

User Guide

Novell® PlateSpin® Recon

3.7.2

June 30, 2011

www.novell.com



Legal Notices

Novell, Inc., makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc., makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. See the [Novell International Trade Services Web page \(http://www.novell.com/info/exports/\)](http://www.novell.com/info/exports/) for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Copyright © 2009-2011 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc.
1800 South Novell Place
Provo, UT 84606
U.S.A.
www.novell.com

Online Documentation: To access the latest online documentation for this and other Novell products, see the [Novell Documentation Web page \(http://www.novell.com/documentation\)](http://www.novell.com/documentation).

Novell Trademarks

For Novell trademarks, see [the Novell Trademark and Service Mark list \(http://www.novell.com/company/legal/trademarks/tmlist.html\)](http://www.novell.com/company/legal/trademarks/tmlist.html).

Third-Party Materials

All third-party trademarks are the property of their respective owners.

Contents

About This Guide	11
1 Introduction	13
1.1 PlateSpin Recon Features	13
1.2 Planning Module Features	14
1.3 Supported Platforms	14
1.3.1 Supported Inventory and Data Collection	15
1.3.2 Supported PlateSpin Portability Suite Consolidation Sources and Targets	17
1.4 Application Scenario	18
1.5 PlateSpin Recon Terminology	21
2 Overview of PlateSpin Recon Client Views	23
2.1 Data Center Explorer	23
2.2 Report Explorer	24
2.3 Projects Explorer	25
2.4 Raters Explorer	25
2.5 Chart Viewer	26
2.6 Job Explorer	26
2.7 Scheduled Job Explorer	27
2.8 Log Explorer	27
3 Working with Data Center Explorer	29
3.1 Understanding the Data Center Explorer	29
3.2 Discovering and Inventorying Data Center Resources	30
3.2.1 Inventory and Discovery Methods	30
3.2.2 Inventory Credentials	42
3.2.3 Inventory Options	44
3.2.4 Discovery Options	45
3.2.5 Viewing Inventory Data and Monitoring Servers	47
3.3 Working with Filters	52
3.4 Working with Groups	52
3.4.1 Creating a Group	52
3.4.2 Copying Servers to a Group	52
3.4.3 Deleting a Server from a Group	53
3.5 Working with Sites	53
3.5.1 Adding a Site	54
3.5.2 Deleting a Site	54
3.5.3 Modifying a Site	54
3.6 Working with PlateSpin Recon Snapshots	55
3.6.1 Exporting a PlateSpin Recon Snapshot	55
3.6.2 Importing a PlateSpin Recon Snapshot	58
3.7 Working with Remote Data	60
3.8 Working with Data Collected by Other PlateSpin Recon Instances	60
3.9 Working with Data Collected by Other Performance Monitoring Tools	62
3.9.1 Required Performance Counters	62
3.10 Using HP Reporter Data with PlateSpin Recon	66

3.10.1	Using Formulas to Convert Metrics to Counters	66
3.10.2	Example Formulas	67
3.10.3	Working with HP Reporter Data	68
3.11	Using Performance Data Import to Import Generic Data into PlateSpin Recon	70
3.11.1	Performance Data Import Tables Data Structure	71
3.11.2	Preparing Data for Import into PlateSpin Recon	72
3.12	Collecting Disk Utilization Counters for Windows Machines	74
3.12.1	Deploying an Add-On	75
3.12.2	Undeploying an Add-On	76

4 PlateSpin Recon Server and Client Tools 79

4.1	Server and Client Options	79
4.2	Credentials Manager	79
4.2.1	Creating a Credential	80
4.2.2	Deleting a Credential	80
4.3	PlateSpin Analyzer	81
4.3.1	Using PlateSpin Analyzer	81
4.3.2	Inventory and Monitoring Tests	83
4.3.3	Conversion Tests	84

5 Generating and Viewing Charts and Reports 85

5.1	Chart Viewer	85
5.1.1	Chart Controls and Commands	86
5.1.2	Expression Types	87
5.1.3	Level of Detail	87
5.1.4	Data Type Statistics	88
5.1.5	Server Instance	88
5.1.6	Date Range Presets	88
5.1.7	Forecasting	88
5.2	Report Explorer	89
5.2.1	Report Explorer Template Nodes	89
5.2.2	Commands	92
5.2.3	Analysis Reports	93
5.2.4	VM Disk Partition Alignment Reports	93
5.2.5	Working with Report Templates	96
5.3	Editing Report Templates	97
5.3.1	Report Template Option Tabs	97
5.3.2	Chargeback Report Templates	103
5.3.3	Consolidation Report Templates	103
5.3.4	Consolidation Project Samples Report Templates	104
5.3.5	Diagnostics Report Templates	104
5.3.6	Disk Report Templates	104
5.3.7	Inventory Report Templates	105
5.3.8	Reclamation Report Templates	105
5.3.9	Software Report Templates	105
5.3.10	Virtual Infrastructure Report Templates	105
5.3.11	Volume Report Templates	106
5.4	View Customization	106
5.4.1	Creating and Editing Report Views	106
5.4.2	Histogram Charts	108
5.4.3	Scatter Charts	109
5.4.4	Usage Charts	110
5.5	Generating Reports	110
5.6	Document Viewer	114
5.6.1	Commands	114

5.6.2	Customizing Reports	115
5.6.3	Viewing Reports	118
5.6.4	Printing and Exporting Reports	119
5.6.5	Copying Servers from a Report to a Group	120
6	Server Consolidation	121
6.1	About Server Consolidation	121
6.2	Creating Projects	122
6.3	Creating and Editing Server Templates	126
6.3.1	Creating a New Server Template	126
6.3.2	Editing an Existing Server Template	130
6.3.3	Copying a Server Template	133
6.4	Identifying Consolidation Candidates	134
6.5	Working with Consolidation Scenarios	134
6.5.1	Determining the Minimum Number of Servers	134
6.5.2	Specifying the Number of Servers	139
6.5.3	Results of Consolidation Scenarios	143
6.5.4	Editing Scenarios	144
6.5.5	Deleting Scenarios	145
6.6	Generating Inventory Reports	145
6.6.1	Finding Servers That Meet the Consolidation Criteria	145
6.6.2	Selecting Candidates for Consolidation	145
6.7	Viewing Consolidation Reports	146
6.7.1	Project Report	147
6.7.2	Scenario Summary Report	147
6.7.3	Workload Assignment Report	149
6.7.4	Unconsolidated Workload Report	150
6.7.5	Unprotected Workload Report	150
6.8	Implementing Scenarios with Portability Suite	151
6.8.1	Implementing Phantom or Existing Target Servers	152
7	IT Chargeback	155
7.1	Understanding Chargeback	155
7.2	Using Raters	155
7.2.1	Rater Resources	156
7.2.2	Rater Charges	157
7.2.3	Virtual Server Rater Methods	157
7.3	Managing Raters	158
7.3.1	Creating a Rater	158
7.3.2	Deleting a Rater	159
7.3.3	Editing a Rater	159
7.3.4	Attaching Raters to Machines	160
7.3.5	Detaching a Rater from a Machine	161
7.4	Chargeback Reporting	161
7.4.1	Generating a Report to Calculate IT Chargeback Cost	161
7.4.2	Deriving the Calculated Cost	161
8	Virtual Infrastructure Management	165
8.1	Virtual Center Inventory and Monitoring	165
8.2	VM Sprawl Reporting	166
8.3	Virtual Machine State Reporting	167

A Troubleshooting	169
B Technical Reference	187
B.1 Features	187
B.2 How PlateSpin Recon Collects Data	187
B.2.1 Discovery	188
B.2.2 Inventory	188
B.2.3 Monitoring	189
B.3 Data Collection	190
B.4 Charts and Formulas	191
B.4.1 Chart Data	191
B.4.2 Utilization	195
B.4.3 Workload	196
B.5 Calculations	198
B.5.1 24-Hour Profile	199
B.5.2 Dual-Core and Multi-Core Machines	200
B.5.3 Hyper-Threading	200
B.5.4 Effective Speed	200
B.5.5 Power and Cooling	201
B.5.6 CO2 Emission	201
B.6 CPU Normalization	201
B.6.1 SPEC Data	202
B.6.2 Normalization Speed	202
B.7 PlateSpin Recon VM CPU Inventory	202
B.8 Inventory and Monitoring Commands for Linux, Solaris and AIX	203
B.8.1 Inventory Commands for Solaris	203
B.8.2 Inventory Commands for Linux	204
B.8.3 Inventory Commands for AIX	205
B.8.4 Monitoring Commands for Linux and Solaris	206
B.8.5 Monitoring Commands for AIX	206
B.8.6 Text Processing	207
B.9 Sudo Configuration to Inventory and Monitor Linux Machines	207
B.10 Managing Credentials	207
C Schedule Options	209
C.1 Later	209
C.2 Recurrence	209
D Quick Reference to PlateSpin Recon Icons and Commands	211
D.1 Data Center Explorer	211
D.1.1 Data Center Icons	211
D.1.2 Data Center Icon Status Indicators	211
D.1.3 Data Center Commands and Associated Icons	212
D.2 Reports Explorer	213
D.3 Projects Explorer	213
D.3.1 Project Explorer Commands	213
D.3.2 Project Explorer Icons	214
D.4 Raters Explorer	214
D.5 Chart Viewer	215
D.6 Job Explorer	216
D.7 Scheduled Job Explorer	216
D.8 Log Explorer	216

E	Best Practices	217
E.1	Consolidation Planning Workflow.	217
E.2	PlateSpin Recon Server System Requirements.	218
E.3	Monitoring Service Configurations	218
E.3.1	Fine-Tuning the Monitoring Service Configuration Parameters Through the PlateSpin Recon Client	218
E.3.2	Fine-Tuning Monitoring the Service Configuration Parameters in the Monitoring Service Configuration File	219
E.4	MS SQL Database Configurations	220
F	Documentation Updates	225
F.1	June 30, 2011	225
F.2	March 22, 2011	225
F.3	March 07, 2011	226
F.4	February 17, 2011	226

About This Guide

This documentation provides information about PlateSpin Recon.

- ♦ Chapter 1, “Introduction,” on page 13
- ♦ Chapter 2, “Overview of PlateSpin Recon Client Views,” on page 23
- ♦ Chapter 3, “Working with Data Center Explorer,” on page 29
- ♦ Chapter 4, “PlateSpin Recon Server and Client Tools,” on page 79
- ♦ Chapter 5, “Generating and Viewing Charts and Reports,” on page 85
- ♦ Chapter 6, “Server Consolidation,” on page 121
- ♦ Chapter 7, “IT Chargeback,” on page 155
- ♦ Chapter 8, “Virtual Infrastructure Management,” on page 165
- ♦ Appendix A, “Troubleshooting,” on page 169
- ♦ Appendix B, “Technical Reference,” on page 187
- ♦ Appendix C, “Schedule Options,” on page 209
- ♦ Appendix D, “Quick Reference to PlateSpin Recon Icons and Commands,” on page 211
- ♦ Appendix E, “Best Practices,” on page 217
- ♦ Appendix F, “Documentation Updates,” on page 225

Audience

This text is intended for IT staff, such as data center administrators and operators, who use PlateSpin Recon in the management and planning of the data center.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please take a moment to provide your comments at the [Novell Documentation Web site](http://www.novell.com/documentation/feedback) (<http://www.novell.com/documentation/feedback>).

Additional Documentation

In addition to this guide, PlateSpin Recon comes with the following documentation:

- ♦ *PlateSpin Recon Getting Started Guide*
- ♦ An online version of the *PlateSpin Recon User Guide* integrated with the product
- ♦ *Release Notes*

For help with third-party vendor products or documentation, please contact your sales or support representative for those products.

Documentation Updates

For the most current version of the PlateSpin Recon 3.7.2 documentation, see the [Novell PlateSpin Recon 3.7.2 documentation Web site](http://www.novell.com/documentation/ps_recon37/) (http://www.novell.com/documentation/ps_recon37/).

Technical Support

Contact Technical Support by:

- ♦ Telephone (North America): +1 800 858 4000. This is a toll-free number.
- ♦ Telephone (global): +1 416 203 4799. This is not a toll-free number.
- ♦ E-mail: support@platespin.com

You can also visit the [PlateSpin Technical Support Web site \(http://www.platespin.com/support/\)](http://www.platespin.com/support/) or [Novell Customer Center Web site \(http://www.novell.com/center/eservice\)](http://www.novell.com/center/eservice).

User Forum

You can also visit the [user forum \(http://forums.novell.com/novell-product-support-forums/platespin/\)](http://forums.novell.com/novell-product-support-forums/platespin/).

Novell PlateSpin Recon provides intelligence, visual analysis, and forecasting for consolidating and optimizing the data center. It collects hardware, software, and service inventory data for server workloads. PlateSpin Recon can remotely gather workload utilization statistics for a clear and concise picture of the application services running in the data center, and to determine how resources are being used.

The PlateSpin Recon Consolidation Planning module determines the optimal fit between server resource supply and workload demand. It helps you visualize the characteristics of server workloads before and after consolidation, determine the servers required for consolidation, and expedite projects by automatically generating consolidation plans. You can also protect servers by allocating them to a suitable virtual recovery environment in case of a service outage.

PlateSpin Recon includes other features such as Virtual Infrastructure Management capabilities, as well as a resource utilization-based IT chargeback system. These can be used to help cope with the challenges presented by the virtualization and consolidation of the data center.

- ♦ [Section 1.1, “PlateSpin Recon Features,” on page 13](#)
- ♦ [Section 1.2, “Planning Module Features,” on page 14](#)
- ♦ [Section 1.3, “Supported Platforms,” on page 14](#)
- ♦ [Section 1.4, “Application Scenario,” on page 18](#)
- ♦ [Section 1.5, “PlateSpin Recon Terminology,” on page 21](#)

1.1 PlateSpin Recon Features

Remote Data Collection: PlateSpin Recon does not require the installation of agent software, eliminating the need to physically touch data center servers. All performance and inventory data is collected via standard OS instrumentation capabilities.

Rich Data Modeling: Helps you make consolidation choices based on sophisticated analysis of resources, workloads, and utilization trends. Tight integration with VMware VirtualCenter provides greater visibility into virtual infrastructures, thereby improving data center management and operations.

Custom Report Creation and Delivery: Helps you define resource and workload parameters and generate custom visual reports for data center assessments and server consolidations. You can quickly identify consolidation candidates based on resource utilization trends, and compare workload characteristics before and after consolidation. Scheduled report delivery via e-mail or FTP ensures easy access to remote data and provides up-to-date information.

Flexible Data Capture and Export: You can export PlateSpin Recon data to a number of formats, including HTML, PDF, Word, CSV, Excel, or images. Raw data can be extracted directly from the database and delivered to business intelligence applications for advanced statistical analysis.

Enterprise-Level Scalability: Each instance of PlateSpin Recon can monitor up to 1,500 servers to ensure enterprise scalability for the world’s largest data centers. Data can be aggregated from multiple PlateSpin Recon data collectors for centralized data warehousing, analysis, and planning, or to accommodate larger enterprise implementations.

Multiple-Data Center Support: You can use PlateSpin Recon in different geographical locations to remotely collect data and provide scheduled updates to a master PlateSpin Recon installation. This allows for centralized analysis and planning for initiatives such as server or data center consolidation.

Flexible Chargeback Reporting: Because virtualization creates a shared pool of computing resources, it can be difficult to manage and monitor virtual resource usage. PlateSpin Recon enables organizations to effectively allocate and share virtual resources across business units and departmental owners. Flexible chargeback reporting capabilities improve virtual infrastructure management and financial accounting by allowing organizations to accurately calculate IT costs based on actual resource usage.

VM Growth Reporting: You can generate reports to monitor the proliferation of virtual machines and avoid the administrative headaches associated with virtual infrastructure sprawl.

1.2 Planning Module Features

Planning: You can automatically generate server consolidation and disaster recovery plans based on detailed workload analysis to ensure the optimal fit between server workloads and virtual resources. Use forecasted data to plan for and accommodate future growth.

Workload Analysis: The PlateSpin Recon Planning Module automatically analyzes five critical workload dimensions—processor, memory, disk space, disk I/O, and network I/O—simultaneously across thousands of servers, providing consolidation plans that maximize utilization while minimizing resource contention.

Scenario Modeling: You can use custom scenarios with user-defined target server specifications, including server templates or existing virtual machine servers to create an optimal consolidation plan.

Power and Cooling Analysis: You can compare and contrast potential power and cooling cost savings and ROI derived from different consolidation scenarios. Custom fields (maintained in a central database) let you input power and cooling requirements for major hardware platforms. This feature lets organizations analyze and cost-justify green computing initiatives.

Time-Based Analysis: You can stagger multiple workloads evenly across virtual hosts and account for hourly peaks and valleys inherent in server utilization trends.

Workload and Utilization Forecasting: You can predict future workloads and resource utilization based on historical trends to better plan for server consolidation and infrastructure growth, as well as enable more proactive systems management. Forecasting data on CPU, disk, memory, and usage trends is presented in easy-to-read charts, reports, and plans.

1.3 Supported Platforms

- ♦ [Section 1.3.1, “Supported Inventory and Data Collection,” on page 15](#)
- ♦ [Section 1.3.2, “Supported PlateSpin Portability Suite Consolidation Sources and Targets,” on page 17](#)

NOTE: For information on PlateSpin Recon Client/Server supported platforms, see “[System Requirements](#)” in the *Novell PlateSpin Recon 3.7.2 Getting Started Guide*.

1.3.1 Supported Inventory and Data Collection

PlateSpin Recon supports discovery, inventory, and data collection (monitoring) for the following operating system families: Microsoft Windows, UNIX/Linux and Novell Netware.

- ♦ [“Supported Microsoft Windows Platforms” on page 15](#)
- ♦ [“Supported UNIX/Linux Platforms” on page 15](#)
- ♦ [“Supported Novell NetWare Platforms” on page 16](#)
- ♦ [“Supported Hypervisors” on page 16](#)

Supported Microsoft Windows Platforms

Table 1-1 *Supported Windows Platforms*

Version	Bit Spec
Windows Server 2000 (SP0, SP1, SP2, SP3, SP4)	32-bit
Windows Server 2003 (SP0, SP1, SP2)	32-bit, 64-bit
Windows Server 2008 (SP1, SP2)	32-bit, 64-bit
Windows Server 2008 R2	64-bit
Windows NT 4.0 Server (SP4, SP6)	32-bit
Windows NT 4.0 Server (SP6a - WMI Core version 1.5 required)	32-bit
Windows XP Professional (SP0, SP1, SP2, SP3)	32-bit, 64-bit

PlateSpin Recon supports the following localized versions of Windows:

- ♦ French
- ♦ German
- ♦ Japanese

Supported UNIX/Linux Platforms

Table 1-2 *Supported UNIX/Linux Platforms*

Version	Bit Spec
Fedora Core 3, 4, 5, 6, 7, 8, 9, 10	32-bit, 64-bit
Novell SUSE Linux Enterprise Server (SLES) 9.x, 10.0, 10 SP1 (10.1), 10 SP2 (10.2), 10 SP3 (10.3), 11	32-bit, 64-bit
Novell SUSE Linux 8.0	32-bit
Novell OpenSUSE 10.1, 10.2, 10.3, 11, 11.1	32-bit, 64-bit
RedHat Linux 7.3, 8.0, 9.0, 9.1, 9.2	32-bit, 64-bit
RedHat Enterprise Linux 2.x AS/ES	32-bit

Version	Bit Spec
RedHat Enterprise Linux 3.0 AS/ES, 4.0 AS/ES, 5.0 AS/ES, 5.1 AS/ES,	32-bit, 64-bit
Ubuntu 7.1, 8.04, 8.10	32-bit, 64-bit
CentOS 2.0	32-bit
CentOS 3.9, 4.6, 5.1, 5.2	32-bit, 64-bit
Sun Solaris 8, 9, 10 (SPARC/x86)	32-bit, 64-bit
Sun Solaris 7 (SPARC)	32-bit
IBM AIX 5.2, 5.3, 6.1	32-bit, 64-bit

Supported Novell NetWare Platforms

Table 1-3 *Supported Novell NetWare Platforms*

Version	Bit Spec
Novell NetWare 5.1 (SP8)	32-bit, 64-bit
Novell NetWare 6.0 (SP5)	32-bit, 64-bit
Novell NetWare 6.5 (SP6)	32-bit, 64-bit

IMPORTANT: Download and install the following Novell Clients on the PlateSpin Recon server computer to enable monitoring of NetWare computers. The client downloads are free. When installing, select *Custom Installation* and install the Novell Modular Authentication Service (NMAS) client.

- ♦ Novell Client 4.91 SP5 (<http://download.novell.com/Download?buildid=qmMAWSRy5q4~>)
- ♦ Novell Client 4.91 Post-SP5 NWFS.SYS (<http://download.novell.com/Download?buildid=m7NbjoMFSxU&donotredirect=true>)

Installation of the Novell Client impacts the standard Windows logon. See the Novell Client documentation for more information.

Only the CPU and disk utilization counters are supported for Netware 5.1 and 6.0. Only the CPU, Disk and memory utilization counters are supported for Netware 6.5 and higher.

Supported Hypervisors

Table 1-4 *Supported Hypervisors*

Version	Bit Spec
Citrix XenServer 4.1, 5.0	32-bit, 64-bit
Windows Server 2008 (SP1, SP2) Hyper-V	64-bit
Windows Server 2008 R2 Hyper-V	64-bit

Version	Bit Spec
Sun Solaris 10 Zones on Sun Solaris 10 Container	32-bit, 64-bit
VMware ESX Server 2.5.3, 2.5.4, 2.5.5, 3.0.0, 3.0.1, 3.0.2, 3.0.3, 3.5.0, 3i, 4.0, 4i	32-bit, 64-bit
VMware ESX Server 4.1,4.1i	64-bit
VMware GSX Server 3.2 (for Windows)	32-bit, 64-bit
VMware 1.x+ (for Windows)	32-bit, 64-bit
Virtual Server 2005 (version 1.0) All editions.	
Virtual Server 2005 R2 (SP1)	
VMware Virtual Center 2.0.1, 2.0.2, 2.5, 4, 4.1	
Xen on Novell SLES 10.0, 10.2, 10.3, 11	32-bit, 64-bit

1.3.2 Supported PlateSpin Portability Suite Consolidation Sources and Targets

Consolidation scenarios for workloads with the supported operating systems and targets can be implemented by using PlateSpin Portability Suite. For detailed information about the Portability Suite supported source and target operating systems, see the [PlateSpin Portability Suite documentation Web site \(http://www.novell.com/documentation/platespin_portabilitysuite_810/\)](http://www.novell.com/documentation/platespin_portabilitysuite_810/).

For more detailed information about source and target operating systems for PlateSpin Portability Suite conversions, see Knowledge Base article [Q7005871 \(http://www.novell.com/support/viewContent.do?externalId=7005871\)](http://www.novell.com/support/viewContent.do?externalId=7005871).

For more information about implementing consolidation scenarios with portability suite, see [Section 6.8, “Implementing Scenarios with Portability Suite,” on page 151](#).

1.4 Application Scenario

PlateSpin Recon and PlateSpin Portability Suite combine to form the only suite that automates the entire capacity planning and server consolidation project from start to finish. The following table illustrates how to assess, design, and implement a solution for optimizing and consolidating the data center.

Table 1-5 *Data Center Optimization and Consolidation Workflow*

Action	Process	Description	Solution
Assess	1. Inventory	<ul style="list-style-type: none">♦ Remotely discover server hardware and software assets across the data center♦ Gather detailed information for each server such as operating system, installed and running applications or services, patch levels, CPU, memory, network and disk resources♦ Connect to the network, then select relevant computers or domains; data on hardware and resources is collected in a centralized database♦ Organize servers by predefined and user-defined groupings for organized analysis	PlateSpin Recon: Remotely and agentlessly identify server resources and workloads in the data center and create a complete picture of hardware and software inventory, utilization levels, and workload sizes

Action	Process	Description	Solution
Assess	2. Collect	<ul style="list-style-type: none"> ♦ Identify workloads and collect utilization data over days, weeks or months to determine workload sizes and appropriate resource matching ♦ Create dynamic reports to identify workload and resource mismatches and identify candidate workloads for consolidation or resource upgrading ♦ Rank and sort server workloads by CPU, network, disk, memory inventory, and/or performance ♦ View summary workload sizing data or drill down to point-in-time granular levels ♦ Collect and analyze data from multiple sites to create a more complete picture of the data center ♦ Export the data to third-party tools or copy visual graphs to business presentations 	PlateSpin Recon: Remotely and agentlessly identify server resources and workloads in the data center and create a complete picture of hardware and software inventory, utilization levels, and workload sizes

Action	Process	Description	Solution
Design	3. Analyze	<ul style="list-style-type: none"> ◆ Identify virtualization candidates ◆ Create scenarios for distributing workloads across servers to maximize utilization ◆ Balance application workloads over time and across hardware resources to minimize resource contention ◆ Use “what-if” modeling to determine different combinations of hardware and virtual hosts required for the project ◆ Generate consolidation or disaster recovery plans with detailed project, scenario, and workload assignment reports and charts ◆ Compare different scenarios based on total cost of ownership, consolidation ratio, rack space, and power needs 	PlateSpin Recon: Generate consolidation scenario plans to optimally combine and allocate workloads to new virtual hosts
Implement	4. Consolidate	<ul style="list-style-type: none"> ◆ Accelerate the server consolidation project by streaming physical servers into virtual environments such as VMware ESX Server, VMware Server, and Microsoft Virtual Server ◆ Automate migrations between dissimilar physical servers or from physical servers to blades 	PlateSpin Portability Suite: Automatically stream data, applications, and operating systems to any physical or virtual platform over the network

Action	Process	Description	Solution
Manage	5. Monitor	<ul style="list-style-type: none"> ♦ Keep track of workload performance ♦ Identify over- and under-utilized workloads for ongoing optimization 	PlateSpin Recon
	6.Chargeback	<ul style="list-style-type: none"> ♦ Assign customized raters to machines ♦ Measure resource usage and bill business owners for usage ♦ Create reports on usage and cost 	
	7. Manage Virtual Infrastructure	<ul style="list-style-type: none"> ♦ Monitor VM creation ♦ Report on VM growth 	

1.5 PlateSpin Recon Terminology

- ♦ **Consolidation candidate:** A workload on the network that you might want to move to a virtual machine.
- ♦ **Data Center assets:** A collective set of physical and virtual machines, and physical and virtual workloads.
- ♦ **Job:** A task performed in PlateSpin Recon.
- ♦ **Protection candidate:** A workload to create standby VM duplicates for disaster recovery.
- ♦ **Rater:** A metric used to assign a billable charge to the use of a resource or part of a resource over a set amount of time.
- ♦ **Scheduled job:** An activity set to occur at a given time and frequency.
- ♦ **Virtual host:** A physical machine that has the virtualization software installed. For example, VMware ESX.
- ♦ **Workload:** A physical or virtual machine.

Overview of PlateSpin Recon Client Views

2

Novell PlateSpin Recon Client provides the following views to help you organize, monitor, view, and manage information about [data center assets](#):

- ♦ [Section 2.1, “Data Center Explorer,” on page 23](#)
- ♦ [Section 2.2, “Report Explorer,” on page 24](#)
- ♦ [Section 2.3, “Projects Explorer,” on page 25](#)
- ♦ [Section 2.4, “Raters Explorer,” on page 25](#)
- ♦ [Section 2.5, “Chart Viewer,” on page 26](#)
- ♦ [Section 2.6, “Job Explorer,” on page 26](#)
- ♦ [Section 2.7, “Scheduled Job Explorer,” on page 27](#)
- ♦ [Section 2.8, “Log Explorer,” on page 27](#)


2.1 Data Center Explorer

The Data Center Explorer is one of several major PlateSpin Recon *views*. It is used to organize, monitor, view, and manage information about [data center assets](#). PlateSpin Recon uses information about them to display dynamic charts, generate performance reports, and generate performance forecasts.

The Data Center Explorer is the central point for:

- ♦ Interacting with physical and virtual machines in the data center
- ♦ Discovering and inventorying hardware and workloads
- ♦ Adding, monitoring, filtering, grouping, or deleting servers

Figure 2-1 *Data Center Explorer*




For more information about working with Data Center Explorer, see [Chapter 3, “Working with Data Center Explorer,”](#) on page 29.

2.2 Report Explorer

The Report Explorer allows you to organize and generate reports on the inventory and monitoring data. The Report Explorer hierarchy consists of template nodes that have various types of report templates.

Figure 2-2 *Report Explorer*




For more information about working with Report Explorer, see [Chapter 5, “Generating and Viewing Charts and Reports,”](#) on page 85.

2.3 Projects Explorer

Projects are displayed in the Project Explorer.

Figure 2-3 *Project Explorer Tree*



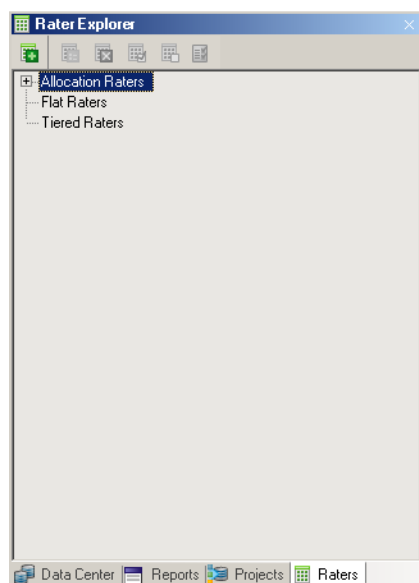
The project, scenario, and consolidated workloads are arranged hierarchically; that is, workloads are listed under a server, servers are listed under a scenario, and scenarios are listed under a project. All selected candidates are listed under workloads inside the project.

For more information about working with Project Explorer, see [Chapter 6, “Server Consolidation,” on page 121](#).

2.4 Raters Explorer

Allows you to set up a chargeback system.

Figure 2-4 *Raters Explorer*



For more information about working with Raters Explorer, see [Section 7.2, “Using Raters,” on page 155](#).

2.5 Chart Viewer

Charts provide a graphical representation of workload and utilization data.

To view a chart for a given server, select the server in the Data Center Explorer or click a cell belonging to the server in a report view.


2.6 Job Explorer

Whenever a task is performed in PlateSpin Recon, a job is created for the task. Information on running, paused or completed jobs is shown in the Job Explorer, located at the bottom of the PlateSpin Recon window, by default.

If a job consists of several tasks, you can expand the job to see all of its tasks and the status of each task by clicking the plus sign next to the job description.

You can also see the steps that make up each task by selecting the task. The steps are displayed in the panel on the right of the Job Explorer.

Figure 2-5 *Task steps displayed in the Job Explorer*



Messages about the progress of a job are captured in a log file.

Right-click any running job entry to display a pop-up menu, or use the toolbar buttons, and you can *Pause*, *Resume*, or *Abort* any running job. The *Archive* command moves jobs to the Archived Jobs list. These jobs can be subsequently viewed by selecting *Tools > Archived Jobs*. You can view the step logs by right-clicking on a step and selecting *View Logs*. The Log Viewer dialog is displayed.


NOTE: The default setting is to purge logs and jobs every 14 days. To change this setting, select *Tools > Options*. The Options dialog box is displayed. Double-click *Advanced* under the *Maintenance* category on the Server page.

2.7 Scheduled Job Explorer

A scheduled job is an activity set to occur at a given time and frequency. For example, it is possible to generate a Network Usage Report every morning at 9:00 AM for the next five days.

A scheduled job is comprised of two components: the action to be performed (the job), and the schedule for performing the activity. The schedule is a reusable component that can be shared amongst multiple jobs, such as creating three reports at 9:00 AM. The Scheduled Job Explorer can be used to view the status of the completed jobs.

Figure 2-6 Job status displayed in the Scheduled Job Explorer



Next Occurrence	Description	Status
10/19/2008 1:00:00 AM	Maintenance Purge (Daily 01:00)	Active

Start Time	Duration	Description	Status	Completed
10/18/2008 1:00 AM	00:00:15	Maintenance Purge	Completed	10/18/2008 1:00 AM
10/17/2008 1:00 AM	00:00:12	Maintenance Purge	Completed	10/17/2008 1:00 AM
10/16/2008 1:00 AM	00:00:09	Maintenance Purge	Completed	10/16/2008 1:00 AM
10/15/2008 1:00 AM	00:00:16	Maintenance Purge	Completed	10/15/2008 1:00 AM
10/14/2008 1:00 AM	00:00:11	Maintenance Purge	Completed	10/14/2008 1:00 AM
10/13/2008 1:00 AM	00:00:05	Maintenance Purge	Completed	10/13/2008 1:00 AM

Errors or warnings, if any, are also displayed. In addition, scheduled jobs can be run manually, or can have their recurrence changed.

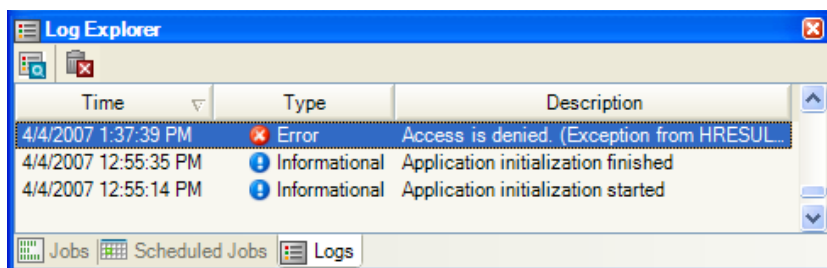
Right-click any running job entry to display a pop-up menu, or use the toolbar buttons, and you can *Suspend*, *Resume*, or *Delete* any scheduled job. The *Edit Recurrence* command allows you to create or edit a scheduled job recurrence pattern. The *Run Now* command allows you to immediately run the scheduled job.

2.8 Log Explorer

The Log Explorer allows you to view general information, warnings, or error messages. Messages that are displayed here are not related to any particular job or server. Right-click any log entry to display a pop-up menu where you can select commands to View Details of a log entry, Delete a log entry, or View XML format of a log entry.

To simultaneously delete multiple logs, select the logs using the Shift or Ctrl keys and press the Delete key or click Delete.

Figure 2-7 Working with the Log Explorer



Working with Data Center Explorer

3

This section provides information about the Novell PlateSpin Recon Data Center Explorer and the features accessible from it.

- ♦ [Section 3.1, “Understanding the Data Center Explorer,” on page 29](#)
- ♦ [Section 3.2, “Discovering and Inventorying Data Center Resources,” on page 30](#)
- ♦ [Section 3.3, “Working with Filters,” on page 52](#)
- ♦ [Section 3.4, “Working with Groups,” on page 52](#)
- ♦ [Section 3.5, “Working with Sites,” on page 53](#)
- ♦ [Section 3.6, “Working with PlateSpin Recon Snapshots,” on page 55](#)
- ♦ [Section 3.7, “Working with Remote Data,” on page 60](#)
- ♦ [Section 3.8, “Working with Data Collected by Other PlateSpin Recon Instances,” on page 60](#)
- ♦ [Section 3.9, “Working with Data Collected by Other Performance Monitoring Tools,” on page 62](#)
- ♦ [Section 3.10, “Using HP Reporter Data with PlateSpin Recon,” on page 66](#)
- ♦ [Section 3.11, “Using Performance Data Import to Import Generic Data into PlateSpin Recon,” on page 70](#)
- ♦ [Section 3.12, “Collecting Disk Utilization Counters for Windows Machines,” on page 74](#)

3.1 Understanding the Data Center Explorer

The Data Center Explorer is one of several major PlateSpin Recon *views*. It is used to organize, monitor, view, and manage information about [data center assets](#). PlateSpin Recon uses information about them to display dynamic charts, generate performance reports, and generate performance forecasts.

In the Data Center Explorer, physical and virtual servers are represented in a hierarchical tree structure.

Table 3-1 *Data Center Explorer Nodes*

Nodes	Description
All	Lists all systems
Discovered	Lists servers that have been discovered but not yet inventoried
Filters	Contains servers listed by built-in criteria, such as domain, operating system, or monitoring status
Groups	Contains manually grouped servers according to any specific data segregation requirements
Sites	Contains servers by data site, including a default local site and others you create to work with data about servers monitored by other instances of PlateSpin Recon, or imported from third-party performance monitoring tools

For more information about the Data Center Explorer icons and commands, see [Appendix D, “Quick Reference to PlateSpin Recon Icons and Commands,”](#) on page 211.

3.2 Discovering and Inventorying Data Center Resources

PlateSpin Recon gets information about [data center assets](#) in two stages.

- ♦ **Discover:** Discovering machines in the network retrieves information such as hostnames, IP addresses, MAC addresses, and port information.
- ♦ **Inventory:** Inventorying the previously discovered machines retrieves deeper data for each server, such as operating system, installed applications, services, processor speed, memory, network, and disk resources.

After machines are inventoried, they can be monitored and included in all report types, regardless of which method is used for inventory.

- ♦ [Section 3.2.1, “Inventory and Discovery Methods,”](#) on page 30
- ♦ [Section 3.2.2, “Inventory Credentials,”](#) on page 42
- ♦ [Section 3.2.3, “Inventory Options,”](#) on page 44
- ♦ [Section 3.2.4, “Discovery Options,”](#) on page 45
- ♦ [Section 3.2.5, “Viewing Inventory Data and Monitoring Servers,”](#) on page 47

PlateSpin identifies a machine through unique attributes such as the IP address. If you change the IP address of the machine after inventorying a machine, you must reinventory the machine.

For information about ports used by PlateSpin Recon to inventory servers, see Knowledge Base article [Q7920571](http://www.novell.com/support/viewContent.do?externalId=7920571) (<http://www.novell.com/support/viewContent.do?externalId=7920571>).






3.2.1 Inventory and Discovery Methods

A common way to use discovery and inventory is to discover a large number of servers on the network. These servers appear under the *Discovered* node in the Data Center Explorer. Next, select the discovered servers and inventory them. Finally, select which inventoried servers you want to monitor. It is also possible to discover, inventory, and start monitoring servers in a single step.

You can add servers to the Data Center Explorer by using the *Inventory* button .


There are seven inventory methods available. Each inventory method is unique and has different requirements.

Table 3-2 *Inventory Methods*

Discover and Inventory	Discover and Optionally Inventory
 Machines (page 31)	 Windows Domain (page 35)
 Virtual Center (page 31)	 IP Range (page 36)
 From a CSV File (page 33)	 Subnet (page 38)
	 SNMP (page 40)

Machines


To discover and inventory physical and virtual servers:

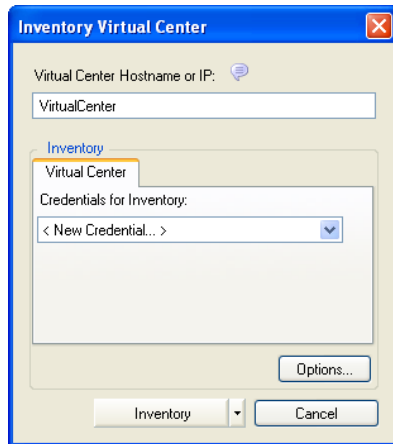
- 1 Click the *Inventory* button  > *Machines* or click the *Tools* menu > *Inventory* > *Machines*.
The Inventory Machines dialog box is displayed.
- 2 In *Hostname or IP of the machine(s) to inventory*, specify the hostname or the IP address of the machines you want to inventory.
You can specify multiple hostnames or IP addresses separated with a comma; however, they need to all use the same credentials.
- 3 Specify credentials for the machines you add in [Step 2](#). For more information on how to specify a new credential, see [Section 3.2.2, “Inventory Credentials,” on page 42](#).
This enables the *Inventory* button.
- 4 (Optional) Start monitoring the servers when they are inventoried and/or add them to existing groups. For more information, see [Section 3.2.3, “Inventory Options,” on page 44](#).
- 5 Configure the schedule when you want to perform the inventory.
 - ♦ To immediately start the inventory, do one of the following:
 - ♦ Click *Now* in the *Inventory* drop-down list.
 - ♦ Click *Schedule* in the *Inventory* drop-down list, then click *Now* > *OK*.
 - ♦ To perform inventory at a later time, set up a schedule:
 1. In the *Inventory* drop-down list, click *Schedule*.
The Scheduling Options dialog box is displayed.
 2. Select one of the following schedules
 - ♦ “[Later](#)” on page 209
 - ♦ “[Recurrence](#)” on page 209
 3. Fill in the fields.
For details about the fields, click the appropriate schedule listed above.
- 6 Click *Inventory*.

Virtual Center

Inventorying a Virtual Center involves collecting information about all the servers and virtual machines in the Virtual Center. This is much more efficient than discovering and inventorying all the servers and virtual machines individually.

To discovery and inventory Virtual Center servers:

- 1 Click the *Inventory* button  > *Virtual Center* or click the *Tools* menu > *Inventory* > *Virtual Center*.
The Inventory Virtual Center dialog box is displayed.



- 2 Specify the hostname or the IP address of the Virtual Center you want to inventory.
- 3 Specify credentials for the machines you add in [Step 2](#). For more information on how to specify a new credential, see [Section 3.2.2, “Inventory Credentials,” on page 42](#).

This enables the *Inventory* button.

- 4 (Optional) Start monitoring the servers when they are inventoried and/or add them to existing groups. For more information, see [Section 3.2.3, “Inventory Options,” on page 44](#).
- 5 Configure the schedule when you want to perform the inventory.
 - ♦ To immediately start the inventory, do one of the following:
 - ♦ Click *Now* in the *Inventory* drop-down list.
 - ♦ Click *Schedule* in the *Inventory* drop-down list, then click *Now > OK*.
 - ♦ To perform inventory at a later time, set up a schedule:
 1. In the *Inventory* drop-down list, click *Schedule*.
The Scheduling Options dialog box is displayed.
 2. Select one of the following schedules
 - ♦ [“Later” on page 209](#)
 - ♦ [“Recurrence” on page 209](#)
 3. Fill in the fields.

For details about the fields, click the appropriate schedule listed above.

- 6 Click *Inventory*.

Virtual Center contents appear in the Data Center Explorer the same way they are organized in the Virtual Center, with all nodes, clusters, resource pools, and virtual machine nodes. The machines do not appear under *All* in the Data Center Explorer until they are inventoried.

When using per core licensing, you must first inventory a Virtual Center before inventorying ESX Servers that belong to it, or the virtual machines on those ESX Servers. For more information, see [Per Core](#) in the [Novell PlateSpin Recon 3.7.2 Getting Started Guide](#).

From a CSV File


Inventorying machines from a CSV file allows you to add servers whose information resides in a comma-separated value file. Use the following format to list the machines in the file:

Address, [Username], [Password], [Description], [Inventory Type]

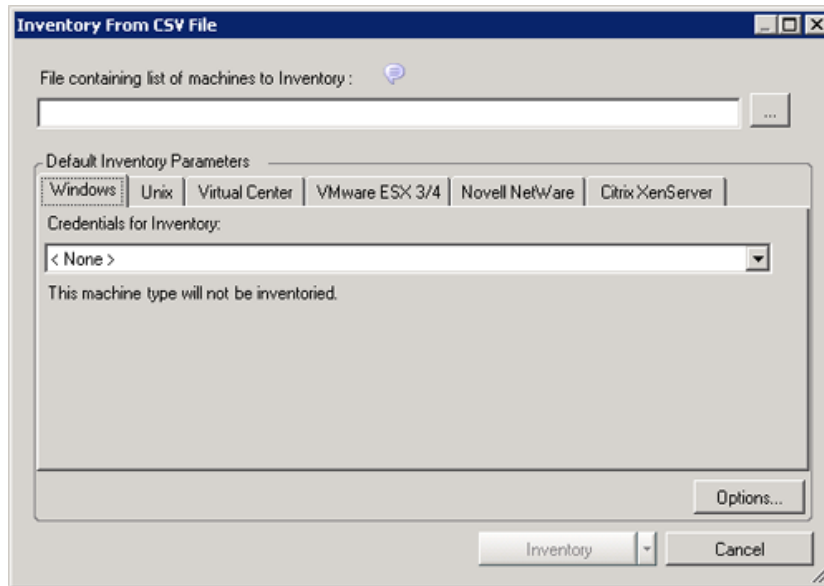
Table 3-3 CSV File Requirements

Field	Required	Description
Address	Yes	IP address or hostname of the machine.
Username	Optional	Username for the credential. If this is left blank, the default credentials are used.
Password	Optional	Password of the credential. This is required only when creating a credential.
Description	Optional	Description for the credential. If the entry for a machine in the CSV file contains username and password, the credential description must be unique.
Inventory type	Optional	Can be any of the following values. Values must match case and spacing as shown below: <ul style="list-style-type: none">♦ UNIX♦ VirtualCenter♦ VMwareESX3♦ VMwareESX4♦ NetWare♦ CitrixXenServer♦ Windows If the value is unspecified, Windows and UNIX are the defaults and these are determined by the formats of the credentials that you specify in the CSV file: <ul style="list-style-type: none">♦ [user],[password] is for UNIX For example, 10.0.0.0,root,password,description♦ [domain\user],[password] is for Windows For example, 10.0.0.0,localhost\administrator,password,description1 This field is also expected to be at the end of the CSV file.

To discover and inventory machines from a CSV file:

- 1 Click the *Inventory* button  > *From CSV File* or click the *Tools* menu > *Inventory* > *From CSV File*.

The Inventory from CSV dialog box is displayed.



- 2 Type the path (or use the adjacent browse button) to locate a CSV file containing servers to inventory.
- 3 (Conditional) If the CSV file does not contain the credentials for all the machines, specify the credentials for the machines in the Default Inventory Parameter panel.

NOTE: The *Default Inventory Parameters* area is disabled if the CSV file contains credentials for all its machines.

For more information on how to specify a new credential, see [Section 3.2.2, “Inventory Credentials,”](#) on page 42.

This enables the *Inventory* button.

- 4 (Optional) Start monitoring the machine when they are inventoried and/or add them to existing groups. For more information, see [Section 3.2.3, “Inventory Options,”](#) on page 44.
- 5 Configure the schedule when you want to perform the inventory.
 - ♦ To immediately start the inventory, do one of the following:
 - ♦ Click *Now* in the *Inventory* drop-down list.
 - ♦ Click *Schedule* in the *Inventory* drop-down list, then click *Now > OK*.
 - ♦ To perform inventory at a later time, set up a schedule:
 1. In the *Inventory* drop-down list, click *Schedule*.
The Scheduling Options dialog box is displayed.
 2. Select one of the following schedules
 - ♦ [“Later”](#) on page 209
 - ♦ [“Recurrence”](#) on page 209
 3. Fill in the fields.
For details about the fields, click the appropriate schedule listed above.
- 6 Click *Inventory*.

