

Operating System	Remarks
Microsoft Windows Server 2016	
Microsoft Windows Server 2012 R2	
Microsoft Windows Server 2012	
Microsoft Windows Server 2008 R2	
Microsoft Windows Server 2008	
Microsoft Windows Server 2003 R2	<code>DoNotReplaceSysFiles</code> must be set to <code>True</code> .
Microsoft Windows Server 2003 with Service Pack 1 (SP1) or later	<code>DoNotReplaceSysFiles</code> must be set to <code>True</code> .

Table 2-6 vCloud: Supported Linux Platforms

Linux Distribution	Versions	Remarks
Red Hat Enterprise Linux (RHEL)	4.x, 5.0 to 5.11, 6.0 to 6.9, and 7.0 to 7.3	<p>PlateSpin Migrate does not support the XFS v5 file system on RHEL 7.3 and on distributions based on RHEL 7.3.</p> <p>For Red Hat Enterprise Linux 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version 2.6.32-642.13.1.el6) for the 6.7 distribution.</p> <p>For Red Hat Enterprise Linux 6.8, Oracle Linux 6.8, and CentOS 6.8 workloads with LVM volumes, PlateSpin Migrate supports incremental replication only for the latest available kernel (version 2.6.32-696.20.1.el6.x86_64) for the 6.8 distribution.</p> <p>Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for RHEL 5. See “Paravirtualized Source Workloads” on page 37.</p>
SUSE Linux Enterprise Server (SLES)	10 and 11 (SP1, SP2, SP3, and SP4)	<p>Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for SLES 10 and 11. See “Paravirtualized Source Workloads” on page 37.</p>
CentOS	See Red Hat Enterprise Linux.	Same level of support as that for workloads running RHEL.
Oracle Linux (OL) (formerly Oracle Enterprise Linux)	See Red Hat Enterprise Linux.	<p>Same level of support for standard kernels as that for workloads running RHEL.</p> <p>Same level of support for Unbreakable Enterprise Kernel (UEK) kernels on supported RHEL distributions for OL 6.7 and later.</p>

Supported Workloads For Migration to Amazon Web Services

PlateSpin Migrate supports the following platforms for migration to Amazon Web Services. See also [Section 2.1.3, “Supported Workload Storage,” on page 32](#) and [Section 2.1.4, “Supported Workload Architectures,” on page 36](#).

For information about migrating workloads to Microsoft Amazon Web Services, see [Chapter 8, “Prerequisites for Migration to Amazon Web Services,” on page 137](#) and [Chapter 27, “Migration to Amazon Web Services,” on page 331](#).

Table 2-7 AWS: Supported Windows Platforms

Operating System	Remarks
Microsoft Windows Server 2016	
Microsoft Windows Server 2012 R2	
Microsoft Windows Server 2012	
Microsoft Windows Server 2008 R2	
Microsoft Windows Server 2008	
Microsoft Windows Server 2003 R2	
Microsoft Windows Server 2003 with Service Pack 1 (SP1) or later	

Table 2-8 AWS: Supported Linux Platforms

Linux Distribution	Versions	Remarks
Red Hat Enterprise Linux (RHEL)	5.1 to 5.11, 6.1 to 6.9, and 7.0 to 7.3	<p>PlateSpin Migrate does not support the XFS v5 file system on RHEL 7.3 and on distributions based on RHEL 7.3.</p> <p>For Red Hat Enterprise Linux 6.7, Oracle Linux 6.7, and CentOS 6.7 workloads with LVM volumes, incremental replication is supported only for the latest available kernel (version 2.6.32-642.13.1.el6) for the 6.7 distribution.</p> <p>For Red Hat Enterprise Linux 6.8, Oracle Linux 6.8, and CentOS 6.8 workloads with LVM volumes, PlateSpin Migrate supports incremental replication only for the latest available kernel (version 2.6.32-696.20.1.el6.x86_64) for the 6.8 distribution.</p> <p>Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for RHEL 5. See “Paravirtualized Source Workloads” on page 37.</p>
SUSE Linux Enterprise Server (SLES)	11 (SP2 and SP3)	<p>Migration of a paravirtualized source workload to a target platform as a fully virtualized workload is supported for SLES 11. See “Paravirtualized Source Workloads” on page 37.</p>
CentOS	See Red Hat Enterprise Linux.	Same level of support as that for workloads running RHEL.
Oracle Linux (OL) (formerly Oracle Enterprise Linux)	See Red Hat Enterprise Linux.	<p>Same level of support for standard kernels as that for workloads running RHEL.</p> <p>Same level of support for Unbreakable Enterprise Kernel (UEK) kernels on supported RHEL distributions for OL 6.7 and later.</p>

2.1.3 Supported Workload Storage

The following workload storage guidelines apply to all migrations:

- ♦ [“Partitioning Schemes”](#) on page 33
- ♦ [“Windows File Systems”](#) on page 33

- ♦ [“Linux File Systems” on page 33](#)
- ♦ [“Data Disks” on page 33](#)
- ♦ [“Linux Disks, Partitions, and Volumes” on page 34](#)
- ♦ [“Linux Live Data Transfer” on page 34](#)
- ♦ [“FC SANs” on page 34](#)
- ♦ [“FCoE SANs” on page 34](#)
- ♦ [“Multipath I/O” on page 35](#)

Partitioning Schemes

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

Windows File Systems

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

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

Linux File Systems

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

NOTE

- ♦ PlateSpin Migrate does not support the XFS v5 file system on RHEL 7.3 and on distributions based on RHEL 7.3.
 - ♦ Migration of encrypted volumes is not supported. If the volumes are encrypted, they must be unlocked (decrypted) for the migration.
-

Data Disks

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

NOTE: The following caveats apply for storage disks:

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

For dynamic disks, the storage does not follow the Same as Source mapping strategy. Both Simple Dynamic Volumes and Spanned Dynamic Volumes will reside on the target workload as Simple Basic Volume disks. The target disk is partitioned as GPT if the total combined size of the dynamic volume's member disks exceeds MBR partition size limits. For more information, see

Microsoft TechNet: Understanding the 2 TB limit in Windows Storage (<https://blogs.technet.microsoft.com/askcore/2010/02/18/understanding-the-2-tb-limit-in-windows-storage/>).

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

Linux Disks, Partitions, and Volumes

- ♦ Migrate supports Linux workloads with `/boot` on the first disk (`sda`).
- ♦ The boot partition of a source Linux workload must have a minimum of 100 MB free space. During the migration process, PlateSpin Migrate uses the free space to create a new `initrd` image with all the required drivers to make the machine ready for the initial boot process.
- ♦ Non-volume storage, such as a swap partition that is associated with the source workload, is recreated in the migrated workload.
- ♦ The layout of volume groups and logical volumes for LVM2 is preserved in the Same as Source mapping strategy so that you can re-create it during migration.
- ♦ LVM raw disk volumes are supported in the Same as Source configurations on Linux workloads.

Linux Live Data Transfer

- ♦ For Linux workloads, Migrate supports only block-based live data transfer with a `blkwatch` driver. For a list of pre-compiled `blkwatch` drivers, see [Section D.2.2, “List of Distributions,” on page 252](#).
- ♦ Some of the supported Linux versions require that you compile the PlateSpin `blkwatch` module for your specific kernel. Those workloads are called out explicitly.
Precompiled `blkwatch` drivers are available for the standard kernel and Unbreakable Enterprise Kernel (UEK) as noted in the [Section D.2.2, “List of Distributions,” on page 252](#). For other Oracle Linux distributions, precompiled drivers are available only for the corresponding Red Hat Compatible Kernel (RHCK).

FC SANs

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

FCoE SANs

Fibre Channel over Ethernet (FCoE) is supported for P2P and P2V migrations for workloads listed in [Table 2-9](#). Migration has been tested using FCoE devices from Qlogic.

Table 2-9 Supported Source Workloads for FCoE

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

FCoE drivers and support functionality are available in the PlateSpin ISO image. See [“Downloading the PlateSpin ISO Images” on page 279](#).

Multipath I/O

PlateSpin Migrate supports migration of a source workload configured for multipath I/O (MPIO) in a Fibre Channel (FC) SAN environment. The target workload can be in the same or different SAN environment. The source workload and target workload must have all SAN disks.

NOTE: The workload must boot from a SAN disk. Workloads with mixed local and SAN disks are not supported, except as otherwise noted in [Table 2-10](#).

MPIO support functionality is available in the PlateSpin ISO image. See [“Downloading the PlateSpin ISO Images” on page 279](#).

See [Table 2-10](#) for a list of platforms that have been tested for migrations in MPIO environments.

Table 2-10 Supported Source Workloads for MPIO

Platform	Versions	Remarks
Microsoft Windows Server	2012 R2 2008 R2	
Microsoft Windows Server in a Failover Cluster	2012 R2	The cluster migration has also been tested using a local system disk with all data disks in a FC SAN.
Red Hat Enterprise Linux (RHEL)	7.2	

MPIO requires additional multipath software to be installed in the operating system, in the form of a Windows Feature or a Linux package or module. You use MPIO management tools to enable MPIO and configure MPIO policies for SAN devices with multiple paths. Refer to vendor documentation for information about setting up hardware to provide multiple paths to a storage device and about installing and configuring MPIO.

See [Table 2-11](#) for information about supported MPIO migration scenarios and target workload expectations.

Table 2-11 Supported MPIO Migration Scenarios

Source Workload	Target Workload	
MPIO Software	Multiple Storage Paths Available	Single Storage Path Available
MPIO software is installed. MPIO is enabled and configured.	MPIO software is automatically reconfigured on the target workload for the target MPIO environment. To disable MPIO, you must manually reconfigure MPIO on the workload.	MPIO software remains, and MPIO is reconfigured for a single path. You can leave the software or remove it manually, as appropriate for your planned network. If you add MPIO hardware after migration is completed, you must manually reconfigure MPIO on the workload.

Source Workload	Target Workload	
MPIO Software	Multiple Storage Paths Available	Single Storage Path Available
MPIO software is installed. MPIO is disabled.	MPIO software remains installed but disabled. To enable MPIO, you must manually configure MPIO on the workload.	MPIO software remains installed but disabled. You can leave the software or remove it manually, as appropriate for your planned network. If you add MPIO hardware after migration is completed, you must manually configure MPIO on the workload.
MPIO software is not installed.	MPIO software is not installed. To enable MPIO, you must manually install and configure MPIO on the workload.	No MPIO-related changes are made for the workload.

2.1.4 Supported Workload Architectures

The following workload architecture guidelines apply to all migrations:

- ♦ [“Protocols” on page 36](#)
- ♦ [“Processors” on page 36](#)
- ♦ [“Cores and Sockets for Target VMs” on page 36](#)
- ♦ [“Virtual CPUs for Target VMs” on page 37](#)
- ♦ [“UEFI and BIOS Firmware” on page 37](#)
- ♦ [“Paravirtualized Source Workloads” on page 37](#)

Protocols

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

Processors

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

- ♦ 64-bit
- ♦ 32-bit

Cores and Sockets for Target VMs

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

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

Some distributions of a guest OS might not honor the cores and cores per socket configuration. For example, guest OSes using SLES 10 SP4 retain their original cores and sockets settings as installed, whereas other SLES and RHEL distributions honor the configuration.

Virtual CPUs for Target VMs

For VM containers using VMware 4.1, PlateSpin Migrate enables you to specify the required number of vCPUs (virtual CPUs) to assign to the target workload. This parameter applies on the initial setup of a workload with an initial replication setting of **Full Replication**. Each vCPU is presented to the guest OS on the VM container as a single core, single socket.

UEFI and BIOS Firmware

Migration of UEFI-based Windows and Linux source workloads is supported for all target platforms. The target workload is converted from UEFI to BIOS for target cloud platforms: Amazon Web Services, Microsoft Azure, and VMware vCloud Director. For other platforms, the target workload is configured as UEFI or BIOS, as supported by the target platform vendor.

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

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

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

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

Paravirtualized Source Workloads

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

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

Only block-based conversions are supported.

Before you migrate a paravirtualized Linux source workload running on Citrix XenServer or KVM to a target platform as fully virtualized guest, do the following:

- ♦ Ensure that both the paravirtual and standard kernel are installed on the paravirtualized source workload.
- ♦ Manually compile the block-based drivers for Xen kernel.

2.1.5 Supported Target Virtualization Platforms

PlateSpin Migrate supports the following target virtualization platforms.

- ♦ [Table 2-12](#) lists supported target VMware platforms for the PlateSpin Migrate Web Interface and Migrate Client. The Migrate Client supports automated migration or the semi-automated migration using the X2P workflow. The Web interface supports automated migration. See:
 - ♦ [Automated Migration to VMware Using Migrate Client](#)
 - ♦ [Migration to VMs on VMware Using X2P Workflow](#)
 - ♦ [Automated Migration to VMware Using Migrate Web Interface](#)

See also [Prerequisites for Migration to VMware](#).

NOTE: For information about creating the target VM disk on VMware containers using Raw Device Mapping (RDM), see [Migration to VMware](#).

- ♦ [Table 2-14](#) lists supported target virtualization platforms for the PlateSpin Migrate Client using the semi-automated X2P workflow.

NOTE

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

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

Platform	Versions	Remarks
VMware vCenter	6.5 6.0 (U1, U2, and U3) 5.5 (U1, U2, and U3) 5.1 (U1, U2, and U3) 5.0 (U1, U2, and U3) 4.1 (U1, U2, and U3)	For vCenter 6.0, Migrate supports VMware Virtual SAN (vSAN) 6.2 storage. For vCenter 5.5, Migrate supports VMware Virtual SAN (vSAN) 5.5 storage. Raw Device Mapping (RDM) for target VMs is supported using the X2P workflow. The following VMware datastore types available in VMware 6.0 and later are not supported: See also Table 2-13, "Supported VMware Datastores," on page 39.

Platform	Versions	Remarks
VMware ESXi	6.5 6.0 (U1, U2, and U3) 5.5 (U1, U2, and U3) 5.1 (U1, U2, and U3) 5.0 (U1, U2, and U3) 4.1 (U1, U2, and U3)	All ESXi versions must have a paid license; migration is unsupported with these systems if they are operating with a free license. Raw Device Mapping (RDM) for target VMs is supported using the X2P workflow. See also Table 2-13, “Supported VMware Datastores,” on page 39.
VMware ESX	4.1 (U1, U2, and U3)	Raw Device Mapping (RDM) for target VMs is supported using the X2P workflow. See also Table 2-13, “Supported VMware Datastores,” on page 39.

Table 2-13 Supported VMware Datastores

Datastore Type	Supported Configurations
VMFS	Supported for all supported versions of VMware vCenter, ESXi, and ESX platforms.
NFS v3	Supported only for supported versions of VMware vCenter 5.5 and ESXi 5.5 platforms.
Other	Other datastore types are not supported, such as Virtual Volumes, NFS v4.1, and vFlash.

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

Platform	Versions	Remarks
Microsoft Windows Server with Hyper-V	Windows Server 2012 R2 Windows Server 2012	Supported for automated workflow or the X2P workflow. See <ul style="list-style-type: none"> ♦ Automated Migration to Hyper-V ♦ Migration to VMs on Hyper-V Using X2P Workflow See also “Prerequisites for Migration to Microsoft Hyper-V” on page 175.

Platform	Versions	Remarks
Citrix XenServer	6.0 to 6.2, and 6.5	<p>Fully virtualized guests are supported.</p> <p>Use the PlateSpin ISO for SUSE Linux Enterprise Server 11 SP3 as the LRD to boot the virtual machine. Citrix XenServer 6.5 and earlier does not support SLES 11 SP4.</p> <p>Supported through the X2P workflow. See Migration to Virtual Machines on Citrix XenServer.</p> <p>See also “Prerequisites for Migration to VMs on Citrix XenServer” on page 179.</p>
SUSE Linux Enterprise Server with Xen	11 SP3 and 11 SP4	<p>Fully virtualized guests are supported.</p> <p>Supported through the X2P workflow. See Migration to Virtual Machines on Xen.</p> <p>See also “Prerequisites for Migration to VMs on Xen” on page 183.</p>
SUSE Linux Enterprise Server (SLES) with KVM	11 SP4 and 12 SP1	<p>Fully virtualized guests are supported.</p> <p>Virtio devices are supported.</p> <p>Supported through the X2P workflow. See Migration to Virtual Machines on KVM.</p> <p>See also “Prerequisites for Migration to VMs on KVM” on page 185.</p>
Red Hat Enterprise Linux (RHEL) with KVM	7.0 and 7.2	<p>Fully virtualized guests are supported.</p> <p>Virtio devices are supported.</p> <p>Supported through the X2P workflow. See Migration to Virtual Machines on KVM.</p> <p>See also “Prerequisites for Migration to VMs on KVM” on page 185.</p>

2.1.6 Supported Target Cloud Platforms

PlateSpin Migrate supports migration of workloads to target cloud platforms. Microsoft Azure and VMware vCloud Director are supported in the Migrate Web Interface. Amazon Web Services is supported in the Migrate Client. See [Table 2-15](#) and [Table 2-16](#).

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

Platform	Versions	Remarks
Microsoft Azure	Azure Cloud global environment Sovereign Azure China environment	Each PlateSpin Server can support only one of the supported Azure environments at a time. See Section 9.9, “Configuring the Target IaaS Environment for Azure Cloud,” on page 159. NOTE: Contact Micro Focus Support and Services for assistance if you want to migrate workloads to other sovereign Azure environments. See also “Prerequisites for Migration to Microsoft Azure” on page 147.
VMware vCloud Director	5.5.x and 5.6.x	See also “Prerequisites for Migration to VMware vCloud Director” on page 163.

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

Platform	Versions	Remarks
Amazon Web Services (AWS)	Amazon EC2 environment	See also “Prerequisites for Migration to Amazon Web Services” on page 137.

2.1.7 Supported International Languages

In addition to English, PlateSpin Migrate provides National Language Support (NLS) for Chinese Simplified (ZH-CN), Chinese Traditional (ZH-TW), French (FR-FR), German (DE-DE), and Japanese (JA-JP).

Localized online documentation is available in these languages, as well as in Spanish (ES-ES) and Brazilian Portuguese (PT-BR).

2.1.8 Supported Web Browsers

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

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

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

To use the Web Interface in one of the supported international languages, see [“Configuring Language Settings for International Versions” on page 99](#).

2.2 Supported Data Transfer Methods

Depending on the selected workload and the migration type, PlateSpin Migrate enables you to select different methods for transferring workload data from the source to the target.

For information on how to select a transfer method, see [“Conversion \(Data Transfer Method\)” on page 300](#).

- ♦ [Section 2.2.1, “File-Level Transfer \(Live\),” on page 42](#)
- ♦ [Section 2.2.2, “Block-Level Transfer \(Live\),” on page 42](#)
- ♦ [Section 2.2.3, “Offline Transfer with Temporary Boot Environment,” on page 43](#)

2.2.1 File-Level Transfer (Live)

The File-Based Live Transfer method, available for Windows workloads, copies data and replicates changes at the file level.

To ensure data consistency, this method leverages the Microsoft Volume Shadow Copy Service (VSS) if available. Many enterprise apps are integrated with VSS; for those which are not, PlateSpin Migrate provides the capability to briefly pause services while the VSS snapshot is captured, to ensure that the data of those applications is captured in a consistent state.

If VSS unavailable (for example, in workloads running Windows Server 2003 with no service packs), PlateSpin Migrate monitors source volumes for changes while transferring data. When the initial transfer is complete, migrate re-sends any files that have changed. If the rate of file system changes is consistently high, data transfer is stopped and a job progress warning is shown.

You can configure your migration job to stop high-transaction services, such as Microsoft SQL Server or Microsoft Exchange Server, during the transfer (see [“Services or Daemons to Stop before Replication or Cutover” on page 304](#)). This has two benefits:

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

This method might be appropriate for moderately active systems and it provides you with the capability to resize your volumes on the target workload.

2.2.2 Block-Level Transfer (Live)

The Block-Based Live Transfer method, available for both Windows and Linux workloads, enables PlateSpin Migrate to transfer data at the block level, providing an exact copy of the source workload.

For Windows workloads, PlateSpin Migrate leverages the Microsoft Volume Snapshot Service (VSS) (Windows 2003 SP1 and later) with applications and services that support VSS.

For Linux workloads, Migrate supports only block-based data transfer with a blkwatch driver. The Migrate distribution includes precompiled blkwatch drivers for workloads running the standard, non-debug kernels of supported Linux distributions. See [Section D.2, “Pre-compiled blkwatch Drivers for Linux Distributions,” on page 252](#).

If your workloads have a non-standard, customized, or newer kernel, you can build a custom blkwatch driver for your specific kernel. See [Knowledgebase Article 7005873 How to Build a Custom Block-Based Linux Kernel Driver](https://www.netiq.com/support/kb/doc.php?id=7005873) (<https://www.netiq.com/support/kb/doc.php?id=7005873>).

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

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

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

2.2.3 Offline Transfer with Temporary Boot Environment

This method enables PlateSpin Migrate to boot your source machine into a temporary pre-execution environment and transfer the data while the source is offline. This method is not applicable with the PlateSpin Migrate Web Interface.

NOTE: The Offline Transfer method lets you migrate the Windows Server 2003 SP0 workloads:

Before you use the Offline Transfer method to migrate a Windows Server 2003 workload, you must do the following:

1. Edit the `boot.ini` file on the workload to set the `/noexecute` parameter to `alwaysoff`.
2. Restart the workload.

The pre-execution environment underlying the Offline transfer method makes use of a Linux RAMDisk (LRD), which contains a minimal set of system files, drivers, and executables, sufficient for an initial, temporary boot. To ensure that the source operating system properly loads and operates in the temporary pre-execution environment, PlateSpin Migrate temporarily modifies its boot files and restores them to their original state after the pre-execution environment has successfully loaded.

The RAMDisk is also used to temporarily boot target physical machines in X2P migrations, as well as to boot target VMs in semi-automated migrations. See “[Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO](#)” on page 204, and “[Registering and Discovering Details for Target Physical Machines with PlateSpin ISO](#)” on page 206.

2.3 Security and Privacy

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

- [Section 2.3.1, “Secure Connections Using TLS 1.2,” on page 44](#)
- [Section 2.3.2, “Security of Workload Data in Transmission,” on page 46](#)
- [Section 2.3.3, “Security of Client-Server Communications,” on page 46](#)
- [Section 2.3.4, “Security of Credentials,” on page 46](#)
- [Section 2.3.5, “User Authorization and Authentication,” on page 46](#)

2.3.1 Secure Connections Using TLS 1.2

PlateSpin Migrate Server supports secure Transport Layer Security (TLS) 1.2 connections if TLS 1.2 is enabled in your environment. Migrate does not have a setting that forces clients to use TLS 1.2 to connect.

By default, PlateSpin Migrate Server accepts the .NET Framework default of TLS 1.0 for source workload connections. You can use Microsoft Windows Registry settings on source Windows workloads to force .NET Framework to choose TLS 1.2 when the workload connects with Migrate Server. A Microsoft update for .NET Framework might be required on the source workload in order to add support for TLS System Default Version settings. A reboot is required.

Configuring Source Workloads to Connect Using TLS 1.2

Microsoft .NET Framework versions earlier than 4.7 do not provide support for applications to use TLS System Default Versions as a cryptographic protocol. They use older versions of the TLS protocol (1.0/1.1), even if more secure protocols are available on the system. This default is not changed by PlateSpin. You can change the TLS system default version for all .NET applications to the more secure TLS 1.2 version.

An administrator user can apply a Microsoft Windows update called *Support for TLS System Default Versions included in the .NET Framework* on the source workload to add support for TLS 1.2 connections. The specific update depends on the operating system on the source workload.

After you apply the update, modify the Windows Registry settings for the `SystemDefaultTlsVersions` and `SchUseStrongCrypto` keys. The `SystemDefaultTlsVersions` settings change the behavior of .NET 3.5 and 2.0 applications. The `SchUseStrongCrypto` settings change the behavior of .NET 4.x applications. For more information, refer to the [Microsoft Security Advisory 2960358](#) article.

NOTE: The Registry settings for .NET security are machine-wide settings that will impact all .NET applications on the source system.

To enable TLS 1.2 as the default cryptographic protocol for source workloads:

- 1 Log in to the source workload as a user with administrative privileges.
- 2 Download and install the Windows update *Support for TLS System Default Versions included in the .NET Framework*, according to the Windows operating system on the source workload.

Windows Operating System	Microsoft .NET Framework Version	Windows Update with TLS 1.2 Support
Windows Server 2016	4.7 or higher	No patch is required.
Windows Server 2012 R2 Windows 8.1	3.5 SP1	Support for TLS System Default Versions included in the .NET Framework version 3.5 SP1 (Microsoft KB Article 3154520)
Windows Server 2012 Windows 8	3.5	Support for TLS System Default Versions included in the .NET Framework version 3.5 (Microsoft KB Article 3154519)
Windows Server 2008 R2 SP1 Windows 7 SP1	3.5.1	Support for TLS System Default Versions included in the .NET Framework version 3.5.1 (Microsoft KB Article 3154518)
Windows Server 2008 SP2	2.0 SP2	Support for TLS System Default Versions included in the .NET Framework version 2.0 SP2 (Microsoft KB Article 3154517)
Windows Server 2003		Windows Server 2003 does not support TLS 1.2.

- 3 Direct .NET Framework on the source workload to use a more secure version of TLS by using the Windows Registry Editor to apply the following Windows Registry settings:

Registry Settings for 64-bit Systems

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\.NETFramework\v2.0.50727]
```

```
"SystemDefaultTlsVersions"=dword:00000001
```

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Microsoft\.NETFramework\v2.0.50727]
```

```
"SystemDefaultTlsVersions"=dword:00000001
```

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\.NetFramework\v4.0.30319]
```

```
"SchUseStrongCrypto"=dword:00000001
```

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Microsoft\.NetFramework\v4.0.30319]
```

```
"SchUseStrongCrypto"=dword:00000001
```

Registry Settings for 32-bit Systems

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\.NETFramework\v2.0.50727]
```

```
"SystemDefaultTlsVersions"=dword:00000001
```

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\.NetFramework\v4.0.30319]
```

```
"SchUseStrongCrypto"=dword:00000001
```

- 4 (Windows Server 2008 R2 and earlier) Enable TLS 1.2 on the sourced workload by using the Windows Registry Editor to apply the following Windows Registry change:

Key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\TLS 1.2\Client
```

```
Value Name: "DisabledByDefault"
```

```
Value Type: DWORD
```

```
Value Data: 0
```

This change is required because Windows Server 2008 R2 and earlier operating systems disable TLS 1.2 by default, while Windows Server 2012 and later operating systems enable TLS 1.2 by default.

- 5 Reboot the source workload.

2.3.2 Security of Workload Data in Transmission

To make the transfer of your workload data more secure, you can configure your migration jobs to encrypt the data in transit to the target. When encryption is enabled, over-the-network data transfer from the source to the target is encrypted by using 128-bit Advanced Encryption Standard (AES). For information about how to enable encryption during data transfer for a migration job, see [Section 26.11, “Conversion \(Data Transfer Method\),” on page 300](#).

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

If FIPS is enabled in a source workload, ensure that the `EnforceFIPSCompliance` parameter is enabled on the PlateSpin Migrate server before you discover the source workload. See [Section 5.3, “Enforcing FIPS Compliance for FIPS-Enabled Source Workloads,” on page 100](#).

2.3.3 Security of Client-Server Communications

Data transmission between the PlateSpin Server and the PlateSpin Migrate Client can be configured to use either HTTP (default) or HTTPS (Secure Hypertext Transfer Protocol). To secure data transmission between the client and the server, enable SSL on your PlateSpin Server host and use HTTPS when specifying the server URL. See [“Connecting to a PlateSpin Migrate Server” on page 59](#).

2.3.4 Security of Credentials

Credentials that you use to access sources and targets in workload migration jobs are:

- Cached, encrypted, and securely stored by the PlateSpin Migrate Client, by using operating system APIs.
- Stored in the PlateSpin Migrate database and are therefore covered by the same security safeguards that you have in place for PlateSpin Server hosts.
- Included within diagnostics, which are accessible to accredited users. You should ensure workload migration projects are handled by authorized staff.

2.3.5 User Authorization and Authentication

PlateSpin Migrate provides a role-based user authorization and authentication mechanism. See [Section 4.1, “Configuring User Authorization and Authentication,” on page 81](#).

NOTE: If you have installed a PlateSpin Migrate Server localized for one language and a PlateSpin Migrate Client localized for a different language, do not use authorization credentials that include any language-specific characters. Using such characters in the login credentials causes miscommunication between the client and the server: the credentials are rejected as invalid.

2.4 Performance

- ♦ [Section 2.4.1, “About Product Performance Characteristics,” on page 47](#)
- ♦ [Section 2.4.2, “Scalability,” on page 47](#)
- ♦ [Section 2.4.3, “Data Compression,” on page 48](#)
- ♦ [Section 2.4.4, “Bandwidth Throttling,” on page 48](#)
- ♦ [Section 2.4.5, “Blackout Window,” on page 48](#)

2.4.1 About Product Performance Characteristics

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

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

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

- ♦ Preparation and network setup
- ♦ Source workload and target machine discovery
- ♦ Target configuration

2.4.2 Scalability

You can set up multiple workload migrations and run them concurrently. Scalability testing performed with VMware ESX hosts suggests the following benchmark recommendations:

- ♦ Multiple migrations to a single VM server: no more than 10
- ♦ Multiple migrations against multiple VM servers (ESX): no more than 40

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

Performance for concurrent migrations depends on the resources on the PlateSpin Migrate Server and the target environment, as well as the available bandwidth. We recommend that you begin with a low load, then increase it and see how the migrations perform in your environment.

We recommend that you keep no more than 50 discovered workloads at a time in the inventory for your PlateSpin Migrate server, depending on its available hardware resources. As you complete workload migrations, you can remove workloads and add others. Use the scheduled start dates to control when migrations begin and how many migration jobs are scheduled to run concurrently.

The Mass Discover CLI enables you to discover multiple workloads from a CSV file. The related migration jobs start according to the schedules you set for them.

2.4.3 Data Compression

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

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

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

Data Compression can be configured per migration job. You can also use the PlateSpin Migrate Client to specify a default compression value to be applied globally. See [Section 6.2, "Configuring Job Values Defaults,"](#) on page 118.

To set the data compression level for the migration job using the PlateSpin Migrate Web Interface, see Compression Level setting in the [Section 26.1.1, "Configuration Workflows Using Migrate Client,"](#) on page 292.

2.4.4 Bandwidth Throttling

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

Bandwidth throttling is a parameter of a workload migration job's configuration properties. To apply bandwidth throttling for the migration job, see ["Bandwidth Throttling during Data Transfer"](#) on page 300.

2.4.5 Blackout Window

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

2.5 Database Server

PlateSpin Migrate includes Microsoft SQL Server Express Edition. The capabilities of SQL Server Express are sufficient for the scalability characteristics described in [Section 2.4.2, "Scalability,"](#) on page 47.

NOTE: Microsoft SQL Server Express has a database size limit of 10 GB and can use only one CPU core at a time and 1 GB memory. For more information about requirements and limitations for SQL Server Express, see the [Microsoft SQL Server 2014 Express documentation \(https://www.microsoft.com/en-us/download/details.aspx?id=42299\)](https://www.microsoft.com/en-us/download/details.aspx?id=42299).

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

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

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

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

2.6 Access and Communication Requirements across your Migration Network

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

- [Section 2.6.1, “Requirements for Discovery,” on page 49](#)
- [Section 2.6.2, “Requirements for Migration,” on page 50](#)
- [Section 2.6.3, “Requirements for Event Messaging,” on page 52](#)
- [Section 2.6.4, “Migrations Across Public and Private Networks through NAT,” on page 53](#)

2.6.1 Requirements for Discovery

The following table lists software, network, and firewall requirements that systems in your environment must meet for the discovery and inventory process. For information about discovery procedures, see [Part IV, “Discovering and Preparing Workloads and Targets,” on page 195](#).

Table 2-17 Network Communication Prerequisites for Discovery Operations

System	Prerequisites
All workloads	Ping (ICMP echo request and response) support

System	Prerequisites
All Windows sources and Hyper-V hosts	<ul style="list-style-type: none"> ♦ Microsoft .NET Framework version 2.0 SP2, 3.5 SP1, or 4.0 ♦ Requires credentials equivalent to built-in Administrator or a domain account Administrator credentials with access to Admin\$ share (membership only in the local Administrators group is insufficient). ♦ The Windows Firewall configured to allow File and Printer Sharing. Use one of these options: <ul style="list-style-type: none"> ♦ Option 1, using Windows Firewall: Use the basic Windows Firewall Control Panel item (firewall.cpl) and select File and printer Sharing in the list of exceptions. - OR - ♦ Option 2, using Windows Firewall with Advanced Security: Use the Windows Firewall with Advanced Security utility (wf.msc) with the following Inbound Rules enabled and set to Allow: <ul style="list-style-type: none"> ♦ File and Printer Sharing (Echo Request - ICMPv4In) ♦ File and Printer Sharing (Echo Request - ICMPv6In) ♦ File and Printer Sharing (NB-Datagram-In) ♦ File and Printer Sharing (NB-Name-In) ♦ File and Printer Sharing (NB-Session-In) ♦ File and Printer Sharing (SMB-In) ♦ File and Printer Sharing (Spooler Service - RPC) ♦ File and Printer Sharing (Spooler Service - RPC-EPMAP) ♦ The Windows Firewall configured to allow Windows Management Instrumentation (WMI-In). ♦ (Conditional) If the volumes are encrypted with the BitLocker disk encryption feature, they must be unlocked.
All Linux sources Citrix XenServer Linux Xen or KVM servers	<ul style="list-style-type: none"> ♦ Secure Shell (SSH) server ♦ Open port 22 (TCP) ♦ 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). ♦ Custom SSH ports are supported; specify the port number during discovery: <i><hostname IP_address>.port_number</i>.
VMware ESX/ESXi Servers	<ul style="list-style-type: none"> ♦ VMware account with an Administrator role ♦ VMware Web services API and file management API (HTTPS / port 443 TCP)
VMware vCenter Servers	The user with access must be assigned the appropriate roles and permissions. Refer to the pertinent release of VMware documentation for more information.

2.6.2 Requirements for Migration

The following table lists firewall requirements that systems in your environment must meet for problem-free operation during workload migration jobs.

Table 2-18 Network Communication Prerequisites for Workload Migration

System	Open Port (Default)	Remarks
PlateSpin Server hosts	Either TCP 80 or TCP 443 TCP	<ul style="list-style-type: none"> Port 80 (TCP) is required for HTTP communication among the PlateSpin Server, sources, and targets. Port 443 (TCP) is required for HTTPS communication (if SSL is used) between the PlateSpin Server and the source or target machines.
All source workloads except those in image deployment jobs.	TCP 3725	<p>Required for targets to initiate communication during file-level data transfer, except for I2X jobs, during which this port needs to be open on the migration target only. For Server Sync jobs, this port is required for both sources and targets.</p> <p>The port number is configurable by setting the FileTransferPort parameter in the PlateSpin Configuration settings for the Migrate server.</p> <p>When the PlateSpin Migrate server is installed on-premise, by default the target workload will connect to the source workload on port 3725 (TCP), although this setting can be reversed (source workload connects to target workload) by changing the SourceListensForConnection parameter setting from True to False.</p> <p>When the PlateSpin Migrate server is deployed in the cloud from the provided cloud-based PlateSpin Migrate server image, the default direction of this connection is reversed automatically: the on-premise source workload will connect to the target workload in the cloud on port 3725 (TCP).</p>
All targets	TCP 3725	<p>Required for:</p> <ul style="list-style-type: none"> File-level Server Sync Image synchronization jobs
All Windows sources and targets	NetBIOS 137 - 139	Required for NetBIOS communications.
All Windows Server Cluster workloads. See "Clusters" on page 25 .		<p>Ensure that the PlateSpin Server can resolve DNS forward lookup and reverse lookup for the IP addresses of the Windows Server Cluster and its cluster nodes. You can update the DNS server or update the local <code>hosts</code> file (<code>%systemroot%\system32\drivers\etc\hosts</code>) on the PlateSpin Server.</p>
All sources	SMB (TCP 139, 445 and UDP 137, 138)	Required for communication and file-level data transfer during offline migration.
All Linux sources Citrix Xen Server Linux Xen or KVM servers	TCP 22	Required for communication during offline migration.

System	Open Port (Default)	Remarks
PlateSpin Server hosts; All Windows sources	TCP 135/445	For DCOM/RPC communication between PlateSpin Server and a source for taking control of and rebooting the workload through WMI. NOTE: WMI (RPC/DCOM) can use TCP ports 135 and 445 as well as random/dynamically assigned ports above 1024.

2.6.3 Requirements for Event Messaging

Table 2-19 shows the protocol and port required for event messaging in a PlateSpin Migration Factory environment. These messages reflect events and state changes and do not contain sensitive information.

Table 2-19 Event Messaging Requirements for Network Protocols and Ports

Traffic	Network Protocol and Port	Other Requirements
Event Messaging	Stomp, port 61613, TCP incoming (not secure)	<p>This port is open by default on the PlateSpin Transformation Manager Appliance, which includes a pre-installed instance of PlateSpin Migrate Connector.</p> <p>You must manually open the port on the following:</p> <ul style="list-style-type: none"> On each PlateSpin Migrate server that you use as a Migration Server resource in a Transformation Manager project <p>For a cloud-based Migrate server, allow inbound connections for STOMP traffic in its Network Security Group.</p> <ul style="list-style-type: none"> On each PlateSpin Migrate Connector host server for standalone Connector instances that are assigned to a Transformation Manager project On firewalls between each Migrate Connector host and the PlateSpin Transformation Manager Appliance On firewalls between each Migrate Connector host and each PlateSpin Migrate server that you use as a Migration Server resource in a Transformation Manager project

2.6.4 Migrations Across Public and Private Networks through NAT

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

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

- ♦ **PlateSpin Server:** In your server's *PlateSpin Server Configuration* tool, record the additional IP addresses assigned to that host:

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

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

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

`10.50.186.147;10.50.186.148`

- ♦ **Source:** As part of that specific migration job, record the additional IP addresses assigned to that workload. See [“Network Identification \(Network Connections\)” on page 315](#).
- ♦ **Target:** When you are attempting to discover a target, such as VMware ESX, specify the public (or external) IP address in the discovery parameters.

2.7 Deciding on the Migration Interface

PlateSpin Migrate includes the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface to allow you to efficiently plan, configure, execute, and test migrations. The PlateSpin Migrate Web Interface supports large scale migration of workloads to VMware containers and to cloud platforms such as Microsoft Azure and VMware vCloud Director. The PlateSpin Migrate Client supports migration of workloads to VMware containers, to the Amazon Web Services cloud platform, physical machines, and virtual machines on other virtual hosts.

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

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

For example:

- ♦ X2P conversions, migration to non-VMware hosts, and migration to Amazon Web Services can be performed only from the PlateSpin Migrate Client.
- ♦ Migration to Microsoft Azure and VMware vCloud Director is possible only from the PlateSpin Migrate Web Interface.
- ♦ Migration to VMware is possible from both the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface.

For a list of migration operations that you can perform using the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface, see [Section 3.4, “Migration Operations Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface,” on page 75](#).

For a list of migration tasks that you can perform using the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface, see [Section 3.5, “Migration Tasks Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface,” on page 76](#).

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



Frequently Asked Questions

This section provides answers to frequently asked questions.

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

Your PlateSpin Migrate product's overall performance, including data transfer speeds and scalability, depend on a variety of factors in your specific environment. See ["Performance" on page 47](#).

How secure is my PlateSpin Migrate product?

PlateSpin Migrate provides several features to help you safeguard your data and increase security. See ["Security and Privacy" on page 43](#).

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

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

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

Yes. See ["Target Discovery in the Migrate Client" on page 200](#).

Can multiple migrations run simultaneously?

Yes. See ["Performance" on page 47](#).

Working With Your PlateSpin Server

This section provides information on typical, usually one-time configuration tasks following product installation. For installation information, see the [PlateSpin Migrate 12.2.1 Installation and Upgrade Guide](#).

- ♦ Chapter 3, “Using the PlateSpin Migrate Tools,” on page 59
- ♦ Chapter 4, “Configuring PlateSpin Users and Access,” on page 81
- ♦ Chapter 5, “Configuring PlateSpin Migrate Server,” on page 91
- ♦ Chapter 6, “Configuring PlateSpin Migrate Client,” on page 117
- ♦ Chapter 7, “Configuring PlateSpin Migrate Web Interface,” on page 125
- ♦ Appendix B, “Rebranding the UI for PlateSpin Migrate Web Interface,” on page 129

3 Using the PlateSpin Migrate Tools

This section introduces the PlateSpin Migrate tools and how you use them to carry out workload migration and management tasks. To interact with the product and perform tasks such as discovery of source workloads and target hosts; setting up, executing, and monitoring jobs; managing license keys; and configuring the default behavior of the server, use either the PlateSpin Migrate Client or the browser-based PlateSpin Migrate Web Interface. To decide which interface to use, see [Section 2.7, “Deciding on the Migration Interface,”](#) on page 53.

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

- [Section 3.1, “Connecting to a PlateSpin Migrate Server,”](#) on page 59
- [Section 3.2, “About the PlateSpin Migrate Client User Interface,”](#) on page 61
- [Section 3.3, “About the PlateSpin Migrate Web Interface,”](#) on page 68
- [Section 3.4, “Migration Operations Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface,”](#) on page 75
- [Section 3.5, “Migration Tasks Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface,”](#) on page 76
- [Section 3.6, “Other PlateSpin Server Management Tools,”](#) on page 78

3.1 Connecting to a PlateSpin Migrate Server

- [Section 3.1.1, “PlateSpin Server Access Using the Migrate Client,”](#) on page 59
- [Section 3.1.2, “PlateSpin Server Access Using the Migrate Web Interface,”](#) on page 61

3.1.1 PlateSpin Server Access Using the Migrate Client

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

- Performs authentication of the specified user account with the PlateSpin Server.
See [Section 4.1, “Configuring User Authorization and Authentication,”](#) on page 81.
- Connects to a specified PlateSpin Server.
- Loads a specified PlateSpin Migrate Network, a collection of discovered source workloads and targets that you work with at one time.

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

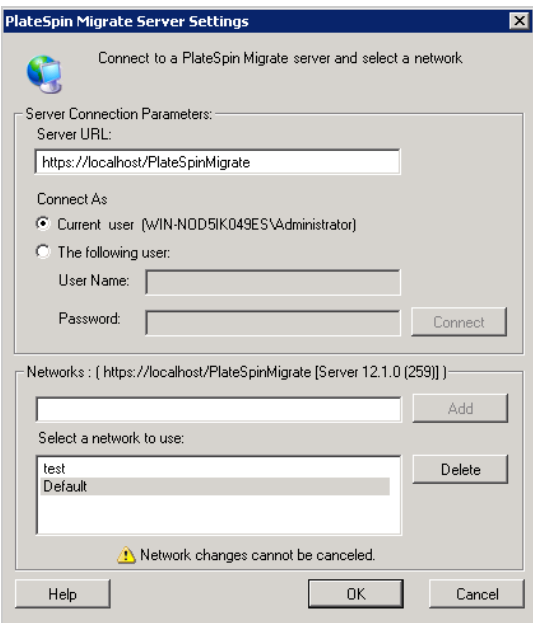
- 1 In the PlateSpin Migrate Client, click **Tools > PlateSpin Server Settings**.

or

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

Server: http://dev3/PlateSpinMigrate	Network: Sample Environment	User: devdomain\john.smith
--------------------------------------	-----------------------------	----------------------------

The PlateSpin Server Settings dialog box opens.



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

Interface Element	Description
Server URL	<p>Type the PlateSpin Server URL in the following format:</p> <pre>http://server_host/platespinmigrate</pre> <p>If SSL is enabled on the PlateSpin Server host, replace <code>http</code> in the URL with <code>https</code>.</p> <p>We recommend that you specify the fully qualified domain name (FQDN) if you are using a domain user account to log into Migrate Server.</p>
Connect As	<p>To connect to a PlateSpin Server, you must have administrative access to the PlateSpin Server host or be a member of one of the PlateSpin Migrate roles. See Section 4.1, “Configuring User Authorization and Authentication,” on page 81.</p>
Networks	<p>To familiarize yourself with PlateSpin Migrate features, use the Sample Environment network. To work with actual source workloads and targets, use the Default network or create your own.</p> <p>To add a network, type the name, then click Add.</p> <p>To remove a network, select it, then click Delete.</p>

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

3.1.2 PlateSpin Server Access Using the Migrate Web Interface

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

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

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

To launch PlateSpin Migrate Web Interface:

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

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

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

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

For information about setting up additional users for PlateSpin, see [Section 4.1, “Configuring User Authorization and Authentication,” on page 81](#).

3.2 About the PlateSpin Migrate Client User Interface

The PlateSpin Migrate Client provides a management tool to manage migrations a variety of virtual host targets, Amazon Web Services cloud targets, physical targets, PlateSpin Image Server targets, and server-sync.

For information about installing the Migrate Client, see “[System Requirements for PlateSpin Migrate Client](#)” and “[Installing the PlateSpin Migrate Client](#)” in the *PlateSpin Migrate 12.2.1 Installation and Upgrade Guide*.

For information about configuration options for the Migrate Client, see [Chapter 6, “Configuring PlateSpin Migrate Client,” on page 117](#).

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

- ♦ [Section 3.2.1, “Navigating the Client Interface,” on page 61](#)
- ♦ [Section 3.2.2, “Servers View,” on page 62](#)
- ♦ [Section 3.2.3, “Jobs View,” on page 67](#)
- ♦ [Section 3.2.4, “Tasks Pane,” on page 67](#)
- ♦ [Section 3.2.5, “Status Bar,” on page 68](#)
- ♦ [Section 3.2.6, “Workload Migration Tasks,” on page 68](#)

3.2.1 Navigating the Client Interface

The PlateSpin Migrate Client window consists of the following elements:

- ♦ **Menu bar:** Reflects the current view and provides command groups for accessing program features and operations.

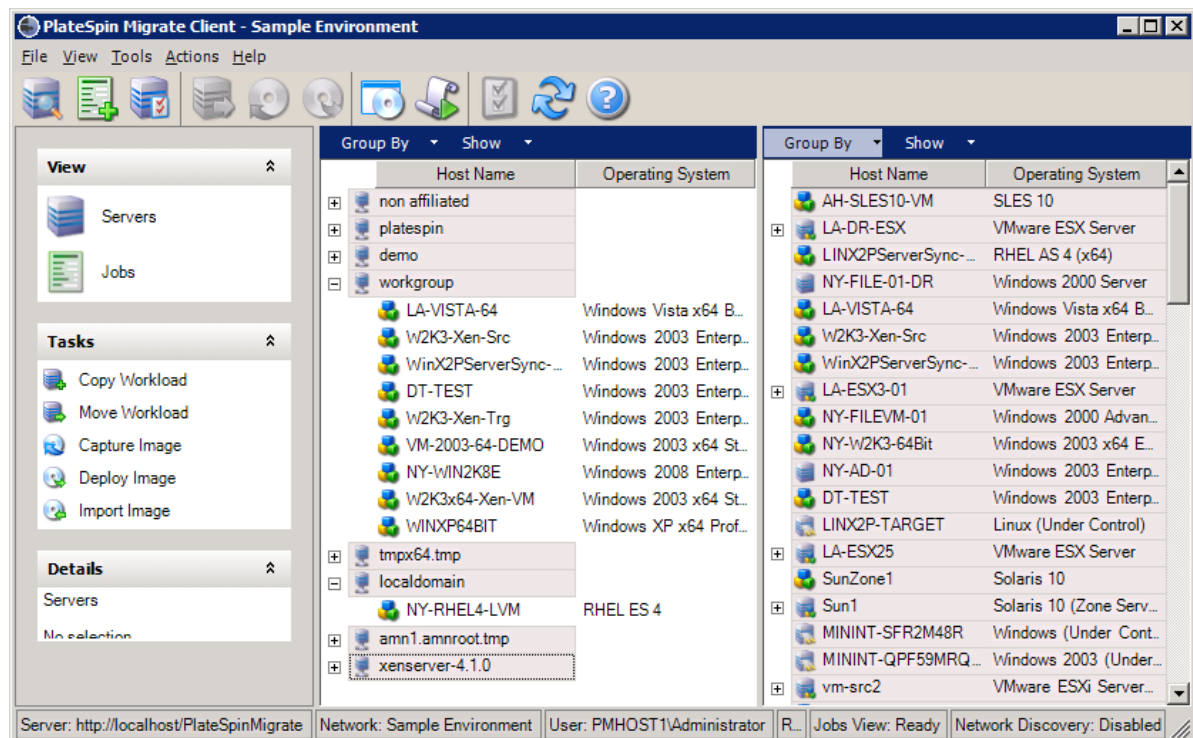
- ♦ **Toolbar:** Reflects the current view and provides visual shortcuts to program features and operations.
- ♦ **Servers View:** The Servers view is the main visual interface to your discovered source workloads and targets. See [Section 3.2.2, “Servers View,” on page 62](#).
- ♦ **Jobs View:** The Jobs view displays all jobs, such as discovery, migration, and image capture. See [Section 3.2.3, “Jobs View,” on page 67](#).
- ♦ **Current view:** The main work area of the interface; lists either machines (when in Servers view mode), or jobs (when in Jobs view mode).
- ♦ **Panes:** Vertically aligned at the left side of the window, panes facilitate the selection of the current view (View pane) or a migration job (Tasks pane). A Details pane reflects the current view and provides summary information about an item selected in the current view.
- ♦ **Tasks Pane:** The Tasks pane of the PlateSpin Migrate Client window contains most essential migration actions. Clicking a task opens the Action window, which you can use to select the migration source, target, and setup method.
- ♦ **Status bar:** At the bottom of the PlateSpin Migrate Client window, the status bar displays the PlateSpin Server that the client is currently connected to, the PlateSpin Migrate Network you are currently working with, the name and role of the current user logged in, and the status of the Automatic Network Discovery feature. See [Section 3.2.5, “Status Bar,” on page 68](#).

3.2.2 Servers View

The Servers view is the main visual interface to your discovered source workloads and targets.

The Servers view consists of two panes that you can customize to suit your needs.

Figure 3-1 PlateSpin Migrate Client’s Servers View



The hierarchical display of items in the Servers view reflects the organization of items on their respective platforms; for example: VMs are shown nested beneath their VM hosts, and PlateSpin Images are beneath their image servers.

In addition, the **Group By** bar enables you to group machines by affiliation to a domain or to a vCenter Server (for VMware ESX server systems). See [“Organizing the Servers View” on page 64](#).

NOTE: The Servers view hierarchy does not reflect advanced VM resource management hierarchies and structures, such as membership in resource pools or affiliation with ESX Distributed Resource Scheduler (DRS) clusters. You can view such information in an item’s properties. See [“Viewing the Properties of Source Workloads and Targets” on page 65](#).

- ♦ [“Distinguishing Target Machines for Semi-Automated \(X2P\) Workflow” on page 63](#)
- ♦ [“Organizing the Servers View” on page 64](#)
- ♦ [“Viewing the Properties of Source Workloads and Targets” on page 65](#)
- ♦ [“List of Machine-Specific Icons in the Servers View” on page 66](#)





Distinguishing Target Machines for Semi-Automated (X2P) Workflow

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

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

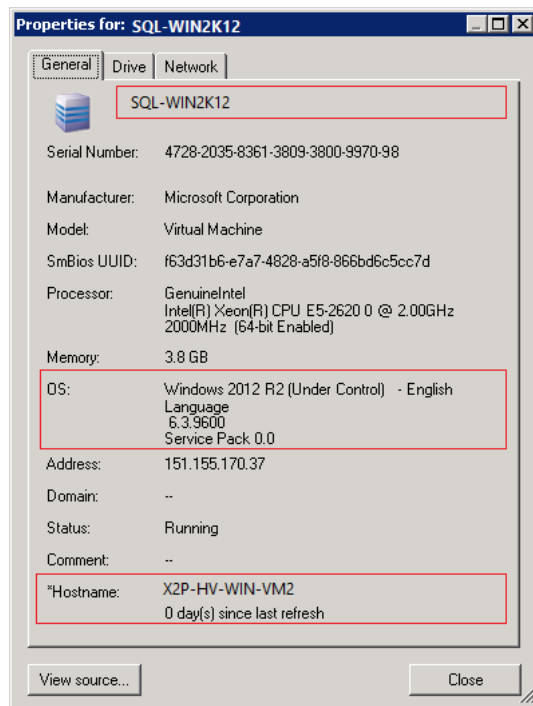
[Figure 3-2](#) provides an example of X2P host names for target workloads with and without an operating system present. Workloads X2P-HV-LX-VM3 and X2P-HV-WIN-VM1 do not have an underlying operating system. The LRD information is displayed as the operating system.

Figure 3-2 X2P Host Name and Operating System Displayed in the Hosts List

 X2P-HV-LX-VM3	Linux (Under Control (64-bit Enabled))
 X2P-HV-LX-VM4 (pgsql.example.com)	SLES 11.4 (x64) (Under Control (64-bit Enabled))
 X2P-HV-WIN-VM1	Linux (Under Control (64-bit Enabled))
 X2P-HV-WIN-VM2 (SQL-WIN2K12)	Windows 2012 R2 (Under Control)

In the Properties dialog for the target workload, the displayed host name is the operating system host name. The registered host name displays at the bottom of the General tab as the ***Hostname** value, as shown in [Figure 3-3](#). The OS value displays the **Under Control** annotation.

Figure 3-3 Properties Dialog for an X2P Target Workload




Organizing the Servers View

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

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

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

Figure 3-4 Servers View Options for Sorting Items by Type



Viewing the Properties of Source Workloads and Targets

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

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

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

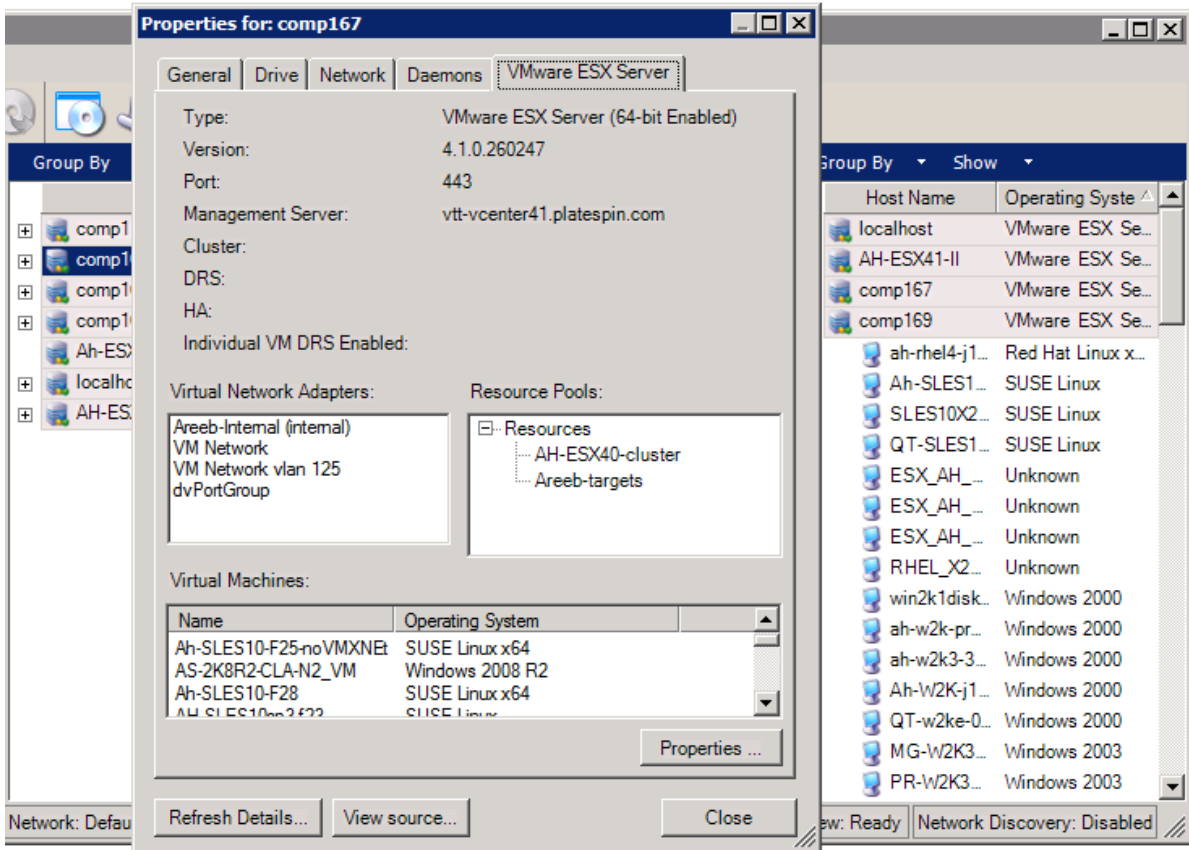
A virtual machine's properties provide information related to the machine's environment on its corresponding virtualization platform, including information about the host, and the amount of allocated memory and processing power.

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

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

The following figure shows the properties of a discovered VMware ESX Server.









Figure 3-5 VMware ESX Server-Specific Information in the System's Properties






List of Machine-Specific Icons in the Servers View

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

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

	Physical machine
	Physical machine in pre-execution environment for offline migration
	Physical machine with workload license
	Virtual machine server
	Virtual machine
	Virtual machine with workload license
	Undiscovered virtual machine
	Virtual machine - Server Sync target


-
-  Virtual machine - Server Sync target with workload license
 -  PlateSpin Image Server
 -  PlateSpin Image
-

3.2.3 Jobs View

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

- ♦ **Jobs:** All jobs submitted for execution.
- ♦ **Saved Jobs:** All saved jobs not yet submitted for execution. See [Section 38.2.1, “Using the Migrate Client,”](#) on page 435.

Figure 3-6 PlateSpin Migrate Client's Jobs View



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

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

3.2.4 Tasks Pane


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

3.2.5 Status Bar

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

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

Figure 3-7 Status Bar of the PlateSpin Migrate Client Window



Double-clicking any of the first three status items opens the PlateSpin Server Settings window. See [Section 3.1, “Connecting to a PlateSpin Migrate Server,” on page 59](#).

Double-clicking the Network Discovery status item turns Automatic Windows Network Discovery on or off. See [Section 19.7, “Discovering Target VMs for Server Sync Jobs,” on page 208](#).

3.2.6 Workload Migration Tasks

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

Task	Description
Copy Workload	Results in a virtual or physical duplicate of a selected physical or virtual workload, except that the new workload is assigned a new network identity. Use this migration task when you intend to keep the source workload operational.
Move Workload	Results in an exact virtual or physical duplicate of a selected physical or virtual workload. Use this migration task when you intend to retire or repurpose the original infrastructure.
Server Sync	Synchronizes a virtual or physical workload with another virtual or physical workload without transferring the entire source volume data over the network.
Capture Image	Creates an image of a physical or virtual workload as a single entity, in PlateSpin Image format.
Deploy Image	Converts a PlateSpin Image into a booted or bootable workload on a physical or virtual machine.

3.3 About the PlateSpin Migrate Web Interface

The PlateSpin Migrate Web Interface provides a web-browser-based management tool to manage automated migrations to target virtual machines on VMware host targets and cloud-based targets. No client installation is required. For information about configuration options for the Web Interface, see [Chapter 7, “Configuring PlateSpin Migrate Web Interface,” on page 125](#).

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

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

- ♦ [Section 3.3.1, “Navigating the Web Interface,” on page 69](#)
- ♦ [Section 3.3.2, “Workloads,” on page 70](#)
- ♦ [Section 3.3.3, “Targets,” on page 74](#)
- ♦ [Section 3.3.4, “Tasks,” on page 74](#)
- ♦ [Section 3.3.5, “Dashboard,” on page 75](#)
- ♦ [Section 3.3.6, “Reports,” on page 75](#)

3.3.1 Navigating the Web Interface

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

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

Navigation Options	Description
Dashboard	Displays the default Dashboard page that provides information about the Migrate licenses, latest tasks, running events, upcoming events, and past events. See Section 3.3.5, “Dashboard,” on page 75 .
Workloads	<p>Displays the Workloads page that lists all the discovered workloads. To add or discover a workload, click Add Workload option on the Dashboard page or the Workloads page. For more information about adding or discovering a workload, see Section 20.6.2, “Workload Discovery in the Migrate Web Interface,” on page 215.</p> <p>You can perform various other tasks such as configuring a workload, preparing a workload for migration, and migrating a workload. See Section 3.3.2, “Workloads,” on page 70.</p>
Targets	<p>Displays the Targets page that lists the already added target containers and lets you add new targets. For more information about adding or discovering a workload, see Section 19.4.2, “Target Discovery in the Web Interface,” on page 202.</p> <p>See Section 3.3.3, “Targets,” on page 74.</p>
Tasks	Displays the Tasks page that lists items requiring user intervention. See Section 3.3.4, “Tasks,” on page 74 .
Reports	Displays the Reports page. See Section 39.1, “Generating Workload and Workload Migration Reports,” on page 441 .

Navigation Options	Description
Settings	<p>Displays the Settings page that lets you configure the following:</p> <ul style="list-style-type: none"> ♦ Workload Tags: See Section 20.9, “Using Tags to Track Logical Associations of Workloads,” on page 218. ♦ Permissions: See Section 7.1, “Managing Security Groups and Workload Permissions,” on page 125. ♦ General Notification Settings: See “Setting Up Event Notifications by Email” on page 102. ♦ Report Notification Settings: See “Setting Up Replication Report Notifications by Email” on page 103. ♦ SMTP: See “Setting up the SMTP Server” on page 101. ♦ Licenses: See “License Activation Using the Web Interface” on page 93.
Downloads	<p>Displays a page that lets you download the following:</p> <ul style="list-style-type: none"> ♦ Migrate Agent: Allows you to download and install the Migrate Agent utility for Windows or Linux. For information about working with the Migrate Agent Utility, see Appendix F, “Migrate Agent Utility,” on page 269. ♦ Migrate Client Setup: Allows you to download and install the PlateSpin Migrate Client. For information about the PlateSpin Migrate Client, see Section 3.2, “About the PlateSpin Migrate Client User Interface,” on page 61. <p>You can also install the PlateSpin Migrate Client using the PlateSpin Migrate Installer. For more information, see Installing the PlateSpin Migrate Client in the PlateSpin Migrate 12.2.1 Installation and Upgrade Guide.</p>
About	Displays information such as the product version, copyright information, license information and also provides links to the Downloads page and the product home page.
Help	Displays the online documentation page.

3.3.2 Workloads

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

- ♦ [“Status for Workloads Managed in Migrate Web Interface”](#) on page 71
- ♦ [“Status for Workloads Managed in Migrate Client”](#) on page 72
- ♦ [“Filtering or Organizing Workloads in the Workloads View”](#) on page 72
- ♦ [“Viewing Details for a Source Workload”](#) on page 73
- ♦ [“Viewing the Command Details for a Source Workload”](#) on page 73
- ♦ [“OS Icons in the Workloads View”](#) on page 74

Status for Workloads Managed in Migrate Web Interface

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

Item	Description
Tasks	Displays a warning icon for a task that might require user attention. For example: if a workload goes offline, then a warning icon displays. Pause over the icon to see more details.
Online	Displays one of the following: <ul style="list-style-type: none">♦ Yes: If the workload is online.♦ No: If the workload is offline.
Workload	Displays the workload name. Click the workload name to configure the workload for migration.
Tag	Displays the tag associated with the workload. For more information about the tags, see Section 7.2, “Managing Workload Tags,” on page 127 and Section 20.9, “Using Tags to Track Logical Associations of Workloads,” on page 218 .
Schedule	Displays the status of the schedule if you have configured a schedule for the workload migration. For example: if the schedule is configured, it displays Active after you have prepared the workload for migration until the end of the migration cycle, unless you pause the schedule. If you click Pause Schedule , then the Paused status displays. If you click Resume Schedule , then Active displays again.
Migration Status	Displays the current status of the workload. For example: <ul style="list-style-type: none">♦ Adding Workload: The process of adding or discovering a workload is in progress.♦ Not Configured: The workload has been discovered but is not yet configured.♦ Migration Configured: The workload has been configured for migration.♦ Preparing Migration: The source workload is being prepared for migration and the target workload is being prepared for replication.♦ Running First Replication: The workload is being replicated for the first time. Click the Migration Status link to view information about the related events.
Last Replication	Displays the date when the workload was last replicated.
Next Replication	Displays the date when the workload is scheduled for the next replication.
Last Test Cutover	Displays the date when the target workload was last tested.

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

Status for Workloads Managed in Migrate Client

The Workloads page displays read-only status for migration jobs managed in the Migrate Client. Event messages for these conditions are also reported to PlateSpin Transformation Manager, where the related job is tracked as an external workload migration.

After you discover details for a workload in the Migrate Client, the Web Interface displays it in the Workloads list with a status of Unconfigured. At this point, you can proceed to manage the workload migration in either the Migrate Client or the Web Interface, depending on your migration goals. See [Section 3.4, “Migration Operations Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface,” on page 75](#).

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

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

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

Filtering or Organizing Workloads in the Workloads View

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

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

For information about how to create tags and associate them with workloads, see [Section 20.9, “Using Tags to Track Logical Associations of Workloads,” on page 218](#).

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

To filter the listing of workloads:

- 1 In the **Workload Status** menu, select one of the following:
 - ♦ All workloads
 - ♦ Replicated
 - ♦ Scheduled
 - ♦ Running Cutover
 - ♦ Running Test Cutover
 - ♦ Running Replication
 - ♦ Failed Workloads
 - ♦ Running Workloads
 - ♦ Not configured
 - ♦ Ready for Replication
 - ♦ Cut Over
- 2 (Optional) In the **Tag** menu, select the tag that is associated with the workloads you want to list, or select **All**.

For information about how to create tags and associate them with workloads, see [Section 20.9, “Using Tags to Track Logical Associations of Workloads,”](#) on page 218.

Viewing Details for a Source Workload

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

- 1 On the Workloads page, click the **Name** link of the workload of interest.
- 2 View the Discovery Details or Migration Details, depending on where it is in the migration lifecycle.
- 3 (Optional) Select the Command Details tab to view information about events for the last command executed on the workload.

Viewing the Command Details for a Source Workload



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

- 1 On the Workloads page, click the **Migration Status** link of the workload of interest.
- 2 On the Command Details page, view information about events for the last command executed on the workload.
- 3 (Optional) If Workload Commands are active for the workload, you can initiate a follow-on action for the migration by clicking the appropriate action.

OS Icons in the Workloads View

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

Table 3-4 Operating System Icons in the Workloads View

	Windows operating systems
	Linux operating systems

3.3.3 Targets

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

- ♦ [Section 2.1.5, “Supported Target Virtualization Platforms,” on page 38](#)
- ♦ [Section 2.1.6, “Supported Target Cloud Platforms,” on page 40](#)

Each platform is identified by the cloud provider or the specific operating system installed on the VMware host server. For more information, see [Chapter 19, “Discovering Target Hosts,” on page 197](#).

3.3.4 Tasks

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

Events are logged whenever some action related to the system or the workload occurs. For example, an event could be the addition of a new workload, the replication of a workload starting or failing, or the detection of the failure of a migrated workload. Some events also email automatic notifications if SMTP is configured. For more information, see [Section 5.4.2, “Notification Service Using Migrate Web Interface,” on page 101](#).

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

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

3.3.5 Dashboard

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

The left pane of the Dashboard page provides a high-level view of the overall state of the PlateSpin Migrate workload inventory, summary of the license information and also lets you add or discover a new workload. For more information about adding or discovering a workload, see [Section 20.6.2, “Workload Discovery in the Migrate Web Interface,” on page 215](#).

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

3.3.6 Reports

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

The following report types are supported:

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

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

Migration Operation	PlateSpin Migrate Client	PlateSpin Migrate Web Interface
Migration to Amazon Cloud		
Physical to Amazon Cloud	✓	x
Virtual to Amazon Cloud	✓	x
Image to Amazon Cloud	x	x
Migration to Microsoft Azure		
Physical to Microsoft Azure	x	✓
Virtual to Microsoft Azure	x	✓

Migration Operation	PlateSpin Migrate Client	PlateSpin Migrate Web Interface
Image to Microsoft Azure	x	x
Migration to VMware vCloud Director		
Physical to VMware vCloud Director	x	✓
Virtual to VMware vCloud Director	x	✓
Image to VMware vCloud Director	x	x
Migration to VMware Hosts		
Physical to VMware (P2V)	✓	✓
Virtual to VMware (V2V)	✓	✓
Image to VMware (I2V)	✓	x
Migration to Other Virtualization Hosts (Microsoft Hyper-V, KVM, Citrix XenServer, Xen)		
Physical to Virtual (P2V)	✓	x
Virtual to Virtual (V2V)	✓	x
Image to Virtual (I2V)	✓	x
Migration to Physical Hosts		
Physical to Physical (P2P)	✓	x
Virtual to Physical (V2P)	✓	x
Image to Physical (I2P)	✓	x
Migration to PlateSpin Image Server		
Physical to Image (P2I)	✓	x
Virtual to Image (V2I)	✓	x

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

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

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

Tasks	PlateSpin Migrate Client	PlateSpin Migrate web Interface
Monitor workload migration workflow	x	✓
Discover Windows standalone workloads	✓	✓
Discover Windows cluster workloads	✓	✓

Tasks	PlateSpin Migrate Client	PlateSpin Migrate web Interface
Discover Linux standalone workloads	✓	✓
Discover Linux cluster workloads	✗	✗
Discover target VMware hosts	✓	✓
Discover target non-VMware hosts	✓	✗
Migrate to physical machines	✓	✗
Migrate to VMware hosts	✓	✓
Migrate to non-VMware hosts	✓	✗
Migrate to Azure Cloud	✗	✓
Migrate to Amazon Web Services	✓	✗
Migrate to VMware vCloud Director	✗	✓
Migrate to image	✓	✗
Migrate Windows workloads with block-based transfer	✗	✓
Migrate Linux workloads with block-based transfer	✓	✓
Migrate Windows workloads with file-based transfer	✓	✓
Migrate Linux workloads with file-based transfer	✓	✗
Migrate Windows clusters with block-based transfer	✓	✓
Migrate workloads using live transfer	✓	✓
Migrate workloads using offline transfer (migrations to physical)	✓	✗
Schedule incremental replication	✗	✓
Migrate staged workloads using imaging	✓	✗
Support post migration scripts	✓	✗
Add new disks during migration	✓	✗
Change disk volume mapping for target workload	✓	✗
Migrate a VM to a vCenter folder	✗	✓
Move a VM to a resource pool	✓	✓
Set compression level	✓	✓
Throttle bandwidth	✓	✓
Set encryption for data transfer	✓	✓
Create tags	✗	✓
View workload migration report	✓	✓

Tasks	PlateSpin Migrate Client	PlateSpin Migrate web Interface
View workload migration status reports	x	✓
Add or remove licenses	✓	✓
Check licenses status	✓	✓
Use security groups	x	✓
Set global defaults for source service	✓	x
Set global defaults for target service	✓	x
Set global defaults for migration job values	✓	x

3.6 Other PlateSpin Server Management Tools

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

- [Section 3.6.1, “PlateSpin Configuration,” on page 78](#)
- [Section 3.6.2, “PlateSpin Migrate Client Command Line Interface,” on page 79](#)
- [Section 3.6.3, “PlateSpin Analyzer,” on page 79](#)
- [Section 3.6.4, “Migrate Agent Utility,” on page 79](#)
- [Section 3.6.5, “PlateSpin ISO,” on page 79](#)

3.6.1 PlateSpin Configuration

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

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

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

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

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

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

- 2 Locate the required server parameter and change its value.
- 3 Save your settings and exit the page.

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

For information about changing the adapter type used during the Target Take Control process of workload migration to a target VM on a Hyper-V Host, see [Section 5.11, “Specifying the Network Adapter Type to Use for Migrations to Hyper-V during Target Take-Control,” on page 109](#).

For information about increasing the upload size limit on post-migration actions, see [Section 5.14, “Increasing the Upload Size Limit for Post-Migration Actions,” on page 114](#).

For information about optimizing data transfer over WAN connections, see [Section 5.14, “Increasing the Upload Size Limit for Post-Migration Actions,”](#) on page 114.

3.6.2 PlateSpin Migrate Client Command Line Interface

The PlateSpin Migrate Client installation includes a command line interface (CLI) tool to help you perform common migrations tasks. Conversion jobs using `.ini` files is supported onto VMware and Hyper-V targets only. See [Appendix I, “Using the PlateSpin Migrate Client Command Line Interface,”](#) on page 461

3.6.3 PlateSpin Analyzer

PlateSpin Migrate Client provides the PlateSpin Analyzer to determine whether discovered Windows machines are suitable for migration jobs. Before you begin any large-scale migration projects, you should identify potential migration problems and correct them beforehand. See [Section 21.5, “Analyzing Suitability of Discovered Windows Workloads For Conversion to Physical Machines,”](#) on page 229.

3.6.4 Migrate Agent Utility

The Migrate Agent utility is a command line utility that you can use to install, upgrade, query, or uninstall the block-based transfer drivers. The utility also enables you to register source workloads with PlateSpin Migrate servers and send details about the workloads to the server via HTTPS (TCP/443). Registration allows you to add workloads that cannot be discovered, such as for Migrate Servers in Microsoft Azure when no VPN is configured between the Migrate server and the source workloads.

Although a reboot is always required when you install, uninstall, or upgrade drivers, the Migrate Agent utility allows you to better control when the action occurs and therefore, when the server reboots. For example, you can use the Migrate Agent utility to install the drivers during scheduled down time, instead of during the first replication. See [Appendix F, “Migrate Agent Utility,”](#) on page 269.

3.6.5 PlateSpin ISO

The PlateSpin ISO file enables you to register target physical machines and target virtual machines with PlateSpin Migrate servers and send details about them to the server via HTTPS (TCP/443). Registration allows you to add target machines that cannot be discovered because they have no operating system installed. See [Appendix G, “PlateSpin ISO Image,”](#) on page 279.

4 Configuring PlateSpin Users and Access

Users have privileges to perform tasks in PlateSpin Migrate based on their assigned PlateSpin user roles: Administrator, Power User, and Operator. In your VMware environment, you can configure PlateSpin user roles to support multitenancy: Virtual Machine Manager, Infrastructure Manager, and User. This section explains the various PlateSpin user roles, role-based privileges, and how to assign users to the roles.

- ♦ [Section 4.1, “Configuring User Authorization and Authentication,” on page 81](#)
- ♦ [Section 4.2, “Configuring PlateSpin Migrate Multitenancy on VMware,” on page 84](#)
- ♦ [Section 4.3, “Configuring Permissions for Workload Access in PlateSpin Migrate Web Interface,” on page 90](#)

4.1 Configuring User Authorization and Authentication

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

NOTE: If you have installed a PlateSpin Migrate Server localized for one language and a PlateSpin Migrate Client localized for a different language, do not use authorization credentials that include any language-specific characters. Using such characters in the login credentials causes miscommunication between the client and the server: the credentials are rejected as invalid.

PlateSpin Migrate’s user auditing functionality is provided through the capability to log user actions. See [Section 6.6, “Managing Migrate Client User Activity Log,” on page 123](#).

- ♦ [Section 4.1.1, “PlateSpin Migrate Roles,” on page 81](#)
- ♦ [Section 4.1.2, “Assigning PlateSpin Migrate Roles to Windows Users,” on page 83](#)

4.1.1 PlateSpin Migrate Roles

A PlateSpin Migrate role is a collection of PlateSpin Migrate privileges that entitle a particular user to perform specific actions. During installation, the PlateSpin Migrate installation program creates three local Windows groups on the PlateSpin Server host:

- ♦ PlateSpin Migrate Administrators
- ♦ PlateSpin Migrate Power Users
- ♦ PlateSpin Migrate Operators

These groups map directly to the three PlateSpin Migrate roles that control user authorization and authentication:

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

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

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

- ♦ For information on PlateSpin Migrate Roles and permission details when you use PlateSpin Migrate Client to perform the workload migration, see [Table 4-1 on page 82](#).
- ♦ For information on PlateSpin Migrate Roles and permission details when you use PlateSpin Migrate Web Interface to perform the workload migration, see [Table 4-2 on page 83](#).

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

Role Details	Administrators	Power Users	Operators
Licensing: Add, delete licenses; transfer workload licenses	Yes	No	No
Machines: Discover, undiscover	Yes	Yes	No
Machines: Delete virtual machines	Yes	Yes	No
Machines: View, refresh, export	Yes	Yes	Yes
Machines: Import	Yes	Yes	No
Machines: Export	Yes	Yes	Yes
PlateSpin Migrate Networks: Add, delete	Yes	No	No
Jobs: Create new job	Yes	Yes	No
Jobs: View, abort, change start time	Yes	Yes	Yes
Imaging: View, start synchronization in existing contracts	Yes	Yes	Yes
Imaging: Consolidate increments, apply increments to base, delete increments, install/delete image servers	Yes	Yes	No
Block-Based Transfer Components: Install, upgrade, remove	Yes	Yes	No

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

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

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

4.1.2 Assigning PlateSpin Migrate Roles to Windows Users

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

4.2 Configuring PlateSpin Migrate Multitenancy on VMware

PlateSpin Migrate includes unique user roles (and a tool for creating them in a VMware data center) that make it possible for non-administrative VMware users (or “enabled users”) to perform Migrate lifecycle operations in the VMware environment. These roles makes it possible for you, as a service provider, to segment your VMware cluster to allow multitenancy: where multiple Migrate containers are instantiated in your data center to accommodate Migrate customers or “tenants” who want to keep their data and evidence of their existence separate from and inaccessible to other customers who also use your data center.

This section includes the following information:

- ♦ [Section 4.2.1, “Defining VMware Roles for Multitenancy,” on page 84](#)
- ♦ [Section 4.2.2, “Assigning Roles In vCenter,” on page 87](#)

4.2.1 Defining VMware Roles for Multitenancy

PlateSpin Migrate requires certain privileges to access and perform tasks in the VMware Infrastructure (that is, VMware “containers”), making the Migrate workflow and functionality possible in that environment. The `PlateSpinRole.xml` file defines the minimum required privileges and aggregates them respectively into three VMware custom roles:

- ♦ PlateSpin Virtual Machine Manager
- ♦ PlateSpin Infrastructure Manager
- ♦ PlateSpin User

This file is included in the PlateSpin Migrate Server installation. An accompanying executable, `PlateSpin.VMwareRoleTool.exe`, accesses the file to enable the creation of these custom PlateSpin roles in a target vCenter environment.

By default, the role definition file (`PlateSpinRole.xml`) and the role definition tool (`PlateSpin.VMwareRoleTool.exe`) are located in the `Install-Directory\PlateSpin Migrate Server\bin\VMwareRolesTool` folder.

This section includes the following information:

- ♦ [“Basic Command Line Syntax” on page 84](#)
- ♦ [“Additional Command Line Parameters and Flags” on page 85](#)
- ♦ [“Tool Usage Example” on page 85](#)
- ♦ [“\(Option\) Manually Defining the PlateSpin Roles in vCenter” on page 86](#)
- ♦ [“Using vCenter to View Privileges for PlateSpin Custom Roles” on page 86](#)

Basic Command Line Syntax

From the location where the role tool is installed, run the tool from the command line, using this basic syntax:

```
PlateSpin.VMwareRoleTool.exe /host=[hostname/IP] /user=[user name] /role=[the role definition file name and location] /create
```

Additional Command Line Parameters and Flags

Apply the following parameters as needed when you use `PlateSpin.VMwareRoleTool.exe` to create or update roles in vCenter:

<code>/create</code>	(mandatory) Creates the roles defined by the <code>/role</code> parameter
<code>/get_all_privileges</code>	Display all server-defined privileges
<code>/get_compatible_roles</code>	Display all roles that are compatible to the role defined by <code>/role</code>
<code>/check_role=[role name]</code>	Check the given role for compatibility with the role defined by <code>/role</code>

Optional Flags

<code>/interactive</code>	Run the tool with interactive options that allow you to choose to create individual roles, check role compatibility, or list all compatible roles. For information about using the tool in interactive mode, see "Diagnosing VMware Role Based Permission Issues in PlateSpin Products" (https://www.novell.com/support/kb/doc.php?id=7018547).
<code>/password=[password]</code>	Provide the VMware password (bypasses the password prompt)
<code>/verbose</code>	Display detailed information

Tool Usage Example

Usage: `PlateSpin.VMwareRoleTool.exe /host=houston_sales /user=pedrom /role=PlateSpinRole.xml /create`

Resulting Actions:

1. The role definition tool runs on the `houston_sales` vCenter server, which has an administrator with the user name `pedrom`.
2. In the absence of the `/password` parameter, the tool prompts for the user password, which you enter.
3. The tool accesses the role definition file, `PlateSpinRole.xml`, which is located in the same directory as the tool executable (there was no need to further define its path).
4. The tool locates the definition file and is instructed (`/create`) to create the roles defined in the contents of that file in the vCenter environment.
5. The tool accesses the definition file and creates the new roles (including the appropriate minimum privileges for defined, limited access) inside vCenter.

The new custom roles are to be [assigned to users later in vCenter](#).

For information about using the tool, see ["Diagnosing VMware Role Based Permission Issues in PlateSpin Products"](https://www.novell.com/support/kb/doc.php?id=7018547). (<https://www.novell.com/support/kb/doc.php?id=7018547>).

(Option) Manually Defining the PlateSpin Roles in vCenter

You use the vCenter client to manually create and assign the PlateSpin custom roles. This requires creating the roles with the enumerated privileges as defined in `PlateSpinRole.xml`. When you create manually, there is no restriction on the name of the role. The only restriction is that the role names you create as equivalents to those in the definition file have all of the appropriate minimum privileges from the definition file.

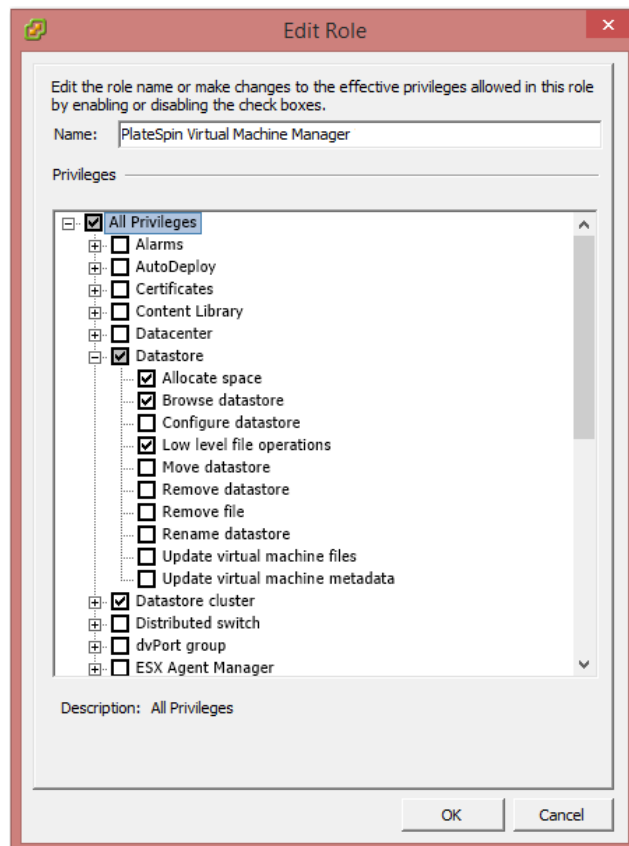
For more information about how to create custom roles in vCenter, see [Managing VMware VirtualCenter Roles and Permissions](http://www.vmware.com/pdf/vi3_vc_roles.pdf) (http://www.vmware.com/pdf/vi3_vc_roles.pdf) in the VMware Technical Resource Center.

Using vCenter to View Privileges for PlateSpin Custom Roles

You use the vCenter client to view the minimal privileges set for the PlateSpin custom roles.

- 1 In vCenter, select a custom role:
 - ♦ PlateSpin Virtual Machine Manager
 - ♦ PlateSpin Infrastructure Manager
 - ♦ PlateSpin User
- 2 Click **Edit** to view the privileges settings in the Edit Role dialog.

For example, the following figure shows some of the privileges set for the PlateSpin Virtual Machine Manager role.



4.2.2 Assigning Roles In vCenter

As you set up a multitenancy environment, you need to provision a single Migrate server per customer or “tenant.” You assign this Migrate server an enabled user with special Migrate VMware roles. This enabled user creates the Migrate container. As service provider, you maintain this user’s credentials and do not disclose them to your tenant customer.

The following table lists the roles you need to define for the enabled user. It also includes more information about the purpose of the role:

vCenter Container for Role Assignment	Role Assignment Specifics	Propagate Instructions	More Information
Root of vCenter inventory tree.	Assign the enabled user the <i>PlateSpin Infrastructure Manager</i> (or equivalent) role.	For security reasons, define the permission as non-propagating.	This role is needed to monitor tasks being performed by the Migrate software and to end any stale VMware sessions.
All data center objects where the enabled user needs access	Assign the enabled user the <i>PlateSpin Infrastructure Manager</i> (or equivalent) role.	For security reasons, define the permission as non-propagating.	This role is needed to allow access to the data center’s datastores for file upload/download. Define the permission as non-propagating.
Each cluster to be added to Migrate as a container, and each host contained in the cluster	Assign the enabled user the <i>PlateSpin Infrastructure Manager</i> (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	To assign to a host, propagate the permission from the cluster object or create an additional permission on each cluster host. If the role is assigned on the cluster object and is propagated, no further changes are necessary when you add a new host to the cluster. However, propagating this permission has security implications.
Each Resource Pool where the enabled user needs access.	Assign the enabled user the <i>PlateSpin Virtual Machine Manager</i> (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	Although you can assign access to any number of Resource Pools in any location in the tree, you must assign the enabled user this role on at least one Resource Pool.
Each VM folder where the enabled user needs access	Assign the enabled user the <i>PlateSpin Virtual Machine Manager</i> (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	Although you can assign access to any number of VM Folders in any location in the tree, you must assign the enabled user this role on at least one folder.


vCenter Container for Role Assignment	Role Assignment Specifics	Propagate Instructions	More Information
<p>Each Network where the enabled user needs access.</p> <p>Distributed Virtual Networks with a dvSwitch and a dvPortgroup</p>	<p>Assign the enabled user the <i>PlateSpin Virtual Machine Manager</i> (or equivalent) role.</p>	<p>Propagation is at the discretion of the VMware administrator.</p>	<p>Although you can assign access to any number of networks in any location in the tree, you must assign the enabled user this role on at least one folder.</p> <ul style="list-style-type: none"> ♦ To assign the correct role to the dvSwitch, propagate the role on the data center (resulting in an additional object receiving the role) or place the dvSwitch in a folder and assign the role on that folder. ♦ For a standard portgroup to be listed as an available network in the Migrate UI, create a definition for it on every host in the cluster.
<p>Each Datastore and Datastore Cluster where the enabled user needs access</p>	<p>Assign the enabled user the <i>PlateSpin Virtual Machine Manager</i> (or equivalent) role.</p>	<p>Propagation is at the discretion of the VMware administrator.</p>	<p>The enabled user must have been assigned this role on at least one Datastore or Datastore Cluster.</p> <p>For Datastore Clusters, the permission must be propagated to the contained datastores. Not providing access to an individual member of the cluster causes both prepare and full replications to fail</p>

The following table shows the role you can assign to the customer or tenant user.

vCenter Container for Role Assignment	Role Assignment Specifics	Propagate Instructions	More Information
Each resource pool(s) and folder(s) where the customer's VMs will be created.	Assign the tenant user the <i>PlateSpin User</i> (or equivalent) role.	Propagation is at the discretion of the VMware administrator.	<p>This tenant is a member of the PlateSpin Administrators group on the PlateSpin Migrate server and is also on the vCenter server.</p> <p>If the tenant will be granted the ability to change the resources used by the VM (that is, networks, ISO images, and so forth), grant this user the necessary permissions on those resources. For example, if you want to allow the customer to change the network where their VM is attached, this user should be assigned the Read-only role (or better) on all of the networks being made accessible to the customer.</p>

The figure below illustrates a Virtual Infrastructure in the vCenter console. The objects labeled in blue are assigned the Infrastructure Manager role. The objects labeled in green are assigned the Virtual Machine Manager role. The tree does not show VM folders, Networks and Datastores. Those objects are assigned the *PlateSpin Virtual Machine Manager* role.

Figure 4-1 Roles assigned in vCenter



Security Implications of Assigning VMware Roles

PlateSpin software uses an enabled user only to perform protection lifecycle operations. From your perspective as a service provider, an end user never has access to the enabled user's credentials and is unable to access the same set of VMware resources. In an environment where multiple Migrate servers are configured to use the same vCenter environment, Migrate prevents possibilities for cross-client access. The major security implications include:

- With the *PlateSpin Infrastructure Manager* role assigned to the vCenter object, every enabled user can see (but not affect) the tasks performed by every other user.

- ♦ Because there is no way to set permissions on datastore folders/subfolders, all enabled users with permissions on a datastore have access to all other enabled users' disks stored on that datastore.
- ♦ With the *PlateSpin Infrastructure Manager* role assigned to the cluster object, every enabled user is able to turn off/on HA or DRS on the entire cluster
- ♦ With the *PlateSpin User* role assigned at the storage cluster object, every enabled user is able to turn off/on SDRS for the entire cluster
- ♦ Setting the *PlateSpin Infrastructure Manager Role* on the DRS Cluster object and propagating this role allows the enabled user to see all VMs placed in the default resource pool and/or default VM folder. Also, propagation requires the administrator to explicitly set the enabled user to have a "no-access" role on every resource pool/VM folder that he or she should not have access to.
- ♦ Setting the *PlateSpin Infrastructure Manager Role* on the vCenter object allows the enabled user to end sessions of any other user connected to the vCenter.

NOTE: Remember, in these scenarios, different enabled users are actually different instances of the PlateSpin software.

4.3 Configuring Permissions for Workload Access in PlateSpin Migrate Web Interface

PlateSpin Migrate Web Interface enables you to set permissions for workload migration management. You configure security groups and assign users and workloads to it. Only members of the security group are permitted to manage the member workloads in that group. See [Section 7.1, "Managing Security Groups and Workload Permissions," on page 125](#).

5 Configuring PlateSpin Migrate Server

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

- [Section 5.1, “PlateSpin Migrate Product Licensing,” on page 91](#)
- [Section 5.2, “Configuring Language Settings for International Versions,” on page 99](#)
- [Section 5.3, “Enforcing FIPS Compliance for FIPS-Enabled Source Workloads,” on page 100](#)
- [Section 5.4, “Configuring the Notification Service,” on page 100](#)
- [Section 5.5, “Configuring Notifications for Events and Migrations,” on page 103](#)
- [Section 5.6, “Enabling Event Messaging for PlateSpin Migration Factory,” on page 105](#)
- [Section 5.7, “Configuring Alternate IP Addresses for PlateSpin Server,” on page 106](#)
- [Section 5.8, “Setting Reboot Method for the Configuration Service,” on page 106](#)
- [Section 5.9, “Configuring the Contact Direction for the Replication Port,” on page 107](#)
- [Section 5.10, “Configuring Behavior for Installing Network Drivers on Target Windows Workloads,” on page 107](#)
- [Section 5.11, “Specifying the Network Adapter Type to Use for Migrations to Hyper-V during Target Take-Control,” on page 109](#)
- [Section 5.12, “Configuring Applications Known to Cause Boot Failure on Windows Target,” on page 110](#)
- [Section 5.13, “Optimizing Data Transfer over WAN Connections,” on page 111](#)
- [Section 5.14, “Increasing the Upload Size Limit for Post-Migration Actions,” on page 114](#)
- [Section 5.15, “Other Use Cases for Custom PlateSpin Server Settings \(Advanced\),” on page 114](#)

5.1 PlateSpin Migrate Product Licensing

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

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

- [Section 5.1.1, “Activating Your Product License,” on page 91](#)
- [Section 5.1.2, “How Migration Licensing Works,” on page 95](#)
- [Section 5.1.3, “Managing License Keys for Workload Migrations,” on page 95](#)
- [Section 5.1.4, “Managing Workload Designations,” on page 98](#)

5.1.1 Activating Your Product License

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

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

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

- ♦ “[License Splitting](#)” on page 92
- ♦ “[License Activation Using Migrate Client](#)” on page 92
- ♦ “[License Activation Using the Web Interface](#)” on page 93

License Splitting

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


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

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

License Activation Using Migrate Client

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

Figure 5-1 License Activation Wizard



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

- ♦ “[Online License Activation](#)” on page 93
- ♦ “[Offline License Activation](#)” on page 93

Online License Activation

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

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

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


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

Offline License Activation

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

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

The Activate License dialog box is displayed:



- 2 Save your hardware ID for use in the next steps.
- 3 Use a computer with Internet access to obtain a license key through the [Web-based license activation utility](http://www.platespin.com/productactivation/ActivateOrder.aspx) (<http://www.platespin.com/productactivation/ActivateOrder.aspx>).
To obtain a license key, you must have a Customer Center account. If you are an existing PlateSpin customer and you don't have a Customer Center account, you must first create one. (See [Create Account](#).) Use your existing PlateSpin username (a valid e-mail address registered with PlateSpin) as input for your Customer Center account username.
- 4 Save your new license key in a location accessible to your PlateSpin Migrate Client.
- 5 In the License Wizard, type the full path to, or browse to and select, the PlateSpin Migrate license file, then click **Next**.

The product is activated based on the selected license.

License Activation Using the Web Interface

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

Figure 5-2 License Activation Using Migrate Web Interface

Dashboard Workloads Targets Tasks Reports Settings Downloads About Help

Workload Tags

Permissions

General Notification Settings

Report Notification Settings

SMTP

Licenses

License Activation Activate

☒ Online Activation (Requires Internet Access)

Email Address:

Activation Code:

☐ Offline Activation (Requires license file)

Your hardware ID is: 5ld7U//meu88FpfJtjngf/vkQUk=

To create a license key file, please access: <http://www.platespin.com/productactivation/ActivateOrder.aspx>

File: No file selected.

Thursday, July 20, 2017 1:15 PM - Pacific Daylight Time

- ◆ “Online License Activation” on page 94
- ◆ “Offline License Activation” on page 94

Online License Activation

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

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

To set up online license activation:

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

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

Offline License Activation

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

- 1 In the PlateSpin Migrate Web Interface, click **Settings > Licenses**, then click **Add license**.
- 2 Click **Offline Activation** and copy the hardware ID displayed in the interface.
- 3 Use a web browser on a computer that has Internet access to navigate to the [PlateSpin Product Activation website](http://www.platespin.com/productactivation/ActivateOrder.aspx) (<http://www.platespin.com/productactivation/ActivateOrder.aspx>). Log in with your Customer Center user name and password.
- 4 Open the [PlateSpin Activate Order page](#) to generate a license key file. You need the following information:
 - ◆ activation code that you received
 - ◆ email address that you provided when you placed your order
 - ◆ hardware ID that you copied in [Step 2](#)

- 5 Save the generated license key file, transfer it to the product host that does not have Internet connectivity, and use it to activate the product.
- 6 In the PlateSpin Migrate Web Interface on the License Activation page, browse to the location of the license key file, then click **Activate**.

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

5.1.2 How Migration Licensing Works

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

The following is a summary of workload assignment behavior by portability task.

Table 5-1 PlateSpin Migrate Workload License Assignment by Migration Type

Task	Workload Assignment Behavior
Copy Workload	A workload license remains with the source.
Move Workload	A workload license is transferred from the source to the target.
Server Sync	Not applicable
Capture Image	A workload license is assigned to the source and remains with it.
Deploy Image	Not applicable

5.1.3 Managing License Keys for Workload Migrations

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

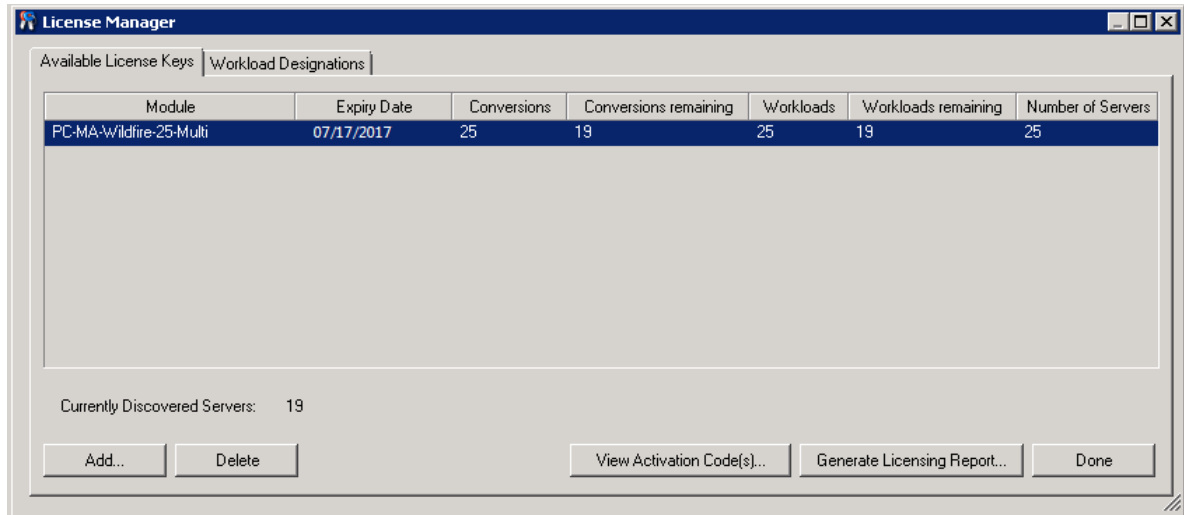
- ♦ [“License Key Management with Migrate Client” on page 96](#)
- ♦ [“License Key Management Using Migrate Web Interface” on page 97](#)

License Key Management with Migrate Client

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

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

Figure 5-3 Available License Keys



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

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

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

Table 5-2 License Manager Command Buttons

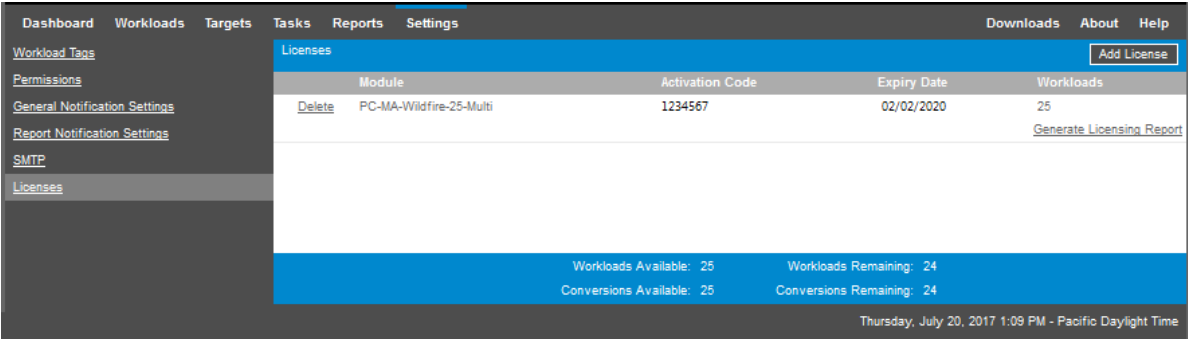
Command	Description
Add	Adds licenses.
Delete	Deletes expired licenses.
View Activation Code(s)	Select a license and click this button to see the activation code and the date it was activated.
Generate Licensing Report	Creates a *.psl file that is used by Technical Support to troubleshoot licensing issues.

License Key Management Using Migrate Web Interface

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

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

Figure 5-4 Available License Keys



The Licenses tab displays the license name (**Module**) along with its activation code, expiry date, and the number entitlements for workload migration.

You can see the number of workload licenses available and remaining at the bottom of the window.

Use the options for related license management tasks:

Table 5-3 Licenses Tab Options


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

5.1.4 Managing Workload Designations

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

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

Figure 5-5 License Manager Workload Designations



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

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

To reset workload licensing:

- 1 On the License Manager's **Workload Designations** tab, select the required workload and click **Transfer Selected Workload**.
The Transfer License dialog box is displayed.
- 2 Use the displayed **Workload Transfer Request** string to obtain a workload transfer code from the [License Entitlement Web portal \(http://www.platespin.com/entitlementmgr/\)](http://www.platespin.com/entitlementmgr/). Log in with credentials associated with your purchase order.
You must have a Customer Center account. If you are an existing PlateSpin customer and you don't have a Customer Center account, you must first create one. (See [Create Account](#).) Use your existing PlateSpin username (a valid e-mail address registered with PlateSpin) as input for your Customer Center account username.
- 3 Return to the License Manager and specify the newly obtained transfer code. Click **Next**.
PlateSpin Migrate resets the selected workload.

5.2 Configuring Language Settings for International Versions

In addition to English, PlateSpin Migrate provides National Language Support (NLS) for the following international languages:

- ♦ Chinese Simplified
- ♦ Chinese Traditional
- ♦ French
- ♦ German
- ♦ Japanese

To manage your PlateSpin Server in one of these supported languages, configure the language code for the operating system on the PlateSpin Migrate Server host and in your Web browser. If you install PlateSpin Migrate Client on a different host machine, configure the operating system on that machine.

- ♦ [Section 5.2.1, “Setting the Language on the Operating System,” on page 99](#)
- ♦ [Section 5.2.2, “Setting the Language in Your Web Browser,” on page 99](#)

5.2.1 Setting the Language on the Operating System

The language of a small portion of system messages generated by PlateSpin Migrate depends on the operating system interface language selected in your PlateSpin Migrate Server host.

To change the operating system language:

- 1 Log in as Administrator on your PlateSpin Migrate Server host or Migrate Client host.
- 2 Start the Regional and Language Options applet (click **Start > Run**, type `intl.cpl`, and press Enter), then click the **Languages** (Windows Server 2003) or **Keyboards and Languages** (Windows Server 2008 and later) tab, as applicable.
- 3 If it is not already installed, install the required language pack. You might need access to your OS installation media.
- 4 Select the required language as the interface language of the operating system. When you are prompted, log out or restart the system.

5.2.2 Setting the Language in Your Web Browser

To use the PlateSpin Migrate Web Interface in one of the supported international languages, the corresponding language must be added in your web browser and moved to the top of the order of preference:

- 1 Access the Languages setting in your web browser.
- 2 Add the required language and move it up the top of the list.
- 3 Save the settings, then start the client application by connecting to your PlateSpin Migrate Server.

NOTE: (For users of Chinese Traditional and Chinese Simplified languages) Attempting to connect to PlateSpin Migrate with a browser that does not have a specific version of Chinese added might result in web server errors. For correct operation, use your browser's configuration settings to add a specific Chinese language (for example, Chinese [zh-cn] or Chinese [zh-tw]). Do not use the culture-neutral Chinese [zh] language.

5.3 Enforcing FIPS Compliance for FIPS-Enabled Source Workloads

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

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

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

- 2 Locate the `EnforceFIPSCompliance` parameter and click **Edit** to change its value to `true`.
- 3 Click **Save**.

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

- 4 Discover the FIPS enabled source workload.

5.4 Configuring the Notification Service

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

- ♦ [Section 5.4.1, “Notification Service Using Migrate Client,” on page 100](#)
- ♦ [Section 5.4.2, “Notification Service Using Migrate Web Interface,” on page 101](#)

5.4.1 Notification Service Using Migrate Client

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

To configure the SMTP settings for the Notification Service:

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

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

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

You can also configure migration progress notifications on a per-migration basis. See [Section 5.5.1, "Notifications Using the Migrate Client,"](#) on page 104.

5.4.2 Notification Service Using Migrate Web Interface

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

- ♦ ["Setting up the SMTP Server" on page 101](#)
- ♦ ["Setting Up Event Notifications by Email" on page 102](#)
- ♦ ["Setting Up Replication Report Notifications by Email" on page 103](#)

Setting up the SMTP Server

- 1 In the Migrate Web Interface, click **Settings > SMTP**.
- 2 Specify the following:
 - ♦ **SMTP Server Address:** The address of the SMTP server.
 - ♦ **Port:** The port at which the SMTP server is listening. By default, it is 25.

- ♦ **Reply Address:** The address from which you want to send email event and progress notifications.
- ♦ **Username and Password:** Provide valid credentials if your SMTP server requires authentication.

3 Click **Save**.

Setting Up Event Notifications by Email

To set up event notifications:

- 1 Configure an SMTP server for PlateSpin Migrate to use. See [“Setting up the SMTP Server” on page 101](#).
- 2 In the PlateSpin Migrate Web Interface, select **Settings > General Notification Settings**.
- 3 Select the **Enable Notifications** check box.
- 4 Click **Edit Recipients**, specify the required email addresses separated by commas and click **OK**.
- 5 Click **Save**.

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

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

Event Types	Remarks
Log Entry Type: Warning	
IncrementalReplicationMissed	Generates when any of the following applies: <ul style="list-style-type: none"> ♦ A replication is manually paused when a scheduled incremental replication is due. ♦ The system attempts to carry out a scheduled incremental replication when a manually-triggered replication is in progress. ♦ The system determines that the target has insufficient free disk space.
FullReplicationMissed	Similar to IncrementalReplicationMissed event.
WorkloadOfflineDetected	Generated when the system detects that a previously online workload is now offline. Applies to workloads whose migration state is not Paused .
Log Entry Type: Error	
FailoverFailed	
FullReplicationFailed	
IncrementalReplicationFailed	
PrepareFailoverFailed	
Log Entry Type: Information	
FailoverCompleted	

Event Types	Remarks
FullReplicationCompleted	
IncrementalReplicationCompleted	
PrepareFailoverCompleted	
TestFailoverCompleted	Generates upon manually marking a Test Cutover operation a success or a failure.
WorkloadOnlineDetected	Generates when the system detects that a previously offline workload is now online. Applies to workloads whose migration state is not Paused .

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

Setting Up Replication Report Notifications by Email

- 1 Set up an SMTP server for PlateSpin Migrate to use. See [“Setting up the SMTP Server” on page 101](#).
- 2 In the PlateSpin Migrate Web Interface, select **Settings > Report Notification Settings**.
- 3 Select the **Enable Report Notifications** check box.
- 4 In the **Report Recurrence** section, click **Edit** and specify the required recurrence pattern for the reports.
- 5 In the **Recipients** section, click **Edit Recipients** to specify the required email addresses separated by commas and click **OK**.
- 6 (Optional) In the **Migrate Access URL** section, specify a non-default URL for your PlateSpin Server.

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

For information on other types of reports that you can generate and view on demand, see [“Generating Workload and Workload Migration Reports” on page 441](#).

5.5 Configuring Notifications for Events and Migrations

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

- [Section 5.5.1, “Notifications Using the Migrate Client,” on page 104](#)
- [Section 5.5.2, “Notifications Using the Web Interface,” on page 104](#)

5.5.1 Notifications Using the Migrate Client

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

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

You specify SMTP server and email account details globally. You can also specify job-specific email addresses. See [Section 5.4, “Configuring the Notification Service,” on page 100](#).

To set up email notifications:

- 1 In the PlateSpin Migrate Client, configure information for the SMTP server for PlateSpin Migrate to use. See [Section 5.4.1, “Notification Service Using Migrate Client,” on page 100](#).
- 2 Start the migration job. For information about starting a migration job, see [Section 26.2, “Initiating a Migration Job,” on page 292](#).
- 3 In the Job Configuration section of the Migration Job window, click **Alerts** and configure the required options.
 - 3a Select **Receive Event Notifications** to receive notifications for `Completed`, `Recoverable Error`, and `Failed` conditions for migration jobs.
 - 3b Select **Receive Progress Notifications** to receive progress notifications via email. Specify the frequency with which you want to receive notifications for the job.
 - 3c (Optional) In **Send To Addresses**, add or remove job-specific email addresses that will receive notifications.
- 4 Click **OK**.

5.5.2 Notifications Using the Web Interface

To set up a list of recipients for event notifications:

- 1 In the PlateSpin Migrate Web Interface, configure information for the SMTP server for PlateSpin Migrate to use. See [“Setting up the SMTP Server” on page 101](#).
- 2 Select **Settings > General Notification Settings**.
- 3 Select the **Enable Notifications** check box.
- 4 In the **Recipients** section, click **Edit Recipients** to specify the required email addresses separated by commas and click **OK**.
- 5 Click **Save**.

To set up a list of recipients for report notifications:

- 1 In the PlateSpin Migrate Web Interface, set up an SMTP server for PlateSpin Migrate to use. See [“Setting up the SMTP Server” on page 101](#).
- 2 Select **Settings > Report Notification Settings**.
- 3 Select the **Enable Report Notifications** check box.
- 4 In the **Report Recurrence** section, click **Edit** and specify the required recurrence pattern for the reports.
- 5 In the **Recipients** section, click **Edit Recipients** to specify the required email addresses separated by commas and click **OK**.

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

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

- 7 Click **Save**.

For information on other types of reports that you can generate and view on demand, see [“Generating Workload and Workload Migration Reports” on page 441](#).

5.6 Enabling Event Messaging for PlateSpin Migration Factory

PlateSpin Migrate provides an event messaging service based on RabbitMQ for use in the PlateSpin Migration Factory environment. Each PlateSpin Migrate server can publish workload migration state change messages to PlateSpin Migrate Connector instances that subscribe to the service on behalf of PlateSpin Transformation Manager projects. For information about how communications work for PlateSpin Migration Factory, see [“PlateSpin Migration Factory Environment”](#) in the [PlateSpin Transformation Manager User Guide](#).

The RabbitMQ message queues are pre-configured and start automatically when you start the PlateSpin service for a PlateSpin Migrate server. No messages are published unless you open port 61613 on the Migrate server to allow registration by subscribers, and a PlateSpin Migrate Connector subscribes.

NOTE: The messaging function starts, stops, and restarts automatically with its parent PlateSpin Migrate server service. Do not modify the default settings for event messaging.

In PlateSpin Transformation Manager, you configure the PlateSpin Migrate server as a Migration Server resource for a project. The project's assigned PlateSpin Migrate Connector subscribes to RabbitMQ event messaging. After RabbitMQ has an active subscriber and there are workload migration activities to report, RabbitMQ begins publishing event messages and registered subscribers can receive them. Migrate Connector passes messages to Transformation Manager only for workloads in the appropriate project.

To enable event messaging for migration jobs on the Migrate server:

- 1 Set up your PlateSpin Migration Factory environment.
See [“PlateSpin Migration Factory Environment”](#) in the [PlateSpin Transformation Manager User Guide](#).
- 2 As the Administrator user, open TCP port 61613 for incoming STOMP traffic on the Migrate server host.
- 3 (Azure) For a cloud-based Migrate Server in Azure, allow inbound connections for STOMP traffic (TCP port 61613) in the Migrate server's Network Security Group.
- 4 Open TCP port 61613 in your network.
See [Section 2.6.3, “Requirements for Event Messaging,” on page 52](#).
- 5 In PlateSpin Transformation Manager, configure the PlateSpin Migrate server as a Migration Server resource for a transformation project.
The PlateSpin Migrate Connector subscriber component registers automatically with RabbitMQ on the PlateSpin Migrate server.

See “[Managing Migration Server Resources](#)” in the *PlateSpin Transformation Manager User Guide*.

- 6 (PTM Automated Mode) In PlateSpin Transformation Manager, configure one or more workload Transformation Plans to use the Migration Server resource you created, or use Auto-Assign to allow it to be considered from among the pool of Migrate servers that you have similarly configured.
- 7 (PTM Planning Mode) In PlateSpin Transformation Manager, import the workloads that you configure for migrations manually in PlateSpin Migrate. The Migrate Connector scans periodically to match the external migrations to the imported workloads, and status information is tracked for them.
- 8 Begin workload migrations.

Whether the execution is automated or manual, the Migrate server generates event messages for workload migration actions performed on that server. RabbitMQ publishes the messages. Migrate Connector receives the messages and passes them to the appropriate project in Transformation Manager, where they are displayed for tracking progress and reporting status.

5.7 Configuring Alternate IP Addresses for PlateSpin Server

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

To add alternate IP addresses for PlateSpin Server:

- 1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/PlateSpinConfiguration/`
- 2 Search to locate the `AlternateServerAddresses` parameter and add IP addresses for the PlateSpin Server.
- 3 Save your settings and exit the page.

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

5.8 Setting Reboot Method for the Configuration Service

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

To skip reboot optimization for a single Windows workload:

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

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

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

To skip reboot optimization for all Windows workloads:

- 1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/PlateSpinConfiguration/`
- 2 Search for the configuration parameter `ConfigurationServiceValues`, then click **Edit** for the parameter.
- 3 Change the setting `SkipRebootOptimization` from `false` to `true`.
- 4 Click **Save**.
- 5 Run the failed Test Cutover or Cutover again for affected Windows workloads.

5.9 Configuring the Contact Direction for the Replication Port

By default, the target workload contacts the source workload to initiate the replication data transfer. When you use the Migrate Agent on the source workload, the source workload contacts the target workload for data transfers. The direction is controlled at the server level. You must reconfigure the replication port direction on the Migrate Server (`SourceListensForConnection=False`).

To configure the direction of contact for replication traffic:

- 1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/PlateSpinConfiguration/`
- 2 Locate the `SourceListensForConnection` parameter and edit its value as **True** or **False**, depending on your migration environment.
 - ♦ **True:** (Default) The target workload contacts the source workload to initiate replication. The source listens for traffic on the replication port (default TCP/3725). The replication port must be open for inbound traffic on the source workload.
 - ♦ **False:** The source workload contacts the target workload to initiate replication. The target listens for traffic on the replication port (default TCP/3725). The replication port must be open for inbound traffic on the target workload.
- 3 Save your settings and exit the page.

5.10 Configuring Behavior for Installing Network Drivers on Target Windows Workloads

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

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

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

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

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

- ♦ If the target machine is booting from SAN

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

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

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

- ♦ [Section 5.10.1, “Understanding Light Networking Parameters,” on page 108](#)
- ♦ [Section 5.10.2, “Configuring Light Networking Parameters,” on page 109](#)

5.10.1 Understanding Light Networking Parameters

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

EnableLightNetworking

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

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

HostNamesForLightNetworking

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

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

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

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

Valid values for the `HostNamesForLightNetworking` parameter are:

NONE

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

<FQDN>

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

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

5.10.2 Configuring Light Networking Parameters

To configure the light networking parameters:

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

5.11 Specifying the Network Adapter Type to Use for Migrations to Hyper-V during Target Take-Control

During the Target Take-Control process for workload migrations, PlateSpin Migrate selects the adapter type used based on the Workload OS and Target Virtual Machine type. For migrations to Microsoft Hyper-V, you can leave it to Migrate to decide, or specify a preferred network adapter type to use as Synthetic or Legacy.

To specify the preferred network adapter type for Hyper-V targets:

- 1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/PlateSpinConfiguration/`
- 2 Locate the `PreferredHyperVNetworkAdapter` parameter and edit its value as **Synthetic** or **Legacy**, depending on your Hyper-V requirement.
- 3 Save your settings and exit the page.

5.12 Configuring Applications Known to Cause Boot Failure on Windows Target

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

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

- ♦ **ApplicationsKnownForBootFailuresOnTarget:** Lists some common applications such as Symantec, Kaspersky Antivirus, Backup Assist, and Carbon Black that are known to cause boot failure on the target. To edit the list of the applications, see [Section 5.12.1, “Editing the List of Applications Known to Cause Boot Failure on Windows Target,” on page 110.](#)
- ♦ **ApplicationsKnownForBootFailuresOnTargetDefaultValue:** Sets whether or not all the applications on the Windows source that are known to cause boot failure on the target be automatically selected for disabling during the conversion. The default value is **false** indicating that the applications are not selected by default.

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

For information about configuring the start-up mode of Windows services on the target, see [Section 26.16, “Service States on Target Windows Workloads,” on page 306.](#)

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

- 1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/PlateSpinConfiguration/`
- 2 Locate the **ApplicationsKnownForBootFailuresOnTarget** parameter and click **Edit**.
- 3 The **Values** option lists applications known to cause boot failure on target. Based on your requirement, add applications or remove existing applications whose boot services you do not want to disable during the conversion.
- 4 Save your settings and exit the page.

5.13 Optimizing Data Transfer over WAN Connections

You can optimize data transfer performance and fine tune it for WAN connections. You do this by modifying configuration parameters that the system reads from settings you make in a configuration tool residing on your PlateSpin Server host. For the generic procedure, see [“PlateSpin Configuration” on page 78](#).

- ♦ [Section 5.13.1, “Tuning Parameters,” on page 111](#)
- ♦ [Section 5.13.2, “Tuning FileTransferSendReceiveBufferSize,” on page 113](#)

5.13.1 Tuning Parameters

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

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

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

Table 5-4 Default and Optimized File Transfer Configuration Parameters

Parameter	Default Value	Maximum Value
AlwaysUseNonVSSFileTransferForWindows2003	False	
FileTransferCompressionThreadsCount	2	N/A
Controls the number of threads used for packet-level data compression. This setting is ignored if compression is disabled. Because the compression is CPU-bound, this setting might have a performance impact.		
FileTransferBufferThresholdPercentage	10	
Determines the minimum amount of data that must be buffered before creating and sending new network packets.		
FileTransferKeepAliveTimeOutMilliSec	120000	
Specifies how long to wait to start sending keep alive messages if TCP times out.		
FileTransferLongerThan24HoursSupport	True	
FileTransferLowMemoryThresholdInBytes	536870912	
Determines when the server considers itself to be in a low memory state, which causes augmentation of some networking behavior.		
FileTransferMaxBufferSizeForLowMemoryInBytes	5242880	
Specifies the internal buffer size used in a low memory state.		

Parameter	Default Value	Maximum Value
FileTransferMaxBufferSizeInBytes	31457280	
Specifies internal buffer size for holding packet data.		
FileTransferMaxPacketSizeInBytes	1048576	
Determines the largest packets that will be sent.		
FileTransferMinCompressionLimit	0 (disabled)	max 65536 (64 KB)
Specifies the packet-level compression threshold in bytes.		
FileTransferPort	3725	
FileTransferSendReceiveBufferSize	0 (8192 bytes)	max 5242880 (5 MB)
<p>Defines the maximum size (in bytes) of the send and receive buffers for TCP connections in the replication network. The buffer size affects the TCP Receive Window (RWIN) size, which sets the number of bytes that can be sent without TCP acknowledgment. This setting is relevant for both file-based and block-based transfers. Tuning the buffer size based on your network bandwidth and latency improves throughput and reduces CPU processing.</p> <p>When the value is set to zero (off), the default TCP window size is used (8 KB). For custom sizes, specify the size in bytes.</p> <p>Use the following formula to determine the proper value:</p> $((\text{LINK_SPEED in Mbps} / 8) * \text{DELAY in sec}) * 1000 * 1024$ <p>For example, for a 100 Mbps link with 10 ms latency, the proper buffer size would be:</p> $(100/8) * 0.01 * 1000 * 1024 = 128000 \text{ bytes}$ <p>For tuning information, see “Tuning FileTransferSendReceiveBufferSize” on page 113.</p>		
FileTransferSendReceiveBufferSizeLinux	0 (253952 bytes)	
<p>Specifies the TCP/IP Receive Window (RWIN) Size setting for file transfer connections for Linux. It controls the number of bytes sent without TCP acknowledgment, in bytes.</p> <p>When the value is set to zero (off), the TCP/IP window size value for Linux is automatically calculated from the <code>FileTransferSendReceiveBufferSize</code> setting. If both parameters are set to zero (off), the default value is 248 KB. For custom sizes, specify the size in bytes.</p> <p>NOTE: In previous release versions, you were required to set this parameter to 1/2 the desired value, but this is no longer required.</p>		
FileTransferShutDownTimeOutInMinutes	1090	
FileTransferTCPTimeOutMilliSec	30000	
Sets both the TCP Send and TCP Receive Timeout values.		
PostFileTransferActionsRequiredTimeInMinutes	60	

5.13.2 Tuning FileTransferSendReceiveBufferSize

The FileTransferSendReceiveBufferSize parameter defines the maximum size (in bytes) of the send and receive buffers for TCP connections in the replication network. The buffer size affects the TCP Receive Window (RWIN) size, which sets the number of bytes that can be sent without TCP acknowledgment. This setting is relevant for both file-based and block-based transfers. Tuning the buffer size based on your network bandwidth and latency improves throughput and reduces CPU processing.

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

To calculate the optimum buffer size:

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

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

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

1b Enter the following at a command prompt:

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

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

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

The packet needs to be fragmented.

1d Note the latency in seconds.

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

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

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

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

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

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

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

3a Determine the MSS:

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

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

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

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

3b Calculate the optimal buffer size:

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

To continue the example:

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

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

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

5.14 Increasing the Upload Size Limit for Post-Migration Actions

PlateSpin Migrate enables you to create custom scripts for post-migration actions and upload them to the PlateSpin Library. You can then associate them with certain migration jobs you configure in the PlateSpin Migrate Client. See [Section 6.5, “Managing Post-Migration Actions \(Windows and Linux\),” on page 121](#).

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

You can increase the upload size limit by modifying the value for `httpRuntime` element’s attribute `maxRequestLength` in the `web.config` file in the `..\Program Files\PlateSpin Migrate Server\Web\` directory on the PlateSpin Server host.

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

To modify the upload size limit for Migrate Client post-migration actions:

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

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

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

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

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

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

5.15 Other Use Cases for Custom PlateSpin Server Settings (Advanced)

[Table 5-5](#) lists configuration keys and values that might address various environmental or functional issues.

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

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

Issue or Use Case	Value Shown in the Config Tool
Discovery/Inventory issues	<pre><add key="UseServiceForCommonInventory" value="true" /> <add key="UseServiceForMigrateInventory" value="false" /> <add key="EarliestWindowsVersionForCommonInventory" value="5.2" /></pre>
Target boot issues related to drivers	<pre><add key="TargetMachineRegistryCleanupLevel" value="None" /></pre>
Controller installation issues on sources (mainly due to environmental constraints)	<pre><add key="InstallControllerUsingService" value="true" /> <add key="RunControllerUnderSameCredsAsRemoteAccess" value="false" /></pre>
Issues related to database size growth	<pre><add key="PowerConvertDBSizeLimitInBytes" value="4294967296" /> <add key="PlateSpinDBCleanupThresholdPercent" value="80" /> <add key="OFXDBCleanupThresholdInBytes" value="4294967296" /></pre>

6 Configuring PlateSpin Migrate Client

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

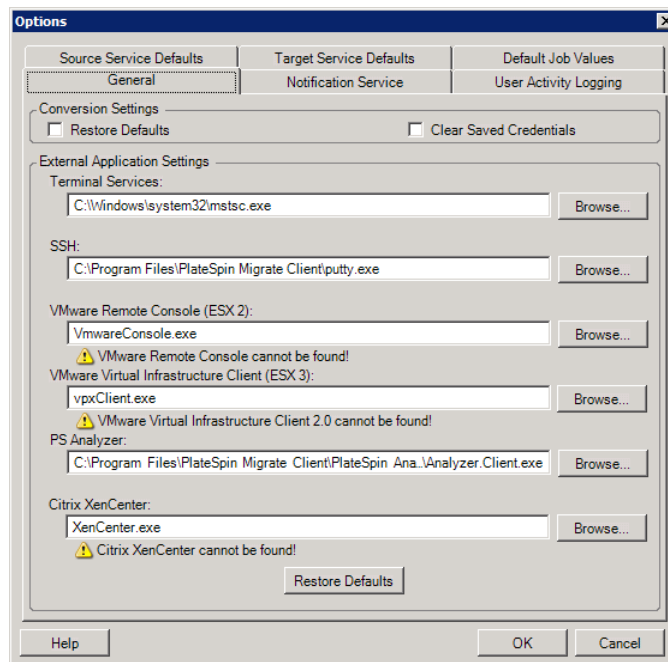
- ♦ [Section 6.1, “Configuring General Options,” on page 117](#)
- ♦ [Section 6.2, “Configuring Job Values Defaults,” on page 118](#)
- ♦ [Section 6.3, “Configuring Source Service Defaults,” on page 119](#)
- ♦ [Section 6.4, “Configuring Target Service Defaults,” on page 120](#)
- ♦ [Section 6.5, “Managing Post-Migration Actions \(Windows and Linux\),” on page 121](#)
- ♦ [Section 6.6, “Managing Migrate Client User Activity Log,” on page 123](#)

6.1 Configuring General Options

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

To configure these general options:

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



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

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

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

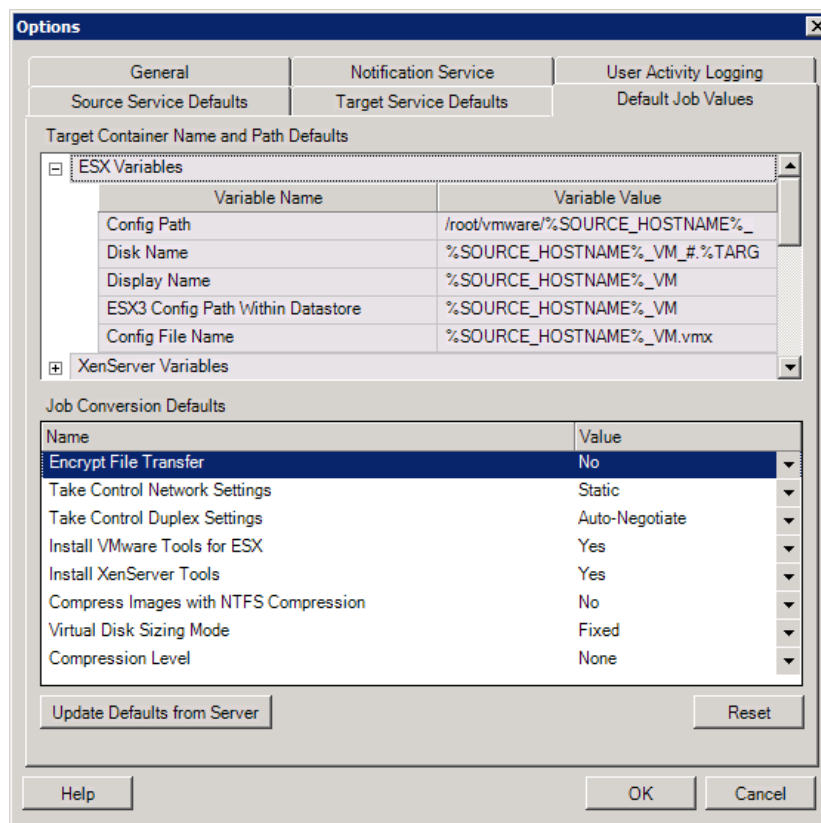
Restore Defaults: Resets the paths to their defaults.

6.2 Configuring Job Values Defaults

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

To configure the default job values:

- 1 Launch the PlateSpin Migrate Client.
- 2 Click **Tools > Options**.
- 3 Click the **Default Job Values** tab.



Target Container Name and Path Details: These variables control naming conventions for target paths and containers. Expand a variable set and click a variable to edit its value. You can edit the following variables:

- ♦ **SOURCE_HOSTNAME:** Host name of the source computer.
- ♦ **IMAGESERVER_LOCATION:** Relevant with image server settings; specify this variable in the **Image Server Variables** section.
- ♦ **SOURCE_VOLUME_SERIAL_NUMBER:** Volume serial number of the source computer.

Job Conversion Defaults: Use this area to set defaults that affect all migration jobs. You can overwrite these defaults when configuring actual migration jobs.

- ♦ **Encrypt File Transfer:** See [Security and Privacy](#).
- ♦ **Install VMware Tools for ESX/XenServer Tools:** See [Virtualization Enhancement Software](#).
- ♦ **Compress Images with NTFS Compression:** See [Section 36.3, “Capturing a Workload to a PlateSpin Image,” on page 417](#). Unrelated to data compression for over-the-network transfer.
- ♦ **Virtual Disk Sizing Mode:** (ESX only)
 - ♦ **Fixed**—space is pre-allocated for the virtual disk
 - ♦ **Dynamic**—the virtual disk is assigned a minimum amount of space, which grows when needed.
- ♦ **Compression Level:** See [Data Compression](#).

Reset: Restores default job values.


Update Defaults from Server: Click to retrieve defaults from the PlateSpin Server if available.

6.3 Configuring Source Service Defaults

PlateSpin Migrate Client enables you to select Windows services and Linux daemons to stop on the source workload during a Live Transfer migration. See [Section 26.15, “Services or Daemons to Stop before Replication or Cutover,” on page 304](#).

To configure the default services on the source:

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



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

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

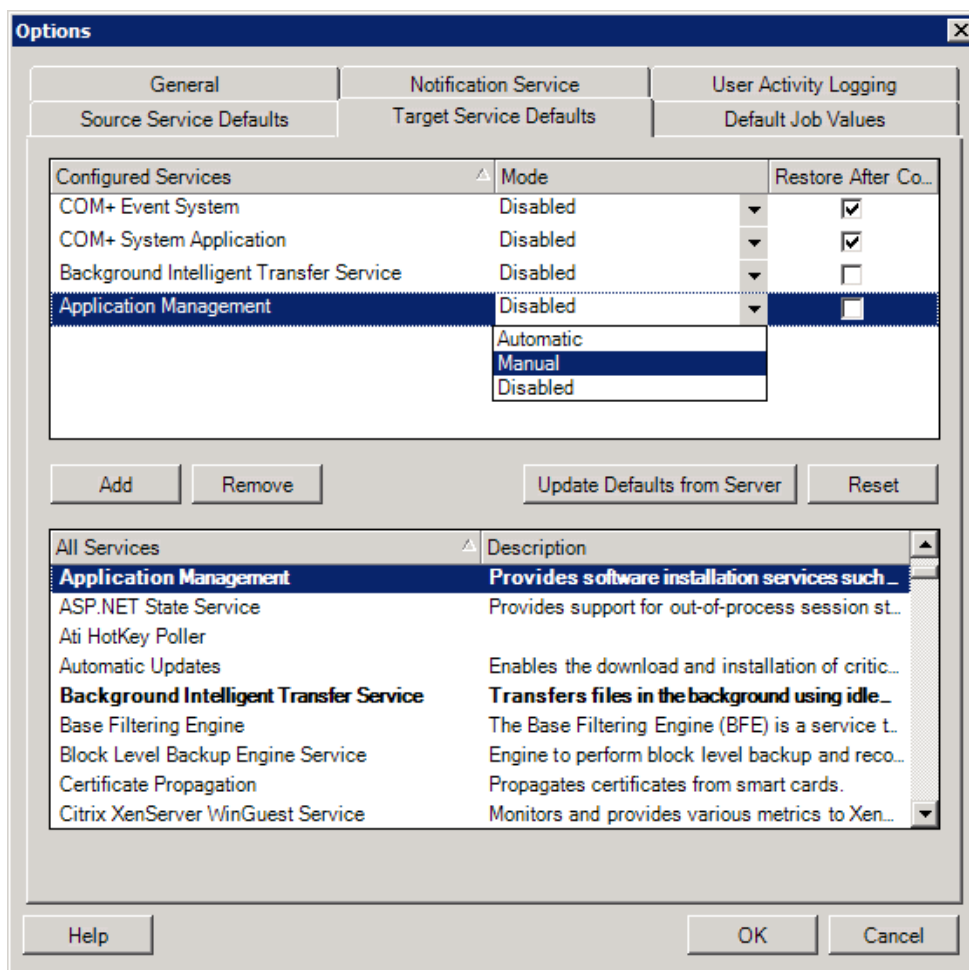
Update Defaults from Server: Retrieves defaults from PlateSpin Server.

6.4 Configuring Target Service Defaults

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

To configure the default services on the target:

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



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

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

Remove: Removes a service.

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

6.5 Managing Post-Migration Actions (Windows and Linux)

PlateSpin Migrate supports the use of scripts to automatically execute custom post-migration tasks on the target workload for certain migration jobs that are performed using PlateSpin Migrate Client. Custom post-migration actions are supported for the following job types:

- One-time Server Sync
- Peer-to-peer workload migration

You configure the action in a batch file, a shell script, or a program executable, then upload them to the PlateSpin Server library of custom actions. You can then associate them with migration jobs you configure in the PlateSpin Migrate Client. At the end of the migration process, PlateSpin Migrate uploads the specified action, along with its dependencies, to the target and executes it.

For the capability to select a post-migration action to run as part of a migration job, you must first save the action and its dependencies in a dedicated directory and add it to the PlateSpin Server library. The maximum size of the directory you upload must not exceed 64 MB. For information about increasing this limit, see [Section 5.14, “Increasing the Upload Size Limit for Post-Migration Actions,” on page 114](#).

To add a post-migration action to the PlateSpin Server library of custom actions:

- 1 Create the action, test it on a sample workload, and save it together with its dependencies in a directory that the PlateSpin Server can access.

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

For Linux operating systems, use tar (or a similar tool) to amalgamate the action’s directory structure into a single file. See [KB Article 7970214 \(https://www.netiq.com/support/kb/doc.php?id=7970214\)](#).

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

The 'Add Action' dialog box is shown with the following details:

- Add Action:**
 - Action Name: X2P_IPconfig
 - Operating System: ☒ Windows, ☐ Linux, ☐ Solaris
- Files:**
 - Action Folder: C:\temp\810\zMOCKUPS (with a 'Browse...' button)
 - File List Table:

File Name	Size	Date Modified
(post-conversion) ACTIONS\X2P_IP\1.bat	2.73 KB	6/12/2009 12:05:18 AM
(post-conversion) ACTIONS\X2P_IP\2.bat	30.29 KB	6/12/2009 12:07:18 AM
(post-conversion) ACTIONS\X2P_IP\3.bat	18.42 KB	6/12/2009 12:08:00 AM
(post-conversion) ACTIONS\X2P_IP\X2P_IP.bat	4.65 KB	6/12/2009 12:08:17 AM
(post-conversion) ACTIONS\X2P_IP\X2P_IP.exe	228.05 KB	6/12/2009 12:12:32 AM
 - Buttons: 'Refresh' and 'Go to Folder' are to the right of the table.
 - File to Execute: (post-conversion) ACTIONS\X2P_IP\X2P_IP.exe (with a 'Set' button)
- Default Options:**
 - Command line arguments: (empty text box)
 - Execution timeout: 0 seconds (with a 'No timeout' checkbox checked)

At the bottom are 'Help', 'OK', and 'Cancel' buttons.

- 4 In the Add Action window, type a name for your custom action, select the target operating system type, then browse to and select the directory that contains the required action with its dependencies.

PlateSpin Migrate populates the list with the contents of the selected folder.

- 5 In the **File Name** column, select the required executable, then click **Set**.
- 6 In the **Default Options** section, specify any required command line arguments and an execution timeout, then click **OK**.

PlateSpin Migrate packages and uploads the library.

The action is now available for selection in migration jobs. See [Section 26.14, “Custom Post-Migration Actions,” on page 303](#).

6.6 Managing Migrate Client User Activity Log

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

- ♦ [Section 6.6.1, “About the Migrate Client User Activity Log,” on page 123](#)
- ♦ [Section 6.6.2, “Configuring Migrate Client User Activity Logging,” on page 124](#)
- ♦ [Section 6.6.3, “Viewing Migrate Client User Activity Log,” on page 124](#)

6.6.1 About the Migrate Client User Activity Log

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

The format of an individual log entry is:

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

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

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

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

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

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

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

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

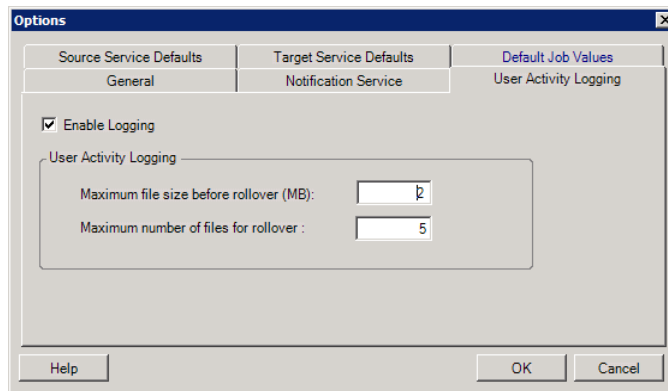
```
PlateSpin.UserActivityLogging.log.1  
PlateSpin.UserActivityLogging.log.2  
PlateSpin.UserActivityLogging.log.3
```

6.6.2 Configuring Migrate Client User Activity Logging

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

To configure User Activity Logging:

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



- 3 Click the **User Activity Logging** tab.
- 4 Specify the following options:

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

- 5 Click **OK**.

6.6.3 Viewing Migrate Client User Activity Log

- 1 Log in to the PlateSpin Migrate Server host as the Administrator.
- 2 Go to the `..\PlateSpin Migrate Server\logs` directory.
- 3 Make a copy of the `PlateSpin.UserActivityLogging.log` file, then open the copy in a text editor.

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

7 Configuring PlateSpin Migrate Web Interface

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

- [Section 7.1, “Managing Security Groups and Workload Permissions,” on page 125](#)
- [Section 7.2, “Managing Workload Tags,” on page 127](#)
- [Section 7.3, “Configuring the Refresh Rates for PlateSpin Migrate Web Interface,” on page 128](#)
- [Section 7.4, “Customizing the UI for PlateSpin Migrate Web Interface,” on page 128](#)

7.1 Managing Security Groups and Workload Permissions

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

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

- [Section 7.1.1, “Prerequisites for Security Groups,” on page 125](#)
- [Section 7.1.2, “Creating Security Groups for Migrate Web Interface,” on page 126](#)
- [Section 7.1.3, “Modifying Security Group Members or Workloads,” on page 126](#)
- [Section 7.1.4, “Deleting a Security Group,” on page 126](#)

7.1.1 Prerequisites for Security Groups

The default users created during Migrate installation are added to every security group you create, by default. For efficient separation of permissions, you must create additional users and assign them to appropriate Workload Migration Roles (Administrator, Power User, or Operator) with permissions that best suit their role in your organization. For more information about the workload migration roles and how to configure them, see [Section 4.1.1, “PlateSpin Migrate Roles,” on page 81](#).

You must also discover workloads to be migrated by using the PlateSpin Migrate Web Interface. After discovery, you can add the workloads to an appropriate security group to be processed by its members for migration configuration and execution, according to the permissions allowed by each user’s assigned roles. See [Section 20.6.2, “Workload Discovery in the Migrate Web Interface,” on page 215](#).

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

- 2 Discover workloads for migration.

7.1.2 Creating Security Groups for Migrate Web Interface

- 1 In the PlateSpin Migrate Web Interface, click **Settings > Permissions**.
- 2 On the Security Groups page, click **Create Security Group**.
- 3 In the **Security Group Name** field, specify a name for the security group.
- 4 (Optional) Click **Add Users** to select the users you want to grant access to this security group and click **OK**.

A PlateSpin Migrate user you recently added to the PlateSpin Server host might not immediately list in the user interface. To list such newly added users, click **Refresh User Accounts**.
- 5 (Optional) In the Migrate Web Interface, add workloads to PlateSpin Migrate that you want to add to the security group.

See [Section 20.6, “Discovering Details for Source Workloads,”](#) on page 213.
- 6 (Optional) Click **Assign Workloads**, select the workloads you want to include in this group, then click **OK**.

Only the users who are members of this security group have access to these workloads.
- 7 Click **Create** to add the new group to the security groups list on the Security Groups page.

7.1.3 Modifying Security Group Members or Workloads

- 1 In the Migrate Web Interface, select **Settings > Permissions**.
- 2 On the Security Groups page, click the security group name, then edit the group information as required:
 - ♦ **Add Users**
 - ♦ **Remove** assigned users

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

7.1.4 Deleting a Security Group

- 1 In the Migrate Web Interface, select **Settings > Permissions**.
- 2 On the Security Groups page, click **Delete** next to the name of the security group you want to delete.

You cannot delete the default All Workloads security group with the default Migrate users.
- 3 Click **OK** to confirm the deletion.

7.2 Managing Workload Tags

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

- ♦ [Section 7.2.1, “Creating a Workload Tag,” on page 127](#)
- ♦ [Section 7.2.2, “Using Workload Tags,” on page 127](#)
- ♦ [Section 7.2.3, “Modifying a Workload Tag,” on page 127](#)
- ♦ [Section 7.2.4, “Deleting a Workload Tag,” on page 127](#)

7.2.1 Creating a Workload Tag

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

To create workload tags:

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

7.2.2 Using Workload Tags

After you create tags, they are available on the Edit Target Details page where you can associate a tag to the appropriate workloads. Use the Tags column on the Workloads view to visually group similar workloads so that you can easily manage operations on these workloads. For information about associating tags with workloads, see [Section 20.9, “Using Tags to Track Logical Associations of Workloads,” on page 218](#).

7.2.3 Modifying a Workload Tag

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

To modify a workload tag:

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

7.2.4 Deleting a Workload Tag

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

To delete a workload tag:

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

7.3 Configuring the Refresh Rates for PlateSpin Migrate Web Interface

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

Table 7-1 Web Interface Default Refresh Intervals

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

- 1 Open the following file in a text editor:

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

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

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

- 3 Save the file.

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

7.4 Customizing the UI for PlateSpin Migrate Web Interface

You can modify the appearance of PlateSpin Migrate Web Interface to match the look and feel of your corporate identity. You can modify colors, logo, and product name. For more information, see [Appendix B, "Rebranding the UI for PlateSpin Migrate Web Interface," on page 129](#).

B Rebranding the UI for PlateSpin Migrate Web Interface

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

- ♦ [Section B.1, “Rebranding the UI Using PlateSpin Configuration Parameters,” on page 129](#)
- ♦ [Section B.2, “Rebranding the Product Name in the Windows Registry,” on page 133](#)

B.1 Rebranding the UI Using PlateSpin Configuration Parameters

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

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

- 1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/PlateSpinConfiguration/`
- 2 Locate the required PlateSpin Server Configuration parameter and click **Edit** to change its value.
- 3 Click **Save**.

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

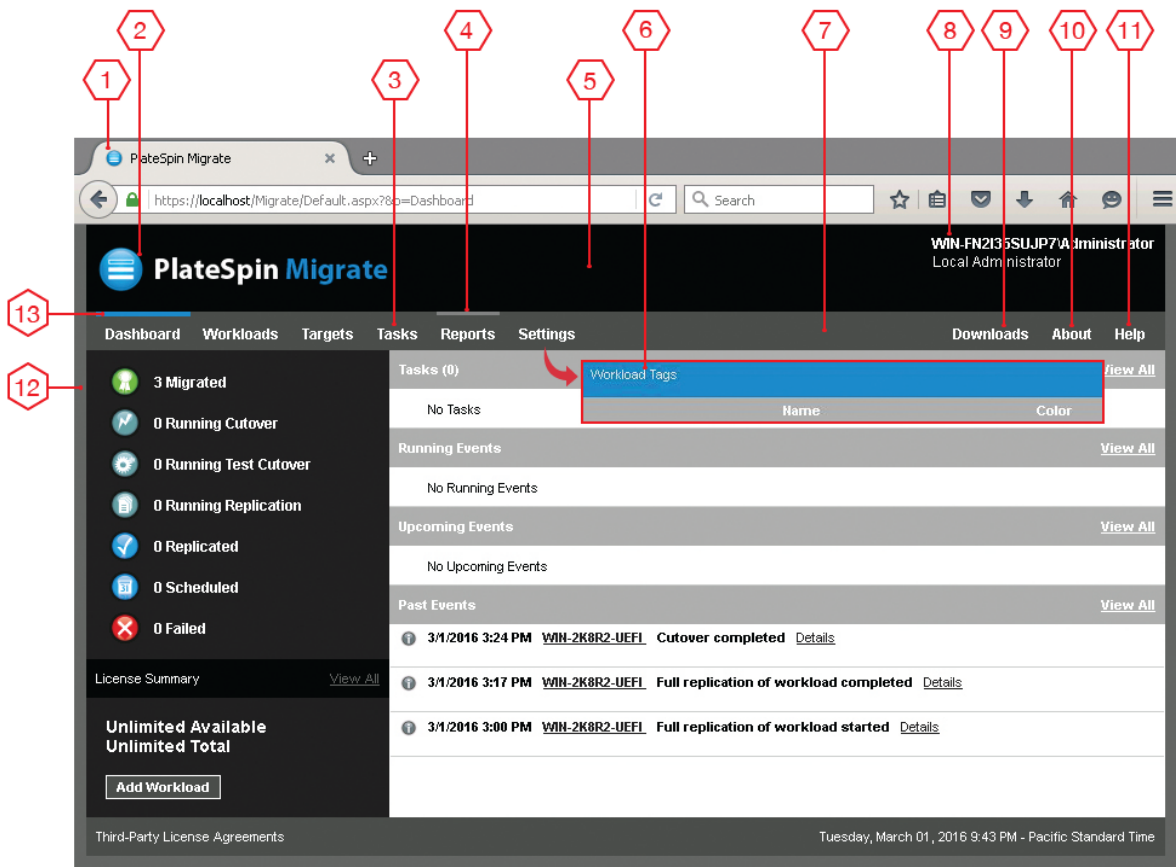
The following sections provide information about the configurable elements in the UI for PlateSpin Migrate Web Interface.

- ♦ [Section B.1.1, “About Configurable UI Elements for PlateSpin Migrate Web Interface,” on page 129](#)
- ♦ [Section B.1.2, “Modifying PlateSpin Configuration Settings for Configurable UI Elements,” on page 130](#)

B.1.1 About Configurable UI Elements for PlateSpin Migrate Web Interface

The look and feel of the PlateSpin Migrate Web Interface is consistent across the various pages. The illustration in [Figure B-1](#) of the PlateSpin Migrate Dashboard identifies the elements that you can modify in the Web Interface UI with numbered callouts. For information about the related parameter for each element, see [Section B.1.2, “Modifying PlateSpin Configuration Settings for Configurable UI Elements,” on page 130](#).

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



B.1.2 Modifying PlateSpin Configuration Settings for Configurable UI Elements

Table B-1 provides information about the setting you must use to modify the corresponding interface element. The ID column in the table lists the ID of the interface element identified in Figure B-1 provided in Section B.1.1, “About Configurable UI Elements for PlateSpin Migrate Web Interface,” on page 129.

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

ID	Setting Name and Description	Default Value
1	<p>WebUIFaviconUrl</p> <p>Location of a valid .ico graphic file. Specify one of the following:</p> <ul style="list-style-type: none"> ♦ A valid URL to the appropriate .ico file on a different machine. For example: https://myserver.example.com/dir1/dir2/icons/mycompany_favicon.ico ♦ A relative path below the root of the local web server where you have uploaded the appropriate .ico file. For example, if you create a path called mycompany\images\icons at the root of the web server to store your custom icon graphics: ~/mycompany/images/icons/ mycompany_favicon.ico In this example, the actual file system path that contains the file is C:\Program Files (x86)\PlateSpin Migrate Server\PlateSpin Forge\web\mycompany\images\icons\mycompany_favicon.ico. 	~/doc/en/favicon.ico ¹
2	<p>WebUILogoUrl</p> <p>Location of product logo graphic file. Specify one of the following:</p> <ul style="list-style-type: none"> ♦ A valid URL to the appropriate graphics file on a different machine. For example: https://myserver.example.com/dir1/dir2/logos/mycompany_logo.png ♦ A relative path below the root of the local web server where you have uploaded the appropriate graphics file. For example, if you create a path called mycompany\images\logos at the root of the web server to store your custom logo images: ~/mycompany/images/logos/mycompany_logo.png In this example, the actual file system path that contains the file is C:\Program Files (x86)\PlateSpin Migrate Server\PlateSpin Forge\web\mycompany\images\logos\mycompany_logo.png. 	~/Resources/protectLogo.png ²
3	<p>WebUISiteNavigationFontColor</p> <p>Color of site navigation link font color in Web UI (RGB hex value)</p>	#FFFFFF
4	<p>WebUISiteNavigationLinkHoverBackgroundColor</p> <p>Color of site navigation link background in hover state (RGB hex value)</p>	#808080

ID	Setting Name and Description	Default Value
5	WebUISiteHeaderBackgroundColor Site header background color (RGB hex value)	#000000
6	WebUISiteAccentFontColor Font color to display with accent color in Web UI (RGB hex value)	#FFFFFF
7	WebUISiteNavigationBackgroundColor Color of site navigation background in Web UI (RGB hex value)	#4D4D4D
8	WebUISiteHeaderFontColor Site header font color in Web UI (RGB hex value)	#FFFFFF
9	WebUIShowDownloadsTab Toggles the visibility of the Downloads tab: <ul style="list-style-type: none"> ♦ True: The Downloads tab is visible on the interface. ♦ False: The Downloads tab is not visible on the interface. 	True
10	WebUIShowAboutTab Toggles the visibility of the About tab: <ul style="list-style-type: none"> ♦ True: The About tab is visible on the interface. ♦ False: The About tab is not visible on the interface. 	True
11	WebUIShowHelpTab Toggle the visibility of the Help tab: <ul style="list-style-type: none"> ♦ True: The Help tab is visible on the interface. ♦ False: The Help tab is not visible on the interface. 	True
12	WebUISiteBackgroundColor Site background color (RGB hex value)	#666666
13	WebUISiteAccentColor Accent color (RGB hex value)	#0088CE

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

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

B.2 Rebranding the Product Name in the Windows Registry

The masthead at the top of the product interface provides space for the corporate logo and the product name. To change the logo, which commonly includes the product name, see [Section B.1, “Rebranding the UI Using PlateSpin Configuration Parameters,”](#) on page 129.

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

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

```
HKEY_LOCAL_MACHINE\SOFTWARE\PlateSpin\MigrateServer\ProductName
```

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

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\PlateSpin\MigrateServer
```

- 4 Double-click the `ProductName` key and change the **Value data** for the key as required and then click **OK**.
- 5 Restart the IIS Server.



Preparing Your Migration Environment

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

- ♦ [Chapter 8, “Prerequisites for Migration to Amazon Web Services,” on page 137](#)
- ♦ [Chapter 9, “Prerequisites for Migration to Microsoft Azure,” on page 147](#)
- ♦ [Chapter 10, “Prerequisites for Migration to VMware vCloud Director,” on page 163](#)
- ♦ [Chapter 11, “Prerequisites for Migration to VMware,” on page 169](#)
- ♦ [Chapter 12, “Prerequisites for Migration to Microsoft Hyper-V,” on page 175](#)
- ♦ [Chapter 13, “Prerequisites for Migration to VMs on Citrix XenServer,” on page 179](#)
- ♦ [Chapter 14, “Prerequisites for Migration to VMs on Xen,” on page 183](#)
- ♦ [Chapter 15, “Prerequisites for Migration to VMs on KVM,” on page 185](#)
- ♦ [Chapter 16, “Prerequisites for Migration to Physical Machines,” on page 187](#)
- ♦ [Chapter 17, “Prerequisites for Migration to an Image,” on page 191](#)
- ♦ [Chapter 18, “Preparing for Synchronization of Workloads with Server Sync,” on page 193](#)

8 Prerequisites for Migration to Amazon Web Services

PlateSpin Migrate Client supports semi-automated migration to your Amazon Web Services (AWS) environment. This section describes the required AWS configuration that you must prepare, such as an AWS account, before you can discover AWS target VMs and configure migrations to them.

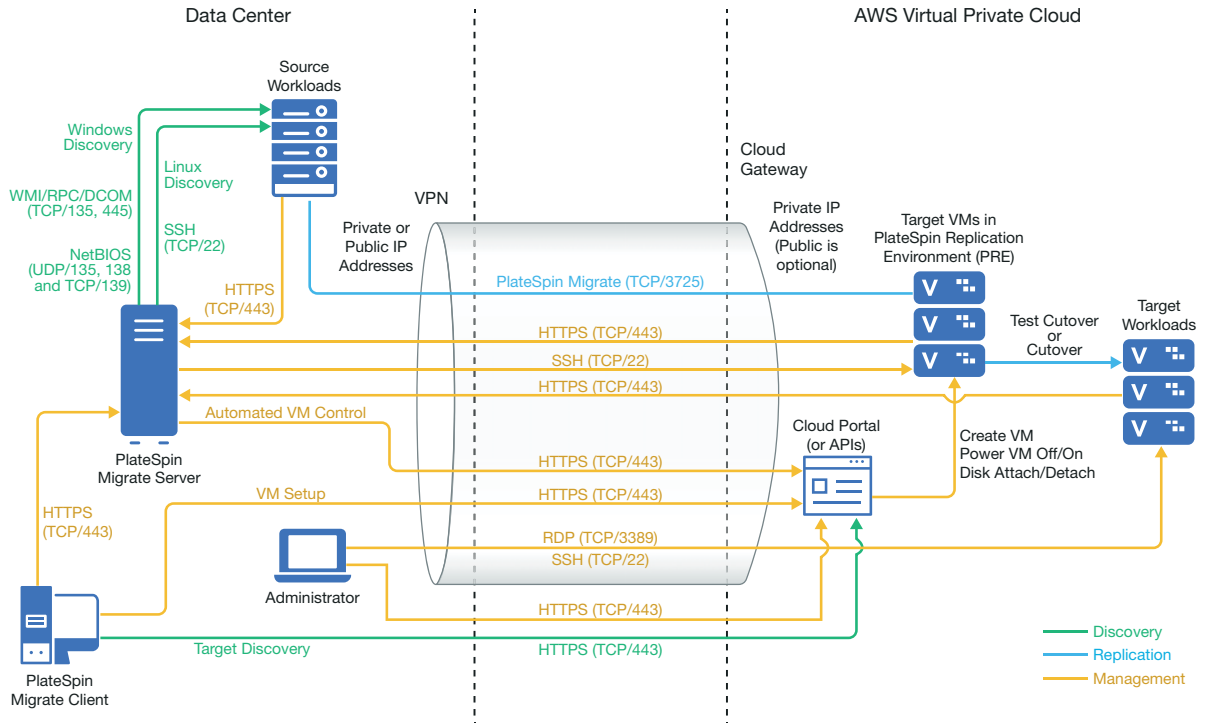
- ♦ [Section 8.1, “Deployment for Migration to Amazon Web Services,” on page 137](#)
- ♦ [Section 8.2, “Requirements for Migrating Workloads to Amazon Web Services,” on page 138](#)
- ♦ [Section 8.3, “Planning For Migrating Workloads to Amazon Web Services,” on page 141](#)
- ♦ [Section 8.4, “Configuring OS License Activation on Windows Targets Migrated to AWS,” on page 142](#)
- ♦ [Section 8.5, “Using the PlateSpin Replication Environment AMIs,” on page 142](#)
- ♦ [Section 8.6, “AWS Networking Guidelines,” on page 143](#)
- ♦ [Section 8.7, “Best Practices For Configuring a Migration Job to Amazon Web Services,” on page 144](#)
- ♦ [Section 8.8, “Checklist for Automated Migration to AWS,” on page 145](#)

8.1 Deployment for Migration to Amazon Web Services

You can deploy a PlateSpin Migrate server on premise in your data center with the source workloads. A site-to-site VPN connection is required between the data center and your account in the AWS cloud. [Figure 8-1](#) shows the location of various components in your AWS migration environment and the communications between them. See [Section 8.3, “Planning For Migrating Workloads to Amazon Web Services,” on page 141](#).

NOTE: [Figure 8-1](#) depicts automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443).

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



8.2 Requirements for Migrating Workloads to Amazon Web Services

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

- ♦ Section 8.2.1, “Minimum AWS Prerequisites,” on page 138
- ♦ Section 8.2.2, “AWS Prerequisites for Using an On Premise Migrate Server,” on page 139

8.2.1 Minimum AWS Prerequisites

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

Table 8-1 Minimum Required Configuration for Your AWS Account

AWS Configuration	Description
AWS Account	To create an AWS account, go to Amazon Web Services Console (http://aws.amazon.com) .
AWS EC2 Subscription	PlateSpin supports only Amazon Virtual Private Cloud (VPC).
Amazon Virtual Private Cloud (VPC)	Create an AWS VPC to launch AWS resources into your virtual network. See Amazon Virtual Private Cloud Documentation .

AWS Configuration	Description
An AWS EC2 Key Pair	A Key pair to log in to your AWS instance. By default, Amazon generates the key pair in the .pem format.
AWS user credentials	<p>You need an AWS Identity and Access Management (IAM) user in your AWS account, with an appropriate IAM role to perform migrations into the VPC using the AWS APIs. Enable Programmatic Access for the user to generate an access key and a secret access key. AWS Management Console Access is optional, but it can be useful for troubleshooting.</p> <p>NOTE: We recommend that administrators regularly rotate access keys for IAM users. See “Rotating Access Keys” in the <i>AWS Identity and Access Management User Guide</i>.</p> <p>For information about setting up the migration user group, policy, and user, see “Create an IAM User for Your Workload Migrations” in the <i>Best Practices for Migrating Servers to Amazon Web Services with PlateSpin Migrate white paper</i>.</p>
Amazon Machine Image (AMI) in your VPC for PlateSpin Migrate Replication Environment	<p>PlateSpin Migrate provides the following AMIs in the Community AMI section of Amazon Web Services Console. For full replication, use one of the following AMIs depending on the source workload OS and your desired OS licensing model on the target workload. For Server Sync, create a private instance of any one of the following AMIs and use the snapshot of the created AMI.</p> <ul style="list-style-type: none"> ♦ PlateSpin Replication Environment - Linux: Applicable for Linux workloads. AWS does not bill you for the OS license on the target workload. ♦ PlateSpin Replication Environment - Windows: Applicable for Windows workloads and is billable. AWS manages the Microsoft software licensing compliance on the target workload and bills you accordingly for the license. ♦ PlateSpin Replication Environment - Windows (BYOL): Applicable for Windows workloads and is non-billable. AWS allows you to bring your own licenses (BYOL), which you have already purchased from Microsoft. You are responsible for complying with Microsoft licensing and AWS does not bill you for the license. <p>For information about using the PlateSpin Replication Environment AMIs, see Section 8.5, “Using the PlateSpin Replication Environment AMIs,” on page 142.</p>

For more information about setting up your AWS cloud account to work with PlateSpin Migrate, see the white paper “Best Practices for Migrating Servers to Amazon Web Services with PlateSpin Migrate” on the [PlateSpin Migrate Resources web page \(https://www.microfocus.com/products/migrate/resources/\)](https://www.microfocus.com/products/migrate/resources/).

8.2.2 AWS Prerequisites for Using an On Premise Migrate Server

- ♦ A PlateSpin Migrate license.
- ♦ PlateSpin Migrate server installed on premise in a network that can properly access the source workloads.

- ◆ PlateSpin Migrate Client installed on the PlateSpin Migrate server or on another machine in the same network as the PlateSpin Migrate server. Migrate Client must be able to connect to the PlateSpin Migrate server and the AWS portal.
- ◆ A site-to-site VPN connection connecting the AWS gateway to your on-premise gateway.

For information, see the following AWS resources:

- ◆ [VPN Connections](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpn-connections.html) (<http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpn-connections.html>)
- ◆ [AWS-Managed VPN Connections](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html) (http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html)
- ◆ For a Windows workload migrated to AWS to be a part of domain, you must configure the DNS server details of the domain in the AWS subnet.
- ◆ An AWS Security Group and the VPC gateway that provides the following inbound and outbound rules. For instructions, see “Create a Security Group” in the Best Practices white paper.

Inbound Rules

- ◆ TCP, port 3725, custom
Provide an address range covering all source workloads.
- ◆ SSH, port 22
Provide the IP address of the PlateSpin Migrate server.
- ◆ RDP, port 3389
Provide the IP address of the machine you plan to use to launch an RDP connect to target workloads.

Outbound Rules

- ◆ TCP, port 3725, custom
Provide an address range covering all source workloads.
- ◆ HTTPS, port 443
Provide the IP address of the PlateSpin Migrate server.

Port 3725 is the default port number for data transfer. By default, the data transfer is initiated from the target workload to the source workload. The port number and direction for initiating the connection are configurable.

- ◆ The minimum network-related prerequisites for a successful migration are:
 - ◆ The source and the target workload must be able to communicate with the PlateSpin Migrate server on port 443. The target workload is the replica of the source workload that will reside in AWS.
 - ◆ The PlateSpin Migrate server must be able to communicate with the AWS API endpoint on port 443.
 - ◆ The PlateSpin Migrate server must be able to communicate with the source workloads on the ports that are used for discovery. See [“Requirements for Discovery” on page 49](#) and [Section 20.6, “Discovering Details for Source Workloads,” on page 213](#).
You can alternatively use the Migrate Agent utility to register source workloads with the Migrate server using HTTPS (TCP/port 443). See [Section 20.7, “Registering Workloads and Discovering Details with Migrate Agent,” on page 215](#).
 - ◆ The cloud-based target workload must be able to communicate (target to source) with the on-premise source workload on port 3725 (TCP) over the site-to-site VPN connection.

The port number is configurable. See port 3725 in [Section 2.6.2, “Requirements for Migration,” on page 50](#).

If you use Migrate Agent for registration and discovery, the default direction of the replication connection must be reversed (source to target) by changing advanced settings on the Migrate server. See [Section 5.9, “Configuring the Contact Direction for the Replication Port,” on page 107](#).

For detailed access and communication requirements across your migration network, See [“Access and Communication Requirements across your Migration Network” on page 49](#).

- ♦ The following Microsoft .NET Framework requirements for source workloads:
 - ♦ Net Framework 4.5 or later is required for migrating Windows Server 2008 R2 workloads.
 - ♦ Net Framework 3.5 or later is required for migrating Windows Server 2003, Windows Server 2003 R2, or Windows Server 2008 workloads.
- ♦ The volume labels on the source workloads must not include any special characters.
- ♦ On the Windows workloads that you want to migrate, you must set the SAN policy to `OnlineAll` in order to ensure that the volumes on the target workload are mounted same as on the source workload.

8.3 Planning For Migrating Workloads to Amazon Web Services

PlateSpin Migrate allows you to use the PlateSpin Migrate Client to migrate Windows and Linux workloads to AWS. For a list of supported workloads, see [“Supported Workloads For Migration to Amazon Web Services” on page 31](#).

NOTE: Migration of AWS source workloads to AWS target containers is not supported.

Consider the following points before you use the PlateSpin Migrate Client to migrate workloads to AWS:

- ♦ Migration of Windows Cluster workloads is not supported.
- ♦ Windows and Linux UEFI workloads are migrated as BIOS workloads.
- ♦ Use PlateSpin Migrate Client to migrate workloads to AWS. The PlateSpin Migrate Web Interface does not support migration of workloads to AWS.
- ♦ PlateSpin Migrate provides a configuration parameter named `AWSActivateWindows` on the PlateSpin Configuration page. By default, this parameter is set to **True** to allow AWS to activate Windows license on a Windows target workload. However, if you want to use your own Microsoft license for which you are responsible for complying with Microsoft licensing, then set this parameter to **False**. For more information about the `AWSActivateWindows` and its related parameters, see [Section 5.7, “Configuring Alternate IP Addresses for PlateSpin Server,” on page 106](#).
- ♦ PlateSpin Migrate supports AWS target instances with up to 26 volumes (EBS volumes) for Windows and 40 volumes (EBS volumes) for Linux with each disk not exceeding 15 volumes.
- ♦ Migrate recommends an AWS instance size that meets or exceeds the source workload's settings for cores, memory, volumes, and NICs. However, you can choose a smaller or larger instance size based on your requirements for the target workload, as limited by the maximum instance sizes available in the AWS region.
- ♦ The size of the disk created on the AWS instance is the size of the source disk partition plus about 1 GB.

8.4 Configuring OS License Activation on Windows Targets Migrated to AWS

PlateSpin Migrate provides the following parameters to configure whether or not AWS should manage the Microsoft software licensing compliance on the target workload:

- ♦ **AWSActivateWindows:** This parameter lets you set whether you want AWS to activate Windows OS license on a Windows target workload or let you use your own licenses for which you are responsible for complying with Microsoft licensing. The default value of this parameter is **True**, which implies that AWS activates the Windows OS license on the target workload provided you specify the correct AMI ID of your private copy of billable AMI when you prepare the AWS target for migration. See [Section 27.2, “Configuring Migration of a Workload to Amazon Web Services,”](#) on page 332.
- ♦ **AWSKMSservers:** This parameter lets you set the AWS KMS Server information that Windows instances use for activation and is applicable only if the value of **AWSActivateWindows** parameter is set to true. The target KMS Server should be in the same AWS Region where the Windows instance is running.
- ♦ **KMSClientSetupKeys:** This parameter lists the commonly used OS version-based Microsoft KMS client setup keys that are used for activating Windows through KMS server. This parameter is applicable only if the value of **AWSActivateWindows** is set to true. If the key for a particular OS is not listed, you can add an entry in the following format:

`OperatingSystemTypeandBranding="Microsoft provided KMS Key"`

Example: For a Windows server with OS type as Windows 2016 and branding as Standard Server, the format is `Windows2016StandardServer="WC2BQ-8NRM3-FDDYY-2BFGV-KHKQY"`

To edit the value of the AWS parameters:

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

8.5 Using the PlateSpin Replication Environment AMIs

PlateSpin Migrate provides the following PlateSpin Replication Environment (PRE) AMIs in the Community AMI section of Amazon Web Services Console:

- ♦ **PlateSpin Replication Environment - Linux:** Applicable for Linux workloads. AWS does not bill you for the OS license on the target workload.
- ♦ **PlateSpin Replication Environment - Windows:** Applicable for Windows workloads and is billable. AWS manages the Microsoft software licensing compliance on the target workload and bills you for the license.
- ♦ **PlateSpin Replication Environment - Windows (BYOL):** Applicable for Windows workloads and is not billable. AWS allows you to bring your own licenses (BYOL), which you have already purchased from Microsoft. You are responsible for complying with Microsoft licensing and AWS does not bill you for the license.

For information on how to use the PRE AMIs, review the following:

- ♦ [Section 8.5.1, “For Full Replication,” on page 143](#)
- ♦ [Section 8.5.2, “For Server Sync,” on page 143](#)

8.5.1 For Full Replication

PlateSpin Migrate requires an ID of the PRE AMI based on the source workload OS (Windows, Linux) and your desired OS licensing model (billable, non-billable) on the target workload.

8.5.2 For Server Sync

PlateSpin Migrate requires a snapshot of any one of the PRE AMIs. You must create a private instance of any one of the PRE AMIs and then create a snapshot of that AMI.

To create a snapshot using the PRE AMI:

- 1 Go to the Amazon Web Services Console and log in to your AWS Console.
- 2 Launch an instance of any of the PRE AMIs.
- 3 Review and configure the instance launch details as required. Then, click **Launch** to assign the key pair to your instance and create the instance.
- 4 Click the instance ID displayed on the Console to go to the launched instance.
- 5 Select the instance, click **Actions** > **Image** > **Create Image** to create an AMI.
- 6 Note down the AMI ID displayed on the Console.
- 7 Go to the AMI and note down the corresponding Snapshot ID of the root device.
- 8 Terminate the instance created in [Step 3](#).
- 9 De-register the AMI that you created in [Step 5](#).

8.6 AWS Networking Guidelines

Consider the following guidelines when you are migrate workloads to AWS:

- ♦ [Section 8.6.1, “Workloads with Multiple NICs,” on page 143](#)
- ♦ [Section 8.6.2, “Private and Public IP Addresses for Workloads Connected on an AWS VPN,” on page 144](#)

8.6.1 Workloads with Multiple NICs

For workloads with multiple NICs, AWS provides only private IP addresses for the target instance. So, you can connect to the target instance using the private IP addresses only.

8.6.2 Private and Public IP Addresses for Workloads Connected on an AWS VPN

Each AWS VM has both a public IP address and a private IP address for communications from machines outside the AWS environment. AWS automatically associates these IP addresses with the primary network interface for the VM.

You can use the Microsoft Remote Desktop client or SSH to remotely connect to the AWS VM. Specify the IP address as follows:

- ♦ **Private IP address:** Use the VM's private IP address if your machine is part of the address space for the AWS VPN.
- ♦ **Public IP address:** Use the VM's public IP address if your machine is not part of the address space for the AWS VPN. A public IP address is not set on the target workload that has multiple NICs.

8.7 Best Practices For Configuring a Migration Job to Amazon Web Services

To help prevent the failure of a migration job to AWS, you must adopt the following best practices when you configure migration jobs:

- ♦ If you use a static IP address for the network, ensure that the address is unique within the supported subnet range.
- ♦ The number of target instances running at any point of time must not exceed the instance limit applicable for your subscription.

NOTE: Every migration job uses a temporary AWS instance for the migration.

8.8 Checklist for Automated Migration to AWS

Task	Description
1. Prepare your VMware migration environment.	Figure 8-1, “On-Premise Migrate Server for Semi-Automated Migration to AWS,” on page 138 Section 8.3, “Planning For Migrating Workloads to Amazon Web Services,” on page 141
2. Discover target containers.	Section 19.4.1, “Target Discovery in the Migrate Client,” on page 200
3. Discover source workloads.	Section 20.6.1, “Workload Discovery in the Migrate Client,” on page 214 -OR- Section 20.7, “Registering Workloads and Discovering Details with Migrate Agent,” on page 215
4. Configure target workload migration.	Section 27.2, “Configuring Migration of a Workload to Amazon Web Services,” on page 332
5. Execute migration.	Chapter 38, “Executing Workload Migrations,” on page 433

9 Prerequisites for Migration to Microsoft Azure

PlateSpin Migrate Web Interface supports automated migration to Microsoft Azure Cloud environments based on your migration goals: Azure global or sovereign Azure China. This section describes the required Azure configuration that you must prepare in the appropriate environment, such as an Azure account, subscriptions, and services, before you can discover Azure target containers and configure migrations to them.

- [Section 9.1, “Deployment for Migration to Azure,” on page 147](#)
- [Section 9.2, “Requirements for Migrating Workloads to Azure,” on page 149](#)
- [Section 9.3, “Planning For Migrating Workloads to Azure,” on page 153](#)
- [Section 9.4, “Azure Networking Guidelines,” on page 154](#)
- [Section 9.5, “Registering a Azure Application to Represent PlateSpin Migrate,” on page 156](#)
- [Section 9.6, “Configuring an Azure User for PlateSpin Migrate to Use,” on page 157](#)
- [Section 9.7, “Enabling Programmatic Deployment of PlateSpin Migrate Replication Environment,” on page 157](#)
- [Section 9.8, “Deploying a Migrate Server Image in Azure,” on page 158](#)
- [Section 9.9, “Configuring the Target IaaS Environment for Azure Cloud,” on page 159](#)
- [Section 9.10, “Managing the Azure User Password for Azure Target Containers,” on page 160](#)
- [Section 9.11, “Checklist for Automated Migration to Azure,” on page 161](#)


9.1 Deployment for Migration to Azure

You can deploy a PlateSpin Migrate server on premise in your data center with the source workloads or in the appropriate Microsoft Azure Cloud environment: Azure global or sovereign Azure China.

For an on-premise Migrate server deployment, you can configure a PlateSpin Migrate server to work with either the Azure global IaaS environment or the sovereign Azure China IaaS environment, but not both environments at the same time. A site-to-site VPN connection is required between the data center and your account in the Azure cloud. [Figure 9-1](#) shows the location of various components in your Azure migration environment and the communications between them. See [Section 9.2.2, “Azure Prerequisites for Using an On Premise Migrate Server,” on page 151](#).


NOTE: [Figure 9-1](#) depicts automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443).

Figure 9-1 On-Premise Migrate Server for Automated Migration to Azure



For a cloud-based Migrate server deployment, the Azure Marketplace in the target Azure environment offers a PlateSpin Migrate Server image that is preconfigured to support its host IaaS environment. Figure 9-2 shows the location of various components in your Azure migration environment and the communications between them. See [Section 9.2.3, “Azure Prerequisites for Using an Azure-Based Migrate Server,”](#) on page 152.

Figure 9-2 Cloud-Based Migrate Server for Automated Migration to Azure



9.2 Requirements for Migrating Workloads to Azure

Based on the location of your PlateSpin Migrate server, review the following sections:

- ♦ [Section 9.2.1, “Minimum Azure Prerequisites,” on page 150](#)
- ♦ [Section 9.2.2, “Azure Prerequisites for Using an On Premise Migrate Server,” on page 151](#)
- ♦ [Section 9.2.3, “Azure Prerequisites for Using an Azure-Based Migrate Server,” on page 152](#)

9.2.1 Minimum Azure Prerequisites

PlateSpin Migrate requires the use of Microsoft Azure Resource Management for migrating workloads into the Microsoft Azure cloud. For migrations to Microsoft Azure Cloud, you must prepare your Azure account, subscriptions, and services in the Azure global environment or the sovereign Azure China environment, based on your migration goals.

[Table 9-1](#) describes the minimum configuration you must perform in the appropriate Azure environment before you can migrate workloads to Azure.

Table 9-1 Minimum Required Configuration for Your Azure Account

Azure Configuration	Description
Microsoft Azure Account	<p>Create a account in the Azure environment where you will migrate workloads.</p> <p>Use the Azure Portal (http://www.azure.com/) for an account for use in the global environment.</p> <p>Use the Azure China Portal (http://www.azure.cn/), to create an account for use in that sovereign environment.</p> <p>In PlateSpin Migrate, you must specify an administrator-level user name and password combination for the Azure account. It can be the Account owner or a user in the Contributor role for the specified Subscription.</p>
Azure Subscription ID	The ID for the Azure Subscription in the specified Azure account that you want to bill for Azure-related costs. An account can have multiple subscriptions.
Contributor user for the subscription created in Azure Active Directory	A user created as a Contributor for the specified subscription in your Azure Active Directory.
A non-administrative user for PlateSpin Migrate interactions with Azure created in Azure Active Directory	See Section 9.6, "Configuring an Azure User for PlateSpin Migrate to Use," on page 157.
Application ID	<p>An ID that represents PlateSpin Migrate as it makes use of the Microsoft Azure API when it replicates or migrates workloads on your behalf to VMs in the target Azure account.</p> <p>See Section 9.5, "Registering a Azure Application to Represent PlateSpin Migrate," on page 156.</p>
Azure Marketplace permission for PlateSpin Migrate Replication Environment images	<p>For each Subscription that you use as a target container, you must permit the programmatic deployments for the PlateSpin Migrate Replication Environment images for the account.</p> <p>See "Enabling Programmatic Deployment of PlateSpin Migrate Replication Environment" on page 157.</p>
Azure Virtual Network and Subnet	<p>You must create least one Virtual Network with a Subnet in the specified Subscription. If you have an site-to-site VPN set up, the subnet must be different than the default Gateway Subnet.</p> <p>Network resources are never created automatically by PlateSpin Migrate, so they always must be set up manually in advance. For instructions, refer to Azure documentation.</p>

Azure Configuration	Description
Azure Storage account	<p>Your VM disks will use the Azure page blob type of general-purpose storage, which can run on Standard (HDD) or Premium (SSD) storage media. A Standard Storage Account can be used for Azure VM sizes that use Standard or Premium storage media. A Premium Storage Account can be used only for Azure VM sizes that use Premium storage media.</p> <p>If no Azure Storage Account is associated with a subscription, PlateSpin Migrate sets up a Standard general-purpose storage account to use as the datastore for the target VM. The datastore name is based on the Azure Resource Group for the Subscription.</p> <p>If you want full control over your Azure Storage Accounts, configure a Standard or a Premium general-purpose storage account for each Azure Subscription before you begin migrating workloads to Azure. Your storage account is shown as a datastore for the target Azure Subscription in the Migrate Web Interface. For information about Azure Storage Accounts, refer to Azure documentation.</p>

For more information about setting up your Azure cloud account to work with PlateSpin Migrate, see the white paper “Best Practices for Migrating Servers to Microsoft Azure with PlateSpin Migrate” on the [PlateSpin Migrate Resources web page](https://www.microfocus.com/products/migrate/resources/) (<https://www.microfocus.com/products/migrate/resources/>).

9.2.2 Azure Prerequisites for Using an On Premise Migrate Server

If you set up an Azure site-to-site VPN (or an Azure Express Route connection) between the premises where your source workloads reside and the target Azure environment, you can deploy your PlateSpin Migrate server on-premises. Before you use PlateSpin Migrate to migrate workloads to Microsoft Azure, ensure that the following cloud access prerequisites are correctly configured and available:

- ♦ A PlateSpin Migrate license.
- ♦ A PlateSpin Migrate server deployed on-premise.
- ♦ A site-to-site VPN connection between your local data center and Microsoft Azure Portal.

For information, see the following Microsoft resources:

- ♦ [Create a Site-to-Site Connection in the Azure Portal](https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-howto-site-to-site-resource-manager-portal) (<https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-howto-site-to-site-resource-manager-portal>)
- ♦ [Create VNet with Site-to-Site VPN Connection Using PowerShell](https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-create-site-to-site-rm-powershell) (<https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-create-site-to-site-rm-powershell>)
- ♦ A default Gateway Subnet.
- ♦ The minimum network-related prerequisites for a successful migration are:
 - ♦ The on-premise source workload and the cloud-based target workload must be able to communicate with the PlateSpin Migrate server through HTTPS (TCP/port 443) over the site-to-site VPN connection.
 - ♦ The on-premise PlateSpin Migrate server must be able to communicate with the Microsoft Azure API endpoint through HTTPS (TCP/port 443).

- ♦ The PlateSpin Migrate server must be able to communicate with the source workloads on the ports that are used for discovery. See [“Requirements for Discovery” on page 49](#) and [Section 20.6, “Discovering Details for Source Workloads,” on page 213](#).

You can alternatively use the Migrate Agent utility to register source workloads with the Migrate server using HTTPS (TCP/port 443). See [Section 20.7, “Registering Workloads and Discovering Details with Migrate Agent,” on page 215](#).

- ♦ The cloud-based target workload must be able to communicate (target to source) with the on-premise source workload on port 3725 (TCP) over the site-to-site VPN connection.

The port number is configurable. See port 3725 in [Section 2.6.2, “Requirements for Migration,” on page 50](#).

If you use Migrate Agent for registration and discovery, the default direction of the replication connection must be reversed (source to target) by changing advanced settings on the Migrate server. See [Section 5.9, “Configuring the Contact Direction for the Replication Port,” on page 107](#).

9.2.3 Azure Prerequisites for Using an Azure-Based Migrate Server

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

- ♦ A PlateSpin Migrate license
- ♦ An Azure Marketplace image of the PlateSpin Migrate server in the target Azure environment

NOTE: The cloud-based Migrate server does not require a site-to-site VPN connection between your local data center and Microsoft Azure Portal. When no VPN is provided between the source network and the cloud-based Migrate server, you can use Migrate Agent to register workloads with the cloud-based Migrate server using secure communications over the public Internet. Internet access and public IP addresses are required. For deployment information, see [Figure 9-2, “Cloud-Based Migrate Server for Automated Migration to Azure,” on page 149](#).

- ♦ Specify **Static** as the allocation method for the public IP address of the Migrate server to ensure that the IP address does not change when the server is restarted.

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

You cannot specify the actual IP address assigned to the public IP resource. Azure allocates and reserves an IP address from a pool of its available IP addresses in the Azure location where you deploy the Migrate server. The address persists through server restarts. Azure releases the IP address only when you delete the resource or change the resource’s allocation method to **Dynamic**.

- ♦ The minimum network-related prerequisites for a successful migration are:
 - ♦ Open TCP port 443 in your network firewall for outbound traffic. The source workload must be able to register (using the Migrate Agent utility) and communicate with the cloud-based PlateSpin Migrate server through HTTPS (TCP/port 443). The PlateSpin Migrate Server uses secure SSL for communications with the workloads you want to migrate.
 - ♦ Open TCP port 3725 in your network firewall for outbound traffic. The on-premise source workload must be able to connect to the cloud-based target workload on TCP port 3725. The PlateSpin Migrate Server uses secure SSL for communications with the workloads you want to migrate.

The direction of the communication (source to target) is automatic, but the port number is configurable. For information about changing the default port setting, see port 3725 in [“Requirements for Migration” on page 50](#).

- ♦ Allow inbound connections for the Network Security Group (NSG) in Azure on TCP port 3725 (or configured port) for target workloads.

For information about creating and configuring a Network Security Group in Azure, see [Filter Network Traffic with Network Security Groups](https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-nsg). (<https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-nsg>).

- ♦ Allow inbound connections in the Network Security Group for HTTPS (TCP port 443) and RDP (TCP port 3389) for the cloud-based Migrate server. The `<Migrate-server-name>-nsg` is created automatically when you deploy the Migrate server in Azure.
- ♦ If you use PlateSpin Transformation Manager with the cloud-based Migrate server, allow inbound connections in the Network Security Group for STOMP communications (TCP port 61613) related to Event Messaging.

NOTE: No messages are published by Event Messaging unless you open port 61613 on the Migrate server host to allow registration by subscribers, and a PlateSpin Migrate Connector subscribes. See [Section 5.6, “Enabling Event Messaging for PlateSpin Migration Factory,” on page 105](#).

- ♦ Install the Migrate Agent on the source workload, then register the workload with the cloud-based PlateSpin Migrate server. See [“Registering Workloads and Discovering Details with Migrate Agent” on page 215](#).

To download the Migrate Agent, launch the PlateSpin Migrate Web Interface and click the **Downloads** tab. For information about installing and using the Migrate Agent, see [“Migrate Agent Utility” on page 269](#).

9.3 Planning For Migrating Workloads to Azure

PlateSpin Migrate enables you to use the PlateSpin Migrate Web Interface to migrate Windows and Linux workloads to Microsoft Azure. For a list of supported workloads, see [“Supported Workloads For Migration to Microsoft Azure” on page 27](#).

NOTE

- ♦ Migration of Azure source workloads to Azure target containers is not supported.
 - ♦ Migration of Windows Cluster workloads to Azure is not supported.
-

Target Azure IaaS Environment

- ♦ For each PlateSpin Migrate server, you must dedicate the server for migrations to the global Azure environment or to a specific sovereign Azure environment. See [Section 9.9, “Configuring the Target IaaS Environment for Azure Cloud,” on page 159](#).

NOTE: PlateSpin Migrate supports sovereign Azure environments that we have tested. Contact [Micro Focus Support and Services](#) for assistance if you want to migrate workloads to other sovereign Azure environments.

Azure Subscription

- ♦ Provide valid credentials for the Azure subscription. See [Section 9.10, “Managing the Azure User Password for Azure Target Containers,”](#) on page 160.

PlateSpin Server Host

- ♦ Ensure that the PlateSpin Server host displays the correct time for the time zone it is in. If the time on the PlateSpin Server host is incorrect, the cutover process fails with a 403 forbidden access error.

OS License for Target Workload

- ♦ You need an OS license for the migrated target workload. For Azure target workloads, you must provide Azure with the license information or Microsoft will charge you for the OS license.

Target Workload

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

- ♦ The PlateSpin Migrate Client does not support migration of workloads to Microsoft Azure. You can use only the PlateSpin Migrate Web Interface to migrate the workloads to Microsoft Azure.
- ♦ Windows and Linux UEFI workloads are migrated as BIOS workloads.
- ♦ Migration of workloads with multiple NICs to Azure is supported for Windows workloads and Linux workloads, up to the number of NICs supported by the Azure VM size.
- ♦ PlateSpin Migrate supports Azure VM sizes with up to 64 data disks. For the maximum VM size in a selected Azure Region, Migrate will use one data disk for the OS disk replication in the PlateSpin Replication Environment. After cutover, this disk becomes the OS disk, and you can add a data disk.
- ♦ Data disks can have a maximum size of 4 TB (4092 GB), depending on the maximum size allowed for the target VM size.
- ♦ The size of the disk created on the Azure VM is the size of the source disk partition plus about 1 GB because of the granularity of disk space on Azure.
- ♦ Migrate initially identifies an Azure VM size in the specified target location that meets or exceeds the source workload's settings for cores, memory, data disks, and NICs. However, you can choose a smaller or larger VM size based on your requirements for the target workload, as limited by the maximum VM sizes available in the selected Azure Region.

9.4 Azure Networking Guidelines

You can create a virtual machine with multiple NICs in Azure virtual networks. Each NIC must be located in one subnet; one subnet can be assigned to multiple NICs. Each NIC has an IP address consistent with its subnet assignment. The IP address and MAC pairing for each NIC persists, even if the order of the NICs changes.

Consider the following guidelines when you are migrating workloads to Microsoft Azure.

- ♦ [Section 9.4.1, “Private or Public IP Addresses for Azure Migration,”](#) on page 155
- ♦ [Section 9.4.2, “Windows Workloads in Azure with Multiple NICs,”](#) on page 155
- ♦ [Section 9.4.3, “Private and Public IP Addresses for Workloads Connected on an Azure VPN,”](#) on page 155

9.4.1 Private or Public IP Addresses for Azure Migration

You can use private IP addresses for workload migration if you have configured an Azure VPN to connect your premise network with your Azure cloud environment. Otherwise, you must enable a public IP address to be assigned to the replication network, cutover network, and test cutover network. If the VM has multiple NICs, only the primary NIC can have a public IP address. The assigned public IP addresses will be in the address space of the specified network and subnet for the designated NIC in each network.

NOTE: PlateSpin requires a public IP address only if an Azure VPN is not available.

If you enable a public IP address for the primary NIC, Azure assigns the NIC both a public IP address and a private IP address. For more information about connecting to the Azure VM, see “[Private and Public IP Addresses for Workloads Connected on an Azure VPN](#)” on page 155.

9.4.2 Windows Workloads in Azure with Multiple NICs

Azure configures the VM with a default gateway that is associated with the primary network interface. Azure removes the gateway information for all secondary NICs, which limits their communications to the same subnet as the primary interface.

For Windows workloads with multiple NICs, you can enable a secondary NIC to communicate outside its own subnet. Use the Windows `route add` command to add a different gateway entry for the secondary NIC in the routing table. See “Configure Windows VMs” in [Create a VM with Multiple NICs](https://azure.microsoft.com/en-us/documentation/articles/virtual-networks-multiple-nics/) (<https://azure.microsoft.com/en-us/documentation/articles/virtual-networks-multiple-nics/>) on the Microsoft Azure website (<https://azure.microsoft.com/>).

9.4.3 Private and Public IP Addresses for Workloads Connected on an Azure VPN

An Azure VM can have one or more NICs attached to it. The primary NIC for the VM can have both a public and private IP address. A private IP address is used for communications from other resources in a virtual network and from machines inside the address space for the Azure VPN that connects your premise network to your Azure cloud environment. A public IP address can be used to communicate with the Internet and with machines outside the Azure cloud environment. Azure automatically associates these IP addresses with the primary network interface for the VM.

You can use the Microsoft Remote Desktop client to connect remotely to the Azure VM. Specify the IP address as follows:

- ♦ **Private IP address:** Use the VM's private IP address if your machine is part of the address space for the Azure VPN.
- ♦ **Public IP address:** Use the VM's public IP address if your machine is not part of the address space for the Azure VPN.

You can alternatively use the **Connect** option in the [Microsoft Azure portal](https://azure.microsoft.com/en-us/features/azure-portal/) (<https://azure.microsoft.com/en-us/features/azure-portal/>) from a machine with an address space that is not part of the Azure VPN. This option automatically launches the Microsoft Remote Desktop client configured to connect to the VM's public IP address for the primary NIC.

NOTE: This portal operation fails if your machine is in the address space of the Azure VPN.

9.5 Registering a Azure Application to Represent PlateSpin Migrate

PlateSpin Migrate uses the Microsoft Azure API to automate workload migrations to Azure. You need to create an Azure application ID for PlateSpin Migrate to use when it uses the Azure API for replicating and migrating workloads to your Azure account.

To register PlateSpin Migrate as an application in Azure:

- 1 Go to the appropriate Azure Portal and log in to your Azure account:
 - ♦ [Azure Portal \(http://www.azure.com/\)](http://www.azure.com/)
 - ♦ [Azure China Portal \(http://www.azure.cn/\)](http://www.azure.cn/)
- 2 In the left column of the Portal menu, click **Azure Active Directory**.
- 3 In the directory menu under **Manage**, select **App registrations**, then click **Add** to open the Create pane.
- 4 In the Create pane, configure the settings for the application:
 - 4a Specify friendly name for the application, such as `PlateSpin Migrate`
The name must be unique in your Azure Active Directory. This is the name that appears in the Applications list.
 - 4b Select **Native** as the Application Type.
 - 4c Specify a valid URL as the **Redirect URI**.
The redirect URI is not used in practice, so you can specify any valid URL that you control.
 - 4d Click **Create**.
- 5 In the Applications list, select the application, then click **Settings** to view the Essentials information, including the Application ID.
- 6 Copy the Application ID value to the clipboard and paste it in a text document where you can access it when you set up target containers for this account.
An application ID is a value in the format of: abc12b34-c5df-6e78-f9a0-bc123456d789.
- 7 Configure permissions for the registered application.
 - 7a At the bottom right of the Settings pane, click **All Settings**.
 - 7b In the Settings menu, under **API Access**, select **Required Permissions**.
 - 7c On the Grant Permissions pane, click **Add**.
 - 7d In the Add Permissions pane, click **Select an API**.
 - 7e In the right pane, select **Windows Azure Service Management API**, then click **Select** at the bottom of the pane.
A green check mark appears next to **Select an API**.
 - 7f In the Add Permissions pane, click **Select Permissions**.
 - 7g In the right pane, select the check box next to **Access Azure Service Management as organization users**, then click **Select** at the bottom of the pane.
A green check mark appears next to **Select Permissions**.
 - 7h At the bottom of the Add Permissions pane, click **Done**.