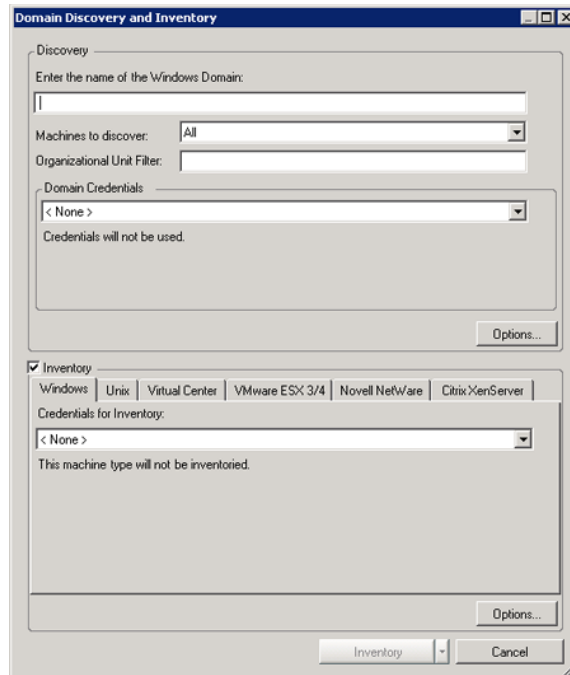
Windows Domain

Inventorying machines by using the Windows Domain option allows you to add all the servers discovered in the specified Windows Domain.

To discover and optionally inventory servers in a specified Windows Domain:

- 1 Click the *Inventory* button  > *Windows Domain* or click the *Tools* menu > *Inventory* > *Windows Domain*.

The Domain Discovery and Inventory dialog box is displayed.



- 2 Type a Windows Domain name.
- 3 In the *Machines to Discover* drop-down list, select which machines to discover. The available options are *All Machines*, *Servers*, or *Workstations* (NT4 Workstation, Windows XP).
- 4 (Conditional) If you want to discover a specific part of a domain, configure the Organizational Unit (OU) filter. For more information, see [“Domain Discovery:” on page 188](#).
- 5 Click the *Options* button to configure discovery options. For more information, see [Section 3.2.4, “Discovery Options,” on page 45](#).
If you want to discover machines only, continue with [Step 6](#). If you want to discover and inventory machine, skip to [Step 7](#).
- 6 (Conditional) If you want to only discover machines without inventorying machine details:
 - 6a Deselect the *Inventory* option.
The *Inventory* option is selected by default.
 - 6b Click *Discover*.
- 7 (Conditional) If you want to discover and inventory machines, additionally configure the following settings:
 - 7a Specify credentials for the machines you added in [Step 3](#).

For more information on how to specify a new credential, see [Section 3.2.2, “Inventory Credentials,” on page 42](#).

This enables the *Inventory* button.

7b (Optional) Start monitoring the machine when they are inventoried and/or add them to existing groups. For more information, see [Section 3.2.3, “Inventory Options,” on page 44](#).

7c Configure the schedule when you want to perform the inventory.


- ♦ To immediately start the inventory, do one of the following:
 - ♦ Click *Now* in the *Inventory* drop-down list.
 - ♦ Click *Schedule* in the *Inventory* drop-down list, then click *Now > OK*.
- ♦ To perform inventory at a later time, set up a schedule:
 1. In the *Inventory* drop-down list, click *Schedule*.
The Scheduling Options dialog box is displayed.
 2. Select one of the following schedules
 - ♦ [“Later” on page 209](#)
 - ♦ [“Recurrence” on page 209](#)
 3. Fill in the fields.

For details about the fields, click the appropriate schedule listed above.

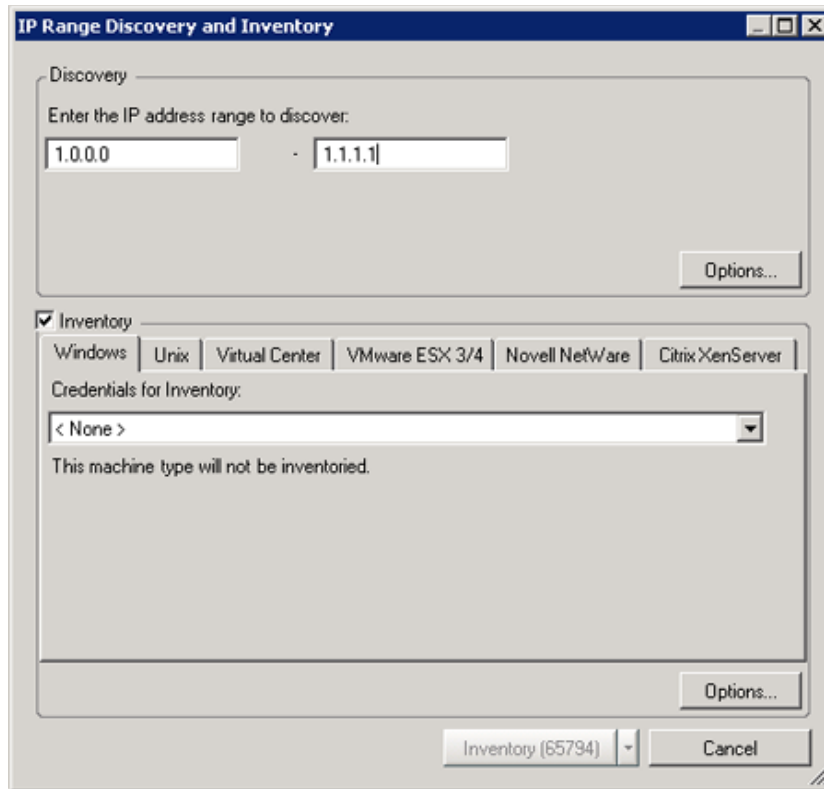
7d Click *Inventory*.

IP Range

To discovery and optionally inventory servers within a specified IP address range:

- 1** Click the *Inventory* button  > *IP Range* or click the *Tools* menu > *Inventory > IP Range*.

The IP Range Discovery and Inventory dialog box is displayed.



- 2 Configure the following settings in the Discovery pane:
 - 2a Specify the IP address range of the servers to discover.
 - 2b Click the *Options* button to configure discovery options. For more information, see [Section 3.2.4, “Discovery Options,” on page 45](#).
If you want to discover machines only, continue with [Step 3](#). If you want to discover and inventory machine, skip to [Step 4](#).
- 3 (Conditional) If you want to only discover machines without inventorying machine details:
 - 3a Deselect the *Inventory* option.
The *Inventory* option is selected by default.
 - 3b Click *Discover*.
The numbers that are displayed in the *Discover* button indicate how many individual IP addresses are contained within the range you entered
- 4 (Conditional) If you want to discover and inventory machines, additionally configure the following settings:
 - 4a Specify credentials for the machines you added in [Step 2](#).
For more information on how to specify a new credential, see [Section 3.2.2, “Inventory Credentials,” on page 42](#).
This enables the *Inventory* button.
 - 4b (Optional) Start monitoring the machine when they are inventoried and/or add them to existing groups. For more information, see [Section 3.2.3, “Inventory Options,” on page 44](#).

4c Configure the schedule when you want to perform the inventory.

- ♦ To immediately start the inventory, do one of the following:
 - ♦ Click *Now* in the *Inventory* drop-down list.
 - ♦ Click *Schedule* in the *Inventory* drop-down list, then click *Now* > *OK*.
- ♦ To perform inventory at a later time, set up a schedule:
 1. In the *Inventory* drop-down list, click *Schedule*.
The Scheduling Options dialog box is displayed.
 2. Select one of the following schedules
 - ♦ “*Later*” on page 209
 - ♦ “*Recurrence*” on page 209
 3. Fill in the fields.


For details about the fields, click the appropriate schedule listed above.

4d Click *Inventory*.


The numbers that are displayed on the *Inventory* button indicate how many individual IP addresses are contained within the range you entered.

Subnet

To discovery and optionally inventory servers in a specified subnet:

- 1 Click the *Inventory* button  > *Subnet* or click the *Tools* menu > *Inventory* > *Subnet*.

The Subnet Discovery and Inventory dialog box is displayed.



The dialog box is titled "Subnet Discovery and Inventory". It has two main sections: "Discovery" and "Inventory".

Discovery Section:

- Network address: 1.0.0.0
- Mask bits: 24
- Subnet mask: 255.255.255.0
- Options... button

Inventory Section:

- Checkmark icon next to "Inventory" label.
- Tabs: Windows, Unix, Virtual Center, VMware ESX 3/4, Novell NetWare, Citrix XenServer.
- Credentials for Inventory: < None >
- Text: This machine type will not be inventoried.
- Options... button


Bottom Section:

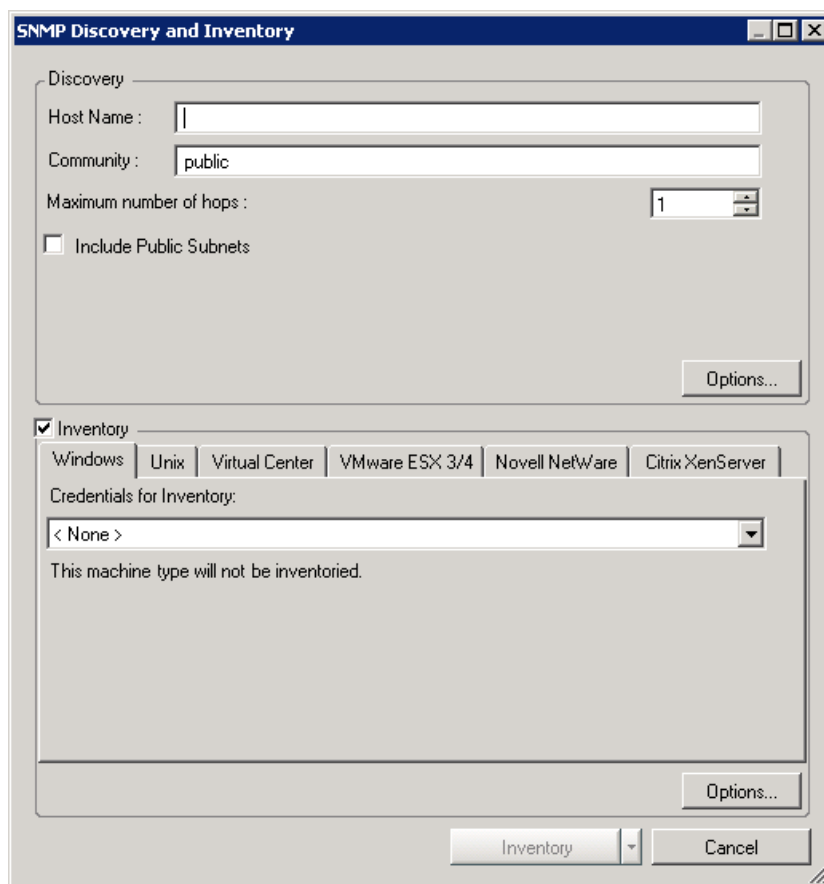
- Inventory [256] button
- Cancel button

- 2** Configure the following settings in the Discovery pane:
- 2a** Specify the network address and the mask bits of the subnet to be discovered.
 - 2b** Click the *Options* button to configure discovery options. For more information, see [Section 3.2.4, “Discovery Options,” on page 45](#).
If you want to discover machines only, continue with [Step 3](#). If you want to discover and inventory machine, skip to [Step 4](#).
- 3** (Conditional) If you want to only discover machines without inventorying machine details:
- 3a** Deselect the *Inventory* option.
The *Inventory* option is selected by default.
 - 3b** Click *Discover*.
The numbers that are displayed in the *Discover* button indicate how many individual IP addresses are contained in the subnet you specified.
- 4** (Conditional) If you want to discover and inventory machines, additionally configure the following settings in the Inventory pane:
- 4a** Specify credentials for the machines you added in [Step 2](#).
For more information on how to specify a new credential, see [Section 3.2.2, “Inventory Credentials,” on page 42](#).
This enables the *Inventory* button.
 - 4b** (Optional) Start monitoring the machine when they are inventoried and/or add them to existing groups. For more information, see [Section 3.2.3, “Inventory Options,” on page 44](#).
 - 4c** Configure the schedule when you want to perform the inventory.
 - ♦ To immediately start the inventory, do one of the following:
 - ♦ Click *Now* in the *Inventory* drop-down list.
 - ♦ Click *Schedule* in the *Inventory* drop-down list, then click *Now > OK*.
 - ♦ To perform inventory at a later time, set up a schedule:
 1. In the *Inventory* drop-down list, click *Schedule*.
The Scheduling Options dialog box is displayed.
 2. Select one of the following schedules
 - ♦ [“Later” on page 209](#)
 - ♦ [“Recurrence” on page 209](#)
 3. Fill in the fields.
For details about the fields, click the appropriate schedule listed above.
 - 4d** Click *Inventory*.
The numbers that are displayed on the *Inventory* button indicate how many individual IP addresses are contained in the subnet you specified.

SNMP

To discover and optionally inventory servers machines by using Simple Network Management Protocol:

- 1 Click the *Inventory* button  > *SNMP* or click the *Tools* menu > *Inventory* > *SNMP*.
The SNMP Discovery and Inventory dialog box is displayed.



- 2 Configure the following settings in the Discovery pane:
 - 2a Specify the hostname.
 - 2b Specify the SNMP community name. By default, it is public.
 - 2c In *Maximum number of Hops*, specify how many hops to make beyond the server from which the discovery is launched. Select whether to include *Public Subnets*.
 - 2d Click the *Options* button to configure discovery options. For more information, see [Section 3.2.4, “Discovery Options,” on page 45](#).
If you want to discover machines only, continue with [Step 3](#). If you want to discover and inventory machine, skip to [Step 4](#).
- 3 (Conditional) If you want to only discover machines without inventorying machine details:
 - 3a Deselect the *Inventory* option.
The *Inventory* option is selected by default.
 - 3b Click *Discover*.

- 4** (Conditional) If you want to discover and inventory machines, additionally configure the following settings in the Inventory pane;
- 4a** Specify credentials for the machines you added in [Step 2](#).
- For more information on how to specify a new credential, see [Section 3.2.2, “Inventory Credentials,” on page 42](#).
- This enables the *Inventory* button.
- 4b** (Optional) Start monitoring the machine when they are inventoried and/or add them to existing groups. For more information, see [Section 3.2.3, “Inventory Options,” on page 44](#).
- 4c** Configure the schedule when you want to perform the inventory.
- ♦ To immediately start the inventory, do one of the following:
 - ♦ Click *Now* in the *Inventory* drop-down list.
 - ♦ Click *Schedule* in the *Inventory* drop-down list, then click *Now > OK*.
 - ♦ To perform inventory at a later time, set up a schedule:
 1. In the *Inventory* drop-down list, click *Schedule*.
The Scheduling Options dialog box is displayed.
 2. Select one of the following schedules
 - ♦ “[Later](#)” on page 209
 - ♦ “[Recurrence](#)” on page 209
 3. Fill in the fields.
For details about the fields, click the appropriate schedule listed above.
- 4d** Click *Inventory*.

3.2.2 Inventory Credentials

To inventory discovered machines, select the *Inventory* check box in the Discovery and Inventory dialog boxes and then provide login credentials for Windows, UNIX, Virtual Center, VMware ESX, Novell NetWare, or Citrix XenServer machines, depending on the type of machines in use.

- ♦ “[Specifying Credentials for Windows or VMware ESX 3/4](#)” on page 42
- ♦ “[Specifying Credentials for UNIX or Citrix XenServer](#)” on page 43
- ♦ “[Specifying Credentials for Virtual Center](#)” on page 43
- ♦ “[Specifying Credentials for Novell NetWare](#)” on page 43

For every machine in the list or range, PlateSpin Recon tries each type of credential provided until it finds one that matches. If the credentials are not valid, the machine is not inventoried.

Specifying Credentials for Windows or VMware ESX 3/4

Do the following in the Discovery and Inventory dialog box:

- 1** In the *Inventory* panel, click the *Windows* or *VMWare ESX 3/4* tab depending on the type of machines in use.

- 2 In the *Credentials for Inventory* drop-down list, select credentials.

To add a new credential, click *<New Credential>*, and configure the Create Stored Credentials dialog box. For detailed information about adding credentials, see [Section 4.2, “Credentials Manager,” on page 79](#).

Specifying Credentials for UNIX or Citrix XenServer

Do the following in the Discovery and Inventory dialog box:

- 1 In the *Inventory* panel, click the *UNIX* or *Citrix XenServer* tab depending on the type of machines in use.
- 2 In the *Credentials for Inventory* drop-down list, select credentials.

To add a new credential, click *<New Credential>*, and configure the Create Stored Credentials dialog box. For detailed information about adding credentials, see [Section 4.2, “Credentials Manager,” on page 79](#).

- 3 Select the *Use Sudo* check box to use super user privileges for connections.

PlateSpin Recon allows inventorying of Linux (RHEL, SLES, Ubuntu, CentOS, SLES with Xen) and Citrix XenServer machines by using sudo user credentials. However, inventorying of the AIX and Solaris machines by using sudo user credentials is not supported. PlateSpin Recon users can inventory machines by using sudo user credentials only if they have the permission to run commands on the target machine. For more information on how to grant permission to perform inventory by using sudo user credentials, see [Section B.9, “Sudo Configuration to Inventory and Monitor Linux Machines,” on page 207](#).

Local policies can sometimes prevent you from establishing an SSH session as a `root` user. A user with sudo privileges can get around this restriction by establishing an SSH session as a non-root user and then using `sudo` to temporarily “switch” to `root` user privileges. To do this, create the credentials using that user’s username and password. PlateSpin Recon uses the `sudo` command to perform any operations that require `root` privileges.

Specifying Credentials for Virtual Center

Do the following in the Discovery and Inventory dialog box:

- 1 In the *Inventory* panel, click the *Virtual Center* tab.
- 2 In the *Credentials for Inventory* drop-down list, select credentials.

To add a new credential, click *<New Credential>*, and configure the Create Stored Credentials dialog box. For detailed information about adding credentials, see [Section 4.2, “Credentials Manager,” on page 79](#).

For inventorying a Virtual Center, credentials must have a minimum read-only user status for the Virtual Center.

Virtual Center automatically includes as users any user with Administrator privileges for the machine where Virtual Center is installed.

Specifying Credentials for Novell NetWare

- 1 In Novell eDirectory on the NetWare server you want to inventory, create a user with the same login credentials as the PlateSpin Recon server administrator.

- 2 Associate the local security policies to the PlateSpin Recon server administrator:
 - 2a On the PlateSpin Recon server, click the desktop *Start* menu > *Settings* > *Control Panel*.
 - 2b Double click *Administrative Tools* > *Local Security Policy*.

The Local Security Settings window is displayed.
 - 2c Click *Local Policies*, then double-click *User Rights Assignment*.
 - 2d Associate the administrator with following policies:
 - ♦ Act as part of the operating system
 - ♦ Create a token object
 - ♦ Log on as a batch job
 - ♦ Log on as a service
 - ♦ Replace a process level token
- 3 Ensure that the recommended Novell Client is installed on the PlateSpin Recon server.

For more information on the recommended Novell Clients, see [“Supported Novell NetWare Platforms”](#) on page 16.
- 4 Through the Novell Client, log in to the PlateSpin Recon server as the administrator user created in [Step 1](#).
- 5 Launch the PlateSpin Recon Client.
- 6 In the *Inventory* drop-down list (located at the top left corner), click *Machines*.

The Inventory Machines dialog box is displayed.
- 7 Click the *Novell NetWare* tab.
- 8 In the *Credentials for Inventory* drop-down list, select <*New Credential*>.

The Create Stored Machines dialog box is displayed.
- 9 Specify the login credentials of the NetWare administrator, and select *Novell NetWare Inventory* as the validation type.

Ensure that you specify the username in the *Username.Context* format.

Ensure that the administrator has permission to write to the Novell SYS: volume. Also, ensure that TCP port 524 is open on the firewall. For more information see [Section 1.3, “Supported Platforms,”](#) on page 14.
- 10 In the *Launch Inventory Process as User* drop-down list, select <*New Credential*>.

The Create Stored Machines dialog box is displayed.
- 11 Fill in the fields:

Username: Specify the administrator username created in [Step 1](#). You must specify the username either in the *hostname\username* or the *domain\username* format. If you choose to specify the username in the *hostname\username* format, ensure that the hostname is the local hostname.

Password: Specify the password of the administrator.

Validation Type: Select *Netware LSA*.

Hostname or IP of machine to validate against: Specify it as localhost.
- 12 Click *Save*.

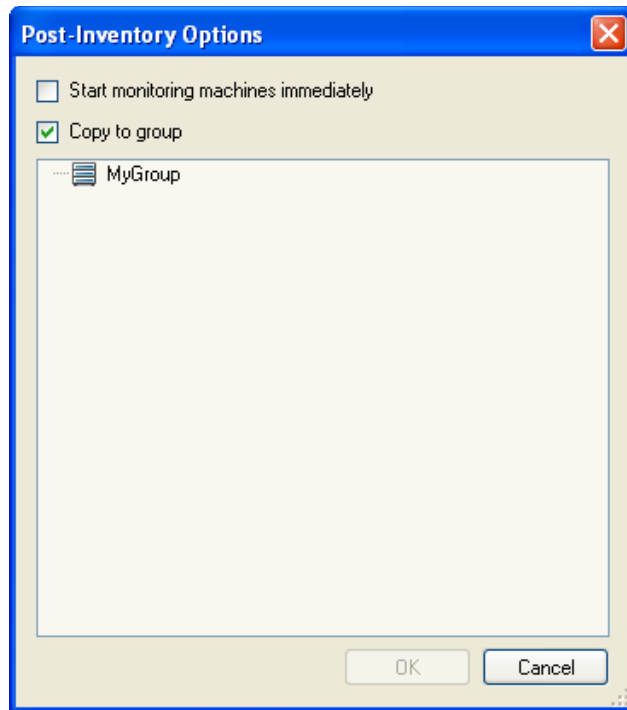
3.2.3 Inventory Options

If you want to start monitoring immediately after inventory or if you want to copy machines to a user-defined group, you must configure the Post-Inventory options. The options are available only if you had selected the *Inventory* check box in Discovery and Inventory dialog boxes.

To configure the Post-Inventory options:

- 1 Click the *Options* button in the *Inventory* area.

The Post-Inventory Options dialog box is displayed.



- 2 Select *Start monitoring machines immediately* to begin monitoring each machine discovered when the inventory is complete.
To start monitoring later, do not select the option.
- 3 Select *Copy to group*, then select a target Group from those displayed in the list box to copy all the inventoried machines to that group.
- 4 Click *OK*.

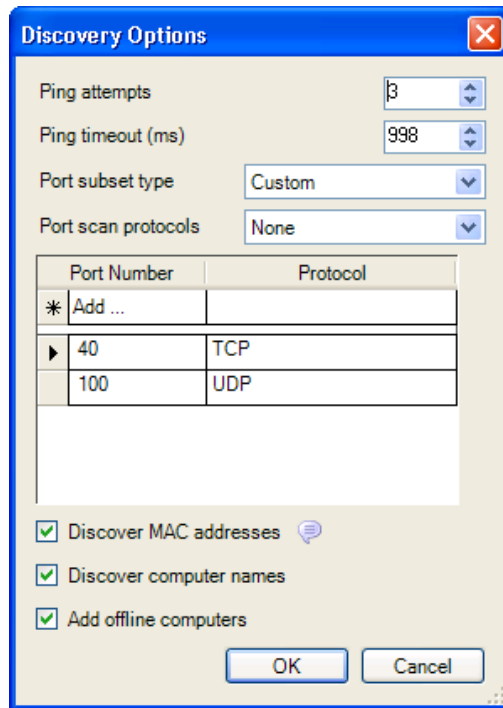
3.2.4 Discovery Options

you can configure additional discovery options when you choose to discovery and inventory in any of the following Inventory methods:

- ♦ Windows Domain
- ♦ IP Range
- ♦ Subnet
- ♦ SNMP

To configure discovery options:

- 1 In the *Discovery* pane, click the *Options* button.
The Discovery Options dialog box is displayed.



- 2 Configure the *Ping Attempts* and *Ping Timeout* options as desired.
- 3 In the *Port subset type* drop-down list, select one of the following subset types:
 - ♦ **Most Common:** Includes ports for common Windows network services.
 - ♦ **All Standard:** Includes ports below 1025 used by standard protocols.
 - ♦ **Custom:** Includes customized ports that you specify.To assign individual scan protocols to specific ports:
 1. In the *Port Number* column, click *Add* and type the port number.
 2. In the *Protocol* column, select the protocol type (TCP or UDP) from the drop-down cell.
 3. Press *Enter*.
 4. (Optional) If you want to add a new port number, repeat [Step 1](#) through [Step 3](#).
- 4 To find MAC addresses of machines in the same ethernet segment as the PlateSpin Recon Server, select *Discover MAC addresses*. The option is selected by default.
- 5 To include the computer names in the discovery poll, select *Discover computer names*. The option is selected by default.
- 6 To discover computers even if they are currently offline, select *Add offline computers*.
- 7 Click *OK*.

3.2.5 Viewing Inventory Data and Monitoring Servers

On completion of an inventory operation, you can view inventory data on servers that are listed in the Data Center Explorer tree. To view dynamic data, servers must be monitored for a period of time.

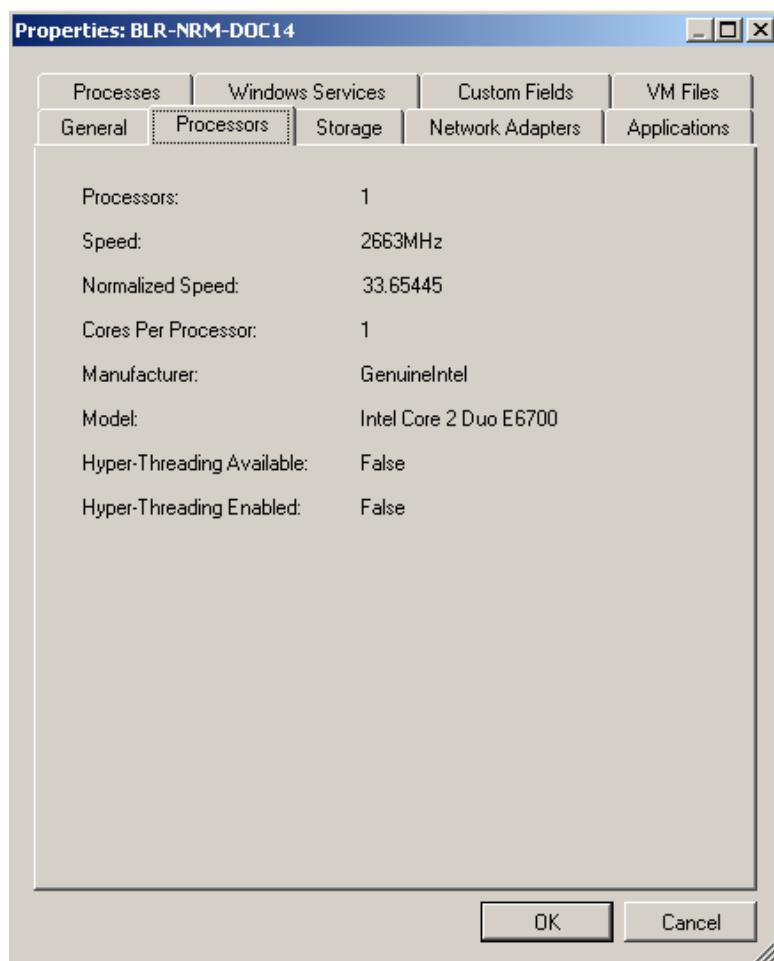
- ♦ “Viewing Properties” on page 47
- ♦ “Refreshing the Inventory” on page 48
- ♦ “Collecting Dynamic Data by Monitoring Servers” on page 48
- ♦ “Viewing Dynamic Data” on page 49
- ♦ “Viewing Logs” on page 49
- ♦ “Custom Fields” on page 49
- ♦ “Monitoring By Using Sudo User Credentials” on page 52

IMPORTANT: When you monitor Windows Server 2003, the memory utilization on the monitored server might increase because the memory usage of the `Svchost.exe` process increases when other users remotely collect performance data about the monitored server. This is a known issue of Windows Server 2003. To troubleshoot the issue, see the [Microsoft Knowledge Base article ID 969701](http://support.microsoft.com/kb/969701) (<http://support.microsoft.com/kb/969701>).

Viewing Properties

Properties can assist with viewing static data for inventoried servers. To view inventory information, right-click a server and click *Properties*.

Figure 3-1 Sample View of the Properties Dialog Box



The tabs available in the Properties dialog box change based on the type of server being viewed.

Refreshing the Inventory

To refresh server inventory information in the Data Center Explorer, right-click a server or multiple servers and click *Refresh Inventory* on the menu.

Inventory Refresh jobs are scheduled and run in the Job Explorer.

Collecting Dynamic Data by Monitoring Servers

- 1 To monitor a single server, right-click the inventoried server, then click *Start Monitoring*.
To monitor multiple servers, select them or their parent group, filter, or site, then click *Start Monitoring*.
A job is created in the Job Explorer. For more details, refer to [Section 2.6, "Job Explorer," on page 26](#).
- 2 After data is collected, click the server to view the corresponding chart in the Chart Viewer. With at least five minutes of data collected, the Chart Viewer can display workloads and utilization data.

Although PlateSpin Recon might not be currently monitoring a server, it is still possible to view charts and reports showing the previously collected server data.

You can monitor multiple servers by selecting them individually or by using groups, filters, or sites.

To stop monitoring, right-click the monitored server and select *Stop Monitoring*. In the confirmation dialog box, click *Yes* to confirm or *No* to exit the dialog box and continue monitoring.

Viewing Dynamic Data

You can view workloads and utilization data in a chart after collecting at least five minutes worth of data. Analysis Reports require a minimum of one hours worth of data. If there is insufficient data, the report is displayed without dynamic data. You can view collected data in a chart by selecting a monitored server.

For more information, see [Section 5.1, “Chart Viewer,” on page 85](#) and [Section 5.2, “Report Explorer,” on page 89](#)

Viewing Logs

Log files reflect the status of data collection for monitored servers.

To inspect a log, right-click a server, click *Logs*, then click *View Logs*.

Custom Fields

Custom fields are available for entering server information that cannot be inventoried by PlateSpin Recon. These non-inventory fields can assist with identifying candidates for planning and consolidation. The information is stored in the visible table name.

You can specify the custom fields in one of the following ways:

- ♦ [“Manually Specifying the Custom Fields” on page 49](#)
- ♦ [“Importing the Custom Fields” on page 50](#)

Manually Specifying the Custom Fields

To specify custom fields for multiple servers:

- 1** Select a server or multiple servers in the Data Center Explorer.
- 2** Right-click the selection and click *Properties* for a single server, or click *Set Custom fields* for multiple servers.
The Properties dialog box is displayed.
- 3** Click the *Custom Fields* tab.
- 4** Select a field to view its definition and parameters.

Properties: DEV

General Processors Disk Drives Network Adapters Applications
Processes Microsoft Virtual Server Windows Services Custom Fields

Rack Units	2
Energy Consumption (W)	330
Heat Dissipation (BTU/hr)	1125
Normalized Speed	4.714753

Heat Dissipation (BTU/hr)
Amount of heat generated by a server. Heat Dissipation (BTU/hr) = Energy...

Reset Reference Values

Company	
Department	
Location City	
Location State	
Location Country	
Function	
Rack	
Billing Code	
Usage	
Lease Expires	
Warranty Expires	

Billing Code
Billing code of the server

OK Cancel

5 Type or select information in any of the fields, including those already filled with data.

6 Click *OK*.

If necessary, click *Reset Reference Values* to return all fields to their default values.

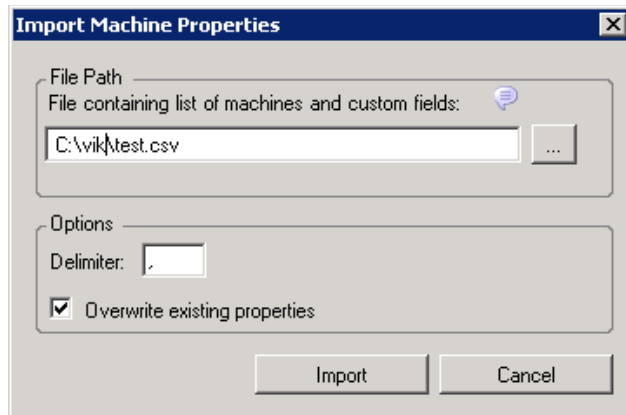
The database is updated for the selected servers. Report templates can filter for particular custom fields.

Importing the Custom Fields

PlateSpin Recon allows you to import the custom fields for all the servers in the local site.

1 In the Data Center Explorer, click *File > Import Machine Properties*.

The Import Machine Properties dialog box is displayed.



2 Type the path (or use the adjacent browse button) to locate a CSV file containing custom fields.

Ensure to follow these guidelines as you work with the CSV file:

- ♦ The first line must contain the label of the custom fields separated by a delimiter such as a comma (.).
- ♦ The first custom field label on the first line must be `Machine`. You can list the remaining custom fields in any order.
- ♦ You must use the following custom field labels only:

Rack Units
 Normalized Speed
 Company
 Department
 Location City
 Location State
 Location Country
 Function
 Rack
 Billing Code
 Usage
 Lease Expires
 Warranty Expires

- ♦ The custom field label names are case sensitive.
- ♦ The values for the custom fields must be specified subsequent to the first line, and must be corresponding to the custom fields.
- ♦ Each line must contain the custom fields for a machine.
- ♦ The value for the `Machine` custom field can be IP address or the machine name. The specified machine name must be unique.
- ♦ The values on each line must be separated by a delimiter such as a comma (.).

A sample CSV file is as follows:

```
Machine,Rack Units,Company,Lease Expires,Warranty Expires
1.x.x.x,130,Novell,11/25/2009,1/1/2010
```

3 Specify the delimiter used in the CSV file. By default, it is a comma (.).

- 4 (Conditional) If you want to replace the existing custom fields with the values contained in the CSV file, select the *Overwrite Existing Properties* option. By default, the option is selected.
- 5 Type or select information in any of the fields, including those already filled with data.
- 6 Click *Import*.

Monitoring By Using Sudo User Credentials

PlateSpin Recon 3.7.2 allows monitoring of Linux (RHEL, SLES, Ubuntu, CentOS, SLES with Xen) and Citrix XenServer machines by using sudo user credentials. However, monitoring of the AIX and Solaris machines by using sudo user credentials is not supported. PlateSpin Recon users can monitor the machines by using sudo user credentials only if they have the permission to run the commands on the target machine. For more information on how to grant permission to perform monitoring by using sudo user credentials, see [Section B.9, “Sudo Configuration to Inventory and Monitor Linux Machines,” on page 207](#).

3.3 Working with Filters

The *Filters* node contains additional nodes that list inventoried servers organized into logical groups based on domain, server make and model, operating system, status, subnet mask, virtual machine servers, and workgroup.

Filters cannot be added or removed, and servers cannot be added to filters.

3.4 Working with Groups

The *Groups* node contains user-defined collections of servers. Use groups to logically group servers by location, function, application, and so on. Groups can be created, deleted, and nested. Servers can belong to more than one group.

- ♦ [Section 3.4.1, “Creating a Group,” on page 52](#)
- ♦ [Section 3.4.2, “Copying Servers to a Group,” on page 52](#)
- ♦ [Section 3.4.3, “Deleting a Server from a Group,” on page 53](#)

3.4.1 Creating a Group

- 1 In the Data Center Explorer, right-click *Groups* and click *Create Group*.
The Create Group dialog box is displayed.
- 2 Type a *Group Name*.
- 3 Click *OK*.

The group is created and ready for servers to be copied to it.

3.4.2 Copying Servers to a Group

- 1 In the Data Center Explorer, select one or more servers.
- 2 Right-click the selection, then click *Copy to Group*.
- 3 Click the desired group on the submenu that is displayed.
The server is displayed under the selected group in the Data Center Explorer.

You can also drag servers from under the *All*, *Filtered* or *Sites* node and drop them on an existing group to copy them to that group.

3.4.3 Deleting a Server from a Group

1 In the Data Center Explorer, do one of the following:

- ♦ Right-click the server you want to delete, then click *Delete*.
- ♦ Select the server you want to delete, then press the Delete key.

A confirmation dialog box is displayed.

2 (Optional) To permanently remove the monitoring data for the server when it is deleted from the group, select *Delete <server> and associated monitoring data from the database*.

3 Click *Delete*.

The server is removed from the group in the Data Center Explorer.

3.5 Working with Sites

The *Sites* node contains logically organized sets of servers for analysis and reporting on data collected from multiple locations.

Each PlateSpin Recon installation includes a default *Local* site that contains machines that you work with using the local instance of PlateSpin Recon.

You can work with data external to the local PlateSpin Recon instance, in which case you must first create a new site of the appropriate site type.

NOTE: Machines imported from a remote site can be viewed in reports and charts after the data has been synchronized. However, you cannot inventory or monitor the machines. If you want to inventory or monitor the machines, you must perform the actions at the original remote site and import the machines again. Subsequently, you need to synchronize the data again.

The following table explains the different types of sites that you can create.

Table 3-4 Site Types

Type	Description
Local	To work with PlateSpin Recon snapshots (*.prs files) exported from other instances of PlateSpin Recon. When it is exported from the local instance of PlateSpin Recon, a snapshot can be imported into the default Local site. For more information, see Section 3.6, "Working with PlateSpin Recon Snapshots," on page 55.
Remote PlateSpin Recon	To connect to and work with data collected by a remote PlateSpin Recon instance. For more information, see Section 3.8, "Working with Data Collected by Other PlateSpin Recon Instances," on page 60.

Type	Description
Remote	To work with data collected by supported third-party performance monitoring tools, such as Microsoft Operations Manager 2005, Windows Performance Monitor (<code>perfmon.exe</code>), HP Reporter, and other generic performance data gathering tools. For more information, see Section 3.9, “Working with Data Collected by Other Performance Monitoring Tools,” on page 62.

Using PlateSpin Recon sites, you can designate a PlateSpin Recon Server as the Master Server and carry out all analysis and reporting tasks within a single PlateSpin Recon environment.

- ♦ [Section 3.5.1, “Adding a Site,”](#) on page 54
- ♦ [Section 3.5.2, “Deleting a Site,”](#) on page 54
- ♦ [Section 3.5.3, “Modifying a Site,”](#) on page 54

3.5.1 Adding a Site

- 1 Right-click anywhere in the Data Center Explorer and choose *Sites > Add Site*.
The Add Site dialog box is displayed.
- 2 Type a unique name for the site.
- 3 In the *Type* drop-down list, ensure that *Local* is selected
- 4 (Optional) Provide a description for the site.
- 5 Ensure that the *Run Import Site Snapshot Wizard* option is deselected.
- 6 Click *OK*.

The site is created and is displayed under the *Sites* node in the Data Center Explorer.

For information on adding Remote or Remote PlateSpin Recon sites, see [Section 3.7, “Working with Remote Data,”](#) on page 60.

3.5.2 Deleting a Site

- 1 In the Data Center Explorer, do one of the following:
 - ♦ Right-click the site you want to delete, then click *Delete*.
 - ♦ Select the site you want to delete, then press the Delete key.

A Confirm Delete dialog box is displayed.

- 2 Click *Yes*.

PlateSpin Recon removes the selected site and associated resources from the Data Center Explorer.

3.5.3 Modifying a Site

- 1 In the Data Center Explorer, right-click the site you want to modify, click *Sites*, then click *Edit Site*.
The Edit Site dialog box is displayed.

- 2 Make the required modifications.
- 3 Click *OK*.

3.6 Working with PlateSpin Recon Snapshots

PlateSpin Recon snapshots (*.prs files) contain data about machines, raters, report templates, generated reports, projects, scenarios, and server templates exported by other instances of PlateSpin Recon, or from the current instance as a backup.

You can also import snapshots exported from previous versions of PlateSpin Recon. In this case, a snapshot upgrade is performed before the data is brought into PlateSpin Recon. You need to consider the following:

- ♦ Importing a large snapshot, for instance if the snapshot was compressed before export, requires two to five times the size of the old snapshot in disk space.
- ♦ Snapshots exported from PlateSpin Recon 3.0 and later are supported.
- ♦ Snapshots from PlateSpin Recon 3.1 are only supported if all hot fixes were applied to PlateSpin Recon before the snapshot was exported.
- ♦ Exporting snapshot of Citrix XenServer is not supported.

You can import a snapshot either into the default Local site, or into a new local site.

- ♦ [Section 3.6.1, “Exporting a PlateSpin Recon Snapshot,” on page 55](#)
- ♦ [Section 3.6.2, “Importing a PlateSpin Recon Snapshot,” on page 58](#)

3.6.1 Exporting a PlateSpin Recon Snapshot

- 1 Right-click in the Data Center Explorer, click *Sites*, then click *Export Snapshot*.

The Export Site Snapshot Wizard is displayed.

- 2 Select export objects you want to include in the snapshot by enabling the check box next to the objects.

NOTE: Even though you can choose to export the Citrix XenServer object, the object will not be included in the snapshot because exporting snapshot of Citrix XenServer is not supported.

When a report template is included in the snapshot, its generated reports can also be included. Similarly, if a project is included in the snapshot, its scenarios can be included.

- 3 Click *Next*.

If you selected to export *Machines* as part of the snapshot, continue with [Step 4](#); else, skip to [Step 6](#).

- 4 In the Step 1b dialog box, select summary options for the data to be exported, such as a data range time frame and whether to include Logs or Credentials.

Export Site Snapshot

Step 1b of 3: Specify the machines filter

Summary Data Granularity

☐ 5 Minute (Used in charts only)

☒ 1 Hour (Used in charts, reports, and projects)

Summary Data Range

☒ All Data

☐ Last Hour(s)

☐ Between and

☐ Monitor Logs

☐ Attached Credentials

5 Click *Next*.

The Step 2 dialog box is displayed.

Export Site Snapshot

Step 2 of 3: Select the file name for the exported snapshot

UNC or File Path on the PlateSpin Recon Server:

Credentials

☒ Compress (file size will be reduced at the cost of slower import/export)

☐ Encrypt Password:

Confirm:

- 6** Specify the filename for the exported snapshot:
- 6a** Specify the UNC (\\<server>\<path>\<filename>) or file path (drive:\<directory>\<filename>) where you want to save the snapshot. You can also click *Browse* to browse for and select a location
 - 6b** Select or create the appropriate credentials to access the location where you want to save the exported snapshot. You need write permission to this location.
 - 6c** Choose whether to *Compress* the file.
 - 6d** Choose whether to *Encrypt* the file. For an encrypted snapshot file, specify and confirm a password for decryption.
- 7** Click *Next*.

The Step 3 dialog box is displayed.

Export Site Snapshot

Step 3 of 3: Choose the delivery types and the schedule for the export job

Delivery Types

Network File ☒ Network File Options...

FTP ☐ FTP Options...

☐ Delete Source File After Delivery

Schedule

☒ Now

☐ Later: 03/20/2008 12:00 AM

☐ Recurrence: Daily 01:00

Help < Back Export Cancel

- 8** Configure the snapshot delivery type and the delivery schedule:
- 8a** If you want the snapshot delivered to a network or FTP location after it is exported, select a delivery type. To configure the delivery method options, click the option button next to the selected delivery method. If you are unsure of the delivery configurations, check with your System Administrator.
The file path must be accessible to the PlateSpin Recon Server.
 - 8b** (Optional) If you want to delete the source snapshot after it is delivered, select the *Delete Source File After Delivery* check box.
 - 8c** Configure the schedule when you want the snapshot to be delivered.
 - ♦ To immediately deliver the snapshot, select *Now*.

- ♦ To deliver the snapshot at a later time, select one of the following schedules
 - ♦ “Later” on page 209
 - ♦ “Recurrence” on page 209

9 Click *Export*.

The data snapshot is ready to be imported into another instance of PlateSpin Recon.

3.6.2 Importing a PlateSpin Recon Snapshot

- 1** Determine the type of local site (the default local site or a new local site) you want to import the required snapshot into. To import into a new local site, you need to create the site before importing the snapshot.
- 2** In the Data Center Explorer, right-click the target local site, click *Sites*, then click *Import Snapshot*.

Or

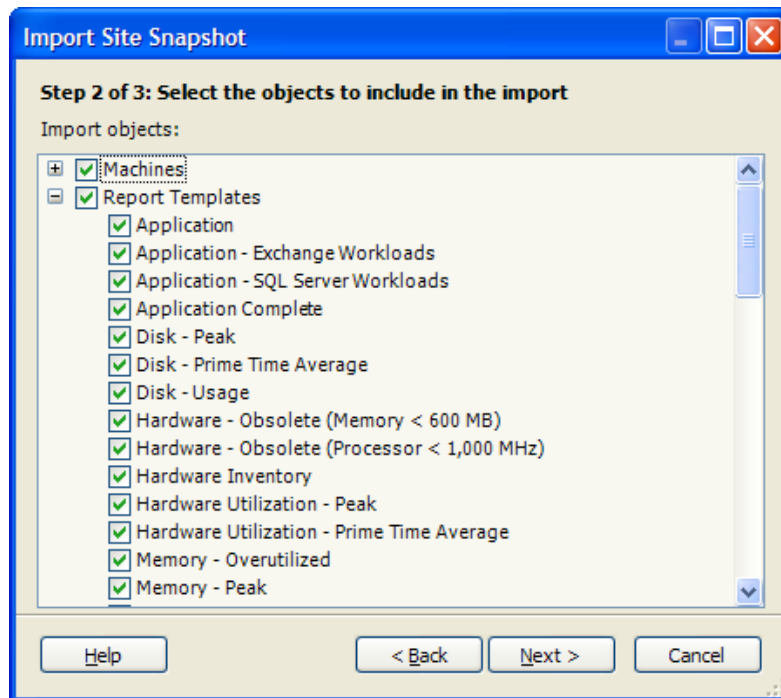
Right-click in the Data Center Explorer, click *Sites*, then click *Import Snapshot*.

The Import Site Snapshot Wizard is displayed.

- 3** In the Step 1 dialog box, configure the following:
 - 3a** Specify or click *Browse* to browse for and select the UNC or the local file path of the snapshot you want to import.
 - 3b** Select or add the credentials needed to access the location.
 - 3c** If the file was encrypted on export, type the *Snapshot Password* to decrypt it.

- 4 Click *Next*.

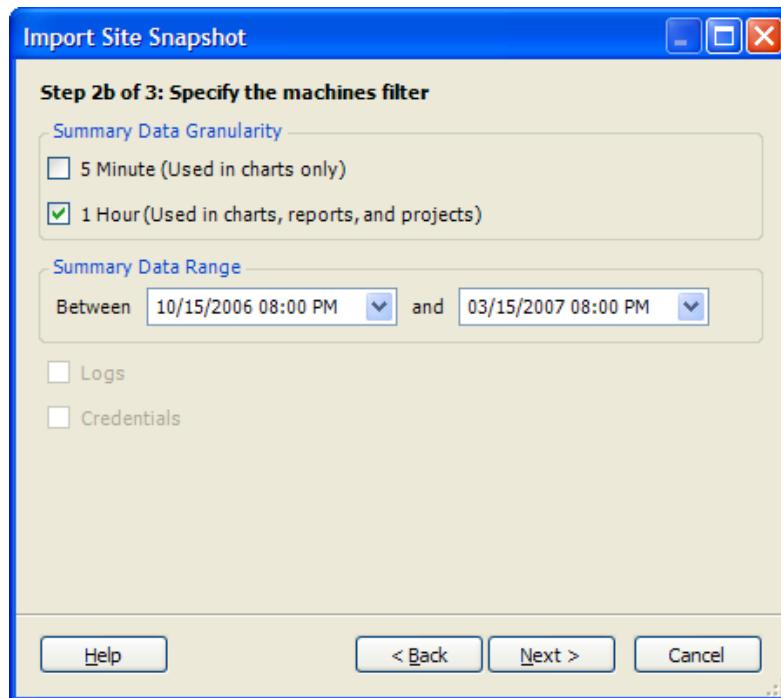
The Step 2 dialog box is displayed.