- 8 [This step must be performed by an Azure global administrator account.] Using an Azure global administrator account, enable the Default Directory
 - 8a In the Portal menu, select **Azure Active Directory**, then click **Enterprise Applications**.
 - 8b Click the new application that you created in [Step 4](#).
 - 8c Under **Security**, click **Permissions**.

Initially, there are no permissions listed in the **Admin Consent** section for the application.
 - 8d Click **Grant admin consent for Default Directory**.
 - 8e A separate browser window opens and prompts you to sign in to administer the application. Sign in using an Azure global administrator account that has permissions to grant admin consent for the application.
 - 8f After authentication succeeds, the **Permissions requested - Accept for your organization** window prompts you to consent to the application permissions. Click **Accept**, then wait for the browser to refresh its content.
 - 8g After the permissions are successfully granted, close the browser window.
- 9 Verify the setup.
 - 9a In Portal menu, select **Azure Active Directory**, then click **Enterprise Applications**.
 - 9b Click the new application that you created in [Step 4](#).
 - 9c Under **Security**, click **Permissions**.
 - 9d Verify that there are two new Permissions listed in the **Admin Consent** section.

9.6 Configuring an Azure User for PlateSpin Migrate to Use

PlateSpin Migrate requires an Azure Active Directory user identity that it can use when it accesses the Migrate Azure API through a specified subscription.

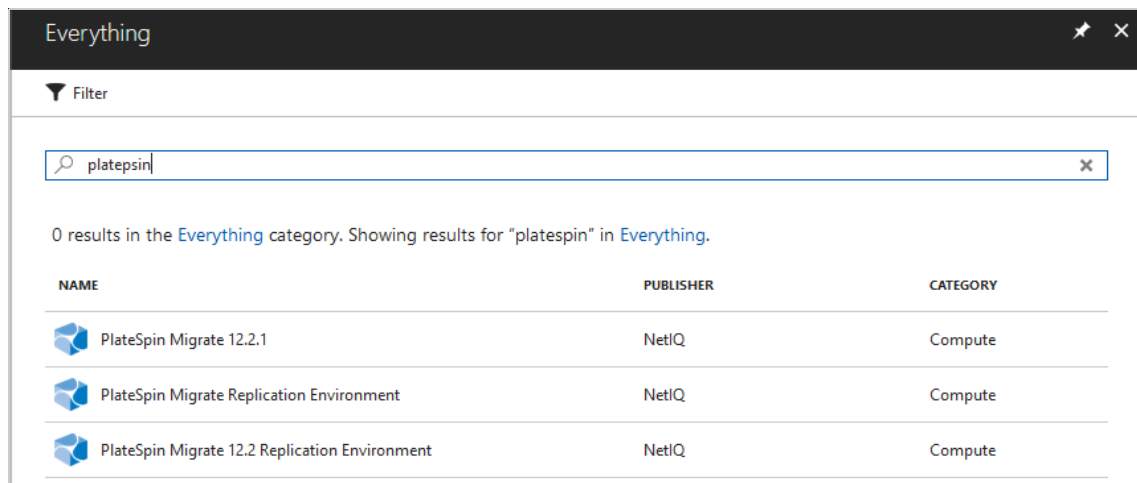
9.7 Enabling Programmatic Deployment of PlateSpin Migrate Replication Environment

PlateSpin Migrate must be able to programmatically deploy a PlateSpin Migrate Replication Environment VM during the replication of workloads to Azure. The required VM image is available in the Azure Marketplace. You must enable programmatic deployment of the image for each subscription that will perform migrations with PlateSpin Migrate.

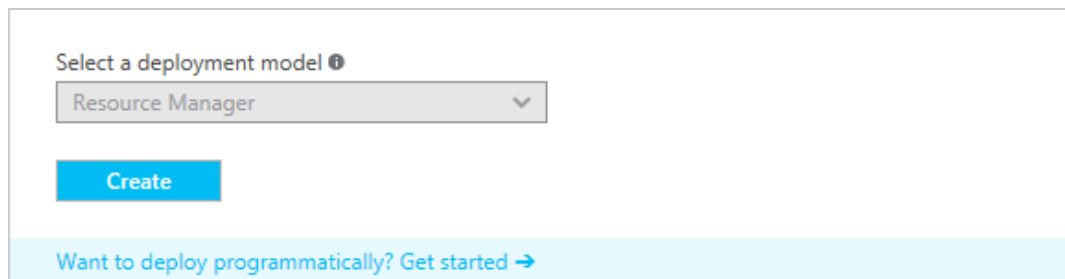
NOTE: All migrations for the target subscription will fail until you enable the subscription to deploy the Replication Environment VM.

To enable programmatic deployment of the Replication Environment VM:

- 1 Go to the appropriate Azure Portal and log in to your Azure account:
 - ♦ [Azure Portal \(http://portal.azure.com/\)](http://portal.azure.com/)
 - ♦ [Azure China Portal \(http://portal.azure.cn/\)](http://portal.azure.cn/)
- 2 In the portal menu, click **New**, then search for PlateSpin images in the Azure Marketplace. Type **platespin** in the Everything filter.
- 3 In the Results panel, select **PlateSpin Migrate 12.x Replication Environment**.



- At the bottom of the PlateSpin Migrate 12.x Replication Environment page under **Select a deployment model**, click **Want to deploy programmatically? Get Started**.



- On the Configure Programmatic Deployment page, read the Terms of Use.
- For each subscription that will perform migrations with PlateSpin Migrate, change its status from **Disable** to **Enable**.
- Click **Save**.

9.8 Deploying a Migrate Server Image in Azure

PlateSpin Migrate offers a PlateSpin Migrate Server image in Azure through the Azure Marketplace in each of the supported Azure environments.

NOTE: For PlateSpin Migrate Servers deployed through the Azure Marketplace, Migrate installs Microsoft SQL Server 2014 Express Edition for the PlateSpin database and generates a password for the SQL administrator (sa) user. For improved security, we recommend that you modify the password for this user after the deployment is complete.

To enable the deployment of the Migrate Server in Azure and assign relevant permissions, do the following:

- Go to the appropriate Azure Portal and log in to your Azure account:
 - [Azure Portal \(http://www.azure.com/\)](http://www.azure.com/)
 - [Azure China Portal \(http://www.azure.cn/\)](http://www.azure.cn/)
- In the portal menu, Select **Marketplace**, then click **New**.

- 3 Under **New**, search for PlateSpin.
- 4 On the PlateSpin Migrate Server page, view the deployment options and related prices.
- 5 Select the appropriate deployment model, then click **Create**.
- 6 Configure **Basics** and **Size** settings.
- 7 In **Settings**, specify **Static** as the assignment method for the public IP address:
 - 7a In the **Settings** panel under **Network**, select **Public IP address**.
 - 7b In the **Choose public IP address** panel, select **Create New**.
 - 7c In the **Create public IP address** panel, specify `migrate-server-ip` as the name and **Static** as the assignment method.
 - 7d Click **OK**.

Azure allocates and reserves an IP address from a pool of its available IP addresses in the Azure location where you deploy the Migrate server.
- 8 Review settings and approve the purchase.
- 9 (Optional, recommended) For improved security, modify the password for the SQL Administrator (`sa`) user on the deployed Migrate server in Azure. For information, contact [Micro Focus Technical Support](#).

9.9 Configuring the Target IaaS Environment for Azure Cloud

Microsoft Azure Cloud provides global and sovereign IaaS environments where you can set up virtual machines. Migrate supports the following target Azure Cloud environments:

- ♦ Global (default)
- ♦ Azure China

NOTE: Contact [Micro Focus Support and Services](#) for assistance if you want to migrate workloads to other sovereign Azure environments.

Only one Azure Cloud environment can be used for target containers on an individual PlateSpin Migrate server. You cannot concurrently configure target locations for global regions and sovereign regions on the same server.

- ♦ **On premise:** PlateSpin Migrate Server is preconfigured to support Azure global. You can modify server settings to work with sovereign Azure China.
- ♦ **Azure Marketplace:** The PlateSpin Migrate Server image is preconfigured to support its host IaaS environment: Azure global or sovereign Azure China.

For each on-premise Migrate server, you can configure which target Azure Cloud environment to use for migrations by using advanced settings on the PlateSpin Configuration web page on your Migrate server host. [Table 9-2](#) and [Table 9-3](#) describe the values to use for the appropriate environment.

Table 9-2 Azure Global Environment Configuration Settings

PlateSpin Configuration Parameters	Values
AzureAuthenticationAddress	https://login.windows.net
AzureCloudAddress	https://management.azure.com/
AzureStorageEndpointSuffix	core.windows.net

Table 9-3 Azure China Environment Configuration Settings

PlateSpin Configuration Parameters	Values
AzureAuthenticationAddress	https://login.chinacloudapi.cn
AzureCloudAddress	https://management.chinacloudapi.cn/
AzureStorageEndpointSuffix	core.chinacloudapi.cn

To set the target Azure Cloud environment for a Migrate Server:

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

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

- 2 Search for and modify values for three Azure settings based on the appropriate Azure Cloud environment:
 - ♦ Azure Global environment. See [Table 9-2, “Azure Global Environment Configuration Settings,” on page 160.](#)
 - ♦ Azure China environment. See [Table 9-3, “Azure China Environment Configuration Settings,” on page 160.](#)

9.10 Managing the Azure User Password for Azure Target Containers

Provide a valid password for the Microsoft Azure user when you add the Azure target container. Ensure that you update the password for the container in PlateSpin Migrate if you modify it in Azure.

Workload migrations can fail in the following conditions:

- ♦ **Invalid password:** If the stored password for the Azure user is invalid, an authentication error occurs when the next connection to Azure is requested.

If the Azure user modifies the password in the Microsoft Azure portal while migration tasks are running, the tasks will fail with an authentication error when the next connection to Azure is requested.

- ♦ **Expired password:** If the stored password for the Azure user expires in Microsoft Azure, a Password Is Expired error occurs when the next connection to Azure is requested.

If the password expires while migration tasks are running, the tasks will fail with a Password Is Expired error when the next connection to Azure is requested.

To resolve failed migrations to Azure for password issues:

- 1 (Conditional) If the Azure user's password expired, log in to the user account in the Microsoft Azure portal, then set a new user password by using the [Azure Self-Service Password Reset](https://azure.microsoft.com/en-us/documentation/articles/active-directory-passwords-getting-started/#step-3-reset-your-azure-ad-password-as-a-user) (<https://azure.microsoft.com/en-us/documentation/articles/active-directory-passwords-getting-started/#step-3-reset-your-azure-ad-password-as-a-user>).
- 2 Log in to the PlateSpin Migrate Web Interface, then go to the Target Containers page.
- 3 Update the password that is stored for the Azure user for any affected Azure target containers.
 - 3a Click the name of the target container to access the container settings, then click **Edit**.
 - 3b Specify a valid password.
 - 3c (Optional) Click **Test Credentials**.
 - 3d Click **Save**.
- 4 Rerun any failed workload migrations to the affected Azure target containers.

9.11 Checklist for Automated Migration to Azure

Task	Description
1. Prepare your VMware migration environment.	Figure 9-1, "On-Premise Migrate Server for Automated Migration to Azure," on page 148 Figure 9-2, "Cloud-Based Migrate Server for Automated Migration to Azure," on page 149 Section 9.3, "Planning For Migrating Workloads to Azure," on page 153
2. Discover target containers.	Section 19.4.2, "Target Discovery in the Web Interface," on page 202
3. Discover source workloads.	Section 20.6.2, "Workload Discovery in the Migrate Web Interface," on page 215 -OR- Section 20.7, "Registering Workloads and Discovering Details with Migrate Agent," on page 215
4. Configure target workload migration.	Section 28.2, "Configuring Migration of a Workload to Microsoft Azure," on page 336
5. Execute migration.	Chapter 38, "Executing Workload Migrations," on page 433

10 Prerequisites for Migration to VMware vCloud Director

PlateSpin Migrate Web Interface supports automated migration to VMware vCloud Director environments. This section describes the required vCloud configuration that you must prepare in the appropriate environment, such as a vCloud Organization, before you can discover vCloud target containers and configure migrations to them.

- [Section 10.1, “Deployment for Migration to VMware vCloud,” on page 163](#)
- [Section 10.2, “Planning For Migrating Workloads to VMware vCloud Director,” on page 165](#)
- [Section 10.3, “Setting up vCloud Organization,” on page 165](#)
- [Section 10.4, “Understanding PlateSpin Replication Environment,” on page 166](#)
- [Section 10.5, “Configuring Advanced PlateSpin Settings for vCloud,” on page 167](#)
- [Section 10.6, “Checklist for Automated Migration to vCloud,” on page 168](#)

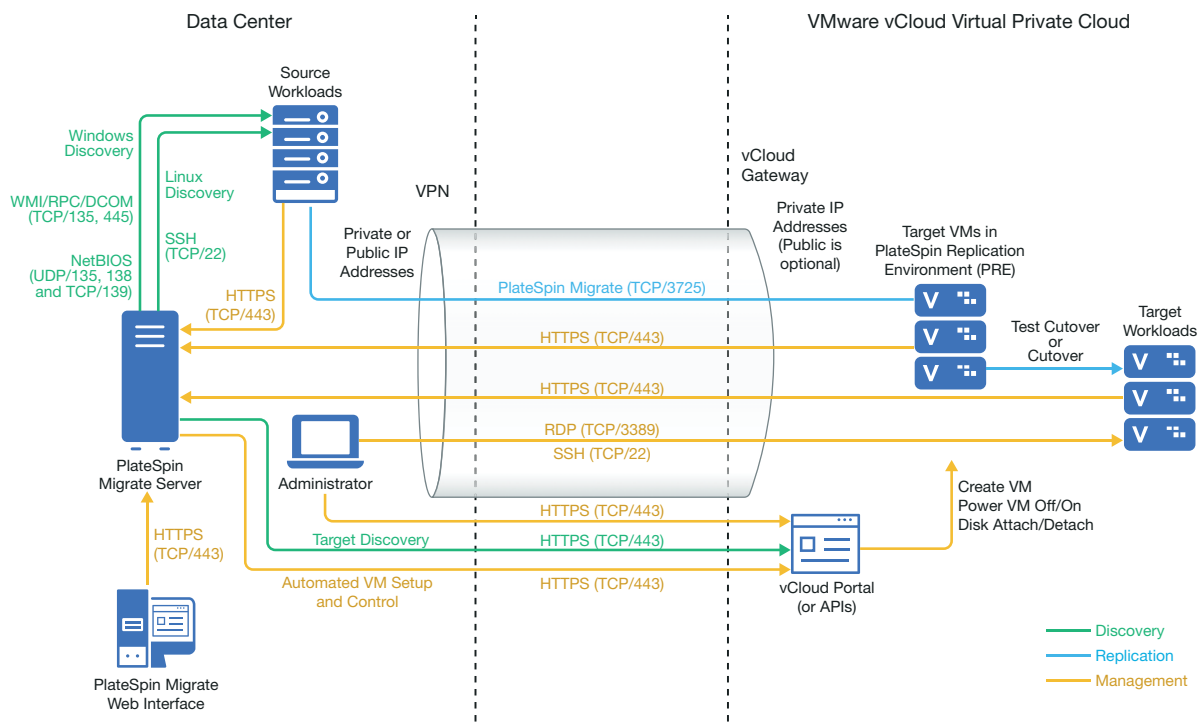
10.1 Deployment for Migration to VMware vCloud

You can deploy a PlateSpin Migrate server on premise in your data center with the source workloads or in the appropriate VMware vCloud Organization.

For an on-premise Migrate server deployment, a site-to-site VPN connection is required between the data center and your account in the vCloud cloud. [Figure 10-1](#) shows the location of various components in your vCloud migration environment and the communications between them. See [Section 10.2, “Planning For Migrating Workloads to VMware vCloud Director,” on page 165](#).

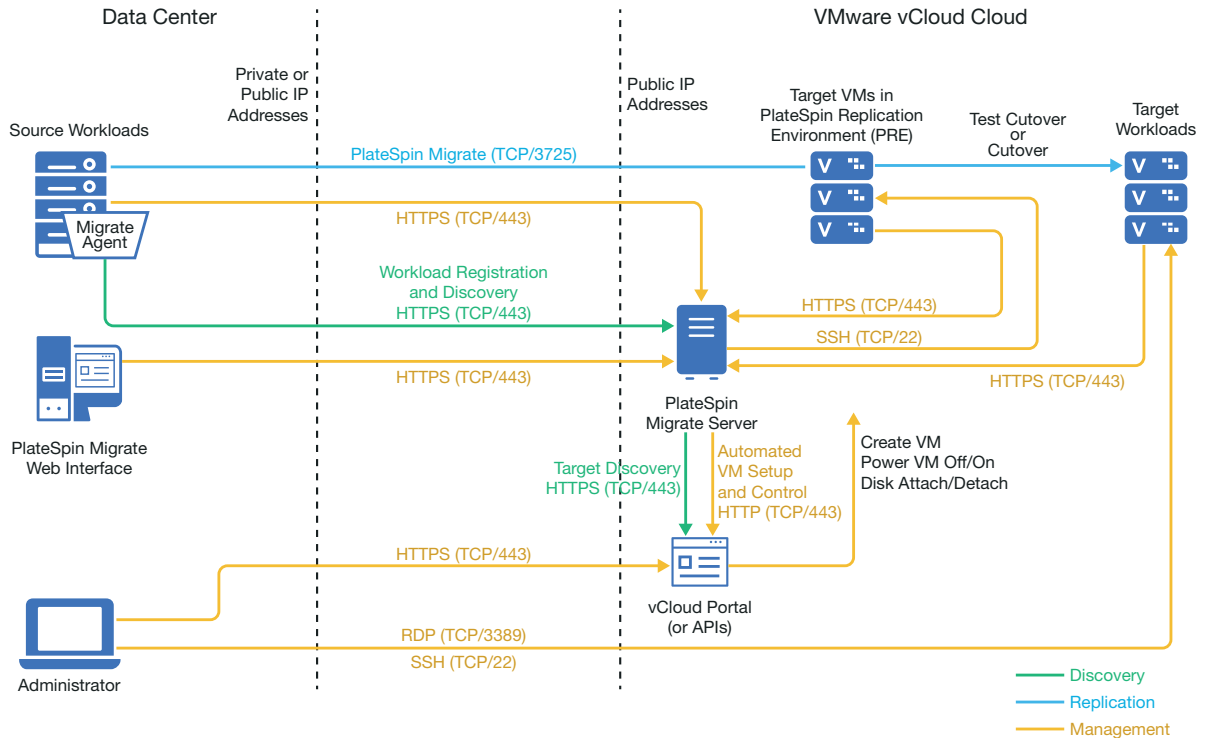
NOTE: [Figure 10-1](#) depicts automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443).

Figure 10-1 On-Premise Migrate Server for Automated Migration to vCloud



For a cloud-based Migrate server deployment, the PlateSpin Migrate Server is available. [Figure 10-2](#) shows the location of various components in your vCloud migration environment and the communications between them. See [Section 10.2, “Planning For Migrating Workloads to VMware vCloud Director,”](#) on page 165.

Figure 10-2 Cloud-Based Migrate Server for Automated Migration to vCloud



10.2 Planning For Migrating Workloads to VMware vCloud Director

PlateSpin Migrate uses the VMware vCloud Director for migrating workloads to VMware vCloud. For a list of supported workloads, see [“Supported Workloads For Migration to VMware vCloud Director”](#) on page 29.

10.3 Setting up vCloud Organization

You must set up a vCloud Organization with at least the following minimum set of resources:

- ☐ Define one or more Organization Virtual Data Center (Org vDC).
- ☐ Define one or more Org vDC Network for the target VM.
- ☐ Create a private Catalog and grant full access permissions for the organization users to access the contents and settings of the catalog.
- ☐ Use Administrator level credentials for discovering the vCloud Organization and performing migrations.
- ☐ Define policies that apply to the target VMs in your Org vDC and ensure the following:
 - ☐ The lease period of the vCloud Organization resources should not expire during migration.
 - ☐ No restrictions are set on VM quota.
 - ☐ No restrictions are set on the number of connections or operations to the vCloud organization.

NOTE: During the Test Cutover, there is cloning for the target VM that consumes twice the storage resources than you need for Cutover. Ensure that the storage quotas for the Org vDC support that need. The additional resources used is temporary, and will be released after the Test Cutover.

For more information, see [VMware vCloud Director Documentation \(https://www.vmware.com/support/pubs/vcd_pubs.html\)](https://www.vmware.com/support/pubs/vcd_pubs.html).

10.4 Understanding PlateSpin Replication Environment

PlateSpin requires a replication environment to migrate workloads to a vCloud Organization. The replication environment is a virtual appliance based on a SLES 11 Operating System and contains all the required PlateSpin tools.

To create this virtual appliance, you must upload a PlateSpin Replication Environment OVF Package to the vCloud organization. You can download this package from the [NetIQ Download \(https://dl.netiq.com\)](https://dl.netiq.com) site for this PlateSpin Migrate release. For more information, see “[Creating the PlateSpin Virtual Appliance in the vCloud Organization](#)” on page 166.

Review the following sections:

- ♦ [Section 10.4.1, “Resources Used in the PlateSpin Replication Environment,” on page 166](#)
- ♦ [Section 10.4.2, “Creating the PlateSpin Virtual Appliance in the vCloud Organization,” on page 166](#)

10.4.1 Resources Used in the PlateSpin Replication Environment

PlateSpin uses the following minimum resources for the Replication Environment Virtual Machine:

Hardware Resource	Details
Virtual CPUs	1
Cores Per Socket	1
RAM	1 GB
Disk	4 GB
Network Adapter of type E1000	1
Virtual Hardware Version	7

10.4.2 Creating the PlateSpin Virtual Appliance in the vCloud Organization

- 1 Ensure that you have set up a vCloud Organization with at least the minimum set of resources. See “[Setting up vCloud Organization](#)” on page 165.
- 2 Download the PlateSpin Replication Environment file named `PlateSpin_Replication_Environment.zip` from the [NetIQ Download site \(https://dl.netiq.com\)](https://dl.netiq.com) for this PlateSpin Migrate release.

- 3 Unzip the `PlateSpin_Replication_Environment.zip` file and extract the contents to a temporary directory. For example, `C:\PlateSpin_Replication_Environment`.
- 4 Use the vCloud Director Web Console to upload the OVF PlateSpin package, which you extracted in the previous step, as a vApp Template to the `PlateSpin` Catalog.

The vCloud Director Web console displays the replication environment in the following hierarchy:

- ♦ *Catalog Name* Catalog (Catalog)
 - ♦ PlateSpin Replication Environment (vApp Template)
 - ♦ PlateSpin Virtual Appliance (VM Template)

10.5 Configuring Advanced PlateSpin Settings for vCloud

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

- ♦ [Section 10.5.1, “Retaining the Cloud Resources For Troubleshooting Migration Errors,” on page 167](#)
- ♦ [Section 10.5.2, “Setting the PlateSpin Replication Environment Password in Clear Text,” on page 167](#)

10.5.1 Retaining the Cloud Resources For Troubleshooting Migration Errors

By default, PlateSpin deletes the target VM and its associated resources when there is an error during migration. If you need these resources for troubleshooting and do not want to delete them, set the `LeaveCloudResourcesOnError` setting to `True`.

10.5.2 Setting the PlateSpin Replication Environment Password in Clear Text

By default, the password required to access the PlateSpin Replication Environment is encrypted. To access the PlateSpin Replication Environment for troubleshooting replication failures, set a password to override its default value. To set a password, edit the value of the `vCloudReplicationEnvironmentPassword` setting. You can then access the PlateSpin Replication Environment as a `root` user with the newly set password.

10.6 Checklist for Automated Migration to vCloud

Task	Description
1. Prepare your VMware migration environment.	Figure 10-1, “On-Premise Migrate Server for Automated Migration to vCloud,” on page 164 Figure 10-2, “Cloud-Based Migrate Server for Automated Migration to vCloud,” on page 165 Section 10.2, “Planning For Migrating Workloads to VMware vCloud Director,” on page 165
2. Discover target containers.	Section 19.4.2, “Target Discovery in the Web Interface,” on page 202
3. Discover source workloads.	Section 20.6.2, “Workload Discovery in the Migrate Web Interface,” on page 215 -OR- Section 20.7, “Registering Workloads and Discovering Details with Migrate Agent,” on page 215
4. Configure target workload migration.	Section 29.2, “Configuring Migration of a Workload to VMware vCloud Director,” on page 349
5. Execute migration.	Chapter 38, “Executing Workload Migrations,” on page 433

11

Prerequisites for Migration to VMware

PlateSpin Migrate supports automated or semi-automated migration to your VMware environment. This section describes the required VMware configuration that you must prepare before you can discover VMware target containers (for automated migration) or target VMs (for semi-automated migrations) and configure migrations to them.

- ♦ [Section 11.1, “Deployment for Migration to VMware,” on page 169](#)
- ♦ [Section 11.2, “Planning for Migration to VMware,” on page 171](#)
- ♦ [Section 11.3, “Checklist for Automated Migration to VMware,” on page 172](#)
- ♦ [Section 11.4, “Checklist for Semi-Automated Migration to Target VMs on VMware,” on page 172](#)

11.1 Deployment for Migration to VMware

[Figure 11-1](#) shows the location of various components in your automated VMware migration environment and the communications between them. Automated migration to VMware target containers is supported by PlateSpin Migrate Client and PlateSpin Migrate Web Interface.

NOTE: [Figure 11-1](#) and [Figure 11-2](#) depict automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443).

Figure 11-1 Automated Migration to VMware

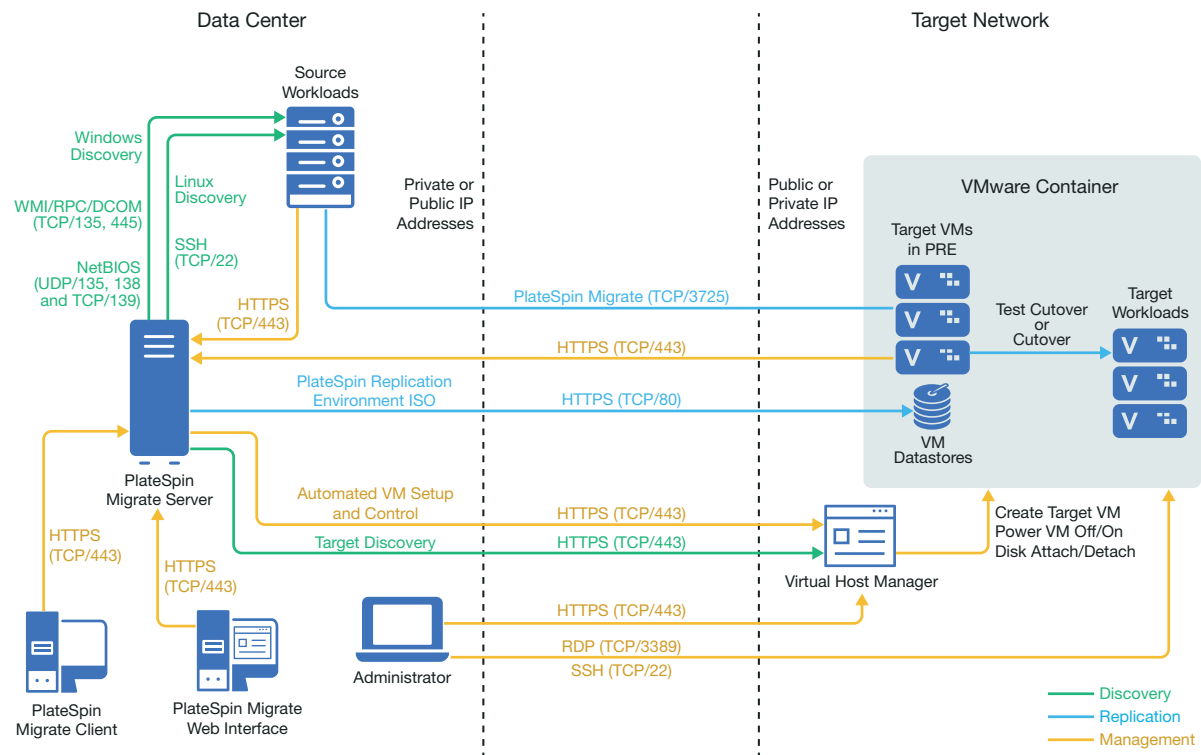



Figure 11-2 shows the location of various components in your semi-automated VMware migration environment and the communications between them. Semi-automated migration to target VMs on VMware is supported by PlateSpin Migrate Client.

Figure 11-2 Semi-Automated Migration to VMs on VMware



11.2 Planning for Migration to VMware

Ensure that your VMware environment meets the following prerequisites for migration to VMware:

- ♦ Use PlateSpin Migrate Client or PlateSpin Migrate Web Interface to migrate workloads to VMware.

See [Table 2-12, “Supported Target VMware Platforms for the Migrate Web Interface and Migrate Client,”](#) on page 38.

- ♦ Your source workload must be supported by PlateSpin Migrate and VMware.

See [Section 2.1.1, “Supported Source Workloads For Migration to Non-Cloud Platforms,”](#) on page 23.

- ♦ Your network environment must meet the requirements for access, discovery, and migration described in [“Access and Communication Requirements across your Migration Network”](#) on page 49.

- ♦ For semi-automated migrations using Migrate Client, ensure that you configure volumes on the target disks with about 50 MB of additional storage space than the source disks.

For information about configuring the migration, see [“Migration to VMware”](#) on page 359.

11.3 Checklist for Automated Migration to VMware

Task	Description
1. Prepare your VMware migration environment.	Figure 11-1, "Automated Migration to VMware," on page 170. Section 11.2, "Planning for Migration to VMware," on page 171
2. Discover target containers.	Section 19.4, "Discovering Details for Target Containers," on page 200
3. Discover source workloads.	Section 20.6.1, "Workload Discovery in the Migrate Client," on page 214 -OR- Section 20.6.2, "Workload Discovery in the Migrate Web Interface," on page 215 -OR- Section 20.7, "Registering Workloads and Discovering Details with Migrate Agent," on page 215
4. Configure target workload migration.	Section 30.2, "Automated Migration to VMware Using Migrate Client," on page 360 -OR- Section 30.4, "Automated Migration to VMware Using Migrate Web Interface," on page 374
5. Execute migration.	Chapter 38, "Executing Workload Migrations," on page 433

11.4 Checklist for Semi-Automated Migration to Target VMs on VMware

Task	Description
1. Prepare your VMware migration environment.	Figure 11-2, "Semi-Automated Migration to VMs on VMware," on page 171 Section 11.2, "Planning for Migration to VMware," on page 171
2. Discover target containers.	Section 19.5, "Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO," on page 204
3. Discover source workloads.	Section 20.6.1, "Workload Discovery in the Migrate Client," on page 214
4. Configure target workload migration.	Section 30.3, "Migration to VMs on VMware Using X2P Workflow," on page 371

Task	Description
5. Execute migration.	Chapter 38, "Executing Workload Migrations," on page 433

12 Prerequisites for Migration to Microsoft Hyper-V

PlateSpin Migrate supports automated or semi-automated migration to your Microsoft Hyper-V environment. This section describes the required Hyper-V configuration that you must prepare before you can discover Hyper-V target containers (for automated migration) or target VMs (for semi-automated migrations) and configure migrations to them.

- ♦ [Section 12.1, “Deployment for Migration to Microsoft Hyper-V,” on page 175](#)
- ♦ [Section 12.2, “Planning for Migration to Microsoft Hyper-V,” on page 177](#)
- ♦ [Section 12.3, “Checklist for Automated Migration to Hyper-V,” on page 178](#)
- ♦ [Section 12.4, “Checklist for Semi-Automated Migration to Target VMs on Hyper-V,” on page 178](#)

12.1 Deployment for Migration to Microsoft Hyper-V

[Figure 12-1](#) shows the location of various components in your automated Hyper-V migration environment and the communications between them.

NOTE: [Figure 12-1](#) and [Figure 12-2](#) depict automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443).

Figure 12-1 Automated Migration to Hyper-V

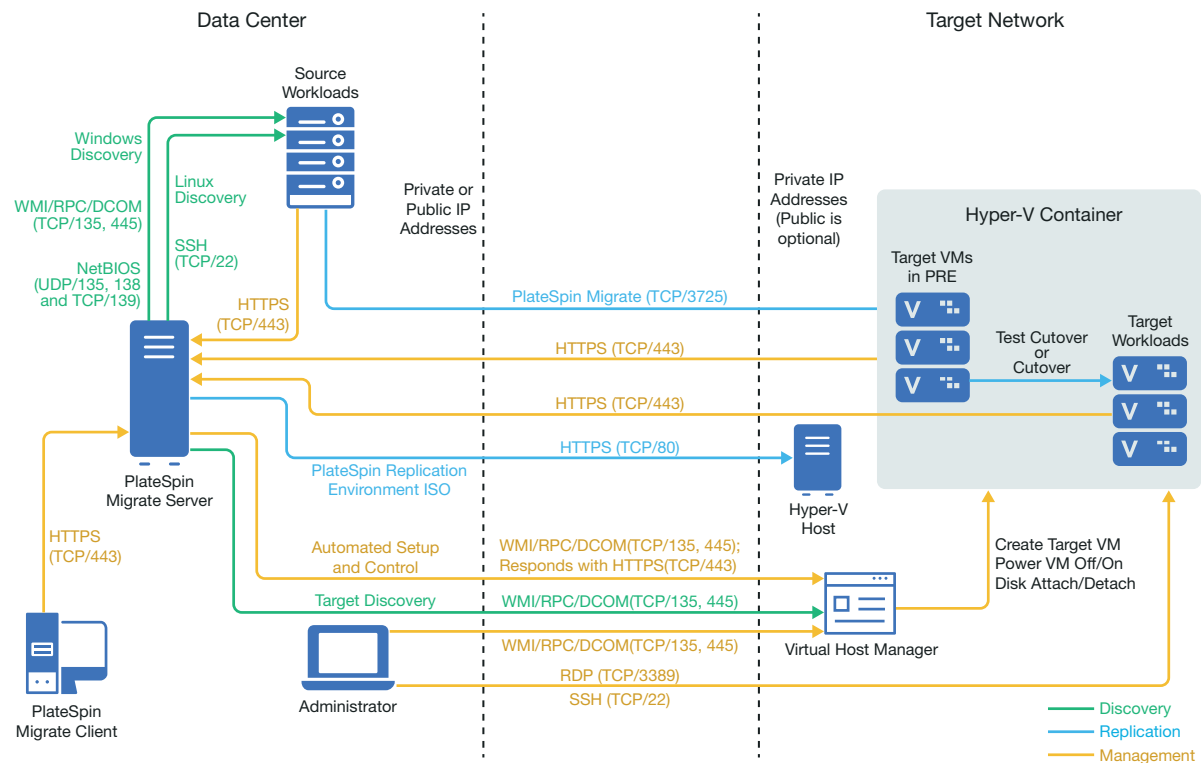
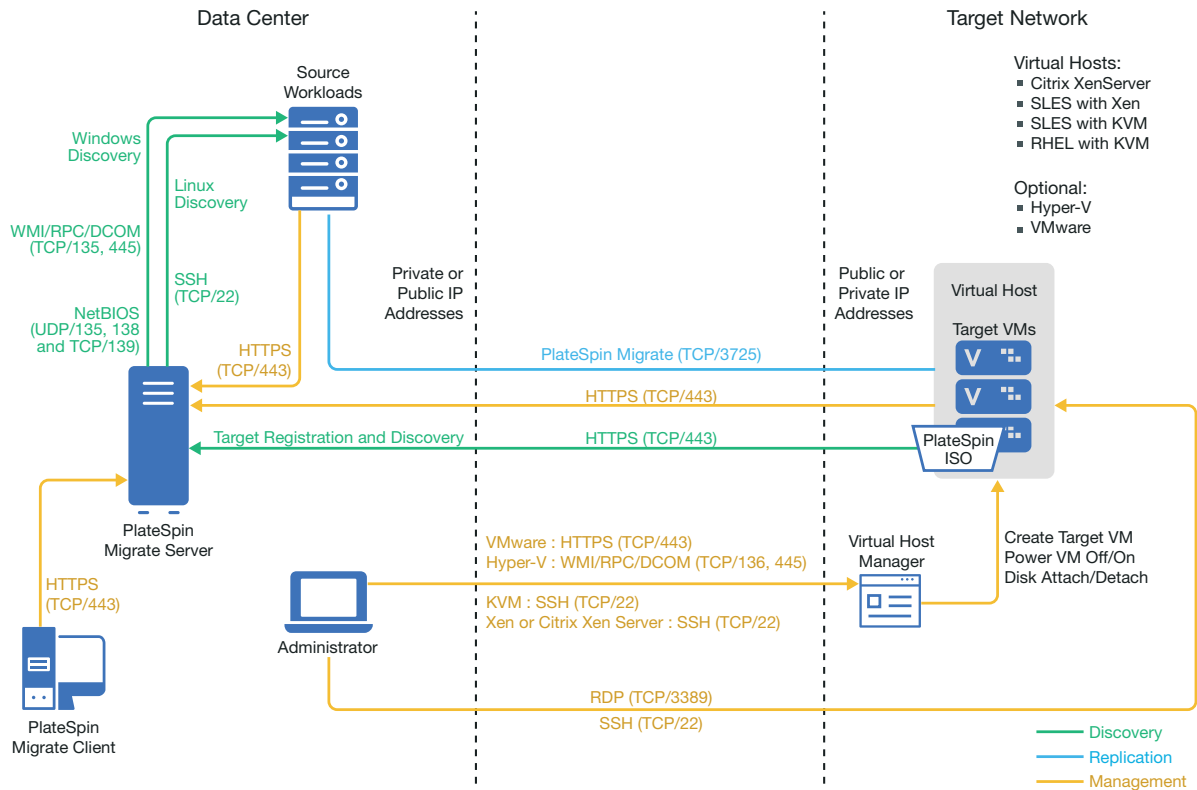


Figure 12-2 shows the location of various components in your semi-automated Hyper-V migration environment and the communications between them.

Figure 12-2 Semi-Automated Migration to VMs on Hyper-V



12.2 Planning for Migration to Microsoft Hyper-V

Ensure that your Microsoft Hyper-V environment meets the following prerequisites for migration to Hyper-V:

- Use PlateSpin Migrate Client to migrate workloads to Microsoft Hyper-V virtual hosts. PlateSpin Migrate Web Interface does not support migration to Hyper-V virtual hosts.
- You can use Hyper-V as the target virtualization platform in fully automated workload virtualization. You can use VMs in Hyper-V as targets for semi-automated (X2P) migrations.
- Your source workload must be supported by PlateSpin Migrate and Hyper-V.
See “Microsoft Windows Server with Hyper-V” in Table 2-14, “Supported Target Virtualization Platforms for the Migrate Client Only,” on page 39.
- For semi-automated (X2P) migrations to VMs on Hyper-V, see also Chapter 25, “Prerequisites for Semi-Automated (X2P) Migrations,” on page 289.
- Your network environment must meet the requirements for access, discovery, and migration described in “Access and Communication Requirements across your Migration Network” on page 49.
- For Hyper-V target VMs with synthetic adapter, you cannot set an MTU value that is less than 1500.
- For semi-automated migrations in the Migrate Client, ensure that you configure volumes on the target disks with about 50 MB of additional storage space than the source disks.

For information about configuring the migration, see “Migration to Microsoft Hyper-V” on page 383.

12.3 Checklist for Automated Migration to Hyper-V

Task	Description
1. Prepare your VMware migration environment.	Figure 12-1, “Automated Migration to Hyper-V,” on page 176. Section 12.2, “Planning for Migration to Microsoft Hyper-V,” on page 177
2. Discover target containers.	Section 19.4, “Discovering Details for Target Containers,” on page 200
3. Discover source workloads.	Section 20.6.1, “Workload Discovery in the Migrate Client,” on page 214 -OR- Section 20.7, “Registering Workloads and Discovering Details with Migrate Agent,” on page 215
4. Configure target workload migration.	Section 31.2, “Automated Migration to Hyper-V,” on page 384
5. Execute migration.	Chapter 38, “Executing Workload Migrations,” on page 433

12.4 Checklist for Semi-Automated Migration to Target VMs on Hyper-V

Task	Description
1. Prepare your VMware migration environment.	Figure 12-2, “Semi-Automated Migration to VMs on Hyper-V,” on page 177 Section 12.2, “Planning for Migration to Microsoft Hyper-V,” on page 177
2. Discover target containers.	Section 19.5, “Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO,” on page 204
3. Discover source workloads.	Section 20.6.1, “Workload Discovery in the Migrate Client,” on page 214
4. Configure target workload migration.	Section 31.3, “Migration to VMs on Hyper-V Using X2P Workflow,” on page 393
5. Execute migration.	Chapter 38, “Executing Workload Migrations,” on page 433

13 Prerequisites for Migration to VMs on Citrix XenServer

PlateSpin Migrate supports semi-automated migration to target VMs on your Citrix XenServer virtual host environment. This section describes the required XenServer configuration that you must prepare before you can discover target VMs and configure migrations to them.


- ♦ [Section 13.1, “Deployment for Migration to Citrix XenServer,” on page 179](#)
- ♦ [Section 13.2, “Planning for Migration to VMs on Citrix XenServer,” on page 180](#)
- ♦ [Section 13.3, “Checklist for Semi-Automated Migration to Target VMs on Citrix XenServer,” on page 181](#)

13.1 Deployment for Migration to Citrix XenServer

[Figure 13-1](#) shows the location of various components in your semi-automated Citrix XenServer migration environment and the communications between them.

NOTE: [Figure 13-1](#) depicts automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443).

Figure 13-1 Semi-Automated Migration to VMs on Citrix XenServer



13.2 Planning for Migration to VMs on Citrix XenServer

Ensure that your Citrix XenServer environment meets the following prerequisites for migration to VMs on Citrix XenServer:

- Use PlateSpin Migrate Client to migrate workloads to virtual machine on Citrix XenServer virtual hosts. PlateSpin Migrate Web Interface does not support migration to XenServer virtual hosts.
- You can use Citrix XenServer as the target virtualization platform in a semi-automated workload migration.
- Your target must be a fully virtualized (not paravirtualized) VM.
- Ensure that the LRD operating system is compatible with the target virtual host. Use the SUSE Linux Enterprise Server (SLES) 11 SP3 version of the PlateSpin Boot OFX ISO file (bootofx.x2p.sles11sp3.iso) for target VMs on Citrix XenServer 6.5 and earlier releases that do not support SLES 11 SP4. See [“PlateSpin ISO Image” on page 279](#).
- Your source workload must be supported by PlateSpin Migrate and Citrix XenServer.
See [“Citrix XenServer” in Table 2-14, “Supported Target Virtualization Platforms for the Migrate Client Only,” on page 39](#).
- Your network environment must meet the requirements for access, discovery, and migration described in [“Access and Communication Requirements across your Migration Network” on page 49](#).
- Configure volumes on the target disks with about 50 MB of additional storage space than the source disks.

For information about configuring semi-automated migration to a virtual machine on XenServer, see [“Migration to Virtual Machines on Citrix XenServer” on page 395](#).

13.3 Checklist for Semi-Automated Migration to Target VMs on Citrix XenServer

Task	Description
1. Prepare your VMware migration environment.	Figure 13-1, “Semi-Automated Migration to VMs on Citrix XenServer,” on page 180 Section 13.2, “Planning for Migration to VMs on Citrix XenServer,” on page 180
2. Discover target containers.	Section 19.5, “Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO,” on page 204
3. Discover source workloads.	Section 20.6.1, “Workload Discovery in the Migrate Client,” on page 214
4. Configure target workload migration.	Section 32.2, “Configuring Migration to a VM on a Citrix XenServer Virtual Host,” on page 396
5. Execute migration.	Chapter 38, “Executing Workload Migrations,” on page 433

14 Prerequisites for Migration to VMs on Xen

PlateSpin Migrate supports semi-automated migration to target VMs on your Xen virtual host environment. This section describes the required Xen configuration that you must prepare before you can discover target VMs and configure migrations to them.

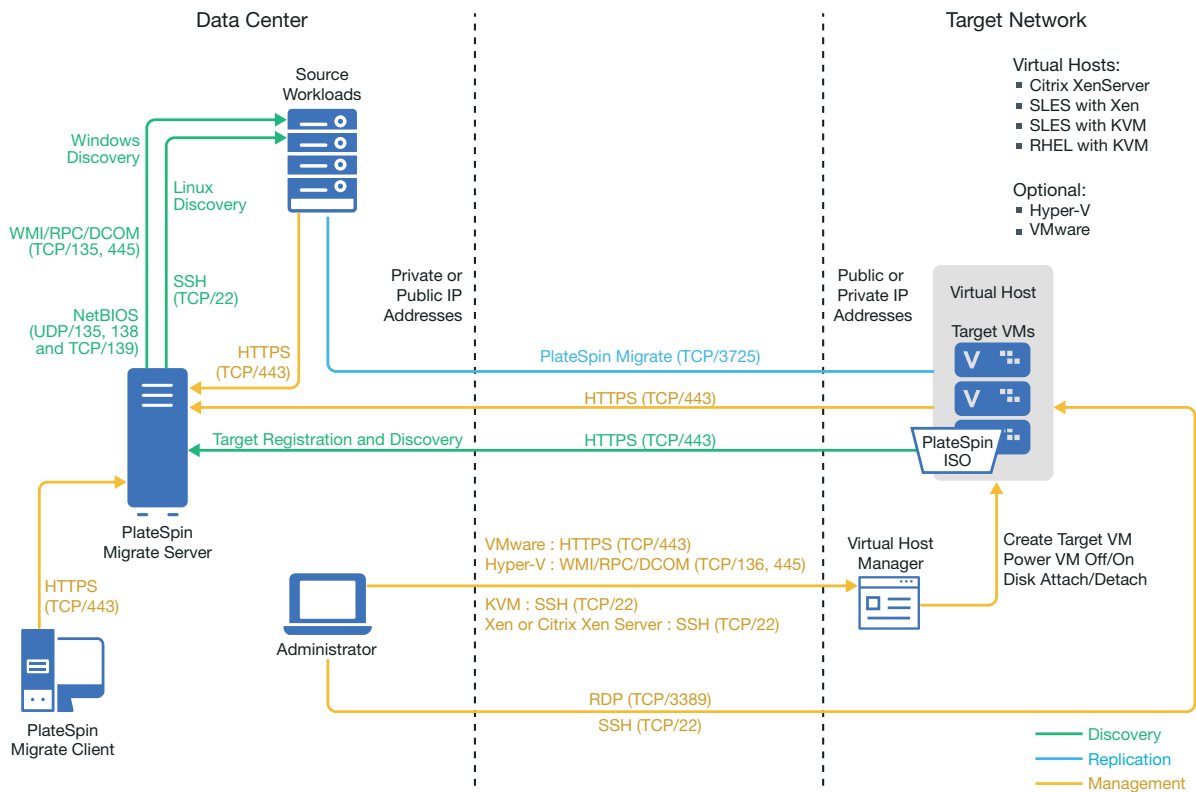
- [Section 14.1, “Deployment for Migration to Xen,” on page 183](#)
- [Section 14.2, “Planning for Migration to VMs on Xen,” on page 184](#)
- [Section 14.3, “Checklist for Semi-Automated Migration to Target VMs on Xen,” on page 184](#)

14.1 Deployment for Migration to Xen

Figure 14-1 shows the location of various components in your semi-automated Xen migration environment and the communications between them.

NOTE: **Figure 14-1** depicts automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443).

Figure 14-1 Semi-Automated Migration to VMs on Xen



14.2 Planning for Migration to VMs on Xen

Ensure that your Xen environment meets the following prerequisites for migration to VMs on Xen:

- ♦ Use PlateSpin Migrate Client to migrate workloads to virtual machines on Xen virtual hosts. PlateSpin Migrate Web Interface does not support migration to Xen virtual hosts.
- ♦ You can use Xen as the target virtualization platform in a semi-automated workload migration.
- ♦ Your target must be a fully virtualized (not paravirtualized) VM.
- ♦ Your source workload must be supported by PlateSpin Migrate and Xen.
See [“SUSE Linux Enterprise Server with Xen”](#) in Table 2-14, [“Supported Target Virtualization Platforms for the Migrate Client Only,”](#) on page 39.
- ♦ Your network environment must meet the requirements for access, discovery, and migration described in [“Access and Communication Requirements across your Migration Network”](#) on page 49.
- ♦ Configure volumes on the target disks with about 50 MB of additional storage space than the source disks.

For information about configuring semi-automated migration to a virtual machine on Xen, see [“Migration to Virtual Machines on Xen”](#) on page 399.

14.3 Checklist for Semi-Automated Migration to Target VMs on Xen

Task	Description
1. Prepare your VMware migration environment.	Figure 14-1, “Semi-Automated Migration to VMs on Xen,” on page 183 Section 14.2, “Planning for Migration to VMs on Xen,” on page 184
2. Discover target containers.	Section 19.5, “Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO,” on page 204
3. Discover source workloads.	Section 20.6.1, “Workload Discovery in the Migrate Client,” on page 214
4. Configure target workload migration.	Section 33.2, “Configuring Migration to a VM on a Xen Virtual Host,” on page 400
5. Execute migration.	Chapter 38, “Executing Workload Migrations,” on page 433

15 Prerequisites for Migration to VMs on KVM

PlateSpin Migrate Client supports semi-automated migration to target VMs on your KVM virtual host environment. This section describes the required KVM configuration that you must prepare before you can discover target VMs and configure migrations to them.


- [Section 15.1, “Deployment for Migration to KVM,” on page 185](#)
- [Section 15.2, “Planning for Migration to VMs on KVM,” on page 186](#)
- [Section 15.3, “Checklist for Semi-Automated Migration to Target VMs on KVM,” on page 186](#)

15.1 Deployment for Migration to KVM

[Figure 15-1](#) shows the location of various components in your semi-automated KVM migration environment and the communications between them.

NOTE: [Figure 15-1](#) depicts automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443).

Figure 15-1 Semi-Automated Migration to VMs on KVM



15.2 Planning for Migration to VMs on KVM

Ensure that your KVM environment meets the following prerequisites for migration to VMs on KVM:

- ♦ Use PlateSpin Migrate Client to migrate workloads to virtual machines on KVM virtual hosts. PlateSpin Migrate Web Interface does not support migration to KVM virtual hosts.
- ♦ You can use KVM as the target virtualization platform in a semi-automated workload migration.
- ♦ Your target must be a fully virtualized (not paravirtualized) VM.
- ♦ Your source workload must be supported by PlateSpin Migrate and KVM.

See the following information in [Table 2-14, “Supported Target Virtualization Platforms for the Migrate Client Only,”](#) on page 39.

- ♦ [“SUSE Linux Enterprise Server \(SLES\) with KVM”](#)
- ♦ [“Red Hat Enterprise Linux \(RHEL\) with KVM”](#)
- ♦ Your network environment must meet the requirements for access, discovery, and migration described in [“Access and Communication Requirements across your Migration Network”](#) on page 49.
- ♦ Configure volumes on the target disks with about 50 MB of additional storage space than the source disks.

For information about configuring semi-automated migration to a virtual machine on KVM, see [“Migration to Virtual Machines on KVM”](#) on page 403.

15.3 Checklist for Semi-Automated Migration to Target VMs on KVM

Task	Description
1. Prepare your VMware migration environment.	Figure 15-1, “Semi-Automated Migration to VMs on KVM,” on page 185 Section 15.2, “Planning for Migration to VMs on KVM,” on page 186
2. Discover target containers.	Section 19.5, “Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO,” on page 204
3. Discover source workloads.	Section 20.6.1, “Workload Discovery in the Migrate Client,” on page 214
4. Configure target workload migration.	Chapter 34, “Migration to Virtual Machines on KVM,” on page 403
5. Execute migration.	Chapter 38, “Executing Workload Migrations,” on page 433

16 Prerequisites for Migration to Physical Machines

PlateSpin Migrate Client supports semi-automated migration to target physical machines. This section describes the required configuration for migrations to physical machines.

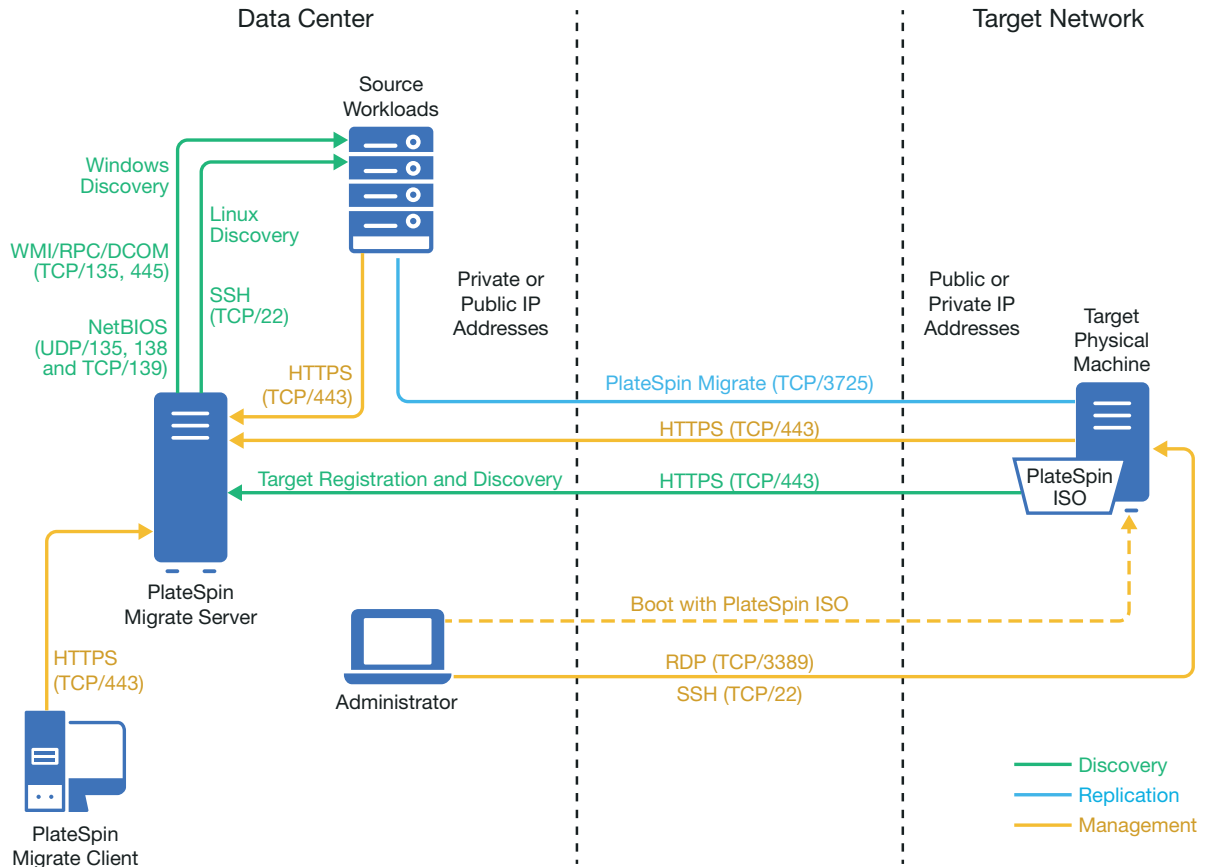
- ♦ [Section 16.1, “Deployment for Migration to Physical Machines,” on page 187](#)
- ♦ [Section 16.2, “Planning for Migration to Physical Machines,” on page 188](#)
- ♦ [Section 16.3, “Best Practices \(X2P\),” on page 189](#)
- ♦ [Section 16.4, “Checklist for Semi-Automated Migration to Physical Machines,” on page 189](#)

16.1 Deployment for Migration to Physical Machines

[Figure 16-1](#) shows the location of various components in your semi-automated physical machine migration environment and the communications between them.

NOTE: [Figure 16-1](#) depicts automated discovery and the network requirements for Windows and Linux workloads. You can alternatively use Migrate Agent on the source workload to register the workload and send its inventory details to PlateSpin Migrate server using HTTPS (TCP/443).

Figure 16-1 Semi-Automated Migration to Physical Machines



16.2 Planning for Migration to Physical Machines

Ensure that your environment meets the following prerequisites for migration to physical machines:

- ♦ Use PlateSpin Migrate Client to migrate workloads to a target physical machine. PlateSpin Migrate Web Interface does not support migration to physical machines.
- ♦ Your physical hardware must be supported by PlateSpin Migrate. See the following information in [“Supported Configurations” on page 23](#):
 - ♦ [Supported Workload Storage](#)
 - ♦ [Supported Workload Architectures](#)
- ♦ Your network environment must meet the requirements for access, discovery, and migration described in [“Access and Communication Requirements across your Migration Network” on page 49](#).
- ♦ Configure volumes on the target disks with about 50 MB of additional storage space than the source disks.

For information about configuring semi-automated migration to a physical machine, see [“Migration to Physical Machines” on page 407](#).

16.3 Best Practices (X2P)

- ♦ When you are migrating a workload from one vendor to a target hardware infrastructure from another vendor (for example, from HP to Dell), or if your source is a virtual machine, ensure that you disable vendor-specific or VM-specific services during the transfer. For example, disable the HP Insight service and the VMware Tools service.

See [Section 26.18, “Windows HAL or Kernel File Replacements,”](#) on page 312.

- ♦ When you are using the Offline transfer method for P2P and V2P migrations, ensure that you select the appropriate Full Duplex speed that matches your network Full Duplex mode.

See [Section 26.22, “Migration Network \(Replication Network\),”](#) on page 318.

- ♦ Ensure that vendor partitions are not being copied from the source.

See [Section 26.23, “Storage Disks and Volumes,”](#) on page 326.

16.4 Checklist for Semi-Automated Migration to Physical Machines

Task	Description
1. Prepare your VMware migration environment.	Figure 16-1, “Semi-Automated Migration to Physical Machines,” on page 188 Section 16.2, “Planning for Migration to Physical Machines,” on page 188
2. Discover target containers.	Section 19.5, “Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO,” on page 204
3. Discover source workloads.	Section 20.6.1, “Workload Discovery in the Migrate Client,” on page 214 -OR- Section 20.7, “Registering Workloads and Discovering Details with Migrate Agent,” on page 215
4. Configure target workload migration.	Section 35.2, “Configuring Migration to a Physical Target (P2P, V2P),” on page 408
5. Execute migration.	Chapter 38, “Executing Workload Migrations,” on page 433

17 Prerequisites for Migration to an Image

For information about capturing a workload to an image, see [Chapter 36, “Workload Migration with a PlateSpin Image,”](#) on page 415.

18 Preparing for Synchronization of Workloads with Server Sync

For information about synchronizing workloads to synchronize just the data that is different between the source and a target, see [Chapter 37, “Synchronizing Workloads with Server Sync,”](#) on page 423.

IV Discovering and Preparing Workloads and Targets

Before you can configure migrations, you must identify your planned target containers and source workloads. You get details about targets and workloads through a discovery and inventory process.

- ♦ [Chapter 19, “Discovering Target Hosts,” on page 197](#)
- ♦ [Chapter 20, “Discovering Source Workloads,” on page 211](#)
- ♦ [Chapter 21, “Preparing Device Drivers,” on page 221](#)
- ♦ [Chapter 22, “Preparing Linux Workloads for Migration,” on page 233](#)
- ♦ [Chapter 23, “Preparing for Migration of Windows Clusters,” on page 237](#)
- ♦ [Appendix C, “Troubleshooting Discovery,” on page 247](#)
- ♦ [Appendix D, “Linux Distributions Supported by Migrate,” on page 251](#)
- ♦ [Appendix E, “Synchronizing Serial Numbers on Cluster Node Local Storage,” on page 267](#)
- ♦ [Appendix F, “Migrate Agent Utility,” on page 269](#)
- ♦ [Appendix G, “PlateSpin ISO Image,” on page 279](#)

19 Discovering Target Hosts

For any workload migration, you must have a discovered source and a discovered target host. A target discovery operation populates the PlateSpin Migrate database with detailed inventory information about the target host and its resources. The inventory provides the data necessary to determine the host's use and to properly configure one or more migrations to the target host.

- [Section 19.1, “About Target Host Discovery,” on page 197](#)
- [Section 19.2, “Network Access Requirements for Target Host Discovery,” on page 198](#)
- [Section 19.3, “Discovery Guidelines for Target Hosts,” on page 198](#)
- [Section 19.4, “Discovering Details for Target Containers,” on page 200](#)
- [Section 19.5, “Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO,” on page 204](#)
- [Section 19.6, “Registering and Discovering Details for Target Physical Machines with PlateSpin ISO,” on page 206](#)
- [Section 19.7, “Discovering Target VMs for Server Sync Jobs,” on page 208](#)
- [Section 19.8, “Refreshing Target Host Details,” on page 208](#)
- [Section 19.9, “Removing Target Containers,” on page 210](#)

19.1 About Target Host Discovery

PlateSpin Migrate Web Interface and PlateSpin Migrate Client provide automated discovery and inventory of [supported target host platforms](#). See [Table 19-1](#) for an overview of the target host discovery capabilities of each tool.

Table 19-1 Supported Target Host Discovery Capabilities

Target Host Discovery	Migrate Client	Web Interface
Cloud Targets		
Amazon Web Services (Cloud Region)	✓	✗
Microsoft Azure (Cloud Location)	✗	✓
VMware vCloud Director (Organization)	✗	✓
VMware Targets		
All VMware ESX hosts affiliated with a VMware vCenter Server	✓	✓
VMware DRS Cluster	✓	✓
VMware ESX Server	✓	✓
Other Targets		
Microsoft Hyper-V virtual host	✓	✗

Target Host Discovery	Migrate Client	Web Interface
Citrix XenServer virtual host	✓	✗
Linux KVM or Xen virtual host	✓	✗
Physical host	✓	✗
Discovery Capabilities		
An individual host server	✓	✓
Multiple virtual host servers at a time	✓	✗
All hosts in a domain	✓	✗

You can view discovered target host containers in the Targets list in either tool:

- ♦ **Web Interface:** The Targets list includes:
 - ♦ All cloud hosts and VMware hosts discovered using the Web Interface
 - ♦ All VMware hosts in the default network discovered using Migrate Client

NOTE: Use the Web Interface to discover target cloud hosts and VMware hosts in non-default networks if you plan to use the Web Interface for migrations to those locations.

All target hosts displayed in the Web Interface Targets list are supported as migration targets using the Web Interface. See [Table 19-1, “Supported Target Host Discovery Capabilities,” on page 197](#).

- ♦ **Migrate Client:** The Targets list includes:
 - ♦ All the discovered VMware target hosts, no matter where you initiated discovery.
 - ♦ All Hyper-V hosts discovered using Migrate Client

For information about the target hosts that the Web interface and the Migrate Client supports, see [Table 19-1, “Supported Target Host Discovery Capabilities,” on page 197](#).

19.2 Network Access Requirements for Target Host Discovery

For information about network access requirements for discovery of target hosts, see [Section 2.6.1, “Requirements for Discovery,” on page 49](#).

19.3 Discovery Guidelines for Target Hosts

For information about the software, network, and firewall requirements that systems in your environment must meet for the discovery and inventory process, see [Section 2.6.1, “Requirements for Discovery,” on page 49](#).

- ♦ [Section 19.3.1, “Target Host Discovery Parameters for Migrate Web Interface,” on page 199](#)
- ♦ [Section 19.3.2, “Target Host Discovery Parameters for Migrate Client,” on page 199](#)

19.3.1 Target Host Discovery Parameters for Migrate Web Interface

Table 19-2 provides guidelines for target type selection, credential format, and syntax for discovery parameters for target hosts using the Migrate Web Interface.

Table 19-2 Guidelines for Migrate Web Interface Target Type and Credentials for Target Hosts

To Discover	Target Type	Credentials
Azure Cloud Location	Microsoft Azure Location	Subscription ID Application ID Azure user with Subscription administrator role
vCloud Organization	VMware Cloud Organization	Organization Administrator credentials (user name and password)
VMware vCenter Cluster	VMware DRS Cluster	VMware vCenter Web service credentials (user name and password)
VMware ESXi Server	VMware ESX Server	ESX account with administrator role OR Windows domain credentials (versions 4 and 4.1 only)

19.3.2 Target Host Discovery Parameters for Migrate Client

Table 19-3 provides guidelines for machine type selection, credential format, and syntax for discovery parameters for target hosts using the Migrate Client.

Table 19-3 Guidelines for Migrate Client Machine Type and Credentials for Target Hosts

To Discover	Machine Type	Credentials	Remarks
AWS target VM (Migrate prepares the AWS target for the source workload you want to migrate)	Linux	Access Key ID Secret Key ID	
VMware ESX hosts affiliated with a VMware vCenter Server	VMware vCenter	VMware vCenter Web service credentials (user name and password) OR Windows domain credentials (versions 4 and 4.1 only)	All subsequent communications with ESX hosts take place through the vCenter Server. To reset the access mechanism to Direct to ESX or back to Via vCenter , refresh the details of the required ESX host and select the required option. See “Refreshing Target Host Details” on page 208.

To Discover	Machine Type	Credentials	Remarks
VMware ESX hosts	VMware ESX	ESX account with administrator role OR Windows domain credentials (versions 4 and 4.1 only)	
Hyper-V hosts	Windows	Local or domain administrator credentials.	For the username, use this format: <ul style="list-style-type: none"> ♦ For domain member machines: <i>authority\principal</i> ♦ For workgroup member machines: <i>hostname\principal</i>
All Linux KVM or Xen virtual hosts	Linux	Root-level username and password	Non-root accounts must be properly configured to use <code>sudo</code> . See KB Article 7920711 (https://www.netiq.com/support/kb/doc.php?id=7920711) .
PlateSpin Image Servers	Windows	Local or domain administrator credentials.	For the username, use this format: <ul style="list-style-type: none"> ♦ For domain member machines: <i>authority\principal</i> ♦ For workgroup member machines: <i>hostname\principal</i>

19.4 Discovering Details for Target Containers

Before you configure a migration job, you must discover and perform an inventory of the target container. The inventory collects information about the host platform and its resources, such as the amount of RAM, number of cores and processors, datastores, networks, and resource groups.

- ♦ [Section 19.4.1, “Target Discovery in the Migrate Client,” on page 200](#)
- ♦ [Section 19.4.2, “Target Discovery in the Web Interface,” on page 202](#)

19.4.1 Target Discovery in the Migrate Client

In Migrate Client, you can discover:

- ♦ An individual virtual machine host server
- ♦ Multiple virtual machine host servers
- ♦ All VMware ESX hosts affiliated with a VMware vCenter Server
- ♦ Hyper-V hosts

Before you begin discovery operations, ensure that PlateSpin Server can communicate with your source workloads and targets. See [Section 2.6.1, “Requirements for Discovery,” on page 49](#).

To discover targets using Migrate Client:

- 1 In the Migrate Client toolbar, click **Discover Details**.

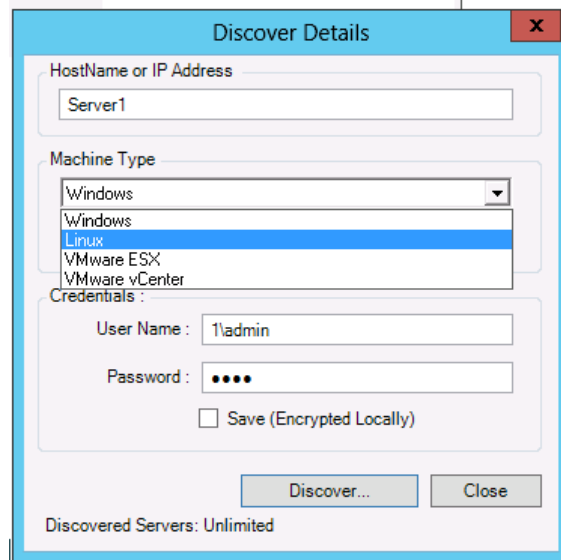
or

In the Servers view, right-click in a blank area, then select **Discover Details**.

- 2 In the **Discover Details** dialog box, type the host name or IP address of the target.

To discover multiple machines, specify multiple host names or IP addresses separated by semicolons. If the target is behind a NAT device, specify its public (external) IP address.

See [“Migrations Across Public and Private Networks through NAT”](#) on page 53.



- 3 Select the machine type for the target container. If you select VMware vCenter, also provide the name of the vCenter cluster.

- ♦ Windows
- ♦ Linux
- ♦ VMware ESX
- ♦ VMware vCenter
- ♦ Microsoft Hyper-V

See [“Discovery Guidelines for Target Hosts”](#) on page 198.

Discovering hosts with Xen hypervisor systems results in these systems being registered as PlateSpin Migrate source workloads (as opposed to VM host targets). For information about using these platforms as workload migration targets, see [“Migration to Virtual Machines on Xen”](#) on page 399.

- 4 Provide administrator credentials for the machine you are discovering.

See [“Discovery Guidelines for Target Hosts”](#) on page 198.

- 5 (Optional) If you want to store these credentials for use during future jobs, enable the **Save (Encrypted Locally)** option.
- 6 Click **Discover** and wait for the process to complete.
- 7 (Optional) If you want to monitor the progress of the job, switch to the Jobs view.

19.4.2 Target Discovery in the Web Interface

To migrate a workload through the Web Interface, you must first add or discover the intended target container and its resources.

PlateSpin Migrate Web Interface supports discovery of virtual and cloud target containers:

- ♦ Microsoft Azure Location
- ♦ VMware vCloud Organization
- ♦ VMware DRS Cluster
- ♦ VMware ESX Server

When you add the target, its associated resources are automatically discovered. You can add one container at a time. All available target containers are listed on the Targets page.

Before you begin discovery operations, ensure that PlateSpin Server can communicate with your source workloads and targets. See section [Section 2.6.1, “Requirements for Discovery,” on page 49](#).

To add a target container:

- 1 In the Migrate Web Interface, click **Targets > Add Target**.
- 2 Select one of the following target types:
 - ♦ Microsoft Azure Location
 - ♦ VMware DRS Cluster
 - ♦ VMware ESX Server
 - ♦ VMware vCloud Organization
- 3 Depending on the type of targets you selected in the previous step, specify the appropriate access information.
 - ♦ **Microsoft Azure Location:** See [Table 19-4](#).
For more information about the options for **Microsoft Azure Location**, see the white paper “Best Practices for Migrating Servers to Microsoft Azure with PlateSpin Migrate” on the [PlateSpin Migrate Resources web page \(https://www.microfocus.com/products/migrate/resources/\)](https://www.microfocus.com/products/migrate/resources/).
 - ♦ **VMware DRS Cluster:** See [Table 19-5](#).
 - ♦ **VMware ESX Server:** See [Table 19-6](#).
 - ♦ **VMware vCloud Organization:** See [Table 19-7](#).

Table 19-4 Options for Microsoft Azure Location Target

Option	Description
Subscription Id	Specify the subscription ID for your Microsoft Azure account.
Application Id	Specify your Azure Application ID required to enable PlateSpin Migrate to use the Azure APIs when it replicates or migrates workloads on your behalf to VMs in the target Azure account.
Username and Password	Specify administrator-level credentials for accessing the parent Microsoft Azure account.
Location Name	<p>Select the location for the Microsoft Azure target.</p> <p>Click Update Location List to refresh the list of available locations in the menu.</p> <p>Locations are now sorted by the geographical region and alphabetically. The mapping is fixed and is based on the current categories that Azure uses. If Microsoft Azure adds new regions after the current release, Migrate displays them dynamically and alphabetically in the Recently Added category.</p>

Table 19-5 Options for VMware DRS Cluster Target

Option	Description
vCenter Hostname or IP	Specify the host name or IP address of the vCenter server.
Cluster Name	Specify the name of the DRS cluster. This is applicable only for VMware DRS Cluster.
Username and Password	Specify administrator-level credentials for accessing the target host.

Table 19-6 Options for VMware ESX Server Target

Option	Description
Hostname or IP	Specify the host name or IP address of the VMware ESX server.
Username and Password	Specify administrator-level credentials for accessing the target host.

Table 19-7 Options for VMware vCloud Organization Target

Option	Description
vCloud Director Server Address	<p>Specify the server host name or the IP address of the vCloud Director server.</p> <p>For example: <code>cloud.example.com</code> or <code>10.10.10.101</code></p>

Option	Description
Organization Name	Specify the name of the organization in the vCloud Director server. The name is case sensitive in vCloud. Type the name exactly as you created it. For example: DemoOrg001
Username and Password	Specify the organization-level administrator credentials for accessing the target host. For example: demouser1 and demopwd

- 4 Click **Test Credentials** to validate the credential values you specified.
- 5 Click **Add** to add and discover details about the target and list it on the Targets page.

19.5 Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO

PlateSpin Migrate Client enables you to migrate a source workload to a target virtual machine on a virtual host, where the VM is regarded as a target physical machine:

- ♦ VMware

Semi-automated migration to VMs on VMware can be done, but fully automated migration to target VMware containers is preferred. Discovery for target VMware containers is available in the Migrate Client and the Migrate Web Interface. See [“Discovering Details for Target Containers”](#).
- ♦ Microsoft Windows Server with Hyper-V

Semi-automated migration to VMs on Hyper-V can be done, but fully automated migration to target Hyper-V containers is preferred. Discovery for target Hyper-V containers is available only in the Migrate Client. See [“Target Discovery in the Migrate Client”](#).
- ♦ Citrix XenServer
- ♦ Xen
- ♦ KVM

For information about supported virtual host platforms, see [Table 2-14, “Supported Target Virtualization Platforms for the Migrate Client Only,”](#) on page 39.

PlateSpin ISO registers the target physical machine with the PlateSpin Migrate server and performs an inventory of the machine to collect information about it, such as the amount of RAM, number of cores and processors, storage disks, and NICs.

- ♦ [Section 19.5.1, “Prerequisites for Discovering Target VMs,”](#) on page 205
- ♦ [Section 19.5.2, “Registering and Discovering Target VMs on Virtual Hosts,”](#) on page 205
- ♦ [Section 19.5.3, “Configuration Information,”](#) on page 206

19.5.1 Prerequisites for Discovering Target VMs

PlateSpin Migrate does not automatically build the target VM for you on the target virtual host. You must manually set up the target virtual machine with guest operating system type and version settings that match your source workload, in accordance with the features and capabilities of the virtualization platform. You must also prepare the PlateSpin ISO file and attach it as a boot CD for the VM.

- 1 Download and prepare the PlateSpin ISO image for use with the target VM. Attended and unattended registration options are possible. See [“Preparing the PlateSpin ISO Image for Target Registration and Discovery” on page 280](#).
- 2 Use the native interface of the required virtualization platform to create a virtual machine. See the following as appropriate for your target VM:
 - ♦ [“Creating and Configuring the Target Virtual Machine \(Hyper-V\)” on page 393](#)
 - ♦ [“Creating and Configuring the Target Virtual Machine \(Citrix XenServer\)” on page 396](#)
 - ♦ [“Creating and Configuring the Target Virtual Machine \(Xen on SLES\)” on page 400](#)
 - ♦ [“Creating and Configuring the Target Virtual Machine \(RHEL KVM\)” on page 404](#)
- 3 Ensure that the VM is configured to restart on reboot and that you attach the PlateSpin ISO file as a boot CD for the VM.

19.5.2 Registering and Discovering Target VMs on Virtual Hosts

After you create and prepare the virtual machine to boot with the PlateSpin ISO, you are ready to register it as a target VM with your PlateSpin Server.

- 1 From the Virtual Machine Manager, power on (or reboot) the virtual machine, then launch the virtual machine console and monitor the boot process.

When the virtual machine completes the boot process, it prompts you for parameters that control the registration of the machine and its profile with PlateSpin Migrate. If you are using the unattended registration process, the required parameters are read from an answer file.

- 2 (Conditional) For 64-bit systems, at the initial boot prompt, type the following:

```
ps64
```

Press Enter.

- 3 At the command line, provide the required information at each individual prompt:
 - ♦ **PlateSpin Server:** Enter the PlateSpin Server URL, using the following format:
`http://Your_PlateSpin_Server/platespinmigrate`
Replace *Your_PlateSpin_Server* with the host name or the IP address of your PlateSpin Server host.
 - ♦ **Credentials (User Name/Password):** Enter the name of an administrator-level user on the PlateSpin Server host, including the domain or machine name. For example:
domain\username, or *localhost\Administrator*. Provide a valid password for the specified user.
 - ♦ **Network Card:** Select the network card that is active, then either enter a temporary static IP address for this NIC or press Enter to dynamically obtain an IP address from a DHCP server.
 - ♦ **Temporary hostname:** Provide a temporary VM name for PlateSpin Migrate Client to use to list the newly registered VM. The workload's target host name you select in the migration job overwrites this name.

- ♦ **SSL encryption:** If your PlateSpin Migrate is installed on a host with SSL encryption enabled, enter `Yes`. If not, enter `No`.
- ♦ **PlateSpin Migrate Network:** Unless you have defined your own PlateSpin Migrate Network in PlateSpin Migrate Client, press `Enter`. If you are working with a non-default PlateSpin Migrate Network, type its name, then press `Enter`.

A controller on your target virtual machine communicates with PlateSpin Server and registers the virtual machine as a physical target for a migration job.

After a few moments, PlateSpin Migrate Client displays the target virtual machine in the Servers view.

NOTE: If registration fails with an authorization error, you might need to synchronize the clocks of the source and the target, modify the LAN Manager Authentication Level on the target, or both. See [Table C-1, “Common Issues and Solutions Related to Discovery Operations,” on page 247](#).

19.5.3 Configuration Information

For information about configuring migration for target VMs on virtual hosts, see the following:

- ♦ [“Migration to VMs on Hyper-V Using X2P Workflow” on page 393](#)
- ♦ [“Migration to Virtual Machines on Citrix XenServer” on page 395](#)
- ♦ [“Migration to Virtual Machines on Xen” on page 399](#)
- ♦ [“Migration to Virtual Machines on KVM” on page 403](#)

19.6 Registering and Discovering Details for Target Physical Machines with PlateSpin ISO

To discover a physical target and inventory its hardware components, you must boot the target machine with the PlateSpin ISO image on a CD or other media from which your target can be booted.

PlateSpin ISO registers the target physical machine with the PlateSpin Migrate server and performs an inventory of the machine to collect information about it, such as the amount of RAM, number of cores and processors, storage disks, and NICs.

- ♦ [Section 19.6.1, “Prerequisites for Discovering Target Physical Machines,” on page 206](#)
- ♦ [Section 19.6.2, “Registering Target Physical Machines Using the PlateSpin ISO,” on page 207](#)
- ♦ [Section 19.6.3, “Configuration Information,” on page 207](#)

19.6.1 Prerequisites for Discovering Target Physical Machines

You must prepare the PlateSpin ISO file and attach it as a boot CD for the physical machine.

- 1 Download and prepare the PlateSpin ISO image for use with the physical machine. Attended and unattended registration options are possible. See [“Preparing the PlateSpin ISO Image for Target Registration and Discovery” on page 280](#).
- 2 Ensure that the physical machine is configured to restart on reboot and that you attach the PlateSpin ISO file as a boot CD.

19.6.2 Registering Target Physical Machines Using the PlateSpin ISO

After you create and prepare the physical machine to boot with the PlateSpin ISO, you are ready to register the target machine with your PlateSpin Server.

- 1 Boot the target machine from the PlateSpin ISO image.
- 2 (Conditional) For 64-bit systems, at the initial boot prompt, type the following:

```
ps64
```


Press Enter.
- 3 At the command line, provide the required information at each individual prompt:
 - ♦ **PlateSpin Server:** Enter the PlateSpin Server URL, using the following format:

```
http://Your_PlateSpin_Server/platespinmigrate
```


Replace *Your_PlateSpin_Server* with the host name or the IP address of your PlateSpin Server host.
 - ♦ **Credentials (User Name/Password):** Enter the name of an administrator-level user on the PlateSpin Server host, including the domain or machine name. For example:
domain\username, or *localhostAdministrator*. Provide a valid password for the specified user.
 - ♦ **Network Card:** Select the network card that is active, then either enter a temporary static IP address for this NIC or press Enter to dynamically obtain an IP address from a DHCP server.
 - ♦ **Temporary hostname:** Provide a temporary VM name for PlateSpin Migrate Client to use to list the newly registered VM. The workload's target host name you select in the migration job overwrites this name.
 - ♦ **SSL encryption:** If your PlateSpin Migrate is installed on a host with SSL encryption enabled, enter *Yes*. If not, enter *No*.
 - ♦ **PlateSpin Migrate Network:** Unless you have defined your own PlateSpin Migrate Network in PlateSpin Migrate Client, press Enter. If you are working with a non-default PlateSpin Migrate Network, type its name, then press Enter.

A controller on your target virtual machine communicates with PlateSpin Server and registers the virtual machine as a physical target for a migration job.

After a few moments, PlateSpin Migrate Client displays the physical target in the Servers view.

NOTE: If registration fails with an authorization error, you might need to synchronize the clocks of the source and the target, modify the LAN Manager Authentication Level on the target, or both. See [Table C-1, "Common Issues and Solutions Related to Discovery Operations," on page 247](#).

19.6.3 Configuration Information

For information about configuring migration to physical machines, see ["Migration to Physical Machines" on page 407](#).

19.7 Discovering Target VMs for Server Sync Jobs

If you want to synchronize two workloads, and if your synchronization target is a virtual machine, you must discover and register an appropriate virtual machine first. For information about the Server Sync feature, see [“Synchronizing Workloads with Server Sync” on page 423](#).

- 1 On your virtual machine host, create a virtual machine with the desired specifications and install the operating system that matches the intended source workload, including the exact service pack.
- 2 Discover the virtual machine host or refresh its details.
- 3 In the Servers view, right-click the newly created virtual machine underneath the virtual machine server, then select **Prepare for synchronization**.
- 4 Specify administrator credentials for the virtual machine server.
- 5 (Optional) If you want to store these credentials for use during future jobs, enable the **Save (Encrypted Locally)** option.
- 6 (Optional) To configure the temporary (Take Control) network settings, such as choosing which virtual network to use from those available on the virtual machine server and configuring TCP/IP settings, click **Configure**, then configure the network settings as required.
- 7 Click **Prepare** and wait for the job to complete.

On completion, the Servers view lists a new Server Sync target underneath the VM host: 

19.8 Refreshing Target Host Details

You should routinely refresh details about your target containers before setting up or executing a migration job.

- ♦ [Section 19.8.1, “Refresh Target Details in the Web Interface,” on page 208](#)
- ♦ [Section 19.8.2, “Refresh Target Details in Migrate Client,” on page 209](#)
- ♦ [Section 19.8.3, “Refresh ESX Host Details in Migrate Client,” on page 209](#)

19.8.1 Refresh Target Details in the Web Interface

PlateSpin Migrate Web Interface enables you to refresh the discovered resources for virtual and cloud target containers:

- ♦ Amazon Cloud Region
- ♦ Microsoft Azure Location
- ♦ VMware DRS Cluster
- ♦ VMware ESX Server
- ♦ VMware vCloud Organization

When you refresh the target, its associated resources are automatically rediscovered and updated. You can refresh one container at a time.

To refresh details for a target container:

- 1 In the PlateSpin Migrate Web Interface, click **Targets**.
- 2 Select a target.

- 3 Click **Refresh**.
- 4 Expand the panels for the associated resources to view the changes.


19.8.2 Refresh Target Details in Migrate Client

Migrate Client allows you to refresh target details for containers discovered using the Migrate Client:

- ♦ Physical servers
- ♦ VMware DRS clusters
- ♦ VMware ESX servers
- ♦ Microsoft Hyper-V virtual hosts
- ♦ All Linux KVM and Xen virtual hosts
- ♦ PlateSpin Image servers

To refresh target details:

- 1 In the Servers view, right-click the required item, then select **Refresh Details**.




- 2 Specify the credentials appropriate for the system being refreshed, then click **Refresh**.
PlateSpin Migrate starts a discovery job, which you can monitor in the Jobs view.

19.8.3 Refresh ESX Host Details in Migrate Client

If you discover ESX hosts through a vCenter Server, all subsequent communications with these VM hosts take place through the vCenter Server.

You can reset this access mechanism to **Direct to ESX** or back to **Via vCenter**:

- 1 Refresh the required ESX host.
See [“Refresh Target Details in Migrate Client” on page 209](#).
The system prompts you to select an access type:



- 2 Select the required option, then click **Refresh**.

PlateSpin Migrate rediscovers the required ESX host by using the specified access mechanism, which it uses for all subsequent communications with that host.

19.9 Removing Target Containers

After you complete all migration jobs for a target container, you can remove (undiscover) the target container. You might also remove a target that will not be used.

IMPORTANT

- ♦ If an object is listed both in the Migrate Client and the Migrate Web Interface, then you must use the Web Interface to remove the object.
 - ♦ Before you delete a target container that is in use for configured jobs, you must ensure that all the affected jobs are completed.
-

To undiscover a workload through the Migrate Client:

- 1 On the Workloads page, right-click the target and select **Undiscover Target**.

To remove a target through the Migrate Web Interface:

- 1 On the Targets page, click **Remove** next to the target you want to remove from Migrate.

20 Discovering Source Workloads

For any workload migration, you must have a discovered source and a discovered target. A workload discovery operation populates the PlateSpin Migrate database with detailed inventory information about a workload that you want to migrate. The workload inventory provides the data necessary to determine the machine's use and to properly configure its migration.

- ♦ [Section 20.1, “About Source Workload Discovery,” on page 211](#)
- ♦ [Section 20.2, “Network Access Requirements for Workload Discovery,” on page 212](#)
- ♦ [Section 20.3, “Discovery Guidelines for Source Workloads,” on page 212](#)
- ♦ [Section 20.4, “Populating the Servers View with a List of Windows Computers in a Domain,” on page 213](#)
- ♦ [Section 20.5, “Discovering Details for All Windows Workloads in a Domain,” on page 213](#)
- ♦ [Section 20.6, “Discovering Details for Source Workloads,” on page 213](#)
- ♦ [Section 20.7, “Registering Workloads and Discovering Details with Migrate Agent,” on page 215](#)
- ♦ [Section 20.8, “Refreshing Source Workload Details,” on page 217](#)
- ♦ [Section 20.9, “Using Tags to Track Logical Associations of Workloads,” on page 218](#)
- ♦ [Section 20.10, “Undiscovering or Removing Source Workloads,” on page 219](#)

20.1 About Source Workload Discovery

PlateSpin Migrate Web Interface and PlateSpin Migrate Client provide automated discovery and inventory of [supported source workloads](#). See [Table 9-1](#) for an overview of the workload discovery capabilities of each tool.

Table 20-1 Supported Source Workload Discovery Capabilities

Source Workload Discovery	Migrate Client	Web Interface
Windows standalone workloads	✓	✓
Windows cluster workloads (to target VMware host)	✓	✓
Linux standalone workloads	✓	✓
Linux cluster workloads	✗	✗
Multiple machines at a time	✓	✗
All machines in a domain	✓	✗

You can view discovered source workloads in the Workloads list in either tool:

- ♦ **Web Interface:** The Workloads list includes:
 - ♦ All source workloads discovered using the Web Interface
 - ♦ Source workloads in the default network discovered using Migrate Client

NOTE: Use the Web Interface to discover source workloads in non-default networks if you plan to migrate them using the Web Interface.

All workloads displayed in the Web Interface Workloads list are supported for migration using the Web Interface. See [Table 20-1](#).

- ♦ **Migrate Client:** The Workloads list includes all discovered source workloads, no matter where you initiated discovery.

Some workloads in the Migrate Client Workloads list might not be supported as migration targets using the Migrate Client. See [Table 20-1](#).

In addition, the Mass Discover CLI enables you to discover multiple workloads from a CSV file. The related migration jobs start according to the schedules you set for them.

20.2 Network Access Requirements for Workload Discovery

For information about network access requirements for discovery for Windows and Linux workloads, see [Section 2.6.1, “Requirements for Discovery,” on page 49](#).

20.3 Discovery Guidelines for Source Workloads

For information about the software, network, and firewall requirements that systems in your environment must meet for the discovery and inventory process, see [Section 2.6.1, “Requirements for Discovery,” on page 49](#).

[Table 20-2](#) provides guidelines for machine type selection, credential format, and syntax for discovery parameters for workloads.

Table 20-2 Guidelines for Machine Type and Credentials for Source Workloads

To Discover	Machine Type	Credentials	Remarks
All Windows workloads	Windows	Local or domain administrator credentials.	For the username, use this format: <ul style="list-style-type: none">♦ For domain member machines: <i>authority\principal</i>♦ For workgroup member machines: <i>hostname\principal</i>
All Linux workloads	Linux	Root-level username and password	Non-root accounts must be properly configured to use <code>sudo</code> . See KB Article 7920711 (https://www.netiq.com/support/kb/doc.php?id=7920711) .
Windows workloads in AWS	Windows		Use the workload's public IP address and the Common Inventory discovery method.

20.4 Populating the Servers View with a List of Windows Computers in a Domain

In the PlateSpin Migrate Client, the Network Discovery feature populates the Server view with all Windows physical machines and virtual machines that are online in a specified domain. PlateSpin Migrate uses the standard Windows network browser function for discovery. Because Linux workloads and virtual machine servers do not advertise to the Windows network browser, they are not automatically detected and do not appear in the list.

Unlike a full discovery with inventory, Network Discovery lists the Windows machines but does not inventory each workload to gather its details. A workload inventory is required for migration jobs. You can use either of the following methods to inventory the workloads:

- ♦ Use **Discover All Servers** to discover details for each of the listed Windows workloads. See [“Discovering Details for All Windows Workloads in a Domain” on page 213](#).
- ♦ Use **Discover Details** to discover details a specific workload. See [“Workload Discovery in the Migrate Client” on page 214](#).

Network Discovery is enabled by default. The option is a toggle between enabled and disabled modes.

To enable or disable Network Discovery:

- 1 In the Migrate Client, double-click **Network Discovery** at the bottom right corner of the Migrate Client window.

20.5 Discovering Details for All Windows Workloads in a Domain

You can use the **Discover All Servers** option in the Servers view to discover and perform an inventory of all Windows workloads in a specified domain. The Network Discovery option must be enabled to detect the Windows servers in the network.

- 1 In Migrate Client, enable the Network Discovery feature.
See [“Populating the Servers View with a List of Windows Computers in a Domain” on page 213](#).
- 2 Expand the list of domains that contain the machines to be inventoried.
- 3 Right-click the domain name, then select **Discover All Servers**.
- 4 Specify domain-level administrator credentials.
- 5 Click **Discover** and wait for the process to complete.
- 6 (Optional) If you want to monitor the progress of the discovery job, switch to the Jobs view.

20.6 Discovering Details for Source Workloads

Before you configure a migration job, you must discover and perform an inventory of the workload. The inventory collects information about the workload such as the server host name, amount of RAM, number of cores and processors, storage disks and volumes, NICs and applications and their start states.

- ♦ [Section 20.6.1, “Workload Discovery in the Migrate Client,” on page 214](#)
- ♦ [Section 20.6.2, “Workload Discovery in the Migrate Web Interface,” on page 215](#)

20.6.1 Workload Discovery in the Migrate Client

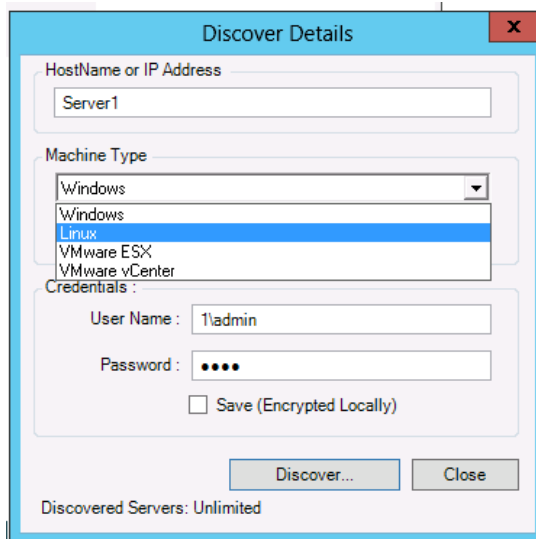
In the PlateSpin Migrate Client, you can use the **Discover Details** option in the Servers view to discover and perform an inventory for physical or virtual machines:

- ♦ An individual Windows workload
- ♦ An individual Linux workload
- ♦ Multiple Windows or Linux workloads at a time

Before starting discovery operations, ensure that PlateSpin Server can communicate with your source workloads. See [Section 2.6.1, “Requirements for Discovery,” on page 49](#).

To discover workloads using Migrate Client:

- 1 On the Migrate Client toolbar, click **Discover Details**.
or
In the Servers view, right-click in a blank area, then select **Discover Details**.
or
In the Servers view, right-click a Windows workload that has been populated through network discovery. then select **Discover Details**.
- 2 In the **Discover Details** dialog box, type the host name or IP address of the source workload.
To discover multiple machines at a time, specify multiple host names or IP addresses separated by semicolons. If the machine is behind a NAT device, specify its public (external) IP address.
See [“Migrations Across Public and Private Networks through NAT” on page 53](#).



- 3 Select the machine type for the source workload
 - ♦ Windows
 - ♦ Linux
- 4 Provide administrator credentials for the machine you are discovering.
See [“Discovery Guidelines for Source Workloads” on page 212](#).
- 5 (Optional) If you want to store these credentials for use during future jobs, enable the **Save (Encrypted Locally)** option.

- 6 Click **Discover** and wait for the process to complete.
- 7 (Optional) If you want to monitor the progress of the job, switch to the Jobs view.

20.6.2 Workload Discovery in the Migrate Web Interface

To migrate a workload through the Web Interface, you must first add (or *discover*) the workload.

PlateSpin Migrate Web Interface supports discovery of a physical, virtual, or cloud-based machine:

- ♦ An individual Windows workload
- ♦ An individual Linux workload

Before you discover a workload, ensure that PlateSpin Server can communicate with your source workloads. See [Section 2.6.1, “Requirements for Discovery,” on page 49](#).