- 5 Select the check boxes next to the objects you want to import.

- 6 Click *Next*.

If the imported objects include machines with monitoring data, the Step 2b dialog box is displayed. Otherwise, proceed to [Step 9](#).



7 Specify the data options for the machines filter.

8 Click *Next*.

The Step 3 dialog box is displayed.

9 In the *Import to Site* drop-down list, select the target site you want to import the required snapshot into.

10 Click *Import*.

After the data is imported, work with it using the Data Center Explorer under the *Sites* node.

3.7 Working with Remote Data

Working with performance monitoring and resource utilization data external to the local instance of PlateSpin Recon allows you to take advantage of the PlateSpin Recon capacity planning, analysis, reporting and chargeback capabilities, regardless of the source of information. Data can be imported on a one-time or recurring basis.

3.8 Working with Data Collected by Other PlateSpin Recon Instances

You can use performance and utilization data collected by other instances of PlateSpin Recon, such as in a setup where multiple PlateSpin Recon instances handle monitoring and data collection, with one acting as a central PlateSpin Recon instance for generating reports and developing consolidation plans.


NOTE: Machines imported from a remote site can be viewed in reports and charts after the data has been synchronized. However, you cannot inventory or monitor the machines. If you want to inventory or monitor the machines, you must perform the actions at the original remote site and import the machines again. Subsequently, you need to synchronize the data again.

To work with data collected by another instance of PlateSpin Recon:

- 1** Create a new Remote PlateSpin Recon site to connect to the remote PlateSpin Recon instance:

- 1a** Right-click in the Data Center Explorer, click *Sites*, then click *Add Site*.

The Add Site dialog box is displayed.



- 1b** Select *Remote PlateSpin Recon* as the new site *Type* and specify the required basic parameters for the new site, including a name and an optional description.

- 1c** Specify the URL of the PlateSpin Recon server in the following format:

`http://hostname_or_IP_address_of_the_Recon_Server[:port]/
platespin_372/recon_372`

For example, `http://10.0.0.0/platespin_372/recon_372`

- 1d** In the *Credentials* drop-down list, select the appropriate credentials or select *<New Credential>* to add new credentials.

- 1e** If you want the data stream encrypted, ensure that the *Encrypt Data Stream* check box is selected.

- 1f** (Optional) To verify connectivity with the remote server, click *Connection Test*.

- 1g** Click *OK*.

The new site is added to the *Sites* node in the Data Center Explorer.

- 2 Right-click the Remote PlateSpin Recon site you created in [Step 1](#), click *Sites*, then click *Import Machines*. The machine nodes are added to the site.
- 3 Right-click the Remote PlateSpin Recon site again, click *Sites*, then click *Synchronize Data*.

3.9 Working with Data Collected by Other Performance Monitoring Tools

You can import and work with performance and utilization data collected by other performance monitoring software, such as Microsoft Operations Manager (MOM) 2005, Microsoft Windows Performance Monitor (`perfmon.exe`), and HP Reporter. You can also use the PlateSpin Recon Performance Data Import feature.

NOTE: Machines imported from other performance monitoring tools can be viewed in reports and charts after the data has been synchronized. However, you cannot inventory or monitor the machines. If you want to inventory or monitor the machines, you must perform the actions at the original remote site and import the machines again. Subsequently, you need to synchronize the data again.

PlateSpin Recon supports summarization of collected data within 5-minute and 60-minute intervals, regardless of how the selected connector summarizes and collects data. PlateSpin Recon requires that data collected by connectors be summarized over intervals of 60 minutes or less. We recommend importing connector data summarized over 60-minute intervals, especially for systems whose utilization data covers a period of longer than one month.

For PlateSpin Recon to generate reliable consolidation plans, we recommend that you collect utilization data for a period of one to three months before importing it into PlateSpin Recon.

For the capability to work with data collected by third-party software, you must:

- ♦ Create a new Remote site of the appropriate type and select an appropriate connector.
- ♦ Inventory the remote machines.
- ♦ Synchronize data in the new site with that collected by the remote monitoring tool.
- ♦ [Section 3.9.1, “Required Performance Counters,” on page 62](#)

3.9.1 Required Performance Counters

You must ensure that the required performance and utilization data is actually gathered for PlateSpin Recon counters in the required categories.

Table 3-5 *Required PlateSpin Recon Counters*

PlateSpin Recon Counter	Expressed in
Processor utilization	Percentage of total processor time
Memory utilization	Total number of free bytes
Disk IO	Bytes per second
Network IF I/O	Bytes per second

NOTE: NetWare machines only require two counters (CPU and disk) to be gathered for successful consolidation. For more information see [Section 1.3, “Supported Platforms,”](#) on page 14.

Where required, PlateSpin Recon provides mechanisms for converting third-party performance and utilization data to the required PlateSpin Recon format.

- ♦ [“Working with Microsoft Operations Manager \(MOM\) 2005 Data”](#) on page 63
- ♦ [“Working with Windows PerfMon Counter Logs”](#) on page 64

Working with Microsoft Operations Manager (MOM) 2005 Data

1 Download and install the PlateSpin Management Pack for MOM 2005 (see Knowledge Base Article [Q7920903](http://www.novell.com/support/viewContent.do?externalId=7920903) (<http://www.novell.com/support/viewContent.do?externalId=7920903>) and collect data by using MOM.

2 Create a new Remote site for connecting to MOM:

2a Right-click in the Data Center Explorer, click *Sites*, then click *Add Site*.

The Add Site dialog box is displayed.

The screenshot shows the 'Add Site' dialog box. The 'Name' field is set to 'MySite' and the 'Type' dropdown is set to 'Remote'. The 'Description' field is empty. The 'Connector Type' dropdown is set to 'Microsoft Operations Manager 2005'. Under the 'Database Connection' section, 'Database Hostname' is 'MyDB', 'Database Instance' is 'OnePoint', 'Connection Timeout' is 300, and 'Query Timeout' is 300. Under the 'Database Credentials' section, the dropdown is set to 'MyDomain\MyName'. At the bottom right, there is a 'Test Settings' button, and at the bottom center, there are 'OK' and 'Cancel' buttons.

- 2b** Select *Remote* as the *Type* of the new site and specify the required basic parameters for the new site, including a name and an optional description.
- 2c** In the *Connector Type* drop-down list, select *Microsoft Operations Manager 2005*.
- 2d** Type the *Database Hostname* for the remote server. *Database Instance* defaults to *OnePoint* and must be changed to reflect the Microsoft Operation Manager 2005 database instance in use.
- 2e** Select the appropriate credentials or select *<New Credential>* to add new credentials.
- 2f** (Optional) To verify connectivity with the remote database, click *Test Settings*
- 2g** Click *OK*.

The new site is added to the *Sites* node in the Data Center Explorer.

3 Inventory the machines:

- 3a** In the Data Center Explorer, right-click the remote site, click *Sites*, then click *Inventory Machines*.

The Inventory Machines dialog box is displayed.

- 3b** Select the check boxes next to the machines you want to inventory, specify the required inventory parameters, then click *OK*.

4 Synchronize site data:

- 4a** In the Data Center Explorer, right-click the required site and select *Sites > Synchronize Data*.

- 4b** Select the check boxes next to the required machines. To select all machines in a site, select the check box next to the site.

- 4c** Select the required data range, scheduling, and recurrence options.

- 4d** Click *OK*.

Working with Windows PerfMon Counter Logs

- 1** Set up your Windows Performance Monitor (*perfmon.exe*) to gather counter data required by PlateSpin Recon. See [Section 3.9.1, “Required Performance Counters,” on page 62](#).
- 2** Export the data to a CSV file. For detailed information, see the Windows documentation.
- 3** Create a new Remote site to work with Windows Performance Monitor logs:
 - 3a** Right-click in the Data Center Explorer, click *Sites*, then click *Add Site*.

The Add Site dialog box is displayed.

Add Site

Name: Type:

Description:

Connector Type:

UNC or File Path on the PowerRecon Server:

Credentials:

- 3b** Select *Remote* as the *Type* of the new site and specify the required basic parameters for the new site, including a name and an optional description.
- 3c** In the *Connector Type* drop-down list, select *Performance Monitor*.
- 3d** Type the path to (or browse to and select) the CSV file that contains data exported from Performance Monitor.
- 3e** Select the required credentials or select *<New Credential>* to add new credentials.
- 3f** (Optional) To verify UNC or local file path access, click *Test Settings*.
- 3g** Click *OK*.
The new site is added to *Sites* node in the Data Center Explorer.
- 4** Inventory the machines:
 - 4a** In the Data Center Explorer, right-click the remote site, click *Sites*, then click *Add Site*.
The Inventory Machines dialog box is displayed.
 - 4b** Select the check boxes next to the machines you want to inventory, specify the required inventory parameters, then click *OK*.

5 Synchronize site data:

- 5a** In the Data Center Explorer, right-click the required site, click *Sites*, then click *Synchronize Data*.
- 5b** Select the check boxes next to the required machines. To select all machines in a site, select the check box next to the site.
- 5c** Select the required data range, scheduling, and recurrence options.
- 5d** Click *OK*.

3.10 Using HP Reporter Data with PlateSpin Recon

PlateSpin Recon features a specialized HP Reporter connector. It enables PlateSpin Recon to connect to an HP Reporter 3.8 database to import performance and resource utilization data for charts, reports and consolidation planning.

You can use any HP Reporter-collected data to generate PlateSpin Recon charts and reports. However, for consolidation planning you must configure Reporter to collect specific utilization data that correlates to data collected by specific PlateSpin Recon counters. For detailed information, see [Section 3.9.1, “Required Performance Counters,” on page 62](#).

You need to specify appropriate formulas to transform the data from Reporter-collected metrics to PlateSpin Recon counters because HP Reporter does not collect data in exactly the same format as PlateSpin Recon.

- ♦ [Section 3.10.1, “Using Formulas to Convert Metrics to Counters,” on page 66](#)
- ♦ [Section 3.10.2, “Example Formulas,” on page 67](#)
- ♦ [Section 3.10.3, “Working with HP Reporter Data,” on page 68](#)

The HP Reporter data summarization interval is specified during the creation of the HP Reporter Metric List.

3.10.1 Using Formulas to Convert Metrics to Counters

You can use a formula to provide the following information to the PlateSpin Recon HP Reporter Connector:

- ♦ Which Reporter metric list contains metrics required by PlateSpin Recon
- ♦ How to convert data from measurement units used by Reporter units to those required by PlateSpin Recon

A valid HP Reporter Connector formula consists of the name of the required HP Reporter metric list, followed by a dot character, followed by an arithmetic expression that converts the metric into a PlateSpin Recon counter. Arithmetic expressions can contain any of the collected Reporter metrics, numeric constants, and arithmetic operations.

NOTE: Only metrics collected by the same metric list can be used in one formula, although different formulas can use different metric lists.

For example, PlateSpin Recon requires memory utilization data expressed in the number of free bytes available to the operating system, while Reporter collects the same data by using two separate metrics:

- ♦ Percentage of memory utilized (GBL_MEM_UTIL)
- ♦ Total size of the memory in gigabytes (GBL_MEM_PHYS)

In this case, the appropriate formula for converting collected metric data into the PlateSpin Recon counter format would be:

$$\text{GBL_MEM_PHYS} * (100 - \text{GBL_MEM_UTIL}) * 1024 * 1024 * 1024 / 100$$

3.10.2 Example Formulas

The following are examples of HP Reporter metrics that you can collect using CODA agents, along with corresponding conversion formulas. Note that To do consolidation planning, PlateSpin Recon requires data for all four counters. Because HP Reporter can collect metric data with several different agents, and the same PlateSpin Recon counters can be calculated using different HP Reporter metrics, this list is not exhaustive and serves only as an example.

Table 3-6 Example HP Reporter Metrics

PlateSpin Recon Counter	HP Reporter Metrics	Conversion Formula
Processor Utilization	GBL_CPU_TOTAL_UTIL	<MetricList>.GBL_CPU_TOTAL_UTIL
	BYCPU_CPU_TOTAL_UTIL	<MetricList>.BYCPU_CPU_TOTAL_UTIL
Memory Utilization	GBL_MEM_PHYS and GBL_MEM_UTIL	<MetricList>.(GBL_MEM_PHYS*(100-GBL_MEM_UTIL))*1024*1024*1024/100
	GBL_MEM_FREE	<MetricList>.(GBL_MEM_FREE)*1024*1024
Disk Utilization	GBL_DISK_PHYS_BYTE_RATE	<MetricList>.GBL_DISK_PHYS_BYTE_RATE
	BYDSK_PHYS_READ_BYTE_RATE	<MetricList>.BYDSK_PHYS_READ_BYTE_RATE
Network Utilization	BYNETIF_OUT_BYTE_RATE and BYNETIF_IN_BYTE_RATE	<MetricList>.BYNETIF_OUT_BYTE_RATE + BYNETIF_IN_BYTE_RATE

Data for certain HP Reporter metrics and PlateSpin Recon counters is collected separately for each device instance. For example, the BYCPU_CPU_TOTAL_UTIL metric gives a separate utilization figure for each CPU in the system. To distinguish among different device instances, include a metric with a device identification string in the list when you set up the HP Reporter metric lists.

The PlateSpin Recon HP Reporter connector identifies the following metrics as containing device instance names:

Table 3-7 *PlateSpin Recon HP Reporter Metrics*

Functional Area	Metric
Processors	BYCPU_ID
Disks	BYDSK_DEVNAME or BYDSK_ID
Network Interface Cards	BYNETIF_NAME or BYNETIF_ID
File Systems	FS_DIRNAME, FS_DEVNAME or FS_DEVNO

3.10.3 Working with HP Reporter Data

- 1 Configure your HP Reporter to collect data in the required performance and utilization measurement categories. For more information, see [Section 3.9.1, “Required Performance Counters,” on page 62](#).
- 2 (Conditional) If you are using Oracle with HP Reporter, install Oracle client software on the PlateSpin Recon Server host. For detailed information, see the Oracle documentation.
- 3 Create a new Remote site for connecting to the HP Reporter database:
 - 3a Right-click in the Data Center Explorer, click *Sites*, then click *Add Site*.
The Add Site dialog box is displayed.

Add Site

Name: Type:

Description:

Connector Type:

Settings **Counters**

Database Connection

Server Name :

Database Name :

Connection Timeout : Query Timeout :

Database Type:

Database Port:

Database Credentials

3b Specify a name and, optionally, a description.

3c In the *Type* drop-down, select *Remote*.

3d In the *Connector Type* drop-down list, select *HP Reporter*.

The HP Reporter-related *Settings* and *Counters* fields are displayed.

3e On the *Settings* tab, provide the following information required to connect to the HP Reporter database:

- ♦ Database hostname
- ♦ Database or database instance name
- ♦ Connection timeout and the Query timeout.

When you synchronize a large data set from HP Reporter, you might need to set the *Connection timeout* and the *Query timeout* values to 0, which means infinite/unlimited timeout.

- ♦ Database type
- ♦ Port number, if a non-default port number is used
- ♦ Valid credentials for accessing the database

3f On the *Counters* tab, for the four required performance and utilization data categories, provide counter formulas for the system to use when converting HP Reporter metrics to PlateSpin Recon counters. For more information, see [Section 3.10.1, “Using Formulas to Convert Metrics to Counters,”](#) on page 66.

3g Click *OK*.

The new site is added to *Sites* node in the Data Center Explorer.

4 Inventory the machines:

4a In the Data Center Explorer, right-click the remote site you created in [Step 3](#), click *Sites*, then click *Inventory Machines*.

The Inventory Machines in Site *site_name* dialog box is displayed.

4b Select the check boxes next to the machines you want to inventory.

4c Select the required credentials or select *<New Credential>* to add new credentials.

4d Click *OK*.

5 Synchronize site data:

5a In the Data Center Explorer, right-click the required site, click *Sites*, then click *Synchronize Data*.

5b Select the check boxes next to the required machines. To select all machines in a site, select the check box next to the site.

5c Select required data range, scheduling, and recurrence options.

5d Click *OK*.

For in-depth information about using PlateSpin Recon in conjunction with HP Reporter, see Knowledge Base article [Q7920889](http://www.novell.com/support/viewContent.do?externalId=7920889) (<http://www.novell.com/support/viewContent.do?externalId=7920889>).

3.11 Using Performance Data Import to Import Generic Data into PlateSpin Recon

PlateSpin has developed a robust Performance Data Import connector for PlateSpin Recon that allows users to import collected performance data from many third-party tools on the market into PlateSpin Recon to take advantage of its charting, reporting, and consolidation planning features.

When collected data has been moved from its proprietary location and format into a set of tables in the PlateSpin Recon monitoring database, PlateSpin Recon can import the data and make use of it like any other performance data. You can also set up a scheduled or recurring import, so the PlateSpin Recon analysis does not have to be a one-time thing.

Almost all performance data gathering tools use a different method and format for storing the gathered data. To enable PlateSpin Recon to successfully import any data, regardless of its original format, you need to use an Extract, Transform, and Load (ETL) tool. This type of tool takes the data, reshapes it, and then loads it into the tables where PlateSpin Recon expects to find it. See [Section 3.11.2, “Preparing Data for Import into PlateSpin Recon,”](#) on page 72.

The workloads in the data being imported need to be of the type that PlateSpin Recon can inventory under normal circumstances. All necessary ports between PlateSpin Recon and the workload being imported must be open.

As with other remote data connectors, data to be imported should contain one or more data points for every hour. Also, to use Consolidation Planning, specific counters need to exist in the data. See [Section 3.9.1, “Required Performance Counters,” on page 62.](#)

- ♦ [Section 3.11.1, “Performance Data Import Tables Data Structure,” on page 71](#)
- ♦ [Section 3.11.2, “Preparing Data for Import into PlateSpin Recon,” on page 72](#)

3.11.1 Performance Data Import Tables Data Structure

There are four tables used for importing of generic performance data:

- ♦ **ImportComputer table:** Contains information about computers that are to be imported into PlateSpin Recon.
- ♦ **ImportCounter table:** Contains information about all counters that were collected for any of the computers in the *ImportComputer* table.
- ♦ **ImportCounterInstance table:** Contains information about named instances of each counter for each computer. The combination of *Computer*, *Counter* and *Instance Name* must be unique.
- ♦ **ImportData table:** Contains monitoring data to be imported.

Each table contains one ID field, which acts as a primary key and is automatically generated by the database server when a new row is inserted into the table.

Table 3-8 *ImportComputer Table*

Column Name	Type	Description
ID	Integer	Generated automatically.
Name	varchar(255)	Network name of computer. This is the address or name of the computer. It allows PlateSpin Recon to address the computer over the network.

Table 3-9 *ImportCounter Table*

Column Name	Type	Description
ID	Integer	Generated automatically.
Name	varchar(255)	Unique name of the counter. This is the name of the counter as it will be shown in PlateSpin Recon charts. Some predefined counter names exist. See Section 3.9.1, “Required Performance Counters,” on page 62.

Table 3-10 *ImportCounterInstance Table*

Column Name	Type	Description
ID	Integer	Generated automatically.

Column Name	Type	Description
Computer	Integer	Reference to the ImportComputer table.
Counter	Integer	Reference to the ImportCounter table.
InstanceName	Integer	Name of the given counter instance on the specific computer.

Table 3-11 *ImportData Table*

Column Name	Type	Description
ID	Integer	Generated automatically.
Instance	Integer	Tells PlateSpin Recon for which computer, counter and counter instance this data value was collected.
ObservationDate	DateTime	Time the data was collected.
ObservationValue	float	Value of the counter.

3.11.2 Preparing Data for Import into PlateSpin Recon

To migrate data from a third-party location and format requires an Extract, Transform and Load (ETL) tool. There are many such tools on the market:

Open Source: CloverETL, Enhydra Octopus, Pentaho Data Integration, and Talend Open Studio.

You can also write proprietary data manipulation scripts, such as SQL queries, although the specifics of these tasks are beyond the scope of this document. For third-party tools, see the documentation that accompanies those tools. For scripts, refer to your in-house reference material on the languages and requirements.

Commercial: Informatica PowerCenter, Inaplex Inaport, and Altova MapForce.

Regardless of the tool you choose, the data transformation workflows are the same:

- ♦ Retrieve hostnames of the workloads you want to import and write them to the ImportComputer table
- ♦ Write the names of the required counters to the ImportCounter table
- ♦ Retrieve and write all performance data related to the above counters

To work with generic imported data:


- 1 Collect data in the required performance and utilization measurement categories. For more information, see [Section 3.9.1, “Required Performance Counters,” on page 62](#).
- 2 Extract, transform, and load the collected performance data into the PlateSpin Recon Performance Data Import tables.

For more information, see [Section 3.11, “Using Performance Data Import to Import Generic Data into PlateSpin Recon,” on page 70](#), [Section 3.11.1, “Performance Data Import Tables Data Structure,” on page 71](#) and [Section 3.11.2, “Preparing Data for Import into PlateSpin Recon,” on page 72](#).

3 Create a new Remote site for connecting to the Performance Data Import tables:

3a Right-click in the Data Center Explorer, click *Sites*, then click *Add Site*.

The Add Site dialog box is displayed.

The image shows a Windows-style dialog box titled "Add Site" with a blue title bar and a red close button. The dialog has a light beige background. It contains several fields: a "Name:" label followed by a text box containing "MySite"; a "Type:" label followed by a dropdown menu showing "Remote"; a "Description:" label followed by a large empty text area; a "Connector Type:" label followed by a dropdown menu showing "Performance Data Import"; and a checkbox labeled "Delete temporary data from import tables when job completes?". At the bottom right, there is a "Test Settings" button, and at the very bottom, "OK" and "Cancel" buttons.

3b Specify a name and, optionally, a description.

3c In the *Type* drop-down, select *Remote*.

3d In the *Connector Type* drop-down list, select *Performance Data Import*.

The *Delete data from tables when done?* check box is displayed.

3e To empty the Performance Data Import tables of the data loaded in [Step 2](#), select the *Delete data from tables when done?* check box.

If you choose not to allow PlateSpin Recon to empty these tables, you must manually do this before performing another import.

3f Click *OK*

The new site is added to *Sites* node in the Data Center Explorer.

4 Inventory the machines:

- 4a** In the Data Center Explorer, right-click the remote site you created in [Step 3](#), click *Sites*, then click *Inventory Machines*.

The Inventory Machines in Site *sitename* dialog box is displayed.

- 4b** Select the check boxes next to the machines you want to inventory.

- 4c** Select the required credentials or select *<New Credential>* to add new credentials.

- 4d** Click *OK*.

5 Synchronize site data:

- 5a** In the Data Center Explorer, right-click the required site, click *Sites*, then click *Synchronize Data*.

- 5b** Select the check boxes next to the required machines. To select all machines in a site, select the check box next to the site.

- 5c** Select required data range, scheduling, and recurrence options.

To perform a recurring import, you need to also set up a matching recurring Extract, Transform, and Load process to refresh the data in the PlateSpin Recon Performance Data Import tables. For detailed information, see your software vendor documentation.

- 5d** Click *OK*.

3.12 Collecting Disk Utilization Counters for Windows Machines

In PlateSpin Recon 3.7.2, the deployment and undeployment of Windows Disk Counters add-on on the target Windows machine is not coupled with the starting and stopping of monitoring of target machines. After you inventory the target machines, you can choose to collect data for additional monitoring counters by deploying the appropriate add-on on the target machine. Subsequently, if you do not want PlateSpin Recon to collect the additional data, you can undeploy the add-on from the target machine. PlateSpin Recon 3.7.2 currently provides only the Windows Disk Counters add-on. By default, the disk counters are not collected as part of monitoring of the target machines.

PlateSpin Recon 3.7.2 provides you the flexibility to choose the granularity at which you want to collect the monitoring data for inventoried Windows machines. PlateSpin Recon 3.7.2 allows you to collect additional disk-level and volume-level counters for the inventoried Windows machines by deploying the Windows Disk Counters add-on. If you do not want PlateSpin Recon to collect the additional data, you can undeploy the add-on from the target machine. You can deploy or undeploy the add-on on Windows machines. Although the deploy and undeploy add-on options are available for the Citrix XenServer and Virtual Center machines, the Windows Disk Counters add-on is deployed or undeployed on the Windows virtual machines that are hosted on the Citrix XenServer and Virtual Center machines.

The Windows Disk Counters add-on collects data for the following additional disk utilization counters:

- ♦ PhysicalDisk_TotalSpaceGB
- ♦ PhysicalDisk_FreeSpaceGB
- ♦ PhysicalDisk_UsedSpaceGB
- ♦ PhysicalDisk_PercentFreeSpace

- ♦ PhysicalDisk_PercentUsedSpace
- ♦ PhysicalDisk_UnpartitionedSpaceGB
- ♦ PhysicalDisk_PartitionedSpaceGB
- ♦ PhysicalDisk_PercentUnpartitionedSpace
- ♦ PhysicalDisk_PercentPartitionedSpace
- ♦ LogicalDisk_UsedMegaBytes
- ♦ LogicalDisk_PercentUsedSpace

When you deploy the add-on on the target Windows machines, the disk utilization and disk allocation data of these machines is available in the Disk Allocation and Volume Utilization reports. For more information on the reports, see [Chapter 5, “Generating and Viewing Charts and Reports,”](#) on page 85.

NOTE: By default, PlateSpin Recon collects data for the additional disk utilization counters for non-Windows machines.

Review the following sections:

- ♦ [Section 3.12.1, “Deploying an Add-On,”](#) on page 75
- ♦ [Section 3.12.2, “Undeploying an Add-On,”](#) on page 76

3.12.1 Deploying an Add-On

You can deploy a PlateSpin Recon add-on on a target Windows machine after you inventory the target machine and before or after you start monitoring of the machine.

To deploy an add-on on a target Windows machine:

- 1 In the Data Center Explorer, right-click the target Windows machine, click *Add-On*, then click *Deploy*.
The Deploy Add-On dialog box is displayed.
- 2 In the *Add-On Name* drop-down list, select the add-on you want to deploy. By default, *Windows Disk Counters Add-On* is selected.
PlateSpin Recon currently supports deployment of the Windows Disk Counters add-on only.
- 3 (Conditional) If the Remote Registry service has dependent services, decide how you want to restart the Remote Registry service and the services that are dependent on the Remote Registry service.
 - ♦ If you want the Remote Registry service and its dependent services to be automatically restarted after the deployment of the add-on, select the *Automatically restart the services which are dependent on the Remote Registry Service* check box.
 - ♦ If you want to manually restart the Remote Registry service and its dependent services at a later time, leave the *Automatically restart the services which are dependent on the Remote Registry Service* check box deselected.

IMPORTANT: If the Remote Registry service does not have dependent services, the deployment of the Windows Disk Counters add-on automatically restarts the Remote Registry service on the target machine.

4 Click *OK*.

After the deployment job is completed, the `HWInputMonPofInfo.dll` is deployed on the target 32-bit Windows machine and the `HWInputMonPofInfo64.dll` is deployed on the target 64-bit Windows machine. The DLLs are registered to `perfmon` that runs on the target machines as an extension DLL for collecting the additional counters. The content of

`PlateSpinRecon_installation_directory\services\monitoring\windows` directory located on the PlateSpin Recon Server is copied to the `%systemroot%\PlateSpinRecon-DiskUtil` directory on the target machine. `%systemroot%` is the complete path of the `WINDOWS` directory such as `c:\WINDOWS` or `c:\WINNT`.

5 (Conditional) If you chose to manually restart the Remote Registry service and the services that are dependent on the Remote Registry service in [Step 3](#), restart the services on the target machine.

6 (Conditional) If you deployed the add-on while the target machine was being monitored, perform the following tasks:

6a Stop monitoring the target machine.

6b Start monitoring the target machine.

IMPORTANT: If you do not stop and start monitoring the target machine, the Windows disk utilization counters data is not collected.

7 To view the deployment status of the add-on, right-click the target machine in the Data Center Explorer, click *Properties*, then click the *General* tab. The deployment status is displayed as a message in *Notes*.

For example, the following message is displayed if the deployment of the add-on is successful:

Windows Disk Counter Add-on is deployed.

After you deploy the Windows Disk Counters add-on on a target Windows machine, if you delete the machine object from the PlateSpin Recon Client, following is the resultant on the machine:

- ♦ If the target Windows machine does not have dependents on the Remote Registry service, all the files pertaining to the Windows Disk Counters add-on are deleted, `perfmon` counters are undeployed, and the Remote Registry service is restarted. However, messages indicating these operations are not displayed in the Job Explorer in the PlateSpin Recon Client.
- ♦ If the target Windows machine has dependents on the Remote Registry service, all the files pertaining to the Windows Disk Counters add-on are deleted and `perfmon` counters are undeployed but the Remote Registry service is not restarted. However, messages indicating these operations are not displayed in the Job Explorer in the PlateSpin Recon Client.

3.12.2 Undeploying an Add-On

1 In the Data Center Explorer, right-click the target Windows machine from which you want to undeploy the PlateSpin Recon add-on, then click *Add-On*, then click *UnDeploy*.

The UnDeploy Add-On dialog box is displayed.

2 In the *Add-On Name* drop-down list, select the add-on you want to undeploy. By default, *Windows Disk Counters Add-On* is selected.

PlateSpin Recon currently supports undeployment of the Windows Disk Counters add-on only.

3 (Conditional) If the Remote Registry service has dependent services, decide how you want to restart the Remote Registry service and the services that are dependent on the Remote Registry service.

- ♦ If you want the Remote Registry service and its dependent services to be automatically restarted after the undeployment of the add-on, select the *Automatically restart the services which are dependent on the Remote Registry Service* check box.
- ♦ If you want to manually restart the Remote Registry service and its dependent services at a later time, leave the *Automatically restart the services which are dependent on the Remote Registry Service* check box deselected.

IMPORTANT: If the Remote Registry service does not have dependent services, the undeployment of the Windows Disk Counters add-on automatically restarts the Remote Registry service on the target machine.

4 Click *OK*.

After the undeployment is completed, the %systemroot%\PlateSpinRecon-DiskUtil directory on the target machine is deleted. The HWInputMonPofInfo.dll and HWInputMonPofInfo64.dll are unregistered from the perfmon running on the target 32-bit Windows machine and the 64-bit Windows machine respectively.

5 (Conditional) If you chose to manually restart the Remote Registry service and the services that are dependent on the Remote Registry service in [Step 3](#), restart the services on the target machine.

6 (Conditional) If you undeployed the add-on while the target machine was being monitored, perform the following tasks:

6a Stop monitoring the target machine.

6b Start monitoring the target machine.

NOTE: If you do not stop and start monitoring the target machine, you encounter the following error:

Category does not exist

For more information on how to troubleshoot the error, see [“Monitoring of a Windows machine displays the “Category does not exist” error in the monitoring log” on page 175.](#)

After you undeploy the add-on from the target Windows machine, the add-on deployment message is removed from the *Properties* page of the target machine (*General* tab > *Notes* section).

PlateSpin Recon Server and Client Tools

4

This section provides information about Novell PlateSpin Recon utilities which enable you to better manage login credentials, analyze [data center assets](#), and monitor operations.

- ♦ [Section 4.1, “Server and Client Options,” on page 79](#)
- ♦ [Section 4.2, “Credentials Manager,” on page 79](#)
- ♦ [Section 4.3, “PlateSpin Analyzer,” on page 81](#)

4.1 Server and Client Options

Server and Client settings are available in the Options dialog box (*Tools > Options*). It is highly recommended that you leave the Server settings as is and *ONLY* change their values when directed to do so by Novell PlateSpin Technical Support.

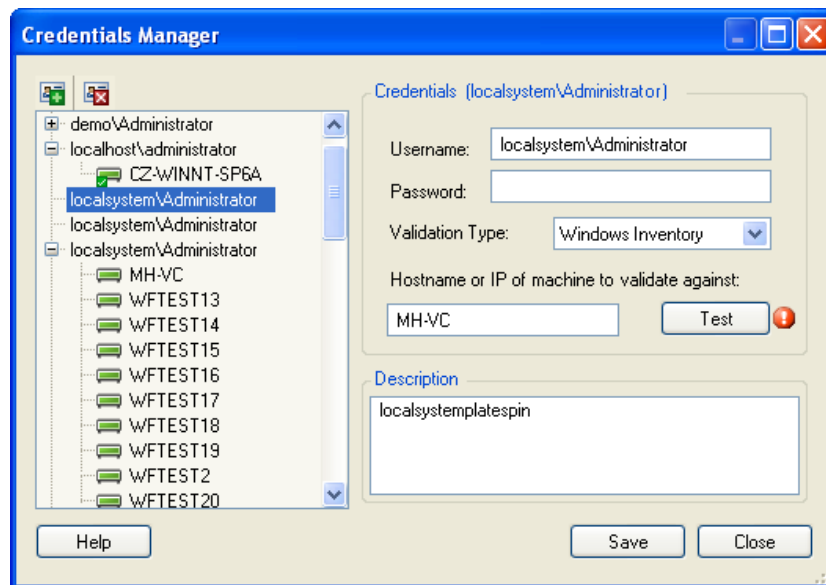
You can configure some Client settings. For instance, it might be necessary to specify a path for various external applications or to change the option to check for a software update.

To reset Server and Client options to their default values, click *Reset All*.

4.2 Credentials Manager

Use the Credentials Manager to create and manage the credentials required to access servers. You can create new credentials, delete saved credentials, or edit credentials in the Credentials Manager dialog box.


Figure 4-1 *Credentials Manager dialog box*



4.2.1 Creating a Credential

- 1 In the PlateSpin Recon Client, click the *Tools* menu > *Credential Manager*.

The Credential Manager window is displayed.

- 2 Click the *New* icon .

- 3 In the Credentials pane, specify the username and password.

NOTE: If you want to add a Windows Inventory credential, you must specify the domain name along with the username in any of the following formats: *domain_name\username* or *username@domain_name*. For example, *platespin\reconuser*.

- 4 In the *Validation Type* drop-down list, select the validation type appropriate to the platform. The available options are Windows Inventory, Unix Inventory, Virtual Center Inventory, VMware ESX3/4 Inventory, Novell Network Inventory, NetWare LSA, Citrix XenServer Inventory, Database or Other.

By default, *Other* is selected

- 5 (Optional) Specify the hostname or the IP address of the machine to validate the credentials against. Click *Test* to validate the specified credentials.
- 6 (Optional) To distinguish the credentials from other credentials with identical user names, specify a description.
- 7 Click *Save* to store the credentials.

The credential is stored in PlateSpin Recon database, and is listed in the left pane of the Credential Manager.

- 8 (Optional) To add a new credential, repeat [Step 1](#) through [Step 7](#).
- 9 Click *Close*.

4.2.2 Deleting a Credential

You can delete a credential only if it is not assigned to a server.


NOTE: If credentials are assigned to servers, a plus sign (+) is displayed next to them in the left pane. To see which machines have those credentials assigned to them, click the plus sign or double-click the credential.

To delete a credential:

- 1 In the PlateSpin Recon Client, click the *Tools* menu > *Credential Manager*.

The Credential Manager window is displayed.

- 2 Delete the credential in any of the following ways:

- ♦ In the left pane, select the credential you want to delete, then click the *Delete* icon  or press the Delete key.
- ♦ In the left pane, right-click the credential you want to delete, then click *Delete*.

4.3 PlateSpin Analyzer

The PlateSpin Analyzer quickly analyzes networks and determines whether PlateSpin Recon can successfully inventory, monitor, or convert the specified Windows and UNIX/Linux servers. It performs two different types of tests: Inventory and Conversion.

- ♦ [Section 4.3.1, “Using PlateSpin Analyzer,” on page 81](#)
- ♦ [Section 4.3.2, “Inventory and Monitoring Tests,” on page 83](#)
- ♦ [Section 4.3.3, “Conversion Tests,” on page 84](#)

4.3.1 Using PlateSpin Analyzer

- 1 Click *Tools > Analyze Servers*.

The PlateSpin Analyzer window is displayed.


- 2 In the *Select Credential* drop-down list, you can choose to add a new credential or to use the credentials and machines specified in a CSV file.

If you choose to create a new credential, continue with [Step 3](#). If you choose to use the CSV file, skip to [Step 4](#).

- 3 (Conditional) If you chose to create a new credential in [Step 2](#), do the following:

- 3a In the Create Stored Credentials dialog box, specify the username and password.

Figure 4-2 Create Stored Credentials dialog box



NOTE: If you want to analyze Windows machine, you must specify the domain name or the host name along with the username in the following formats: *domain_name\username* or *hostname\username*.

- 3b In the *Validation Type* drop-down box, select *Windows Test* or *Unix Test*.

PlateSpin Analyzer does not support analysis of Unix or Linux servers, but they can be used for credential validation.

3c (Optional) Type a description for the credentials.

3d Click *Save*.

NOTE: If *Save* is disabled, verify that the credentials are entered in the correct format.

3e In the *Domain Name* drop-down list, select one of the available options: *Domain Name*, *IP/Host Name*, *IP Range*.

3f Type a value in the empty field appropriate for the selection in [Step 3e](#). For example, if you selected *Domain Name*, enter your domain name.

3g Continue with [Step 5](#).

4 (Conditional) If you chose to use the credentials and machines specified in a CSV file in [Step 2](#), type the complete path of the CSV file in the empty field. You can use the adjacent browse button to browse for and select the file.

Ensure to follow these guidelines as you work with the CSV file:

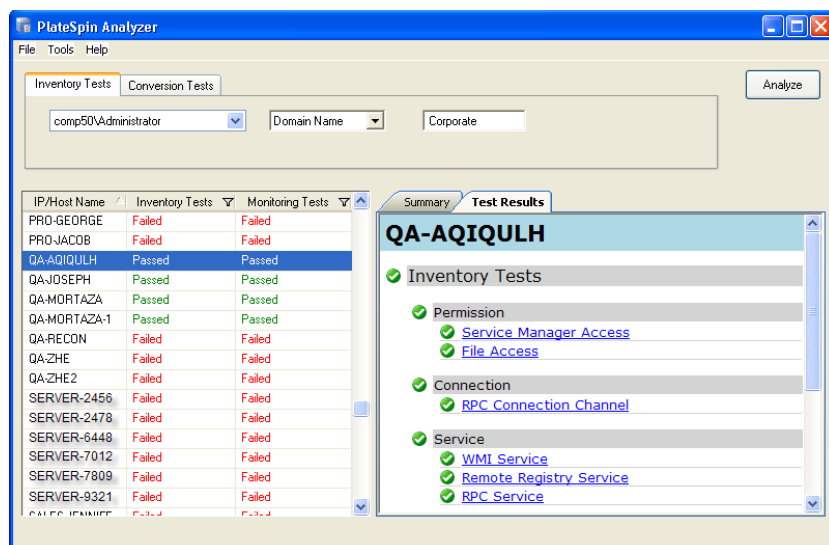
- ♦ The values must be specified in the following order:
 - Hostname or IP address
 - <Hostname or IP address or domain name>\Username
 - Password
 - Description
 - Platform
 - TRUE or FALSE
- ♦ It is mandatory to specify the TRUE or FALSE value for UNIX and Windows platforms in the CSV file. However, the value is considered only for the UNIX platform, and is ignored for the Windows platform. If you specify the value as TRUE for a UNIX platform, then the sudo privileges are assigned to the specified user.
- ♦ Each line must contain the custom fields for a machine.
- ♦ The value for the Machine custom field can be IP address or the hostname; however, it must be unique.
- ♦ The values on each line must be separated by a delimiter such as a comma (,).

A sample CSV file is as follows:

```
DEV-SINISAR, platespin\user1, abc, Description test, Windows, FALSE
```

5 Click *Analyze* and wait for the process to complete.

Figure 4-3 PlateSpin Analyzer - Analysis



When the analysis is complete, the left pane displays either *Passed* or *Failed*. To view result details, click the *Test Results* tab in the right pane. If any of the tests returned a failed result, you must resolve the cause before PlateSpin Recon can successfully inventory the host. A failed result typically indicates that there are insufficient permissions or that the required services are not running.

4.3.2 Inventory and Monitoring Tests

Inventory tests are divided into three main groups, each containing a subset of tests.

Table 4-1 Inventory Test groups

Main Test Group	Test Group Subset	Description
Permission	Service Manager Access	Verifies that the user has permission to open the Service Manager.
	File Access	Verifies that the correct access permission is granted on the System32 directory of the machine being inventoried.
Connection	RPC Connection Channel	Determines whether an RPC connection channel can be established by resolving the DNS/NETBIOS name.
Service	WMI Service	Determines whether the WMI (Windows Management Instrumentation) service is enabled and running on the client machine(s) being inventoried.
	Remote Registry Service	Verifies that the Remote Registry service is enabled on the client machine(s).
	RPC Service	Verifies that the RPC (Remote Procedure Call) Service is enabled on the client machine(s).

Monitoring tests are divided into two main groups, each containing a subset of tests:

Table 4-2 *Monitoring test groups*

Service	RPC Service	Verifies that the RPC service is running on the client machine(s).
	DCOM Service	Verifies that the DCOM service is running on the client machine(s).
Port	TCP Port 135	Determines whether TCP port 135 and 445 are available. To obtain more conclusive information, consider running Port Scanner for all TCP ports above 1024.
	TCP Port 445	

4.3.3 Conversion Tests

Conversion Tests are similar to those performed by the PlateSpin Portability Suite PlateSpin Analyzer tool. For more information, refer to Knowledge Base article [Q7920478](http://www.novell.com/support/viewContent.do?externalId=7920478) (<http://www.novell.com/support/viewContent.do?externalId=7920478>).

Generating and Viewing Charts and Reports

5

When PlateSpin Recon has collected inventory and monitoring data, the data can be displayed for analysis in charts or customizable reports. The Chart Viewer and the Report Explorer display collected information in a variety of formats.

- ♦ [Section 5.1, “Chart Viewer,” on page 85](#)
- ♦ [Section 5.2, “Report Explorer,” on page 89](#)
- ♦ [Section 5.3, “Editing Report Templates,” on page 97](#)
- ♦ [Section 5.4, “View Customization,” on page 106](#)
- ♦ [Section 5.5, “Generating Reports,” on page 110](#)
- ♦ [Section 5.6, “Document Viewer,” on page 114](#)

5.1 Chart Viewer

Charts provide a graphical representation of workload and utilization data. To view a chart for a given server, select the server in the Data Center Explorer or click a cell belonging to the server in a report view. You can also select several servers in the Data Center Explorer (hold down Ctrl or Shift while you select the servers) to render charts for multiple servers. In this case, a chart is not automatically rendered, but a *Create Chart* button is displayed in the Chart Viewer. Click *Create Chart* to render the chart.

Review the following sections to work with charts:

- ♦ [Section 5.1.1, “Chart Controls and Commands,” on page 86](#)
- ♦ [Section 5.1.2, “Expression Types,” on page 87](#)
- ♦ [Section 5.1.3, “Level of Detail,” on page 87](#)
- ♦ [Section 5.1.4, “Data Type Statistics,” on page 88](#)
- ♦ [Section 5.1.5, “Server Instance,” on page 88](#)
- ♦ [Section 5.1.6, “Date Range Presets,” on page 88](#)
- ♦ [Section 5.1.7, “Forecasting,” on page 88](#)















All times displayed in the Chart Viewer are local times. Times shown in reports are in UTC time.



5.1.1 Chart Controls and Commands

Figure 5-1 Chart Controls



Table 5-1 Chart Viewer Commands

Command	Description
	Print Preview. Preview chart before printing.
	Print. Print the chart.
	Copy To Clipboard. Copy the chart to the clipboard.
	Line. Display straight lines between points of continuous data.
	Curve. Display curved lines between points of continuous data.
	Candlestick. View the variability of sample data. This button appears only for the candle chart.
	Stacked Bar. View cumulative workload data from multiple servers. This is enabled for Expression types other than Utilization.
	Stacked Area. View cumulative workload data from multiple servers. This is enabled for Expression types other than Utilization.
	Legend Box. Give a visual cue and description for each data object. The legend is very useful for quickly glancing through the data on a chart.
	Point Labels. Display labels for the values in a chart.
	Zoom. Switch between normal and magnified view.
	Zoom In/Out. Adjust the granularity of the chart for the collected data. See Section 5.1.3, "Level of Detail," on page 87 .
	Navigate Beginning/Back/Forward/End. Scroll the visible data.
	Historical Data Settings. Select the historical data time range upon which to base forecasting.

Command	Description
	Forecast Settings. Select the forecasting time period and algorithm. This button is only enabled for daily or weekly Series charts or for Profile charts.
	Refresh All Data From Server. Refresh the view.

5.1.2 Expression Types

Expressions are individual data series that can be viewed on the currently displayed chart. Select an available Type/Expression combination and the chart is rendered to display the new metrics. The list of available expressions varies depending on the expression type and machine selected.

Table 5-2 *Expression types*

Type	Expressions Available	Description
Workload	For a complete list of expressions available for Workload, see Section B.4.3, "Workload," on page 196.	The operating system, applications, and data stack running on a server, with the quantified dimensions of Processor Speed, Memory, Disk Space, Network, and Time.
Utilization	For a complete list of expressions available for Utilization, see Section B.4.2, "Utilization," on page 195.	Characterizes the fit between workload and resources (% utilization).
Counter	Counter values are named for the Windows Performance Monitor counters used to retrieve them from Windows systems. The name consists of the Performance Monitor category joined by an underscore to the counter name. For example, Memory/PagesPerSecond becomes Memory_PagesPerSecond. For a complete list of the available counters, see "Platform Counters" on page 191.	Primitive measure of performance collected from the operating system.

5.1.3 Level of Detail

From the drop-down list, select from the following levels of granularity (in terms of time) for collected data:

- ♦ **5 Min:** Averages of data collected during five-minute intervals. (Series charts only.)
- ♦ **1 Hour:** Summaries of data collected during one-hour intervals. (Series and Candle charts only.)

- ♦ **1 Day:** Characterizes workloads over a day. This comprises 24 values that correspond to each hour of the day (0 to 23). Each value is calculated by taking the peak value from all values observed for that hour. (All chart types.)
- ♦ **1 Week:** Characterizes workloads over a week. (All chart types.)

5.1.4 Data Type Statistics

The following data type statistics are collected for each hour and are available for selection when a 1 hour or higher zoom level is selected: This option is available only for series charts.

- ♦ **AVG:** Average of all values.
- ♦ **MAX:** Maximum sustained load.
- ♦ **MIN:** Minimum sustained load.
- ♦ **STDEV:** Standard deviation.