To discover a workload:

- 1 In the PlateSpin Migrate Web Interface, click **Workloads > Add Workload**.
Alternatively, you can click the **Add Workload** option on the Dashboard page.
- 2 Specify the host name or the IP address of the workload you want to add.
- 3 Select the type of workload.
- 4 Specify the credentials to connect to the workload.
- 5 Click **Add Workload** to discover the workload and list it on the Workloads page.

20.7 Registering Workloads and Discovering Details with Migrate Agent

Migrate Agent is a command line utility that enables you to register source workloads with PlateSpin Migrate servers and send details about the workloads to the server via HTTPS (TCP/443). Registration allows you to add workloads that cannot be discovered, such as:

- ♦ When you deploy Migrate server in the cloud without a site-to-site VPN
- ♦ When corporate network or policy restrictions prohibit opening ports for automated discovery

Migrate Agent enables you to migrate a Windows workload without opening any inbound ports, such as SMB or NetBIOS. Only HTTPS and a replication port (TCP/3725 is the default) are needed outbound for the source workloads.

When you use the Migrate Agent on the source workload, the source workload contacts the target workload for data transfers. The direction is controlled at the server level. You must reconfigure the replication port direction on the Migrate Server (`SourceListensForConnection=False`). See [Section 5.9, “Configuring the Contact Direction for the Replication Port,” on page 107](#).

You must install the Migrate Agent utility on each source workload. When you use the `register` option, Migrate Agent performs discovery locally on the workload and sends its details to the Migrate Server through HTTPS (TCP/443). After you register the workload, use the Migrate Web Interface to configure the workload migration to the target cloud where the Migrate Server instance is deployed.

Registered workloads differ from discovered workloads in the following ways:

- ♦ Registered source workloads do not store the source credentials on the Migrate Server.

- ♦ You must use Migrate Agent to install, upgrade, and remove the Windows PlateSpin drivers from registered source workloads.
- ♦ After you delete the contract for a registered source workload, you must manually remove the OFX Controller from the workload. See [“Cleaning Up Linux Workloads” on page 446](#).

For information about the Migrate Agent, see [“Migrate Agent Utility” on page 269](#).

- ♦ [Section 20.7.1, “Windows Workload Registration and Discovery with Migrate Agent,” on page 216](#)
- ♦ [Section 20.7.2, “Linux Workload Registration and Discovery with Migrate Agent,” on page 217](#)

20.7.1 Windows Workload Registration and Discovery with Migrate Agent

- 1 Log in as Administrator to the source Windows workload.
- 2 Ensure that TCP port 443 is open on the workload.
- 3 Download Migrate Agent Utility for Windows. Save the `MigrateAgent.cli.exe` file to a convenient location on the workload.
See [“Migrate Agent Utility for Windows” on page 270](#).
- 4 In an Administrator Prompt, navigate to the location where you saved the file, then view the command Help by entering:

```
MigrateAgent.cli.exe help
```

- 5 Register the workload with the appropriate Migrate Server cloud instance. Enter

```
MigrateAgent.cli.exe /register /psserver=ps_dns_or_ipaddr <username> /  
password=<password>
```

Provide the credentials for an administrator-level user of the PlateSpin Migrate Server who has the permissions needed to add a workload. You can use the `/password=` option with the password, use the `-pwdfile=` option with a path to a file that contains the password, or do not specify the password in the command sequence. If you exclude the password from the command line, the script will prompt for it. The password is obscured as you type it and it does not appear in the process list.

For example:

```
Migrate.Agent.cli.exe /register /psserver=10.10.10.101 /username=jsmith /  
password=jspwd
```

NOTE: If you modify the public IP address of the Migrate Server, you must run the following command on each of the source Windows workloads that are configured for the server to modify the IP address.

```
MigrateAgent.cli.exe /config /setting=psserver:<new-ps-dns-or-ipaddr>
```

For example:

```
MigrateAgent.cli.exe /config /setting=psserver:10.10.20.202
```

20.7.2 Linux Workload Registration and Discovery with Migrate Agent

The Migrate Agent Utility for Linux requires the source machine to have GNU C Library (glibc) 2.11.3 or higher installed.

- 1 Log in as the `root` user to the source Linux workload.
- 2 Ensure that TCP port 443 is open on the workload.
- 3 Download the Migrate Agent Utility for Linux. Extract the downloaded file to the `/MigrateAgent` directory,

See [“Migrate Agent Utility for Linux” on page 272](#).

- 4 In a terminal, navigate to the `/MigrateAgent` directory, then view the command Help by entering:

```
./MigrateAgent -h
```

- 5 Register the workload with the appropriate Migrate Server cloud instance. Enter

```
./MigrateAgent register [-h] <ps_dns_or_ipaddr> <ps_username> [[-p  
<user_password>] | [-pf <passwordfile_path>]]
```

Specify the IP address or DNS name of the PlateSpin Migrate Server instance in the cloud. Provide the credentials for an administrator-level user of the PlateSpin Migrate Server who has the permissions needed to add a workload. You can use the `-p` option with the password, use the `-pf` option with a path to a file that contains the password, or do not specify the password in the command sequence. If you exclude the password from the command line, the script will prompt for it. The password is obscured as you type it and it does not appear in the process list.

For example:

```
./MigrateAgent register 10.10.10.101 jsmith -p jspwd
```

NOTE: If you modify the public IP address of the Migrate Server, you must run the following command on each of the source Linux workloads that are configured for the server to modify the IP address.

```
./MigrateAgent configure <ps_dns_or_ipaddr> <new-ps-dns-or-ipaddr>
```

For example:

```
./MigrateAgent configure 10.10.10.101 10.10.20.202
```

20.8 Refreshing Source Workload Details

If you make changes on the source workload before the migration begins, you might need to re-discovery the workload details. In the Migrate Client, you can refresh discovery details. In the Migrate Web Interface, you must remove and re-add the workload.


- ♦ [Section 20.8.1, “Refresh Workload Details in Migrate Client,” on page 218](#)
- ♦ [Section 20.8.2, “Removing and Re-Adding Workloads in the Web Interface,” on page 218](#)

20.8.1 Refresh Workload Details in Migrate Client

PlateSpin Migrate Client allows you to refresh workload details. You should routinely refresh your source workloads and targets before setting up a migration job.

To refresh a source workload details:

- 1 In the Servers view, right-click the required item, then select **Refresh Details**.



- 2 Specify the credentials appropriate for the system being refreshed, then click **Refresh**.
PlateSpin Migrate starts a discovery job, which you can monitor in the Jobs view.

20.8.2 Removing and Re-Adding Workloads in the Web Interface

PlateSpin Migrate Web Interface does not support refreshing details for the discovered workloads. To update details about a discovered workload, you must remove the workload, and then add and discover its details again. Configuration details are lost if the workload is in a configured state when you remove it. If a migration license is in use, it is removed from the workload and returned to the license pool. For information about removing the workload, see [Section 20.10, “Undiscovering or Removing Source Workloads,”](#) on page 219.

20.9 Using Tags to Track Logical Associations of Workloads

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

For information about creating, modifying, or deleting workload tags, see [Section 7.2, “Managing Workload Tags,”](#) on page 127.

After you create tags, they are available at the bottom of the Edit Target Details page where you can assign a tag to the appropriate workloads. The Workloads page includes a **Tag** column where the single tag you associate with a workload is displayed. You can sort on this column to group similar workloads together. This enables you to easily locate and run operations on the tagged workloads at the same time.

NOTE: When you export a workload with a tag setting to a new server, the tag settings persist.

To associate a tag with a workload during Configure Migration:

- 1 In the Migrate Web Interface, click **Workloads**.
- 2 In the workload list, select the workload you want to tag and click **Configure Migration**.

- 3 Configure the workload.
- 4 In the Tag section at the bottom of the Edit Target Details page, select the tag name you want to associate with the workload
- 5 Click **Save**.

To add or modify a tag associated with configured workload:

- 1 In the Migrate Web Interface, click **Workloads**.
- 2 In the workload list, click the workload you want to tag to open the Target Details page.
- 3 Click **Edit**.
- 4 In the Tag section at the bottom of the Edit Target Details page, select the tag name you want to associate with the workload.
- 5 Click **Save**.

To disassociate a tag from a workload:

- 1 In the Migrate Web Interface, click **Workloads**.
- 2 In the workload list, select the workload for which you want to remove the tag and click **Configure Migration**.
- 3 In the Tag section of the configuration page, select the empty string and click **Save**.

20.10 Undiscovering or Removing Source Workloads

After you complete all migration jobs for a source workload and the cutover completes successfully, you can remove (undiscover) the source workload.

IMPORTANT

- ♦ Before you delete an object that is in use for configured jobs, you must ensure that all the affected jobs are completed.
- ♦ If block-level transfer is enabled, remove the block-based transfer driver from the source workload:
 - ♦ **Windows:** Select to uninstall the block-based transfer driver.
 - ♦ **Linux:** Manually uninstall the blkwatch driver from the source. See [Block-level data transfer software](#) in [Cleaning Up Linux Workloads](#).

A reboot of the source workload is required after the driver is removed.

To undiscover a workload through the Migrate Client:

- 1 On the Workloads page, right-click the workload object and select **Undiscover Server**.
- 2 (Block-level transfer) Remove the block-based driver from the source workload, then reboot the source workload.

To remove a workload through the Migrate Web Interface:

- 1 On the Workloads page, select the workload, then click **Remove Workload**.
- 2 (Block-level transfer) Remove the block-based driver from the source workload, then reboot the source workload.

21 Preparing Device Drivers

PlateSpin Analyzer ships with a library of device drivers, and during migration jobs it installs the appropriate drivers for the target. If you require specific drivers for your target infrastructure, you might need to add (upload) drivers to the PlateSpin Migrate driver database.

To determine if the required drivers are available for conversion of Windows workloads to physical machines, you can use the PlateSpin Analyzer function in PlateSpin Migrate Client. PlateSpin Analyzer can help identify missing or incompatible drivers. See [Section 21.5, “Analyzing Suitability of Discovered Windows Workloads For Conversion to Physical Machines,”](#) on page 229.

- ♦ [Section 21.1, “Packaging Device Drivers for Windows Systems,”](#) on page 221
- ♦ [Section 21.2, “Packaging Device Drivers for Linux Systems,”](#) on page 221
- ♦ [Section 21.3, “Uploading Drivers to the PlateSpin Migrate Device Driver Database,”](#) on page 222
- ♦ [Section 21.4, “Using the Plug and Play \(PnP\) ID Translator Feature,”](#) on page 224
- ♦ [Section 21.5, “Analyzing Suitability of Discovered Windows Workloads For Conversion to Physical Machines,”](#) on page 229

21.1 Packaging Device Drivers for Windows Systems

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

- 1 Prepare all interdependent driver files (*.sys, *.inf, *.dll, etc.) for your target infrastructure and device. If you have obtained manufacturer-specific drivers as a .zip archive or an executable, extract them first.
- 2 Save the driver files in separate folders, with a discrete folder per device.

The drivers are now ready for upload. See [Section 21.3, “Uploading Drivers to the PlateSpin Migrate Device Driver Database,”](#) on page 222.

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

- ♦ All 64-bit Windows systems
 - ♦ 32-bit versions of Windows Server 2008 and Windows 7
-

21.2 Packaging Device Drivers for Linux Systems

To package your Linux device drivers for uploading to the PlateSpin Migrate driver database, you can use a custom utility included in your Linux ISO boot image.

- 1 Find a Linux workstation that has the same kernel version as the source machine. Source machine itself is one of the best choices. On the Linux workstation, create a directory for your device driver files. All the drivers in the directory must be for the same kernel and architecture.
- 2 Download the boot image and mount it.

For example, assuming that the ISO has been copied under the `/root` directory, issue these commands:

```
# mkdir /mnt/ps bootofx.x2p.iso
# mount -o loop /root/ /mnt/ps
```

- 3 From the `/tools` subdirectory of the mounted ISO image, copy the `packageModules.tar.gz` archive into a another working directory and extract it.

For example, with the `.gz` file is inside your current working directory, issue this command:

```
tar -xvzf packageModules.tar.gz
```

- 4 Enter the working directory and execute the following command:

```
./PackageModules.sh -d <path_to_driver_dir> -o <package name>
```

Replace `<path_to_driver_dir>` with the actual path to the directory where you saved you driver files, and `<package name>` with the actual package name, using the following format:

Drivename-driverversion-dist-kernelversion-arch.pkg

For example, `bnx2x-1.48.107-RHEL4-2.6.9-11.EL-i686.pkg`

The package is now ready for upload. See [Section 21.3, “Uploading Drivers to the PlateSpin Migrate Device Driver Database,”](#) on page 222.

21.3 Uploading Drivers to the PlateSpin Migrate Device Driver Database

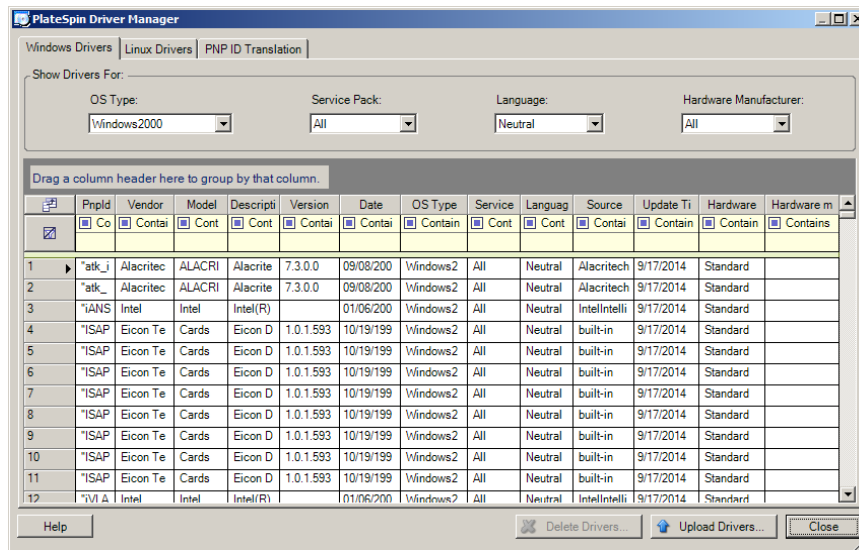
Use the PlateSpin Driver Manager to upload device drivers to the driver database.

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


- ♦ [Section 21.3.1, “Device Driver Upload Procedure \(Windows\),”](#) on page 222
- ♦ [Section 21.3.2, “Device Driver Upload Procedure \(Linux\),”](#) on page 223

21.3.1 Device Driver Upload Procedure (Windows)

- 1 Obtain and prepare the required device drivers.
See [Packaging Device Drivers for Windows Systems](#).
- 2 Click **Tools > Manage Device Drivers** and select the **Windows Drivers** tab:



3 Click **Upload Drivers**.



4 Select the Hardware Manufacturer.

For most X2P migrations, select **Standard** as the **Hardware Manufacturer** option, unless your drivers are designed specifically for any of the target environments listed.

5 Select the Storage Type.

IMPORTANT: If you select the **Storage Type** as **FCoE**, then you must ensure that all the drivers applicable for the FCoE storage device are in the same folder.

6 Browse to the folder that contains the required driver files, and select applicable OS type, language, and hardware manufacturer options

7 Click **Upload** and confirm your selections when prompted.


The system uploads the selected drivers to the driver database.

21.3.2 Device Driver Upload Procedure (Linux)

1 Obtain and prepare the required device drivers.

See [Packaging Device Drivers for Linux Systems](#).

- 2 Click **Tools > Manage Device Drivers** and select the **Linux Drivers** tab:



- 3 Click **Upload Drivers**, browse to the folder that contains the required driver package (*.pkg), and click **Upload All Drivers**.

The system uploads the selected drivers to the driver database.

21.4 Using the Plug and Play (PnP) ID Translator Feature


“Plug and Play” (PnP) refers to Windows operating system functionality that supports connectivity, configuration, and management with native plug and play devices. In Windows, the feature facilitates discovery of PnP compliant hardware devices attached to a PnP compliant bus. PnP compliant devices are assigned a set of Device Identification Strings by their manufacturer. These strings are programmed into the device when it is built. These strings are fundamental to how PnP works: they are part of the Windows' information source used to match the device with a suitable driver.

When the PlateSpin Server discovers workloads and their available hardware, the discovery includes these PnP IDs and the storage of that data as part of the workload's details. PlateSpin uses the IDs to determine which, if any, drivers need to be injected during a conversion operation. The PlateSpin Server maintains a database of PnP IDs for the associated drivers of each of the supported operating systems. Because Windows and Linux use different formats for PnP IDs, a Windows workload discovered by the Migrate Linux RAM disk contains Linux-style PnP IDs.


These IDs are formatted consistently, so PlateSpin can apply a standard transformation to each of them to determine its corresponding Windows PnP ID. The translation occurs automatically within the PlateSpin product. The feature enables you or a support technician to add, edit or remove custom PnP mappings.

Follow these steps to use the PnP ID Translation feature:


- 1 Launch the PlateSpin Driver Manager tool and connect to the PlateSpin Server.
- 2 In the Driver Manager tool, select the PNP ID Translation tab to open the **PNP ID Translation** list, which includes the currently known custom PnP ID mappings.




- 3 On the list page, click **Add** to display the Create PNP ID Mapping dialog box.



- 4 In the **Linux PNP ID** field, add a Linux PnP ID.
 - 4a (Conditional) If you know it, type the Linux PnP ID you want to use.
or
 - 4b (Conditional) Select an ID from a previously discovered workload:
 - 4b1 Adjacent to the **Linux PnP ID** field, click **Select** to open the Select Linux PnP ID dialog box.



- 4b2** On the dialog box, click **Select Machine** to display a list of the machines previously discovered by the PlateSpin Linux RAM disk.
- 4b3** Highlight one of the devices in the list, then click **Select** to populate the list in the Select Linux PnP ID dialog box.



- 4b4** Select a device on the list, then click **OK** to apply the standard transformation to the PnP ID and display it in the Create PnP ID Mapping dialog box.

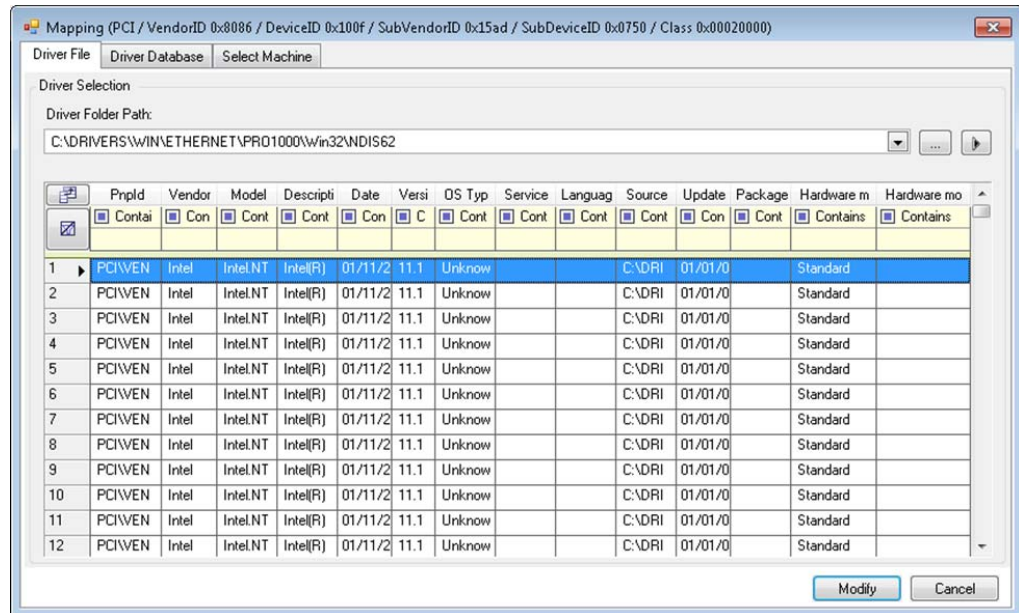
5 In the **Windows PNP ID** field, add a Windows PnP ID:

5a (Conditional) If you know it, type the Windows PnP ID you want to use.

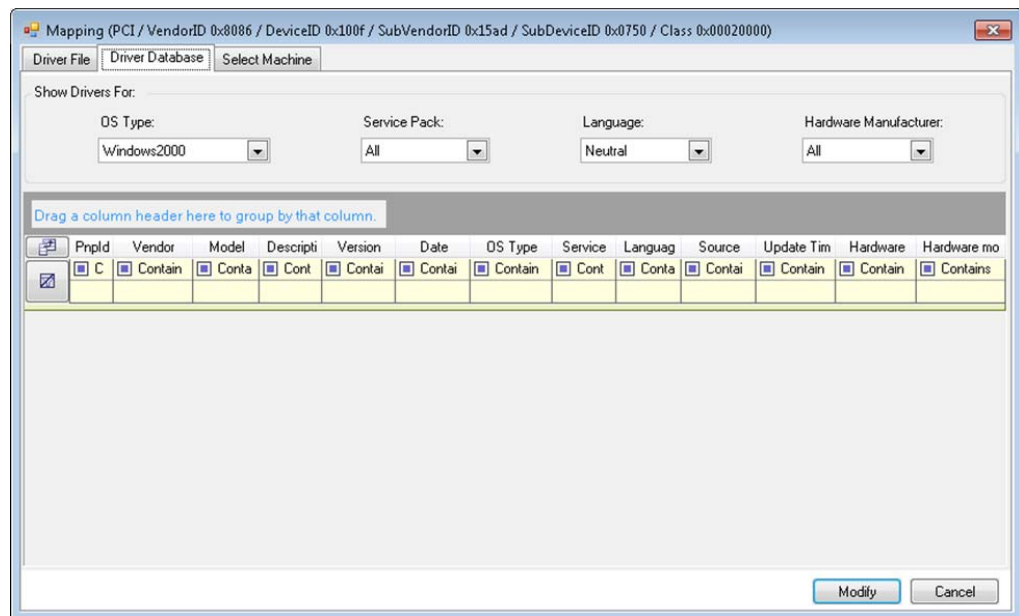
or

5b (Conditional) Adjacent to the **Windows PNP ID** field, click **Select** to open a mapping tool that presents three methods for helping you map a the Windows PnP ID:

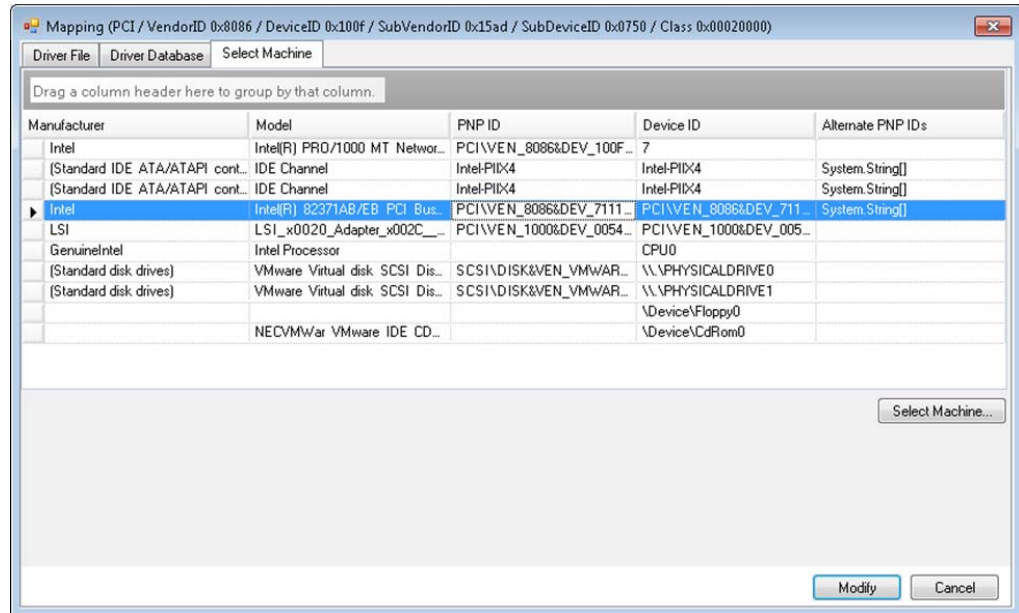
- Under the **Driver File** tab, browse to and select a Windows driver file (that is, a file with the *.inf extension), select the desired PnP ID, then click **Modify**.



- Under the **Driver Database** tab, browse to and select the existing driver database, select the correct PnP ID, then select **Modify**.




- Under the **Select Machine** tab, click **Select Machine**, then, from the list of Windows machines discovered using live discovery, select a machine, click **OK** to display its devices, select the desired PnP ID, then click **Modify**.



IMPORTANT: Selecting a Windows PnP ID that does not have an associated driver package installed might result in a failure at conversion time.

- In the Create PnP Id Mapping dialog box, confirm that the correct Linux PnP ID and the correct Windows PnP are selected, then click **OK** to display the PNP ID Translation page of the PlateSpin Driver Manager.



- 7 (Optional) To modify or remove the mapping in the PNP ID Translation list, select the mapping pattern, then click **Remove** or **Modify**, depending on the operation you want to perform.

Remove simply deletes the mapping (after displaying a confirmation dialog box).

To modify,

- 7a Click **Modify** to open the Create PNP id Mapping dialog box.
- 7b Repeat [Step 5](#) to modify the Windows PnP ID.

NOTE: You cannot select or modify the Linux PnP ID.

21.5 Analyzing Suitability of Discovered Windows Workloads For Conversion to Physical Machines

Before you begin any large-scale migration projects, you should identify potential migration problems and correct them beforehand. The PlateSpin Migrate Client provides the PlateSpin Analyzer utility to validate the following:

- ♦ Compatibility of target hardware for migration to physical targets
- ♦ Availability of drivers in the driver database for the physical server hardware
- ♦ Compatibility of source hardware for offline migration

NOTE: PlateSpin Analyzer currently supports only Windows workloads.

- ♦ [Section 21.5.1, “About PlateSpin Analyzer Tests and Results,” on page 230](#)
- ♦ [Section 21.5.2, “PlateSpin Analyzer in the Migrate Client,” on page 231](#)

21.5.1 About PlateSpin Analyzer Tests and Results

For target hardware support, PlateSpin Analyzer checks whether hardware drivers are in the driver repository for the following conversion types:

- ♦ Physical to physical (P2P)
- ♦ Image to physical (I2P)
- ♦ Virtual to physical (V2P)

[Table 21-1](#) describes the purpose of each test.

Table 21-1 *PlateSpin Analyzer Tests*

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

[Table 21-2](#) describes the status messages in the test results.

Table 21-2 *Status Messages in PlateSpin Analyzer Test Results*

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

For more information about using PlateSpin Analyzer and understanding the results, see [KB Article 7920478 \(https://www.netiq.com/support/kb/doc.php?id=7920478\)](#).

21.5.2 PlateSpin Analyzer in the Migrate Client

To open PlateSpin Analyzer:

- 1 On the **Tools** menu, click **Analyze Servers**.
The PlateSpin Analyzer window opens.
- 2 Select the required PlateSpin Migrate Network and the required machines to analyze.
- 3 (Optional) To reduce the analysis time, limit the scope of machines to a specific language.
- 4 (Optional) To analyze machines in the inventory of a different PlateSpin Server, click **Connect**, then specify the required PlateSpin Server URL and valid credentials.
- 5 Click **Analyze**.

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

Analyzed servers are listed in the left pane. Select a server to view test results in the right pane.

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

The **Test Results** tab provides the test results about a selected machine.

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

22 Preparing Linux Workloads for Migration

Perform the tasks in this section to prepare your Linux workloads for migration using PlateSpin Migrate

- ♦ [Section 22.1, “Verifying Block-Based Drivers for Linux,” on page 233](#)
- ♦ [Section 22.2, “Adding Drivers to the PlateSpin ISO Image,” on page 233](#)
- ♦ [Section 22.3, “Configuring LVM Snapshots for Linux Volume Replication,” on page 233](#)
- ♦ [Section 22.4, “Using Custom Freeze and Thaw Scripts for Linux Block-Level Migrations,” on page 234](#)
- ♦ [Section 22.5, “Preparing Paravirtualized Linux Source Workload,” on page 235](#)

22.1 Verifying Block-Based Drivers for Linux

Verify that a `blkwatch` module is available for the workload’s Linux distribution. For a list of preconfigured drivers, see [Appendix D, “Linux Distributions Supported by Migrate,” on page 251](#).

If you plan to protect a supported Linux workload that has a non-standard, customized, or newer kernel, rebuild the PlateSpin `blkwatch` module, which is required for block-level data replication.

See [Knowledgebase Article 7005873 \(https://www.netiq.com/support/kb/doc.php?id=7005873\)](https://www.netiq.com/support/kb/doc.php?id=7005873).

22.2 Adding Drivers to the PlateSpin ISO Image

The PlateSpin ISO image contains a large library of device drivers sufficient to boot most common targets. However, occasionally you might want to use your own, such as lesser-known, vendor-specific, or custom-developed drivers for Linux workloads.

You can modify the PlateSpin ISO image to add your vendor-specific or custom-developed drivers. See [Section G.3, “Injecting Additional Device Drivers into the PlateSpin ISO Image,” on page 280](#).

22.3 Configuring LVM Snapshots for Linux Volume Replication

We recommend that you prepare snapshots for block-level data transfer. Ensure that each volume group has sufficient free space for snapshots (at least 10% of the sum of all partitions). If snapshots are not available, PlateSpin Migrate locks and releases each block in turn on the source workload for data transfer.

The `blkwatch` driver leverages LVM snapshots if they are available. Copying blocks from the snapshot helps avoid potential open file conflicts.

For LVM storage, see [Knowledgebase Article 7005872 \(https://www.netiq.com/support/kb/doc.php?id=7005872\)](https://www.netiq.com/support/kb/doc.php?id=7005872).

22.4 Using Custom Freeze and Thaw Scripts for Linux Block-Level Migrations

For Linux workload migrations, PlateSpin Migrate supports the use of *freeze* and *thaw* shell scripts to provide an additional means of control over your Linux block-level migration process.

Migrate executes these scripts during Linux workload migrations, at the beginning and end of block-level data transfer sessions. Specifically, they interject in the migration process in the following fashion:

1. First pass of all volumes without snapshots:
 - ♦ Regular (non-LVM) volumes
 - ♦ LVM without enough space to take a snapshot
2. *Freeze script*
3. Take snapshots
4. Second pass of all non-snapshot volumes
5. *Thaw script*
6. Transfer volume snapshots

You can use this capability to complement the automated daemon control feature provided through the user interface. See [Section 26.15, “Services or Daemons to Stop before Replication or Cutover,” on page 304](#).

For example, you might want to use this feature to cause an application to flush its data to disk so that the workload remains in a more consistent state during a Live Transfer migration.

To use the feature, do the following before setting up your migration job:

- 1 Create the following files:
 - ♦ `platespin.freeze.sh` is a shell script to contain the `freeze` logic.
 - ♦ `platespin.thaw.sh` is a shell script to contain the `thaw` logic.
 - ♦ `platespin.conf` is a text file that defines any required arguments, along with a timeout value.

The required format for the contents of the `platespin.conf` file is:

```
[ServiceControl]
```

```
(optional) FreezeArguments=<arguments>
```

```
(optional) ThawArguments=<arguments>
```

```
(optional) TimeOut=<timeout>
```

Replace `<arguments>` with the required command arguments, separated by a space, and `<timeout>` with a timeout value in seconds. If unspecified, the default timeout is used (60 seconds).

- 2 Save the scripts, along with the `.conf` file, on your Linux source workload, in the following directory:
`/etc/platespin/`

22.5 Preparing Paravirtualized Linux Source Workload

Before you migrate a paravirtualized Linux source workload running on Citrix XenServer or KVM to a target platform as fully virtualized guest, do the following:

- ♦ Ensure that both the paravirtualized kernel and the standard kernel are installed on the paravirtualized source workload.
- ♦ Manually compile the block-based drivers for Xen kernel.
- ♦ Use block-based migration.

See [“Paravirtualized Source Workloads” on page 37](#).

23 Preparing for Migration of Windows Clusters

You can migrate Microsoft Windows Cluster business services to a target VMware vCenter container or to a physical machine. For information about supported Microsoft Windows Clusters, see “Clusters” in [Section 2.1.1, “Supported Source Workloads For Migration to Non-Cloud Platforms,”](#) on page 23.

You can use PlateSpin Migrate Client or PlateSpin Migrate Web Interface to migrate Windows Clusters to VMware vCenter containers. You can also use PlateSpin Migrate Client to migrate Windows Clusters to physical machines. The prerequisites for migration are the same.

NOTE: The Windows cluster management software provides the failover and failback control for the resources running on its cluster nodes. This document refers to this action as a *cluster node failover* or a *cluster node failback*.

- [Section 23.1, “Planning Your Cluster Workload Migration,”](#) on page 237
- [Section 23.2, “Configuring Windows Active Node Discovery,”](#) on page 242
- [Section 23.3, “Configuring the Block-Based Transfer Method for Clusters,”](#) on page 243
- [Section 23.4, “Adding Resource Name Search Values,”](#) on page 243
- [Section 23.5, “Quorum Arbitration Timeout,”](#) on page 244
- [Section 23.6, “Setting Local Volume Serial Numbers,”](#) on page 244
- [Section 23.7, “PlateSpin Cutover,”](#) on page 244
- [Section 23.8, “PlateSpin Cluster Migration,”](#) on page 245
- [Section 23.9, “Migrating Windows Clusters with the Web Interface,”](#) on page 245
- [Section 23.10, “Migrating Windows Clusters with the Migrate Client,”](#) on page 245

23.1 Planning Your Cluster Workload Migration

When active node discovery is enabled (the default) for the PlateSpin environment, migration of a Windows cluster is achieved through incremental replications of changes on the active node streamed to a virtual one node cluster. If you disable active node discovery, each node of a Windows cluster can be discovered and migrated as a standalone node.

Before you configure Windows clusters for migration, ensure that your environment meets the prerequisites and that you understand the conditions for migrating cluster workloads.

- [Section 23.1.1, “Requirements for Cluster Migration,”](#) on page 238
- [Section 23.1.2, “Block-Based Transfer for Clusters,”](#) on page 239
- [Section 23.1.3, “Impact of Cluster Node Failover on Replication,”](#) on page 240
- [Section 23.1.4, “Cluster Node Similarity,”](#) on page 241
- [Section 23.1.5, “Migration Setup,”](#) on page 242

23.1.1 Requirements for Cluster Migration

The scope of support for cluster migration is subject to the conditions described in [Table 23-1](#). Consider these requirements when you configure migration for clusters in your PlateSpin environment.

Table 23-1 Cluster Migration Requirements

Requirement	Description
Discover the active node as a Windows Cluster	<p>The PlateSpin global configuration setting <code>DiscoverActiveNodeAsWindowsCluster</code> determines whether Windows clusters are migrated as clusters or as separate standalone machines:</p> <ul style="list-style-type: none">♦ True (Default): The active node is discovered as a Windows cluster.♦ False: Individual nodes can be discovered as standalone machines. <p>See Section 23.2, “Configuring Windows Active Node Discovery,” on page 242.</p>
Resource name search values	<p>The PlateSpin global configuration setting <code>MicrosoftClusterIPAddressNames</code> determines the cluster resource names that can be discovered in your PlateSpin environment. You must configure search values that help to differentiate the name of the shared Cluster IP Address resource from the name of other IP address resources on the cluster.</p> <p>See Section 23.4, “Adding Resource Name Search Values,” on page 243.</p>
Windows Cluster Mode	<p>The PlateSpin global configuration setting <code>WindowsClusterMode</code> determines the method of block-based data transfer for incremental replications:</p> <ul style="list-style-type: none">♦ Default: Driverless synchronization.♦ SingleNodeBBT: Driver-based block-based transfer. <p>See the following:</p> <ul style="list-style-type: none">♦ “Block-Based Transfer for Clusters” on page 239♦ “Configuring the Block-Based Transfer Method for Clusters” on page 243
Active node host name or IP address	<p>You must specify the host name or IP address of the cluster’s active node when you perform an Add Workload operation. Because of security changes made by Microsoft, Windows clusters can no longer be discovered by using the virtual cluster name (that is, the shared cluster IP address).</p>
Resolvable host name	<p>The PlateSpin Server must be able to resolve the host name of each of the nodes in the cluster by their IP address.</p> <p>NOTE: DNS forward lookup and reverse lookup are required to resolve the host name by its IP address.</p>
Quorum resource	<p>A cluster’s quorum resource must be co-located on the node with the cluster’s resource group (service) being migrated.</p>
Similarity of cluster nodes	<p>In the default Windows Cluster Mode, driverless sync can continue from any node that becomes active if the nodes are similar. If they do not match, replications can occur only on the originally discovered active node.</p> <p>See “Cluster Node Similarity” on page 241.</p>

Requirement	Description
PowerShell 2.0	Windows PowerShell 2.0 must be installed on each node of the cluster.

23.1.2 Block-Based Transfer for Clusters

Block-based transfer for clusters works differently than for standalone servers. The initial replication either makes a complete copy (full) or uses a driverless synchronization method performed on the active node of the cluster. Subsequent incremental replications can use a driverless method or driver-based method for block-based data transfer.

NOTE: PlateSpin Migrate does not support file-based transfer for clusters.

The PlateSpin global configuration setting `WindowsClusterMode` determines the method of block-based data transfer for incremental replications:

- ♦ **Default:** Driverless synchronization.
- ♦ **SingleNodeBBT:** Driver-based block-based transfer. Use only with Fibre Channel SANs.

WARNING: Do not attempt to use SingleNodeBBT on clusters with shared iSCSI drives. It renders the cluster unusable.

Table 23-2 describes and compares the two methods.

Table 23-2 Comparison Block-Based Data Transfer Methods for Incremental Replication

Consideration	Default BBT	Single-Node BBT
Data transfer method	Uses driverless synchronization with an MD5-based replication on the currently active node.	Uses a BBT driver installed on the originally discovered active node.
Performance	Potentially slow incremental replications.	Significantly improves performance for incremental replications.
Drivers	<ul style="list-style-type: none"> ♦ No BBT driver to install. ♦ No reboot is required on the source cluster nodes. 	<ul style="list-style-type: none"> ♦ Use the Migrate Agent utility to install a BBT driver on the originally discovered active node of the cluster. ♦ Reboot the node to apply the driver. This initiates a failover to another node in the cluster. After the reboot, make the originally discovered node the active node again. ♦ The same node must remain active for replications to occur and to use single-node block-based transfer. ♦ After you install the BBT driver, either a full replication or a driverless incremental replication must occur before the driver-based incremental replications can begin.

Consideration	Default BBT	Single-Node BBT
Supported Windows Clusters	Works with any supported Windows Server clusters.	Works with Windows Server 2008 R2 and later clusters. Other supported Windows clusters use the driverless synchronization method for replication.
First incremental replication	Uses driverless sync on the active node.	Uses driver-based block-based transfer on the originally discovered active node if a full replication was completed after the BBT driver was installed. Otherwise, it uses driverless sync on the originally discovered active node.
Subsequent incremental replication	Uses driverless sync on the active node.	Uses driver-based block-based transfer on the originally discovered active node. If a cluster switches nodes, the driverless sync method is used for the first incremental replication after the originally active node becomes active again. See "Impact of Cluster Node Failover on Replication" on page 240.

23.1.3 Impact of Cluster Node Failover on Replication

[Table 23-3](#) describes the impact of cluster node failover on replication and the required actions for the Migrate administrator.

Table 23-3 *Impact of Cluster Node Failover on Replication*

Cluster Node Failover or Failback	Default BBT	Single-Node BBT
Cluster node failover occurs during the first full replication	Replication fails. The first full replication must complete successfully without a cluster node failover. <ol style="list-style-type: none"> 1. Remove the cluster from Migrate. 2. (Optional) Make the originally discovered active node the active node again. 3. Re-add the cluster using the active node. 4. Re-run the first full replication. 	

Cluster Node Failover or Failback	Default BBT	Single-Node BBT
Cluster node failover occurs during a subsequent full replication or a subsequent incremental replication	<p>The replication command aborts and a message displays indicating that the replication needs to be re-run.</p> <p>If the new active node's profile is similar to the failed active node, the migration contract remains valid.</p> <ol style="list-style-type: none"> 1. Re-run the replication on the now-active node. <p>If the new active node's profile is not similar to the failed active node, the migration contract is valid only on the originally active node.</p> <ol style="list-style-type: none"> 1. Make the originally discovered active node the active node again. 2. Re-run the replication on the active node. 	<p>The replication command aborts and a message displays indicating that the replication needs to be re-run. The migration contract is valid only on the originally discovered active node.</p> <ol style="list-style-type: none"> 1. Make the originally discovered active node the active node again. 2. Re-run the replication on the active node. <p>This first incremental replication after a cluster failover/failback event automatically uses driverless sync. Subsequent incremental replications will use the block-based driver as specified by single-node BBT.</p>
Cluster node failover occurs between replications	<p>If the new active node's profile is similar to the failed active node, the migration contract continues as scheduled for the next incremental replication. Otherwise, the next incremental replication command fails.</p> <p>If a scheduled incremental replication fails:</p> <ol style="list-style-type: none"> 1. Make the originally discovered active node the active node again. 2. Run an incremental replication. 	<p>Incremental replication fails if the active node switches between replications.</p> <ol style="list-style-type: none"> 1. Ensure that the originally discovered active node is again the active node. 2. Run an incremental replication. <p>This first incremental replication after a cluster failover/failback event automatically uses driverless sync. Subsequent incremental replications will use the block-based driver as specified by single-node BBT.</p>

23.1.4 Cluster Node Similarity

In the default Windows Cluster Mode, the cluster nodes must have similar profiles to prevent interruptions in the replication process. The profiles of cluster nodes are considered similar if all of the following conditions are met:

- ♦ Serial numbers for the nodes' local volumes (System volume and System Reserved volume) must be the same on each cluster node.

NOTE: Use the customized *Volume Manager* utility to change the local volume serial numbers to match each node of the cluster. See [“Synchronizing Serial Numbers on Cluster Node Local Storage” on page 267](#).

If the local volumes on each node of the cluster have different serial numbers, you cannot run a replication after a cluster node failover occurs. For example, during a cluster node failover, the active node Node 1 fails, and the cluster software makes Node 2 the active node. If the local drives on the two nodes have different serial numbers, the next replication command for the workload fails.

- ♦ The nodes must have the same number of volumes.
- ♦ Each volume must be exactly the same size on each node.
- ♦ The nodes must have an identical number of network connections.

23.1.5 Migration Setup

To configure migration for a Windows cluster, follow the normal workload migration workflow. Ensure that you provide the host name or IP address of the cluster's active node.

23.2 Configuring Windows Active Node Discovery

You can discover Windows Server clusters as clusters or as individual standalone machines, depending on the PlateSpin global configuration setting `DiscoverActiveNodeAsWindowsCluster`.

To discover Windows clusters as clusters, set the `DiscoverActiveNodeAsWindowsCluster` parameter to `True`. This is the default setting. Cluster discovery, inventory, and workload migration use the host name or IP address of a cluster's active node, instead of using its cluster name and an administration share. You do not configure separate workloads for the cluster's non-active nodes. For other cluster workload migration requirements, see ["Requirements for Cluster Migration" on page 238](#).

To discover all Windows clusters as individual standalone machines, set the `DiscoverActiveNodeAsWindowsCluster` parameter to `False`. This setting allows the PlateSpin Server to discover all nodes in a Windows failover cluster as standalone machines. That is, it inventories a cluster's active node and non-active nodes as a regular, cluster-unaware Windows workloads.

To enable or disable cluster discovery:

- 1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/PlateSpinConfiguration/`
Replace `Your_PlateSpin_Server` with the DNS host name or IP address of your PlateSpin Migrate Server.
- 2 Search for `DiscoverActiveNodeAsWindowsCluster`, then click **Edit**.
- 3 In the **Value** field, select **True** to enable cluster discovery, or select **False** to disable cluster discovery.
- 4 Click **Save**.

23.3 Configuring the Block-Based Transfer Method for Clusters

Incremental replications for Windows clusters can use a driverless method (Default) or driver-based method (SingleNodeBBT) for block-based data transfer, depending on the PlateSpin global configuration setting `WindowsClusterMode`. For more information, see [“Block-Based Transfer for Clusters” on page 239](#).

To configure `WindowsClusterMode`:

- 1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/PlateSpinConfiguration/`
Replace *Your_PlateSpin_Server* with the DNS host name or IP address of your PlateSpin Migrate Server.
- 2 Search for `WindowsClusterMode`, then click **Edit**.
- 3 In the **Value** field, select **Default** to use driverless synchronization for incremental replication, or select **SingleNodeBBT** to use block-based drivers for incremental replication.
- 4 Click **Save**.

23.4 Adding Resource Name Search Values

To help identify the active node in a Windows failover cluster, PlateSpin Migrate must differentiate the name of the shared Cluster IP Address resource from the names of other IP address resources on the cluster. The shared Cluster IP Address resource resides on the cluster's active node.

The global parameter `MicrosoftClusterIPAddressNames` on the PlateSpin Server Configuration page contains a list of search values to use in discovery for a Windows cluster workload. When you add a Windows cluster workload, you must specify the IP address of the cluster's currently active node. PlateSpin Migrate searches the names of the cluster's IP address resources on that node to find one that *starts with* the specified characters of any value in the list. Thus, each search value must contain enough characters to differentiate the shared Cluster IP Address resource on a specific cluster, but it can be short enough to apply to discovery in other Windows clusters.

For example, a search value of `Clust IP Address` or `Clust IP` matches the resource names *Clust IP Address* for 10.10.10.201 and *Clust IP Address* for 10.10.10.101.

The default name for the shared Cluster IP Address resource is `Cluster IP Address` in English, or the equivalent if the cluster node is configured in another language. The default search values in the `MicrosoftClusterIPAddressNames` list include the resource name `Cluster IP Address` in English and each of the [supported languages](#).

Because the resource name of the shared Cluster IP Address resource is user-configurable, you must add other search values to the list, as needed. If you change the resource name, you must add a related search value to the `MicrosoftClusterIPAddressNames` list. For example, if you specify a resource name of `Win2012-CLUS10-IP-ADDRESS`, you should add that value to the list. If you have multiple clusters using the same naming convention, an entry of `Win2012-CLUS` matches any resource name that starts with that sequence of characters.

To add search values in the `MicrosoftClusterIPAddressNames` list:

- 1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/PlateSpinConfiguration/`
Replace *Your_PlateSpin_Server* with the DNS host name or IP address of your PlateSpin Migrate Server.
- 2 Search for `MicrosoftClusterIPAddressNames`, then click **Edit**.
- 3 In the **Value** field, add one or more search values to the list.
- 4 Click **Save**.

23.5 Quorum Arbitration Timeout

You can set the `QuorumArbitrationTimeMax` registry key for Windows Server failover clusters in your PlateSpin environment by using the global parameter `FailoverQuorumArbitrationTimeout` on the PlateSpin Server Configuration page. The default timeout is 60 seconds, in keeping with the Microsoft default value for this setting. See *QuorumArbitrationTimeMax* (<https://msdn.microsoft.com/en-us/library/aa369123%28v=vs.85%29.aspx?f=255&MSPPErr=-2147217396>) on the Microsoft Developer Network website. The specified timeout interval is honored for quorum arbitration at failover and fallback.

To set the quorum arbitration timeout for all Windows failover clusters:

- 1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/PlateSpinConfiguration/`
Replace *Your_PlateSpin_Server* with the DNS host name or IP address of your PlateSpin Migrate Server.
- 2 Search for `FailoverQuorumArbitrationTimeout`, then click **Edit**.
- 3 In the **Value** field, specify the maximum number of seconds to allow for quorum arbitration.
- 4 Click **Save**.

23.6 Setting Local Volume Serial Numbers

You can use the *Volume Manager* utility to change the local volume serial numbers to match in each node of the cluster. See “[Synchronizing Serial Numbers on Cluster Node Local Storage](#)” on page 267.

23.7 PlateSpin Cutover

When the PlateSpin cutover operation is complete and the virtual one-node cluster comes online, you see a multi-node cluster with one active node (all other nodes are unavailable).

To perform a PlateSpin cutover (or to test the PlateSpin cutover on) a Windows cluster, the cluster must be able to connect to a domain controller. To leverage the test failover functionality, you need to migrate the domain controller along with the cluster. During the test, bring up the domain controller, followed by the Windows cluster workload (on an isolated network).

23.8 PlateSpin Cluster Migration

A PlateSpin cluster migration operation requires a full replication for Windows Cluster workloads. After PlateSpin cluster migration is complete, you must reattach the shared storage and rebuild the cluster environment before you can rejoin additional nodes to the newly restored cluster.

For information about rebuilding the cluster environment after a PlateSpin migration, see [Knowledgebase Article 7016770](#).

After PlateSpin cluster migration is complete for a Windows Server 2003 or Windows Server 2003 R2 cluster, you must restart the cluster service on the target.

23.9 Migrating Windows Clusters with the Web Interface

Use the PlateSpin Migrate Web Interface to add the active node by specifying the IP address of the active node, configure migration for the active node, and then run the migration.

23.10 Migrating Windows Clusters with the Migrate Client

You can use a **Move** job to migrate the essential services of a cluster that results in a functional single-node cluster in a virtual machine in VMware or a physical machine.

The workflow of migrating a Windows cluster is similar to that of migrating a standalone server:

- 1 Discover the active node by specifying the IP address of the active node.
- 2 In the Servers view, use drag-and-drop to start a migration job, then configure the job's parameters.
- 3 (Conditional: successful migration) If the migration job completes successfully, perform a [Server Sync operation](#) on the active node.

NOTE: If the active node in the cluster fails over before you can perform a Server Sync operation, perform a full migration using the *new* active node, and then perform a Server Sync on this new node.

- 4 (Conditional: failover prior to migration) If a cluster failover occurs prior to the completion of file transfer, the migration job aborts. If this happens, refresh the source and retry the migration job.

NOTE: If you select **Shut down** for the source's post-migration end state, a shutdown of all source nodes of the cluster results.

C Troubleshooting Discovery

Table C-1 provides information to help you troubleshoot common problems that might occur during workload discovery or target discovery.

- ♦ Section C.1, “Common Discovery Issues and Solutions,” on page 247
- ♦ Section C.2, “Modifying the OFX Controller Heartbeat Startup Delay (Windows Workloads),” on page 249

C.1 Common Discovery Issues and Solutions

Table C-1 Common Issues and Solutions Related to Discovery Operations

Problems or Messages	Solutions
Application has generated an error occurs during registration of physical server	<p>This error occurs if the physical server is unable to contact the PlateSpin Server. A common cause is incorrect information entered during the registration process.</p> <p>To restart the registration process:</p> <ol style="list-style-type: none">1. Enter <code>RegisterMachine.bat</code>.2. Ping to confirm basic connectivity with the PlateSpin Server.
My physical server has completed the registration process, but is not seen in PlateSpin Migrate Client.	<p>The full registration process can take some time to complete. After the second command prompt window has closed on the physical server, wait a few minutes before clicking the Refresh button in PlateSpin Migrate Client.</p>
Problems discovering source and target servers	<p>KB Article 7920291 (https://www.netiq.com/support/kb/doc.php?id=7920291) contains troubleshooting checklists for discovering the following:</p> <ul style="list-style-type: none">♦ Linux servers and VMware ESX Servers♦ Windows-based source and target servers <p>The article also has instructions for troubleshooting WMI connections and checking if DCOM is enabled.</p>
Package <...> Not Found occurs during discovery of existing Windows servers	<p>Check for required IIS configuration and network settings.</p> <p>See “Installing Prerequisite Software” in the <i>PlateSpin Migrate 12.2.1 Installation and Upgrade Guide</i>.</p>

Problems or Messages	Solutions
Could not find file \\{servername}\admin\$\{randomID}.xml	<p>This error might occur on Windows Server 2003 hosts.</p> <p>In some cases, either of these troubleshooting steps addresses the issue:</p> <ul style="list-style-type: none"> ◆ Ensure that the Admin\$ share on the PlateSpin Server host is accessible. If not, enable it and try the discovery again. <p>- OR -</p> <ul style="list-style-type: none"> ◆ Do the following: <ol style="list-style-type: none"> 1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at: <code>https://Your_PlateSpin_Server/PlateSpinConfiguration/</code> 2. Locate and edit the ForceMachineDiscoveryUsingService entry and change it to True. 3. Save the value and retry the discovery.
Unable to connect neither to the SSH server running on <IP_address> nor to VMware Virtual Infrastructure web-services at <ip_address>/sdk	<p>This message has a number of possible causes:</p> <ul style="list-style-type: none"> ◆ The workload is unreachable. ◆ The workload does not have SSH running. ◆ The firewall is on and the required ports have not been opened. ◆ The workload's specific operating system is not supported. <p>For network and access requirements for a workload, see “Access and Communication Requirements across your Migration Network” on page 49</p>
Access denied	<p>This authentication problem indicates either an invalid user name or password. For information on proper workload access credentials, see Table 20-2, “Guidelines for Machine Type and Credentials for Source Workloads,” on page 212.</p>

Related KB Articles are listed in [Table C-2](#).

Table C-2 KB Articles for Discovery Issues

ID	Description
7920339 (https://www.netiq.com/support/kb/doc.php?id=7920339)	ERRMSG: Discovery fails with The request failed with HTTP status 407 message
7920862 (https://www.netiq.com/support/kb/doc.php?id=7920862)	ERRMSG: Recoverable Error: ControllerConnectionBroken during discovery
7920291 (https://www.netiq.com/support/kb/doc.php?id=7920291)	ERRMSG: Server details discovery problems
7021574 (https://www.netiq.com/support/kb/doc.php?id=7021574)	ERRMSG: X2P Target Discovery Failed: Linux job did not complete successfully
For more discovery-related TIDs in the Knowledgebase	Search on “discovery” for the PlateSpin Migrate product (https://www.netiq.com/support/kb/?q=discovery&product=PlateSpin_Migrate).

C.2 Modifying the OFX Controller Heartbeat Startup Delay (Windows Workloads)

To avoid discovery failures caused by timing issues, a default heartbeat startup delay of 15 seconds (15000 ms) is set on the OFX Controller. The setting is configurable by adding the `HeartbeatStartupDelayInMS` registry key on the source workload. This registry key is not configured by default.

To enable a heartbeat delay of shorter or longer duration:

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

Path for a 64-bit source workload:

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

Path for a 32-bit source workload:

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

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

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

- 4 Restart the source workload.

D Linux Distributions Supported by Migrate

PlateSpin Migrate software includes pre-compiled versions of the `blkwatch` driver for many non-debug Linux distributions (32-bit and 64-bit). This section includes the following information:

- ♦ [Section D.1, “Analyzing Your Linux Workload,” on page 251](#)
- ♦ [Section D.2, “Pre-compiled blkwatch Drivers for Linux Distributions,” on page 252](#)

D.1 Analyzing Your Linux Workload

Prior to determining whether PlateSpin Migrate has a `blkwatch` driver for your distribution, you need to learn more about the kernel of your Linux workload so that you can use it as a search term against the list of supported distributions. This section includes the following information:

- ♦ [Section D.1.1, “Determining the Release String,” on page 251](#)
- ♦ [Section D.1.2, “Determining the Architecture,” on page 251](#)

D.1.1 Determining the Release String

You can determine the release string of the kernel of your Linux workload by running the following command at the workload’s Linux terminal:

```
uname -r
```

For example, if you run `uname -r`, you might see the following output:

```
3.0.76-0.11-default
```

If you search the list of distributions, you see there are two entries that match this string:

- ♦ `SLES11SP3-GA-3.0.76-0.11-default-x86`
- ♦ `SLES11SP3-GA-3.0.76-0.11-default-x86_64`

The search results indicate that the product has drivers for both 32-bit (x86) and 64-bit (x86_64) architectures.

D.1.2 Determining the Architecture

You can determine the architecture of your Linux workload by running the following command at the workload’s Linux terminal:

```
uname -m
```

For example, if you run `uname -m`, you might see the following output:

```
x86_64
```

With this information, you can determine that the workload has 64-bit architecture.

D.2 Pre-compiled blkwatch Drivers for Linux Distributions

PlateSpin Migrate provides precompiled block-based Linux Kernel drivers, called block watch (blkwatch) drivers, for many non-debug Linux distributions. The driver must be built for the specific kernel running on the Linux system. You can search the [List of Distributions](#) to determine if the release string and architecture of your Linux workload kernel matches a supported distribution in the list. If you find your release string and architecture, PlateSpin Migrate has a pre-compiled version of the blkwatch driver.

If your search is unsuccessful, you can create a custom blkwatch driver by following the steps found in the [KB Article 7005873](https://www.netiq.com/support/kb/doc.php?id=7005873) (<https://www.netiq.com/support/kb/doc.php?id=7005873>).

Self-compiled drivers are supported only for the Linux major and minor kernel versions that appear in the [List of Distributions](#), or a patched version thereof. If the major and minor kernel version in the release string of your Linux workload kernel matches a major and minor kernel version in the list, your self-compiled driver will be supported.

D.2.1 List Item Syntax

Each item in the list is formatted using the following syntax:

```
<Distro>-<Patch>-<Kernel_Release_String>-<Kernel_Architecture>
```

So, for a SLES 9 SP1 distribution with a kernel release string of 2.6.5-7.139-bigsm for 32-bit (x86) architecture, the item is listed in a format like this:

```
SLES9-SP1-2.6.5-7.139-bigsm-x86
```

D.2.2 List of Distributions

Oracle Linux 6 U7

NOTE: Blkwatch drivers for kernel version 2.6.32-573 do not support incremental replication for workloads with LVM volumes. Update the kernel, then see RHEL 6 U7 for drivers for kernel 2.6.32-642.

```
OEL6-U7-2.6.32-573.el6.i686-x86
```

```
OEL6-U7-2.6.32-573.el6.x86_64-x86_64
```

```
OEL6-U7_UK-2.6.39-400.250.7.el6uek.i686-x86
```

```
OEL6-U7_UK-3.8.13-68.3.4.el6uek.x86_64-x86_64
```

Oracle Linux 6 U8

```
OEL6-U8-2.6.32-642.el6.i686-x86
```

```
OEL6-U8-2.6.32-642.el6.x86_64-x86_64
```

```
OEL6-U8_UK-2.6.39-400.278.2.el6uek.i686-x86
```

```
OEL6-U8_UK-4.1.12-37.4.1.el6uek.x86_64-x86_64
```

Oracle Linux 6 U9

```
OEL6-U9_UK-2.6.39-400.294.3.el6uek.i686-x86
```

```
OEL6-U9_UK-4.1.12-61.1.28.el6uek.x86_64-x86_64
```

Oracle Linux 7 GA

```
OEL7-GA-3.10.0-123.el7.x86_64-x86_64
```

```
OEL7-GA_UK-3.8.13-35.3.1.el7uek.x86_64-x86_64
```

Oracle Linux 7 U1

OEL7-U1-3.10.0-229.el7.x86_64-x86_64
OEL7-U1_U EK-3.8.13-55.1.6.el7uek.x86_64-x86_64

Oracle Linux 7 U2

OEL7-U2-3.10.0-327.el7.x86_64-x86_64
OEL7-U2_U EK-3.8.13-98.7.1.el7uek.x86_64-x86_64

Oracle Linux 7 U3

OEL7-U3_U EK-4.1.12-61.1.18.el7uek.x86_64-x86_64

Red Hat Enterprise Linux 4 GA

RHEL4-GA-2.6.9-5.EL-x86
RHEL4-GA-2.6.9-5.EL-x86_64
RHEL4-GA-2.6.9-5.ELhugemem-x86
RHEL4-GA-2.6.9-5.ELsmp-x86
RHEL4-GA-2.6.9-5.ELsmp-x86_64

Red Hat Enterprise Linux 4 U1

RHEL4-U1-2.6.9-11.EL-x86
RHEL4-U1-2.6.9-11.EL-x86_64
RHEL4-U1-2.6.9-11.ELhugemem-x86
RHEL4-U1-2.6.9-11.ELsmp-x86
RHEL4-U1-2.6.9-11.ELsmp-x86_64

Red Hat Enterprise Linux 4 U2

RHEL4-U2-2.6.9-22.EL-x86
RHEL4-U2-2.6.9-22.EL-x86_64
RHEL4-U2-2.6.9-22.ELhugemem-x86
RHEL4-U2-2.6.9-22.ELsmp-x86
RHEL4-U2-2.6.9-22.ELsmp-x86_64

Red Hat Enterprise Linux 4 U3

RHEL4-U3-2.6.9-34.EL-x86
RHEL4-U3-2.6.9-34.EL-x86_64
RHEL4-U3-2.6.9-34.ELhugemem-x86
RHEL4-U3-2.6.9-34.ELlargesmp-x86_64
RHEL4-U3-2.6.9-34.ELsmp-x86
RHEL4-U3-2.6.9-34.ELsmp-x86_64

Red Hat Enterprise Linux 4 U4

RHEL4-U4-2.6.9-42.EL-x86
RHEL4-U4-2.6.9-42.EL-x86_64
RHEL4-U4-2.6.9-42.ELhugemem-x86
RHEL4-U4-2.6.9-42.ELlargesmp-x86_64
RHEL4-U4-2.6.9-42.ELsmp-x86
RHEL4-U4-2.6.9-42.ELsmp-x86_64

Red Hat Enterprise Linux 4 U5

RHEL4-U5-2.6.9-55.EL-x86
RHEL4-U5-2.6.9-55.EL-x86_64
RHEL4-U5-2.6.9-55.ELhugemem-x86
RHEL4-U5-2.6.9-55.ELlargesmp-x86_64
RHEL4-U5-2.6.9-55.ELsmp-x86
RHEL4-U5-2.6.9-55.ELsmp-x86_64

Red Hat Enterprise Linux 4 U6

RHEL4-U6-2.6.9-67.EL-x86
RHEL4-U6-2.6.9-67.EL-x86_64
RHEL4-U6-2.6.9-67.ELhugemem-x86
RHEL4-U6-2.6.9-67.ELlargesmp-x86_64
RHEL4-U6-2.6.9-67.ELsmp-x86
RHEL4-U6-2.6.9-67.ELsmp-x86_64

Red Hat Enterprise Linux 4 U7

RHEL4-U7-2.6.9-78.EL-x86
RHEL4-U7-2.6.9-78.EL-x86_64
RHEL4-U7-2.6.9-78.ELhugemem-x86
RHEL4-U7-2.6.9-78.ELlargesmp-x86_64
RHEL4-U7-2.6.9-78.ELsmp-x86
RHEL4-U7-2.6.9-78.ELsmp-x86_64

Red Hat Enterprise Linux 4 U8

RHEL4-U8-2.6.9-89.EL-x86
RHEL4-U8-2.6.9-89.EL-x86_64
RHEL4-U8-2.6.9-89.ELhugemem-x86
RHEL4-U8-2.6.9-89.ELlargesmp-x86_64
RHEL4-U8-2.6.9-89.ELsmp-x86
RHEL4-U8-2.6.9-89.ELsmp-x86_64

Red Hat Enterprise Linux 4 U9

RHEL4-U9-2.6.9-100.EL-x86
RHEL4-U9-2.6.9-100.EL-x86_64
RHEL4-U9-2.6.9-100.ELhugemem-x86
RHEL4-U9-2.6.9-100.ELlargesmp-x86_64
RHEL4-U9-2.6.9-100.ELsmp-x86
RHEL4-U9-2.6.9-100.ELsmp-x86_64

Red Hat Enterprise Linux 5 GA

RHEL5-GA-2.6.18-8.el5-x86
RHEL5-GA-2.6.18-8.el5-x86_64
RHEL5-GA-2.6.18-8.el5PAE-x86

Red Hat Enterprise Linux 5 U1

RHEL5-U1-2.6.18-53.el5-x86
RHEL5-U1-2.6.18-53.el5-x86_64
RHEL5-U1-2.6.18-53.el5PAE-x86

Red Hat Enterprise Linux 5 U2

RHEL5-U2-2.6.18-92.el5-x86
RHEL5-U2-2.6.18-92.el5-x86_64
RHEL5-U2-2.6.18-92.el5PAE-x86

Red Hat Enterprise Linux 5 U3

RHEL5-U3-2.6.18-128.el5-x86
RHEL5-U3-2.6.18-128.el5-x86_64
RHEL5-U3-2.6.18-128.el5PAE-x86

Red Hat Enterprise Linux 5 U4

RHEL5-U4-2.6.18-164.el5-x86
RHEL5-U4-2.6.18-164.el5-x86_64

RHEL5-U4-2.6.18-164.el5PAE-x86

Red Hat Enterprise Linux 5 U5

RHEL5-U5-2.6.18-194.el5-x86

RHEL5-U5-2.6.18-194.el5-x86_64

RHEL5-U5-2.6.18-194.el5PAE-x86

Red Hat Enterprise Linux 5 U6

RHEL5-U6-2.6.18-238.el5-x86

RHEL5-U6-2.6.18-238.el5-x86_64

RHEL5-U6-2.6.18-238.el5PAE-x86

Red Hat Enterprise Linux 5 U7

RHEL5-U7-2.6.18-274.el5-x86

RHEL5-U7-2.6.18-274.el5-x86_64

RHEL5-U7-2.6.18-274.el5PAE-x86

Red Hat Enterprise Linux 5 U8

RHEL5-U8-2.6.18-308.el5-x86

RHEL5-U8-2.6.18-308.el5-x86_64

RHEL5-U8-2.6.18-308.el5PAE-x86

Red Hat Enterprise Linux 5 U9

RHEL5-U9-2.6.18-348.el5-x86

RHEL5-U9-2.6.18-348.el5-x86_64

RHEL5-U9-2.6.18-348.el5PAE-x86

Red Hat Enterprise Linux 5 U10

RHEL5-U10-2.6.18-371.el5-x86

RHEL5-U10-2.6.18-371.el5-x86_64

RHEL5-U10-2.6.18-371.el5PAE-x86

Red Hat Enterprise Linux 5 U11

RHEL5-U11-2.6.18-398.el5-x86

RHEL5-U11-2.6.18-398.el5-x86_64

RHEL5-U11-2.6.18-398.el5PAE-x86

Red Hat Enterprise Linux 6 GA

RHEL6-GA-2.6.32-71.el6.i686-x86

RHEL6-GA-2.6.32-71.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U1

RHEL6-U1-2.6.32-131.0.15.el6.i686-x86

RHEL6-U1-2.6.32-131.0.15.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U2

RHEL6-U2-2.6.32-220.el6.i686-x86

RHEL6-U2-2.6.32-220.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U3

RHEL6-U3-2.6.32-279.el6.i686-x86

RHEL6-U3-2.6.32-279.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U4

RHEL6-U4-2.6.32-358.el6.i686-x86

RHEL6-U4-2.6.32-358.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U5

RHEL6-U5-2.6.32-431.el6.i686-x86

RHEL6-U5-2.6.32-431.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U6

RHEL6-U6-2.6.32-504.el6-x86
RHEL6-U6-2.6.32-504.el6-x86_64

Red Hat Enterprise Linux 6 U7

NOTE: Blkwatch drivers for kernel version 2.6.32-573 do not support incremental replication for workloads with LVM volumes. Update the kernel, then use drivers for kernel 2.6.32-642.

RHEL6-U7-2.6.32-573.el6.i686-x86
RHEL6-U7-2.6.32-573.el6.x86_64-x86_64
RHEL6-RHSA201700361-2.6.32-642.13.1.el6.i686-x8
RHEL6-RHSA201700361-2.6.32-642.13.1.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U8

RHEL6-U8-2.6.32-642.el6.i686-x86
RHEL6-U8-2.6.32-642.el6.x86_64-x86_64

Red Hat Enterprise Linux 6 U9

RHEL6-U9-2.6.32-696.el6.i686-x86
RHEL6-U9-2.6.32-696.el6.x86_64-x86_64

Red Hat Enterprise Linux 7 GA

RHEL7-GA-3.10.0-123.el7.x86_64-x86_64

Red Hat Enterprise Linux 7 U1

RHEL7-U1-3.10.0-229.el7.x86_64-x86_64

Red Hat Enterprise Linux 7 U2

RHEL7-U2-3.10.0-327.el7.x86_64-x86_64

Red Hat Enterprise Linux 7 U3

RHEL7-U3-3.10.0-514.el7.x86_64-x86_64

SUSE Linux Enterprise Server 9 GA

SLES9-GA-2.6.5-7.97-bigsmp-x86
SLES9-GA-2.6.5-7.97-default-x86
SLES9-GA-2.6.5-7.97-default-x86_64
SLES9-GA-2.6.5-7.97-smp-x86
SLES9-GA-2.6.5-7.97-smp-x86_64

SUSE Linux Enterprise Server 9 SP 1

SLES9-SP1-2.6.5-7.139-bigsmp-x86
SLES9-SP1-2.6.5-7.139-default-x86
SLES9-SP1-2.6.5-7.139-default-x86_64
SLES9-SP1-2.6.5-7.139-smp-x86
SLES9-SP1-2.6.5-7.139-smp-x86_64

SUSE Linux Enterprise Server 9 SP 2

SLES9-SP2-2.6.5-7.191-bigsmp-x86
SLES9-SP2-2.6.5-7.191-default-x86
SLES9-SP2-2.6.5-7.191-default-x86_64
SLES9-SP2-2.6.5-7.191-smp-x86
SLES9-SP2-2.6.5-7.191-smp-x86_64

SUSE Linux Enterprise Server 9 SP 3

SLES9-SP3-2.6.5-7.244-bigsmp-x86
SLES9-SP3-2.6.5-7.244-default-x86
SLES9-SP3-2.6.5-7.244-default-x86_64

SLES9-SP3-2.6.5-7.244-smp-x86
SLES9-SP3-2.6.5-7.244-smp-x86_64

SUSE Linux Enterprise Server 9 SP 4

SLES9-SP4-2.6.5-7.308-bigsmp-x86
SLES9-SP4-2.6.5-7.308-default-x86
SLES9-SP4-2.6.5-7.308-default-x86_64
SLES9-SP4-2.6.5-7.308-smp-x86
SLES9-SP4-2.6.5-7.308-smp-x86_64

SUSE Linux Enterprise Server 10 GA

SLES10-GA-2.6.16.21-0.8-bigsmp-x86
SLES10-GA-2.6.16.21-0.8-default-x86
SLES10-GA-2.6.16.21-0.8-default-x86_64
SLES10-GA-2.6.16.21-0.8-smp-x86
SLES10-GA-2.6.16.21-0.8-smp-x86_64
SLES10-GA-2.6.16.21-0.8-xen-x86
SLES10-GA-2.6.16.21-0.8-xen-x86_64
SLES10-GA-2.6.16.21-0.8-xenpae-x86

SUSE Linux Enterprise Server 10 SP 1

SLES10-SP1-2.6.16.46-0.12-bigsmp-x86
SLES10-SP1-2.6.16.46-0.12-default-x86
SLES10-SP1-2.6.16.46-0.12-default-x86_64
SLES10-SP1-2.6.16.46-0.12-smp-x86
SLES10-SP1-2.6.16.46-0.12-smp-x86_64
SLES10-SP1-2.6.16.46-0.12-xen-x86
SLES10-SP1-2.6.16.46-0.12-xen-x86_64
SLES10-SP1-2.6.16.46-0.12-xenpae-x86

SUSE Linux Enterprise Server 10 SP 2

SLES10-SP2-2.6.16.60-0.21-bigsmp-x86
SLES10-SP2-2.6.16.60-0.21-default-x86
SLES10-SP2-2.6.16.60-0.21-default-x86_64
SLES10-SP2-2.6.16.60-0.21-smp-x86
SLES10-SP2-2.6.16.60-0.21-smp-x86_64
SLES10-SP2-2.6.16.60-0.21-xen-x86
SLES10-SP2-2.6.16.60-0.21-xen-x86_64
SLES10-SP2-2.6.16.60-0.21-xenpae-x86

SUSE Linux Enterprise Server 10 SP 2 LTSS U2

SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-bigsmp-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-default-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-default-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-smp-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-smp-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xen-x86
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xen-x86_64
SLES10-SP2_LTSS_U2-2.6.16.60-0.42.54.1-xenpae-x86

SUSE Linux Enterprise Server 10 SP 3

SLES10-SP3-2.6.16.60-0.54.5-bigsmp-x86
SLES10-SP3-2.6.16.60-0.54.5-default-x86

SLES10-SP3-2.6.16.60-0.54.5-default-x86_64
SLES10-SP3-2.6.16.60-0.54.5-smp-x86
SLES10-SP3-2.6.16.60-0.54.5-smp-x86_64
SLES10-SP3-2.6.16.60-0.54.5-xen-x86
SLES10-SP3-2.6.16.60-0.54.5-xen-x86_64
SLES10-SP3-2.6.16.60-0.54.5-xenpae-x86

SUSE Linux Enterprise Server 10 SP 3 LTSS U1

SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-bigsmp-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-default-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-default-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-smp-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-smp-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xen-x86
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xen-x86_64
SLES10-SP3_LTSS_U1-2.6.16.60-0.113.1-xenpae-x86

SUSE Linux Enterprise Server 10 SP 3 LTSS U2

SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-bigsmp-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-default-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-default-x86_64
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-smp-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-smp-x86_64
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-xen-x86
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-xen-x86_64
SLES10-SP3_LTSS_U2-2.6.16.60-0.123.1-xenpae-x86

SUSE Linux Enterprise Server 10 SP 4

SLES10-SP4-2.6.16.60-0.85.1-bigsmp-x86
SLES10-SP4-2.6.16.60-0.85.1-default-x86
SLES10-SP4-2.6.16.60-0.85.1-default-x86_64
SLES10-SP4-2.6.16.60-0.85.1-smp-x86
SLES10-SP4-2.6.16.60-0.85.1-smp-x86_64
SLES10-SP4-2.6.16.60-0.85.1-xen-x86
SLES10-SP4-2.6.16.60-0.85.1-xen-x86_64
SLES10-SP4-2.6.16.60-0.85.1-xenpae-x86

SUSE Linux Enterprise Server 10 SP 4 U4

SLES10-SP4_U4-2.6.16.60-0.93.1-bigsmp-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-default-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-default-x86_64
SLES10-SP4_U4-2.6.16.60-0.93.1-smp-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-smp-x86_64
SLES10-SP4_U4-2.6.16.60-0.93.1-xen-x86
SLES10-SP4_U4-2.6.16.60-0.93.1-xen-x86_64
SLES10-SP4_U4-2.6.16.60-0.93.1-xenpae-x86

SUSE Linux Enterprise Server 10 SP 4 U5

SLES10-SP4_U5-2.6.16.60-0.97.1-bigsmp-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-default-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-default-x86_64
SLES10-SP4_U5-2.6.16.60-0.97.1-smp-x86

SLES10-SP4_U5-2.6.16.60-0.97.1-smp-x86_64
SLES10-SP4_U5-2.6.16.60-0.97.1-xen-x86
SLES10-SP4_U5-2.6.16.60-0.97.1-xen-x86_64
SLES10-SP4_U5-2.6.16.60-0.97.1-xenpae-x86

SUSE Linux Enterprise Server 10 SP 4 U6

SLES10-SP4_U6-2.6.16.60-0.99.1-bigsmp-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-default-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-default-x86_64
SLES10-SP4_U6-2.6.16.60-0.99.1-smp-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-smp-x86_64
SLES10-SP4_U6-2.6.16.60-0.99.1-xen-x86
SLES10-SP4_U6-2.6.16.60-0.99.1-xen-x86_64
SLES10-SP4_U6-2.6.16.60-0.99.1-xenpae-x86

SUSE Linux Enterprise Server 10 SP 4 U7

SLES10-SP4_U7-2.6.16.60-0.101.1-bigsmp-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-default-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-default-x86_64
SLES10-SP4_U7-2.6.16.60-0.101.1-smp-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-smp-x86_64
SLES10-SP4_U7-2.6.16.60-0.101.1-xen-x86
SLES10-SP4_U7-2.6.16.60-0.101.1-xen-x86_64
SLES10-SP4_U7-2.6.16.60-0.101.1-xenpae-x86

SUSE Linux Enterprise Server 10 SP 4 U8

SLES10-SP4_U8-2.6.16.60-0.103.1-bigsmp-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-default-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-default-x86_64
SLES10-SP4_U8-2.6.16.60-0.103.1-smp-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-smp-x86_64
SLES10-SP4_U8-2.6.16.60-0.103.1-xen-x86
SLES10-SP4_U8-2.6.16.60-0.103.1-xen-x86_64
SLES10-SP4_U8-2.6.16.60-0.103.1-xenpae-x86

SUSE Linux Enterprise Server 10 SP 4 LTSS U1

SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-bigsmp-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-default-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-default-x86_64
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-smp-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-smp-x86_64
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-xen-x86
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-xen-x86_64
SLES10-SP4_LTSS_U1-2.6.16.60-0.105.1-xenpae-x86

SUSE Linux Enterprise Server 10 SP 4 LTSS U2

SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-bigsmp-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-default-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-default-x86_64
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-smp-x86
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-smp-x86_64
SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-xen-x86

SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-xen-x86_64

SLES10-SP4_LTSS_U2-2.6.16.60-0.107.1-xenpae-x86

SUSE Linux Enterprise Server 11 GA

SLES11-GA-2.6.27.19-5-default-x86

SLES11-GA-2.6.27.19-5-default-x86_64

SLES11-GA-2.6.27.19-5-pae-x86

SUSE Linux Enterprise Server 11 SP 1

SLES11-SP1-2.6.32.12-0.6-default-x86

SLES11-SP1-2.6.32.12-0.6-default-x86_64

SLES11-SP1-2.6.32.12-0.6-pae-x86

SUSE Linux Enterprise Server 11 SP 1 U14

SLES11-SP1_U14-2.6.32.54-0.3-default-x86

SLES11-SP1_U14-2.6.32.54-0.3-default-x86_64

SLES11-SP1_U14-2.6.32.54-0.3-pae-x86

SUSE Linux Enterprise Server 11 SP 1 U15

SLES11-SP1_U15-2.6.32.59-0.3-default-x86

SLES11-SP1_U15-2.6.32.59-0.3-default-x86_64

SLES11-SP1_U15-2.6.32.59-0.3-pae-x86

SUSE Linux Enterprise Server 11 SP 1 U16

SLES11-SP1_U16-2.6.32.59-0.7-default-x86

SLES11-SP1_U16-2.6.32.59-0.7-default-x86_64

SLES11-SP1_U16-2.6.32.59-0.7-pae-x86

SUSE Linux Enterprise Server 11 SP 1 LTSS U1

SLES11-SP1_LTSS_U1-2.6.32.59-0.9-default-x86

SLES11-SP1_LTSS_U1-2.6.32.59-0.9-default-x86_64

SLES11-SP1_LTSS_U1-2.6.32.59-0.9-pae-x86

SUSE Linux Enterprise Server 11 SP 1 LTSS U2

SLES11-SP1_LTSS_U2-2.6.32.59-0.13-default-x86

SLES11-SP1_LTSS_U2-2.6.32.59-0.13-default-x86_64

SLES11-SP1_LTSS_U2-2.6.32.59-0.13-pae-x86

SUSE Linux Enterprise Server 11 SP 2 GA

SLES11SP2-GA-3.0.13-0.27-default-x86

SLES11SP2-GA-3.0.13-0.27-default-x86_64

SLES11SP2-GA-3.0.13-0.27-pae-x86

SLES11SP2-GA-3.0.13-0.27-xen-x86

SLES11SP2-GA-3.0.13-0.27-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U1

SLES11SP2-U1-3.0.26-0.7-default-x86

SLES11SP2-U1-3.0.26-0.7-default-x86_64

SLES11SP2-U1-3.0.26-0.7-pae-x86

SLES11SP2-U1-3.0.26-0.7-xen-x86

SLES11SP2-U1-3.0.26-0.7-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U2

SLES11SP2-U2-3.0.31-0.9-default-x86

SLES11SP2-U2-3.0.31-0.9-default-x86_64

SLES11SP2-U2-3.0.31-0.9-pae-x86

SLES11SP2-U2-3.0.31-0.9-xen-x86

SLES11SP2-U2-3.0.31-0.9-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U3

SLES11SP2-U3-3.0.34-0.7-default-x86

SLES11SP2-U3-3.0.34-0.7-default-x86_64

SLES11SP2-U3-3.0.34-0.7-pae-x86

SLES11SP2-U3-3.0.34-0.7-xen-x86

SLES11SP2-U3-3.0.34-0.7-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U4

SLES11SP2-U4-3.0.38-0.5-default-x86

SLES11SP2-U4-3.0.38-0.5-default-x86_64

SLES11SP2-U4-3.0.38-0.5-pae-x86

SLES11SP2-U4-3.0.38-0.5-xen-x86

SLES11SP2-U4-3.0.38-0.5-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U5

SLES11SP2-U5-3.0.42-0.7-default-x86

SLES11SP2-U5-3.0.42-0.7-default-x86_64

SLES11SP2-U5-3.0.42-0.7-pae-x86

SLES11SP2-U5-3.0.42-0.7-xen-x86

SLES11SP2-U5-3.0.42-0.7-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U6

SLES11SP2-U6-3.0.51-0.7.9-default-x86

SLES11SP2-U6-3.0.51-0.7.9-default-x86_64

SLES11SP2-U6-3.0.51-0.7.9-pae-x86

SLES11SP2-U6-3.0.51-0.7.9-xen-x86

SLES11SP2-U6-3.0.51-0.7.9-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U7

SLES11SP2-U7-3.0.58-0.6.2-default-x86

SLES11SP2-U7-3.0.58-0.6.2-default-x86_64

SLES11SP2-U7-3.0.58-0.6.2-pae-x86

SLES11SP2-U7-3.0.58-0.6.2-xen-x86

SLES11SP2-U7-3.0.58-0.6.2-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U8

SLES11SP2-U8-3.0.58-0.6.6-default-x86

SLES11SP2-U8-3.0.58-0.6.6-default-x86_64

SLES11SP2-U8-3.0.58-0.6.6-pae-x86

SLES11SP2-U8-3.0.58-0.6.6-xen-x86

SLES11SP2-U8-3.0.58-0.6.6-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U9

SLES11SP2-U9-3.0.74-0.6.6-default-x86

SLES11SP2-U9-3.0.74-0.6.6-default-x86_64

SLES11SP2-U9-3.0.74-0.6.6-pae-x86

SLES11SP2-U9-3.0.74-0.6.6-xen-x86

SLES11SP2-U9-3.0.74-0.6.6-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U10

SLES11SP2-U10-3.0.74-0.6.8-default-x86

SLES11SP2-U10-3.0.74-0.6.8-default-x86_64

SLES11SP2-U10-3.0.74-0.6.8-pae-x86

SLES11SP2-U10-3.0.74-0.6.8-xen-x86
SLES11SP2-U10-3.0.74-0.6.8-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U11

SLES11SP2-U11-3.0.74-0.6.10-default-x86
SLES11SP2-U11-3.0.74-0.6.10-default-x86_64
SLES11SP2-U11-3.0.74-0.6.10-pae-x86
SLES11SP2-U11-3.0.74-0.6.10-xen-x86
SLES11SP2-U11-3.0.74-0.6.10-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U12

SLES11SP2-U12-3.0.80-0.5-default-x86
SLES11SP2-U12-3.0.80-0.5-default-x86_64
SLES11SP2-U12-3.0.80-0.5-pae-x86
SLES11SP2-U12-3.0.80-0.5-xen-x86
SLES11SP2-U12-3.0.80-0.5-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U13

SLES11SP2-U13-3.0.80-0.7-default-x86
SLES11SP2-U13-3.0.80-0.7-default-x86_64
SLES11SP2-U13-3.0.80-0.7-pae-x86
SLES11SP2-U13-3.0.80-0.7-xen-x86
SLES11SP2-U13-3.0.80-0.7-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U14

SLES11SP2-U14-3.0.93-0.5-default-x86
SLES11SP2-U14-3.0.93-0.5-default-x86_64
SLES11SP2-U14-3.0.93-0.5-pae-x86
SLES11SP2-U14-3.0.93-0.5-xen-x86
SLES11SP2-U14-3.0.93-0.5-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U15

SLES11SP2-U15-3.0.101-0.5-default-x86
SLES11SP2-U15-3.0.101-0.5-default-x86_64
SLES11SP2-U15-3.0.101-0.5-pae-x86
SLES11SP2-U15-3.0.101-0.5-xen-x86
SLES11SP2-U15-3.0.101-0.5-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U16

SLES11SP2-U16-3.0.101-0.7.15-default-x86
SLES11SP2-U16-3.0.101-0.7.15-default-x86_64
SLES11SP2-U16-3.0.101-0.7.15-pae-x86
SLES11SP2-U16-3.0.101-0.7.15-xen-x86
SLES11SP2-U16-3.0.101-0.7.15-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 U17

SLES11SP2-U17-3.0.101-0.7.17-default-x86
SLES11SP2-U17-3.0.101-0.7.17-default-x86_64
SLES11SP2-U17-3.0.101-0.7.17-pae-x86
SLES11SP2-U17-3.0.101-0.7.17-xen-x86
SLES11SP2-U17-3.0.101-0.7.17-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 LTSS U1

SLES11SP2-LTSS_U1-3.0.101-0.7.19-default-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-default-x86_64

SLES11SP2-LTSS_U1-3.0.101-0.7.19-pae-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-xen-x86
SLES11SP2-LTSS_U1-3.0.101-0.7.19-xen-x86_64

SUSE Linux Enterprise Server 11 SP 2 LTSS U2

SLES11SP2-LTSS_U2-3.0.101-0.7.21-default-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-default-x86_64
SLES11SP2-LTSS_U2-3.0.101-0.7.21-pae-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-xen-x86
SLES11SP2-LTSS_U2-3.0.101-0.7.21-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 GA

SLES11SP3-GA-3.0.76-0.11-default-x86
SLES11SP3-GA-3.0.76-0.11-default-x86_64
SLES11SP3-GA-3.0.76-0.11-pae-x86
SLES11SP3-GA-3.0.76-0.11-xen-x86
SLES11SP3-GA-3.0.76-0.11-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U1

SLES11SP3-U1-3.0.82-0.7-default-x86
SLES11SP3-U1-3.0.82-0.7-default-x86_64
SLES11SP3-U1-3.0.82-0.7-pae-x86
SLES11SP3-U1-3.0.82-0.7-xen-x86
SLES11SP3-U1-3.0.82-0.7-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U2

SLES11SP3-U2-3.0.93-0.8-default-x86
SLES11SP3-U2-3.0.93-0.8-default-x86_64
SLES11SP3-U2-3.0.93-0.8-pae-x86
SLES11SP3-U2-3.0.93-0.8-xen-x86
SLES11SP3-U2-3.0.93-0.8-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U3

SLES11SP3-U3-3.0.101-0.8-default-x86
SLES11SP3-U3-3.0.101-0.8-default-x86_64
SLES11SP3-U3-3.0.101-0.8-pae-x86
SLES11SP3-U3-3.0.101-0.8-xen-x86
SLES11SP3-U3-3.0.101-0.8-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U4

SLES11SP3-U4-3.0.101-0.15-default-x86
SLES11SP3-U4-3.0.101-0.15-default-x86_64
SLES11SP3-U4-3.0.101-0.15-pae-x86
SLES11SP3-U4-3.0.101-0.15-xen-x86
SLES11SP3-U4-3.0.101-0.15-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U5

SLES11SP3-U5-3.0.101-0.21-default-x86
SLES11SP3-U5-3.0.101-0.21-default-x86_64
SLES11SP3-U5-3.0.101-0.21-pae-x86
SLES11SP3-U5-3.0.101-0.21-xen-x86
SLES11SP3-U5-3.0.101-0.21-xen-x86_64

SUSE Linux Enterprise Server 11 SP 3 U6

SLES11SP3-U6-3.0.101-0.29-default-x86

```
SLES11SP3-U6-3.0.101-0.29-default-x86_64
SLES11SP3-U6-3.0.101-0.29-pae-x86
SLES11SP3-U6-3.0.101-0.29-xen-x86
SLES11SP3-U6-3.0.101-0.29-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 3 U7

```
SLES11SP3-U7-3.0.101-0.31-default-x86
SLES11SP3-U7-3.0.101-0.31-default-x86_64
SLES11SP3-U7-3.0.101-0.31-pae-x86
SLES11SP3-U7-3.0.101-0.31-xen-x86
SLES11SP3-U7-3.0.101-0.31-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 3 U8

```
SLES11SP3-U8-3.0.101-0.35-default-x86
SLES11SP3-U8-3.0.101-0.35-default-x86_64
SLES11SP3-U8-3.0.101-0.35-pae-x86
SLES11SP3-U8-3.0.101-0.35-xen-x86
SLES11SP3-U8-3.0.101-0.35-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 4 GA

```
SLES11SP4-GA-3.0.101-63-default-x86
SLES11SP4-GA-3.0.101-63-default-x86_64
SLES11SP4-GA-3.0.101-63-pae-x86
SLES11SP4-GA-3.0.101-63-xen-x86
SLES11SP4-GA-3.0.101-63-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 4 U1

```
SLES11SP4-U1-3.0.101-65-default-x86
SLES11SP4-U1-3.0.101-65-default-x86_64
SLES11SP4-U1-3.0.101-65-pae-x86
SLES11SP4-U1-3.0.101-65-xen-x86
SLES11SP4-U1-3.0.101-65-xen-x86_64
```

SUSE Linux Enterprise Server 11 SP 4 U2

```
SLES11SP4-U2-3.0.101-68-default-x86
SLES11SP4-U2-3.0.101-68-default-x86_64
SLES11SP4-U2-3.0.101-68-pae-x86
SLES11SP4-U2-3.0.101-68-xen-x86
SLES11SP4-U2-3.0.101-68-xen-x86_64
```

D.2.3 Other Linux Distributions That Use blkwatch Drivers

PlateSpin Migrate supports other Linux distributions listed in [Table D-1](#) if the distribution is based on a supported release version of Red Hat Enterprise Linux or SUSE Linux Enterprise Server. You can use the pre-compiled blkwatch driver for the supported Linux Distribution.

Table D-1 *Blkwatch Driver Support for Other Linux Distributions*

Other Linux Distribution	Based on a Supported Release Version for RHEL or SLES	Notes
CentOS	Red Hat Enterprise Linux	
Oracle Linux (OL) (formerly Oracle Enterprise Linux (OEL))	Red Hat Enterprise Linux	<p>Blkwatch drivers are available for the standard kernel and the Unbreakable Enterprise Kernel (UEK) as noted in the Section D.2.2, “List of Distributions,” on page 252. For other Oracle Linux distributions, precompiled drivers are available only for the corresponding Red Hat Compatible Kernel (RHCK).</p> <p>Workloads using the Oracle Linux Unbreakable Enterprise Kernel are not supported in Migrate 12.1 and lower versions.</p>

For a list of supported kernel distributions, see [“List of Distributions” on page 252](#).

E Synchronizing Serial Numbers on Cluster Node Local Storage

This section details the procedure you can use to change local volume serial numbers to match each node of the Windows cluster that you want to migrate. The information includes the use of the Volume Manager utility (`VolumeManager.exe`) to synchronize serial numbers on cluster node local storage.

To download and run the utility:

- 1 From the [Micro Focus Downloads \(https://www.microfocus.com/support-and-services/download/\)](https://www.microfocus.com/support-and-services/download/) site, search for the PlateSpin Migrate product, then click **Submit Query**.
- 2 On the Products tab, select PlateSpin Migrate12.2.1 to go to the release-specific download page, then click **proceed to download**.
- 3 On the download page, click **download** on the `VolumeManager.exe` line or select the comparable download manager link, then save the file.
- 4 Copy the downloaded file to an accessible location on each cluster node.
- 5 On the active node of the cluster, open an administrative command prompt, navigate to the location of the downloaded utility, and run the following command:

```
VolumeManager.exe -l
```

A listing of the local volumes and their respective serial numbers is displayed. For example:

```
Volume Listing:
```

```
-----
```

```
DriveLetter (*) VolumeId="System Reserved" SerialNumber: AABB-CCDD
```

```
DriveLetter (C:) VolumeId=C:\ SerialNumber: 1122-3344
```

Make note of these serial numbers or keep them displayed for later comparison.

- 6 Verify that all local storage serial numbers of the active node match the local storage serial numbers on each of the other nodes in the cluster.
 - 6a On each cluster node, run the `VolumeManager.exe -l` command to obtain its volume serial numbers.
 - 6b Compare the local storage serial numbers of the active node ([Step 5](#)) against the local storage serial numbers of the node ([Step 6a](#)).
 - 6c (Conditional) If there are any differences in the serial numbers between the active node and this node, take note of the serial number you want to propagate on this node and run the following command to set, and then to verify the serial number:

```
VolumeManager -s <VolumeId> <serial-number>
```

Following are two examples of how this command could be used:

- ♦ `VolumeManager -s "System Reserved" AAAA-AAAA`
- ♦ `VolumeManager -s C:\ 1111-1111`

- 6d** When you have successfully changed all of the volume serial numbers on a node of the cluster, you need to restart that node.
- 6e** Repeat [Step 6a](#) through [Step 6d](#) for each node of the cluster.
- 7** (Conditional) If the cluster has already been migrated in a PlateSpin environment, we recommend running a full replication on the active node to ensure that any changes are propagated to the database.

F Migrate Agent Utility

Migrate Agent is a command line utility that you can use to install, upgrade, query, or uninstall the block-based transfer drivers. The utility also enables you to register source workloads with PlateSpin Migrate servers and send details about the workloads to the server via HTTPS (TCP/443, outbound).

You can use the Migrate Agent utility for registration and discovery instead of automated discovery in any live migration scenario. Using Migrate Agent is required to register and discover details about source workloads in scenarios where automated discovery is not possible, such as:

- ♦ When you deploy Migrate server in the cloud without deploying a site-to-site VPN between your network and your cloud environment.

Automated discovery is not available over the public Internet.

- ♦ When corporate network or policy restrictions prohibit opening inbound ports on source workloads.

For information about inbound ports required for automated discovery of Windows and Linux workloads, see [Section 2.6.1, “Requirements for Discovery,” on page 49](#).

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

- ♦ [Section F.1, “Requirements for Migrate Agent Utility,” on page 269](#)
- ♦ [Section F.2, “Migrate Agent Utility for Windows,” on page 270](#)
- ♦ [Section F.3, “Migrate Agent Utility for Linux,” on page 272](#)
- ♦ [Section F.4, “Using Migrate Agent with Block-Based Transfer Drivers,” on page 275](#)
- ♦ [Section F.5, “Using Migrate Agent to Register Workloads,” on page 277](#)

F.1 Requirements for Migrate Agent Utility

When you use the Migrate Agent on the source workload, the source workload contacts the target workload for data transfers.

- ♦ Ensure that you configure source workloads to support outbound traffic:
 - ♦ HTTPS port (TCP/443)
 - ♦ Replication port (TCP/3725 is the default)
- ♦ Reconfigure the replication port direction on the Migrate Server (`SourceListensForConnection=False`) using the PlateSpinConfiguration page. See [Section 5.9, “Configuring the Contact Direction for the Replication Port,” on page 107](#).
- ♦ A reboot is required when you install, uninstall, or upgrade block-based transfer drivers.

NOTE: Migrate Agent is supported only for live migrations. You can use Migrate Agent for automated migrations performed in the Migrate Web Interface and Migrate Client. It is not supported for use with semi-automated (X2P) migrations.

F.2 Migrate Agent Utility for Windows

To download the Migrate Agent utility for Windows to the source workload:

- 1 Log in to the source Windows computer as the Administrator user.
- 2 In a web browser, launch the PlateSpin Migrate Web Interface and log in.
- 3 Click the **Downloads** tab.
- 4 Click the Migrate Agent application link for the Windows target platform, then save the compressed `MigrateAgent.cli.exe` file.
- 5 Extract the contents of the file to access the executable file.
- 6 (Optional) View the Migrate Agent Help by entering

```
MigrateAgent.cli.exe -h
```

The syntax for running the Migrate Agent utility for Windows is:

```
MigrateAgent.cli.exe {command} [command_option] [/psserver=%IP%]
```

Table F-1 describes the commands, command options, and switch available for the `MigrateAgent.cli.exe` command on Windows.

Table F-1 Migrate Agent Utility for Windows Commands, Command Options, and Switch

Usage	Description
Commands	
h ? help	Displays usage and options for the command.
logs view-logs	Opens the application log directory.
reg register /reg /psserver=%IP% /username=<username> [[/password=<password>] [/pwdfile=<path-to-password-file>]] If you do not specify the password or a path to a file that contains the password, you will be prompted for the password. The password is obscured as you type it and it does not appear in the process list. Example: MigrateAgent.cli.exe /register /psserver=10.10.10.101 /username=jsmith /password=jspwd	<p>Registers this machine as a workload on the specified server. It also checks for driver upgrades from the specified PlateSpin Server.</p> <p>Enables you to add workloads that cannot be discovered. Registered workloads differ from discovered workloads in the following ways:</p> <ul style="list-style-type: none">♦ Registered source workloads do not store the source credentials.♦ You must use Migrate Agent to install, upgrade, and remove Block-based Transfer (BBT) drivers from registered source workloads.♦ After you delete the contract for a registered source workload, you must manually remove the OFX controller from the workload. To remove the OFX controller from the workload, see KB Article (https://www.netiq.com/support/kb/doc.php?id=7018453).
status /status [/psserver=%IP%]	<p>Shows installation status for the PlateSpin controller and drivers on this workload.</p> <p>If you specify the PlateSpin Server, it checks for driver upgrades from the server.</p>

Usage	Description
din driver-install /din [/psserver=%IP%]	Installs the PlateSpin drivers. If you specify the PlateSpin Server, it checks for driver upgrades from the server.
dup driver-upgrade /dup [/psserver=%IP%]	Upgrades the PlateSpin drivers. If you specify the PlateSpin Server, it checks for driver upgrades from the server.
dun driver-uninstall [/dun /psserver=%IP%]	Uninstalls the PlateSpin drivers.
con config /con /setting=<setting_name>:<value> Example: migrateagent.cli.exe /config / setting=psserver:10.10.10.202	Specifies the name of the setting and its value to change in the configuration file on this workload. The psserver option stops the OFX Controller (ofxcontroller) service, modifies the OfxController.exe.config file with the new IP address, and restarts the service. If you modify the public IP address of the PlateSpin Server, you must run this command on each of the source workloads that are configured for the server.
Switch	
/psserver=%IP%	Specifies the IPv4 address of the PlateSpin Server. Downloads the block-based transfer drivers from the specified server when you invoke the status, driver-install, or driver-upgrade options.
Command Options	
username /username= <i>value</i>	Specifies the PlateSpin Server user name for an administrator-level user with rights to add a workload.
password pwd p /password= <i>value</i>	Specifies the password for the specified PlateSpin Server user name. If you exclude the password from the command line, the script will prompt for it. The password is obscured as you type it and it does not appear in the process list. Do not combine this option with the pwdfile option.
pwdfile pf /pwdfile= <i>value</i>	Specifies the path to a file that contains the password for the specified PlateSpin Server user name. Do not combine this option with the password option.
setting /setting=<setting_name>:<value>	Specifies the setting name and value of the configuration setting to modify. Supported setting names are: psserver altAddress heartbeat

F.3 Migrate Agent Utility for Linux

Migrate Agent Utility requires the source machine to have GNU C Library (glibc) 2.11.3 or higher installed.

Ensure that you download the application with the appropriate architecture for your source Linux machines. The file name is case sensitive.

- ♦ **64-bit:** `MigrateAgent-x86_64.tar.gz`
- ♦ **32-bit:** `MigrateAgent-x86.tar.gz`

To download and install the Migrate Agent utility for Linux on the source workload:

- 1 Log in to the source Linux workload as the `root` user.
- 2 Use either of the following methods to get the `MigrateAgent-arch.tar.gz` file.

Replace *arch* with the appropriate architecture (`x86_64` or `x86`).

- ♦ Download the zipped file from the Web Interface:
 1. In a web browser, launch the PlateSpin Migrate Web Interface and log in.
`https://<Your_PlateSpin_Server>/Migrate`
Replace *Your_PlateSpin_Server* with the DNS name or IP address of your PlateSpin Migrate server.
 2. Click the **Downloads** tab.
 3. Click the Migrate Agent application link for the appropriate Linux platform (`x86_64` or `x86`), then save the `MigrateAgent-arch.tar.gz` file.

-OR-

- ♦ Use the `wget` command to copy the file from the PlateSpin Server:
 1. Launch a terminal, then enter

```
wget --no-check-certificate --http-user=username --http-password=password  
https://Your_PlateSpin_Server/Migrate/Downloads/MigrateAgent-arch.tar.gz
```

Replace *Your_PlateSpin_Server* with the DNS name or IP address of your PlateSpin Migrate server. Replace *arch* with `x86_64` or `x86`.

- 3 Open the `MigrateAgent-arch.tar.gz` file in Archive Manager, then extract the `MigrateAgent` directory and its contents to the `root` directory (`/`).

Alternatively, in a shell prompt, enter

```
tar xvf MigrateAgent-arch.tar.gz
```

Replace *arch* with `x86_64` or `x86`.

- 4 Change directory to the `/MigrateAgent` directory, then list its contents. In a terminal, enter:

```
cd MigrateAgent  
ls
```

The directory contains a `commands` file and the `MigrateAgent` script file.

- 5 View the command Help by entering:

```
./MigrateAgent -h
```

The syntax for running the Migrate Agent utility is:


```
./MigrateAgent [Command] [-h]
```

Table F-2 describes the options and arguments available for the `MigrateAgent` command on Linux.

Table F-2 *Migrate Agent Utility for Linux Command Options and Arguments*

Usage	Description
Commands	
<pre>register <server> <user> [[-p password] [-pf <password-file-path>]]</pre> <p>For <i>server</i>, specify the DNS name or IP address of your PlateSpin Migrate server.</p> <p>For <i>user</i>, specify a valid PlateSpin Server user name for an administrator-level user with rights to add a workload.</p> <p>For the password, do one of the following:</p> <ul style="list-style-type: none">♦ Use the <code>-p</code> option and type the password in the command for the specified PlateSpin user name. <code>-p mypassword</code>♦ Use the <code>-pf</code> option to specify the path to a file that contains the password for the specified PlateSpin user name. <code>-pf /tmp/jsmith-password-file.txt</code>♦ Do not specify the password in the command. You will be prompted to enter the password at the command line. <p>Example:</p> <pre>./MigrateAgent register 10.10.10.101 jsmith - p jspwd</pre>	<p>Registers this machine as a workload on the specified server. It also checks for driver upgrades from the specified PlateSpin Server.</p> <p>Enables you to add workloads that cannot be discovered. Registered workloads differ from discovered workloads in the following ways:</p> <ul style="list-style-type: none">♦ Registered source workloads do not store the source credentials.♦ You must use Migrate Agent to install, upgrade, and remove the Linux blkwatch drivers from registered source workloads.♦ After you delete the contract for a registered source workload, you must manually remove the OFX Controller from the workload. See Section 40.2, “Cleaning Up Linux Workloads,” on page 446.
<pre>status [<server>]</pre> <p>For <i>server</i>, specify the DNS name or IP address of your PlateSpin Migrate server.</p>	<p>Shows installation status for the PlateSpin controller and drivers.</p> <p>If you specify the PlateSpin Server, it checks for driver upgrades from the server.</p>
<pre>driver-install [<server>]</pre> <p>For <i>server</i>, specify the DNS name or IP address of your PlateSpin Migrate server.</p>	<p>Installs the appropriate PlateSpin blkwatch driver.</p> <p>If you specify the PlateSpin Server, it checks for driver upgrades from the server.</p>
<pre>driver-upgrade [<server>]</pre> <p>For <i>server</i>, specify the DNS name or IP address of your PlateSpin Migrate server.</p>	<p>Upgrades the installed PlateSpin blkwatch driver.</p> <p>If you specify the PlateSpin Server, it checks for driver upgrades from the server.</p>
<pre>driver-uninstall</pre>	<p>Uninstalls the installed PlateSpin blkwatch driver from the source Linux workload.</p>

Usage	Description
<p><code>configure <server> <new-server></code></p> <p>For <i>server</i>, specify the DNS name or IP address of your PlateSpin Migrate server.</p> <p>For <i>new-server</i>, specify the new DNS name or IP address of the PlateSpin Migrate server.</p> <p>Example:</p> <pre>./MigrateAgent configure 10.10.10.10 10.10.20.20</pre>	<p>Stops the OFX Controller (<code>ofxcontroller</code>) service, modifies the OFX Controller configuration file with the new address, and restarts the service. If you modify the public IP address of the PlateSpin Server, you must run this command on each of the source workloads that are configured for the server.</p>
Command Options	
<i>server</i>	<p>Specifies the DNS name or IP address of the PlateSpin Migrate Server.</p> <p>Downloads the blkwatch drivers from the specified server when you invoke the <code>status</code>, <code>driver-install</code>, or <code>driver-upgrade</code> options.</p>
<i>user</i>	Specifies the PlateSpin Server user name for an administrator-level user with rights to add a workload.
Options	
<code>-h, --help</code>	Displays usage and options for the command.
<code>-p, --password</code>	Specifies the password for the PlateSpin Server user name.
<code>-p <user_password></code>	<p>If you exclude the password from the command line, the script will prompt for it. The password is obscured as you type it and it does not appear in the process list.</p> <p>Do not combine this option with the <code>passwordfile</code> option.</p>
<code>-pf, --passwordfile</code>	Specifies the path to a file that contains the password for the specified PlateSpin Server user name.
<code>-pf <passwordfile_path></code>	Do not combine this option with the <code>password</code> option.

Usage	Description
Logging	
logging.json	<p>Contains the logging configuration settings in JSON format for logging Migrate Agent utility actions.</p> <p>To view logging settings, use the <code>cat</code> command:</p> <pre>cat MigrateAgent/logging.json</pre> <p>You can edit the file in a text editor. Set the level of the logging by changing the "level:" value from "DEBUG" to "INFO" or "ERROR". For example:</p> <pre>"level": "DEBUG"</pre> <p>or</p> <pre>"level": "INFO"</pre> <p>or</p> <pre>"level": "ERROR"</pre> <p>Logged messages are written by default to the <code>MigrateAgent.log</code> file in the <code>MigrateAgent</code> directory. You can modify the log file name setting in the <code>logging.json</code> file.</p>
MigrateAgent.log	<p>Contains the logged messages for the <code>MigrateAgent</code> command. To view the log, use the <code>cat</code> command.</p> <pre>cat MigrateAgent.log</pre>

F.4 Using Migrate Agent with Block-Based Transfer Drivers

A copy of the block-based transfer drivers is bundled with the Migrate Agent utility. You can alternatively specify the `/psserver=` command line switch in order to download the drivers from the PlateSpin Server when you invoke the `status`, `driver-install`, or `driver-upgrade` options. This is useful when the server is patched with a new driver package, but the Migrate Agent command line utility is not patched.

NOTE: To avoid confusion, the recommended method of using the Migrate Agent is to install, uninstall, or upgrade the drivers and then reboot prior to doing a replication.

You should reboot the system each time you install, upgrade, or uninstall the drivers. The reboot forces the running driver to stop and the new driver to be applied on system restart. If you do not reboot the system prior to replication, the source continues to act as if the operation has not been completed. For example, if you install drivers without rebooting the system, the source acts as if no driver is installed during replication. Similarly, if you upgrade the drivers without rebooting, the source continues to use the already running driver during replication until you reboot the system.

If the version of the installed driver is different than the version of the running driver, the `status` option will remind the user to reboot. For example:

```

C:\MigrateAgent\MigrateAgent.cli.exe status
Step 1 of 2: Querying the PlateSpin controller service
Done
Step 2 of 2: Querying the installed PlateSpin driver version
Done

The task completed successfully
PlateSpin Controller Service Status
The PlateSpin Controller service is not installed

PlateSpin Driver Status
Installed Driver Version: 8.0.0.11
Running Driver Version: Not running. Reboot to load the driver.
Upgrade Available: No

```

PlateSpin creates a task to warn the user that a reboot is necessary in order to complete the driver installation or upgrade. The notification appears in the Tasks list (Figure F-1). During replication, the notification appears on the Command Details page (Figure F-2).

Figure F-1 Reboot Notification Task


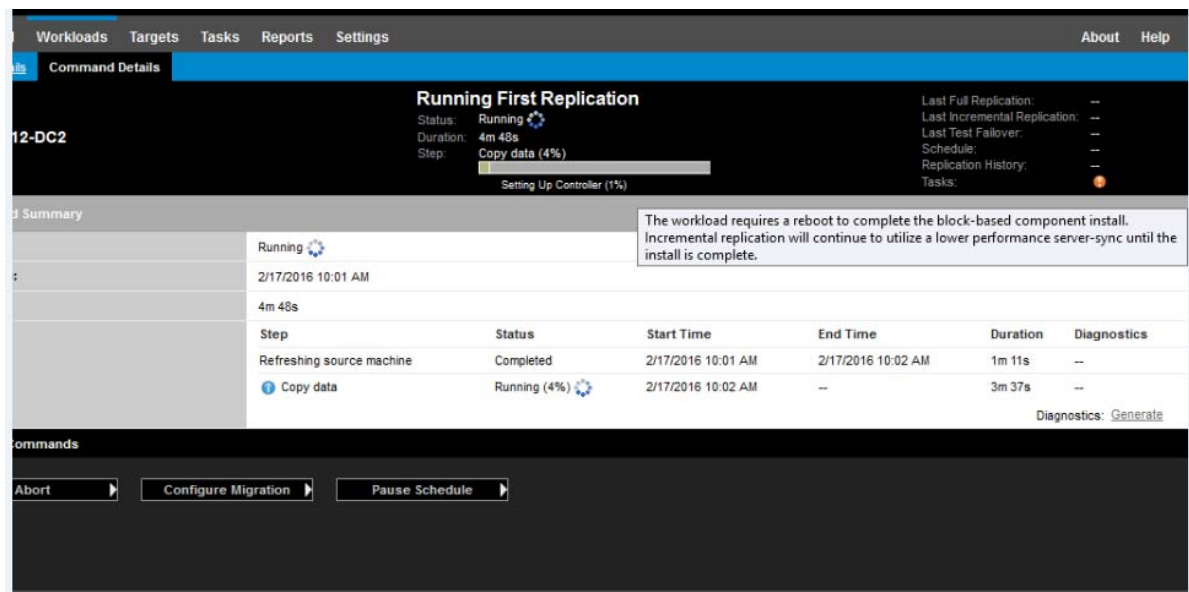
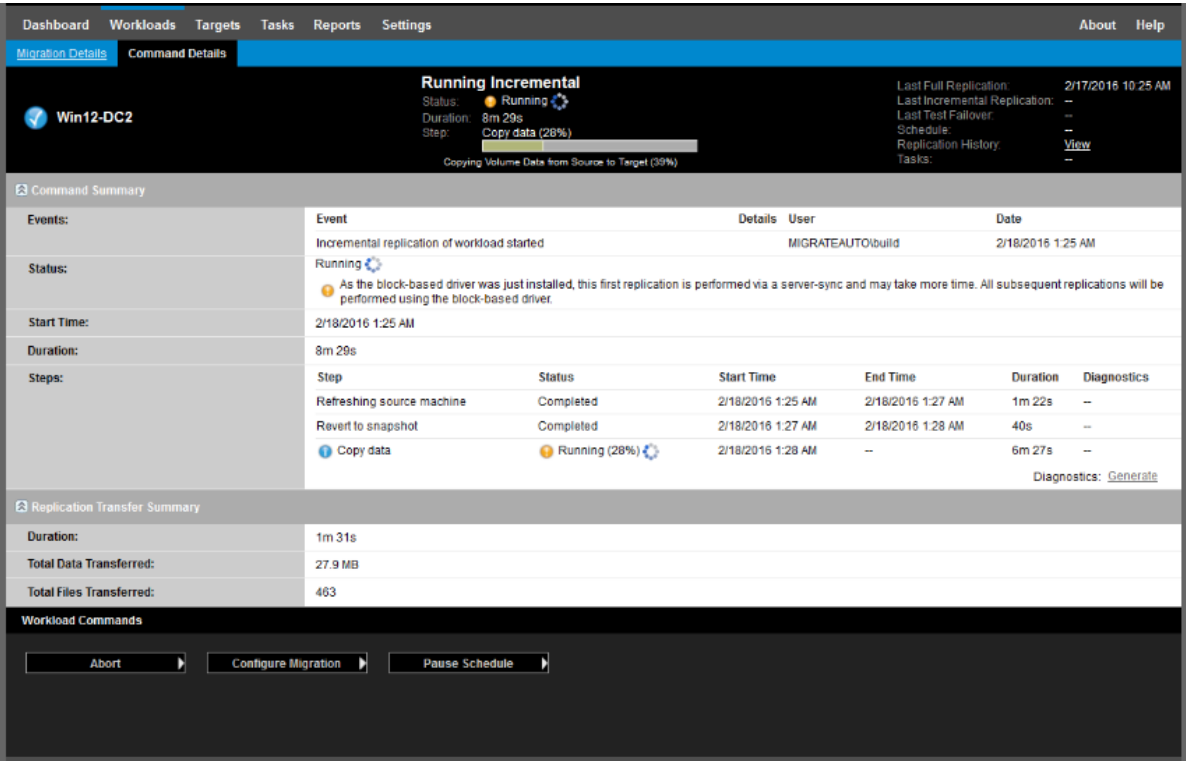


Figure F-2 Reboot Notification During Replication



Rebooting the source machine applies and starts the installed or upgraded drivers. If the driver was recently installed, after the reboot, one full replication or a server-sync replication is required in order to ensure that all of a source's changes are captured. This server-sync replication will be represented to the user in the Status field as a warning (Figure F-3). Subsequent incremental replications will complete without warning.

Figure F-3 Server-Sync Required Notification



F.5 Using Migrate Agent to Register Workloads

Migrate Agent enables you to register workloads with a PlateSpin Migrate Server:

- In a cloud environment without a VPN set up between the Migrate server and the source workloads
- In any deployment scenario where manual registration is preferred over opening ports for automated discovery using WMI/RPC/DCOM (TCP/135, 445) for Windows workloads and SSH (TCP/22) for Linux workloads

See [Section 20.7, “Registering Workloads and Discovering Details with Migrate Agent,”](#) on page 215.

G PlateSpin ISO Image

The PlateSpin ISO image file enables you to boot BIOS or UEFI firmware-based target physical machines and virtual machines during semi-automated migrations and semi-automated Server Sync operations. The semi-automated migration is used to transfer the workload to a physical machine or virtual machine that has been registered in PlateSpin Migrate. This registration occurs when you boot the target machine with the PlateSpin ISO image and register with the PlateSpin Server by following the prompts. It also discovers the target's hardware details and sends them to the server.

- ♦ [Section G.1, “Downloading the PlateSpin ISO Images,” on page 279](#)
- ♦ [Section G.2, “Preparing the PlateSpin ISO Image for Target Registration and Discovery,” on page 280](#)
- ♦ [Section G.3, “Injecting Additional Device Drivers into the PlateSpin ISO Image,” on page 280](#)
- ♦ [Section G.4, “Adding Registration Information to the PlateSpin ISO for Unattended Registration of Physical or Virtual Machines,” on page 281](#)
- ♦ [Section G.5, “Using PlateSpin ISO,” on page 281](#)

G.1 Downloading the PlateSpin ISO Images

You can download the PlateSpin ISO image from the PlateSpin Migrate software download page at [Micro Focus Downloads \(https://www.microfocus.com/support-and-services/download/\)](https://www.microfocus.com/support-and-services/download/). Search for downloads for the current product and version:

Product: PlateSpin Migrate

Version: 12.2.1

Dates: All dates

Three compressed .iso files are contained in `PhysicalTarget.zip` at the download site. Each ISO file uses the SUSE Linux Enterprise Server (SLES) operating system for the Linux RAMDisk (LRD). The LRD contains a minimal set of system files, drivers, and executables, sufficient for an initial, temporary boot. See [Table G-1](#) for information about the operating system version used for the LRD and which PlateSpin ISO file to use for Fibre Channel over Ethernet (FCoE) or Multipath I/O (MPIO) support. You can use the SLES 11 SP3 for target virtual machines if the virtual host does not support SLES 11 SP4.

Table G-1 *PlateSpin ISO Image Files*

File	LRD OS	FCoE Functionality	MPIO Functionality
<code>bootofx.x2p.iso</code>	SLES 11 SP4	Yes	No
<code>bootofx.x2p.mpio.iso</code>	SLES 11 SP4	No	Yes
<code>bootofx.x2p.sles11sp3.iso</code>	SLES 11 SP3	No	No

G.2 Preparing the PlateSpin ISO Image for Target Registration and Discovery

- 1 Download the PlateSpin ISO image from [Micro Focus Downloads](#) and extract the contents. See [Downloading the PlateSpin ISO Images](#).
- 2 (Optional) Inject additional device drivers for Linux workloads into the PlateSpin ISO image, complete the steps in [Injecting Additional Device Drivers into the PlateSpin ISO Image](#).
- 3 (Optional) For an unattended registration, modify the PlateSpin ISO to provide the appropriate responses from an answer file. See [Adding Registration Information to the PlateSpin ISO for Unattended Registration of Physical or Virtual Machines](#).
- 4 Save the PlateSpin ISO image:
 - ♦ **Physical Machine:** Burn the PlateSpin ISO image on a CD or save it to the required media, from which your target can boot.
 - ♦ **Virtual Machine:** Save the PlateSpin ISO image on the virtual host for a target VM in a location where you can use it to boot the target machine.
- 5 Use native tools to prepare the target machine to boot from the PlateSpin ISO image.
Ensure that the machine is configured to restart on reboot and that you attach the PlateSpin ISO file as a boot CD for the VM.

For information about registering the target machine, see the following:

- ♦ “[Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO](#)” on [page 204](#)
- ♦ “[Registering and Discovering Details for Target Physical Machines with PlateSpin ISO](#)” on [page 206](#)

G.3 Injecting Additional Device Drivers into the PlateSpin ISO Image

The PlateSpin ISO image contains a large library of device drivers sufficient to boot most common targets. However, occasionally you might want to use your own, such as lesser-known, vendor-specific or custom-developed drivers for Linux workloads.

To inject drivers into the PlateSpin ISO image for Linux workloads:

- 1 Download and extract the PlateSpin ISO images. See [Downloading the PlateSpin ISO Images](#).
- 2 Obtain or compile the required *.ko driver files.

IMPORTANT: The `rebuildiso.sh` script that helps you rebuild the ISO file is currently supported only on 64-bit version of the OS. Ensure that the drivers are valid for the kernel (3.0.93-0.8-default) included with the ISO file.

- 3 Mount the ISO image in any Linux machine (root credentials required). Use the following command syntax:

```
mount -o loop <path-to-ISO> <mount_point>
```
- 4 Copy the `rebuildiso.sh` script, located in the `/tools` subdirectory of the mounted ISO file, into a temporary working directory.
- 5 Create another working directory for the required driver files and save them in that directory.

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

```
./rebuildiso.sh -i <ISO_file> -d <driver_dir> -m32|-m64
```

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

NOTE: To rebuild Migrate LRD ISO, a minimum of `genisoimage 1.1.11` is required. By default, operating systems such as RHEL 7, CentOS 7, and Ubuntu 14.04.1 have the required `genisoimage` version.

7 Unmount the ISO file (execute the command `umount <mount_point>`).

G.4 Adding Registration Information to the PlateSpin ISO for Unattended Registration of Physical or Virtual Machines

PlateSpin Migrate provides a mechanism for automating the registration and discovery of details for a target physical or virtual machine. You must first update the PlateSpin ISO image with specific registration information before booting the target.

For details, see [KB Article 7013485](https://www.netiq.com/support/kb/doc.php?id=7013485) (<https://www.netiq.com/support/kb/doc.php?id=7013485>).

G.5 Using PlateSpin ISO

After you have prepared the PlateSpin ISO for your environment, you can use the file to register and discover target physical machines or target virtual machines in a semi-automated migration or Server Sync operation. See the following procedures in “[Discovering Target Hosts](#)”:

- ♦ “[Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO](#)” on [page 204](#)
- ♦ “[Registering and Discovering Details for Target Physical Machines with PlateSpin ISO](#)” on [page 206](#)

V Configuring Workloads

After you discover targets and workloads, you are ready to prepare for migration by configuring migration jobs for your workloads.

- ♦ [Chapter 24, “Prerequisites for Automated Migrations,” on page 285](#)
- ♦ [Chapter 25, “Prerequisites for Semi-Automated \(X2P\) Migrations,” on page 289](#)
- ♦ [Chapter 26, “Configuration Essentials,” on page 291](#)
- ♦ [Chapter 27, “Migration to Amazon Web Services,” on page 331](#)
- ♦ [Chapter 28, “Migration to Microsoft Azure,” on page 335](#)
- ♦ [Chapter 29, “Migration to VMware vCloud Director,” on page 349](#)
- ♦ [Chapter 30, “Migration to VMware,” on page 359](#)
- ♦ [Chapter 31, “Migration to Microsoft Hyper-V,” on page 383](#)
- ♦ [Chapter 32, “Migration to Virtual Machines on Citrix XenServer,” on page 395](#)
- ♦ [Chapter 33, “Migration to Virtual Machines on Xen,” on page 399](#)
- ♦ [Chapter 34, “Migration to Virtual Machines on KVM,” on page 403](#)
- ♦ [Chapter 35, “Migration to Physical Machines,” on page 407](#)
- ♦ [Chapter 36, “Workload Migration with a PlateSpin Image,” on page 415](#)
- ♦ [Chapter 37, “Synchronizing Workloads with Server Sync,” on page 423](#)

24 Prerequisites for Automated Migrations

PlateSpin Migrate Client and PlateSpin Migrate Web Interface enable you to automate migration of workloads to target virtualization containers and target cloud containers.

- ♦ [Section 24.1, “Supported Source Workloads for Automated Migration,” on page 285](#)
- ♦ [Section 24.2, “Supported Target Platforms for Automated Migrations,” on page 286](#)
- ♦ [Section 24.3, “Preparing Targets for Automated Migration,” on page 286](#)
- ♦ [Section 24.4, “Network Connections and Bandwidth,” on page 287](#)
- ♦ [Section 24.5, “Automated Workflow,” on page 287](#)

24.1 Supported Source Workloads for Automated Migration

In an automated migration, PlateSpin Migrate builds the target virtual machine on the destination container based on the target workload details you configure for the conversion. Automation supports source workloads based on the destination target platform. For information about source workloads for supported virtualization and cloud platforms, see [Table 24-2](#).

Table 24-1 *Supported Source Workloads for Automated Migrations*

Target Platform	Migrate Client	Migrate Web Interface
Amazon Web Services	Table 2-7, “AWS: Supported Windows Platforms,” on page 31 Table 2-8, “AWS: Supported Linux Platforms,” on page 32	Not supported
Microsoft Azure	Not supported	Table 2-3, “Azure: Supported Windows Platforms,” on page 28 Table 2-4, “Azure: Supported Linux Platforms,” on page 28
VMware vCloud Director	Not supported	Table 2-5, “vCloud: Supported Windows Platforms,” on page 29 Table 2-6, “vCloud: Supported Linux Platforms,” on page 30
VMware	Table 2-1, “Non-Cloud Platforms: Supported Windows Workloads,” on page 24 Table 2-2, “Non-Cloud Platforms: Supported Linux Workloads,” on page 26	Table 2-1, “Non-Cloud Platforms: Supported Windows Workloads,” on page 24 Table 2-2, “Non-Cloud Platforms: Supported Linux Workloads,” on page 26

Target Platform	Migrate Client	Migrate Web Interface
Hyper-V	Table 2-1, “Non-Cloud Platforms: Supported Windows Workloads,” on page 24 Table 2-2, “Non-Cloud Platforms: Supported Linux Workloads,” on page 26	Not supported

24.2 Supported Target Platforms for Automated Migrations

In an automated migration, PlateSpin Migrate prepares the virtual machine on the target platform before the replications begin. You can schedule when the first full replication begins. The Prepare Workload step must be executed prior to the scheduled start time. For information about supported virtualization and cloud platforms, see [Table 24-2](#).

Table 24-2 *Supported Target Platforms for Automated Migrations*

Target Platform	Migrate Client	Migrate Web Interface
Amazon Web Services	Table 2-16, “Supported Target Cloud Platform for the Migrate Client,” on page 41	Not supported
Microsoft Azure	Not supported	See “Microsoft Azure” in Table 2-15, “Supported Target Cloud Platforms for the Migrate Web Interface,” on page 41
VMware vCloud Director	Not supported	See “VMware vCloud Director” in Table 2-15, “Supported Target Cloud Platforms for the Migrate Web Interface,” on page 41
VMware	Table 2-12, “Supported Target VMware Platforms for the Migrate Web Interface and Migrate Client,” on page 38	Table 2-12, “Supported Target VMware Platforms for the Migrate Web Interface and Migrate Client,” on page 38
Hyper-V	See “Hyper-V” in Table 2-14, “Supported Target Virtualization Platforms for the Migrate Client Only,” on page 39	Not supported

24.3 Preparing Targets for Automated Migration

In an automated migration, PlateSpin requires information about the target container where it will create the virtual machines. You must prepare your target environment for discovery, and discover the target. For information about configuring the target container environment for use with PlateSpin Migrate, see [Table 24-3](#). For discovery of target containers, see [“Discovering Details for Target Containers” on page 200](#).

Table 24-3 Prerequisites for Target Containers

Target Platform	Migrate Client	Migrate Web Interface
Amazon Web Services	“Prerequisites for Migration to Amazon Web Services” on page 137	Not supported
Microsoft Azure	Not supported	“Prerequisites for Migration to Microsoft Azure” on page 147
VMware vCloud Director	Not supported	“Prerequisites for Migration to VMware vCloud Director” on page 163
VMware	“Prerequisites for Migration to VMware” on page 169	“Prerequisites for Migration to VMware” on page 169
Hyper-V	“Prerequisites for Migration to Microsoft Hyper-V” on page 175	Not supported

24.4 Network Connections and Bandwidth

Before you execute replications for an automated migration:

- ◆ Ensure that your network access and ports are properly configured. See [“Requirements for Migration” on page 50](#).
- ◆ Ensure that you test the connection to see if there are any connection or bandwidth issues, and resolve them. For information about optimizing throughput on the connection, see [“Using the iPerf Network Test Tool to Optimize Network Throughput for PlateSpin Products” on page 487](#).

24.5 Automated Workflow

Refer to the checklist to understand the automated workflow:

- ◆ [“Checklist for Automated Migration to AWS” on page 145](#)
- ◆ [“Checklist for Automated Migration to Azure” on page 161](#)
- ◆ [“Checklist for Automated Migration to vCloud” on page 168](#)
- ◆ [“Checklist for Automated Migration to VMware” on page 172](#)
- ◆ [“Checklist for Automated Migration to Hyper-V” on page 178](#)

For information about configuring automated migration to a target container, see:

- ◆ [“Configuring Full Migration of a Workload to Amazon Web Services” on page 332](#)
- ◆ [“Configuring Migration of a Workload to Microsoft Azure” on page 336](#)
- ◆ [“Configuring Migration of a Workload to VMware vCloud Director” on page 349](#)
- ◆ [“Automated Migration to VMware Using Migrate Client” on page 360](#)
- ◆ [“Automated Migration to VMware Using Migrate Web Interface” on page 374](#)
- ◆ [“Automated Migration to Hyper-V” on page 384](#)

25 Prerequisites for Semi-Automated (X2P) Migrations

PlateSpin Migrate Client enables you to migrate workloads to physical machines (X2P). You use PlateSpin ISO to register the target physical machine with the PlateSpin Migrate server and report details about it. This manual process of target preparation and discovery is referred to as the *X2P workflow*.

- [Section 25.1, “Supported Source Workloads for X2P Migrations,” on page 289](#)
- [Section 25.2, “Supported Target Platforms for X2P Migrations,” on page 289](#)
- [Section 25.3, “X2P Workflow for VMs,” on page 289](#)

25.1 Supported Source Workloads for X2P Migrations

You can also use the X2P workflow to migrate workloads to a virtual machines that you set up on a supported virtual host. You configure the VM with guest operating system type and version settings that match your source workload, in accordance with the features and capabilities of the target virtualization platform. For information about source workloads for supported virtualization platforms, see:

- [Table 2-1, “Non-Cloud Platforms: Supported Windows Workloads,” on page 24](#)
- [Table 2-2, “Non-Cloud Platforms: Supported Linux Workloads,” on page 26](#)

25.2 Supported Target Platforms for X2P Migrations

PlateSpin Migrate Client supports using the X2P workflow for migrations to physical machines and to any supported virtual host, even if an automated alternative is available. For information about supported virtualization platforms, see [Section 2.1.5, “Supported Target Virtualization Platforms,” on page 38](#).

25.3 X2P Workflow for VMs

To migrate a workload to a VM on a virtual host:

- 1 Use the native interface of the required virtualization platform to set up the target virtual machine with guest operating system type and version settings that match your source workload, in accordance with the features and capabilities of the target virtualization platform.
- 2 Begin booting the newly created virtual machine by using the appropriate PlateSpin ISO image, load the appropriate driver, if needed, then continue the boot process.

This special boot process discovers and registers the target virtual machine as a PlateSpin Migrate physical machine target. See [“Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO” on page 204](#).

- 3 Use the PlateSpin Migrate Client to create and execute an X2P migration job.
- 4 Upon completion of the migration job, install virtualization enhancement software specific to the target virtualization platform.

For information about configuring semi-automated migration to a virtual machine running on virtualization hosts that PlateSpin Migrate regards as a physical machine:

- ♦ [“Migration to VMs on VMware Using X2P Workflow” on page 371](#)
- ♦ [“Migration to VMs on Hyper-V Using X2P Workflow” on page 393](#)
- ♦ [“Migration to Virtual Machines on Citrix XenServer” on page 395](#)
- ♦ [“Migration to Virtual Machines on Xen” on page 399](#)
- ♦ [“Migration to Virtual Machines on KVM” on page 403](#)

26 Configuration Essentials

When you configure a workload for migration, the workload type and target determine the configuration options available. This section describes the essentials for configuration of each parameter.

- ♦ [Section 26.1, “Configuration Workflows,” on page 291](#)
- ♦ [Section 26.2, “Initiating a Migration Job,” on page 292](#)
- ♦ [Section 26.3, “Saving a Migration Configuration,” on page 295](#)
- ♦ [Section 26.4, “Editing a Migration Job,” on page 295](#)
- ♦ [Section 26.5, “Migrate License Key,” on page 296](#)
- ♦ [Section 26.6, “Credentials for Source Workloads and Target Hosts,” on page 297](#)
- ♦ [Section 26.7, “Migration Schedule,” on page 298](#)
- ♦ [Section 26.8, “Blackout Window for Data Transfer,” on page 298](#)
- ♦ [Section 26.9, “Compression during Data Transfer,” on page 299](#)
- ♦ [Section 26.10, “Bandwidth Throttling during Data Transfer,” on page 300](#)
- ♦ [Section 26.11, “Conversion \(Data Transfer Method\),” on page 300](#)
- ♦ [Section 26.12, “Encrypt Data Transfer,” on page 301](#)
- ♦ [Section 26.13, “Virtualization Enhancement Software,” on page 302](#)
- ♦ [Section 26.14, “Custom Post-Migration Actions,” on page 303](#)
- ♦ [Section 26.15, “Services or Daemons to Stop before Replication or Cutover,” on page 304](#)
- ♦ [Section 26.16, “Service States on Target Windows Workloads,” on page 306](#)
- ♦ [Section 26.17, “Daemon States on Target Linux Workloads,” on page 310](#)
- ♦ [Section 26.18, “Windows HAL or Kernel File Replacements,” on page 312](#)
- ♦ [Section 26.19, “Post-Cutover End States for Source and Target Workloads,” on page 313](#)
- ♦ [Section 26.20, “Target Workload Settings for VMs,” on page 314](#)
- ♦ [Section 26.21, “Network Identification \(Network Connections\),” on page 315](#)
- ♦ [Section 26.22, “Migration Network \(Replication Network\),” on page 318](#)
- ♦ [Section 26.23, “Storage Disks and Volumes,” on page 326](#)

26.1 Configuration Workflows

Refer to the migration configuration sections for a step-by-step walk through the migration configuration for the various migration job types.

- ♦ [Section 26.1.1, “Configuration Workflows Using Migrate Client,” on page 292](#)
- ♦ [Section 26.1.2, “Configuring Workflows Using Migrate Web Interface,” on page 292](#)

26.1.1 Configuration Workflows Using Migrate Client

The PlateSpin Migrate Client supports migration of workloads to Amazon Web Services, VMware containers, Microsoft Hyper-V, Citrix XenServer, Xen, KVM, physical machines, images, and server-sync.

- ♦ [Migration to Amazon Web Services](#)
- ♦ [Migration to VMware](#)
- ♦ [Migration of Windows Clusters](#)
- ♦ [Migration to Microsoft Hyper-V](#)
- ♦ [Migration to Virtual Machines on Citrix XenServer](#)
- ♦ [Migration to Virtual Machines on Xen](#)
- ♦ [Migration to Virtual Machines on KVM](#)
- ♦ [Migration to Physical Machines](#)
- ♦ [Workload Migration with a PlateSpin Image](#)
- ♦ [Synchronizing Workloads with Server Sync](#)

26.1.2 Configuring Workflows Using Migrate Web Interface

The PlateSpin Migrate Web Interface supports large scale migration of workloads to VMware containers and cloud platforms such as Microsoft Azure, VMware vCloud Director, and Amazon Web Services.

- ♦ [Migration to Microsoft Azure](#)
- ♦ [Migration to VMware vCloud Director](#)
- ♦ [Automated Migration to VMware Using Migrate Web Interface](#)
- ♦ [Preparing for Migration of Windows Clusters](#)

26.2 Initiating a Migration Job

After workload discovery, the migration job for the workload is in a unconfigured state. Migration jobs are not automatically initiated with default settings. You must initiate the migration job by starting configuration for the migration.

- ♦ [Section 26.2.1, “Prerequisites for Migration Jobs,” on page 292](#)
- ♦ [Section 26.2.2, “Initiate a Migration Job Using Migrate Client,” on page 293](#)
- ♦ [Section 26.2.3, “Initiate a Migration Job Using the Migrate Web Interface,” on page 294](#)

26.2.1 Prerequisites for Migration Jobs

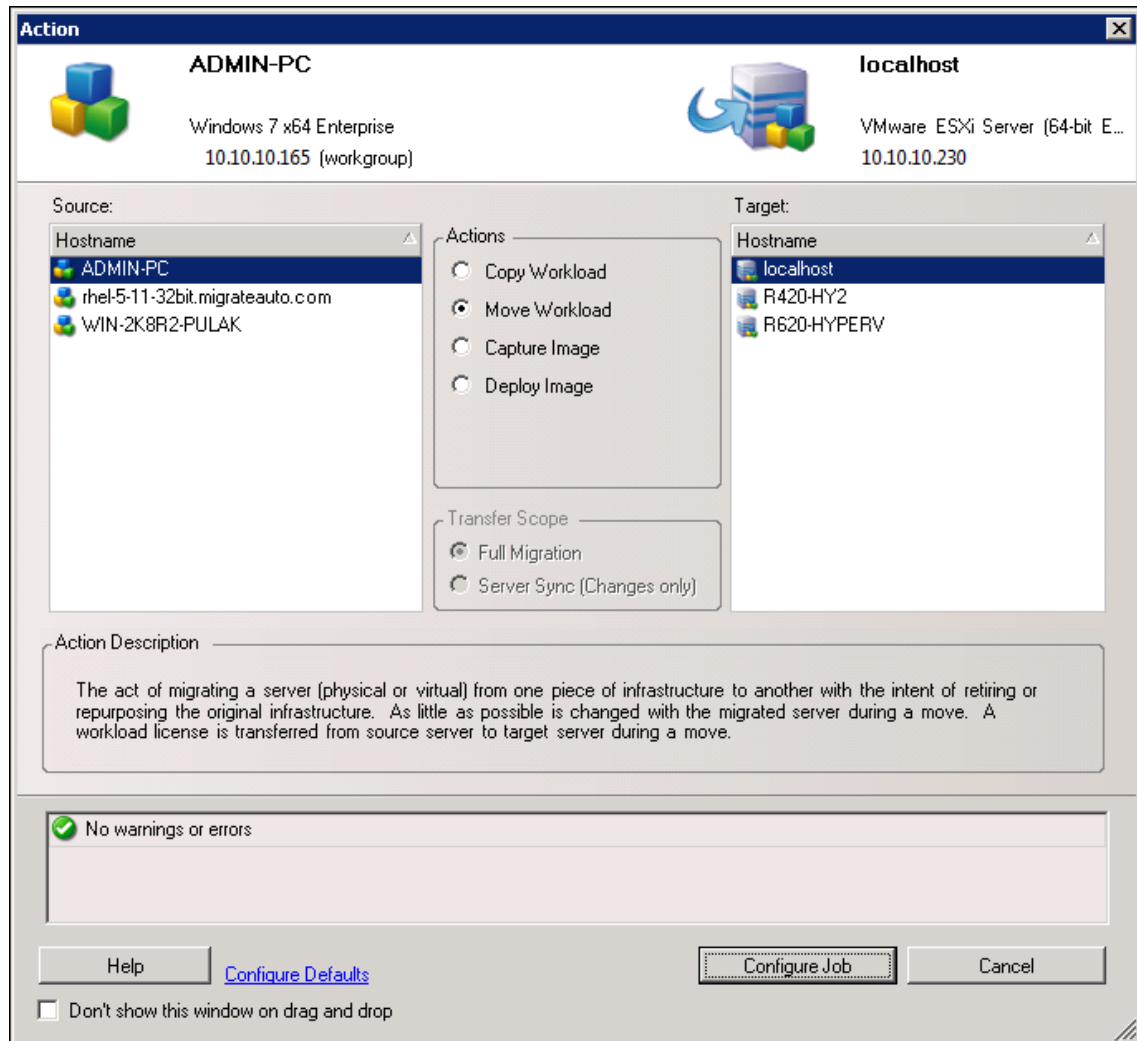
For any migration job, ensure that you have completed the following tasks:

- ♦ You must have discovered details for the source workload and the target host. See [Part IV, “Discovering and Preparing Workloads and Targets,” on page 195](#).
- ♦ Ensure that the credentials for the source workload and target host are valid.

26.2.2 Initiate a Migration Job Using Migrate Client

To start a migration job for a Workload:

- 1 In the Migrate Client, open the Action window. Use any of the following methods:
 - ♦ Drag a discovered source and drop it on a discovered target.
 - ♦ Click a task in the Tasks pane.
 - ♦ Click the New Job toolbar.
 - ♦ In the Jobs view, right-click a source and select a command from the context menu. Available commands depend on the type of source.



The Source and Target panes display workloads and targets applicable to the selected type of a migration job under **Actions**:

- ♦ Copy Workload
- ♦ Move Workload
- ♦ Capture Image
- ♦ Deploy Image

For **Transfer Scope**, the **Full Transfer** and **Server Sync** options are enabled in the following circumstances:

- ♦ The system detects an existing operating system on the target
- ♦ The operating system profile of the target matches that of the source workload

See [“Synchronizing Workloads with Server Sync” on page 423](#).

- 2 Check the validation messages at the bottom of the window.
- 3 To start configuring your migration job, click **Configure Job**.
- 4 (Optional) For convenience, to avoid displaying the **Action** window on drag-and-drop, select **Don't show this window on drag and drop** before you proceed. Subsequent drag-and-drops actions bypass the Action window and directly open a Conversion Job window.

To restore the job migration startup behavior, restore the application defaults. See [“Configuring General Options” on page 117](#).

- 5 Configure the migration as appropriate for the workload and target host.
 - ♦ [Migration to Amazon Web Services](#)
 - ♦ [Automated Migration to VMware Using Migrate Client](#)
 - ♦ [Preparing for Migration of Windows Clusters](#)
 - ♦ [Migration to Microsoft Hyper-V](#)
 - ♦ [Migration to Virtual Machines on Citrix XenServer](#)
 - ♦ [Migration to Virtual Machines on Xen](#)
 - ♦ [Migration to Virtual Machines on KVM](#)
 - ♦ [Migration to Physical Machines](#)
 - ♦ [Workload Migration with a PlateSpin Image](#)
 - ♦ [Synchronizing Workloads with Server Sync](#)

26.2.3 Initiate a Migration Job Using the Migrate Web Interface

- 1 In the PlateSpin Migrate Web Interface, click **Workloads**.
- 2 On the Workloads page, select the workload to migrate.
- 3 Click **Configure Migration**.
- 4 Specify the **Initial Transfer Method** for replication based on the scope of data you want to transfer from the source to the target:
 - ♦ **Full Replication:** Migrate replicates the full volume from the source to the target.
 - ♦ **Incremental Replication:** Migrate replicates only differences in data from the source to the target, provided the workloads have similar operating system and volume profiles.
- 5 Select a discovered target host, then click **Configure Migration**.
- 6 Configure the Target Workload Details as appropriate for the workload and target host.
 - ♦ [Migration to Microsoft Azure](#)
 - ♦ [Migration to VMware vCloud Director](#)
 - ♦ [Automated Migration to VMware Using Migrate Web Interface](#)
 - ♦ [Preparing for Migration of Windows Clusters](#)
- 7 Click one of the following:
 - ♦ **Save & Prepare**

- ♦ **Save**
- ♦ **Cancel**

26.3 Saving a Migration Configuration


After you configure a workload for migration, you can save the migration configuration for execution at a later time.

- ♦ [Section 26.3.1, “Using the Migrate Client,” on page 295](#)
- ♦ [Section 26.3.2, “Using the Migrate Web Interface,” on page 295](#)

26.3.1 Using the Migrate Client

To save a migration configuration:

- 1 Set up a migration job and configure the options.
- 2 On the Edit Migration Details page, click the arrow at the right side of the **Save** button to expand the Save menu, then select **Save as** or **Save with NTFS Encryption**.



26.3.2 Using the Migrate Web Interface

To save a migration configuration:

- 1 Set up a migration job and configure the options.
- 2 Do one of the following:
 - ♦ Click **Save & Prepare** to save the migration and to begin preparations for the target VM replication environment on the target host.
 - ♦ Click **Save** to save the migration for subsequent changes or later execution.

26.4 Editing a Migration Job

You can save an incomplete configuration for a migration job, then add or change settings later.

- ♦ [Section 26.4.1, “Edit Migration Job Using Migrate Client,” on page 295](#)
- ♦ [Section 26.4.2, “Edit Migration Job Using Migrate Web Interface,” on page 296](#)

26.4.1 Edit Migration Job Using Migrate Client

- 1 In the Jobs view, locate the required job.
- 2 Open the Migration Job window.
- 3 Modify the settings as appropriate.
- 4 Click **OK**.

26.4.2 Edit Migration Job Using Migrate Web Interface

- 1 On the Workloads page, click the name link of the workload to migrate.
- 2 On the Migration Details page, click **Edit**.
- 3 Modify the settings as appropriate.
- 4 Click **Save**.

26.5 Migrate License Key

By default, PlateSpin Migrate automatically selects the best license key for a particular migration job. For information about product licensing and license key management, see [Section 5.1, “PlateSpin Migrate Product Licensing,” on page 91](#).

- ♦ [Section 26.5.1, “License Key in Migrate Client,” on page 296](#)
- ♦ [Section 26.5.2, “License Key in Migrate Web Interface,” on page 296](#)

26.5.1 License Key in Migrate Client

If you have multiple license keys, PlateSpin Migrate client enables you to select a specific license key to apply to a particular migration job, assuming its workload licenses are available (neither expired nor exhausted). Certain licenses cannot be selected if they are invalid for the current migration. Licenses can be invalid for reasons such as:

- ♦ There are no remaining migrations for the license.
- ♦ The license does not allow X2V migrations and the current migration is a P2V.
- ♦ The license does not support live transfer migrations and the current migration is marked for live transfer.

To view or modify the license key selected for a migration job:

- 1 Start the migration job. For information about starting a migration job, see [Section 26.2, “Initiating a Migration Job,” on page 292](#).
- 2 In the Job Configuration section of the Migration Job window, click **License**.
- 3 To manually choose a different key, deselect **Automatically select the best license key during the conversion**, then select the required license key from the menu.
- 4 Click **OK**.

The selected license key is displayed on the **Licenses** tab and the description is updated accordingly.

26.5.2 License Key in Migrate Web Interface

If multiple license keys are available, PlateSpin Migrate Web Interface consumes workload licenses associated with the license keys in order of their start date until all workloads associated with the key are consumed. You cannot specify the key to be used by each workload.

26.6 Credentials for Source Workloads and Target Hosts

When you configure a migration job, you can validate the provided credentials and save them for future migration jobs that use the same source and target. If you modify the password on the workload or target host, you must also modify the credentials stored in PlateSpin Migrate.

- ♦ [Section 26.6.1, “About Credentials,” on page 297](#)
- ♦ [Section 26.6.2, “Credentials in Migrate Client,” on page 297](#)
- ♦ [Section 26.6.3, “Credentials in Migrate Web Interface,” on page 297](#)

26.6.1 About Credentials

For a migration job to execute properly, you must provide valid credentials for your source and target. For more information about credentials format, see:

- ♦ [Section 19.3, “Discovery Guidelines for Target Hosts,” on page 198](#)
- ♦ [Section 20.3, “Discovery Guidelines for Source Workloads,” on page 212](#)

26.6.2 Credentials in Migrate Client

To modify source and target credentials:

- 1 In the Jobs view, select the required workload or target.
- 2 In the Job Configuration section of the Migration Job window, click **Access**.
- 3 Specify the credentials.
- 4 Click **OK**.

26.6.3 Credentials in Migrate Web Interface

To modify target credentials:

- 1 In the Migrate Web Interface, click **Targets**, then click the target name.
- 2 On the Target Details page, click **Edit**.
- 3 On the Edit Target Details page, specify the new user name and password.
- 4 Click **Save**.

To modify source workload credentials:

- 1 In the Migrate Web Interface, click **Workloads**, then click the workload name.
- 2 On the Workload Details page, click **Edit**.
- 3 On the Edit Target Workload Details page, go to **Migration Settings > Source Credentials**.
- 4 Specify the new user name and password for the source workload.
- 5 Click **Save**.

26.7 Migration Schedule

The migration schedule enables you to specify whether to start the first replication manually or on a specific date and a specific time.

- ♦ [Section 26.7.1, “Migration Schedule Using Migrate Client,” on page 298](#)
- ♦ [Section 26.7.2, “Migration Schedule Using Migrate Web Interface,” on page 298](#)

26.7.1 Migration Schedule Using Migrate Client

To schedule the migration start date and time:

- 1 In the Jobs view, locate the required job.
- 2 In the Job Configuration section of the Migration Job window, click **Schedule**.
- 3 Select **Run at a later time**, then specify the date and start time for the first replication.
- 4 Click **OK**.

26.7.2 Migration Schedule Using Migrate Web Interface

To schedule the migration start date and time:

- 1 On the Edit Migration Details page, go to **Schedule Settings > Full Replication**, then click **Edit**.
- 2 Click **Start**, then set the date and time when you want to start the first full replication.
You can type the date (dd/mm/yyyy) or click the Calendar icon to select the date. The default run time is 12:00:00 a.m. (hh:mm:ss a.m. or p.m.).
- 3 Click **Close** to return to the Edit Migration Details page.
- 4 Click **Save**.

26.8 Blackout Window for Data Transfer

The blackout window suspends scheduled replications from starting during a specified period of time and pattern. It helps you to reserve network bandwidth for users or mission critical communications during peak traffic periods. You can also use it to prevent conflicts for other data backup or snapshot activities. For example, suspend replications during peak network utilization hours or to prevent conflicts between VSS-aware software and the PlateSpin VSS block-level data transfer component.

The default setting is None. No blackout window is scheduled.

NOTE: The blackout start and end times are based on the system clock on the PlateSpin Server.

- ♦ [Section 26.8.1, “Blackout Window Using the Migrate Client,” on page 298](#)
- ♦ [Section 26.8.2, “Blackout Window Using the Migrate Web Interface,” on page 299](#)

26.8.1 Blackout Window Using the Migrate Client

PlateSpin Migrate Client does not provide an option to configure a blackout window for data transfer.

26.8.2 Blackout Window Using the Migrate Web Interface

To set or modify a blackout window:

- 1 On the Edit Migration Details page, go to **Schedule Settings** > **Blackout Window**, then click **Edit**.
- 2 Specify the start and end time for the blackout period.
The blackout start and end times are based on the system clock on the PlateSpin Server.
- 3 Select **Daily**, **Weekly**, or **Monthly** to enable a blackout window, then set the recurrence pattern.
- 4 Click **Close** to return to the Edit Migration Details page.
- 5 Click **Save**.

26.9 Compression during Data Transfer

The Compression Level setting controls whether data is compressed during transmission between the source and target workloads, and the level of data compression applied. See [Section 2.4.3, “Data Compression,” on page 48](#).

Select one of the following options:

- ♦ **None:** No compression.
- ♦ **Fast:** Consumes the least CPU resources on the source, but yields a lower compression ratio.
- ♦ **Optimal:** (Default) Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option.
- ♦ **Maximum:** Consumes the most CPU resources on the source, but yields a higher compression ratio.
- ♦ [Section 26.9.1, “Compression Using Migrate Client,” on page 299](#)
- ♦ [Section 26.9.2, “Compression Using Migrate Web Interface,” on page 299](#)

26.9.1 Compression Using Migrate Client

To enable and use compression for data transfer:

- 1 In the Jobs view, locate the required job.
- 2 In the Network Configuration section of the Migration Job window, select **Enable Compression**.
- 3 Specify the appropriate compression level: **Fast**, **Optimal**, or **Maximum**.
- 4 Click **OK**.

26.9.2 Compression Using Migrate Web Interface

To enable and use compression for data transfer:

- 1 On the Edit Migration Details page, go to **Schedule Settings** > **Compression Level**.
- 2 Specify the appropriate compression level: **Fast**, **Optimal**, or **Maximum**.
- 3 Click **Save**.

26.10 Bandwidth Throttling during Data Transfer

Bandwidth throttling enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. Throttling helps to prevent migration traffic from congesting your production network and to reduce the overall load of your PlateSpin Server. You can specify a throughput rate for each migration job. See [“Bandwidth Throttling” on page 48](#).

NOTE: Throttling time is local to the source workload.

- ♦ [Section 26.10.1, “Bandwidth Throttling Using Migrate Client,” on page 300](#)
- ♦ [Section 26.10.2, “Bandwidth Throttling Using Migrate Web Interface,” on page 300](#)

26.10.1 Bandwidth Throttling Using Migrate Client

To enable and use bandwidth throttling for data transfer:

- 1 In the Jobs view, locate the required job.
- 2 In the Network Configuration section of the Migration Job window, view **Bandwidth Throttling**.
- 3 Select the **Enable Throttling** option, specify the required maximum value in Mbps, and optionally a time period during which to enforce the throttling.

If no time interval is defined, bandwidth is throttled to the specified rate at all times by default. If time interval is defined and the migration job executes outside this interval, data is transferred at full speed.
- 4 Click **OK**.

26.10.2 Bandwidth Throttling Using Migrate Web Interface

To enable and use bandwidth throttling for data transfer:

- 1 On the Edit Migration Details page, go to **Schedule Settings > Bandwidth Throttling**.
- 2 Specify the maximum bandwidth to consume in Mbps as the **Throttling Rate**.
A value of Off disables bandwidth throttling.
- 3 Specify one of the following throttling patterns:
 - ♦ **Always:** Always throttle data transfer for the replications. No throttling pattern is specified.
 - ♦ **Custom:** Specify the start and end time and days of the week to throttle data transfer for the replications running in that window.
- 4 Click **Save**.

26.11 Conversion (Data Transfer Method)

Conversion options enable you to specify:

- ♦ How data is transferred from source to target. PlateSpin Migrate supports multiple transfer methods, and their availability depends on your workload and migration job type.

See [“Supported Data Transfer Methods” on page 42](#).

- ♦ The scope of workload data to transfer from the source to the target (**Full Migration** and **Changes only**). Applicable only to Server Sync jobs.

See [“Synchronizing Workloads with Server Sync”](#) on page 423.

26.11.1 Conversion Using Migrate Client

To specify the transfer options for a migration job:

- 1 In the Jobs view, locate the required job.
- 2 In the Job Configuration section of the Migration Job window, click **Conversion**.
- 3 Select the scope and method of data transfer.
- 4 Click **OK**.

26.11.2 Data Transfer Using Migrate Web Interface

- 1 On the Edit Migration Details page, go to **Migration Settings > Transfer Method**.
- 2 Specify the appropriate data transfer method.
- 3 Click **Save**.

26.12 Encrypt Data Transfer

The Encrypt Data Transfer option determines whether to encrypt the data for transmission from the source workload to the target workload. See [Section 2.3, “Security and Privacy,”](#) on page 43.

- ♦ [Section 26.12.1, “Encrypt Data Transfer Using Migrate Client,”](#) on page 301
- ♦ [Section 26.12.2, “Encrypt Data Transfer Using Migrate Web Interface,”](#) on page 301

26.12.1 Encrypt Data Transfer Using Migrate Client

To enable and use encryption for data transfer:

- 1 In the Jobs view, locate the required job.
- 2 In the Network Configuration section of the Migration Job window, click **Encryption**.
- 3 Select **Encrypt Data Transfer**.
- 4 Click **OK**.

26.12.2 Encrypt Data Transfer Using Migrate Web Interface

To enable and use encryption for data transfer for Windows workloads:

- 1 On the Edit Migration Details page, go to **Migration Settings > Data Transfer**.
- 2 Select **Encrypt Data Transfer**.
- 3 Click **Save**.

To enable and use encryption for data transfer for Linux workloads:

- 1 On the Edit Migration Details page, go to **Migration Settings > Transfer Encryption**.

- 2 Select **Encrypt Data Transfer**.
- 3 Click **Save**.

26.13 Virtualization Enhancement Software

For migrations between different virtualization hosts, PlateSpin Migrate provides a mechanism to automatically uninstall virtualization enhancement software, such as VMware Tools.

When converting a workload on a VMware platform that has an earlier version of VMware Tools installed, PlateSpin Migrate identifies the presence of obsolete software and adds a VMware Tools Cleanup step in the migration job.

You must provide administrator credentials to uninstall VMware Tools. The credentials provided must match the administrator-level user account that was logged in during the installation of VMware Tools.

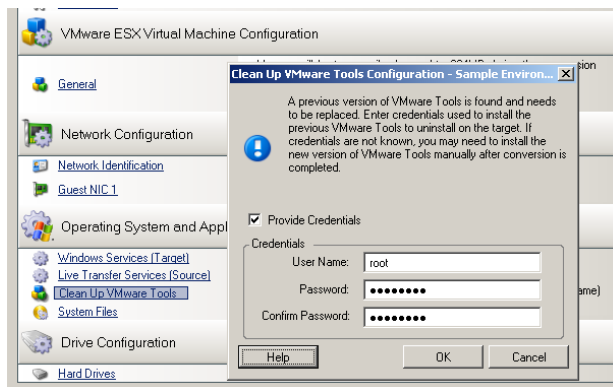
When the earlier version is uninstalled, PlateSpin Migrate proceeds with the installation of the new version of VMware Tools.

NOTE: If you are downgrading a virtual machine that has VMware Tools installed, or if you are converting a virtual machine to another VMware target that has an older version of VMware Tools, the installation of VMware Tools during the configuration of the target will fail.

26.13.1 Replace VMware Tools using Migrate Client

To configure a job to remove or replace VMware Tools during the migration:

- 1 In the Jobs view, select the required workload.
- 2 In the Operating System and Application Configuration section of the Migration Job window, click **Clean up VMware Tools**.



- 3 Depending on the target, PlateSpin Migrate identifies existing instances of VMware Tools and prompts to either replace or remove them, as applicable:
 - ♦ **For non-VMware targets:** The job configuration interface prompts you to uninstall VMware Tools. Provide the same administrator-level credentials used to install the software. If the credentials are unknown, VMware Tools remains on the target machine after migration.

- ♦ **For VMware targets:** The job configuration interface prompts you to replace VMware Tools. Provide the same administrator-level credentials used to install the obsolete version of VMware Tools. If the credentials are unknown, install the new version of VMware Tools manually after the migration completes.

4 Click **OK**.

26.13.2 Replace VMware Tools using Migrate Web Interface

To remove or replace VMware Tools during a migration:

- 1 On the Edit Target Workload Details page, go to **Target Workload Settings > VM Tools**.
- 2 To install the VM tools, select the **Install VM Tools option**. This option is selected by default.
- 3 On the Edit Target Workload Details page, go to **Target Workload Test Settings > VM Tools**.
- 4 To install the VM tools, select the **Install VM Tools option**. This option is selected by default.
- 5 Click **Save**.

26.14 Custom Post-Migration Actions

PlateSpin Migrate Client enables you to execute a custom action on your target. You must define and save your custom actions and their dependencies in advance. See [“Managing Post-Migration Actions \(Windows and Linux\)” on page 121](#).

NOTE: Post-migration actions are supported for peer-to-peer and one-time Server Sync migrations only.

When you are setting up a migration job, select the required action, any required command line parameters, and a timeout as required. You must also provide valid credentials for the target workload. If the target workload credentials are unknown, you can use the credentials of the source workload.

To specify a custom post-migration action for your migration job:

- 1 Start the migration job. For information about starting a migration job, see [Section 26.2, “Initiating a Migration Job,” on page 292](#).
- 2 In the Virtual Machine Configuration section of the Migration Job window, click **Post-Migration**.

The screenshot shows a configuration window for post-migration actions. It has a title bar and a main content area. At the top, there's a 'Select Action' dropdown menu with 'IPconfig_X2V' selected. Below this, the window is split into two columns. The left column is titled 'Execution Parameters' and contains a 'Command Line' text box, an 'Execution Timeout' spinner set to '5' with 'seconds' next to it, and a checked checkbox labeled 'No Timeout'. The right column is titled 'Credentials' and contains 'Username' and 'Password' text boxes, and a checked checkbox labeled 'Use Source Credentials'.

3 Specify the following options:

- ♦ **Select Action:** From the drop-down list, select a custom action previously saved in your library of post-migration actions.
- ♦ **Execution Parameters:** Specify any required command line parameters for the action. If required, specify a timeout.
- ♦ **Credentials:** Provide administrator credentials for the target. If they are the same as those for the source, and if they have been saved, select [Use Source Credentials](#).

26.15 Services or Daemons to Stop before Replication or Cutover

For Live Transfer of data, PlateSpin Migrate provides a mechanism to stop selected services or daemons during the migration. This ensures that data on your source is captured in a consistent state.

If your source workload is running Microsoft SQL Server or Microsoft Exchange Server software, you can configure your migration job to automatically copy the database files of these servers. If you do not require the migration to include the volume containing the databases, consider not stopping these services.

If your source workload includes I/O-intensive application services that might inhibit the ability of the file transfer process to keep up with the changes, consider stopping them during a Live Transfer migration.

After the completion of the migration, services that you select to stop during a Live Transfer migration are automatically restarted on the source, unless you explicitly configure your migration job to power off the source on completion.

For Linux systems, consider using the custom `freeze` and `thaw` scripting capability. See [Section 22.4, “Using Custom Freeze and Thaw Scripts for Linux Block-Level Migrations,” on page 234](#).


TIP: You can globally configure your preferences for stopping selected Windows services during VSS File-based or VSS Block-based Live Transfer performed using the PlateSpin Migrate Client. See [“Configuring Source Service Defaults” on page 119](#).

- ♦ [Section 26.15.1, “Services and Daemons to Stop Using Migrate Client,” on page 304](#)
- ♦ [Section 26.15.2, “Services and Daemons to Stop using Migrate Web Interface,” on page 305](#)

26.15.1 Services and Daemons to Stop Using Migrate Client

To specify which services or daemons you want the system to stop during Live Transfer:

- 1 In the Jobs view, select the required workload.
- 2 In the Operating System and Application Configuration section of the Migration Job window, click [Live Transfer Services/Daemons \(Source\)](#).
- 3 To indicate that you want SQL Server and Exchange Server database files copied during the migration, click [Advanced](#) (applicable to Windows systems only).



- 4 Click **OK**.

26.15.2 Services and Daemons to Stop using Migrate Web Interface

To stop Windows services:

- 1 On the Edit Target Workload Details page, go to **Migration Settings > Services to Stop before Any Replication**.
- 2 Select the services to stop for replication.
We recommend that all the non-VSS compliant services or antivirus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.
- 3 On the Edit Target Workload Details page, go to **Migration Settings > Services to Stop before Cutover with Replication**.
- 4 Select the Windows services that should be permanently stopped on the source workload for cutover with any replication. The services stopped on the source workload during the replication process are not restored afterwards. This does not apply for Test Cutover.
- 5 Click **Save**.

To stop Linux Daemons:

- 1 On the Edit Target Workload Details page, go to **Migration Settings > Daemons to Stop before Any Replication**.
- 2 Select the Linux daemons that you want to be temporarily stopped on the source workload before replication. These daemons will be restored after replication completes.

- 3 On the Edit Target Workload Details page, go to **Migration Settings > Daemons to Stop before Cutover with Replication**.
- 4 Select the Linux services that should be permanently stopped on the source workload for Cutover with any Replication. The daemons stopped on the source workload during the replication process are not restored after Cutover. The stopped daemons are restored after a Test Cutover.
- 5 Click **Save**.

26.16 Service States on Target Windows Workloads

In scenarios such as the following, you might want to change the start-up mode of the services on target Windows workloads:

- ♦ If you do not want a certain Windows service to continue running on a virtualized workload, then configure the job to disable the service on the target workload.
- ♦ If you require that a service on the target starts based on a request from some other service, you can set the start-up mode of the required service to manual.
- ♦ If you want to configure a job to restore the original start-up mode of the service post the migration. For example, you might want to disable a virus scanner during the migration, but restore the start-up mode of the scanner after the migration completes.
- ♦ Some applications on a source workload are known to cause boot failure on the target workload if the corresponding application services are not disabled during the conversion. The **ApplicationsKnownForBootFailuresOnTarget** parameter on the PlateSpin Server Configuration page lists such applications that are likely to cause boot failure on target workload. You can edit this list to add or remove the applications from the list.

A global setting, **ApplicationsKnownForBootFailuresOnTargetDefaultValue**, on the PlateSpin Server Configuration page sets whether the services of all such applications listed in the **ApplicationsKnownForBootFailuresOnTarget** parameter must be selected by default so that the corresponding application services can be disabled on the target during the conversion.

For information about the configuring applications known to cause boot failure on Windows target, see [Section 5.12, “Configuring Applications Known to Cause Boot Failure on Windows Target,” on page 110](#).

For information on modifying or disabling the service state on the target, review the following sections:

- ♦ [Section 26.16.1, “Service States using Migrate Client,” on page 306](#)
- ♦ [Section 26.16.2, “Service States using Migrate Web Interface,” on page 308](#)

26.16.1 Service States using Migrate Client

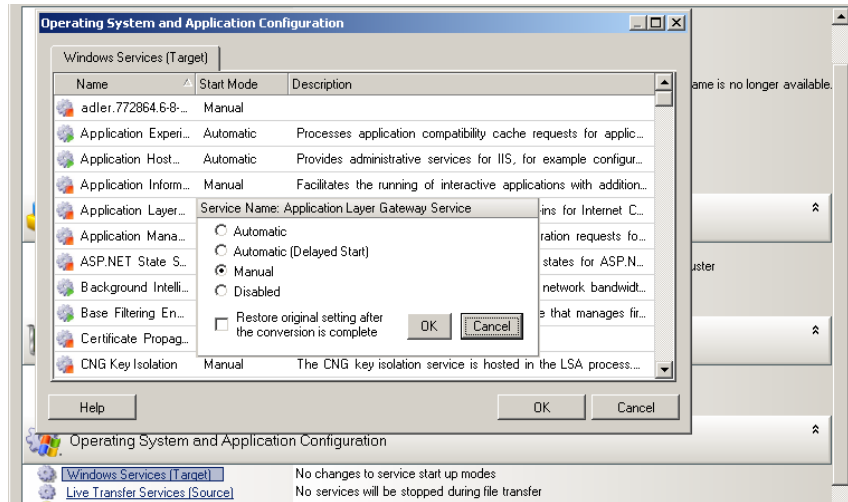
You can specify the preferred run states for services on target Windows workloads that will be enabled after cutover or test cutover. Windows service states options are:

- ♦ Automatic
- ♦ Automatic (Delayed Start)
- ♦ Manual
- ♦ Disabled

Modifying the Windows Service State on the Target Post Migration

To configure post-migration startup mode of Windows services:

- 1 Start the migration job. For information about starting a migration job, see [Section 26.2, “Initiating a Migration Job,” on page 292](#).
- 2 In the Operating System and Application Configuration section of the Migration Job window, click **Windows Services (Target)** and then click an item in the **Start Mode** column.



- 3 Select the desired startup mode.
- 4 To restore the original setting after conversion is complete, select the check box.
- 5 Click **OK**.

Disabling the Windows Boot Service State on the Target Post Migration

- 1 Start the migration job. For information about starting a migration job, see [Section 26.2, “Initiating a Migration Job,” on page 292](#).
- 2 In the Operating System and Application Configuration section of the Migration Job window, click **Windows Services (Target)** and then click **More Options**.

PlateSpin Migrate reviews the existing applications on the source to check if any of the applications listed in the **ApplicationsKnownForBootFailuresOnTarget** configuration parameter is installed on the source. PlateSpin Migrate lists all such applications, which are known to cause boot failure on the target during conversion in the Application Known For Boot Failure panel.

These applications are selected by default if the value of the **ApplicationsKnownForBootFailuresOnTargetDefaultValue** parameter on the PlateSpin Configuration page is set to **true**.

Target Workload Settings

Hostname: ☐ No Change ☐ Set to:

Domain / Workgroup: ☒ Workgroup: WORKGROUP ☐ Join Workgroup: ☐ Join Domain:

Domain Credentials: Username:
Password:
Confirm Password:

Network Connections:

Modify start-up modes on the replicated workload:

Service Name	Mode
App Readiness	Manual
Application Experience	Manual
Application Identity	Manual
Application Information	Manual
Application Layer Gateway Service	Manual

[Apply](#) [Cancel](#) [More Options](#)

Service States on Target VM:

- 4 Click **Apply**.

Disabling the Windows Boot Service State on the Target Post Migration

- 1 On the Edit Migration Details, go to **Migration Settings > Boot Services to Disable on Target**.
- 2 Click **Add Services**.

PlateSpin Migrate reviews the existing applications on the source to check if any of the applications listed in the **ApplicationsKnownForBootFailuresOnTarget** configuration parameter is installed on the source. PlateSpin Migrate lists all such applications, which are known to cause boot failure on the target during conversion in the Application Known For Boot Failure panel.

These applications are selected by default if the value of the **ApplicationsKnownForBootFailuresOnTargetDefaultValue** parameter on the PlateSpin Configuration page is set to **true**.

Migration Settings

Transfer Method:

File Based

Block Based

Use block-based transfer driver

Do not use block-based transfer driver

Encrypt Data Transfer

Source Credentials:

User Name:

Administrator

Password:

Test Credentials

CPU:

Sockets

3

Cores Per Socket

2

Total CPUs

6

Virtual Machine Name:

WIN-ML4R12MVKPB_VM

Configuration File Datastore:

datastore1 (53.1 GB free of 923.0 GB)

Virtual Machine Configuration Path:

/WIN-ML4R12MVKPB_VM/WIN-ML4R12MVKPB_VM.vmx

Disks:

Disk Name	Disk Index	Datastore	Disk Path	Thin Disk
Disk 1	0	datastore1 (53.1 GB)	/WIN-ML4R12MVKPB_VM/WIN-ML4R12MVKPB_VM.vmx	
Disk 2	1	datastore1 (53.1 GB)	/WIN-ML4R12MVKPB_VM/WIN-ML4R12MVKPB_VM.vmx	
Disk 3	2	datastore1 (53.1 GB)	/WIN-ML4R12MVKPB_VM/WIN-ML4R12MVKPB_VM.vmx	
Disk 4	3	datastore1 (53.1 GB)	/WIN-ML4R12MVKPB_VM/WIN-ML4R12MVKPB_VM.vmx	

Volumes:

Include	Name	Used Space	Free Space	Disk Name
<input checked="" type="checkbox"/>	C: (NTFS - Boot)	14.4 GB	5.52 GB	Disk 1
<input checked="" type="checkbox"/>	E: (NTFS)	55.3 MB	4.94 GB	Disk 2
<input checked="" type="checkbox"/>	F: (NTFS)	31.9 MB	465.06 MB	Disk 3
<input checked="" type="checkbox"/>	G: (NTFS)	155.4 MB	2 TB	Disk 4
<input checked="" type="checkbox"/>	\\?Volume{23164ac4-a2d7-11e5-bf83-806e0f8e0903} (NTFS - System)	29.7 MB	70.26 MB	Disk 1

Replication Network for Target:

Select Boot Services to be disabled:

Disable

Service Name

Mode

1394 OHCI Compliant Host Controller

Manual

1394 OHCI Compliant Host Controller (Legacy)

Manual

ACPI Power Meter Driver

Manual

adp94xx

Manual

adpahci

Manual

Applications Known For Boot Failure:

Symantec

Replication Networks for Source:

Services to Stop before any Replication:

Services to Stop for Cutover with Replication:

Boot Services to disable on Target:

- Modify the selection of the applications in the Application Known For Boot Failure panel depending on whether or not you want to disable the boot services of the applications on the target. Selecting an application sets the start-up mode of the corresponding boot service on the target as **Disabled**.
- In the Select Boot Services to be disabled panel, review the modified boot services configuration. Ensure that the settings are correctly configured to prevent any operating system issues.
- Click **Apply**.

26.17 Daemon States on Target Linux Workloads

You can specify the preferred run states for daemons on target Linux workloads that will be enabled after cutover or test cutover. Linux daemons states options are enabled or disabled at the following runlevels and system boot:

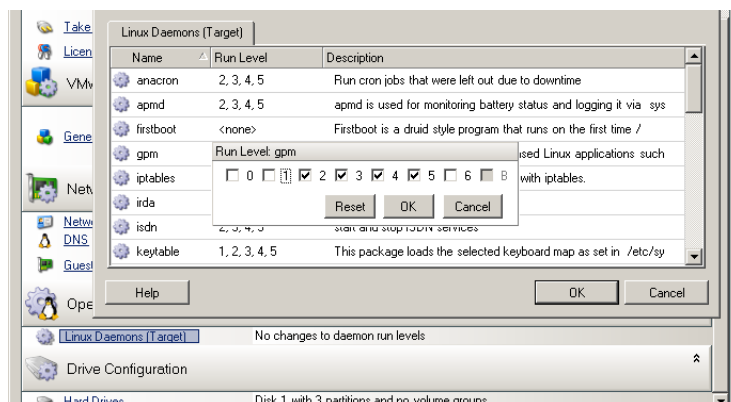
0	Shutdown
1	Single-user mode

2	Unused (user-defined)
3	Full multi user-mode (no GUI)
4	Unused (user-defined)
5	Full multi-user mode with display manager (GUI)
6	Reboot
Boot	Start at power on

26.17.1 Daemon States using Migrate Client

To configure the post-migration run level of Linux daemons:

- 1 Start the migration job. For information about starting a migration job, see [Section 26.2, “Initiating a Migration Job,”](#) on page 292.
- 2 In the Operating System and Application Configuration section of the Migration Job window, click **Linux Daemons (Target)** and then click an item in the **Run Level** column



- 3 Select the desired run levels. Click **OK**.

26.17.2 Daemon States using Migrate Web Interface

To set start states for Linux daemons on the target VM:

- 1 On the Edit Target Workload Details page, go to **Target Workload Settings > Daemon States on Target VM**.
- 2 Select Linux daemons' start conditions on the target VM. Enable the daemon to start by selecting the check boxes at the appropriate runlevels (0 to 6) and Boot.
- 3 Click **Save**.

26.18 Windows HAL or Kernel File Replacements


When you use PlateSpin Migrate Client to migrate Windows workloads with system files (such as a HAL or kernel files) that are incompatible with the target infrastructure, PlateSpin Migrate uses an appropriate file from its library and saves a backup copy of the source file (*.bak) on the target, in the same system directory.

You can use Migrate Client to view the HAL or kernel files that PlateSpin Migrate identifies as those requiring replacement.

To view the files selected for replacement during migration:

- 1 In the Jobs view, select the required workload.
- 2 In the Operating System and Application Configuration section of the Migration Job window, click **System Files**.

Files selected for replacement during migration are listed.



- 3 Click **OK**.

The following warnings might display at the bottom of the dialog box:

Driver Cache is empty	Indicates that you might need to place the necessary files into the local driver cache on the source Windows server (..\Windows\Driver Cache).
The driver cache contains a higher version	PlateSpin Migrate has a partial match with its matrix but the driver cache contains a later version of one or more system files than the one that PlateSpin Migrate will use.
File <filename> will be replaced with lower version	PlateSpin Migrate has not found a match for the system files in its matrix. It will replace the system files with a version that is earlier than the ones that were discovered as the source machine's original system files.
File <filename> will be replaced with higher version	PlateSpin Migrate has not found a match for the system files in its matrix. It will replace the system files with a version that is later than the ones that were discovered as the source machine's original system files.

If warnings appear on the screen, click **More Help** (only available if warnings exist) to learn more.

See also [KB Article 7920815 FAQ: Understanding the System Files Information Screen](https://www.netiq.com/support/kb/doc.php?id=7920815) (<https://www.netiq.com/support/kb/doc.php?id=7920815>).

26.19 Post-Cutover End States for Source and Target Workloads

After a successful cutover, PlateSpin Migrate shuts down or starts the source workload and target workload, depending on the nature of the migration. For example, if the migration goal is to copy the workload, you might want both the source and target workload to be running after cutover. If you are moving a workload, you might want to stop the source workload after cutover and leave the target workload running.

- ♦ [Section 26.19.1, “Workload End States Using the Migrate Client,” on page 313](#)
- ♦ [Section 26.19.2, “Workload End States Using the Migrate Web Interface,” on page 313](#)

26.19.1 Workload End States Using the Migrate Client

To specify non-default post-cutover end states for your source and target:

- 1 In the Jobs view, select the required workload.
- 2 In the Job Configuration section of the Migration Job window, click **End States**.
- 3 Configure the appropriate settings:
 - ♦ **Source Machine End State:** Specify whether to shut down the source workload after a successful cutover. For a workload move, the shut down is selected by default.
 - ♦ **Target Machine End State:** Specify whether to power on, power off, or suspend the target workload after a successful cutover.
- 4 Click **OK**.

26.19.2 Workload End States Using the Migrate Web Interface

To specify post-cutover end states for the source and target workloads after a cutover with replication:

- 1 On the Workloads page, select the prepared workload that you want to migrate.
- 2 Click **Run Migration**.
- 3 On the Workload Commands page, specify the full or incremental replication method.
- 4 For Post-Replication Cutover, enable **Run cutover after successful replication**.
- 5 Specify the appropriate run state for the source and target workload by enabling or disabling the following settings:
 - ♦ Shut down source after cutover
 - ♦ Shut down target after cutover
- 6 Click **Execute**.

PlateSpin Migrate starts the replication for the workload, executes the cutover, then shuts down the source or target as configured.

26.20 Target Workload Settings for VMs

For jobs that involve workload virtualization, PlateSpin Migrate provides a mechanism for specifying target VM configuration options, such as providing a target VM name and a configuration file path, selecting a datastore to use, and allocating virtual memory, in accordance with the features and capabilities of the selected virtualization platform.

If you have resource pools configured on your target virtualization platform, you can select a resource pool for your VM to be assigned to.

NOTE: If your target VMware ESX server is part of a fully automated Distributed Resource Scheduler (DRS) cluster (a cluster with its VM migration automation level set to **Fully Automated**), the newly created target VM's automation level is changed to **Partially Automated** for the duration of the migration. This means that your target VM might power up on a different ESX server from the one initially selected, but migration is prevented from automatic execution.

- ♦ [Section 26.20.1, “Target VM Configuration in Migrate Client,” on page 314](#)
- ♦ [Section 26.20.2, “Target VM Configuration in Migrate Web Interface,” on page 314](#)

26.20.1 Target VM Configuration in Migrate Client

To modify target VM configuration options:

- 1 In the Jobs view, select the required workload.
- 2 In the Virtual Machine Configuration section of the Migration Job window, click **General**.
- 3 Specify the values for the configuration options and click **OK**.

PlateSpin Migrate displays target virtual machine configuration options specific to the selected target and also provides access to advanced configuration options. See:

- ♦ [Section 30.2.1, “Target VM Configuration: VMware ESXi 5 and Later,” on page 368](#)
- ♦ [Section 30.2.2, “Target VM Configuration: VMware ESX 4.1,” on page 369](#)
- ♦ [Section 31.2.1, “Target VM Configuration: Microsoft Hyper-V,” on page 390](#)
- ♦ [Section 32.2.5, “Target VM Configuration: Citrix XenServer,” on page 398](#)

26.20.2 Target VM Configuration in Migrate Web Interface

Migrate Web Interface displays target virtual machine configuration options specific to the selected target. You can specify different values as needed for the target workload test settings.

- 1 On the Edit Target Workload Details page, go to **Target Workload Settings**.
- 2 Modify the target VM settings as appropriate for the target container platform:
 - ♦ **Azure:** [Target Workload Settings](#)
 - ♦ **vCloud:** [Target Workload Settings](#)
 - ♦ **VMware:** [Target Workload Settings](#)
- 3 (Optional) Go to **Target Workload Test Settings**, then modify the target VM test settings as appropriate for the target container platform:
 - ♦ **Azure:** [Target Workload Test Settings](#)
 - ♦ **vCloud:** [Target Workload Test Settings](#)

- ♦ **VMware:** [Target Workload Test Settings](#)

4 Click **Save**.

26.21 Network Identification (Network Connections)

PlateSpin Migrate enables you to manage the network identity and domain registration of your migration target workload and specify related preferences as part of a migration job. By default, a job is configured to preserve a source workload's network identity and domain registration. You can modify the default configuration to suit the objectives of your migration job.

Proper configuration of migration target's network identity is especially important when you are migrating a workload to a different domain, planning to take it off a domain, or if you intend to change the host name of a workload while it is in the domain.

- ♦ [Section 26.21.1, "Network Identification Using Migrate Client," on page 315](#)
- ♦ [Section 26.21.2, "Network Connections Using Migrate Web Interface," on page 317](#)

26.21.1 Network Identification Using Migrate Client

To configure a target workload's network identity options:

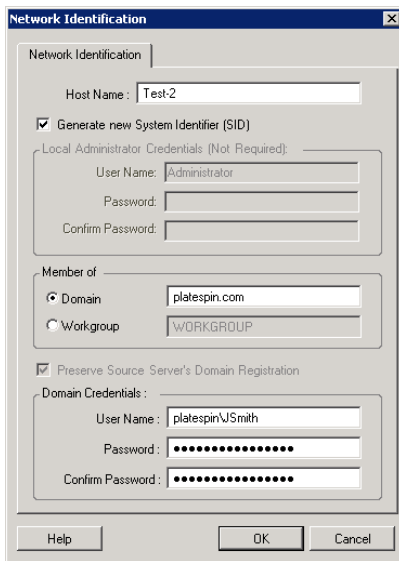
- 1 In the Jobs view, select the required workload.
- 2 In the Network Configuration section of the Migration Job window, click **Network Identification**.
- 3 Specify the options and then click **OK**.

Configuration options vary depending on whether the target machine is Windows or Linux. For information about the configuration options, see the following sections:

- ♦ ["Managing the Identity of Windows Workloads" on page 315](#)
- ♦ ["Managing the Network Identity of Linux Workloads" on page 317](#)

Managing the Identity of Windows Workloads

Use these settings to configure the network identity of your target Windows workload.



Host Name: Specify the desired host name for the target machine.

Generate New SID: When this option is selected, the target workload is assigned a new System Identifier (SID). Credentials are required only for Windows 2008, and must be the credentials for the local (embedded) Administrator account. If this account has been locally renamed on the source, provide the new name.

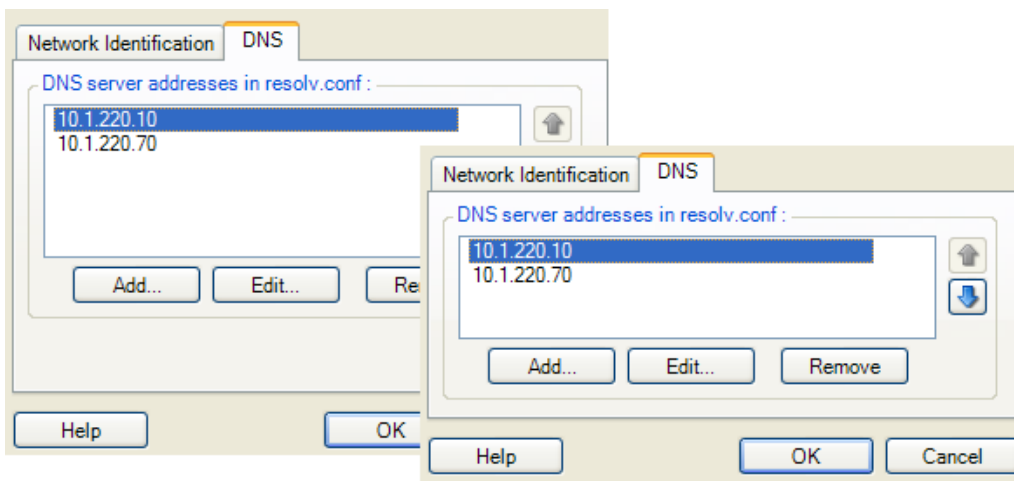
Member of (Domain / Workgroup): Select the required option and type the name of the domain or workgroup that you want the target machine to join.

Preserve Source Server's Domain Registration: Preserves domain registration and ensures that the source server domain registration remains intact during migration. If you disable this option, the source machine's domain account is transferred to the target machine. The source server still appears to be on the domain, but does not have a valid connection.

Domain Credentials: If the target machine is to be part of a domain, specify valid credentials for a user account with permission to add servers to the domain, such as a member of the Domain Admins group or Enterprise Admins group.

Managing the Network Identity of Linux Workloads

Use these settings to configure the network identity of your target Linux workload and DNS server addresses as required.



Network Identification tab: Specify the desired host name for the target server.

DNS tab: Use the **Add**, **Edit**, and **Remove** buttons to manage DNS server entries for the new virtual machine.

26.21.2 Network Connections Using Migrate Web Interface

Migrate Web Interface displays target network configuration options specific to the selected target. You can specify different network values as needed for the target workload test settings.

- 1 On the Edit Target Workload Details page, go to **Target Workload Settings > Network Connections**
- 2 Modify the Network Connections settings as appropriate for the target workload on the target platform:

Parameter	Description
IP Address	Specify DHCP, or provide an IP address for each network connection.
DNS Servers	If you choose static, specify information about your DNS servers.

- ♦ **Azure:** [Target Workload Settings > Network Connections](#)

For Azure, configure these additional settings:

Parameter	Description
Include	<p>If the workload has multiple NICs, select Include for each NIC to be migrated.</p> <ul style="list-style-type: none">♦ At least one NIC is required.♦ The number of NICs to migrate cannot exceed the maximum number of NICs supported by the selected cloud instance. <p>The available NICs apply to the NICs in Target Workload Test Settings.</p>

Parameter	Description
Network and Subnet	For each NIC, specify the network to use and a subnet in that network.
Primary Connection	If you have multiple NICs, specify one of the included NICs to use as the primary connection. The default Primary Connection is the first NIC in the list.
Public IP	If you do not use an Azure VPN, the primary NIC requires a public IP address that is automatically assigned by a Azure.
Resource Group	Type or select a resource group to use for the NIC. The Azure Resource Group setting is the default.

- ♦ **vCloud:** [Target Workload Settings > Network Connection](#)
 - ♦ **VMware:** [Target Workload Settings > Network Connections](#)
- 3 (Optional) Go to **Target Workload Test Settings > Network Connections**, then modify the target VM test settings as appropriate for the target container platform:
- ♦ **Azure:** [Target Workload Test Settings > Network Connections](#)
 - ♦ **vCloud:** [Target Workload Test Settings > Network Connections](#)
 - ♦ **VMware:** [Target Workload Test Settings > Network Connections](#)
- 4 Click **Save**.

26.22 Migration Network (Replication Network)

For each workload migration job, you must properly configure workload networking to enable communications between the source workloads and the target workloads or PlateSpin Replication Environment during the migration process. The network configuration of a target workload must be appropriate for its end state.

- ♦ [Section 26.22.1, “Migration Network Using Migrate Client,” on page 318](#)
- ♦ [Section 26.22.2, “Replication Network Using Migrate User Interface,” on page 324](#)

26.22.1 Migration Network Using Migrate Client

Temporary Networking: Also called *Take Control Network Settings*; they apply to source and target workloads booted into a temporary pre-execution environment. See [“Offline Transfer with Temporary Boot Environment” on page 43](#).

- ♦ [“Temporary \(Take Control\) Network Settings” on page 319](#)
- ♦ [“TCP/IP and Advanced Network Settings” on page 323](#)

Temporary (Take Control) Network Settings

Temporary (Take Control) Network Settings control how source workloads, targets, and the PlateSpin Server communicate among each other during the migration. If required, you can manually specify a temporary network address to your source and target, or configure them to use a DHCP-assigned IP address during the migration.

During Windows and Linux workload migrations, the Temporary Network Settings control the PlateSpin Server's communication with the source and target workloads that are booted into a temporary pre-execution environment. See [“Offline Transfer with Temporary Boot Environment” on page 43](#).

To configure Temporary (Take Control) network settings:

- 1 Start the migration job. For information about starting a migration job, see [Section 26.2, “Initiating a Migration Job,” on page 292](#).
- 2 In the Job Configuration section of the Migration Job window, click **Take Control**.
- 3 To access network interface mapping and TCP/IP settings, click **Configure** in the source and target areas as applicable.
- 4 Click **OK**.

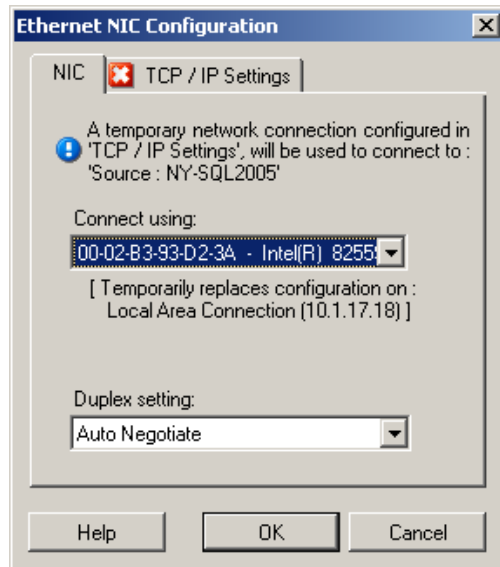
Configuration options for the Temporary networking vary and depend on whether the network interface is virtual or physical, and whether it is connecting a Windows or a Linux workload.

- ♦ [“Temporary \(Take Control\) Network Settings: Physical Network Interfaces” on page 320](#)
- ♦ [“Temporary \(Take Control\) Network Settings: Virtual Network Interfaces” on page 320](#)
- ♦ [“Target Post-Migration Networking” on page 321](#)

Target Take Control network settings are used only during an Offline migration process. On completion, target network settings are read from settings you specify for Target Post-Migration Networking. See [“Target Post-Migration Networking” on page 321](#).

Temporary (Take Control) Network Settings: Physical Network Interfaces

These settings apply only to source physical machines. For target physical machines, Temporary (Take Control) network settings are configured during the boot process that uses the PlateSpin ISO image. See [“Registering and Discovering Details for Target Physical Machines with PlateSpin ISO”](#) on page 206.




Connect using: If multiple network adapters are present, select the adapter that can communicate with both the PlateSpin Server and the target.

Duplex setting: Use the drop-down list to select network card duplexing. It must match the duplex setting for the switch to which the network interface is connected. When the source is connected to switch ports that are set to 100 Mbit full duplex and cannot be changed to auto negotiation, select **Force NIC to Full Duplex**.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See [“TCP/IP and Advanced Network Settings”](#) on page 323.

Temporary (Take Control) Network Settings: Virtual Network Interfaces

These settings apply to both source and target Take Control network settings.



Map to Virtual Network: From the drop-down list, select the virtual switch or network to use for communication during an Offline migration. If multiple virtual network adapters are present, select the adapter that can communicate with both the PlateSpin Server and the source machine. This network can differ from the network on which the target virtual machine will run after the migration.

VLAN ID: (Applicable for target machine on a Hyper-V server only) Enable this option to specify the virtual network ID to be used on the target machine. If you do not specify this ID, then the virtual network ID of the source machine is used by default.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See [“TCP/IP and Advanced Network Settings” on page 323](#).