NOTE: Standard deviation is not calculated for Virtual Center counters.

- ♦ **Peak AVG:** : (1 Day or 1 Week only) Peak average represents the averages of the values for the highest period for each 12-hour interval or each day in the reporting period.

5.1.5 Server Instance

Click the drop-down list to see all the instances of the currently selected resource (expression). Select each instance to see chart data for that resource. If multiple servers were used to render the chart, all instances of the resources for all servers are displayed. The instances are displayed as legends in the Series and Profile charts, and in the drop-down list located in the tool bar for candle charts.



5.1.6 Date Range Presets

Select from the drop-down list *Last 3 days*, *Last 7 days*, *Last Month*, or *All Data*. Select *<Custom>* to open the Advanced Date Filter Settings dialog box for more specific date range options.

5.1.7 Forecasting

Forecasting is available for series and profile charts, for the 1 day and 1 week level of detail. At least three days of monitoring data are necessary to use forecasting in charts.

To add forecasting to a chart view:

- 1 Display a *Profile* chart.
- 2 Click the *Date Range* drop-down list and select a time period to base the forecast on.
- 3 (Optional) Click *Historical Data Settings*  for advanced date filter settings.
The Advanced Date Filter Settings dialog box is displayed. Choose any desired settings and click *OK*.
- 4 Click the *Forecast Period* drop-down list and select a time period for the forecast.
- 5 (Optional) Click *Forecast Settings*  for advanced custom forecast settings.

The Advanced Forecast Settings dialog box is displayed. Choose any desired settings and click *OK*.

- 6 Click the *Forecast > Trend* drop-down list and select a forecast trend or regression method.

5.2 Report Explorer

The Report Explorer is used to organize and generate reports.

- ♦ [Section 5.2.1, “Report Explorer Template Nodes,” on page 89](#)
- ♦ [Section 5.2.2, “Commands,” on page 92](#)
- ♦ [Section 5.2.3, “Analysis Reports,” on page 93](#)
- ♦ [Section 5.2.4, “VM Disk Partition Alignment Reports,” on page 93](#)
- ♦ [Section 5.2.5, “Working with Report Templates,” on page 96](#)

5.2.1 Report Explorer Template Nodes

The Report Explorer hierarchy consists of template nodes that have various types of report templates.

Table 5-3 *Report Explorer Template Nodes*

Template Node	Report Templates	Node Description
Analysis	Application - Exchange Workloads	For reporting dynamic data such as hardware utilization and workload sizes.
	Application - SQL Server Workloads	
	Disk - Average 30 Day Forecast	
	Disk - Peak	
	Hardware Utilization - Average 30 Day Forecast	
	Hardware Utilization - Peak	
	Memory - Average 30 Day Forecast	
	Memory - Overutilized	
	Memory - Peak	
	Memory - Underutilized	
	Network -Average 30 Day Forecast	
	Network - Peak	
	Processor - Average 30 Day Forecast	
	Processor - Overutilized	
	Processor - Peak	
	Processor - Underutilized	
	Resource Utilization Trend	
	Volume Utilization - Average	
	Volume Utilization - Peak	
	Workload Size - Average 30 Day Forecast	
	Workload Size - Peak	
Chargeback	Cost Allocation	For reporting the transfer of charges incurred by the IT department back to the business units that are using services. See Section 7.4, “Chargeback Reporting,” on page 161.
Consolidation	Project	For reporting on projects, scenarios, and servers in the Consolidation Planning Module. See Section 6.7, “Viewing Consolidation Reports,” on page 146 for information on how to use templates within this node.
	Scenario Summary	
	Workload Assignment	

Template Node	Report Templates	Node Description
Consolidation Project Samples	Good Consolidation Candidates	Commonly used sample consolidation reports.
	Inventory Summary	
	Poor Consolidation Candidates	
	Utilization Summary	
	Workload List with Conditional Formatting	
Diagnostics	Error Report	For diagnosing errors that cause monitoring data to be missing from a chart.
	Missing Data - Detailed	
	Missing Data - Summary	
Disk	Disk Allocation	The Disk Allocation report is generated based on the last collected monitored data. For reporting the size and type of physical disks on machines.
	Disk Inventory	
Inventory	Hardware - Obsolete	For reporting static data such as operating system, processor speed, memory, disk space, network, power, and cooling.
	Hardware Inventory	
	Power and Cooling	
	Virtual Machine Server Candidates	
	Virtual Machine State	
	VM Files Report	
	VM Snapshots Report	
	VM Sprawl Report	
Reclamation	Disk Reclamation	For reporting on elements that can be reclaimed to ensure efficient use of resources.
	Memory Reclamation	
	Powered On VMs Reclamation	
	Volume Reclamation	
Software	Application > Application Inventory	For reporting installed applications, services, and running processes.
	Application Complete > Application Complete Inventory	
	Processes > Process Inventory	
	Service > Service Inventory	
	Service Complete > Service Complete Inventory	








Template Node	Report Templates	Node Description
Virtual Infrastructure	Cluster Configuration	For reporting on virtual machines and virtual machine hosts.
	Resource Pool Configuration	
	VI Configuration	
	VirtualCenter Report	
	VM Cluster	
	VM Details Report	
	VM Host Partition Alignment	
	VM Partition Alignment	
	VM Resource Pool	
	VM Virtual Infrastructure	
Volume	Volume Inventory	For reporting the space and file system associated with each logical volume on machines.

As shown in [Table 5-3 on page 90](#), within each report node are one or more templates that you can use as is or that you can customize with sorting, filtering, and view selection options. You can add additional Analysis, Chargeback, Disk, Inventory, Software, and Volume report templates.

5.2.2 Commands

The following Report Explorer commands are accessible by right-clicking a report template or a generated report:

Table 5-4 Report Explorer Commands

Commands	Tooltip/Description
	Create Report Template. Create a new report template based on the report node selected.
	Copy Report Template. Copy a report template.
	Edit Report Template. Edit a report by setting group, views, filters, summaries, and forecasts for a report template or a report.
	Generate Report. Generate a report based on the report template selected.
	View Report. View an existing (previously generated) report.
	Send Report. Send a generated report via e-mail or to an FTP server or network location.
	Delete. Delete the selected report or template.

5.2.3 Analysis Reports

Analysis reports use dynamic data such as hardware utilization and workload sizes. To view workloads or utilization of servers that are being monitored, generate an Analysis report. The report is displayed with historical data, which is useful for getting a single value (maximum, minimum, peak, average) for expressions over a specified period of time. Use this data to identify over-utilized or under-utilized server resources.

NOTE: Data must be collected for at least one hour before Analysis reports are generated. If there is insufficient data, the report is displayed without dynamic data.

To get detailed historical data for each data point, click the report cell to view the related detailed chart.

Machines without monitoring data are included in Analysis reports by default to allow complete analysis of all inventoried machines.

Figure 5-2 Sample Analysis Report

Hardware Utilization - Peak							
Machine	Operating System	Memory			Processor		
Host	Type	Size (MB)	Used (%)		Total Speed (MHz)	%	
			PEAK	Hour		PEAK	Hour
DEV-YONINAJ	Windows XP	2048.0	91.3	14	3990.0	55.2	13
DEV-FREETS	Windows 2003	2048.0	36.8	16	5984.0	2.8	16
DEV-NADERS	Windows 2003	2048.0	43.4	18	3990.0	3.8	19
DEV-CYRILM	Windows 2003	2040.0	15.2	14	3740.0	3.4	13
DEV-TONYVL	Windows 2003	3584.0	61.8	16	6384.0	25.3	18
DEV-MICHAELB	Windows XP	2048.0	94.1	18	6384.0	12.2	13
DEV-MIRAG	Windows 2003	2048.0	93.6	19	5984.0	27.8	19
DEV-EDWARDRC	Windows 2003	2048.0	37.4	15	3990.0	1.6	15
DEV-TONYVP	Windows 2003	2048.0	85.2	21	3990.0	17.1	18
DEV-MIRCEAA	Windows XP	2048.0	56.5	20	3994.0	26.5	16
DEV-ARIO	Windows 2003 (64bit)	4032.0	71.6	15	5320.0	31.3	13
DEV-WILSONT	Windows 2003	2048.0	90.6	20	3740.0	39.9	10
DEV-GEORGED	Windows 2003	3072.0	73.2	19	3990.0	24.1	13

5.2.4 VM Disk Partition Alignment Reports

PlateSpin Recon 3.7.2 now provides the following two new reports that show the alignment status of the disk partitions:

- ♦ [“VM Host Partition Alignment Report” on page 93](#)
- ♦ [“VM Partition Alignment Report” on page 95](#)

VM Host Partition Alignment Report

The VM Host Partition Alignment Report displays the disk partition alignment status of ESX servers that are supported by PlateSpin Recon. The following table describes the data that is generated and displayed by the VM Host Partition Alignment report:

Category	Fields	Additional Details
Operating System	Version	
Disk	Disk Type	

Category	Fields	Additional Details
Partition	Storage Type	
	Alignment Boundary Size (KB)	Alignment Boundary Size defines the starting offset of a partition. For example, the EMC CLARiiON storage vendor recommends a starting block of 128 to align the partition to the 64-KB boundary. The 64-KB boundary value is used by default according to VMware Recommendations (http://www.vmware.com/pdf/esx3_partition_align.pdf) . However, every disk or disk array does not have the alignment boundary at 64 KB. To know the alignment boundary of the disk of your machine, refer to your storage vendor recommendations or refer to the system guides.
	Name	
	File System	
	Starting Offset (bytes)	
	Alignment Status	The disk partition alignment status is one of the following: <ul style="list-style-type: none"> ♦ Aligned ♦ Unaligned ♦ Indeterminate: Indicates that the starting offset value of the partition has not been inventoried properly or the value is set to zero (0).

The VM Host Partition Alignment report displays the alignment status for external disks that might have different boundary sizes; however, the alignment status might not be correct. To obtain the correct alignment status, ensure that the boundary size of the disks in the report template is correct.

To view or change the boundary size of the disks in the report template:

- 1 In the Report Explorer, browse to *Virtual Infrastructure*, right-click the *VM Host Partition Alignment* report template, then click *Edit Report Template*.

The Edit Report Template dialog box is displayed.

- 2 Click the *Custom Attributes* tab.

- 3 Verify and modify the Alignment Boundary Size of the partition disk, if necessary.

The default alignment boundary size is 64 KB. To know the alignment boundary size of an external disk, contact your storage administrator or the disk vendor.

- 4 Click *OK*.

VM Partition Alignment Report

The VM Partition Alignment Report displays the disk partition alignment status of the Windows and Linux virtual machines that run on the ESX host machine supported by PlateSpin Recon. The disk partition alignment status is applicable for the disks that reside on the VMFS partition of the ESX host machine. The host machine can be inventoried or not. The data of the inventoried virtual machines is grouped under their host, and the data of virtual machines whose host is not inventoried is grouped as VMs without Host.

The following table describes the data that is generated and displayed by the VM Partition Alignment report:

Category	Fields	Additional Details
Operating System	Type	
	Edition	
	Version	
Disk	Disk Type	
	Storage Type	
	Alignment Boundary Size (KB)	<p>Alignment Boundary Size defines the starting offset of a partition. For example, the EMC CLARiiON storage vendor recommends a starting block of 128 to align the partition to the 64-KB boundary.</p> <p>The 64-KB boundary value is used by default according to VMware Recommendations (http://www.vmware.com/pdf/esx3_partition_align.pdf). However, every disk or disk array does not have the alignment boundary at 64 KB. To know the alignment boundary of the disk of your machine, refer to your storage vendor recommendations or refer to the system guides.</p>
Partition	Name	
	File System	
	Starting Offset (bytes)	
	Alignment Status	<p>The disk partition alignment status is one of the following:</p> <ul style="list-style-type: none">♦ Aligned♦ Unaligned♦ Indeterminate: Indicates that the starting offset value of the partition has not been inventoried properly or the value is set to zero (0).

Although reporting of the alignment status for external disks attached to the virtual machines that run on ESX host is not supported, the VM Partition Alignment report also displays the alignment status for external disks. However, the alignment status of the external disks might not be correct if the boundary size of the disks are different. To obtain the correct alignment status, ensure that the boundary size of the disks in the report template is correct.

To view or change the boundary size of the disks in the report template:

- 1 In the Report Explorer, browse to *Virtual Infrastructure*, right-click the *VM Partition Alignment* report template, then click *Edit Report Template*.

The Edit Report Template dialog box is displayed.

- 2 Click the *Custom Attributes* tab.

- 3 Verify and modify the Alignment Boundary Size of the partition disk, if necessary.

The default alignment boundary size is 64 KB. To know the alignment boundary size of an external partition disk, contact your storage administrator or the disk vendor.

- 4 Click *OK*.

5.2.5 Working with Report Templates

- ♦ [“Creating a Report Template” on page 96](#)
- ♦ [“Copying a Report Template” on page 96](#)
- ♦ [“Deleting a Report Template” on page 96](#)

Creating a Report Template

- 1 Right-click the report type node in the Report Explorer for the type of report that you want to create, then select *Create Report Template* from the menu.

The Create Report Template dialog box is displayed.

- 2 Define the template settings, using the procedures and guidelines for editing an existing template. See [Section 5.3, “Editing Report Templates,” on page 97](#).

- 3 Click *OK*.

The new template is displayed under the report type node in the Report Explorer.

Copying a Report Template

- 1 Right-click the report template you want to copy in the Report Explorer and select *Copy Report Template* from the menu.

The Edit Report Template dialog box is displayed with (COPY) added to the template name.

- 2 (Optional) Edit the template name.

- 3 Click *OK*.

The new copy of the template is displayed in the Report Explorer.

Deleting a Report Template

- 1 Right-click the report template you want to delete in the Report Explorer and select *Delete* from the menu.

The Confirm Delete dialog box is displayed.

2 Click *OK*.

The report template is removed from the Report Explorer.

5.3 Editing Report Templates

PlateSpin Recon includes a library of standard, predefined templates available for generating reports. Templates specify the content to include in a report and the form in which the content is presented.

- ◆ [Section 5.3.1, “Report Template Option Tabs,” on page 97](#)
- ◆ [Section 5.3.2, “Chargeback Report Templates,” on page 103](#)
- ◆ [Section 5.3.3, “Consolidation Report Templates,” on page 103](#)
- ◆ [Section 5.3.4, “Consolidation Project Samples Report Templates,” on page 104](#)
- ◆ [Section 5.3.5, “Diagnostics Report Templates,” on page 104](#)
- ◆ [Section 5.3.6, “Disk Report Templates,” on page 104](#)
- ◆ [Section 5.3.7, “Inventory Report Templates,” on page 105](#)
- ◆ [Section 5.3.8, “Reclamation Report Templates,” on page 105](#)
- ◆ [Section 5.3.9, “Software Report Templates,” on page 105](#)
- ◆ [Section 5.3.10, “Virtual Infrastructure Report Templates,” on page 105](#)
- ◆ [Section 5.3.11, “Volume Report Templates,” on page 106](#)

PlateSpin Recon reporting capabilities allow for the customization of most report templates to meet specific business requirements. You can configure the content and summaries included in most reports, as well as the time frame, report presentation, and a number of other elements. Users can save modified templates and generate reports from them.

For all reports except Consolidation reports, you can copy the template and then edit each copy. This preserves the original. Use the modified templates to generate alternate reports.

To edit a report template, right-click the template and select *Edit Report Template* from the menu. The Edit Report Template dialog box is displayed.

5.3.1 Report Template Option Tabs

The options available for customizing report templates vary depending on the type of report template selected.

- ◆ [“Group Tab” on page 98](#)
- ◆ [“Views Tab” on page 98](#)
- ◆ [“Field Filter Tab” on page 98](#)
- ◆ [“Row Filter Tab” on page 100](#)
- ◆ [“Columns Tab” on page 100](#)
- ◆ [“Time Period Tab” on page 100](#)
- ◆ [“Options Tab” on page 100](#)
- ◆ [“Date Filter Tab” on page 100](#)
- ◆ [“Summaries Tab” on page 101](#)

- ♦ “Forecast Tab” on page 101
- ♦ “Custom Attributes” on page 102

Group Tab

Allows you to select filters or groups of servers to include in the report. You must select distinct groups or filter conditions within the *Filters*, *Groups*, or *Sites* nodes. The *Filters*, *Groups*, or *Sites* nodes are the first-level logical groups. If you select only a logical group in a report template, then the generated report will be empty.

For more information about the groups, see [Section 3.4, “Working with Groups,” on page 52](#).

Views Tab

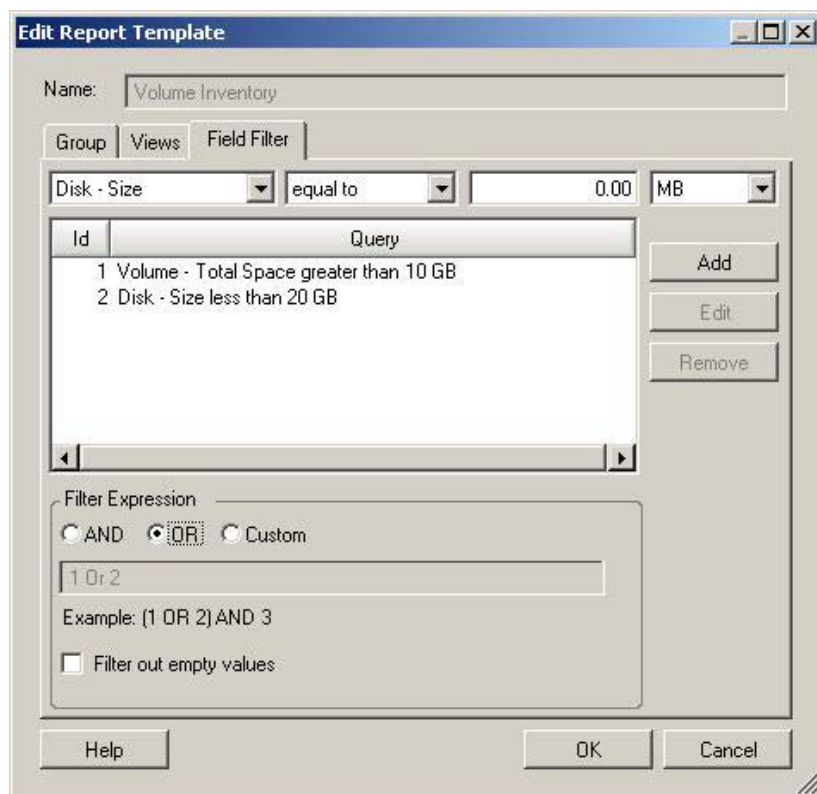
The *Views* tab displays the views available for the report, including default and previously created views. Use the arrow buttons to reorder the views. Use the check boxes in the *Visible* column to determine which views to display in the report.

To add another view to a report, click *Add*. To edit an existing view, select it and click *Edit*. For more information, see [Section 5.4.1, “Creating and Editing Report Views,” on page 106](#). To delete a view, select it and click *Delete*.

Field Filter Tab

You can filter servers based on the constraints you set and use expressions to specify how the constraints should be used and interpreted to narrow or expand the report contents.

Figure 5-3 Field Filter Tab



To add a query statement:

- 1 Enter query constraints in the fields and drop-down lists at the top of the Field Filter page.
There might not be a units field or drop-down list, depending on the resource selected in the first drop-down list.
- 2 Click *Add*.
The query statement is added to the list of queries.

To edit a query statement:

- 1 Select the query statement.
- 2 Click *Edit*.
The query settings are copied to the fields at the top of the page.
- 3 Edit the settings as desired.
- 4 Click *Add*.
The selected query statement is updated with the new settings.

To remove a query statement:

- 1 Select the query statement you want to delete.
- 2 Click *Remove*.
The selected query statement is deleted.

The filter expression determines how the query statements are evaluated.

Filters affect which servers are made available for a report views.

Table 5-5 *Filter Expressions*

Filter Expression	Description
AND	All query statements must be true for a server to be added to the gathered pool of servers. For example, if there are several query statements, the first one is evaluated to gather some servers. Then the second query statement is evaluated against the gathered servers, removing any from the pool that don't agree with the new query statement. This continues through all the query statements.
OR	All query statements are evaluated individually. For example, if there are several query statements, the first one is evaluated to gather some servers. Then the second query statement is evaluated adding any new servers not gathered by the first statement to the pool of gathered servers. This continues through all the query statements.
Custom	A Custom expression allows you to specify how to evaluate the query statements. A mixture of OR and AND criteria can be used. You can use parentheses to control the evaluation. Use the query statement ID numbers in your expression. If Custom is selected, the accompanying field cannot be empty.

Select the *Filter out empty values* check box if you want to exclude machines that have null values for the fields included in the specified query.

Row Filter Tab

Used to determine which rows are included or excluded from the report. This is different from *Field Filter* tabs, which control which machines are included or excluded from the report, but the guidelines for managing Field Filter criteria can also be applied here. See [“Field Filter Tab” on page 98](#).

This option is available only for the Application Complete and Software Complete reports.

Columns Tab

Allows you to customize the columns displayed in the reports generated, based on these templates. See [Step 1 on page 115](#) for more information on the settings in this page.

This option is available only for the software report templates.

Time Period Tab

Allows you to specify the time period to use for these reports. The settings on this page are the same as the top part of the Date Filter page. See [“Date Filter Tab” on page 100](#) and [Section 5.3.5, “Diagnostics Report Templates,” on page 104](#).

This option is available only for the diagnostic report templates.

Options Tab

Allows you to specify the counter setting report mode for these reports. See [Section 5.3.5, “Diagnostics Report Templates,” on page 104](#).

This option is available only for the missing data diagnostic report templates.

Date Filter Tab

Allows you to set the date and time range for calculating summary data values in a report. For example, select the Last 1 Month(s) collection range, Monday to Friday, and 9:00 AM to 5:00 PM to view workload data summarized over business hours for all servers that are monitored:

Figure 5-4 Date Filter Tab

Summaries Tab

Summarizes average (AVG), maximum (MAX), minimum (MIN) and standard deviation (STDEV) data types. AVG data can be calculated by Average, Peak Hour, or Total of the data collected.

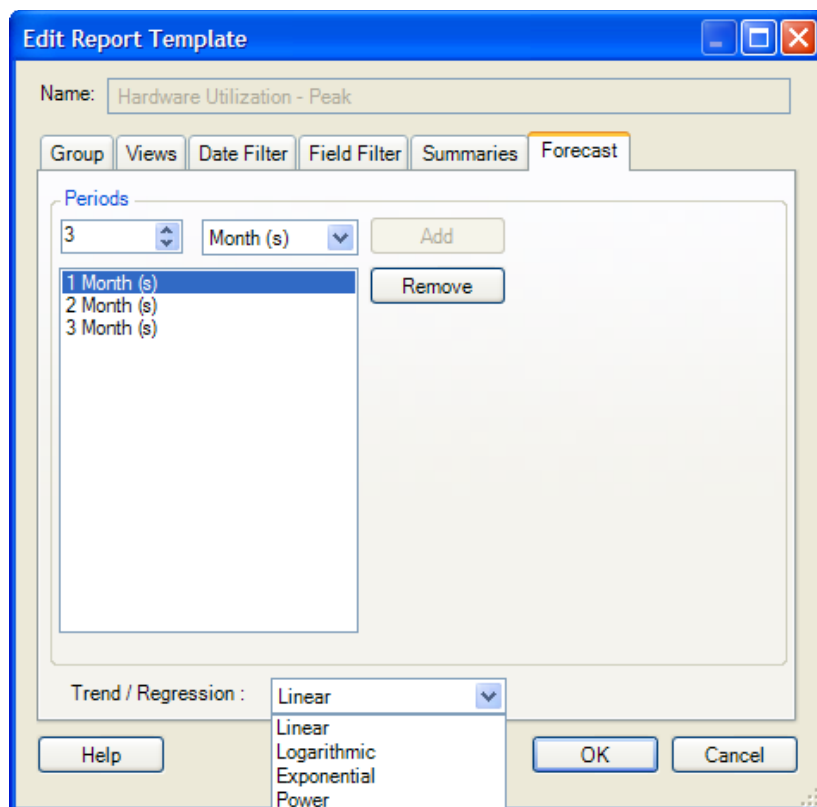
Calculation types:

- ♦ **Average:** Average of all hourly values across business days and business hours.
- ♦ **Peak Hour:** A calculation based on the highest hourly value.
- ♦ **Total:** Sums all the data points in the time frame. For example, Disk IOs per second, for a 1-hour period equates to 10 disk IOs per second X 60 seconds X 60 minutes = 36000. Select the *Summary Rows* column in the report to display the columns containing calculated data.

Forecast Tab

Projects workloads and utilization into the future based on summary data for a historical period. This lets data center architects anticipate future needs and plan accordingly.

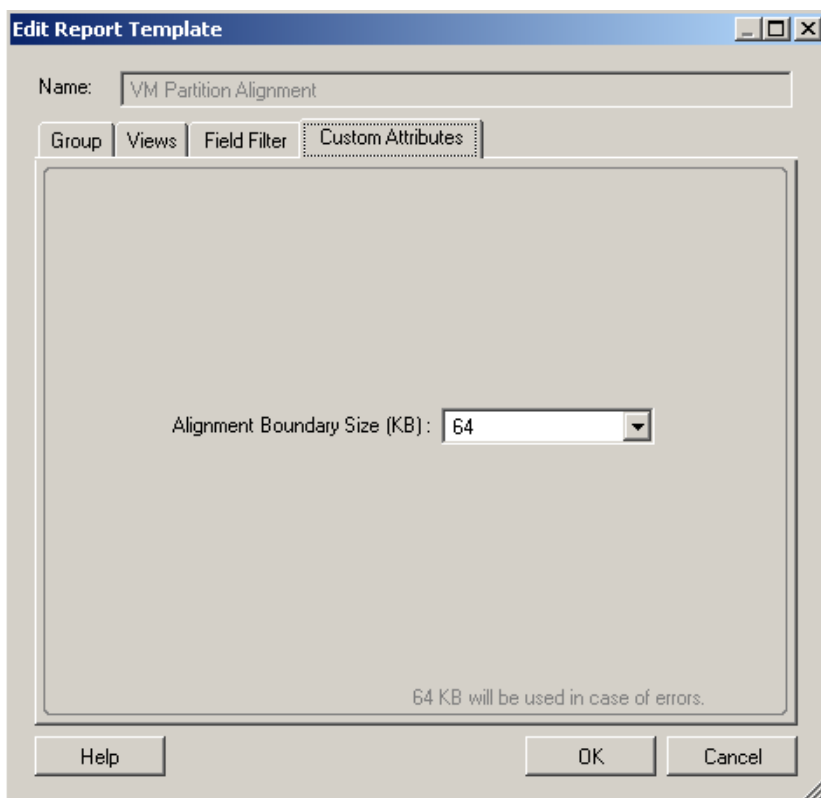
Figure 5-5 *Forecast Tab*



Custom Attributes

Allows you to select or specify the size in KB to which the disk partition needs to be aligned. The default value is 64 KB. The alignment boundary size that you specify is used to calculate the disk partition offset value. This tab is available only for VM Partition Alignment and VM Host Partition Alignment report templates.

Figure 5-6 Custom Attributes Tab



5.3.2 Chargeback Report Templates

Edit Chargeback report templates in the same way as Analysis report templates. See [Section 5.3.1, “Report Template Option Tabs,”](#) on page 97 for information about the *Group*, *Views*, and *Field Filter* tabs.

Chargeback reports support only table views. See “[Creating and Editing a Table View](#)” on page 107 for information about table views. See [Section 5.2.5, “Working with Report Templates,”](#) on page 96 for information about the *Date Filter*, *Summaries* and *Forecast* tabs.

5.3.3 Consolidation Report Templates

Project, Scenario Summary, and Workload Assignment reports are based on editable templates. Unlike Inventory and Analysis report templates, there is only one version of a given Consolidation report template available at a time. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*.

For more detailed information on the other configuration options of the consolidation report template, see [Section 5.6.2, “Customizing Reports,”](#) on page 115.

5.3.4 Consolidation Project Samples Report Templates

Consolidation Project Samples reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*. For information on the option tabs, see [Section 5.3.1, “Report Template Option Tabs,” on page 97](#).

For more detailed information on the other configuration options of the Consolidation Project Samples report templates, see [Section 5.6.2, “Customizing Reports,” on page 115](#).

5.3.5 Diagnostics Report Templates

Edit Diagnostics report templates by changing options in the Edit Report Template dialog box. To access this dialog box, right-click a Diagnostics report template in the Report Explorer, and on the menu, click *Edit Report Template*.

The available dialog box tabs depend on the *Diagnostics* report template selected:

- ♦ **Group:** See [“Group Tab” on page 98](#).
- ♦ **Views:** See [Section 5.4.1, “Creating and Editing Report Views,” on page 106](#).
- ♦ **Time Period:** Select a time period upon which to base the report. Any intervals during that time period that are missing data are included in the report.
- ♦ **Options:** For Missing Data reports only. Select a Report Mode to determine which intervals to include in the report:
 - ♦ **Online:** Includes intervals during which all counters are missing data.
 - ♦ **CPM:** Includes intervals during which CPM counters are missing data.
 - ♦ **Full:** Includes intervals during which any counter is missing data.

5.3.6 Disk Report Templates

Disk reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*. For more information on the option tabs, see [Section 5.3.1, “Report Template Option Tabs,” on page 97](#).

For more detailed information on the other configuration options of the Disk report templates, see [Section 5.6.2, “Customizing Reports,” on page 115](#).

5.3.7 Inventory Report Templates

Inventory reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*. For more information on the option tabs, see [Section 5.3.1, “Report Template Option Tabs,” on page 97](#).

For more detailed information on the other configuration options of the Inventory report templates, see [Section 5.6.2, “Customizing Reports,” on page 115](#).

5.3.8 Reclamation Report Templates

Reclamation reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*.

For more detailed information on the other configuration options of the Reclamation report templates, see [Section 5.6.2, “Customizing Reports,” on page 115](#).

5.3.9 Software Report Templates

Software reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*. For more information on the option tabs, see [Section 5.3.1, “Report Template Option Tabs,” on page 97](#).

For more detailed information on the other configuration options of the Software report templates, see [Section 5.6.2, “Customizing Reports,” on page 115](#).

5.3.10 Virtual Infrastructure Report Templates

Virtual Infrastructure reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*. For more information on the option tabs, see [Section 5.3.1, “Report Template Option Tabs,” on page 97](#).

For more detailed information on the other configuration options of the Virtual Infrastructure report templates, see [Section 5.6.2, “Customizing Reports,” on page 115](#).

5.3.11 Volume Report Templates

Volume reports are based on editable templates. The templates can be copied or edited. Edit the template as desired before viewing the report.

- 1 In the Report Explorer, right-click the desired report template, then click *Edit Report Template*. The appropriate Default Settings dialog box is displayed.
- 2 Select the required report columns, data sorting options, and summary rows, then click *Save*. For more information on the option tabs, see [Section 5.3.1, “Report Template Option Tabs,” on page 97](#).

For more detailed information on the other configuration options of the Volume report templates, see [Section 5.6.2, “Customizing Reports,” on page 115](#).

5.4 View Customization

This section discusses the customization of report views, specifically table and chart views.

- ♦ [Section 5.4.1, “Creating and Editing Report Views,” on page 106](#)
- ♦ [Section 5.4.2, “Histogram Charts,” on page 108](#)
- ♦ [Section 5.4.3, “Scatter Charts,” on page 109](#)
- ♦ [Section 5.4.4, “Usage Charts,” on page 110](#)

5.4.1 Creating and Editing Report Views

- ♦ [“Adding a View to a Report” on page 106](#)
- ♦ [“Editing a Report View” on page 107](#)
- ♦ [“Removing a View from a Report” on page 107](#)
- ♦ [“Creating and Editing a Table View” on page 107](#)
- ♦ [“Creating and Editing a Chart View” on page 107](#)

Adding a View to a Report

- 1 From the Report Explorer, right-click the report template you want to customize and select *Edit Report Template*.
The Edit Report Template dialog box is displayed.
- 2 Click the *Views* tab.
The Views page is displayed, showing any default or previously created views.
- 3 Click *Add*.
The Create View (Step 1: General Settings) dialog box is displayed.
- 4 Type a unique *Name* for the new view.
- 5 For *Chart* views, proceed to [“Creating and Editing a Chart View” on page 107](#). For *Table* views, proceed to [“Creating and Editing a Table View” on page 107](#).

Editing a Report View

- 1 Select a view on the *Views* page of the Edit Report Template dialog box.
- 2 Click *Edit*.
The Create View (Step 2: Customization) page is displayed for the view type selected.
- 3 For *Chart* views, proceed to [“Creating and Editing a Chart View” on page 107](#). For *Table* views, proceed to [“Creating and Editing a Table View” on page 107](#).

Removing a View from a Report

- 1 Select a view on the *Views* page of the Edit Report Template dialog box.
- 2 To temporarily remove the view, deselect the *Visible* check box. To permanently remove the view, click *Delete*.
- 3 Click *OK*.

Creating and Editing a Table View

- 1 Select *Table* in the *Type* area and click *Next*.
The Step 2: Customization page is displayed.
- 2 On the *General* tab, select the column titles to display in the report.
See [Section 5.6.2, “Customizing Reports,” on page 115](#) for information on arranging columns, conditional formatting, column sorting, and summary rows.
- 3 Use the *Forecast* tab to create forecast columns in the report view.
Create forecast columns and display: Clear the check box if forecast columns are not required. When selected, this option enables the following column options:
 - ♦ **Original Values:** Displays the original values.
 - ♦ **Change (%):** Displays the percentage of change from the original values.
 - ♦ **Change (Units):** Displays the change from original values in units.Forecasting is only supported in table views.
- 4 Use the *Groups* tab to group rows in the table view. For information on report groups, see [“Adding or Modifying Groups” on page 117](#).
- 5 Use the *Matrix* tab to add Matrix columns to the table view. For information on report matrixes, see [“Adding or Modifying Matrixes” on page 117](#).
- 6 Click *Finish* to add the view to the list of available views in the Edit Report Template dialog box.

Creating and Editing a Chart View

- 1 Select *Chart* in the *Type* area.
- 2 Select a chart type from the drop-down list (histogram, scatter chart, or usage chart. For analysis reports only).
- 3 Select an expression from the drop-down list. The expressions listed vary depending on the chart type selected.

- 4 Click *Next* to display the view. For the histogram chart type, the Settings dialog box is displayed. Use it to set various options as described in [Section 5.4.2, “Histogram Charts,” on page 108](#).
- 5 Click *OK*.
The Create View dialog box (Step 2: Customization page) is displayed, showing the view.
 - ♦ For a *Scatter* chart, the Chart View Settings dialog box is also displayed. Use it to set various data, summary, and sort options as described in [Section 5.4.3, “Scatter Charts,” on page 109](#).
 - ♦ For an Analysis report, the *Usage* chart type is also available. Select *Usage* and click *Next*. The Usage View Settings dialog box is displayed. Use it to set various options as described in [Section 5.4.4, “Usage Charts,” on page 110](#).
- 6 Click *Customize* to return to the chart-appropriate Settings dialog box and modify the specified options.
- 7 Use the various buttons at the top of the chart in the Create View dialog box to print or modify the chart.
- 8 Click *Finish*.

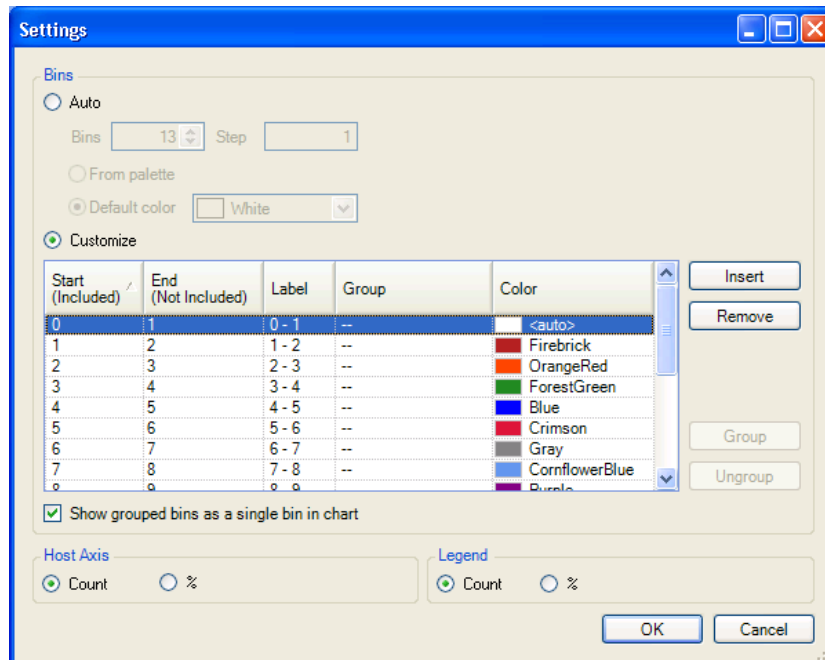
The view is displayed in the list of available views in the Edit Report Template dialog box.

5.4.2 Histogram Charts

In the Create View (Step 2: Customization) dialog box, click *Customize* to modify the contents of the histogram.

The Settings dialog box is displayed.

Figure 5-7 Settings Dialog Box



The default setting for the histogram includes five bins with a Step size of 1.

Auto

Select *Auto* to specify a different number of Bins or Step sizes to be used in the histogram. Click *From palette* to automatically generate colors for the histogram bars from the palette. Click *Default color* to choose one default color for all histogram bars.

Customize

Select *Customize* to modify the upper or lower boundaries of any bin needing modification. Optionally, specify custom labels and colors for the corresponding histogram bars.

Click the *Insert* button to insert a new bin below the currently highlighted bin. The appropriate *Start*, *End*, and *Label* columns are automatically adjusted to accommodate the new bin, which can be customized.

Click the *Remove* button to remove the currently highlighted bin. The appropriate *Start*, *End*, and *Label* columns are automatically adjusted.

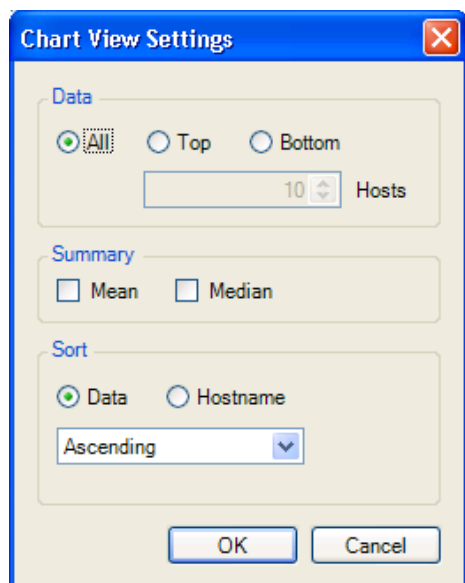
Select any number of consecutive or non-consecutive bins and click the *Group* button to group them together. Enter the group's name in the *Group* column for any one of the bins. The name is automatically assigned to all members of the group.

Click the check box to display the grouped bins as a single bin in the chart.

5.4.3 Scatter Charts

In the Create View (Step 2: Customization) dialog box, click the *Customize* button to modify the contents of the scatter chart. The Chart View Settings dialog box is displayed.

Figure 5-8 Chart View Settings



In the *Data* area, choose how much data to include in the chart. Select *All* to include all data, or select *Top* or *Bottom* to limit the chart to the highest or lowest values of data. For *Top* or *Bottom*, specify the number of values, up to 25, to include.

In the *Summary* area, select *Mean* and/or *Median* to include those lines in the chart. These values are based on all data, not just displayed values.

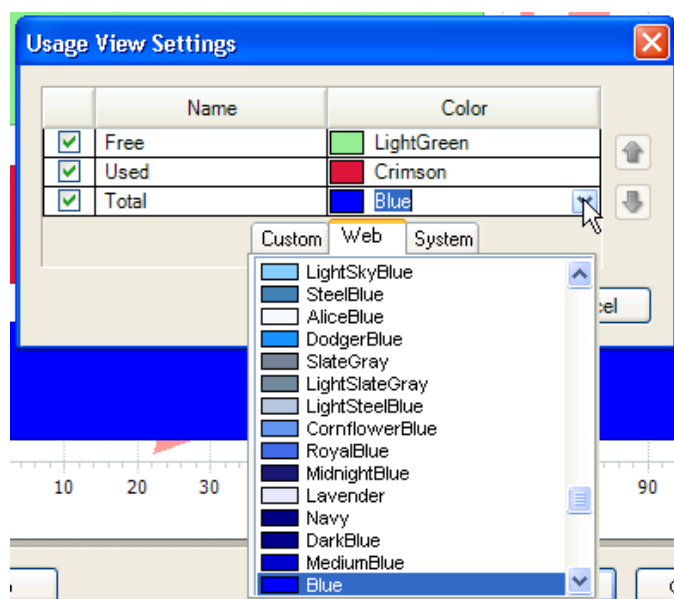
In the *Sort* area, choose to sort the chart in ascending or descending order according to data quantity, or alphabetically by hostname.

Click *OK* to display the chart.

5.4.4 Usage Charts

Usage charts are unique to Analysis report templates. In the Create View (Step 2: Customization) dialog box, click *Customize* to modify the contents of the usage chart. The Usage View Settings dialog box is displayed.


Figure 5-9 Usage View Settings



Use the check boxes to determine which bars to include in the chart. If desired, use the drop-down lists in the *Color* section to assign colors to the bars in the chart.

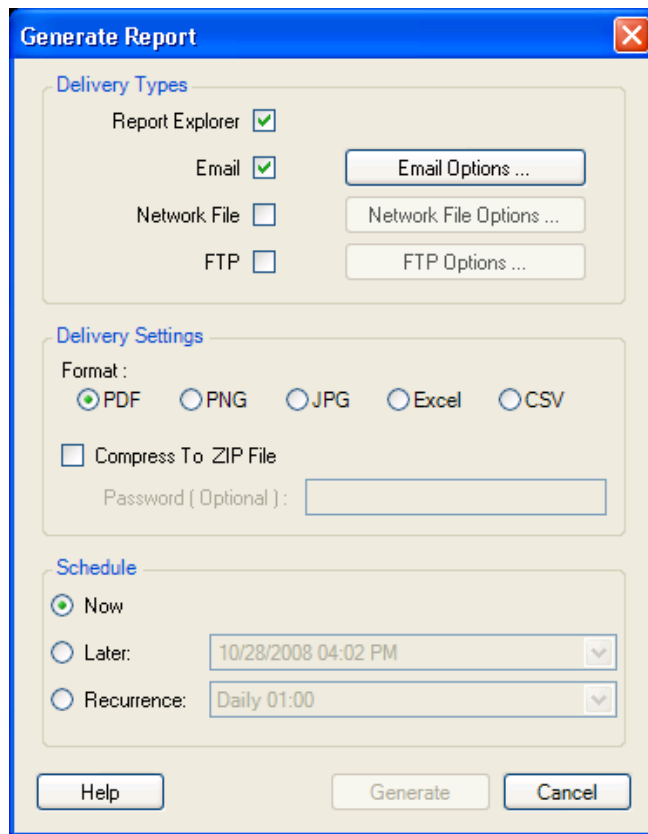
To move a bar to the left or to the right in the report, select it and use the up- or down-arrow. Click *OK* to display the chart.

5.5 Generating Reports


- 1 Select the report template and click *Generate Report* .

The Generate Report dialog box is displayed.

Figure 5-10 *Generate Report Dialog Box*



- 2 Select a delivery types for the report. By default, the report is delivered to the Report Explorer. Other delivery options include:
 - ♦ Sending it by e-mail
 - ♦ Saving it to a network file
 - ♦ Uploading it to an FTP server
- 3 Configure delivery options:
 - ♦ For the Email delivery option, click *Email Options*. The Email Delivery Options dialog box is displayed.
 - ♦ Provide the requested information. Only one recipient is required in the *To* field, although if you enter more than one, you need to separate the addresses with a semicolon. The *CC* and *BCC* fields can be left blank. Any default values supplied for the *Subject* and *Name* can be changed. The *From* Address is mandatory, but the *Reply To* address can be left blank. The *SMTP Server* and *SMTP Port* are mandatory. *Credentials* are only optional if they are not required for your SMTP server, which is unlikely.



The dialog box is titled "Email Delivery Options" and contains several sections for configuring email delivery. The "Recipients" section has fields for "To:", "Cc:", and "Bcc:", with the "To:" field containing "YourName@YourDomain.com;YourName2@YourDomain.com". The "Message" section has a "Subject:" field containing "Recon Report". The "From" section has fields for "Name:" (PlateSpin Recon), "Email:" (MyName@MyDomain.com), and "Reply To:". The "SMTP Settings" section has a "SMTP Server:" field (smtp.MyDomain.com) and a "SMTP Port:" spinner set to 25. The "Credentials" section has a dropdown menu showing "platespin\Administrator". At the bottom are "OK" and "Cancel" buttons.

Email Delivery Options

Recipients

To : YourName@YourDomain.com;YourName2@YourDomain.com

Cc :

Bcc :

Message

Subject : Recon Report

From

Name : PlateSpin Recon

Email : MyName@MyDomain.com

Reply To :

SMTP Settings

SMTP Server: smtp.MyDomain.com

SMTP Port: 25

Credentials

platespin\Administrator

OK Cancel

- ♦ For the *FTP* and *Network File* delivery types, use the adjacent option buttons to open dialog boxes and specify an FTP path or UNC/local folder path. Choose to have later recurrences overwrite the report saved at that location, or append the date and time to the specified filename so that previous reports are preserved.

FTP Delivery Options

File Options

Path : Z:\FTP\

When transferring files :

☒ Always append date and time to filename

☐ Always overwrite existing files

FTP Server

Host : ftp.acme.com

Port : 21

☐ Use Passive Mode

Protocol

☒ FTP ☐ FTPS ☐ SFTP

Credentials

MB-WIN2K3SP1\Administrator

OK Cancel

Network File Delivery Options

UNC or Local Folder Path :

Z:\PowerRecon Reports\

Credentials

MB-WIN2K3SP1\Administrator

When transferring files :

☒ Always append date and time to filename

☐ Always overwrite existing files

OK Cancel

- ♦ For delivery options other than Report Explorer, the default format for the report is *PDF*. Additionally, there are several other available delivery formats.
- 4 If desired, select the option to compress the report in a ZIP file. Optionally, the file can be password-protected.
- 5 You can also set options in the *Schedule* area to specify whether to generate the report immediately, at a later time, or at a recurring interval. For information about scheduling options, see [Appendix C, “Schedule Options,” on page 209](#).
- 6 Click *Generate*.

When the report is generated, it is delivered using the delivery types specified. If this includes Report Explorer, the report is cached for quick viewing at any time. An item is added to the Report Explorer beneath the report template used to generate the report, named with the date and time the report was generated.