Target Post-Migration Networking

Target post-migration network settings defined in a migration job control the network configuration of a target after the migration is complete. This applies to both physical and virtual network interfaces.

During workload migration, the target workload’s post-migration network settings are configured while the workload is booted into a pre-execution environment.

To configure target post-migration network settings:


- 1 Start the migration job. For information about starting a migration job, see [Section 26.2, “Initiating a Migration Job,” on page 292](#).
- 2 In the Network Configuration section of the Migration Job window, do one of the following:
 - ♦ **For target virtual machines:** click **Guest NIC**.
 - ♦ **For target physical machines:** click **Network Connection**.
- 3 Configure the options as required and click **OK**.

The Configuration options for the target post-migration network settings vary and depend on whether the network interface is virtual or physical, and whether it is connecting a Windows or a Linux workload. For more information about the options, review the following sections:

- ♦ [“Post-Migration Networking for Physical Network Interfaces \(Windows and Linux\)” on page 322](#)
- ♦ [“Post-Migration Networking for Virtual Network Interfaces \(Windows and Linux\)” on page 322](#)

Post-Migration Networking for Physical Network Interfaces (Windows and Linux)

Use these settings to configure the post-migration network settings of a workload being migrated to physical hardware.




Connect using: If multiple network adapters are present, select the adapter that can communicate with the PlateSpin Server.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See [“TCP/IP and Advanced Network Settings” on page 323](#).

Post-Migration Networking for Virtual Network Interfaces (Windows and Linux)

By default, PlateSpin Migrate configures a migration job to create a virtual NIC for each NIC found on the source. For post-migration connectivity, ensure that the target virtual NIC is mapped to the appropriate virtual network on the target virtualization platform.



Include in Conversion: When this option is selected, PlateSpin Migrate creates a virtual NIC for a source NIC.

Map to Virtual Network: Select the virtual network that will be used on the target VM. Choose a virtual network that allows the target VM to communicate with the server.

Start connected: Enable this option to connect the virtual network interface when starting the ESX target machine.

VLAN ID: (Applicable for target machine on a Hyper-V server only) Enable this option to specify the virtual network ID to be used on the target machine. If you do not specify this ID, then the virtual network ID of the source machine is used by default.

TCP/IP Settings tab: Click the tab to access TCP/IP and advanced network settings. See [“TCP/IP and Advanced Network Settings” on page 323](#).

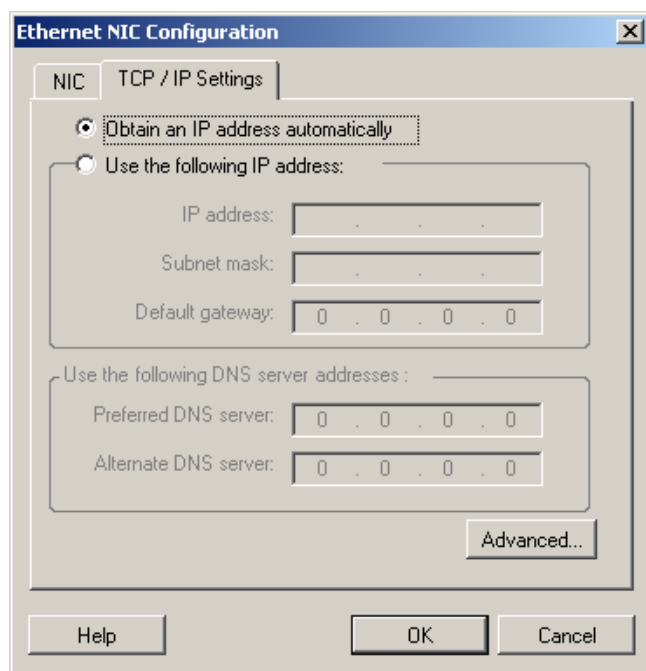
TCP/IP and Advanced Network Settings

PlateSpin Migrate provides a standard network configuration interface to both source and target network settings, and for both Temporary and target post-migration networking. Configuration settings vary slightly, depending on the operating system.

- ♦ [“TCP/IP and Advanced Network Settings \(Windows\)” on page 323](#)
- ♦ [“TCP/IP and Advanced Network Settings \(Linux\)” on page 324](#)

TCP/IP and Advanced Network Settings (Windows)

The following are standard TCP/IP and advanced network settings for Windows workloads:



Obtain an IP address automatically: When this option is selected, the workload uses an IP address automatically assigned by a DHCP server during the migration process.

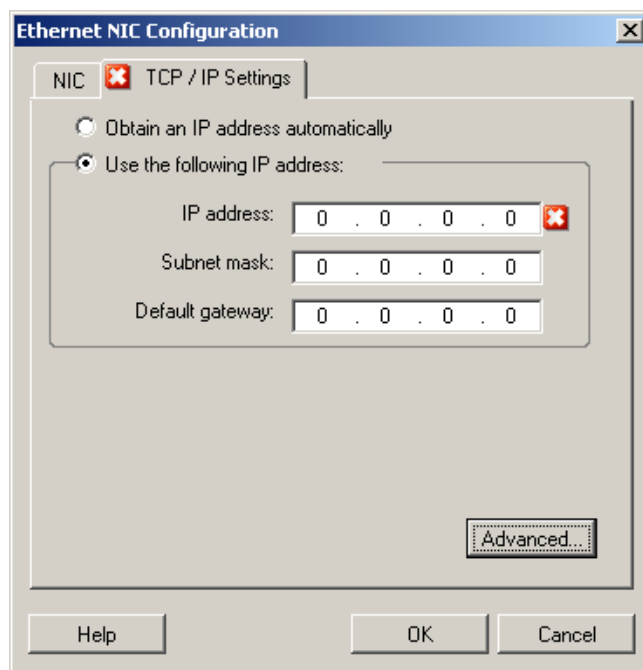
Use the following IP address: Select this option to specify a static IP address.

Use the following DNS server addresses: If required, specify preferred and alternative DNS server addresses.

Advanced: Click this button to access advanced TCP/IP configuration settings, then specify or edit default gateway, DNS server, and WINS server information as required.

TCP/IP and Advanced Network Settings (Linux)

The following are standard TCP/IP and advanced network settings for Linux workloads:



Obtain an IP address automatically: When this option is selected, the workload uses an IP address automatically assigned by a DHCP server during the migration process.

Use the following IP address: Select this option to specify a static IP address.

Advanced: Click this button to access DNS configuration settings, then specify preferred and alternate DNS server addresses as required. You can also indicate whether you want DNS addresses copied to the `resolv.conf` file located in your target's `/etc` directory.

26.22.2 Replication Network Using Migrate User Interface

To specify the Replication Network for migration to Azure:

- 1 In the Web Interface, select the Workload to go to the Target Configuration page, then click **Edit**.
- 2 Navigate to **Target Workload Settings > Network Connections**, then specify the Primary NIC. Migrate uses the Primary NIC as the Replication NIC.

- 3 Under **Migration Settings** in **Replication Network for Target**, specify the replication network settings:
 - 3a Select a network and subnet to use for replication traffic.
 - 3b If you do not use an Azure VPN, click **Edit**, then select **Create Public IP**.

When no VPN is present in the deployment, the replication NIC requires a public IP address that is automatically assigned by Azure.
 - 3c Specify a resource group to use for the replication network.

The **Azure Resource Group** setting is the default. To specify a different resource group, click **Edit** and do one of the following:

 - ♦ Type the name to use when PlateSpin creates a new resource group.
 - ♦ Select an existing resource group from the list.
 - 3d Select one of the following:
 - ♦ **DHCP**: Obtain an IP address automatically assigned by a DHCP server.
 - ♦ **Static**: Specify a static private IP address, a subnet mask, and a gateway IP address.
- 4 In **Replication Networks for Source**, specify one or more network interfaces (NIC or IP address) on the source workload to use for replication traffic that are valid for communications with the replication environment.

To specify the Replication Network for migration to vCloud:

- 1 In the Web Interface, select the Workload to go to the Target Configuration page, then click **Edit**.
- 2 Under **Migration Settings** in **Replication Network for Target**, specify a network interface (NIC or IP address) on the target to use for replication traffic.
- 3 Under **Migration Settings** in **Replication Networks for Source**, specify one or more network interfaces (NIC or IP address) on the source to use for replication traffic.
 - ♦ **DHCP**: Obtain an IP address automatically assigned by a DHCP server.
 - ♦ **Static - Manual**: Specify a static IP address.
 - ♦ **Static - IP Pool**: Select this option to automatically issue IP address from the IP pool.

For Windows workloads that have more than one NIC, select the connection for each NIC.

For this setting, you can also specify an MTU value that the PlateSpin Migrate Linux RAM Disk (LRD) replication network can use. Setting a low value helps to avoid jabber over networks. For example: a VPN.

The default value is an empty string. When networking is configured in the LRD, it allows the network device to set its own default, which is usually 1500. However, if you specify a value, PlateSpin Migrate adjusts the MTU when it configures the network interface.

To specify the Replication Network for migration to VMware:

- 1 In the Web Interface, select the Workload to go to the Target Configuration page, then click **Edit**.
- 2 Under **Migration Settings** in **Replication Network for Target**, specify a network interface (NIC or IP address) on the target to use for replication traffic.
- 3 Under **Migration Settings** in **Replication Networks for Source**, specify one or more network interfaces (NIC or IP address) on the source to use for replication traffic.

26.23 Storage Disks and Volumes

PlateSpin Migrate provides mechanisms for configuring your migration job to handle your workload volumes and their physical or virtual layout in the target infrastructure. For information about the supported storage, see [Section 2.1.3, “Supported Workload Storage,” on page 32](#).

Storage layout and volume configuration settings depend on the job configuration mode (Advanced or Wizard), migration type, target virtualization platform, and source operating system.


The following topics provide additional information:

- ♦ [Section 26.23.1, “Storage Disks and Volumes Using Migrate Client,” on page 326](#)
- ♦ [Section 26.23.2, “Storage and Volume Using Migrate Web Interface,” on page 330](#)

26.23.1 Storage Disks and Volumes Using Migrate Client

To access drive configuration options:

- ♦ In the **Drive Configuration** of the Migration Job windows, click **Hard Drives**.




Settings vary depending on the target system.

- ♦ [“Windows Drive Configuration” on page 326](#)
- ♦ [“Linux Drive and LVM Volume Configuration” on page 327](#)
- ♦ [“Target VM-Specific P2V/V2V Drive Configuration” on page 330](#)
- ♦ [“Volume Mapping in Server Sync” on page 330](#)

Windows Drive Configuration

Use these settings to select the volumes to copy during the migration:



Copy: Select the volumes to be copied during the migration.

New Free Space: To resize the volume during the migration, specify the desired amount of free space. PlateSpin Migrate automatically adjusts **New Size**.

New Size: To resize the volume during the migration, specify the desired size. PlateSpin Migrate automatically adjusts **New Free Space**.

To Disk: Select which hard drive the volume will be copied to on the physical target machine.

Preserve Partitions: Click this column to determine if an existing vendor partition should remain intact during the migration. If the partitions are not selected, PlateSpin Migrate permanently removes the partitions from the server.

Linux Drive and LVM Volume Configuration

Use these settings to select the volumes and non-volume source spaces to copy and size during the migration. If LVM is installed on the source, a **Volume Group** tab provides you with corresponding options.

- ♦ [“Handling Linux Disks and Volume Groups” on page 328](#)
- ♦ [“Linux Drive and LVM Volume Configuration \(Settings Tab\)” on page 328](#)
- ♦ [“Linux Drive and LVM Volume Configuration \(Volume Groups Tab\)” on page 329](#)

Handling Linux Disks and Volume Groups

The PlateSpin Migrate Client provides you with Linux-specific user interface elements that provide you with options to properly handle your Linux storage.

Note the following sequence of steps that you must take for properly configuring and mapping newly-added disks and volume groups.

- 1 After adding a new disk, go to the **Volume Groups** tab and map the required volume group name by selecting the **Include** option.

See [Linux Drive and LVM Volume Configuration \(Volume Groups Tab\)](#).

- 2 Specify Size in Allocation for Volume Group Box
- 3 For each added disk, specify the required size in the corresponding **Allocation for Volume Group** field.

After the system focus shifts away from the field, the size of the newly-added disk is updated dynamically.

Linux Drive and LVM Volume Configuration (Settings Tab)

Use these settings to select source volumes to copy, non-volume source spaces to re-create and size, and target disks to repartition and populate.

Drive Configuration

Settings

Volume Groups

Virtual disks to create:

Add

Remove Unused Disks

Disk	Datastore	Size	Thin	File Name
Virtual disk 0	newLUN-VC...	10 GB	<input type="checkbox"/>	/linux_VM/linux_VM_1.vmdk

Select volumes to copy and size:

Inclu...	Volume	Free Space	Size	New Free Space	New SI...	Disk/Volume Group/EVMS Volu...
<input checked="" type="checkbox"/>	/	6.6 GB	9 GB	6.6 GB	9 GB	Disk 0

Select non-volume storage to recreate and size:

Include	Type	Partition	Size	Is Swap	New Size	Disk/Volume Group/EVMS Volume
<input checked="" type="checkbox"/>		/dev/sda1	1 GB	<input checked="" type="checkbox"/>	1 GB	Disk 0

Help

OK

Cancel

Include: Select the volumes or non-volume source spaces to be copied or re-created and sized during the migration.

New Free Space: To resize the volume during the migration, enter the desired amount of free space. PlateSpin Migrate automatically adjusts **New Size**.


New Size: To resize the volume during the migration, enter the desired size. PlateSpin Migrate automatically adjusts **New Free Space**.

Disk/Volume Group: Select which hard drive or volume group the volume will be copied to on the physical target machine.

Preserve Partitions: For each disk, click the corresponding cell in this column to select existing vendor partitions to preserve during the migration. If the partitions are not selected, PlateSpin Migrate permanently removes them from the server.

Linux Drive and LVM Volume Configuration (Volume Groups Tab)

Use these settings to manage volume groups.



Add Volume Group: Creates a volume group on the target machine that is not present on the source machine.

Rename Volume Group: Renames a volume group that is being copied from the source to the target.

Delete Volume Group: Deletes a volume group so that it is not created on the target machine. The volumes assigned to the volume group can be reassigned to other locations by using the **Settings** tab (by default, they are assigned to disk).

Allocation for Volume Group: To allocate space on disks to a volume group, select the volume group, then select the disks to include in it. Specify the amount of space to be allocated to it on each included disk.

Target VM-Specific P2V/V2V Drive Configuration

When you configure a peer-to-peer virtualization job, the job configuration window provides access to settings specific to the target virtualization platform.

PlateSpin Migrate displays target virtual machine drive configuration settings specific to the selected target:

- ♦ [Section 30.2.3, “Drive Configuration: VMware ESX,” on page 371](#)
- ♦ [Section 31.2.2, “Drive Configuration: Hyper-V,” on page 392](#)

Volume Mapping in Server Sync

When you are using Server Sync to synchronize two Windows or Linux workloads, PlateSpin Migrate Client provides you with the capability to specify the required mapping between source volumes and existing volumes on the target. See [Section 37.4, “Server Sync Volume Mapping,” on page 427](#).

26.23.2 Storage and Volume Using Migrate Web Interface

- 1 On the Edit Target Workload Details page, go to **Target Workload Settings > Migration Settings**.
- 2 Configure the following options:

Setting Name	Description
Disks	Specify the path to the hard disk on the target virtual machine.
Volumes	Select volumes to be included in the target for migration.
NTFS Cluster Size	(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see the Microsoft Support KB Article 140365 .
Non-volume Storage	(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups	(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.
Volume Groups	(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the Converted Logical Volumes section of the settings.
Converted Logical Volumes	(For Linux Workloads) Specify one or more LVM logical volumes to be migrated for a Linux workload.

- 3 Click **Save**.

27 Migration to Amazon Web Services

- ♦ [Section 27.1, “Planning for Migration to Amazon Web Services,” on page 331](#)
- ♦ [Section 27.2, “Configuring Migration of a Workload to Amazon Web Services,” on page 332](#)
- ♦ [Section 27.3, “Configuring Full Migration of a Workload to Amazon Web Services,” on page 332](#)

27.1 Planning for Migration to Amazon Web Services

Before you begin migrations to your cloud environment in Amazon Web Services, ensure that your migration environment meets the following guidelines:

Supported Cloud Platforms

- ♦ See [“Supported Target Cloud Platforms” on page 40](#).

Supported Workloads

- ♦ See [“Supported Workloads For Migration to Amazon Web Services” on page 31](#), as appropriate for the target Azure cloud environment.

Network Access and Communications

- ♦ See [“Access and Communication Requirements across your Migration Network” on page 49](#).

Prerequisites

- ♦ See [Chapter 8, “Prerequisites for Migration to Amazon Web Services,” on page 137](#).

Targets and Workloads

- ♦ **Target AWS EC2 cloud account (automated):** See [Section 19.4.1, “Target Discovery in the Migrate Client,” on page 200](#).
- ♦ **Source Workloads:** See [Section 20.6.2, “Workload Discovery in the Migrate Web Interface,” on page 215](#).

Additional Information

- ♦ [Amazon Elastic Compute Cloud Documentation \(https://aws.amazon.com/documentation/ec2/\)](https://aws.amazon.com/documentation/ec2/)
- ♦ [AWS Managed VPN \(http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html\)](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html) in the Amazon Virtual Private Cloud User Guide.
- ♦ [Your Customer Gateway \(http://docs.aws.amazon.com/AmazonVPC/latest/NetworkAdminGuide/Introduction.html\)](http://docs.aws.amazon.com/AmazonVPC/latest/NetworkAdminGuide/Introduction.html) in the Amazon Virtual Private Cloud Network Administrator Guide.

27.2 Configuring Migration of a Workload to Amazon Web Services

Depending on the scope of data you want to transfer from the source to target, see one of the following:

- ♦ **Full Migration:** See [Section 27.3, “Configuring Full Migration of a Workload to Amazon Web Services,” on page 332](#).
- ♦ **Server Sync:** See [Section 37.3, “Server Sync to Amazon Web Services,” on page 426](#).

27.3 Configuring Full Migration of a Workload to Amazon Web Services

- 1 Launch the PlateSpin Migrate Client.
- 2 Discover your source workload. See [Chapter 20, “Discovering Source Workloads,” on page 211](#).
- 3 Right-click the discovered source workload and click **Prepare AWS Target**.
- 4 Perform the following steps:
 - 4a To log in to your AWS account for the first time, specify the following:
 - ♦ **AWS Access Key ID and AWS Secret Access Key:** The two parts of an access key that is a valid set of security credentials required to access AWS account.
 - ♦ **AWS Region:** The geographic area to which your AWS account has access.

During subsequent logins, these options are automatically populated with the values that you specified during your earlier login.
 - 4b Click **Update Region List** to update the supported AWS Regions.
 - 4c Click **Login**.
 - 4d (Conditional) If you are logging in for the first time, the Replication Environment Details dialog displays. Specify the following and then click **Save**:
 - ♦ **Linux AMI ID:** Specify the ID of the PlateSpin Replication Environment - Linux AMI.
 - ♦ **Windows AMI ID:** Specify the ID of the PlateSpin Replication Environment - Windows or the PlateSpin Replication Environment - Windows(BYOL) AMI depending on whether you want AWS to activate the OS license or you want to bring your own Microsoft license.
 - ♦ **Snapshot ID:** The snapshot ID of the PRE AMI that you noted down when you created a snapshot from the private instance of AMI. See [“To create a snapshot using the PRE AMI:” on page 143](#).
 - ♦ **Keypair Name:** Specify the name of the key pair associated with the AMI snapshot.
 - ♦ **Keypair File Path:** Specify the location where you have stored the key pair (.pem) file.
 - 4e (Conditional) If you are logging in for the first time, the Migrate Server Details dialog displays. Specify the Migrate server IP and the credentials to register the target instance with the Migrate server, then click **Save**.
 - 4f In the Prepare AWS Target dialog, ensure that the transfer scope is set to **Full Migration** and do the following:
 - 4f1 Select the VPC you want to use.

The values for the **Security Group** and **Subnet** options are automatically populated based on the VPC you select. The **Instance Type** option is automatically populated based on the source workload. You can select different values based on your requirements.

- 4f2** Specify the host name to register the target with the Migrate server. By default, it displays `sourcename_RepEnv`.
- 4f3** To configure the Network Interfaces, click **Configure**:
 - 4f3a** Based on your requirement, you can add or remove NICs.
 - 4f3b** Select one of the following options
 - ♦ **Obtain settings automatically:** To enable the network settings to be automatically assigned by a DHCP server, you select this option.
 - ♦ **Use the following settings:** To specify the network settings, select this option. The IP address must be unique within the supported subnet.
- 4f4** The Target Instance Storage section displays the disk size and device name for all the disks available on the source. The disks are sorted based on the device name. You can choose to delete the existing data disks or click **Add Disk** to add more disks based on your requirements.
- 4g** (Optional) Click the settings icon to modify the replication details or the Migrate server details.
- 4h** Click **Prepare**.
- 4i** On completion of the registration job, you are prompted to view the logs. To view the logs, click **Yes**.
- 4j** Click **Close**.
- 5** In the PlateSpin Migrate Client, the registered target instance is listed with the name as specified in [Step 4f2](#). Drag the source workload to this target to start a migration job.
- 6** Configure the required parameters of the job. See [Chapter 26, “Configuration Essentials,” on page 291](#).
- 7** Run the job.

28 Migration to Microsoft Azure

- ♦ [Section 28.1, “Planning for Migration to Microsoft Azure,” on page 335](#)
- ♦ [Section 28.2, “Configuring Migration of a Workload to Microsoft Azure,” on page 336](#)

28.1 Planning for Migration to Microsoft Azure

Before you begin migrations to your cloud environment in Microsoft Azure, ensure that your migration environment meets the following guidelines:

Supported Cloud Platforms

- ♦ See [“Supported Target Cloud Platforms” on page 40](#).

Supported Workloads

- ♦ See [“Supported Workloads For Migration to Microsoft Azure” on page 27](#), as appropriate for the target Azure cloud environment.

Network Access and Communications

- ♦ See [“Access and Communication Requirements across your Migration Network” on page 49](#).

Prerequisites

- ♦ See [Chapter 9, “Prerequisites for Migration to Microsoft Azure,” on page 147](#).

Targets and Workloads

- ♦ **Target Azure cloud subscription (automated):** See [Section 19.4.2, “Target Discovery in the Web Interface,” on page 202](#).
- ♦ **Source Workloads:** See [Section 20.6.2, “Workload Discovery in the Migrate Web Interface,” on page 215](#).

Additional Information

- ♦ See [“Create a Site-to-Site Connection in the Azure Portal”](#) in the *Microsoft Azure VPN Gateway Documentation*.
- ♦ See [“Create a VNet with a Site-to-Site VPN Connection Using PowerShell”](#) in the *Microsoft Azure VPN Gateway Documentation*.

28.2 Configuring Migration of a Workload to Microsoft Azure

When you add or discover a workload, the workload is listed on the Workloads page and the status is set as **Not Configured**. Before you migrate the workload, you must configure the workload for migration:

- 1 Launch the PlateSpin Migrate Web Interface.
- 2 If you have not configured a Microsoft Azure Location as a migration target, click **Targets > Add Target**, and then configure the target Azure container.
See [Section 3.3.3, “Targets,” on page 74](#).
- 3 On the Workloads page, select the workload you want to configure.
- 4 Click **Configure Migration**.
- 5 Specify the **Initial Transfer Method** for replication based on the scope of data you want to transfer from the source to the target:
 - ♦ **Full Replication:** Migrate replicates the full volume from the source to the target.
 - ♦ **Incremental Replication:** Migrate replicates only differences in data from the source to the target, provided the workloads have similar operating system and volume profiles.

NOTE: PlateSpin Migrate does not support Incremental Replication for the initial replication of data to existing target workloads in Azure Cloud. However, you can schedule Incremental Replications for subsequent replication of data. See [Incremental Recurrence](#) in [Step 8](#).

- 6 Select an existing Microsoft Azure Location target to which you want to migrate the source workload.
To verify availability of Premium Storage for the target location, refer to the [Microsoft Azure Products Available by Region \(https://azure.microsoft.com/en-us/regions/services/\)](https://azure.microsoft.com/en-us/regions/services/).
- 7 Click **Configure Migration**.

8 Configure the following settings:

Setting Name	Description
Schedule Settings	
Incremental Recurrence	<p>Specify the time and pattern when you want to run incremental replications after the first full replication, or start each incremental replication manually.</p> <p>The default setting is None. The incremental replications are unscheduled.</p> <p>To set or modify the incremental recurrence time and pattern:</p> <ol style="list-style-type: none"> 1. Click Edit. 2. For Begin the recurrence schedule, set the date and time when you want to begin the scheduled incremental replications. You can type the date (dd/mm/yyyy) or click the Calendar icon to select the date. By default, the run time is 12:00:00 a.m. (hh:mm:ss a.m. or p.m.). 3. For Recurrence run setting, set the pattern to follow for scheduled incremental replications: <ul style="list-style-type: none"> ♦ Daily: The replication takes place on the specified daily intervals or on weekdays every week for a period of 60 days from the time the replication starts. ♦ Weekly: The replication takes place at specified intervals for a period of 8 weeks from the time the replication starts. ♦ Monthly: The replication takes place at specified intervals for a period of 2 months from the time the replication starts. <p>NOTE:</p> <ul style="list-style-type: none"> ♦ Scheduled incremental replications are skipped until the first full replication is complete. ♦ Scheduled incremental replications take place for a maximum period of 60 days from the time that the scheduled incremental replication runs begin.
Full Replication	<p>Specify when you want the first full replication to run, or start the first full replication manually. The first full replication is a one-time event, but the run is attempted daily as scheduled until the first replication begins and completes successfully.</p> <p>The default setting is None. The first full replication is unscheduled.</p> <p>NOTE: You must prepare the workload prior to the scheduled time or the manual start. The full replication cannot run unless the target VM exists and the workload preparation is complete. If they are not ready, Migrate skips the scheduled full replication and retries it at the scheduled time on the next day.</p> <p>To set or modify the schedule for the first full replication:</p> <ol style="list-style-type: none"> 1. Click Edit. 2. Click Start, then set the date and time when you want to start the first full replication. You can type the date (dd/mm/yyyy) or click the Calendar icon to select the date. By default, the run time is 12:00:00 a.m. (hh:mm:ss a.m. or p.m.).

Setting Name	Description
Blackout Window	<p>Specify a replication blackout window that suspends scheduled replication activities for a specified period of time and pattern. For example, suspend replications during peak network utilization hours or to prevent conflicts between VSS-aware software and the PlateSpin VSS block-level data transfer component.</p> <p>The default setting is None. No blackout window is scheduled.</p> <p>To set or modify a blackout window:</p> <ol style="list-style-type: none"> 1. Click Edit. 2. Specify the start and end time for the blackout period. The blackout start and end times are based on the system clock on the PlateSpin Server. 3. Select Daily, Weekly, or Monthly to enable a blackout window, then set the recurrence pattern.
Compression Level	<p>This setting controls whether data is compressed during transmission between the source and target workloads, and the level of data compression applied. See Section 2.4.3, "Data Compression," on page 48.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> ♦ None: No compression. ♦ Fast: Consumes the least CPU resources on the source, but yields a lower compression ratio. ♦ Optimal: (Default) Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option. ♦ Maximum: Consumes the most CPU resources on the source, but yields a higher compression ratio.
Bandwidth Throttling	<p>Bandwidth throttling enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. Throttling helps to prevent migration traffic from congesting your production network and to reduce the overall load of your PlateSpin Server. You can specify a throughput rate for each migration job.</p> <p>Throttling is disabled by default with a Throttling Rate value of Off.</p> <p>To throttle replications to a specified rate:</p> <ol style="list-style-type: none"> 1. Specify a maximum throughput value in Mbps for data transfer for the workload. 2. Specify the throttling pattern: <ul style="list-style-type: none"> ♦ Always: Always throttle data transfer for the replications. ♦ Custom: Specify the time and days to throttle data transfer for the replications running in that window. Throttling time is local to the source workload.

Setting Name	Description
Migration Settings	
Transfer Method	<p>(For Windows Workloads) Select a data transfer mechanism and security through encryption. See Section 2.2, "Supported Data Transfer Methods," on page 42.</p> <p>To enable encryption, select the Encrypt Data Transfer option. See Section 2.3, "Security and Privacy," on page 43.</p> <p>NOTE: The Offline Transfer with Temporary Boot Environment transfer method is not applicable for the Web interface.</p>
Transfer Encryption	<p>(For Linux Workloads) To enable encryption, select the Encrypt Data Transfer option. See Section 2.3, "Security and Privacy," on page 43.</p>
Source Credentials	<p>Specify the credentials required for accessing the workload. See Section 20.3, "Discovery Guidelines for Source Workloads," on page 212.</p>

Setting Name	Description
Cloud Instance Size	<p>Select the cloud instance size appropriate for your workload and the Storage account type for the target container.</p> <p>IMPORTANT: The Cloud Instance Size must be of the same storage type as the target account: Standard Storage or Premium Storage. Otherwise, you receive a validation error. To verify the availability of Premium Storage for the target location, refer to the Microsoft Azure Products Available by Region.</p> <p>By default, Migrate selects a cloud instance size that supports the same Storage account type and that most closely matches your source workload for the following components:</p> <ul style="list-style-type: none"> ◆ Total number of cores ◆ Amount of memory ◆ Number of data disks ◆ Number of NICs <p>The default instance either meets or exceed the settings for each of these components on the source workload. However, you can choose a smaller instance size based on your requirements:</p> <ul style="list-style-type: none"> ◆ The target VM uses the allowed CPU and memory for the instance size. To reduce the number of CPUs or amount of memory on the target workload: <ol style="list-style-type: none"> 1. Select a smaller cloud instance size with fewer CPUs or less memory that best fits your needs. ◆ The target VM uses up to the maximum allowed number of data disks for the instance size. To migrate only some of the data disks: <ol style="list-style-type: none"> 1. Select a smaller cloud instance size with fewer data disks that best fits your needs. 2. Deselect the volumes that should not be migrated until the number of disks for migration fits the selected instance. ◆ The target VM uses up to the maximum allowed number of NICs for the instance size. To migrate only some of the NICs: <ol style="list-style-type: none"> 1. Select a cloud instance size with fewer NICs that best fits your needs. At least one NIC is required. 2. Under Target Workload Settings, deselect the NICs that should not be migrated until the number of NICs for migration fits the selected instance. <p>NOTE: The number of data disks consumed by volumes on the target VM cannot exceed the maximum number of data disks supported by the selected cloud instance.</p> ◆ In the Cloud Instance Size list, the Supports Premium Storage column indicates the storage account type of the instance: Standard Storage (No) or Premium Storage (Yes). Ensure that your new instance size supports the same storage account type as the target container.

Setting Name	Description
Azure Resource Group	<p>Specify a resource group to use for the target VM resources. Do one of the following:</p> <ul style="list-style-type: none"> ◆ Allow PlateSpin to create a new resource group with the default name: <code><hostname>-VM-Resources</code> ◆ Type the name to use when PlateSpin creates a new resource group. ◆ Select an existing resource group from the list.
Virtual Machine Name	Specify a display name for the new virtual machine.
Disks	Specify the path to the hard disk on the target virtual machine.
Volumes	Select volumes to be included in the target for migration.
NTFS Cluster Size	(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see the Microsoft Support KB Article 140365
Non-volume Storage	(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups	(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.
Volume Groups	(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the Converted Logical Volumes section of the settings.
Converted Logical Volumes	(For Linux Workloads) Select LVM logical volumes to be included in the target for migration.
Replication Network for Target	<p>The replication NIC is the primary NIC that you specify in Target Workload Settings > Network Connections.</p> <ol style="list-style-type: none"> 1. Select a network and subnet to use for replication traffic. 2. If you do not use an Azure VPN, the replication NIC requires a public IP address that is automatically assigned by Azure. Click Edit, then select Create Public IP. 3. Specify a resource group to use for the replication network. The Azure Resource Group setting is the default. To specify a different resource group, click Edit and do one of the following: <ul style="list-style-type: none"> ◆ Type the name to use when PlateSpin creates a new resource group. ◆ Select an existing resource group from the list. 4. Select one of the following: <ul style="list-style-type: none"> ◆ DHCP: Obtain an IP address automatically assigned by a DHCP server. ◆ Static: Specify a static private IP address, a subnet mask, and a gateway IP address. The IP address must be unique within the supported subnet.

Setting Name	Description
Replication Networks for Source	<p>Specify one or more network interfaces (NIC or IP address) on the source workload to use for replication traffic that are valid for communications with the replication environment.</p> <p>If the network for the NIC you specify is not part of your Azure VPN, ensure that the NIC has a public IP address.</p>
Services to Stop Before Any Replication	(For Windows Workloads) We recommend that all the non-VSS compliant services or antivirus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.
Services to Stop for Cutover with Replication	(For Windows Workloads) Select the Windows services that should be permanently stopped on the source workload for cutover with any replication. The services stopped on the source workload during the replication process are not restored afterwards. This does not apply for Test Cutover.
Daemons to Stop before Any Replication	(For Linux Workloads) Select the Linux services that you want to be temporarily stopped on the source workload before replication. These services will be restored back after replication completes.
Daemons to Stop for Cutover with Replication	(For Linux Workloads) Select the Linux services that should be permanently stopped on the source workload for Cutover with any Replication. The services stopped on the source workload during the replication process are not restored after Cutover. The stopped services are restored after a Test Cutover.
Target Workload Settings	
(These settings are applied during the Run Cutover.)	
Hostname	<p>Do one of the following:</p> <ul style="list-style-type: none"> ♦ To retain the same host name, select No Change. ♦ To change the host name, select Set To and specify the new name. <p>NOTE: An incremental replication is required if you change the host name at cutover.</p>

Setting Name	Description
Domain / Workgroup	<p>(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays:</p> <ul style="list-style-type: none"> ♦ Workgroup: <i>Workgroup_name</i> where <i>Workgroup_name</i> is the workgroup name to which the source belongs. ♦ Domain: <i>Domain_name</i> where <i>Domain_name</i> is the domain name to which the source belongs. <p>NOTE: An incremental replication is required if you change the domain or workgroup name at cutover.</p> <p>Do one of the following depending on where you want the target workload to join:</p> <ul style="list-style-type: none"> ♦ When the source workload belongs to a workgroup: Assume that the source workload belongs to a workgroup named WorkGroup1. <ul style="list-style-type: none"> ♦ For the target workload to join the same workgroup (WorkGroup1), retain the following existing selection: Workgroup: Workgroup1 ♦ For the target workload to join a different workgroup (say WorkGroup2), select Join Workgroup and specify the name as WorkGroup2. ♦ For the target workload to join a domain, select Join Domain and specify the domain name you want the target to join. ♦ When the source workload belongs to a domain: Assume that the source workload belongs to a domain named Domain1. <ul style="list-style-type: none"> ♦ For the target workload to join a workgroup, click Join Workgroup and specify the name of the workgroup you want the target to join. ♦ For the target workload to join the same domain (Domain1) with the domain registration settings preserved, retain the following existing selection: Domain: Domain1 ♦ For the target workload to join the same domain (Domain1) without preserving the domain registration settings, select Join Domain and specify the domain name as Domain1. ♦ For the target workload to join a different domain, select Join Domain and specify the domain name you want the target to join.
Domain Credentials	<p>(For Windows Workloads) If you select Join Domain, specify the domain administrator credentials.</p>

Setting Name	Description
Network Connections	<ol style="list-style-type: none"> For workloads that have more than one NIC, select Include for each NIC to be migrated. Deselect Include to exclude a NIC. <ul style="list-style-type: none"> At least one NIC is required. The number of NICs to migrate cannot exceed the maximum number of NICs supported by the selected cloud instance. For each included NIC, select a network and subnet. If you have multiple NICs, specify one of the included NICs to use as the primary connection. Click Edit for the appropriate NIC, then select Primary Connection. The default Primary Connection is the first NIC in the list. For more information, see Section 9.4, "Azure Networking Guidelines," on page 154. If you do not use an Azure VPN, the primary NIC requires a public IP address that is automatically assigned by a Azure. For the primary NIC, click Edit, then select Create Public IP. For each included NIC: <ol style="list-style-type: none"> Specify a resource group to use for the NIC. The Azure Resource Group setting is the default. To specify a different resource group, click Edit and do one of the following: <ul style="list-style-type: none"> Type the name to use when PlateSpin creates a new resource group. Select an existing resource group from the list. Select one of the following: <ul style="list-style-type: none"> DHCP: Obtain an IP address automatically assigned by a DHCP server. Static: Specify a static IP address, a subnet mask, and a gateway IP address. The IP address must be unique within the supported subnet.
DNS Servers	<p>Specify the DNS Servers for the target workloads. This is applicable only if you select Static in the Network Connections option:</p> <ul style="list-style-type: none"> Primary DNS server: Specify the primary DNS server address. Alternative DNS server: Specify an alternate DNS server address. Additional DNS server: To specify additional DNS server addresses: <ol style="list-style-type: none"> Click Advanced. Specify the DNS server address. Click Add to add the server in the DNS Server Addresses list. Click OK.
Services States on Target VM	(For Windows Workloads) Select Windows services' start conditions on the target VM. Start options are Automatic , Manual , Disabled , and Automatic (Delayed Start) .
Daemons States to Change	(For Linux Workloads) Select Linux daemons' start conditions on the target VM. Enable the daemon to start by selecting the check boxes at the appropriate runlevels (0 to 6) and Boot.

Setting Name	Description
Target Workload Test Settings	
(These settings are applied during the Test Cutover)	
Hostname	<p>Do one of the following:</p> <ul style="list-style-type: none"> ♦ To retain the same host name, select No Change. ♦ To change the host name, select Set To and specify the new name. <p>NOTE: An incremental replication is not required if you change the host name at test cutover.</p>
Domain / Workgroup	<p>(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays:</p> <ul style="list-style-type: none"> ♦ Workgroup: <i>Workgroup_name</i> where <i>Workgroup_name</i> is the workgroup name to which the source belongs. ♦ Domain: <i>Domain_name</i> where <i>Domain_name</i> is the domain name to which the source belongs. <p>NOTE: An incremental replication is not required if you change the domain or workgroup name at test cutover.</p> <p>Do one of the following depending on where you want the target workload to join:</p> <ul style="list-style-type: none"> ♦ When the source workload belongs to a workgroup: Assume that the source workload belongs to a workgroup named WorkGroup1. <ul style="list-style-type: none"> ♦ For the target workload to join the same workgroup (WorkGroup1), retain the following existing selection: Workgroup: Workgroup1 ♦ For the target workload to join a different workgroup (say WorkGroup2), select Join Workgroup and specify the name as WorkGroup2. ♦ For the target workload to join a domain, select Join Domain and specify the domain name you want the target to join. ♦ When the source workload belongs to a domain: Assume that the source workload belongs to a domain named Domain1. <ul style="list-style-type: none"> ♦ For the target workload to join a workgroup, click Join Workgroup and specify the name of the workgroup you want the target to join. ♦ For the target workload to join the same domain (Domain1) with the domain registration settings preserved, retain the following existing selection: Domain: Domain1 ♦ For the target workload to join the same domain (Domain1) without preserving the domain registration settings, select Join Domain and specify the domain name as Domain1. ♦ For the target workload to join a different domain, select Join Domain and specify the domain name you want the target to join.
Domain Credentials	(For Windows Workloads) If you select Join Domain , specify the domain administrator credentials.

Setting Name	Description
Network Connections	<p>Available NICs match the included NICs in Target Workload Settings > Network Connections.</p> <ol style="list-style-type: none"> For each included NIC, select a network and subnet. If you have multiple NICs, specify one of the included NICs to use as the primary connection. Click Edit for the appropriate NIC, then select Primary Connection. <p>The default Primary Connection is the first NIC in the list. For more information, see Section 9.4, "Azure Networking Guidelines," on page 154.</p> <ol style="list-style-type: none"> If you do not use an Azure VPN, the primary NIC requires a public IP address that is automatically assigned by a Azure. For the primary NIC, click Edit, then select Create Public IP. For each included NIC: <ol style="list-style-type: none"> Specify a resource group to use for the NIC. The Azure Resource Group setting is the default. To specify a different resource group, click Edit and do one of the following: <ul style="list-style-type: none"> Type the name to use when PlateSpin creates a new resource group. Select an existing resource group from the list. Select one of the following: <ul style="list-style-type: none"> DHCP: Obtain an IP address automatically assigned by a DHCP server. Static: Specify a static IP address, a subnet mask, and a gateway IP address. The IP address must be unique within the supported subnet.
DNS Servers	<p>Specify the DNS Servers for the target workloads. This is applicable only if you select Static in the Network Connections option:</p> <ul style="list-style-type: none"> Primary DNS server: Specify the primary DNS server address. Alternative DNS server: Specify an alternate DNS server address. Additional DNS server: To specify additional DNS server addresses: <ol style="list-style-type: none"> Click Advanced. Specify the DNS server address. Click Add to add the server in the DNS Server Addresses list. Click OK.
Services States on Target VM	(For Windows Workloads) Select Windows services that must be automatically stopped on the target VM.
Daemons States to Change	(For Linux Workloads) Select Linux daemons that must be automatically stopped on the target VM.
Tag	
Tag	Select a tag to assign to the workload. For more information about tags, see Section 20.9, "Using Tags to Track Logical Associations of Workloads," on page 218 .

9 (Optional) To change the target, click **Change Target**.

NOTE: If you change the target, all the settings you specified will be cleared.

10 Do one of the following:

- ♦ Click **Save** to save the settings.
- ♦ Click **Save and Prepare** to save the settings and start preparing the workload migration.
- ♦ Click **Cancel** to exit.

29 Migration to VMware vCloud Director

- ♦ [Section 29.1, “Planning for Migration to VMware vCloud Director,” on page 349](#)
- ♦ [Section 29.2, “Configuring Migration of a Workload to VMware vCloud Director,” on page 349](#)

29.1 Planning for Migration to VMware vCloud Director

Before you begin migrations to your cloud environment in VMware vCloud Director, ensure that your migration environment meets the following guidelines:

Supported Cloud Platforms

- ♦ See [“Supported Target Cloud Platforms” on page 40](#).

Supported Workloads

- ♦ See [“Supported Workloads For Migration to VMware vCloud Director” on page 29](#), as appropriate for the target Hyper-V platform.

Network Access and Communications

- ♦ See [“Access and Communication Requirements across your Migration Network” on page 49](#).

Prerequisites

- ♦ See [“Prerequisites for Migration to VMware vCloud Director” on page 163](#).

Targets and Workloads

- ♦ **Target VMware vCloud Organization (automated):** See [“Target Discovery in the Web Interface” on page 202](#).
- ♦ **Source Workloads:** Use either of the following discovery methods:
 - ♦ [“Workload Discovery in the Migrate Web Interface” on page 215](#)
 - ♦ [“Registering Workloads and Discovering Details with Migrate Agent” on page 215](#)

Additional Information

- ♦ [“Working with Virtual Machines”](#) in the *VMware vCloud Director 5.6 Documentation Center*.

29.2 Configuring Migration of a Workload to VMware vCloud Director

When you add or discover a workload, the workload is listed on the Workloads page and the status is set as **Not Configured**. Before you migrate the workload, you must configure the workload for migration:

- 1 Launch the PlateSpin Migrate Web Interface.
- 2 On the Workloads page, select the workload you want to configure.

3 Click **Configure Migration**.

4 Select one of the following based on the scope of data you want to transfer from the source to the target:

- ♦ **Full Replication:** A full volume of data transfer takes place from the source to the target.
- ♦ **Incremental Replication:** Only differences are transferred from the source to the target, provided they have similar operating system and volume profiles.

NOTE: PlateSpin Migrate does not support Incremental Replication for the initial replication of data to existing target workloads in VMware vCloud Director. However, you can schedule Incremental Replications for subsequent replication of data. See [Incremental Recurrence](#) in [Step 8](#).

5 Select a VMware vCloud Organization, which you previously configured as a target, to which you want to migrate the source data. See [Section 3.3.3, “Targets,” on page 74](#).

6 Click **Configure Migration**.

7 Configure the following settings. Ensure that the IP address for the source workload, the replication network for target, the cutover network, and the test cutover network are all different.

Setting Name	Description
Schedule Settings	
Incremental Recurrence	<p>Specify the following:</p> <ul style="list-style-type: none">♦ Start of Recurrence: The date when you want to start the replication. You can specify the date or click the calendar icon to select the date. By default, the time is 12:00 a.m.♦ Recurrence Pattern: The pattern to follow for the recurrence of the replication. For example:<ul style="list-style-type: none">♦ To use incremental recurrence everyday, select Daily.♦ To never use incremental recurrence, select None. <p>NOTE</p> <ul style="list-style-type: none">♦ Scheduled incremental replications are skipped until the first full replication is complete.♦ When you schedule incremental recurrence, the replication takes place for a maximum period of 60 days from the starting time of replication. For example:<ul style="list-style-type: none">♦ If you select Daily, then the replication takes place for 60 days from the time the replication starts.♦ If you select Weekly, then the replication takes place for 8 weeks from the time the replication starts.♦ If you select Monthly, then the replication takes place for 2 months from the time the replication starts.

Setting Name	Description
Full Replication	<p>Do one of the following:</p> <ul style="list-style-type: none"> ♦ To specify a schedule for the replication, click Start and specify the date when you want to start the full replication. ♦ To start full replication manually without setting a schedule, click None. <p>NOTE: You must prepare the workload prior to the scheduled time. The full replication cannot run unless the target VM exists and the workload preparation is complete. Migrate skips the scheduled full replication and retries it at the next scheduled time.</p>
Blackout Window	<p>Use these settings to force a replication blackout. The replication blackout suspends scheduled replications during peak utilization hours or prevents conflicts between VSS-aware software and the PlateSpin VSS block-level data transfer component.</p> <p>To specify a blackout window, click Edit and do the following:</p> <ul style="list-style-type: none"> ♦ Specify the start and end time for the blackout period. ♦ Select one of the blackout recurrence pattern such as daily, weekly, or monthly. If you do not want to force a replication blackout, select None. <p>NOTE: The blackout start and end times are based on the system clock on the PlateSpin Server.</p>
Compression Level	<p>These settings control whether data is compressed during transmission between the source and target workloads, and the level of data compression applied. See Section 2.4.3, "Data Compression," on page 48. Select one of the following options:</p> <ul style="list-style-type: none"> ♦ Fast: Consumes the least CPU resources on the source, but yields a lower compression ratio. ♦ Optimal: Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option. ♦ Maximum: Consumes the most CPU resources on the source, but yields a higher compression ratio.
Bandwidth Throttling	<p>These settings control the bandwidth throttling. PlateSpin Migrate enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. You can specify a throughput rate for each migration job. Throttling provides a way to prevent migration traffic from congesting your production network and to reduce the overall load of your PlateSpin Server.</p> <p>To throttle replications to a specified rate, specify the required throughput value in Mbps and the time pattern.</p>
Migration Settings	
Transfer Method	<p>(For Windows Workloads) Select a data transfer mechanism and security through encryption. See Section 2.2, "Supported Data Transfer Methods," on page 42.</p> <p>To enable encryption, select the Encrypt Data Transfer option. See Section 2.3, "Security and Privacy," on page 43.</p> <p>NOTE: The Offline Transfer with Temporary Boot Environment transfer method is not applicable for the Web interface.</p>

Setting Name	Description
Transfer Encryption	(For Linux Workloads) To enable encryption, select the Encrypt Data Transfer option. See Section 2.3, "Security and Privacy," on page 43 .
Source Credentials	Specify the credentials required for accessing the workload. See Section 20.3, "Discovery Guidelines for Source Workloads," on page 212 .
CPU	<p>(For migration to vCloud and VM containers using VMware 5.1, 5.5, and 6.0 with a minimum VM hardware Level 8) Specify the number of sockets and the number of cores per socket for the target workload. It automatically calculates the total cores. This parameter applies on the initial setup of a workload with an initial replication setting of Full Replication.</p> <p>NOTE: The maximum number of cores the workload can use is subject to external factors such as the guest operating system, the VM hardware version, VMware licensing for the ESXi host, and ESXi host compute maximums for vSphere (see ESXi/ESX Configuration Maximums (VMware KB 1003497) (https://kb.vmware.com/kb/1003497)).</p> <p>Some distributions of a guest OS might not honor the cores and cores per socket configuration. For example, guest OSES using SLES 10 SP4 retain their original cores and sockets settings as installed, whereas other SLES and RHEL distributions honor the configuration.</p>
Organization Virtual Data Center	(For migration to vCloud) Select a virtual data center associated with your organization.
vApp	Specify a name for the VMware vApp.
Virtual Machine Name	Specify a display name for the new virtual machine.
Disks	Specify the path to the hard disk on the target virtual machine.
Volumes	Select volumes to be included in the target for migration.
NTFS Cluster Size	(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see Microsoft Support KB Article 140365 .
Non-volume Storage	(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups	(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.
Volume Groups	(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the Converted Logical Volumes section of the settings.
Converted Logical Volumes	(For Linux Workloads) Specify one or more LVM logical volumes to be migrated for a Linux workload.
Replication Network for Target	Specify a network interface (NIC or IP address) on the target to use for replication traffic.

Setting Name	Description
Replication Networks for Source	<p>Specify one or more network interfaces (NIC or IP address) on the source to use for replication traffic.</p> <ul style="list-style-type: none"> ♦ DHCP: Obtain an IP address automatically assigned by a DHCP server. ♦ Static - Manual: Specify a static IP address. ♦ Static - IP Pool: Select this option to automatically issue IP address from the IP pool. <p>For Windows workloads that have more than one NIC, select the connection for each NIC.</p> <p>For this setting, you can also specify an MTU value that the PlateSpin Migrate Linux RAM Disk (LRD) replication network can use. Setting a low value helps to avoid jabber over networks. For example: a VPN.</p> <p>The default value is an empty string. When networking is configured in the LRD, it allows the network device to set its own default, which is usually 1500. However, if you specify a value, PlateSpin Migrate adjusts the MTU when it configures the network interface.</p>
Services to Stop Before Any Replication	(For Windows Workloads) We recommend that all the non-VSS compliant services or antivirus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.
Services to Stop for Cutover with Replication	(For Windows Workloads) Select the Windows services that should be permanently stopped on the source workload for cutover with any replication. The services stopped on the source workload during the replication process are not restored afterwards. This does not apply for Test Cutover.
Daemons to Stop Before Any Replication	(For Linux Workloads) Select the Linux services that you want to be temporarily stopped on the source workload before replication. These services will be restored back after replication completes.
Daemons to Stop for Cutover with Replication	(For Linux Workloads) Select the Linux services that should be permanently stopped on the source workload for Cutover with any Replication. The services stopped on the source workload during the replication process are not restored after Cutover. The stopped services are restored after a Test Cutover.
Target Workload Settings	
(These settings are applied during the Run Cutover.)	
VM Memory	Specify the amount of memory allocated to the target workload.
VM Tools	To install the VM tools, select the Install VM Tools option . This option is selected by default.
Hostname	<p>Do one of the following:</p> <ul style="list-style-type: none"> ♦ To retain the same host name, select No Change. ♦ To change the host name, select Set To and specify the new name. <p>NOTE: An incremental replication is required if you change the host name at cutover.</p>

Setting Name	Description
System Identifier (SID) - (This setting is applicable only for Windows Server 2008 and Windows Server 2003)	<p>Before you generate a new SID for the Windows Server 2003 target workload computer, you must do the following:</p> <ul style="list-style-type: none"> ◆ Enable the SID generation: <ol style="list-style-type: none"> 1. Open a web browser and go to: <code>https://hostname or IP_address/platespinconfiguration</code> Replace <i>hostname</i> or <i>IP_address</i> with the DNS host name or IP address of your PlateSpin Migrate Server. If SSL is not enabled, use <code>http</code> in the URL. 2. On the PlateSpin Server Configuration page, set alwaysGenerateNewSid to True. ◆ Ensure that the host name of the source and target workloads are different. <p>To generate a new system identifier for the target workload, select Generate New System Identifier (SID) in the Target Workload Test Settings section of the Web Interface. For Windows Server 2008, you must specify the local Administrator account credentials. If this account has been locally renamed on the source, provide the new name.</p>

Setting Name	Description
Domain / Workgroup	<p>(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays:</p> <ul style="list-style-type: none"> ♦ Workgroup: <i>Workgroup_name</i> where <i>Workgroup_name</i> is the workgroup name to which the source belongs. ♦ Domain: <i>Domain_name</i> where <i>Domain_name</i> is the domain name to which the source belongs. <p>NOTE: An incremental replication is required if you change the domain or workgroup at cutover.</p> <p>Do one of the following depending on where you want the target workload to join:</p> <ul style="list-style-type: none"> ♦ When the source workload belongs to a workgroup: Assume that the source workload belongs to a workgroup named WorkGroup1. <ul style="list-style-type: none"> ♦ For the target workload to join the same workgroup (WorkGroup1), retain the following existing selection: Workgroup: Workgroup1 ♦ For the target workload to join a different workgroup (say WorkGroup2), select Join Workgroup and specify the name as WorkGroup2. ♦ For the target workload to join a domain, select Join Domain and specify the domain name you want the target to join. ♦ When the source workload belongs to a domain: Assume that the source workload belongs to a domain named Domain1. <ul style="list-style-type: none"> ♦ For the target workload to join a workgroup, click Join Workgroup and specify the name of the workgroup you want the target to join. ♦ For the target workload to join the same domain (Domain1) with the domain registration settings preserved, retain the following existing selection: Domain: Domain1 ♦ For the target workload to join the same domain (Domain1) without preserving the domain registration settings, select Join Domain and specify the domain name as Domain1. ♦ For the target workload to join a different domain, select Join Domain and specify the domain name you want the target to join.
Domain Credentials	<p>(For Windows Workloads) If you select Join Domain, specify the domain administrator credentials.</p>
Network Connections	<p>Select the local area connection and then select one of the following:</p> <ul style="list-style-type: none"> ♦ DHCP: Obtain an IP address automatically assigned by a DHCP server. ♦ Static: Specify a static IP address. <p>For Windows workloads that have more than one NIC, select the connection for each NIC.</p>

Setting Name	Description
DNS Servers	<p>Specify the DNS Servers for the target workloads. This is applicable only if you select Static in the Network Connections option:</p> <ul style="list-style-type: none"> ♦ Primary DNS server: Specify the primary DNS server address. ♦ Alternative DNS server: Specify an alternate DNS server address. ♦ Additional DNS server: To specify additional DNS server addresses: <ol style="list-style-type: none"> 1. Click Advanced. 2. Specify the DNS server address. 3. Click Add to add the server in the DNS Server Addresses list. 4. Click OK.
Services States on Target VM	(For Windows Workloads) Select Windows services that must be automatically stopped on the target VM.
Daemons States on Target VM	(For Linux Workloads) Select Linux daemons that must be automatically stopped on the target VM.
Target Workload Test Settings (These settings are applied during the Test Cutover)	
VM Memory	Specify the amount of memory allocated to the target workload.
VM Tools	To install the VM tools, select the Install VM Tools option . This option is selected by default.
Hostname	<p>Do one of the following:</p> <ul style="list-style-type: none"> ♦ To retain the same host name, select No Change. ♦ To change the host name, select Set To and specify the new name. <p>NOTE: An incremental replication is not required if you change the host name at test cutover.</p>
System Identifier (SID) - (This Setting is applicable only for Windows Server 2008 and Windows Server 2003)	<p>Before you generate a new SID for the Windows Server 2003 target workload computer, you must do the following:</p> <ul style="list-style-type: none"> ♦ Enable the SID generation: <ol style="list-style-type: none"> 1. Open a web browser and go to: <code>https://hostname or IP_address/platespinconfiguration</code> Replace <i>hostname</i> or <i>IP_address</i> with the DNS host name or IP address of your PlateSpin Migrate Server. If SSL is not enabled, use <code>http</code> in the URL. 2. On the PlateSpin Server Configuration page, set alwaysGenerateNewSid to True. ♦ Ensure that the host name of the source and target workloads are different. <p>To generate a new system identifier for the target workload, select Generate New System Identifier (SID) in the Target Workload Test Settings section of the Web Interface. For Windows Server 2008, you must specify the local Administrator account credentials. If this account has been locally renamed on the source, provide the new name.</p>

Setting Name	Description
Domain / Workgroup	<p>(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays:</p> <ul style="list-style-type: none"> ♦ Workgroup: <i>Workgroup_name</i> where <i>Workgroup_name</i> is the workgroup name to which the source belongs. ♦ Domain: <i>Domain_name</i> where <i>Domain_name</i> is the domain name to which the source belongs. <p>NOTE: An incremental replication is not required if you change the domain or workgroup at test cutover.</p> <p>Do one of the following depending on where you want the target workload to join:</p> <ul style="list-style-type: none"> ♦ When the source workload belongs to a workgroup: Assume that the source workload belongs to a workgroup named WorkGroup1. <ul style="list-style-type: none"> ♦ For the target workload to join the same workgroup (WorkGroup1), retain the following existing selection: Workgroup: Workgroup1 ♦ For the target workload to join a different workgroup (say WorkGroup2), select Join Workgroup and specify the name as WorkGroup2. ♦ For the target workload to join a domain, select Join Domain and specify the domain name you want the target to join. ♦ When the source workload belongs to a domain: Assume that the source workload belongs to a domain named Domain1. <ul style="list-style-type: none"> ♦ For the target workload to join a workgroup, click Join Workgroup and specify the name of the workgroup you want the target to join. ♦ For the target workload to join the same domain (Domain1) with the domain registration settings preserved, retain the following existing selection: Domain: Domain1 ♦ For the target workload to join the same domain (Domain1) without preserving the domain registration settings, select Join Domain and specify the domain name as Domain1. ♦ For the target workload to join a different domain, select Join Domain and specify the domain name you want the target to join.
Domain Credentials	<p>(For Windows Workloads) If you select Join Domain, specify the domain administrator credentials.</p>
Network Connections	<p>Select the network connection and then select one of the following:</p> <ul style="list-style-type: none"> ♦ DHCP: Obtain an IP address automatically assigned by a DHCP server. ♦ Static: Specify a static IP address.

Setting Name	Description
DNS Servers	<p>Specify the DNS Servers for the target workloads. This is applicable only if you select Static in the Network Connections option:</p> <ul style="list-style-type: none"> ♦ Primary DNS server: Specify the primary DNS server address. ♦ Alternative DNS server: Specify an alternate DNS server address. ♦ Additional DNS server: To specify additional DNS server addresses: <ol style="list-style-type: none"> 1. Click Advanced. 2. Specify the DNS server address. 3. Click Add to add the server in the DNS Server Addresses list. 4. Click OK.
Services States on Target VM	(For Windows Workloads) Select Windows services that must be automatically stopped on the target VM.
Daemons States to Change	(For Linux Workloads) Select Linux daemons that must be automatically stopped on the target VM.
Tag	
Tag	Select a tag to assign to the workload. For more information about tags, see Section 20.9, "Using Tags to Track Logical Associations of Workloads," on page 218 .

8 (Optional) To change the target, click **Change Target**.

NOTE: If you change the target, all the settings you specified will be cleared.

9 Do one of the following:

- ♦ Click **Save** to save the settings.
- ♦ Click **Save and Prepare** to save the settings and start preparing the workload migration.
- ♦ Click **Cancel** to exit.

30 Migration to VMware

For migration of workloads to a VMware virtual host, PlateSpin Migrate provides automated setup of the target virtual machine on a specified ESX host, in accordance with the features and capabilities of the selected virtualization platform. In addition to the migration settings, you specify settings for the target VM that Migrate will create, such as:

- ♦ Target VM name and configuration file path
- ♦ Datastore to use from available resources on the target virtual host
- ♦ Network settings
- ♦ Virtual memory allocation

NOTE: Raw Device Mapping (RDM) for target VMs on VMware is supported only by using the X2P workflow. When you use the X2P workflow for migrating a workload to VMware, you must set up the VMware Tools for the target workload before you perform the conversion. See [Section 30.3.3, “Setting Up VMware Tools for the Target Workload,” on page 372](#).

NOTE: If your target VMware ESX server is part of a fully automated Distributed Resource Scheduler (DRS) cluster (a cluster with its VM migration automation level set to **Fully Automated**), the newly created target VM's automation level is changed to **Partially Automated** for the duration of the migration. This means that your target VM might power up on a different ESX server from the one initially selected, but migration is prevented from automatic execution.

Use the guidelines in this section to configure migration to VMware.

- ♦ [Section 30.1, “Planning for Migration to VMware,” on page 359](#)
- ♦ [Section 30.2, “Automated Migration to VMware Using Migrate Client,” on page 360](#)
- ♦ [Section 30.3, “Migration to VMs on VMware Using X2P Workflow,” on page 371](#)
- ♦ [Section 30.4, “Automated Migration to VMware Using Migrate Web Interface,” on page 374](#)

30.1 Planning for Migration to VMware

Before you begin migrations to virtual machines on VMware, ensure that your migration environment meets the following guidelines:

Supported VMware Platforms

- ♦ See “VMware vCenter” in [Table 2-14, “Supported Target Virtualization Platforms for the Migrate Client Only,” on page 39](#).

Supported Workloads

- ♦ See [Section 2.1.1, “Supported Source Workloads For Migration to Non-Cloud Platforms,” on page 23](#), as appropriate for the target Hyper-V platform.

Network Access and Communications

- ♦ See [“Access and Communication Requirements across your Migration Network”](#) on page 49.

Prerequisites

- ♦ See [Chapter 11, “Prerequisites for Migration to VMware,”](#) on page 169.

Target Discovery

Using Migrate Client

- ♦ **Target VMware virtual host (automated):** See [Section 19.4.1, “Target Discovery in the Migrate Client,”](#) on page 200

Using Migrate Web Interface

- ♦ **Target VMware virtual host (automated):** See [Section 19.4.2, “Target Discovery in the Web Interface,”](#) on page 202.

Using PlateSpin ISO

- ♦ **Target VMs on VMware virtual host (semi-automated):** See [Section 19.5, “Registering and Discovering Details for Target VMs on Virtual Hosts with PlateSpin ISO,”](#) on page 204.

Workload Discovery

Using Migrate Client

- ♦ **Source Workloads:** See [Section 20.6.1, “Workload Discovery in the Migrate Client,”](#) on page 214.

Using Migrate Web Interface

- ♦ **Source Workloads:** See [Section 20.6.2, “Workload Discovery in the Migrate Web Interface,”](#) on page 215.

Using Migrate Agent

- ♦ **Source Workloads:** See [Section 20.7, “Registering Workloads and Discovering Details with Migrate Agent,”](#) on page 215.

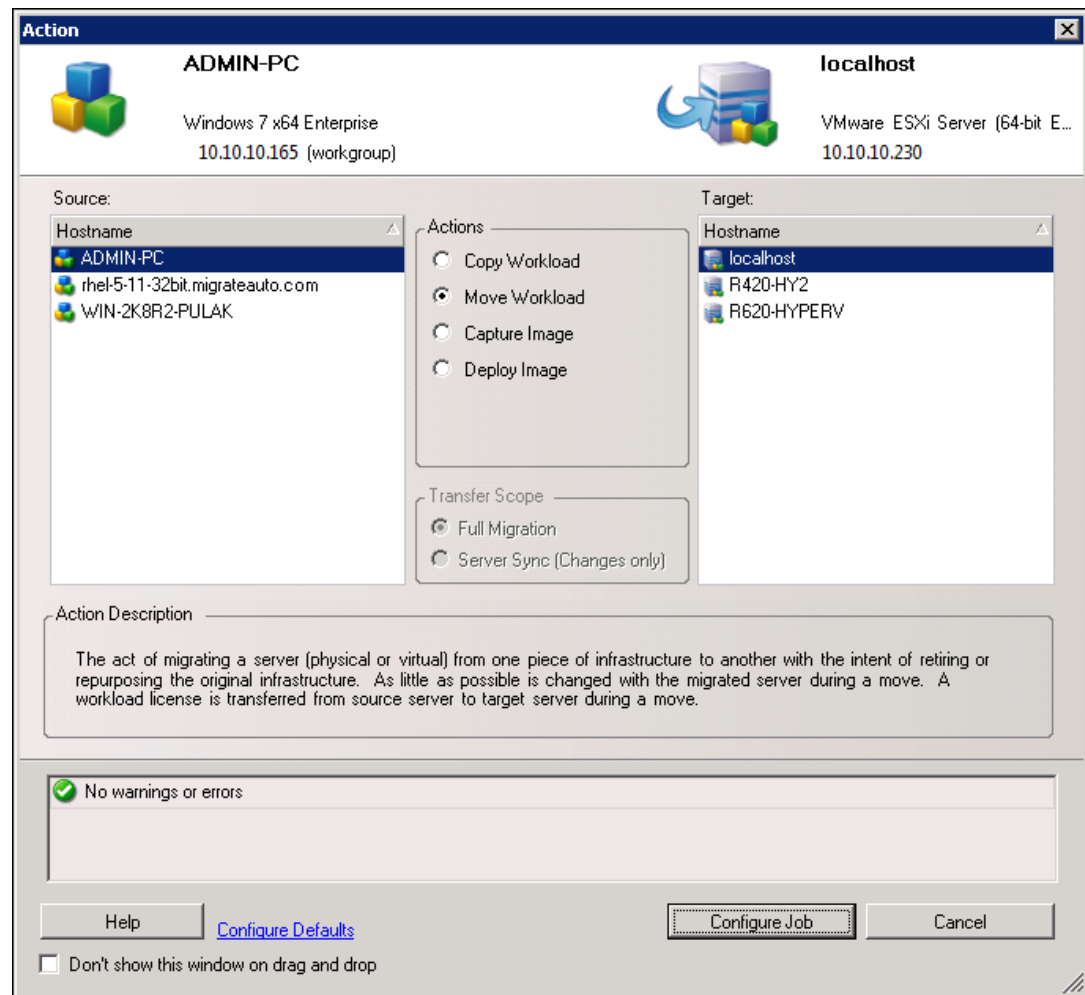
Additional Information

- ♦ *vSphere Virtual Machine Administration* (<https://docs.vmware.com/en/VMware-vSphere/6.5/vsphere-esxi-vcenter-server-65-virtual-machine-admin-guide.pdf>)

30.2 Automated Migration to VMware Using Migrate Client

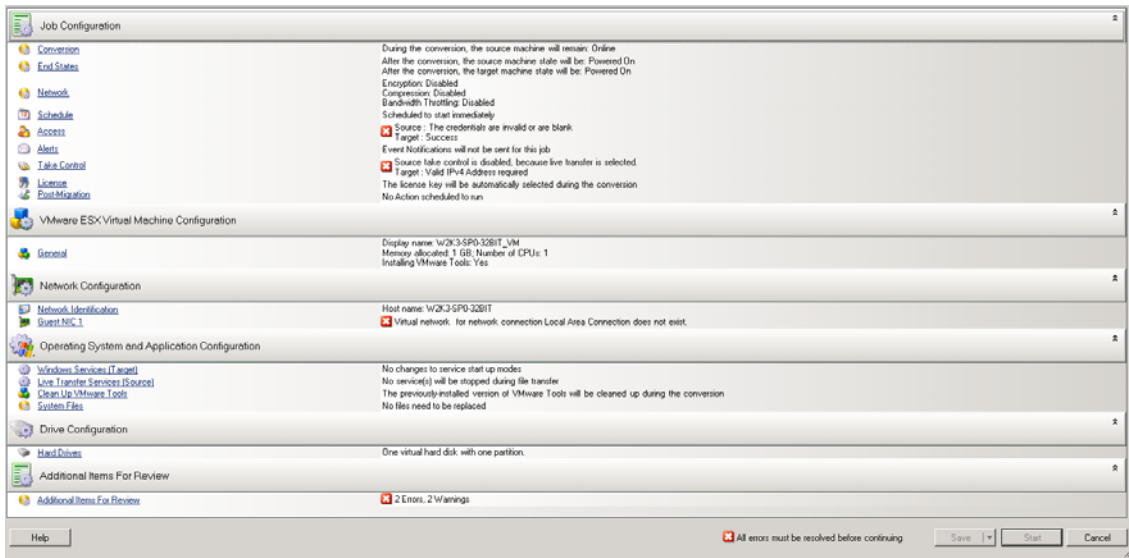
- 1 Discover or refresh your source workload and your target VM host.
See [“Discovering and Preparing Workloads and Targets”](#) on page 195.
- 2 In the Migrate Client, initiate a peer-to-peer workload migration.
 - 2a Expand the Tasks options, then select the conversion type, depending on your goals for the migration:
 - ♦ **Copy Workload**
 - ♦ **Move Workload**

The Source and Target panes display workloads and targets applicable to the selected type of a migration job.



- 2b** In the Source pane, select the workload you want to migrate.
- 2c** In the Target pane, select target host for the migration.
- 2d** Check the validation messages at the bottom of the window.
- 2e** Click **Configure Job** to access the Peer-to-Peer Migration Job window.

Figure 30-1 Peer-to-Peer Migration Job Window



3 In the Job Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
License	
License Key	<p>PlateSpin Migrate automatically selects the best license key for a migration job. If you have multiple license keys, you can specify the license key to use for the workload, assuming licenses are available (neither expired nor exhausted).</p> <p>To specify an alternate key to use:</p> <ol style="list-style-type: none"> 1. Deselect Automatically select the best license key during the conversion, then select the appropriate license key from the menu. 2. Click OK. <p>The selected license key is displayed on the License tab and its description is updated.</p>
Conversion	
Transfer Scope	Specify the scope of workload data to transfer from the source to the target as Full Migration or Server Sync (Changes Only) .
Transfer Method	Specify how data is transferred from source to target. The availability depends on your workload and migration job type. See Section 2.2, "Supported Data Transfer Methods," on page 42.
End State	
Source Machine End State	Specify whether to shut down the source workload after a successful cutover. For a workload move, the shut down is selected by default.
Target Virtual Machine End State	Specify whether to power on, power off, or suspend the target workload after a successful cutover.

Setting Name	Description
Network	
Compression	<p>Specify whether to compress data during transmission between the source and target workloads, and the level of data compression to apply. See Section 2.4.3, “Data Compression,” on page 48. Select one of the following options:</p> <ul style="list-style-type: none"> ♦ Fast: Consumes the least CPU resources on the source, but yields a lower compression ratio. ♦ Optimal: Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option. ♦ Maximum: Consumes the most CPU resources on the source, but yields a higher compression ratio.
Encryption	Select Encrypt Data Transfer to encrypt the data as it is transferred from source to target. See Section 2.3, “Security and Privacy,” on page 43 .
Bandwidth Throttling	<p>Specify whether to throttle bandwidth for data transfer traffic between the source and target machines. To enable throttling, select the Enable Bandwidth Throttling option, specify the required maximum value in Mbps, and optionally a time period during which to enforce the throttling. If specified, the from and to time values are based on the source workload’s system time.</p> <p>If no time interval is defined, bandwidth is throttled to the specified rate at all times by default. If time interval is defined and the migration job executes outside this interval, data is transferred at full speed.</p>
IP Addresses	<p>Specify additional IP addresses for source workloads to enable communication in environments that use network address translation (NAT).</p> <p>For information on how to specify additional IP addresses for your PlateSpin Server, see Section 2.6.4, “Migrations Across Public and Private Networks through NAT,” on page 53.</p>
Schedule	
Schedule	<p>Specify when to start the migration job:</p> <ul style="list-style-type: none"> ♦ Run immediately ♦ Run at a later time <p>Use the calendar menu to specify the date and time to begin the migration.</p> <p>NOTE: You must prepare the workload prior to the scheduled time. The full replication cannot run unless the target VM exists and the workload preparation is complete. Migrate skips the scheduled full replication and retries it at the next scheduled time.</p>
Access Settings	
Source Credentials	<p>(Windows) Specify the account user name with local or domain-level administrative privileges and a valid password. Use this format:</p> <ul style="list-style-type: none"> ♦ For domain member machines: <i>authority\principal</i> ♦ For workgroup member machines: <i>hostname\principal</i> <p>(Linux) Specify the root or root-level user name and a valid password.</p>

Setting Name	Description
Target Credentials	<p>(VMware DRS Cluster) Specify VMware vCenter Web service user name and password.</p> <p>(VMware ESX Server) Specify one of the following:</p> <ul style="list-style-type: none"> ♦ ESX account with administrator role <p>OR</p> <ul style="list-style-type: none"> ♦ Windows domain credentials (versions 4 and 4.1 only)

Alerts

Receive Event Notifications	Specify whether to send email notifications for event conditions. You must configure an SMTP server to use this feature.
Receive Progress Notifications	If you enable Event notifications, you can optionally receive progress notifications at a specified interval.
Send to Addresses	Add or remove valid email addresses for recipients of the notifications.

Take Control Settings

Target Virtual Machine	Under Target Virtual Machine, click Configure, then specify the options for the virtual network and the TCP/IP settings for the replication NIC, then click OK .
------------------------	---

Post-Migration

Action	Specify a pre-configured action from the PlateSpin Migrate library. See Section 6.5, "Managing Post-Migration Actions (Windows and Linux)," on page 121 .
Execution Parameters	Specify the command line command to run the selected action. You can specify a timeout for the execution.
Credentials	Specify the user name and password to use for the post-migration tasks. You can optionally use the source credentials.

- 4** In the Virtual Machine Configuration section of the Migration Job window, click **General**, then configure the following settings:

Setting Name	Description
VMware ESX Virtual Machine	
Virtual Machine Name	Specify a name to use for the target VM as it appears in the VMware.
Datastore	Select a datastore associated with your VM for storing VM configuration files.
Path	<p>Type the path to use for the target VM file, including the VM file name. For example:</p> <p><i>/hostname-VM/hostname-VM.vmx</i></p>
Virtual Machine Memory Allocation	Specify the amount of virtual memory in GB.
Install VMware Tools	Specify whether to install the latest VMware tools on the target VM. If they are installed on the source, they will be uninstalled and reinstalled using the version appropriate for the platform of the VMware target host.
Virtual Devices	Specify preferences for the virtual devices.

Setting Name	Description
Advanced	<p>(For expert users) Specify preferences for the Resource Pool, number of CPUs, and CPU scheduling affinity based on their availability on the target VMware server. Each vCPU is presented to the guest OS on the VM container as a single core, single socket.</p> <p>(For migration to VM container that is part of a DRS Cluster) Specify the Resource Pool location where the migrated VM is to be created.</p>

PlateSpin Migrate displays target virtual machine configuration options specific to the selected target and also provides access to advanced configuration options. For information about host-specific configuration options, see:

- ♦ [Target VM Configuration: VMware ESXi 5 and Later](#)
- ♦ [Target VM Configuration: VMware ESX 4.1](#)

5 In the Network Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Network Configuration	
Network Identification Settings for Windows	
Host Name	Specify the desired host name for the target machine.
Generate New SID	When this option is selected, the target workload is assigned a new System Identifier (SID). Credentials are required only for Windows 2008 systems, and must be the credentials for the local (embedded) Administrator account. If this account has been locally renamed on the source, provide the new name.
Member of Domain / Workgroup	Select the required option and type the name of the domain or workgroup that you want the target machine to join.
Preserve Source Server's Domain Registration	Preserves domain registration and ensures that the source server domain registration remains intact during migration. If you disable this option, the source machine's domain account is transferred to the target machine. The source server still appears to be on the domain, but does not have a valid connection.
Domain Credentials	If the target machine is to be part of a domain, specify valid credentials for a user account with permission to add servers to the domain, such as a member of the Domain Admins group or Enterprise Admins group.
Network Identification Settings for Linux	
Host Name	On the Network Identification tab, specify the desired host name for the target machine.
DNS	Use the Add , Edit , and Remove buttons to manage DNS server entries for the new virtual machine.

- 6 In the Operating System and Applications Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Operating System and Application Configuration	
Windows Services (Target)	<p>Select Windows services' start conditions on the target VM after cutover. Start options are Automatic, Manual, Disabled, and Automatic (Delayed Start).</p> <p>To modify the settings:</p> <ol style="list-style-type: none"> 1. Click the Status column for the service, then select from the Windows start options. 2. When you are done setting services start states, click OK.
Live Transfer Services (Source)	<p>Specify the Windows services to stop on the source workload during live data transfers.</p> <p>We recommend that all the non-VSS compliant services or antivirus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.</p> <p>To modify the settings:</p> <ol style="list-style-type: none"> 1. Select Stopped next to the service to be stopped for live data transfer. 2. When you are done setting services to stop, click OK.
Linux Daemons (Target)	<p>Specify the start states for daemons on the target VM after cutover.</p> <p>To modify the settings:</p> <ol style="list-style-type: none"> 1. Click the Run Level column for the daemon, then select from run levels 0 through 6 and Boot (B), then click OK. 2. When you are done setting daemon start states, click OK.
Live Transfer Daemons (Source)	<p>Specify the daemons to stop on the source workload during live data transfers.</p> <p>To modify the settings:</p> <ol style="list-style-type: none"> 1. Select Stopped next to the daemon to be stopped for live data transfer. 2. When you are done setting daemons to stop, click OK.

- 7 In the Drive Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Drive Configuration	
Hard Drives	Specify drive and volume configurations to be migrated.
Disks	Specify the path to the hard disk on the target virtual machine.
Volumes	Select volumes to be included in the target for migration.
NTFS Cluster Size	(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see the Microsoft Support KB Article 140365 .

Setting Name	Description
Non-volume Storage	(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups	(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.
Volume Groups	(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the Converted Logical Volumes section of the settings.
Converted Logical Volumes	(For Linux Workloads) Specify one or more LVM logical volumes to be migrated for a Linux workload.

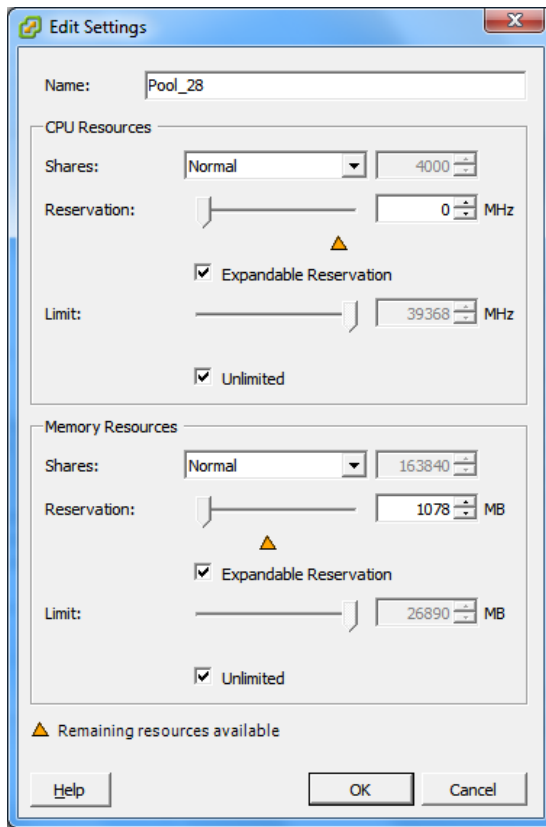
PlateSpin Migrate displays storage configuration options specific to the selected target. For information about host-specific configuration options, see:

- ♦ [Drive Configuration: VMware ESX](#)

- 8 In the Additional Items for Review section of the Migration Job window, review errors and messages about the workload configuration. You must resolve errors before you can submit the migration job.
- 9 Click **OK**.

30.2.1 Target VM Configuration: VMware ESXi 5 and Later

The following are configuration options specific to VMware vSphere 5 and later (applicable to all VMs under the containing resource pool).



Name: Specify the display name for the new virtual machine.


CPU Resources

- ♦ **Shares:** CPU shares for this virtual machine with respect to the parent's total. Peer VMs share resources according to their relative share values bounded by the **Reservation** and **Limit**. Select **Low**, **Normal**, or **High**, which specify share values respectively in a 1:2:4 ratio. Select **Custom** to give each virtual machine a specific number of shares, which express a proportional weight.
- ♦ **Reservation:** Guaranteed CPU allocation for this VM.
Expandable Reservation: Select this option to specify that more than the specified reservation is allocated if resources are available in a parent.
- ♦ **Limit:** Upper limit for this virtual machine's CPU allocation.
Unlimited: Select this option to specify no upper limit.

Memory Resources: (these are similar to CPU resource settings, but apply to memory resources)

30.2.2 Target VM Configuration: VMware ESX 4.1

The following are configuration options specific to VMware ESX systems prior to vSphere 5. To access settings that control resource pools, the number of CPUs, and CPU scheduling affinity, click [Advanced](#).



Virtual Machine Name: Specify the display name for the new virtual machine.

Datastore: Select the datastore where you want to create the *.vmtx file.

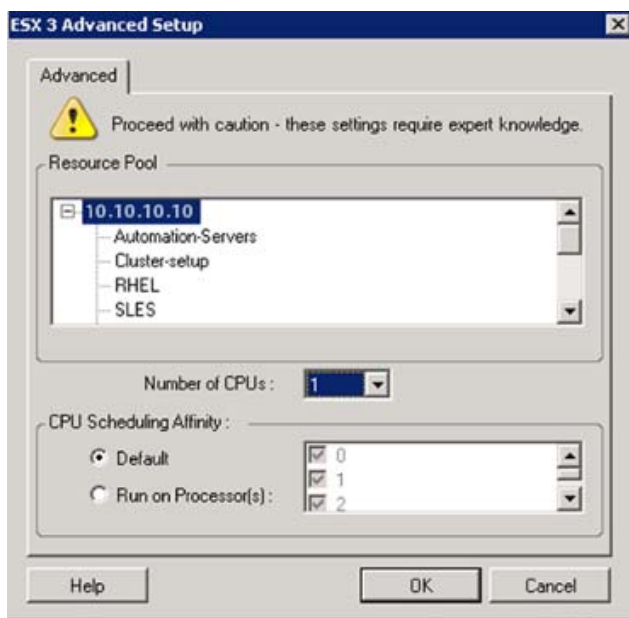
Configuration File Path: Specify a name and the directory path for the virtual machine's *.vmtx configuration file.

Virtual Machine Memory Allocation: Specify a value for the amount of virtual RAM to be assigned to the virtual machine.

Install VMware Tools: Enable this option to install VMware tools during the migration process (recommended).

SCSI Drives: Select either **BusLogic** or **LSILogic** (the recommended option).

Advanced: Click this button to view or modify advanced VM configuration settings.



Resource Pool: If required, assign your target VM to a resource pool. When no resource pool is specified, the VM is assigned to the root resource pool.

Number of CPUs: Select the required number of CPUs to assign to the target VM. For example, you can convert a single-processor workload to a multi-processor VM, or a multi-processor workload to a single-processor VM.

CPU Scheduling Affinity: Represents which ESX Server processors the virtual machine can run on (if your ESX Server is a multiprocessor system). Specify the required processor or select **Default** (recommended).

For details, see your VMware documentation.

30.2.3 Drive Configuration: VMware ESX

The following are drive configuration settings specific to VMware ESX:

Disk	Datastore	Size	File Name
Virtual disk 0	storage1	4 GB	/NY-RHEL4-LVM_VM/NY-RHEL4-LVM_VM_1.vmdk

Include	Volume	Free Space	Size	New Free Space	New Size	Disk/Volume Group
<input checked="" type="checkbox"/>	/	2.4 GB	3.2 GB	2.4 GB	3.2 GB	VolGroup00
<input checked="" type="checkbox"/>	/boot	81.2 MB	98.7 MB	81.2 MB	98.7 MB	Disk 0
<input checked="" type="checkbox"/>	/home	88.1 MB	98.7 MB	88.1 MB	98.7 MB	Disk 0

Include	Type	Partition	Size	Is Swap	Disk/Volu...	New Size
<input checked="" type="checkbox"/>		/dev/VolGroup00/Lo...	512 MB	<input checked="" type="checkbox"/>	Vol...	512 MB

Datastore: Select the datastore volume on the ESX server where you want to place the vmdk files.

Copy: Select the volumes to be copied during the migration.

New Free Space: To resize the volume during the migration, specify the desired amount of free space. PlateSpin Migrate automatically adjusts New Size.

New Size: To resize the volume during the migration, specify the desired size. PlateSpin Migrate automatically adjusts New Free Space.

Disk/Volume Group: Assign the volume to a disk or, if LVM is enabled, to a volume group. The volume will be copied to this disk or volume group on the target machine.

Create: Select any non-volume disk partitions that should be created on the target machine (for example, a Linux swap partition).

New Size: To resize the non-volume partition during the migration, specify the desired size.

30.3 Migration to VMs on VMware Using X2P Workflow

Raw Device Mapping (RDM) for target VMs on VMware is only supported by using the X2P workflow. When you use the X2P workflow for migrating a workload to VMware, you must set up the VMware Tools for the target workload before you perform the conversion.

Use the guidelines in this section to configure migration to VMs on VMware virtual hosts.

- [Section 30.3.1, “Downloading and Saving the PlateSpin ISO Image \(VMware\),” on page 372](#)
- [Section 30.3.2, “Creating and Configuring the Target Virtual Machine \(VMware\),” on page 372](#)

- ♦ [Section 30.3.3, “Setting Up VMware Tools for the Target Workload,” on page 372](#)
- ♦ [Section 30.3.4, “Registering the Virtual Machine with PlateSpin Server \(VMware\),” on page 373](#)
- ♦ [Section 30.3.5, “Migrating Your Source Workload to the Target Virtual Machine \(VMware\),” on page 374](#)

30.3.1 Downloading and Saving the PlateSpin ISO Image (VMware)

- 1 Download and prepare the PlateSpin ISO image for use with the target VM. Attended and unattended registration options are possible.

See [“Preparing the PlateSpin ISO Image for Target Registration and Discovery” on page 280](#).

- 2 Save the ISO image in a location that VMware server can access. For example: `c:\temp`.

This ensures that the PlateSpin ISO image is available to the target VM as a bootable CD-ROM image.

30.3.2 Creating and Configuring the Target Virtual Machine (VMware)

- 1 Log on to the VMware server using the vSphere client and then use the New Virtual Machine Wizard to create a new virtual machine with the following settings:
 - ♦ **Name and Location:** Specify a name for your new target and accept the default location.
 - ♦ **Operating System Type and Version:** Specify the operating system type and version settings that matches the source workload. The wizard uses this information to set appropriate default values, such as the amount of memory needed, and resource limits for the VM.
 - ♦ **Assign Memory:** Assign at least 384 MB of RAM to the VM.
 - ♦ **Connect Virtual Hard Disk:** Ensure that the disk size of every disk is about 50 MB more than the corresponding disk on your source workload.
 - ♦ **Installation Options:** Configure the VM to boot from an ISO image file, and point the wizard to the downloaded PlateSpin ISO image.
 - ♦ **Summary:** Configure the VM to not start upon creation (deselect the **Start the virtual machine after it is created** option).
- 2 Set up VMware tool for the target workload. See [Section 30.3.3, “Setting Up VMware Tools for the Target Workload,” on page 372](#).

30.3.3 Setting Up VMware Tools for the Target Workload

VMware Tools setup packages are automatically copied to the target during conversion so that the configuration service can install the tools on the target VM when the target VM contacts the PlateSpin Server. However, if you choose to migrate workloads to VMware using the X2P workflow, you must set up the VMware tools for the target workload before you perform the conversion. Perform the following steps to prepare your environment for setting the VMware tools for the target workload:

- 1 Retrieve the VMware Tools packages from an ESX host:
 - 1a Secure copy (`scp`) the `windows.iso` image from the `/usr/lib/vmware/isoimages` directory on an accessible ESX host to a local temporary folder.
 - 1b Open the ISO and extract its setup packages, saving them to an accessible location:
 - ♦ **VMware 5.x and later:** The setup packages are `setup.exe` and `setup64.exe`.

- ♦ **VMware 4.x:** The setup packages are `VMware Tools.msi` and `VMware Tools64.msi`.
- 2 Create OFX packages from the setup packages you extracted:
 - 2a Zip the package you want, making sure that the setup installer file is at the root of the `.zip` archive.
 - 2b Rename the `.zip` archive to `1.package` so that it can be used as an OFX package.

NOTE: If you want to create an OFX package for more than one of the setup packages, remember that each setup package must have its own unique `.zip` archive.

Because each package must have the same name (`1.package`), if you want to save multiple `.zip` archives as OFX packages, you need to save each in its own unique subdirectory.

- 3 Copy the appropriate OFX package (`1.package`) to the `%ProgramFiles%\PlateSpin Migrate Server\Packages\%GUID%` directory on the PlateSpin Server.

The value of `%GUID%` depends on the version of your VMware ESX host and its VMware Tools architecture, as shown in [Table 30-1](#). Use the appropriate GUID value to copy the package to the correct directory.

Table 30-1 GUIDs for the VMware Tools Directory Names

VMware Server Version	VMware Tools Architecture	GUID
4.0	x86	D052CBAC-0A98-4880-8BCC-FE0608F0930F
4.0	x64	80B50267-B30C-4001-ABDF-EA288D1FD09C
4.1	x86	F2957064-65D7-4bda-A52B-3F5859624602
4.1	x64	80B1C53C-6B43-4843-9D63-E9911E9A15D5
5.0	x86	AD4FDE1D-DE86-4d05-B147-071F4E1D0326
5.0	x64	F7C9BC91-7733-4790-B7AF-62E074B73882
5.1	x86	34DD2CBE-183E-492f-9B36-7A8326080755
5.1	x64	AD4FDE1D-DE86-4d05-B147-071F4E1D0326
5.5	x86	660C345A-7A91-458b-BC47-6A3914723EF7
5.5	x64	8546D4EF-8CA5-4a51-A3A3-6240171BE278
6.0	x86	311E672E-05BA-4CAF-A948-B26DF0C6C5A6
6.0	x64	D7F55AED-DA64-423F-BBBE-F1215529AD03
6.5	x86	D61C0FCA-058B-42C3-9F02-898F568A3071
6.5	x64	5D3947B7-BE73-4A00-A549-B15E84B98803

30.3.4 Registering the Virtual Machine with PlateSpin Server (VMware)

After you create the virtual machine and prepare it to boot with the PlateSpin ISO, you are ready to register it as a target VM with your PlateSpin Server. See [“Registering and Discovering Target VMs on Virtual Hosts” on page 205](#).

30.3.5 Migrating Your Source Workload to the Target Virtual Machine (VMware)

- 1 Use PlateSpin Migrate Client to start an X2P migration job with your source workload being the job's migration source and the target being the new VM on VMware.
See [“Migration to Physical Machines” on page 407](#).
- 2 For host-specific target VM configuration options for the Virtual Machine Configuration dialog, see:
 - ♦ [“Target VM Configuration: VMware ESXi 5 and Later” on page 368](#)
 - ♦ [“Target VM Configuration: VMware ESX 4.1” on page 369](#)
- 3 For host-specific storage configuration options, see [Section 30.2.3, “Drive Configuration: VMware ESX,” on page 371](#).
- 4 Monitor the migration job in Jobs view in PlateSpin Migrate Client.
When the job reaches the **Configure Target Machine** step, the virtual machine's console returns to the boot prompt of the PlateSpin ISO image.
- 5 Shut down the virtual machine and reconfigure it to boot from disk rather than from the boot image.
- 6 Power on the virtual machine.
The migration job resumes, reboots the target, and completes the workload configuration.

30.4 Automated Migration to VMware Using Migrate Web Interface

- 1 Launch the PlateSpin Migrate Web Interface.
- 2 On the Workloads page, select the workload you want to configure.
- 3 Click **Configure Migration**.
- 4 Select one of the following based on the scope of data you want to transfer from the source to the target:
 - ♦ **Full Replication:** A full volume of data transfer takes place from the source to the target.
 - ♦ **Incremental Replication:** Only differences are transferred from the source to the target, provided they have similar operating system and volume profiles.
- 5 Select the VM host, which you previously configured as a target, to which you want to migrate the source data. Select.
If the target you need is not yet configured, click **Add Targets**, configure the target, then try again to configure the workload. See [Chapter 19, “Discovering Target Hosts,” on page 197](#).
- 6 Click **Configure Migration**.