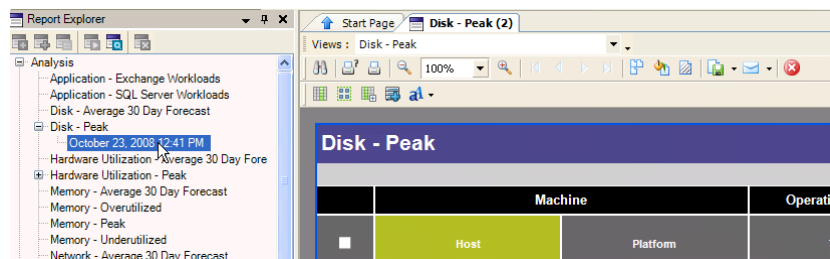
NOTE: Only reports that have tables are successfully exported to Excel format. Reports that contain histograms or charts cannot be exported to the Excel format.

To view a report, double-click the generated instance. The report is displayed in the Document Viewer. For information, see [Section 5.5, “Generating Reports,” on page 110](#). To delete a report, right-click the report and select *Delete Report* on the menu.

5.6 Document Viewer

To display a report in the Document Viewer, generate a report from a report template and then double-click the instance of the report or double-click a report instance you previously generated. See [Section 5.5, “Generating Reports,” on page 110](#).

Figure 5-11 Report Displayed in the Document Viewer

















A new tabbed page is added to the Document Viewer, showing the report through its first available view. To switch to a different view, click the *Views* drop-down list and select a view. For more information on views, see [Section 5.4.1, “Creating and Editing Report Views,” on page 106](#).

- ♦ [Section 5.6.1, “Commands,” on page 114](#)
- ♦ [Section 5.6.2, “Customizing Reports,” on page 115](#)
- ♦ [Section 5.6.3, “Viewing Reports,” on page 118](#)
- ♦ [Section 5.6.4, “Printing and Exporting Reports,” on page 119](#)
- ♦ [Section 5.6.5, “Copying Servers from a Report to a Group,” on page 120](#)

5.6.1 Commands

Table 5-6 Document Viewer Commands

Commands	Tooltip/Description
	Columns. Select columns content for viewing.
	Groups. Add, or modify existing grouping in a report.
	Matrix. Add or modify existing matrixes in a report.


Commands	Tooltip/Description
	Copy to Group. Copy selected servers to a group.
	Font Size. Change font size (small, medium, or large).
	Print/Quick Print. Show the Print Options dialog box to allow you to pick a printer and its settings, or just print the report immediately with the current default printer and settings.
	Search. Opens a Find dialog box for searching.
	Zoom Out, Zoom Percent, Zoom In. Zoom in or out incrementally/decrementally, or specify a zoom setting.
	Multiple Pages. Reduce the view zoom setting so you can view multiple pages of the report simultaneously.
	Color Background. Show a dialog box to allow you to set the background color of the report.
	Watermark. Show a dialog box where you can set a watermark for the report, such as DRAFT, TOP SECRET or an image.
	Export Document. Export to PDF, HTML, Text, CSV, MHT, Excel, RTF, or Image formats.
	Send via Email. Send the report through e-mail as an attachment. The file is saved locally first and then your default mail client is launched with the saved file attached. Possible file formats are PDF, MHT, RTF, Excel, CSV, Text, or an image format.
	Exit. Close the report preview page.


5.6.2 Customizing Reports

After a report is generated, it can be further adjusted and customized. This includes revising which columns to include in the report, changing the font, and adding or modifying grouping or matrixes. These customizations are displayed in the Document Viewer and are included in the report when it is exported, but are not saved to the template.

- ♦ [“Customizing Columns” on page 115](#)
- ♦ [“Changing Report Font Size” on page 116](#)
- ♦ [“Adding or Modifying Groups” on page 117](#)
- ♦ [“Adding or Modifying Matrixes” on page 117](#)

Customizing Columns

- 1 Click *Columns*  in the Document Viewer.
The Report Column Selection dialog box is displayed.



- 2 Select a column and use the arrow buttons to rearrange how it appears in the report (the higher the column in the tree hierarchy, the more to the left it appears in the report). Conditional formatting can be selected for numeric fields. Select custom colors and breakpoints, or use the default settings.

The report will highlight the following:


- ♦ All the values less than the *Low Breakpoint* with the *Low Color*
- ♦ All the values between the *Low* and *High Breakpoints* with the *Mid Color*
- ♦ All the values greater than the *High Breakpoint* with the *High Color*.

Selecting *Transparent* leaves the values without highlights.

- 3 Sort the report by selecting a criterion from the *Sort by* drop-down list. Select *Ascending* or *Descending* order. For grouped reports, this controls the order within the group, because the whole report is sorted by the grouping column.


Optionally, add summary rows to the table by selecting from the check boxes in the *Summary Rows* area. PlateSpin Recon calculates and displays the *Count*, *Sum*, *Maximum*, *Average*, or *Minimum* for numeric columns.

Changing Report Font Size


- 1 Click **Font**  in the Document Viewer. A drop-down menu is displayed, showing the following options:
 - ♦ Small Font
 - ♦ Medium Font
 - ♦ Large Font

Click the desired font size. The changes are applied to the report displayed in the Document Viewer.

Adding or Modifying Groups

- 1 Click *Groups*  in the Document Viewer.

The Report Group Settings dialog box is displayed.




The **Report Group Settings** dialog box is shown. It has a title bar with standard window controls. The main area is divided into several sections:

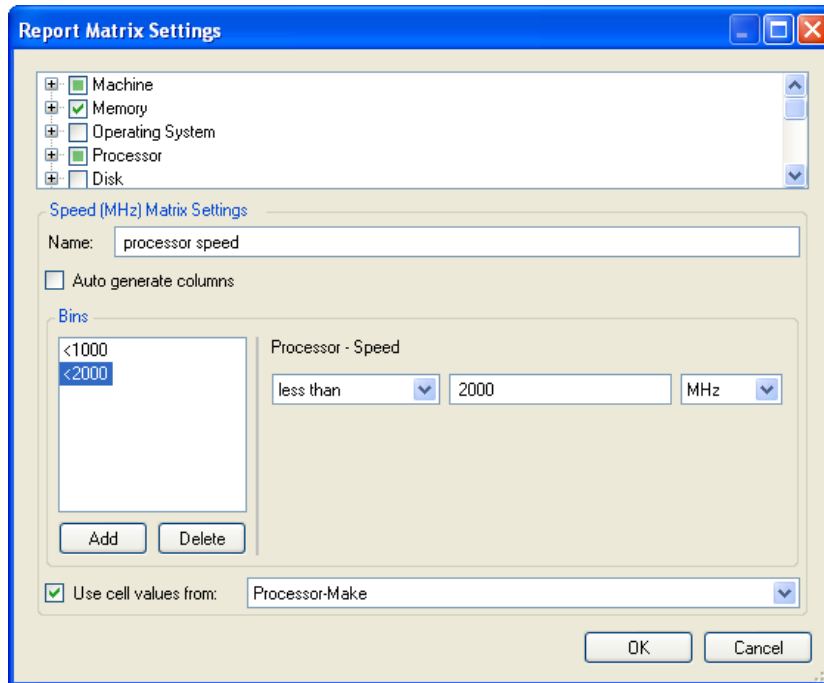
- MyGroup**: A list box containing the text "MyGroup". To its right are **Add** and **Delete** buttons.
- Custom MyGroup Bins**: A section with a dropdown menu currently showing "Memory - Size". Below it is a checkbox labeled "Automatically generate bins based on values".
- New Bin**: A list box on the left of this section. To its right is a dropdown menu showing "custom values", a text input field, and a unit dropdown menu showing "MB". Below these is the instruction: "Enter number ranges separated by commas. For example: 1-5, 8, 9-12".
- Sort**: A section with two radio buttons: "Ascending" (selected) and "Descending".
- Summary**: A section with four checkboxes: "Sum" (checked), "Maximum", "Average" (checked), and "Minimum".
- At the bottom right are **OK** and **Cancel** buttons.

- 2 Click *Add* next to the upper list box to enter a group name.
- 3 In the Custom <groupname> Bins area, select the element to group by from the drop-down list.
- 4 Do either of the following:
 - ♦ Select the option *Automatically generate bins based on values*.
 - ♦ Click *Add* in the lower area to create and name a custom bin. Define an expression for each custom bin.
- 5 Select the desired sorting order for the group.
- 6 Select the summary rows to include in the table. These summarize the values for each bin.
- 7 To modify a group or a bin, select its name and edit it or its associated settings. Alternatively, click the appropriate *Delete* button to delete the group or bin.
- 8 Click *OK*.

Adding or Modifying Matrixes

- 1 Click *Matrix*  in the Document Viewer.

The Report Matrix Settings dialog box is displayed.



- 2 In the upper list box, select elements for creating a matrix.
- 3 For each selected element, enter a *Name* for its matrix.
- 4 Do either of the following:
 - ♦ Check *Auto generate columns* to create a column for each unique value of the element.
 - ♦ Click *Add* to add custom columns. For each column, specify a name and define an expression.
- 5 Check *Use cell values from* to specify a value to place in the appropriate matrix cells. If this box is deselected, the values in the matrix columns are either True or False.
- 6 To remove a matrix from a report, deselect its check box. To modify a matrix, select it and modify its associated settings. Delete a bin by selecting its name and clicking *Delete*.


5.6.3 Viewing Reports

- ♦ [“Diagnostics Reports” on page 118](#)
- ♦ [“Multiple Report Pages” on page 118](#)



Diagnostics Reports

Click a row in a Missing Data - Summary report to view the corresponding Missing Data - Detailed and Error reports.

Multiple Report Pages

To view multiple pages at once, click *Multiple Pages* . Select the number of pages across and the number of pages vertically to view simultaneously. The more pages viewed, the smaller the text. The options available are from 1X1 through 5X10 pages.

5.6.4 Printing and Exporting Reports

Click *Print*  to open the Print dialog box, or click *Quick Print*  to immediately print the report, using the current default printer settings.


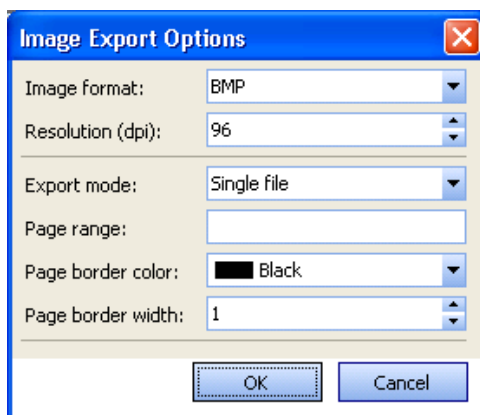
Click *Export Document*  to export data to any one of formats, such as PDF, HTML, Text, CSV, MHT, Excel, RTF, or an image format. Select the desired format from the drop-down list. Click *Export Document* to automatically export data in the format most recently selected in the session. If a format has not yet been selected in this session, the default format is PDF.

Image Formats

After selecting to export a document as an *Image*, the Image Export Options dialog box is displayed.

Figure 5-12 Image Exports Options Dialog Box



Use this dialog box to select from several image types to export, including `bmp`, `emf`, `wmf`, `gif`, `jpg`, `png`, and `tiff`. You can also specify the following options to control the actions output:

Resolution (dpi): Image resolution you want to use, keeping in mind that the higher the resolution, the bigger the exported file.

Export mode: Allows you to choose how you want to export the image:

- ♦ **single file:** One long image with a single set of column headers at the top.
- ♦ **single file page-by-page:** One long image made up of individual pages, each with a set of column headers at the top.
- ♦ **different files:** A set of separate files, each with a single page image in it. When this option is selected, numbers are added to the end of the filename you specify for each separate file, incrementing until all the pages are exported. For example, if you specified `MyImage.bmp` for a two-page report, the two files exported would be named `MyImage01.bmp` and `MyImage02.bmp`.

Page range: Range of pages in the report to export. For example, to export the first five pages and the tenth page of a 20 page report, specify `1-5,10`.


Page border color: The color to use for the report page border.

Page border width: The width of the page border.

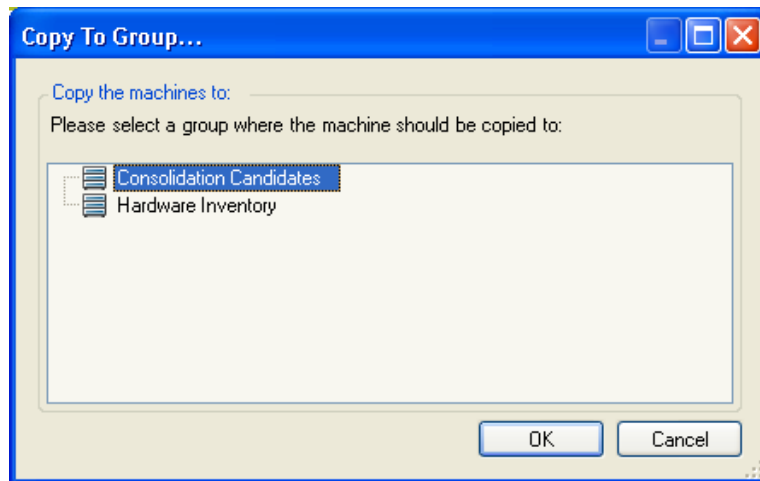
5.6.5 Copying Servers from a Report to a Group

Sometimes it is useful to copy the displayed servers from a report to a Data Center Explorer group. Not all reports display servers, but those types of reports do not display the Copy To Group tool. The groups must already exist. See [Section 3.4, “Working with Groups,” on page 52](#).

To copy servers to a group:

- 1 Click the check boxes next to the servers in the report you want to add to a group.
- 2 Click *Copy To Group* .

The Copy To Group dialog box is displayed.



- 3 Select the target group and click *OK*. To exit without copying, click *Cancel*.

Server Consolidation

6

This section provides information about the PlateSpin Recon Server Consolidation and Protection features.

- ♦ [Section 6.1, “About Server Consolidation,” on page 121](#)
- ♦ [Section 6.2, “Creating Projects,” on page 122](#)
- ♦ [Section 6.3, “Creating and Editing Server Templates,” on page 126](#)
- ♦ [Section 6.4, “Identifying Consolidation Candidates,” on page 134](#)
- ♦ [Section 6.5, “Working with Consolidation Scenarios,” on page 134](#)
- ♦ [Section 6.6, “Generating Inventory Reports,” on page 145](#)
- ♦ [Section 6.7, “Viewing Consolidation Reports,” on page 146](#)
- ♦ [Section 6.8, “Implementing Scenarios with Portability Suite,” on page 151](#)

6.1 About Server Consolidation

The PlateSpin Recon Consolidation Planning Module allows for exploration of scale-up and scale-out consolidation scenarios with candidate workloads. Scenarios can assist with determining the minimum number of target servers and balancing workloads across these servers.

The key benefits of the Consolidation Planning Module include:

- ♦ Maximizing server utilization by distributing the number of workloads.
- ♦ Minimizing resource contention by balancing software workloads across five dimensions: Processor, Memory, Disk, Network, and Time.
- ♦ Eliminating the need for extra servers with advanced workload sizing and optimization.
- ♦ Identifying servers to be protected in a disaster recovery plan.

The following are the key stages to assessing server consolidation:

- 1 Discover and list servers in the PlateSpin Recon Client. For more information, see [Section 3.2, “Discovering and Inventorying Data Center Resources,” on page 30](#).
- 2 Inventory servers and collect static data. For more information, see [Section 3.2, “Discovering and Inventorying Data Center Resources,” on page 30](#).
- 3 Monitor servers and collect dynamic data. For more information, see [Section 3.2.5, “Viewing Inventory Data and Monitoring Servers,” on page 47](#).
- 4 Create a consolidation project. For more information, see [Section 6.2, “Creating Projects,” on page 122](#).
- 5 Select servers for consolidation. For more information, see [Section 6.4, “Identifying Consolidation Candidates,” on page 134](#).
- 6 Identify the optimal consolidation scenario. For more information, see [Section 6.5, “Working with Consolidation Scenarios,” on page 134](#).

6.2 Creating Projects

A project can consist of any combination of consolidation candidates, protection candidates, workloads, and scenarios. Consolidation candidates are workloads on the network that you might want to move to virtual machines. Protection candidates are workloads you might want to create standby VM duplicates of for disaster recovery.

You create and compare scenarios in a project. The Project Report consists of scenario comparison results, which help in determining the best consolidation scenario for the environment.

Before you create a project, identify and group consolidation candidates. The identified candidates allow you to create and compare consolidation scenarios in a project.


NOTE: Collecting at least 24 hours of data for each workload before starting the project is recommended.

To create a new project:

1 Do any of the following:

- ♦ Click *Consolidation* , then click *New Project*.
- ♦ In the Project Explorer, click .

The Project Wizard (Step 1 of 5: Project Name and Type) is displayed.



2 Type a unique project name for the project.

3 Select one of the following project types:

- ♦ **Server Consolidation:** The project produces a recommendation for the best placement of workloads with respect to concurrent execution and CPU and memory headspace requirements.
- ♦ **Workload Protection:** The project produces a recommendation for deployment of workloads for disaster recovery. Specify existing virtual servers or use server templates to build a scenario.

- ♦ **Both:** The project produces a recommendation for both consolidation and protection. For example, optimally consolidate servers, and provision the remaining headroom for data protection.

A workload cannot be provisioned for both consolidation and protection in the same project.

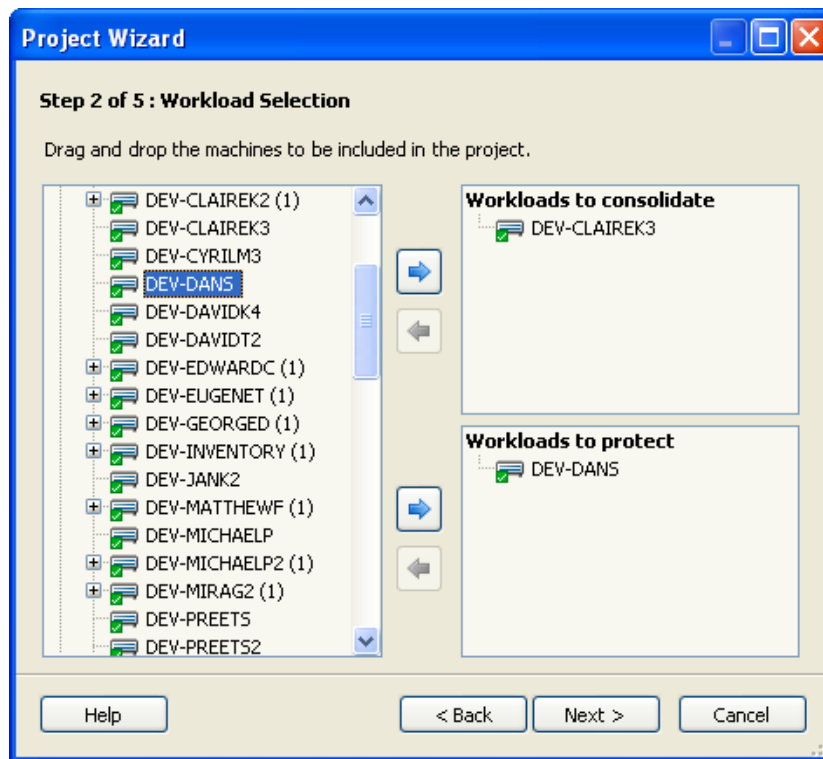
4 Click *Next*.

The Step 2 of 5: Workload Selection dialog box is displayed.

5 Identify the workloads you want to include in the project.

6 Identify the workloads you want to include in the project, and include them in the Workloads to Consolidate panel or the Workloads to Protect panel in any of the following ways:

- ♦ Drag workloads from the tree on the left and drop them into the Workloads to Consolidate panel or the Workloads to Protect panel on the right.
- ♦ Select servers in the right pane, then click the arrow pointing to the panel to which you want to add them.



If the existing [virtual hosts](#) have been inventoried by PlateSpin Recon, they can be added to the project as potential target servers. The existing servers take precedence over the server template that you provide when generating scenarios. Select any virtual hosts that exist in your environment to be used for consolidation or protection. Drop the target servers into either the Protection or Consolidation panel.

NOTE: PlateSpin Recon always considers a Solaris 10 machine as a target server for consolidation because the machine has the global zone enabled by default. Therefore, you cannot consider a Solaris 10 machine as a consolidation source.

To remove a server from the Workloads to Consolidate panel or the Workloads to Protect panel, select the server in the panel, then click the arrow pointing towards the server list.

Servers cannot be targeted for consolidation and protection within the same project.

7 Click *Next*.

The Step 3 of 5: Workload Utilization Date and Time Range dialog box is displayed.

The screenshot shows the 'Project Wizard' dialog box, specifically 'Step 3 of 5: Workload Utilization Date and Time Range'. The 'Date Range' section has three radio buttons: 'All Data' (selected), 'Last' (with a value of 24 and a unit dropdown set to 'Hour(s)'), and 'Between' (with two date and time dropdowns: '07/02/2008 02:07 PM' and '07/09/2008 02:07 PM'). The 'Business Days' section has checkboxes for all days of the week, all of which are checked. There is an 'Advanced...' button to the right. At the bottom are 'Help', '< Back', 'Next >', and 'Cancel' buttons.

8 Set the date and time range to summarize the workload data.

Date Range and *Business Days* are required for calculating the 24H profile of workloads.

9 (Optional) To manually set workload scaling factors, click *Advanced*.

The Workload Scaling dialog box is displayed.

10 Use the slide controls to adjust values or type the desired number.

The screenshot shows the 'Workload Scaling' dialog box. It contains a text box explaining that workload scaling can be used to adjust workload size to account for variability and projected growth, with a scale from peak hour average to two times the sum of peak hour average standard deviation. Below this are five rows of controls for 'Processor:', 'Memory:', 'Disk Space:', 'Disk IO:', and 'Network IO:'. Each row has a slider between 'Decrease' and 'Increase' labels, and a numeric input field. The values in the input fields are: Processor (1.00), Memory (1.00), Disk Space (1.25), Disk IO (1.00), and Network IO (1.00). To the right of each input field are labels: 'AVG + 0.00 STDEV' for Processor, Memory, and Network IO, and 'AVG' for Disk Space. At the bottom are 'OK' and 'Cancel' buttons.

Workloads can be scaled across five dimensions to account for workload variability and projected growth. Server resources allocated to workloads on target servers are based on scaled workloads. By scaling workloads upward, more server resources might be required, potentially resulting in lower consolidation ratios. When scaling pointers are moved, workload scaling factors change.

Workloads cannot be scaled below their peak hourly averages.

11 Click *OK*.

12 Click *Next*.

The Step 4 of 5: Forecast Options dialog box is displayed.

13 Select forecast options and a trend type.

Workloads can be forecasted to account for future growth.

The screenshot shows the 'Project Wizard' dialog box, specifically 'Step 4 of 5: Forecast Options'. The dialog has a blue title bar with standard window controls. The main area is divided into two sections. The top section, titled 'Periods', contains a numeric spinner set to '90', a dropdown menu set to 'Day (s)', and 'Add' and 'Remove' buttons. Below this is a list box containing '30 Day (s)', '60 Day (s)', and '90 Day (s)'. The bottom section, titled 'Trend / Regression:', features a dropdown menu currently showing 'Linear'. The dropdown list also includes 'Linear', 'Logarithmic', 'Exponential', and 'Power'. At the bottom of the dialog are four buttons: 'Help', '< Back', 'Next >', and 'Cancel'.


14 Click *Next*.

The Step 5 of 5: Power and Cooling dialog box is displayed.

15 Specify the values for the *Project Type*, or deselect the *Enable Power and Cooling* check box if these values are not to be included.

Energy Cost: Cost of energy per kilowatt hour.

Server Derating Factor: Percentage of total energy consumption by the average server. For example, a server rated at 800 watts uses 400 watts with a server derating factor of 50%.



16 Click *Create*.

The new project is displayed in the Project Explorer.




When creating a project, a 24H Profile of workloads is computed. This includes calculating Processor Speed (MHz), Memory (MB), Disk Space Used (GB), Disk IO (MB/s), and Network IO (MB/s) over a 24-hour period for each workload. The 24H Profile comprises 24 values of workloads that correspond to each hour of the day (0 to 23). Each value in the 24H Profile is calculated by taking the peak value from all values observed for that hour.

6.3 Creating and Editing Server Templates

Server templates are used to define servers for use in consolidation scenarios. You create server templates, each with their own components and utilization thresholds. A common use for server templates is for scale-up or scale-out consolidation scenarios. The created server templates are used as target servers for consolidating workloads.

- ♦ [Section 6.3.1, “Creating a New Server Template,” on page 126](#)
- ♦ [Section 6.3.2, “Editing an Existing Server Template,” on page 130](#)
- ♦ [Section 6.3.3, “Copying a Server Template,” on page 133](#)

6.3.1 Creating a New Server Template

- 1 Click *Consolidation*  , then click *Server Templates* .

The Server Template Manager dialog box is displayed.

- 2 Click .

The New Server Template dialog box is displayed.

New Server Template

General Components

Name: New Server

Make: HP

Model: ProLiant DL580

TCO: 11000 Rack Units: 3

Power: 1163 W

Virtualization: Microsoft Virtual Server

Help OK Cancel

3 On the General page, specify the following:

User Interface	Required	Description
Name	Yes	Name of the template to be created.
Make	No	Server manufacturer.
Model	No	Server model.
TCO	No	Total cost of ownership.
Rack Units	No	Number of rack units.
Power	No	Power consumption in watts.

User Interface	Required	Description
Virtualization	Yes	Virtualization platform: <ul style="list-style-type: none"> ♦ Citrix XenServer ♦ Microsoft Hyper-V ♦ Microsoft Virtual Server ♦ None ♦ Solaris 10 Zones ♦ Virtual Iron ♦ VMware ESX ♦ VMware ESX 3 ♦ VMware ESX 4 ♦ VMware GSX ♦ VMware Server ♦ SLES XenServer

- 4 Click the *Components* tab, then specify the following on the Components page:

The screenshot shows the 'New Server Template' dialog box with the 'Components' tab selected. The dialog is divided into two main sections: 'Settings' and 'Thresholds'.

Settings:

- Processor Count: 1
- Processor Speed: 3333 MHz
- Cores Per Processor: 1
- CPU Model: Core Xeon
- Hyper-Threading: ☐
- Normalized Speed: 41.10
- Memory Size: 32768 MB
- Disk Space: ☐ Fixed (0 GB), ☒ Unlimited

Thresholds:

- Processor: 80 %
- Memory: 90 %
- Disk Space: 85 %
- Disk IO: 50 MB/s
- Network IO: 10 MB/s

Buttons at the bottom: Help, OK, Cancel.

User Interface	Required	Description
Settings		
Processor Count	Yes	Total processor count.
Processor Speed	Yes	Instructions per second the server can execute in MHz or GHz. The minimum processor speed you can use is 266 MHz.
Memory	Yes	Memory size for server in MB or GB.
Cores Per Processor	Yes	Number of cores per processor.
CPU Model	Yes	Model of CPU: <ul style="list-style-type: none"> ♦ Pentium IV ♦ Core ♦ Xeon ♦ Core Xeon ♦ Athlon ♦ Opteron ♦ Intel Nehalem
NOTE: The CPU model for any processor in the Pentium series is shown as Pentium IV. This is based on the normalization speed that is calculated for a series of processors.		
Hyper-Threading	No	Specifies if hyper-threading is available.
Normalized Speed	Readonly	A unitless measure of processor capacity. This provides more realistic comparisons across diverse processor architectures.
Memory Size	Yes	The amount of memory to use. The minimum memory size you can use is 256 MB.
Disk Space	Yes	Disk space in GB or TB that is required for the server. Select <i>Unlimited</i> if you are connecting target servers to a SAN or NAS, where disk space is not a limiting factor. If you use <i>Fixed</i> , the minimum amount of disk space you can specify is 1 GB.
Thresholds		
Processor	No	Processor utilization. By default, the utilization is set to 80%.
Memory	No	Memory utilization. By default, the utilization is set to 90%.
Disk Space	No	Maximum disk utilization threshold. By default, the utilization is set to 85%.
Disk IO	No	Maximum disk I/O allowed. By default, the Max IO is set to 50 MB/s.

User Interface	Required	Description
Network IO	No	Maximum network IO. By default, the network IO is set to 10 MB/s.

Thresholds specify the level of utilization the Consolidation Planning Module can assume for the server templates. This affects the outcome of consolidation scenarios; lower utilization thresholds result in more servers being required to host the workloads of a scenario.




For example: workloads requiring 180 GB of memory can be consolidated onto a server with 200 GB of memory if the utilization threshold is set at 90%. However, if the memory utilization threshold is set at 50%, only 100 GB of memory can be used, resulting in fewer workloads being consolidated.

Utilization thresholds are initially set to the defaults listed in the table above and can be customized for specific needs.

- 5 Click *OK*.



The newly created server template is added to the Server Template list.

6.3.2 Editing an Existing Server Template

- 1 Click *Consolidation*  , then click *Server Templates* .

The Server Template Manager dialog box is displayed.

- 2 Do any of the following:

- ♦ Select the server template you want to edit, then click .
- ♦ Right-click the server template you want to edit, then click  *Edit*.

The Edit Server Template dialog box is displayed.

Edit Server Template

General Components

Name: My Old Server Template

Make: Computers Inc.

Model: A-999

TCO: 0 Rack Units: 0

Power: 0 W

Virtualization: None

Help OK Cancel

- 3 On the General page, make any desired modifications by using the following information:

User Interface	Required	Description
Name	Yes	Name of the template to be created.
Make	No	Server manufacturer.
Model	No	Server model.
TCO	No	Total cost of ownership.
Rack Units	No	Number of rack units.
Power	No	Power consumption in watts.

User Interface	Required	Description
Virtualization	Yes	Virtualization platform: <ul style="list-style-type: none"> ♦ Microsoft Virtual Server ♦ Virtual Iron ♦ VMware ESX ♦ VMware ESX 3 ♦ VMware ESX 4 ♦ VMware GSX ♦ VMware Server ♦ Citrix XenServer ♦ Microsoft Hyper-V ♦ SLES XenServer ♦ None

- 4 Click the *Components* tab, then make any desired modifications by using the following information:

User Interface	Required	Description
Settings		
Processor Count	Yes	Total processor count.
Processor Speed	Yes	Instructions per second the server can execute in MHz or GHz. The minimum processor speed you can use is 266 MHz.
Memory	Yes	Memory size for server in MB or GB.
Cores Per Processor	Yes	Number of cores per processor.
CPU Model	Yes	Model of CPU: <ul style="list-style-type: none"> ♦ Pentium IV ♦ Core ♦ Xeon ♦ Core Xeon ♦ Athlon ♦ Opteron
Hyper-Threading	No	Specifies if hyper-threading is available.
Normalized Speed	Readonly	A unitless measure of processor capacity. This provides more realistic comparisons across diverse processor architectures.
Memory Size	Yes	The amount of memory to use. The minimum memory size you can use is 256 MB.

User Interface	Required	Description
Disk Space	Yes	Disk space in GB or TB that is required for the server. Select <i>Unlimited</i> if you are connecting target servers to a SAN or NAS, where disk space is not a limiting factor. If you use <i>Fixed</i> , the minimum amount of disk space you can specify is 1 GB.
Thresholds		
Processor	No	Processor utilization. By default, the utilization is set to 80%.
Memory	No	Memory utilization. By default, the utilization is set to 90%.
Disk Space	No	Maximum disk utilization threshold. By default, the utilization is set to 85%.
Disk IO	No	Maximum disk I/O allowed. By default, the Max IO is set to 50 MB/s.
Network IO	No	Maximum network IO. By default, the network IO is set to 10 MB/s.


Thresholds specify the level of utilization the Consolidation Planning Module can assume for the server templates. This affects the outcome of consolidation scenarios; lower utilization thresholds result in more servers being required to host the workloads of a scenario.

For example: workloads requiring 180 GB of memory can be consolidated onto a server with 200 GB of memory if the utilization threshold is set at 90%. However, if the memory utilization threshold is set at 50%, only 100 GB of memory can be used, resulting in fewer workloads being consolidated.



Utilization thresholds are initially set to the defaults listed in the table above and can be customized for specific needs.

- 5 Click *OK*.

6.3.3 Copying a Server Template

- 1 Click *Consolidation*  , then click *Server Templates* .

The Server Template Manager dialog box is displayed.

- 2 Do any of the following:
 - ♦ Select the server template you want to copy, then click .
 - ♦ Right-click the server template you want to copy, then click  *Copy*.

The New Server Template dialog box is displayed with prepopulated template information.

- 3 Rename the template and make any other required modifications.
- 4 Click *OK*.

The template is copied and is displayed in the Server Template Manager listing.

6.4 Identifying Consolidation Candidates

This activity involves selecting servers for consolidation. This requires you to generate inventory and analysis reports so you can find a set of servers that meet your consolidation requirements, then you can select your consolidation candidates from that set of servers. For more information, see [Section 5.5, “Generating Reports,” on page 110](#).

6.5 Working with Consolidation Scenarios

Scenarios allow the quick creation and evaluation of one or more consolidation alternatives. Generating multiple scenarios for comparison is recommended. This helps to determine the optimal consolidation scenario.


There are two types of scenarios:

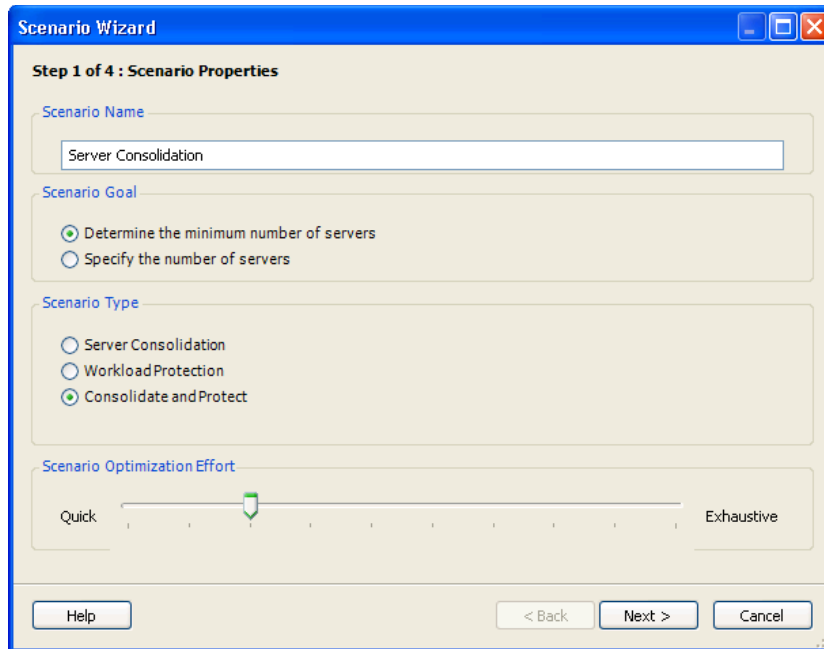
- ♦ **Determine the minimum number of servers:** PlateSpin Recon determines how many of a specific type of target server are required to balance the workloads. With this type of scenario, only one server template can be used.
- ♦ **Specify the number of servers:** You set the number and type of target servers for consolidation. PlateSpin Recon calculates the optimal distribution of workloads for the chosen targets. With this type of scenario, multiple server templates can be used.

Any inventoried machine running any virtualization platform supported by PlateSpin Recon is considered a target server.

- ♦ [Section 6.5.1, “Determining the Minimum Number of Servers,” on page 134](#)
- ♦ [Section 6.5.2, “Specifying the Number of Servers,” on page 139](#)
- ♦ [Section 6.5.3, “Results of Consolidation Scenarios,” on page 143](#)
- ♦ [Section 6.5.4, “Editing Scenarios,” on page 144](#)
- ♦ [Section 6.5.5, “Deleting Scenarios,” on page 145](#)

6.5.1 Determining the Minimum Number of Servers

- 1 In the Project Explorer tab, select a project, then click *New Scenario* .
The Scenario Wizard (Step 1 of 4: Scenario Properties) is displayed.



The available scenario types are determined by the *Project Type* specified while creating the project.

2 Specify a unique name for the scenario.

3 In the Scenario Goal pane, select *Determine the minimum number of servers*.

This goal determines the minimum number of a specific type of target servers to balance workloads. You can use only one server template with this option.

4 Select a scenario type.

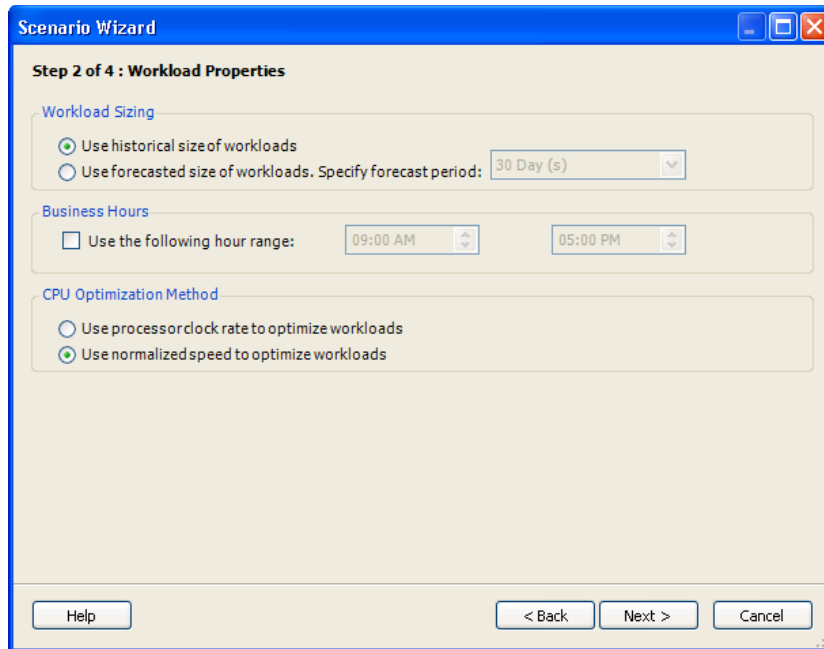
This step is relevant only if the project upon which the scenario is based has the goal configured as *Both* (server consolidation and workload protection).

5 Adjust the *Scenario Optimization Effort* slider.

Quick is faster but not as intensive, and *Exhaustive* takes longer but is more thorough.

6 Click *Next*.

The Step 2 of 4: Workload Properties dialog box is displayed.



7 Specify Workload Sizing.

- ♦ **Use historical size of workloads:** Creates the scenario based on the historical size of the workloads.
- ♦ **Use forecasted size of workloads (specify forecast period):** Creates the scenario based on the forecasted size of the workloads after the specified forecast period.

For scenarios based on protection only projects, these values are dimmed.

8 To average workload and utilization expressions over a specified time range, select the *Use the following hour range* check box and set the start and end times for the business hours in your environment.

Hours that you specify here are used for all reports created for this scenario or for any node under this scenario. Report templates do not have any direct control over business/peak hours. Consolidation and Planning charts use both the specified business hours and all 24-hour data points.

9 Specify the CPU optimization method you want to use.

- ♦ **Use processor clock rate to optimize workloads:** Optimizes the CPU based on the actual processor clock.
- ♦ **Use normalized speed to optimize workloads:** Optimizes the CPU based on a normalized value. See [Section B.6, “CPU Normalization,” on page 201](#) for more information.

Consolidation charts based on this scenario differ depending on the method you choose. If you select normalized speed, the processor utilization and capacity values in consolidation charts created based on this scenario will also use the normalized values.

In any reports, Normalized Speed columns are not displayed by default, even if it's a scenario summary or workload assignment report and the optimization method selected is *Use normalized speed to optimize workloads*.

For scenarios based on protection only projects, these values are dimmed.

10 Click *Next*.

The Step 3 of 4: Workload Selection dialog box is displayed.

Scenario Wizard

Step 3 of 4 : Workload Selection

Select the workloads to be used in this scenario

Consolidate Protect

	Name	Processor (Normalized)	Memory (MB)	Disk Space (GB)	Disk IO (MB/s)	Network IO (MB/s)
<input checked="" type="checkbox"/>	DEV-HERBB	0.17	676.73	75.22	0.01	0.12
<input checked="" type="checkbox"/>	DEV-MARTINS	3.03	1430.31	36.37	0.96	0

Help < Back Next > Cancel

11 Select workloads or servers, depending on the scenario type.

Scenario Wizard

Step 3 of 4 : Workload Selection

Select the workloads to be used in this scenario

Consolidate Protect

	Name	Processor (MHz)	Memory (MB)	Disk Space (GB)	Disk IO (MB/s)	Network IO (MB/s)
<input checked="" type="checkbox"/>	IT-LAPTEMP	15.06	357	30.36	0.01	0.08

Help < Back Next > Cancel

- 12** (Conditional) If the *Scenario Type* includes server consolidation, select the servers to be consolidated in this scenario.

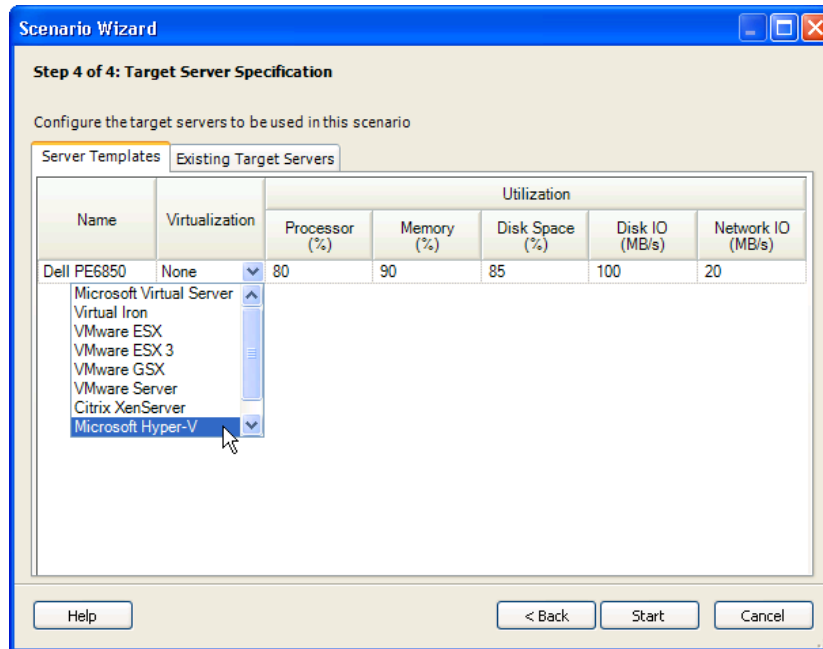
By default, all servers specified in the project for consolidation are selected.

- 13** (Conditional) If the *Scenario Type* includes workload protection, select workloads on the *Protect* tab.

By default, all workloads specified in the project for protection are selected.

- 14** Click *Next*.

The Step 4 of 4: Target Server Specification dialog box is displayed.



- 15** In the *Name* drop-down list, select the server template you want to use.

or

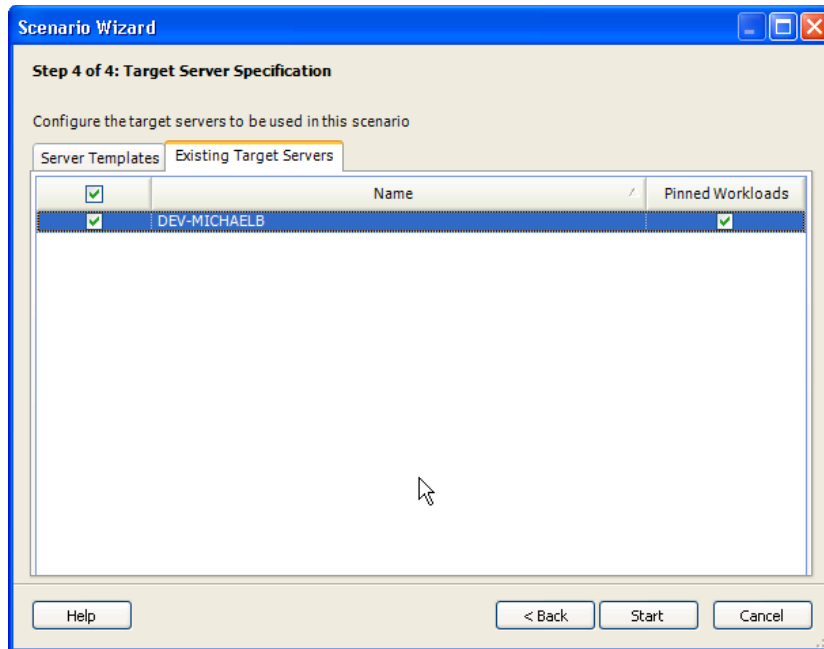
Create a new server template by selecting *New Server*. (You can use only one server template.)

You must specify at least one server template, even if you want to consolidate only to existing target servers.

- 16** (Optional) In the *Virtualization* drop-down list, select the Virtualization platform to use.

- 17** (Optional) Click the *Existing Target Servers* tab to use servers included in the project.

If there are no existing target servers specified in the project, this option is unavailable.



18 Select the existing target servers to use in the scenario, then select the *Pinned Workloads* check box to force workloads currently on a particular target server to remain on that server after consolidation.

19 Click *Start*.

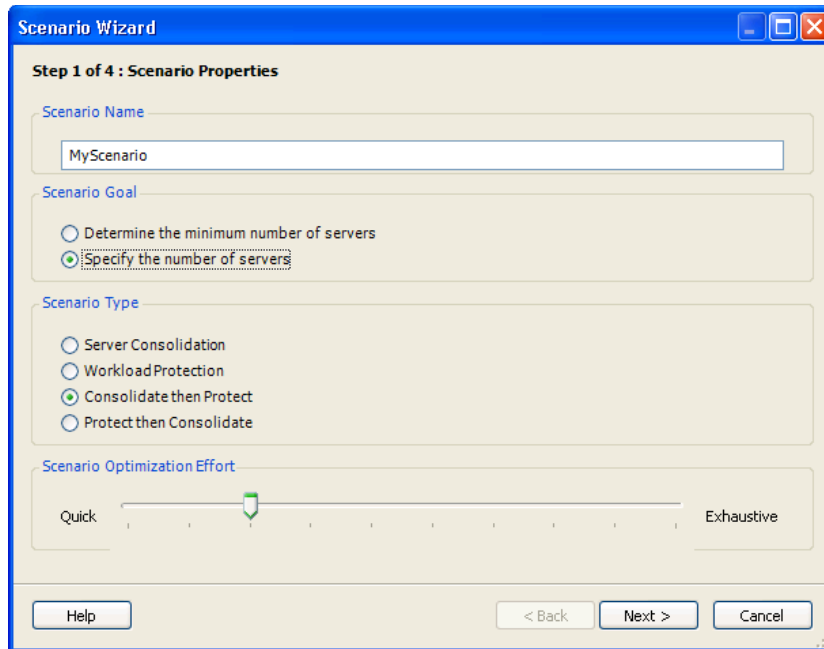
The Create Scenario job runs.

You can watch the progress in the Job Explorer window. When the scenario is created, it is displayed under Scenarios in the base project.


6.5.2 Specifying the Number of Servers

1 Click *New Scenario* .

The Scenario Wizard (Step 1 of 4: Scenario Properties) is displayed.