7 Configure the following settings:

Setting Name	Description
Schedule Settings	
Incremental Recurrence	<p>Specify the following:</p> <ul style="list-style-type: none"> ♦ Start of Recurrence: The date when you want to start the replication. You can specify the date or click the calendar icon to select the date. By default, the time is 12:00 a.m. ♦ Recurrence Pattern: The pattern to follow for the recurrence of the replication. For example: <ul style="list-style-type: none"> ♦ To use incremental recurrence everyday, select Daily. ♦ To never use incremental recurrence, select None. <p>NOTE</p> <ul style="list-style-type: none"> ♦ Scheduled incremental replications are skipped until the first full replication is complete. ♦ When you schedule incremental recurrence, the replication takes place for a maximum period of 60 days from the starting time of replication. For example: <ul style="list-style-type: none"> ♦ If you select Daily, then the replication takes place for 60 days from the time the replication starts. ♦ If you select Weekly, then the replication takes place for 8 weeks from the time the replication starts. ♦ If you select Monthly, then the replication takes place for 2 months from the time the replication starts.
Full Replication	<p>Do one of the following:</p> <ul style="list-style-type: none"> ♦ To specify a schedule for the replication, click Start and specify the date when you want to start the full replication. ♦ To start full replication manually without setting a schedule, click None. <p>NOTE: You must prepare the workload prior to the scheduled time. The full replication cannot run unless the target VM exists and the workload preparation is complete. Migrate skips the scheduled full replication and retries it at the next scheduled time.</p>
Blackout Window	<p>Use these settings to force a replication blackout. The replication blackout suspends scheduled replications during peak utilization hours or prevents conflicts between VSS-aware software and the PlateSpin VSS block-level data transfer component.</p> <p>To specify a blackout window, click Edit and do the following:</p> <ul style="list-style-type: none"> ♦ Specify the start and end time for the blackout period. ♦ Select one of the blackout recurrence pattern such as daily, weekly, or monthly. If you do not want to force a replication blackout, select None. <p>NOTE: The blackout start and end times are based on the system clock on the PlateSpin Server.</p>

Setting Name	Description
Compression Level	<p>These settings control whether data is compressed during transmission between the source and target workloads, and the level of data compression applied. See Section 2.4.3, “Data Compression,” on page 48. Select one of the following options:</p> <ul style="list-style-type: none"> ♦ Fast: Consumes the least CPU resources on the source, but yields a lower compression ratio. ♦ Optimal: Consumes optimal CPU resources on the source and yields an optimal compression ratio. This is the recommended option. ♦ Maximum: Consumes the most CPU resources on the source, but yields a higher compression ratio.
Bandwidth Throttling	<p>These settings control the bandwidth throttling. PlateSpin Migrate enables you to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. You can specify a throughput rate for each migration job. Throttling provides a way to prevent migration traffic from congesting your production network and to reduce the overall load of your PlateSpin Server.</p> <p>To throttle replications to a specified rate, specify the required throughput value in Mbps and the time pattern.</p>
Migration Settings	
Transfer Method	<p>(For Windows Workloads) Select a data transfer mechanism and security through encryption. See Section 2.2, “Supported Data Transfer Methods,” on page 42.</p> <p>To enable encryption, select the Encrypt Data Transfer option. See Section 2.3, “Security and Privacy,” on page 43.</p> <p>NOTE: The Offline Transfer with Temporary Boot Environment transfer method is not applicable for the Web interface.</p>
Transfer Encryption	<p>(For Linux Workloads) To enable encryption, select the Encrypt Data Transfer option. See Section 2.3, “Security and Privacy,” on page 43.</p>
Source Credentials	<p>Specify the credentials required for accessing the workload. See Section 20.3, “Discovery Guidelines for Source Workloads,” on page 212.</p>
CPU	<p>(For migration to vCloud and VM containers using supported versions of VMware 5.1 and later with a minimum VM hardware Level 8) Specify the number of sockets and the number of cores per socket for the target workload. It automatically calculates the total cores. This parameter applies on the initial setup of a workload with an initial replication setting of Full Replication.</p> <p>NOTE: The maximum number of cores the workload can use is subject to external factors such as the guest operating system, the VM hardware version, VMware licensing for the ESXi host, and ESXi host compute maximums for vSphere (see ESXi/ESX Configuration Maximums (VMware KB 1003497) (https://kb.vmware.com/kb/1003497)).</p> <p>Some distributions of a guest OS might not honor the cores and cores per socket configuration. For example, guest OSes using SLES 10 SP4 retain their original cores and sockets settings as installed, whereas other SLES and RHEL distributions honor the configuration.</p>

Setting Name	Description
Number of CPUs	(For migration to VM containers using VMware 4.1) Specify the required number of vCPUs (virtual CPUs) to assign to the target workload. This parameter applies on the initial setup of a workload with an initial replication setting of Full Replication. Each vCPU is presented to the guest OS on the VM container as a single core, single socket.
Resource Pool for Target VM	(For migration to VM container that is part of a DRS Cluster) Specify the Resource Pool location where the migrated VM is to be created.
VM Folder for Target VM	(For migration to VM container that is part of a DRS Cluster) Specify the VM folder location where the migrated VM is to be created.
Virtual Machine Name	Specify a display name for the new virtual machine.
Configuration File Datastore	Select a datastore associated with your VM for storing VM configuration files.
Virtual Machine Configuration Path	Specify the path to the configuration file on the target virtual machine.
Disks	Specify the path to the hard disk on the target virtual machine.
Volumes	Select volumes to be included in the target for migration.
NTFS Cluster Size	(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see the Microsoft Support KB Article 140365 .
Non-volume Storage	(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups	(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.
Volume Groups	(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the Converted Logical Volumes section of the settings.
Converted Logical Volumes	(For Linux Workloads) Specify one or more LVM logical volumes to be migrated for a Linux workload.
Replication Network for Target	Specify a network interface (NIC or IP address) on the target to use for replication traffic.
Replication Networks for Source	Specify one or more network interfaces (NIC or IP address) on the source to use for replication traffic.
Services to Stop Before Any Replication	(For Windows Workloads) We recommend that all the non-VSS compliant services or anti-virus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.
Services to Stop for Cutover with Replication	(For Windows Workloads) Select the Windows services that should be permanently stopped on the source workload for cutover with any replication. The services stopped on the source workload during the replication process are not restored afterwards. This does not apply for Test Cutover.

Setting Name	Description
Daemons to Stop Before Any Replication	(For Linux Workloads) Select the Linux daemons that you want to be temporarily stopped on the source workload before replication. These daemons will be restored after replication completes.
Daemons to Stop for Cutover with Replication	(For Linux Workloads) Select the Linux daemons that should be permanently stopped on the source workload for Cutover with any Replication. The daemons stopped on the source workload during the replication process are not restored after Cutover. The stopped daemons are restored after a Test Cutover.
Target Workload Settings (These settings are applied during the Run Cutover.)	
VM Memory	Specify the amount of memory allocated to the target workload.
VM Tools	To install the VM tools, select the Install VM Tools option . This option is selected by default.
Hostname	Do one of the following: <ul style="list-style-type: none"> ♦ To retain the same host name, select No Change. ♦ To change the host name, select Set To and specify the new name.
System Identifier (SID) - (This setting is applicable only for Windows Server 2008 and Windows Server 2003)	<p>Before you generate a new SID for the Windows Server 2003 target workload computer, you must do the following:</p> <ul style="list-style-type: none"> ♦ Enable the SID generation: <ol style="list-style-type: none"> 1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at: <code>https://Your_PlateSpin_Server/PlateSpinConfiguration/</code> Replace <i>Your_PlateSpin_Server</i> with the DNS host name or IP address of your PlateSpin Migrate Server. If SSL is not enabled, use <code>http</code> in the URL. 2. On the PlateSpin Server Configuration page, set alwaysGenerateNewSid to True. ♦ Ensure that the host name of the source and target workloads are different. <p>To generate a new system identifier for the target workload, select Generate New System Identifier (SID) in the Target Workload Test Settings section of the Web Interface. For Windows Server 2008, you must specify the local Administrator account credentials. If this account has been locally renamed on the source, provide the new name.</p>

Setting Name	Description
Domain / Workgroup	<p>(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays:</p> <ul style="list-style-type: none"> ♦ Workgroup: <i>Workgroup_name</i> where <i>Workgroup_name</i> is the workgroup name to which the source belongs. ♦ Domain: <i>Domain_name</i> where <i>Domain_name</i> is the domain name to which the source belongs. <p>Do one of the following depending on where you want the target workload to join:</p> <ul style="list-style-type: none"> ♦ When the source workload belongs to a workgroup: Assume that the source workload belongs to a workgroup named WorkGroup1. <ul style="list-style-type: none"> ♦ For the target workload to join the same workgroup (WorkGroup1), retain the following existing selection: Workgroup: Workgroup1 ♦ For the target workload to join a different workgroup (say WorkGroup2), select Join Workgroup and specify the name as WorkGroup2. ♦ For the target workload to join a domain, select Join Domain and specify the domain name you want the target to join. ♦ When the source workload belongs to a domain: Assume that the source workload belongs to a domain named Domain1. <ul style="list-style-type: none"> ♦ For the target workload to join a workgroup, click Join Workgroup and specify the name of the workgroup you want the target to join. ♦ For the target workload to join the same domain (Domain1) with the domain registration settings preserved, retain the following existing selection: Domain: Domain1 ♦ For the target workload to join the same domain (Domain1) without preserving the domain registration settings, select Join Domain and specify the domain name as Domain1. ♦ For the target workload to join a different domain, select Join Domain and specify the domain name you want the target to join.
Domain Credentials	<p>(For Windows Workloads) If you select Join Domain, specify the domain administrator credentials.</p>
Network Connections	<p>Select the local area connection and then select one of the following:</p> <ul style="list-style-type: none"> ♦ DHCP: Obtain an IP address automatically assigned by a DHCP server. ♦ Static: Specify a static IP address. <p>For Windows workloads that have more than one NIC, select the connection for each NIC.</p>

Setting Name	Description
DNS Servers	<p>Specify the DNS Servers for the target workloads. This is applicable only if you select Static in the Network Connections option:</p> <ul style="list-style-type: none"> ♦ Primary DNS server: Specify the primary DNS server address. ♦ Alternative DNS server: Specify an alternate DNS server address. ♦ Additional DNS server: To specify additional DNS server addresses: <ol style="list-style-type: none"> 1. Click Advanced. 2. Specify the DNS server address. 3. Click Add to add the server in the DNS Server Addresses list. 4. Click OK.
Services States on Target VM	(For Windows Workloads) Select Windows services' start conditions on the target VM. Start options are Automatic, Manual, Disabled, and Automatic (Delayed Start).
Daemons States to Change	(For Linux Workloads) Select Linux daemons' start conditions on the target VM. Enable the daemon to start by selecting the check boxes at the appropriate runlevels (0 to 6) and Boot.

Target Workload Test Settings

(These settings are applied during the Test Cutover)

VM Memory	Specify the amount of memory allocated to the target workload.
VM Tools	To install the VM tools, select the Install VM Tools option . This option is selected by default.
Hostname	<p>Do one of the following:</p> <ul style="list-style-type: none"> ♦ To retain the same host name, select No Change. ♦ To change the host name, select Set To and specify the new name.
System Identifier (SID) - (This Setting is applicable only for Windows Server 2008 and Windows Server 2003)	<p>Before you generate a new SID for the Windows Server 2003 target workload computer, you must do the following:</p> <ul style="list-style-type: none"> ♦ Enable the SID generation: <ol style="list-style-type: none"> 1. Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at: https://Your_PlateSpin_Server/PlateSpinConfiguration/ Replace <i>Your_PlateSpin_Server</i> with the DNS host name or IP address of your PlateSpin Migrate Server. If SSL is not enabled, use <code>http</code> in the URL. 2. On the PlateSpin Server Configuration page, set alwaysGenerateNewSid to True. ♦ Ensure that the host name of the source and target workloads are different. <p>To generate a new system identifier for the target workload, select Generate New System Identifier (SID) in the Target Workload Test Settings section of the Web Interface. For Windows Server 2008, you must specify the local Administrator account credentials. If this account has been locally renamed on the source, provide the new name.</p>

Setting Name	Description
Domain / Workgroup	<p>(For Windows Workloads) Depending on whether the source workload belongs to workgroup or domain, one of the following displays:</p> <ul style="list-style-type: none"> ♦ Workgroup: <i>Workgroup_name</i> where <i>Workgroup_name</i> is the workgroup name to which the source belongs. ♦ Domain: <i>Domain_name</i> where <i>Domain_name</i> is the domain name to which the source belongs. <p>Do one of the following depending on where you want the target workload to join:</p> <ul style="list-style-type: none"> ♦ When the source workload belongs to a workgroup: Assume that the source workload belongs to a workgroup named WorkGroup1. <ul style="list-style-type: none"> ♦ For the target workload to join the same workgroup (WorkGroup1), retain the following existing selection: Workgroup: Workgroup1 ♦ For the target workload to join a different workgroup (say WorkGroup2), select Join Workgroup and specify the name as WorkGroup2. ♦ For the target workload to join a domain, select Join Domain and specify the domain name you want the target to join. ♦ When the source workload belongs to a domain: Assume that the source workload belongs to a domain named Domain1. <ul style="list-style-type: none"> ♦ For the target workload to join a workgroup, click Join Workgroup and specify the name of the workgroup you want the target to join. ♦ For the target workload to join the same domain (Domain1) with the domain registration settings preserved, retain the following existing selection: Domain: Domain1 ♦ For the target workload to join the same domain (Domain1) without preserving the domain registration settings, select Join Domain and specify the domain name as Domain1. ♦ For the target workload to join a different domain, select Join Domain and specify the domain name you want the target to join.
Domain Credentials	<p>(For Windows Workloads) If you select Join Domain, specify the domain administrator credentials.</p>
Network Connections	<p>Select the network connection and then select one of the following:</p> <ul style="list-style-type: none"> ♦ DHCP: Obtain an IP address automatically assigned by a DHCP server. ♦ Static: Specify a static IP address.

Setting Name	Description
DNS Servers	<p>Specify the DNS Servers for the target workloads. This is applicable only if you select Static in the Network Connections option:</p> <ul style="list-style-type: none"> ♦ Primary DNS server: Specify the primary DNS server address. ♦ Alternative DNS server: Specify an alternate DNS server address. ♦ Additional DNS server: To specify additional DNS server addresses: <ol style="list-style-type: none"> 1. Click Advanced. 2. Specify the DNS server address. 3. Click Add to add the server in the DNS Server Addresses list. 4. Click OK.
Services States on Target VM	(For Windows Workloads) Select Windows services that must be automatically stopped on the target VM.
Daemons States to Change	(For Linux Workloads) Select Linux daemons that must be automatically stopped on the target VM.
Tag	
Tag	Select a tag to assign to the workload. See Section 7.2, “Managing Workload Tags,” on page 127 .

8 (Optional) To change the target, click **Change Target**.

NOTE: If you change the target, all the settings you specified will be cleared.

9 Do one of the following:

- ♦ Click **Save** to save the settings.
- ♦ Click **Save and Prepare** to save the settings and start preparing the workload migration.
- ♦ Click **Cancel** to exit.

31

Migration to Microsoft Hyper-V

For migration of workloads to a Microsoft Hyper-V virtual host, PlateSpin Migrate provides automated setup of the target virtual machine on a specified Hyper-V host, in accordance with the features and capabilities of the selected virtualization platform. In addition to the migration settings, you specify settings for the target VM that Migrate will create, such as:

- ♦ Target VM name and configuration file path
- ♦ Datastore to use from available resources on the target virtual host
- ♦ Network settings
- ♦ Virtual memory allocation

NOTE: For migration to virtual hosts running Windows Server with Hyper-V, semi-automated workload virtualization is available. See [Section 31.3, “Migration to VMs on Hyper-V Using X2P Workflow,” on page 393](#).

Use the guidelines in this section to configure migration to Hyper-V virtual hosts.

- ♦ [Section 31.1, “Planning for Migration to Hyper-V,” on page 383](#)
- ♦ [Section 31.2, “Automated Migration to Hyper-V,” on page 384](#)
- ♦ [Section 31.3, “Migration to VMs on Hyper-V Using X2P Workflow,” on page 393](#)

31.1 Planning for Migration to Hyper-V

Before you begin migrations to virtual machines on Hyper-V virtual hosts, ensure that your migration environment meets the following guidelines:

Supported Hyper-V Platforms

- ♦ See “Microsoft Windows Server with Hyper-V” in [Table 2-14, “Supported Target Virtualization Platforms for the Migrate Client Only,” on page 39](#).

Supported Workloads

- ♦ See [Section 2.1.1, “Supported Source Workloads For Migration to Non-Cloud Platforms,” on page 23](#), as appropriate for the target Hyper-V platform.

Network Access and Communications

- ♦ See [“Access and Communication Requirements across your Migration Network” on page 49](#).

Prerequisites

- ♦ See [Chapter 12, “Prerequisites for Migration to Microsoft Hyper-V,” on page 175](#).

Targets and Workloads

- ♦ **Target Hyper-V virtual host (automated):** See [“Target Discovery in the Migrate Client” on page 200](#)

- ♦ **Target VM on a Hyper-V virtual host (semi-automated):** See [“Registering and Discovering Target VMs on Virtual Hosts”](#) on page 205.
- ♦ **Source Workloads:** Use either of the following discovery methods:
 - ♦ [“Workload Discovery in the Migrate Client”](#) on page 214
 - ♦ [“Registering Workloads and Discovering Details with Migrate Agent”](#) on page 215

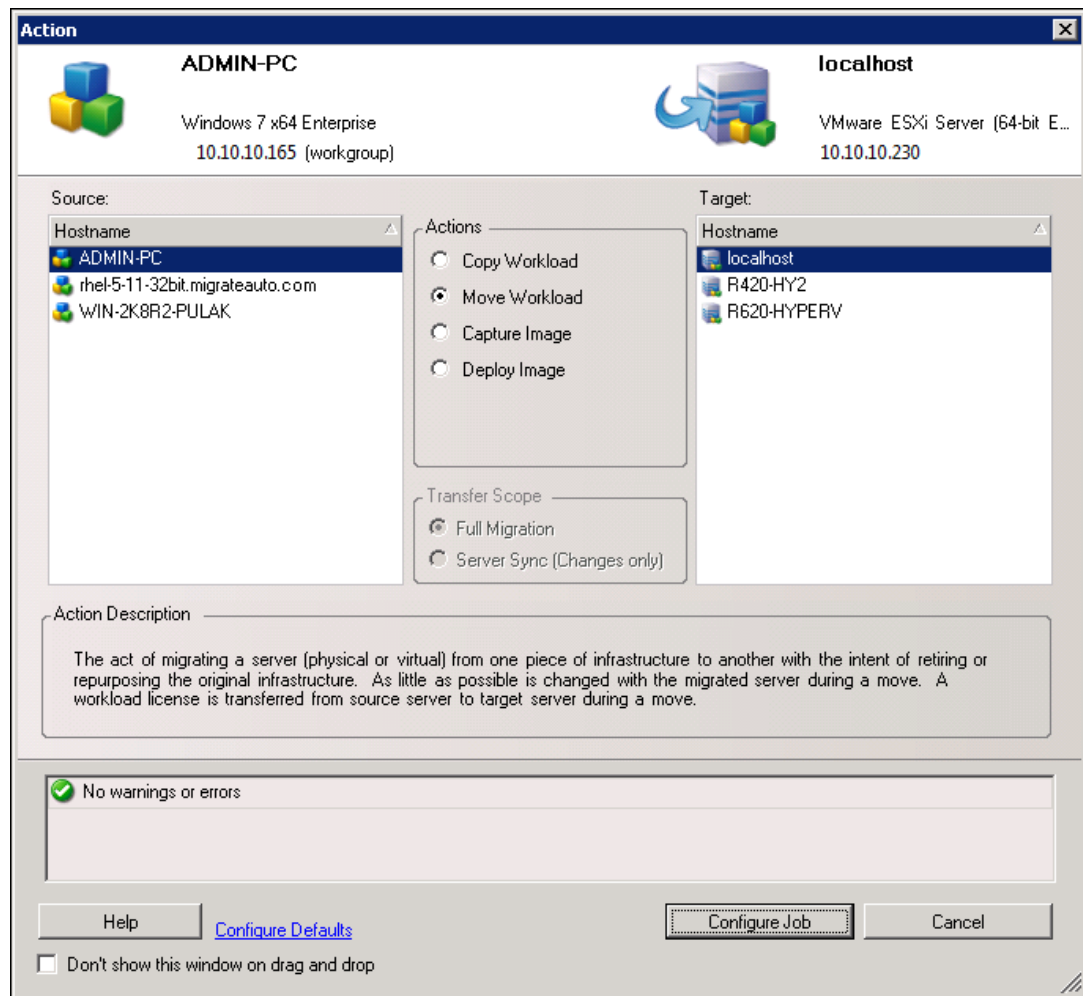
Additional Information

- ♦ [Microsoft Hyper-V Getting Started Guide](http://technet.microsoft.com/en-us/library/cc732470.aspx) (<http://technet.microsoft.com/en-us/library/cc732470.aspx>)[https://technet.microsoft.com/en-us/library/mt169373\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/mt169373(v=ws.11).aspx)
- ♦ [Hyper-V](https://technet.microsoft.com/en-us/library/mt169373(v=ws.11).aspx) ([https://technet.microsoft.com/en-us/library/mt169373\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/mt169373(v=ws.11).aspx))

31.2 Automated Migration to Hyper-V

- 1 Discover or refresh your source workload and your target VM host.
See [“Discovering and Preparing Workloads and Targets”](#) on page 195.
- 2 In the Migrate Client, initiate a peer-to-peer workload migration.
 - 2a Expand the Tasks options, then select the conversion type, depending on your goals for the migration:
 - ♦ **Copy Workload**
 - ♦ **Move Workload**

The Source and Target panes display workloads and targets applicable to the selected type of a migration job.



- 2b In the Source pane, select the workload you want to migrate.
- 2c In the Target pane, select target host for the migration.
- 2d Check the validation messages at the bottom of the window.
- 2e Click **Configure Job** to access the Peer-to-Peer Migration Job window.
- 3 In the Job Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
License	
License Key	<p>PlateSpin Migrate automatically selects the best license key for a migration job. If you have multiple license keys, you can specify the license key to use for the workload, assuming licenses are available (neither expired nor exhausted).</p> <p>To specify an alternate key to use:</p> <ol style="list-style-type: none"> 1. Deselect Automatically select the best license key during the conversion, then select the appropriate license key from the menu. 2. Click OK. <p>The selected license key is displayed on the License tab and its description is updated.</p>

Setting Name	Description
Conversion	
Transfer Scope	Specify the scope of workload data to transfer from the source to the target as Full Migration or Server Sync (Changes Only) .
Transfer Method	Specify how data is transferred from source to target. The availability depends on your workload and migration job type. See "Supported Data Transfer Methods" on page 42.
End State	
Source Machine End State	Specify whether to shut down the source workload after a successful cutover. For a workload move, the shut down is selected by default.
Target Virtual Machine End State	Specify whether to power on, power off, or suspend the target workload after a successful cutover.
Network	
Compression	Specify whether to compress data during transmission between the source and target workloads, and the level of data compression to apply: Full , Optimal , Maximum . See "Compression during Data Transfer" on page 299.
Encryption	Select Encrypt Data Transfer to encrypt the data as it is transferred from source to target. See "Security and Privacy" on page 43.
Bandwidth Throttling	<p>Specify whether to throttle bandwidth for data transfer traffic between the source and target machines. To enable throttling, select the Enable Bandwidth Throttling option, specify the required maximum value in Mbps, and optionally a time period during which to enforce the throttling. If specified, the from and to time values are based on the source workload's system time.</p> <p>If no time interval is defined, bandwidth is throttled to the specified rate at all times by default. If time interval is defined and the migration job executes outside this interval, data is transferred at full speed.</p>
IP Addresses	<p>Specify additional IP addresses for source workloads to enable communication in environments that use network address translation (NAT).</p> <p>For information on how to specify additional IP addresses for your PlateSpin Server, see "Migrations Across Public and Private Networks through NAT" on page 53.</p>
Schedule	
Schedule	<p>Specify when to start the migration job:</p> <ul style="list-style-type: none"> ♦ Run immediately ♦ Run at a later time <p>Use the calendar menu to specify the date and time to begin the migration.</p> <p>NOTE: You must prepare the workload prior to the scheduled time. The full replication cannot run unless the target VM exists and the workload preparation is complete. Migrate skips the scheduled full replication and retries it at the next scheduled time.</p>

Setting Name	Description
Access Settings	
Source Credentials	<p>(Windows) Specify the account user name with local or domain-level administrative privileges and a valid password. Use this format:</p> <ul style="list-style-type: none"> ♦ For domain member machines: <i>authority\principal</i> ♦ For workgroup member machines: <i>hostname\principal</i> <p>(Linux) Specify the root or root-level user name and a valid password.</p>
Target Credentials	Provide Windows Domain or Administrator credentials.
Alerts	
Receive Event Notifications	Specify whether to send email notifications for event conditions. You must configure an SMTP server to use this feature. See "Notification Service Using Migrate Client" on page 100.
Receive Progress Notifications	If you enable Event notifications, you can optionally receive progress notifications at a specified interval.
Send to Addresses	Add or remove valid email addresses for recipients of the notifications.
Take Control Settings	
Target Virtual Machine	Under Target Virtual Machine, click Configure, then specify the options for the virtual network and the TCP/IP settings for the replication NIC, then click OK .
Post-Migration	
Action	Specify a pre-configured action from the PlateSpin Migrate library. See Section 6.5, "Managing Post-Migration Actions (Windows and Linux)," on page 121.
Execution Parameters	Specify the command line command to run the selected action. You can specify a timeout for the execution.
Credentials	Specify the user name and password to use for the post-migration tasks. You can optionally use the source credentials.

- 4 In the Virtual Machine Configuration section of the Migration Job window, click **General**, then configure the required settings.

PlateSpin Migrate displays target virtual machine configuration options specific to the selected target and also provides access to advanced configuration options. For information about host-specific configuration options, see ["Target VM Configuration: Microsoft Hyper-V"](#).

- 5 In the Network Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Network Configuration	
Network Identification Settings for Windows	
Host Name	Specify the desired host name for the target machine.

Setting Name	Description
Generate New SID	When this option is selected, the target workload is assigned a new System Identifier (SID). Credentials are required only for Windows 2008 systems, and must be the credentials for the local (embedded) Administrator account. If this account has been locally renamed on the source, provide the new name.
Member of Domain / Workgroup	Select the required option and type the name of the domain or workgroup that you want the target machine to join.
Preserve Source Server's Domain Registration	Preserves domain registration and ensures that the source server domain registration remains intact during migration. If you disable this option, the source machine's domain account is transferred to the target machine. The source server still appears to be on the domain, but does not have a valid connection.
Domain Credentials	If the target machine is to be part of a domain, specify valid credentials for a user account with permission to add servers to the domain, such as a member of the Domain Admins group or Enterprise Admins group.
Network Identification Settings for Linux	
Host Name	On the Network Identification tab, specify the desired host name for the target machine.
DNS	Use the Add , Edit , and Remove buttons to manage DNS server entries for the new virtual machine.

6 In the Operating System and Applications Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Operating System and Application Configuration	
Windows Services (Target)	<p>Select Windows services' start conditions on the target VM after cutover. Start options are Automatic, Manual, Disabled, and Automatic (Delayed Start).</p> <p>To modify the settings:</p> <ol style="list-style-type: none"> 1. Click the Status column for the service, then select from the Windows start options. 2. When you are done setting services start states, click OK.
Live Transfer Services (Source)	<p>Specify the Windows services to stop on the source workload during live data transfers.</p> <p>We recommend that all the non-VSS compliant services or antivirus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.</p> <p>To modify the settings:</p> <ol style="list-style-type: none"> 1. Select Stopped next to the service to be stopped for live data transfer. 2. When you are done setting services to stop, click OK.

Setting Name	Description
Linux Daemons (Target)	Specify the start states for daemons on the target VM after cutover. To modify the settings: <ol style="list-style-type: none"> Click the Run Level column for the daemon, then select from run levels 0 through 6 and Boot (B), then click OK. When you are done setting daemon start states, click OK.
Live Transfer Daemons (Source)	Specify the daemons to stop on the source workload during live data transfers. To modify the settings: <ol style="list-style-type: none"> Select Stopped next to the daemon to be stopped for live data transfer. When you are done setting daemons to stop, click OK.

- 7** In the Drive Configuration section of the Migration Job window, configure the following settings. For options specific Hyper-V, see [Section 31.2.2, “Drive Configuration: Hyper-V,” on page 392](#).


Setting Name	Description
Drive Configuration	
Hard Drives	Specify drive and volume configurations to be migrated.
Disks	Specify the path to the hard disk on the target virtual machine.
Volumes	Select volumes to be included in the target for migration.
NTFS Cluster Size	(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see the Microsoft Support KB Article 140365 .
Non-volume Storage	(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups	(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.
Volume Groups	(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the Converted Logical Volumes section of the settings.
Converted Logical Volumes	(For Linux Workloads) Specify one or more LVM logical volumes to be migrated for a Linux workload.

- 8** In the Additional Items for Review section of the Migration Job window, review errors and messages about the workload configuration. You must resolve errors before you can submit the migration job.

- 9** Click **OK**.

31.2.1 Target VM Configuration: Microsoft Hyper-V

The following are configuration options specific to Hyper-V 2012 systems.



Virtual Machine Name: Specify the display name for the new virtual machine.

Datastore: Select the datastore where you want to create the *.vmtx file.

Configuration File Path: Specify a name and the directory path for the virtual machine's *.vmtx configuration file.

Virtual Machine Memory Allocation: Specify a value for the amount of virtual RAM to be assigned to the virtual machine.

Virtual Machine Generation Type: Select one of the following generations type for the new virtual machine:

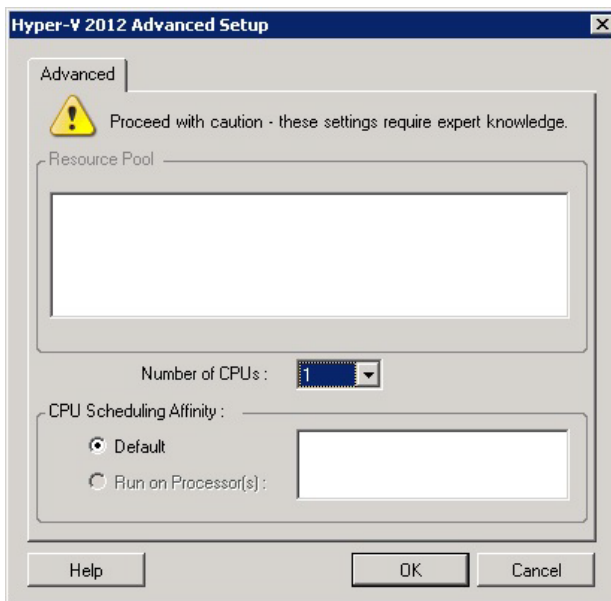
- ♦ **Generation 1:** Select this option to deploy the target virtual machine with Hyper-V BIOS architecture.
- ♦ **Generation 2:** Select this option to deploy the target virtual machine with Hyper-V UEFI architecture

This option is displayed only for Hyper-V machines.

Install Integration Services: Enable this option to install integration services during the migration process (recommended).

SCSI Drives: Select either **BusLogic** or **LSILogic** (the recommended option).

Advanced: Click this button to view or modify advanced VM configuration settings.



Number of CPUs: Select the required number of CPUs to assign to the target VM. For example, you can convert a single-processor workload to a multi-processor VM, or a multi-processor workload to a single-processor VM.

NOTE: For Generation 1, you can create four legacy network cards and eight synthetic network cards (if integration service is enabled). For Generation 2, you can create eight 8 synthetic network cards.

CPU Scheduling Affinity: Represents which Hyper-V Server processors the virtual machine can run on (if your Hyper-V Server is a multiprocessor system). Specify the required processor or select **Default** (recommended).

For details, see your Hyper-V documentation.

31.2.2 Drive Configuration: Hyper-V

The following are drive configuration settings specific to Hyper-V:

Drive Configuration

Settings

Virtual Disks to create :

Disk	Datastore	File Name	Size	Thin	Disk Type
1	C:	\Hyper-V\WIN-EFI-12R2_VM\Virtual Ha...\WIN-EFI-12R2_VM_1.vhdx	39.8 GB	<input type="checkbox"/>	IDE
2	C:	\Hyper-V\WIN-EFI-12R2_VM\Virtual Ha...\WIN-EFI-12R2_VM_2.vhdx	2 GB	<input type="checkbox"/>	SCSI
3	C:	\Hyper-V\WIN-EFI-12R2_VM\Virtual Ha...\WIN-EFI-12R2_VM_3.vhdx	3 GB	<input type="checkbox"/>	SCSI

Select volumes to copy and size :

Copy ?	Volume	Free Space	Size	New Free Space	New Size	To Disk
<input checked="" type="checkbox"/>	C: (NTFS - Boot)	26.6 GB	34.6 GB	26.6 GB	34.6 GB	1
<input checked="" type="checkbox"/>	E: (NTFS)	4.2 GB	4.9 GB	4.2 GB	4.9 GB	1
<input checked="" type="checkbox"/>	F: (NTFS)	1.3 GB	2 GB	1.3 GB	2 GB	2
<input checked="" type="checkbox"/>	G: (NTFS)	1 GB	1 GB	1 GB	1 GB	3
<input checked="" type="checkbox"/>	H: (NTFS)	2 GB	2 GB	2 GB	2 GB	3

Select non-volume partitions to recreate and size :

Create ?	Type	Partition	Size	New Size	To Disk
<input type="checkbox"/>	0x7 (NTFS (HPFS))	\disk0\partition1\	99 MB	99 MB	1
<input type="checkbox"/>	0x0 (Empty)	\disk0\partition2\	128 MB	128 MB	1
<input type="checkbox"/>	0x0 (Empty)	\disk2\partition0\	32 MB	32 MB	3

Help OK Cancel

Datastore: Select the datastore volume on the Hyper-V server where you want to place the .vhdx and .vhd files.

Disk Type: A Generation 1 disk containing the System/Boot volume should be on an IDE disk. (You can create a maximum of three IDE disks.)

NOTE: For a Generation 1 disk, the values of second and third disk are chained. For example, if you select the third disk (from the top of the **Disk Type** list) as **IDE**, the second disk autoselects as **IDE**. If you select the second disk as a **SCSI** then the third disk autoselects to **SCSI**.

Copy?: Select the volumes to be copied during the migration.

New Free Space: To resize the volume during the migration, specify the desired amount of free space. PlateSpin Migrate automatically adjusts New Size.

New Size: To resize the volume during the migration, specify the desired size. PlateSpin Migrate automatically adjusts New Free Space.

To Disk: Assign the volume to a disk or, if LVM is enabled, to a volume group. The volume is copied to this disk or volume group on the target machine.

Create?: Select any non-volume disk partitions that should be created on the target machine (for example, a Linux swap partition).

New Size: To resize the non-volume partition during the migration, specify the desired size.

31.3 Migration to VMs on Hyper-V Using X2P Workflow

For migration of workloads to a Hyper-V virtual host using X2P workflow, PlateSpin Migrate requires that you manually set up the target virtual machine with guest operating system type and version settings that match your source workload, in accordance with the features and capabilities of the Hyper-V virtualization platform. Use the PlateSpin ISO to register the target machine with the PlateSpin Server and send machine details. Use PlateSpin Migrate Client to configure, execute, and manage the migration job.

Use the guidelines in this section to configure migration to VMs on Hyper-V virtual hosts.

- ♦ [Section 31.3.1, “Downloading and Saving the PlateSpin ISO Image \(Hyper-V\),” on page 393](#)
- ♦ [Section 31.3.2, “Creating and Configuring the Target Virtual Machine \(Hyper-V\),” on page 393](#)
- ♦ [Section 31.3.3, “Registering the Virtual Machine with PlateSpin Server \(Hyper-V\),” on page 394](#)
- ♦ [Section 31.3.4, “Migrating Your Source Workload to the Target Virtual Machine \(Hyper-V\),” on page 394](#)
- ♦ [Section 31.3.5, “Post-Migration Steps \(Hyper-V\),” on page 394](#)

31.3.1 Downloading and Saving the PlateSpin ISO Image (Hyper-V)

- 1 Download and prepare the PlateSpin ISO image for use with the target VM. Attended and unattended registration options are possible.

See [“Preparing the PlateSpin ISO Image for Target Registration and Discovery” on page 280](#).

- 2 Save the ISO image in a location that Hyper-V server can access. For example: `c:\temp`.

This ensures that the PlateSpin ISO image is available to the target VM as a bootable CD-ROM image.

31.3.2 Creating and Configuring the Target Virtual Machine (Hyper-V)

- 1 In the Hyper-V Manager, use the New Virtual Machine Wizard to create a new virtual machine with the following settings:
 - ♦ **Name and Location:** Specify a name for your new target and accept the default location.
 - ♦ **Operating System Type and Version:** Specify the operating system type and version settings that matches the source workload. The wizard uses this information to set appropriate default values, such as the amount of memory needed, and resource limits for the VM.
 - ♦ **Assign Memory:** Assign at least 384 MB of RAM to the VM.
 - ♦ **Connect Virtual Hard Disk:** Ensure that the disk size of every disk is about 50 MB more than the corresponding disk on your source workload.
 - ♦ **Installation Options:** Configure the VM to boot from an ISO image file, and point the wizard to the downloaded PlateSpin ISO image.
 - ♦ **Summary:** Configure the VM to not start upon creation (deselect the **Start the virtual machine after it is created** option).
- 2 After creating the VM, remove the default NIC and replace it with a generic one, called *Legacy Network Adapter*.

This is required because the New Virtual Machine Wizard creates a NIC of a custom Microsoft type, which is currently unsupported by PlateSpin Migrate.

- 3 Connect the newly added NIC (*Legacy Network Adapter*) to the external virtual network.

31.3.3 Registering the Virtual Machine with PlateSpin Server (Hyper-V)

After you create the virtual machine and prepare it to boot with the PlateSpin ISO, you are ready to register it as a target VM with your PlateSpin Server. See [“Registering and Discovering Target VMs on Virtual Hosts” on page 205](#).

31.3.4 Migrating Your Source Workload to the Target Virtual Machine (Hyper-V)

- 1 Use PlateSpin Migrate Client to start an X2P migration job with your source workload being the job’s migration source and the target being the new VM on Hyper-V.
See [“Migration to Physical Machines” on page 407](#).
- 2 For host-specific target VM configuration options for the Virtual Machine Configuration dialog, see [“Target VM Configuration: Microsoft Hyper-V” on page 390](#).
- 3 For host-specific storage configuration options, see [Section 31.2.2, “Drive Configuration: Hyper-V,” on page 392](#).
- 4 Monitor the migration job in Jobs view in PlateSpin Migrate Client.
When the job reaches the **Configure Target Machine** step, the virtual machine’s console returns to the boot prompt of the PlateSpin ISO image.
- 5 Shut down the virtual machine and reconfigure it to boot from disk rather than from the boot image.
- 6 Power on the virtual machine.
The migration job resumes, reboots the target, and completes the workload configuration.

31.3.5 Post-Migration Steps (Hyper-V)

Install Hyper-V Integration Services (virtualization enhancement software). For more information, see your [Microsoft Hyper-V Getting Started Guide](#).

32 Migration to Virtual Machines on Citrix XenServer

For migration of workloads to a Citrix XenServer virtual host, PlateSpin Migrate requires that you manually set up the target virtual machine with guest operating system type and version settings that match your source workload, in accordance with the features and capabilities of the XenServer virtualization platform. Use the PlateSpin ISO to register the target machine with the PlateSpin Server and send machine details. Use PlateSpin Migrate Client to configure, execute, and manage the migration job.

Use the guidelines in this section to configure migration to VMs on Citrix XenServer virtual hosts.

- ♦ [Section 32.1, “Planning for Migration to Citrix XenServer,” on page 395](#)
- ♦ [Section 32.2, “Configuring Migration to a VM on a Citrix XenServer Virtual Host,” on page 396](#)

32.1 Planning for Migration to Citrix XenServer

Before you begin migrations to virtual machines on Citrix XenServer virtual hosts, ensure that your migration environment meets the following guidelines:

Supported Citrix XenServer Platforms

- ♦ See “[Citrix XenServer](#)” in [Table 2-14, “Supported Target Virtualization Platforms for the Migrate Client Only,” on page 39](#).

Supported Workloads

- ♦ See “[Supported Source Workloads For Migration to Non-Cloud Platforms](#)” on page 23, as appropriate for the target Citrix XenServer platform.

Network Access and Communications

- ♦ See “[Access and Communication Requirements across your Migration Network](#)” on page 49.

Prerequisites

- ♦ See “[Prerequisites for Migration to VMs on Citrix XenServer](#)” on page 179.

Targets and Workloads

- ♦ **Target VM on a Citrix XenServer virtual host (semi-automated):** See “[Registering and Discovering Target VMs on Virtual Hosts](#)” on page 205.
- ♦ **Source Workloads:** Use either of the following discovery methods:
 - ♦ “[Workload Discovery in the Migrate Client](#)” on page 214
 - ♦ “[Registering Workloads and Discovering Details with Migrate Agent](#)” on page 215

Additional Information

- ♦ [Citrix XenServer 6.1.0 Administrator's Guide](http://docs.vmd.citrix.com/XenServer/6.1.0/1.0/en_gb/reference.html) (http://docs.vmd.citrix.com/XenServer/6.1.0/1.0/en_gb/reference.html)

32.2 Configuring Migration to a VM on a Citrix XenServer Virtual Host

You can use Citrix XenServer as the target virtualization platform in a semi-automated workload virtualization.

This section includes the following topics:

- ♦ [Section 32.2.1, “Downloading and Preparing the PlateSpin ISO Image \(Citrix XenServer\),” on page 396](#)
- ♦ [Section 32.2.2, “Creating and Configuring the Target Virtual Machine \(Citrix XenServer\),” on page 396](#)
- ♦ [Section 32.2.3, “Registering the Virtual Machine with PlateSpin Server \(Citrix XenServer\),” on page 397](#)
- ♦ [Section 32.2.4, “Migrating Your Source Workload to the Target Virtual Machine \(Citrix XenServer\),” on page 397](#)
- ♦ [Section 32.2.5, “Target VM Configuration: Citrix XenServer,” on page 398](#)

32.2.1 Downloading and Preparing the PlateSpin ISO Image (Citrix XenServer)

- 1 Download and prepare the PlateSpin ISO image for use with the target VM. Attended and unattended registration options are possible.

See [“Preparing the PlateSpin ISO Image for Target Registration and Discovery” on page 280](#).

- 2 Save the downloaded image file in the following directory on the Citrix XenServer host:

```
/var/lib/xen/images
```

This ensures that the PlateSpin ISO image is available to the target VM as a bootable CD-ROM image.

32.2.2 Creating and Configuring the Target Virtual Machine (Citrix XenServer)

- 1 On Citrix XenServer, use the Virtual Machine Manager Wizard or the Create Virtual Machines program shortcut to create a new virtual machine.

Ensure that the new virtual machine is created with the following settings:

- ♦ **Virtualization method:** Fully virtualized.
- ♦ **Operating System Type and Version:** Specify the operating system type and version settings that matches the source workload. The wizard uses this information to set appropriate default values (such as the amount of memory needed) and resource limits for the VM.
- ♦ **Memory:** Assign at least 384 MB of RAM to the VM. This ensures that the VM has sufficient resources during the migration and improves transfer speed. If the virtual machine requires less memory after the migration, reduce the assigned memory after the migration completes.

- ♦ **Disks:** Assign disks such that the disk size of every disk is about 50 MB more than the corresponding disk on your source workload. The storage can be either a raw SAN LUN or a virtual disk. Also, create a Virtual CD-ROM assigned to the downloaded PlateSpin ISO image.
- 2 Ensure that the VM is configured to restart on reboot by exporting the VM's settings from the xend database to a text file and making sure that the `on_reboot` parameter is set to `restart`. If not, shut down the VM, update the settings, and re-import them into the xend database.
For detailed instructions, see the [XenServer 6.1.0 Virtual Machine User's Guide \(http://support.citrix.com/article/CTX134587\)](http://support.citrix.com/article/CTX134587).
- 3 From the Virtual Machine Manager, launch the virtual machine console and monitor the boot process.
When the virtual machine completes the boot process, it prompts you for parameters that control the registration of the machine and its profile with PlateSpin Migrate. If you are using the unattended registration process, the required parameters are read from an answer file.

32.2.3 Registering the Virtual Machine with PlateSpin Server (Citrix XenServer)


After you create the virtual machine and prepare it to boot with the PlateSpin ISO, you are ready to register it as a target VM with your PlateSpin Server. See [“Registering and Discovering Target VMs on Virtual Hosts” on page 205](#).

32.2.4 Migrating Your Source Workload to the Target Virtual Machine (Citrix XenServer)

- 1 Use PlateSpin Migrate Client to start an X2P migration job with your source workload being the job's migration source and the target being the new VM on the Citrix XenServer hypervisor.
See [“Migration to Physical Machines” on page 407](#).
- 2 For host-specific target VM configuration options for the Virtual Machine Configuration dialog, see [“Target VM Configuration: Citrix XenServer” on page 398](#).
- 3 Monitor the migration job in the Jobs view in PlateSpin Migrate Client.
When the job reaches the **Configure Target Machine** step, the virtual machine's console returns to the boot prompt of the PlateSpin ISO image.
- 4 Shut down the virtual machine, reconfigure it to boot from disk rather than from the boot image, and deselect the **VS Tools Installed** option.
- 5 Power on the virtual machine.
The migration job resumes, reboots the target, and completes the workload configuration.

32.2.5 Target VM Configuration: Citrix XenServer

The following are configuration options specific to Citrix XenServer.



Virtual Machine Name: Specify the display name for the new virtual machine.

Number of CPUs: Select the number of CPUs to assign to the target VM. For example, you can convert a single-processor workload to a multi-processor VM, or a multi-processor workload to a single-processor VM.

Virtual Machine Memory Allocation: Specify a value for the amount of virtual RAM to be assigned to the virtual machine.

Install XenServer Tools: Enable this option to install XenServer Tools during the migration process (recommended).

33 Migration to Virtual Machines on Xen

For migration of workloads to a Xen virtual host, PlateSpin Migrate requires that you manually set up the target virtual machine with guest operating system type and version settings that match your source workload, in accordance with the features and capabilities of the Xen virtualization platform. Use the PlateSpin ISO to register the target machine with the PlateSpin Server and send machine details. Use PlateSpin Migrate Client to configure, execute, and manage the migration job.

Use the guidelines in this section to configure migration to VMs on Xen virtual hosts.

- ♦ [Section 33.1, “Planning for Migration to Xen,” on page 399](#)
- ♦ [Section 33.2, “Configuring Migration to a VM on a Xen Virtual Host,” on page 400](#)

33.1 Planning for Migration to Xen

Before you begin migrations to virtual machines on Xen virtual hosts, ensure that your migration environment meets the following guidelines:

Supported Xen Platforms

- ♦ See “[SUSE Linux Enterprise Server with Xen](#)” in Table 2-14, “Supported Target Virtualization Platforms for the Migrate Client Only,” on page 39.

Supported Workloads

- ♦ See “[Supported Source Workloads For Migration to Non-Cloud Platforms](#)” on page 23, as appropriate for the target Xen platform.

Network Access and Communications

- ♦ See “[Access and Communication Requirements across your Migration Network](#)” on page 49.

Prerequisites

- ♦ See “[Prerequisites for Migration to VMs on Xen](#)” on page 183.

Targets and Workloads

- ♦ **Target VM on a XEN virtual host (semi-automated):** See [Section 19.5.2, “Registering and Discovering Target VMs on Virtual Hosts,” on page 205.](#)
- ♦ **Source Workloads:** Use either of the following discovery methods:
 - ♦ “[Workload Discovery in the Migrate Client](#)” on page 214
 - ♦ “[Registering Workloads and Discovering Details with Migrate Agent](#)” on page 215

Additional Information

- ♦ [SUSE Linux Enterprise Server 11 SPX Virtualization with Xen](https://www.suse.com/documentation/sles11/singlehtml/book_xen/book_xen.html) (https://www.suse.com/documentation/sles11/singlehtml/book_xen/book_xen.html)

33.2 Configuring Migration to a VM on a Xen Virtual Host

You can use the Xen Hypervisor on SUSE Linux Enterprise Server 11 as the target virtualization platform in a semi-automated workload virtualization.

This section includes the following topics:

- ♦ [Section 33.2.1, “Downloading and Preparing the PlateSpin ISO Image \(Xen on SLES\),” on page 400](#)
- ♦ [Section 33.2.2, “Creating and Configuring the Target Virtual Machine \(Xen on SLES\),” on page 400](#)
- ♦ [Section 33.2.3, “Registering the Virtual Machine with PlateSpin Server \(Xen on SLES\),” on page 401](#)
- ♦ [Section 33.2.4, “Migrating Your Source Workload to the Target Virtual Machine \(Xen on SLES\),” on page 401](#)
- ♦ [Section 33.2.5, “Post-Migration Steps \(Xen on SLES\),” on page 401](#)

33.2.1 Downloading and Preparing the PlateSpin ISO Image (Xen on SLES)

- 1 Download and prepare the PlateSpin ISO image for use with the target VM. Attended and unattended registration options are possible.

See [“Preparing the PlateSpin ISO Image for Target Registration and Discovery” on page 280](#).

- 2 Save the prepared PlateSpin ISO image in the following directory:

```
/var/lib/xen/images
```

This ensures that the PlateSpin ISO image is available to the target VM as a bootable CD-ROM image.

33.2.2 Creating and Configuring the Target Virtual Machine (Xen on SLES)

- 1 On SLES 11, use the Virtual Machine Manager Wizard or the Create Virtual Machines program shortcut to create a new virtual machine.

Ensure that the new virtual machine is created with the following settings:

- ♦ **Virtualization method:** Fully virtualized.
- ♦ **Operating System Type and Version:** Specify the operating system type and version settings that matches the source workload. The wizard uses this information to set appropriate default values (such as the amount of memory needed) and resource limits for the VM.
- ♦ **Memory:** Assign at least 384 MB of RAM to the VM. This ensures that the VM has sufficient resources during the migration and improves transfer speed. If the virtual machine requires less memory after the migration, reduce the assigned memory after the migration completes.

- ♦ **Disks:** Assign disks such that the disk size of every disk is about 50 MB more than the corresponding disk on your source workload. The storage can be either a raw SAN LUN or a virtual disk. Also, create a Virtual CD-ROM assigned to the downloaded PlateSpin ISO image.
- 2 Ensure that the VM is configured to restart on reboot by exporting the VM's settings from the xend database to a text file and making sure that the `on_reboot` parameter is set to `restart`. If not, shut down the VM, update the settings, and re-import them into the xend database.
For detailed instructions, see your [SLES 11 documentation \(https://www.suse.com/documentation/sles11/\)](https://www.suse.com/documentation/sles11/).
- 3 From the Virtual Machine Manager, launch the virtual machine console and monitor the boot process.
When the virtual machine completes the boot process, it prompts you for parameters that control the registration of the machine and its profile with PlateSpin Migrate. If you are using the unattended registration process, the required parameters are read from an answer file.

33.2.3 Registering the Virtual Machine with PlateSpin Server (Xen on SLES)

After you create the virtual machine and prepare it to boot with the PlateSpin ISO, you are ready to register it as a target VM with your PlateSpin Server. See “[Registering and Discovering Target VMs on Virtual Hosts](#)” on page 205.

33.2.4 Migrating Your Source Workload to the Target Virtual Machine (Xen on SLES)

- 1 Use PlateSpin Migrate Client to start an X2P migration job with your source workload being the job's migration source and the target being the new VM on the Xen hypervisor.
See “[Migration to Physical Machines](#)” on page 407.
- 2 Monitor the migration job in the PlateSpin Migrate Client's Jobs view.
When the job reaches the **Configure Target Machine** step, the virtual machine's console returns to the boot prompt of the PlateSpin ISO image.
- 3 Shut down the virtual machine, reconfigure it to boot from disk rather than from the boot image, and deselect the **VS Tools Installed** option.
- 4 Power on the virtual machine.
The migration job resumes, reboots the target, and completes the workload configuration.

33.2.5 Post-Migration Steps (Xen on SLES)

Install SUSE Drivers for Xen (virtualization enhancement software). For more information, see the following online document:

[SUSE Linux Enterprise Server 11 SPX Virtualization with Xen \(https://www.suse.com/documentation/sles11/singlehtml/book_xen/book_xen.html\)](https://www.suse.com/documentation/sles11/singlehtml/book_xen/book_xen.html)

34 Migration to Virtual Machines on KVM

For migration of workloads to a KVM virtual host, PlateSpin Migrate requires that you manually set up the target virtual machine with guest operating system type and version settings that match your source workload, in accordance with the features and capabilities of the KVM virtualization platform. Use the PlateSpin ISO to register the target machine with the PlateSpin Server and send machine details. Use PlateSpin Migrate Client to configure, execute, and manage the migration job.

Use the guidelines in this section to configure migration to VMs on KVM virtual hosts.

- ♦ [Section 34.1, “Planning for Migration to KVM,” on page 403](#)
- ♦ [Section 34.2, “Configuring Migration to a VM on a KVM Virtual Host,” on page 404](#)

34.1 Planning for Migration to KVM

Before you begin migrations to virtual machines on KVM virtual hosts, ensure that your migration environment meets the following guidelines:

Supported KVM Platforms

- ♦ See the following information in [Table 2-14, “Supported Target Virtualization Platforms for the Migrate Client Only,” on page 39](#):
 - ♦ [“SUSE Linux Enterprise Server \(SLES\) with KVM”](#)
 - ♦ [“Red Hat Enterprise Linux \(RHEL\) with KVM”](#)

Supported Workloads

- ♦ See [“Supported Source Workloads For Migration to Non-Cloud Platforms” on page 23](#), as appropriate for the target KVM platform.

Network Access and Communications

- ♦ See [“Access and Communication Requirements across your Migration Network” on page 49](#).

Prerequisites

- ♦ See [“Prerequisites for Migration to VMs on KVM” on page 185](#).

Targets and Workloads

- ♦ **Target VM on a KVM virtual host (semi-automated):** See [“Registering and Discovering Target VMs on Virtual Hosts” on page 205](#).
- ♦ **Source Workloads:** Use either of the following discovery methods:
 - ♦ [“Workload Discovery in the Migrate Client” on page 214](#)
 - ♦ [“Registering Workloads and Discovering Details with Migrate Agent” on page 215](#)

Additional Information

- ♦ *SUSE Linux Enterprise Server 11 SPX Virtualization with KVM* (https://www.suse.com/documentation/sles11/singlehtml/book_kvm/book_kvm.html)
- ♦ *Red Hat Enterprise Linux 7.X Virtualization Deployment and Administration Guide* (https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Virtualization_Deployment_and_Administration_Guide/index.html)

34.2 Configuring Migration to a VM on a KVM Virtual Host

You can use KVM as the target virtualization platform in a semi-automated workload virtualization.

- ♦ Section 34.2.1, “Downloading and Preparing the PlateSpin ISO Image (KVM),” on page 404
- ♦ Section 34.2.2, “Creating and Configuring the Target Virtual Machine (RHEL KVM),” on page 404
- ♦ Section 34.2.3, “Registering the Virtual Machine with PlateSpin Server (RHEL KVM),” on page 405
- ♦ Section 34.2.4, “Migrating Your Source Workload to the Target Virtual Machine (RHEL KVM),” on page 405

34.2.1 Downloading and Preparing the PlateSpin ISO Image (KVM)

- 1 Download and prepare the PlateSpin ISO image for use with the target VM. Attended and unattended registration options are possible.

See “Preparing the PlateSpin ISO Image for Target Registration and Discovery” on page 280.

- 2 Save the ISO image in a location that the KVM virtual host can access.

This ensures that the PlateSpin ISO image is available to the target VM as a bootable CD-ROM image.

34.2.2 Creating and Configuring the Target Virtual Machine (RHEL KVM)

- 1 On RHEL KVM, use the Virtual Machine Manager Wizard or the Create Virtual Machines program shortcut to create a new virtual machine.

Ensure that the new virtual machine is created with the following settings:

- ♦ **Virtualization method:** Fully virtualized.
- ♦ **Operating System Type and Version:** Specify the operating system type and version settings that matches the source workload. The wizard uses this information to set appropriate default values (such as the amount of memory needed) and resource limits for the VM.
- ♦ **Memory:** Assign at least 384 MB of RAM to the VM. This ensures that the VM has sufficient resources during the migration and improves transfer speed. If the virtual machine requires less memory after the migration, reduce the assigned memory after the migration completes.

- ♦ **Disks:** Assign disks such that the disk size of every disk is about 50 MB more than the corresponding disk on your source workload. The storage can be either a raw SAN LUN or a virtual disk. Also, create a Virtual CD-ROM assigned to the downloaded PlateSpin ISO image.
- 2 Ensure that the VM is configured to restart on reboot.
 - 3 From the Virtual Machine Manager, launch the virtual machine console and monitor the boot process.

When the virtual machine completes the boot process, it prompts you for parameters that control the registration of the machine and its profile with PlateSpin Migrate. If you are using the unattended registration process, the required parameters are read from an answer file.

34.2.3 Registering the Virtual Machine with PlateSpin Server (RHEL KVM)

After you create the virtual machine and prepare it to boot with the PlateSpin ISO, you are ready to register it as a target VM with your PlateSpin Server. See [“Registering and Discovering Target VMs on Virtual Hosts” on page 205](#).

34.2.4 Migrating Your Source Workload to the Target Virtual Machine (RHEL KVM)

- 1 Use PlateSpin Migrate Client to start an X2P migration job with your source workload being the job's migration source and the target being the new VM on the RHEL KVM hypervisor.
See [“Migration to Physical Machines” on page 407](#).
- 2 Monitor the migration job in the Jobs view in PlateSpin Migrate Client.
When the job reaches the **Configure Target Machine** step, the virtual machine's console returns to the boot prompt of the PlateSpin ISO image.
- 3 Shut down the virtual machine, reconfigure it to boot from disk rather than from the boot image.
- 4 Power on the virtual machine.
The migration job resumes, reboots the target, and completes the workload configuration.

35

Migration to Physical Machines

PlateSpin Migrate supports semi-automated migration to physical machines. You prepare the target machine to meet migration needs, and then use PlateSpin Migrate to automate the data migration. Use the guidelines in this section to configure migration to physical machines.

- ♦ [Section 35.1, “Planning for Migration to Physical Machines,” on page 407](#)
- ♦ [Section 35.2, “Configuring Migration to a Physical Target \(P2P, V2P\),” on page 408](#)

35.1 Planning for Migration to Physical Machines

Before you begin migrations to physical machines, ensure that your migration environment meets the following guidelines:

Supported Physical Hardware

- ♦ See the following information in [“Supported Configurations” on page 23](#):
 - ♦ [Supported Workload Storage](#)
 - ♦ [Supported Workload Architectures](#)

Supported Workloads

- ♦ See [“Supported Source Workloads For Migration to Non-Cloud Platforms” on page 23](#).

Network Access and Communications

- ♦ See [“Access and Communication Requirements across your Migration Network” on page 49](#).

Prerequisites

- ♦ See [“Prerequisites for Migration to Physical Machines” on page 187](#).

Targets and Workloads

- ♦ **Target physical host (semi-automated):** See [“Registering and Discovering Details for Target Physical Machines with PlateSpin ISO” on page 206](#).
- ♦ **Source Workloads:** Use either of the following discovery methods:
 - ♦ [“Workload Discovery in the Migrate Web Interface” on page 215](#)
 - ♦ [“Registering Workloads and Discovering Details with Migrate Agent” on page 215](#)

35.2 Configuring Migration to a Physical Target (P2P, V2P)

To initiate a peer-to-peer workload migration to a physical machine:

- 1 (Recommended) Use PlateSpin Analyzer to ensure that:
 - ♦ Your source operating system and hardware are supported by PlateSpin Migrate.
 - ♦ PlateSpin Migrate's X2P device driver database contains device drivers that your target requires for the operating system being ported.

See [“Analyzing Suitability of Discovered Windows Workloads For Conversion to Physical Machines”](#) on page 229.

- 2 Discover your source workload.

Use either of the following discovery methods:

- ♦ [Section 20.6.1, “Workload Discovery in the Migrate Client,”](#) on page 214
- ♦ [Section 20.7, “Registering Workloads and Discovering Details with Migrate Agent,”](#) on page 215

- 3 (Conditional) If drivers for the physical target are not available in the PlateSpin Migrate's X2P device driver database, upload the required drivers to the database.

See [Chapter 21, “Preparing Device Drivers,”](#) on page 221.

- 4 Register your target physical machine with PlateSpin Migrate by booting it with the PlateSpin Boot OFX ISO.

See [Section 19.6, “Registering and Discovering Details for Target Physical Machines with PlateSpin ISO,”](#) on page 206.

- 5 Launch Migrate Client, then start a peer-to-peer workload migration.

The Source and Target panes display workloads and targets applicable to the selected type of a migration job.

See [Section 26.2, “Initiating a Migration Job,”](#) on page 292.

- 5a Under Tasks, select the conversion type, depending on your goals for the migration:

- ♦ **Copy Workload**
- ♦ **Move Workload**

In the Action dialog, the Transfer Scope is set to **Full Migration**.

- 5b In the Source pane, select the workload you want to migrate.

- 5c In the Target pane, select target physical machine for the migration.

- 5d Read the validation messages at the bottom of the window.

- 5e Click **Configure Job** to access the Peer-to-Peer Migration Job window.

- 6 Configure the required parameters of the job.

See [Chapter 26, “Configuration Essentials,”](#) on page 291.

Setting Name	Description
License	
Licenses License Key	<p>PlateSpin Migrate automatically selects the best license key for a migration job. If you have multiple license keys, you can specify the license key to use for the workload, assuming licenses are available (neither expired nor exhausted).</p> <p>To specify an alternate key to use:</p> <ol style="list-style-type: none"> 1. Deselect Automatically select the best license key during the conversion, then select the appropriate license key from the menu. 2. Click OK. <p>The selected license key is displayed on the License tab and its description is updated.</p>
Conversion	
Transfer Scope	Set by default to Full Migration .
Transfer Method	Specify how data is transferred from source to target. The availability depends on your workload and migration job type. See Section 2.2, "Supported Data Transfer Methods," on page 42.
End State	
Source Machine End State	Specify whether to shut down the source workload after a successful cutover. For a workload move, Shutdown is selected by default.
Target Virtual Machine End State	Specify whether to power on, power off, or suspend the target workload after a successful cutover.
Network	
Compression	Specify whether to compress data during transmission between the source and target workloads, and the level of data compression to apply: Fast , Optimal , or Maximum . Compression is disabled by default. See Section 26.9, "Compression during Data Transfer," on page 299.
Encryption	Select Encrypt Data Transfer to encrypt the data as it is transferred from source to target. See Section 2.3, "Security and Privacy," on page 43.
Bandwidth Throttling	<p>Select Enable Throttling to control the amount of available bandwidth consumed by direct source-to-target communication over the course of a workload migration. Specify the required throughput value in Mbps and the time pattern. Bandwidth throttling is disabled by default. See Section 26.10, "Bandwidth Throttling during Data Transfer," on page 300.</p> <p>Time-based throttling is based on the source server time.</p>
Advanced Additional Source Machine Addresses	<p>Specify additional IP addresses for source workloads to enable communication in environments that use network address translation (NAT).</p> <p>See Section 2.6.4, "Migrations Across Public and Private Networks through NAT," on page 53.</p>

Setting Name	Description
Schedule	
Schedule	<p>Specify when to start the migration job:</p> <ul style="list-style-type: none"> ♦ Run immediately ♦ Run at a later time <p>Use the calendar menu to specify the date and time to begin the migration.</p> <p>NOTE: You must prepare the target machine prior to the scheduled time. The full replication cannot run unless the target machine is available. Migrate skips the scheduled full replication and retries it at the next scheduled time.</p>
Access Settings	
Source Credentials	<p>(Windows) Specify the account user name with local or domain-level administrative privileges and a valid password. Use this format:</p> <ul style="list-style-type: none"> ♦ For domain member machines: <i>authority\principal</i> ♦ For workgroup member machines: <i>hostname\principal</i> <p>(Linux) Specify the <code>root</code> or root-level user name and a valid password.</p>
Target Credentials	
Alerts	
Receive Event Notifications	Specify whether to send email notifications for event conditions. You must configure an SMTP server to use this feature. See Section 5.4.1, "Notification Service Using Migrate Client," on page 100.
Receive Progress Notifications	If you enable Event notifications, you can optionally receive progress notifications at a specified interval.
Send to Addresses	Add or remove valid email addresses for recipients of the notifications.
Take Control Settings	
Target Virtual Machine	Under Target Virtual Machine, click Configure , then specify the options for the virtual network and the TCP/IP settings for the replication NIC, then click OK .
Post-Migration	
Action	Specify a pre-configured action from the PlateSpin Migrate library. See Section 6.5, "Managing Post-Migration Actions (Windows and Linux)," on page 121.
Execution Parameters	Specify the command line command to run the selected action. You can specify a timeout for the execution.
Credentials	Specify the user name and password to use for the post-migration tasks. You can optionally use the source credentials.

7 (Target VMs using X2P workflow) In the Virtual Machine Configuration section of the Migration Job window, click **General**, then configure the required settings.

PlateSpin Migrate displays target virtual machine configuration options specific to the selected target and also provides access to advanced configuration options for some platforms. For information about host-specific configuration options, see:

- ♦ ["Target VM Configuration: VMware ESXi 5 and Later"](#)

- ◆ “Target VM Configuration: VMware ESX 4.1”
- ◆ “Target VM Configuration: Microsoft Hyper-V”
- ◆ “Target VM Configuration: Citrix XenServer”

Setting Name	Description
Virtual Machine Name	Specify a name to use for the target VM as it appears in the virtual host environment.
Number of CPUs	Select the number of CPUs to assign to the target VM. For example, you can convert a single-processor workload to a multi-processor VM, or a multi-processor workload to a single-processor VM.
Virtual Machine Memory Allocation	Specify the amount of virtual memory.

8 In the Network Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Network Configuration	
Network Identification Settings for Windows	
Host Name	Specify the desired host name for the target machine.
Generate New SID	When this option is selected, the target workload is assigned a new System Identifier (SID). Credentials are required only for Windows 2008 systems, and must be the credentials for the local (embedded) Administrator account. If this account has been locally renamed on the source, provide the new name.
Member of Domain / Workgroup	Select the required option and type the name of the domain or workgroup that you want the target machine to join.
Preserve Source Server's Domain Registration	Preserves domain registration and ensures that the source server domain registration remains intact during migration. If you disable this option, the source machine's domain account is transferred to the target machine. The source server still appears to be on the domain, but does not have a valid connection.
Domain Credentials	If the target machine is to be part of a domain, specify valid credentials for a user account with permission to add servers to the domain, such as a member of the Domain Admins group or Enterprise Admins group.
Network Identification Settings for Linux	
Host Name	On the Network Identification tab, specify the desired host name for the target machine.
DNS	Use the Add , Edit , and Remove buttons to manage DNS server entries for the new virtual machine.

- 9 In the Operating System and Applications Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Operating System and Application Configuration	
Windows Services (Target)	<p>Select Windows services' start conditions on the target VM after cutover. Start options are Automatic, Manual, Disabled, and Automatic (Delayed Start).</p> <p>To modify the settings:</p> <ol style="list-style-type: none">1. Click the Status column for the service, then select from the Windows start options.2. When you are done setting services start states, click OK.
Live Transfer Services (Source)	<p>Specify the Windows services to stop on the source workload during live data transfers.</p> <p>We recommend that all the non-VSS compliant services or antivirus are stopped temporarily on the source while the VSS snapshot is being captured on the source. Select the Windows services that you want to be temporarily stopped on the source workload while the VSS snapshot is being captured on the source. These services are restored as soon as the VSS snapshot creation completes.</p> <p>To modify the settings:</p> <ol style="list-style-type: none">1. Select Stopped next to the service to be stopped for live data transfer.2. When you are done setting services to stop, click OK.
Linux Daemons (Target)	<p>Specify the start states for daemons on the target VM after cutover.</p> <p>To modify the settings:</p> <ol style="list-style-type: none">1. Click the Run Level column for the daemon, then select from run levels 0 through 6 and Boot (B), then click OK.2. When you are done setting daemon start states, click OK.
Live Transfer Daemons (Source)	<p>Specify the daemons to stop on the source workload during live data transfers.</p> <p>To modify the settings:</p> <ol style="list-style-type: none">1. Select Stopped next to the daemon to be stopped for live data transfer.2. When you are done setting daemons to stop, click OK.

- 10 In the Drive Configuration section of the Migration Job window, configure the following settings:

Setting Name	Description
Drive Configuration	
Hard Drives	Specify drive and volume configurations to be migrated.
Disks	Specify the path to the hard disk on the target virtual machine.
Volumes	Select volumes to be included in the target for migration.
NTFS Cluster Size	(For File-Based Windows Workloads) Specify the cluster size for the NTFS volume. For information about the default cluster size for an NTFS volume, see the Microsoft Support KB Article 140365 .

Setting Name	Description
Non-volume Storage	(For Linux Workloads) Specify a non-volume storage, such as a swap partition, that is associated with the source workload. This storage is re-created in the migrated workload.
Disks For Volume Groups	(For Linux Workloads) Specify the datastore name and the path where the virtual disk must be created on the target machine. You can choose to retain the path specified by default.
Volume Groups	(For Linux Workloads) Specify the LVM volume groups to be migrated with the LVM logical volumes listed in the Converted Logical Volumes section of the settings.
Converted Logical Volumes	(For Linux Workloads) Specify one or more LVM logical volumes to be migrated for a Linux workload.

- 11 (Target VMs using X2P workflow) PlateSpin Migrate displays storage configuration options specific to the selected target and also provides access to advanced configuration options for some platforms. For information about host-specific configuration options, see:
 - ♦ [“Drive Configuration: VMware ESX”](#)
 - ♦ [“Drive Configuration: Hyper-V”](#)
- 12 In the Additional Items for Review section of the Migration Job window, review errors and messages about the workload configuration. You must resolve errors before you can submit the migration job.
- 13 Click **OK**.

36 Workload Migration with a PlateSpin Image

This section provides information about using the PlateSpin Image volume archiving feature (Windows only).

- ♦ [Section 36.1, “About PlateSpin Images,” on page 415](#)
- ♦ [Section 36.2, “Designating a PlateSpin Image Server,” on page 415](#)
- ♦ [Section 36.3, “Capturing a Workload to a PlateSpin Image,” on page 417](#)
- ♦ [Section 36.4, “Deploying a PlateSpin Image,” on page 418](#)
- ♦ [Section 36.5, “Managing PlateSpin Images,” on page 419](#)

36.1 About PlateSpin Images

One of three fundamental workload infrastructures supported by PlateSpin Migrate, a PlateSpin Image is an image of a supported Windows workload consisting of volume data along with configuration specifics of the source server’s hardware, operating system, and network identity.

Image configurations are maintained in an XML (`config.xml`) file with each image having one or more sets of associated volume data.

PlateSpin Images and the image server’s `config.xml` configuration file are stored on the designated PlateSpin Image Server host in the following directory:

```
..\Program Files\PlateSpin Image Server
```

In addition to volume data directly captured during an X2I migration, PlateSpin Images support existing or raw volume data.

Like peer-to-peer migrations, image deployment allows for key workload configuration options, such as those for managing the workload’s disk layout, volume sizes, network identity, and domain or workgroup affiliation.

36.2 Designating a PlateSpin Image Server

To work with PlateSpin Images, you must first designate a machine as an image server by installing the PlateSpin Image Server software on it. You can install a PlateSpin Image Server instance either on a dedicated host or on your PlateSpin Server host. For information about storing PlateSpin Images on a NAS (Network Attached Storage) device or a remote share, see [KB Article 7921021 \(https://www.netiq.com/support/kb/doc.php?id=7921021\)](https://www.netiq.com/support/kb/doc.php?id=7921021).

NOTE: Although co-location of the PlateSpin Server with a PlateSpin Image Server instance on the same host is supported, the recommended setup is to install a PlateSpin Image Server on a dedicated host, which simplifies troubleshooting related to imaging functionality.

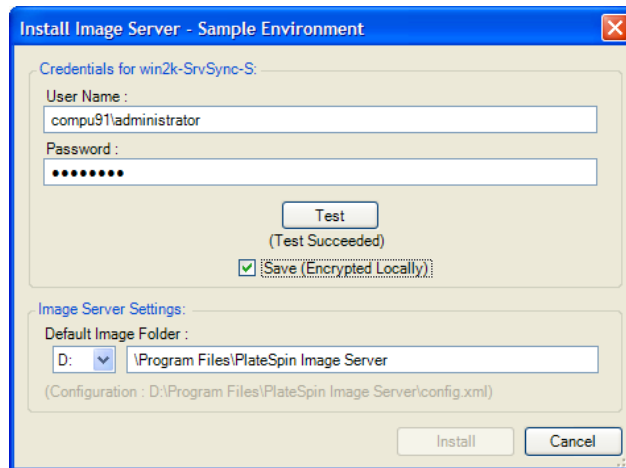
Dedicated PlateSpin Image Server hosts must meet the following requirements:

Table 36-1 PlateSpin Image Server Host Requirements


Requirement	Details
Operating System	Any of the following, running on dedicated hardware or in a virtual machine: <ul style="list-style-type: none"> ♦ Microsoft Windows Server 2012 R2 ♦ Microsoft Windows Server 2012 ♦ Microsoft Windows Server 2008 R2
Disk Space	Minimum 100 MB for basic controller software. Additional space requirements depend on the number and size of workload images that you intend to store on a given image server.
Software	♦ Microsoft .NET Framework 3.5 SP1

To designate a machine as a PlateSpin Image Server:

- 1 Discover the system you plan to designate as a PlateSpin Image Server.
- 2 In the Servers view, right-click the discovered server and select **Install Image Server**.




- 3 Provide administrator credentials for the selected host and specify the desired directory for image files.
- 4 Click **Install**.

PlateSpin Image Server software installs a controller on the selected host and configures it to run as a PlateSpin Image Server. On completion, the Servers view lists a new PlateSpin Migrate item: 

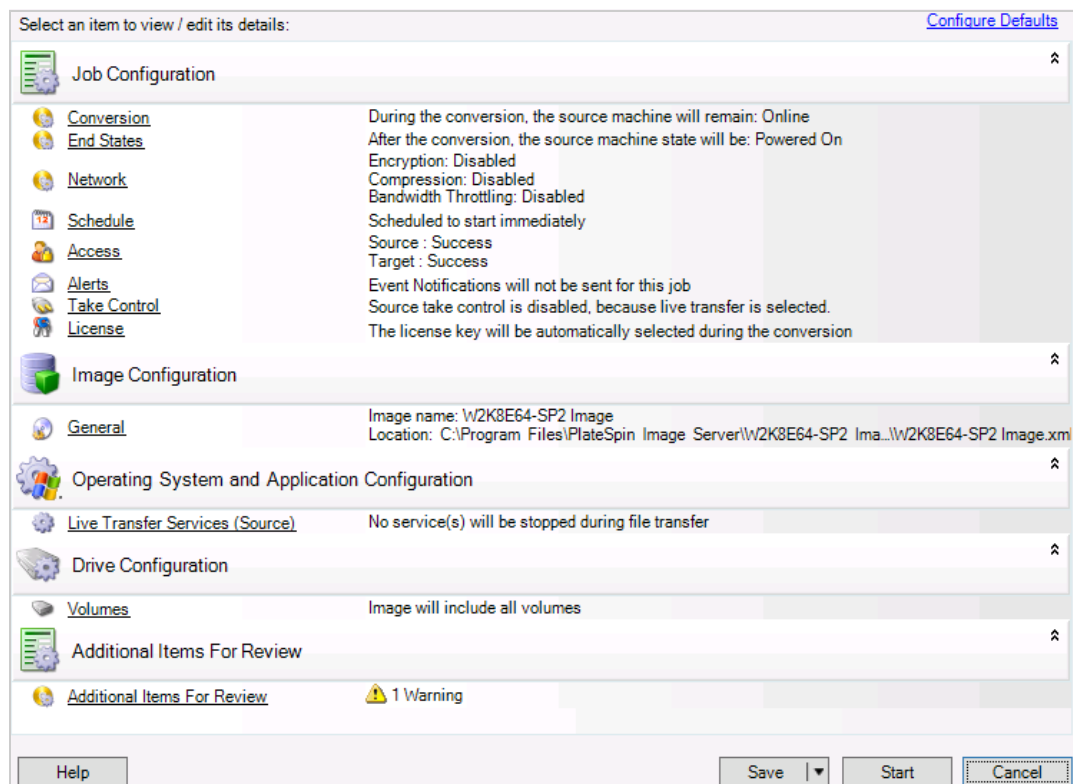
36.3 Capturing a Workload to a PlateSpin Image

Use this procedure to capture a physical or virtual workload as a PlateSpin Image.

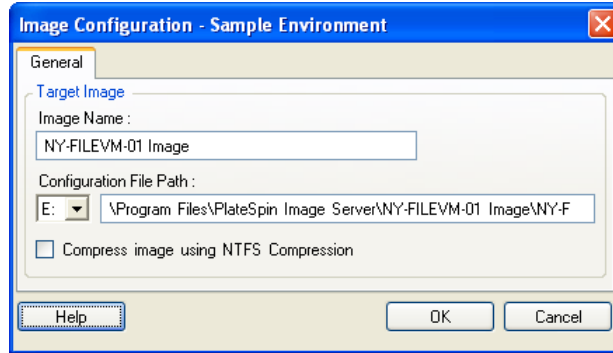
- 1 Discover, or refresh the details of, your source workload and your PlateSpin Image Server.
- 2 Start a new Capture Image job by using one of the following methods:
 - ♦ In the Servers view, right-click the source workload, then select **Capture Image**. In the Action window, select the source workload and the target image server.
 - ♦ In the Tasks pane, click **Capture Image**. In the Action window, select the source workload and the target image server.
 - ♦ In the Servers view, drag the source workload and drop it on the image server. If you configured PlateSpin Migrate to bypass the Action window on drag-and-drop, the Create Image dialog box prompts you to specify whether you want to create a new image or use existing volume data.




- 3 Select **Create Image**, then click **OK**.



- 4 Specify the required settings for the migration job by clicking the links in each category:
 - ♦ **Job Configuration:** Specify the required transfer method and operational continuity settings for your source and target (**General**), scheduling options (**Schedule**), source and target credentials (**Credentials**), job status and progress notification options, temporary network settings (**Take Control**), and the required license key to use (**License Key**).
 - ♦ **Image Configuration:** Specify the image name, the path to the location where the you want the image to be stored, and whether or not to use NTFS compression (under Image Configuration, click **General**).



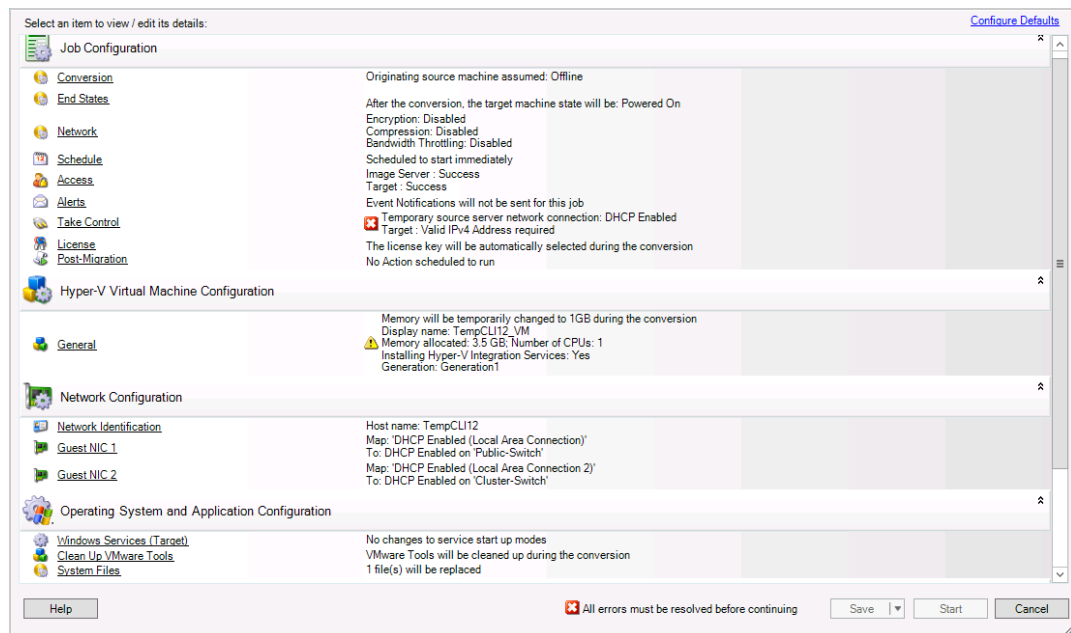
- ♦ **Operating System and Application Configuration:** If you selected the Live Transfer method, specify how you want PlateSpin Migrate to handle operating system and application services on your source (**Live Transfer Services**).
- ♦ **Drive Configuration:** Select the volumes that you want PlateSpin Migrate to include in the image and specify the path for the package file (under Drive Configuration, click **Volumes**).



36.4 Deploying a PlateSpin Image

Use this procedure to deploy a PlateSpin Image on a supported physical machine or virtualization platform.

- 1 Drag and drop the required PlateSpin Image to a discovered target physical machine or VM host.



2 Specify the required settings for the migration job by clicking the links in each category.

Migration jobs are auto-configured to create the target machine with the same settings as the source server. Depending on the objectives of the migration, you can:

- ♦ Modify the **Network Identification** settings to configure the host name and domain/workgroup registration of the target machine.
- ♦ Modify the **Guest NIC** settings to configure the TCP/IP properties for the network adapters on the target machine.
- ♦ Modify the **Drive Configuration** settings to select the volumes to copy during the migration.

3 If the intended target is a virtual machine, specify the required virtual machine parameters and select the options you require, such as memory allocation, or automatic installation of VMware Tools or VMAdditions.

4 Review and address errors and warnings.

5 Click **Start** to deploy the image.

36.5 Managing PlateSpin Images

- ♦ [Section 36.5.1, “Moving Images from One PlateSpin Image Server to Another,” on page 420](#)
- ♦ [Section 36.5.2, “Automating Image Operations,” on page 420](#)
- ♦ [Section 36.5.3, “Browsing and Extracting Image Files,” on page 420](#)

36.5.1 Moving Images from One PlateSpin Image Server to Another

- 1 Copy the image directory from the old PlateSpin Image Server host's file system to a location on the new PlateSpin Image Server host.
- 2 Update the new PlateSpin Image Server's `config.xml` file to identify the path to and the name of the image that was moved from the old PlateSpin Image Server.
- 3 Refresh the new image server's PlateSpin Migrate Clients details in the 's Servers view.

For more information, see [KB Article 7920189 \(https://www.netiq.com/support/kb/doc.php?id=7920189\)](https://www.netiq.com/support/kb/doc.php?id=7920189).

36.5.2 Automating Image Operations

You can use the ImageOperations command line utility, included with PlateSpin Migrate, to automate several tasks related to images, such as regularly moving multiple base images, along with related increments, between PlateSpin Image Servers.

The utility provides the capability to automate the following operations:

- ♦ **Register:** Associate an image or image increments with a specified image server.
- ♦ **Unregister:** Disassociate a registered image from a specified image server.
- ♦ **Gather:** Assemble a package of a PlateSpin Image and its volumes into a specified subdirectory.

To use the ImageOperations command line utility:

- 1 On your PlateSpin Image Server host, open a command interpreter (`cmd.exe ..\Program Files\PlateSpin Image Server`) and change the current directory to `\ImageOperations`.
- 2 Type `ImageOperations` followed by the required command and parameters, then press Enter. For command syntax and usage details, type `ImageOperations`, then press Enter.
- 3 When you have finished, refresh the image server's details in the Servers view.

36.5.3 Browsing and Extracting Image Files

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

To do this, you can use the PlateSpin Image Browser utility, which enables you to browse, search, sort, and extract files from different sources:

- ♦ An image file
- ♦ A specific image increment file

You can work with both base images and image increments by loading different files:

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

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

- ♦ “Starting the Image Browser and Loading Image Files” on page 421
- ♦ “Sorting and Searching Items in the Image Browser Interface” on page 421
- ♦ “Extracting Items” on page 422
- ♦ “Browsing and Extracting Image Files at the Command Line” on page 422

Starting the Image Browser and Loading Image Files

- 1 Start the Image Browser program (`ImageBrowser.exe`), located in one of the following directories:

- ♦ On your PlateSpin Server host:

`..\PlateSpin Migrate Server\bin\ImageOperations`

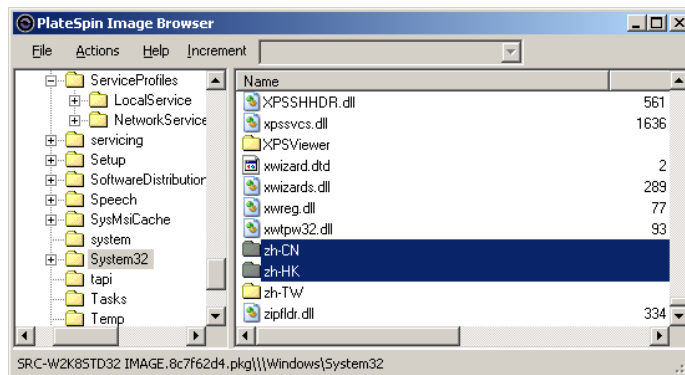
- ♦ On your PlateSpin Image Server host:

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

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

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

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



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

Sorting and Searching Items in the Image Browser Interface

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

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

To search for an item:

- 1 Load the required image or image increment. See [“Starting the Image Browser and Loading Image Files” on page 421](#).

- 2 In the left pane, select a volume or a subdirectory.

- 3 On the **Actions** menu, click **Search**.

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

The Image Browser Search window opens.

- 4 Specify the name of the file you are searching. If you are using a regular expression, select the corresponding option.

- 5 Click **Search**.

The results are shown in the right pane.

Extracting Items

- 1 Load the required image or image increment. See [“Starting the Image Browser and Loading Image Files” on page 421](#).

- 2 Locate and select the required file or directory. You can select multiple files and directories in the right pane.

- 3 On the **Actions** menu, click **Extract**.

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

The Browse for Folder dialog box opens.

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

The selected items are extracted to the specified destination.

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

Browsing and Extracting Image Files at the Command Line

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

To start the utility:

- 1 On your PlateSpin Image Server host, open a command interpreter (`cmd.exe ..\Program Files\PlateSpin Image Server`) and change the current directory to `\ImageOperations`.

- 2 At the command prompt, type `ImageBrowser.Console`, then press Enter.

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

37 Synchronizing Workloads with Server Sync

The Server Sync feature enables you to reduce the scope of data that is transferred from your source to your target to just data that is different between a source and a target, effectively synchronizing their volume contents.

For example, when setting up a job for a workload migration operation, you can choose to update an existing physical or virtual machine to match the state of your source workload without transferring volume data in its entirety. PlateSpin Migrate compares the target physical or virtual workload with the selected source and transfers only data that is different between the two, overwriting files on the target with those on the source workload.

Server Sync is useful in situations where the size of volume data or network conditions are prohibitive for a direct source-to-target virtualization over the network.

- ♦ [Section 37.1, “Server Sync to a Virtual Target,” on page 423](#)
- ♦ [Section 37.2, “Server Sync to a Physical Target,” on page 426](#)
- ♦ [Section 37.3, “Server Sync to Amazon Web Services,” on page 426](#)
- ♦ [Section 37.4, “Server Sync Volume Mapping,” on page 427](#)

37.1 Server Sync to a Virtual Target

- 1 Discover your source workload.

See [“Discovering Details for Source Workloads” on page 213](#).

- 2 Create a target virtual machine by using one of the following methods:

- ♦ Do an initial migration of your workload to a virtual machine. See [Chapter 26, “Configuration Essentials,” on page 291](#).
- OR -
- ♦ Using your virtualization platform’s native interface, manually install a virtual machine with the same operating system profile as that of your source.

NOTE: When you are creating a virtual target for Server Sync, you should also manually install the appropriate virtualization enhancement tools, such as VMware Tools or XenServer Tools.

- OR -

- ♦ (Windows only) Capture your workload to a PlateSpin Image, and deploy it to a virtual machine on your virtualization platform. See [“Capturing a Workload to a PlateSpin Image” on page 417](#).
- 3 (Conditional) Because the Server Sync option is disabled for a Hyper-V VM, it is necessary to use the following steps, as documented in [KB 7010748 \(https://www.netiq.com/support/kb/doc.php?id=7010748\)](https://www.netiq.com/support/kb/doc.php?id=7010748):

NOTE: Hyper-V automated server sync is available.

3a After booting the target VM with the LRD ISO (`bootofx.x2p.iso`) wait for the Migrate Server URL prompt, Then press Alt+F7 to launch the debug console.

3b From the debug console, run the following command to determine which devices are `/`, `/boot` and `swap`:

```
fdisk -l
```

3c Using the information obtained from the debug console, mount the appropriate devices as under:

```
mount /dev/%root device% /  
mount /dev/%boot device% /boot
```

3d Press Alt+F1 to switch to the server command line.

3e At the command line, provide the required information at each individual prompt:

- ♦ **PlateSpin Server:** Use the following format:

```
http://server_host/platespinmigrate
```

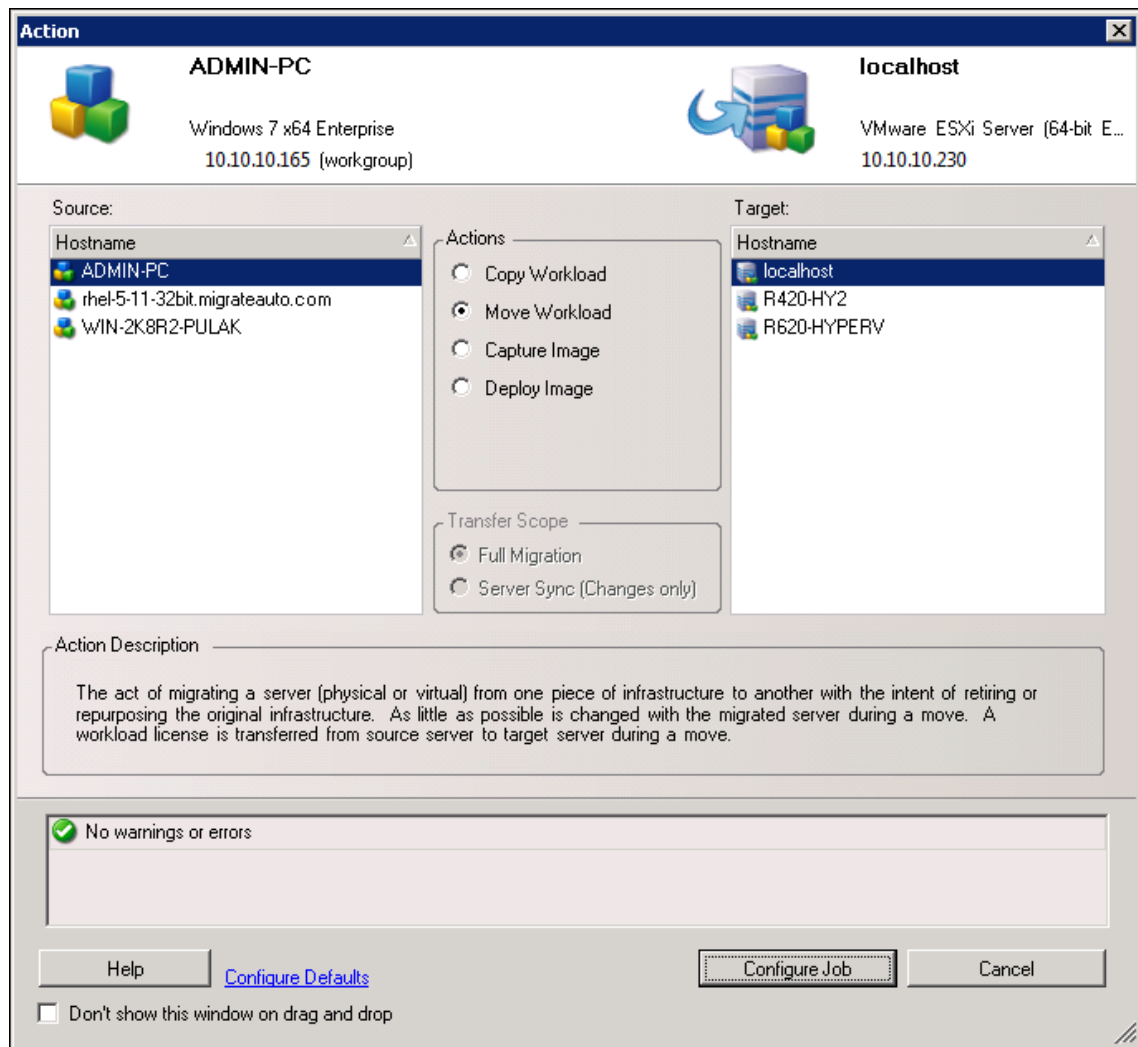
Replace *server_host* with the actual PlateSpin Server host's name or IP address.

- ♦ **Credentials (User Name/Password):** Enter the name of an administrator-level user on the PlateSpin Server host, including the domain or machine name. For example: *domain\username*, or *localhostAdministrator*. Provide a valid password for the specified user.
- ♦ **Network Card:** Select the network card that is active, then either enter a temporary static IP address for this card or press Enter to use a DHCP server.
- ♦ **Temporary hostname:** Provide a temporary VM name for PlateSpin Migrate Client to use to list the newly registered VM. The workload's target host name you select in the migration job overwrites this name.
- ♦ **SSL encryption:** If your PlateSpin Migrate is installed on a host with SSL encryption enabled, enter *Yes*. If not, enter *No*.
- ♦ **PlateSpin Migrate Network:** Unless you have defined your own PlateSpin Migrate Network in PlateSpin Migrate Client, press Enter. If you are working with a non-default PlateSpin Migrate Network, type its name, then press Enter.

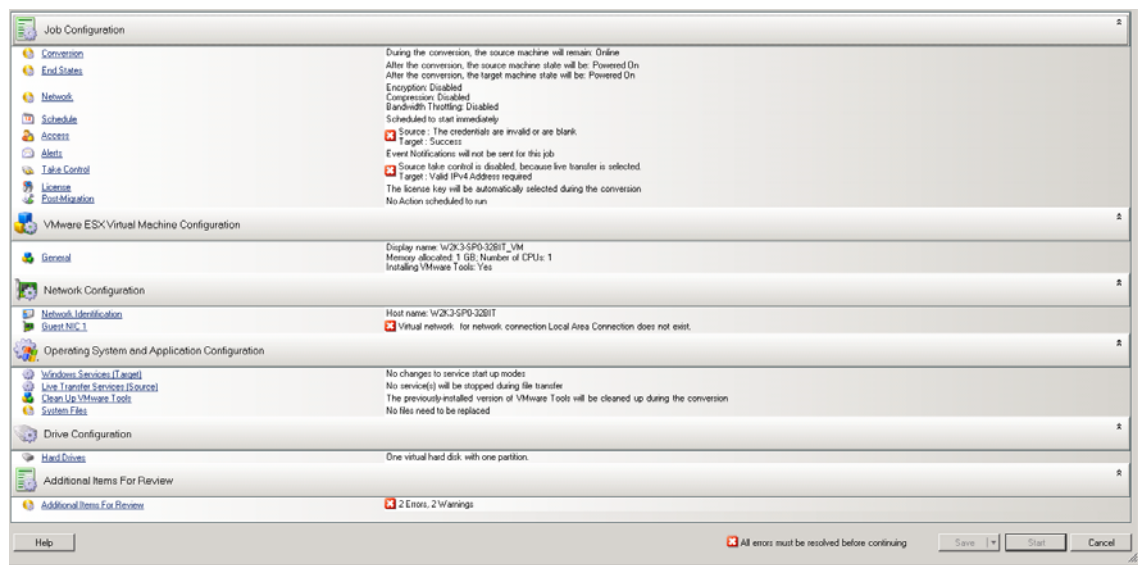
A controller on your target virtual machine communicates with PlateSpin Server and registers the virtual machine as a physical target for a migration job.

- 4** In the Servers view, drag your source workload and drop it on the required target (Server Sync target or discovered physical machine under control).

The system validates the selected source and target and, if it detects matching operating systems on them, provides you with two **Transfer Scope** options, **Full Migration** and **Server Sync**:



5 Select the **Server Sync** option, then click **Configure Job**.



- 6 In the job configuration window, specify the parameters of the job as dictated by the purpose of the operation, address any warnings and errors, and ensure that you map the required volumes on the source to those on the target (see [“Server Sync Volume Mapping” on page 427](#)).