- 2 Type a unique scenario name.
- 3 In the Scenario Goal pane, select *Specify the number of servers* to set the number and type of target servers for consolidation.
You can use multiple server templates with this option.
- 4 Select a scenario type.
If some workloads must be left unconsolidated because the number of templates specified is inadequate, the algorithm fits as many of the consolidation workloads onto those templates as it can if you choose *Consolidate then Protect*. The reverse is true if you choose *Protect then Consolidate*.
- 5 Adjust the *Scenario Optimization Effort* slider.
Quick is faster but not as intensive, and *Exhaustive* takes longer but is more thorough.
- 6 Click *Next*.
The Step 2 of 4: Workload Properties dialog box is displayed.



7 Specify Workload Sizing.

- ♦ **Use historical size of workloads:** Creates the scenario based on the historical size of the workloads.
- ♦ **Use forecasted size of workloads (specify forecast period):** Creates the scenario based on the forecasted size of the workloads after the specified forecast period.

For scenarios based on protection only projects, these values are dimmed.

8 To average workload and utilization expressions over a specified time range, select the *Use the following hour range* check box and set the start and end times for the business hours in your environment.

Hours that you specify here are used for all reports created for this scenario or for any node under this scenario. Report templates do not have any direct control over business/peak hours. Consolidation and Planning charts use both the specified business hours and all 24-hour data points.

9 Specify the CPU optimization method you want to use.

9a Use processor clock rate to optimize workloads: Optimizes the CPU based on the actual processor clock. For more information, see [Section B.6, “CPU Normalization,” on page 201](#).

9b Use normalized speed to optimize workloads: Optimizes the CPU based on a normalized value.

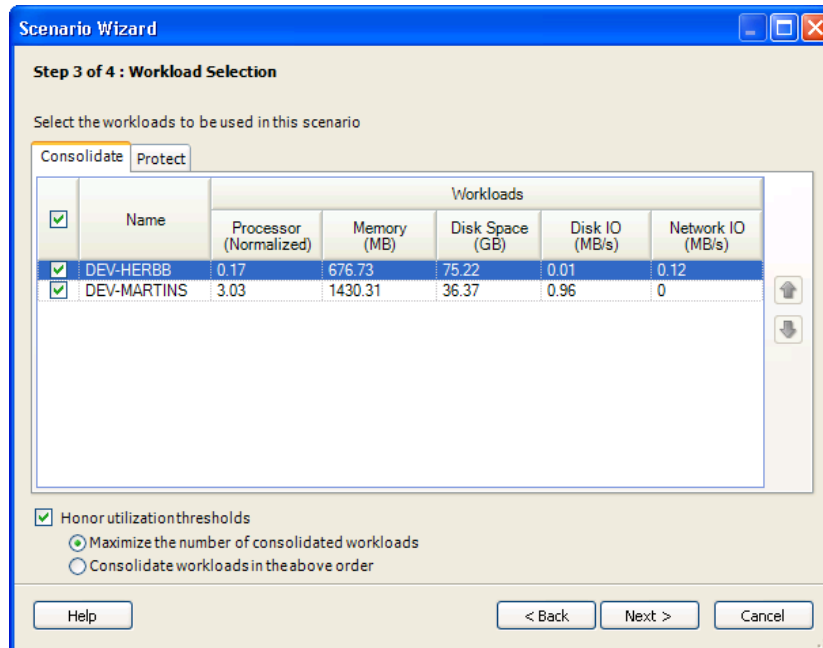
Consolidation charts based on this scenario differ depending on the method you choose. If you select normalized speed, the processor utilization and capacity values in consolidation charts created based on this scenario will also use the normalized values.

In any reports, Normalized Speed related columns are not displayed by default, even if it's a scenario summary/workload assignment report and the optimization method selected is *Use normalized speed to optimize workloads*.

For scenarios based on protection only projects, these values are dimmed.

- 10 Click *Next*.

The Step 3 of 4: Workload Selection dialog box is displayed.



- 11 (Conditional) If the *Scenario Type* includes server consolidation, select the servers to be consolidated in this scenario.

By default, all servers specified in the project for consolidation are selected

- 12 (Conditional) If the *Scenario Type* includes workload protection, select workloads on the *Protect* tab.

By default, all workloads specified in the project for protection are selected.

If you are running a scenario with a set number of target servers, in some cases the set number of target servers do not have sufficient resources to accommodate all workloads.

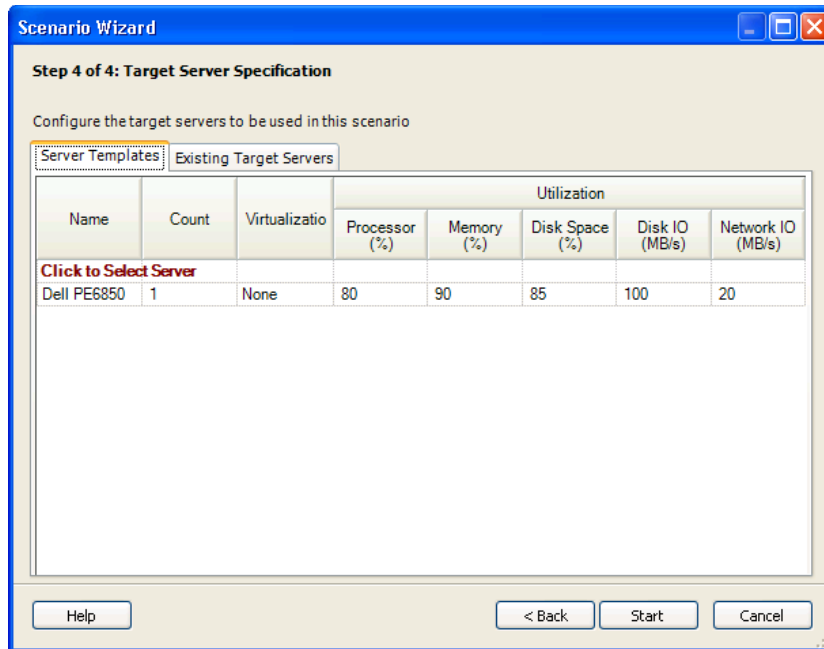
By default, *Honor utilization thresholds* is selected. This maintains the total utilization of a target server below the set thresholds. The sum of workloads do not violate the utilization thresholds for a specified server.

- 13 Define consolidation priorities by selecting one of the following options:

- ♦ **Maximize the number of consolidated workloads:** Workloads are selected for consolidation with the goal of consolidating the largest possible number of workloads. The order shown is not relevant and the up and down arrows are disabled.
- ♦ **Consolidate workloads in the above order:** Prioritize workloads for consolidation only, in the given order. This ensures that certain workloads are consolidated. To move workloads into the desired order of consolidation priority, use the up and down arrows.

- 14 Click *Next*.

The Step 4 of 4: Target Server Specification dialog box is displayed.

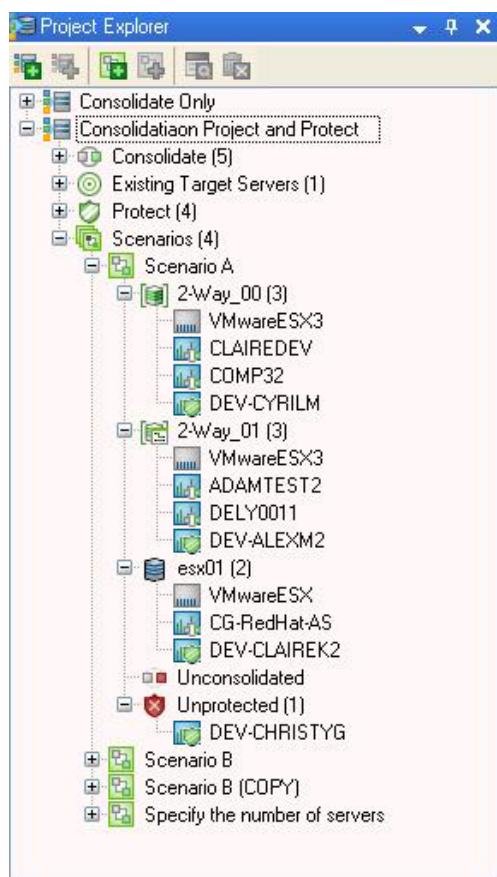


- 15 In the *Name* drop-down list, select the Server template you want to use.
or
Create a new server template by selecting *Click to Select Server*.
- 16 Specify the count for each server template.
- 17 (Conditional) If you want to use the existing server targets specified in the project, click the *Existing Target Servers* tab and select the servers.
This option is available only if a project has existing target servers.
- 18 Click *Start*.
The Create Scenario job runs.
You can watch the progress in the Job Explorer window. When the scenario is created, it is displayed under Scenarios in the base project.

6.5.3 Results of Consolidation Scenarios

The result of any scenario depends on the workloads used and the server templates selected. Higher consolidation ratios are achieved with larger target servers. The created scenario is displayed in the Project Explorer. For a description of the icons in the graphic below, see the listing of Project Explorer icons in [Table D-6 on page 214](#).

Figure 6-1 Project Explorer



Virtualization Overhead: If server templates configured with a virtualization platform are selected as input to the scenario, the overhead related to the selected virtualization platform is shown as an additional workload under each target server and is named after the virtualization platform. Virtualization overhead workloads are not movable.


Unconsolidated Workloads: Lists all servers that were not consolidated during the scenario creation process. Workloads are rejected if they cannot be added to a server without exceeding utilization thresholds. These servers are listed under *Unconsolidated* in the scenario. Moving unconsolidated workloads to a server can result in a violation of utilization thresholds.

Unprotected Workloads: Lists all servers that were not protected during the scenario creation. Workloads are rejected if they cannot be added to a server without exceeding utilization thresholds. These servers are listed under *Unprotected* in the scenario. Moving unprotected workloads to a server can result in a violation of utilization thresholds.

6.5.4 Editing Scenarios

You can move workloads between target servers or from *Unconsolidated* or *Unprotected* to a target server by dragging and dropping in the Project Explorer. Alternatively, right-click the workload and click *Unconsolidate Workload* on the menu.

6.5.5 Deleting Scenarios

- 1 In the Project Explorer, select a scenario.
- 2 Click .

6.6 Generating Inventory Reports

- [Section 6.6.1, “Finding Servers That Meet the Consolidation Criteria,” on page 145](#)
- [Section 6.6.2, “Selecting Candidates for Consolidation,” on page 145](#)

6.6.1 Finding Servers That Meet the Consolidation Criteria

- 1 In the Report Explorer, right-click the *Hardware Inventory Report* template under *Inventory*, then select *Edit Report Template*.
The Edit Report Template dialog box is displayed.
- 2 Click the *Group* tab, select a filter or group to use when you generate reports based on this template.
- 3 Click the *Views* tab and select a view. Use the buttons adjacent to the list to add or edit a view. See [Section 5.4, “View Customization,” on page 106](#)
- 4 Click the *Field Filter* tab. Use the drop-down lists to specify query parameters. See [“Field Filter Tab” on page 98](#).
- 5 Click *OK*.
- 6 In the Report Explorer, right-click *Hardware Inventory* and select *Generate Report*.
- 7 Schedule the report for *Now*, and accept the default *Delivery Type* of the *Report Explorer*.
- 8 Click *Generate*.

The report instance, shown as the date and time the report was generated, is displayed under the report template in the Report Explorer and the report is previewed in the Document Viewer, showing the servers that meet your consolidation requirements.

6.6.2 Selecting Candidates for Consolidation

- 1 In the Report Explorer, right-click the *Analysis* report template for *Hardware Utilization - Peak* and select *Edit Report Template* on the menu.
The Edit Report Template dialog box is displayed.
- 2 On the *Group* tab, make the same selection you made in [Step 1 on page 145](#).
- 3 On the *Views* tab, select the table view and click *Edit*.
The Create View dialog box is displayed.
- 4 On the *General* tab, select the *Summary Rows* (applied to columns containing calculated data) you want to create in the report.
- 5 Click *Finish*.
The Edit Report Template dialog box is displayed again.
- 6 On the *Date Filter* tab, set the date and time range to view utilization expressions for all servers that were monitored.

- 7 On the *Field Filter* tab, use the drop-down lists to specify the same query parameters you selected in [Step 1 on page 145](#).
- 8 On the *Summaries* tab, select *AVG*.
- 9 Click the *Calculation type* drop-down list and a type:
 - ♦ **Average:** Average of all hourly values across business days and business hours.
 - ♦ **Peak Hour:** A calculation based on the highest hourly value.
 - ♦ **Total:** Sums all the data points in the time frame. For example, Disk IOs per second for a 1-hour period equates to 10 disk IOs per second X 60 seconds X 60 minutes = 36000. Select the *Summary Rows* to create in the report of the columns containing calculated data.
- 10 Click *OK*.
- 11 In the Report Explorer, under Analysis, right-click the previously edited Hardware Utilization - Peak report template and click *Generate Report* on the menu. Accept the default *Delivery Type* of the *Report Explorer*.
- 12 Click *Generate*.

The report instance, shown as the date and time the report was generated, is displayed under the report template in the Report Explorer and the report is previewed in the Document Viewer.

Comparing and contrasting the two reports generated in [Step 1 on page 145](#) and [Step 1 on page 145](#) will help you determine your consolidation needs.

To help with this, while viewing a generated report, click a cell to display a customizable chart for that server/workload metric. For more information see [Section 5.1, “Chart Viewer,” on page 85](#).

6.7 Viewing Consolidation Reports

Consolidation reports are used for analysis. Unlike other PlateSpin Recon reports, which are generated through the Report Explorer, consolidation reports are generated through the Project Explorer. Like other PlateSpin Recon reports, generated consolidation reports are viewed in the Document Viewer. In addition, consolidation reports cannot be saved and must be generated whenever necessary.

The following types of consolidation reports are available:

- ♦ Project Report
- ♦ Scenario Summary Report
- ♦ Workload Assignment Report
- ♦ Unconsolidated Report
- ♦ Unprotected Report

Project, scenario summary, and workload assignment reports are template based. You can edit the templates through the Report Explorer. For information about editing consolidation report templates, see [Section 5.3.3, “Consolidation Report Templates,” on page 103](#).

For detailed information about the different types of consolidation reports, review the following sections:

- ♦ [Section 6.7.1, “Project Report,” on page 147](#)
- ♦ [Section 6.7.2, “Scenario Summary Report,” on page 147](#)

- ♦ Section 6.7.3, “Workload Assignment Report,” on page 149
- ♦ Section 6.7.4, “Unconsolidated Workload Report,” on page 150
- ♦ Section 6.7.5, “Unprotected Workload Report,” on page 150

6.7.1 Project Report

A Project Report provides a comparison of consolidation scenarios.

To view this report, double-click a project. The Project Report is displayed in the Document Viewer.

Figure 6-2 Sample Project Report

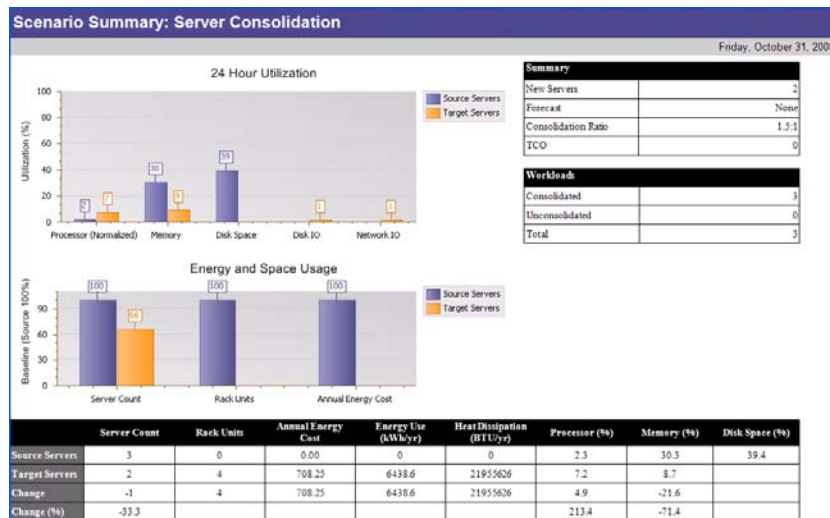
Project: Consolidation and Protection											
Friday, October 31, 2008											
Scenario	Utilization (%)			Workloads			New Servers				
Name	Processor (Spent)	Processor (Normalized)	Memory	Consolidated	Protected	Count	Consolidation Ratio	Protection Ratio	TCD	Rack Units	Power (W)
Before Consolidation	2.41	2.30	30.32	0	0	—	—	—	—	—	—
Server Consolidation	11.11	7.21	8.67	2	1	2	1.00:1.1	0.6:1	0	4	1470
Summary Statistics											
Max	11.10	7.20	8.70	2	1	2	1.0:1	0.6:1	0	4	1470
Avg	11.10	7.20	8.70	2	1	2	1.0:1	0.6:1	0	4	1470
Min	11.10	7.20	8.70	2	1	2	1.0:1	0.6:1	0	4	1470

6.7.2 Scenario Summary Report

The Scenario Summary Report helps in assessing the value of using a particular infrastructure for server consolidation.

Double-click the selected scenario to view the Scenario Summary report in the Document Viewer and the 24H Profile Stacked Utilization chart in the Chart Viewer.

Figure 6-3 Sample Scenario Summary Report



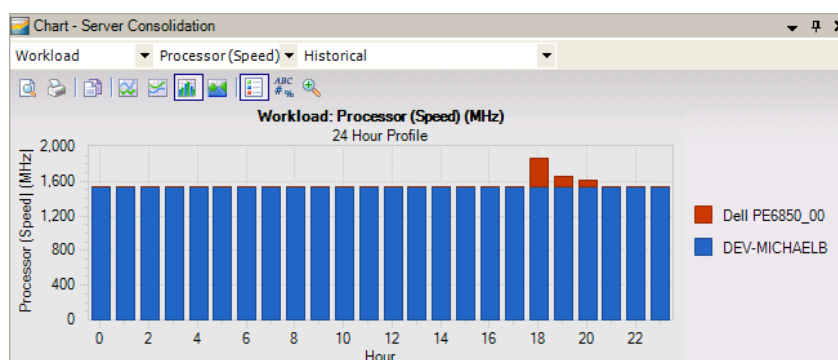
If there are multiple target servers, a summary statistics table is displayed with SUM, AVG, MAX, and MIN values.

Figure 6-4 Sample Report Summary Statistics Table

Machine		Processor	Memory	Disk		Network
Name	Workloads	% MHz used (%)	Used (%)	Used (%)	IO (%)	IO (%)
DEV-MICHAELB	0	24.0	27.0	--	0.0	0.0
Dell PE6850_00	3	3.2	7.5	--	1.0	1.0
Summary Statistics						
Sum	3					
Max	3	24.0	27.0	--	1.0	1.0
Avg	2	13.6	17.3	--	0.5	0.5
Min	0	3.2	7.5	--	0.0	0.0

If you click one of the server cells, the corresponding 24H Profile stack chart with all workloads assigned to the target server is displayed.

Figure 6-5 Sample 24H Profile Stack Chart



24H Profile Chart Appearance

Twenty-four hour (24H) Profile charts in the Data Center Explorer and in Consolidation Planning are slightly different. In the Data Center Explorer, the 24H profile that is shown. select several monitored machines, click *Create Chart* and then select *Profile* as the type of chart) in workload profile charts is for each individual resource instance separately and shows a different color for each instance and resource. These values are stacked for each given hour. This gives the maximum value for every instance for a specific hour over the monitoring period. These stacked values are then summed.


Figure 6-6 Summed Stack Values



This highlights the hour with the maximum utilization for a specific instance of the resource. This approach is particularly useful for determining maximum utilization of any individual resource instance.

In Consolidation Planning, the 24H chart has the values of the instances and resources summed for each time moment and then the maximum of this sum for every hour is identified.

Figure 6-7 Consolidation Planning 24H Chart



This is needed to determine the summary resource consumption of a workload.

6.7.3 Workload Assignment Report

In the Workload Assignment report, you can see which workloads are assigned to servers and how workloads are utilized in each of these assigned servers. You can modify workload assignments according to requirements and consider different kinds of server options while creating a project.

Workload Assignment reports are available for individual target servers or for all target servers included in a scenario.

- ♦ [“Individual Server” on page 149](#)
- ♦ [“All Servers” on page 150](#)

Individual Server

To view the Workload Assignment report and 24H Profile chart, double-click a target server in the *Project Explorer*.

NOTE: A 24H Profile cannot be created for workloads when there is no one-hour summary data.

If you click one of the servers in the Assignment report, the corresponding 24H Profile stack chart with all workloads assigned to the target server is displayed. The utilization values displayed are for the target servers.

Figure 6-8 Assignment Report Generated 24H Profile Stack Chart



All Servers

To view the Workload Assignment report for all the servers in the current scenario:

- 1 Right-click a scenario, select *View Report > View Workload Assignment Report* from the drop-down menu

6.7.4 Unconsolidated Workload Report

In the Unconsolidated Workload Report, you can see which of the consolidation candidate workloads are not included in the consolidation plan, as well as an explanation of the reason why PlateSpin Recon could not include them.

To view this report, right-click a scenario and select *View Report > View Unconsolidated Report*. The Unconsolidated Report is displayed in the Document Viewer.

Figure 6-9 Sample Unconsolidated Report

Unconsolidated Report		
B384 Project > B384 Scenario2 > Unconsolidated		April 6, 2007
Machine		
<input type="checkbox"/>	Machine Name	Reason
<input type="checkbox"/>	DEV-DAVIDT	There are not enough resources in any of the available containers
<input type="checkbox"/>	DEV-NADERS	There are not enough resources in any of the available containers
<input type="checkbox"/>	DEV-BORISB	There are not enough resources in any of the available containers
<input type="checkbox"/>	DEV-DAVIDK	There are not enough resources in any of the available containers

6.7.5 Unprotected Workload Report

In the Unprotected Report, you can see which of the consolidation candidate workloads were not included in the protection plan, as well as an explanation of the reason why PlateSpin Recon could not include them.

To view this report, right-click a scenario and select *View Report > View Unprotected Report*. The Unconsolidated Report is displayed in the Document Viewer.

Figure 6-10 Sample Unprotected Report

Unprotected Report		
Project 2 Protect > Scenario A > Unprotected		Monday, April 09, 2007
Machine		
<input type="checkbox"/>	Machine Name	Reason
<input type="checkbox"/>	DEV-DSBUILD	There are not enough resources in any of the available containers
<input type="checkbox"/>	FOGBUGZ	There are not enough resources in any of the available containers
<input type="checkbox"/>	VM-SRC3	There are not enough resources in any of the available containers
<input type="checkbox"/>	VM-SRC6	There are not enough resources in any of the available containers
<input type="checkbox"/>	YONINAVM	There are not enough resources in any of the available containers

6.8 Implementing Scenarios with Portability Suite

When there are identified scenarios and servers to deploy, you can implement the consolidation and protection by using PlateSpin Portability Suite. You can implement an entire scenario, separate containers within the scenario, or individual servers in the scenario.

PlateSpin Recon 3.7 supports PlateSpin Portability Suite 8.0.0. PlateSpin Recon 3.7.2 additionally supports PlateSpin Portability Suite 8.1.2 and 8.1.3.

To implement a scenario:

- 1 In the Project Explorer, right-click the scenario, then select *Implement using Portability Suite*.
The Implementing scenario <name> using Portability Suite Step 1 of 2 dialog box is displayed.

The screenshot shows a Windows-style dialog box titled "Implementing scenario MyScenario (2) using PortabilitySuite". It is labeled "Step 1 of 2: PortabilitySuite Settings". The dialog contains two main sections: "PortabilitySuite Settings" and "Destination".


PortabilitySuite Settings: This section has a label "Specify the PortabilitySuite Server you will be using for the consolidation." Below it are three fields: "PortabilitySuite Server URL" with the value "http://localhost/PortabilitySuite", "Network Name" with the value "Default", and "PortabilitySuite Credential" with a dropdown menu showing "gsx01\administrator".

Destination: This section has a label "Specify the folder relative to the server where the PortabilitySuite jobs should be saved. Note that the folder should be accessible by the account assigned to the PlateSpin Recon Service." Below it is a "Directory" field with the value "c:\temp".

At the bottom of the dialog are four buttons: "Help", "< Back", "Next >", and "Cancel".

- 2 Confirm the *Portability Suite Server URL* and *Network name*.
- 3 Use the drop-down list to select or create the credentials that Portability Suite uses to discover and convert the servers within the scenario or container.
- 4 Specify the directory where the Portability Suite jobs must be saved.
- 5 Click *Next*.

The Step 2 of 2 dialog box is displayed.



6 Click *Not Mapped*.

A table drops down where a match column is displayed.

If you are attempting to consolidate multiple containers, the mapping of container to target server is one-to-one. Unavailable containers that are already mapped to a server are dimmed. When they are either unmapped or the scenario they are in is deleted, they are once again made available.

7 Click *Start*.

When the job is initiated, PlateSpin Recon:

- ♦ Authenticates the servers.
- ♦ Creates and runs server discoveries for Portability Suite.
- ♦ Creates and saves conversion jobs in the file location specified in [Step 1 on page 151](#). The default location is `c:\temp`.

When the job completes, you can copy the saved conversion jobs to the Portability Suite machine and load it by using PlateSpin Portability Suite Client window (*File > Load Job*).

6.8.1 Implementing Phantom or Existing Target Servers

The steps to implement at the “phantom” (server template), or existing target server level, are identical to implementing a scenario: right-click the phantom or existing target virtual machine server and select *Implement Using Portability Suite*.

If a phantom server represents an existing target virtual machine server, you can map it to that existing server. When the job is created through Portability Suite, the source and target servers are auto-configured.

To map a phantom to a target virtual machine server and deploy it:

1 Right-click the phantom and select *Edit Mapping*.

The Server Mapping dialog box is displayed, listing the available servers that contain virtual machines. Data is listed for each machine, and the match column indicates whether or not the machine is a match to the phantom.

Test	Results
CPU	<p>Failure: The processing power of the destination server is less than the processing power of the source server.</p> <p>Warning: The destination server has more processing power than is necessary. A warning is displayed if the processing power of the destination server is 500 Mhz greater than the processing power of the source server.</p> <p>Pass: The CPU match is acceptable.</p>
Disk	<p>Failure: The disk space on the destination server is less than the disk space on the source server.</p> <p>Warning: The disk space on the destination server is higher than necessary. A warning is displayed if the destination server has more than 50 GB of surplus space compared to the source server.</p> <p>Pass: The disk match is acceptable.</p>
Memory	<p>Failure: The destination server has less memory than the source server.</p> <p>Warning: The destination server has more memory than is necessary. A warning is displayed if the memory on the destination server is 512 MB greater than the memory on the source server.</p> <p>Pass: The memory match is acceptable.</p>

- 2 In the Server Mapping dialog box, select the target virtual machine server to map the phantom to, then click *OK*.
- 3 Right-click the newly mapped phantom and select *Implement Using Portability Suite*.
- 4 Complete the Implementing Server Using Portability Suite Wizard as described in [“To implement a scenario:” on page 151](#).
- 5 Click *Start*.

This section includes information on using the PlateSpin Recon flexible Chargeback feature.

- ♦ [Section 7.1, “Understanding Chargeback,” on page 155](#)
- ♦ [Section 7.2, “Using Raters,” on page 155](#)
- ♦ [Section 7.3, “Managing Raters,” on page 158](#)
- ♦ [Section 7.4, “Chargeback Reporting,” on page 161](#)

7.1 Understanding Chargeback

The PlateSpin Recon flexible chargeback reporting capabilities allow organizations to accurately calculate IT costs based on actual and virtualized resource usage. IT departments can then allocate their costs to different business units and departments.

Even when resources are part of a virtualized pool of computing resources shared across many workloads, PlateSpin Recon not only allows you to track, assign, and report on used resources, but it can also generate reports and deliver them to the appropriate business units.

To set up a chargeback system:

- 1 Create raters to assign appropriate charges for usage.
For more information on how to create a rater, see [Section 7.3.1, “Creating a Rater,” on page 158](#)
- 2 Attach the raters to particular machines to be tracked.
For more information on how to attach raters to machines, see [Section 7.3.4, “Attaching Raters to Machines,” on page 160](#)
- 3 Generate Chargeback reports.
For more information on how to generate a Chargeback report, see [Section 5.5, “Generating Reports,” on page 110](#).

NOTE: It is useful to organize machines using Chargeback into their own group. For more information, see [Section 3.4, “Working with Groups,” on page 52](#).

7.2 Using Raters

A rater is a metric used to assign a billable charge to the use of a resource—or part of a resource—over a set amount of time. This converts resource usage data into currency values. There are three types of raters:

- ♦ **Allocation Raters:** These raters can only be attached to virtual machines. They take a cost for a virtual machine server resource and assign an appropriate percentage of that cost to each virtual machine the rater is attached to, based on the usage of that resource.
- ♦ **Flat Raters:** These raters assign a value to a resource at a fixed rate. For example, a flat rater charge for used disk space could be \$1 per gigabyte of disk space used per month.

- ♦ **Tiered Raters:** These raters assign values to a resource with different flat rates for different levels of usage. For example, a tiered rater charge for memory used could be \$0 for 0-512 MB and \$.05 per MB for use per day above 512 MB.

Review the following sections:

- ♦ [Section 7.2.1, “Rater Resources,” on page 156](#)
- ♦ [Section 7.2.2, “Rater Charges,” on page 157](#)
- ♦ [Section 7.2.3, “Virtual Server Rater Methods,” on page 157](#)

7.2.1 Rater Resources

Raters calculate charges either according to the maximum use of a resource, or according to the average use during a specific time period. The exception is when they charge against hourly data, where charges are calculated using the average use only.

Table 7-1 *Rater Resources*

Resource	Description
Uptime (Hours)	The charge is based on the number of hours the network entity has been monitored by PlateSpin Recon during the time specified in the report. <i>Measurement and Time Period are dimmed for Flat and Tiered Raters. Method is also dimmed for Allocation Raters.</i>
VC_ProcessorTime (hour)	The hourly average usage is computed and the sum of all the hour averages during the time specified in the report is used for the charge. <i>Measurement and Time Period are dimmed.</i>
Processor (MHz)	Selecting the Hour <i>Time Period</i> forces the use of the Average value, even if Maximum is selected.
Memory Used(MB)	
Network (MB/sec)	
Disk (MB/sec)	
Disk Space Used (GB)	
Disk Writes (MB/sec)	
Total Processor (Count)	Static counter. The charge is made against a static value. <i>Measurement and Time Period are dimmed.</i>
Total Memory (MB)	Static counter. The charge is made against a static value. <i>Measurement and Time Period are dimmed.</i>
Total Disk (GB)	Static counter. The charge is made against a static value. <i>Measurement and Time Period are dimmed.</i>

NOTE: Aside from *VC_ProcessorTime*, described in the preceding table, VC_specific counters are not required nor useful for chargeback, and as a result are not available when creating raters. Use the standard memory counters (*Memory Used* and *Total Memory*) as needed.

7.2.2 Rater Charges

Flat and Tiered raters: Charge represents the amount of cost assigned to a machine for every measure of usage. For example, for every 3 GB of disk space used per week, apply the charge. For Tiered raters, a Charge value must be entered for each tier.

Allocation raters: Charge represents the total cost of a virtual machine server resource. For example, a rater could include a charge of \$1000 per month for virtual machine server disk space. Every virtual machine to which the rater is attached is assigned a percentage of the \$1000, proportional to its usage of the selected resource. For example, based on the above rater, if a virtual machine server contains 3 virtual machines that use 50%, 25%, and 10% of its disk space, respectively (either Virtual Server Usage or Virtual Server Capacity), they are assigned \$500, \$250, and \$100 per month in chargeback costs.

IMPORTANT: When using Chargeback with Solaris Zones, allocation raters calculate utilization charges against the host, not the zone, using the formula:

- ♦ $\text{charge} = \text{unit price} * (\text{zone total resource} / \text{host total resource})$

If the zone and the host have the same number of resources—CPU, disk or memory—then the result is as expected; but if they differ, unexpected results could be seen in the Cost Allocation report.

For example, if the host has two CPUs but the zone only has one and the charge in the rater for the Processor resource is specified as \$1, then with the above formula the resulting charge would be \$0.50.

If the host has a considerably larger number of resources compared to the zone, especially when looking at memory and disk resources, the resulting charge could be very small. Since the Chargeback Cost Allocation report displays only two decimal places, the conversion could theoretically lose these small values and the charge on the report could appear as zero.

7.2.3 Virtual Server Rater Methods

The Virtual Server Rater Method determines whether the cost assigned to any virtual machine the rater is attached to is calculated based on its percent of usage of the virtual machine server total Virtual Server Usage for that resource, or based on its percent of usage of the virtual machine server total Virtual Server Capacity for that resource.

For example, if a virtual machine server has 1000 GB of disk space, but only 800 are actually used, a virtual machine using 200 GB of disk space is using 20% of Virtual Server Capacity, or 25% of Virtual Server Usage. For a Charge of \$1000/month, the virtual machine would be charged \$200 if Virtual Server Capacity is selected, but \$250 if Virtual Server Usage is selected.

7.3 Managing Raters



The Rater Explorer window is used to manage raters. You can add, edit, or delete raters, as well as attach or detach raters to or from discovered machines. The Rater Explorer also allows you to view current raters and their attachments.

A machine can have multiple raters and a rater can be used with multiple machines. Ratings are performed on a one machine to one rater basis (that is, a rater only considers data from one machine at a time).


- ♦ [Section 7.3.1, “Creating a Rater,” on page 158](#)
- ♦ [Section 7.3.2, “Deleting a Rater,” on page 159](#)
- ♦ [Section 7.3.3, “Editing a Rater,” on page 159](#)
- ♦ [Section 7.3.4, “Attaching Raters to Machines,” on page 160](#)
- ♦ [Section 7.3.5, “Detaching a Rater from a Machine,” on page 161](#)

7.3.1 Creating a Rater

1 In the Rater Explorer, do any of the following:

- ♦ Click the  icon.
- ♦ Right-click the Rater type you want to create, then click  *New*.

The New Rater dialog box is displayed.



The "New Rater" dialog box is shown with a blue title bar and standard Windows window controls. It contains three sections: "General", "Measure", and "Charge". In the "Rater Type" section at the top, "Flat Rater" is selected with a radio button. The "General" section has a "Name" field containing "A1 Rater" and an empty "Description" field. The "Measure" section has a "Resource" dropdown set to "Processor (MHz)", "Measurement" radio buttons for "Average" (selected) and "Maximum", "Time Period" radio buttons for "Hour" (selected), "Day", "Week", and "Month", and a "Factor" field set to "1.00". The "Charge" section has a "Charge" field set to "1.00" and a currency dropdown set to "Canadian Dollar (\$)". At the bottom are "Help", "OK", and "Cancel" buttons.

2 Select a Rater type.

The Rater type is automatically selected if you chose to create a new rater by right-clicking the Rater type in [Step 1](#). For more information on the Rater types, see [Section 7.2, “Using Raters,” on page 155](#).

3 Specify a name for the Rater.

4 (Optional) Specify a description.

5 In the *Resource* drop-down list, select a Rater resource.

Resources are also known as counters, because their purpose is to count usage of a resource. For more information on resources, see [Section 7.2.1, “Rater Resources,” on page 156](#).

6 Select the *Measurement* and *Time Period* to use, if they are not dimmed.

For *Measurement*, selecting *Average* means that the average value over the *Time Period* is used. Selecting *Maximum* means that out of all the averages within the *Time Period*, the maximum average is used. For example, there are 24 averages per day for hourly data and the maximum average out of all the 24 averages is used against the charge value.

7 In the *Factor* option, specify a factor, which is the multiple of the resource units selected that you want to use for the charge.

For example, you could charge \$2.00 per 3 GB of disk space used. In this case, the *Factor* would be 3.

8 In the *Charge* option, specify the Rater charges (monetary value) and select a currency type from the drop-down list.

For more information, see [Section 7.2.2, “Rater Charges,” on page 157](#).

9 (Conditional) For Allocation raters (virtual machines only), select a Virtual Server Rater method in the *Method* drop-down list.



For more information, see [Section 7.2.3, “Virtual Server Rater Methods,” on page 157](#).

10 Click *OK*.

The new rater is displayed in the Rater Explorer window under the appropriate rater type.



7.3.2 Deleting a Rater

In the Rater Explorer Window, do any of the following:

- ♦ Select the rater you want to delete, click the  icon or Press Delete.
- ♦ Right-click the rater you want to delete, then click  *Delete*.

7.3.3 Editing a Rater

1 In the Rater Explorer, do any of the following:

- ♦ Select the rater you want to edit, then click the  icon.
- ♦ Right-click the Rater type you want to edit, then click  *Edit*.

The Edit Rater dialog box is displayed.

2 Edit the rater values as desired.

NOTE: You cannot edit the Rater type. If you want to change a rater type, you must delete the rater, then create a new rater with the same name and the desired type.

3 Click *OK*.

A confirmation prompt window is displayed.

4 Click *OK*.

If you have changed the rater name, the rater with the old name is removed from the Rater Explorer window, and the rater with the new name is added to the Rater Explorer window.



Deleting a rater that is attached to a machine does not affect the machine, except to exclude it from any chargeback calculations unique to the deleted rater.

7.3.4 Attaching Raters to Machines

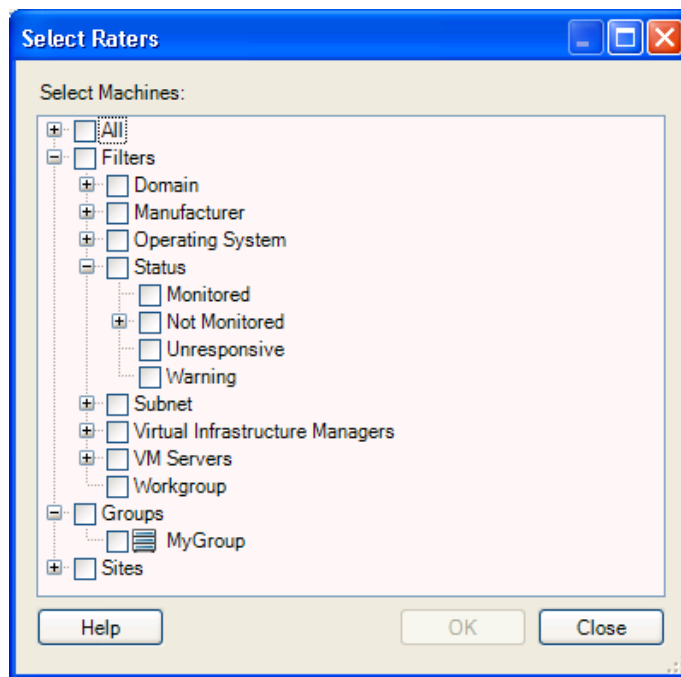
A machine needs to be inventoried and in most cases monitored to enable a rater to access its resources. You then need to attach the rater to the machine to have the machine cost calculated and made available for reporting.

To calculate a charge, PlateSpin Recon must have monitoring data for all counters except Total Processor, Total Memory, and Total Disk, which use inventory data.

1 In the Rater Explorer, do any of the following:

- ♦ Select the rater you want to attach to a machine, then click the  icon.
- ♦ Right-click the Rater type you want to edit, then click  *Attach*.

The Select Raters dialog box is displayed.



You can also use the Data Center Explorer window to attach a rater to a machine. Right-click the machine you want to add the rater to, then click *Chargeback* in the menu. The Raters for *machine_name* window is displayed where you can add or remove raters by selecting them and using the *Add* and *Remove* buttons. Adding and removing raters from here is the same as attaching or detaching them from the Rater Explorer window.

2 Select the check box next to the machine you want to add the rater to.

You can select more than one machine.



You can attach Raters to groups, sites, or filters by selecting them in the Select Raters dialog box.

3 Click *OK*.

The machines you attached the rater to are displayed in the Rater Explorer window beneath that rater.

7.3.5 Detaching a Rater from a Machine

In the Rater Explorer, do any of the following:

- ♦ Select the machine you want to detach the rater from, then click the  icon.
- ♦ Right-click the machine you want to detach the rater from, then click  *Detach*.

The rater is immediately detached and the machine is no longer displayed beneath the rater in the Rater Explorer window.

7.4 Chargeback Reporting

Use the Cost Allocation report template under the Chargeback node in the Report Explorer window to calculate chargeback values for a particular time period and to generate a report that includes chargeback data. In addition to the columns available for inventory reports, Chargeback reports can include raters attached to machines, their types, their calculated costs, and a description.

- ♦ [Section 7.4.1, “Generating a Report to Calculate IT Chargeback Cost,” on page 161](#)
- ♦ [Section 7.4.2, “Deriving the Calculated Cost,” on page 161](#)

7.4.1 Generating a Report to Calculate IT Chargeback Cost

To generate a report that calculates an IT chargeback cost for a particular business unit, ensure that the *Chargeback - Calculated Cost* column is included in the report. In the *Groups* tab of the Create View dialog box, create a grouping by a column such as *Machine - Department* or *Machine - Billing Code*.

The generated report includes a total chargeback cost for every department, billing code, or any other column used for grouping.

For more information on generating reports, see [Chapter 5, “Generating and Viewing Charts and Reports,” on page 85](#).

7.4.2 Deriving the Calculated Cost

The *Calculated Cost* in the Cost Allocation report is derived based on the Rater a machine is attached to and the date range specified.

Figure 7-1 Sample Rater

Edit Rater

Rater Type: ☐ Allocation Rater ☒ Flat Rater ☐ Tiered Rater

General

Name:

Description:

Measure

Resource:

Time Period: ☐ Hour ☐ Day ☒ Week ☐ Month

Measurement: ☒ Average ☐ Maximum

Factor:

Charge

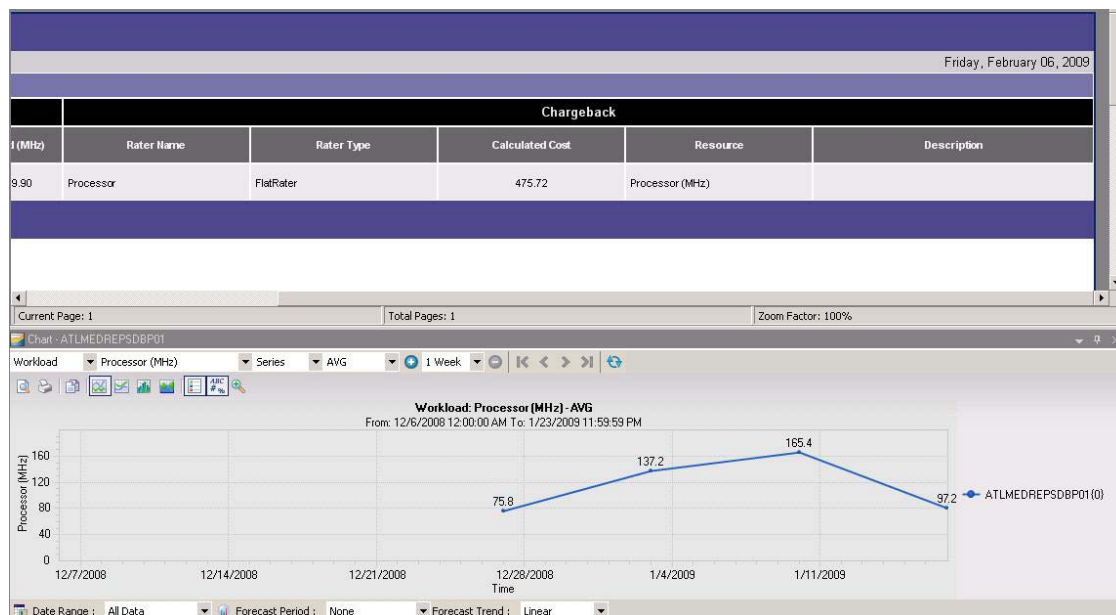
Charge:

Help OK Cancel

A machine attached to the above Rater would be charged \$1.00 for every week of 1 MHz processor use.

When the Cost Allocation report is run for this machine, PlateSpin Recon adds the weekly processor use together for each week in the report and multiplies it by the charge specified in the Rater. The following figure shows a sample processor use chart and the chargeback Cost Allocation report that would result.

Figure 7-2 Sample Chart and Correlating Chargeback Report



The four data points are added together, multiplied by the unit cost, giving you the Calculated Cost. The slight difference between this calculation and what is shown in the report is due to the chart rounding values to a single decimal place. In this case:

- ♦ $(75.8 + 137.2 + 165.4 + 97.2) * 1.00 = 475.6$

The PlateSpin Recon Virtual Infrastructure Management tools facilitate the administration of Virtual Infrastructure in the data center. These include Virtual Center monitoring and the ability to set up recurrent inventory jobs to have the most up-to-date information about virtual machines being managed by a Virtual Center. VM Sprawl reports can also be generated on a one-time or recurrent basis to provide information such as the age, last modification, and last boot time of the virtual machines in the data center.

- ♦ [Section 8.1, “Virtual Center Inventory and Monitoring,” on page 165](#)
- ♦ [Section 8.2, “VM Sprawl Reporting,” on page 166](#)
- ♦ [Section 8.3, “Virtual Machine State Reporting,” on page 167](#)

8.1 Virtual Center Inventory and Monitoring


Direct Virtual Center inventory and monitoring enables multiple virtual machines to be more efficiently inventoried from one single point instead of inventorying them individually.

Monitoring through the Virtual Center provides different information than monitoring a virtual machine directly, because it is a picture of the virtual machine resource usage as detected by the Virtual Center and ESX host. This is more accurate than a virtual machine reporting on its own resource utilization. For example, a virtual machine using 90% of its allotted processing capabilities reports its processor usage as 90%. In reality, it might have been allocated only 30% of the host processor capabilities. This difference can only be detected by monitoring through the Virtual Center, which reports what percentage of the host processor a virtual machine is using.

This difference has become more and more useful as data centers place emphasis on managing and optimizing virtual machine usage as opposed to merely consolidating physical servers onto virtual machines.

When a Virtual Center is inventoried, its associated machines are displayed in the Data Center Explorer in a hierarchy that reflects the VMware Virtual Infrastructure Client hierarchy, including all data centers, folders, clusters and pools.

Figure 8-1 Data Center Explorer and VMware Virtual Infrastructure Client Comparison



Inventorying a Virtual Center discovers all the machines that are part of the Virtual Center, so that they can then be inventoried and monitored. When a virtual machine is inventoried through the Virtual Center, it shows up as inventoried both as a node in the Virtual Center, and as an inventoried machine under All in the Data Center Explorer. Either node can subsequently be used to initiate monitoring.

You can right-click a cluster in a Virtual Center to view its properties, which are set in the Virtual Center.

When a virtual machine that is part of a Virtual Center is monitored, certain values, such as Processor Time, are collected both directly from the virtual machine and through the Virtual Center. Values that are being collected through the Virtual Center are prefixed with VC, for example VC_ProcessorTime.

For information about how to inventory a Virtual Center, see [“Virtual Center” on page 31](#)

8.2 VM Sprawl Reporting

The VM Sprawl report provides information about all virtual machines found by PlateSpin Recon, along with information about their servers.

NOTE: If a VM is under a cluster and the parent host ESX machine is not inventoried by PlateSpin Recon, the VM is displayed under the Cluster group of this report.

You can use this report to keep track of the number and location of virtual machines resident in the data center. VM Sprawl report information is obtained from the virtual machines' host servers.

Citrix XenServer pools can also be inventoried by PlateSpin Recon, and these pools are treated similarly to a Virtual Center. The VM Sprawl report can also be used to report on these pools.

The VM Sprawl Report template is located in the Inventory Reports section of the Reports Explorer. For more information about Inventory reports, see [Section 5.2.5, “Working with Report Templates,” on page 96](#).

8.3 Virtual Machine State Reporting

The Virtual Machine State report shows only the machines displayed under the VM Servers filter. It is found under the Inventory Reports section of the Reports Explorer. For more information about Inventory reports, see [Section 5.2.5, “Working with Report Templates,” on page 96](#).

Troubleshooting

A

This section provide solutions to the issues you might encounter when using Novell PlateSpin Recon:

- ♦ “Refresh Inventory or Monitoring fails for a machine whose operating system has been reformatted” on page 170
- ♦ “Disk partitions for Citrix XenServer are not displayed” on page 171
- ♦ “PlateSpin Recon fails to inventory the FC SAN disk connected to a Citrix XenServer 4.0 machine” on page 171
- ♦ “Inventory of Solaris or Linux machines might fail with an error” on page 171
- ♦ “Disk IO metrics are not available for ESX 3.0” on page 171
- ♦ “The Queue Length metrics are not collected for Virtual Center and ESX” on page 172
- ♦ “Monitoring a Windows machine fails with an error related to access to the registry keys” on page 172
- ♦ “Disk partition counters report errors for the AIX machines” on page 172
- ♦ “The loop7 partitions of Citrix XenServer are inventoried and monitored as a new disk” on page 172
- ♦ “Unable to detect the processor model for a Windows operating system installed on Connoi hardware” on page 173
- ♦ “Issues when attempting to add servers to the PlateSpin Recon Inventory” on page 173
- ♦ “Importing a machine from PlateSpin Recon snapshot fails” on page 173
- ♦ “A warning message is logged in to the monitoring log while monitoring the Windows XP machines” on page 173
- ♦ “Monitoring of data collection for the newly added disks or volumes is delayed for Windows machines” on page 174
- ♦ “The data for the counters related to disk utilization might not be collected for Solaris servers” on page 174
- ♦ “The data of a few disk counters for the disconnected disks on Solaris is incorrectly displayed” on page 174
- ♦ “Inventory of an ESX machine fails if its virtual machines have duplicate SMBIOS IDs” on page 174
- ♦ “Monitoring of a Windows machine might fail” on page 175
- ♦ “Generating of the Disk Analysis report or the raters based on disk usage for Windows machines might fail” on page 175
- ♦ “Monitoring of a Windows machine displays the “Category does not exist” error in the monitoring log” on page 175
- ♦ “W3P Service consumes high memory or is terminated with system memory out of exception” on page 177
- ♦ “Disk Space % in scenarios is shown as 0 even though Disk Space (GB) has a value when you view the chart for a workload” on page 177

- ♦ “Monitoring of a virtual machine that runs on a Windows 2008 SP1 Hyper-V displays an error for the % Total Run Time counter” on page 178
- ♦ “Monitoring of a Windows virtual machine that runs on a Citrix XEN server might displays an error for the Xen_NetworkInterface_BytesTotalPerSecond and Xen_PhysicalDisk_DiskBytesPerSecond counters” on page 178
- ♦ “Deploying or undeploying of the Windows Disk Counters add-on fails if you select to perform it on Windows Server 2008 SP2 Hyper-V host and its guest virtual machine” on page 178
- ♦ “Unable to generate Cost allocation report with Allocation rater” on page 178
- ♦ “Monitoring of a Linux machine by using sudo user credentials might fail” on page 179
- ♦ “The value of HyperV_NetworkInterface_BytesTotalPerSecond is incorrect for virtual machines that are configured to a legacy network adapter and run on Windows Server 2008 SP1/SP2/R2 Hyper-V host” on page 179
- ♦ “Restarting of Remote Registry service after deployment or undeployment of the Windows Disk Counters Add-On fails with errors” on page 180
- ♦ “The deployment or undeployment of the Windows Disk Counters add-on might result in the error, “The process cannot access the file because it is being used by another process”” on page 180
- ♦ “Refresh Inventory and Monitoring of machines inventoried by using the CSV file might fail” on page 181
- ♦ “Performing Stop Monitoring and Start Monitoring jobs in immediate successions on a large set of UNIX machines might result in an error” on page 181
- ♦ “PlateSpin Recon Client might incorrectly display the status of a job as running even after the job has been completed” on page 181
- ♦ “Stopping of the PlateSpin Recon Monitoring service might take considerable time or restarting of the PlateSpin Recon Monitoring service might fail” on page 182
- ♦ “Logs Explorer might display the warning, “Client data may be out of date”” on page 182
- ♦ “The PlateSpin Recon Monitoring service is terminated because of the out of memory exception” on page 182
- ♦ “Loading of the PlateSpin Recon job from the Portability Suite might display a credential-specific error” on page 183
- ♦ “Inventory of Solaris machine fails” on page 183
- ♦ “Creating scenarios for forecasted projects might fail with the out of memory exception” on page 183
- ♦ “Generating Analysis reports from the PostgreSQL database might take a considerable time” on page 183
- ♦ “The storage type for disks attached to AIX machines is displayed as Unknown” on page 184
- ♦ “Importing of a snapshot that contains Citrix Xen virtual machines fails” on page 184
- ♦ “The operating system type for Windows 2003 R2 and Windows 2008 R2 devices is displayed as Unknown” on page 184

Refresh Inventory or Monitoring fails for a machine whose operating system has been reformatted

Source: PlateSpin Recon; Inventory and Monitoring.

- Explanation: If you reformat the operating system of a machine after it has been inventoried, the Refresh Inventory or Monitoring for the machine fails.
- Possible Cause: The credentials and the monitoring plug-in associated with the machine correspond to the original operating system.
- Action: Perform the following tasks in the Data Center Explorer of the PlateSpin Recon Client:
- 1 Delete the existing object of the machine whose operating system has been reformatted.
 - 2 Inventory or monitor the machine.
- For more information on how to inventory or monitor a machine, see [Section 3.2, “Discovering and Inventorying Data Center Resources,”](#) on page 30.

Disk partitions for Citrix XenServer are not displayed

- Source: PlateSpin Recon; Inventory.
- Explanation: The disk partitions of an inventoried Citrix XenServer are not displayed in the *Storage* tab of the *Properties* page of the server.
- Possible Cause: Citrix XenServer is formatted in to logical volumes by default and not in to partitions. PlateSpin Recon currently does not have LVM inventorying and reporting capability.
- Action: None.

PlateSpin Recon fails to inventory the FC SAN disk connected to a Citrix XenServer 4.0 machine

- Source: PlateSpin Recon; Inventory.
- Explanation: PlateSpin Recon does not inventory the FC SAN disk attached to a Citrix XenServer 4.0 machine. Consequently, the disk is not displayed in the *Storage* tab of the *Properties* page of Citrix XenServer.
- Possible Cause: PlateSpin Recon does not have supported operating system commands and Xen APIs to identify FC SAN Disk.
- Action: None.

Inventory of Solaris or Linux machines might fail with an error

- Source: PlateSpin Recon; Inventory.
- Explanation: The inventory of a Solaris or Linux machine might fail with an error
- Possible Cause: The `/tmp` directory on the Solaris or Linux machine does not have sufficient free space.
- Action: Before inventorying a Solaris or Linux machine, ensure that the `/tmp` directory on the machine has a minimum of 50 MB of free space.

Disk IO metrics are not available for ESX 3.0

- Source: PlateSpin Recon.

Explanation: The Disk IO% counter is not displayed for ESX 3.0 in the Chart Viewer.

Possible Cause: ESX 3.0 does not support the required counters used in the calculation of Disk IO%.

Action: None.

The Queue Length metrics are not collected for Virtual Center and ESX

Source: PlateSpin Recon; Monitoring.

Explanation: The QueueLength metrics are not collected for any type of storage volumes (local or shared) for ESX monitored directly or through vCenter.

Possible Cause: Virtual Center and ESX do not support the required counters for the QueueLength metrics.

Action: None.

Monitoring a Windows machine fails with an error related to access to the registry keys

Source: PlateSpin Recon; Monitoring.

Explanation: When monitoring a Windows machine fails, the monitor log contains the following error:

`Access to the registry key key_number is denied.`

Possible Cause: The registry key is being read by some other application at the same time.

Action: Ignore the errors. PlateSpin Recon automatically collects the data after the registry key is released by the other application.

Disk partition counters report errors for the AIX machines

Source: PlateSpin Recon; Monitoring.

Explanation: The data is not collected for the following counters for the AIX machines:

- ◆ PhysicalDisk_PartitionedSpaceGB
- ◆ PhysicalDisk_UnpartitionedSpaceGB
- ◆ PhysicalDisk_PercentPartitionedSpace
- ◆ PhysicalDisk_PercentUnpartitionedSpace counters

Additionally, a warning message is displayed in the logs. The warning message also includes the solution for the issue.

Possible Cause: PlateSpin Recon internally uses the `lspv` command to fetch the disk partition counters for AIX. When the command is run on an AIX machine, the machine might hang or slow down. This is a known issue with AIX.

Action: Perform the tasks as suggested in the warning message.

The loop7 partitions of Citrix XenServer are inventoried and monitored as a new disk

Source: PlateSpin Recon; Inventory and Monitoring

Explanation: When you mount the XenServer tools ISO to the DVD drive at `/dev/xvd` on a VM of a Citrix XenServer host to install the XenServer tools, the Citrix XenServer host creates a disk partition with the same size as that of the mounted ISO. When you inventory or monitor the Citrix XenServer, PlateSpin Recon considers the disk partition as a new disk, and reports it as `loop7` in the machine properties and charts.

Action: None.

Unable to detect the processor model for a Windows operating system installed on Connoi hardware

Source: PlateSpin Recon; Inventory and Monitoring.

Action: None.

Issues when attempting to add servers to the PlateSpin Recon Inventory

Source: PlateSpin Recon; Inventory.

Explanation: You might encounter issues when attempting to add servers to the PlateSpin Recon Inventory, including the following error messages:

`Network Path Not found`

`Access Denied`

`The RPC Server is unavailable`

`Failed. The Network location cannot be reached`

Action: Refer to Knowledge Base article [Q7920525](http://www.novell.com/support/viewContent.do?externalId=7920525) (<http://www.novell.com/support/viewContent.do?externalId=7920525>): Troubleshooting problems when adding servers to the PlateSpin Recon Inventory.

Importing a machine from PlateSpin Recon snapshot fails

Source: PlateSpin Recon; Snapshots.

Explanation: When attempting to import a machine from a PlateSpin Recon snapshot, you might encounter the following error message:

`Cannot import MachineObject 'MACHINE-NAME [MACHINE-NAME]'`

Action: Refer to Knowledge Base article [Q7920900](http://www.novell.com/support/viewContent.do?externalId=7920900) (<http://www.novell.com/support/viewContent.do?externalId=7920900>): Cannot import MachineObject error when importing a machine from a snapshot.

A warning message is logged in to the monitoring log while monitoring the Windows XP machines

Source: PlateSpin Recon; Monitoring.

Explanation: While monitoring a Windows XP machine, the following warning is logged in to the monitoring log:

`The interface is unknown counter_name.`

Possible Cause: The Remote Registry service is in the Stop state on the target Windows machine that is being monitored.

Action: Refer to the [Microsoft Knowledge Base article ID 923416 \(http://support.microsoft.com/kb/923416\)](http://support.microsoft.com/kb/923416).

Monitoring of data collection for the newly added disks or volumes is delayed for Windows machines

Source: PlateSpin Recon; Monitoring.

Explanation: If you add a new disk or a volume to a Windows machine that is being monitored, the data of the new disk or volume is not immediately collected and reflected in the PlateSpin Recon charts.

Possible Cause: PlateSpin Recon reports the data for only the disks and volumes that exists on the target Windows machine when you perform Start Monitoring on the machine.

Action: Stop and restart the monitoring of the Windows machine to which you have added a new disk or volume:

- 1 In the Data Center Explorer, navigate to target Windows machine.
- 2 Right-click the target machine, then click *Stop Monitoring*.
- 3 Right-click the target machine again, then click *Start Monitoring*.

The data for the counters related to disk utilization might not be collected for Solaris servers

Source: PlateSpin Recon; Monitoring.

Possible Cause: The external disks connected to the Solaris machine are not labeled.

Action: Assign a label to the connected disks on the Solaris machine.

The data of a few disk counters for the disconnected disks on Solaris is incorrectly displayed

Source: PlateSpin Recon; Monitoring.

Explanation: The data of the following counters for the disconnected disks on a Solaris machine is incorrectly displayed as zero (0):

- ♦ DiskBytesPerSecond
- ♦ DiskAverageQueueLength
- ♦ DiskTransfersPerSecond
- ♦ DiskPercentIdleTime
- ♦ Disk IO%

Possible Cause: The external disks are not properly disconnected.

Action: Clean the file system on the Solaris machine in one of the following ways:

- ♦ Run the `devfsadm -C` command.
- ♦ Reboot the machine.

Inventory of an ESX machine fails if its virtual machines have duplicate SMBIOS IDs

Source: PlateSpin Recon; Inventory and Monitoring.

Possible Cause: PlateSpin Recon uses SMBIOS IDs to establish the uniqueness of the machines. If the virtual machines of an ESX machine have same SMBIOS IDs, the inventory of the ESX machine fails while importing the data.

Action: Ensure that the virtual machines of an ESX machine have unique SMBIOS IDs.

Monitoring of a Windows machine might fail

Source: PlateSpin Recon; Monitoring.

Explanation: When monitoring of a Windows machine fails, the following error message is logged into the monitoring logs:

The network path was not found

Possible Cause: Unable to access the Windows machine as a network-shared resource.

Action: Enable the network shared resource on the Windows machine.

Possible Cause: The Remote Registry service is not running on the target Windows machine.

Action: Manually start the Remote Registry service on the Windows machine.

Generating of the Disk Analysis report or the raters based on disk usage for Windows machines might fail

Source: PlateSpin Recon; Monitoring.

Explanation: Generating of the Disk Analysis report or the raters based on disk usage for Windows machines might fail with the following message that is logged in the monitoring logs:

The component Disk Space Used(GB) does not have enough hourly summary data for the machine *IP_address*.

Possible Cause: The target Windows machine does not have the Disk Utilization add-on deployed.

Action: Do the following:

- 1 Deploy the Disk Utilization add-on on the target Windows machine.
For more information about deploying the Disk Utilization add-on, see [Section 3.12, “Collecting Disk Utilization Counters for Windows Machines,”](#) on page 74.
- 2 Monitor the Windows machine for a minimum duration of one hour.
- 3 Generate the Disk Analysis report or the rater.

Monitoring of a Windows machine displays the “Category does not exist” error in the monitoring log

Source: PlateSpin Recon; Monitoring.

Possible Cause: After performing the Undeploy Add-On task on the Windows machine, you did not restart monitoring on the machine by performing Stop Monitoring and Start Monitoring actions.

Action: Restart monitoring on the Windows machine by first performing the Stop Monitoring action, and subsequently the Start Monitoring action.

Possible Cause: The Windows Disk Counters Add-On is manually undeployed from the machine and not through the *Undeploy Add-On* option in the Recon PlateSpin Client.

Action: Do the following in the PlateSpin Recon Client:

- 1 Right-click the Windows machine object on which you have manually undeployed the Windows Disk Counters Add-On.
- 2 Click *Undeploy Add-On*.
- 3 Right-click the Windows machine object.
- 4 Click *Stop Monitoring*.
- 5 Right-click the Windows machine object.
- 6 Click *Start Monitoring*.

Possible Cause: When you run the following command on the PlateSpin Recon server, it results in system errors.

```
net view \\IPAddress_of_remote_machine
```

Action: To troubleshoot the system errors, see the Microsoft Knowledge Base articles appropriate to the system error code.

Possible Cause: Data collection from the PlateSpin Recon option, Windows Disk Counters Add-On, has been disabled by perfmon.

Action: Do the following on the Windows machine:

- 1 In the Registry Editor, navigate to
HKLM\SYSTEM\CurrentControlSet\Services\ReconWinMonitoring\Performance.
- 2 Change the value of the registry key *Disable Performance Counters* to 0.
- 3 Restart the Remote Registry service.

Possible Cause: Perfmon of the target Windows machine is corrupted.

Action: Rebuild the Performance counters by referring to [Microsoft Knowledge Base article - How to manually rebuild Performance Counter Library values \(http://support.microsoft.com/kb/300956\)](http://support.microsoft.com/kb/300956).

Possible Cause: The event logs of the target windows machine show Perflib errors with event id 1015.

```
Event ID 1015 Source Perflib
Event ID 1015
Source Perflib
Type Error
Description The timeout waiting for the performance data
collection function"<function name>" in the "<dll name>"
Library to finish has expired. There may be a problem with
this extensible counter or the service it is collecting data
from or the system may have been very busy when this call was
attempted.
```


Action: By default, the system uses the same collect timeout period of 10 seconds for all services. Override the period for an failing service:

- 1 In the Registry Editor, navigate to
HKLM\SYSTEM\CurrentControlSet\Services\<service-name>\Performance.
- 2 Increase the value of Collect Timeout.
- 3 Restart the Remote Registry service.

Possible Cause: You do not have sufficient privileges to read the perfmon counters on the remote Windows machine.

Action: Request the administrator of the remote machine to grant you sufficient privileges to read the perfmon counters through the following User Rights policies: Profile System Performance and Profile a Single Process.

Possible Cause: While configuring the deployment of the Windows Disk Counters add-on, you chose not to select the *Automatically restart the services which are dependent on the Remote Registry Service* check box but you did not manually restart the Remote Registry service and the services that are dependent on the Remote Registry service after the add-on was deployed.

Action: On the target Windows machine where you deployed the Windows Disk Counters add-on, manually restart the Remote Registry service and its dependent services.

W3P Service consumes high memory or is terminated with system memory out of exception

Source: PlateSpin Recon.

Possible Cause: If the PlateSpin Recon Client is open without being used for more than 48 hours, the W3P service caches data for the Recon client session and this results in the W3P Service out of memory.

Action: If the W3P service is terminated, close the PlateSpin Recon Client, start the W3P service, then launch the PlateSpin Recon Client.

Action: If the W3P service is not terminated but the W3P service memory consumption is high, relaunch the PlateSpin Recon Client.

Disk Space % in scenarios is shown as 0 even though Disk Space (GB) has a value when you view the chart for a workload

Source: PlateSpin Recon.

Explanation: If you select a server template that has unlimited total disk space to create a scenario, the following are the known behaviors:

- ♦ The chart for the Disk Space (%) utilization, viewed on the server template and the assigned workloads to that template, always shows zero (0) for y.
- ♦ The workload assignment report on the server template shows empty (--) values for the Disk Used (%) column.

Action: None.

Monitoring of a virtual machine that runs on a Windows 2008 SP1 Hyper-V displays an error for the % Total Run Time counter

Source: PlateSpin Recon; Monitoring.

Explanation: Monitoring of a virtual machine that runs on a Windows 2008 SP1 Hyper-V reports the following error message for the % Total Run Time counter in the monitoring logs for the node:

```
Cannot get instance information for counter % Total Run Time  
in category Hyper-V Hypervisor Virtual Processor on computer  
w2k8vm204
```

Possible Cause: The data collection for the % Total Run Time counter, which is listed in the Hyper-VHypervisor Virtual Processor counter category, causes the error because the Microsoft Performance Library API does not work as expected.

Action: None.

Monitoring of a Windows virtual machine that runs on a Citrix XEN server might displays an error for the Xen_NetworkInterface_BytesTotalPerSecond and Xen_PhysicalDisk_DiskBytesPerSecond counters

Source: PlateSpin Recon; Monitoring.

Explanation: When you monitor a Windows virtual machine that runs on a Citrix XEN server, you might encounter the following issues:

- ♦ The HANDLE_INVALID exception when monitoring the following counters:
 - ♦ Xen_NetworkInterface_BytesTotalPerSecond
 - ♦ Xen_PhysicalDisk_DiskBytesPerSecond.
- ♦ The data is not collected for the Xen_Processor_PercentProcessorTime counter even though there is no exception.

Possible Cause: The XEN API does not work as expected.

Action: None.

Deploying or undeploying of the Windows Disk Counters add-on fails if you select to perform it on Windows Server 2008 SP2 Hyper-V host and its guest virtual machine

Source: PlateSpin Recon; Deploying or Undeploying of Windows Disk Counters.

Possible Cause: The deployment or undeployment of the Windows Disk Counters add-on on a Windows host internally triggers the deployment or undeployment action for its Windows virtual machines also. When you simultaneously select Windows host and its Windows virtual machines as deployment or undeployment candidates, two deployment or undeployment jobs are triggered for the Windows virtual machines and this might result in an error.

Action: Deployment or undeployment of the Windows Disk Counters add-on must be performed separately for the host machine and the guest virtual machine.

Unable to generate Cost allocation report with Allocation rater

Source: PlateSpin Recon; Reports.

Possible Cause: For Allocation rater, charge represents the total cost of a virtual machine server resource. It requires the host machine of the virtual machine to be inventoried and monitored. If the host machine is not inventoried and monitored, the Cost Allocation report displays the following generic error message:

Error was encountered during calculation for chargeback.

All the appropriate warning messages are logged in the Cost Allocation report's job log file located in the
Recon_server_installation_directory\logs\workflow directory.

Action: Inventory and monitor the host machine of the virtual machine.

Monitoring of a Linux machine by using sudo user credentials might fail

Source: PlateSpin Recon; Monitoring.

Explanation: Monitoring of a Linux machine by using sudo user credentials might fail with the following error message:

sudo: sorry, you must have a tty to run sudo

Possible Cause: The */etc/sudoers* configuration file on the target machine contains the *Defaults requiretty* setting.

Action: Do the following:

- 1** At the Linux console prompt, execute *visudo*.
This opens the */etc/sudoers* configuration file in the vi editor.
- 2** Comment out the following line:
Defaults requiretty
- 3** In the Data Center Explorer view of Recon PlateSpin Client, right click the Linux machine, then click *Stop Monitoring*.
- 4** Right-click the Linux machine, then click *Start Monitoring*.

The value of HyperV_NetworkInterface_BytesTotalPerSecond is incorrect for virtual machines that are configured to a legacy network adapter and run on Windows Server 2008 SP1/SP2/R2 Hyper-V host

Source: PlateSpin Recon; Monitoring.

Explanation: If a virtual machine is configured with a legacy network adapter and runs on Windows Server 2008 SP1/SP2/R2 Hyper-V host, the data for the HyperV_NetworkInterface_BytesTotalPerSecond counters is incorrectly displayed as zero (0).

Possible Cause: The Microsoft Performance Library API does not work as expected for the Bytes Received/sec or Bytes Sent/sec counters of the HyperV_NetworkInterface_BytesTotalPerSecond category on HyperV servers.

Action: None.

Restarting of Remote Registry service after deployment or undeployment of the Windows Disk Counters Add-On fails with errors

Source: PlateSpin Recon; Deployment or undeployment of Windows Disk Counters Add-On.

Explanation: When you deploy or undeploy the Windows Disk Counters add-on, the Remote Registry service is automatically restarted on the target machine. The service restart process might fail with errors. The error messages are displayed in the `job.log`.

Action: If any error is reported while restarting Remote Registry service after deploying or undeploying the Windows Disk Counters add-on, you must undeploy the Disk Performance counters DLL from the target machine `perfmon`:

- 1 On the target machine, open the Registry Editor.

- 2 Navigate to the

`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ReconWinMonitoring` key. If the key exists, it indicates that the `%SystemRoot%\PlateSpinRecon-DiskUtil` directory exists. Do the following:

- 2a From the command prompt of the target machine, go to `%systemroot%\PlateSpinRecon-DiskUtil`.

- 2b Execute the following command:

```
HWInputMon.exe -Uninstall false
```

- 2c Delete the

`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ReconWinMonitoring` key.

- 2d Delete the `%SystemRoot%\PlateSpinRecon-DiskUtil` directory.

The deployment or undeployment of the Windows Disk Counters add-on might result in the error, "The process cannot access the file because it is being used by another process"

Source: PlateSpin Recon; Deployment or undeployment of Windows Disk Counters Add-On.

Explanation: If the Windows virtual machines are directly listed in the group and are as well as listed under the Citrix XenServer or Virtual Center host machines included in the same group, deployment or undeployment of the Windows Disk Counters add-on on a group might result in the following error:

The process cannot access the file because it is being used by another process

Possible Cause: The Windows virtual machines directly receive the Windows Disk Counters add-on deployment or undeployment job from the group and also from the Citrix XenServer or Virtual Center host. In this scenario, one of the jobs succeed and the other one fails.

Action: Ignore the error because the Windows Disk Counters add-on is successfully deployed or undeployed on the Windows virtual machine through the job that has succeeded.

Refresh Inventory and Monitoring of machines inventoried by using the CSV file might fail

Source: PlateSpin Recon; Inventory.

Possible Cause: If you inventory and monitor machines by using a CSV file, the credential description that you specify in the file must be unique for the machines that contain username and password. If the credential description is not unique, then incorrect credentials might be assigned to a few machines and consequently refresh inventory and monitoring actions on those machines fail.

Action: Before you inventory and monitor machines by using a CSV file, ensure that the credential description specified in the CSV file is unique for the machines that contain username and password.

Action: If you have already inventoried machines by using an incorrect CSV file, do one of the following:

- ♦ Delete the problematic machines from the PlateSpin Recon Client, specify a unique credential description for the machines listed in the CSV file, then inventory the machines again.
- ♦ Attach the correct credentials to the problematic machines.
 1. In the Data Center Explorer, right-click the problematic machine, then click *Attach Credentials*.
The Attach Credentials dialog box is displayed.
 2. Specify the credentials for the machine, then click *OK*.

Performing Stop Monitoring and Start Monitoring jobs in immediate successions on a large set of UNIX machines might result in an error

Source: PlateSpin Recon; Monitoring.

Explanation: If you perform Stop Monitoring and Start Monitoring in immediate successions on a large set of UNIX machines, the monitoring scripts (*.sh) might not exist on the target UNIX machines and the following error message is displayed in the Job Explorer:

```
/bin/sh: ./lininfo.shSLESPlateSpinRecon_server_name: No such  
file or directory counter_name
```

Possible Cause: When you initiate the Stop Monitoring job on a large set of UNIX machines, the job might take considerable time to remove the monitoring scripts (*.sh) from the target machines even though the status of the job is displayed as completed in the Job Explorer. If you immediately initiate the Start Monitoring job for these machines, the monitoring data collection might fail because the monitoring scripts (*.sh) required to monitor the machines might also be incorrectly removed by the previous Stop Monitoring job.

Action: Stop and start monitoring the problematic machines.

PlateSpin Recon Client might incorrectly display the status of a job as running even after the job has been completed

Source: PlateSpin Recon; Inventory and Monitoring.

Explanation: PlateSpin Recon Client might fail to update the status of a job and would constantly display the status as Running even after the job is completed.

Action: Close the PlateSpin Recon Client, and relaunch it.

Stopping of the PlateSpin Recon Monitoring service might take considerable time or restarting of the PlateSpin Recon Monitoring service might fail

Source: PlateSpin Recon

Possible Cause: When you try to stop the PlateSpin Recon Monitoring service, the service commits all cached data in to the database. This process might take 10 to 15 minutes. During this period, the Stop action for the service might appear to be unresponsive.

When you try to restart the PlateSpin Recon Monitoring service, the service might fail to restart because the restart process first tries to stop the service, which takes considerable time to complete.

Action: If you had initiated the stop action for the PlateSpin Recon Monitoring service, you do not need to do anything because the service becomes responsive after 10 to 15 minutes and is subsequently stopped.

If you had initiated the restart action for the PlateSpin Recon Monitoring service, you must start the service after it becomes responsive (the service actions are available for selection).

WARNING: You must not try to terminate the PlateSpin Recon Monitoring Service process because it might lead to data inconsistency.

Logs Explorer might display the warning, “Client data may be out of date”

Source: PlateSpin Recon Client.

Explanation: If the PlateSpin Recon Client is open for a considerable time, the following warning is displayed in the Logs Explorer of the PlateSpin Recon Client:

Client data may be out of date.

Possible Cause: The Web service session that is established on launching the PlateSpin Recon client has expired.

Action: Close the current instance of the PlateSpin Recon Client and relaunch the client.

The PlateSpin Recon Monitoring service is terminated because of the out of memory exception

Source: PlateSpin Recon.

Possible Cause: PlateSpin Recon is monitoring a large number of unresponsive or unreachable machines.

Action: Stop monitoring the unresponsive or unreachable machines.

Loading of the PlateSpin Recon job from the Portability Suite might display a credential-specific error

Source: PlateSpin Recon; Portability Suite.

Explanation: Loading of the PlateSpin Recon job might fail with the following credential-specific error:

```
Unable to login to VMware Virtual Infrastructure
web... Please ensure that the username and password are
correct.
```

The error is displayed in the PlateSpin Portability Suite.

Possible Cause: The password required to access the target ESX server has not been specified.

Action: Specify the password of the target ESX server by using the Portability Suite:

1. In the Peer-to-Peer Conversion Job window, click the *Job Configuration* panel > the *Credentials* link.
2. Specify the password.
3. Click *OK*.

Inventory of Solaris machine fails

Source: PlateSpin Recon; Inventory

Possible Cause: The password of the Solaris machine contains special characters such as (

Action: Remove the special characters from the password of the Solaris machine.

Creating scenarios for forecasted projects might fail with the out of memory exception

Source: PlateSpin Recon; Scenarios and Projects.

Explanation: If you try to create a scenario from a project that consists of a large number of machines with a large number of forecast days, the process might fail with the out of memory exception.

Action: Reduce the number of machines or the forecast days in the project. For more information about the project and scenarios recommendations, see [Appendix E, “Best Practices,” on page 217](#).

Generating Analysis reports from the PostgreSQL database might take a considerable time

Source: PlateSpin Recon; Reports.

Action: Before generating Analysis reports from the PostgreSQL database, do the following:

- 1 Open the `PlateSpin.PowerRecon.Monitoring.Database.config` file using a Text editor.

The `PlateSpin.PowerRecon.Monitoring.Database.config` file is located in the

`PlateSpin_Recon_installation_directory\configs\` directory on the PlateSpin Recon Server.

- 2 Change the value of `Pooling` to `True`.
- 3 Restart the PlateSpin Recon Monitoring Service.

The storage type for disks attached to AIX machines is displayed as Unknown

Source: PlateSpin Recon.

Possible Cause: PlateSpin Recon currently does not support collecting of the storage type data for the disks attached to AIX machines. Hence, the storage type for disks attached to AIX machines is displayed as *Unknown* in the Properties page of the AIX machine (in the *Storage* tab > *View by Disk* view) and in the Reports.

Action: None.

Importing of a snapshot that contains Citrix Xen virtual machines fails

Source: PlateSpin Recon; Snapshots.

Explanation: If you try to import a snapshot that contains Citrix Xen virtual machines, the import operation might fail with the *Stealthiest* exception.

Possible Cause: You have inventoried first the Citrix Xen virtual machines and subsequently the Citrix Xen host before creating the snapshot of the machines.

Action: Before creating a snapshot with Citrix Xen virtual machines, you must follow one of the following inventory process:

- ♦ Inventory the Citrix Xen virtual machines only without inventorying the host machine.
- ♦ Inventory first the Citrix Xen host machine and subsequently the Citrix Xen virtual machines.

The operating system type for Windows 2003 R2 and Windows 2008 R2 devices is displayed as Unknown

Source: PlateSpin Recon.

Explanation: When you perform an inventory of Windows 2003 R2 and Windows 2008 R2 devices, the Inventory Summary report displays the Operating System type for these devices as *Unknown*.

Action: In the PlateSpin Recon Report Explorer, do the following:

- 1 In *Consolidation Project Samples*, right-click *Inventory Summary* and select *Edit Report Template*.
The Edit Report Template window is displayed.
- 2 Navigate to the Views panel, select *OS Summary* view and click *Edit*.
The Create View window is displayed.
- 3 Click *Customize*.
The Settings window is displayed.
- 4 Add the required operating system to the list of operating systems and click *OK*.

NOTE: Ensure that the Internal name you give to the operating system is either Windows 2003 R2 or Windows 2008 R2.

Technical Reference

B

This section contains technical information to help you customize PlateSpin Recon.

- ♦ [Section B.1, “Features,” on page 187](#)
- ♦ [Section B.2, “How PlateSpin Recon Collects Data,” on page 187](#)
- ♦ [Section B.3, “Data Collection,” on page 190](#)
- ♦ [Section B.4, “Charts and Formulas,” on page 191](#)
- ♦ [Section B.5, “Calculations,” on page 198](#)
- ♦ [Section B.6, “CPU Normalization,” on page 201](#)
- ♦ [Section B.7, “PlateSpin Recon VM CPU Inventory,” on page 202](#)
- ♦ [Section B.8, “Inventory and Monitoring Commands for Linux, Solaris and AIX,” on page 203](#)
- ♦ [Section B.9, “Sudo Configuration to Inventory and Monitor Linux Machines,” on page 207](#)
- ♦ [Section B.10, “Managing Credentials,” on page 207](#)

This section is for system administrators and virtualization architects who require advanced knowledge of PlateSpin Recon.

It is assumed that readers of this document are familiar with:

- ♦ Administration of Windows and Linux operating systems
- ♦ Administration of VMware ESX, VMware Server, and Microsoft Virtual servers

B.1 Features

- ♦ Hardware and software asset inventory
- ♦ Secure on-site data collection and analysis
- ♦ Windows and Linux support
- ♦ Management reports
- ♦ Integration with Microsoft Operations Manager 2005 and the remote PlateSpin Recon server
- ♦ Advanced project and scenario modeling
- ♦ Integration with the PlateSpin Portability Suite

B.2 How PlateSpin Recon Collects Data

PlateSpin Recon has three sequential stages to its data collection.

- ♦ [Section B.2.1, “Discovery,” on page 188](#)
- ♦ [Section B.2.2, “Inventory,” on page 188](#)
- ♦ [Section B.2.3, “Monitoring,” on page 189](#)

B.2.1 Discovery

Domain Discovery: PlateSpin Recon uses Windows Active Directory via LDAP to scan the network for a list of the machines on the specified domain. By default, this includes only online machines, but there is an option to include offline machines as well.

An Organizational Unit (OU) filter can also be specified, narrowing the area of the domain that PlateSpin Recon will poll during discovery. An Organization Unit is a container within a domain where computers can reside for segmentation. For example, if your domain has OU containers setup for each department, you can tell PlateSpin Recon to just look for machines within a specific department within the domain.

PlateSpin Recon only uses OU filters during discovery. Machines discovered in this way are unaffected during inventory and monitoring should machines be moved out of their previous OU containers. For more information on Organization Units and if they are in use in your domains, check with your System Administrator.

Subnet, IP Range Scan: For each machine in the subnet or IP range, PlateSpin Recon pings the machine. If it replies, it is considered a discovered machine.

Another option is to port scan through TCP, UDP or both. PlateSpin Recon tries to connect to ports and records which ports are being used. This option must be used with caution because network security might consider this an attack.

B.2.2 Inventory

- ♦ [“Linux, Solaris, and ESX 2.x” on page 188](#)
- ♦ [“ESX 3.x and Virtual Center” on page 188](#)
- ♦ [“Microsoft Windows Inventory” on page 188](#)

Linux, Solaris, and ESX 2.x

- ♦ PlateSpin Recon sends the `getplatform` script, which returns the architecture and `glibc` version of the machine being inventoried.
- ♦ Based on `getplatform`, PlateSpin Recon uses the SCP protocol to transfer a platform-specific inventory binary and libraries to the `/tmp` directory of the machine being inventoried.
- ♦ Over `ssh`, PlateSpin Recon executes the binary, streaming the command file over `stdin`.
- ♦ Logs and progress files are streamed back from the inventoried machine to the PlateSpin Recon Server using `stderr` while the machine XML is streamed over `stdout`.

ESX 3.x and Virtual Center

- ♦ PlateSpin Recon runs the executable locally on the PlateSpin Recon Server.
- ♦ The executable accesses ESX 3 or Virtual Center Web services, which provide the necessary inventory data.

Microsoft Windows Inventory

- ♦ If you are inventorying a Windows machine, you must make sure that WMI is installed and running on the machine. For Windows NT, you must manually install WMI components. For more information on downloading and installing WMI on Windows NT, see [Windows](#)

Management Instrumentation (WMI) CORE 1.5 (Windows NT 4.0) (<http://www.microsoft.com/downloads/en/details.aspx?displaylang=en&FamilyID=C174CFB1-EF67-471D-9277-4C2B1014A31E>) at the Microsoft Download Center.

To establish connection with the target Windows machine and to run the Inventory executable on the target machine, PlateSpin Recon by default uses WMI. If WMI fails, PlateSpin Recon uses the Remote Service as a failover. You can configure PlateSpin Recon to always use the Remote Service instead of WMI of the target machine:

1. In PlateSpin Recon Client, click the *Tools* menu.
2. Press the Ctrl key and click *Options*.
The *Server* page of the Options dialog box is displayed by default.
3. In the *Inventory* category, click the plus sign (+) next to *Advanced*.
4. Change the value of *Install Remote Service* to *True*.
5. Click *OK*.

PlateSpin Recon copies the inventory executable to the ADMIN\$ share on the target machine.

B.2.3 Monitoring

To understand the PlateSpin Recon Monitoring process, review the following sections:

- ♦ “Linux, Solaris, AIX, and ESX 2.x” on page 189
- ♦ “ESX 3.x and Virtual Center” on page 189
- ♦ “Microsoft Windows” on page 189

Linux, Solaris, AIX, and ESX 2.x

- ♦ PlateSpin Recon sends a script (`lininfo.sh`, `lininfo.sh`, `solinfo.sh`, `aixinfo.sh`, or `esxinfo.sh`) to the machine being inventoried.
- ♦ The script is run through `ssh`.
- ♦ The `ssh` server must be enabled for monitoring to function.
- ♦ Logs are streamed back to the PlateSpin Recon Server over `stderr`.
- ♦ Performance data is streamed back over `stdout`.

ESX 3.x and Virtual Center

- ♦ PlateSpin Recon calls ESX 3 or Virtual Center Web services, which provide the necessary performance data.

Microsoft Windows

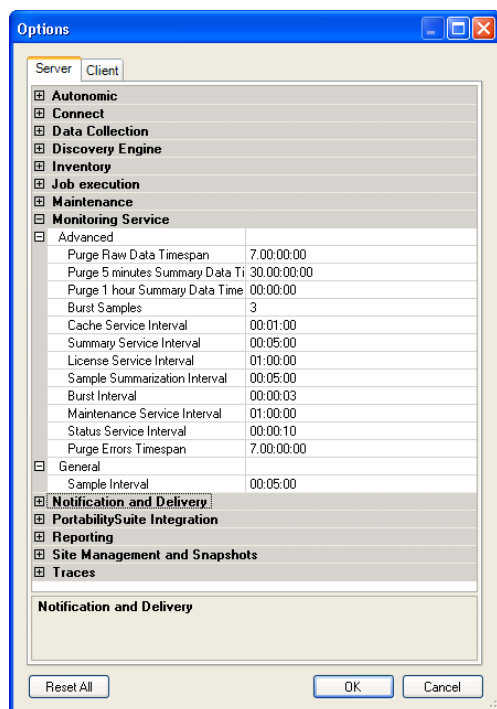
- ♦ PlateSpin Recon uses the Windows Performance Counter API to retrieve performance data. It does not use WMI.
- ♦ The Remote Registry service must be enabled for Windows monitoring to function.

B.3 Data Collection

By default, every 5 minutes, PlateSpin Recon takes three samples, three seconds apart, from all counters. The three samples are averaged together to produce the five-minute data that is stored in the database.

These values can be modified by clicking the *Tools* menu, then holding down the Ctrl key while selecting *Options* from the menu. On the *Server* tab, expand the Monitoring Service section:


Figure B-1 Options Window



Burst interval: The interval between different burst samples

Burst samples: The number of samples collected at a burst interval before an average is calculated

Sample interval: The interval between different samples



B.4 Charts and Formulas

- ♦ [Section B.4.1, “Chart Data,” on page 191](#)
- ♦ [Section B.4.2, “Utilization,” on page 195](#)
- ♦ [Section B.4.3, “Workload,” on page 196](#)

B.4.1 Chart Data

The five-minute data collected as described in [Section B.3, “Data Collection,” on page 190](#) is stored in the database for thirty days. The one-hour data is a summary of all the five-minute data collected during a particular hour. It is stored permanently in the database. All other chart data is calculated from the five-minute and one-hour data.

The twenty-four hour profile chart provides a value for each hour in a twenty-four hour period. This value is the largest average value calculated for the particular hour over the selected date range. This produces a utilization profile for the hourCounter.

Counter values are named for the Windows Performance Monitor counters used to retrieve them from Windows systems. The name consists of the Performance Monitor category joined by an underscore to the counter name. For example, Memory/PagesPerSecond becomes Memory_PagesPerSecond.

For Windows systems, these values are taken directly from Performance Monitor, but for other platforms, scripts are used to calculate the values.

- ♦ [“Platform Counters” on page 191](#)
- ♦ [“Collecting Counters for ESX Machines Monitored through the Virtual Center Server for Consolidated Planning” on page 194](#)
- ♦ [“Prerequisite for Collecting the XEN Counters for VMs Running on SLES with the Xen Kernel” on page 194](#)
- ♦ [“Prerequisite for Collecting Counters for Citrix Xen Hosts” on page 195](#)

Platform Counters

The following table lists the counters collected by PlateSpin Recon for the supported platforms:

Table B-1 Platform Counters

Name	Platforms
LogicalDisk_FreeMegabytes	All platforms except for ESX 2.5.x, NetWare, Windows 2000, and Windows NT
LogicalDisk_UsedMegabytes	All platforms except for ESX 2.5.x and NetWare
LogicalDisk_PercentUsedSpace	All platforms except for ESX 2.5.x and NetWare
LogicalDisk_PercentFreeSpace	All platforms except for ESX 2.5.x, NetWare, Windows 2000, and Windows NT
Memory_AvailableBytes	Windows, Windows NT, Solaris, Linux, ESX 2.5.x, ESX 3.x, ESX 4, Citrix XenServer

Name	Platforms
Memory_PagesPerSecond	Windows, Windows NT
Memory_PercentAvailableMemory	ESX 3.x, ESX 4
NetworkInterface_BytesTotalPerSecond	Windows, Solaris, Linux, ESX 2.5.x, ESX 3.x, ESX 4, Citrix XenServer
NetworkInterface_PacketsPerSecond	Windows, Solaris, Linux, ESX 2.5.x, ESX 3.x, ESX 4
PhysicalDisk_CurrentDiskQueueLength	All platforms except for ESX 3.x, ESX 4.0, and NetWare
PhysicalDisk_DiskBytesPerSecond	Windows, Windows NT, Solaris, Linux, ESX 2.5.x, ESX 3.x, ESX 4, Citrix XenServer
PhysicalDisk_DiskTransfersPerSecond	Windows, Windows NT, Solaris, Linux, ESX 2.5.x
PhysicalDisk_DiskWriteBytesPerSecond	Windows
PhysicalDisk_PercentIdleTime	All platforms except for ESX 2.x, ESX 3.0, and NetWare
PhysicalDisk_TotalSpaceGB	All platforms except for ESX 2.x and NetWare
PhysicalDisk_UsedSpaceGB	All platforms except for ESX 2.x and NetWare
PhysicalDisk_FreeSpaceGB	All platforms except for ESX 2.x and NetWare
PhysicalDisk_PercentUsedSpace	All platforms except for ESX 2.x and NetWare
PhysicalDisk_PercentFreeSpace	All platforms except for ESX 2.x and NetWare
PhysicalDisk_PartitionedSpaceGB	All platforms except for ESX 2.x and NetWare
PhysicalDisk_UnpartitionedSpaceGB	All platforms except for ESX 2.x and NetWare
PhysicalDisk_PercentPartitionedSpace	All platforms except for ESX 2.x and NetWare
PhysicalDisk_PercentUnpartitionedSpace	All platforms except for ESX 2.x and NetWare
Processor_PercentProcessorTime	Windows, Windows NT, Solaris, Linux, ESX 2.5.x, ESX 3.x, ESX 4, Citrix XenServer
System_ProcessorQueueLength	Windows, Windows NT, Solaris, Linux, ESX 2.5.x
VC_Processor	All VMs, ESXs, Resource Pool, and clusters monitored through VMware vCenter.
VC_ProcessorTime	All VMs and ESXs monitored through VMware vCenter.
VC_ProcessorPercentReady	All VMs monitored through VMware vCenter
VC_ProcessorPercentWait	All VMs monitored through VMware vCenter
VC_MemoryPercentUsed	All VMs, ESXs, Resource Pool, and clusters monitored through VMware vCenter.
VC_MemoryAvailableBytes	All VMs, ESXs, and Resource Pool monitored through VMware vCenter.

Name	Platforms
VC_NetBytesTotalPerSecond	All VMs and ESXs monitored through VMware vCenter.
VC_NetPacketsPerSecond	All VMs and ESXs monitored through VMware vCenter.
VC_DiskBytesPerSecond	All VMs and ESXs monitored through VMware vCenter.
VC_MemoryBalloon	All VMs, ESXs, Resource Pool, and clusters monitored through VMware vCenter.
VC_MemorySwapUsed	All ESXs and clusters monitored through VMware vCenter.
VC_EffectiveMemory	All clusters enabled with DRS and monitored through VMware vCenter.
VC_EffectiveCPU	All clusters enabled with DRS and monitored through VMware vCenter.
VC_DiskReadBytesPerSecond	All VMs and ESXs monitored through VMware vCenter.
VC_DiskWriteBytesPerSecond	All VMs and ESXs monitored through VMware vCenter
VC_DiskReadCountPerSecond	All VMs and ESXs monitored through VMware vCenter
VC_DiskWriteCountPerSecond	All VMs and ESXs monitored through VMware vCenter
VC_DiskBytesPerRead	All VMs and ESXs monitored through VMware vCenter.
VC_DiskBytesPerWrite	All VMs and ESXs monitored through VMware vCenter.
VC_DiskPercentIdleTime	ESX 3.5 and later monitored through VMware vCenter except vCenter 2.0.
VC_DiskTotalSpaceGB	All ESXs monitored through VMware vCenter.
VC_DiskUsedSpaceGB	All ESXs monitored through VMware vCenter.
VC_DiskFreeSpaceGB	All ESXs monitored through VMware vCenter.
VC_DiskPercentUsedSpace	All ESXs monitored through VMware vCenter.
VC_DiskPercentFreeSpace	All ESXs monitored through VMware vCenter.
VC_DiskPartitionedSpaceGB	All ESXs monitored through VMware vCenter.
VC_DiskUnpartitionedSpaceGB	All ESXs monitored through VMware vCenter.
VC_DiskPercentPartitionedSpace	All ESXs monitored through VMware vCenter.
VC_DiskPercentUnpartitionedSpace	All ESXs monitored through VMware vCenter.
VC_LogicalDiskUsedMegabytes	All ESXs monitored through VMware vCenter.