For target machine on a Hyper-V server, enable the **VLAN ID** option to specify the virtual network ID to be used on the target machine. If you do not specify this ID, then the virtual network ID of the source machine is used by default.

When you have finished, click **Start**.

PlateSpin Migrate starts the job and lists it in the Jobs view.

37.2 Server Sync to a Physical Target

- 1 Discover your source workload.

See [“Discovering Details for Source Workloads” on page 213](#).

- 2 Discover your physical target by using the appropriate PlateSpin ISO boot image.

See [“Registering and Discovering Details for Target Physical Machines with PlateSpin ISO” on page 206](#).

- 3 In the Servers view, drag your source workload and drop it on the required target (Server Sync target or discovered physical machine under control).

The system validates the selected source and target and, if it detects matching operating systems on them, it provides you with two **Transfer Scope** options, **Full Migration** and **Server Sync**, similar to the [“Server Sync to a Virtual Target” on page 423](#) (see [Step 4](#)).

- 4 Select the **Server Sync** option, then click **Configure Job**.

- 5 In the job configuration window, specify the parameters of the job as dictated by the purpose of the operation, address any warnings and errors, and ensure that you map the required volumes on the source to those on the target.

- 6 When you have finished, click **Start**.

PlateSpin Migrate starts the job and lists it in the Jobs view.

37.3 Server Sync to Amazon Web Services

- 1 Launch the PlateSpin Migrate Client.

- 2 Discover your source workload. See [Section 20.6, “Discovering Details for Source Workloads,” on page 213](#).

- 3 Right-click the discovered source workload and click **Prepare AWS Target**.

- 4 Perform the following steps:

4a To log in to your AWS account for the first time, specify the following:

- ♦ **AWS Access Key ID and AWS Secret Access Key:** The two parts of an access key that is a valid set of security credentials required to access AWS account.
- ♦ **AWS Region:** The geographic area to which your AWS account has access.

During subsequent logins, these options are automatically populated with the values that you specified during your earlier login.

4b Click **Update Region List** to update the supported AWS Regions.

4c Click **Login**.

- 4d** (Conditional) If you are logging in for the first time, the Replication Environment Details dialog displays. Specify the following and then click **Save**:
- ♦ **Linux AMI ID:** Specify the ID of the PlateSpin Replication Environment - Linux AMI.
 - ♦ **Windows AMI ID:** Specify the ID of the PlateSpin Replication Environment - Windows or the PlateSpin Replication Environment - Windows(BYOL) AMI depending on whether you want AWS to activate the OS license or you want to bring your own Microsoft license.
 - ♦ **Snapshot ID:** The snapshot ID of the PRE AMI that you noted down when you created a snapshot from the private instance of AMI. See [“To create a snapshot using the PRE AMI:” on page 143](#).
 - ♦ **Keypair Name:** Specify the name of the key pair associated with the AMI snapshot.
 - ♦ **Keypair File Path:** Specify the location where you have stored the key pair (.pem) file.
- 4e** (Conditional) If you are logging in for the first time, the Migrate Server Details dialog displays. Specify the Migrate server IP and the credentials to register the target instance with the Migrate server, then click **Save**.
- 4f** In the Prepare AWS Target dialog, do the following:
- 4f1** Select **ServerSync**.
 - 4f2** Select the appropriate target instance.
 - 4f3** Specify the host name to register the target with the Migrate server.
- The Target Instance Storage section displays all the disks available on the selected target. The disks are sorted based on the device name.
- 4g** (Optional) Click the settings icon to modify the replication details or the Migrate server details.
- 4h** Click **Prepare**.
- 4i** On completion of the registration job, you are prompted to view the logs. To view the logs, click **Yes**. Click **Close**.
- 5** In the PlateSpin Migrate Client, the registered target instance is listed with the host name you specified in [Step 4f3](#). Drag the source workload to this target to start a migration job.
- The system validates the selected source and target and, if it detects matching operating systems on them, it provides you with two **Transfer Scope** options, **Full Migration** and **Server Sync**.
- 6** Select the **Server Sync** option, then click **Configure Job**.
- 7** In the job configuration window, specify the parameters of the job as dictated by the purpose of the operation, address any warnings and errors, and ensure that you map the required volumes on the source to those on the target.
- 8** Click **Start**. PlateSpin Migrate starts the job and lists it in the Jobs view.

37.4 Server Sync Volume Mapping

When you are using Server Sync to synchronize two Windows or Linux workloads, PlateSpin Migrate Client provides you with the capability to specify the required mapping between source volumes and existing volumes on the target. See [“Synchronizing Workloads with Server Sync” on page 423](#).

To access volume configuration options in a Server Sync job:

- 1** In the Jobs view, select the required workload.

2 In the **Drive Configuration** section of the Migration Job window, do one of the following:

- ♦ **On Windows:** Click **Volume Mapping**.
- ♦ **On Linux:** Click **Drives and Volumes**.

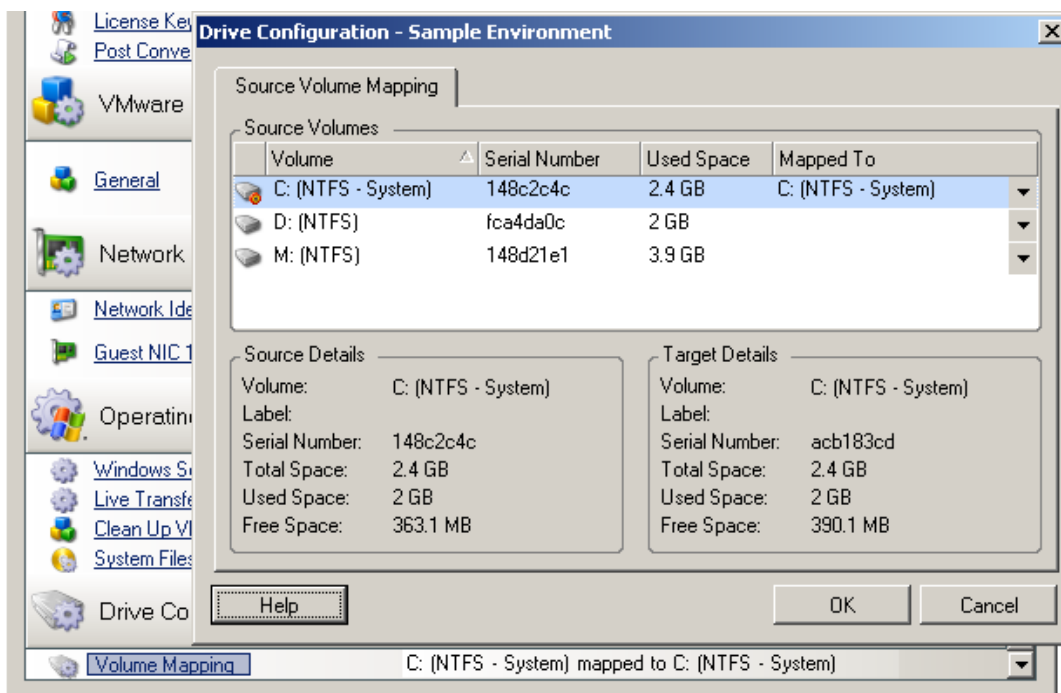
3 Configure the Server Sync volume configuration options.

The following topics provide information about Server Sync volume configuration options specific to Windows and Linux workloads.

- ♦ [Section 37.4.1, “Server Sync Volume Configuration \(Windows\),” on page 428](#)
- ♦ [Section 37.4.2, “Server Sync Volume Configuration \(Linux\),” on page 429](#)

37.4.1 Server Sync Volume Configuration (Windows)


A Server Sync job for Windows workloads provides detailed drive and volume information for both the source and the target, and enables you to specify the required mapping.



Mapped To: Map each volume on the source to an existing volume on the target.

37.4.2 Server Sync Volume Configuration (Linux)

A Server Sync job for Linux workloads provides detailed mount point and volume information for both the source and the target, and enables you to specify the required mapping.



Mapped To: Map each volume on the source to an existing volume on the target.

VI

Executing Migrations

After you configure migration settings for the workload, you are ready execute the migration. Ensure that the target VMs are prepared for migration, then begin replicating data to the target. You can monitor the health of migration jobs and generate reports about them.

- ♦ [Chapter 38, “Executing Workload Migrations,” on page 433](#)
- ♦ [Chapter 39, “Generating Reports,” on page 441](#)
- ♦ [Chapter 40, “Post-Migration Cleanup of Source Workloads,” on page 445](#)
- ♦ [Appendix H, “Troubleshooting PlateSpin Migrate,” on page 447](#)

38 Executing Workload Migrations

After you discover and configure workloads for migration, you execute and monitor the migration by performing the migration tasks described in this section. Use the PlateSpin Migrate Web Interface or PlateSpin Migrate Client as appropriate for the migration types and target containers. See [“Migration Tasks Matrix for PlateSpin Migrate Client and PlateSpin Migrate Web Interface” on page 76](#).

- ♦ [Section 38.1, “Preparing a Migration,” on page 433](#)
- ♦ [Section 38.2, “Starting Migration Execution \(First Replication\),” on page 434](#)
- ♦ [Section 38.3, “Scheduling Migration Execution \(First Replication\),” on page 435](#)
- ♦ [Section 38.4, “Starting Incremental Replications,” on page 436](#)
- ♦ [Section 38.5, “Scheduling Incremental Replications,” on page 437](#)
- ♦ [Section 38.6, “Viewing Properties for an In-Progress or Completed Migration,” on page 438](#)
- ♦ [Section 38.7, “Canceling an In-Progress Migration,” on page 438](#)
- ♦ [Section 38.8, “Restarting or Shutting Down the Source Workload,” on page 439](#)


38.1 Preparing a Migration


After you configure a workload for migration, PlateSpin Migrate uses the migration settings to install any required data transfer software on the source workload and create a target workload on the target container.

- ♦ [Section 38.1.1, “Using the Migrate Client,” on page 433](#)
- ♦ [Section 38.1.2, “Using the Migrate Web Interface,” on page 434](#)

38.1.1 Using the Migrate Client


When you start a migration job from the PlateSpin Migrate Client, PlateSpin Migrate validates the job type, the source, the target, and the selected parameters, and might generate errors and warnings.

 Error markers show configurations that you need to change before the migration job can start.

 Warning markers alert you to settings that should be confirmed prior to starting the migration.

In a default PlateSpin Migrate configuration, validation messages display at the bottom of the Action window. However, If you have configured PlateSpin Migrate to bypass the Action window on drag-and-drop, errors and warnings are displayed in a separate window:

Figure 38-1 Migration Validation Window



To force this window to open only on errors, select **Show me only when validation errors occur**.

38.1.2 Using the Migrate Web Interface

To immediately prepare the workload for migration:

- 1 On the Edit Migration Details page, click **Save and Prepare**.

To prepare a preconfigured workload for migration:

- 1 On the Workloads page, select the preconfigured workload you want to migrate.
- 2 Click **Prepare Migration**.

38.2 Starting Migration Execution (First Replication)

After the migration preparation completes successfully, the migration is ready for execution. Execution begins with the first replication. The first replication is a full replication with a Full Replication contract type or an incremental data synchronization for a pre-existing target workload with an Incremental Replication contract type.

The first replication is unscheduled by default. You can manually start the first replication. You can alternatively schedule the date and time to run the first replication. See [Section 38.3, “Scheduling Migration Execution \(First Replication\),” on page 435](#).

NOTE: You must prepare the source and target workload prior to the manual start. The full replication cannot run unless the target workload exists and the workload preparation is complete. See [Section 38.1, “Preparing a Migration,” on page 433](#).

- ♦ [Section 38.2.1, “Using the Migrate Client,” on page 435](#)
- ♦ [Section 38.2.2, “Using the Migrate Web Interface,” on page 435](#)

38.2.1 Using the Migrate Client

To manually start the first replication:

- 1 In the Jobs view, locate the prepared workload that you want to migrate.
- 2 Right-click the job and select **Start**.

PlateSpin Migrate starts the first full replication for the workload.

38.2.2 Using the Migrate Web Interface

To manually start the first replication:

- 1 On the Workloads page, select the prepared workload that you want to migrate.
- 2 Click **Run Migration**.
- 3 On the Workload Commands page, do one of the following depending on the migration contract type you configured for the workload:
 - ♦ **Full Replication:** Select **Full Replication** as the replication method.
 - ♦ **Incremental Replication:** Select **Incremental Replication** as the replication method.
- 4 (Optional) Set the following options as appropriate if you want to cut over the workload after a successful manual replication:
 - ♦ Run cutover after successful replication
 - ♦ Shut down source after cutover
 - ♦ Shut down target after cutover
- 5 Click **Execute**.

PlateSpin Migrate starts the first replication for the workload.

38.3 Scheduling Migration Execution (First Replication)

After the migration preparation completes successfully, the migration is ready for execution. Execution begins with the first replication, which might be a full replication or a data synchronization for a pre-existing target workload.

The default schedule setting is None. The first replication is unscheduled. You can schedule the start date and time to run the first replication. You can alternatively start the first replication manually. See [Section 38.2, “Starting Migration Execution \(First Replication\),” on page 434](#).

The first replication for a scheduled migration execution is a one-time event, but the run is attempted daily as scheduled until the first replication begins and completes successfully.


NOTE: You must prepare the workload prior to the scheduled time or the manual start. The first replication cannot run unless the target workload exists and the workload preparation is complete. If they are not ready, Migrate skips the scheduled replication and retries it at the scheduled time on the next day.

- ♦ [Section 38.3.1, “Using the Migrate Client,” on page 436](#)
- ♦ [Section 38.3.2, “Using the Migrate Web Interface,” on page 436](#)

38.3.1 Using the Migrate Client

To modify the start date and time for the first replication:

- 1 In the Jobs view, locate the required job.
- 2 Right-click the job and select **Change Start Time** to open the Change Job Start Time dialog box.



- 3 Specify the required start date and time, then click **OK**.
PlateSpin Migrate reschedules the job and executes it at the specified time.

38.3.2 Using the Migrate Web Interface

To modify the start date and time for the first replication:

- 1 On the Workloads page, locate and click the workload.
- 2 On the Migration Details page, click **Edit**.
- 3 On the Edit Migration Details page, go to **Schedule Settings > Full Replication**, then click **Edit**.
- 4 Click **Start**, then set the date and time when you want to start the first full replication.
You can type the date (dd/mm/yyyy) or click the Calendar icon to select the date. The default run time is 12:00:00 a.m. (hh:mm:ss a.m. or p.m.).
- 5 Click **Close** to return to the Edit Migration Details page, then click **Save**.

38.4 Starting Incremental Replications

After the first replication completes successfully, you can start each incremental replication manually. You can alternatively schedule the time and pattern to run incremental replications that occur after the first replication. See [Section 38.5, “Scheduling Incremental Replications,” on page 437](#).

- ♦ [Section 38.4.1, “Using the Migrate Web Interface,” on page 436](#)

38.4.1 Using the Migrate Web Interface

To start an incremental replication manually:

- 1 On the Workloads page, locate and select the workload.
- 2 Click **Run Migration**.
- 3 On the Workload Commands page, select **Incremental Replication** as the replication method.

- 4 (Optional) Set the following options as appropriate if you want to cut over the workload after a successful manual replication:
 - ♦ Run cutover after successful replication
 - ♦ Shut down source after cutover
 - ♦ Shut down target after cutover
- 5 Click **Execute**.

PlateSpin Migrate starts the incremental replication for the workload.

38.5 Scheduling Incremental Replications

After you configure and save a workload migration, you can modify the time and pattern to run incremental replications that occur after the first replication. You can alternatively start each incremental replication manually. See [Section 38.4, “Starting Incremental Replications,” on page 436](#).

NOTE:

- ♦ Scheduled incremental replications are skipped until the first full replication is complete.
- ♦ Scheduled incremental replications take place for a maximum period of 60 days from the time that the scheduled incremental replication runs begin.

-
- ♦ [Section 38.5.1, “Using the Migrate Web Interface,” on page 437](#)

38.5.1 Using the Migrate Web Interface

To schedule the incremental replication recurrence time and pattern:

- 1 On the Workloads page, locate and click the workload.
- 2 On the Migration Details page, click **Edit**.
- 3 On the Edit Migration Details page, go to **Schedule Settings > Incremental Recurrence**, then click **Edit**.

The default Incremental recurrence setting is None. The incremental replications are unscheduled.
- 4 For **Begin the recurrence schedule**, set the date and time when you want to begin the scheduled incremental replications.

You can type the date (dd/mm/yyyy) or click the Calendar icon to select the date. The default run time is 12:00:00 a.m. (hh:mm:ss a.m. or p.m.).
- 5 For **Recurrence run setting**, set the pattern to follow for scheduled incremental replications:
 - ♦ **Daily:** The replication takes place on the specified daily intervals or on weekdays every week for a period of 60 days from the time the replication starts.
 - ♦ **Weekly:** The replication takes place at specified intervals for a period of 8 weeks from the time the replication starts.
 - ♦ **Monthly:** The replication takes place at specified intervals for a period of 2 months from the time the replication starts.
- 6 Click **Close** to return to the Edit Migration Details page, then click **Save**.

38.6 Viewing Properties for an In-Progress or Completed Migration

After you add a workload to PlateSpin Migrate, the Configuration page displays the properties of the workload's migration configuration throughout the migration lifecycle.

- ♦ [Section 38.6.1, “Using the Migrate Client,” on page 438](#)
- ♦ [Section 38.6.2, “Using the Migrate Web Interface,” on page 438](#)

38.6.1 Using the Migrate Client

To view properties for a workload migration:

- 1 In the Jobs view, locate the required job.
- 2 Right-click the job and select **View**.
Migrate Client opens the job configuration window.
- 3 View the workload migration configuration parameters and settings in read-only mode.

38.6.2 Using the Migrate Web Interface

To view properties for a workload migration:

- 1 On the Workloads page, locate and click the workload.
Migrate Web Interface opens the Migration Details page.
- 2 View the workload migration configuration parameters and settings in read-only mode.

38.7 Canceling an In-Progress Migration

You might need to cancel an in-progress workload migration that has become non-responsive.

- ♦ [Section 38.7.1, “Using the Migrate Client,” on page 438](#)
- ♦ [Section 38.7.2, “Using the Migrate Web Interface,” on page 438](#)

38.7.1 Using the Migrate Client

- 1 In the Jobs view, locate the required job.
- 2 Right-click the job and select **Abort**.

38.7.2 Using the Migrate Web Interface

To view properties for a workload migration:

- 1 On the Workloads page, locate and click the stalled workload.
- 2 View the replication or cutover status.
- 3 Click **Abort**.

38.8 Restarting or Shutting Down the Source Workload

PlateSpin Migrate Client enables you to restart or shut down a source workload if the migration job is not active.

To shut down or restart the source workload from the Migrate Client:

- 1 In the Jobs view, locate the required job.
- 2 Right-click the job and select **Restart Source** or **Shutdown Source** as applicable.

To automate the startup state of source and target workloads, specify the required post-migration state in your migration job. See [“Post-Cutover End States for Source and Target Workloads” on page 313](#).

39 Generating Reports

You can generate reports about discovered workloads and the workload migrations by using PlateSpin Migrate Client or PlateSpin Migrate Web Interface. For information about generating a Licensing Report, see [“License Key Management with Migrate Client” on page 96](#).

- ♦ [Section 39.1, “Generating Workload and Workload Migration Reports,” on page 441](#)
- ♦ [Section 39.2, “Generating Diagnostic Reports,” on page 442](#)

39.1 Generating Workload and Workload Migration Reports


You can generate detailed reports of running and completed migration jobs. A migration report records the tasks performed during the job.

- ♦ [Section 39.1.1, “Generate Reports using the Migrate Client,” on page 441](#)
- ♦ [Section 39.1.2, “Generate Reports using the Web Interface,” on page 442](#)

39.1.1 Generate Reports using the Migrate Client

To generate a job report:

- 1 In the Jobs view, locate the required job.
- 2 Right-click the job and select **Report**.



A Web browser window displays the requested report.

39.1.2 Generate Reports using the Web Interface

PlateSpin Migrate Web provides reports that give analytical insight into your workload migration contracts over time. See [Table 39-1](#) for a list of available reports. The reports open in the Web Interface. You can print a report using browser options, or export it to XML.

Table 39-1 Reports Available in PlateSpin Migrate Web Interface

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

To generate a report:

- 1 In your PlateSpin Migrate Interface, click **Reports**.
A list of the report types is displayed.
- 2 Click the name of the required report type.
- 3 Select one or more workloads for which you want to create the report.
- 4 Configure the time period for which you want to view the report.
- 5 Do one of the following:
 - ♦ Click **Printable View** to view the report in your web browser.
 - ♦ Click **Export to XML**, then save the XML file to your computer.

39.2 Generating Diagnostic Reports

- ♦ [Section 39.2.1, “Using the Migrate Client,” on page 442](#)
- ♦ [Section 39.2.2, “Using the Migrate Web Interface,” on page 443](#)

39.2.1 Using the Migrate Client

PlateSpin Migrate provides a tool that can produce a diagnostics report for any running or completed job.

To view a diagnostics report:

- 1 In the Jobs view, right-click the required job and select **Run Diagnostics**.

The diagnostics report is displayed in a browser window. This process might take a few moments.

The diagnostics report lists several statistics:

- ♦ All the operations involved in the job. Click any operation to view its XML representation.
- ♦ The status of each operation.
- ♦ The controller that ran the operation. Click the controller to view its XML representation, or click **Logs** to view its event log.

In addition, the report contains links to:

- ♦ The XML representations of the source machine, original target machine, and the target VM host.
- ♦ The root operation for the job, as well as a variety of logs and reports.

You can send diagnostics reports directly to Technical Support. Follow the instructions provided on the report.

39.2.2 Using the Migrate Web Interface

In the Migrate Web Interface, after you have executed a command, you can generate detailed diagnostic reports about the command's details.

- 1 Click **Command Details**, then click the **Generate** link in the lower right of the panel.
After a few moments, the page refreshes and displays a **Download** link above the **Generate** link.
- 2 Click **Download**.
A `.zip` file contains the comprehensive diagnostic information about the current command.
- 3 Save the file, then extract the diagnostics to view them.
- 4 Have the `.zip` file ready if you need to contact Technical Support.

40 Post-Migration Cleanup of Source Workloads

Source workloads might occasionally need a post-migration cleanup. For example, you might need to clean up a source workload of all PlateSpin software components after an unsuccessful migration.

- ♦ [Section 40.1, “Cleaning Up Windows Workloads,” on page 445](#)
- ♦ [Section 40.2, “Cleaning Up Linux Workloads,” on page 446](#)

40.1 Cleaning Up Windows Workloads

The following are instructions for cleaning up Windows workloads by component and use case.

Component	Use Case	Removal Instructions
File-based Transfer Component	All Migrations	At root level for each volume migrated, remove all files named <code>PlateSpinCatalog*.dat</code>
Workload discovery software	All migrations	<ol style="list-style-type: none">1. In the Servers view, undiscover the source (right-click, then select Undiscover).2. In the source workload's Windows directory:<ul style="list-style-type: none">♦ Remove all files named <code>machinediscovery*</code>.♦ Remove the subdirectory named <code>platespin</code>.
Controller software	All migrations	<ol style="list-style-type: none">1. In the Servers view, undiscover the source (right-click, then select Undiscover).2. Open a command prompt and change the current directory to:<ul style="list-style-type: none">♦ <code>\Program Files\platespin*</code> (32-bit systems)♦ <code>\Program Files (x86)\platespin</code> (64-bit systems)3. Run the following command: <code>ofxcontroller.exe /uninstall</code>4. Remove the <code>platespin*</code> directory

40.2 Cleaning Up Linux Workloads

The following are instructions for cleaning up Linux workloads by component and use case.

Component	Use Case	Removal Instructions
Controller software	Offline migrations	In the source workload's file system, under <code>/boot</code> , remove the <code>ofx</code> directory with its contents.
	All live migrations	<ul style="list-style-type: none">♦ Kill these processes:<ul style="list-style-type: none">♦ <code>kill -9 ofxcontrollerd</code>♦ <code>kill -9 ofxjobexec</code>♦ Remove the OFX controller rpm package: <code>rpm -e ofxcontrollerd</code>♦ In the source workload's file system, remove the <code>/usr/lib/ofx</code> directory with its contents.
Block-level data transfer software	All block-level migrations	<ol style="list-style-type: none">1. Check if the driver is active: <code>lsmod grep blkwatch</code> If the driver is still loaded in memory, the result should contain a line, similar to the following: <code>blkwatch_7616 70924 0</code>2. (Conditional) If the driver is still loaded, remove it from memory: <code>rmmmod blkwatch_7616</code>3. Remove the driver from the boot sequence: <code>blkconfig -u</code>4. Remove the driver files by deleting the following directory with its contents: <code>/lib/modules/[Kernel_Version]/Platespin</code>5. Delete the following file: <code>/etc/blkwatch.conf</code>
LVM snapshots	Block-level migrations using LVM snapshots	<ol style="list-style-type: none">1. In the Jobs view, generate a Job Report for the failed job, then note the name of the snapshot.2. Remove the snapshot device by using the following command: <code>lvremove <i>snapshot_name</i></code>



Troubleshooting PlateSpin Migrate

This section provides a series of topics about troubleshooting PlateSpin Migrate.

For information about common problems that occur during discovery of workloads or targets, see [Appendix C, “Troubleshooting Discovery,” on page 247](#).

- [Section H.1, “Migration of Workloads to Azure Cloud,” on page 447](#)
- [Section H.2, “Migration of Workloads to VMware,” on page 449](#)
- [Section H.3, “Peer-to-Peer Migrations \(Windows\),” on page 449](#)
- [Section H.4, “PlateSpin Images,” on page 450](#)
- [Section H.5, “Shrinking the PlateSpin Migrate Databases,” on page 451](#)
- [Section H.6, “Troubleshooting the Configuration Service,” on page 451](#)
- [Section H.7, “Validation Warning for Bandwidth Throttling,” on page 456](#)
- [Section H.8, “PlateSpin OFX Controller Does Not Start on a Virtual Machine Source,” on page 456](#)
- [Section H.9, “Target Windows Machine Becomes Unbootable on Second Boot,” on page 456](#)
- [Section H.10, “Two or More Volumes Have the Same Volume Serial Number,” on page 457](#)

H.1 Migration of Workloads to Azure Cloud

Use information in this section to help troubleshoot common problems that might occur during migration of workloads to Microsoft Azure Cloud.

- [Section H.1.1, “Assigning a Reserved IP Address to a Migrate Server in Azure,” on page 447](#)
- [Section H.1.2, “Computer Name Not Available for Azure VM in Azure Portal,” on page 448](#)
- [Section H.1.3, “Outbound Email Stuck after Migrating Microsoft Exchange Server 2016 to Azure Cloud,” on page 448](#)

H.1.1 Assigning a Reserved IP Address to a Migrate Server in Azure

In Azure, the Dynamic assignment method is the default setting for the public IP address. The IP address can change every time the server is stopped and started. You should modify the setting to use the Static assignment method. Using a reserved IP address ensures that Azure allocates and reserves an IP address for the life of the resource.

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

To apply a reserved IP address to an existing Migrate Server in Azure that has a dynamic IP address:

- 1 Specify **Static** as the assignment method for the public IP address of the Migrate Server resource:
 - 1a Go to the [Microsoft Azure Portal \(https://portal.azure.com\)](https://portal.azure.com) and log in to your Azure account.
 - 1b Open Resources, select the Migrate Server resource, then select **Stop**.
 - 1c In the information for the Migrate Server, select the public IP address.
 - 1d In the **Public IP Address Configuration** panel under **Settings**, select **Configuration**.
 - 1e Specify **Static** as the assignment method for the Public IP address.
 - 1f Click **Save**.

Azure allocates and reserves an IP address from a pool of its available IP addresses in the Azure location where you deploy the Migrate server.
 - 1g Start the Migrate Server resource.

Heartbeat communications for existing migration jobs will be broken until you modify the server IP address stored in the OFX Controller configuration file on the source workload.
- 2 For each source workload that has already been configured for migration on the Migrate Server, use the Migrate Agent to set the new IP address:

```
migrateagent.cli.exe config /setting=psserver:<new_ipaddress_or_dns_name>
```

The `psserver` option stops the OFX Controller (`ofxcontroller`) service, modifies the `OfxController.exe.config` file with the new address, and restarts the service. Heartbeat communications now work with the server's new IP address.

H.1.2 Computer Name Not Available for Azure VM in Azure Portal

After a successful cutover of a workload to Azure Cloud, the Azure Portal Web UI does not display the target computer name.

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

H.1.3 Outbound Email Stuck after Migrating Microsoft Exchange Server 2016 to Azure Cloud

Issue: After you migrate a Microsoft Exchange 2016 server to Microsoft Azure, the user's outgoing messages get stuck in the `Drafts` folder of their Microsoft Outlook application.

Fix: After you migrate a Microsoft Exchange Server workload to Microsoft Azure, ensure that you modify the Exchange internal and external DNS settings to use **Microsoft Hyper-V Network Adapter**. See [KB Article 7021909 \(https://www.netiq.com/support/kb/doc.php?id=7021909\)](https://www.netiq.com/support/kb/doc.php?id=7021909).

H.2 Migration of Workloads to VMware

Use information in this section to help troubleshoot common problems that might occur during migration of workloads to VMware.

- ♦ [Section H.2.1, “Outbound Email Stuck after Migrating Microsoft Exchange Server 2016 to VMware,” on page 449](#)

H.2.1 Outbound Email Stuck after Migrating Microsoft Exchange Server 2016 to VMware

Issue: After you migrate a Microsoft Exchange 2016 server to VMware, the users’ outgoing messages get stuck in their `Drafts` folder.

Fix: After you migrate a Microsoft Exchange Server workload to VMware, ensure that you modify the Exchange internal and external DNS settings to use **VMXNET 3**. See [KB Article 7021909 \(https://www.netiq.com/support/kb/doc.php?id=7021909\)](https://www.netiq.com/support/kb/doc.php?id=7021909).

H.3 Peer-to-Peer Migrations (Windows)

[Table H-1](#) provides information to help you troubleshoot common problems that might occur during Windows peer-to-peer migrations.

Table H-1 Common Issues and Solutions Related to Peer-to-Peer Migrations (Windows)

Problems or Messages	Solutions
One of the following errors displays during offline migration: <ul style="list-style-type: none">♦ Waiting for Controller to start (Failed)♦ Controller Connection Not Established♦ Controller Connection Broken♦ Unable to start the Heartbeat Service	<p>This indicates one of the following problems:</p> <ul style="list-style-type: none">♦ The network settings for the temporary IP addresses under Job Configuration > Advanced might not be configured properly.♦ There was a possible network outage that prevented the source/target machine from communicating with the PlateSpin Server.♦ The source/target machine was not able to fully boot into the pre-execution environment. <p>To diagnose the exact cause of failure, check the state of the system where the controller failed to start. Commands such as <code>ipconfig</code> and <code>ping</code> are available to verify basic network connectivity.</p>
File transfer hangs at 1% or progresses at a slow pace	<p>By default, a link type of AUTO is used on the source server during a migration. If the source server is connected to a switch port that is forced to 100/FULL, the Force Full Duplex option must be enabled when configuring the migration. If this option is set incorrectly, a duplex mismatch occurs on the network.</p>

Problems or Messages	Solutions
Unable to determine suitable boot partition	<p>When converting existing source servers, the boot volume must pass the following checks:</p> <ul style="list-style-type: none"> ♦ It must be on a basic disk ♦ It must have 175 MB of free space ♦ It must be a primary partition ♦ If any of these are not true for the system volume, the migration fails while attempting to take control of the source server.
Job remains in a Scheduled state for a long period and then changes to Recoverable error (all sub-steps display NotStarted status)	<p>There is a problem with the Operations Framework Controller on the PlateSpin Server. Use the Windows services plug-in to confirm that the Controller is running. See KB Article 7920862 (https://www.netiq.com/support/kb/doc.php?id=7920862) for other troubleshooting instructions.</p>
Troubleshooting failures at the Configuring Operating System stage (also applicable to Configure Target Machine or Configuring Virtual Machine migration steps)	<p>Generally, failures during the configuration step indicate that a time-out occurred when attempting to configure the target physical or virtual machine. Although the migration job appears to have failed, the overall migration is probably successful and the configuration service running on the target will likely continue its operations.</p> <p>KB Article 7920327 (https://www.netiq.com/support/kb/doc.php?id=7920327) contains a detailed troubleshooting checklist and lists information required if technical support is necessary.</p>
Live Transfer is unavailable	<p>Either an unsupported file system or operating system exists on the server.</p>

Related KB Articles:

ID	Description
7920862 (https://www.netiq.com/support/kb/doc.php?id=7920862)	ERRMSG: PlateSpin Migrate Job remains at a "Scheduled" or "Recoverable Error" state
7920810 (https://www.netiq.com/support/kb/doc.php?id=7920810)	INFO: Restore job stalls - "The configuration service in the target machine"
2790341 (https://www.netiq.com/support/kb/doc.php?id=7920341)	INFO: What ports does PlateSpin Migrate use during discovery, migration and file transfer?

H.4 PlateSpin Images

[Table H-2](#) provides information to help you troubleshoot common problems that might occur for PlateSpin Images.

Table H-2 Common Issues and Solutions Related to PlateSpin Images

Problems or Messages	Solutions
Cannot see PlateSpin Images on PlateSpin Image Server	<p>If the Servers view is configured to group servers by machine, discovered image servers cannot be expanded. To display the images, reconfigure the Servers View so the servers are grouped by domain instead of machine.</p>

Problems or Messages	Solutions
Failed to mount image. The volume does not contain a recognized file system	This error message might appear when you are importing or deploying volume data while installing a PlateSpin Image Server on Windows Server 2003. To resolve the error, use the Windows services plug-in on the PlateSpin Image Server. Modify the logon properties for the PlateSpin Migrate Operations Management Controller service to use an account with local administrative privileges. Restart the service after making this change.
Security descriptors are not intact on deployed server when you are using volume data from a Symantec Ghost image	When you are creating a PlateSpin Image using raw volume data that was extracted from a Ghost Image, the security descriptors are not preserved on the VM. This is because the extracted files inherit permissions of their parent folder.
Related KB Articles:	
ID	Description
7920879 (https://www.netiq.com/support/kb/doc.php?id=7920879)	ERRMSG: The file cannot be accessed by the system

H.5 Shrinking the PlateSpin Migrate Databases

When the PlateSpin Migrate databases (OFX and PortabilitySuite) reach a predetermined capacity, cleanup on those databases occurs at regular intervals. If there is a need to further regulate the size or content of those databases, Migrate provides a PlateSpin Database Cleanup utility (PlateSpin.DBCleanup.exe) to further clean up and shrink those databases. [KB Article 7006458 \(https://www.netiq.com/support/kb/doc.php?id=7006458\)](https://www.netiq.com/support/kb/doc.php?id=7006458) explains the location of the tool and the options available for it, should you decide to use it for offline database operations.

H.6 Troubleshooting the Configuration Service

After Test Cutover or Cutover, an error occurs on the target VM because of non-specific Configuration Service issues. The common error message is:

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

Troubleshooting tips in this section explain common Configuration Service issues and some alternative ways to resolve them.

- ♦ [Section H.6.1, “Understanding What Is Causing the Problem,” on page 452](#)
- ♦ [Section H.6.2, “What Can Be Done to Resolve the Problem,” on page 452](#)
- ♦ [Section H.6.3, “Additional Troubleshooting Tips,” on page 455](#)

H.6.1 Understanding What Is Causing the Problem

The Configuration Service error indicates that the PlateSpin Server is unable to communicate with the Configuration Service on the Target VM. Analyze your system to determine the possible root cause of the problem.

- ♦ [“Target VM Fails to Boot” on page 452](#)
- ♦ [“Network Is Not Set Up Correctly” on page 452](#)
- ♦ [“Unable to Read or Write Status Messages to Floppy Devices” on page 452](#)

Target VM Fails to Boot

The operating system must be loaded in the target VM in order for the Configuration Service to start up normally. A failure to boot indicates that there could be a driver conflict, a boot loader error, or possible disk corruption.

We recommend that you open a service ticket with Micro Focus Customer Care if the operating system fails to boot on the target VM.

Network Is Not Set Up Correctly

The network must be set up correctly in order for the Configuration Service on the target workload to communicate with the PlateSpin Server.

Ensure that you have configured your network in a way that the target workload can communicate with the PlateSpin Server.

Unable to Read or Write Status Messages to Floppy Devices

The Configuration Service must be able to communicate with the floppy devices for VMware VMs in order to read and write status messages for the PlateSpin Server.

On the target VM, verify that the machine is able to communicate with the floppy devices:

- 1 On the VM, open the log file (C:\windows\platespin\configuration\data\log.txt).
- 2 Any of the following messages might be an indication that the floppy is inaccessible:

```
Failed (5) to write to file \\?\Volume{<guid-number>}\log.zip

CopyFile \\?\Volume{<guid-number>}\windows\platespin\configuration\data\result.txt
to \\?\Volume{<guid-number>}\result.txt failed

The output floppy was not accessible after the timeout period
```

H.6.2 What Can Be Done to Resolve the Problem

To resolve a Configuration Service error, you can try any of the solutions in this section.

- ♦ [“Skip the Target VM Reboot Optimizations” on page 453](#)
- ♦ [“Reduce the Read/Write Traffic to Floppy Devices” on page 453](#)
- ♦ [“Change the Startup Type to Increase the Delay” on page 454](#)
- ♦ [“Configure Conflicting Services to Not Run Automatically at Startup” on page 455](#)

Skip the Target VM Reboot Optimizations

Migrate tries to minimize the number of reboots that occur on the target VM by default in order to speed up the Cutover process. It is possible that allowing the additional reboots will improve the target VM's ability to communicate with the PlateSpin Server.

To skip reboot optimizations:

- 1 Log in as Administrator to the PlateSpin Migrate Web Interface, then open the PlateSpin Server Configuration page at:
`https://Your_PlateSpin_Server/PlateSpinConfiguration/`
- 2 Search for the parameter **ConfigurationServiceValues**.
- 3 Edit the **ConfigurationServiceValues** parameter and set the **SkipRebootOptimization** option to `true`.
- 4 Click **Save**.
- 5 Run an incremental or full replication.
The replication also propagates the modified configuration settings to the target VM.
- 6 Run the Test Cutover or Cutover again for affected workloads.

Reduce the Read/Write Traffic to Floppy Devices

You can decrease the number of times the PlateSpin Server attempts to read from and write to the VMware input or output floppy devices if the diagnostic log shows the following error:

```
Information:1:Attempting floppy download
```

followed by

```
Verbose:1:Failed to copy file from remote URL
```

-or-

```
Exception: The remote server returned an error: (500) Internal Server Error
```

This error is caused by VMware locking the resource. It indicates that the PlateSpin Server is detaching and reattaching the floppy each time it checks the status. Locking can cause the target VM to fail to read and write to the floppy device. See [Using the VMware vCenter Server 4.x, 5.x and 6.0 Datastore Browser to Download or Copy a Powered-On Virtual Machine's .vmx and .nvram Files Fails \(1019286\)](https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1019286) (https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1019286).

If you experience floppy device locking issues, you can increase values for the Configuration Service polling settings on the PlateSpin Server:

vmwareConfigServicePollStartDelay

This parameter determines how long to wait before the PlateSpin Server starts polling for target workload status. The default value is 120 seconds (2 minutes).

vmwareConfigServicePollIntervalInMilliseconds

This parameter determines how frequently the PlateSpin Server attempts to communicate with the target workload and to read or write to the VMware floppy devices. The poll interval default is 30000 ms (30 seconds).

vmwareConfigServicePollStartTimeout

This parameter determines how long the PlateSpin Server waits after it starts the target VM before it displays an error in the Web Interface. The default value is 420 seconds (7 minutes).

vmwareConfigServicePollUpdateTimeout

This parameter determines how long the PlateSpin Server waits after each polling interval before displaying an error in the Web Interface. The default value is 300 seconds (5 minutes).

Higher values for these parameters reduce the frequency that the PlateSpin Server attempts to read from and write to the VMware floppy devices on target VMs.

To reduce read and write traffic for VMware floppy devices:

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

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

- 2 Search for the Configuration Service polling parameters, modify their settings as appropriate, then click **Save**.

For example:

```
vmwareConfigServicePollStartDelay = 180 (3 minutes)
vmwareConfigServicePollIntervalInMilliseconds = 300000 (5 minutes)
vmwareConfigServicePollStartTimeout = 1200 (20 minutes)
vmwareConfigServicePollUpdateTimeout = 900 (15 minutes)
```

or

```
vmwareConfigServicePollStartDelay = 300 (5 minutes)
vmwareConfigServicePollIntervalInMilliseconds = 480000 (8 minutes)
vmwareConfigServicePollStartTimeout = 1200 (20 minutes)
vmwareConfigServicePollUpdateTimeout = 900 (15 minutes)
```

- 3 Run an incremental or full replication.
The replication also propagates the modified configuration settings to the target VM.
- 4 Run the Test Cutover or Cutover again for affected workloads.

Change the Startup Type to Increase the Delay

The Configuration Service might be coming up before resources are accessible. You can change the Configuration Service startup type to have increase the delay.

To change the startup type:

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

`https://Your_PlateSpin_Server/PlateSpinConfiguration/`

- 2 Search for the parameter **windowsConfigServiceStartType**.
- 3 Change the **windowsConfigServiceStartType** value to **AutoDelay**.

Options for **windowsConfigServiceStartType** are:

- ♦ **GroupDelay** is the default value and adds the Configuration Service to the end of the **ServiceGroupOrder** in the registry.

- ♦ **AutoDelay** will maximize the amount of time the service waits before starting (2 minutes after boot). Also modify the **ServicesPipeTimeoutForWindowsConfigService** parameter value in [Step 4](#).
 - ♦ **NoDelay** is the most efficient option and starts the service as soon as Windows can. However, it is not recommended because of the potential issues connecting to resources.
- 4 (AutoDelay) Change the **ServicesPipeTimeoutForWindowsConfigService** parameter setting to 180 seconds to account for the 120 seconds that the service will take to start up after boot when AutoDelay is set for **windowsConfigServiceStartType** in [Step 3](#).
 - 5 Click **Save**.
 - 6 Run an incremental or full replication.
The replication also propagates the modified configuration settings to the target VM.
 - 7 Run the Test Cutover or Cutover again for affected workloads.

Configure Conflicting Services to Not Run Automatically at Startup

During a Cutover action, a Windows service interferes with the mounting of floppy drivers.

Determine which Windows Services are configured to start up at reboot. Some services are known to interfere with the Configuration Service writing to a floppy, such as Wireless Configuration and some antivirus software. You should configure these services to not run automatically on Test Cutover or Cutover, then run the Test Cutover or Cutover again.

You can also try to disable all non-essential services for Test Cutover and Cutover on the Configuration page, then run the Test Cutover or Cutover again.

H.6.3 Additional Troubleshooting Tips

If the Configuration Service cannot contact the PlateSpin Server, diagnostics will tell only part of the picture. You must also get logs from the target VM:

- ♦ **Windows workloads:** The Configuration Service logs are found in the `C:\windows\platespin\configuration\data` folder.
 - ♦ The `log.txt` file contains all of the logging information, but the `Config.ini` file is useful in understanding what is to be configured.
 - ♦ The `result.txt` file contains the status of the Configuration Service run.
 - ♦ If the target VM cannot read from the input floppy device, it will not have the merged `Config.ini` file, which might include custom network configuration information for the test Cutover network environment.
 - ♦ If the `Config.ini` file has no network related information (such as a `[NIC0]`), the target VM network adapter might have special characters in the name.
It is a known issue that the `Config.ini` file might not be accurate until it is merged with the one from the floppy device.
 - ♦ The target VM tries a reboot if it cannot connect to either the output floppy or input floppy (one time only). You will see a `config.ini.floppyreboot` file if this is the case.
- ♦ **Linux workloads:** The Configuration Service logs are found in the `/tmp` folder.
 - ♦ The main log files are named `file*.platespin.fileLogger`.
We recommend examining any configuration folders in `/tmp`. Tar the configuration folders along with the `file*.platespin.fileLogger` files to send to Micro Focus Customer Care.
 - ♦ Other config files to check for include the following:

```
/tmp/Ofx.RunCommand.Output*  
/tmp/*DiskHelper*  
/tmp/*VmTools*
```

- ♦ The configuration file is `/usr/lib/psconfigservice/data/config.conf`.
- ♦ The end result log file is `/usr/lib/psconfigservice/data/result.txt`.

H.7 Validation Warning for Bandwidth Throttling

Issue: After you configure migration for a workload with no warnings or validation errors, you might get a warning message if you then set or modify the value for **Bandwidth Throttling**, even if the setting is valid.

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

H.8 PlateSpin OFX Controller Does Not Start on a Virtual Machine Source

Issue: If you configure Migrate to install the block-based component during the first replication, PlateSpin OFX Controller might not start on the source workload during the Install Block-Based Components step. The Service Manager reports this problem if the VM is running so slowly that the OFX Controller startup event times out.

Workaround: Manually start PlateSpin OFX Controller on the source workload. To avoid the problem, you can configure the workload to install the block-based component during Prepare Workload instead of First Replication, or you can increase the Memory and CPU resources of the source VM to improve its startup performance.

H.9 Target Windows Machine Becomes Unbootable on Second Boot

Issue: The target Windows machine becomes unbootable during the second boot.

When PlateSpin Migrate executes the Configuration Service on a target Windows machine, the normal networking tasks performed during the second boot can be problematic in the following scenarios:

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

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

- ♦ If the target machine is booting from SAN.

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

Workaround: PlateSpin Migrate provides two light networking configuration settings for the PlateSpin Server that optimizes the network configuration process on the target machine during the second boot and helps avoid situations that can cause a target machine to become unbootable. Light networking is useful for P2P, V2V, and C2C migrations as well as for X2V semi-automated migrations where the networking hardware on the target VM is manually configured to match the source machine. See [Section 5.10, “Configuring Behavior for Installing Network Drivers on Target Windows Workloads,”](#) on page 107.

H.10 Two or More Volumes Have the Same Volume Serial Number

Issue: When you attempt to set up a migration job for a Windows server, the following error is displayed:

```
[Source] Two or more volumes have the same serial number. Change the serial numbers so that they are unique and rediscover the machine.
```

Workaround: This problem can occur if the Volume Serial Numbers for two or more volumes are the same. PlateSpin Migrate requires the serial numbers to be unique.

To resolve this issue, modify the serial numbers for the data volumes as appropriate, and then rediscover the machine. For information about how to use Windows native tools to modify the serial numbers, see [KB Article 7921101](#).

VII

Additional PlateSpin Tools

PlateSpin Migrate provides additional tools to support your migration efforts.

- ♦ [Appendix I, “Using the PlateSpin Migrate Client Command Line Interface,” on page 461](#)
- ♦ [Appendix J, “Using the iPerf Network Test Tool to Optimize Network Throughput for PlateSpin Products,” on page 487](#)

Using the PlateSpin Migrate Client Command Line Interface

The PlateSpin Migrate Client installation includes a command line interface (CLI) tool to help you perform common migrations tasks. Conversion jobs using `.ini` files is supported onto VMware and Hyper-V targets only. Using this tool, you can

- Discover and subsequently refresh a host or target server to populate the Migrate Server with server information.
- Migrate (also known as "convert") heterogeneous workloads across x86 server and desktop infrastructure in the data center.
- Prepare the target host for its new workload and then, after a conversion, synchronize the host and the target.
- Install an image server, capture an image, deploy an image, or incrementally migrate an image.
- Check the status of a job as it is running, and if necessary, abort it.

This section includes information that can help you use the CLI effectively. The content includes:

- [Section I.1, "Where Is the Tool Located?," on page 461](#)
- [Section I.2, "Before You Use the Tool," on page 461](#)
- [Section I.3, "Configurable .ini Files \(Jobs\) You Can Use with the Tool," on page 464](#)

I.1 Where Is the Tool Located?

The CLI tool, `PlateSpin.Migrate.Console.exe`, is installed with the PlateSpin Migrate Client at the following location:

- **32-bit host:** `C:\Program Files\PlateSpin Migrate Client\CommandLine\PlateSpin.Migrate.Console.exe`
- **64-bit host:** `C:\Program Files(x86)\PlateSpin Migrate Client\CommandLine\PlateSpin.Migrate.Console.exe`

I.2 Before You Use the Tool

This section includes the following information:

- [Section I.2.1, "Pre-configuring the Migrate Server Values for CLI," on page 462](#)
- [Section I.2.2, "Becoming Familiar with the Commands," on page 462](#)

I.2.1 Pre-configuring the Migrate Server Values for CLI

Before you begin using the command line utility, you need to ensure that the Migrate Server is properly configured. You can check the configuration in the `PlateSpin.Migrate.Console.config` file, located in the same path as the command line utility. After the Migrate installation, the following config file should already be populated with values.

```
<?xml version="1.0" encoding="utf-8 ?>"
<configuration>
  <appSettings>
    <add key="MigrateServerURL" value="https://localhost/PlateSpinMigrate/" />
    <add key="ServerDomain" value="" />
    <add key="psuser" value="administrator" />
    <add key="pspassword" value="encoded_password" />
    <add key="encoded" value="yes" />
  </appSettings>
</configuration>
```

The tool uses these values as it executes commands. You need to reconcile the values in the file with the settings for the Migrate Server with which you want to connect.

The value for the `pspassword` key is blank by default and you must specify an encoded password as the value. To encode the password, use the `encode` command. For more information about commands, see [Section I.2.2, “Becoming Familiar with the Commands,” on page 462](#).

If you choose to provide encoded passwords for source workload and target container, set the value for the `encoded` key in the following line of the `PlateSpin.Migrate.Console.config` file to `yes`, otherwise set value to `no`.

```
<add key="encoded" value="no" />
```

I.2.2 Becoming Familiar with the Commands

You can display the commands supported in the tool by running it with the `Help` option or with the `?` option from the command prompt, like this:

```
C:\Program Files\PlateSpin Migrate
Client\CommandLine>PlateSpin.Migrate.Console.exe Help
```

The tool displays a matrix that includes information similar to what is included in the following table:

Table I-1 Commands available from the Migrate CLI tool

Command	Description
run	Runs a configured <code>.ini</code> file as a scheduled job. When the you add the <code>/wait=no</code> parameter and the job starts to run, its Job ID is displayed in the interface.
query	Runs a query on the job (when you specify a Job ID) to display its current status.
discover	Runs an operation that inventories the details of a supported workload or target computer in preparation for a migration or “conversion” job.
refresh	Refreshes a discovered server.
unDiscover	Undiscovers a server.
imageserver	Performs imaging operations on a workload (that is, <i>install server</i> , <i>uninstall server</i> , <i>update tools</i>) on a server.

Command	Description
abort	Aborts a scheduled job.
licenseInfo	Displays the license information of the migrate server.
serversync	Prepares the server for the Server Sync operation and then runs a serversync job using the configuration file.
encode	Encodes the text input or the data in the text file.
massdiscover	<p>Performs mass discovery of source workloads and targets. The discovered workloads and targets are displayed both in the PlateSpin Migrate Client and the PlateSpin Migrate Web Interface</p> <p>To mass discover workloads and targets, you must first list the workloads and targets that you want to discover in a CSV file. To create this CSV file, refer to the sample CSV file located at \PlateSpin Migrate Client\CommandLine\Sample INI\MassDiscovery.csv.</p>

When you run any of these commands, you must include its required parameter(s) in the command line. You can also include some optional parameters, when needed. For example, `savejob=` parameter saves the job in default location.

To display a list of these parameters at the command prompt, run the command without any parameter. For example, if you run the `discover` command without parameters, like this:

```
C:\Program Files\PlateSpin Migrate
Client\CommandLine>PlateSpin.Migrate.Console.exe discover
```

the command line interface displays these following:

```
[discover]                                discovers a server
Required Parameters:
  /machineAddress=    machine address to discover
  /userName=          the username to use
  /password=          the password to use
  /type=              type like windows, linux,vmware_esx,vmware_vcenter,
Optional Parameters:
  /network=           network name to connect to
  /address=           server address to connect to
  /psuser=            Username used for accessing PlateSpin Migrate server as
user different from the one logged on this computer
  /pspassword=        Password used for accessing Platespin Migrate server for
the user different from the one logged on this computer
  /wait=              wait for completion of job [yes,no]
  /clusterName=       clustername to be discovered
  /verbose=           verbose mode for output [on,off]
  /output=            the output file
  /format=            the ouptput format to display in [text,html,xml]
  /sslcertificatewarnings= Whether to Ignore or Enforce SSL Certificate
Warnings [Ignore| Enforce]
```

NOTE: You should become familiar with the different CLI commands and their respective required and optional parameters.

Command Line Syntax

If you were to run the discover command (which is also a job), you would use a syntax similar to this example, at the command prompt:

```
C:\Program Files\PlateSpin Migrate
Client\CommandLine>PlateSpin.Migrate.Console.exe discover /
machineaddress=10.10.8.100 /username=administrator /password=password /
type=windows /wait=no
```

Note that all required parameters and one optional parameter are included in this example.

When the discover command (job) starts, the CLI tool displays its job ID, similar to this example:

```
8be8d306-7665-4869-9795-a9dbb3ce1471
```

You can leverage this ID to learn the status of the job, just by using the query command, like this:

```
C:\Program Files\PlateSpin Migrate
Client\CommandLine>PlateSpin.Migrate.Console.exe query /id=8be8d306-7665-4869-
9795-a9dbb3ce1471
```

The query command yields a status report that includes all of the details of the job. This is the same kind of information you might see from the Migrate Client Jobs view.

I.3 Configurable .ini Files (Jobs) You Can Use with the Tool

When you install the PlateSpin Migrate Client, the installation creates a separate directory for a number of preconfigured jobs (actually, .ini files) that can do the following:

- ♦ Workload conversion (that is, a migration operation)
- ♦ Server Sync
- ♦ Imaging capture and deployment of image target

You execute a job by using the run command at the command line. The values in the files are the optional parameters that run along with the job. Each of these functions has a “default” .ini file version that runs with basic settings, and one or more “platform-specific” .ini file(s) that run with custom settings:

- ♦ Conversion-Default.ini
- ♦ Conversion-Windows.ini (customized)
- ♦ Conversion-Linux.ini (customized)
- ♦ ServerSync-Default.ini
- ♦ ServerSync-Windows.ini (customized)
- ♦ ServerSync-Linux.ini (customized)
- ♦ CaptureImage-Default.ini
- ♦ CaptureImage.ini (customized)
- ♦ DeployImage-Default.ini
- ♦ DeployImage.ini (customized)
- ♦ IncrementalImaging-Default.ini
- ♦ IncrementalImaging.ini (customized)

This section includes more details about these jobs in the following subsections:

- ♦ [Section I.3.1, “Conversion Jobs,” on page 465](#)
- ♦ [Section I.3.2, “ServerSync Jobs,” on page 472](#)
- ♦ [Section I.3.3, “Imaging Jobs,” on page 476](#)

I.3.1 Conversion Jobs

The CLI tool supports converting Windows and Linux workloads (source) to Hyper-V, vCenter, or ESX servers (target). There are two types of `.ini` files, one for a basic job configuration, and one for custom configurations. While the job is running you can abort the job or check its status.

Before you start a conversion job, ensure that you run the `discover` command on the source computer and then on the target container. The following is example syntax for running the `discover` command:

```
discover /machineaddress=10.10.10.10 /username=administrator /  
password=anything@123 /type=vmware_vcenter
```

The tables in this section are named by the respective conversion jobs `.ini` files they represent. The table contents include the file section names within the `.ini` and the available settings you can configure according to your conversion needs:

IMPORTANT: For conversions to Hyper-V Generation 1 or BIOS machines, you must ensure that the boot volume is always mapped to Disk1 irrespective of the number of disks on the target machine. So, in the `.ini` file, you must ensure that the `MapTo=` setting in the `[Volume]` section that has `VolumeToCopy` mapped to boot volume is set to Disk1.

Sample of the settings in `Conversion-Windows.ini` file:

```
[Volume1]  
  
VolumeToCopy=boot_volume  
  
FreeSpace=  
  
MapTo=Disk1  
  
[Volume2]  
  
VolumeToCopy=non_boot_volume  
  
FreeSpace=  
  
MapTo=Disk2
```

Conversion-Default.ini

Table I-2 Details of `Conversion-Default.ini`

File Sections and Default Settings	Comment
[Type]	
Conversion=X2V	{required} This value must be used for every conversion.

File Sections and Default Settings	Comment
[JobConfig]	
Default=true	
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
TakeControl=static/dhcp	{conditional} Use this value only if the source is Windows Server 2003. Otherwise, it is not required.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
DNS=	
[TargetContainer]	
Address=	<p>{required} Specify the IP address for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery. ♦ If Hyper-V is discovered, specify the Hyper-V IP Address.
UserName=	<p>{required} Specify the username for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered via VCenter, specify the vCenter username. ♦ If ESX is discovered via Direct ESX discovery, specify the ESX root username. ♦ If Hyper-V is discovered, specify the Hyper-V username.
Password=	<p>{required} Specify the password for the target container depending on how it is discovered. For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered via VCenter, specify the vCenter password. ♦ If ESX is discovered via Direct ESX discovery, specify the ESX root password. ♦ If Hyper-V is discovered, specify the Hyper-V password.
[NewMachine]	
DisplayName=	{required} Specify the name you want to display in the target container console.
HostName=	{required} Host name of the target machine.

Conversion-Windows.ini

You can skip system volume.

Table I-3 Details of *Conversion-Windows.ini*

File Sections and Default Settings	Comment
[Type]	
Conversion=X2V	{required} This value must be used for every conversion.
[JobConfig]	
Default=false	
[Transfer]	
TransferType=VSSFileBased/VSSblockBased/FileBased	Possible settings shown. If the Windows source machine support VSS snapshotting, use the VSS setting, if it does not support VSS, use the Filebased setting.
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the TransferType setting. If that setting is Filebased and you want to perform an offline conversion, this setting must be set to false.
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
EndState=ShutDown/Donothing/Reboot	Possible settings shown.
TakeControl=static/dhcp	{conditional} Use this value only if the source is Windows Server 2003. Otherwise, it is not required.
TakeControlAddress=	
VirtualNetwork=	For offline conversions, specify the MAC address of the source workload.
SubnetMask=	
DefaultGateway=	
DNS=	
[TargetContainer]	

File Sections and Default Settings	Comment
Address=	<p>{required} Specify the IP address for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery. ♦ If Hyper-V is discovered, specify the Hyper-V IP Address.
UserName=	<p>{required} Specify the username for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered via VCenter, specify the vCenter username. ♦ If ESX is discovered via Direct ESX discovery, specify the ESX root username. ♦ If Hyper-V is discovered, specify the Hyper-V username.
Password=	<p>{required} Specify the password for the target container depending on how it is discovered. For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered via VCenter, specify the vCenter password. ♦ If ESX is discovered via Direct ESX discovery, specify the ESX root password. ♦ If Hyper-V is discovered, specify the Hyper-V password.
VirtualNetwork=	Specify the target container virtual network name you want to use.
TakeControl=static/dhcp	Specify <i>static</i> or <i>dhcp</i> , depending on your networking configuration.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
DNS=	
[NewMachine]	
DisplayName=	{required} Specify the name you want to display in the target container console.

File Sections and Default Settings	Comment
DataStore=	<p>Specify the name of datastore you want to use for configuration files.</p> <ul style="list-style-type: none"> ♦ On ESX: datastore1 ♦ On Hyper-V: E:
ConfigPath=	<ul style="list-style-type: none"> ♦ On ESX: Specify the complete path where you want to create the .vmx file. For example: <i>/folder_name/vmx_file_name</i> The .vmx file is created in the specified folder within the datasource. ♦ On Hyper-V: Specify the path to the folder where you want to create the configuration file. For example: <i>Drive:\folder_name\config_file_name</i>
Memory=	Specify the amount of RAM you want for the target computer. The setting can be in MB or GB and must be specified with integers (no decimal values).
InstallTools=true/false	Possible settings shown. Default is true.
NumberOfCPU=	Specify the number of CPUs you want for the target computer.
HostName=	{required} Specify the target host name.
WorkGroup=	{optional} Specify the workgroup name you want to join.
Domain=	
DomainUserName=	
DomainUserPassword=	
EndState=VMPowerOFF/VMPowerON/VMSuspend	Possible settings shown.
ScsiType=	(On VMware) Specify the SCSI Adapter type. If you do not specify a type or specify an unsupported adapter type, the default adapter type is used.
ResourcePool=	(On VMware) Specify the ResourcePool name in the vCenter. If the resource pool is nested, then use \ to separate names. For example, windows\local.
UseThinDisks=	To use thin disks, specify true . Else, specify false .
BootMode=	<p>(On Hyper-V for Windows workload) Specify the boot mode supported on the target machine. For example:</p> <ul style="list-style-type: none"> ♦ If the target machine is Windows Server 2012, specify either BIOS or UEFI. ♦ If the target machine is Windows Server 2008, specify BIOS.

File Sections and Default Settings	Comment
[EthernetNic1]	You can repeat this section of the .ini file for every NIC at the target container. For example, the second NIC section would be named [EthernetNic2] . Configuration settings would be specified for each NIC section in the file.
DHCPEnabled=true/false	Specify true for DHCP and false for static IP.
VirtualNetwork=	Specify the target container virtual network name you want to use.
Address=	Specify the IP address for the target machine.
SubnetMask=	
DefaultGateway=	
DNS=	Specify one or more DNS names separated by commas.
[DriveGeneral]	If you have multiple disks at the source, you can specify them here. You can specify as many disks as there are at the source.
DataStore1=	Specify the datastore on the target container. For example: <ul style="list-style-type: none"> ♦ On ESX: datastore1 ♦ On Hyper-V: E:
Disk1=	Specify the path to the configuration file on the target container. For example: <ul style="list-style-type: none"> ♦ On ESX: /win2k8r2/win2k8r2.vmdk ♦ On Hyper-V: \win2k8r2\win2k8r2.vhdx
DataStore2=	
Disk2=	
[Volume1]	You can repeat this section of the .ini file for every volume at the target container. For example, the second volume section would be named [Volume2] . Configuration settings would be specified for each volume section in the file.
VolumeToCopy=	Specify the volume to copy to the target.
MapTo=	Specify the disk to map.
FreeSpace=	Specify the amount of free space, in MB or GB, available on the target for File-Based conversion.

Conversion-Linux.ini

The sections in the `Conversion-Windows.ini` and in the `Conversion-Linux.ini` file are identical, except for the settings in **[Transfer]** section, along with the settings for workgroup and domain configuration. The differences for the Linux source job are shown in the following table.

Table I-4 *Conversion-Linux.ini: Differences in Setting Details of the [Transfer] section*

File Sections and Default Settings (differences only)	Comment
[Transfer]	
TransferType=BlockBased/FileBased	Possible settings shown. Linux does not support VSS.
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the TransferType setting. If that setting is Filebased and you want to perform an offline conversion, this setting must be set to false.
[Source]	
VirtualNetwork=	For offline conversions, specify the MAC address of the source workload.
[NewMachine]	
ScsiType=	(On VMware) Specify the Scsi Adapter type. If you do not specify a type or specify an unsupported adapter type, the default adapter type is used.
ResourcePool=	(On VMware) Specify the ResourcePool name in the vCenter. If the resource pool is nested, then use \ to separate names. For example, windows\local.
UseThinDisks=	To use thin disks, specify true . Else, specify false .
[EthernetNic1]	
DNS=	Specify one or more DNS names separated by commas.
[LVMGroup]	
Group1= Add enteries depending on the number of groups you want. If you have two groups, then add the following: Group1= Group2=	Name of the LVM group in the source.
[Volume1]	
FreeSpace=	Specify the amount of free space, in MB or GB, available on the target for File-Based conversion.

I.3.2 ServerSync Jobs

Use `serversync` command to perform the Server Sync operation. There are two types of `.ini` files, one for a basic job configuration, and one for custom configurations. While the job is running you can abort the job or check its status. If you specify the required settings, it will start the job. Then, when it runs, the job populates the other values with default settings.

The tables in this section are named by the respective Server Sync jobs `.ini` files they represent. The table contents include the file section names within the `.ini` and the available settings you can configure according to your conversion needs:

ServerSync-Default.ini

Table I-5 Details of *ServerSync-Default.ini*

File Sections and Default Settings	Comment
[Type]	
Conversion=Sync2V	{required} This value must be used for every Server Sync operation.
[JobConfig]	
Default=true	
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
TakeControl=static/dhcp	{conditional} Use this value only if the source is Windows Server 2003 or for offline conversion.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
DNS=	
[TargetContainer]	
Address=	<p>{required} Specify the IP address for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none">♦ If ESX is discovered, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery.♦ If Hyper-V is discovered, specify the Hyper-V IP Address.

File Sections and Default Settings	Comment
UserName=	<p>{required} Specify the username for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered via VCenter, specify the vCenter username. ♦ If ESX is discovered via Direct ESX discovery, specify the ESX root username. ♦ If Hyper-V is discovered, specify the Hyper-V username.
Password=	<p>{required} Specify the password for the target container depending on how it is discovered. For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered via VCenter, specify the vCenter password. ♦ If ESX is discovered via Direct ESX discovery, specify the ESX root password. ♦ If Hyper-V is discovered, specify the Hyper-V password.
[ExistingTargetMachine]	
DisplayName=	{required} Specify the display name of the target machine where you want to sync.
HostName=	{required}

ServerSync-Windows.ini

For prepare for Sync, the ServerSync command uses target container and network details from TargetContainer and machine name from ExistingTargetMachine file sections.

Table I-6 Details of *ServerSync-Windows.ini*

File Sections and Default Settings	Comment
[Type]	
Conversion=Sync2V	{required} This value must be used for every Server Sync operation.
[JobConfig]	
Default=false	
[Transfer]	
TransferType=VSSFileBased/VSSblockBased/FileBased	Possible settings shown. If the Windows source machine support VSS snapshotting, use the vss settings, if it does not support VSS, use the Filebased setting.

File Sections and Default Settings	Comment
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the <code>TransferType</code> setting. If that setting is <code>Filebased</code> and you want to perform an offline conversion, this setting must be set to <code>false</code> .
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
EndState=ShutDown/Donothing/Reboot	Possible settings shown.
TakeControl=static/dhcp	{conditional} Use this value only if the source is Windows Server 2003. Otherwise, it is not required.
TakeControlAddress=	
VirtualNetwork=	For offline conversions, specify the MAC address of the source workload.
SubnetMask=	
DefaultGateway=	
DNS=	
[TargetContainer]	
Address=	<p>{required} Specify the IP address for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery. ♦ If Hyper-V is discovered, specify the Hyper-V IP address.
UserName=	<p>{required} Specify the username for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered via VCenter, specify the vCenter username. ♦ If ESX is discovered via Direct ESX discovery, specify the ESX root username. ♦ If Hyper-V is discovered, specify the Hyper-V username.

File Sections and Default Settings	Comment
Password=	<p>{required} Specify the password for the target container depending on how it is discovered. For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered via VCenter, specify the vCenter password. ♦ If ESX is discovered via Direct ESX discovery, specify the ESX root password. ♦ If Hyper-V is discovered, specify the Hyper-V password.
VirtualNetwork=	Specify the target container virtual network name you want to use.
TakeControl=static/dhcp	Specify <code>static</code> or <code>dhcp</code> depending on your networking configuration.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
DNS=	
[ExistingTargetMachine]	
DisplayName=	{required} Specify the display name of the target machine where you want to sync.
HostName=	.
InstallTools=true/false	.
WorkGroup=	Specify the workgroup name if you want to join workgroup.
Domain=	.
DomainUserName=	.
DomainUserPassword=	.
EndState=VMPowerOFF/VMPowerON/VMSuspend	Possible settings shown.
[EthernetNic1]	<p>You can repeat this section of the .ini file for every NIC at the target container. For example, the second NIC section would be named [EthernetNic2]. Configuration settings would be specified for each NIC section in the file.</p>
DHCPEnabled=true/false	Specify true for DHCP and false for static IP.
VirtualNetwork=	Specify the target container virtual network name you want to use.
Address=	Specify the IP address for the target machine.
SubnetMask=	
DefaultGateway=	

File Sections and Default Settings	Comment
DNS=	

ServerSync-Linux.ini

The sections in the `ServerSync-Windows.ini` and in the `ServerSync-Linux.ini` file are identical, except for the settings in `[Transfer]` section, along with the settings for the workgroup and domain configuration. For prepare for Sync, the `ServerSync` command uses target container and network details from `TargetContainer` and machine name from `ExistingTargetMachine` file sections.

The differences for the Linux source job are shown in the following table.

Table I-7 *ServerSync-Linux.ini: Differences in Setting Details of the [Transfer] section*

File Sections and Default Settings (differences only)	Comment
<code>[Transfer]</code>	
<code>TransferType=BlockBased/FileBased</code>	Possible settings shown. Linux does not support VSS.
<code>LiveTransferEnabled=true/false</code>	Possible settings shown. This setting is dependent on the <code>TransferType</code> setting. If that setting is <code>Filebased</code> and you want to perform an offline conversion, this setting must be set to <code>false</code> .
<code>[Source]</code>	
<code>VirtualNetwork=</code>	For offline conversions, specify the MAC address of the source workload.

I.3.3 Imaging Jobs

The CLI tool supports several imaging operations (for example, install, uninstall, and update tools) through its `imageserver` command. Before you start an `imageserver` job, ensure that you run the `discover` command on the source computer and then on the target container.

In addition to the `imageserver` job, the CLI tool supports imaging Windows workloads (source) to the target. There are two types of imaging `.ini` files, one for a basic job configuration, and one for custom configurations. While the job is running you can abort the job or check its status.

The tables in this section are named by the respective imaging jobs `.ini` files they represent. The table contents include the file section names within the `.ini` and the available settings you can configure according to your conversion needs:

CaptureImage-Default.ini

Table I-8 *Details of CaptureImage-Default.ini*

File Sections and Default Settings	Comment
<code>[Type]</code>	

File Sections and Default Settings	Comment
Conversion=X2I	{required} This value must be used for every image capture.
[JobConfig]	
Default=true	
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
TakeControl=static/dhcp	{conditional} Use this value only if the source is Windows Server 2003. Otherwise, it is not required.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
[ImageConfiguration]	
ImageDisplayName=	{required} Specify the display name of the image in the image server.
[TargetContainer]	
Address=	{required} Specify IP address of image server.
UserName=	{required} Specify the username of the image server.
Password=	{required} Specify the password of the image server.

CaptureImage.ini

You can skip system volume for capture image.

Table I-9 Details of *CaptureImage.ini*

File Sections and Default Settings	Comment
[Type]	
Conversion=X2I	{required} This value must be used for every image capture.
[JobConfig]	
Default=false	
[Transfer]	
TransferType=VSSFileBased/FileBased	Possible settings are shown. If the Windows source machine support VSS snapshotting, use the <code>VSSFileBased</code> setting, if it doesn't support VSS, use the <code>Filebased</code> setting.

File Sections and Default Settings	Comment
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the <code>TransferType</code> setting. If that setting is <code>Filebased</code> and you want to perform an offline conversion, this setting must be set to <code>false</code> .
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
EndState=ShutDown/Donothing/Reboot	Possible settings are shown.
TakeControl=static/dhcp	{conditional} Use this value only if the source is Windows Server 2003. Otherwise, it is not required.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
DNS=	
[ImageConfiguration]	
ImageDisplayName=	Specify a name that you want to appear in the Image Server.
CompresionEnabled=true/false	Specify whether or not to use NTFS file compression. Default is <code>false</code> .
ConfigurationPath=	Specify the absolute path of the configuration file.
[TargetContainer]	
Address=	{required} Specify IP address of image server.
UserName=	{required} Specify the username of the image server.
Password=	{required} Specify the password of the image server.
[Volume1]	
VolumeToCopy=	Specify the name of the volume you want to capture.
MapTo=	Specify the path where you want to create the package file for the volume.

DeployImage-Default.ini

Table I-10 Details of *DeployImage-Default.ini*

File Sections and Default Settings	Comment
[Type]	
Conversion=I2V	{required} This value must be used for every image deployment.
[JobConfig]	
Default=true	
[Source]	
Address=	{required} Specify an IP address for the image container.
UserName=	{required} Specify a username credential for the image container.
Password=	{required} Specify a password credential for the image container.
ImageDisplayName=	Specify the name of the image that you want to deploy.
[TargetContainer]	
Address=	<p>{required} Specify the IP address for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none">♦ If ESX is discovered, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery.♦ If Hyper-V is discovered, specify the Hyper-V IP Address.
UserName=	<p>{required} Specify the username for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none">♦ If ESX is discovered via VCenter, specify the vCenter username.♦ If ESX is discovered via Direct ESX discovery, specify the ESX root username.♦ If Hyper-V is discovered, specify the Hyper-V username.

File Sections and Default Settings	Comment
Password=	{required} Specify the password for the target container depending on how it is discovered. For example: <ul style="list-style-type: none"> ♦ If ESX is discovered via VCenter, specify the vCenter password. ♦ If ESX is discovered via Direct ESX discovery, specify the ESX root password. ♦ If Hyper-V is discovered, specify the Hyper-V password.
[NewMachine]	
DisplayName=	{required} Specify the name you want to display in the target container.
HostName=	Specify the target host name.

DeployImage.ini

You can skip system volume.

IMPORTANT: For deploying image to Hyper-V Generation 1 or BIOS machines, you must ensure that the boot volume is always mapped to Disk1 irrespective of the number of disks on the target machine. So, in the .ini file, you must ensure that the MapTo= setting in the [Volume] section that has VolumeToCopy mapped to boot volume is set to Disk1.

Sample of the settings in Conversion-Windows.ini file:

```
[Volume1]
VolumeToCopy=boot_volume
FreeSpace=
MapTo=Disk1
[Volume2]
VolumeToCopy=non_boot_volume
FreeSpace=
MapTo=Disk2
```

Table I-11 Details of DeployImage.ini

File Sections and Default Settings	Comment
[Type]	
Conversion=I2V	{required} This value must be used for every image deployment.
[JobConfig]	
Default=false	

File Sections and Default Settings	Comment
[Source]	
Address=	{required} Specify an IP address for the image container.
UserName=	{required} Specify a username credential for the image container.
Password=	{required} Specify a password credential for the image container.
ImageDisplayName=	Specify a name for the image that you want to deploy.
[TargetContainer]	
Address=	<p>{required} Specify the IP address for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered, specify the IP Address of the ESX irrespective of whether the ESX is discovered via VCenter or via Direct ESX discovery. ♦ If Hyper-V is discovered, specify the Hyper-V IP Address.
UserName=	<p>{required} Specify the username for the target container depending on how it is discovered.</p> <p>For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered via VCenter, specify the vCenter username. ♦ If ESX is discovered via Direct ESX discovery, specify the ESX root username. ♦ If Hyper-V is discovered, specify the Hyper-V username.
Password=	<p>{required} Specify the password for the target container depending on how it is discovered. For example:</p> <ul style="list-style-type: none"> ♦ If ESX is discovered via VCenter, specify the vCenter password. ♦ If ESX is discovered via Direct ESX discovery, specify the ESX root password. ♦ If Hyper-V is discovered, specify the Hyper-V password.
TakeControl=static/dhcp	Specify <i>static</i> or <i>dhcp</i> , depending on your networking configuration.
VirtualNetwork=	
TakeControlAddress=	
SubnetMask=	

File Sections and Default Settings	Comment
DefaultGateway=	
[NewMachine]	
DisplayName=	{required} Specify the name you want to display in target container.
DataStore=	Specify the name of datastore you want to use. <ul style="list-style-type: none"> ♦ On ESX: datastore1 ♦ On Hyper-V: E:
ConfigPath=	<ul style="list-style-type: none"> ♦ On ESX: Specify the complete path where you want to create the .vmx file. For example: /folder_name/vmx_file_name The .vmx file is created in the specified folder within the datasource. ♦ On Hyper-V: Specify the path to the folder where you want to create the configuration file. For example: Drive:\folder_name\config_file_name
Memory=	Specify the amount of RAM you want for the target computer in MB or GB.
WorkGroup=	Specify the name of the workgroup you want to join.
InstallTools=true/false	Possible settings shown. Default is true.
NumberOfCPU=	Specify the number of CPUs you want for the target computer.
Memory=	.Specify the amount of RAM you want for the target computer.
Domain=	
DomainUserName=	
DomainUserPassword=	
HostName=	
EndState=VMPowerOFF/VMPowerON/VMSuspend	Possible settings shown.
ScsiType=	(On VMware) Specify the Scsi Adapter type. If you do not specify a type or specify an unsupported adapter type, the default adapter type is used.
ResourcePool=	(On VMware) Specify the ResourcePool name in the vCenter. If the resource pool is nested, then use \ to separate names. For example, windows\local.
UseThinDisks=	To use thin disks, specify true . Else, specify false .

File Sections and Default Settings	Comment
BootMode=	(On Hyper-V for Windows workload) Specify the boot mode supported on the target machine. For example: <ul style="list-style-type: none"> ♦ If the target machine is Windows Server 2012, specify either BIOS or UEFI. ♦ If the target machine is Windows Server 2008, specify BIOS.
[EthernetNic1]	If you have two (or more) NICs at the target, you can specify them and their configurations
DHCPEnabled=true/false	Specify true for DHCP and false for static IP.
VirtualNetwork=	Specify the target container virtual network name you want to use
Address=	
SubnetMask=	
DefaultGateway=	
DNS=	Specify one or more DNS names separated by commas.
[DriveGeneral]	If you have multiple disks at the source, you specify them here (create more if needed). The <code>.ini</code> file shows examples of two disks being specified. You can specify as many disks as there are at the source.
DataStore1=	Specify the datastore on the target container. For example: <ul style="list-style-type: none"> ♦ On ESX: <code>datastore1</code> ♦ On Hyper-V: <code>E:</code>
Disk1=	Specify the path to the configuration file on the target container. For example: <ul style="list-style-type: none"> ♦ On ESX: <code>/win2k8r2/win2k8r2.vmdk</code> ♦ On Hyper-V: <code>\win2k8r2\win2k8r2.vhdx</code>
DataStore2=	
Disk2=	
[Volume1]	
VolumeToCopy=	Specify the volume to copy to the target.
MapTo=	Specify the disk to map.
FreeSpace=	Specify the amount of free space, in MB or GB, available on the target for File-Based conversion.

IncrementalImaging-Default.ini

Table I-12 Details of *IncrementalImaging-Default.ini*

File Sections and Default Settings	Comment
[Type]	
Conversion=Sync2I	{“Existing Image”: required} Every incremental image capture uses this setting.
[JobConfig]	
Default=true	
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
TakeControl=static/dhcp	{conditional} Use this value only if the source is Windows Server 2003. Otherwise, it is not required.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
ImageDisplayName=	Specify an image name that already exists in the image server.
[TargetContainer]	
Address=	{required} Specify IP address of image server.
UserName=	{required} Specify the username of the image server.
Password=	{required} Specify the password of the image server.

IncrementalImaging.ini

Table I-13 Details of *IncrementalImaging.ini*

File Sections and Default Settings	Comment
[Type]	
Conversion=Sync2I	{“Existing Image”: required} Every incremental image capture uses this setting.
[JobConfig]	
Default=false	
[Transfer]	

File Sections and Default Settings	Comment
LiveTransferEnabled=true/false	Possible settings shown. This setting is dependent on the <code>TransferType</code> setting. If that setting is <code>Filebased</code> and you want to perform an offline conversion, this setting must be set to <code>false</code> .
[Source]	
Address=	{required} Specify an IP address for the source workload.
UserName=	{required} Specify a username credential for the source workload.
Password=	{required} Specify a password credential for the source workload.
TakeControl=static/dhcp	{conditional} Use this value only if the source is Windows Server 2003. Otherwise, it is not required.
TakeControlAddress=	
SubnetMask=	
DefaultGateway=	
EndState=ShutDown/Donothing/Reboot	Possible settings shown.
ImageDisplayName=	Specify an image name that already exists in image server.
[TargetContainer]	
Address=	{required} Specify IP address of image server.
UserName=	{required} Specify the username of the image server.
Password=	{required} Specify the password of the image server.

J Using the iPerf Network Test Tool to Optimize Network Throughput for PlateSpin Products

Before you execute replication, ensure that you test the connection to see if there are any connection or bandwidth issues, and resolve them. This section describes how to use the open source iPerf Network Test tool to optimize throughput on the connection.

- ♦ [Section J.1, “Introduction,” on page 487](#)
- ♦ [Section J.2, “Calculations,” on page 488](#)
- ♦ [Section J.3, “Setup,” on page 489](#)
- ♦ [Section J.4, “Methodology,” on page 490](#)
- ♦ [Section J.5, “Expectations,” on page 491](#)

J.1 Introduction

In the interest of helping PlateSpin administrators in their efforts to achieve better network throughput when using PlateSpin products, the iPerf Network Test tool is provided on the PlateSpin LRD (Linux RAM Disk) take-control environment. As stated in the iPerf documentation: “The primary goal of iPerf is to help in tuning TCP connections over a particular path. The most fundamental tuning issue for TCP is the TCP window size, which controls how much data can be in the network at any one point.”

The purpose of this README is to describe a basic method for network tuning and testing as it relates to using PlateSpin products. First, you calculate a theoretical optimum TCP window size. Then you use the iPerf tool to validate and fine-tune this calculated size and measure the resulting throughput. Using this method is also useful in determining the real achievable throughput for a given network.

Both the iPerf tool and PlateSpin products are actually using the *TCP send/receive buffer size* in order to affect the eventual internal choice of *TCP window size*. Going forward, these terms will be used interchangeably.

NOTE: There are many factors that affect network throughput. A wealth of information is available on the Internet that can aid in understanding. One such resource is the [Network Throughput Calculator \(http://wintelguy.com/wanperf.pl\)](http://wintelguy.com/wanperf.pl), which can help in calculating the expected maximum TCP throughput given applicable customer network characteristics. We strongly recommend that this online calculator be used in order to correctly set expectations regarding throughput.

J.2 Calculations

Tuning of the TCP window size is based on a number of factors, including network link speed and network latency. For our purposes relating to PlateSpin products, the initial choice of TCP window size for tuning is based on standard calculations (widely available on the Internet and elsewhere) as follows:

$$\text{WinSizeInBytes} = ((\text{LINK_SPEED}(\text{Mbps}) / 8) * \text{DELAY}(\text{sec})) * 1000 * 1024$$

For example, for a 54 Mbps link with 150 ms latency, the proper initial window size would be:

$$(54/8) * 0.15 * 1000 * 1024 = 1,036,800 \text{ bytes}$$

For a 1000 Mbps link with 10 ms latency, the proper initial window size would be:

$$(1000/8) * 0.01 * 1000 * 1024 = 1,280,000 \text{ bytes}$$

In order to get a latency value for the network, use `ping` from the command prompt (Windows) or the terminal (Linux). Although the `ping` round-trip time (RTT) is arguably different than the actual latency, the value obtained is sufficiently close for use in this method.

The following is a sample output from a Windows `ping` command, where the latency is observed to be 164 ms on average:

```
ping 10.10.10.232 -n 5
```

```
Pinging 10.10.10.232 with 32 bytes of data:
Reply from 10.10.10.232: bytes=32 time=154ms TTL=61
Reply from 10.10.10.232: bytes=32 time=157ms TTL=61
Reply from 10.10.10.232: bytes=32 time=204ms TTL=61
Reply from 10.10.10.232: bytes=32 time=153ms TTL=61
Reply from 10.10.10.232: bytes=32 time=153ms TTL=61
```

```
Ping statistics for 10.10.10.232:
    Packets: Sent = 5, Received = 5, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 153ms, Maximum = 204ms, Average = 164ms
```

The following is a sample output from a Linux `ping` command, where the latency is observed to be 319 ms on average:

```
ping 10.10.10.232 -c 5
```

```
PING 10.10.10.232 (10.10.10.232) 56(84) bytes of data.
64 bytes from 10.10.10.232: icmp_seq=1 ttl=62 time=0.328 ms
64 bytes from 10.10.10.232: icmp_seq=2 ttl=62 time=0.280 ms
64 bytes from 10.10.10.232: icmp_seq=3 ttl=62 time=0.322 ms
64 bytes from 10.10.10.232: icmp_seq=4 ttl=62 time=0.349 ms
64 bytes from 10.10.10.232: icmp_seq=5 ttl=62 time=0.316 ms

--- 10.10.10.232 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 3998ms
rtt min/avg/max/mdev = 0.280/0.319/0.349/0.022 ms
```

In practice, you should use the `-n` or `-c` option to specify a larger number of ping packets in order to more closely measure the latency value.

J.3 Setup

The iPerf tool runs in either server mode or client mode.

The basic usage syntax for `iperf` server mode is:

```
iperf -s -w <win_size>
```

The basic usage syntax for `iperf` client mode is:

```
iperf -c <server_ip> -w <win_size>
```

Our intent is to measure and tune the network between a source and target workload. In many cases, these can be the actual source and targets in use. It is possible to complete the testing using a different workload for either source or target, provided that the substitute has the same network characteristics as the original, such as NIC, network connection, and so on.

NOTE: Ensure that you are not testing the throughput from the PlateSpin server to either the source or the target, as this traffic is minimal, and does not represent the traffic that occurs during a migration or replication.

While it is possible to use a live workload (either Windows or Linux) as the target/iperf server, the following steps provide the environment most similar to what happens at migration/replication time, and is strongly recommended.

To set up and run `iperf` on the target:

- 1 Boot the target using the LRD.
- 2 In the LRD console, use the helper terminal (accessible via Alt-F2) to do the following:
 - 2a Set up networking using option 5.
 - 2b Mount the CD media using option 6.
- 3 In the LRD console, switch to the debug terminal (accessible via Alt-F7) to go to the location of the iPerf tool:

```
cd /mnt/cdrom/LRDTools/iperf_2.0.X/linux
```

- 4 Run the iPerf tool in server mode. Enter

```
./iperf -s -w <win_size>
```

To set up and run `iperf` on the source:

- 1 Mount the LRD ISO by using software or physical media.
- 2 Open a command prompt (Windows) or terminal (Linux) and go to the location of the iPerf tool:

```
cd <media>/LRDTools/iperf_2.0.X/
```

- 3 As determined by the source operating system, go to the `windows` or `linux` subdirectory:

```
cd windows
```

-OR-

```
cd linux
```

- 4 Run the iPerf tool in client mode. Enter

```
iperf -c <target_ip> -w <win_size>
```

NOTE: You can download and use `iperf3` for the calculations, which is helpful in certain scenarios where `iperf2` is unable to generate useful throughput numbers. Although the command syntax and output from `iperf3` differs slightly, it should be fairly straightforward to adapt and interpret the newer output, if necessary.

J.4 Methodology

Starting with the initial `win_size` calculated in the [Calculations](#) section, record the output from several iterations of the iPerf tool using the calculated value as well as slightly larger and smaller values. We recommend that you increase and decrease the `win_size` by increments of about 10 percent of the original value.

For example, for the example of 1,280,000 bytes above, you might increase or decrease `win_size` in increments of about 100,000 bytes.

NOTE: The `-w` option of `iperf` allows specifying units such as K (kilobytes) or M (megabytes).

Using the same example, you could use `-w` values of 1.28M, 1.38M, 1.18M, and so on as the `win_size` in Step 4. Of course, it is assumed that only the run step is repeated for each iteration of the iPerf tool.

Sample output from an `iperf` client iteration looks similar to the following:

```
iperf.exe -c 10.10.10.232 -w 1.1M

-----
Client connecting to 10.10.10.232, TCP port 5001
TCP window size: 1.10 MByte
-----
[296] local 10.10.10.224 port 64667 connected with 10.10.10.232 port 5001
[ ID] Interval      Transfer    Bandwidth
[296]  0.0-10.2 sec  11.3 MBytes  9.29 Mbits/sec
```

Sample output from the referenced target server looks similar to the following:

```
./iperf -s -w .6M

-----
Server listening on TCP port 5001
TCP window size: 1.20 MByte (WARNING: requested 614 Kbyte)
-----
[  4] local 10.10.10.232 port 5001 connected with 10.10.10.224 port 64667
[  4] 0.0-10.2 sec  11.3 MBytes  9.29 Mbits/sec
```

NOTE:

- ♦ The client disconnects from the server after a single iteration, while the server continues to listen until stopped by using Ctrl-C.
 - ♦ The window size specified for a Linux server is 1/2 the desired value, because Linux doubles the requested TCP buffer size as a matter of course.
-

Use several iterations to determine the optimal value for the TCP window size. Remember to use only 1/2 the desired value when specifying the `-w` option for `iperf` on Linux.

Increased throughput indicates that you are getting closer to an optimal TCP window size. Finally, as you get closer to an optimal value, use longer iterations in order to more closely simulate real running conditions. To achieve a longer iteration, use the `-t <time_in_seconds>` option to `iperf`. This option needs to be specified only on the client side.

For example:

```
iperf.exe -c 10.10.10.232 -w 1.25M -t 60
```

After an optimal value has been determined, configure this value in the `FileTransferSendReceiveBufferSize` parameter for the appropriate PlateSpin server at:

https://<my_ps_server>/PlatespinConfiguration/

This global value applies to all workloads on the PlateSpin server, so care should be taken to group workloads and their respective networks in a sensible manner across available PlateSpin servers.

J.5 Expectations

Modifying the TCP window size indirectly with the TCP send/receive buffer size can be a very effective method for increasing network throughput in some scenarios. Two to three or even more times the original throughput can sometimes be achieved. However, it is important to remember that the network characteristics can (and often do) change over time because of changes in usage patterns, hardware, software, or other infrastructure.

We strongly recommend that you use this method to calculate the optimum value at the same time of day and under the same network usage patterns that are intended to be in use during the planned live migration or replication tasks. We also recommend that you recalculate the setting periodically in order to account for changing network conditions.



Documentation Updates

This guide has been updated since the General Availability of PlateSpin Migrate 12.2.1.

- ♦ [Appendix K, “Documentation Updates,” on page 495](#)

K

Documentation Updates

This section contains information on documentation content changes that were made in the English version of the *PlateSpin Migrate User Guide* since the General Availability of PlateSpin Migrate 12.2.1.

- ♦ [Section K.1, “October 2018,” on page 495](#)
- ♦ [Section K.2, “September 2018,” on page 495](#)
- ♦ [Section K.3, “March 2018,” on page 495](#)
- ♦ [Section K.4, “February 2018,” on page 496](#)
- ♦ [Section K.5, “January 2018,” on page 497](#)

K.1 October 2018

Location	Update
Section 9.5, “Registering a Azure Application to Represent PlateSpin Migrate,” on page 156	Updated procedure to add a step and update an existing step.

K.2 September 2018

Location	Update
Table 2-13, “Supported VMware Datastores,” on page 39	This table is new. It was added to clarify VMware datastore support.
Section 2.3.3, “Security of Client-Server Communications,” on page 46	Added information about Windows Registry settings to use on source workloads to force .NET Framework to require TLS 1.2 connections.

K.3 March 2018

Location	Update
Table 9-2, “Azure Global Environment Configuration Settings,” on page 160 Table 9-3, “Azure China Environment Configuration Settings,” on page 160	Corrected the PlateSpin Configuration Parameter from <code>AzureStorageEndpointStorage</code> to <code>AzureStorageEndpointSuffix</code> .
Figure 9-1, “On-Premise Migrate Server for Automated Migration to Azure,” on page 148	Corrected the arrow direction for HTTPS traffic to correctly point from the Web Interface to the Migrate Server.

Location	Update
Figure 10-1, “On-Premise Migrate Server for Automated Migration to vCloud,” on page 164	Corrected the arrow direction for HTTPS traffic to correctly point from the Web Interface to the Migrate Server.

K.4 February 2018

Location	Update
Table 2-2, “Non-Cloud Platforms: Supported Linux Workloads,” on page 26 Table 2-4, “Azure: Supported Linux Platforms,” on page 28 Table 2-6, “vCloud: Supported Linux Platforms,” on page 30 Table 2-8, “AWS: Supported Linux Platforms,” on page 32	For Red Hat Enterprise Linux 6.8, Oracle Linux 6.8, and CentOS 6.8 workloads with LVM volumes, PlateSpin Migrate supports incremental replication only for the latest available kernel (version 2.6.32-696.20.1.el6.x86_64) for the 6.8 distribution.
Table 2-12, “Supported Target VMware Platforms for the Migrate Web Interface and Migrate Client,” on page 38	For vCenter 6.0, Migrate supports VMware Virtual SAN (vSAN) 6.2 storage.
“Data Disks” on page 33	Software RAID: PlateSpin Migrate supports hardware RAID; however, PlateSpin Migrate does not support software RAID. This is applicable for both Windows and Linux workloads.
Section 5.10, “Configuring Behavior for Installing Network Drivers on Target Windows Workloads,” on page 107	This section is new.
“AWS user credentials” in Table 8-1, “Minimum Required Configuration for Your AWS Account,” on page 138	NOTE: We recommend that administrators regularly rotate access keys for IAM users. See “Rotating Access Keys” in the <i>AWS Identity and Access Management User Guide</i> .
Section 9.8, “Deploying a Migrate Server Image in Azure,” on page 158	For PlateSpin Migrate Servers deployed through the Azure Marketplace, Migrate installs Microsoft SQL Server 2014 Express Edition for the PlateSpin database and generates a password for the SQL administrator (sa) user. For improved security, we recommend that you modify the password for this user after the deployment is complete. For information, contact Micro Focus Technical Support .

K.5 January 2018

Location	Update
Section 2.6.3, “Requirements for Event Messaging,” on page 52	If you use PlateSpin Transformation Manager with a cloud-based Migrate server, allow inbound connections in the Network Security Group for STOMP communications (TCP port 61613) related to Event Messaging.
Step 3 in Section 5.6, “Enabling Event Messaging for PlateSpin Migration Factory,” on page 105	
Section 9.2.3, “Azure Prerequisites for Using an Azure-Based Migrate Server,” on page 152	
Section 9.2.3, “Azure Prerequisites for Using an Azure-Based Migrate Server,” on page 152	When no VPN is provided between the source network and the cloud-based Migrate server, you can use Migrate Agent to register workloads with the cloud-based Migrate server using secure communications over the public Internet. Internet access and public IP addresses are required. For deployment information, see Figure 9-2, “Cloud-Based Migrate Server for Automated Migration to Azure,” on page 149 .