Name	Platforms
VC_LogicalDiskFreeMegabytes	All ESXs monitored through VMware vCenter.
VC_LogicalDiskPercentUsedSpace	All ESXs monitored through VMware vCenter.
VC_LogicalDiskPercentFreeSpace	All ESXs monitored through VMware vCenter.
Xen_Memory_AvailableBytes	Xen on Novell SLES 10 and SLES11, Citrix XenServer
Xen_Processor_PercentProcessorTime	Xen on Novell SLES 10 and SLES 11, Citrix XenServer
Xen_NetworkInterface_ByteTotalPerSecond	Citrix XenServer
Xen_PhysicalDisk_DiskBytesPerSecond	Citrix XenServer

NOTE: If any of the Virtual Center counters listed in [Table B-1 on page 191](#) are not available on the Virtual Center, PlateSpin Recon does not collect data for such counters and will not display any error message for it.

Collecting Counters for ESX Machines Monitored through the Virtual Center Server for Consolidated Planning

ESX machines are inventoried and monitored through the VC Server. However, the data does not contain all counters of the ESX machine required for consolidated planning because of the limitations of the VC Server. You need to directly collect additional information from the ESX machine.

- 1 In the Data Center Explorer, navigate to *Filters > VM Servers > VMware ESX Servers*.
- 2 Right-click the ESX machine whose counters you want to collect, then click *Stop Monitoring*.
- 3 Right-click the ESX machine again, then click *Attach Credentials*.
The Attach Credentials dialog box is displayed.
- 4 Specify the credentials for the machine, then click *OK*.
- 5 Right-click the ESX machine, click *Refresh Inventory*, then click *Now*.
- 6 Right-click the ESX machine, then click *Start Monitoring*.

Prerequisite for Collecting the XEN Counters for VMs Running on SLES with the Xen Kernel

If you want to collect the Xen counters for VMs running on SLES with the Xen kernel, do the following on the SLES host machine before you start monitoring it:

- 1 Edit `/etc/xen/xend-config.sxp` to add the following line:
`(xen-api-server ((9363 none)(unix none)))`
- 2 Restart the Xen daemon by entering the following command:
`/etc/init.d/xend restart`

Prerequisite for Collecting Counters for Citrix Xen Hosts

To collect counters for the Citrix Xen hosts, you must attach the credential to all the hosts.

- 1 Ensure that SSH is enabled on port 22 on the Citrix Xen hosts.
- 2 In the Data Center Explorer, right-click a Citrix Xen host, then click *Stop Monitoring*.
- 3 Right-click the Citrix Xen host again, then click *Attach Credentials*.
The Attach Credentials dialog box is displayed.
- 4 Specify the credentials for the machine, then click *OK*.
- 5 Right-click the Citrix Xen host, then click *Start Monitoring*.
- 6 Repeat [Step 2](#) through [Step 5](#) for all the Citrix Xen hosts.

B.4.2 Utilization

Utilization formulas represent the percentage of the available resources being used by a given workload. The calculations are based on counter data.

Table B-2 Utilization Formulas

Name	Formula	Platforms
Disk IO (%)	MAX(0,(100-PERF("PhysicalDisk_PercentIdleTime")))	All platforms except for NetWare, ESX 2.5.x, ESX 3.0.
Disk Space Free (%)	PERF("PhysicalDisk_PercentFreeSpace")	All platforms except for NetWare, ESX 2.5.x
Disk Space Used (%)	PERF("PhysicalDisk_PercentUsedSpace")	All platforms except for NetWare, ESX 2.5.x
Memory Free (%)	100*(PERF("Memory_AvailableBytes"))/INV("TotalMemory")	Windows, Solaris, Linux, ESX 2.5.x, ESX 3.x, ESX 4, Citrix XenServer
Memory Used (%)	100*(INV("TotalMemory")-PERF("Memory_AvailableBytes"))/INV("TotalMemory")	Windows, Solaris, Linux, ESX 2.5.x, ESX 3.x, ESX 4, Citrix XenServer
Processor (%)	PERF("Processor_PercentProcessorTime")	Windows, Solaris, Linux, ESX 2.5.x, ESX 3.x, ESX 4, Citrix XenServer
Disk Space Partitioned (%)	PERF("PhysicalDisk_PercentPartitionedSpace")	All platforms except for NetWare, ESX 2.5.x
Disk Space Unpartitioned (%)	PERF("PhysicalDisk_PercentUnpartitionedSpace")	All platforms except for NetWare, ESX 2.5.x
Logical Disk Used Space (%)	PERF("LogicalDisk_PercentUsedSpace")	All platforms except for NetWare, ESX 2.5.x

Name	Formula	Platforms
Logical Disk Free Space (%)	PERF("LogicalDisk_PercentFreeSpace")	All platforms except for NetWare, ESX 2.5.x, Windows NT, and Windows 2000
VC_Memory Free (%)	100*(PERF("VC_MemoryAvailableBytes"))/INV("TotalMemory")	All
VC_Memory Used(%)	100*(INV("TotalMemory")-PERF("VC_MemoryAvailableBytes"))/INV("TotalMemory")	All
VC_ProcessorTime (Hour)	PERF("VC_ProcessorTime")	All
VC_Disk IO (%)	MAX(0,(100-PERF("VC_DiskPercentIdleTime")))	ESX 3.5 or later monitored through vCenter 2.5 or later
VC_Disk Space Free (%)	PERF("VC_DiskPercentFreeSpace")	All ESXs monitored through vCenter
VC_Disk Space Partitioned (%)	PERF("VC_DiskPercentPartitionedSpace")	All ESXs monitored through vCenter
VC_Disk Space Unpartitioned (%)	PERF("VC_DiskPercentUnpartitionedSpace")	All ESXs monitored through vCenter
VC_Disk Space Used (%)	PERF("VC_DiskPercentUsedSpace")	All ESXs monitored through vCenter
VC_Logical Disk Used Space (%)	PERF("VC_LogicalDisk_PercentUsedSpace")	All ESXs monitored through vCenter
VC_Logical Disk Free Space (%)	PERF("LogicalDisk_PercentFreeSpace")	All ESXs monitored through vCenter

NOTE: Utilization expressions are not available for Virtual Center clusters.

B.4.3 Workload

Workload formulas measure workload size; that is, the amount of processing power and raw resources being used for a workload, independent of the total hardware and resources available.

- ♦ PERF: Dynamic performance data
- ♦ INV: Static inventory data

Table B-3 Workload Formulas

Name	Formula	Platforms
Disk (MB/sec)	(PERF("PhysicalDisk_DiskBytesPerSecond"))/(1024*1024)	All
Disk (Transfers/sec)	PERF("PhysicalDisk_DiskTransfersPerSecond")	Windows, Windows NT, Solaris, Linux, ESX 2.5.x

Name	Formula	Platforms
Disk Queue Length	PERF("PhysicalDisk_CurrentDiskQueueLength")	Windows, Windows NT, Solaris, Linux, ESX 2.5.x, Citrix XenServer
Disk Space Free (GB)	PERF("PhysicalDisk_FreeSpaceGB")	All platforms except for NetWare and ESX 2.5.x
Disk Space Used (GB)	PERF("PhysicalDisk_UsedSpaceGB")	All platforms except for NetWare and ESX 2.5.x
Disk Writes (MB/sec)	(PERF("PhysicalDisk_DiskWriteBytesPerSecond"))/(1024*1024)	Windows
Logical Disk Used Space (MB)	PERF("LogicalDisk_UsedMegabytes")	All platforms except for NetWare, ESX 2.5.x
Logical Disk Free Space (MB)	PERF("LogicalDisk_FreeMegabytes")	All platforms except for NetWare, ESX 2.5.x, Windows NT, and Windows 2000
Memory (Pages/sec)	PERF("Memory_PagesPerSecond")	Windows, Windows NT
Memory Free (MB)	(PERF("Memory_AvailableBytes"))/(1024*1024)	All
Memory Used (MB)	(INV("TotalMemory")-PERF("Memory_AvailableBytes"))/(1024*1024)	All
Network (MB/sec)	(PERF("NetworkInterface_BytesTotalPerSecond"))/(1024*1024)	Windows, Solaris, Linux, ESX 2.5.x, ESX 3.x, Citrix XenServer
Network (Packets/sec)	PERF("NetworkInterface_PacketsPerSecond")	Windows, Solaris, Linux, ESX 2.5.x, ESX 3.x, ESX 4
Processor (MHz)	PERF("Processor_PercentProcessorTime")*INV("ProcessorCapacity")/100	Windows, Windows NT, Linux, ESX 2.5.x, ESX 3.x, ESX 4, Citrix XenServer, Solaris
Processor Queue Length	PERF("System_ProcessorQueueLength")	Windows, Windows NT, Solaris, Linux, ESX 2.5.x
VC_Disk (MB/sec)	(PERF("PhysicalDisk_DiskBytesPerSecond"))/(1024*1024)	All
VC_MemoryFree (MB)	(PERF("Memory_AvailableBytes"))/(1024*1024)	All
VC_MemoryUsed (MB)	(INV("TotalMemory")-PERF("Memory_AvailableBytes"))/(1024*1024)	All
VC_Network (MB/sec)	(PERF("VC_NetBytesTotalPerSecond"))/(1024*1024)	All
VC_Network (Packets/sec)	PERF("VC_NetPacketsPerSecond")	All

Name	Formula	Platforms
VC_Processor (MHz)	PERF("VC_Processor")	All
VC_Disk Space Partitioned (GB)	PERF("VC_DiskPartitionedSpaceGB")	ESX monitored through vCenter
VC_Disk Space Used(GB)	PERF("VC_DiskUsedSpaceGB")	ESX monitored through vCenter
Disk Space Partitioned (GB)	PERF("PhysicalDisk_PartitionedSpaceGB")	All platforms except for NetWare and ESX 2.5.x
VC_Total Disk (GB)	PERF("VC_DiskTotalSpaceGB")	All ESXs monitored through vCenter
Total Processor (Count)	INV("ProcessorCount")	
VC_Processor (Normalized)	PERF("VC_Processor")/INV("ProcessorCapacity")*INV("ProcessorCapacityInSpecUnits")	
Total Memory (MB)	INV("TotalMemory")/(1024*1024)	
VC_Disk Space Unpartitioned (GB)	PERF("VC_DiskUnpartitionedSpaceGB")	All ESXs monitored through vCenter
Disk Space Unpartitioned (GB)	PERF("PhysicalDisk_UnpartitionedSpaceGB")	All platforms except for NetWare and ESX 2.5.x
Processor (Normalized)	PERF("Processor_PercentProcessorTime")*INV("ProcessorCapacityInSpecUnits")/100	
Total Disk (GB)	PERF("PhysicalDisk_TotalSpaceGB")	All platforms except for NetWare and ESX 2.5.x
VC_Disk Space Free (GB)	PERF("VC_DiskFreeSpaceGB")	All ESXs monitored through vCenter
VC_Logical Disk Used Space (MB)	PERF("VC_LogicalDiskUsedMegabytes")	All ESXs monitored through vCenter
VC_Logical Disk Free Space (MB)	PERF("VC_LogicalDiskFreeMegabytes")	All ESXs monitored through vCenter

NOTE: Workload expressions are not available for Virtual Center clusters.

B.5 Calculations

- ♦ [Section B.5.1, "24-Hour Profile," on page 199](#)
- ♦ [Section B.5.2, "Dual-Core and Multi-Core Machines," on page 200](#)
- ♦ [Section B.5.3, "Hyper-Threading," on page 200](#)
- ♦ [Section B.5.4, "Effective Speed," on page 200](#)

- ♦ [Section B.5.5, “Power and Cooling,” on page 201](#)
- ♦ [Section B.5.6, “CO2 Emission,” on page 201](#)

B.5.1 24-Hour Profile

The 24-hour profile consists of 24 values that correspond to each hour of the day (0 to 23). Each value is calculated by taking the peak value for that specific hour over the number of monitoring days.

Example

Consider a system monitored for 3 days. The following chart presents the data set for its Processor workload, and the calculated 24-hour profile:

Table B-4 *Processor Workload Data Set*

Hour	Day 1	Day 2	Day 3	24-Hour Profile
0	200	300	500	500
1	300	300	250	300
2	300	400	450	450
3	700	250	350	750
4	500	500	500	500
5	600	400	500	600
6	500	550	750	750
7	400	600	800	800
8	500	300	250	500
9	450	350	640	640
10	200	240	340	340
11	300	600	+60	650
12	300	-	-	300
13	550	340	850	850
14	440	500	600	600
15	800	540	600	800
16	750	260	430	750
17	500	640	450	640
18	600	200	350	600
19	300	250	650	650
20	200	300	200	300

Hour	Day 1	Day 2	Day 3	24-Hour Profile
21	250	350	400	400
22	600	250	350	600
23	500	650	750	750

B.5.2 Dual-Core and Multi-Core Machines

PlateSpin Recon considers and treats dual-core or multi-core machine as two-processor or x-processor machines.

B.5.3 Hyper-Threading

The effect of hyper-threading is controlled by two configuration parameters: Aggregate Instance Function (SUM, AVG) and Hyper-threading Normalization Factor (1 – 2). The default value for aggregate instance function is AVG, and the normalization factor is set to 1.25. These values can be configured in the advanced settings under *Tools > Options*.

PlateSpin Recon collects each individual processor instance's utilization values at a given time and then calculates the processor utilization based on the aggregate function. If the aggregate function is SUM, then the physical processor utilization is the sum of all logical processor instances' utilization. If the aggregate function is AVG, then the physical processor utilization would be average of all logical processor instances' utilization.

Example

Consider a system with single-core processor with hyper-threading enabled at p MHz speed. For a given period of time, PlateSpin Recon calculates the percentile utilization values for 24H profile values of instance X and Y as follows.

$\{X\} = \{X1, X2, Xi, \dots, X24\}$, Xi is the utilization of the processor instance X at hour i.

$\{Y\} = \{Y1, Y2, Yi, \dots, Y24\}$, Yi is the utilization of the processor instance Y at hour i.

Table B-5 Aggregate Function Calculations

(Aggregate Function, Normalization Factor)	(SUM, f)	(AVG, f)
Reports: Peak %	$\text{Max } \{ X + Y \}$	$\text{Max } \{ (X + Y) / 2 \}$
Reports: Peak (MHz)	$\text{Max } \{ X + Y \} * (1/100) * p * f$	$\text{Max } \{ (X+Y)/2 \} * (1/100) * p * f$
CPM: Workload calculation	$\{X+Y\} * (1/100) * p * f$	$\{(X+Y)/2\} * (1/100) * p * f$

B.5.4 Effective Speed

PlateSpin Recon takes both number of cores and hyper-threading into account when calculating the effective speed of a machine.

Example:

Server template specification: 2 x 3000 GHz, dual core, and hyper-threading enabled

Number of processors: $P = 2 * 2 = 4$

Processor capacity: $C = 3000$ GHz

Hyper-threading normalization factor: $n = 1.25$ (default value)

Scaling factor: $A(p) = P / \{ 1 + \alpha (P-1) \}$, where $\alpha = 0.05$

$A(p) = 4 / \{ 1 + 0.05 * 3 \} = 3.4782$

Effective speed = $A(p) * C * n = 3.4782 * 3000 * 1.25 = 13043$ MHz

B.5.5 Power and Cooling

Energy Use (kWh/yr) = Power Usage (W) x Server Derating Factor x Number of hours in a year (8760) / 1000

Heat Dissipation (BTU/yr) = Energy Use (kWh/yr) x 3.41

Annual Energy Cost = Energy Cost per kWh x Energy Use (kWh/yr)

B.5.6 CO2 Emission

Energy Use (kWh/yr) = Power Usage (W) x Server Derating Factor x Number of hours in a year (8760) / 1000

CO2 Emission (lb/yr) = Energy Use (kWh/yr) x 1.54

The constant factor of 1.54 has been considered as per the standards.

B.6 CPU Normalization

CPU normalization allows the actual speed of a CPU to be considered for operations such as consolidation planning, instead of relying on the CPU clock frequency, which can be misleading. Different CPU manufacturers and models provide a wide array of performance results for CPUs with the same clock frequency. For example, there is a considerable difference between a Core 2 and a Pentium IV CPU with a clock frequency of 2.0 GHz. CPU normalization uses the SPEC data industry standard benchmarks to give real-world performance values. This is an important factor in realistic consolidation planning. See [Section B.6.1, “SPEC Data,” on page 202](#) for more information.

In consolidation planning, when a scenario is created, the user is given an option to select between clock frequency and normalized speed as the unit of measurement for processor utilization and capacity. The consolidation planning algorithms use the appropriate unit of measurement when assigning workloads to containers, based on these selections. For more information see [Section 6.5, “Working with Consolidation Scenarios,” on page 134](#).

NOTE: The data required to perform consolidation planning based on normalized speed is actually produced when a project is created. Projects from a pre-3.5 version of PlateSpin Recon, even if upgraded, cannot use normalized speed in consolidation planning (the option is disabled in the scenario creation wizard). To use normalization speeds with these projects, they must be re-created from scratch.

- ♦ [Section B.6.1, “SPEC Data,” on page 202](#)
- ♦ [Section B.6.2, “Normalization Speed,” on page 202](#)

B.6.1 SPEC Data

PlateSpin Recon CPU normalized speed makes use of SPECint and SPECfp (98, 2000, 2006) CPU benchmarking data. The PlateSpin Recon CPU Normalization rating for specific CPUs are an average of data available from <http://www.spec.org> (<http://www.spec.org>) for those CPUs as of June 2008. For the latest SPECf benchmark results, visit <http://www.spec.org> (<http://www.spec.org>).

B.6.2 Normalization Speed

You can change default normalization speed for inventoried machines from the Custom Fields dialog box. You can also view normalization speed in server templates. Although you cannot directly change these values, you can change other values and see how the changes affect the normalization speed. For more information, see [Section 6.3, “Creating and Editing Server Templates,” on page 126](#).

To view or edit normalization speed for an inventoried machine:

- 1 In the Data Center Explorer, right-click a machine and select *Properties*.
The Properties dialog box for that machine is displayed.
- 2 Click the *Custom Fields* tab.
In the top section you can see the default *Normalized Speed* for the machine.
- 3 Click the *Normalized Speed* value and edit as desired.

IMPORTANT: The new value that you specify is automatically changed to the default value when you refresh inventory of servers. You must manually reset the values after you refresh inventory of the servers.

To reset the value to its default, click *Reset Reference Values*.

- 4 Click *OK*.

CPU Normalization values can also be used as column values in relevant reports, although the columns are not shown by default. To display normalization-oriented columns, see [Section 5.3, “Editing Report Templates,” on page 97](#) and [Section 5.4.1, “Creating and Editing Report Views,” on page 106](#).

B.7 PlateSpin Recon VM CPU Inventory

When PlateSpin Recon inventories Virtual Machines with one or more CPUs allocated to them, the cores and CPUs are displayed as follows:

Cores: Regardless of the VM guest machine configuration, PlateSpin Recon shows the actual number of cores from the hardware host. This may contradict with third-party tools, such as CPUID.

CPUs: On ESX hosts, there are three possible virtual CPU configurations (1,2 and 4) for VMs. The PlateSpin Recon inventory of the CPUs depends on the VM configuration and the cores on the host machine. ESX allocates as many cores as needed to accommodate the user requirement. PlateSpin Recon reports the minimum number of CPUs required to provide the indicated number of virtual CPUs.

The following table shows some example configurations and what PlateSpin Recon reports for each.

Table B-6 *Sample Virtual CPU Configurations and What PlateSpin Recon Reports*

Host Cores	Guest Virtual CPUs	PlateSpin Recon Reported Cores	PlateSpin Recon Reported CPUs
1	1	1	1
	2	1	2
	4	1	4
2	1	2	1
	2	2	1
	4	2	2
4	1	4	1
	2	4	1
	4	4	1

B.8 Inventory and Monitoring Commands for Linux, Solaris and AIX

This section contains lists of commands used by PlateSpin Recon during inventory and monitoring.

- ♦ [Section B.8.1, “Inventory Commands for Solaris,” on page 203](#)
- ♦ [Section B.8.2, “Inventory Commands for Linux,” on page 204](#)
- ♦ [Section B.8.3, “Inventory Commands for AIX,” on page 205](#)
- ♦ [Section B.8.4, “Monitoring Commands for Linux and Solaris,” on page 206](#)
- ♦ [Section B.8.5, “Monitoring Commands for AIX,” on page 206](#)
- ♦ [Section B.8.6, “Text Processing,” on page 207](#)

B.8.1 Inventory Commands for Solaris

Table B-7 *Inventory Commands for Solaris*

Command	Description
kstat	Get network connection speed.

Command	Description
chmod	Change file access permissions for the inventory script that is placed on the machine being inventoried.
rm	Remove the inventory script that was placed on the machine.
prtdiag	Get CPU cache information and machine model for Sparc.
psrinfo	Get processor information such as family, model, step, etc., which are related to x86 mainly.
isainfo	Get CPU instruction set information such as sse, sse2, mmx, etc., which are related to x86 mainly.
prtconf	Get configured device information, which is related to L2 cache information of x86, network adapter resolving, and ACPI support.
ifconfig	Get network related information such as IP address, DHCP, net mask, etc.
ndd	Get network connection speed.
uname	Determine kernel version.
mount	Get boot loader.
pkginfo	Get installed programs.
ps	Get process information.
who	Get run level.
hostname	Get host name.

B.8.2 Inventory Commands for Linux

Table B-8 *Inventory Commands for Linux*

Command	Description
chkconfig	Get service(daemon) information.
df	Get size information for the mounted volumes
mount	Get volume information.
e2label	Get disk label.
route	Get gateway.
ifconfig	Get network related information such as IP address, net address, and so on.
egrep	Get DHCP information.
getcfg	Get DHCP for SUSE.
rpm	Get installed programs.

Command	Description
lspci	Get PCI device information such as vendor name, device model in human readable format.
iwconfig	Get wireless network adapter NIC list.
iwlist	Get wireless bit rate.
ethtool	Get speed and link status for network connection.
dmesg	Get real MAC address.
kudzu	Get network device information for Redhat Linux.
hwinfo	Get network adapter information for SUSE Linux.
fdisk	Determine if a partition is active or not.
runlevel	Get current run level.
hostname	Get host name and DNS domain name.

B.8.3 Inventory Commands for AIX

Table B-9 *Inventory Commands for AIX*

Command	Description
rpm	Get a list of installed programs.
ps	Get service (daemon) information.
mount	Get list of mounted volumes.
df	Get size information for mounted volumes.
lslv	Get SerialNumber and Label for mounted volumes.
hostname	Get hostname.
namerslv	Get DNS domain name.
getconf	Determine whether a 64-bit kernel.
uname	Determine OS release version.
ifconfig	Get configured network adapters and their IP address.
netstat	Get MAC address.
entstat	Get speed of network adapter.
prtconf	Get processor information and default gateway.
lspv	Get list of disks and their partitions.
bootinfo	Get disk size.
lscfg	Get disk model, manufacturer and type and network adapter model.

Command	Description
lsitab	Get daemon information.
getconf	Get machine memory and bootloader information.
lsattr	Get machine serial number and model,hyperthreading details and subnet mask.

B.8.4 Monitoring Commands for Linux and Solaris

Table B-10 *Monitoring Commands for Linux and Solaris*

Command	Platform	Description
psrinfo	Solaris	count processor number.
mpstat	Solaris	Get processor time per cpu instance.
vmstat	Solaris	Get processor queue length and memory available.
ifconfig	Solaris, Linux	Get NIC list.
sed	Solaris, Linux	Edit input stream.
tail	Solaris, Linux	Edit input stream.
netstat	Solaris	Get network package transfer rate.
kstat	Solaris	Get NIC speed.
prtconf	Solaris	Get module and instance binding for NIC.
iostat	Solaris	Get disk io information such as queue and speed.
expr	Solaris, Linux	evaluate express.

B.8.5 Monitoring Commands for AIX

Table B-11 *AIX Monitoring Commands*

Command	Description
vmstat	Get Available Memory in Bytes
ifconfig	Get NIC information
entstat	Get Network Packets, Bytes Send/Received
mpstat	Get Percentage processor time
iostat	Get Disk transfer per second, disk idle, disk byte per second
sar	Get Disk Queue length, Percentage processor used
lspv	Get Disk partitioned or Unpartitioned space, disk list

Command	Description
bootinfo	Get disk size
df	Get the volume level information
oslevel	Get the version of AIX

B.8.6 Text Processing

grep, cat, cut, awk, ls

B.9 Sudo Configuration to Inventory and Monitor Linux Machines

PlateSpin Recon users can inventory and monitor the machines by using sudo user credentials only if they have the permission to run commands on the target machine. To grant permission for the users to perform inventory and monitoring:

- 1 Log in as root in to the Linux machine for which you want to grant permission to a sudo user.
- 2 Open the `sudoers` configuration file by executing the following command:

```
visudo
```

- 3 Append the following line to the file:

```
nonroot_username ALL=(root) NOPASSWD: ALL
```

In the command, replace `nonroot_username` with the sudo username.

IMPORTANT: Make sure that the password settings for a sudo user is not overridden by a higher group-level setting.

B.10 Managing Credentials

The workload credentials that you provide in the [Credential Manager](#) are saved in the PlateSpin Recon database. The password is encrypted by using the 1024-bit RSA encryption. The encryption keys are saved in a PlateSpin Recon DLL file. The keys are encrypted by using 128-bit AES (Rijndael) encryption and are stored as binary data in the DLL file. When you monitor and inventory a UNIX or Linux workload from the PlateSpin Recon Server, the password is sent over the SSH connection. For a Windows workload, PlateSpin Recon uses Windows authentication and the password is not sent across the network.

Schedule Options

C

Certain PlateSpin Recon jobs can be scheduled to occur at a later date, or on a recurring basis. These can be viewed through the Scheduled Job Explorer. Jobs can include inventorying, report generation, and snapshot export.

Following are the schedule options:

- ♦ [Section C.1, “Later,” on page 209](#)
- ♦ [Section C.2, “Recurrence,” on page 209](#)

C.1 Later

The Later scheduling option lets you to run the job at a specific date and time. By default, the current date and time are displayed.

Select Later to run the job at a specific date and time (either typed manually in mm/dd/yyyy hh:mm AM|PM format, or selected from the pop-up calendar)

To change the date, do any of the following:

- ♦ Click the drop-down list, then select the desired date in the calendar.
- ♦ Manually enter the date in the mm/dd/yy format.

To change the time, manually specify the time in the hh:mm AM or PM format.

C.2 Recurrence

Select *Recurrence* to run the job on a recurring basis.

- 1 From the drop-down list, select a previously entered recurrence schedule, or select *<New Recurrence>*.

The Create Recurrence dialog box is displayed.

Create Recurrence

Recurrence Name:

Recurrence pattern

☒ Daily
 ☐ Weekly
 ☐ Monthly
 ☐ Yearly

☒ Every 1 day(s)
 ☐ Every weekday

Range of recurrence

Start: 2007/03/15 12:11:30
 ☒ No end date
 ☐ End after: 1 occurrences
 ☐ End by: 3/15/2007

OK Cancel

2 Specify a unique name.

This name is used to refer to the recurrence schedule in the *Recurrence* drop-down list, so that it can be reused for other jobs.

3 In the *Recurrence pattern* area, select a rate of recurrence, and use options to specify when, and how often jobs are run.

Quick Reference to PlateSpin Recon Icons and Commands

D

This section provides a quick reference to the icons and commands used in PlateSpin Recon Client.











- ♦ [Section D.1, “Data Center Explorer,” on page 211](#)
- ♦ [Section D.2, “Reports Explorer,” on page 213](#)
- ♦ [Section D.3, “Projects Explorer,” on page 213](#)
- ♦ [Section D.4, “Raters Explorer,” on page 214](#)
- ♦ [Section D.5, “Chart Viewer,” on page 215](#)
- ♦ [Section D.6, “Job Explorer,” on page 216](#)
- ♦ [Section D.7, “Scheduled Job Explorer,” on page 216](#)
- ♦ [Section D.8, “Log Explorer,” on page 216](#)

D.1 Data Center Explorer

- ♦ [Section D.1.1, “Data Center Icons,” on page 211](#)
- ♦ [Section D.1.2, “Data Center Icon Status Indicators,” on page 211](#)
- ♦ [Section D.1.3, “Data Center Commands and Associated Icons,” on page 212](#)









D.1.1 Data Center Icons

Table D-1 Data Center Explorer Icons

	Physical Server		Virtual Center Folder
	Virtual Machine		Virtual Center Cluster
	Groups		Virtual Center Pool
	Virtual Center		Primary Site
	Virtual Center Data Center		Secondary Site











D.1.2 Data Center Icon Status Indicators



Table D-2 Data Center Explorer Icon Status Indicators

	Monitored		Virtual Machine Powered On
	Monitored virtual machine		Virtual Machine Powered Off
	Not monitored		Unresponsive/error
	Not inventoried virtual machine		Warning

D.1.3 Data Center Commands and Associated Icons

Table D-3 Data Center Explorer Commands and Associated Icons







Commands	Submenu	Description
 Inventory		Inventories the selected servers. The selection must be discovered but not yet inventoried. To inventory servers that have already been inventoried, use <i>Refresh Inventory</i> .
 Start Monitoring		Monitors the selected servers. The selection can be a node (<i>All</i> , <i>Filters</i> , <i>Groups</i> , or <i>Sites</i>) in the Data Center Explorer pane, or a particular Filter, Group, or Site to monitor all child servers.
 Stop Monitoring		Stops monitoring the selected servers. The selection can be a node (<i>All</i> , <i>Filters</i> , <i>Groups</i> , or <i>Sites</i>) in the Data Center Explorer pane, or a particular Filter, Group, or Site to stop monitoring all child servers.
 Logs	View Logs	View logs of selected servers.
	Acknowledge Logs	If there are some errors associated with a machine, an error icon is displayed on that machine. If the errors are acceptable, you can use <i>Acknowledge Logs</i> to remove the error icon from the display of the machine.
 Create Group		Create a new group of servers.
 Copy to Group	<group_name>	Copies the selected servers to the designated group.
 Rename Group		Renames the selected group.
 Create Template		Opens the New Server Template dialog box.
 Attach Credentials		Opens the Attach Credentials dialog box.
Launch	Terminal Services	Opens the selected application (if it is installed). Paths are configurable by selecting <i>Options</i> from the <i>Tools</i> menu, then clicking the <i>Client</i> tab.
	SSH	
	VMware Remote Console	
	VMware Virtual Infrastructure Client	
	VMware Web Access	
 Delete		Delete the server or group of servers.
Chargeback		Opens the Raters for <i>name</i> dialog box.
View XML		Displays the selected machine details as an XML document. Only the first machine is displayed if there are multiple selections.

Commands	Submenu	Description
 Properties		Opens a Properties dialog box for the selected server.
 Set Custom Fields		Opens a Properties dialog box (Custom Fields page) for the selected servers when multiple servers are selected.

To apply Data Center Explorer commands to servers, right-click a server and select the required command from the menu.

D.2 Reports Explorer

Table D-4 *Report Explorer Commands*





Commands	Tooltip/Description
	Create Report Template. Create a new report template based on the report node selected.
	Copy Report Template. Copy a report template.
	Edit Report Template. Edit a report by setting group, views, filters, summaries, and forecasts for a report template or a report.
	Generate Report. Generate a report based on the report template selected.
	View Report. View an existing (previously generated) report.
	Send Report. Send a generated report via e-mail or to an FTP server or network location.
	Delete. Delete the selected report or template.






D.3 Projects Explorer

- ♦ [Section D.3.1, “Project Explorer Commands,” on page 213](#)
- ♦ [Section D.3.2, “Project Explorer Icons,” on page 214](#)

D.3.1 Project Explorer Commands














Table D-5 *Project Explorer Commands*

Command	Action
	Create a new project.
	Copy a project
	View project settings
	Create a new scenario.

Command	Action
	Copy a scenario
	View report/View Project report/View Workload Assignment report.
	Delete project/scenario/workload, depending on the selection made.
	Implement by using the Portability Suite
	Edit Mapping





D.3.2 Project Explorer Icons



Table D-6 *Project Explorer Icons*

Icon	Description
	Consolidation Project
	Consolidation Scenarios
	Server
	Existing Target Servers
	Virtual Machine
	Workloads
	Virtualization Overhead
	Server Template is mapped to a Virtual Server in the database
	Server is selected for Consolidation
	Unconsolidated Server
	Server is selected for Protection
	Server is Unprotected
	Unconsolidated Workloads

D.4 Raters Explorer

















Table D-7 *Project Explorer Commands*

Command	Action
	Create a new rater.
	Edit a rater.
	Delete a rater.
	Attach a rater to a machine.

Command	Action
	Detach a rater from a machine.
	View the properties of a rater.






D.5 Chart Viewer

Table D-8 *Chart Viewer Commands*

Command	Description
	Print Preview. Preview chart before printing.
	Print. Print the chart.
	Copy To Clipboard. Copy the chart to the clipboard.
	Line. Display straight lines between points of continuous data.
	Curve. Display curved lines between points of continuous data.
	Candlestick. View the variability of sample data. This button appears only for the candle chart.
	Stacked Bar. View cumulative workload data from multiple servers. This is enabled for Expression types other than Utilization.
	Stacked Area. View cumulative workload data from multiple servers. This is enabled for Expression types other than Utilization.
	Legend Box. Give a visual cue and description for each data object. The legend is very useful for quickly glancing through the data on a chart.
	Point Labels. Display labels for the values in a chart.
	Zoom. Switch between normal and magnified view.
	Zoom In/Out. Adjust the granularity of the chart for the collected data. See Section 5.1.3, "Level of Detail," on page 87 .
	Navigate Beginning/Back/Forward/End. Scroll the visible data.
	Historical Data Settings. Select the historical data time range upon which to base forecasting.
	Forecast Settings. Select the forecasting time period and algorithm. This button is only enabled for daily or weekly Series charts or for Profile charts.
	Refresh All Data From Server. Refresh the view.






D.6 Job Explorer

Table D-9 *Job Explorer Toolbar Commands*

Commands	Tooltip/Description
	Suspend: Suspends a running job.
	Resume: Resumes a paused job.
	Abort: Abort the currently running job.
	Archive: Archives the selected job.
	View Logs: Opens a log viewer for the selected step.



D.7 Scheduled Job Explorer

Table D-10 *Scheduled Job Explorer Toolbar Commands*

Commands	Tooltip/Description
	Pause: Pauses an active or scheduled job.
	Resume: Resumes a paused job.
	Edit Recurrence: Opens the <i>Create Recurrence</i> dialog box where you can create or edit a scheduled job recurrence pattern.
	Run Now: Runs the scheduled job immediately.
	Delete: Deletes the scheduled job.

D.8 Log Explorer

Table D-11 *Log Explorer Toolbar Commands*

Commands	Tooltip/Description
	View Details: Opens Log Details for the selected log entry.
	Delete: Deletes the selected log entry. To simultaneously delete multiple logs, select the logs using the Shift or Ctrl keys and press the Delete key or click Delete.

Best Practices

E

These Best Practices are designed for data center operators and administrators who use PlateSpin Recon to monitor workloads and do consolidation planning based on collected data. The information provided here has been compiled by the PlateSpin support team based on experience with hundreds of Real-Time Consolidation Accelerator users.

Where applicable, guidelines in Best Practices refer to PlateSpin Support Knowledge Base articles. To access the articles, use your PlateSpin download credentials. If you have forgotten your password, you can request a password reminder e-mail at: <http://www.platespin.com/downloads/forgotpassword.aspx> (<http://www.platespin.com/downloads/forgotpassword.aspx>).

- ♦ [Section E.1, “Consolidation Planning Workflow,” on page 217](#)
- ♦ [Section E.2, “PlateSpin Recon Server System Requirements,” on page 218](#)
- ♦ [Section E.3, “Monitoring Service Configurations,” on page 218](#)
- ♦ [Section E.4, “MS SQL Database Configurations,” on page 220](#)

E.1 Consolidation Planning Workflow

Use this generic consolidation planning workflow:

Table E-1 *Best Practices Usage*

Guideline	Knowledge Base Article
Step 1: Inventory the servers.	
Step 2: Monitor the servers and collect data.	
Step 3: Analyze the collected data and create consolidation projects and scenarios by using PlateSpin Recon charts and the Consolidation Planning Module.	
When you have finished, consolidate the servers by streaming physical servers into virtual environments by using PlateSpin PowerConvert.	
For early detection of potential inventory and monitoring issues, use PlateSpin Recon PlateSpin Analyzer to scan your network before starting the inventory.	Q7920825 (http://www.novell.com/support/viewContent.do?externalId=7920825) : Using PlateSpin Analyzer and understanding its results.
If possible, use Microsoft Operations Manager (MOM) 2005 and set up PlateSpin Recon to connect and collect monitoring data directly from it. This feature allows you to reduce your project timeline from weeks to days.	Q7920903 (http://www.novell.com/support/viewContent.do?externalId=7920903) : PlateSpin Recon with MOM 2005

E.2 PlateSpin Recon Server System Requirements

To perform PlateSpin Recon operations on more than 300 machines, it is recommended that the PlateSpin Recon Server meets the following system requirements:

- ♦ **Hardware Requirements:**

Memory: 4 GB or more

Processor: 3.2 GHZ or more with quad core

- ♦ **Software Requirement:**

Operating system: Windows 2008 SP2 (32-bit)

E.3 Monitoring Service Configurations

- ♦ PlateSpin Recon 3.7.2 provides the Windows Disk Counters add-on to monitor the disk utilization of a server. The counters are registered in the perfmon of the target server and are processed through the Registry services. Hence, you must not monitor the PlateSpin Recon Server for the optimal performance of the server.
- ♦ If the *Unresponsive* filter group in the PlateSpin Recon Client (*Data Center Explorer > Filters > Status*) contains more than 20 inaccessible machines for more than two hours, you must stop monitoring all the inaccessible machines. You can start monitoring these machines after they can be accessed from the PlateSpin Recon Server. To verify if the machines are accessible, ping them from the server.
- ♦ To improve the performance of the PlateSpin Recon Server when you want to monitor more than 500 machines, you can fine-tune the values of a few monitoring service configuration parameters. The values of some of the parameters can be changed through the PlateSpin Recon Client and the value of other parameters can be changed by editing the Monitoring service configuration file.
 - ♦ [Section E.3.1, “Fine-Tuning the Monitoring Service Configuration Parameters Through the PlateSpin Recon Client,” on page 218](#)
 - ♦ [Section E.3.2, “Fine-Tuning Monitoring the Service Configuration Parameters in the Monitoring Service Configuration File,” on page 219](#)

E.3.1 Fine-Tuning the Monitoring Service Configuration Parameters Through the PlateSpin Recon Client

- 1 In the PlateSpin Recon Client, click the *Tools* menu.
- 2 Press the Ctrl key and click *Options*.

The *Server* page of the Options dialog box is displayed by default.

- 3** In the *Monitoring Service* category, fine-tune the value of the configuration parameters as recommended in the following table:

Configuration Parameter	Description	Default Value	Recommendation
Sample Interval	Indicates the time interval between subsequent sampling of data from remote machines.	5 minutes	Increase the sample interval value.
Burst Samples	Indicates the number of data samples to be collected per sample interval.	3	Reduce the burst value.
Burst Interval	Indicates the time between burst samples.	20 seconds	Increase the interval to 60 seconds.
Cache service Interval	Indicates the time interval to commit the collected data in to the Monitoring database.	1 minute	Increase the interval to 2 minutes.

E.3.2 Fine-Tuning Monitoring the Service Configuration Parameters in the Monitoring Service Configuration File

- 1 On the PlateSpin Recon 3.7.2 Server, open the `PlatespinRecon3.7.2_Server_installation_directory\Services\Monitoring\PlateSpin.Monitoring.Service.exe.config` file using a text editor.
- 2 Change the value of the configuration parameters as recommended in the following table:

Configuration Parameter	Description	Default Value	Recommendation
ConnectionThreadMaxNumber	Indicates the number of parallel connections that can be made.	400	The default value is applicable for 500 machines. If you encounter thread execution timeout issues when you monitor more than 500 machines, change the value to 300.
CacheServiceChunkSize	Indicates the number of records per transaction to be inserted in to the History table.	4000	If you want to monitor more than 500 machines, change the value to 2000. This reduces delays in data insertion in to the History table.
QUEUEPRUNNINGFACTOR	When you multiply the value of this parameter with the value of the Sample Interval parameter, the result determines the time interval to purge the old requests from the queue.	3 minutes	By default, 15 minutes is the time interval to purge the old requests from the queue (3 minutes * 5 minutes = 15 minutes). Change the value of the parameter to 2 for requests older than 10 minutes to be deleted for improving the performance.

3 Restart the PlateSpin Recon Monitoring Service.

E.4 MS SQL Database Configurations

If your database is an MS SQL database, refer to [Table E-2, “Fine-tuning MS SQL Database Configuration,” on page 221](#) to fine-tune the database performance when you want to monitor more than 500 machines. The table provides instructions to perform the tasks by using Microsoft SQL Server Management Studio; however, you can use any other database management tool to perform the tasks.

Table E-2 Fine-tuning MS SQL Database Configuration

Task	Description	How to Perform the Task	When to Perform the Task
Perform Read commit snapshot ON on the PlateSpin Recon Monitoring database	Enables fast data read queries.	<ol style="list-style-type: none"> 1. Stop the PlateSpin Recon Monitoring Service. 2. In the Microsoft SQL Server Management Studio, run the following query on the PlateSpin Recon Monitoring Database: <pre>ALTER DATABASE database_name SET READ_COMMITTED _SNAPSHOT ON</pre> 3. Start the PlateSpin Recon Monitoring Service. 	
Rebuild index on the CounterInstance table for IX_CounterInstance	Enables fast join operations required during summarization.	<ol style="list-style-type: none"> 1. Stop the PlateSpin Recon Monitoring Service. 2. In the Microsoft SQL Server Management Studio, browse to <i>PlateSpin Recon Monitoring database > Tables > Counter Instance > Indexes</i>. 3. Right-click the <i>IX_CounterInstance</i> table, then click <i>Rebuild</i>. 4. Verify the properties of <i>Index</i> to ensure that fragmentation % is close to 0. If fragmentation is greater than 50%, rebuild the index. 5. Start the PlateSpin Recon Monitoring Service. 	You can perform this task once a week.

Task	Description	How to Perform the Task	When to Perform the Task
Reorder Index on PR_HistoryData table	Enables faster data insertions in to the History table and requires less pages need to locked because it controls fragmentation level below 50%.	<ol style="list-style-type: none"> 1. Stop the PlateSpin Recon Monitoring Service. 2. In the Microsoft SQL Server Management Studio, browse to <i>PlateSpin Recon Monitoring database > Tables > PR_HistroyData > Indexes</i>. 3. Right-click <i>IX_PR_HistoryData</i> , then click <i>Properties</i>. 4. On the <i>General</i> page, move up <i>ObservationDate</i> to be listed first in <i>Index Key Columns</i>. 5. Start the PlateSpin Recon Monitoring Service. 	You must perform this task after installing PlateSpin Recon.
Rebuild index on PR_HistoryData table for IIX_PR_HistoryData index	Enables faster data insertions and deletions from the History table.	<ol style="list-style-type: none"> 1. Stop the PlateSpin Recon Monitoring Service. 2. In the Microsoft SQL Server Management Studio, browse to <i>PlateSpin Recon Monitoring database > Tables > PR_HistroyData > Indexes</i>. 3. Right-click <i>IX_PR_HistoryData</i> , then click <i>Rebuild</i>. 4. Verify % fragmentation. If fragmentation is greater than 50%, rebuild the index 5. Start the PlateSpin Recon Monitoring Service. 	You can perform this task once a week.

Task	Description	How to Perform the Task	When to Perform the Task
Restrict Transaction log file size on the PlateSpin Recon Monitoring database to 8 GB	Keeps transaction log file size in control when large number of machines are monitored, and enables database to commit data more frequently because the transaction log file cannot contain large data.	<ol style="list-style-type: none"> 1. Stop the PlateSpin Recon Monitoring Service. 2. In the Microsoft SQL Server Management Studio, right -click the <i>PlateSpin Recon Monitoring database</i>, then click <i>Properties</i>. 3. On the <i>Files</i> page, click the <i>AutoGrowth</i> option for the Monitoring Log. 4. Ensure that the <i>Restricted File Growth</i> is not more than 8 GB. 5. Start the PlateSpin Recon Monitoring Service. 	Currently, the transaction log file size is by default 2000 GB. You must restrict this when you want to monitor a large number of nodes.

Documentation Updates

F

This section contains information on documentation content changes that were made in this *User Guide* after the initial release of Novell PlateSpin Recon 3.7.2. The information can help you to keep current on updates to the documentation.

The documentation is provided on the Web in two formats: HTML and PDF. The HTML and PDF documentation are both kept up-to-date with the documentation changes listed in this section.

If you need to know whether a copy of the PDF documentation you are using is the most recent, the PDF document contains its publish date on the front title page.

The documentation was updated on the following date:

- ♦ [Section F.1, “June 30, 2011,” on page 225](#)
- ♦ [Section F.2, “March 22, 2011,” on page 225](#)
- ♦ [Section F.3, “March 07, 2011,” on page 226](#)
- ♦ [Section F.4, “February 17, 2011,” on page 226](#)

F.1 June 30, 2011

Updates were made to the following section:

Location	Update
Chapter 1, “Introduction,” on page 13	<ul style="list-style-type: none">♦ Added support for SUSE Linux Enterprise Server (SLES) 10 SP3 in Supported UNIX/Linux Platforms section.♦ Added support for SUSE Linux Enterprise Server (SLES) 10 SP2 and SP3 in Supported Hypervisors sections.

F.2 March 22, 2011

Updates were made to the following section:

Location	Update
Appendix A, “Troubleshooting,” on page 169	Updated the following scenario: “Creating scenarios for forecasted projects might fail with the out of memory exception” on page 183

F.3 March 07, 2011

Updates were made to the following section:

Location	Update
Section 3.2.5, "Viewing Inventory Data and Monitoring Servers," on page 47	Added following the note: IMPORTANT: When you monitor Windows Server 2003, the memory utilization on the monitored server might increase because the memory usage of the <code>Svchost.exe</code> process increases when other users remotely collect performance data about the monitored server. This is a known issue of Windows Server 2003. To troubleshoot the issue, see the Microsoft Knowledge Base article ID 969701 (http://support.microsoft.com/kb/969701) .

F.4 February 17, 2011

Updates were made to the following section:

Location	Update
Appendix B, "Technical Reference," on page 187	Added Section B.10, "Managing Credentials," on page 207 .